

ERC Working Papers in Economics 26/04

September / 2026

PATENTS, PAY, AND PEOPLE: HOW INNOVATION REHAPES THE WORKFORCE

Çiğdem Ekiz

UNU-MERIT

Maastricht University, Maastricht, The Netherlands

E-mail: ekiz@merit.unu.edu

Eren Gürer

Department of Economics

Middle East Technical University, Ankara, Türkiye

E-mail: egurer@metu.edu.tr

Erol Taymaz

Department of Economics

Middle East Technical University, Ankara, Türkiye

E-mail: etaymaz@metu.edu.tr

Patents, Pay and People: How Innovation Reshapes the Workforce^{*}

Çiğdem Ekiz[†] Eren Güler[‡] Erol Taymaz[§]

September 11, 2026

Abstract

We examine how innovation shapes firm-level workforce dynamics using rich administrative data from Türkiye, a developing economy characterized by medium-technology, incremental innovation. We find that patent-based innovation increases both firm size and average daily wages, with the largest employment gains among workers earning just above the median wage. To uncover the mechanisms behind these outcomes, we construct worker transition matrices. Relative to comparable non-innovative firms, innovative firms (i) retain incumbent workers at higher rates, particularly at the top of the wage distribution, (ii) promote retained workers into higher wage bins more frequently, and (iii) hire more new workers, disproportionately into higher wage bins. Finally, occupational composition analysis shows that employment growth is concentrated among technical and production occupations, consistent with the medium-technology, incremental nature of innovation in our sample.

Keywords: innovation, patents, wages, employment, worker transitions

JEL Classification: J21, J31, O31

^{*}We are grateful to the Republic of Türkiye Ministry of Industry and Technology for granting access to the Entrepreneurship Information System.

[†]Maastricht University, UNU-MERIT. E-Mail: ekiz@merit.unu.edu. ORCID ID: 0000-0001-8778-3447

[‡]Middle East Technical University. E-Mail: egurer@metu.edu.tr. ORCID ID: 0000-0001-8238-1967

[§]Middle East Technical University. E-Mail: etaymaz@metu.edu.tr. ORCID ID: 0000-0001-7525-6674

1 Introduction

A substantial body of empirical work examines the firm-level implications of innovative activity. Evidence on the employment effects of innovation is mixed: some studies find that innovation is associated with employment declines at the firm level (Graetz and Michaels, 2018; Acemoglu and Restrepo, 2020), while others find it to be associated with employment growth (Van Reenen, 1997; Van Roy et al., 2018; Domini et al., 2021). Nonetheless, the empirical literature generally finds that innovation is associated with higher average wages. See, for example, Van Reenen (1996), Kline et al. (2019) among many others.

The specific mechanisms underlying rising average wages at the firm-level are less well understood. If innovation reduces employment at the firm level, for instance, by automating certain tasks, innovative firms may lay off low-wage and even some medium- and high-wage workers, raising average wages among those retained (Graetz and Michaels, 2018; Acemoglu and Restrepo, 2020). Alternatively, if the employment effects of innovation are skill-biased, innovative firms may lay off low-skilled workers while hiring high-skilled ones, again pushing average wages upward (Machin and Van Reenen, 1998).¹

If innovation does not reduce employment at the firm level, several other possibilities arise. Innovation may raise the wages of all incumbent workers through rent-sharing, as firms distribute the surplus it generates across their workforce. These gains, however, need not be evenly distributed, and may instead accrue only to high-skilled incumbents (Kline et al., 2019; Alam and Diegmann, 2026). Alternatively, it may lead firms to hire more high-wage workers (Domini et al., 2022), or even to replace their existing workforce with a new, more highly skilled one. In each of these cases, innovation again leads to higher average wages at the firm level.

Our main contribution in this paper is to provide, to the best of our knowledge, the first evidence on these potential mechanisms using administrative matched employer–employee data from a developing country where innovation is dominated by incremental, medium-technology activity, Türkiye. This setting lets us examine how firm-level innovation translates into workforce dynamics in a context distinct from the advanced, high-technology innovation (e.g., automation, robotization) that most existing studies focus on.

We draw on the Entrepreneurship Information System (EIS), a comprehensive administrative matched employer–employee dataset for Türkiye covering 2006–2023. EIS links worker registers, firm registers, balance sheet data, and patent records through unique firm identifiers. We proxy innovation using first-time patent applications, construct treatment

¹A related possibility is job polarization, whereby innovation raises employment at both the high- and low-skill ends of the distribution at the expense of middle-skill jobs (Acemoglu and Autor, 2011). Even under polarization, average wages can still rise if the growth in high-skill employment dominates the growth in low-skill employment.

and control groups through a matching strategy, and estimate causal effects using a stacked difference-in-differences event-study design.

Our results show that innovation raises both firm size and average daily wages at the firm-level. To examine how these employment gains are distributed across the wage distribution, we define seven wage bins relative to the statutory minimum wage. Employment rises in nearly every bin, with the smallest (and only marginally significant) increase in the lowest wage bin, while the largest increase occurs in Bin 3, which comprises workers earning between 1.5 and 2 times the minimum wage and places most of them just above the median wage.

To uncover the mechanisms behind these patterns, we construct worker transition matrices separately for incumbent workers and new hires. Tracking incumbent workers at treated (innovative) and control (comparable non-innovative) firms from three years before the patent application, we find that treated firms retain incumbents at higher rates than control firms at every point of the wage distribution, with the treated–control retention gap largest at the top. Among retained workers, those at treated firms are more likely to move up to higher wage bins across most of the distribution. Treated firms also hire considerably more new workers by the end of the post-treatment period, with new hiring concentrated in middle-to-upper wage bins.

Finally, turning to occupational composition, we find that employment growth is especially pronounced among technical and associate-professional staff, together with production occupations such as machinists, welders, toolmakers, and other metal and machinery tradesmen, as well as assembly-line and machine operators, a pattern consistent with the incremental, production-oriented character of patenting activity in our sample.

Related literature. A large literature examines the firm-level implications of innovative activity, with findings that vary across contexts. Some empirical studies find that innovation increases employment (Van Reenen, 1997; Lachenmaier and Rottmann, 2011; Castillo et al., 2014; Meschi et al., 2016; Van Roy et al., 2018; Domini et al., 2021; Damioli et al., 2024), while others find the opposite (Acemoglu and Restrepo, 2020; Bordot, 2022). By contrast, most studies examining the implications of innovation for wages document a positive association (Van Reenen, 1996; Toivanen and Väänänen, 2012; Castillo et al., 2014; Cirillo, 2014; Aghion et al., 2018; Bogliacino et al., 2018; Aghion et al., 2019; Cirera and Martins-Neto, 2023).

Most closely related to our study are papers that use integrated employer–employee data to track workers over time within innovating firms, inferring patterns of workforce dynamics from their trajectories. We review this literature below, distinguishing between studies of automation and studies of patent-based innovation.

A first strand studies automation and robotics, using matched employer–employee data

to document how technology adoption reshapes workforce composition. Bessen et al. (2025), Humlum (2019), and Dauth et al. (2021) track individual workers at automating firms using Dutch, Danish, and German administrative data, respectively, and find that incumbents face elevated separation risk or transition into new tasks, while new hiring shifts toward higher-skill profiles. Domini et al. (2021) and Domini et al. (2022), using French firm-level data, cannot track individual workers directly but instead observe “jobs”. Based on these job-level data, Domini et al. (2021) infer hiring and separation patterns around spikes in automation investment and find lower separation risk for incumbents alongside higher hiring rates. Domini et al. (2022) attribute wage growth due to investment in automation and AI-related capital mainly to the hiring of new, high-wage employees.

A second strand focuses on patent-based innovation. Kline et al. (2019) exploit variation in initial USPTO patent decisions, and find that patent allowances generate substantial firm-level surplus, a considerable share of which is passed through to incumbent workers as higher earnings. These gains are concentrated among long-tenure workers and in the upper half of the within-firm earnings distribution, and are accompanied by higher retention rates among these workers. Also focusing on the US, Kogan et al. (2023) use patent texts to construct firm-level measures of labor-saving and labor-augmenting technological exposure, and find that while labor-saving technologies reduce the earnings of incumbent workers, labor-augmenting technologies raise employment and wage bills, with newly hired workers earning more than the incumbents they replace.²

Taken together, the existing evidence on the specific mechanisms through which innovation shapes workforce dynamics is drawn almost entirely from developed economies (Kline et al., 2019; Humlum, 2019; Dauth et al., 2021; Domini et al., 2021, 2022; Kogan et al., 2023; Bessen et al., 2025). This paper contributes to this literature by examining a developing economy where innovation is predominantly medium-technology, incremental, and production-oriented: Türkiye.

Our paper also relates to a broader literature that examines how innovation reshapes firm-level labor outcomes without tracking individual workers’ transitions directly. Rather than following the same worker across firms and wage bins, these studies generally infer distributional consequences from aggregate measures, such as skill-specific wage premia, within-firm wage inequality, or the share of employment accounted for by high-skilled workers. This evidence is complementary to ours and we review some of such studies below.

Van Reenen (1996), Chennells and Van Reenen (1999), and Pianta and Tancioni (2008), using firm- and industry-level data, show that innovation raises wages on average but

²Also linking matched employer–employee records to patent data, Grinza and Quatraro (2019) study the reverse channel, i.e., how worker replacements affect firms’ patent activity, and find, using Italian longitudinal data, that turnover among incumbent workers significantly dampens innovation performance.

concentrates gains among skilled workers. Aghion et al. (2018) document how patent-based innovation gains are distributed among inventors, entrepreneurs, and blue- and white-collar workers using Finnish data. Alam and Diegmann (2026), using German employer–employee records and quasi-random patent examiner assignment, show that managers in publicly traded firms disproportionately benefit, in terms of wages, from patent-induced shocks.

A related body of work shows that innovation reshapes workforce composition along skill lines. Balasubramanian and Sivadasan (2011), linking US census and patent data, find that patenting firms shift toward higher skill intensity. Harrison et al. (2014) show, at the firm level, that product innovation creates jobs while process innovation displaces them, with effects that vary by skill. Meschi et al. (2016) document similar patterns in Turkish manufacturing, where skill-biased technological change widened both employment and wage gaps between skilled and unskilled workers. By contrast, Domini et al. (2022) find no evidence that investment in automation and AI-related capital increases within-firm wage or gender wage inequality.

The remainder of the paper proceeds as follows. Section 2 describes the data. Section 3 defines the treatment and control groups and introduces the details of the matching procedure. Section 4 presents the econometric specification and examines the effects of innovation on wages and employment across the wage distribution. Section 5 analyzes worker transitions for incumbent workers and new hires. Section 6 examines occupational composition, and Section 7 concludes.

2 Data

Our primary data source is the Entrepreneurship Information System (EIS), a comprehensive administrative matched employer–employee dataset covering 2006–2023, maintained by the Republic of Türkiye Ministry of Industry and Technology and accessible to researchers through its secure data center. EIS brings together administrative sub-datasets compiled by several government bodies, including the Turkish Statistical Institute, the Ministry of Treasury and Finance, the Social Security Institution, the Ministry of Trade, and the Turkish Patent and Trademark Office, spanning firm registers, worker registers, buyer-supplier transactions, balance sheet statements, export-import records, and patent data, all linkable through unique firm identifiers. Given its administrative nature and comprehensive coverage, EIS represents the most reliable available source for studying labor market dynamics in Türkiye.

For this study, we draw on four of these sub-datasets: worker registers from the Social Security Institution, firm registers from the Turkish Statistical Institute, balance sheet data from the Ministry of Treasury and Finance, and patent data from the Turkish Patent and

Trademark Office. The worker registers cover all formally employed workers in Türkiye over the period 2006–2023, though unique worker identifiers are available only from the first quarter of 2012 onwards. Records are available only for March, June, September, and December until 2018, and for every month thereafter. To ensure consistency across the full sample period, we retain only observations from March, June, September, and December and treat these months as representing quarters. Each worker-quarter observation includes monthly earnings, days worked, age, gender, and occupation classified under the ISCO-08 scheme.³ We keep only the highest-paying job for workers holding multiple jobs in a given month. Observations with zero days (2.5% of the sample) worked are excluded, as are those with reported earnings below 85 percent of the statutory minimum wage (0.06% of the sample), which we attribute to recording errors.

From the worker-level records, we construct the following firm-quarter variables. First, mean log daily wage is the logarithm of average daily wages across workers, where daily wage is computed as monthly earnings divided by days worked. Second, we construct wage bins relative to the statutory minimum wage and define employment in each wage bin as the number of workers whose wages fall within the bin’s bounds. The total number of employees is the sum across all bins. All wages are deflated using a minimum wage index with 2006 Q1 as the base period. Finally, for each firm-quarter, we count the number of workers in each one-digit ISCO occupation group, which we use as the dependent variable in the occupation-level analyses.

The firm registers, compiled annually by the Turkish Statistical Institute, provide firm-level information on industry affiliation at the NACE Rev.2 four-digit level⁴ and firm age, both of which enter the matching procedure described in Section 3. The balance sheet data, collected annually by the Ministry of Treasury and Finance, contain detailed financial statement items. We use these records to construct firm-level export intensity, defined as the ratio of export revenues to net sales. Net sales, deflated using the consumer price index, and export intensity also serve as matching variables.

The patent data from the Turkish Patent and Trademark Office record both patent applications and grants for all registered firms over the period 2006–2023. The underlying records are available at a daily frequency and are aggregated to the quarterly level to align with the worker register data. We use patent applications as a proxy for innovation activity. In additional analyses, we also use granted patents data with the purpose of replicating selected results separately for the subsamples of granted and rejected patent applications.

Together, these four sub-datasets, worker registers, firm registers, balance sheet data,

³ISCO-08 (International Standard Classification of Occupations, 2008) is the occupational classification system developed by the International Labour Organization, organizing jobs into a hierarchical structure based on the nature of the work performed and the required skill level.

⁴NACE Rev.2 is the European standard classification of economic activities. The four-digit level corresponds to the most disaggregated class.

and patent records, form the empirical foundation of our analysis, enabling us to track employment, wages, and innovation activity at the firm-quarter level over the period 2012–2023.

3 Treatment and Control Groups

3.1 Innovative (Treated) Firms

We utilize patent applications to proxy for innovation, i.e., a firm is assumed to be “innovative” once it applies for a patent. This strategy is implemented in several studies, including (Griliches (1990); Czarnitzki and Hussinger (2004); Kogan et al. (2017)).

Our data also include information on granted patents, which allows us to repeat some of our baseline analyses separately for the subsamples of granted and rejected applications. However, we treat all patent applications, regardless of outcome, as innovative activity in our main analyses, for two reasons. First, only slightly more than half of applications are eventually granted, so using grants alone would sharply reduce the sample. Second, a grant requires absolute novelty (new to the world), whereas we are interested in firm-level innovative activity: filing an application indicates that a firm has engaged in a technology-development effort that is new to the firm, even if not novel globally. Indeed, our results for rejected applications are qualitatively similar to those for granted applications, just smaller in magnitude.

We estimate the effects of patent applications, i.e., innovation, on various outcomes using a difference-in-differences (DiD) strategy. We define the pre-treatment period as the 24 quarters preceding a firm’s patent application and the post-treatment period as the 16 quarters following it, including the application quarter. In what follows, let t denote the number of quarters relative to treatment, with $t \in \{-24, -23, \dots, 0, \dots, 15\}$, where $t = 0$ denotes the treatment quarter or the first quarter after treatment. We opt for a longer pre-treatment window to account for anticipation effects: as we show in the following sections, patent applications appear to be anticipated by firms as early as three years in advance. We discuss this issue further below.


This choice of window length restricts our estimation sample to firms that applied for a patent between 2012 Q1 and 2020 Q1, yielding 33 treatment cohorts. As detailed in the Data section, our worker registers, which provide our main outcome variables such as wages and employment, are available only between 2006 Q1 and 2023 Q4. Given this coverage, only firms with patent applications falling within the 2012 Q1–2020 Q1 window can be assigned the full 24-quarter pre-treatment and 16-quarter post-treatment periods required by our design.

We define the treatment event as a firm’s first patent application. Specifically, a firm is

classified as a first-time patent applicant if it filed no patent applications between 2006 Q1 and 2011 Q4 and submitted its first patent application in 2012 Q1 or later. If the firm filed several patent applications after 2012 Q1, we only consider the first one. Figure 1 presents the number of first-time patent applications in each quarter defined accordingly.

We refer to all firms with a patent application as the "Full Sample". These firms do not necessarily report worker registers for all 24 pre-treatment and 16 post-treatment quarters, resulting in an unbalanced panel. For our event-study design, we require non-missing wage and employment observations across all 40 quarters surrounding treatment. We therefore further define the "Effective Sample" as the subset of firms that report worker registers for all 40 quarters surrounding treatment. This dynamically balanced panel is the treatment sample we use in our estimations and comprises a total of 1,844 firms that applied for a patent between 2012 Q1 and 2020 Q1.

Figure 1: Number of first-time patent applications over time




Notes: The full sample (solid blue line) shows the number of firms filing their first patent application between 2012 Q1 and 2020 Q1. The effective sample (dashed red line) includes firms that continuously report both balance sheet and worker register data in every year from six years before to four years after their first patent application, yielding a dynamically balanced panel.

3.2 What Type of Innovation?

Next, we turn to understanding the activities of the firms in our effective treatment sample and the types of patents they apply for. We find that 67.5% of the firms filing a patent application for the first time operate in the manufacturing industry, i.e., NACE2 codes 10 through 32. Figure 2 presents the distribution of firms across NACE2 industries in our

effective sample. These top 10 sectors shown in the figure account for two-thirds of all first-time patent applications, and seven of them belong to manufacturing.

Figure 2: Sectoral Distribution of First-Time Patent Applicants by NACE2



Notes: This figure shows the distribution across NACE2 sectors of the 1,844 firms in our effective treatment sample. The figure reports the top 10 NACE2 sectors with the highest number of patent applications. Brief sector descriptions are as follows: 28 – Mfg. of machinery and equipment; 46 – Wholesale trade; 25 – Mfg. of fabricated metal products; 22 – Mfg. of rubber and plastic products; 13 – Mfg. of textiles; 27 – Mfg. of electrical equipment; 29 – Mfg. of motor vehicles and trailers; 10 – Mfg. of food products; 47 – Retail trade; 62 – Computer programming and consultancy.

To better understand the type of innovation taking place in these manufacturing industries, we examine the patent portfolios of several large, well-established firms using data from the Turkish Patent and Trademark Office (TÜRKPATENT). Because these firms are well established, it is unlikely that they filed no patent applications between 2006 Q1 and 2011 Q4. As a result, they are unlikely to qualify as first-time patent applicants under our definition and therefore do not appear in our effective sample. Nevertheless, reviewing their patent portfolios provides a useful illustration of the type of innovation pursued in these industries.

Across these manufacturing industries, we observe both product and process innovation. In NACE2 28 (machinery and equipment), one firm patented a hydraulic cylinder that adjusts the angle of an attachment on construction equipment, while another patented a new method for positioning and filling the fuel tank on an excavator. In NACE2 13 (textiles), one firm patented a method for producing a foam-coated woven fabric, while another

introduced a multifunctional first-aid bandage. In NACE2 27 (electrical equipment), one firm patented a wireless control method for household appliances, while another patented a new production method for a combined washer-dryer device. Finally, in NACE2 25 (fabricated metal products), a representative patent describes an insulated sandwich panel compatible with solar panel installations.

As these examples illustrate, firms in these industries pursue both product and process innovation. Ultimately, however, most of these are medium-technology industries, and the innovation being patented reflects incremental improvements to existing products and production processes rather than breakthrough technologies.

3.3 Control Group Construction via Matching

We next build a control group since patent applicants are not a random sample of firms. They may already be larger or growing faster than average before filing a patent application. If so, their growth could slow down afterward simply due to regression to the mean. To rule this out, we match treated firms to similar firms that did not apply for a patent.

We employ a hybrid, staggered matching strategy, running the same matching algorithm separately for each of the 33 treatment cohorts. Specifically, to match treated firms in a given cohort to an appropriate control firm, we retain all untreated firms with non-missing worker register data over the same 40-quarter window. We implement exact matching on NACE2 sector together with propensity score matching on the logarithms of firm age, net sales, mean employee age, mean daily wage, and number of employees, as well as export intensity (defined as export value divided by net sales), all measured at $t = -12$, that is, roughly three years prior to treatment, and just before the onset of anticipation effects associated with innovation. See Section 4.2 for more on this anticipation window. The matching procedure begins with the earliest cohort, i.e., firms that applied for a patent in 2012 Q1, and proceeds sequentially through later cohorts. A firm matched to a control in an earlier cohort is not eligible to be matched again in a later cohort.

This procedure yields 1,771 successful matches out of the 1,844 treated firms in our effective sample, a success rate of approximately 96%. Table 1 reports the means and standard deviations of the variables used in the propensity score matching step. NACE2 is omitted from this table, as exact matching guarantees that treated and control firms operate within the same sector by construction.

The table also includes summary statistics for all unique firms observed during the potential matching period (N unique firms = 1,308,071). Treated firms differ markedly from this broader population: they are larger, pay higher wages, and export more intensively. This confirms that patent applicants are not representative of the average firm, and justifies our matching approach.

Table 1: Covariate Balance Between Treated and Control Firms

	Treated		Control		All Firms	
	Mean	SD	Mean	SD	Mean	SD
Log Firm Age	2.67	0.69	2.63	0.75	2.73	1.13
Export Intensity	0.16	0.24	0.16	0.28	0.04	0.18
Log Net Sales	15.52	2.02	15.54	2.11	12.84	1.78
Log Mean Employee Age	3.51	0.10	3.51	0.12	3.54	0.19
Log Mean Daily Wage	3.34	0.45	3.34	0.50	2.98	0.27
Log Number of Employees	3.76	1.69	3.71	1.67	1.51	1.22
Number of unique firms	1,771		1,771		1,308,071	

Notes: We first perform exact matching within firms' NACE 2-digit industries. We then implement one-to-one propensity score matching using the variables reported in the table. Export intensity is defined as the ratio of export value to total sales. The remaining variables are self-explanatory. SD denotes standard deviation.

4 Wages and Employment

This section examines the implications of innovation for firms' wages and employment. We begin by presenting our econometric methodology and examining how average wages and firm size evolve following innovation. We then define seven wage bins relative to the statutory minimum wage and investigate the effects of innovation on employment across these wage bins.

4.1 Econometric Specification

Our strategy for identifying the implications of innovation, proxied by patent applications, relies on event-study analyses (except for the transition analyses of Section 5). Recall that the treatment and control groups are constructed as described in Section 3, and that these groups constitute a dynamically balanced panel: all units are observed for 40 quarters around their (cohort-specific) quarter of treatment, regardless of the calendar quarter in which treatment occurs. Because treatment timing is staggered across cohorts, the calendar quarters in which different cohorts (and their matched controls) are observed do not necessarily overlap, even though the number of quarters observed relative to treatment is the same across cohorts.

Recall also that matching is implemented separately for each cohort based on firms' pre-treatment attributes. Consequently, the matching quarter, which differs across treatment cohorts, corresponds to the hypothetical pre-treatment quarter for the control units. As a

result, comparisons based on calendar quarters would evaluate control units from other cohorts relative to a calendar quarter that has a different distance to treatment time than that of the treated unit. Such comparisons would be inconsistent with our matching procedure. We therefore rely on the stacked difference-in-differences method in the spirit of Cengiz et al. (2019), which compares treated and control units across quarters defined relative to treatment rather than across calendar quarters. Appendix A discusses this issue in detail and explains why we favor this approach over alternative staggered-adoption estimators.

Formally, we estimate the following event-study specification:

$$y_{i,t} = \alpha_i + \gamma_t + \sum_{\substack{\tau=-24 \\ \tau \neq -12}}^{15} \beta_\tau (\text{Treated}_i \times 1[t = \tau]) + \varepsilon_{i,t} \quad (1)$$

where $y_{i,t}$ denotes the outcome of firm i in quarter t , with t measured relative to the firm's cohort-specific patent-application (treatment) quarter, $t \in \{-24, -23, \dots, 0, \dots, 15\}$. Variables α_i and γ_t are firm and relative-quarter fixed effects, respectively, and Treated_i is an indicator equal to one if firm i belongs to the treated group. The coefficients of interest, β_τ , capture the effect of innovation on the outcome at each quarter relative to treatment. We omit $t = -12$, roughly three years prior to treatment, as the baseline category, since this is both the quarter used for matching and the quarter immediately preceding the start of anticipation effects (see Section 4.2). Standard errors are clustered at the firm level.

4.2 Average Wage and Firm Size


We begin our analysis by estimating the effect of innovation on firm size and average log daily wage, using the event-study strategy introduced in Section 4.1. Both variables are constructed from worker register data, as described in Section 2.

Results are presented in Figure 3. The left panels show the event-study estimates, while the right panels show raw means of the outcomes. Both panels reveal the same pattern: trends in firm size and average log daily wage are indistinguishable across treatment and control groups until roughly three years before treatment, at which point they begin to diverge. We interpret this divergence as evidence that innovation is anticipated well before the patent application itself, with firms beginning to adjust their behavior roughly three years in advance. Such an anticipation period is consistent with the documented lag between R&D investment and observable innovation outcomes (Desyllas and Hughes, 2010; Garcia

et al., 2026).⁵

Accordingly, in all of our event-study charts, the dashed vertical line at $t = -11$ marks the onset of anticipation effects, while the solid vertical line at $t = 0$ marks the quarter of patent application. As noted in Section 4.1, treated firms are matched to control firms just before this anticipation window, at $t = -12$, which we also use as the omitted baseline period in our event-study estimates.

Figure 3: Innovation effects on firm size and mean wage



Notes: The sample consists of 1,771 treated and 1,771 control firms, yielding 141,680 firm-quarter observations. Treated firms are those that applied for a patent between 2012 Q1 and 2020 Q1 and continuously report worker registers every quarter for 24 quarters before and 16 quarters after the patent application (dynamically balanced panel). Control firms are obtained from a hybrid matching procedure. Further details on the construction of the treatment and control groups are provided in Section 3. Event-study coefficients in the left panels are estimated using the stacked difference-in-differences method introduced in Section 4.1. The right panels report the evolution of raw means of outcomes (firm size and log daily wage) for the treatment (solid blue line) and control (dashed red line) groups.

Overall, Figure 3 shows that both firm size and average log daily wage respond positively

⁵In the Turkish context, this three-year window also aligns closely with the duration of the most prominent public R&D support scheme for industrial firms, the TÜBİTAK 1501 Industrial R&D Projects Support Programme, under which projects may run for a maximum of 36 months. See <https://tubitak.gov.tr/tr/destekler/sanayi/ulusal-destek-programlari/1501-tubitak-sanayi-ar-ge-projeleri-destekleme-programi>.

to innovation. By the third year after patent application, firm size increases by close to 0.3 log points among treated firms, while average daily wage rises by approximately 0.04 log points. The wage result is consistent with the prior literature discussed in our literature review, where innovation is typically found to raise wages. The positive employment effect further suggests that patent applications in Türkiye are, on average, not associated with labor displacement at the firm-level.

As a further analysis, we re-estimate this specification separately for the subsamples of firms whose patent applications were granted (1003 firms) and those whose applications were rejected (768 firms). Results are presented in Appendix Figures C1 and C2. As shown, the effects are larger in magnitude for firms whose patent applications were granted, but the same qualitative pattern also holds for firms whose applications were rejected. This suggests that the observed effects are not just driven by the granting of a patent, but rather by the innovation process around the application itself.

4.3 Construction of Wage Bins

Having shown that innovation raises both firm size and average wages, we now turn to how these increases are distributed along the wage distribution. To do so, we define seven wage bins relative to the statutory minimum wage (MW), which form the basis of the empirical analysis in the remainder of this section. To ensure that a given bin implies a comparable level of affluence across years, we define bins relative to the statutory minimum wage rather than in nominal or real terms. Specifically, in each quarter, a worker's wage is assigned to one of the seven bins according to its ratio to that quarter's minimum wage, as shown in Table 2.

Table 2 also reports the distribution of workers across wage bins for treated and control firms three years before treatment at $t = -12$, which we defined in Section 4.2 as the quarter immediately preceding the onset of anticipation effects of innovation. Although the treated and control distributions differ slightly, workers in the first two bins remain below the median in both groups and can therefore be regarded as low-skilled and/or medium- to low-skilled. Starting from Bin 3, which exceeds the 60th percentile of the distribution in both groups, we move into the medium- to high-skilled range.

Table 2: Wage Bin Definitions and Distribution at $t = -12$: Treated vs. Control

Bin	Wage Range (relative to MW)	Treated		Control	
		Fraction	Cumulative	Fraction	Cumulative
1	[1.0, 1.1)	0.190	0.190	0.223	0.223
2	[1.1, 1.5)	0.269	0.460	0.229	0.452
3	[1.5, 2.0)	0.173	0.632	0.167	0.618
4	[2.0, 2.5)	0.110	0.742	0.099	0.717
5	[2.5, 3.5)	0.103	0.845	0.110	0.827
6	[3.5, 6.0)	0.083	0.928	0.092	0.919
7	≥ 6.0	0.072	1.000	0.081	1.000
N		197,619		192,957	

Notes: This table reports wage bin definitions, expressed relative to the statutory minimum wage (MW) of the relevant quarter, together with the distribution of workers across these bins for treated and control firms at $t = -12$, i.e., three years prior to treatment and immediately before the onset of anticipation effects. Each bin includes its lower bound and excludes its upper bound. Fraction denotes the share of workers in a given bin, and Cumulative denotes the cumulative share of workers in that bin or below. N denotes the total number of worker-quarter observations underlying these calculations.

4.4 Employment across Wage Bins


We have thus far shown that innovation, proxied by patent applications, positively affects average daily wage and firm size in Türkiye. We next turn to the specific mechanisms underlying these outcomes. For average daily wage to rise, it is sufficient that employment share above the mean wage increases. But where exactly does employment share increase along the wage distribution, just above the middle, or at the top? It is even possible that employment at the lower end of the distribution declines, provided this is offset by a stronger increase at the upper end.

To address these questions, we estimate the effects of patent applications on employment in the seven wage bins defined in Table 2. We follow the same event-study specification described in Section 4.1, where the outcome variables are the logarithms of employment levels within each wage bin. Because many firm-bin cells have zero employment, and the logarithm of zero is undefined, we follow the recommendation of Chen and Roth (2024) to assign an explicit value to $\log(0)$. Specifically, we set $\log(0) = -1$, which is, in essence, equivalent to assigning a log-difference value of one to the extensive margin, that is, to a bin moving from zero to one employee.⁶


⁶We also assign -0.1 to $\log(0)$ as an alternative. The results, not reported in the manuscript, are very similar.

Figure 4: Innovation effects on log employment across wage bins


Bin 1



Bin 3



Bin 6




Notes: The sample consists of 1,771 treated and 1,771 control firms, yielding 141,680 firm-quarter observations. Treated firms are those that applied for a patent between 2012 Q1 and 2020 Q1 and continuously report worker registers every quarter for 24 quarters before and 16 quarters after the patent application (dynamically balanced panel). Control firms are obtained from a hybrid matching procedure. Further details on the construction of the treatment and control groups, as well as on the wage bins (Bin 1–Bin 7), are provided in Section 3. Event-study coefficients in the left panels are estimated using the stacked difference-in-differences method introduced in Section 4.1. The right panels report the evolution of raw mean log bin sizes for the treatment (solid blue line) and control (dashed red line) groups. Results for the remaining bins are reported in Appendix Figure C3.

Figure 4 shows the event-study estimates (left-hand panels) alongside the raw means of log employment for the treated and control groups (right-hand panels), for three exemplary

bins-bins 1, 3, and 6-while event-study results for the remaining bins are presented in Appendix Figure C3. The effect of innovation on employment in the first wage bin only becomes statistically significantly positive in the third year after treatment. Effects on bins 3 and 6, by contrast, turn positive and significant shortly after the start of the anticipation window, with the coefficient increasing smoothly thereafter.

Figure 5: Log average change in employment across wage bins 9-16 quarters after innovation



Notes: Bar heights represent the average log change in employment for each wage bin in innovative firms relative to similar non-innovative firms, calculated as the mean of the last eight point estimates (i.e., $t = 8$ to $t = 15$, or the 9th to 16th quarters after treatment) from the corresponding event-study specifications reported in Figures 4 and C3.

To bring the remaining wage bins into the picture and facilitate easy comparison across all seven, we average the event-study coefficients over 9 to 16 quarters after treatment (the last eight quarters in event-study charts) and plot them in the bar chart in Figure 5. The resulting value can be interpreted as the average rise in employment within a given wage bin, three to four years after patent application, in innovative firms relative to similar but non-innovative firms. The smallest rise in employment occurs in the lowest wage bin, at 0.10 log points, and, as seen in the corresponding event-study chart, this relative increase (relative to the control group) is only marginally significant. The largest average relative increase occurs in the third wage bin, at 0.31 log points. As shown in Table 2, this bin

comprises workers earning between 1.5 and 2 times the statutory minimum wage, placing most of them above the median wage and, in some cases, above the 60th percentile. The average relative employment rise in the remaining bins hovers around 0.19-0.25 log points.

What does this result tell us? First, there is no employment loss anywhere along the wage distribution: employment rises almost everywhere, with the possible exception of the first wage bin. Second, the strongest employment gains occur among medium-to-high-wage workers, that is, those earning just above the median wage.

As an additional analysis, we re-estimate the wage-bin specification separately for the subsamples of firms whose patent applications were granted and those whose applications were rejected. For each subsample, we again average the event-study coefficients over 9 to 16 quarters after treatment for each wage bin; the corresponding event-study estimates are not reported, but the resulting bar chart is presented in Appendix Figure C4. The pattern observed in the full sample is preserved in both subsamples, though the magnitude of the effect is stronger among firms whose applications were granted.

5 Worker Transitions

While the results of the previous section shed light on employment gains across the wage distribution in innovative firms, several questions remain open. What are the sources of these employment changes? Do innovative firms promote incumbent workers to higher wages? Do they retain incumbent workers in place while hiring new, high wage workers? Or do they, at least in part, replace incumbent workers following innovation? This section attempts to answer these questions.


For this purpose, we construct transition matrices for treated and control firms and compare them. Specifically, we construct these matrices separately for incumbent workers and new workers. Incumbent workers are defined as those employed at the firm in the quarter just before the anticipation window, $t = -12$. New workers are defined as those who do not work for the firm at $t = -12$ but are employed by the firm at a selected later date, e.g., at $t = 15$.

Note that, in order to construct these matrices, we must be able to track workers starting from $t = -12$, twelve quarters before the patent application quarter. Recall from Section 2 that our worker register data include worker identifiers only as of 2012 Q1. Thus, the earliest cohort we can include in the transition matrices is the one whose pre-anticipation period, $t = -12$, corresponds to 2012 Q1, that is, the cohort treated in 2015 Q1. As a result, this section is restricted to firms treated as of 2015 Q1, whereas the previous section included all firms treated as of 2012 Q1.

5.1 Incumbent Workers

We start by exploring the transitions of incumbent workers. There are 197,619 workers in treated firms and 192,957 workers in control firms at $t = -12$. We examine where these workers transition to by $t = 15$ (16 quarters after treatment), broken down by wage bin. For example, a worker employed in bin j at $t = -12$ may be working in bin i of the same firm, in bin i of a different firm, or may be in nonemployment at $t = 15$. The full transition matrices for the treatment and control groups are presented in Appendix Table B2. In the main text, we highlight only important patterns.

Figure 6: Worker Transitions from $t = -12$ to $t = 15$ by Wage Bin



Notes: This figure shows, separately by wage bin, the fraction of workers employed at treated and control firms at $t = -12$ (immediately before the anticipation effect) who transition into *Nonemployment*, remain at the *Same Firm*, or move to a *Different Firm* by $t = 15$ (four years after treatment). Within each wage bin, the three fractions sum to one. Wage bins (Bin 1–Bin 7) are defined in Section 3. Red solid lines with circles denote treated firms; blue dashed lines with squares denote control firms. The associated full transition matrix is presented in Appendix Table B2.

We begin by illustrating, in Figure 6, where incumbent workers starting in any wage bin j (on the x-axis) transition to in the treated and control groups. The first panel shows the

fraction of workers in each wage bin who transition into nonemployment, while the other two panels show, respectively, the fraction who remain at the same firm and the fraction who move to a different firm (irrespective of the destination wage bin). These fractions for each wage bin sum up to one.


Figure 6 reveals one clear pattern: as evident from the Same Firm panel, the line for treated firms lies above that for control firms across all wage bins. This means that innovative firms retain incumbent workers at higher rates than non-innovative firms, at every point of the wage distribution. The treated–control gap in retention rates is largest at the top of the distribution, in bin 7: treated firms retain more than 45% of these workers, compared to around 30% at control firms.

We next ask what happens to incumbent workers who remain at the same firm: do they stay in their original wage bin, get promoted to a higher one, or move down? To answer this, we construct a heat map, Figure 7. Specifically, for both treated and control groups, we calculate, among the workers who start in wage bin j at $t = -12$ and remain at the same firm, the fraction who are in wage bin i at $t = 15$. We then take the difference between the treated and control fractions for every j -to- i transition to construct Figure 7.

Consider the following example to illustrate the construction of this heat map. Among workers in treated firms who are in wage bin 1 at $t = -12$ and remain at the same firm, around 33.11% stay in bin 1 and about 39.28% move to bin 2 by $t = 15$ (these fractions, together with transitions into the remaining bins, sum to one). Among workers in control firms, the corresponding fractions are 41.82% and 38.73%, respectively. Taking the difference, the bin-1-to-bin-1 cell is about -8.71 percentage points, meaning that the share of workers who start and remain in bin 1 is 8.71 percentage points lower in treated firms than in control firms. Similarly, the bin-1-to-bin-2 cell is about 0.55 percentage points, meaning that the share of workers who move from bin 1 to bin 2 is 0.55 percentage points *higher* in treated firms than in control firms. By construction, each column sums to zero. Positive cells are shaded in blue and negative cells in red, with the intensity of the shade reflecting the magnitude.

Figure 7 reveals an interesting pattern. Below the diagonal, corresponding to transitions into higher wage bins, the cells are almost uniformly blue for all bins except bin 5. This indicates that the rate of moving up into higher wage bins is greater in treated firms than in control firms. For instance, as noted above, treated firms retain 8.71 percentage points fewer of their bin-1 workers in bin 1; of this gap, 0.55 points instead move to bin 2, 4.10 points to bin 3, and 2.15 points to bin 4. Looking at the second column, 4.48 points fewer of the workers who started in bin 2 remain in bin 2 at treated firms. The rate of moving into bin 1 is lower, and 6.21 points more move up to bin 3.

Figure 7: Treated–Control Differences in Wage Bin Transitions ($t = -12$ to $t = 15$) Among Stayers



Notes: This figure focuses exclusively on workers who remained at the same firm between $t = -12$ and $t = 15$ (“stayers”). Wage bins (Bin 1–Bin 7) are defined in Section 3. The cell in row i , column j reports the difference, in percentage points, between (i) the fraction of treated stayers who started in bin j at $t = -12$ that ended in bin i at $t = 15$, and (ii) the analogous fraction for control stayers who started in bin j (Treated minus Control). Blue cells indicate positive differences (treated stayers are relatively more likely to make that transition than control stayers), while red cells indicate negative differences (treated stayers are relatively less likely to make that transition). Diagonal cells, highlighted with black borders, correspond to workers who remained in the same wage bin between $t = -12$ and $t = 15$. The associated full transition matrix is presented in Appendix Table B2.

Interpreting the remaining columns similarly, workers who start in bins 1, 2, 3, or 6 are consistently more likely to move up to a higher wage bin at treated firms relative to control firms, while workers who start in bins 4 or 7 are consistently more likely to remain in their starting bin, although even within bin 4, the rate of moving up to bin 5 is still higher at treated firms. The only exception is bin 5: workers who start in bin 5 are less likely to remain in that bin at treated firms, with 3.72 percentage points more moving down to bins 3 or 4, even though 1.55 percentage points are still more likely to move up all the way to bin 7.

Taken together, these results point to two clear patterns in innovative firms relative to non-innovative firms: (i) innovative firms are more likely to retain their workers at every

wage bin, and especially so at the very top of the distribution; and (ii) workers that stay are more likely to move up the wage ladder.

We construct the same transition matrices for the period from $t = -12$ to $t = 0$. The resulting matrix is presented in Table B1, with the corresponding figures shown in Figure C5 and Figure C6. The results are broadly similar to those presented in the main text.

5.2 New Hires


We next turn to new hires, i.e., workers who are employed at the firm at $t = 15$ but not at $t = -12$. There are 176,310 such workers in treated firms, compared to 132,055 in control firms. This considerably higher number of new hires in treated firms shows that, in addition to retaining incumbent workers at higher rates, treated firms also hire more new workers, both in absolute terms and relative to their size at $t = -12$ (when treated and control firms employed a comparable 197,619 and 192,957 workers, respectively, as noted in the previous section).

We construct transition matrices for these new workers by source, nonemployment or a different firm, and by wage bin. Specifically, these matrices show whether a new worker in bin j at $t = 15$ came from nonemployment at $t = -12$ or from bin i of a different firm. The full matrices for the treated and control groups are presented in Appendix Table B4. As with incumbent workers, we highlight only the important patterns here.

We start by illustrating, in Figure 8, the number of new hires at $t = 15$ by wage bin, for the treated and control groups. The left panel shows new hires in absolute terms, as the log number of new hires, while the right panel shows the number of new hires into each bin relative to the size of that bin at $t = -12$.

The left panel of Figure 8 shows that treated firms hire more new workers than control firms in absolute terms across all wage bins. The right panel tells a similar story: the number of new hires relative to bin size at $t = -12$ is higher in treated firms in all bins (especially in bin 3) except bin 2. At first glance, the similar hiring rates into bin 2 across treated and control firms may appear at odds with the result in Section 4.4, where the average increase in employment 9 to 16 quarters after treatment is 19 log points (more than 19%) higher in the treated group relative to the control group. However, this is not at odds with that result: as Figure 6 shows, treated firms retain their incumbent workers in bin 2 at a considerably higher rate than control firms, so the overall employment gain in bin 2 is driven by retention rather than new hiring.

Figure 8: Destination of New Hires by $t = 15$, by Wage Bin




Notes: This figure shows the wage bin destination of workers newly hired at treated and control firms by $t = 15$, defined as workers present at $t = 15$ but not employed at the firm at $t = -12$. The left panel plots the log number of new hires in each wage bin. The right panel plots the number of new hires in each wage bin relative to the number of workers employed in that bin at $t = -12$. Wage bins (Bin 1–Bin 7) are defined in Section 3. Red solid lines with circles denote treated firms; blue dashed lines with squares denote control firms. The associated full transition matrix is presented in Appendix Table B4.

We next turn to the sources of new workers by wage bin at $t = 15$ (on the x-axis), that is, whether they come from nonemployment or from a different firm. The left panel of Figure 9 shows the fraction of new workers coming from nonemployment, while the right panel shows the fraction coming from a different firm. Within each bin, these fractions sum to one. Visual inspection reveals no distinguishable difference between the treated and control groups. We also construct a heat map analogous to Figure 7 for new workers, to investigate whether there is a systematic difference in the transitions of new workers coming from other firms, that is, whether treated firms tend to hire new workers by promoting them into higher wage bins or by placing them into lower ones relative to their previous firm. We find no such systematic difference, and do not report this heat map in the interest of space. Overall, the main finding of this section is that innovative firms hire more new workers into upper wage bins relative to non-innovative firms.

We also construct the equivalent of the new-worker transition matrix for the anticipation window, i.e., for workers who do not work for the firm at $t = -12$ but do so at $t = 0$. This matrix and the associated figures are presented in Appendix Table B3 and Figures C7 and C8.

Figure 9: Source of New Hires by $t = 15$, by Wage Bin



Notes: This figure shows, separately by wage bin, the source of workers newly present at treated and control firms by $t = 15$ (i.e., workers present at $t = 15$ but not employed at the firm at $t = -12$). The left panel plots the fraction of these new hires who were in *Nonemployment* at $t = -12$; the right panel plots the fraction who were employed at a *Different Firm* at $t = -12$. Within each wage bin, the two fractions sum to one. Wage bins (Bin 1–Bin 7) are defined in Section 3. Red solid lines with circles denote treated firms; blue dashed lines with squares denote control firms. The associated full transition matrix is presented in Appendix Table B4.

5.3 What Do Transitions Tell Us?

It is useful to bring these transition matrix results together. Recall that our purpose in this section is to uncover the mechanisms underlying the two stylized facts documented earlier: the increase in average daily wages and in total employment at innovative firms. Naturally, different mechanisms are at work at different points of the wage distribution. Bin 3, for example, grows larger in treated firms primarily because these firms hire considerably more new workers into this bin relative to control firms. The upper bins in treated firms grow through a combination of mechanisms, including higher retention rates among incumbent workers and greater promotion from lower bins.

These bin-specific mechanisms are, admittedly, somewhat ad hoc and should not be over-generalized. The broader message, however, is as follows. Innovative firms experience increases in average daily wages and firm size relative to non-innovative firms through three main channels: (i) they retain incumbent workers at higher rates, with the treated–control retention gap being largest at the top of the wage distribution, (ii) they generally promote incumbent workers into higher wage bins more frequently, and (iii) they hire more new workers, disproportionately into higher wage bins. These channels appear to play a prominent role in explaining these stylized facts.

In what follows, we take a closer look at the implications of these transition patterns for wage inequality.

Wage inequality. We next examine how wage inequality within treated and control firms evolves over the same horizon. While we do not pursue a causal test, we conduct a descriptive analysis to gain some insight. We compare the variance of log daily wages among all workers in treated and control firms between $t = -12$ (just before the anticipation window) and $t = 15$ (16 quarters after treatment). In treated firms, the variance of log daily wages falls from 0.342 to 0.285. In control firms, it falls from 0.356 to 0.286. Wage inequality thus falls by a similar magnitude in treated and control firms, suggesting that innovation neither increases nor decreases within-firm wage inequality. This descriptive comparison therefore provides no clear evidence that innovation has a visible effect on within-firm wage inequality. Unlike the other analyses in this section, we use all treated firms here rather than restricting to those treated as of 2015 Q1.

6 Occupational Composition

We next examine whether the occupational composition of innovative firms changes following innovation. The EIS data include information on workers' 4-digit occupations under the ISCO-08 classification starting in 2014. However, this variable is missing for approximately 40% of worker-quarter observations.

To address these gaps, we impute occupation values from other records of the same individual. The imputation follows a four-step strategy. We first search within the same establishment-year for the lowest-skill occupation on record. If none is found, we broaden the search to the entire firm within that year, again taking the lowest-skill occupation available. If this fails to find any record belonging to that individual as well, we turn to any occupation reported for the worker within the same quarter, and finally, within the same calendar year. When none of these searches yield a value, we assign the individual's modal occupation across their entire work history. This procedure allows us to recover occupation information as far back as 2012 Q1, the earliest quarter for which worker identifiers are available.

Following the event-study convention established in Section 4, we require six years of pre-treatment and four years of post-treatment observations. Since our occupation data begin in 2012 Q1, this quarter must serve as the first pre-treatment period for the earliest cohort included in the analysis, namely, the cohort treated in 2018 Q1. As a result, the regressions in this section are restricted to firms treated between 2018 Q1 and 2020 Q1 (587 firms) and control firms that are matched with these firms.

Our dependent variable is the firm-level employment in each 1-digit ISCO-08 class. We first present some descriptive statistics regarding the occupation composition of our treated and control groups in Table 3 for the quarter just before anticipation, $t = -12$. The "Mean Workers" panel provides the average number of workers in each occupation, while the

“Employing Share” panel presents the fraction of firms that employ at least one worker in the respective occupation class.

Table 3: Occupation Structure of the Sample

ISCO-08 Code	Occupation	Mean Workers		Employing Share (%)	
		Treated ($N = 587$)	Control ($N = 587$)	Treated ($N = 587$)	Control ($N = 587$)
1	Managers	8	5	16.52	16.87
2	Professionals	19	18	28.96	23.85
3	Technicians & Assoc. Prof.	15	15	28.45	27.26
4	Clerical Support Workers	11	14	27.09	25.72
5	Service & Sales Workers	6	15	17.04	23.17
6	Skilled Agr., Forestry & Fishery	1	2	3.07	2.04
7	Craft & Trades Workers	34	30	35.60	36.63
8	Plant & Machine Operators	45	43	28.45	30.83
9	Elementary Occupations	27	33	34.75	31.69

Notes: “Mean Workers” is the mean number of workers employed in each occupation group across firms. “Employing Share” is the percentage of firms employing at least one worker in the given occupation group. Statistics are computed over 587 treated and 587 control firms for the quarter immediately preceding innovation anticipation, $t = -12$. The sample is smaller than indicated in Section 3 because occupation data are available in our dataset only from 2012 onward. As we require 24 quarters (6 years) of pre-treatment observations, our sample is restricted to events as of 2018 Q1, which reduces the number of treatment firms.


Recall from the discussion around Figure 10 that a larger fraction of our effective sample consists of manufacturing firms. It therefore comes as no surprise that “ISCO7: Craft & Trade Workers”, a category that includes machinists, welders, toolmakers, and metal and machinery tradesmen, together with “ISCO8: Plant & Machine Operators,” which largely comprises assembly-line and machine operators, and “ISCO9: Elementary Occupations,” populated mainly by manufacturing laborers, packers, and cleaning and security staff, emerge as the most crowded occupation groups in Table 3. Firms in our sample employ, on average, around or more than 30 workers in each of these three categories.

While average employment across the treated and control groups is similar in most occupations, Table 3 shows that the control group employs noticeably more workers, on average, in “ISCO5: Service & Sales Workers.” At the same time, the share of firms employing at least one worker in this category is not that different across the treated and control groups, suggesting that there is no systematic differences in whether this occupation is staffed at all. Recall that our matching procedure introduced in Section 3.3 does not incorporate any information on occupation. Thus, such minor discrepancies in occupational composition are not surprising.

We then estimate the event-study specification introduced in Section 4.1, this time using

firm-level employment in each of the nine 1-digit ISCO occupations as the dependent variable. The resulting event-study coefficients, along with the associated raw means, are reported in Appendix Figure C9. In the main text, we summarize these results with a bar chart showing the mean of the coefficients from $t = 8$ to $t = 15$, that is, the average increase in log employment in treated firms relative to control firms, 9 to 16 quarters after treatment, for each occupation.

Figure 10: Log average change in employment across occupations 9-16 months after treatment



Notes: Bar heights represent the average log change in employment by occupation, calculated as the mean of the last eight point estimates (i.e., $t = 8$ to $t = 15$, or the 9th to 16th quarters after treatment) from the corresponding event-study specifications reported in Figure C9. Occupations are classified according to the one-digit ISCO-08 classification, the definitions of which are provided in Table 3.

Results are presented in Figure 10. Employment rises more in the treated group than in the control group across almost all occupations, consistent with our earlier finding that total employment increases in innovative firms relative to non-innovative ones. This increase, however, is especially pronounced for ISCO3, ISCO7, and ISCO8, a pattern that aligns closely with our discussion in Section 3.2 on the type of innovation undertaken by our treated firms. Anecdotal evidence suggests that the patent applications in our treated sample largely reflect incremental, medium-technology innovations. Such innovations appear to stimulate production, in turn driving an expansion of technical and production

staff.

A limitation of our occupational composition analysis is worth noting. As discussed earlier in this section, a large share of worker-quarter observations have missing occupation data, particularly before 2017, requiring substantial imputation. This imputation appears to generate an artificial dip in mean employment across nearly all occupations around 2015–2016—the pre-treatment period for the 2018–2019 cohorts that dominate our occupation regressions. During these years, workers in several occupations appear to become temporarily concentrated in a single occupation code at many firms, before redistributing back to their original occupations by 2017–2018 as occupation data become more accurate. This manifests as a fall in mean log employment across occupations toward $t = -12$, as seen in Appendix Figure C9.

We do not view this as a major concern for our analysis. First, the pattern is common to both treated and control firms and affects them similarly, as shown in Appendix Figure C9. Second, it is confined to the pre-treatment period and does not affect the post-treatment coefficient means reported in Figure 10.

7 Conclusion

This paper examines how innovation affects firms’ workforce dynamics, using rich matched employer–employee administrative data from Türkiye. We treat firms that apply for a patent as innovative, and compare their outcomes to those of similar non-innovative firms.

Our first finding is that patent-based innovation in Türkiye is associated with employment and wage growth at the firm-level rather than worker displacement: treated firms grow larger and pay higher average wages than matched controls. This pattern fits the incremental, medium-technology character of the innovation in our sample, which is dominantly concentrated in manufacturing and appears to expand production rather than replace workers.

Second, looking at how these gains are distributed across the wage distribution, we find that employment rises almost everywhere, with the largest increase among workers earning between 1.5 and 2 times the minimum wage, just above the median. So the average wage increase at innovative firms is not driven by a narrow group at the top. It rather reflects a broader employment increase concentrated in the middle-to-upper part of the distribution.

Third, and central to the paper, we use worker transition matrices to uncover the underlying mechanisms. Relative to comparable non-innovative firms, we find that (i) innovative firms retain incumbent workers at higher rates across all wage bins, with the gap widest at the top; (ii) among stayers, workers at innovative firms are more likely to move up to a higher wage bin; and (iii) innovative firms hire substantially more new workers,

disproportionately into higher wage bins. Together, these three channels account for the rise in both average wages and employment.

Fourth, the occupational results support the same story. Employment growth is concentrated among technicians and associate professionals, craft and trades workers, and plant and machine operators. These are occupations one would expect to expand when incremental product and process innovation increases production.

These results indicate that supporting innovation, and patenting activity in particular, has additional benefits in the context we study, beyond its usual contribution to growth and development through innovation itself. First, they suggest that supporting early-stage patenting activity in a middle-income economy where innovation is dominantly medium-technology, incremental, and concentrated in manufacturing is unlikely to cause labor displacement at the firm level. If anything, it seems to help firms grow and retain their workers. Second, since employment gains are concentrated among middle-wage workers, innovation support programs do not make the wage distribution less equitable.

References

- Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier.
- Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244.
- Aghion, P., Akcigit, U., Hyytinen, A., and Toivanen, O. (2018). On the returns to invention within firms: Evidence from Finland. In *AEA Papers and Proceedings*, volume 108, pages 208–212. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.
- Aghion, P., Bergeaud, A., Blundell, R. W., and Griffith, R. (2019). The innovation premium to soft skills in low-skilled occupations. Working Paper 3489777, SSRN.
- Alam, A. and Diegmann, A. (2026). Patents, firm rents, and worker compensation: Causal evidence from quasi-random patent allocation. Working Paper 12666, CESifo.
- Balasubramanian, N. and Sivadasan, J. (2011). What happens when firms patent? New evidence from US economic census data. *The Review of Economics and Statistics*, 93(1):126–146.
- Bessen, J., Goos, M., Salomons, A., and Van den Berge, W. (2025). Automatic reaction—what happens to workers at firms that automate? *The Review of Economics and Statistics*, 107(1):125.
- Bogliacino, F., Guarascio, D., and Cirillo, V. (2018). The dynamics of profits and wages: Technology, offshoring and demand. *Industry and Innovation*, 25(8):778–808.
- Bordot, F. (2022). Artificial intelligence, robots and unemployment: Evidence from OECD countries. *Journal of Innovation Economics & Management*, 37(1):117–138.
- Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *Review of Economic Studies*, 91(6):3253–3285.
- Callaway, B. and Sant’Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.
- Castillo, V., Maffioli, A., Rojo, S., and Stucchi, R. (2014). The effect of innovation policy on smes’ employment and wages in Argentina. *Small Business Economics*, 42(2):387–406.
- Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, 134(3):1405–1454.

- Chen, J. and Roth, J. (2024). Logs with zeros? Some problems and solutions. *The Quarterly Journal of Economics*, 139(2):891–936.
- Chennells, L. and Van Reenen, J. (1999). Has technology hurt less skilled workers? A survey of the micro-econometric evidence. Working Paper W99/27, Institute for Fiscal Studies.
- Cirera, X. and Martins-Neto, A. S. (2023). Do innovative firms pay higher wages? Micro-level evidence from Brazil. *Research Policy*, 52(1):104645.
- Cirillo, V. (2014). Patterns of innovation and wage distribution. Do “innovative firms” pay higher wages? Evidence from Chile. *Eurasian Business Review*, 4(2):181–206.
- Czarnitzki, D. and Hussinger, K. (2004). The link between R&D subsidies, R&D spending and technological performance. Working Paper 04-56, ZEW discussion papers.
- Damioli, G., Van Roy, V., Vértesy, D., and Vivarelli, M. (2024). Drivers of employment dynamics of AI innovators. *Technological Forecasting and Social Change*, 201:123249.
- Dauth, W., Findeisen, S., Suedekum, J., and Woessner, N. (2021). The adjustment of labor markets to robots. *Journal of the European Economic Association*, 19(6):3104–3153.
- Desyllas, P. and Hughes, A. (2010). Do high technology acquirers become more innovative? *Research Policy*, 39(8):1105–1121.
- Domini, G., Grazzi, M., Moschella, D., and Treibich, T. (2021). Threats and opportunities in the digital era: Automation spikes and employment dynamics. *Research Policy*, 50(7):104137.
- Domini, G., Grazzi, M., Moschella, D., and Treibich, T. (2022). For whom the bell tolls: The firm-level effects of automation on wage and gender inequality. *Research Policy*, 51(7):104533.
- Garcia, R. d. C., Araujo, V., and Ferreira, S. (2026). How R&D grant size affects innovation in small firms. *Small Business Economics*, pages 1–27.
- Graetz, G. and Michaels, G. (2018). Robots at work. *Review of Economics and Statistics*, 100(5):753–68.
- Griliches, Z. (1990). Patent statistics as economic indicators: A survey. *Journal of Economic Literature*, 28(4):1661–1707.
- Grinza, E. and Quatraro, F. (2019). Workers’ replacements and firms’ innovation dynamics: New evidence from Italian matched longitudinal data. *Research Policy*, 48(9):103804.

- Harrison, R., Jaumandreu, J., Mairesse, J., and Peters, B. (2014). Does innovation stimulate employment? A firm-level analysis using comparable micro-data from four European countries. *International Journal of Industrial Organization*, 35:29–43.
- Humlum, A. (2019). Robot adoption and labor market dynamics. *Princeton University*.
- Kline, P., Petkova, N., Williams, H., and Zidar, O. (2019). Who profits from patents? Rent-sharing at innovative firms. *The Quarterly Journal of Economics*, 134(3):1343–1404.
- Kogan, L., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2023). Technology and labor displacement: Evidence from linking patents with worker-level data. Working Paper 31846, National Bureau of Economic Research.
- Kogan, L., Papanikolaou, D., Seru, A., and Stoffman, N. (2017). Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics*, 132(2):665–712.
- Lachenmaier, S. and Rottmann, H. (2011). Effects of innovation on employment: A dynamic panel analysis. *International Journal of Industrial Organization*, 29(2):210–220.
- Machin, S. and Van Reenen, J. (1998). Technology and changes in skill structure: Evidence from seven OECD countries. *The Quarterly Journal of Economics*, 113(4):1215–1244.
- Meschi, E., Taymaz, E., and Vivarelli, M. (2016). Globalization, technological change and labor demand: A firm-level analysis for Turkey. *Review of World Economics*, 152(4):655–680.
- Pianta, M. and Tancioni, M. (2008). Innovations, wages, and profits. *Journal of Post Keynesian Economics*, 31(1):101–123.
- Toivanen, O. and Väänänen, L. (2012). Returns to inventors. *Review of Economics and Statistics*, 94(4):1173–1190.
- Van Reenen, J. (1996). The creation and capture of rents: Wages and innovation in a panel of UK companies. *The Quarterly Journal of Economics*, 111(1):195–226.
- Van Reenen, J. (1997). Employment and technological innovation: Evidence from UK manufacturing firms. *Journal of Labor Economics*, 15(2):255–284.
- Van Roy, V., Vértesy, D., and Vivarelli, M. (2018). Technology and employment: Mass unemployment or job creation? Empirical evidence from European patenting firms. *Research policy*, 47(9):1762–1776.


A Why Stacked Difference-in-Differences?

Here we explain why we rely on the stacked difference-in-differences approach rather than alternative staggered-adoption estimators, such as Callaway and Sant’Anna (2021) or Borusyak et al. (2024).

Recall that matching is implemented based on firms’ attributes prior to their respective quarter of treatment. This means that firms belonging to different treatment cohorts are matched to their control firms in different calendar quarters. Consequently, if we were to use an estimator based on comparisons across calendar quarters, such as Callaway and Sant’Anna (2021) or Borusyak et al. (2024), we would be making undesirable comparisons.

This issue is illustrated in the left-hand-side of Figure A1. Cohort 1 is matched to its control group in calendar time 3 (labeled $M1$), with treatment starting in time 4, while Cohort 2 is matched to its control group in calendar time 5 (labeled $M2$), with treatment starting in time 6. Suppose we want to compare outcomes just before and just after treatment. We would then, for example, compare Treated-1’s outcomes before and after treatment, i.e., calendar times 3 and 4, to those of Control-1 in the same quarters. This comparison is valid, since Treated-1 and Control-1 are matched at calendar time 3. However, under a calendar-time comparison, we would also compare Treated-1 to Control-2 at calendar times 3 and 4. This latter comparison is inconsistent with our matching procedure (thus, undesirable), since times 3 and 4 do not correspond to the (hypothetical) just-before and just-after treatment periods for Control-2. The unit Control-2 is matched at calendar time 5 and thus calendar time 5 correspond to its hypothetical pre-treatment time.

Figure A1: From Calendar Time to Treatment Time: The Stacked DiD Approach



Notes: This figure is illustrative and intended to motivate the stacked difference-in-differences approach. $M1$ and $M2$ denote the calendar quarters in which Cohort 1 and Cohort 2 are matched to their respective control groups (Treated-1 to Control-1, Treated-2 to Control-2). The left panel shows $M1$ and $M2$ falling in different calendar quarters; the right panel shows both aligning at $t = -1$ once time is re-indexed relative to treatment.

The stacked difference-in-differences method resolves this issue by comparing treated and control units across quarters defined relative to treatment, $t \in \{-24, -23, \dots, 0, \dots, 15\}$, rather than across calendar quarters. This re-indexing is illustrated in the right-hand-side

of Figure A1: once calendar time is replaced with treatment time, the matching quarter for every cohort falls at the same relative position, $t = -1$, immediately preceding treatment at $t = 0$. This allows us to make comparisons relative to the treatment quarter, as intended by our matching procedure: all treated units, just before and just after treatment, are now compared to control units at the same relative quarters, that is, the quarter in which they are matched and the quarter immediately following it.

B Transition Matrices

Table B1: Incumbent Workers' Transition Matrix ($t = -12$ to $t = 0$): Employment Outcomes by Pre-Anticipation Wage Bin

	Treated (N = 197,619)							Control (N = 192,957)						
	Bin 1 37,576	Bin 2 53,233	Bin 3 34,119	Bin 4 21,695	Bin 5 20,285	Bin 6 16,385	Bin 7 14,326	Bin 1 42,945	Bin 2 44,185	Bin 3 32,155	Bin 4 19,109	Bin 5 21,204	Bin 6 17,703	Bin 7 15,656
N														
Nonemployment (% of Bin Size)	30.36	21.76	17.38	16.05	15.54	16.31	17.04	34.82	24.91	18.81	18.82	19.95	21.41	29.29
<i>Same Firm (% of Bin Size)</i>														
Bin 1	19.71	5.68	1.08	0.19	0.08	0.05	0.03	19.12	5.49	0.74	0.30	0.14	0.10	0.08
Bin 2	14.28	28.44	10.80	2.61	0.58	0.28	0.09	11.18	25.08	11.60	7.11	0.81	0.16	0.10
Bin 3	4.86	13.60	25.97	14.99	5.77	0.66	0.31	3.70	13.67	27.62	13.20	4.34	0.84	0.20
Bin 4	1.31	3.13	13.27	25.40	14.80	2.09	0.61	0.99	2.19	10.66	19.80	12.35	2.55	0.26
Bin 5	0.68	2.27	6.58	16.64	29.67	14.30	1.58	0.55	1.10	3.89	13.32	30.35	14.21	0.74
Bin 6	0.19	0.32	1.49	3.52	12.55	38.93	12.54	0.13	0.20	0.95	3.88	11.54	38.47	18.52
Bin 7	0.06	0.08	0.30	0.56	1.69	8.98	52.58	0.03	0.09	0.33	0.39	0.95	5.15	34.41
<i>Different Firm (% of Bin Size)</i>														
Bin 1	15.81	9.62	6.31	4.25	3.30	1.98	1.06	16.63	10.05	6.74	5.44	3.99	2.34	1.26
Bin 2	7.61	8.21	5.79	3.57	2.33	1.11	0.60	7.55	9.27	7.05	4.26	2.63	1.33	0.71
Bin 3	3.16	4.19	5.51	4.34	2.75	1.37	0.76	2.96	4.61	6.28	4.85	3.00	1.43	0.75
Bin 4	1.05	1.48	2.72	3.23	2.67	1.26	0.65	1.29	1.97	2.94	4.25	3.05	1.53	0.97
Bin 5	0.64	0.89	2.00	2.86	4.19	3.19	1.25	0.76	1.00	1.72	2.83	3.56	2.86	2.10
Bin 6	0.23	0.27	0.62	1.47	3.20	6.57	2.92	0.24	0.30	0.55	1.27	2.62	5.29	3.34
Bin 7	0.05	0.08	0.16	0.33	0.87	2.91	7.99	0.06	0.06	0.10	0.27	0.71	2.31	7.28
<i>Chi-Squared Test: Treated vs. Control (by Bin)</i>														
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00							

Notes: This table reports the destination of incumbent employees as of the quarter immediately following treatment ($t = 0$), conditional on the wage bin in which they were employed twelve quarters before treatment ($t = -12$, prior to anticipation effects), separately for employees of treated and control firms. The number in parentheses next to each group label (Treated, Control) is the total number of employees in that group, irrespective of the bin in which they were located at $t = -12$. Each column corresponds to a starting wage bin at $t = -12$; N denotes the number of employees in that bin (column) at $t = -12$, and all percentage rows are expressed as a percentage of this bin size. By $t = 0$, employees are categorized as nonemployed, employed at the same firm in a given wage bin (*Same Firm*), or employed at a different firm in a given wage bin (*Different Firm*); within a given column, these percentages sum to 100. The bottom row reports, separately for each pre-anticipation wage bin, the p-value of a chi-squared test comparing the (absolute) distribution of outcomes between treated and control employees who started in that bin.

Table B2: Incumbent Workers' Transition Matrix ($t = -12$ to $t = 15$): Employment Outcomes by Pre-Anticipation Wage Bin

	Treated (N = 197,619)							Control (N = 192,957)						
	Bin 1	Bin 2	Bin 3	Bin 4	Bin 5	Bin 6	Bin 7	Bin 1	Bin 2	Bin 3	Bin 4	Bin 5	Bin 6	Bin 7
N	37,576	53,233	34,119	21,695	20,285	16,385	14,326	42,945	44,185	32,155	19,109	21,204	17,703	15,656
Nonemployment (% of Bin Size)	41.40	33.13	30.39	29.84	30.41	34.07	36.22	42.81	36.20	30.16	30.42	34.08	37.45	51.28
<i>Same Firm (% of Bin Size)</i>														
Bin 1	8.33	2.46	0.73	0.16	0.09	0.04	0.01	8.89	2.68	0.43	0.19	0.08	0.04	0.06
Bin 2	9.88	14.58	6.67	2.10	0.98	0.24	0.04	8.23	13.33	7.13	4.82	0.95	0.18	0.09
Bin 3	4.17	10.81	13.09	9.39	4.73	0.70	0.10	2.65	7.15	12.37	8.02	3.77	1.06	0.12
Bin 4	1.47	3.85	8.91	12.71	8.06	2.80	0.29	0.78	3.06	7.38	7.98	6.23	2.88	0.20
Bin 5	0.85	2.33	7.24	12.65	16.43	9.84	1.26	0.51	1.73	5.84	10.71	15.90	9.70	1.14
Bin 6	0.39	0.59	2.52	5.00	11.48	21.03	23.43	0.15	0.61	2.42	5.73	11.22	21.69	16.45
Bin 7	0.06	0.10	0.46	0.93	2.81	6.90	20.80	0.04	0.13	0.28	0.81	1.90	4.93	13.86
<i>Different Firm (% of Bin Size)</i>														
Bin 1	15.40	10.39	6.95	4.90	3.73	2.42	1.07	17.00	10.67	7.86	6.49	4.35	2.55	1.12
Bin 2	9.06	9.33	6.59	4.58	2.88	1.86	0.84	9.68	11.52	9.29	5.87	4.22	2.68	1.10
Bin 3	4.61	6.21	6.69	5.60	3.82	2.19	0.79	4.62	6.74	7.54	5.94	4.13	2.08	0.87
Bin 4	2.23	3.02	3.87	4.01	3.26	2.15	0.76	2.33	3.06	4.08	4.74	3.19	1.87	0.75
Bin 5	1.35	2.11	3.64	4.14	4.42	3.48	1.47	1.42	1.98	3.31	4.35	4.23	3.33	1.19
Bin 6	0.67	0.85	1.64	2.78	4.41	6.16	3.14	0.73	0.90	1.47	2.92	3.90	4.96	3.11
Bin 7	0.13	0.23	0.61	1.18	2.50	6.13	9.78	0.15	0.24	0.45	0.99	1.86	4.62	8.66
<i>Chi-Squared Test: Treated vs. Control (by Bin)</i>														
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00							

Notes: This table reports the destination of incumbent employees as of the last treatment quarter ($t = 15$), conditional on the wage bin in which they were employed twelve quarters before treatment ($t = -12$, prior to anticipation effects), separately for employees of treated and control firms. The number in parentheses next to each group label (Treated, Control) is the total number of employees in that group, irrespective of the bin in which they were located at $t = -12$. Each column corresponds to a starting wage bin at $t = -12$; N denotes the number of employees in that bin (column) at $t = -12$, and all percentage rows are expressed as a percentage of this bin size. By $t = 15$, employees are categorized as nonemployed, employed at the same firm in a given wage bin (*Same Firm*), or employed at a different firm in a given wage bin (*Different Firm*); within a given column, these percentages sum to 100. The bottom row reports, separately for each pre-anticipation wage bin, the p-value of a chi-squared test comparing the (absolute) distribution of outcomes between treated and control employees who started in that bin.

Table B3: New Workers' Transition Matrix ($t = 0$): Origins at $t = -12$ by Contemporaneous Wage Bin

	Treated (N = 111,749)							Control (N = 92,543)						
	Bin 1	Bin 2	Bin 3	Bin 4	Bin 5	Bin 6	Bin 7	Bin 1	Bin 2	Bin 3	Bin 4	Bin 5	Bin 6	Bin 7
N	29,633	33,724	20,327	10,870	8,564	5,615	3,016	23,907	28,348	19,081	7,346	7,000	4,392	2,469
%	26.52	30.18	18.19	9.73	7.66	5.02	2.70	25.83	30.63	20.62	7.94	7.56	4.75	2.67
Nonemployment (% of Bin Size)	60.27	52.25	46.82	44.01	40.50	38.49	41.81	64.46	54.96	47.20	43.67	40.20	43.65	41.15
<i>Different Firm (% of Bin Size)</i>														
Bin 1	26.51	25.11	21.67	18.70	13.46	7.32	3.32	23.59	23.42	22.04	18.98	15.17	7.79	3.12
Bin 2	8.68	14.41	14.60	12.68	9.32	4.77	1.72	8.05	13.55	16.64	14.08	12.40	4.80	1.74
Bin 3	2.78	5.43	9.86	10.83	10.23	5.56	1.96	2.40	4.99	8.54	10.45	10.51	6.03	2.11
Bin 4	0.88	1.65	3.97	6.64	8.31	6.46	2.49	0.82	1.70	3.14	6.22	7.47	6.01	2.39
Bin 5	0.56	0.85	2.19	5.14	10.78	11.90	5.04	0.46	0.98	1.66	4.53	9.01	10.86	5.22
Bin 6	0.23	0.22	0.72	1.52	6.12	18.22	12.77	0.18	0.30	0.58	1.72	4.00	14.80	13.20
Bin 7	0.08	0.09	0.18	0.48	1.28	7.28	30.90	0.04	0.11	0.21	0.35	1.23	6.06	31.07
<i>Chi-Squared Test: Treated vs. Control (by Bin)</i>														
p-value	0.00	0.00	0.00	0.04	0.00	0.00	1.00							

Notes: This table reports the origins of new workers at $t = 0$, defined as employees who are observed at the treated or control firm at $t = 0$ but were not present at that firm twelve quarters earlier ($t = -12$, prior to anticipation effects). Columns correspond to the wage bin of the new worker at $t = 0$; N denotes the total number of such workers in each bin, and % denotes that bin's share of all new workers within the treated or control group. Row entries give the percentage of workers in a given $t = 0$ wage bin who, at $t = -12$, were nonemployed or employed at a different firm in the indicated wage bin (*Different Firm*); these percentages sum to 100 within each column. By construction, same-firm origins are excluded since workers in this table were not employed at the firm at $t = -12$. The number in parentheses next to each group label (Treated, Control) is the total number of new workers across all bins. The bottom row reports, separately for each $t = 0$ wage bin, the p-value of a chi-squared test comparing the (absolute) distribution of origins between treated and control new workers in that bin.


Table B4: New Workers' Transition Matrix ($t = 15$): Origins at $t = -12$ by Contemporaneous Wage Bin

	Treated (N = 176,310)							Control (N = 132,055)						
	Bin 1	Bin 2	Bin 3	Bin 4	Bin 5	Bin 6	Bin 7	Bin 1	Bin 2	Bin 3	Bin 4	Bin 5	Bin 6	Bin 7
N	33,257	49,844	39,916	20,740	16,815	10,074	5,664	25,809	45,216	24,678	12,033	12,985	7,395	3,939
%	18.86	28.27	22.64	11.76	9.54	5.71	3.21	19.54	34.24	18.69	9.11	9.83	5.60	2.98
Nonemployment (% of Bin Size)	69.67	62.42	55.57	53.26	49.58	46.98	38.29	70.64	61.38	55.22	53.20	53.69	46.13	38.87
<i>Different Firm (% of Bin Size)</i>														
Bin 1	20.03	20.97	19.77	17.53	14.81	9.99	5.54	18.98	20.12	20.07	16.82	14.15	10.56	4.27
Bin 2	6.74	10.43	12.58	11.56	10.54	7.33	3.27	7.16	12.10	12.31	11.77	10.17	8.34	2.64
Bin 3	2.16	3.96	7.03	8.39	8.65	6.94	4.15	1.94	4.12	6.92	8.71	8.41	8.05	3.45
Bin 4	0.78	1.35	2.66	4.34	6.20	6.11	4.10	0.67	1.32	3.09	4.41	5.14	5.94	3.63
Bin 5	0.43	0.63	1.70	3.10	5.88	8.48	7.91	0.41	0.71	1.76	3.45	4.74	8.60	7.01
Bin 6	0.14	0.18	0.56	1.39	3.50	9.67	15.22	0.14	0.20	0.53	1.32	2.83	8.48	15.21
Bin 7	0.05	0.05	0.12	0.41	0.83	4.51	21.52	0.06	0.06	0.11	0.32	0.87	3.91	24.93
<i>Chi-Squared Test: Treated vs. Control (by Bin)</i>														
p-value	0.01	0.00	0.09	0.28	0.00	0.00	0.00							

Notes: This table reports the origins of new workers at $t = 15$, defined as employees who are observed at the treated or control firm at $t = 15$ but were not present at that firm twelve quarters earlier ($t = -12$, prior to anticipation effects). Columns correspond to the wage bin of the new worker at $t = 15$; N denotes the total number of such workers in each bin, and % denotes that bin's share of all new workers within the treated or control group. Row entries give the percentage of workers in a given $t = 15$ wage bin who, at $t = -12$, were nonemployed or employed at a different firm in the indicated wage bin (*Different Firm*); these percentages sum to 100 within each column. By construction, same-firm origins are excluded since workers in this table were not employed at the firm at $t = -12$. The number in parentheses next to each group label (Treated, Control) is the total number of new workers across all bins. The bottom row reports, separately for each $t = 15$ wage bin, the p-value of a chi-squared test comparing the (absolute) distribution of origins between treated and control new workers in that bin.


C Additional Figures

Figure C1: Innovation effects on firm size and mean wage (rejected patent applications)



Notes: The sample consists of 768 treated and 768 control firms, yielding 61,440 firm-quarter observations. Treated firms are those whose patent application was rejected, filed between 2012 Q1 and 2020 Q1, and that continuously report worker registers every quarter for 24 quarters before and 16 quarters after the patent application (dynamically balanced panel). Control firms are obtained from a hybrid matching procedure. Further details on the construction of the treatment and control groups are provided in Section 3. Event-study coefficients in the left panels are estimated using the stacked difference-in-differences method introduced in Section 4.1. The right panels report the evolution of raw means of outcomes (firm size and log daily wage) for the treatment (solid blue line) and control (dashed red line) groups.


Figure C2: Innovation effects on firm size and mean wage (granted patent applications)




Notes: The sample consists of 1,003 treated and 1003 control firms, yielding 80,240 firm-quarter observations. Treated firms are those whose patent application was granted, filed between 2012 Q1 and 2020 Q1, and that continuously report worker registers every quarter for 24 quarters before and 16 quarters after the patent application (dynamically balanced panel). Control firms are obtained from a hybrid matching procedure. Further details on the construction of the treatment and control groups are provided in Section 3. Event-study coefficients in the left panels are estimated using the stacked difference-in-differences method introduced in Section 4.1. The right panels report the evolution of raw means of outcomes (firm size and log daily wage) for the treatment (solid blue line) and control (dashed red line) groups.

Figure C3: Innovation effects on log employment across remaining wage bins

Bin 2



Bin 4



Bin 5





Figure C3: Innovation effects on log employment across remaining wage bins (continued)

Bin 7




Notes: The sample consists of 1,771 treated and 1,771 control firms, yielding 141,680 firm-quarter observations. Treated firms are those that applied for a patent between 2012 Q1 and 2020 Q1 and continuously report worker registers every quarter for 24 quarters before and 16 quarters after the patent application (dynamically balanced panel). Control firms are obtained from a hybrid matching procedure. Further details on the construction of the treatment and control groups, as well as on the wage bins (Bin 1–Bin 7), are provided in Section 3. Event-study coefficients in the left panels are estimated using the stacked difference-in-differences method introduced in Section 4.1. The right panels report the evolution of raw mean log bin sizes for the treatment (solid blue line) and control (dashed red line) groups.

Figure C4: Log average change in employment across wage bins 9-16 quarters after innovation: granted vs. rejected patent applications




Notes: Bar heights represent the average log change in employment for each wage bin in innovative firms whose patent application was either rejected or granted, relative to similar non-innovative firms, calculated as the mean of the last eight point estimates (i.e., $t = 8$ to $t = 15$, or the 9th to 16th quarters after treatment) from the corresponding event-study specifications for the rejected- and granted-applications subsamples, respectively (not reported). Red bars correspond to firms whose patent application was rejected; blue bars correspond to firms whose patent application was granted.

Figure C5: Worker Transitions from $t = -12$ to $t = 0$ by Wage Bin



Notes: This figure shows, separately by wage bin, the fraction of workers employed at treated and control firms at $t = -12$ (immediately before the anticipation effect) who transition into *Nonemployment*, remain at the *Same Firm*, or move to a *Different Firm* by $t = 0$ (time of treatment). Within each wage bin, the three fractions sum to one. Wage bins (Bin 1–Bin 7) are defined in Section 3. Red solid lines with circles denote treated firms; blue dashed lines with squares denote control firms. The associated full transition matrix is presented in Appendix Table B1.

Figure C6: Treated–Control Differences in Wage Bin Transitions ($t = -12$ to $t = 0$) Among Stayers




Notes: This figure focuses exclusively on workers who remained at the same firm between $t = -12$ and $t = 0$ (“stayers”). Wage bins (Bin 1–Bin 7) are defined in Section 3. The cell in row i , column j reports the difference, in percentage points, between (i) the fraction of treated stayers who started in bin j at $t = -12$ that ended in bin i at $t = 0$, and (ii) the analogous fraction for control stayers who started in bin j (Treated minus Control). Blue cells indicate positive differences (treated stayers are relatively more likely to make that transition than control stayers), while red cells indicate negative differences (treated stayers are relatively less likely to make that transition). Diagonal cells, highlighted with black borders, correspond to workers who remained in the same wage bin between $t = -12$ and $t = 0$. The associated full transition matrix is presented in Appendix Table B1.

Figure C7: Destination of New Hires by $t = 0$, by Wage Bin



Notes: This figure shows the wage bin destination of workers newly hired at treated and control firms by $t = 0$, defined as workers present at $t = 0$ but not employed at the firm at $t = -12$. The left panel plots the log number of new hires in each wage bin. The right panel plots the number of new hires in each wage bin relative to the number of workers employed in that bin at $t = -12$. Wage bins (Bin 1–Bin 7) are defined in Section 3. Red solid lines with circles denote treated firms; blue dashed lines with squares denote control firms. The associated full transition matrix is presented in Appendix Table B3.


Figure C8: Source of New Hires by $t = 0$, by Wage Bin




Notes: This figure shows, separately by wage bin, the source of workers newly present at treated and control firms by $t = 0$ (i.e., workers present at $t = 0$ but not employed at the firm at $t = -12$). The left panel plots the fraction of these new hires who were in *Nonemployment* at $t = -12$; the right panel plots the fraction who were employed at a *Different Firm* at $t = -12$. Within each wage bin, the two fractions sum to one. Wage bins (Bin 1–Bin 7) are defined in Section 3. Red solid lines with circles denote treated firms; blue dashed lines with squares denote control firms. The associated full transition matrix is presented in Appendix Table B3.

Figure C9: Innovation effects on employment across occupations


ISCO 01: Managers



ISCO 02: Professionals



ISCO 03: Technicians & Assoc. Prof.



ISCO 04: Clerical Support Workers






Figure C9: Innovation effects on employment across occupations (continued)


ISCO 05: Service & Sales Workers



ISCO 06: Skilled Agr., Forestry % Fishery



ISCO 07: Craft & Trades Workers



ISCO 08: Plant & Machine Operators





Figure C9: Innovation effects on employment across occupations (continued)

ISCO 09: Elementary Occupations



Notes: The sample consists of 587 treated and 587 control firms, yielding 46,960 firm-quarter observations. Treated firms are those that applied for a patent between 2018 Q1 and 2020 Q1 and continuously report worker registers every quarter for 24 quarters before and 16 quarters after the patent application (dynamically balanced panel). Control firms are obtained from a hybrid matching procedure. Further details on the construction of the treatment and control groups are provided in Section 3. Occupations are classified according to the one-digit ISCO-08 classification. Event-study coefficients in the left panels are estimated using the stacked difference-in-differences method introduced in Section 4.1. The right panels report the evolution of raw mean log employment by occupation for the treatment (solid blue line) and control (dashed red line) groups.