NET WORK LISER

Discussion Paper Series

IZA DP No. 19040

October 2026

Rent Sharing under Innovation Tax Breaks

Sabien Dobbelaere
Vrije Universiteit Amsterdam,
Tinbergen Institute
and IZA@LISER

Ana Popovici
Vrije Universiteit Amsterdam
and Tinbergen Institute

The IZA Discussion Paper Series (ISSN: 2365-9793) ("Series") is the primary platform for disseminating research produced within the framework of the IZA@LISER Network, an unincorporated international network of labour economists coordinated by the Luxembourg Institute of Socio-Economic Research (LISER). The Series is operated by LISER, a Luxembourg public establishment (établissement public) registered with the Luxembourg Business Registers under number J57, with its registered office at 11, Porte des Sciences, 4366 Esch-sur-Alzette, Grand Duchy of Luxembourg.

Any opinions expressed in this Series are solely those of the author(s). LISER accepts no responsibility or liability for the content of the contributions published herein. LISER adheres to the European Code of Conduct for Research Integrity. Contributions published in this Series present preliminary work intended to foster academic debate. They may be revised, are not definitive, and should be cited accordingly. Copyright remains with the author(s) unless otherwise indicated.

IZA DP No. 19040 October 2026

Rent Sharing under Innovation Tax Breaks*

Abstract

We study whether firms share gains from innovation tax incentives with incumbent workers and whether such rent sharing is associated with increased worker retention. The Dutch innovation box reduces tax rates on profits from in-house innovation. Linking first-time beneficiaries to matched employer-employee data, we combine nearest-neighbour matching with staggered difference-in-differences. Part of the additional after-tax surplus is passed on to incumbent workers: wages rise by 3.27% on average, with increases of 2.05% for workers least involved in R&D and 4.14% for tertiary-educated STEM workers. Descriptively, treated non-R&D incumbents are 3.47 percentage points less likely to switch employers following treatment.

JEL classification

H25, J31, J62, O31, O38

Keywords

innovation policy, rent sharing, worker mobility

Corresponding author

Ana Popovici
a.popovici@vu.nl

* The authors would like to thank Eric Bartelsman, Ana Figueiredo, and participants at the Vrije Universiteit Amsterdam economics seminar for their comments and suggestions. The paper also benefited from insightful comments and suggestions at the 28th Colloquium on Personnel Economics, the Economic Statistics Centre of Excellence Conference on Economic Measurement, the Paris School of Economics Workshop in Labour Economics, and the 38th European Association of Labour Economists (EALE) Conference. The authors are also grateful to Statistics Netherlands (CBS) for providing access to confidential microdata.

1 Introduction

Innovation tax breaks increase firms' after-tax returns to successful innovation, enlarging the surplus that may be shared with incumbent workers through periodic wage renegotiations. Such rent sharing may extend beyond workers directly involved in R&D to those who have accumulated valuable firm-specific production, organisational, and process knowledge. Firms may have an incentive to share part of the additional surplus with these workers to preserve valuable worker-firm matches. Thus, wage gains may be accompanied by subsequent changes in worker retention. From a job-ladder perspective, higher wages at the current employer make outside job offers less attractive, potentially reducing worker mobility. Understanding these mechanisms is important for evaluating not only who benefits from innovation policy but also whether the distribution of increased after-tax returns is associated with changes in labour mobility and the reallocation of workers across firms.

We examine these mechanisms using the Dutch innovation box, an ex-post output-based innovation tax incentive that substantially increases the after-tax profits generated by in-house innovation. During our treatment window, the policy reduced the corporate tax rate on qualifying innovation profits from 25% to 5%. Importantly, the benefit is realized only once an innovation generates sufficient profits. The policy therefore differs from ex-ante input-based R&D subsidies that lower the cost of innovation before its commercial success is known. By increasing retained profits, the innovation box provides a suitable setting to investigate whether firms share part of the additional firm surplus with their workers through wage increases, and whether these wage gains are associated with changes in worker mobility.

The Dutch institutional setting makes wage responses and potentially associated subsequent worker retention particularly relevant. The wage-setting system combines sectoral collective bargaining with firm-level discretion, allowing scope for firm-specific pay responses to innovation tax breaks. Another feature of the Dutch labour market is its comparatively strong employment protection and high dismissal costs relative to many OECD countries (OECD, 2020). Replacing an incumbent worker can therefore be costly not only because a productive match and firm-specific knowledge are lost, but also because separation and subsequent recruitment entail adjustment costs. These labour market frictions may influence firms' willingness to renegotiate wages. The innovation box raises retained profits and thereby strengthens firms' internal financial position. This capital-side channel may reinforce firms' incentives to remain on the profitability path created by successful innovation. Our empirical analysis focuses on the labour market side of this mechanism: whether the additional after-tax surplus is shared with incumbent workers and whether wage gains are accompanied by stronger worker attachment to the firm.

We link administrative information on all firms first benefiting from the innovation box between 2010 and 2018 to population-wide matched employer-employee data from Statistics Netherlands. The worker data cover 2007-2023 and allow us to follow wages, hours, and employer transitions over a long horizon. We focus on incumbent workers, defined as workers continuously employed at the treated firm from three years before treatment through the first year in which the firm receives the tax benefit. These workers predate the realized tax benefit and are therefore particularly well suited for studying whether rents generated by additional retained profits are shared with the existing workforce.

Innovation box beneficiaries are not a random sample of firms, and their workers differ systematically from workers at non-beneficiary firms. We therefore combine a nearest-neighbour matching

procedure with a staggered difference-in-differences design. Treated incumbents are matched to observationally similar incumbent workers at firms that never benefit from the innovation box using rich pre-treatment information on worker characteristics, wage histories, and firm characteristics. On the resulting matched panels, we estimate dynamic treatment effects using the interaction-weighted estimator of Sun and Abraham (2021). This estimator accommodates heterogeneous treatment effects across treatment cohorts or exposure length, and avoids the contamination problems of conventional two-way fixed-effects. Our main analysis follows approximately 728,000 incumbent workers, including about 378,000 treated workers.

We find clear evidence of rent sharing from innovation tax breaks. After firms begin benefiting from the innovation box, incumbent workers' annual base wages increase by 3.27% on average relative to matched control workers. The wage effect emerges gradually and persists over the post-treatment horizon. Importantly, the gains are not limited to the workers who are more likely to have directly generated the innovation. Workers with below-upper-secondary education, whom we classify as least likely to be directly involved in R&D, experience an average wage increase of 2.05%. Tertiary-educated workers with a STEM field of study, who are most likely to be involved in innovative activities, experience a larger increase of 4.14%. Across various worker dimensions (including age, gender, parental migration background, education level, and field of study), the wage effect remains positive. The incidence of innovation rents thus extends beyond inventors and highly educated R&D workers to the broader incumbent workforce.

We complement our rent-sharing analysis with descriptive evidence on worker mobility. On average, incumbent workers at treated firms are 2.19 percentage points less likely to switch employers following treatment, with the reduction driven by non-R&D workers (3.47 percentage points). For these workers, lower mobility becomes visible with a lag, around two years after treatment, following the gradual increase in base wages. By contrast, we find no significant mobility response for R&D workers, despite their larger wage gains. Non-R&D workers who leave treated firms also experience larger wage gains upon switching employers than comparable leavers from control firms. These patterns are consistent with a job-ladder interpretation: higher wages at the current employer may make outside offers less attractive, so workers who do leave require larger wage improvements. While these descriptive patterns do not establish that wage gains cause increased retention, they suggest that tax-break-induced rent sharing may be associated with more stable employment relationships, particularly for non-R&D incumbents who have accumulated valuable firm-specific knowledge.

Our work contributes to three strands of literature that are typically studied separately. First, we add to the literature on innovation policy. Existing studies have mainly asked whether tax incentives increase R&D spending, patenting, or other innovation outcomes, and whether preferential tax regimes induce profit shifting (Agrawal et al., 2020; Cappelen et al., 2012; Chen et al., 2021; Dechezleprêtre et al., 2023; Schwab, 2021; Haufler and Schindler, 2023). We instead study real labour market outcomes of an output-based innovation incentive, an incidence margin that remains comparatively underexplored (Bornemann et al., 2023; Chen and Hshieh, 2024). Our results show that part of the fiscal gain accruing to successful innovative firms is passed on to their existing workforce.

Second, we contribute to the literature on rent sharing and innovation. Previous research shows that workers can capture part of the rents generated by innovation and other firm-level successes (such as productivity gains), with particularly strong effects for inventors and highly skilled workers

(Aghion et al., 2023, 2018; Bessen et al., 2025; Kline et al., 2019). We show that rent sharing following a policy-induced increase in retained profits is broader: significant wage gains extend to workers unlikely to be directly involved in R&D. This finding points to the role of firm-specific knowledge accumulated throughout the organization, not only by the workers formally engaged in inventive activity.

Third, we connect rent sharing to the literature on labour market reallocation. Studies on corporate tax incidence show that changes in firms' tax liabilities can affect worker compensation through bargaining over after-tax rents (Arulampalam et al., 2012; Dwenger et al., 2019; Fuest et al., 2018). We extend this literature by examining whether wage gains induced by an innovation tax break are accompanied by changes in worker mobility. Our findings suggest that such pass-through may be associated with preserving valuable worker-firm matches, particularly for non-R&D workers who have accumulated valuable firm-specific knowledge over time. While these mobility results are descriptive, they point to a potential link between tax-break-induced rent sharing and the reallocation of labour across firms.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting of the Dutch labour market and the innovation box. Section 3 presents the matched employer-employee data and the construction of the incumbent-worker panels. Section 4 discusses the empirical strategy. Section 5 reports the main rent-sharing results and descriptive evidence on worker mobility. Section 6 presents robustness checks. Section 7 discusses the policy implications and concludes.

2 Institutional Setting

2.1 Labour Market Institutions

The extent to which firms share rents with their workers depends on the institutional context of wage setting. The Dutch system combines coordination at the national and sectoral levels with scope for wage setting at the firm level. Collective bargaining covers roughly three quarters of employees, and agreements are concluded both at the sector and the firm level. In 2019, 75% of employees covered by a collective agreement were covered by a sectoral agreement and 25% by a firm-level agreement. The system comprised more than 500 company-specific agreements and almost 200 sector-wide agreements (Jansen, 2021; International Labour Organization, 2022).

Sectoral agreements commonly specify occupation- and seniority-related wage scales and general wage increases, which are an important source of wage growth for many workers (Hartog and Salverda, 2018). At the same time, Dutch wage setting has become increasingly decentralised. Boumans (2022) documents a long-run shift towards customised firm-level wage setting with more decentralisation, which encourages remuneration instruments such as one-off payments, profit sharing, and flexible remuneration schemes. Consequently, Ibsen and Keune (2018) characterise the Netherlands as a system of "organised decentralisation": sectoral agreements continue to provide a framework and minimum conditions, while leaving room for terms and conditions to be determined or further elaborated at the company level (e.g. through works councils). This combination of broad collective-bargaining coverage and firm-level discretion is particularly relevant for our setting. The innovation box generates a firm-specific increase in after-tax surplus, and the institutional framework leaves scope for firm-specific pay responses through which part of this surplus may be shared

with incumbent workers.

Employment protection also affects the margin through which firms respond to this additional surplus. Dutch employment relationships are segmented between permanent ("open-ended") and fixed-term ("temporary") contracts. Employment protection legislation, including protection against dismissal, is very strong for the former but not for the latter (OECD, 2020). Temporary employment is widespread: temporary contracts accounted for around 20% of employees in the Netherlands, compared with around 12% on average across OECD countries in 2019 (OECD, 2021). Moreover, Hartog and Salverda (2018) report that the share of flexible jobs, comprising jobs with a temporary contract, increased from around 16% in 2001-2004 to 26% in 2016. For firms employing incumbents on open-ended contracts, separation can therefore entail institutional costs in addition to the loss of firm-specific knowledge and the costs of recruiting and training replacements. These frictions may strengthen firms' incentives to raise wages to preserve valuable worker-firm matches, providing an institutional link between rent sharing and our descriptive evidence on retention.

2.2 Innovation Tax Breaks

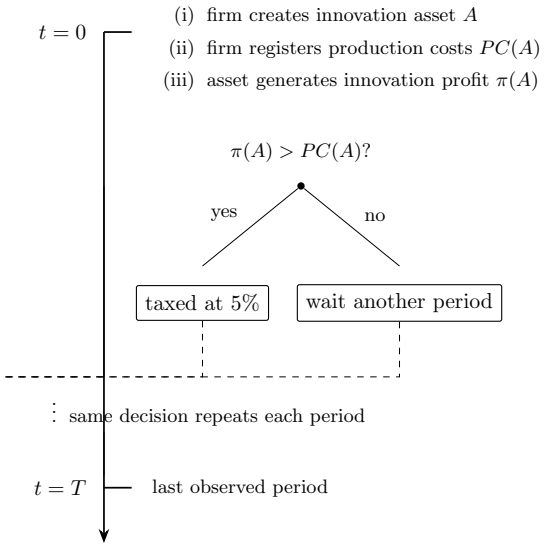
The Dutch innovation box policy reduces the effective corporate tax rate on profits generated from in-house innovation from the statutory rate of 25% to a preferential rate of 5%.[1] We follow all firms that benefited from the policy between 2010 and 2018. By increasing the after-tax surplus, the policy creates a suitable setting for assessing whether retained profits are passed on to workers. The retained profits generated by the innovation box are large. For every euro of qualifying innovation profit, the firm keeps $0.25 - 0.05 = 0.20$ euros, an 80% reduction in the tax burden on that profit stream. The gain is tied directly to the profitability of registered intangible assets rather than to any contemporaneous change in the firm's production technology or workforce composition. We consider the extensive margin of firms benefiting from the innovation box.[2]

The innovation box policy is available to all firms operating in the Netherlands. To qualify, a firm must register its eligible intangible assets with the Dutch Enterprise Agency (RVO) when filing its tax return. Only assets developed in-house qualify. In practice, qualifying activities span a wide range of product and process innovations, giving the policy broad sectoral coverage. Before any profit is taxed at the preferential 5% rate, the firm must fully recover the initial production costs associated with the registered asset. Once production costs are recovered, all subsequent qualifying profits are taxed at 5% for as long as the registered asset is profitable. Figure 1a illustrates the timing of the policy. If the firm's qualifying profits already exceed the production-cost threshold at $t = 0$, the first tax benefit accrues immediately. Otherwise, the benefit is deferred until the threshold is crossed. In practice, deferral is rare: 98.2% of firms receive their first tax break in the same year they apply. Despite the large tax benefit, policy uptake has been limited. Between 2010 and 2018, a total of 10,203 unique firms (approximately 2% of the Dutch population of firms) benefited from the innovation box. For accounting details and the policy's operational timeline in practice, we refer to Appendix A.1.

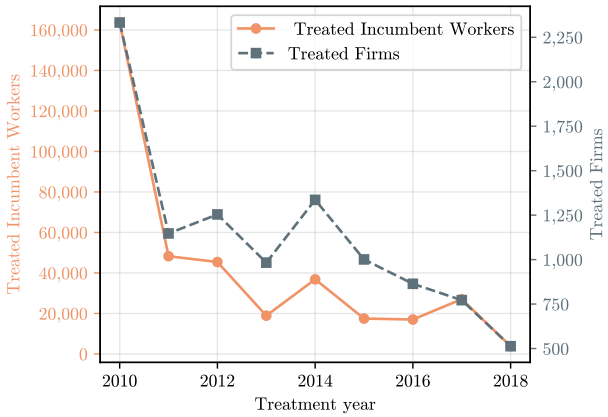
[1] The policy was introduced in its current form in 2010, replacing the narrower Patent Box regime of 2007. The preferential rate was subsequently raised to 7% in 2018 and to 9% in 2021, leaving the core mechanism unchanged throughout our 2010-2018 analysis window. This ex-post profit-side instrument coexists with an ex-ante cost-side instrument, the WBSO (acronym for "The Wage Tax and Social Insurance Act") wage tax subsidy for R&D personnel. We provide a detailed comparison of these instruments and their potential interactions in Appendix A.1.

[2] Data limitations prevent us from examining the intensive margin.

Figure 1: Policy Timing and the Participants



(a) Timing of innovation box treatment



(b) Treated incumbent workers and firms by year

Notes: Figure 1a illustrates how innovation box treatment status is determined at the firm level in each period. A firm's innovation profit ($\pi(A)$) is compared with its unrecovered production costs ($PC(A)$). Once innovation profit exceeds these costs, the firm qualifies for the reduced 5% tax rate. Otherwise, eligibility is reassessed in the following period. Figure 1b plots the number of incumbent workers (orange) and firms first receiving the innovation box benefit (blue) by treatment cohort between 2010 and 2018. Treated incumbent workers are defined as workers observed at the same treated firm for four consecutive years, $t-3$ through t, where t denotes the firm's treatment year. The sample underlying Figure 1b comprises 378,409 treated incumbent workers employed at 10,203 treated firms.

Innovation box beneficiaries are not a random draw from the population of Dutch firms. In Appendix A.3, we document how treated firms differ from the broader population of Dutch firms. At first, beneficiaries are larger and have more automation expenditures compared to the average Dutch firms. However, the population of beneficiaries converges toward the general population of Dutch firms over time. Firms in later treatment cohorts are smaller and spend less on automation: from 2015 onward, firm size no longer predicts treatment, and from 2017 onward, neither do automation expenditures. This selection is not a concern for the central question of our paper because our analysis focuses on how the additional profits retained by innovation box beneficiaries affect their incumbent workers.

Figure 1b plots the annual flow of first-time beneficiary firms. The number of benefiting firms decreases over time. In the early years of the policy, a substantial stock of eligible firms have entered treatment, whereas in later years the annual inflow of new beneficiaries gradually stabilises. One possible explanation is offered by complicated accounting standards: firms might find it hard to estimate whether they would qualify from innovation or might not have the manpower to manage the application procedure. This explanation is based on discussions with industry professionals specialised in consulting firms through their application procedures. Overall, many innovative firms that would quality for the policy, do not participate. Innovation box beneficiaries account for only 16.7% of innovative firms in the Netherlands.[3] Because firms enter innovation box treatment in the same year they apply and the benefiting pool comprises only a subset of innovative firms in the Netherlands, we are able to capture rent sharing from tax breaks rather than from changes in production processes and innovation. Any changes in production processes would have been implemented prior to the first benefit. Since treatment is assigned at the firm level, workers do not select into treatment. We therefore focus on individual rather than firm-level wage outcomes, examining how innovation box tax breaks affect the wage trajectories of incumbent workers.

Figure 1b also shows the number of affected incumbent workers in each treatment cohort. We define a treated incumbent workers as a worker who has worked at a treated firm for four consecutive years (three years preceding firm treatment and the treatment year itself). The number of affected workers closely tracks the number of beneficiary firms. In total, we identify 378,409 workers who may potentially receive a share of the additional profits retained by firms benefiting from the innovation box. Using administrative data tracking all employment records, we have access to a large pool of workers at innovative firms that do not benefit from the tax break. These workers form a subset of our control group from which we select matched incumbents.

3 Data and Estimation Samples

3.1 Matched Employer-Employee Data

We use administrative microdata from Statistics Netherlands to construct the population wide matched employer-employee data. A detailed description of the data structure and datasets can be found in Appendix B.1. Payslip information comes from monthly employee records (Statistics Netherlands, 2024c) and is available over the period 2006-2023. We aggregate to the annual level,

[3] Innovative firms are defined as firms benefiting from the ex-ante input-side policy instrument, the WBSO, rather than through self-reported R&D expenditures. We provide more details on the calculation of the population of innovative firms in Appendix A.2.

keeping only workers who worked more than 520 hours in a given year at a firm and earned at least 90% of the annual minimum wage at the same firm. We trim outliers at the 1% level for both base wages and hours worked, grouped by year. We drop workers on internships and Social Employment Act contracts. For all remaining workers, we retrieve demographic data (Statistics Netherlands, 2024a) and their highest obtained level of education (Statistics Netherlands, 2024b). While the payslip and demographic data are present for the full population of Dutch workers, education data are compiled from an annual survey. The coverage has improved over the years but remains incomplete and biased toward highly educated workers. For that reason, we impute missing values using a two-step procedure: (1) for workers with at least one non-missing observation, we use their first non-missing observed value; (2) for workers without non-missing observations across the panel, we predict their education level and whether their field of study is STEM or non-STEM using an inverse Mincer equation. We provide imputation procedure details in Appendix B.2 and robustness checks on this imputation in Section 6.

We restrict the sample to workers at firms with more than two employees [4], where firm size is taken from the Business Demography Framework (BDK) dataset. Following Abowd et al. (1999), we construct one spell per worker-year-firm combination, imposing the following restriction: if a worker has two spells in a given year, we choose the spell with the largest number of hours worked. Because the payslip dataset starts in 2006, we cannot observe worker tenure before this point, therefore we construct worker tenure by treating 2006 as each worker's earliest possible start date, which left-censors tenure for any worker already employed in 2006. We follow workers from 2007 to 2023. This leaves us with a matched employer-employee panel of approximately 70 million observations covering 7 million unique employees. Appendix B.1 provides additional information on the matched employer-employee data.

3.2 Incumbent Worker Panels

From the matched employer-employee data, we select all possible incumbent workers within our treatment window and split this pool into treated incumbents and all possible control incumbents. We define treated incumbent workers as those present at the firm for four consecutive years, up to the treatment year. We focus on them because these are the workers most likely to have built the firm-specific knowledge and bargaining position needed to capture a share of post-tax break rents. Due to the staggered adoption of treatment, defining incumbent samples across cohorts is not trivial. Each treatment cohort needs its own pre-treatment window, and a worker's status must be pinned down relative to their own firm's treatment year rather than to calendar time. We follow the approach used in the job-loss literature for defining pre-displacement worker cohorts (Britto et al., 2022; Dos Santos et al., 2026; Jarosch, 2023). For any treatment year t between 2010 and 2018, we select all workers employed at the firm in $t-3$ who remained continuously employed there through t. This restriction yields yearly incumbent pools of 2.5 to 3 million observations, from which we construct a treated incumbent pool (workers at firms treated in year t) and a control incumbent pool (all incumbents at never-treated firms).

[4] We exclude firms with two or fewer employees, as very small business units in the source data often correspond to sole proprietorships rather than firms with a standard employer-employee relationship. The number of employees is calculated as the annual average number of employees at the business unit. We decided to keep workers at small firms because effects of the policy on start-ups and young firms are important. For this reason, we run extensive results on firm heterogeneity in Section 5.1.

The incumbent worker definition implicitly assumes that incumbents were not hired in anticipation of the policy, a concern that applies most directly to R&D workers themselves. We address this concern further by distinguishing workers in terms of most likely and least likely to have been directly involved in the innovation process. We use education and field of study as a proxy for making this distinction, following established ISCED (UNESCO Institute for Statistics, 2015) and OECD (OECD, 2025) standards.[5] We rely on these proxies because occupation data is not registered in any of our available microdata.

Treated incumbents differ systematically from a random draw of incumbents: they earn more, work more hours, and are more likely to be male. These differences stay constant over all treatment and control cohorts. Because of this, we employ a matching strategy when constructing our final analysis panels (see Section 3.3). Since only a small share of incumbents are treated in any given year, the large pool of potential controls allows us to identify closely matched workers.

3.3 Estimation Samples

We build three estimation samples, constructed from the same matching procedure but differing in the education and field-of-study restriction applied to workers. Sample A (All Workers) includes all matched treated and control incumbents. Sample B (non-R&D) restricts to workers with below-upper-secondary education. Sample C (R&D) restricts to tertiary-educated workers in STEM fields of study. We use this Sample A/B/C naming convention consistently for the remainder of the paper. Sample B workers, with below-upper-secondary education, are the least likely to be R&D employees and therefore the least likely to have been hired in anticipation of the policy. We consider Sample B estimates as the cleanest causal estimate of the policy's effect (see Section 4). Sample C workers, tertiary-educated in STEM fields, are the group most likely to be doing R&D, and therefore the group whose wage response is potentially most exposed to anticipatory hiring.

To match treated and control workers, we develop a tailored matching algorithm that accounts for staggered treatment adoption. An extensive description of the algorithm is provided in Appendix B.4. We employ a 1-nearest-neighbour (1-NN) approach, which offers more flexibility than propensity score matching when working with a large number of features (Caliendo and Kopeinig, 2008; Dehejia and Wahba, 2002; Rosenbaum and Rubin, 1983). Matching is performed separately for each treatment cohort, and control workers are always drawn from incumbents at never-treated firms. Within a cohort, matching is with replacement, so two treated workers can be assigned the same control worker, and a control worker matched in one cohort is not used again in later cohorts.

Control incumbents work at firms that never benefit from the innovation box, but that does not mean they work at firms that do not innovate. Take-up of the innovation box is low, about 2% of Dutch firms, and many firms with qualifying innovations never claim it (see Section 2.2 and Appendix A.2), so the control group from which we select matched incumbents is sizable (see Section 3.2). What separates the two groups is the firm-level tax treatment of innovation profits. Comparing incumbents at firms receiving the tax break with observationally similar incumbents at firms that do not allows us to isolate the effect of the tax break on incumbent workers. This motivates our focus on individual incumbent workers rather than firm-level aggregates. Our research design addresses selection and changes in workforce composition. We follow the same workers who were employed

[5] We use three broad education categories, from lower to higher: below upper secondary, upper secondary and tertiary. For the field of study, we distinguish STEM from non-STEM studies following the OECD (2025) surveys.

at the firm three years before the first benefit year, so changes in workforce composition cannot drive our estimates.

Each treated incumbent is compared with a control incumbent matched on individual and firm characteristics. Our preferred specification matches on the following individual characteristics: base wages, average within-firm base wage growth, previous-period base wages, base hours worked, overtime hours worked, tenure, gender, parental migration background, age, education level, and field of study. We additionally match on the following firm characteristics: firm size, firm age, firm location, industry, and firm quality, which is defined as the firm fixed effect obtained from an AKM (Abowd et al., 1999) estimation.[6] Robustness of the matching results to the choice of matching features is discussed in Appendix B.4. Treatment is assigned at the firm level, so incumbents do not select into treatment, and worker fixed effects absorb all time-invariant differences between them.

Our matching algorithm performs well. In Appendix B.4, we present extensive tests and results of the matching procedure. Using our preferred matching specification, the samples are balanced on all matching characteristics except for firm size. This is because some innovation box users are very large corporations. As a heterogeneity check in our main results, we exclude these large corporations and find that this does not strongly affect our estimates (see Section 5.1). Kernel density functions comparing treated, matched control, and unmatched control workers show that the distributions of treated and matched control workers are closely aligned across all matching characteristics, especially for firm characteristics such as industry, location, and age. For individual characteristics, the kernel density functions of treated and matched control workers are nearly identical. These matching results reflect the large control pools from which workers are selected.[7]

We capture rent sharing through increases in annual base wages. The base wage is the fixed, contractually negotiated component of pay recorded in the payslip data, and is the component most directly tied to the periodic wage renegotiations through which rent sharing is institutionally transmitted (Section 2.1). Other outcomes (bonuses, overtime income, base hours, and overtime hours) are reported as secondary outcomes, to further test how the incidence of the innovation tax break is transmitted.

Table 1 reports descriptive statistics and standardized mean differences (SMD) for selected characteristics between treated and matched control workers for the three estimation samples on which we estimate our main results. Our main analysis sample, Sample A, includes 7.2 million observations from 2007 to 2023, following 728,480 incumbent workers: 378,409 treated and 350,071 matched control. Sample B, the non-R&D workers, comprises 1,734,690 observations, following 165,032 incumbent workers: 89,632 treated and 75,400 matched control. Sample C, R&D workers, comprises 606,033 observations, following 61,805 incumbent workers: 39,385 treated and 22,420

[6] Using the universe of Dutch matched employer-employee data, we estimate a standard AKM wage decomposition, regressing log earnings on worker and firm fixed effects, year effects, and a flexible polynomial of workers' age, following Dobbelaere et al. (2025). Worker and firm effects are separately identified through workers who move across employers. The worker fixed effect captures time-invariant worker characteristics, such as ability, that are rewarded similarly across firms. The firm fixed effect captures the wage premium received by workers employed at a given firm. Following Davidson et al. (2014) and Xu (2025), we interpret this firm fixed effect as a proxy for firm quality that characterises the economic opportunities available across employers.

[7] We could also have matched workers within industries and firm-size categories. However, this would have substantially reduced the pool of potential control workers, compromising the quality of the matches between treated and control incumbents. Because our analysis is targeted to incumbent workers, we choose to make workers as similar as possible, and supplement our main analysis with detailed firm heterogeneity results in Section 5.1 and Appendix C.2.

matched control.

Table 1: Descriptive Statistics: Estimation Samples

	Sample A. Full Sample			Sample B. non-R&D (Below-Upper-Sec.)			Sample C. R&D (Tertiary STEM)		
	Treated	Matched Control	SMD	Treated	Matched Control	SMD	Treated	Matched Control	SMD
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Base Wage (annual, €)	36,591	36,859	-0.02	30,059	29,961	0.01	45,058	45,317	-0.02
Bonuses (annual, €)	5,958	5,415	0.08	4,367	3,851	0.16	7,370	7,290	0.01
Overtime Income (annual, €)	589.1	465.7	0.08	987.5	928.5	0.03	359.5	227.4	0.10
Base Hours (annual)	1,855	1,876	-0.07	1,953	1,970	-0.09	1,915	1,927	-0.05
Overtime Hours (annual)	26.38	21.22	0.08	47.50	45.02	0.03	12.89	8.50	0.10
Age (years)	42.35	42.12	0.02	46.77	46.37	0.04	38.31	38.36	-0.01
Gender (% male)	76.85	72.46	0.10	87.20	86.86	0.01	88.52	88.38	0.00
Parental Migration Background (%)	14.25	13.50	0.02	18.79	17.24	0.04	11.69	12.08	-0.01
Firm									
Firm Age: share adult (%)	59.14	63.57	-0.09	64.20	68.97	-0.10	57.68	58.55	-0.02
Firm Size: share large (%)	69.73	49.80	0.42	61.21	36.66	0.51	67.25	53.06	0.29
Manufacturing (%, SBI)	36.92	14.93	0.52	52.95	26.83	0.55	37.51	17.38	0.46
Services (%, SBI)	51.29	45.55	0.11	37.17	45.69	-0.17	52.56	56.16	-0.07
Observations	378,409	350,071		89,632	75,400		39,385	22,420	
Total panel size	7,239,116			1,734,690			606,033		

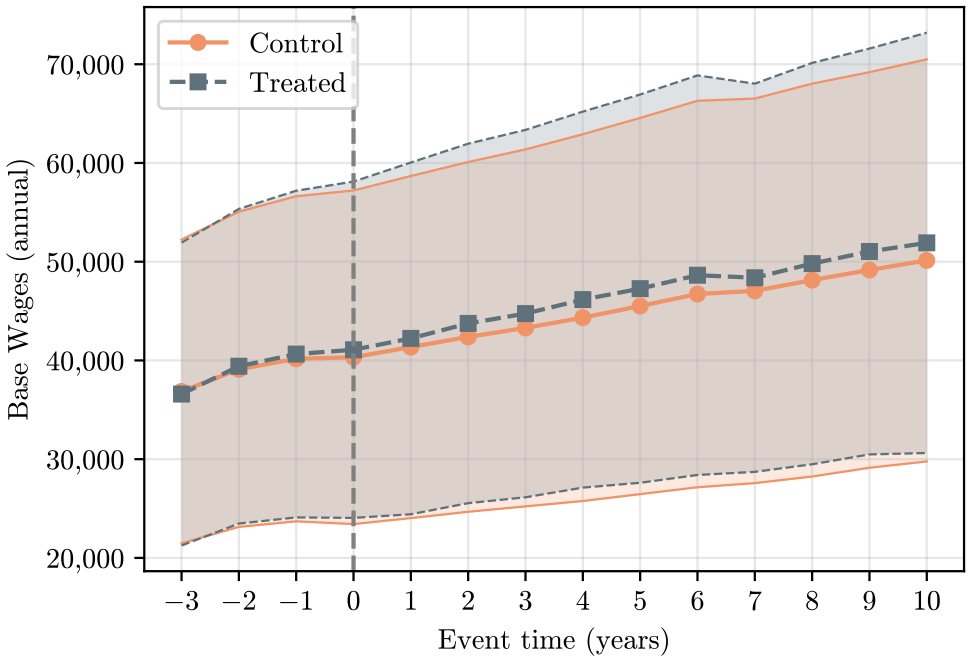
Notes: This table reports average characteristics of treated and matched control incumbent workers, measured three years before treatment (the matching year, i.e. $t-3$). Sample A reports the full matched sample, Sample B restricts to workers with below-upper-secondary education, and Sample C restricts to tertiary-educated workers in STEM fields of study. *Parental Migration Background* refers to workers with at least one foreign-born parent. *Firm Age* reports the share of workers at adult firms, defined as firms which are at least 10 years old, and *Firm Size* reports the share of workers at large firms (more than 250 employees). *Manufacturing* and *Services* refer to the firm's SBI industry classification. Columns (3), (6), and (9) report standardized mean differences (SMD) between treated and matched control workers. Total panel size is the full multi-year panel underlying each sample's event-study estimates.

Mean base wages differ substantially across the three samples by construction. Sample B workers earn less on average than Sample C workers, consistent with the non-R&D sample (B) including less-educated, lower-paid incumbents and the R&D sample (C) including tertiary-educated STEM workers. Within each sample, treated and matched control workers are closely aligned on base wages and hours worked, with SMDs of at most 0.09 in absolute value. Individual demographic characteristics are similarly well balanced across all three samples, with SMDs of at most 0.10. The remaining difference concerns firm size: treated incumbents more often work at large firms, which we address with the firm heterogeneity analysis in Section 5.1. Because our main outcome is the base wage, we are particularly interested in its evolution over time. Figure 2 plots raw base wage trends for Sample A and shows that wage trends are near-identical for treated and matched control groups before treatment and diverge only after treatment. This descriptive pattern supports our matching procedure and methodological choice.

4 Empirical Strategy

We estimate the potential rent-sharing effects of the policy using a staggered difference-in-differences (DiD) design with matched treatment and control panels. The matched control workers provide the

Figure 2: Raw Pattern of Base Wages (Treated and Matched Controls)



Notes: The figure displays average base wages for treated and control incumbent workers before and after the treatment year.

counterfactual to treated workers absent the policy benefit. We closely follow the implementation of standard staggered-adoption DiD estimators (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021), employing the interaction-weighted (IW) estimator proposed by Sun and Abraham (2021).[8] We choose this estimator because it allows for heterogeneous treatment effects across cohorts, which is empirically natural in our setting: firms that received the treatment in 2010 may differ from those entering in 2015 in ways that affect whether and how much of the surplus they share with their workers. The Sun and Abraham (2021) estimator also restricts the control group to never-treated workers, which sacrifices some efficiency but avoids contamination from treatment effects among already-treated cohorts. We discuss these choices in detail below.

The matching algorithm described in Section 3.3 and the IW estimator are complementary. Matching operates on the pre-treatment observable characteristics of workers and firms three years before the treatment year, restricting the sample to the most comparable treated and never-treated pairs. Conditional on this matched sample, the IW estimator then recovers cohort-specific causal effects using only the never-treated matches as controls. This two-step approach strengthens parallel trends through matching while avoiding contamination from inappropriate timing overlaps through the IW estimator's exclusion restriction.

[8]For a comprehensive overview of recent advances in difference-in-differences methods with staggered treatment adoption and heterogeneous treatment effects, see Roth et al. (2023).

4.1 Staggered Difference-in-Difference

Let Y_{it} denote the labour outcome of interest for worker i in year t, and let G_i be the year in which the firm employing worker i first benefits from the tax break. We define event time relative to the firm's first benefit year as

$$\ell \;=\; t - G_i,$$

where $\ell = 0$ marks the first year the firm receives the treatment, $\ell < 0$ denotes pre-treatment years, and $\ell > 0$ denotes post-treatment years. We restrict the event window to $\ell \in [-3, 10]$. The long post-treatment window is feasible because we observe complete worker employment spells up to 2023. The pre-treatment window is dictated by the incumbency definition. As noted before, we define a treated incumbent worker as an individual who has been continuously employed at the same firm for at least three consecutive years prior to G_i and who remains employed at that firm in year G_i. This restriction ensures that our treatment group consists of workers whose labour outcomes predate the policy benefit and who carry the accumulated process knowledge that gives the firm an incentive to reward them[9]. Because treatment is assigned at the firm level, incumbent workers do not directly select into treatment.

Workers are categorised into cohorts e based on the first year their employer benefits from the innovation box, such that all workers in cohort e share $G_i = e$. Following Sun and Abraham (2021), we estimate the following IW specification:

$$Y_{it} \;=\; \sum_{\ell \neq 0} \sum_{e < \infty} \delta_{e,\ell} \Big(\mathbf{1}\{G_i = e\} \cdot \mathbf{1}\{t - e = \ell\} \Big) \;+\; \alpha_i \;+\; \lambda_t \;+\; \varepsilon_{it}, \tag{1}$$

where $\delta_{e,\ell}$ is the cohort-specific average treatment effect on the treated at event time ℓ for cohort e, α_i are worker fixed effects absorbing all time-invariant worker characteristics, and λ_t are year fixed effects absorbing common aggregate shocks. All coefficients are normalised relative to $\ell = 0$, the first year of the firm's tax benefit.[10]

We form the aggregate IW estimator $\hat{\nu}_\ell$ by taking a cohort-share-weighted average of the cohort-specific estimates $\hat{\delta}_{e,\ell}$, as implemented in the `sunab()` function of the `fixest` package in R (Bergé, 2023). The IW estimator is consistent for this weighted average of $\delta_{e,\ell}$ under the assumptions discussed in Section 4.2 (see infra). A key advantage over the conventional two-way fixed effects (TWFE) estimator with leads and lags is that the IW estimator is free of contamination: the coefficient $\hat{\nu}_\ell$ at relative period ℓ is not polluted by treatment effects from other relative periods $\ell' \neq \ell$, a bias that arises under heterogeneous treatment effects in the TWFE specification (Sun and Abraham, 2021). The estimated $\hat{\nu}_\ell$ therefore captures the average labour effect for incumbent workers who remain employed at the treated firm ℓ periods after the firm first benefits from the innovation box.

[9] Because our panel starts in 2007 and the first cohort is treated in 2010, we are limited on defining incumbents at periods earlier than $t - 3$. However, we perform detailed robustness analyses on the definition of incumbents at $t - 1$ and $t - 2$ in Section 6.

[10] The conventional normalisation is $\ell = -1$. We normalise at $\ell = 0$ because the tax benefit accrues to the firm within the filing year, and the adjustment of wages and retention decisions in response to that benefit plausibly requires time: firms must first realise the tax saving, incorporate it into their financial planning, and then adjust compensation. In addition, it preserves all three pre-treatment periods for testing parallel trends. We rerun our analysis with normalization at $\ell = -1$ in Appendix D.4, with no substantial difference in results. We see that the standard normalization further justifies our choices, as the effects are not immediate and appear after period $\ell = 2$. Differences in magnitude are also not significant.

As noted above, the IW estimator uses only never-treated workers as controls for each cohort-specific estimate $\hat{\delta}_{e,\ell}$ (Sun and Abraham, 2021). Discarding not-yet-treated observations reduces efficiency, but eliminates contamination from units whose own future treatment effects might otherwise bleed into the control group. In our setting, this is a deliberate choice: we have a large pool of never-treated workers, and we prioritise internal validity over efficiency.[11]

4.2 Identifying Assumptions

Our object of interest is the effect of the innovation tax break on the wages of incumbent workers. Treatment is the first year in which the firm at which the incumbent workers is employed pays a preferential corporate income tax rate on its innovation profits. $\delta_{e,\ell}$ measures how the base wages of treated incumbents evolve relative to matched incumbents at firms that never receive this rate. Thus, the key difference between treated and control incumbents is the tax treatment of their employer's innovation profits. The incumbents are employed at the firm from $t-3$ onward, before the tax benefit was realised, and they do not choose their employer's tax status. The IW estimator requires three assumptions for a causal interpretation of the estimated $\delta_{e,\ell}$ as the effect of the tax break. We state each formally, discuss its plausibility in our setting, and describe the empirical tests we use to assess these assumptions.

First, baseline outcomes should follow parallel trends. Absent treatment, workers at treated firms would have experienced the same average evolution of wages over time as matched control workers. We address this through a flexible matching algorithm (see Appendix B.4) that constructs a control group of workers at never-treated firms who are observationally similar to treated workers on a rich set of worker and firm characteristics measured three years before treatment. Matching before estimating the event study substantially narrows the gap between treated and control units before any treatment effects arise, strengthening the credibility of parallel trends for the matched sample. Additionally, we restrict our estimation to large firms in a heterogeneity analysis (see Section 5.1). Large firms are not expanding as rapidly and follow different wage expansion paths. Thus, they are less likely to hire in anticipation of the policy, fall under similar bargaining agreements and be already innovative. As a final verification, we inspect the estimated $\hat{\nu}_\ell$ for $\ell \in [-3,-1]$, as well as cohort-by-cohort pre-trends (see Appendix D.2).

Second, there should be no anticipatory behaviour. Workers do not have information about their employer's corporate tax position. The innovation box benefit accrues to the firm through its annual tax filing, and individual employees are not informed about the firm's application or expected benefit. At the firm level, however, managers may adjust hiring or investment plans in the years before registration. Our analysis conditions on workers already employed before treatment, so firm-level anticipation in hiring does not confound our incumbent-worker estimates. We are limited by data availability in the pre-treatment period, as we only observe firms at most three periods before treatment. Because of this limitation, we distinguish between R&D workers and non-R&D workers, in addition to the full sample. There is no reason to believe that non-R&D workers will be hired in anticipation of the treatment.

[11] The trade-off between efficiency (using not-yet-treated units as controls, as in Callaway and Sant'Anna 2021) and internal validity (restricting to never-treated units, as in Sun and Abraham 2021) in discussed in Roth et al. (2023). Our setting favours internal validity. Because we are working with large worker panels, rather than with firms there exists a large sample pools of counterfactual never-treated workers that are closely matched to treated workers.

Lastly, there should be no interference between treated workers. The potential outcome of worker i must not depend on the treatment status of other workers. In our setting, this means that the treatment received by one firm does not affect the wage outcomes of workers at other firms. Because we study rent sharing (base wages effects), interference between workers at different firms is unlikely since the innovation box only increases the surplus available within a treated firm. One potential violation is general-equilibrium spillovers. If treated firms use their surplus to poach workers from untreated firms, the latter may raise wages to retain their employees, contaminating the control group. A second potential violation might be that treated firms might retain talented workers, reducing mobility out of treated firms and thereby limiting the outside options available to workers at control firms. We assess the relevance of both potential violations by examining descriptive mobility patterns across the treatment and control groups in Section 5.2. Note, however, that labour mobility in the Netherlands is comparatively low (OECD, 2025), which limits the scope for contamination.

5 Results

5.1 Rent Sharing

The innovation tax break leads to rent sharing among incumbent workers, as shown in Table 2. After their firm first benefits from the innovation box, incumbents' annual base wages rise by 3.27% relative to matched control incumbents, about €1,400 per worker per year. Rent sharing extends beyond workers who are likely to have produced the innovation. Base wages rise by 2.05% for non-R&D workers in Sample B, and by 4.14% for R&D workers in Sample C. Studies of innovation rents typically find gains concentrated among inventors and highly skilled workers (Kline et al., 2019; Aghion et al., 2018). Since we study the pass-through from the tax break, we find that the effects extend across the whole workforce. Because the innovation box still rewards successful innovation by increasing after-tax returns to innovation, part of the wage gains for R&D workers may still reflect returns to their own inventive efforts. Workers with less than upper-secondary education are unlikely to contribute directly to the innovation rewarded by the innovation box. We therefore interpret the Sample B estimate as a cleaner estimate of rent sharing induced by tax breaks.

Figure 3 shows the dynamics of base wage increases (rent sharing). The effect is small and insignificant in the first year after entry into treatment and grows to about 4% from $\ell = 4$ onward. A gradual rise in base wages is consistent with firms passing on a lasting increase in retained profits through successive wage negotiations (Arulampalam et al., 2012; Fuest et al., 2018; Dwenger et al., 2019), rather than paying out a one-time innovation windfall. Thus, we interpret this lagged effect as supporting evidence of rent sharing from innovation tax breaks.

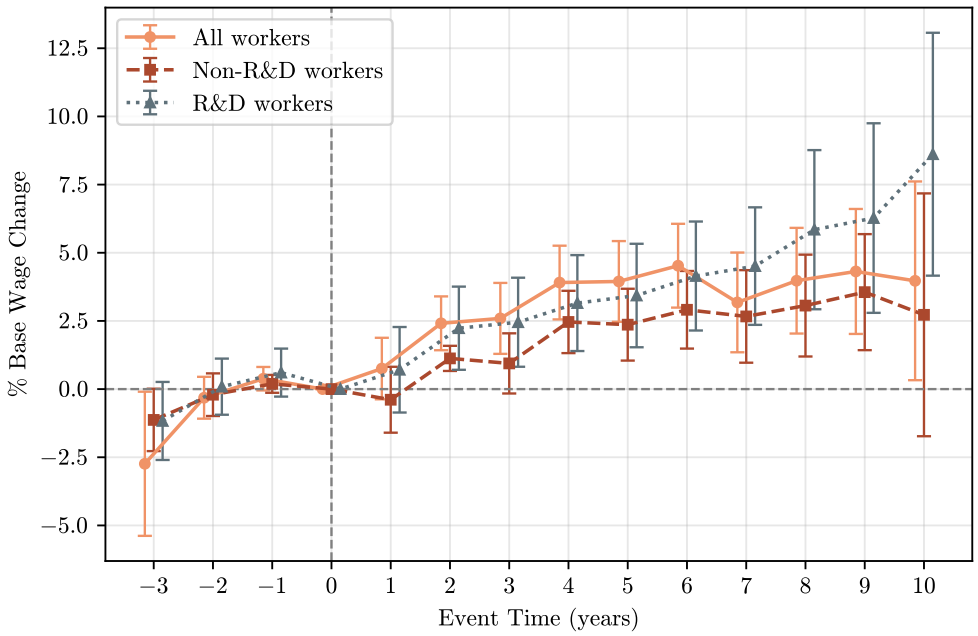
The response to innovation tax breaks runs through base wages rather than bonuses or hours worked. Bonuses do not respond to the innovation box treatment. One-off payments are more commonly used for passing on a one-time innovation benefit (Toivanen and Väänänen, 2012), whereas a permanent rise in base wages is consistent with a lasting increase in retained profits. Base hours rise by 0.58% in Sample A, about 11 hours per year, far too little to account for the 3.27% rise in base wages. Overtime income and overtime hours fall slightly, by less than 0.1% of the control mean in every sample. The base wage response is therefore not a reallocation of hours or overtime pay. Appendix C.1 reports all event-study coefficients, in percent and in euros.

Table 2: Innovation Box Tax Break Effects on Incumbent Workers' Pay and Hours

	Base Wage (%) (1)	Bonuses (%) (2)	Overtime Income (%) (3)	Base Hours (%) (4)	Overtime Hours (%) (5)
Sample A: Full Sample					
IW-ATT	3.275***	0.385	−0.057**	0.579***	−0.068**
	(0.671)	(0.530)	(0.029)	(0.178)	(0.033)
Observations	7,239,116	7,239,116	7,239,116	7,239,116	7,239,116
Sample B: non R&D workers					
IW-ATT	2.050***	0.356	−0.072*	0.297**	−0.045
	(0.633)	(0.396)	(0.043)	(0.135)	(0.044)
Observations	1,734,690	1,734,690	1,734,690	1,734,690	1,734,690
Sample C: R&D workers					
IW-ATT	4.143***	1.658*	−0.030	0.633	−0.070***
	(1.107)	(0.861)	(0.032)	(0.385)	(0.027)
Observations	606,033	606,033	606,033	606,033	606,033
Worker Fixed Effects	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes

Notes: Each panel is a separate set of regressions estimated on its own subsample. Coefficients are expressed as a percentage of that panel's own control-group mean: income outcomes (base wage, bonuses, overtime income) are a percentage of the control group's mean annual base wage, and hours outcomes (base hours worked, overtime hours) are a percentage of the control group's mean annual base hours worked, with both means computed separately within each panel's own subsample. Standard errors, clustered at the firm level, are in parentheses. Statistical significance is indicated as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 3: Event Study on Base Wages: Estimates of Rent Sharing



Notes: This figure shows the event-study coefficients estimating the effect of the innovation box policy on annual base wages, expressed as a percentage of each control group's mean annual base wage. The error bars denote 95% confidence intervals, based on standard errors clustered at the firm level. Event time is centered at $t = 0$, the first fiscal year that the firm receives the tax benefit. $\ell = 0$ is the omitted reference period. IW-ATT for each group is 3.295^{***} (Sample A - All workers), 2.050^{***} (Sample B - Non-R&D workers), 4.143^{***} (Sample C - R&D workers).

We next examine whether the increase in base wages is driven by particular firms or workers. Table 3 reports the results. The effect on base wages is positive and significant at the 5% level in every subgroup, which we discuss in detail below.

We first examine how the effect differs by firm heterogeneity. Two reasons motivate this. First, rent-sharing elasticities vary with firm size (Card et al., 2018), and small firms respond more strongly to tax incentives (Dechezleprêtre et al., 2023), so the pass-through from innovation tax breaks to incumbent wages might be largest at small firms. Second, matching leaves a residual imbalance in firm size (cfr. Table 1). Comparing treated and control incumbents within the same size class removes this imbalance. Large firms are the most informative group for investigating tax-break-induced rent sharing. Company-specific collective agreements, which leave room for firm-level pay responses, are most common at large employers (Jansen, 2021; International Labour Organization, 2022), and large never-treated firms are the controls most likely to hold profitable innovations since take-up of the box partly depends on idiosyncratic factors. Within large firms, the comparison therefore comes closest to holding innovation and the institutional setting constant while varying the tax regime.

The effects are largest at small firms (4.66%) and at small, young firms (5.35%), as predicted. But small, young, and medium-sized firms all show significant pre-treatment coefficients, so we cannot separate these larger effects from pre-existing wage growth. This is explained by smaller firms being on an expansion path, in which early employees, often owners, follow different wage-growth and bargaining paths than employees at established firms (see Appendix C.2.1). There are no expansion effects at large firms.

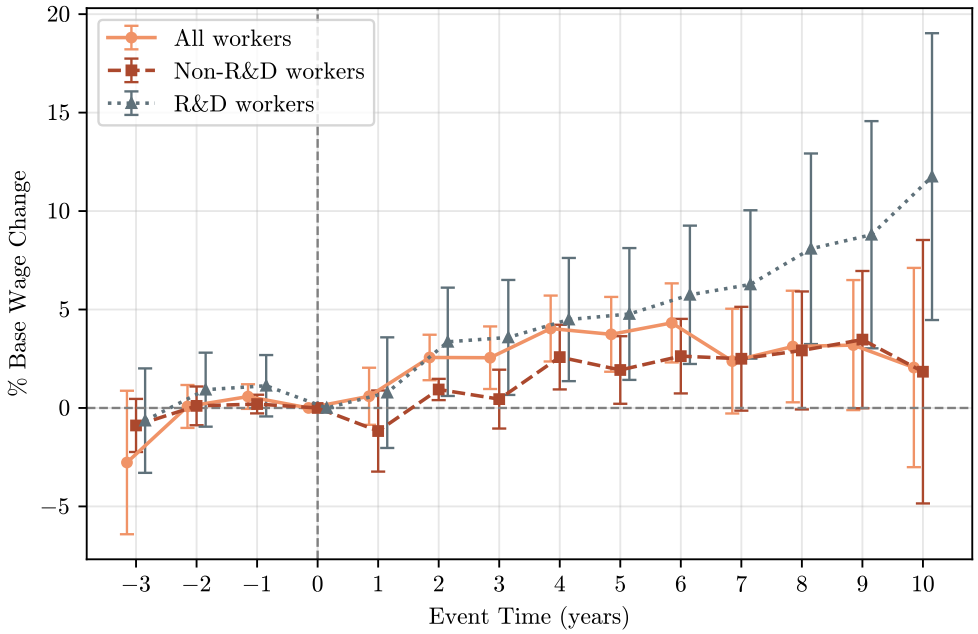
Table 3: Heterogeneity in the Effect on Base Wages: Estimates of Rent Sharing across Firm and Worker Characteristics

	IW-ATT (%) (1)	Observations (2)
Firm size and age		
Young (<10 years)	3.014**	2,292,608
	(1.222)	
Small and young	5.348***	387,480
	(0.434)	
Small (<50 employees)	4.659***	1,209,178
	(0.231)	
Medium (50–250 employees)	1.480***	1,597,262
	(0.220)	
Large (>250 employees)	2.794***	4,399,758
	(0.925)	
Large, excl. >5,000 employees	3.212***	2,993,151
	(0.841)	
Industry		
Manufacturing	2.939***	2,094,544
	(0.977)	
Services	2.440***	3,266,725
	(0.801)	
Worker characteristics		
Education		
Below upper secondary (Sample B)	2.050***	1,734,690
	(0.633)	
Upper secondary	2.557***	2,821,670
	(0.522)	
Tertiary	4.552***	2,603,895
	(1.116)	
Field of study		
STEM field	3.432***	1,417,599
	(0.892)	
Non-STEM field	2.821***	5,815,971
	(0.580)	
Gender		
Male	3.118***	5,469,847
	(0.629)	
Female	3.241***	1,769,247
	(0.778)	
Parental migration background		
Native-born parents	3.407***	6,217,460
	(0.670)	
Foreign-born parent(s)	2.586***	1,017,893
	(0.718)	
Worker fixed effects	Yes	
Year fixed effects	Yes	

Notes: Each row is a separate estimation of equation (1) on the indicated subgroup of Sample A. All characteristics are measured at $t-3$. Column (1) reports the IW-ATT on annual base wages as a percentage of the subgroup's control-group mean base wage. Below-upper-secondary workers form Sample B. Standard errors, clustered at the firm level, are in parentheses. Statistical significance is indicated as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 4 splits large firms into Samples B and C. Non-R&D workers in Sample B at large firms gain 1.76% and R&D workers in Sample C 5.80%. As in the full sample, the effect at large firms builds up over several years. This pattern is consistent with firms raising pay over successive wage rounds to retain workers whose firm-specific knowledge is valuable to them. We consider 1.76% as our most conservative estimate of rent sharing from innovation tax breaks.

Figure 4: Event Study on Base Wages: Estimates of Rent Sharing for Large Firms



Notes: This figure shows the event-study coefficients estimating the effect of the innovation box policy on annual base wages, expressed as a percentage of each control group's mean annual base wage for large firms only (> 250 employees at $t-3$). The error bars denote 95% confidence intervals, based on standard errors clustered at the firm level. Event time is centered at $t=0$, the first fiscal year that the firm receives tax benefit. $\ell=0$ is the omitted reference period. IW-ATT for each group is 2.794^{***} (Sample A - All workers in large firms), 1.765^{*} (Sample B - Non-R&D workers in large firms), 5.802^{***} (Sample C - R&D workers in large firms).

Rent sharing is not restricted to one industry. Splitting the sample into manufacturing and services, the two largest groups of beneficiaries, gives IW-ATTs of 2.94% and 2.44%, respectively. Because treated and control incumbents in each regression are employed in the same broad industry, this split also mitigates concerns that differences in sectoral wage agreements drive the estimated effect.

Finally, the gains reach every group of workers we can distinguish. Selection into occupations is shaped by education, field of study, gender, and parental migration background (Goldin, 2014), so we re-estimate the effect separately along each of these dimensions. The IW-ATT is positive in all ten subgroups, including Samples B and C, and each remains significant at the 5% level. All estimates lie between 2.05% and 4.55%, and within each dimension we cannot reject equal effects (see Appendix C.2.3). The tax break therefore raises pay across the incumbent workforce rather than for a group defined by education, field of study, gender, or parental migration background.

The broad-based wage response suggests that rent sharing induced by innovation tax breaks operates through a different mechanism than the sharing of innovation rents. Studies of innovation rents find gains concentrated among inventors and highly skilled workers, who are rewarded for

producing the innovation (Aghion et al., 2018; Kline et al., 2019). The tax break, by contrast, raises the after-tax surplus the firm retains, and evidence on corporate tax incidence shows that such surplus is passed on to workers through wage increases (Arulampalam et al., 2012; Fuest et al., 2018; Dwenger et al., 2019). This is because a tax break has no direct productivity component, allowing us to isolate rent sharing in a way that innovation or productivity shocks cannot. Firm-specific knowledge is not confined to R&D workers: incumbents throughout the organisation accumulate production, organisational, and process knowledge. When the tax break increases the after-tax surplus, firms share the increased surplus with these workers through successive wage negotiations to retain the workforce that sustains their innovation and profitability. This explains why base wages rise gradually and across the whole incumbent workforce, rather than for inventors alone.

5.2 Descriptive Worker Mobility

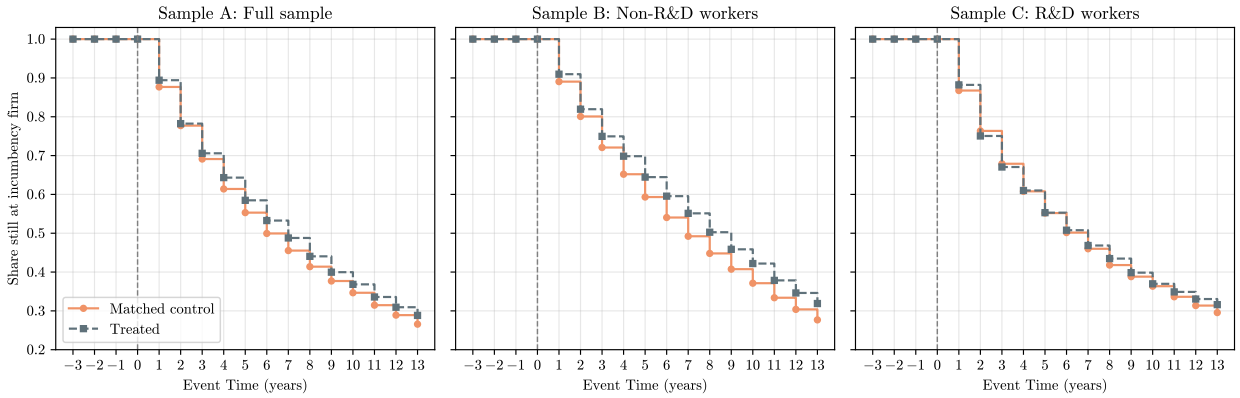
Existing studies show that firms share innovation rents with their workers, partly to retain the workers who produced the innovation (Kline et al., 2019). We are instead interested in the pass-through from the tax break to wages, and the two mechanisms have different implications for mobility. If our estimates captured sharing from innovation rents, we would expect retention effects concentrated among R&D workers. Tax-break-induced rent sharing instead works through wage renegotiations: the tax break raises the firm's after-tax profit without changing production or which workers are valuable, and incumbent workers bargain over this larger surplus (Fuest et al., 2018). In such case, we would expect wage gains without differential retention for R&D workers. We therefore complement our rent-sharing estimates in Section 5.1 with a descriptive analysis of worker mobility. We find no differential retention for R&D workers, consistent with rent sharing from innovation tax breaks rather than sharing from innovation rents.

The mobility analysis also serves as a check on the no-interference assumption discussed in Section 4.2. If treated firms used the tax break to poach workers from control firms or to retain workers who would otherwise have moved to control firms, flows between the two groups would contaminate the comparison. We find no evidence of such flows.

First, we check whether treated incumbents stay at their incumbency firm longer than matched controls. Figure 5 plots Kaplan-Meier survival curves. Five years after the first benefit year, 58.5% of treated and 55.3% of control incumbents in Sample A are still at the incumbency firm. The gap is largest for non-R&D workers in Sample B (64.5% versus 59.3% after five years), while the curves for R&D workers in Sample C are almost identical (55.3% versus 55.2%). These raw patterns already point to our mechanism. We find no evidence of retention effects among R&D workers, but mobility effects in the full sample might be driven by non-R&D workers.

We test whether the patterns observed in Figure 5 are statistically significant. Figure 6 reports event-study estimates for Switch_{it}, an indicator equal to one from the first year in which an incumbent is observed at a firm other than the incumbency firm. Following the job-loss literature (Britto et al., 2022; Dos Santos et al., 2026; Jarosch, 2023), we estimate a TWFE specification with worker and cohort-by-year fixed effects, in which control workers inherit the first benefit year of the treated worker they are matched to. Details can be found in Appendix C.3.2. We find no significant mobility response for R&D workers. In the full sample, incumbent workers at treated firms are 2.19 percentage points less likely to switch employers following treatment, a pattern driven by non-R&D workers, whose switching probability declines by 3.47 percentage points. This pattern

Figure 5: Survival at the Incumbency Firm



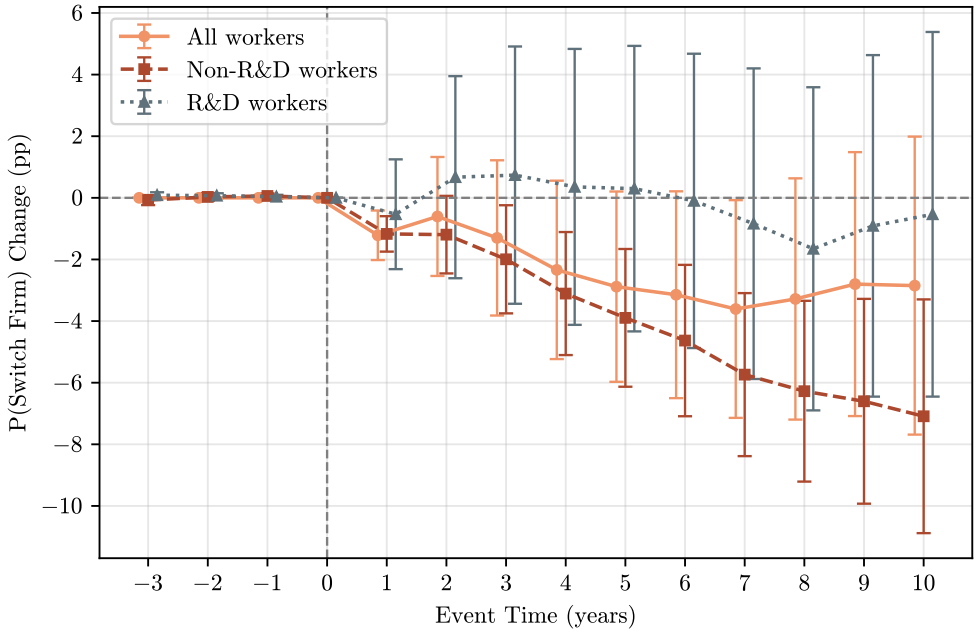
Notes: This figure presents Kaplan–Meier estimates of the share of incumbents who remain employed at their incumbency firm, by years since the firm's first benefit year ($\ell = 0$).

contrasts with what we would expect from the sharing of innovation rents, suggesting that our rent-sharing results stem from innovation tax breaks rather than innovative activity within the firm. Non-R&D workers hold firm-specific production and process knowledge, and they face more fluid labour markets (Jovanovic, 1979; Lamo et al., 2011). The firm shares the tax-induced surplus with its workers through firm-wide wage renegotiations. Because non-R&D workers' decisions to quit are more responsive to the wage premium offered by their current employer, the same degree of rent sharing translates into a stronger retention response. Since these workers are unlikely to have contributed directly to the innovation, their higher retention is unlikely to reflect rewards for inventive effort. Instead, these descriptive findings are consistent with a job-ladder interpretation, whereby internal wage gains make outside offers less attractive to non-R&D incumbents who have accumulated firm-specific knowledge.

As a last check, we analyse what happens to treated incumbents who leave (see Appendix C.3.3). We find that they realise larger wage gains at the move than control leavers. We measure the wage change as the log base wage in the first year at the new employer minus the log base wage in the last year at the incumbency firm. Treated switchers gain 0.8 log points more than control switchers in Sample A (all workers) and 1.1 log points more in Sample B (non-R&D workers). In Sample C (R&D workers), the difference is -0.2 log points and insignificant, in line with the absence of a retention response. We interpret the results for non-R&D workers as further evidence consistent with a job-ladder interpretation, whereby workers accept outside offers only if they exceed the value of their current job (Bagger et al., 2014). By raising incumbents' wages, rent sharing raises reservation wages. Because treatment changes who leaves, these are descriptive differences between selected groups, not causal effects. In Appendix C.3.3, we provide additional descriptive evidence on worker movements between treatment and control groups.

Overall, the mobility results support our rent-sharing findings and the underlying tax-break mechanism. R&D workers show the largest wage gain in our main estimates, yet their retention does not change, contrary to what we would expect from the sharing of innovation rents. R&D workers are valuable innovators in both groups of treated and control firms, so innovation-related retention incentives difference out. What remains is the additional surplus generated by the tax break, which R&D workers capture through higher wages but which does not change their incentive to stay. Non-R&D workers respond to the same rent sharing by staying longer, because their quit

Figure 6: Event Study on Probability of Switching Employers from Incumbency Firm



Notes: This figure shows the event-study coefficients estimating the effect of the innovation box policy on the probability of switching employers from the incumbency firm P(switch firm). The error bars denote 95% confidence intervals, based on standard errors clustered at the firm level. Event time is centered at $t = 0$, first fiscal year the firm receives tax benefit. $\ell = 0$ is the omitted reference period. ATT for each group is -2.192^{*} (Sample A - All workers), -3.472^{***} (Sample B - Non-R&D workers), -0.191 (Sample C - R&D workers).

decisions are more sensitive to wage gains. Their higher retention is consistent with tax-break-induced rent sharing and accounts for the modest retention response observed in the full sample.

6 Robustness

Two-Way Fixed Effects Estimation. Our main results use the interaction-weighted (IW) estimator of Sun and Abraham (2021) because it is robust to the bias that arises under staggered adoption with heterogeneous treatment effects across treatment cohorts or exposure length. We first verify that this choice matters in the direction the literature predicts by re-estimating the pooled ATT with a standard TWFE specification. Because treatment is staggered and effects grow over the post-treatment horizon, TWFE is biased under such heterogeneity (Sun and Abraham, 2021), typically inflating the estimated effect relative to a properly weighted dynamic average. Consistent with this, the TWFE estimate of the base wage ATT is just slightly larger than our preferred IW-ATT estimate (3.85% vs. 3.27%). However, there is no significant difference between the two estimators. We believe that these similar results follow from our staggered matching strategy. These results are presented in Appendix D.1.

Cohort-by-Cohort Analysis. Our main analysis pools treatment cohorts using the IW estimator. To assess whether the pooled effect is driven by a small number of cohorts, we re-estimate the event study for base wages separately by treatment cohort in each of the three estimation samples (see Appendix D.2). We find evidence of rent sharing through higher base wages across all treatment cohorts. For non-R&D workers in Sample B, the cohort-specific effect is positive in all nine cohorts, ranges from 1.0% to 2.9%, and is significant in seven of them. The 2010 cohort is particularly informative because it is the year of the policy introduction and also forms the largest cohort. For non-R&D workers in this cohort, wages show no pre-trend and rise by 1.82% after treatment, consistent with our pooled estimates. We also analyze pre-treatment effects per cohort. There are some pre-treatment effects concentrated in the 2012–2015 cohorts, which coincides with a shift in the composition of the treated pool toward smaller firms (cfr. Appendix A.3). As discussed in Section 5.1, small firms are on a different wage-growth path. Restricting the cohort-by-cohort analysis to large firms eliminates most of these pre-treatment effects: six of the nine cohorts of non-R&D workers at large firms show no significant pre-treatment effects, and the estimated effect remains positive in eight of the nine cohorts. Finally, the cohort-by-cohort mobility estimates confirm the pooled descriptive results discussed in Section 5.2.

Incumbent Definition. Our main analysis defines incumbents as workers employed at the firm from $t-3$ through t inclusive. This choice is dictated by data availability. Our payslip and firm data begin in 2006. Because tenure is used as a predictor in both the education imputation and the worker matching, 2007 is the earliest year available for our analysis. Given that the first treatment cohort enters in 2010, $t-3$ is the longest feasible incumbency window that preserves all treatment cohorts. We can, however, test the sensitivity of our results by redefining incumbents from $t-2$ or $t-1$ onward. Each relaxed definition yields a superset of the stricter one and therefore a considerably larger sample. This test also provides a check on the parallel-trends assumption, as a shorter pre-treatment tenure requirement includes workers hired closer to the firm's treatment year, whom our $t-3$ definition is designed to exclude. Relaxing the definition to $t-1$ introduces a pre-trend in the base-wage event study over two periods, and relaxing it to $t-2$ introduces a pre-trend confined to one period. Both relaxed definitions yield somewhat smaller effects than

our preferred definition: 2.66% under $t-1$ and 3.05% under $t-2$, compared with 3.27% under $t-3$. This is consistent with more recent hires having accumulated less firm-specific knowledge and therefore sharing less in the firm's increased after-tax surplus. Appendix D.3 reports the estimated average treatment effects and event-study figures underlying this comparison.

Normalization Timing. Our main specification uses the first benefit year, $\ell = 0$, as the reference period because the tax benefit accrues within the filing year and wages plausibly respond with a lag. We re-estimate the event study with the conventional reference period, $\ell = -1$. The results support our choice. The coefficient on the first benefit year is small and insignificant in all three samples (-0.38% in Sample A, -0.19% in Sample B, and -0.60% in Sample C), so wages do not respond in the year the firm first benefits. Because $\ell = 0$ now enters the post-treatment average with an effect close to zero, the pooled estimates are somewhat smaller but remain significant. Appendix D.4 presents the full set of event-study estimates.

Placebo Treatment. Matching does not guarantee that our never-treated control pool is free of workers who resemble treated incumbents in unobserved ways that also correlate with wage growth. To test for this, we restrict the sample to matched control workers only and randomly assign each a placebo treatment date, drawn from the same distribution of event times used in the main analysis. Because none of these workers is ever employed by a firm benefiting from the innovation box, the placebo treatment should yield no effect if our control pool is free from contamination. We find no significant ATT under this placebo assignment, confirming that the estimated rent-sharing effect in the main analysis is not an artifact of the control group itself.

6.1 No Education Imputation

Our main analysis imputes education and field of study for workers whose values are missing from the education survey, as described in Appendix B.2. This imputation is central to our design: education and field of study are the basis on which we construct the non-R&D and R&D samples (Samples B and C, respectively), our proxy for distinguishing workers unlikely to be involved in R&D from those most likely to be. If the imputation were correlated with treatment status, that is, if imputed values were systematically less accurate for workers at treated firms, our bounding strategy, and not just our sample size, would be compromised. We test this directly by dropping all workers whose education and field of study are not directly observed and re-running the matching and estimation procedure on this strictly observed subsample. The full-sample effect is essentially unchanged at 3.32%. The full set of results is reported in Appendix D.6.

7 Policy Implications and Conclusion

7.1 Implications for Innovation Policy

Innovation boxes are part of a broader set of innovation policies that are widely used across OECD countries (Einhoff and Paunov, 2025). Their effectiveness remains contested. Innovation boxes may induce profit-shifting responses (Haufler and Schindler, 2023) and existing evidence shows only modest increases in additional innovation spending (Agrawal et al., 2020; Cappelen et al., 2012;

Chen et al., 2021; Dechezleprêtre et al., 2023; Schwab, 2021). We contribute to this debate by studying a comparatively underexplored dimension of this ex-post output-based policy instrument (Bornemann et al., 2023; Chen and Hshieh, 2024): its impact on the workforce of beneficiary firms. Our findings demonstrate that the additional retained profits generated by innovation tax breaks are partially passed on to incumbent workers. However, these benefits to incumbent workforces should be assessed together with the effects on firms, the fiscal cost of the policy, and the extent to which the policy encourages socially desirable innovation.

The gradual rise in base wages is consistent with rent sharing through successive wage renegotiations. The accompanying decline in employer switching is concentrated among non-R&D incumbents and becomes visible only after a lag. These descriptive patterns are consistent with a job-ladder mechanism in which higher wages at the current employer make outside offers less attractive, particularly for workers with valuable firm-specific knowledge. The Dutch institutional setting, with comparatively high dismissal costs (OECD, 2020), may further shape firms' willingness to renegotiate wages. At the same time, the monetary gains accruing to workers appear modest relative to the associated tax expenditure. Previous policy analysis (Bijlsma et al., 2023) estimates the forgone tax revenue from the innovation box at €2,390 million in 2022. Our upper-bound estimate of nominal rent sharing amounts to approximately €530 million.[12] This implies that roughly 78% of the forgone tax revenue is not passed on to incumbent workers but remains within the firm or is allocated to other uses. Over the years, the tax advantage provided by the policy was reduced: the preferential tax rate first increased to 7% and later to 9%. When considering further changes to the tax regime, policymakers should therefore weigh the benefits accruing to firms and workers against the associated fiscal cost (such as forgone tax revenue).

Beyond the direct fiscal cost, the broader welfare implications of innovation policy should also be considered. Our analysis focuses on only one group of beneficiaries: incumbent workers. Although these workers are central to firms' production and innovation processes, their outcomes provide only a partial view of the overall effects of the policy. Our wage estimates are consistent with tax-break-induced rent sharing through periodic wage renegotiations, and our descriptive mobility results suggest an association with worker retention for non-R&D workers. However, our analysis does not allow to assess whether the innovation box raises aggregate innovation activity or improves overall welfare. A more comprehensive assessment would require examining firms' hiring responses, effects on workers' lifetime earnings, and the general equilibrium effects of rent sharing.

7.2 Conclusion

Innovation policy plays an important role in European competitiveness and citizens' welfare (European Commission, 2022). In this paper, we study the labour market effects of a large tax break on innovation profits in the Netherlands, focusing on tax-break-induced rent sharing through periodic wage renegotiations. Understanding the incidence of output-based innovation incentives is important for assessing how the additional retained profits generated by such policies are distributed and whether they generate real economic effects rather than merely affecting where profits are booked for tax purposes (Haufler and Schindler, 2023). While previous research finds only limited evidence of innovation tax breaks substantially increasing R&D spending or patenting (Chen et al.,

[12] This calculation multiplies an average increase in incumbents' annual base wages of approximately €1,400 per worker by the 378,409 treated workers in our main sample, with the estimated effect expressed in 2022 values.

2021; Dechezleprêtre et al., 2023; Schwab, 2021), their effects on workers remain comparatively underexplored (Bornemann et al., 2023; Chen and Hshieh, 2024). Motivated by firms' incentives to renegotiate wages for workers who have accumulated firm-specific production, organisational and process knowledge, we examine whether firms share the additional after-tax surplus with incumbent workers who were already present during the innovation process.

We find that firms share retained profits broadly across their incumbent workforce rather than exclusively with workers most directly involved in R&D. Depending on the likelihood of workers' participation in the innovation process, annual base wages increase by between 2.05% and 4.14% after firms begin benefiting from the tax break. Descriptively, these wage gains are accompanied by lower worker mobility: on average, incumbent workers at treated firms are 2.19 percentage points less likely to switch employers following treatment. This effect is driven by workers less likely to be directly involved in innovation, whose probability of switching employers decreases by 3.47 percentage points on average following treatment. By contrast, we find no significant mobility response for R&D workers. This pattern is consistent with a job-ladder interpretation in which internal wage gains may make outside offers less attractive for non-R&D incumbents with accumulated firm-specific knowledge.

Taken together, our findings show that the incidence of innovation tax incentives extends beyond firms and their shareholders. Part of the additional after-tax surplus is passed on to incumbent workers through wage gains. These gains are accompanied by a reduction in the likelihood that worker-firm matches are dissolved, and, hence, reduced worker mobility for non-R&D incumbents. Evaluations of innovation tax incentives should therefore consider not only their effects on innovation and tax revenues, but also the distribution of the resulting rents within firms and their potential implications for labour reallocation across firms.

References

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.

Aghion, P., Akcigit, U., Hyytinen, A., and Toivanen, O. (2018). On the returns to invention within firms: Evidence from Finland. *AEA Papers and Proceedings*, 108:208–212.

Aghion, P., Bergeaud, A., and Van Reenen, J. (2023). The impact of regulation on innovation. *American Economic Review*, 113(11):2894–2936.

Agrawal, A., Rosell, C., and Simcoe, T. (2020). Tax credits and small firm R&D spending. *American Economic Journal: Economic Policy*, 12(2):1–21.

Arulampalam, W., Devereux, M. P., and Maffini, G. (2012). The direct incidence of corporate income tax on wages. *European Economic Review*, 56(6):1038–1054.

Bagger, J., Fontaine, F., Postel-Vinay, F., and Robin, J.-M. (2014). Tenure, experience, human capital, and wages: A tractable equilibrium search model of wage dynamics. *American Economic Review*, 104(6):1551–1596.

Bergé, L. (2023). *fixest: Fast Fixed-Effects Estimations*. R package version 0.11.2.

Bessen, J., Goos, M., Salomons, A., and Van den Berge, W. (2025). What happens to workers at firms that automate? *Review of Economics and Statistics*, 107(1):125–141.

Bijlsma, M., Brouwer, E., de Jong, G., Busschots, T., den Hertog, P., Bongers, F., Smeitink, A., and Massop, M. (2023). Evaluatie van de innovatiebox 2010-2019: Doelgroepbereik, doeltreffendheid en doelmatigheid. Report 2023-125, SEO Economisch Onderzoek and Dialogic, Amsterdam / Utrecht. Commissioned by the Dutch Ministry of Economic Affairs and Climate Policy and the Ministry of Finance.

Bornemann, T., Laplante, S. O. K., and Osswald, B. (2023). The effect of intellectual property boxes on innovative activity and tax benefits. *Journal of the American Taxation Association*, 45(2):7–33.

Boumans, S. (2022). Neoliberalisation of industrial relations: The ideational development of dutch employers' organisations between 1976 and 2019. *Economic and Industrial Democracy*, 43(4):1610–1631.

Britto, D. G., Pinotti, P., and Sampaio, B. (2022). The effect of job loss and unemployment insurance on crime in Brazil. *Econometrica*, 90(4):1393–1423.

Caliendo, M. and Kopeinig, S. (2008). Some practical guidance for the implementation of propensity score matching. *Journal of Economic Surveys*, 22(1):31–72.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Cappelen, Å., Raknerud, A., and Rybalka, M. (2012). The effects of R&D tax credits on patenting and innovations. *Research Policy*, 41(2):334–345.

Card, D., Cardoso, A. R., Heining, J., and Kline, P. (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, 36(S1):S13–S70.

Chen, J. and Hshieh, S. (2024). The labor effects of r&d tax incentives: Evidence from VC-backed startups. *Review of Finance*, 28(5):1451–1482.

Chen, Z., Liu, Z., Suárez Serrato, J. C., and Xu, D. Y. (2021). Notching R&D investment with corporate income tax cuts in China. *American Economic Review*, 111(7):2065–2100.

Davidson, C., Heyman, F., Matusz, S., Sjöholm, F., and Zhu, S. C. (2014). Globalization and imperfect labor market sorting. *Journal of International Economics*, 94(2):177–194.

Dechezleprêtre, A., Einiö, E., Martin, R., Nguyen, K.-T., and Van Reenen, J. (2023). Do tax incentives increase firm innovation? An RD design for R&D, patents, and spillovers. *American Economic Journal: Economic Policy*, 15(4):486–521.

Dehejia, R. H. and Wahba, S. (2002). Propensity score-matching methods for nonexperimental causal studies. *Review of Economics and statistics*, 84(1):151–161.

Dobbelaere, S., McCormack, G., Prinz, D., and Sóvágó, S. (2025). Firm consolidation and labor market outcomes. *Journal of Economic Behavior & Organization*, 235:107036.

Dos Santos, A. G. F. V., Marie, O., and Markiewicz, A. (2026). The unequal job security scars of displacement. *Journal of Public Economics*, 255:105562.

Dwenger, N., Rattenhuber, P., and Steiner, V. (2019). Sharing the burden? Empirical evidence on corporate tax incidence. *German Economic Review*, 20(4):e107–e140.

Einhoff, J. and Paunov, C. (2025). Innovation policy transformed? Unveiling a new paradigm through natural language processing. Technical Report 2025/21, OECD Publishing, Paris.

European Commission (2022). A new European innovation agenda. Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions.

Fuest, C., Peichl, A., and Siegloch, S. (2018). Do higher corporate taxes reduce wages? micro evidence from Germany. *American Economic Review*, 108(2):393–418.

Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review*, 104(4):1091–1119.

Hartog, J. and Salverda, W. (2018). The labor market in the Netherlands, 2001-2016. *IZA World of Labor*, (418).

Haufler, A. and Schindler, D. (2023). Attracting profit shifting or fostering innovation? On patent boxes and R&D subsidies. *European Economic Review*, 155:104446.

Ibsen, C. L. and Keune, M. (2018). Organised decentralisation of collective bargaining: Case studies of Germany, Netherlands and Denmark. OECD Social, Employment and Migration Working Papers 217, OECD Publishing, Paris.

International Labour Organization (2022). Statistics on industrial relations. Data retrieved from ILOSTAT.

Jansen, N. (2021). Collective bargaining in the Netherlands. COLBAR-EUROPE report 4, COLBAR-EUROPE, WageIndicator Foundation and University of Amsterdam, Amsterdam.

Jarosch, G. (2023). Searching for job security and the consequences of job loss. *Econometrica*, 91(3):903–942.

Jovanovic, B. (1979). Job matching and the theory of turnover. *Journal of Political Economy*, 87(5, Part 1):972–990.

Kline, P., Petkova, N., Williams, H., and Zidar, O. (2019). Who profits from patents? Rent-sharing at innovative firms. *Quarterly Journal of Economics*, 134(3):1343–1404.

Lamo, A., Messina, J., and Wasmer, E. (2011). Are specific skills an obstacle to labor market adjustment? *Labour Economics*, 18(2):240–256.

OECD (2020). Recent trends in employment protection legislation. Chapter 3, OECD employment outlook 2020: Worker security and the COVID-19 crisis, OECD Publishing, Paris.

OECD (2021). Share of temporary employment. data retrieved from oecd statline, https://stats.oecd.org/Index.aspx?DataSetCode=TEMP_I.

OECD (2025). Harmful tax practices - Consolidated peer review results: Update (as of January 2025). Technical report, OECD/G20 Inclusive Framework on BEPS, Action 5. Results approved by the Inclusive Framework on 31 January 2025.

OECD (2025). Survey of adult skills 2023 technical report. Technical report, OECD Publishing.

Rosenbaum, P. R. and Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55.

Roth, J., Sant'Anna, P. H., Bilinski, A., and Poe, J. (2023). What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics*, 235(2):2218–2244.

Schwab, T. (2021). Thinking outside the box: The cross-border effect of tax cuts on R&D. *Journal of Public Economics*, 204:104536.

Statistics Netherlands (2024a). Gbapersoontab - persoonskenmerken uit de basisregistratie personen. The Hague. Microdata accessed via remote access.

Statistics Netherlands (2024b). Hoogsteopltab - hoogst behaald en hoogst gevolgd opleidingsniveau en -richting van personen. The Hague. Microdata accessed via remote access.

Statistics Netherlands (2024c). Spolis - banen van werknemers, jaarbestand. The Hague. Microdata accessed via remote access.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

Toivanen, O. and Väänänen, L. (2012). Returns to inventors. *Review of Economics and Statistics*, 94(4):1173–1190.

UNESCO Institute for Statistics (2015). International standard classification of education: Fields of education and training 2013 (isced-f 2013) - detailed field descriptions. Technical Report UIS/2015/INS/6, UNESCO Institute for Statistics, Montreal, Quebec.

Xu, X. (2025). On the move: Local opportunities and youth migration. *Fiscal Studies*, 46(4):451–473.

Supplementary Appendix for:
Rent Sharing under Innovation Tax Breaks

Sabien Dobbelaere* Ana Popovici†

October 2026

*Vrije Universiteit Amsterdam, Tinbergen Institute, Luxembourg Institute of Socio-Economic Research (LISER). sabien.dobbelaere@vu.nl.

†Vrije Universiteit Amsterdam and Tinbergen Institute. a.popovici@vu.nl.

A. Policy Appendix

A.1 Policy History and Operational Timeline

We describe the full timeline of innovation policies in the Netherlands. In 2007, the Netherlands introduced a Patent Box regime. Its limited scope and complex structure hindered widespread adoption among firms. The policy focused on profits derived from registered intellectual property (IP), rather than broader intangible assets. This limited uptake, combined with the Dutch government's ambition to position the country at the forefront of innovation, led to a policy revision in 2010 (Mohnen et al., 2017). The revised policy, the innovation box, offered greater clarity and predictability, and adoption increased substantially, with firms of all sizes participating. Table A.1 summarises how the policy changed in 2010. Additionally, it provides a detailed look into how firms can register their S&O[1] certificates to benefit from the policy.

Firms aiming to benefit from the innovation box had to re-register their intangible assets. The scope of the policy increased substantially, as there was no longer a cap on the benefit a firm could receive. Our data record the first year in which a firm benefits from the innovation box, not earlier use of the Patent Box, so a former Patent Box user that re-registered in 2010 enters our sample as treated in 2010.

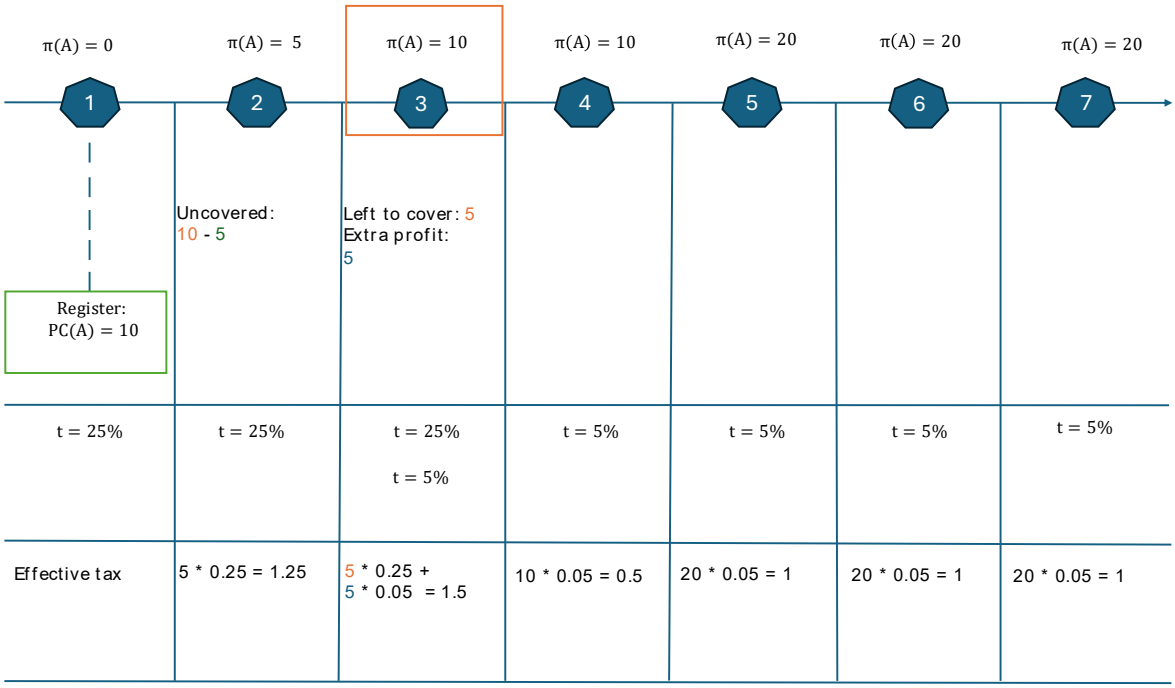
The policy changed in three further ways during our observation window (Bijlsma et al., 2023). First, the preferential rate rose from 5% to 7% on January 1, 2018 and to 9% on January 1, 2021. The 2018 cohort therefore entered at the 7% rate, and for firms treated in 2010-2017 the tax saving per euro of qualifying profit fell from 20 cents to 18 cents in 2018 and to 16 cents in 2021. Second, firms can use a simplified flat-rate method, designed for small innovators and described below, from 2013 onward. Third, the regime was aligned with the OECD modified nexus approach in 2017. None of these changes altered our underlying mechanism whereby higher retained profits might incentivize firms to share the additional after-tax surplus with their workers through successive wage negotiations. Once registered production costs are recovered, profits from the qualifying asset are taxed at the preferential rate.

We outline the policy's operational timeline in practice. Consider a transport firm that, at time $t = 1$, registers production costs of $PC(A) = 10$ with the Dutch Enterprise Agency (RVO) for developing a route-optimisation algorithm A. No innovation-related profits are generated at $t = 1$. At $t = 2$, the innovation yields profits of $\pi(A) = 5$. Since $\pi(A) < PC(A) = 10$ at this point, the entire profit is taxed at the statutory rate of 25%. At $t = 3$, the innovation generates profits of $\pi(A) = 10$. Because 5 units of the production cost remained unrecovered from $t = 2$, the profit realised at $t = 3$ is split. Five units are taxed at the statutory 25% rate, and the remaining five units at the reduced innovation box rate of 5%. At $t = 4$, all production costs have been fully recovered. From that point onward, all innovation-related profits qualify for the reduced 5% tax rate. In our data, this firm is first treated at $t = 3$, the first year in which part of its profit is taxed at the preferential rate. Figure A.1 illustrates this stepwise mechanism.

Identifying and attributing the exact profits derived from registered production costs is non-trivial. Standardised accounting rules determine the portion of profits attributable to innovation,

[1] An S&O ("speur- en ontwikkelingswerk") certificate is a subset of R&D. It is a stricter definition for what qualifies as innovation in the context of the policy. It is used specifically for calculating eligibility for the policy.

Figure A.1: Innovation Box Policy Diagram



Notes: The figure illustrates how the effective tax burden evolves as innovation-related profits ($\pi(A)$) rise. Initially, firms register qualifying development costs ($PC(A) = 10$) with the Dutch Enterprise Agency (RVO). Profits up to this threshold are taxed at the statutory corporate tax rate of 25%. Once these costs are fully recovered, any excess profit qualifies for the reduced innovation box rate of 5%. For instance, when $\pi(A) = 5$, the entire amount is taxed at 25%. At $\pi(A) = 10$, half is taxed at 25% and half at 5%. For subsequent profits, such as $\pi(A) = 20$, beginning in period 5, the full amount is taxed at the reduced 5% rate.

Table A.1: Patent Box versus Innovation Box

	Patent Box (2007–2009)	**Innovation Box (2010–present)**
Calculatory basis	Net income attributable to patented intangible assets (after deduction of cost base)	Net income from patented and/or S&O-certified assets
Effective tax rate	10%	5% (no ceiling), 7% from 2018, 9% from 2021
Eligibility	Only in-house patented assets	In-house patents and S&O-certified innovations
Cost treatment	Production costs must be activated and deducted at full rate	Production costs must be recovered before profits enter the box
Scope of benefits	Only royalties and other attributable revenues	Broader: royalties, sales income, licensing revenues, etc.
Transferability	Only for entities subject to corporate income tax, not transferable	Only for entities subject to corporate income tax, not transferable
Other IP rights	No (strictly patents)	Yes (including S&O without patent, under certain conditions)
Complexity	High, disagreements common with tax authorities	Improved clarity and predictability, ruling available
Ceiling / cap	Yes: max 4× production costs	No ceiling, full benefit above threshold

Notes: Differences between the Patent Box (2007–2009) and the innovation box (from 2010) in the Netherlands (Mohnen et al., 2017; Bijlsma et al., 2023).

as only this portion qualifies for the innovation box. Mohnen et al. (2017) describe several methods for this attribution. The per-asset method allocates profits to individual intangible assets. The peeling method distributes profits across core functions, such as S&O, marketing, production, and entrepreneurship, in cooperation with the tax authorities. The cost-plus method applies a standardised mark-up to S&O expenditures. From 2013 onward, firms whose asset results from S&O work covered by an S&O certificate in the current or the two preceding years can instead use a simplified flat-rate method, under which 25% of profit, up to €25,000 per year, is treated as qualifying profit. With a 25% statutory rate and a 5% preferential rate, the flat-rate method therefore saves at most €5,000 in tax per year. Firms that expect larger qualifying profits use one of the attribution methods above, for which no cap applies.

A.2 The Innovation Box and the WBSO

The innovation box rewards successful innovation. As such, it is an ex-post output-based policy instrument. However, the start-up phase of innovation is equally important. The ex-ante input-based incentive programme *Wet Bevordering Speur- en Ontwikkelingswerk* (WBSO) has been in place since 1994 and provides reductions in wage tax and employer contributions to the national insurance scheme for S&O personnel. By lowering the labour cost of S&O personnel, the WBSO may affect firms' hiring policies. Table A.2 compares the two instruments in detail.

The two instruments interact. Eligibility for the innovation box requires either a patent or an S&O certificate issued under the WBSO (Table A.2), so many beneficiaries also use the WBSO. Our control incumbents work at firms that never benefit from the innovation box, and some of these firms use the WBSO. Our estimates therefore compare innovation box beneficiaries with

non-beneficiaries, whatever their use of the WBSO, and do not measure the effect of the WBSO itself.

We focus on all workers at the firm, not just workers directly involved in innovation processes. While the interaction between hiring and innovation decisions is important, we do not study hiring. We focus on incumbent workers, who already work at the treated firms before treatment. Appendix C.2 reports results by firm size and age, the dimensions along which spillovers between hiring and pay decisions are most likely.

We compute the share of innovation box beneficiaries among all innovative firms in the Netherlands. Rather than relying on survey data (such as the R&D survey), we define innovative firms as those benefiting from the WBSO. There are 61,135 unique WBSO beneficiaries between 2004 and 2020, and 37,377 during our treatment period, 2010-2018. Almost 70% of innovation box beneficiaries also use the WBSO. Our 10,203 treated firms therefore make up only 16.7% of all innovative firms. Even if we take only WBSO beneficiaries during our treatment period, this share increases to about 27%. This leaves more than 70% of innovative firms from which matched incumbent workers can be sampled. Taken together with the probit estimates and our matching algorithm, this strengthens our identification.

Table A.2: Input versus Output Innovation Policy Instruments in the Netherlands

	Innovation box	**WBSO**
Policy lever	**Output-based** — reduces tax on *profits* from innovation	**Input-based** — reduces tax on *costs* of innovation
Timeline	2010–present	1994–present
Tax instrument	Reduced CIT rate on qualifying profit (5% until 2017, 7% in 2018-2020, 9% from 2021)	Wage tax credit on S&O labour costs, additional deduction for self-employed
Calculatory basis	Net income from patented and/or S&O-certified intangible assets	Qualifying S&O wage bill (+ flat supplement for other S&O costs)
Qualifying condition	Patent or S&O certificate *and* realised profit	S&O certificate (RVO), no profit requirement
Benefit timing	*Ex post* — only when innovation generates taxable income	*Ex ante* — immediate reduction of S&O labour costs
Eligible firms	CIT-paying entities, SME and large firm tracks	All firms with S&O employees, incl. start-ups & loss-making firms
Primary beneficiary	Profitable, commercialising firms	Firms in active S&O phase, regardless of profitability
Risk borne by firm	Low — benefit only materialises with success	High — subsidises input even if S&O fails
Additionality mechanism	Encourages profit retention and IP development in the Netherlands	Directly lowers marginal cost of hiring S&O labour

Notes: Differences between the WBSO and the innovation box in the Netherlands (Mohnen et al., 2017; Bijlsma et al., 2023). S&O stands for *speur- en ontwikkelingswerk*, and CIT for corporate income tax.

A.3 Treated Firms

The innovation box data come from internal tax records of the Dutch Enterprise Agency (RVO). They cover all firms that applied for the innovation box between 2010 and 2018 and record whether,

and in which year, each firm first benefited from the preferential rate. We observe 10,203 firms that benefit in this period. We define a firm's treatment year as its first benefit year and keep the firm treated from then on, whether or not it continues to report qualifying profit in later years. Our estimates are therefore effects of first entry into the innovation box. Firms that first benefit after 2018 are not in our data and can appear among the never-treated firms from which we draw control incumbents (see Appendix B.4).

Table A.3 describes treated firms by size, ownership, industry, and province. Most treated firms are small and domestically owned. A quarter of all treated firms operate in manufacturing and almost half in services. The largest shares are located in Zuid-Holland, Noord-Brabant, Noord-Holland, and Gelderland. Wages are higher at large and foreign-owned beneficiaries than at small and domestic ones.

Treated firms differ from the broader population of Dutch firms. We examine these differences by estimating how firm characteristics affect the probability of treatment. Let D_i be a binary indicator equal to one if firm i is treated and zero otherwise. We model the probability of treatment with a probit specification

$$\Pr(D_i = 1 \mid X_i) = \Phi\big(X_i'\gamma\big), \tag{A.1}$$

where X_i is a vector of firm characteristics measured three years before the first benefit year, and γ is the parameter vector to be estimated. Table A.4 reports the results. Column (1) compares treated firms with a random sample of never-treated firms observed in 2007–2015. Despite government efforts to promote the innovation box among small and medium-sized enterprises (SMEs), beneficiaries are systematically different. They are larger, younger and more often domestically owned. They also spend more on automation (our proxy for intangible capital), pay higher wages, and are more likely to operate in manufacturing.

To explore how selection evolves over time, we re-estimate the probit separately by treatment year, each time comparing that year's new beneficiaries with a random sample of never-treated firms (columns 2010 to 2018). Later cohorts look more like the average firm. The firm-size coefficient is positive and significant for the 2010, 2011, 2012, and 2014 cohorts, and small and insignificant for 2013 and from 2015 onward. The automation coefficient is positive and significant through 2016 and loses significance in 2017 and 2018. This trend may reflect two underlying dynamics. First, government outreach aimed at increasing SME participation may have been effective. Second, the number of new beneficiaries declines over time (see Figure 1b in the main text). These two forces may interact. As the policy matures and initial demand is saturated, outreach may have succeeded in attracting the average Dutch firm.

Taken together, the probit estimates confirm that innovation box beneficiaries form a non-random and evolving subset of Dutch firms. We address this selection by matching treated and control incumbents within each treatment cohort on worker and firm characteristics measured three years before treatment (see Appendix B.4).

Table A.3: Treated Firms: Summary Statistics by Firm Characteristics

	Mean Base Wage	SD Base Wage	Mean Employment	SD Employment	Firms (%)
Firm size					
Large	42,786.03	10,388.21	1,082.19	2,329.38	9.00
Medium	40,005.09	8,303.51	109.96	51.56	30.12
Small	38,505.55	9,861.12	19.03	12.52	60.88
Foreign ownership					
Domestic	38,295.35	9,032.03	120.33	766.05	77.05
Foreign	43,752.31	10,520.45	246.40	797.95	22.95
Country of headquarters					
Netherlands	40,321.38	9,535.13	188.77	1,010.49	75.48
Belgium	38,919.79	7,984.15	111.60	156.99	2.21
Canada	44,589.03	9,813.95	365.53	737.33	0.54
Switzerland	42,421.20	8,999.92	243.07	789.68	0.93
Germany	42,426.22	10,088.09	192.22	322.08	3.89
Finland	43,476.63	10,525.06	191.05	259.91	0.53
France	43,412.08	10,060.24	554.68	1,713.44	2.56
Ireland	42,751.34	9,181.47	156.20	168.15	0.58
Japan	44,590.86	11,559.63	305.76	1,151.62	1.02
Sweden	43,444.18	9,858.11	113.65	203.60	1.40
United Kingdom	44,149.62	10,980.74	221.03	552.06	3.63
United States	45,570.39	10,816.68	199.68	470.83	7.23
Industry					
Agriculture	33,579.31	8,612.05	44.33	92.25	2.22
Construction	40,233.86	8,904.83	223.70	513.92	4.75
Health	35,106.68	8,132.21	656.35	1,778.80	0.62
Manufacturing	37,670.17	7,696.43	129.52	406.35	32.29
Services	40,516.21	10,442.22	144.20	952.40	59.41
Water	37,444.95	7,271.25	169.95	424.37	0.71
Province					
Groningen	36,981.87	7,834.52	80.51	179.02	2.62
Friesland	36,236.64	7,494.78	96.04	570.70	2.60
Drenthe	36,200.69	8,281.19	83.66	224.95	2.08
Overijssel	37,480.01	8,456.95	93.43	212.15	8.95
Flevoland	38,323.43	9,108.57	164.32	1,119.99	2.07
Gelderland	37,740.50	8,432.37	99.53	369.81	13.10
Utrecht	41,195.38	10,063.61	281.69	1,645.94	10.01
Noord-Holland	40,923.63	10,475.51	189.94	1,063.71	16.31
Zuid-Holland	40,907.74	10,178.35	144.49	639.69	18.91
Zeeland	36,694.73	8,068.61	57.88	93.41	1.55
Noord-Brabant	39,192.44	9,363.39	131.03	465.31	17.15
Limburg	38,962.57	9,275.68	124.05	292.60	4.65

Notes: Mean and standard deviation of annual base wages (euros) and of employment, by firm characteristic, for the 10,203 firms that first benefit from the innovation box in 2010–2018. Firms (%) is the share of treated firms in each category.

Table A.4: Determinants of First Treatment of the Innovation Box: Probit Estimates

	Dependent variable: firm first benefits from the innovation box (covariates at $t-3$)									
	All	2010	2011	2012	2013	2014	2015	2016	2017	2018
Intercept	-4.765***	-7.961***	-6.155***	-7.111***	-5.988***	-6.131***	-5.649***	-4.533***	-4.605***	-5.287***
	(0.332)	(0.806)	(1.099)	(1.162)	(1.099)	(1.009)	(1.042)	(1.226)	(1.203)	(1.413)
Foreign	-0.229***	-0.320**	-0.078	-0.341**	-0.394**	-0.273*	-0.010	-0.397**	0.396**	-0.846***
	(0.050)	(0.125)	(0.153)	(0.167)	(0.161)	(0.153)	(0.149)	(0.178)	(0.166)	(0.273)
Firm age	-0.006***	-0.003	-0.008**	-0.005	-0.001	-0.010***	-0.003	-0.007*	0.004	-0.012**
	(0.001)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)	(0.005)
Firm size	0.129***	0.162***	0.148**	0.312***	0.040	0.174***	0.045	-0.002	-0.067	0.077
	(0.020)	(0.049)	(0.062)	(0.070)	(0.061)	(0.059)	(0.065)	(0.075)	(0.075)	(0.096)
Automation Costs (log)	0.145***	0.164***	0.248***	0.185***	0.151***	0.086**	0.149***	0.167***	0.056	0.007
	(0.014)	(0.034)	(0.046)	(0.048)	(0.042)	(0.042)	(0.046)	(0.051)	(0.052)	(0.059)
Mean Base Wages (log)	0.938***	1.967***	1.027***	1.433***	1.199***	1.441***	1.439***	1.030***	0.834***	1.178***
	(0.087)	(0.218)	(0.283)	(0.300)	(0.276)	(0.260)	(0.270)	(0.328)	(0.304)	(0.361)
Manufacturing	0.936***	0.835***	1.669***	0.755*	1.454***	0.737**	-0.144	0.469	1.373***	0.597
	(0.118)	(0.272)	(0.451)	(0.414)	(0.407)	(0.309)	(0.377)	(0.426)	(0.504)	(0.463)
Services	0.000	0.034	0.717	-0.193	0.522	-0.114	-0.786**	-0.319	0.604	-0.150
	(0.115)	(0.267)	(0.447)	(0.407)	(0.401)	(0.297)	(0.367)	(0.415)	(0.490)	(0.436)
Construction	-0.218*	0.249	0.575	-0.421	-0.101	-0.676**	-1.513***	-0.401	0.310	-0.221
	(0.126)	(0.302)	(0.483)	(0.439)	(0.439)	(0.339)	(0.421)	(0.439)	(0.521)	(0.475)
Log-likelihood	-3,629.548	-595.864	-388.273	-358.942	-394.694	-437.093	-354.061	-304.343	-277.738	-195.408
AIC	7,277.097	1,209.729	794.546	735.884	807.387	892.186	726.122	626.685	573.477	408.816
BIC	7,338.399	1,255.007	836.175	777.234	849.251	934.324	767.298	666.198	613.445	448.114
Observations	60,390	11,310	7,540	7,310	7,740	7,980	7,170	5,960	6,270	5,820

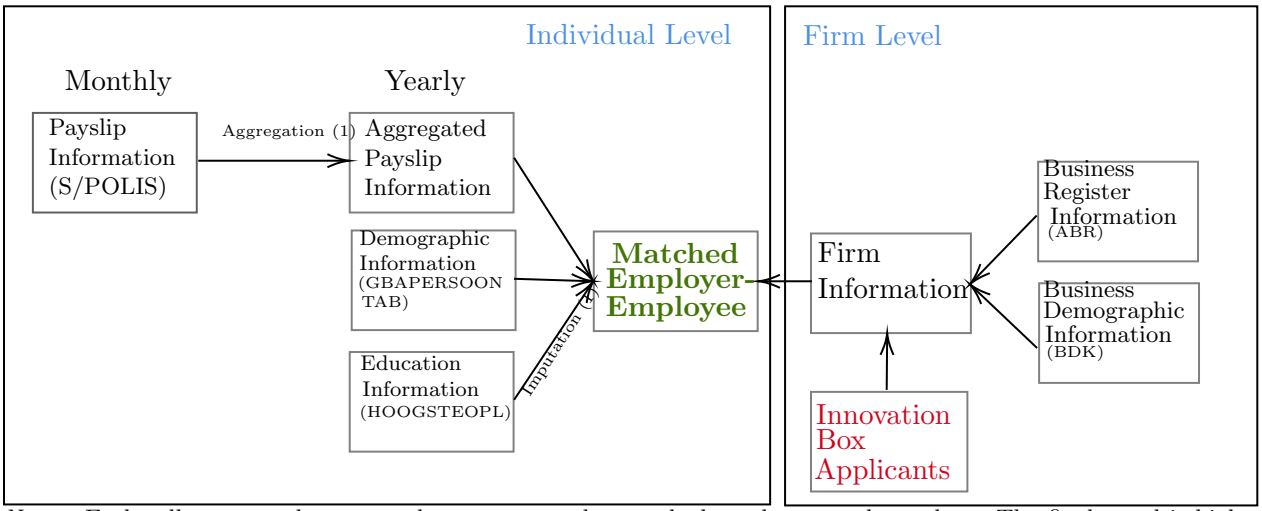
Notes: Probit coefficients from estimating equation (A.1). Column *All* pools all treated firms, with covariates measured three years before their first benefit year, and a random sample of never-treated firms observed in 2007–2015. Columns 2010 to 2018 restrict the treated firms to the indicated cohort and compare the cohort-specific treated sample with a random sample of never-treated firms. Automation costs and mean base wages are in logs, and manufacturing, services, and construction are industry dummies. Robust standard errors, clustered at the firm level, in parentheses. Statistical significance is indicated as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B. Data Appendix

B.1 Data Sources and the Matched Employer-Employee Panel

Figure B.1 presents the structure of the underlying datasets and how they are combined to construct the estimation samples. We draw on two types of records. Individual-level data are identified by encrypted personal identifiers, and firm-level data by encrypted business-unit identifiers. Depending on the source, raw individual-level data are available at monthly or annual frequency, and we aggregate monthly data to the annual level where necessary. From these matched employer-employee records, we construct our incumbent panels.

Figure B.1: Data Structure and Workflow



Notes: Each cell names a dataset used to construct the matched employer-employee data. The final panel is highlighted in green. Where applicable, the original dataset names used by Statistics Netherlands are provided in parentheses below each dataset. The innovation box data are highlighted in red.

Payslip information comes from monthly employee records (Statistics Netherlands, 2024c), available from 2006 to 2023. We apply the following restrictions in every year. A worker-year enters the panel if the worker works more than 520 hours at the firm and earns at least 90% of the annual minimum wage at the same firm. We drop internships and contracts under the Social Employment Act (*Wet sociale werkvoorziening*). When a worker has more than one spell in a year, we keep the spell with the most hours, following Abowd et al. (1999). We keep firms with more than two employees, with firm size taken from the Business Demography Framework (BDK), because very small units often correspond to sole proprietorships rather than to employer-employee relationships. Finally, we trim outliers at the 1% level of base wages and of hours worked, by year. Because the restrictions apply year by year, they also apply after treatment. An incumbent who falls below the hours or earnings threshold in a given year does not enter the wage panel in that year.

We take personal characteristics (age, gender, and parental migration background) from the GBAPERSOONTAB (Statistics Netherlands, 2024a), which covers the full population of the Netherlands for the entire span of our panel. We first construct a harmonised annual panel covering all individuals in the Netherlands using this source, which we then merge with the cleaned payslip data to build the matched employer–employee panel. We take each individual's highest obtained educa-

tion level from the HOOGSTEOPLTAB (Statistics Netherlands, 2024b). Appendix B.2 describes how complete missing values of education and personal characteristics.

We sum all pay components from the monthly records to annual values. The base wage is the annual contractual wage for regular hours, excluding overtime pay and special payments. Because it is an annual amount in euros, it moves with base hours as well as with the hourly rate. Appendix C.1 separates the two. Bonuses are the special remunerations (*bijzondere beloningen*) recorded separately from the base wage. These are payments made at irregular intervals, which include the annual holiday allowance and any thirteenth-month payment as well as bonuses and profit-sharing payments. Because nearly all Dutch employees receive a holiday allowance, most workers have positive bonuses (the 25^{th} percentile in Table B.1 is €2,591), and the variable should not be interpreted as discretionary performance pay alone. Overtime income and overtime hours are pay and hours for work beyond contractual hours, and base hours are the contractual hours paid at the base wage. Tenure is measured from 2006, the first year of the payslip data, so it is left-censored for workers already employed in 2006.

We obtain firm-level information from the Business Demography Framework (BDK) and Business Register Information (ABR). These datasets cover the universe of Dutch firms and provide detailed information on the country of headquarters, location within the Netherlands, firm size, and firm age. For a subset of firms, we supplement these data with information on automation costs from the Production Statistics Survey, which we only use in the aforementioned probit regressions (see Appendix A.3).

Table B.1 reports summary statistics for the matched employer-employee panel of 72,693,414 worker-year observations, covering 2006 to 2023, after imputing education and personal characteristics (see Appendix B.2). Of these, 70,303,071 observations fall in the 2007-2023 estimation window. Workers earn a mean annual base wage of 38.0 thousand euros and work 1,758 base hours per year on average. The mean age is 44, and 83% of workers have native-born parents. The sample is skewed towards male workers (59%), a consequence of our hours restriction. Bonuses average €5,911 but are very dispersed (standard deviation of €11,796). Overtime is rare, with median overtime pay and hours of zero. The education and field-of-study shares are post-imputation. Appendix B.2 describes the imputation procedure, and Appendix D.6 shows that our main estimate holds when we drop imputed observations.

B.2 Imputation of Education and Personal Characteristics

The education data are not available for all individuals in the Netherlands, and coverage is skewed towards more highly educated individuals in earlier years, improving in later years. We match education codes registered by the statistical agency to the Standaard Onderwijsindeling 2016 (SOI 2016)[2] and to the International Standard Classification of Education (ISCED) (Centraal Bureau voor de Statistiek, 2020) to retrieve each individual's education level and field of study.

Missing values also arise in other demographic variables, such as gender. Because our sample is the employed labour force, with a mean age of 44, these characteristics rarely change within a

[2]The SOI 2016 distinguishes three main education levels. These are low (e.g. early childhood, primary, and lower secondary education), medium (upper secondary education), and high (higher education, i.e.' bachelor's, master's, and doctoral levels). For additional information, see the full classification at Statistics Netherlands (CBS).

Table B.1: Matched Employer-Employee Data: Summary Statistics

	Mean	SD	P25	Median	P75	Observations
Wages, hours and worker characteristics						
Base wage (annual)	38,012	17,773	25,527	33,600	45,342	72,693,414
Bonuses (annual)	5,911	11,796	2,591	4,257	6,965	72,693,414
Overtime income (annual)	392.6	1,584.4	0.0	0.0	0.0	72,693,414
Base hours (annual)	1,758	346	1,557	1,872	2,052	72,693,414
Overtime hours (annual)	16.96	64.58	0.00	0.00	0.00	72,693,414
Age (years)	43.98	11.34	35.00	44.00	53.00	72,340,890
Tenure (days)	2,119	1,513	975	1,741	2,921	72,693,414
Worker composition	Share (%)					Observations
Education						
Below upper secondary	15.93					11,521,428
Upper secondary	43.60					31,540,721
Tertiary	40.47					29,278,741
Field of study						
STEM	14.82					10,722,498
Non-STEM	85.18					61,618,392
Parental Migration Background						
Native-born parents	83.33					60,280,926
Foreign-born parent(s)	16.67					12,059,964
Gender						
Male	59.26					42,867,476
Female	40.74					29,473,414

Notes: Statistics computed over the full matched employer-employee panel, 2006-2023 (72,693,414 worker-year observations), including the 2006 base year used to construct tenure. Base wage, bonuses, and overtime income are annual amounts in euros, and hours are annual hours. Tenure is measured in days from 2006. Education categories and fields of study follow the classification in Centraal Bureau voor de Statistiek (2020). The reported shares for these two variables are calculated after imputation.

worker's observed panel, so we impute missing values using each worker's first non-missing value, carried forward and backward to all years with missing data. Imputing missing education data requires a more involved procedure because of both the scale of missingness and its concentration among specific groups of workers. We implement a two-step procedure.

Step 1. Personal imputation. For each worker above age 25, we impute missing education levels using that worker's first observed education level, carried forward (and backward) to all years in which education is missing. Because education rarely changes once a worker has entered the labour force, this approach recovers a substantial share of missing values, with the largest gains in 2006–2010, when raw coverage was weakest. Table B.2 reports missing values before and after this step.

Table B.2: Missing Values Before and After Step 1 of the Imputation Procedure

	Missing before		Missing after		Recovered
	N	%	N	%	% of missing
Birth year	699,726	0.65	396,910	0.37	43.28
Gender	699,843	0.65	397,165	0.37	43.25
Parental migration background	865,343	0.81	820,798	0.77	5.15
Education code (raw)	42,799,231	39.90	33,694,653	31.41	21.27
Education level (SOI 2016)	42,866,234	39.97	33,711,370	31.43	21.36
Field of study (ISCED-F 2013)	43,088,618	40.17	33,731,051	31.45	21.72

Notes: Number and share of worker-year records with missing values before and after Step 1, which carries each worker's first observed value forward and backward. *Education code (raw)* is the education record as registered, and *Education level (SOI 2016)* is its mapping into the SOI 2016 levels. *Field of study* follows ISCED-F 2013. Records still missing after Step 1 are imputed in Step 2.

Step 2. Inverse-Mincer ordered probit. For workers with no observed education level in any year, we impute education using an ordered probit model of skill level on earnings, hours, and worker and firm characteristics, following the inverse-Mincer procedure of Dobbelaere et al. (2025). The education level is encoded as an ordinal outcome following standard OECD classifications,

$$\text{skill}_{it} = \begin{cases} 0 & \text{Below upper secondary} \\ 1 & \text{Upper secondary} \\ 2 & \text{Tertiary} \end{cases} \tag{B.1}$$

and modelled via the latent-variable representation:

$$\text{skill}^*_{it} = X'_{it}\beta + \varepsilon_{it}, \qquad \varepsilon_{it} \sim N(0,1), \tag{B.2}$$

$$\text{skill}_{it} = \begin{cases} 0 & \text{if skill}^*_{it} \leq c_1 \\ 1 & \text{if } c_1 < \text{skill}^*_{it} \leq c_2 \\ 2 & \text{if skill}^*_{it} > c_2, \end{cases} \tag{B.3}$$

where c_1 and c_2 are estimated threshold parameters. The vector X_{it} includes base wage, hours worked, tenure, age, gender, parental migration background, firm size category, firm age category, industry group (manufacturing, services), and year fixed effects.

Estimating the ordered probit on the full panel (72,693,414 worker-year observations) is computationally infeasible. We therefore estimate the model on a random subsample of 100,000 individuals, retaining all years per individual to preserve the panel structure, and apply the fitted model to predict education for all observations with missing values in the full panel. Table B.3 reports the estimates. The estimation sample contains 663,144 worker-year observations.

Because the model includes contemporaneous earnings, imputed education may in principle respond to the treatment. We avoid this by assigning workers to Samples B and C using only their education and field of study at $t-3$, three years before treatment (Appendix B.4). For workers imputed in Step 2, the prediction at $t-3$ uses only covariates measured at $t-3$. For workers imputed in Step 1, the carried value is an observed education record, which does not depend on wages. Post-treatment wage changes therefore cannot move workers between samples.

Base wage is positively associated with predicted skill level, while hours worked, tenure, and age are negatively associated, consistent with experience and age proxying for cohorts with lower formal qualifications in this setting. Female workers and workers at medium-sized and large firms are associated with higher predicted skill, while workers with native-born parents, workers in manufacturing and services, and workers at old firms are associated with lower predicted skill.

Because the observed education sample is skewed towards younger, more highly educated workers, we do not assign predicted classes using the model's own threshold cutoffs c_1, c_2 directly. Instead, following Dobbelaere et al. (2025), we map predicted class probabilities to discrete education categories using cumulative shares from the observed (non-imputed) education distribution. Let $\hat{p}_{0,it}$, $\hat{p}_{1,it}$, $\hat{p}_{2,it}$ denote the predicted probabilities of each education category for observation it, and let q_0 and $q_0 + q_1$ denote the cumulative shares of below-upper-secondary and below-or-upper-secondary education, respectively, among non-imputed observations. The predicted education category is then

assigned as:

$$\widehat{\text{skill}}_{it} = \begin{cases} 0 & \text{if } \hat{p}_{0,it} \geq q_0 \\ 1 & \text{if } \hat{p}_{0,it} < q_0 \text{ and } \hat{p}_{0,it} + \hat{p}_{1,it} \geq q_0 + q_1 \\ 2 & \text{otherwise.} \end{cases} \tag{B.4}$$

This step re-aligns the predicted class shares to match the observed (non-imputed) education distribution.

Table B.3: Ordered Probit Estimates: Education Level

	Education level	
	Coefficient	SE
Base wage (1,000 euros)	0.048***	(0.000)
Hours worked (100)	-0.115***	(0.001)
Tenure (years)	-0.022***	(0.000)
Age	-0.027***	(0.000)
Female	0.347***	(0.004)
Native-born parents	-0.117***	(0.002)
Medium firm	0.172***	(0.005)
Large firm	0.173***	(0.004)
Adult firm	0.036***	(0.003)
Old firm	-0.328***	(0.018)
Manufacturing	-0.321***	(0.006)
Services	-0.079***	(0.004)
Threshold 1	-3.715***	(0.016)
Threshold 2 (log increment)	0.909***	(0.002)
Year fixed effects	Yes	
Log-likelihood	-444,905.415	
AIC	889,872.830	
BIC	890,227.297	
Observations	663,144	

Notes: Dependent variable is the ordinal education category (0 = below upper secondary, 1 = upper secondary, 2 = tertiary). Base wage, hours worked, and tenure are rescaled to thousands of euros, hundreds of hours, and years, respectively, for readability. The estimation sample is drawn from a random subsample of 100,000 individuals, with all available years retained per individual. Statistical significance is indicated as follows: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

In the 2007-2023 panel (70,303,071 worker-year observations), 23,584,852 observations (33.6%) still have missing education after Step 1 and are imputed in Step 2. Table B.4 reports the post-imputation distribution of education level by year, which shows a gradual upskilling trend consistent with known patterns in Dutch educational attainment over this period.

Field of study imputation. Workers classified as below upper secondary (whether observed or imputed in the preceding step) do not have a meaningful field of study, since this education level does not entail a post-secondary discipline. These workers are assigned directly to non-STEM fields, without model-based imputation, as their fields of study fall outside science, technology, engineering, and mathematics (STEM).

Table B.4: Post-Imputation Distribution of Education Level by Year

Year	Below Upper Secondary	Upper Secondary	Tertiary
2006	0.221	0.453	0.327
2007	0.207	0.447	0.347
2008	0.204	0.449	0.347
2009	0.191	0.452	0.357
2010	0.185	0.454	0.361
2011	0.175	0.455	0.370
2012	0.165	0.455	0.379
2013	0.159	0.448	0.401
2014	0.146	0.446	0.408
2015	0.146	0.444	0.410
2016	0.143	0.443	0.414
2017	0.139	0.446	0.416
2018	0.139	0.445	0.416
2019	0.137	0.444	0.419
2020	0.132	0.443	0.425
2021	0.128	0.439	0.433
2022	0.132	0.434	0.434
2023	0.135	0.427	0.438

Notes: Shares of workers in each education category by year in the full panel after imputation. The year 2006 is included because it serves as the base year for constructing worker tenure (Appendix B.1) and as the omitted category in the year fixed effects of the imputation model. It is part of the panel but not of the 2007-2023 estimation window.

For workers with upper-secondary or tertiary education, we impute missing field of study using a binary logistic regression, since field of study is a binary (rather than ordinal) outcome at this level of aggregation:

$$\Pr(\text{STEM}_{it} = 1 \mid X_{it}) = \frac{\exp(X_{it}'\gamma)}{1 + \exp(X_{it}'\gamma)}, \tag{B.5}$$

where X_{it} includes the same continuous and categorical predictors as the education-level model, plus the (by then imputed) education level itself, encoded as a dummy distinguishing upper secondary from tertiary (the reference category). The model is estimated on the subsample of workers with observed field of study and applied to all upper-secondary and tertiary workers with missing field of study in the full panel. Training accuracy is 0.792, close to the sample's non-STEM base rate of 80-85% (see Table B.5), so the model adds little discriminative power beyond predicting the majority class. This does not drive our results. Table D.5 in Appendix D shows that dropping imputed observations entirely leaves our main estimate essentially unchanged.

As with education level, predicted STEM status is not assigned using a naive 0.5-probability threshold. Instead, we use the observed share of non-STEM workers among those with non-missing field of study, $q_{\text{Non-STEM}}$, as the classification threshold on the predicted probability of STEM:

$$\widehat{\text{STEM}}_{it} = \mathbb{1}\left[\hat{p}_{\text{STEM},it} \geq q_{\text{Non-STEM}}\right]. \tag{B.6}$$

Table B.5 reports the post-imputation distribution of field of study conditional on education level, confirming the deterministic non-STEM assignment for below-upper-secondary workers and yields a plausible split between STEM and non-STEM fields among higher-educated workers, with

approximately 20% of upper-secondary workers and 14.5% of tertiary workers classified as STEM. Appendix D.6 in Appendix D reports our results without the imputation. Dropping all observations for which we do not observe education leaves the main estimate essentially unchanged.

Table B.5: Post-Imputation Distribution of Field of Study by Education Level

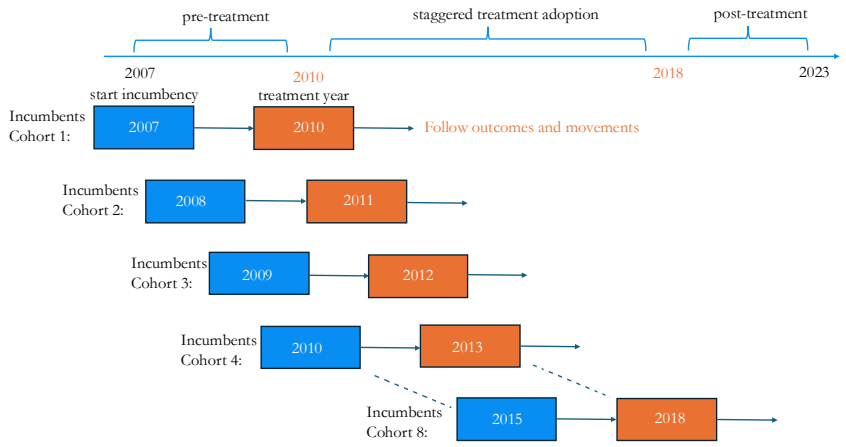
Education Level	Non-STEM	STEM
Below Upper Secondary	1.000	0.000
Upper Secondary	0.798	0.202
Tertiary	0.855	0.145

Notes: Shares of workers classified as non-STEM versus STEM, by education level, in the full panel after imputation. By construction, all below-upper-secondary workers are classified as non-STEM.

B.3 Incumbent Panels

Figure B.2 illustrates how we build incumbent cohorts. For each treatment year t between 2010 and 2018, incumbents are workers observed at the same firm in every year from $t-3$ through t. Treated incumbents work at firms first treated in t, and control incumbents work at never-treated firms. We follow both groups from t onward, up to 2023.

Figure B.2: Creation of Incumbency Cohorts



Notes: Construction of the incumbent cohorts. Each row is a treatment cohort. Incumbents are observed at the same firm from three years before the treatment year (blue) through the treatment year (orange), and their outcomes and moves are followed from the treatment year until 2023. The figure shows the cohorts treated in 2010 to 2013 and the last cohort, treated in 2018. The cohorts treated in 2014 to 2017 are built in the same way.

Treated incumbents are not a random draw from the pool of incumbents. Figure B.3 compares treated incumbents to a random draw of all incumbents, year by year, on the characteristics we later match on, and the first column of Table B.6 summarises the same comparison as standardised mean differences (SMD). Treated incumbents earn a higher base wage in every year (SMD 0.12) and receive higher bonuses in most years. They also work more hours (SMD 0.30) and more overtime

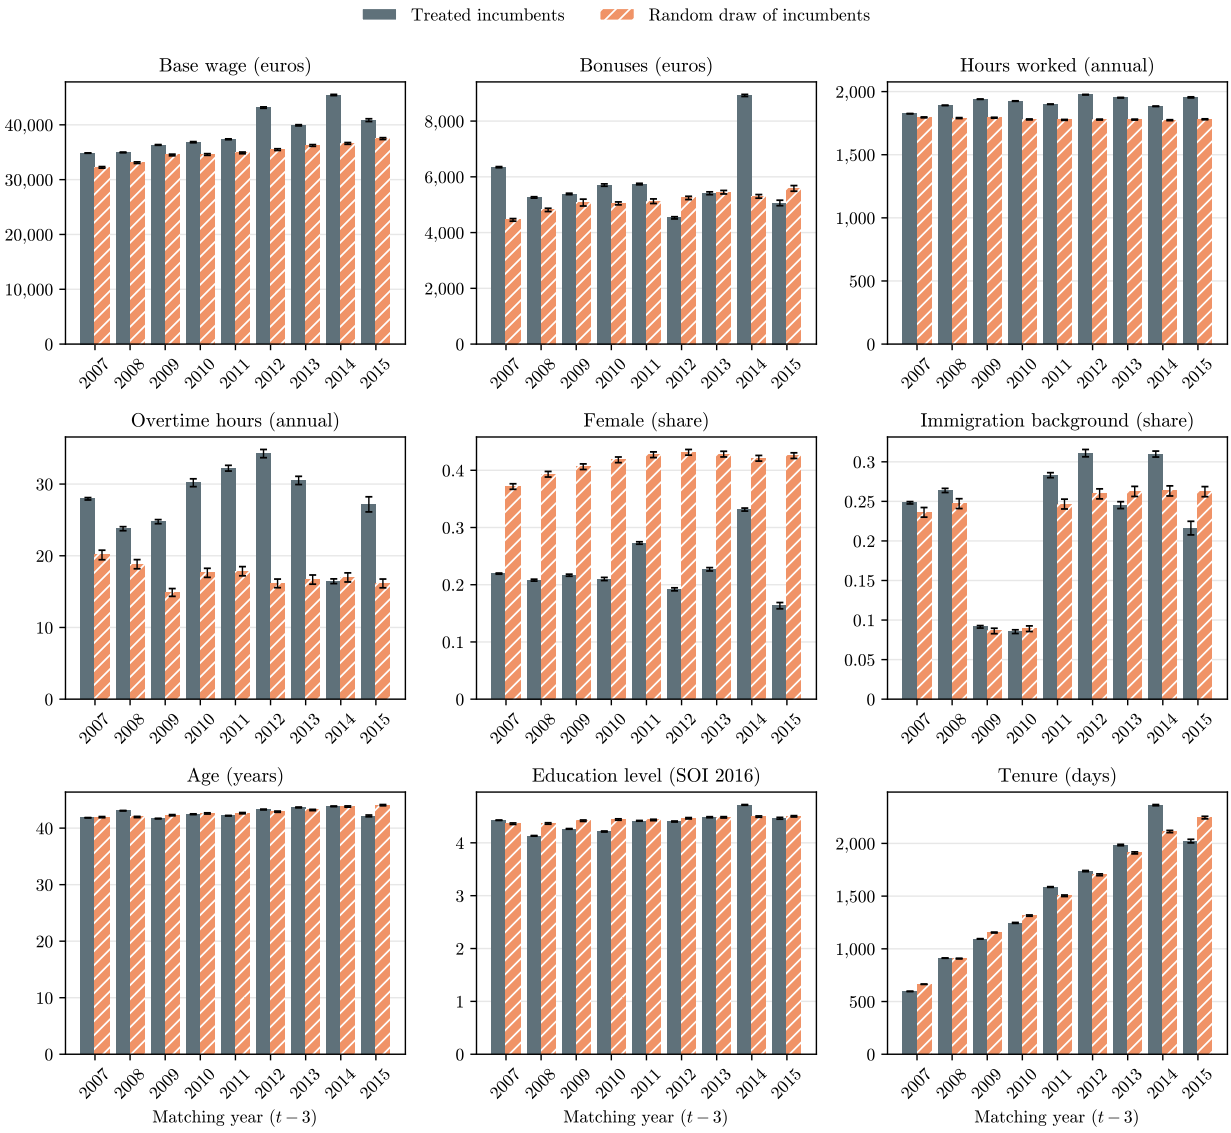
hours (SMD 0.15), and they are less often female in every year (SMD −0.41). Education level and age, by contrast, are already fairly close between the two groups before any matching is performed (SMD −0.12 and −0.04, respectively).

These gaps are large and persistent across the sample period, and go in the same direction on nearly every characteristic that plausibly correlates with wage growth. A random incumbent is therefore not a credible counterfactual for a treated incumbent. This is why we construct a matched control group, described in Appendix B.4, that is comparable to treated incumbents on these same characteristics.

B.4 Matching Algorithm and Estimation Samples

Appendix B.3 shows that treated incumbents earn more, work more overtime, and are more likely to be male than a random draw of incumbents. We address this by constructing a matched control group for every treatment cohort. Because treatment adoption is staggered between 2010 and 2018, we cannot match once on a single cross-section. Each cohort of treated incumbents must be matched on pre-treatment characteristics measured relative to its own treatment year, and a worker matched into one cohort must be removed from consideration in all later cohorts. We use nearest-neighbour matching with one neighbour per treated worker (1-NN), rather than propensity score matching, since it scales better to the large number of matching covariates we use (Section 3.3). Matching is performed separately within each treatment year, and control incumbents are always drawn from workers at never-treated firms. Within a cohort, a control worker can be the nearest neighbour of more than one treated worker, so matching is with replacement within a cohort. Across cohorts, no worker is used twice. Every matched worker, treated or control, is removed from the pool for all later cohorts. Algorithm B.1 states the procedure formally.

Figure B.3: Treated versus Random-Draw Incumbents, by Year



Notes: Each panel compares the mean characteristic of treated incumbents (dark grey) with that of a random draw of all incumbents (orange, hatched), by matching year $(t-3)$, 2007–2015. From left to right and top to bottom, the panels show base wage and bonuses in annual euros, hours worked and overtime hours in annual hours, the share of female workers, parental migration background, age, education level on the SOI 2016 scale, and tenure at the firm in days. The error bars denote ± 1 standard error of the mean.

Algorithm B.1 Incumbent Worker Matching Across Staggered Treatment Cohorts

Require: Firm treatment dictionary $\mathcal{F} : b \mapsto y_b^*$ (first treatment year); yearly cross-sectional worker panels $\{P_t\}_{t=2007}^{2018}$, each containing worker–firm pairs (w, b) observed in year t; set of all ever-treated firm IDs $\mathcal{B}^{\text{treat}}$; matching covariates $\mathbf{X}$; incumbency window $\delta = 3$

Ensure: Matched worker pairs $\mathcal{M}$, partitioned by cohort year

$\mathcal{M} \leftarrow \emptyset$

$\mathcal{E} \leftarrow \emptyset$ ▷ Global exclusion set — workers already matched on either side

for $y \in \{2010, 2011, \ldots, 2018\}$ **do**

— Step 1: Identify incumbent workers at the same firm across 4 years —

$\mathcal{I}_y \leftarrow \big\{(w, b) \;\big|\; (w, b) \in P_y \cap P_{y-1} \cap P_{y-2} \cap P_{y-3}\big\}$ ▷ Workers present at the same firm b in all of y, $y-1$, $y-2$, $y-3$

— Step 2: Treated incumbents —

$\mathcal{B}_y \leftarrow \{b \mid \mathcal{F}(b) = y\}$ ▷ Firms first treated in year y

$\mathcal{T}_y \leftarrow \big\{(w, b) \in \mathcal{I}_y \;\big|\; b \in \mathcal{B}_y,\ w \notin \mathcal{E}\big\}$

— Step 3: Control incumbents —

$\mathcal{C}_y \leftarrow \big\{(w, b) \in \mathcal{I}_y \;\big|\; b \notin \mathcal{B}^{\text{treat}},\ w \notin \mathcal{E}\big\}$

— Step 4: Match on pre-treatment characteristics —

Collect $\mathbf{X}_{y-\delta}$ for workers in $\mathcal{T}_y$ and $\mathcal{C}_y$, measured in panel $P_{y-\delta}$ ▷ $y - 3$ when $\delta = 3$

Run 1-NN matching (controls may be reused within cohort y) from $\mathcal{T}_y$ to $\mathcal{C}_y$ using $\mathbf{X}_{y-\delta}$, yielding pairs $\mathcal{M}_y$

— Step 5: Update exclusion set —

$\mathcal{E} \leftarrow \mathcal{E} \cup \big\{w \mid w \text{ appears in } \mathcal{M}_y\big\}$ ▷ Both treated and control workers excluded from future cohorts

$\mathcal{M} \leftarrow \mathcal{M} \cup \mathcal{M}_y$

end for

return $\mathcal{M}$

The algorithm proceeds cohort by cohort. For each treatment year t, we first identify the pool of incumbents, that is, workers observed at the same firm in every year from $t-3$ through t (Step 1). We split this pool into treated incumbents at firms first treated in year t, and control incumbents at any never-treated firm, excluding any worker already matched in an earlier cohort (Steps 2–3). We then run 1-NN matching from treated to control incumbents on characteristics measured three years before treatment, at $t-3$ (Step 4), and add every matched worker, treated and control, to a global exclusion set so that no worker is reused in a later cohort (Step 5). Repeating this across all treatment years from 2010 to 2018 yields the pooled set of matched pairs that make up Sample A. Because control pools are large relative to each treated cohort, reuse of controls within a cohort is limited. The 378,409 treated incumbents in Sample A are matched to 350,071 distinct control workers, so each control is used 1.08 times on average. Each distinct control worker enters the estimation panels once, without weights.

Samples B and C are not constructed through separate matching procedures. Instead, they are subsets of the matched Sample A, defined according to workers' education and field of study measured at $t-3$. Sample B requires below-upper-secondary education, and Sample C requires tertiary education in a STEM field.

Choice of matching covariates. Table B.6 reports the standardised mean differences that motivate our choice of matching covariates. We test five progressively richer sets of features.

Features (1) matches only on individual-level characteristics. Features (2) adds education level and field of study, Features (3) adds firm size and firm age, and Features (4) adds firm province and industry. Features (5), our preferred specification, further adds an AKM firm fixed effect (Abowd et al., 1999). We build up the covariate set this way because Appendix A.3 shows that innovation box firms differ systematically from the broader population of Dutch firms, so a specification that matches only on worker characteristics would leave firm-side selection unaddressed.

Table B.6: Covariate Balance: Treated vs. Random and Matched Control Groups

	Treated & Random Control	Treated & Matched Control				
		Features (1)	Features (2)	Features (3)	Features (4)	Features (5)
Base wage (t)	0.123	**-0.022**	**-0.022**	**-0.023**	**-0.023**	**-0.023**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Base wage ($t-1$)	0.155	**0.004**	**0.004**	**0.004**	**0.003**	**0.003**
	(0.000)	(0.158)	(0.158)	(0.148)	(0.164)	(0.165)
Base wage growth rate	**0.027**	**-0.005**	**-0.005**	**-0.005**	**-0.005**	**-0.005**
	(0.000)	(0.054)	(0.054)	(0.056)	(0.057)	(0.057)
Hours worked	0.296	**-0.072**	**-0.072**	**-0.072**	**-0.074**	**-0.074**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Overtime hours	0.151	**0.073**	**0.073**	**0.073**	**0.073**	**0.073**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Tenure	-0.506	**-0.063**	**-0.063**	**-0.063**	**-0.065**	**-0.065**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Gender	-0.408	**-0.096**	**-0.096**	**-0.097**	**-0.098**	**-0.098**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Parental migration background	**-0.015**	**-0.024**	**-0.024**	**-0.023**	**-0.019**	**-0.019**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Age	**-0.039**	**0.021**	**0.021**	**0.022**	**0.021**	**0.021**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Education level	-0.119	**-0.039**	**-0.039**	**-0.039**	**-0.038**	**-0.037**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Field of study	0.265	0.166	0.166	0.166	0.167	0.167
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Firm size	-0.423	-0.563	-0.563	-0.555	-0.552	-0.552
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Firm age	**-0.005**	**0.016**	**0.016**	**0.020**	**0.018**	**0.018**
	(0.179)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Firm province	**0.029**	**0.000**	**0.000**	**0.000**	**-0.010**	**-0.010**
	(0.000)	(0.959)	(0.959)	(0.903)	(0.000)	(0.000)
Firm industry	0.308	0.165	0.165	0.166	0.147	0.147
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Firm FE (AKM)	0.135	**0.028**	**0.028**	**0.028**	**0.025**	**0.026**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

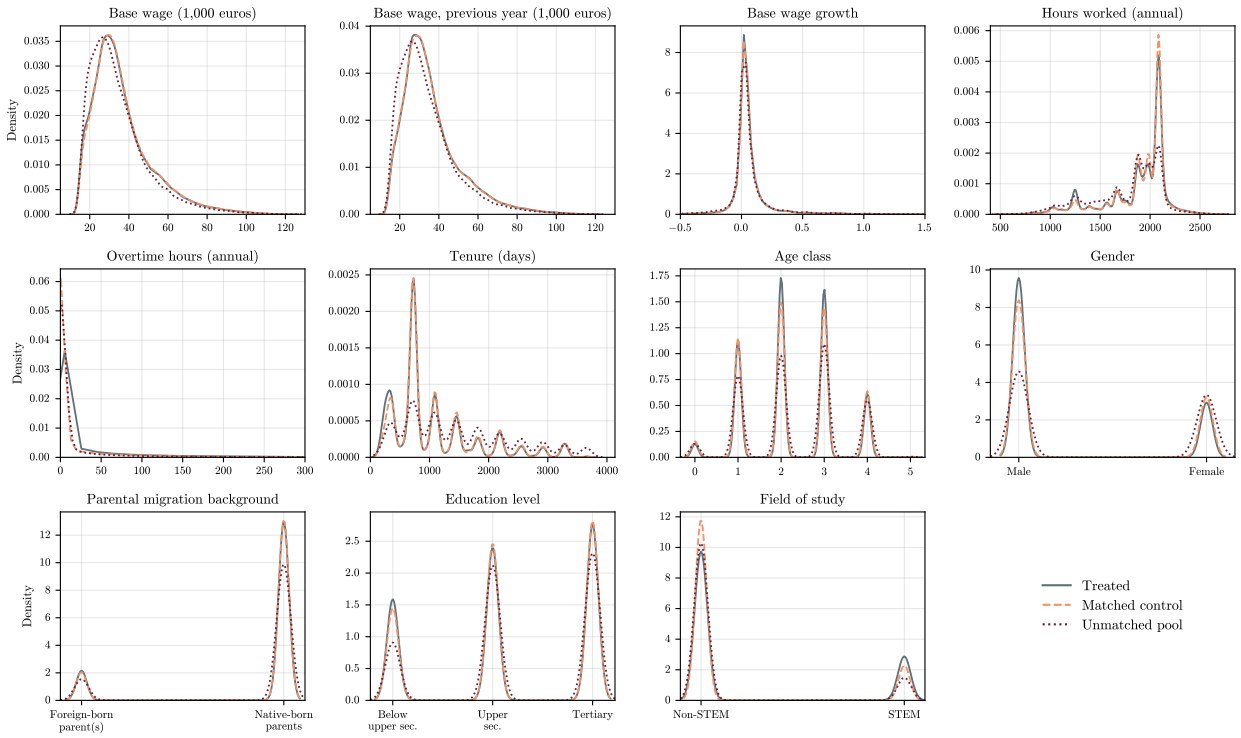
Notes: This table reports standardised mean differences (SMD) in characteristics between treated and control groups before and after matching. Column (1) compares treated workers to a random draw of all non-treated workers. Columns (2)–(6) compare treated workers to matched control workers under five progressively richer sets of matching covariates. Features (1) matches on individual-level characteristics (base wage, base wage in the previous year, base wage growth, hours worked, overtime hours, gender, parental migration background, age, tenure). Features (2) adds education level and field of study, Features (3) further adds firm size and firm age, Features (4) additionally includes firm province and industry, and Features (5) further adds AKM firm fixed effects. Matching is nearest-neighbour matching on the listed covariates. Controls may be reused within a cohort but not across cohorts. p-values (in parentheses) are from two-sample t-tests of equality of means between treated and control groups. **Bold** values indicate $|\text{SMD}| < 0.1$, the conventional threshold for covariate balance. Horizontal lines within the matched columns mark the boundary at which each additional set of features is introduced into the matching procedure. The "Treated & Random Control" SMD values are taken from the Features (1) specification. Minor variation in this column across specifications reflects changes in the random seed across matching runs.

Individual-level characteristics are well balanced even under the leanest specification. Base wage, base wage growth, hours worked, overtime hours, tenure, gender, parental migration background, and age fall below the $|\text{SMD}| < 0.1$ threshold already under Features (1), and adding richer covariates does not meaningfully change this. Education level is also balanced, but field of study is not (SMD of 0.167). Firm characteristics are harder to balance. Firm age, firm province, and the AKM firm fixed effect are balanced, but firm size and industry are not. The SMD for firm industry falls from 0.308 to 0.147, while the SMD for firm size remains large (-0.552 under Features (5)). This mirrors the trade-off described in Section 3.3 in the main text. Tightening the match on firm

size and industry would shrink the pool of usable control incumbents. We therefore accept some residual imbalance in firm-level characteristics to achieve the closest possible match on worker-level characteristics, as our research question focuses on workers rather than firms.

Kernel density functions. Figures B.4 and B.5 show balance across the full distribution of every matching characteristic, not just at the mean. Each panel plots the kernel density of a characteristic in the matching year $(t-3)$ for treated incumbents, matched control incumbents, and the full, unmatched pool of incumbents at never-treated firms. For worker characteristics (Figure B.4), the densities of treated and matched control incumbents lie almost on top of each other, while the unmatched pool differs visibly. Before matching, incumbents at never-treated firms earn less, are more often women, and have a more dispersed tenure distribution. After matching, these differences disappear. This holds for the level of base wages, for base wages in the previous year, and for base wage growth, so controls are selected on their recent wage path rather than on a single year's wage. A difference remains only for field of study. Treated incumbents more often hold a STEM degree than matched controls, in line with its SMD of 0.167 (see Table B.6).

Figure B.4: Kernel Density Functions for the Worker Matching Characteristics

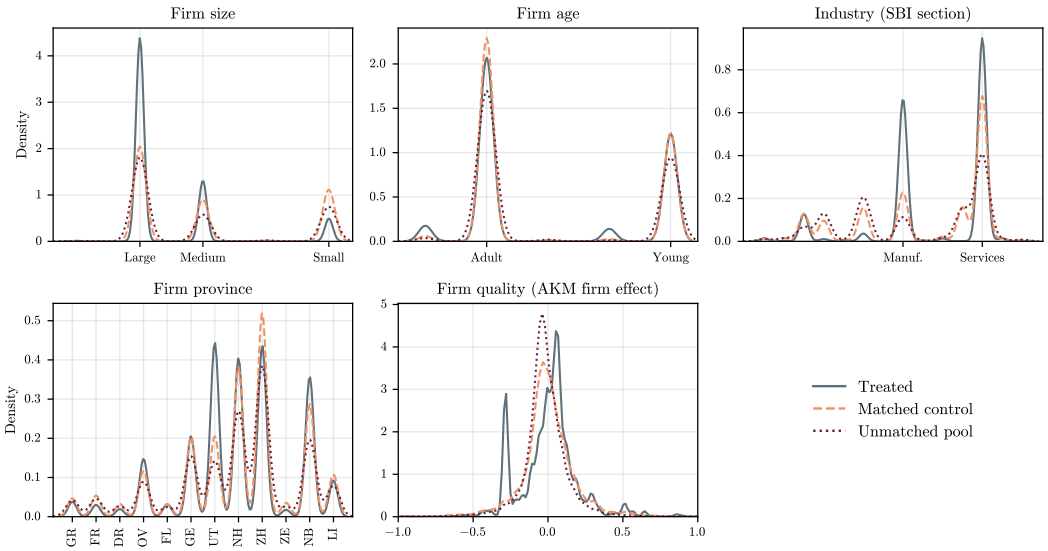


Notes: Kernel density functions of the worker characteristics in the matching year $(t-3)$ for treated incumbents (solid dark grey), matched control incumbents (dashed orange), and the full unmatched pool of incumbents at never-treated firms (dotted burgundy). Age class, gender, parental migration background, education level, and field of study are categorical, and their densities are estimated on the category codes. Age classes are ordered from youngest (0) to oldest (5). For categorical characteristics, the height of a peak is not proportional to the share of workers in that category, because the smoothing differs across groups. Table 1 in the main text reports the shares.

Firm characteristics are harder to balance (see Figure B.5), because treated firms differ most from the rest of the economy and we prioritise a close match on worker characteristics. Firm age is well balanced, with 59% of treated(defined as firms at least 10 years old) and about a third of

both groups at young firms. Firm location is also close. The provincial distribution of matched controls follows that of treated incumbents more closely than the unmatched pool does. Firm quality, measured by the AKM firm effect, is balanced at the mean (SMD 0.026), although the distribution for treated incumbents is more concentrated, with peaks at particular firms. Firm size and industry are not balanced. Treated incumbents are concentrated at large firms (70% of treated against 50% of matched control incumbents) and in manufacturing (37% against 15%, see Table 1 in the main text). This size difference is unlikely to affect our main results, as only 9% of firms in our sample are large and replacement occurs only within, rather than across, treatment cohorts. We assess this concern through an extensive firm-level heterogeneity analysis. Matching moves the control group towards the treated group on industry, raising the manufacturing share from about 8% in the unmatched pool to 15%, but not on firm size, where matched controls are slightly less often at large firms than the unmatched pool (50% against about 57%). We therefore do not rely on matching alone for these two characteristics. Instead, we estimate the effect within firm size classes and within industries, so that treated and control incumbents in each regression work at firms of the same size or in the same industry (see Appendices C.2.1 and C.2.2).

Figure B.5: Kernel Density Functions for the Firm Matching Characteristics

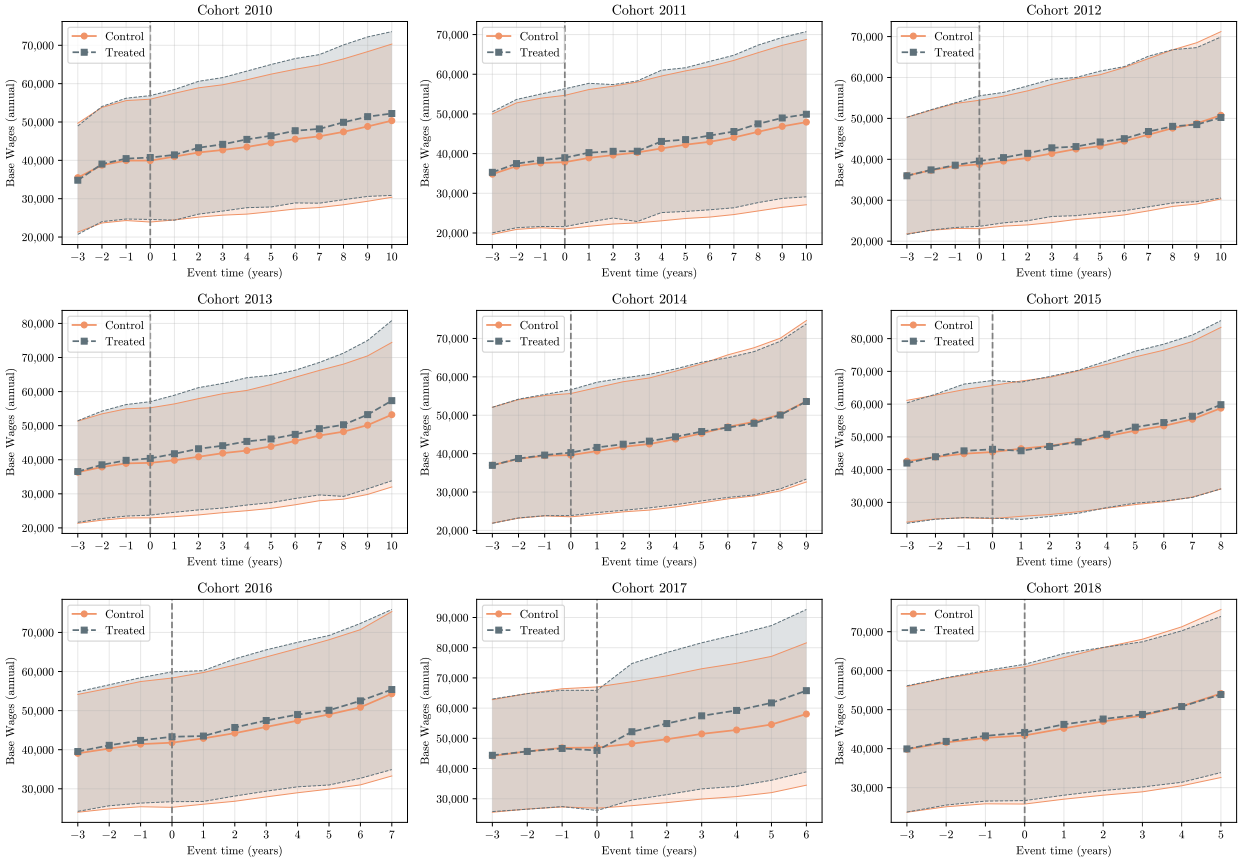


Notes: Kernel density functions of the firm characteristics in the matching year ($t-3$) for treated incumbents (solid dark grey), matched control incumbents (dashed orange), and the full unmatched pool of incumbents at never-treated firms (dotted burgundy). Firm size, firm age, industry, and province are categorical, and their densities are estimated on the category codes. Only the main categories are labelled. Firm size classes are large (more than 250 employees), medium (50–250) and small (fewer than 50). Industries are sections of the Dutch standard industrial classification (SBI), with manufacturing and services labelled. Provinces are ordered by CBS province code, from Groningen (GR) to Limburg (LI). For categorical characteristics, the height of a peak is not proportional to the share of workers in that category. Table 1 in the main text reports the shares. The AKM firm effect is the firm fixed effect from a wage decomposition on the universe of Dutch employer-employee data.

As a further check, Figure B.6 plots treated and matched control base wages separately for each cohort entering treatment between 2010 and 2018. In every cohort, treated and control wages are close before treatment. After treatment, treated wages rise above control wages in five of the nine cohorts, most clearly in 2010, 2011, 2013, 2016 and 2017. In the 2012, 2014, 2015, and 2018 cohorts, the two series stay close. The pooled pattern of raw base wage trends for treated and matched

control groups shown in Figure 2 in the main text is therefore not driven by a single cohort.

Figure B.6: Raw Base Wage Trends by Treatment Cohort



Notes: Each panel shows the mean annual base wage of treated incumbents (grey, dashed, squares) and matched control incumbents (orange, solid, circles) for one treatment cohort (2010-2018) by event time. Shaded bands denote ± 1 standard deviation around the mean. The dashed vertical line marks the cohort's treatment year ($\ell = 0$). Event time is shown up to $\ell = 10$.

C. Results Appendix

C.1 Rent Sharing: Event-Study Estimates

Table C.1 reports the full event-time coefficients underlying the interaction-weighted average treatment effects on the treated (IW-ATTs) in Table 2 in the main text and the event study in Figure 3 in the main text, for Samples A, B, and C and all five outcomes. Table C.2 reports the same estimates in euros per year for income outcomes and hours per year for hours outcomes.

The tables show rent sharing through base wage increases. In euros, the base wage IW-ATT is €1,403 per year in Sample A, €696 in Sample B, and €2,300 in Sample C. The other pay and hours outcomes move little. Bonuses do not respond significantly in any sample. The largest point estimate, 1.66% in Sample C, is not significant. Overtime income and overtime hours fall slightly, by 0.06% and 0.07% of the respective control means in Sample A.

The coefficients at long horizons come from fewer cohorts. Because the panel ends in 2023, all nine treatment cohorts contribute to event times up to $\ell = 5$, the 2010–2016 cohorts to $\ell = 7$, and only the 2010–2013 cohorts, which hold 73% of treated incumbents, to $\ell = 10$ (Appendix C.3.1). Hence, up to $\ell = 5$ the coefficients are based on all cohorts, and the effect already reaches 3.95% at $\ell = 5$ in Sample A, 2.36% in Sample B, and 3.43% in Sample C. The further growth beyond $\ell = 5$ can partly reflect the changing set of cohorts.

C.2 Heterogeneity

This section reports the subsample estimates behind the heterogeneity results in Section 5.1, by firm size and age, by industry, and by worker characteristics. Table C.3 brings the IW-ATTs together, in percent and in euros, with tests of equal effects within each dimension.

C.2.1 Firm Size and Age

We split Sample A into six firm size and age groups, fixed at $t - 3$, the matching year. These are young firms (less than 10 years old), small and young firms (less than 10 years old and fewer than 50 employees), small firms (fewer than 50 employees), medium-sized firms (50–250 employees), large firms (more than 250 employees), and large firms excluding the largest corporations (250–5,000 employees). Treated and control incumbents in each group work at firms of the same size or age class, so these splits also remove the residual size imbalance left after matching (Appendix B.4). Additionally, the firm heterogeneity analysis controls for how different regulatory frameworks affect firms of different sizes.

Table C.4 reports the IW-ATTs together with the pre-treatment coefficients, and Figure C.1 plots the event studies. The effects are largest among small firms (4.66%) and small, young firms (5.35%). However, both groups, as well as medium-sized firms, also show significant pre-treatment wage growth. Their coefficients at $\ell = -3$ and $\ell = -2$ are negative and significant at the 5% level. Young firms show a large negative coefficient at $\ell = -3$ (−5.64%, significant at the 5% level). We

Table C.1: Event-Study Estimates and IW-ATT by Sample

	Base wage (1)	Bonuses (2)	Overtime income (3)	Base hours (4)	Overtime hours (5)
Sample A: Full sample					
IW-ATT	3.275*** (0.671)	0.385 (0.530)	−0.057** (0.029)	0.579*** (0.178)	−0.068** (0.033)
$t-3$	−2.739** (1.349)	−0.962* (0.546)	0.037 (0.047)	−1.846 (1.293)	0.092 (0.059)
$t-2$	−0.317 (0.392)	−0.229 (0.402)	0.047* (0.028)	0.542*** (0.189)	0.060** (0.030)
$t-1$	0.381* (0.220)	−0.166 (0.214)	−0.029 (0.031)	0.901*** (0.186)	0.006 (0.023)
$t+1$	0.754 (0.576)	−0.089 (0.304)	0.036* (0.022)	−0.816* (0.495)	0.010 (0.019)
$t+2$	2.413*** (0.504)	−0.196 (0.566)	0.004 (0.024)	0.447*** (0.170)	−0.030 (0.025)
$t+3$	2.593*** (0.664)	0.644 (0.976)	0.009 (0.036)	0.383 (0.395)	−0.005 (0.052)
$t+4$	3.908*** (0.690)	−0.401 (0.690)	−0.008 (0.030)	0.764*** (0.212)	−0.033 (0.037)
$t+5$	3.948*** (0.755)	−0.205 (0.726)	−0.034 (0.032)	0.423* (0.249)	−0.060 (0.042)
$t+6$	4.525*** (0.784)	0.083 (0.703)	−0.112** (0.044)	0.931*** (0.253)	−0.129*** (0.043)
$t+7$	3.179*** (0.933)	1.064 (0.723)	−0.141*** (0.052)	0.482 (0.440)	−0.144*** (0.048)
$t+8$	3.975*** (0.989)	0.490 (0.729)	−0.124*** (0.047)	1.021*** (0.290)	−0.143*** (0.049)
$t+9$	4.314*** (1.169)	0.346 (0.794)	−0.145*** (0.049)	1.338*** (0.277)	−0.158*** (0.053)
$t+10$	3.971** (1.859)	0.100 (0.811)	−0.180** (0.075)	0.679 (0.635)	−0.150** (0.070)
Observations			7,239,116		
Sample B: Non-R&D workers					
IW-ATT	2.050*** (0.633)	0.356 (0.396)	−0.072* (0.043)	0.297** (0.135)	−0.045 (0.044)
$t-3$	−1.131* (0.584)	−0.130 (0.450)	0.089 (0.094)	−0.278 (0.341)	0.134 (0.094)
$t-2$	−0.205 (0.399)	0.272 (0.594)	0.051 (0.057)	0.433*** (0.127)	0.073 (0.049)
$t-1$	0.192 (0.167)	−0.024 (0.193)	−0.013 (0.042)	0.598*** (0.126)	0.036 (0.030)
$t+1$	−0.390 (0.617)	−0.057 (0.246)	0.045 (0.034)	−0.866 (0.580)	−0.001 (0.028)
$t+2$	1.126*** (0.235)	0.299 (0.212)	−0.022 (0.041)	0.311*** (0.110)	−0.054 (0.037)
$t+3$	0.942* (0.563)	0.318 (0.226)	−0.046 (0.054)	0.019 (0.437)	−0.043 (0.061)
$t+4$	2.462*** (0.583)	−0.142 (0.483)	−0.077 (0.050)	0.578*** (0.164)	−0.071 (0.052)
$t+5$	2.363*** (0.672)	0.079 (0.618)	−0.136** (0.058)	0.322 (0.212)	−0.099* (0.055)
$t+6$	2.909*** (0.726)	0.303 (0.384)	−0.204*** (0.064)	0.535* (0.321)	−0.151*** (0.056)
$t+7$	2.665*** (0.866)	0.472 (0.658)	−0.182*** (0.067)	0.329 (0.317)	−0.123** (0.059)
$t+8$	3.063*** (0.953)	0.433 (0.684)	−0.193*** (0.072)	0.690*** (0.229)	−0.118* (0.063)
$t+9$	3.556*** (1.086)	0.167 (0.751)	−0.159** (0.072)	0.908*** (0.219)	−0.089 (0.069)
$t+10$	2.724 (2.272)	0.139 (0.690)	0.003 (0.097)	−0.253 (0.765)	0.059 (0.091)
Observations			1,734,690		
Sample C: R&D workers					
IW-ATT	4.143*** (1.107)	1.658* (0.861)	−0.030 (0.032)	0.633 (0.385)	−0.070*** (0.027)
$t-3$	−1.168 (0.730)	−0.546 (0.433)	0.068** (0.031)	−0.403 (0.785)	0.070** (0.030)
$t-2$	0.089 (0.525)	−0.159 (0.338)	0.104*** (0.026)	1.005* (0.524)	0.096*** (0.027)
$t-1$	0.605 (0.449)	0.003 (0.308)	0.022 (0.027)	1.370*** (0.515)	0.027 (0.024)
$t+1$	0.709 (0.800)	0.276 (0.557)	0.007 (0.025)	−0.817 (0.680)	−0.026* (0.013)
$t+2$	2.233*** (0.778)	0.269 (0.946)	0.000 (0.024)	0.305 (0.428)	−0.038** (0.018)
$t+3$	2.452*** (0.833)	0.916 (0.966)	−0.020 (0.033)	0.504 (0.424)	−0.036 (0.032)
$t+4$	3.154*** (0.897)	0.776 (0.767)	−0.013 (0.035)	0.695 (0.438)	−0.052* (0.029)
$t+5$	3.432*** (0.968)	1.652* (0.989)	−0.007 (0.040)	0.518 (0.452)	−0.056* (0.034)
$t+6$	4.147*** (1.019)	2.177** (1.008)	0.007 (0.041)	0.934** (0.471)	−0.062* (0.033)
$t+7$	4.511*** (1.099)	2.974** (1.253)	−0.037 (0.047)	1.127*** (0.378)	−0.074* (0.041)
$t+8$	5.846*** (1.489)	1.908 (1.210)	−0.015 (0.050)	1.024*** (0.395)	−0.071 (0.043)
$t+9$	6.273*** (1.773)	2.016** (1.018)	−0.024 (0.056)	1.003*** (0.382)	−0.089* (0.049)
$t+10$	8.614*** (2.272)	2.329* (1.197)	−0.146* (0.077)	1.384*** (0.454)	−0.172*** (0.064)
Observations			606,033		

Notes: Each column within a sample reports a separate estimation of equation (1) in the main text using the interaction-weighted estimator of Sun and Abraham (2021), with $\ell = 0$ as the omitted reference period. The shaded IW-ATT row is the interaction-weighted average treatment effect over all post-treatment periods. Income outcomes (base wage, bonuses, overtime income) are expressed as a percentage of the sample's control-group mean annual base wage (€42,850 in Sample A, €33,960 in Sample B, €55,522 in Sample C), and hours outcomes (base hours, overtime hours) as a percentage of the sample's control-group mean annual hours worked (1,871, 1,952, and 1,935 hours). All nine treatment cohorts contribute to event times up to $t+5$, and $t+10$ is estimated from the 2010-2013 cohorts. Standard errors, clustered at the firm level, in parentheses. Statistical significance is indicated as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table C.2: Event-Study Estimates and IW-ATT by Sample: Raw Effects

	Base wage (1)	Bonuses (2)	Overtime income (3)	Base hours (4)	Overtime hours (5)
Sample A: Full sample					
IW-ATT	1,403.3*** (287.6)	165.1 (227.3)	−24.3** (12.3)	10.8*** (3.3)	−1.3** (0.6)
$t-3$	−1,173.7** (577.9)	−412.4* (233.9)	15.7 (20.2)	−34.5 (24.2)	1.7 (1.1)
$t-2$	−135.9 (167.8)	−98.1 (172.2)	20.1* (12.0)	10.1*** (3.5)	1.1** (0.6)
$t-1$	163.4* (94.1)	−71.1 (91.6)	−12.2 (13.4)	16.9*** (3.5)	0.1 (0.4)
$t+1$	322.9 (246.6)	−38.2 (130.4)	15.6* (9.4)	−15.3* (9.3)	0.2 (0.4)
$t+2$	1,033.8*** (215.8)	−84.2 (242.5)	1.5 (10.4)	8.4*** (3.2)	−0.6 (0.5)
$t+3$	1,111.1*** (284.4)	275.9 (418.2)	4.0 (15.6)	7.2 (7.4)	−0.1 (1.0)
$t+4$	1,674.4*** (295.5)	−172.0 (295.6)	−3.4 (12.7)	14.3*** (4.0)	−0.6 (0.7)
$t+5$	1,691.6*** (323.7)	−87.9 (311.2)	−14.5 (13.9)	7.9* (4.7)	−1.1 (0.8)
$t+6$	1,939.0*** (336.0)	35.7 (301.2)	−48.1** (19.0)	17.4*** (4.7)	−2.4*** (0.8)
$t+7$	1,362.2*** (399.8)	455.8 (309.9)	−60.3*** (22.3)	9.0 (8.2)	−2.7*** (0.9)
$t+8$	1,703.4*** (423.7)	210.1 (312.5)	−53.3*** (20.3)	19.1*** (5.4)	−2.7*** (0.9)
$t+9$	1,848.8*** (501.0)	148.5 (340.3)	−62.1*** (20.8)	25.0*** (5.2)	−2.9*** (1.0)
$t+10$	1,701.4** (796.7)	42.9 (347.4)	−77.1** (31.9)	12.7 (11.9)	−2.8** (1.3)
Observations			7,239,116		
Sample B: Non-R&D workers					
IW-ATT	696.3*** (215.0)	120.9 (134.6)	−24.5* (14.8)	5.8** (2.6)	−0.9 (0.9)
$t-3$	−384.0* (198.2)	−44.3 (152.8)	30.1 (32.0)	−5.4 (6.7)	2.6 (1.8)
$t-2$	−69.7 (135.4)	92.3 (201.6)	17.3 (19.3)	8.5*** (2.5)	1.4 (1.0)
$t-1$	65.3 (56.8)	−8.3 (65.4)	−4.3 (14.4)	11.7*** (2.5)	0.7 (0.6)
$t+1$	−132.5 (209.4)	−19.2 (83.4)	15.3 (11.7)	−16.9 (11.3)	0.0 (0.5)
$t+2$	382.4*** (79.9)	101.4 (71.9)	−7.5 (13.8)	6.1*** (2.1)	−1.1 (0.7)
$t+3$	320.0* (191.1)	108.0 (76.7)	−15.7 (18.2)	0.4 (8.5)	−0.8 (1.2)
$t+4$	836.1*** (197.9)	−48.1 (163.9)	−26.2 (17.1)	11.3*** (3.2)	−1.4 (1.0)
$t+5$	802.4*** (228.2)	26.8 (209.9)	−46.3** (19.6)	6.3 (4.1)	−1.9* (1.1)
$t+6$	987.8*** (246.6)	102.9 (130.5)	−69.4*** (21.7)	10.4* (6.3)	−3.0*** (1.1)
$t+7$	905.1*** (294.1)	160.4 (223.3)	−61.7*** (22.8)	6.4 (6.2)	−2.4** (1.1)
$t+8$	1,040.2*** (323.5)	147.1 (232.2)	−65.5*** (24.4)	13.5*** (4.5)	−2.3* (1.2)
$t+9$	1,207.5*** (368.7)	56.8 (255.0)	−54.1** (24.5)	17.7*** (4.3)	−1.7 (1.3)
$t+10$	925.1 (771.5)	47.3 (234.5)	1.1 (32.9)	−4.9 (14.9)	1.1 (1.8)
Observations			1,734,690		
Sample C: R&D workers					
IW-ATT	2,300.3*** (614.4)	920.3* (477.9)	−16.6 (17.6)	12.2 (7.5)	−1.4*** (0.5)
$t-3$	−648.7 (405.3)	−303.0 (240.7)	38.0** (17.4)	−7.8 (15.2)	1.4** (0.6)
$t-2$	49.5 (291.6)	−88.3 (187.8)	57.8*** (14.7)	19.5* (10.1)	1.9*** (0.5)
$t-1$	335.9 (249.3)	1.8 (170.9)	12.3 (14.8)	26.5*** (10.0)	0.5 (0.5)
$t+1$	393.6 (444.2)	153.0 (309.5)	3.8 (13.9)	−15.8 (13.2)	−0.5* (0.3)
$t+2$	1,239.9*** (432.0)	149.4 (525.2)	−0.1 (13.4)	5.9 (8.3)	−0.7** (0.4)
$t+3$	1,361.5*** (462.5)	508.3 (536.1)	−11.3 (18.1)	9.8 (8.2)	−0.7 (0.6)
$t+4$	1,751.4*** (498.2)	430.9 (426.1)	−7.1 (19.2)	13.5 (8.5)	−1.0* (0.6)
$t+5$	1,905.4*** (537.5)	917.5* (549.1)	−3.8 (22.0)	10.0 (8.7)	−1.1* (0.7)
$t+6$	2,302.7*** (566.0)	1,208.9** (559.9)	3.8 (22.6)	18.1** (9.1)	−1.2* (0.6)
$t+7$	2,504.5*** (610.4)	1,651.2** (695.7)	−20.3 (26.1)	21.8*** (7.3)	−1.4* (0.8)
$t+8$	3,246.1*** (826.6)	1,059.1 (671.6)	−8.6 (27.8)	19.8*** (7.6)	−1.4 (0.8)
$t+9$	3,482.7*** (984.6)	1,119.2** (565.1)	−13.2 (31.2)	19.4*** (7.4)	−1.7* (0.9)
$t+10$	4,782.7*** (1,261.6)	1,292.9* (664.4)	−81.0* (42.7)	26.8*** (8.8)	−3.3*** (1.2)
Observations			606,033		

Notes: Same estimates as in Table C.1, in raw units. Income outcomes (base wage, bonuses, overtime income) are in euros per year, and hours outcomes (base hours, overtime hours) are in hours per year. Standard errors, clustered at the firm level, in parentheses. Statistical significance is indicated as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table C.3: Heterogeneity in the IW-ATT on Base Wages: Percentage and Level Effects

Dimension	Subgroup	Mean base wage (€)	IW-ATT (% of own mean) Estimate	(s.e.)	IW-ATT (€/year)	Observations	Equality p-value
Firm size	Small	41,200	4.66***	(0.23)	1,920	1,209,178	<0.01
	Medium	42,400	1.48***	(0.22)	630	1,597,262	
	Large	43,800	2.79***	(0.92)	1,220	4,399,758	
Industry	Manufacturing	42,800	2.94***	(0.98)	1,260	2,094,544	0.69
	Services	44,100	2.44***	(0.80)	1,080	3,266,725	
Education	Below upper secondary	34,000	2.05***	(0.63)	700	1,734,690	0.15
	Upper secondary	38,200	2.56***	(0.52)	980	2,821,670	
	Tertiary	53,700	4.55***	(1.12)	2,450	2,603,895	
Field of study	STEM	45,200	3.43***	(0.89)	1,550	1,417,599	0.57
	Non-STEM	42,400	2.82***	(0.58)	1,200	5,815,971	
Gender	Male	45,500	3.12***	(0.63)	1,420	5,469,847	0.90
	Female	35,800	3.24***	(0.78)	1,160	1,769,247	
Parental migration background	Native-born parents	43,100	3.41***	(0.67)	1,470	6,217,460	0.40
	Foreign-born parent(s)	41,000	2.59***	(0.72)	1,060	1,017,893	

Notes: IW-ATT on base wages from separate estimations of equation (1) in the main text on each subgroup of Sample A. Below-upper-secondary workers form Sample B. The ATT is expressed as a percentage of the subgroup's control-group mean annual base wage and in euros per year. Mean base wage is rounded to the nearest €100 and the euro-denominated ATT to the nearest €10. The *equality p-value* is from a Wald test that the ATTs within a dimension are equal, treating the subgroup estimates as independent. Because subgroups share employers, their estimates are likely positively correlated, which makes the test conservative. Standard errors, clustered at the firm level, in parentheses. Statistical significance is indicated as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

therefore cannot separate the larger effects at small and young firms from pre-existing wage growth. This is consistent with these firms being on an expansion path, in which early employees, often including owners, follow different wage-growth and bargaining paths than employees at established firms.

At large firms, none of the pre-treatment coefficients is significant at the 5% level, and the IW-ATT is 2.79%, close to the full-sample estimate of 3.27%. Excluding the largest corporations, the IW-ATT is 3.21%. The effect is therefore present, and of similar size, where both pre-trends and the size imbalance are least of a concern. Table C.4 also reports the number of observations in each category. We see that large differences in sample sizes are not driving the observed results.

Table C.5 reports the event-time coefficients on base wages for large firms in Samples A, B, and C, plotted in Figure 4 in the main text. The IW-ATT is 1.76% for non-R&D workers (significant only at the 10% level) and 5.80% for R&D workers. None of the pre-treatment coefficients in the three large-firm samples is significant at the 5% level (see Table C.5). We consider these effects a clean estimates of rent sharing from the innovation box tax break.

C.2.2 Industry

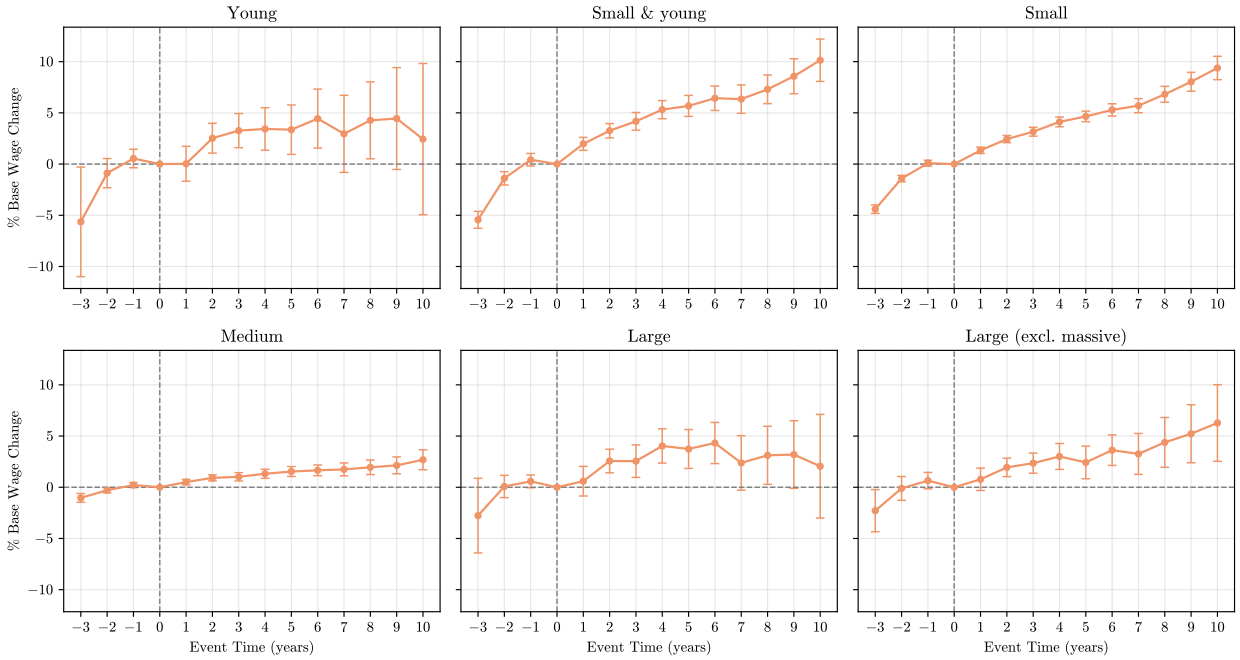
Sectoral collective agreements set a large part of Dutch wage growth (see Section 2.1 in the main text), and treated incumbents work more often in manufacturing than their matched controls (see

Table C.4: Base Wage IW-ATT by Firm Size and Age at Treatment

		Pre-treatment coefficients				Mean base
	IW-ATT	$\ell = -3$	$\ell = -2$	$\ell = -1$	Observations	wage (€)
Young	3.01**	−5.64**	−0.88	0.54	2,292,608	43,488
	(1.22)	(2.73)	(0.73)	(0.46)		
Small & young	5.35***	−5.44***	−1.39***	0.42	387,480	43,322
	(0.43)	(0.42)	(0.33)	(0.31)		
Small	4.66***	−4.40***	−1.41***	0.09	1,209,178	41,239
	(0.23)	(0.21)	(0.16)	(0.15)		
Medium	1.48***	−1.03***	−0.29**	0.22*	1,597,262	42,402
	(0.22)	(0.22)	(0.13)	(0.12)		
Large	2.79***	−2.77	0.08	0.58*	4,399,758	43,750
	(0.92)	(1.86)	(0.56)	(0.32)		
Large, excl. >5,000 employees	3.21***	−2.29**	−0.11	0.65	2,993,151	43,482
	(0.84)	(1.05)	(0.59)	(0.41)		

Notes: IW-ATT estimates and pre-treatment event-time coefficients on base wages, obtained by estimating equation (1) in the main text separately for each firm size and age group in Sample A (defined at $t-3$), with $\ell = 0$ as the reference period. Estimates are expressed as a percentage of the group's control-group mean annual base wage, reported in the last column. Young firms are less than 10 years old. Small firms have fewer than 50 employees, medium-sized firms 50–250, and large firms more than 250. Standard errors, clustered at the firm level, in parentheses. Statistical significance is indicated as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure C.1: Event-Study Estimates on Base Wage by Firm Size and Age at Treatment



Notes: Each panel presents event-time coefficients on base wages, estimated using equation (1) in the main text for the corresponding firm size or age group reported in Table C.4, with $\ell = 0$ as the omitted reference period. The error bars denote 95% confidence intervals. Coefficients are expressed as a percentage of the group's control-group mean annual base wage. Vertical dashed lines mark the treatment year, and horizontal dashed lines mark zero.

Table C.5: Event-Study Estimates on Base Wage, Large Firms

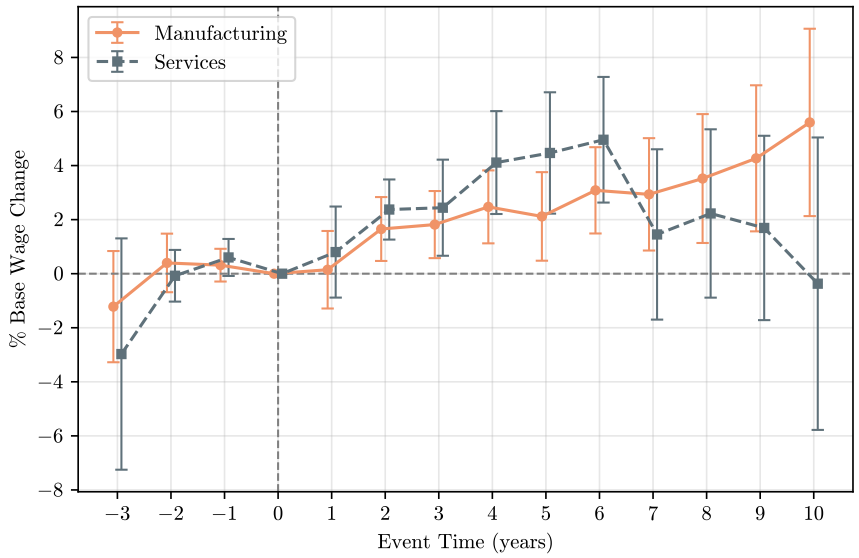
	Sample A Full sample	Sample B Non-R&D workers	Sample C R&D workers
IW-ATT	2.794*** (0.925)	1.765* (0.918)	5.802*** (1.878)
$t-3$	−2.769 (1.859)	−0.888 (0.688)	−0.643 (1.352)
$t-2$	0.076 (0.556)	0.106 (0.502)	0.929 (0.958)
$t-1$	0.579* (0.319)	0.198 (0.241)	1.129 (0.795)
$t+1$	0.594 (0.739)	−1.172 (1.051)	0.778 (1.432)
$t+2$	2.564*** (0.589)	0.938*** (0.275)	3.360** (1.403)
$t+3$	2.552*** (0.811)	0.449 (0.762)	3.581** (1.489)
$t+4$	4.034*** (0.854)	2.583*** (0.837)	4.490*** (1.595)
$t+5$	3.739*** (0.969)	1.929** (0.875)	4.773*** (1.707)
$t+6$	4.320*** (1.024)	2.631*** (0.965)	5.747*** (1.791)
$t+7$	2.380* (1.357)	2.500* (1.344)	6.273*** (1.921)
$t+8$	3.121** (1.445)	2.921* (1.528)	8.085*** (2.468)
$t+9$	3.194* (1.684)	3.471* (1.777)	8.798*** (2.942)
$t+10$	2.053 (2.581)	1.841 (3.412)	11.749*** (3.714)
Observations	4,399,758	895,099	382,426
Mean base wage (€)	43,750	33,129	45,896

Notes: Event-time coefficients on base wages from estimations of equation (1) in the main text for workers at large firms (more than 250 employees at $t-3$), with $\ell = 0$ as the omitted reference period. Coefficients are expressed as a percentage of the sample's control-group mean annual base wage, reported in the last row. The shaded IW-ATT row is the interaction-weighted average treatment effect over all post-treatment periods. Standard errors, clustered at the firm level, in parentheses. Statistical significance is indicated as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 1 in the main text). We therefore also estimate the effect separately for incumbents at manufacturing firms and at services firms, so that treated and control workers in each regression work in the same broad sector. Rent sharing is present in both sectors and of similar size. The IW-ATT is 2.94% in manufacturing and 2.44% in services, and we cannot reject that the two are equal ($p = 0.69$, Table C.3).

The dynamics differ (see Figure C.2). In services, the effect rises faster, to 5.0% at $\ell = 6$, and then falls and becomes imprecise. In manufacturing, the effect keeps growing, to 5.6% at $\ell = 10$. The long-horizon estimates rest on the 2010-2013 cohorts only (see Appendix C.1), so we do not interpret this difference further. Earlier cohorts contain more manufacturing firms (see Appendix A.3). Sectoral differences are therefore not the driver of the rent-sharing results. Rent sharing is present and comparable in magnitude across sectors.

Figure C.2: Event-Study Estimates on Base Wage, by Industry



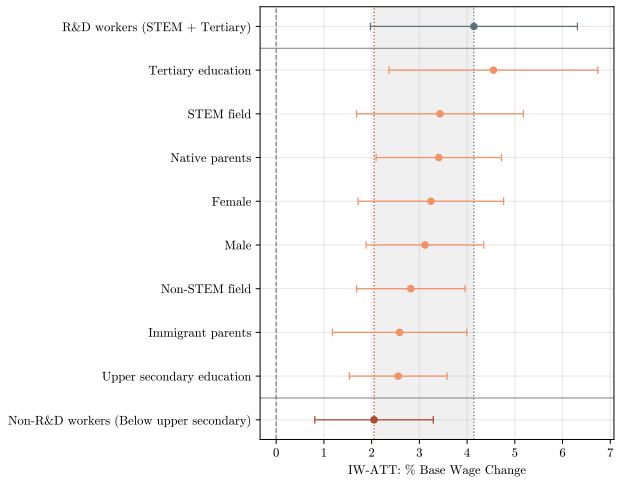
Notes: Event-time coefficients on base wages, obtained by estimating equation (1) in the main text separately for incumbent workers at manufacturing and services firms, with $\ell = 0$ as the omitted reference period. The error bars denote 95% confidence intervals. Coefficients are expressed as a percentage of the subsample's control-group mean annual base wage (€42,750 in manufacturing, €44,077 in services). The IW-ATT is 2.94% in manufacturing and 2.44% in services, with standard errors of 0.98 and 0.80 (see Table C.3). The vertical dashed line marks the treatment year, and the horizontal dashed line marks zero.

C.2.3 Worker Characteristics

This section reports the heterogeneity results by worker characteristics summarised in Section 5.1. Figure C.3 collects the IW-ATTs on base wages for ten subgroups, and Figure C.4 plots event-study estimates across four dimensions – education, field of study, gender, and parental migration background. Table C.3 reports the IW-ATTs. The observed effects reflect the full-sample estimates.

The base wage effect emerges gradually and grows over the first six years after treatment in every subgroup, consistent with wage adjustment through periodic renegotiation. Beyond $\ell = 6$, the estimates rest on fewer cohorts and the confidence intervals widen. For women, non-STEM

Figure C.3: Base Wage IW-ATT by Worker Characteristics



Notes: Each row presents the IW-ATT for base wages, together with its 95% confidence interval, obtained by estimating equation (1) in the main text separately for the indicated subgroup. Estimates are expressed as percentages of the corresponding control group's mean annual base wage. Subgroups are ordered by their IW-ATT estimates.

workers, and workers with foreign-born parents, the estimates fall back over these years.

Together with Samples B and C, these subgroups make up the ten estimates in Figure C.3. All ten IW-ATTs are significant at the 5% level. We also test whether the differences between subgroups are statistically significant. Within each worker dimension, we cannot reject equal effects at conventional levels (Table C.3).

Tertiary-educated workers show the largest point estimates at every post-treatment event time, followed by upper-secondary and then below-upper-secondary workers. The gap between tertiary and the other two groups widens over time, and the tertiary confidence interval also widens the most by the end of the window. The IW-ATTs are 2.05%, 2.56%, and 4.55%. Despite the consistent ordering, we cannot reject equality across the three groups ($p = 0.15$).

We compare workers with STEM and non-STEM degrees at all three education levels. By construction, all below-upper-secondary workers are classified as non-STEM (Appendix B.2). The two series are close through $\ell = 6$, after which the STEM estimates keep rising while the non-STEM estimates level off. The IW-ATTs are 3.43% for STEM and 2.82% for non-STEM workers ($p = 0.57$ for equality). Field of study alone therefore does not drive the gap between Samples B and C. The education split suggests that educational attainment is the more important dimension.

Women's estimates lie about one percentage point above men's from $\ell = 1$ to $\ell = 6$ and below them afterwards, when they also become imprecise. The IW-ATTs are 3.12% for men and 3.24% for women. We do not observe significant wage differences between the two.

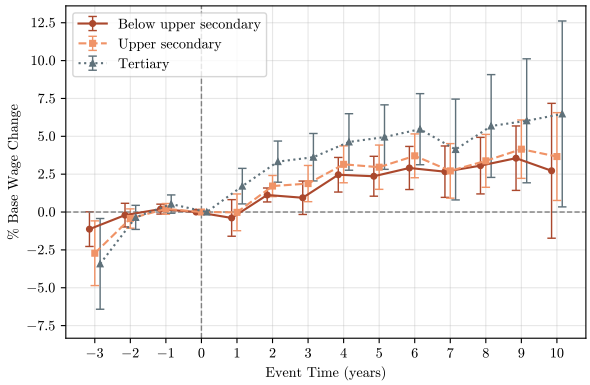
Workers with native-born parents and workers with at least one foreign-born parent follow nearly identical paths through $\ell = 6$. From $\ell = 7$ onward, the estimates for workers with foreign-born parents fall and become noisier, with wide confidence intervals and a visible dip at the final event time. The IW-ATTs are 3.41% and 2.59%.

Lastly, the heterogeneity results support our interpretation that the estimates capture rent shar-

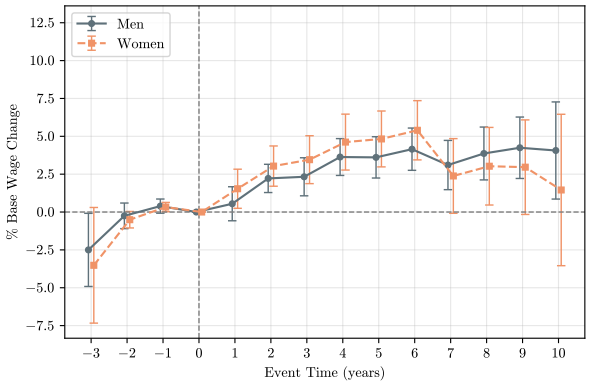
ing from the innovation box tax break. The effects differ only by firm size and age, characteristics that shape how firms set wages, and not along the other dimensions that we consider. This is consistent with what we would expect if our estimates capture rent sharing from the tax break, which is passed on to workers through wage renegotiations. Our estimates capture the cumulative pass-through over several rounds of renegotiation, and both the timing of these rounds and the wage paths around them differ between small and large firms. Large, established firms renegotiate wages regularly and follow stable wage paths, so the pass-through stands out clearly against baseline wage growth. Small and young firms follow steeper and more volatile wage paths, which makes the pass-through harder to separate from ordinary wage growth. We therefore interpret our estimates as rent sharing from the innovation tax break, identified most cleanly for large firms.

The absence of differences across industries further supports this interpretation. If wages were renegotiated through sector-wide bargaining, the pass-through would likely vary with each industry's bargaining structure. Instead, it is similar across industries, consistent with renegotiation taking place within firms' own established pay scales. This is also the channel through which a firm-specific tax break can affect wages since sector-wide agreements do not respond to one firm's tax position.

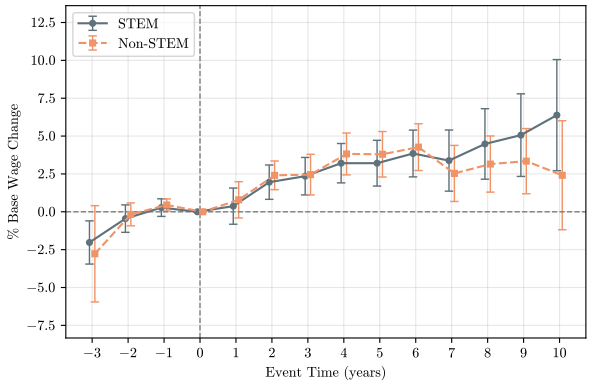
Figure C.4: Event-Study Estimates on Base Wages, by Worker Characteristics



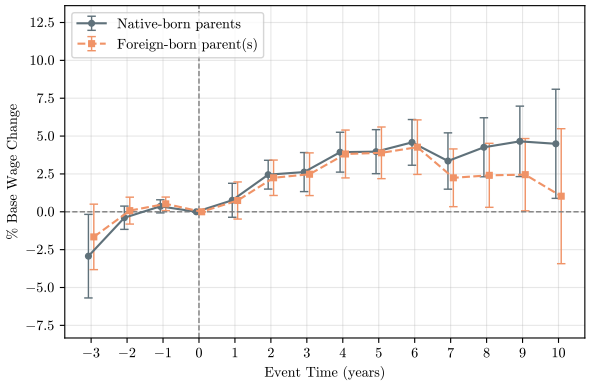
(a) Education



(b) Gender



(c) Field of Study



(d) Parental Migration Background

Notes: Each panel reports event-time coefficients on base wages, obtained by estimating equation (1) in the main text for the indicated subgroups of Sample A, with $\ell = 0$ as the omitted reference period. Below-upper-secondary workers form Sample B. The error bars denote 95% confidence intervals. Coefficients are expressed as a percentage of each subgroup's control-group mean annual base wage (Table C.3). All panels share the same vertical axis. Standard errors are clustered at the firm level.

C.3 Mobility

C.3.1 Retention at the Incumbency Firm

Table C.6 reports the results from the Kaplan–Meier curves (see Figure 5 in Section 5.2 of the main text) at selected horizons, together with the number of incumbents still at the incumbency firm. These are the workers who can enter the post-treatment wage panel of Appendix C.1, which follows incumbents while they are employed at the incumbency firm. Table C.6 shows how large this margin is. Five years after the first benefit year, the retention gap implies that 5.4% of treated stayers in Sample A and 8.0% in Sample B would have left at the control group's retention rate. In Sample C, the corresponding share is 0.2%, and it stays below 2% at every horizon.

Table C.6: Retention at the Incumbency Firm by Event Time

	Years since first benefit year (ℓ)					
	1	2	3	5	7	10
Sample A: Full sample						
Treated (%)	89.4	78.2	70.6	58.5	48.8	36.8
Matched control (%)	87.7	77.7	69.1	55.3	45.5	34.7
Difference (pp)	1.7	0.5	1.5	3.2	3.3	2.2
Treated still at firm	338,365	296,035	267,115	221,289	182,570	115,000
Controls still at firm	306,933	272,069	241,935	193,604	157,570	98,167
Sample B: Non-R&D workers						
Treated (%)	91.0	81.9	75.0	64.5	55.1	42.2
Matched control (%)	89.0	80.1	72.1	59.3	49.2	37.1
Difference (pp)	1.9	1.9	2.9	5.1	5.9	5.1
Treated still at firm	81,535	73,451	67,181	57,772	48,988	33,240
Controls still at firm	67,143	60,376	54,346	44,728	36,779	24,150
Sample C: R&D workers						
Treated (%)	88.2	75.1	67.1	55.3	46.8	37.0
Matched control (%)	86.8	76.4	67.9	55.2	46.0	36.4
Difference (pp)	1.5	−1.3	−0.8	0.1	0.8	0.6
Treated still at firm	34,742	29,559	26,410	21,778	18,168	11,289
Controls still at firm	19,451	17,119	15,222	12,374	10,159	6,253
Cohorts observed	2010–18	2010–18	2010–18	2010–18	2010–16	2010–13
Their share of treated (%)	100	100	100	100	92	73

Notes: Kaplan-Meier estimates of the share of incumbents still employed at their incumbency firm ℓ years after the firm's first benefit year, and the number of treated and matched control incumbents still at the firm in that year. The bottom rows list the treatment cohorts observed at each event time, given that the panel ends in 2023, and their share of all treated incumbents in Sample A.

C.3.2 Event-Study Estimates

The mobility panel contains the same treated and matched control incumbents as the wage panel, but follows them through all subsequent employment spells up to 2023. The outcome, Switch_{it}, equals one from the first year in which worker i is observed working in a firm other than the

incumbency firm, and zero otherwise. Once a worker has switched, it stays equal to one. Exits from the payslip data are not switches, and years in which a worker is not observed in the payslip data do not enter the panel. Appendix C.3.3 describes these other exits, which do not reflect a switch to another firm.

We follow the job-loss literature (Britto et al., 2022; Dos Santos et al., 2026; Jarosch, 2023) in our descriptive analysis of mobility. Control workers inherit the first benefit year G_i of the treated worker they are matched to, so event time $\ell = t - G_i$ is defined for every worker. We estimate:

$$\text{Switch}_{it} \;=\; \sum_{\substack{\ell=-3 \\ \ell\neq 0}}^{10} \beta_\ell \left(\text{Treated}_i \cdot \mathbb{1}\{t - G_i = \ell\}\right) \;+\; \alpha_i \;+\; \lambda_{G_i t} \;+\; \varepsilon_{it}, \tag{C.1}$$

where Treated_i equals one for treated incumbents, α_i are worker fixed effects, and $\lambda_{G_i t}$ are cohort-by-year fixed effects. Within a cohort, calendar year and event time coincide, so $\lambda_{G_i t}$ absorbs the switching profile of control workers separately for each cohort, and β_ℓ measures the difference in the cumulative probability of having switched firms between treated and matched control incumbents of the same cohort, ℓ years after the first benefit year. The results are presented in Figure 6 of the main text.

Incumbents are employed at the incumbency firm from $t-3$ through t by construction, so $\text{Switch}_{it} = 0$ for every worker before the first benefit year. The pre-treatment coefficients in Table C.7 are therefore zero or close to zero by construction and carry no information on parallel trends in mobility. Matching balances baseline attachment to the employer instead. Treated and control incumbents are matched on tenure at $t-3$ (SMD -0.065, Table B.6). By 2023, 23.6% of matched control incumbents in Sample A and 17.7% in Sample B have moved to another firm (Table C.8). As a rough benchmark, the pooled ATTs of -2.19 and -3.47 percentage points correspond to 9% and 20% of these cumulative shares. R&D workers in Sample C are the most mobile group, with 31.3% of control and 30.9% of treated incumbents moving to another firm by 2023. The absence of a retention effect in Sample C therefore does not reflect low baseline mobility. This is consistent with R&D workers being in demand at treated and control firms alike (Section 5.2). Our estimates therefore capture rent sharing not as a retention device, but as firm-wide wage renegotiation. The tax break enlarges the firm's surplus, and the firm passes part of it on to all incumbent workers through its established pay scales.

C.3.3 Where Do Leavers Go?

Table C.8 splits incumbents by what we observe after they leave the incumbency firm. By the end of the panel, 59.5% of treated and 61.6% of matched control incumbents in Sample A have left. We distinguish leavers who are next observed at another firm from those who are never observed again in the payslip data before the panel ends in 2023. The remaining leavers (10.9% of treated and 10.3% of control leavers in Sample A) have gaps in their post-incumbency employment spells. The payslip data cover employees only, so workers who are never observed again may have retired, become unemployed, entered long-term sick leave or disability, become self-employed, or emigrated.

As a share of all incumbents, 31.2% of treated and 31.7% of control workers are never observed again, so treated workers are slightly more likely to remain in the data. Among leavers, the share of never observed again is higher for treated workers (52.4% versus 51.5%), but this reflects the

Table C.7: Event-Study Estimates: Probability of Switching Firms

	Sample A: Full sample (1)	Sample B: Non-R&D workers (2)	Sample C: R&D workers (3)
Pooled ATT	-2.19^{*}	-3.47^{***}	-0.19
	(1.27)	(0.95)	(1.96)
$t-3$	0.00	-0.07	0.08
	(0.00)	(0.08)	(0.49)
$t-2$	0.00	0.03	0.07
	(0.00)	(0.06)	(0.41)
$t-1$	0.00	0.06	0.04
	(0.00)	(0.04)	(0.32)
t	Reference period (omitted)		
$t+1$	-1.22^{***}	-1.77^{***}	-0.53
	(0.41)	(0.29)	(0.91)
$t+2$	-0.60	-1.95^{***}	0.66
	(0.99)	(0.64)	(1.67)
$t+3$	-1.30	-2.00^{**}	0.74
	(1.29)	(0.90)	(2.13)
$t+4$	-2.34	-3.11^{***}	0.36
	(1.48)	(1.02)	(2.28)
$t+5$	-2.88^{*}	-3.90^{***}	0.30
	(1.58)	(1.14)	(2.36)
$t+6$	-3.15^{*}	-4.63^{***}	-0.10
	(1.71)	(1.25)	(2.44)
$t+7$	-3.61^{**}	-5.74^{***}	-0.84
	(1.80)	(1.35)	(2.57)
$t+8$	-3.58^{*}	-6.28^{***}	-1.65
	(2.00)	(1.50)	(2.68)
$t+9$	-2.80	-6.60^{***}	-0.91
	(2.19)	(1.70)	(2.83)
$t+10$	-2.85	-7.09^{***}	-0.53
	(2.47)	(1.94)	(3.02)
Observations	7,923,037	1,773,502	710,533
R^2	0.585	0.573	0.621
Adjusted R^2	0.544	0.522	0.583
RMSE	0.226	0.171	0.245
F-statistic	852.22	390.77	145.40

Notes: Event-time coefficients β_ℓ from estimations of equation (C.1) in the main text and the pooled ATT (see Figure 6 in the main text) on Switch_{it}, an indicator equal to one from the first year an incumbent is observed working in a firm other than the incumbency firm, in percentage points ($\times 100$). $\ell = 0$, the first benefit year, is the omitted reference period. Pre-treatment coefficients are zero or close to zero by construction. Standard errors, clustered at the incumbency-firm level, in parentheses. Statistical significance is indicated as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

smaller number of treated leavers. Fewer treated workers move to another firm, while the share of all incumbents who leave the data is almost the same in both groups. Because treated and control workers are matched on age, retirement at the statutory pension age cannot explain differences between the two groups. Remaining candidates include early retirement, exit into unemployment, and exit into sick leave. Previous literature has examined the effects of rent sharing on retirement decisions and health outcomes (Hampton and Totty, 2023; Frimmel et al., 2018). Our table is purely descriptive, and pursuing these channels further is beyond the scope of this paper.

The lower separation rate of treated workers therefore comes from moves to other employers. Treated incumbents in Sample A are 1.7 percentage points less likely to be observed at another firm, consistent with the retention estimates in Section 5.2 in the main text. The pattern is more pronounced for non-R&D workers (Sample B). Their shares of never observed again differ by only 0.6 percentage points (39.2% versus 39.8%), while moves to another firm are 3.2 percentage points less common among treated workers (14.5% versus 17.7%), similar in size to the retention effect of 3.47 percentage points in Figure 6 in the main text. Among R&D workers (Sample C), treated and control workers are equally likely to move to another firm (30.9% versus 31.3%), in line with the absence of a retention effect for this group.

Lastly, we check the characteristics of the firms to which incumbent workers move after leaving their incumbency employer. Table C.9 compares treated and matched control incumbents who move to another firm, with all outcomes measured at the worker's first switch. Each entry is the difference in means between treated and control switchers. Because treatment affects who switches, these are descriptive differences, not causal effects. Treated switchers are 16.0 percentage points more likely than control switchers to join a firm that already benefits from the innovation box in Sample A (19.0 in Sample B and 9.6 in Sample C), and correspondingly less likely to join a never-treated firm. This largely reflects where treated and control workers are employed. The wage change at the move is the log base wage in the first year at the destination firm minus the log base wage in the last year at the incumbency firm. Treated switchers gain 0.8 log points more than control switchers in Sample A and 1.1 log points more in Sample B, with no difference in Sample C (see Section 5.2 of the main text).

Table C.8: Leavers: Moves to Other Firms and Exits from the Data

	Treated (%)	Matched control (%)	Difference (pp)	(SE)
Sample A: Full sample				
Share of all incumbents				
Left the incumbency firm	59.5	61.6	-2.09^{***}	(0.11)
Moved to another firm	21.8	23.6	-1.73^{***}	(0.10)
Never observed again	31.2	31.7	-0.53^{***}	(0.11)
Neither category	6.5	6.3	0.17^{***}	(0.06)
Share of leavers				
Moved to another firm	36.7	38.2	-1.56^{***}	(0.15)
Never observed again	52.4	51.5	0.91^{***}	(0.15)
Neither category	10.9	10.3	0.65^{***}	(0.09)
Incumbents (leavers)	378,409 (225,285)	350,071 (215,722)		
Sample B: Non-R&D workers				
Share of all incumbents				
Left the incumbency firm	55.6	60.4	-4.78^{***}	(0.24)
Moved to another firm	14.5	17.7	-3.18^{***}	(0.18)
Never observed again	39.2	39.8	-0.62^{**}	(0.24)
Neither category	1.9	2.9	-0.98^{***}	(0.08)
Share of leavers				
Moved to another firm	26.1	29.3	-3.20^{***}	(0.29)
Never observed again	70.4	65.9	4.55^{***}	(0.30)
Neither category	3.5	4.8	-1.35^{***}	(0.13)
Incumbents (leavers)	89,632 (49,845)	75,400 (45,535)		
Sample C: R&D workers				
Share of all incumbents				
Left the incumbency firm	58.4	59.6	-1.21^{***}	(0.41)
Moved to another firm	30.9	31.3	-0.46	(0.39)
Never observed again	19.2	20.2	-1.04^{***}	(0.33)
Neither category	8.3	8.0	0.28	(0.23)
Share of leavers				
Moved to another firm	52.9	52.6	0.31	(0.54)
Never observed again	32.9	34.0	-1.07^{**}	(0.51)
Neither category	14.2	13.4	0.76^{**}	(0.37)
Incumbents (leavers)	39,385 (22,985)	22,420 (13,356)		

Notes: Leavers are incumbents who are no longer employed at their incumbency firm at the end of the panel in 2023. *Moved to another firm* means the leaver is next observed at a different employer. *Never observed again* means the leaver does not appear in the payslip data in any later year up to 2023. *Neither category* covers leavers who meet neither definition, workers with gaps in their employment. Shares of all incumbents use all treated or matched control incumbents as the denominator, and shares of leavers use leavers only. Differences are treated minus control, in percentage points. Heteroskedasticity-robust standard errors in parentheses, equivalent to a two-sample test of equal proportions. Statistical significance is indicated as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table C.9: Destination Outcomes of Incumbent Leavers: Treated vs. Matched Control Switchers

	Sample A: Full sample	Sample B: Non-R&D workers	Sample C: R&D workers
	(1)	(2)	(3)
Destination firm's treatment status			
Already-treated firm	0.160^{***}	0.190^{***}	0.096^{***}
	(0.002)	(0.004)	(0.006)
Not-yet-treated firm	0.007^{***}	0.019^{***}	-0.001
	(0.001)	(0.002)	(0.003)
Never-treated firm	-0.168^{***}	-0.209^{***}	-0.094^{***}
	(0.002)	(0.005)	(0.006)
Wage change at the move			
$\Delta \log(\text{wage})$	0.008^{***}	0.011^{***}	-0.002
	(0.001)	(0.002)	(0.003)
Observations	165,058	26,378	19,173

Notes: Each entry is the coefficient $\hat{\beta}$ from a separate OLS regression of the indicated outcome, measured at the worker's first switch, on a constant and a treatment dummy, estimated on incumbent switchers in the indicated panel. $\hat{\beta}$ equals the difference in means between treated and matched control switchers. Heteroskedasticity-robust standard errors in parentheses. Statistical significance is indicated as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

D. Robustness Appendix

This section reports the methodological checks discussed in Section 6 in the main text. Table D.1 collects the IW-ATTs for Sample A under each check, next to the baseline. We discuss each robustness check in detail below.

Table D.1: Robustness: Average Treatment Effects under Alternative Specifications

	Baseline	Reference	Incumbents from		Placebo	TWFE
		period $\ell = -1$	$t-2$	$t-1$	treatment	
	(1)	(2)	(3)	(4)	(5)	(6)
Base wage	3.27***	2.50***	3.05***	2.66***	0.24	3.85***
	(0.67)	(0.58)	(0.66)	(0.60)	(0.23)	(0.73)
Bonuses	0.39	0.47	0.53	0.52	0.23	0.17
	(0.53)	(0.53)	(0.51)	(0.48)	(0.17)	(0.69)
Overtime income	−0.06**	−0.01	−0.05**	−0.08***	0.04	−0.05*
	(0.03)	(0.04)	(0.03)	(0.03)	(0.03)	(0.03)
Base hours	0.58***	−0.31**	0.45***	0.11	0.37***	0.86***
	(0.18)	(0.14)	(0.16)	(0.15)	(0.12)	(0.32)
Overtime hours	−0.07**	−0.06*	−0.08**	−0.11***	0.04*	−0.10**
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.04)
Observations	7,239,116	7,239,116	8,329,950	9,400,837	3,422,354	7,239,116
Worker and year fixed effects			Yes			
Estimator	IW	IW	IW	IW	IW	TWFE

Notes: Average treatment effects on the treated for Sample A, in percent. Column (1) is the baseline specification (incumbents employed from $t-3$, reference period $\ell = 0$). Column (2) uses $\ell = -1$ as the reference period. Columns (3) and (4) define incumbents as workers employed at the firm from $t-2$ and $t-1$ onward. Column (5) estimates the model on matched control workers only, with randomly assigned placebo treatment dates. Columns (1)–(5) report the interaction-weighted ATT from estimations of equation (1) in the main text, and column (6) reports the static two-way fixed effects estimate (Appendix D.1). Income outcomes are expressed as a percentage of the Sample A control-group mean annual base wage (€42,850), and hours outcomes as a percentage of the control-group mean annual hours worked (1,871 hours in columns (1)–(5) and 1,889 hours in column (6)). Standard errors, clustered at the firm level, in parentheses. Statistical significance is indicated as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

D.1 Two-Way Fixed Effects

As an alternative to the interaction-weighted (IW) estimator, we estimate the effect on the Sample A wage panel using the two-way fixed effects (TWFE) specification employed in our descriptive mobility analysis. (equation (C.1)). Control workers inherit the first benefit year G_i of the treated worker they are matched to, so event time $\ell = t - G_i$ is defined for every worker. We estimate:

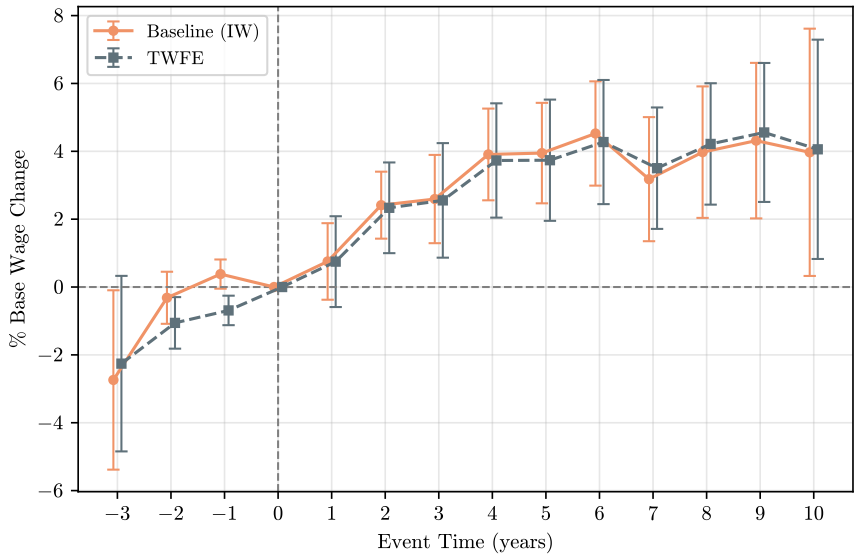
$$Y_{it} \;=\; \sum_{\substack{\ell=-3 \\ \ell\neq 0}}^{10} \beta_\ell \left(\text{Treated}_i \cdot \mathbb{1}\{t - G_i = \ell\}\right) \;+\; \alpha_i \;+\; \lambda_{G_i t} \;+\; \varepsilon_{it}, \tag{D.1}$$

where Treated_i equals one for treated incumbents, α_i are worker fixed effects, and $\lambda_{G_i t}$ are cohort-by-year fixed effects. For the pooled ATT, we replace the event-time interactions with a single interaction, $\text{Treated}_i \cdot \mathbb{1}\{t \geq G_i\}$. Standard errors are clustered at the firm level.

The cohort-by-year fixed effects compare each treated cohort only with its own matched controls, observed over the same calendar years. The specification therefore never uses earlier-treated workers as controls for later-treated ones, the comparisons that bias static TWFE estimates under staggered adoption (Goodman-Bacon, 2021). It differs from the IW estimator in two respects. First, the IW estimator identifies common year effects from all never-treated controls, whereas equation (D.1) allows a separate year profile for each cohort, estimated from that cohort's matched controls only. Second, the two pooled estimates weight cohorts and event times differently. The IW-ATT weights each cohort-by-event-time effect by its share of treated observations, while the pooled TWFE coefficient uses regression weights, which depend on cohort size and on the number of post-treatment years each cohort contributes.

Figure D.1 plots the event-study coefficients from equation (D.1) next to the baseline IW estimates. After treatment, the two sets of estimates are nearly identical. The effect on base wages is 2.33% under TWFE and 2.41% under IW at $\ell = 2$, 3.73% and 3.91% at $\ell = 4$, and 4.06% and 3.97% at $\ell = 10$. The pooled TWFE estimate is 3.85% with a standard error of 0.73, compared with an IW-ATT of 3.27% with a standard error of 0.67 (Table D.1, columns 1 and 6). Since the event-time coefficients barely differ, the gap between the pooled estimates reflects the different weights the two estimators place on event times rather than a different wage response. Before treatment, the TWFE coefficients at $\ell = -2$ and $\ell = -1$ are -1.06% and -0.69%. They are statistically significant but about a quarter of the size of the post-treatment effect from $\ell = 4$ onward, and the wage path breaks clearly after the first benefit year. This indicates that the effects are heterogeneous across cohorts and that the Sun and Abraham (2021) estimator should be our preferred estimator.

Figure D.1: Event Study on Base Wages: Two-Way Fixed Effects and Baseline Estimates



Notes: Event-study coefficients on base wages for Sample A from the baseline interaction-weighted estimator (equation (1) in the main text, solid orange circles) and from the two-way fixed effects specification in equation (D.1) (dashed grey squares), with 95% confidence intervals. Coefficients are expressed as a percentage of the Sample A control-group mean annual base wage (€42,850), and $\ell = 0$, the first benefit year, is the reference period. Standard errors are clustered at the firm level.

D.2 Cohort-by-Cohort Estimates

Our main specification pools all treatment cohorts $e \in \{2010, \ldots, 2018\}$ and recovers cohort-specific dynamic treatment effects $\hat{\delta}_{e,\ell}$ using the interaction-weighted (IW) estimator of Sun and Abraham (2021). As an additional check on these pooled estimates, we re-estimate the model separately for each treatment cohort e, using only that cohort's treated incumbent workers and their matched never-treated controls (Appendix B.4, Algorithm B.1). Each cohort-specific regression compares the cohort's treated incumbents only with their own matched controls, observed over the same calendar years. The cohort-by-cohort estimates therefore keep the within-cohort alignment that matching creates, in the spirit of stacked designs. The estimates are also free of comparisons between differently timed treatment cohorts. The support our choice of using IW-estimator.

For each cohort e, let $\mathcal{S}_e$ denote the matched estimation sample consisting of (i) treated incumbent workers i whose employer first receives the innovation box in year $G_i = e$, and (ii) their nearest-neighbour matched control workers. For worker $i \in \mathcal{S}_e$ observed in year t, we define event time relative to the cohort's benefit year as $\ell = t - e$ for treated workers, and assign control workers to the reference period, $\ell = 0$.

For each outcome Y_{it} and cohort e, we estimate:

$$Y_{it} = \sum_{\substack{\ell=-3 \\ \ell \neq 0}}^{10} \beta_{e,\ell} \mathbb{1}\{\ell_{it} = \ell\} + \alpha_i + \lambda_t + \varepsilon_{it}, \qquad i \in \mathcal{S}_e, \tag{D.2}$$

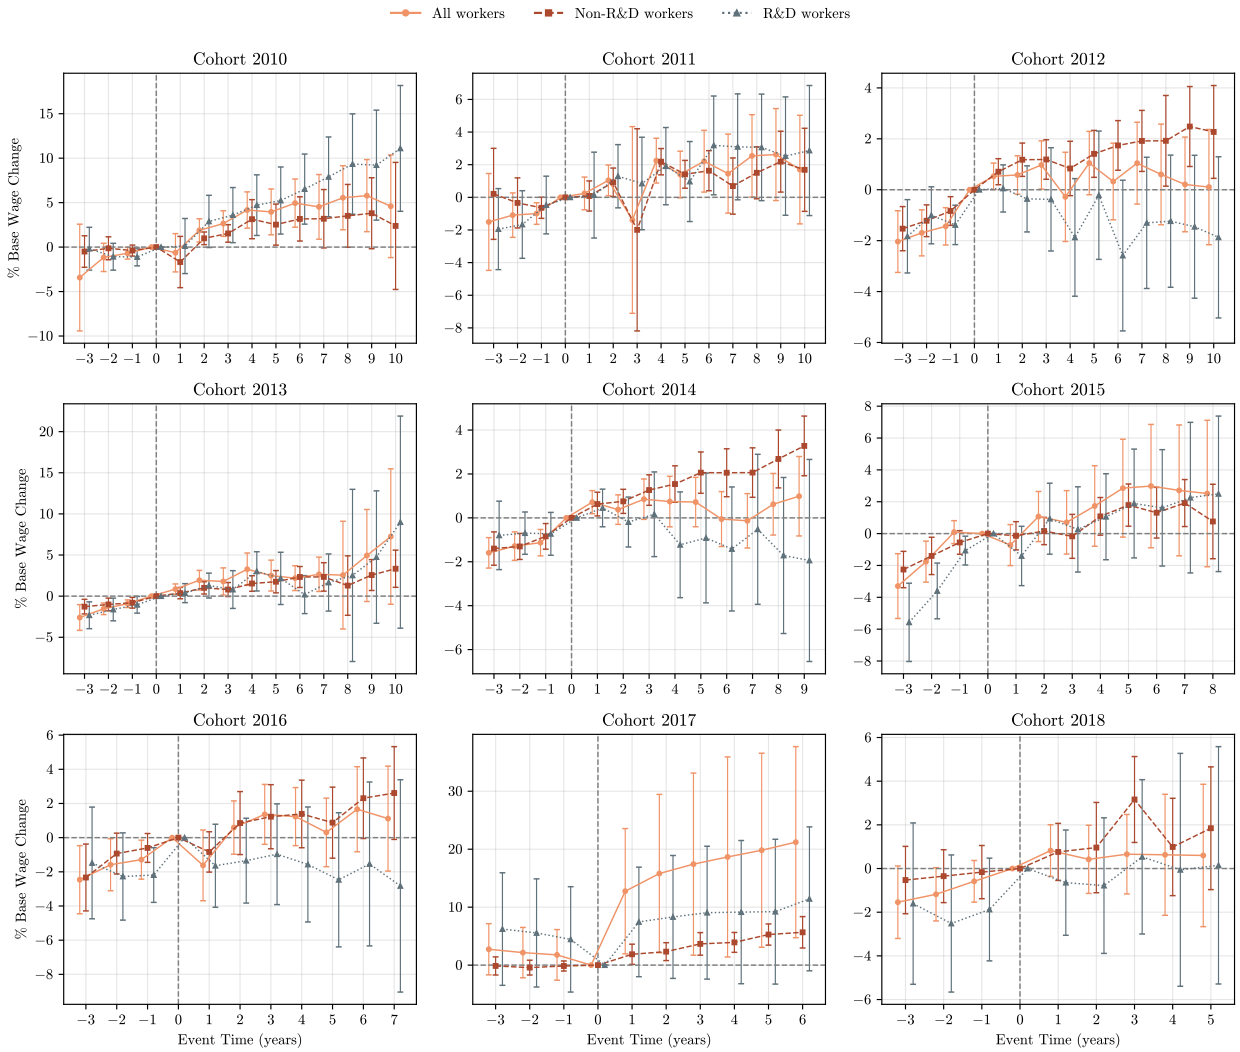
where α_i are worker fixed effects and λ_t are calendar-year fixed effects. $\ell = 0$, the first benefit year, is the omitted reference category. Worker fixed effects absorb all time-invariant differences between treated and matched control workers, including any residual selection on levels not fully addressed by matching. We estimate equation (D.2) separately for base wages and the probability of switching firms for each treatment cohort, and cluster standard errors at the firm level. Matched control workers inherit the event time of their treated matches, as in equation (C.1). Figure D.2 plots the cohort-specific coefficients on base wages for Samples A, B, and C. Figure D.3 plots effects only for large firms. Figure D.4 plots the effects on the probability of switching firms.

We first describe the estimates in Figure D.2 in detail. In Sample A, the pre-treatment coefficients of most cohorts are negative and rise towards zero before treatment (Figure D.2). The rise is steepest and most precisely estimated for the 2012–2015 cohorts, which drive the pre-treatment movement in the pooled estimates (Appendix C.1). Post-treatment effects vary across cohorts. They are largest and most precisely estimated for the 2010 cohort, small (about 1% or less) for the 2012, 2014, and 2018 cohorts, and large but imprecise for the 2017 cohort. Sample B shows the same pattern at a smaller scale. Sample C is the noisiest of the three (Figure D.2). Its effect comes mainly from the 2010 cohort, whose pre-treatment coefficients are small and whose post-treatment effects grow to about 9–11% between $\ell = 8$ and $\ell = 10$. Most later cohorts show small and insignificant effects, and the 2015 cohort shows a clear rising pre-trend.

The 2010 cohort coincides with the introduction of the policy. The innovation box took its current form in 2010, and a large stock of firms whose innovations were already profitable entered in that first year, whereas later inflows could increasingly reflect innovations that have only recently become profitable, or application lags. For the 2010 cohort, the timing of treatment is therefore set by the introduction of the policy rather than by the success of the firm's innovation. If incumbents

were sharing rents from innovations that predate the tax benefit, their wages would already be rising relative to matched controls before 2010. The cohort-by-cohort estimates show that they are not. For non-R&D workers in the 2010 cohort, the three pre-treatment coefficients lie between −0.5% and 0% and are not statistically different from zero. The average treatment effect is around 1.85%, which rests within the bounds of the effects of the main analysis.

Figure D.2: Base Wage: Cohort-by-Cohort Event Study

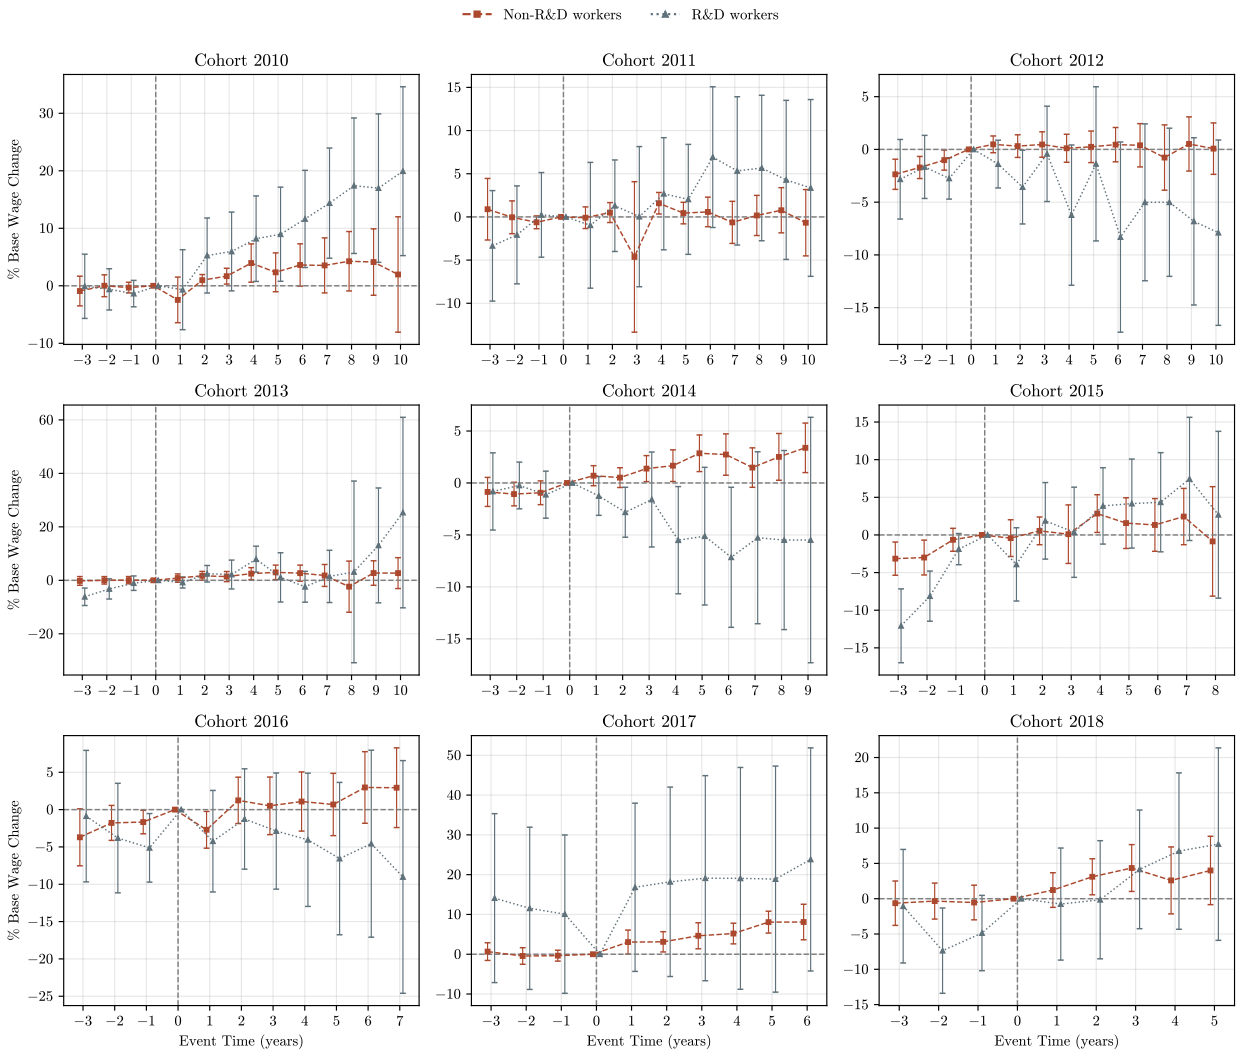


Notes: Each panel plots $\hat{\beta}_{e,\ell}$ from estimations of equation (D.2) on base wages for one treatment cohort e, with 95% confidence intervals, for Sample A (full sample), Sample B (non-R&D workers), and Sample C (R&D workers). Coefficients are expressed as a percentage of the control-group mean annual base wage of each sample. The vertical dashed line marks the cohort's own treatment year ($\ell = 0$), and the horizontal dashed line marks zero. Vertical axes differ across panels.

Figure D.3 restricts the sample to incumbents at large firms. This removes most of the pre-treatment movement. In Sample B, the 2010, 2011, 2013, 2014, 2017, and 2018 cohorts have flat pre-treatment coefficients, while the 2012 and 2015 cohorts keep a rising pre-trend, and the 2016 cohort a less precisely estimated one. Post-treatment effects remain positive in most cohorts. In Sample C, confidence intervals are wide throughout, reflecting the small number of tertiary-educated STEM incumbents at large firms in any single cohort. As in the full Sample A, the effect

comes mainly from the 2010 cohort. The cohort-by-cohort results for large companies corroborate our intuition that pre-trends reflect wage expansion paths that are different at small firms and that might not reflect rent-sharing though subsequent wage renegotiations.

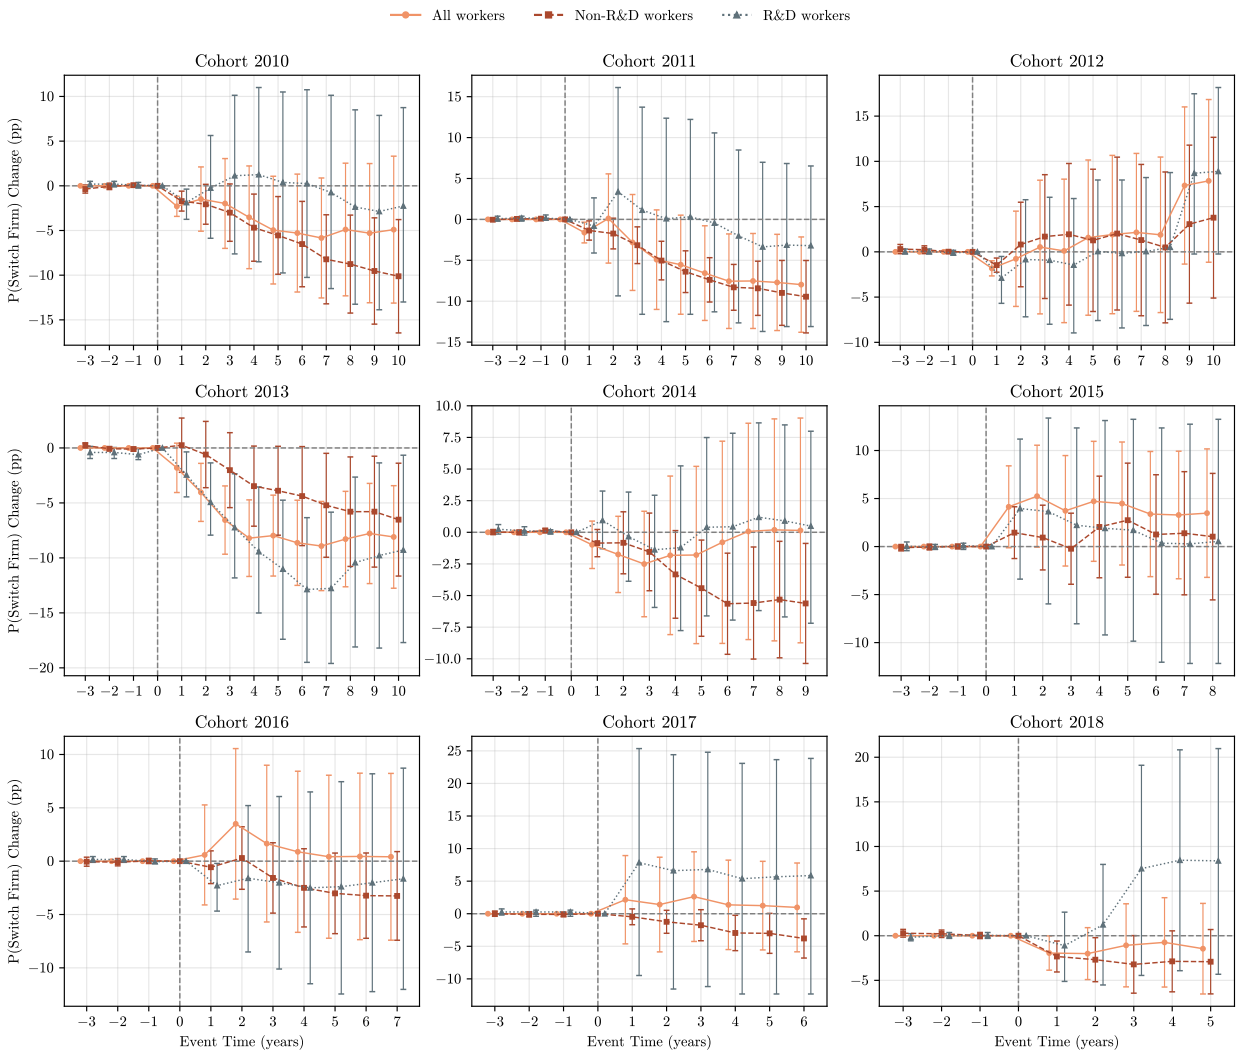
Figure D.3: Base Wage: Cohort-by-Cohort Event Study, Large Firms



Notes: Cohort-specific event-study coefficients on base wages, as in Figure D.2, for Samples B and C restricted to large firms (more than 250 employees at $t-3$).

Lastly, we inspect the probability of switching firms by cohort. In Sample B, the post-treatment switching coefficients are negative in seven of the nine cohorts (Figure D.4). The 2012 and 2015 cohorts are the exceptions, with positive but insignificant coefficients. In Sample A, the 2010, 2011, 2013, 2014, and 2018 cohorts show declining switching probabilities after treatment, the 2012 and 2015 cohorts show positive coefficients, and the 2016 and 2017 cohorts show small positive coefficients close to zero (Figure D.4). In Sample C, confidence intervals are wide, and only the 2013 cohort shows a large and significant decline (Figure D.4). The retention effect is therefore broad-based among non-R&D workers, in line with the pooled estimates in Section 5.2 in the main text.

Figure D.4: Probability of Switching Firms: Cohort-by-Cohort Event Study



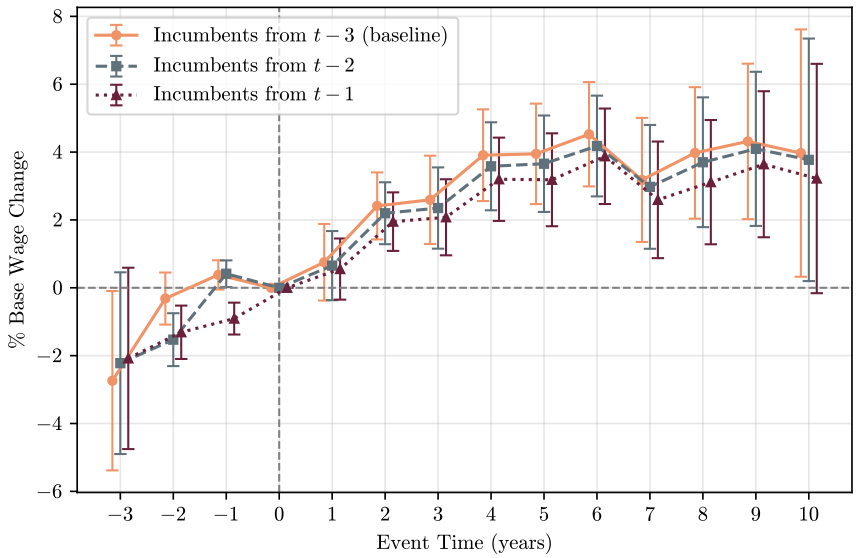
Notes: Each panel plots the cohort-specific coefficients on Switch_{it} (in percentage points) for one treatment cohort e, with 95% confidence intervals, for Sample A (full sample), Sample B (non-R&D workers), and Sample C (R&D workers). The vertical dashed line marks the cohort's own treatment year, and the horizontal dashed line marks zero. Vertical axes differ across panels.

D.3 Incumbent Definition

Our main analysis defines incumbents as workers employed at the firm from $t-3$ through t. Relaxing the definition to employment from $t-2$ or from $t-1$ onward yields supersets of the main sample, with 8.3 and 9.4 million worker-year observations instead of 7.2 million. Figure D.5 plots the base wage event studies under the three definitions, and Table D.2 reports the event-time coefficients and IW-ATTs for all outcomes. We rerun the matching for each new definition, so incumbents are matched on their characteristics at $t-2$ and $t-1$, respectively.

The base wage IW-ATT falls monotonically as the tenure requirement is relaxed, from 3.27% under the $t-3$ definition to 3.05% under $t-2$ and 2.66% under $t-1$ (Table D.1). This ordering is consistent with workers with longer tenure having accumulated more firm-specific knowledge and sharing more in the firm's rents. The base hours effect also falls, from 0.58% under $t-3$ to 0.45% under $t-2$ and 0.11% under $t-1$. The relaxed definitions introduce pre-trends. Under $t-2$, the coefficient at $\ell=-2$ is -1.53% and the coefficient at $\ell=-1$ is 0.41%, both significant. Under $t-1$, the coefficients at $\ell=-2$ and $\ell=-1$ are -1.31% and -0.91%, both significant at the 1% level, so treated wages rise relative to controls throughout the two years before treatment. These pre-trends are expected because the relaxed definitions admit workers hired shortly before the firm's treatment year, whose wage growth in their first years at the firm likely differs from that of longer-tenured workers, and whom our $t-3$ definition is designed to exclude. These results support our choice to define incumbents from $t-3$.

Figure D.5: Event-Study Estimates on Base Wage under Alternative Incumbent Definitions



Notes: Event-time coefficients on base wages from estimations of equation (1) in the main text for incumbents defined as workers employed at the firm from $t-3$ (baseline), $t-2$, and $t-1$ onward, with $\ell=0$ as the omitted reference period. The error bars denote 95% confidence intervals. Coefficients are expressed as a percentage of the Sample A control-group mean annual base wage (€42,850). The vertical dashed line marks the treatment year, and the horizontal dashed line marks zero.

Table D.2: Event-Study Estimates under Alternative Incumbent Definitions

	Base wage (1)	Bonuses (2)	Overtime income (3)	Base hours (4)	Overtime hours (5)
Incumbents from $t-2$					
IW-ATT	3.046*** (0.655)	0.525 (0.509)	−0.055** (0.027)	0.450*** (0.161)	−0.075** (0.032)
$t-3$	−2.222 (1.366)	−0.790 (0.556)	0.085* (0.048)	−1.433 (1.292)	0.142** (0.060)
$t-2$	−1.529*** (0.397)	−0.778** (0.382)	0.051* (0.027)	−0.624*** (0.223)	0.070** (0.029)
$t-1$	0.414** (0.201)	−0.118 (0.218)	−0.017 (0.029)	1.005*** (0.174)	0.017 (0.022)
$t+1$	0.653 (0.520)	−0.030 (0.290)	0.043** (0.022)	−0.807* (0.444)	0.011 (0.018)
$t+2$	2.197*** (0.466)	−0.011 (0.532)	0.008 (0.023)	0.376** (0.160)	−0.026 (0.024)
$t+3$	2.351*** (0.611)	0.816 (0.952)	0.013 (0.036)	0.306 (0.360)	−0.007 (0.052)
$t+4$	3.581*** (0.662)	−0.273 (0.647)	−0.007 (0.029)	0.603*** (0.197)	−0.040 (0.036)
$t+5$	3.657*** (0.725)	−0.109 (0.680)	−0.033 (0.032)	0.261 (0.226)	−0.066 (0.042)
$t+6$	4.179*** (0.757)	0.332 (0.670)	−0.113*** (0.043)	0.728*** (0.229)	−0.138*** (0.042)
$t+7$	2.975*** (0.931)	1.079 (0.695)	−0.141*** (0.050)	0.311 (0.408)	−0.156*** (0.047)
$t+8$	3.700*** (0.976)	0.578 (0.714)	−0.121*** (0.045)	0.790*** (0.261)	−0.154*** (0.047)
$t+9$	4.095*** (1.159)	0.488 (0.776)	−0.139*** (0.045)	1.130*** (0.244)	−0.167*** (0.052)
$t+10$	3.772** (1.824)	0.349 (0.801)	−0.176*** (0.068)	0.502 (0.572)	−0.167** (0.067)
Observations			8,329,950		
Incumbents from $t-1$					
IW-ATT	2.661*** (0.604)	0.518 (0.479)	−0.075*** (0.026)	0.114 (0.153)	−0.106*** (0.031)
$t-3$	−2.081 (1.364)	−0.615 (0.564)	0.078 (0.049)	−1.590 (1.293)	0.137** (0.061)
$t-2$	−1.311*** (0.401)	−0.603 (0.382)	0.041 (0.027)	−0.718*** (0.218)	0.059** (0.030)
$t-1$	−0.908*** (0.240)	−0.548** (0.213)	−0.033 (0.028)	−0.300 (0.205)	0.004 (0.022)
$t+1$	0.552 (0.460)	0.177 (0.284)	0.029 (0.021)	−0.856** (0.383)	0.000 (0.017)
$t+2$	1.950*** (0.440)	0.138 (0.505)	−0.002 (0.023)	0.124 (0.152)	−0.040* (0.023)
$t+3$	2.078*** (0.572)	0.854 (0.898)	−0.007 (0.033)	0.030 (0.329)	−0.033 (0.047)
$t+4$	3.199*** (0.626)	−0.158 (0.612)	−0.027 (0.028)	0.265 (0.194)	−0.068** (0.034)
$t+5$	3.184*** (0.698)	−0.087 (0.640)	−0.051 (0.031)	−0.183 (0.251)	−0.095** (0.039)
$t+6$	3.876*** (0.718)	0.335 (0.645)	−0.134*** (0.041)	0.358 (0.220)	−0.174*** (0.041)
$t+7$	2.592*** (0.877)	0.858 (0.660)	−0.166*** (0.048)	−0.146 (0.384)	−0.198*** (0.047)
$t+8$	3.114*** (0.934)	0.373 (0.669)	−0.153*** (0.043)	0.252 (0.292)	−0.204*** (0.047)
$t+9$	3.641*** (1.098)	0.395 (0.738)	−0.174*** (0.044)	0.721*** (0.243)	−0.220*** (0.052)
$t+10$	3.221* (1.725)	0.083 (0.756)	−0.220*** (0.065)	0.152 (0.540)	−0.232*** (0.067)
Observations			9,400,837		

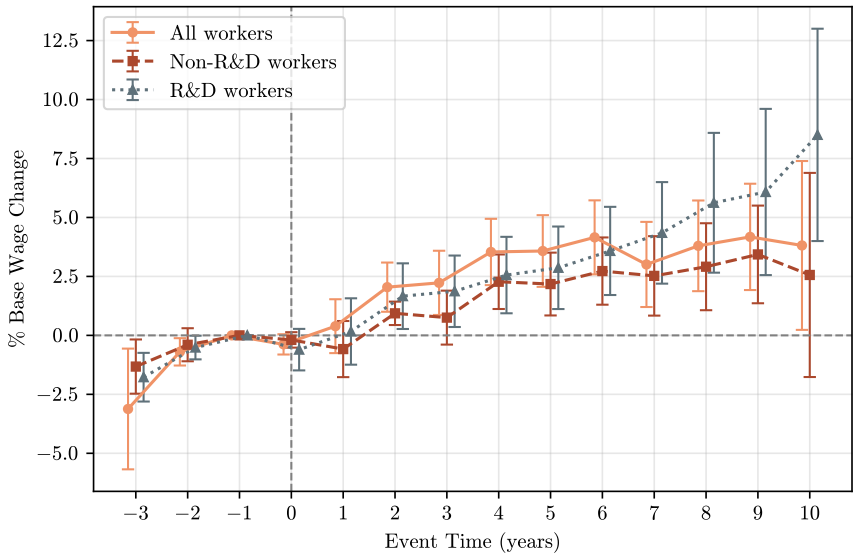
Notes: Each column within a panel reports results from a separate regression estimated with the interaction-weighted estimator of Sun and Abraham (2021), with $\ell = 0$ as the omitted reference period. Incumbents are workers employed at the firm from $t-2$ (upper panel) or $t-1$ (lower panel) onward. The shaded IW-ATT row is the interaction-weighted average treatment effect over all post-treatment periods. Income outcomes are expressed as a percentage of the control-group mean annual base wage (€42,850), hours outcomes as a percentage of the control-group mean annual hours worked (1,871 hours), both taken from Sample A. Standard errors, clustered at the firm level, in parentheses. Statistical significance is indicated as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

D.4 Reference Period at $\ell = -1$

Our main specification omits $\ell = 0$, the first benefit year, as the reference period, because the tax benefit accrues within the filing year and wages plausibly respond with a lag (see Section 4.1 in the main text). The conventional choice is $\ell = -1$. Figure D.6 and Table D.3 report the event study with $\ell = -1$ omitted instead. The change shifts every coefficient by about the $\ell = -1$ coefficient of the main specification (0.38% in Sample A) and adds $\ell = 0$ to the post-treatment average.

The IW-ATT falls from 3.27% to 2.50% in Sample A, from 2.05% to 1.61% in Sample B, and from 4.14% to 3.08% in Sample C, roughly a quarter smaller in each sample, and remains significant at the 1% level in all three samples. The dynamics are unchanged. The coefficient at $\ell = 0$ is small in all three samples (−0.38%, −0.19%, and −0.60%) and not significant at the 5% level, and the effect emerges at $\ell = 2$ and grows over the window. This result justifies our choice of $\ell = 0$ as the reference period.

Figure D.6: Event-Study Estimates on Base Wage with Reference Period $\ell = -1$



Notes: Event-time coefficients on base wages from estimations of equation (1) in the main text for Sample A (full sample), Sample B (non-R&D workers), and Sample C (R&D workers), with $\ell = -1$ as the omitted reference period. The error bars denote 95% confidence intervals. Coefficients are expressed as a percentage of each sample's control-group mean annual base wage (€42,850, €33,960, and €55,522). The vertical dashed line marks the treatment year, and the horizontal dashed line marks zero.

D.5 Placebo Treatment

Matching does not guarantee that our never-treated control pool is free of workers who resemble treated incumbents in unobserved ways that also correlate with wage growth. To test for this, we restrict the Sample A panel to matched control workers and assign each a placebo treatment date, drawn from the distribution of treatment dates in the main analysis. Because none of these workers is ever employed at a firm that receives the innovation box, the placebo treatment should produce no wage effect if the control pool is free of such contamination.

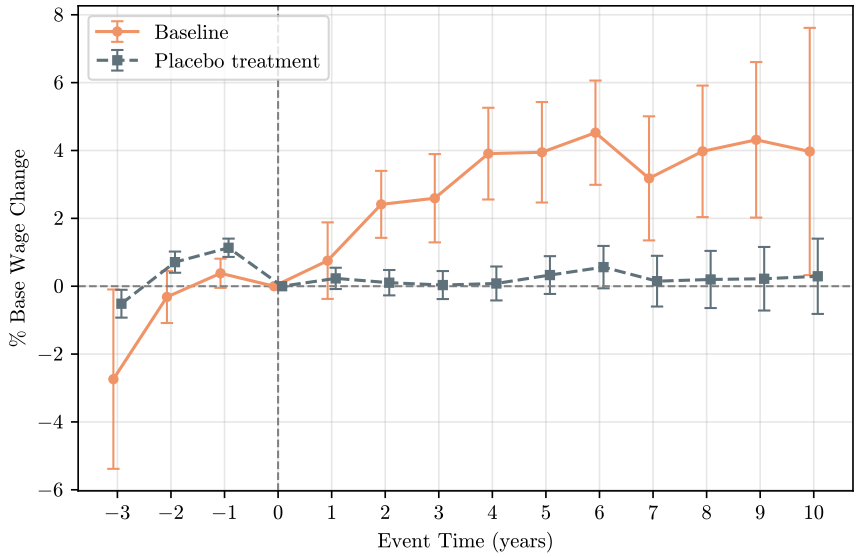
Table D.3: Event-Study Estimates and ATT by Sample, Reference Period $\ell = -1$

	Base wage (1)	Bonuses (2)	Overtime income (3)	Base hours (4)	Overtime hours (5)
Sample A: Full sample					
IW-ATT	2.497*** (0.579)	0.475 (0.534)	−0.013 (0.037)	−0.314** (0.144)	−0.060* (0.035)
$t-3$	−3.120** (1.306)	−0.796* (0.462)	0.065 (0.057)	−2.747** (1.276)	0.086 (0.064)
$t-2$	−0.698** (0.296)	−0.063 (0.319)	0.075* (0.040)	−0.359*** (0.093)	0.054 (0.033)
t	−0.381* (0.220)	0.166 (0.214)	0.029 (0.031)	−0.901*** (0.186)	−0.006 (0.023)
$t+1$	0.390 (0.583)	0.075 (0.375)	0.066 (0.043)	−1.699*** (0.465)	0.004 (0.032)
$t+2$	2.044*** (0.532)	−0.033 (0.623)	0.033 (0.041)	−0.442*** (0.135)	−0.036 (0.033)
$t+3$	2.227*** (0.694)	0.804 (1.055)	0.039 (0.053)	−0.502 (0.391)	−0.011 (0.062)
$t+4$	3.537*** (0.717)	−0.238 (0.736)	0.021 (0.047)	−0.127 (0.198)	−0.039 (0.044)
$t+5$	3.577*** (0.775)	−0.047 (0.761)	−0.003 (0.049)	−0.463* (0.238)	−0.065 (0.049)
$t+6$	4.156*** (0.801)	0.246 (0.762)	−0.081* (0.048)	0.053 (0.237)	−0.133*** (0.044)
$t+7$	3.007*** (0.921)	1.211 (0.778)	−0.102* (0.054)	−0.195 (0.438)	−0.146*** (0.053)
$t+8$	3.797*** (0.981)	0.641 (0.774)	−0.084 (0.056)	0.367 (0.276)	−0.143** (0.056)
$t+9$	4.176*** (1.150)	0.496 (0.810)	−0.101* (0.053)	0.723*** (0.264)	−0.154*** (0.056)
$t+10$	3.813** (1.826)	0.218 (0.876)	−0.128** (0.058)	0.087 (0.633)	−0.139** (0.061)
Observations			7,239,116		
Sample B: Non-R&D workers					
IW-ATT	1.608*** (0.533)	0.319 (0.311)	−0.048 (0.056)	−0.320** (0.144)	−0.074 (0.048)
$t-3$	−1.323** (0.587)	−0.106 (0.363)	0.101 (0.087)	−0.875** (0.346)	0.098 (0.089)
$t-2$	−0.397 (0.357)	0.296 (0.470)	0.064 (0.048)	−0.165 (0.130)	0.037 (0.042)
t	−0.192 (0.167)	0.024 (0.193)	0.013 (0.042)	−0.598*** (0.126)	−0.036 (0.030)
$t+1$	−0.581 (0.607)	−0.031 (0.236)	0.057 (0.056)	−1.463** (0.582)	−0.038 (0.042)
$t+2$	0.935*** (0.253)	0.324 (0.256)	−0.010 (0.058)	−0.287* (0.155)	−0.091* (0.047)
$t+3$	0.753 (0.583)	0.342 (0.255)	−0.034 (0.072)	−0.578 (0.459)	−0.080 (0.071)
$t+4$	2.274*** (0.590)	−0.117 (0.487)	−0.065 (0.067)	−0.018 (0.161)	−0.108* (0.060)
$t+5$	2.177*** (0.679)	0.101 (0.567)	−0.124* (0.071)	−0.272 (0.213)	−0.136** (0.063)
$t+6$	2.727*** (0.727)	0.322 (0.390)	−0.192** (0.075)	−0.058 (0.322)	−0.188*** (0.063)
$t+7$	2.522*** (0.859)	0.495 (0.610)	−0.172** (0.079)	−0.236 (0.320)	−0.163** (0.069)
$t+8$	2.911*** (0.941)	0.449 (0.656)	−0.182** (0.085)	0.129 (0.219)	−0.157** (0.074)
$t+9$	3.432*** (1.056)	0.172 (0.691)	−0.147* (0.081)	0.364* (0.210)	−0.126* (0.075)
$t+10$	2.562 (2.208)	0.112 (0.715)	0.020 (0.090)	−0.777 (0.774)	0.034 (0.086)
Observations			1,734,690		
Sample C: R&D workers					
IW-ATT	3.083*** (0.933)	1.437 (0.955)	−0.043 (0.043)	−0.635*** (0.214)	−0.082** (0.035)
$t-3$	−1.773*** (0.526)	−0.549 (0.373)	0.046 (0.030)	−1.773*** (0.516)	0.043 (0.030)
$t-2$	−0.516** (0.256)	−0.162 (0.273)	0.082*** (0.031)	−0.365*** (0.135)	0.069*** (0.026)
t	−0.605 (0.449)	−0.003 (0.308)	−0.022 (0.027)	−1.370*** (0.515)	−0.027 (0.024)
$t+1$	0.165 (0.718)	0.273 (0.720)	−0.014 (0.045)	−2.120*** (0.550)	−0.052* (0.028)
$t+2$	1.663** (0.710)	0.263 (1.075)	−0.022 (0.040)	−1.031*** (0.222)	−0.065** (0.029)
$t+3$	1.870** (0.773)	0.925 (1.094)	−0.042 (0.045)	−0.846*** (0.244)	−0.062 (0.043)
$t+4$	2.557*** (0.827)	0.792 (0.921)	−0.035 (0.045)	−0.675*** (0.223)	−0.079** (0.036)
$t+5$	2.866*** (0.893)	1.664 (1.174)	−0.028 (0.050)	−0.812*** (0.259)	−0.082** (0.041)
$t+6$	3.581*** (0.954)	2.216* (1.195)	−0.014 (0.050)	−0.373 (0.255)	−0.087** (0.039)
$t+7$	4.346*** (1.098)	3.021** (1.478)	−0.051 (0.058)	0.297 (0.323)	−0.094* (0.052)
$t+8$	5.622*** (1.512)	1.953 (1.390)	−0.029 (0.062)	0.204 (0.314)	−0.093* (0.053)
$t+9$	6.080*** (1.798)	2.127* (1.237)	−0.031 (0.069)	0.186 (0.301)	−0.104* (0.059)
$t+10$	8.501*** (2.297)	2.501* (1.433)	−0.149* (0.083)	0.649* (0.341)	−0.183*** (0.070)
Observations			606,033		

Notes: Same specification as Table C.1, with $\ell = -1$ instead of $\ell = 0$ as the omitted reference period. The shaded IW-ATT row is the interaction-weighted average treatment effect over all post-treatment periods, including $\ell = 0$. Income outcomes are expressed as a percentage of the sample's control-group mean annual base wage (€42,850 in Sample A, €33,960 in Sample B, €55,522 in Sample C), and hours outcomes as a percentage of the sample's control-group mean annual hours worked (1,871, 1,952, and 1,935 hours in Samples A, B and C, respectively). Standard errors, clustered at the firm level, in parentheses. Statistical significance is indicated as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The placebo IW-ATT on base wages is 0.24%, about a fourteenth of the baseline estimate, and no post-treatment coefficient is significant at the 5% level (Figure D.7 and Table D.4). The control group therefore shows none of the post-treatment wage growth we find for treated incumbents. Two further features of the placebo estimates are worth noting. The pre-treatment coefficients at $\ell = -2$ and $\ell = -1$ are positive and significant (0.71% and 1.13%), opposite in sign to the $\ell = -3$ coefficient of the main specification. And the placebo ATT on base hours is 0.37%, close to the baseline estimate of 0.58%. The latter is a further reason to interpret the base hours estimates in Table 2 in the main text with caution. Part of the base hours response is also present among workers who never receive treatment. This confirms that the effect does not influence base and overtime hours.

Figure D.7: Event-Study Estimates on Base Wage: Placebo Treatment



Notes: Event-time coefficients on base wages from estimations of equation (1) in the main text, with $\ell = 0$ as the omitted reference period. The error bars denote 95% confidence intervals. The baseline series is the Sample A estimate. The placebo series is estimated on matched control workers only, each assigned a placebo treatment date drawn from the distribution of treatment dates in the main analysis. Coefficients are expressed as a percentage of the Sample A control-group mean annual base wage (€42,850). The vertical dashed line marks the (placebo) treatment year, and the horizontal dashed line marks zero.

D.6 No Education Imputation

Section 6.1 in the main text reports that dropping all workers whose education and field of study are not directly observed, and re-running the full matching and estimation procedure on this strictly observed subsample, leaves the average treatment effect on base wages essentially unchanged at 3.32%, against 3.27% in the main specification. Table D.5 reports the full event-study results. The observed-education subsample contains 4.1 million worker-year observations, 57% of the main panel, and over-represents highly educated workers (Appendix B.2).

The dynamics match those of the main specification. The effect emerges after the first benefit year and lies around 4% from $\ell = 4$ onward. Every post-treatment coefficient from $\ell = 2$ to $\ell = 9$ is significant at the 5% level. The pre-treatment coefficients show the same pattern as in Sample

Table D.4: Event-Study Estimates: Placebo Treatment

	Base wage (1)	Bonuses (2)	Overtime income (3)	Base hours (4)	Overtime hours (5)
Placebo treatment (matched control workers only)					
IW-ATT	0.236 (0.225)	0.228 (0.171)	0.035 (0.028)	0.367*** (0.122)	0.044* (0.025)
$t-3$	−0.516** (0.209)	−0.298* (0.179)	0.025 (0.033)	−0.172 (0.151)	0.023 (0.028)
$t-2$	0.709*** (0.159)	−0.045 (0.171)	0.036 (0.027)	0.932*** (0.120)	0.051** (0.025)
$t-1$	1.133*** (0.138)	0.062 (0.142)	0.035 (0.023)	1.096*** (0.111)	0.030 (0.020)
$t+1$	0.231 (0.160)	0.063 (0.161)	0.036 (0.024)	0.124 (0.129)	0.041* (0.021)
$t+2$	0.104 (0.191)	−0.037 (0.188)	0.046 (0.029)	0.021 (0.140)	0.057** (0.026)
$t+3$	0.034 (0.211)	0.210 (0.184)	0.068** (0.035)	−0.033 (0.150)	0.058* (0.030)
$t+4$	0.081 (0.255)	0.625** (0.273)	0.027 (0.036)	0.047 (0.166)	0.036 (0.032)
$t+5$	0.329 (0.284)	0.023 (0.217)	0.021 (0.041)	0.296* (0.165)	0.027 (0.036)
$t+6$	0.563* (0.319)	0.087 (0.218)	0.029 (0.042)	0.625*** (0.176)	0.046 (0.039)
$t+7$	0.149 (0.381)	0.338 (0.264)	0.010 (0.047)	0.368* (0.211)	0.012 (0.041)
$t+8$	0.199 (0.430)	0.582 (0.430)	0.065 (0.054)	0.743*** (0.225)	0.067 (0.046)
$t+9$	0.220 (0.478)	0.972** (0.448)	0.049 (0.057)	0.905*** (0.239)	0.059 (0.049)
$t+10$	0.293 (0.566)	0.500 (0.584)	0.117* (0.071)	0.884*** (0.271)	0.089 (0.056)
Observations			3,422,354		

Notes: Each column reports results from a separate regression estimated using the interaction-weighted estimator of Sun and Abraham (2021) on matched control workers only, with randomly assigned placebo treatment dates and $\ell = 0$ as the omitted reference period. The shaded IW-ATT row is the interaction-weighted average treatment effect over all post-treatment periods. Income outcomes are expressed as a percentage of the control-group mean annual base wage (€42,850), hours outcomes as a percentage of the control-group mean annual hours worked (1,871 hours), both taken from Sample A. Standard errors, clustered at the firm level, in parentheses. Statistical significance is indicated as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

A, with −2.82% at $\ell = -3$, which is not significant, and small coefficients at $\ell = -2$ and $\ell = -1$. The other outcomes also line up with the main estimates. Overtime income and overtime hours fall slightly, and the effects on base hours and bonuses are not significant.

Table D.5: Event-Study Estimates without Education Imputation

	Base wage	Overtime income	Bonuses	Base hours	Overtime hours
IW-ATT	3.32***	-0.07**	0.55	0.35*	-0.09***
	(0.79)	(0.03)	(0.53)	(0.21)	(0.03)
$t-3$	-2.82*	0.08**	-1.19**	-1.75	0.13***
	(1.48)	(0.04)	(0.59)	(1.59)	(0.05)
$t-2$	-0.37	0.08***	-0.42	0.73***	0.10***
	(0.37)	(0.03)	(0.35)	(0.24)	(0.03)
$t-1$	0.42*	-0.00	-0.26	1.08***	0.02
	(0.24)	(0.03)	(0.22)	(0.23)	(0.02)
t		Reference period (omitted)			
$t+1$	1.02*	0.02	-0.07	-0.72*	0.00
	(0.52)	(0.02)	(0.32)	(0.41)	(0.02)
$t+2$	2.49***	-0.01	-0.33	0.31	-0.04
	(0.56)	(0.03)	(0.64)	(0.21)	(0.03)
$t+3$	2.66***	-0.02	0.92	0.18	-0.04
	(0.75)	(0.03)	(1.21)	(0.44)	(0.04)
$t+4$	3.85***	-0.02	-0.29	0.51**	-0.06*
	(0.80)	(0.03)	(0.74)	(0.25)	(0.03)
$t+5$	3.93***	-0.05	-0.09	0.17	-0.09**
	(0.88)	(0.03)	(0.71)	(0.26)	(0.04)
$t+6$	4.49***	-0.10***	0.18	0.68***	-0.14***
	(0.95)	(0.04)	(0.76)	(0.26)	(0.04)
$t+7$	2.77**	-0.14***	1.36*	0.05	-0.17***
	(1.16)	(0.04)	(0.70)	(0.53)	(0.04)
$t+8$	3.87***	-0.12***	0.85	0.72**	-0.17***
	(1.19)	(0.04)	(0.68)	(0.32)	(0.05)
$t+9$	4.19***	-0.15***	0.75	1.01***	-0.19***
	(1.45)	(0.05)	(0.68)	(0.32)	(0.05)
$t+10$	4.07*	-0.21***	0.59	0.42	-0.21***
	(2.24)	(0.07)	(0.76)	(0.68)	(0.07)
Observations	4,145,688	4,145,688	4,145,688	4,145,688	4,145,688
R^2	0.8932	0.6136	0.5505	0.6351	0.6492
F-statistic	3807.28	723.25	557.90	792.69	842.93
RMSE	6373.91	908.58	7335.14	169.66	31.98
Adjusted R^2	0.8806	0.5681	0.4977	0.5922	0.6079
Within R^2	0.0145	0.0009	0.0040	0.0061	0.0016

Notes: Event-time coefficients and IW-ATT from estimations of equation (1) in the main text after after excluding workers whose education or field of study is not directly observed and re-running the matching algorithm, with $\ell = 0$ as the omitted reference period. Income outcomes (base wage, overtime income, bonuses) are expressed as a percentage of the control group's mean annual base wage (€43,961), and hours outcomes (base hours, overtime hours) as a percentage of the control group's mean annual hours worked (1,869 hours). Standard errors, clustered at the firm level, in parentheses. Statistical significance is indicated as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

References

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.

Bijlsma, M., Brouwer, E., de Jong, G., Busschots, T., den Hertog, P., Bongers, F., Smeitink, A., and Massop, M. (2023). Evaluatie van de innovatiebox 2010-2019: Doelgroepbereik, doeltreffendheid en doelmatigheid. Report 2023-125, SEO Economisch Onderzoek and Dialogic, Amsterdam / Utrecht. Commissioned by the Dutch Ministry of Economic Affairs and Climate Policy and the Ministry of Finance.

Britto, D. G., Pinotti, P., and Sampaio, B. (2022). The effect of job loss and unemployment insurance on crime in Brazil. *Econometrica*, 90(4):1393–1423.

Centraal Bureau voor de Statistiek (2020). *Standaard Onderwijsindeling 2016. Editie 2019/'20*. Centraal Bureau voor de Statistiek (CBS), Den Haag/Heerlen.

Dobbelaere, S., McCormack, G., Prinz, D., and Sóvágó, S. (2025). Firm consolidation and labor market outcomes. *Journal of Economic Behavior & Organization*, 235:107036.

Dos Santos, A. G. F. V., Marie, O., and Markiewicz, A. (2026). The unequal job security scars of displacement. *Journal of Public Economics*, 255:105562.

Frimmel, W., Horvath, T., Schnalzenberger, M., and Winter-Ebmer, R. (2018). Seniority wages and the role of firms in retirement. *Journal of Public Economics*, 164:19–32.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.

Hampton, M. and Totty, E. (2023). Minimum wages, retirement timing, and labor supply. *Journal of Public Economics*, 224:104924.

Jarosch, G. (2023). Searching for job security and the consequences of job loss. *Econometrica*, 91(3):903–942.

Mohnen, P., Vankan, A., and Verspagen, B. (2017). Evaluating the innovation box tax policy instrument in the Netherlands, 2007-13. *Oxford Review of Economic Policy*, 33(1):141–156.

Statistics Netherlands (2024a). Gbapersoontab - persoonskenmerken uit de basisregistratie personen. The Hague. Microdata accessed via remote access.

Statistics Netherlands (2024b). Hoogsteopltab - hoogst behaald en hoogst gevolgd opleidingsniveau en -richting van personen. The Hague. Microdata accessed via remote access.

Statistics Netherlands (2024c). Spolis - banen van werknemers, jaarbestand. The Hague. Microdata accessed via remote access.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.