

First version: July, 2026
This version: September, 2026



CAF-WORKINGPAPER
#2026/11

Automation and the Margins of Firm and Worker Adjustment in Argentina

Irene Brambilla*** | Octavio Bertín* | Jessica Bracco* | Guillermo Falcone* | Azul Menduiña*

*CEDLAS-IIE-UNLP.
**CONICET.

Irene Brambilla:
irene.brambilla@gmail.com

Octavio Bertín:
octaviobertin2001@gmail.com

Jessica Bracco:
braccojessica@gmail.com

Guillermo Falcone:
guillermofalcone@gmail.com

Azul Menduiña:
azulmenduina@gmail.com

CAF - WORKING PAPER #2026/11

This version: September 11, 2026

Automation and the Margins of Firm and Worker Adjustment in Argentina

Irene Brambilla^{1,2} Octavio Bertín¹ Jessica Bracco¹ Guillermo Falcone¹
Azul Menduiña¹

Abstract

We examine how automation reshapes firm labor demand and incumbent-worker trajectories in Argentina over 2007–2022. We link customs records on automation-capital imports to administrative employer–employee and simplified self-employment tax data, identify adoption as discrete import spikes, and estimate stacked event studies comparing adopters and their incumbents with not-yet-adopting firms. Five years after adoption, employment is 19 percent lower and the wage bill 20 percent lower, while firm-level wages barely move; the contraction runs through hiring rather than dismissals. Incumbents’ reported earnings stay flat, but their formal attachment weakens—fewer months in registered wage employment, more exits from the registry, and more transitions into simplified self-employment. Automation costs here appear less as payroll wage cuts than as separations, reallocation, and movement out of registered wage employment.

JEL codes: J23, J31, J46, J63, O33

Keywords: automation, worker displacement, firm labor demand, formal employment, self-employment, Argentina

¹CEDLAS-IEE-UNLP. ²CONICET.

Small sections of text that are less than two paragraphs may be quoted without explicit permission as long as this document is acknowledged. Findings, interpretations and conclusions expressed in this publication are the sole responsibility of its author(s) and cannot be, in any way, attributed to CAF, its Executive Directors or the countries they represent. CAF does not guarantee the accuracy of the data included in this publication and is not, in any way, responsible for any consequences resulting from its use.

©2022 Corporación Andina de Fomento

We thank Pablo Brassiolo and seminar participants at the CAF Workshop for helpful comments and suggestions. We gratefully acknowledge financial support from CAF – Development Bank of Latin America and the Caribbean – for this project. The views expressed herein are those of the authors and do not necessarily reflect those of CAF. All remaining errors are our own.

CAF - DOCUMENTO DE TRABAJO #2026/11

Esta versión: 11 de septiembre de 2026

Automatización y los márgenes de ajuste de empresas y trabajadores en Argentina

Irene Brambilla^{1,2} Octavio Bertín¹ Jessica Bracco¹ Guillermo Falcone¹
Azul Menduiña¹

Abstract

Examinamos cómo la automatización transforma la demanda laboral de las empresas y las trayectorias de los trabajadores incumbentes en Argentina entre 2007 y 2022. Vinculamos registros aduaneros de importaciones de bienes de capital de automatización con datos administrativos de empleo registrado y del régimen simplificado de trabajo independiente (monotributo), identificamos la adopción como picos discretos de importación y estimamos estudios de eventos apilados. Cinco años después de la adopción, el empleo es 19 por ciento menor y la masa salarial 20 por ciento menor, mientras que los salarios apenas se mueven; la contracción opera vía contrataciones y no despidos. Los ingresos declarados de los incumbentes permanecen estables, pero su vínculo formal se debilita: menos meses en el empleo registrado, más salidas del registro y más transiciones al trabajo independiente simplificado. Los costos de la automatización aparecen menos como recortes salariales y más como separaciones y reasignación.

JEL codes: J23, J31, J46, J63, O33

Keywords: automatización, desplazamiento de trabajadores, demanda laboral, empleo formal, trabajo independiente, Argentina

¹CEDLAS-III-UNLP. ²CONICET.

Pequeñas secciones del texto, menores a dos párrafos, pueden ser citadas sin autorización explícita siempre que se cite el presente documento. Los resultados, interpretaciones y conclusiones expresados en esta publicación son de exclusiva responsabilidad de su(s) autor(es), y de ninguna manera pueden ser atribuidos a CAF, a los miembros de su Directorio Ejecutivo o a los países que ellos representan. CAF no garantiza la exactitud de los datos incluidos en esta publicación y no se hace responsable en ningún aspecto de las consecuencias que resulten de su utilización.

1 | INTRODUCTION

Automation can make firms more productive, more competitive, and in some cases larger. But it also reopens a fundamental question in labor economics: do workers at automating firms share in those gains, or do they bear the costs of adjustment? A large literature studies robots and automation at the industry, region, and firm levels ([Acemoglu and Restrepo, 2020](#); [Graetz and Michaels, 2018](#); [Koch et al., 2021](#); [Bonfiglioli et al., 2024](#)). Yet the central welfare question is ultimately about workers. Even when automation improves firm performance, incumbent workers may not benefit proportionally; they may instead bear the costs through separations, wage changes, or weaker labor-market attachment.

The effect of automation on workers presents an additional layer of concern in developing economies. In those settings, unemployment and wage reductions are not the only margins of adjustment; instead, workers may also move into informality, low-productivity self-employment, or intermittent participation, often with limited income support during the transition. Existing evidence for Latin America suggests that adverse shocks are often absorbed through transitions across employment states, with informal employment acting as a buffer for workers leaving formal jobs ([Brambilla et al., 2023](#); [Dix-Carneiro and Kovak, 2019](#); [Dix-Carneiro et al., 2021](#)). But that margin of adjustment need not be benign: it may come with lower earnings, weaker protection, and a more fragile attachment to stable formal jobs. A separation induced by automation may therefore trigger not just a temporary earnings loss, but a persistent transition to low-quality work. In that sense, automation may have much sharper distributional consequences in economies where informality is pervasive and a central margin of adjustment.

In this paper, we study how automation reshapes firm labor demand and worker trajectories in Argentina over 2007–2022. We combine three administrative sources: transaction-level customs data on imports of automation-related capital goods, linked employer–employee records from the Sistema Integrado Previsional Argentino (SIPA), the country’s formal social-security registry, and records from AFIP, Argentina’s tax authority, on monotributo, the country’s simplified self-employment tax regime. This allows us to identify firms experiencing directly observed, import-based automation events, measure their pre-event incumbent workforce, and then follow those workers across all subsequent formal employers. Rather than relying on regional or local labor market exposure measures, this matched employer–employee approach enables us to study the dynamics at the firm level, determining how automation affects employment relationships, reported earnings, worker reallocation, formal labor-market attachment, and subsequent transitions into self-employment, non-employment, or informality.

We define automation events as discrete spikes in firm imports of automation-related capital goods. In Argentina, where advanced automation equipment is primarily sourced from abroad, such spikes provide a useful observable proxy for large episodes of automation adoption. The capital goods that we consider are industrial robots, machining centers, numerically controlled machinery, automatic data-processing equipment, and industrial control instruments. We use a baseline definition of automation imports that includes all those categories and also a narrow definition based on industrial robots alone. Regardless of the particular definition used, an important caveat is that it captures investment in automation made by firms that import capital goods directly while it excludes firms that purchase automation equipment through integrators, distributors, or related firms ([Bessen et al., 2025](#)). Our event-study design, which we describe below, partly addresses this concern by comparing firms with different timing of observed direct-import adoption, rather than comparing direct importers with arbitrary non-importing firms. Indirect adoption before the observed spike remains possible, but if future direct importers have already begun adjusting

through third-party adoption, their outcomes would tend to move before treatment and make our estimates more conservative. In the next sections, we discuss these measurement issues in detail, including the restrictions we impose to limit the risk of capturing imports intended for resale. Against this background, we interpret our estimates as the effects of directly observed, import-based automation events, not as effects for the full universe of automation investments.

Our empirical strategy builds on a simple insight, also emphasized by [Bessen et al. \(2025\)](#): automation investment is lumpy. Firms do not adopt automation smoothly over time; they make large, discrete purchases of capital at identifiable moments. We therefore define an automation event as the first year in which a firm's automation imports, measured relative to its average wage bill, exceed three times their own average in other years, and we use the staggered timing of those first spikes in a stacked difference-in-differences event-study design ([Cengiz et al., 2019](#); [Wing et al., 2024](#)). This shifts identification away from cross-sectional differences between firms and toward the timing of discrete adoption episodes within a common event-time framework. Around each event, we observe outcomes for three years before the event and follow firms and incumbent workers for five years afterward. Given our 2007–2022 data coverage, this implies that the estimation sample includes events occurring between 2010 and 2016. At the worker level, we focus on workers with stable pre-event attachment to the firm and trace their trajectories across all subsequent formal-sector employers. Because the design compares adopters with firms that adopt later, what identification requires is not that adoption be random but that its *timing* be unpredictable among eventual adopters. We show that it largely is: recent changes in employment, wages, exports, hiring, and separations have essentially no power to predict whether a firm's first spike occurs this year rather than later.

The results point to five facts. First, registered labor demand at automating firms contracts sharply. Employment and the wage bill move in parallel with those of later adopters before the event and begin to fall after the first post-event year; by event year +5, employment is 19 percent lower and the wage bill 20 percent lower than at comparable firms that adopt only later. Second, that contraction runs through a reduction in hiring rather than through dismissals. Separating the two gross flows, we find that hires fall by 21 points of pre-event employment by event year +5, while separations never rise significantly above the control path at any horizon. The gap opens because adopters stop replacing the workers who leave, not because they shed the workers they have—so a substantial part of the incidence falls on the workers the firm never hires, who are invisible in any analysis based on incumbents alone. Third, firm-level wage outcomes barely change: average and median wages end the window less than one percent below the control path, are statistically indistinguishable from zero, and are much smaller than the 19 percent employment decline. The wage-bill contraction is thus a quantity adjustment, not a price adjustment. Fourth, for incumbent workers the reported-wage margin likewise shows no robust response: total annual earnings and average monthly wages, both across all registered wage jobs and at the original employer, remain flat and imprecisely estimated. Fifth, worker adjustment appears instead on the attachment margin. By the end of the window, incumbent workers spend 0.14 fewer months per year in registered wage employment and are 1.0 percentage point more likely to have left the registry—both distinguishable from zero only at the 10 percent level—while the sharpest and most precisely estimated margin is a 0.49 percentage-point rise in the probability of being observed only in the simplified self-employment regime. Scaled to a cohort of 1,000 incumbents, the estimates imply 25 additional departures from the original employer by event year +5, of which 15 move to another registered wage job and 10 leave formal wage employment altogether—half of them reappearing in simplified self-employment and half no longer observed in any of our records. These effects are modest

in levels, but persistent: automation does not generate mass exit from the formal system, but it moves a small non-negligible share of incumbents away from standard registered wage employment.

Three further findings qualify this average picture. First, the near-zero average earnings response is not uniform: incumbents of small adopters see annual earnings fall by 4 percent—the only group in our data with a clear earnings decline—while at large adopters adjustment runs through separations and transitions into simplified self-employment instead. Reported earnings are nowhere the principal margin of adjustment; where they respond at all, it is in one corner of the firm-size distribution. Second, displacement is concentrated at the top of the pre-event wage distribution rather than the bottom: the probability of leaving the original employer rises monotonically across wage quartiles, reaching 6 percentage points in the highest, while incumbents in the lowest quartile show earnings gains. Third, the results are insensitive to the spike threshold that defines an event, but not to the technology that does: when we narrow the basket of automation goods to industrial robots alone, the worker-level displacement margins disappear and incumbent earnings drift upward instead. The robots-only estimates rest on far fewer adopters and are correspondingly imprecise, but the contrast suggests that automation is a bundle of technologies whose consequences for incumbents differ, and that the displacement we document is a property of broad automation adoption rather than of robot adoption narrowly defined.

Our paper is most closely related to [Bonfiglioli et al. \(2024\)](#) and [Bessen et al. \(2025\)](#). [Bonfiglioli et al. \(2024\)](#) study the firm-level effects of robot imports and emphasize that demand shocks can make robot imports and employment positively correlated: expanding firms are more likely to automate, and they are also more likely to hire. This creates a selection problem in which raw adoption-employment correlations can understate displacement or even have the wrong sign. Our setting differs in a way that is useful for this concern. During much of our sample period, Argentine firms faced substantial frictions in importing capital goods and accessing credit, including foreign-exchange and import-authorization constraints. Because the automation capital we study is largely imported in Argentina, the exact timing of observed adoption is likely shaped not only by firms' desired investment date but also by when they can finance and clear the import transaction. This does not make adoption randomly assigned, but it weakens the simple story in which observed robot imports merely mark the peak of a firm-specific demand boom. Consistent with that interpretation, we observe no systematic employment growth at adopting firms before the automation event, and the post-event decline is not mechanically driven by mean reversion from an unusually high pre-period. Moreover, any remaining positive demand shock around adoption would work against finding an employment decline: it would raise employment at treated firms and partly offset the labor-saving effect of automation. The fact that we still estimate a large post-event contraction in registered labor demand is therefore hard to reconcile with a purely positive demand-driven adoption story.

[Bessen et al. \(2025\)](#) provide the closest worker-level comparison. They link automating firms to individual workers in the Netherlands and show that automation increases separations and lowers cumulative wage income, mainly through fewer days worked. We build on their firm-worker perspective, but the Argentine setting changes both the observed margins and the interpretation of worker losses. Like them, we find weaker attachment among incumbents, both to the original employer and to registered employment more broadly, though the margins differ in precision: our estimate of the separation margin is sizeable but imprecise, and the sharpest signal is the move into simplified self-employment. Unlike in their Dutch evidence, however, reported wage earnings among workers who remain observed in registered wage employment do not show a comparable negative path. The adjustment margin we observe is instead a change in employment status: workers spend

fewer months in registered wage employment, are more likely to exit the formal registry, are more likely to appear only in the simplified self-employment tax regime,¹ consistent with a shift toward weaker and potentially lower-quality formal attachment. This difference is informative. Argentina's registered wage-employment sector in our sample period was characterized by collective bargaining, union wage schedules, wage floors, payroll taxes, and regulated employment relationships. These institutions may make large worker-specific changes in reported wages less common or less visible in administrative annual earnings data. The point applies in both directions. For workers who remain at the automating firm, one might expect productivity gains, survivor selection, or retention motives to raise compensation. In a highly regulated payroll environment, however, such increases may not appear fully as registered wages if firms rely on nonreported bonuses or off-payroll payments. Evidence from Brazil is consistent with this margin: [Feinmann \(2024\)](#) survey formal employees and document widespread payments under the table, with 24 percent of formal employees reporting that they receive part of their pay off the books.² We therefore interpret the small reported-wage response cautiously. It does not imply that compensation is irrelevant; it means that the margins we observe most clearly are employment quantities and labor-market states: separations, movement across employers, exits from registered wage employment, and transitions into the simplified self-employment regime.

The paper also contributes to the literature on worker adjustment in developing economies and Latin America ([Brambilla et al., 2023](#); [Faber, 2020](#); [Carbonero et al., 2023](#); [Dix-Carneiro and Kovak, 2019](#); [Dix-Carneiro et al., 2021](#)). Existing evidence for the region has mostly relied on regional exposure measures and local labor-market outcomes. We instead observe firm-level adoption episodes directly and trace incumbent workers through the formal employer–employee registry and self-employment-tax records before and after those events. This provides a different lens on automation in Latin America: rather than inferring exposure from local labor-market composition, we can measure how workers attached to adopting firms fare inside and outside the standard registered wage-employment regime. It also links our results to evidence that informality and self-employment act as margins of adjustment in developing economies. In our setting, formal wage cuts are not the main observed response; displacement appears through transitions out of the original employer, out of the formal registry, into the simplified self-employment regime, and possibly into unregistered work or non-employment for workers who disappear from our administrative records. The combination of these registries therefore lets us distinguish continued attachment to registered wage employment from a move toward a more weakly attached, non-wage, or unobserved labor-market state.

Finally, the paper reconnects to the broader literature on technological change, robots, and automation. In the task-based tradition, [Autor et al. \(2003\)](#), [Acemoglu and Autor \(2011\)](#), [Autor \(2015\)](#), and [Acemoglu and Restrepo \(2018\)](#) emphasize how new technologies reshape tasks, labor demand, and wage inequality. A growing empirical literature studies these issues using industry, regional, plant, and firm-level data, including [Graetz and Michaels \(2018\)](#), [Acemoglu and Restrepo \(2020\)](#), [Dauth et al. \(2021\)](#), [Cheng et al. \(2019\)](#), [Carbonero et al. \(2023\)](#), [Koch et al. \(2021\)](#), [Humlum \(2022\)](#), [Aghion et al. \(2020\)](#), [Bonfiglioli et al. \(2024\)](#), and [Bessen et al. \(2025\)](#). A central lesson from the firm-level branch of this literature is

¹This regime is not informal employment in a legal sense. For our purposes, it is informative because it captures a move away from the standard registered wage-employment relationship, while workers who exit the formal registry and do not appear in monotributo records may be unemployed, out of the labor force, or working in unregistered employment.

²We do not observe off-payroll compensation in the formal social-security records, and we do not use the Brazilian evidence as a direct estimate for Argentina. The point is narrower: in highly regulated Latin American labor markets, reported payroll wages need not be the only margin of compensation adjustment, either upward or downward.

that firms that automate differ systematically from firms that do not: automating firms are often larger, more productive, and on different trajectories even before adoption (Koch et al., 2021; Humlum, 2022). Measurement also varies across studies, including import records, survey indicators, third-party automation expenditures, and broader proxies based on industrial-equipment investment (Humlum, 2022; Koch et al., 2021; Bonfiglioli et al., 2024; Bessen et al., 2025; Aghion et al., 2020). Our contribution is not the use of trade records per se, but the combination of import-based adoption events with linked employer–employee data in a setting with limited domestic production of automation capital and where we can follow incumbent workers around those events.

The remainder of the paper is organized as follows. Section 2 describes the data and the measurement of automation imports. Section 3 defines automation events in the data and presents descriptive evidence on automating firms and workers. Section 4 presents our empirical strategy. Section 5 reports the results, robustness checks, and heterogeneity analyses. Section 6 concludes.

2 | DATA AND MEASUREMENT

We use three linked administrative sources. Customs transaction records identify firm-level imports of automation-related capital goods and cover 2007–2022. The Sistema Integrado Previsional Argentino (SIPA), Argentina’s formal social-security registry, provides linked employer–employee records on registered employment relationships and wages, also for 2007–2022. Records from *monotributo*, the simplified tax regime for small-scale self-employment, identify workers who leave registered wage employment and register under this regime; these records are available for 2011–2018.

2.1 | Customs Records and Automation Imports

Our first source is Argentine customs data, which cover the universe of firm-level import transactions and report detailed product codes (HS 6–10 digits), declared values, quantities, and countries of origin. We match customs records to firms using the tax identifier (CUIT).

As in Acemoglu and Restrepo (2022), we identify automation from the product composition of firms’ imports, but we do not adopt their basket unchanged. Their classification was built to compare aggregate technology adoption across countries; ours has to capture the automation-related capital that Argentine firms actually buy abroad. Four of their automation categories carry over directly: industrial robots, automatic machine tools, numerically controlled metalworking machinery, and automatic regulating instruments. We leave aside automatic welding machines, conveyors, and weaving, knitting and other dedicated textile machinery. We add programmable and numerical controllers, and computing and data-processing equipment.

We order these goods from the narrowest to the broadest notion of automation, and carry that ordering through the paper. Definition 1 comprises industrial robots alone; Definition 2 adds machining centres; Definition 3 adds numerically controlled metalworking machinery; Definition 4 adds computing and data-processing equipment; and Definition 5, our baseline, adds automatic control instruments and programmable controllers. Table B.1 lists the codes in each group. Because the boundary of the basket is a judgement rather than a fact, Section 5.6.2 re-estimates every result under all five definitions.

Figure 1 shows, by year, the number of Argentine firms importing automation-related goods and the total value of those imports. These series help characterize the prevalence and aggregate relevance of automation-related capital inflows in Argentina over 2007–2022.

These automation importers are a small and distinctive slice of Argentine importing activity. Over 2007–2022, 16,016 firms imported at least one automation-related product, out of 65,244 firms that imported any good—roughly one in four importers—and the 8,503 firms we later classify as having an automation spike account for 13 percent of all importers. By value the share is much smaller: automation goods represent about 1.6 percent of total Argentine imports over the period. The margin we study is therefore not a large share of aggregate trade. It is a concentrated and lumpy investment made by a particular set of firms, and it is that firm-level lumpiness, rather than aggregate exposure, that our design exploits.




FIGURE 1 Automation-Related Imports in Argentina, 2007–2022

Notes: Panel (a) plots the number of formal-registry firms importing at least one automation-related product each year. Panel (b) plots the total real value of automation-related imports (in millions of real USD) by year. *Source:* Own computations based on Argentine customs records.

While our measure of the technology shock is constructed using imports, the real object of interest is adoption. Three gaps separate them. First, coverage: firms that automate through domestic distributors, integrators, or related-party procurement never appear as direct importers. Second, attribution: the goods a firm imports need not be deployed within that firm. Third, dating: the date of an import need not be the date at which equipment is installed and put to work. [Bessen et al. \(2025\)](#) bring evidence to bear on the first two by comparing an import-based measure with an expenditure-based measure of automation costs in Dutch data. They find that the two are correlated across firms but only weakly correlated within firms over time, and that direct imports are concentrated in a narrower set of firms and sectors. The within-firm result is the one that matters most for a design like ours, which reads adoption off the timing of a firm’s own imports.

The Argentine setting is well suited to an import-based measure. Argentina is a net importer of automation technology: domestic production of advanced automation capital is limited, and the country is not a significant exporter of these goods, so the re-export problem that complicates the Dutch mapping in [Bessen et al. \(2025\)](#) does not arise here. Each of the three gaps is therefore narrower than it would be elsewhere, and the design narrows them further. First, coverage binds least: firms that automate exclusively through domestic distributors or related-party procurement never enter the sample at all, so they do not contaminate the treated–control contrast. The concern is confined to firms that do both, and even for those, an import several times larger than the firm’s normal flow marks an economically meaningful moment in its investment path. The comparison also uses firms whose own first spike arrives later, rather than firms that never import automation goods, so that permanent differences between direct importers and non-importers do not drive the

estimates.³ Second, we address attribution directly, by excluding sectors whose firms import these goods for resale rather than for use in production, and by building the design around large, lumpy changes in a firm's own imports rather than small continuous purchases: it is routine purchasing, not discrete investment, that tracks true automation spending most weakly within a firm. Finally, dating works in our favor. Equipment clears customs before it is installed and put to work, so our event date runs ahead of actual deployment: five years after the import is less than five years of use. This understates the effects rather than inflating them, and it leaves the pre-event period safely before adoption. We therefore read our estimates as the effects of directly observed, lumpy import-based automation events.

This timing wedge is a feature of the Argentine import regime rather than an incidental one. Throughout most of the sample period, importing capital goods required non-automatic licensing—DJAI between 2012 and 2015, SIMI thereafter—and clearance was slow even by regional standards: World Bank Enterprise Survey data show that clearing imports through customs took 19.3 days in Argentina in 2017 against a Latin American average of 13.9, up from 7.5 days in 2006, while obtaining an import licence took 77 days in 2010 against a regional average of 27.2. These frictions do more than lengthen the interval between decision and import: they add noise to it. Our design does not require that automation itself be random—it compares adopters with firms that adopt later, and so already conditions on eventual adoption. What it requires is that the timing of adoption among adopters not be chosen in response to the firm's labor-market conditions. A firm that intended to import in 2013 may clear customs only in 2015, for reasons that have to do with licence processing rather than with its hiring or wage path. Displacement runs in both directions: a firm may equally bring a purchase forward, taking advantage of a period in which restrictions ease or of a licence that clears sooner than expected. Timing is thus uncertain on both sides of the intended date. Relative to a setting where capital goods enter automatically, the year in which an adopting firm's spike appears carries a larger administrative component—something that works to our advantage.⁴ We return to this in Section 4.

2.2 | Administrative Records

Our main labor-market source is the Sistema Integrado Previsional Argentino (SIPA), the country's social-security registry. The registry records formal employer–employee relationships and monthly information on each match.⁵ For each worker, we observe the employing firm, the months worked, labor earnings, and basic characteristics such as gender, age, region, and sector. The registry does not record hours worked, so wages throughout the paper refer to monthly earnings per registered job rather than to an hourly rate. Because workers can be matched to more than one employer within a year, the registry allows us to follow workers across firms and to construct outcomes based on all formal jobs, not only the job held at the automating firm.

From these records we construct worker-level outcomes that speak to different margins of adjustment: total real earnings at the automating firm, total real earnings across all

³Some future direct importers may have begun adopting indirectly through third parties before their observed spike. Under the labor-saving interpretation of automation, this would move control outcomes in the same direction as treated outcomes and bias the treated–control contrast toward zero.

⁴A firm awaiting delivery might begin adjusting before the equipment arrives. Any such anticipation would show up as pre-event movement in the outcomes, which we examine directly in Section 5.

⁵We exclude observations corresponding to workers hired through intermediary agencies, as well as worker-months whose reported wage falls below a minimum threshold of ARS 10 (in January 2007 prices), with the threshold updated over time using inflation. Wages below this threshold are effectively zero and likely reflect reporting errors or non-meaningful employment relationships rather than genuine labor earnings.

employers, average real wages across all jobs, the number of months observed in the formal registry, and the number of employers. The earnings measures combine pay and attachment: a worker who spends fewer months in registered employment earns less over the year even at an unchanged wage. At the firm level, the registry gives us total employment, the wage bill, and the wage distribution among registered employees.

The registry's main limitation in our setting is that it observes only the formal labor market. We measure reallocation within registered employment well, but when a worker disappears from the registry we do not directly observe whether they become unemployed, move into self-employment, or work informally. In Argentina, as in other Latin American economies, this matters more than it would in higher-income economies, given the size of the informal sector. Two features of our analysis follow from this limitation. First, months in the formal registry are themselves a central outcome. Second, we link the registry to *monotributo*, the country's simplified tax scheme for small-scale self-employment, administered by the federal tax authority (AFIP); we observe its taxpayer records for 2011–2018, a shorter window than the rest of our data. Linking them to the employer–employee panel lets us separate workers who leave registered wage employment and register as self-employed from those who disappear from our administrative data altogether.⁶

3 | AUTOMATION EVENTS AND SAMPLE DESCRIPTIVES

3.1 | Defining Automation Events in the Data

Building on the customs-based measure from Section 2, we define events as in Bessen et al. (2025), who in turn draw on the lumpiness of automation investment (Davis and Haltiwanger, 1999).⁷ Firms rarely import automation capital smoothly over time; adoption tends to occur through discrete purchase episodes. Following their formulation, we express each firm's real automation imports as a share of its average real wage bill, which states each purchase in units of the firm's own payroll. Let

$$A_{jt} = \frac{M_{jt}^A}{\overline{WB}_j},$$

where M_{jt}^A denotes real automation-related imports by firm j in year t , and $\overline{WB}_j$ is the firm's average real wage bill over the sample. We then define a firm as having an *automation spike* in year t if this normalized import measure exceeds γ times its own average value in all other years:

⁶This distinction should be interpreted with caution. A worker may leave the employer–employee registry without appearing in the simplified-regime records and still be working informally. The measure is therefore useful for detecting continued registered activity after formal-sector exit, but it is an imperfect proxy for total employment. Registration under this regime also carries thinner social insurance than dependent employment: contributions are lower and the associated pension and health coverage less generous. For some independent workers—particularly those with strong outside options—self-employment could be a choice rather than a fallback, but for many a move out of registered wage employment and into the regime represents a decline in job quality and in social protection.

⁷Our baseline automation-import measure covers twenty-nine tariff codes spanning industrial robots, machining centres, CNC machine tools, computing equipment, automatic control instruments, and programmable and numerical controllers. It is therefore considerably broader than industrial robots: broader automation goods account for 77.7 percent of events, whereas robot-dominated events account for 2.5 percent. Appendix C documents the product composition and economic scale of these events.

$$\text{spike}_{jt} = \mathbb{1} \{A_{jt} \geq \gamma \bar{A}_{j,-t}\}, \quad (1)$$

where $\bar{A}_{j,-t}$ is the average of firm j 's normalized automation imports across all other years in which the firm appears in the registry, including years in which it imports no automation goods. As in [Bessen et al. \(2025\)](#), the normalization is a matter of units rather than of classification: because $\bar{WB}_j$ does not vary within a firm, it enters both sides of the inequality, and the condition is equivalent to comparing a firm's automation imports against their own average in other years. Following them, we set $\gamma = 3$ in the baseline specification. Figure 2 reports the annual event profile under three alternative thresholds.




FIGURE 2 Number of Automation Events by Year and Threshold

Notes: The figure plots the number of firms experiencing their first automation spike in each year, under three alternative thresholds ($\gamma \in \{2, 3, 4\}$); the baseline uses $\gamma = 3$. The sample comprises firms observed in the formal registry that import automation-related goods at least once; 8,503 of them experience at least one spike over the full 2007–2022 panel. Event years 2010–2016 are those used as treated cohorts in the stacked design; 2017–2019 are shown for context and fall outside the estimation window. *Source:* Own computations based on Argentine customs and SIPA records.

When a firm satisfies equation 1 in more than one year, we define the *automation event* as the year of its first spike.⁸ Figure 3 shows a sharp rise in normalized automation imports exactly at the event year and low values elsewhere, which is consistent with a dynamic pattern in firms' automation investments: an initial sizable purchase followed by smaller automation-related imports in subsequent years.

⁸ Approximately one-third of treated firms experience more than one spike. Among these firms, the vast majority exhibit two (77%) or three (21%) spikes, with only a very small fraction experiencing four or more. The average number of spikes among firms with multiple spikes is 2.25.




FIGURE 3 Normalized Automation Imports around the Event Year

Notes: The figure plots mean real automation imports, normalized by the firm's average real wage bill, by year relative to the automation event ($t = 0$). The sample is restricted to firms present in the formal registry in every year of the panel. The pronounced spike at $t = 0$, and the fact that post-event imports settle at a level above the pre-event one, are consistent with automation investment arriving as a discrete episode that shifts the firm onto a higher technology path. *Source:* Own computations based on Argentine customs and SIPA records.

3.2 | Firms and Workers in the Analysis Sample

For the firm-level analysis, we construct an annual panel by linking the full population of registered employers in SIPA, described in Section 2, to firm-level customs records through tax identifiers. SIPA provides yearly measures of formal employment, wages, and the wage bill, while the customs records identify whether and when each firm imports automation-related capital goods. We classify a firm as an adopter when it first satisfies the automation-spike rule defined in Section 3.1, and take the year of its first spike as the treatment year. The linked panel allows us to characterize adopting firms and follow their formal labor outcomes around the automation event.

For the worker-level analysis, we focus on individuals who were employed by an automation-spike firm for at least 27 of the 36 months preceding the event—equivalently, at least 75% of the three pre-event years. We refer to these individuals as *incumbent workers*. This requirement selects workers with a stable pre-event attachment to the firm. We further restrict the sample to workers for whom the automating firm was the main source of labor earnings in that pre-event window, and we exclude the roughly 30,000 individuals who meet the incumbent definition at more than one spike firm, since their post-event trajectories cannot be cleanly assigned to a single event.

Once a worker enters the sample, we follow that individual regardless of subsequent attachment to the original employer. Three margins are visible in the data: continued registered wage employment (whether at the automating firm or elsewhere in SIPA); transition into *monotributo*, Argentina's simplified tax regime for self-employment and small-scale activity; and exit from the formal registry altogether. This lets us separate displacement from the original employer, reallocation within registered wage employment, and movements across broader formal-employment states.

The remainder of this subsection describes the firms and workers in the linked administrative data. We first characterize firms that import automation-related capital goods,

distinguish firms with discrete import spikes from other importers, and examine which importers are observed as formal employers in SIPA. We then describe workers at automation-importing firms and the incumbent workers attached to adopting firms before the event.

3.2.1 | Automation-Importing Firms in Customs and SIPA

We begin with the customs records, which allow us to characterize firms that import automation-related capital goods and the timing of those purchases. We divide importing firms into three mutually exclusive groups: automation importers that experience at least one spike, automation importers that never experience a spike, and firms that import other goods but no automation-related products. Table 1 compares these groups over 2007–2022. Of the 16,016 automation importers observed during this period, 8,503 experience at least one spike and 7,513 do not. Although the non-spike group is smaller, its average cumulative automation-import value over the panel is USD 1.5 million per firm, compared with USD 0.4 million among spike firms.

Spike and non-spike automation importer differ only modestly in the number of years in which they import automation goods (3.5 versus 3.9 years). The spike classification therefore does not primarily distinguish firms by the frequency of their automation imports, but by how those purchases are distributed over time. Conditional on positive automation imports, annual values are more variable within spike firms: the average within-firm coefficient of variation is 1.02, compared with 0.86 among non-spike firms. This pattern is consistent with spike firms concentrating their purchases in relatively large investment episodes, whereas non-spike importers distribute their automation investment more evenly across import years.

TABLE 1 Importer-Firm Profiles, 2007-2022

	Automation importers		Other importers
	Spike	Non-spike	
Number of firms	8,503	7,513	49,228
Share registered in SIPA (%)	100.0	85.0	81.1
Average years importing	9.3	8.5	4.3
Average import value (USD mill.)	43.0	58.6	2.2
Average years importing automation	3.5	3.9	-
Average automation-import value (USD mill.)	0.4	1.5	-
Average within-firm CV of automation imports	1.02	0.86	-
Share exporting (%)	61.9	44.9	28.6
Average years exporting	5.3	2.7	1.5
Average export value (USD mill.)	72.6	36.8	2.5

Notes: The table reports firm-level descriptive statistics using customs administrative data for 2007-2022. “Automation-spike importers” are firms experiencing discrete spikes in automation imports. “Automation non-spike importers” are firms that import automation-related goods in at least one year but do not experience a discrete automation-import spike. “Other importers” are firms that import some goods in at least one year but never import automation-related goods. The categories are exhaustive and mutually exclusive within the importer universe. The share registered in SIPA is the percentage of firms that appear in SIPA at least once during the panel, that is, firms that register at least one formal employee at some point in 2007-2022. Average years importing or exporting refer to the average number of years in which the firm imports or exports in at least one month. Import value, automation-import value, and export value are firm-level averages of total values over the panel, expressed in millions of U.S. dollars. The within-firm coefficient of variation is computed using annual automation import values over years with positive automation imports only; the table reports the average coefficient of variation across firms. *Source:* Own computations based on Argentine customs and SIPA records.

Appendix Figure A.1 provides a complementary view of non-spike importers by distinguishing large regular importers from small or occasional importers. Although large regular importers constitute a relatively small group, they account for a disproportionate share of aggregate automation-import value. The spike classification thus isolates lumpy purchase episodes rather than the firms with the largest cumulative automation expenditures.

Table 1 also connects the customs universe to the formal-employer records by reporting whether each importer ever appears in SIPA. By construction, all automation-spike firms are observed in SIPA, since the spike measure is defined among customs importers linked to the formal registry. Coverage is also high among the remaining groups: 85 percent of non-spike automation importers and 81.1 percent of other importers appear as formal employers at least once during 2007–2022.⁹ We next use the linked customs–SIPA panel to place these importing firms within the broader universe of formal employers. Table 2 restricts attention to this matched universe and compares automation-spike importers, non-spike automation importers, other importers, and firms that never import.

⁹ Among the 16,016 automation importers, 2,033 (12.7 percent) have at least one year in which they import automation goods without appearing as formal employers in SIPA. These unmatched importer-year observations account, on average, for less than 1 percent of annual automation-import value. More than half of these firms never appear in SIPA during the panel and import automation goods in fewer than two years on average.

TABLE 2 Formal-Employer Profiles by Importer Type, 2007-2022

	Non-importers	Automation importers		Other importers
		Spike	Non-spike	
Number of firms	1,456,726	8,503	6,384	39,936
Average months SIPA-registered	68.8	158.3	133.3	121.5
Average monthly employees	6.6	221.3	106.7	39.0
Average monthly real wages	8,020	17,815	17,065	12,857
Share in Buenos Aires/CABA (%)	47.5	35.1	50.2	50.1
Share in multiple regions (%)	6.6	45.8	36.6	27.4
Manufacturing (%)	7.1	64.3	17.0	32.7
Market services (%)	62.1	18.5	78.9	52.6
Primary sector (%)	9.3	6.2	1.4	6.7
Public/social services (%)	16.0	5.8	1.6	4.5
Construction (%)	5.4	5.2	1.1	3.6

Notes: The table reports firm-level descriptive statistics using SIPA administrative data for 2007-2022. The sample includes firms that appear in SIPA at least once, that is, firms that report at least one formal employee during the period. “Non-importers” are firms that do not import goods in any year during the sample period. “Automation-spike importers” are firms experiencing discrete spikes in automation imports. “Automation non-spike importers” are firms that import automation-related goods in at least one year but do not experience a discrete automation-import spike. “Other importers” are firms that import some goods in at least one year but never import automation-related goods. The categories are exhaustive and mutually exclusive. Average months SIPA-registered is the average number of months in which firms report formal employment. Average monthly employees is the average monthly number of formal employees. Average monthly real wages are expressed in December 2015 prices. Sector classifications follow the firm’s main activity. Firms are also classified according to their geographic presence. The first percentage reports the share of firms operating exclusively in CABA or Buenos Aires Province, while the second reports the share of firms operating in more than one administrative jurisdiction (with CABA treated as a separate jurisdiction for this purpose). *Source:* Own computations based on Argentine customs and SIPA records.

Automation-spike firms are substantially larger than other formal employers. They employ an average of 221 workers per month, compared with 6.6 at non-importing firms and 39.0 at importers of non-automation goods. Non-spike automation importers are also comparatively large, averaging 107 workers per month. Automation-related imports are therefore concentrated among relatively large formal employers, regardless of whether their purchases satisfy the spike criterion.

Wages display the same pattern. Average monthly real wages per worker are 17,815 pesos at spike firms and 17,065 pesos at non-spike automation importers, more than twice the 8,020-peso average among non-importing firms. Both groups of automation importers are also more persistently observed in SIPA: spike firms appear with registered employment for 158 months on average and non-spike firms for 133 months, compared with 69 months among non-importing firms. Direct automation importers are therefore large, high-wage, and persistently registered formal employers rather than marginal firms in the formal economy.

Automation-spike firms also have a distinctive sectoral and geographic profile. Manufacturing accounts for 64.3 percent of automation-spike firms, compared with 7.1 percent of non-importing firms, while market services account for 18.5 percent of spike firms, sub-

stantially below their share among non-importers.¹⁰ Appendix Figures A.2 and A.3 provide the detailed sectoral and geographic distributions. Within manufacturing, spikes are most common in fabricated metal products, machinery and equipment, rubber and plastic products, food products, and electrical equipment; within market services, they are concentrated in warehousing and transportation support, computer programming and consulting, architectural and engineering activities, and administrative activities. Among firms that can be assigned to a single province, the largest numbers are located in Buenos Aires Province and CABA, followed by Córdoba and Santa Fe, although spike events are observed across much of the country. Table 2 also shows that 46 percent of spike firms operate in multiple jurisdictions. These patterns show that automation spikes are concentrated in Argentina's main industrial regions but occur across much of the country.

3.2.2 | Workers at Automation-Importing Firms

Table 3 divides workers into four mutually exclusive groups according to their employers' import status and, for workers at spike firms, their pre-event attachment. These groups are: workers never employed by a firm while it imports any goods; incumbent workers at automation-spike firms who satisfy the 27-of-36-month attachment requirement and for whom the spike firm is the principal employer; workers employed by automation-importing firms who are not classified as exposed to a spike¹¹; and workers employed by other importing firms. Incumbent workers at automation-spike firms have the highest wages and strongest formal-sector attachment among the four groups. Their average monthly real wage is 24,219 pesos, compared with 18,866 pesos among workers in the non-spike automation group, 12,665 pesos among workers at other importers, and 10,135 pesos among non-importing workers. They are observed in SIPA for 153 months on average, compared with 102, 86, and 66 months, respectively. Because the incumbent definition selects workers with stable pre-event attachment, these differences should be interpreted as descriptive characteristics of the four groups rather than as differences attributable solely to employer import status.

¹⁰Note that, by construction, the share of service firms is relatively low for automation-spike firms. Automation-import spikes are not identified for wholesale and retail trade firms, as many firms in this sector may act as intermediaries rather than directly adopting the automation technologies themselves. Consequently, trade-sector firms with automation imports are classified simply as automation importers rather than automation-spike firms. For the econometric analysis, however, wholesale and retail trade firms are excluded altogether.

¹¹This group includes workers at non-spike automation importers and workers at spike firms who do not satisfy the incumbent definition.

TABLE 3 Worker Profiles by Employer Importer Type, 2007-2022

	Non- importing	Automation		Other
	workers	Spike	Non-spike	importers
Number of workers	12,906,812	1,117,211	3,441,712	3,626,738
Average months SIPA-registered	66.3	153.0	101.5	86.4
Average monthly real wages	10,135	24,219	18,866	12,665
Share women (%)	44.1	32.5	30.1	37.4
Share digital workers (%)	2.0	3.6	4.2	2.3

Notes: The table reports worker-level descriptive statistics using SIPA administrative data for 2007-2022. “Non-importing workers” are workers who are never employed by a firm importing any goods at the time of their employment. “Automation non-spike workers” are workers who, in at least one month, are employed by firms importing automation-related goods at that time but are not exposed to an automation spike. “Automation-spike workers” are workers employed for at least 27 of the 36 months before the event at firms experiencing automation spikes, for whom this employment also represents their primary source of formal labor income. “Other-importer workers” are workers who, in at least one month, are employed by firms importing non-automation goods. The categories are exhaustive and mutually exclusive. Workers employed by firms that experience an automation spike but who do not satisfy the 27-of-36-month pre-event attachment requirement are not classified as automation-spike workers; they are classified as automation workers because they are not considered exposed to the shock in the incumbent-worker design. Digital workers are identified from occupation titles associated with information technology, programming, electronics and telecommunications, data processing, robot operation, and electrical or electrotechnical work. The share of women and the share of digital workers are computed over the subset of workers for whom this information is available. *Source:* Own computations based on Argentine customs and SIPA records.

The gender and occupational composition of both automation-worker groups also differs from that of non-importing workers. Women account for 32.5 percent of incumbent workers at spike firms and 30.1 percent of workers in the non-spike automation group, compared with 44.1 percent among non-importing workers. Digital workers account for 3.6 and 4.2 percent of the two automation groups, respectively, compared with 2.0 percent among non-importing workers.¹²

3.2.3 | Pre-Event Characteristics of Adopting Firms and Incumbent Workers

Table 4 characterizes adopting firms and their incumbent workers before the automation event. Firm characteristics are measured over the three years preceding the first spike, while the worker column describes incumbents identified from their attachment to the adopting firm during the same period.

Before the event, adopting firms are observed in SIPA for an average of 29.8 of the 36 months and employ 222 workers per month, with average monthly real wages of 17,742 pesos. Averaging workforce shares across firms, women hold 18.3 percent of positions and

¹²Occupation is not recorded in SIPA itself. It comes from a separate administrative source covering 2021–2023, in which occupation is observed only when an employment relationship is created or modified. We classify a worker as digital if that individual ever appears in a digital or technology-related occupation within that window—information technology and programming, electronics and telecommunications, data processing, robot operation, and electrical or electrotechnical work—and apply the classification to all of that worker’s spells, including those that predate it. Occupation is available for roughly 60 percent of automation-spike workers, and the assumption becomes less credible the farther a spell is from 2021–2023, so digital shares should be read as an approximation rather than a contemporaneous measure of task content.

digital workers 4.8 percent. Pre-event automation imports are limited: 28.6 percent of firms import automation goods during this period, and the average import value is USD 0.04 million. At the same time, 45 percent of adopters export before the spike, with average exports of USD 11.4 million. Thus, although their direct automation imports are modest before the event, a sizeable share of adopting firms is already engaged in international markets.

TABLE 4 Pre-Event Characteristics of Adopting Firms and Incumbent Workers

	Automation-spike firms	Automation-spike workers
Average months SIPA-registered	29.8	-
Average monthly employees	221.9	-
Average monthly real wages	17,742	-
Average percentage of digital jobs	4.8	-
Average percentage of women	18.3	-
Share: Automation-importing (%)	28.6	-
Average automation-import value (USD mill.)	0.04	-
Share: Exporting (%)	45.0	-
Average export value (USD mill.)	11.4	-
Average months at spike firm	-	35.2
Average monthly real wages at spike firm	-	24,959
Average age	-	41.1
Share: More than 5 years at spike firm (%)	-	96.1
Share: More than 10 years at spike firm (%)	-	62.0
Share: More than 5 years in SIPA (%)	-	99.2
Share: More than 10 years in SIPA (%)	-	79.5

Notes: The table reports characteristics of automation-spike firms and incumbent workers in the three years before the automation spike. “Automation-spike workers” are workers employed for at least 27 of the 36 months before the event at firms experiencing automation spikes, for whom this employment also represents their primary source of formal labor income. Average monthly employees is the average, across firms, of monthly formal positions in the firm. Average monthly real wages are expressed in December 2015 prices. The average percentage of digital jobs is the average, across firms, of the share of positions held by digital workers. Digital workers are identified from occupation titles associated with information technology, programming, electronics and telecommunications, data processing, robot operation, and electrical or electrotechnical work. The average percentage of women is the average, across firms, of the share of positions held by women. For workers, average monthly real wages are the average of monthly registered wages, average age is measured before the event, and tenure at the spike firm is measured in the worker’s original employer. Shares with more than 5 or 10 years of tenure are computed over the set of workers for whom the sample window allows tenure to be observed at least 5 or 10 years before the spike. *Source:* Own computations based on Argentine customs and SIPA records.

By construction, the incumbent-worker sample is strongly attached to the adopting employer before the event. These workers are observed at the spike firm for an average of 35.2 of the 36 pre-event months, have an average age of 41.1, and earn average monthly real wages of 24,959 pesos at that firm. Among workers for whom five years of prior tenure can be observed, 96.1 percent have at least five years at the spike firm; among those for whom ten years can be observed, 62.0 percent have at least ten years. Looking across all

employers in SIPA, 99.2 percent have formal-sector histories extending beyond five years and 79.5 percent beyond ten. Although the incumbent definition mechanically ensures strong recent attachment, it does not impose these much longer employment histories. The sample therefore consists predominantly of established formal-sector workers rather than recent hires in short-lived employment spells.

This pre-event profile provides an important benchmark for interpreting the worker-level analysis. Any subsequent weakening of formal attachment would represent a departure from long-standing registered wage employment, rather than the continuation of a tenuous pre-existing connection to the formal sector. The implications of such a transition depend in part on the employment alternatives available outside registered wage jobs, which we describe next.

3.3 | Informality in the Argentine Labor Market

Argentina's labor market contains several quantitatively important states outside registered wage employment. Monotributo records allow us to observe one form of registered self-employment, but workers absent from both SIPA and monotributo may be unregistered wage employees, other self-employed workers, unemployed, or inactive. To establish the scale of these states, we use the Permanent Household Survey (EPH, INDEC), harmonized by SEDLAC (CEDLAS and the World Bank), for 2007–2022. The survey records registration status only for salaried workers: a wage job is registered if it carries pension rights. It provides no equivalent classification for own-account workers and small employers, so we retain self-employment as a single category rather than assigning it a registration status. Among individuals aged 15–64, we therefore distinguish registered and unregistered salaried workers, the self-employed, unpaid family workers, the unemployed, and the inactive (Table 5). We measure informality in two ways: as the unregistered share of salaried employment and as the share of total employment outside registered wage jobs.

Employment outside registered wage jobs is quantitatively as important as registered wage employment. Averaging over 2007–2022, registered salaried workers account for 50 percent of employment, unregistered salaried workers for 26 percent, and the self-employed for 23 percent, with the remainder consisting of unpaid family workers. Consequently, 34 percent of salaried workers are unregistered, while registered wage employment covers only 31 percent of the working-age population. The registered-wage share of employment increased through the mid-2010s but subsequently receded as salaried informality and self-employment expanded, leaving the share of employment outside registered wage jobs approximately unchanged between 2007 and 2022 (Appendix Figure A.4).

Informality varies sharply across sectors and worker characteristics, and the two measures generally identify the same high- and low-informality groups (Figure 4). Across sectors (panel (a)), the unregistered share of salaried employment is below 11 percent in mining, utilities, education, finance, and public administration, but reaches 59 percent in agriculture, 67 percent in construction, and 78 percent in domestic service. Manufacturing—the sector most exposed to automation imports in our data—lies between these extremes: 28 percent of its salaried workers are unregistered, while 45 percent of all manufacturing employment falls outside registered wage jobs. Similar gradients appear across worker characteristics (panel (b)). The unregistered share is 42 percent in the Northwest and 37 percent in the Northeast, compared with 19 percent in Patagonia; 53 percent among workers with no more than eight years of schooling, compared with 16 percent among those with at least fourteen; and 59 percent among workers aged 15–24, compared with 31 percent at ages 25–40 and 27 percent at ages 41–64.

These descriptive patterns establish the scale of employment outside registered wage


jobs, but they do not identify the destinations of workers who leave adopting firms. Unregistered wage employment and self-employment together account for nearly half of all jobs, so disappearing from SIPA need not imply unemployment or labor-force exit. Because the administrative data observe only registered wage employment and monotributo, we interpret exit from SIPA as weakened attachment to registered wage employment, not as direct evidence of a transition into informality.

TABLE 5 Labor-market states of the working-age population, 2007–2022

Labor-market state	% of employment	% of 15–64
Registered salaried	50.4	31.3
Unregistered salaried	26.1	16.2
Self-employed	23.0	14.3
Unpaid family worker	0.6	0.4
Unemployed	–	5.6
Inactive	–	32.3

Notes: Individuals aged 15–64, period average of semiannual estimates, 2007–2022 (2015H2–2016H1 not surveyed). A salaried job is registered if it carries pension rights. The self-employed—own-account workers and small employers—are reported as one category: the survey does not record a registration status for them.

Source: EPH (INDEC), harmonized by SEDLAC (CEDLAS–World Bank).



(a) By sector

(b) By worker characteristics

FIGURE 4 Informality

Notes: The figure reports two measures of informality: the unregistered share of salaried employment and the share of total employment outside registered wage jobs. Panel (a) reports both by sector; panel (b) reports them by region, education, and age. *Source:* EPH (INDEC), harmonized by SEDLAC (CEDLAS–World Bank), 2011–2019.

Overall, the descriptive evidence shows that the empirical analysis studies a selected but economically relevant margin of automation adoption. Direct automation importers are larger, more formal, more export-oriented, and more manufacturing-intensive than the average firm. Within that group, spike firms are not simply the firms that import the most automation goods every year; they are firms with discrete import episodes, which are better suited to an event-study design. Their workers are high-wage, experienced, and strongly attached to SIPA before the event. These facts discipline the interpretation of the results: the

estimates below speak to directly observed, import-based automation events among formal employer firms, and to the trajectories of incumbent workers who held stable registered jobs before adoption.

4 | EMPIRICAL STRATEGY

4.1 | Stacked Difference-in-Differences

Our empirical design exploits variation in the timing of the first directly observed automation event across firms, estimated within a stacked difference-in-differences framework (Cengiz et al., 2019; Wing et al., 2024). This estimator is preferable to a conventional two-way fixed-effects event study under staggered treatment timing because it avoids the contamination that arises when already-treated units serve as controls for later-treated units (Goodman-Bacon, 2021; de Chaisemartin and d'Haultfœuille, 2020; Sun and Abraham, 2021). Other estimators address this problem as well (Callaway and Sant'Anna, 2021; de Chaisemartin and d'Haultfœuille, 2020), but the stacked approach has the practical advantage of making the implicit comparisons explicit: each cohort of treated units is matched against a clearly defined set of not-yet-treated controls within a common event-time window.

We construct the stacked sample as follows. For each event year s with at least one first-time automation spike, we build a separate *sub-experiment*. Treated units are firms whose first event occurs in year s ; controls are firms whose first event occurs strictly after the post-treatment window, so that they remain untreated throughout the full $[s - 3, s + 5]$ event window. Within each sub-experiment, we require units to be observed in every event-time period from $e = -3$ to $e = 5$ and drop units that do not meet this balance requirement. We then vertically concatenate all admissible sub-experiments. Because our data cover 2007–2022 and each sub-experiment must contain both treated units and clean controls, the baseline stack includes event years from 2010 through 2016.

To illustrate the construction, consider the sub-experiment with event year $s = 2014$. A treated firm is one whose first observed automation spike occurs in 2014, and we use its outcomes from 2011 ($e = -3$) through 2019 ($e = 5$). A valid control is a firm whose first spike occurs in 2020 or later, leaving it untreated through 2019. A firm whose first spike occurs in 2017 is not a valid control for this sub-experiment because it becomes treated within the event window. The same later-adopting firm may serve as a clean control in more than one earlier sub-experiment and subsequently appear as treated in its own sub-experiment.

The worker-level sample inherits the same cohort structure. For each event year s , treated workers are incumbents of treated firms at $e = -1$, while control workers are incumbents of clean-control firms at $e = -1$. In both groups, incumbents are workers who were attached to the relevant firm for at least 27 of the 36 months preceding the event—equivalently, at least 75 percent of the three pre-event years—and for whom that firm was the principal employer. Once a worker enters the sample, we follow that individual across all subsequent formal-registry states regardless of continued attachment to the original employer. Post-event outcomes therefore capture realized trajectories across the formal sector, including transitions to monotributo or out of the formal registry altogether.

4.2 | Estimation

We estimate the same saturated event-study specification in the firm and worker stacks. Let u index a unit—a firm or an incumbent worker—within sub-experiment s , and let $e = t - s$ denote event time. Define $D_{us} = 1$ when unit u belongs to the treated group in sub-experiment s , and let $I_h = \mathbb{1}\{e = h\}$ denote an event-time indicator. With $e = -1$ as the

omitted reference period, the estimating equation is

$$Y_{use} = \alpha + \rho D_{us} + \sum_{\substack{h=-3 \\ h \neq -1}}^5 \gamma_h I_h + \sum_{\substack{h=-3 \\ h \neq -1}}^5 \delta_h (D_{us} \times I_h) + \varepsilon_{use}. \quad (2)$$

The treatment-group main effect ρ captures average level differences between treated and control units, and the γ_h coefficients trace the event-time path common to both groups. The parameters of interest are the interaction coefficients δ_h . For $h < -1$, they are pre-treatment placebo contrasts used to assess differential trends and anticipation. For $h \geq 0$, they measure the treated–control change at event time h relative to the year immediately preceding the automation event.

We estimate equation (2) by weighted least squares using corrective sample weights, or Q-weights. Without this correction, sub-experiments with disproportionately large control groups receive too much implicit weight, and treated and control trends are aggregated using different cohort shares. Treated units receive weight one. A control unit in sub-experiment s receives

$$Q_{us} = \begin{cases} 1, & \text{if } D_{us} = 1, \\ \frac{N_s^D / N^D}{N_s^C / N^C}, & \text{if } D_{us} = 0, \end{cases} \quad (3)$$

where N_s^D and N_s^C are the numbers of treated and control units in sub-experiment s , and $N^D = \sum_s N_s^D$ and $N^C = \sum_s N_s^C$ are the corresponding totals across the trimmed stack. We compute these weights separately for the firm and worker samples because their treated and control cohort shares differ.

Let Ω denote the set of sub-experiments satisfying the event-window and clean-control requirements. Under the identifying assumptions stated below, weighted least squares makes each post-event interaction coefficient equal to the trimmed aggregate ATT of [Wing et al. \(2024\)](#):

$$\delta_h = \sum_{s \in \Omega} \frac{N_s^D}{N^D} \text{ATT}(s, s + h). \quad (4)$$

Each cohort-specific effect is therefore weighted by that cohort's share of all treated units in the trimmed stack. Because the same set of cohorts and treated-cohort weights is used at every event time, changes in δ_h across the post-treatment period reflect treatment dynamics rather than changes in the composition of the cohorts being averaged.

At the firm level, u indexes a firm-by-sub-experiment observation and treatment is assigned by the firm's own first automation spike. At the worker level, u indexes a worker-by-sub-experiment observation and $D_{us} = 1$ when the worker's original employer belongs to the treated cohort. Standard errors are clustered at the firm level in both analyses, since treatment is assigned at the firm level and workers attached to the same original employer share common shocks. Clustering by firm also allows observations from the same firm to be dependent across the multiple sub-experiments in which that firm may appear.

Worker outcomes split into two groups based on how years without registered wage employment are treated. Log earnings and log monthly-wage outcomes are defined only for worker-years with positive registered-sector earnings. Years with zero earnings enter through the complementary attachment and labor-market-state outcomes—months in the

formal registry, exit indicators, and monotributo status—which are defined over the full incumbent-worker panel.

4.3 | Identifying Assumptions

Identification rests first on parallel untreated trends within each sub-experiment. In the absence of the automation event, firms treated in year s and clean-control firms treated only after $s + 5$ must have followed the same average outcome trajectory; the analogous condition must hold for their incumbent workers. We assess parallel trends graphically through the pre-event coefficients δ_h for $h < -1$. Coefficients centered near zero support the design, although they cannot rule out differential shocks beginning at or immediately before adoption.

The leading threat to parallel trends is endogenous automation timing, emphasized by Bonfiglioli et al. (2024). If firms automate in response to positive demand shocks that also drive hiring, raw adoption–employment correlations will understate, or may even reverse the sign of, the labor-saving effect. Our design addresses permanent differences between automating firms and arbitrary non-adopters by comparing current adopters only with firms that adopt later. It does not, however, eliminate time-varying shocks that determine the exact year of adoption. In Argentina, non-automatic import licensing, foreign-exchange restrictions, real-exchange-rate volatility, and credit constraints create a wedge between a firm’s desired investment date and the date at which imported automation capital clears customs. These frictions weaken a direct mapping between a contemporaneous firm-specific demand boom and the observed import spike. For this reason, the absence of systematic pre-event employment growth is a key credibility check for the design: it provides evidence against the concern that adoption timing is closely aligned with firm-specific demand shocks that also affect employment.

Parallel trends and no anticipation are arguably more plausible at the worker level than at the firm level. Firms choose when to automate and may adjust employment before the observed import date. Incumbent workers, by contrast, were attached to the firm for most of the three years preceding the spike and are unlikely to have selected their employer based on precise knowledge of its future automation timing. For this reason, we view the worker-level estimates as the most compelling causal evidence in the paper; the firm-level estimates remain informative about adjustment inside adopting firms but are more exposed to endogenous-timing concerns.

Two additional concerns remain. First, if firms reorganize staffing before the automation import clears customs, those changes should appear in the pre-event coefficients; the three-year pre-period is intended to reveal economically meaningful anticipation of this kind. Second, control firms in the same industry or local labor market may be indirectly affected by treated firms through product-market competition, input linkages, or worker reallocation. We cannot rule out these spillovers. If automation at treated firms also reduces outcomes at clean-control firms, the treated–control contrast will generally be attenuated; spillovers operating in the opposite direction could instead amplify it.

5 | RESULTS

We present the results in five steps. First, we ask whether the exact timing of a firm’s first automation spike is predictable from the firm’s own recent trajectory—a direct test of the identifying variation. Then, we document how adopting firms adjust registered labor demand and wages, overall and by initial firm size, and open the firm-level adjustment into

its gross-flow components: hiring and separations. Next we turn to incumbent workers and the margins along which they absorb the shock, and then to how those margins vary across worker groups. We close with robustness to the treatment definition—the spike threshold γ and the HS basket of automation goods.

5.1 | Validating the Design: Is Spike Timing Predictable?

The stacked design does not require automation adoption to be random. It requires something weaker: that, among firms that eventually adopt automation technology, the *year* in which the first import spike arrives is not systematically preceded by differential changes in firm outcomes. The descriptive statistics in Section 3 showed that adopters are positively selected relative to non-adopters—they are larger, pay higher wages, and survive longer in the registry—but that selection margin is handled by the design itself, which, as in [Bessen et al. \(2025\)](#), compares current adopters only with firms that adopt later.

The residual threat is timing: if firms spike precisely when they are expanding (or shrinking), the event-study coefficients would partly reflect that trajectory rather than the effect of automation. Table 6 takes this concern to the data. It reports the results of regressing an indicator equal to one in the year of the firm's first automation spike on the firm's own recent changes: the growth of log employment and of the log average wage, the growth of $\log(1 + \text{exports})$, and changes in gross hiring and separation rates. The sample contains only eventual adopters: each firm contributes the years in which it appears in the registry, up to and including its first spike, and post-spike years are dropped. Candidate years begin in 2009—the first year for which changes over the two preceding years can be computed in our 2007–2022 panel—so firms whose first spike falls earlier are excluded. All regressions include firm and calendar-year fixed effects, so the coefficients compare each firm's spike year with the same firm's other pre-spike years. Columns (1)–(5) enter each predictor separately; column (6) enters them jointly.

The table reports the test under four specifications: Panel A uses changes from $t - 2$ to $t - 1$ on all observations; Panel B repeats it excluding the bottom and top percentile of each predictor; Panels C and D measure changes over the full $t - 3$ to $t - 1$ pre-window, with and without outliers, respectively. The answer is that spike timing is essentially unpredictable from the firm's own recent trajectory. Point estimates are economically negligible throughout, and the within-firm R^2 never exceeds 0.0003. Across the four panels, no predictor is robustly significant. The one coefficient in Panel A that attains significance at the 5 percent level—the univariate change in log employment, -0.0085 (s.e. 0.0043)—is a useful illustration of how little predictive content these regressions carry. Taken at face value, it implies that a firm whose employment fell by 10 log points between $t - 2$ and $t - 1$ becomes only 0.09 percentage points more likely to spike in year t , against a baseline annual hazard of roughly 14–15 percent—a proportional shift of well under one percent. And even more, it does not survive within the table: the coefficient shrinks to -0.0007 (s.e. 0.0065) once the predictors enter jointly (Panel A, column 6), becomes statistically indistinguishable from zero once the tails of the predictor distribution are trimmed (Panel B: -0.0029 , s.e. 0.0068), and is also not significant when the change is measured over the full pre-window $t - 3$ to $t - 1$ (Panel C: -0.0034 , s.e. 0.0040).

The only other coefficient that reaches 5 percent significance anywhere in the table is the change in the *hiring* rate in Panels B and D, but its sign is positive: firms are slightly more likely to spike after a year of rising hiring. If anything, then, adoption follows expansions rather than contractions. Any such selection would bias the post-event estimates *toward zero*: firms adopting during an expansion would, absent adoption, have grown faster than their controls. The employment declines we document below are therefore, if

anything, understated. Finally, when the predictors enter the regression jointly, they are never significant: the F-statistic on the five predictors is 1.17 ($p = 0.32$) in Panel A and 1.22 ($p = 0.30$) in Panel C. In short, the margin on which the design rests—*when* a firm's first spike arrives—is not predicted by anything in the firm's recent trajectory.

The statistical unpredictability of timing is exactly what the institutional environment described in Section 4.3 would produce. During most of the sample period, importing capital goods into Argentina was governed by non-automatic licensing regimes (DJA between 2012 and 2015, SIMI thereafter), with approval lags that were long and variable, foreign-exchange access was rationed, and imported capital was hard to finance. The exact year in which a large automation purchase finally cleared customs was therefore only partly under the firm's control, placing a wedge between the decision to automate—which selection into the eventual-adopter sample already conditions on—and the observed timing of the spike that identifies the event study.

5.2 | Firm-Level Adjustment

Figure 5 shows the firm-level response to automation of four labor-demand measures. *Employment* counts the number of distinct individuals employed by the firm at any point during the year. A worker employed for one month and a worker employed for all twelve months each count as one employee. *Total positions* counts the cumulative number of monthly job positions during the year. Thus, a worker employed throughout all twelve months contributes twelve positions, while a worker employed for a single month contributes one position. *Average monthly positions* is obtained by dividing total positions by the number of months in which the firm is active in the SIPA records, capturing both the level and the duration of employment spells throughout the year. Lastly, *wage bill* is the firm's total real payroll across all registered positions in the year. Across all four measures, automation is followed by a clear contraction in registered labor demand. The decline is gradual rather than instantaneous: the estimates are close to zero in the event year and in event year +1, employment is 4 percent lower by event year +2, and the gap then widens steadily to 19 percent for employment and 20 percent for the wage bill by event year +5 (log-point estimates of -0.207 and -0.221). The result is therefore not an artifact of a particular employment measure. It is a broad reduction in the scale of registered wage employment at the automating firm.

Two features of these paths deserve emphasis before we turn to wages. First, the pre-event window shows what the design requires: adopters and not-yet-adopters on parallel paths before the event. Six of the eight pre-event coefficients in Figure 5 are statistically indistinguishable from zero, and all are quantitatively small. The two exceptions—total positions and the wage bill at $t = -3$ —are small (about 3 log points), vanish by $t = -2$, and do not appear in employment itself. These paths rule out the two patterns that would threaten the design: adopters were not expanding into the event, as they would be if spikes were timed to demand booms (Bonfiglioli et al., 2024), and they were not already contracting, which would make the post-event decline the continuation of a prior trend. Both patterns are also absent from the timing regressions of Table 6. Second, the post-event decline unfolds gradually over five years rather than as a discrete step, consistent with reorganization of production rather than a one-time accounting event—and hard to reconcile with mean reversion from an unusually good pre-event year, which would produce a sharp drop and recovery rather than a widening gap.

These estimates are gaps between adopters and comparable later adopters: a coefficient of -0.21 at $t = 5$ says that, five years after the event, adopters employ 19 percent fewer registered workers than the firms they would otherwise resemble—not necessarily that

TABLE 6 Do pre-spike changes in firm characteristics predict automation timing?

Panel A. Changes from $t - 2$ to $t - 1$: all observations						
	(1)	(2)	(3)	(4)	(5)	(6)
Change in log employment, $t-2$ to $t-1$	-0.0085** (0.0043)					-0.0007 (0.0065)
Change in log mean wage, $t-2$ to $t-1$		-0.0032 (0.0085)				-0.0150 (0.0106)
Change in $\log(1 + \text{exports})$, $t-2$ to $t-1$			0.0001 (0.0005)			-0.0004 (0.0006)
Change in hiring rate, $t-2$ to $t-1$				0.0003 (0.0002)		0.0002 (0.0002)
Change in separation rate, $t-3$ to $t-2$					-0.0120 (0.0100)	-0.0112 (0.0110)
Observations	43868	43868	48284	36753	36789	36650
Within R-squared	0.0001	0.0000	0.0000	0.0000	0.0000	0.0001
Panel B. Changes from $t - 2$ to $t - 1$: dropping bottom and top 1%						
	(1)	(2)	(3)	(4)	(5)	(6)
Change in log employment, $t-2$ to $t-1$	-0.0029 (0.0068)					-0.0022 (0.0126)
Change in log mean wage, $t-2$ to $t-1$		-0.0108 (0.0118)				-0.0186 (0.0156)
Change in $\log(1 + \text{exports})$, $t-2$ to $t-1$			-0.0003 (0.0006)			-0.0010 (0.0007)
Change in hiring rate, $t-2$ to $t-1$				0.0117** (0.0048)		0.0109 (0.0068)
Change in separation rate, $t-3$ to $t-2$					-0.0159 (0.0131)	-0.0144 (0.0154)
Observations	42476	42451	46537	35722	35409	32416
Within R-squared	0.0000	0.0000	0.0000	0.0002	0.0000	0.0003
Panel C. Changes from $t - 3$ to $t - 1$: all observations						
	(1)	(2)	(3)	(4)	(5)	(6)
Change in log employment, $t-3$ to $t-1$	-0.0034 (0.0040)					0.0047 (0.0055)
Change in log mean wage, $t-3$ to $t-1$		-0.0022 (0.0080)				-0.0049 (0.0101)
Change in $\log(1 + \text{exports})$, $t-3$ to $t-1$			0.0004 (0.0005)			-0.0002 (0.0006)
Change in hiring rate, $t-3$ to $t-1$				0.0005* (0.0003)		0.0004 (0.0003)
Change in separation rate, $t-4$ to $t-2$					-0.0149 (0.0115)	-0.0138 (0.0123)
Observations	36801	36801	42492	30423	30468	30352
Within R-squared	0.0000	0.0000	0.0000	0.0001	0.0001	0.0002
Panel D. Changes from $t - 3$ to $t - 1$: dropping bottom and top 1%						
	(1)	(2)	(3)	(4)	(5)	(6)
Change in log employment, $t-3$ to $t-1$	-0.0031 (0.0058)					0.0012 (0.0092)
Change in log mean wage, $t-3$ to $t-1$		0.0073 (0.0106)				0.0134 (0.0143)
Change in $\log(1 + \text{exports})$, $t-3$ to $t-1$			0.0002 (0.0006)			-0.0003 (0.0007)
Change in hiring rate, $t-3$ to $t-1$				0.0146*** (0.0054)		0.0130* (0.0071)
Change in separation rate, $t-4$ to $t-2$					-0.0157 (0.0150)	-0.0092 (0.0165)
Observations	35839	35705	41048	29561	29293	26908
Within R-squared	0.0000	0.0000	0.0000	0.0003	0.0000	0.0002

Notes: Each column reports a linear probability model regressing an indicator equal to one in the year of the firm's first automation spike on recent changes in the firm's own outcomes. The sample contains only firms that eventually adopt: each firm contributes the years in which it is observed in the registry, up to and including its first spike; post-spike years are dropped, candidate years begin in 2009 (the first year with two prior years of data), and firms whose first spike falls before 2009 are excluded. Changes are measured from $t - 2$ to $t - 1$ in Panels A and B, and from $t - 3$ to $t - 1$ in Panels C and D; separation windows are shifted back one year because a separation dated s is an outflow after year s . Panels B and D drop the bottom and top 1 percent of each predictor. Columns 1–5 enter each predictor separately; column 6 enters all five jointly (joint F-test: 1.17, $p = 0.32$ in Panel A; 1.22, $p = 0.30$ in Panel C). All regressions include firm and calendar-year fixed effects, and standard errors are clustered by firm. The sample covers the same sectors as the event studies. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Source: Own computations based on Argentine customs and SIPA records.

they shrank in absolute terms. And a gap like this can open in two ways: because adopters dismiss relatively more workers, or because they hire relatively fewer of them. The event studies alone cannot tell these apart. Section 5.3 separates the two margins directly, and Section 5.4 follows the incumbent workers themselves. The distinction matters because each margin implies a different incidence: if adopters dismiss, the cost of automation falls on the workers already at the firm; if they stop hiring, it falls on those the firm never takes on—workers who are invisible in any incumbent-based analysis.




FIGURE 5 Firm-level labor-demand responses to automation

Notes: Each panel plots the treated $\times$ event-time interaction coefficients from the stacked event study, with 95 percent confidence intervals; outcomes are in logs. Event time $t = -1$ is the omitted period. Estimates use corrective weights, and standard errors are clustered by firm. *Source:* Own computations based on Argentine customs and SIPA records.

Figure 6 turns to firm-level wage measures: average and median wages, and wages at the 25th and 75th percentiles of the within-firm distribution, all estimated on firms with at least ten employees at $t = -1$.¹³ The wage estimates are mostly statistically indistinguishable from zero—the exception is the 75th percentile, which ends the window at -1.7 percent (s.e. 0.8)—and, more importantly, they are small. At event year $+5$ the average wage is -0.8 percent and the median -0.7 percent (s.e. 0.8 each), much smaller than the 19 percent decline in employment and the 20 percent decline in the wage bill.

The wage-bill contraction is therefore a quantity adjustment, not a price adjustment. One caveat applies: firm-level wages describe the workers the firm employs in each year, not the same workers over time. If the workers who leave or are never replaced earn more or less than those who stay, average and percentile wages will move even if no individual

¹³ Within-firm wage summaries are poorly defined in very small firms: with a handful of employees, percentiles and averages track the pay of one or two workers, and a single hire or departure can move them mechanically. The threshold is measured at $t = -1$, before the event, so the restriction cannot select firms on their post-event employment path; firms that later shrink below ten employees remain in the sample.

wage changes. The small drifts in Figure 6 can thus reflect modest cuts for stayers, a changing workforce, or both. The worker-level analysis in Section 5.4 speaks directly to that distinction; for now, the point is that the firm-side response is concentrated on employment rather than on salaries.




FIGURE 6 Firm-level wage responses to automation

Notes: Each panel plots the treated $\times$ event-time interaction coefficients from the stacked event study, with 95 percent confidence intervals; outcomes are in logs. The panels report average and median wages (top) and the 25th and 75th percentiles of the within-firm wage distribution (bottom), for firms with at least ten employees at $t = -1$. Event time $t = -1$ is the omitted period. Estimates use corrective weights, and standard errors are clustered by firm. *Source:* Own computations based on Argentine customs and SIPA records.

In Figure 7 we test whether the firm-level pattern hides heterogeneity by initial size, splitting adopters at the median of employment at $t = -1$.¹⁴ The employment contraction is broadly shared: smaller and larger adopters reduce registered employment by similar proportions, -0.199 and -0.174 log points at event year $+5$, respectively. The median-wage panels reinforce the message of Figure 6: wage movements are small and imprecisely estimated in both groups, with, if anything, a slight downward drift among smaller adopters (-1.8 percent at $e + 5$, not statistically distinguishable from zero) and a flat profile among larger ones. We conclude that the automation event reduces registered employment across the adopter size distribution, without a robust wage-rate response in either group.

¹⁴Splitting the sample requires recomputing corrective weights: the weights depend on each sub-experiment's share of treated units, and those shares can change within a subgroup. Here, and in every heterogeneous regression in the paper, the weights are therefore recalculated on the subgroup's own estimation sample.




FIGURE 7 Firm-level heterogeneity by initial firm size

Notes: Each panel plots the treated $\times$ event-time interaction coefficients from the stacked event study, with 95 percent confidence intervals; outcomes are in logs. The panels report log employment and log median wages by initial employment group (small = bottom half, large = top half of employment at $t = -1$), with corrective weights recomputed within each group's effective sample. Event time $t = -1$ is the omitted period. Standard errors are clustered by firm. *Source:* Own computations based on Argentine customs and SIPA records.

5.3 | The Mechanism: Hiring versus Separations

Two distinct margins can produce the contraction observed in Figure 5: firms can dismiss workers, or they can simply stop replacing the ones who leave. Figure 8 separates these gross margins, estimating the baseline specification on annual hires and separations, each divided by the firm's employment at $t = -1$, so that post-event changes in firm size do not mechanically alter the ratio. Hires and separations are adjacent-year worker flows: a hire is a worker present in year t but not in $t - 1$, a separation a worker present in t but not in $t + 1$, so hiring is undefined in a firm's first observed year and separations in its last. Panels (c) and (d) repeat the exercise by initial firm size, recomputing the weights within each group.

The answer is unambiguous: the contraction operates through hiring. Relative to not-yet-adopting firms, hires fall from event year +2 onward and the shortfall grows to 21 points of pre-event employment by year +5 (-0.215 , s.e. 0.066)—the same order of magnitude as the employment decline itself. Separations, by contrast, never rise significantly: the point estimates are small and positive in the first post-event years ($+1.7$ and $+1.6$ points, neither significant) and decline below zero afterwards, reaching -7.3 points (s.e. 2.2) by year +5—mechanically consistent with a shrinking workforce: with the denominator fixed at $t = -1$, a smaller pool of workers present each year can generate fewer separations.¹⁵ The only pre-event feature is a small, parallel dip in both flows two to three years out: treated firms hire slightly less and separate slightly less than controls, so it is worker turnover, not net employment, that is marginally lower before adoption.

Splitting adopters by size shows the hiring margin at work in both groups, at very different magnitudes. Among larger adopters the hiring shortfall is precisely estimated and moderate (-6.9 points of pre-event employment at +5, s.e. 1.4), while among smaller adopters it is much larger but noisy (-33.5 points, s.e. 14.0). Separations again fail to rise for either group; for smaller adopters they end the window significantly below the control path.

The event study design measures by construction adopters relative to not-yet-adopters. It is also worth looking, more simply, at what hiring and separations do around the spike in absolute terms among adopting firms themselves. Appendix Figures D.2 report the raw

¹⁵Dividing instead by previous-year employment approximately nets out this pool effect, and separations are then flat—Appendix Figure D.1.

levels among treated firms only, with no comparison group and no causal interpretation: the average adopter hires 47 workers in the year before the event, 51 in the adoption year, and 37 by event year +5. Separations rise from 38 per year before the event to 42 around it and remain at 41 at the end of the window. The same panels split by initial size show that small adopters *increase* both hiring and separations in the first post-event years—turnover rises while net size stagnates—whereas large adopters reduce hiring steadily from 81 to 62 per year.

What comes next is, by construction, a different analysis—and it is where the paper's central evidence lies. Rather than aggregating flows across firms, the next section fixes a specific population—incumbents with stable, multi-year attachment before the event—and compares them with otherwise-similar incumbents at firms that adopt only later, following each worker across every subsequent employer, into self-employment, or out of the formal registry.




FIGURE 8 Gross worker flows at automating firms: hiring and separations

Notes: Panels (a)–(b) report stacked event-study estimates for total hires and separations divided by employment at $t = -1$. Panels (c)–(d) split by initial size (bottom vs. top half of employment at $t = -1$), with the weights recomputed within each group. Hires and separations are adjacent-year worker flows (a hire is present in t but not $t - 1$; a separation present in t but not $t + 1$), undefined at the panel borders, so flows in a firm's first and last observed year are dropped. Event time $t = -1$ is the omitted period. Estimates use corrective weights, and standard errors are clustered by firm. *Source:* Own computations based on Argentine customs and SIPA records.

5.4 | Worker-Level Adjustment

We now turn the focus to incumbent workers: individuals attached to the firm at $t = -1$ and who had the same firm as their principal employer for at least 27 of the 36 months preceding the event. We follow these workers throughout the entire formal registry regardless of whether they remain at the original employer or not. Figure 9 reports the earnings and

wage margins, conditional on a worker-year with positive registered wage earnings. In the worker-weighted baseline, total annual earnings across all registered wage jobs remain essentially flat after the original employer automates: the coefficients fluctuate within a band of roughly ± 2 percent and end the window at +1.4 percent (s.e. 2.3), always not statistically significant. Average monthly wages across all jobs deliver the same message (+1.8 percent at $t + 5$, s.e. 2.1). The original-employer panels reinforce this conclusion: total earnings and monthly wages at the original employer dip by 1.5–2 percent in mid-window (event years +1 through +3), with none of the estimates statistically distinguishable from zero, and return to zero by the end of the window. Among incumbent workers who remain observed in registered wage employment, reported earnings display no sizeable or robust movement in either direction relative to comparable workers at not-yet-adopting firms. This sharpens the firm-level wage-rate-versus-quantity distinction: the small negative drift in firm-level wages does not translate into a reported-wage decline for the specific incumbents we can follow. This aggregate null, however, masks offsetting responses across subgroups: earnings rise for low-wage incumbents and fall for those at small adopters, as Section 5.5 shows.



FIGURE 9 Worker-level earnings and wage outcomes

Notes: Each panel plots the treated $\times$ event-time interaction coefficients from the stacked event study, with 95 percent confidence intervals; outcomes are in logs, so zero-earnings worker-years enter the analysis through the attachment and state outcomes of Figure 10 rather than through these panels. Event time $t = -1$ is the omitted period. The specification uses worker weights, all sub-experiments, base fixed effects, and the full incumbent-worker sample; standard errors are clustered by anchor firm. *Source:* Own computations based on Argentine customs and SIPA records.

Figure 10 shows that the margin that does move is attachment to registered employment. The movements are not large, and some are less precisely estimated, but they are consistent in sign across measures. Months in the formal registry decline gradually to -0.14 months

per year by event year +5 (s.e. 0.08), and the probability of exiting the registry altogether rises to 1.0 percentage point (s.e. 0.6); both margins are statistically distinguishable from zero only at the 10 percent level at the end of the window. The sharpest result is the transition into the simplified self-employment regime: the probability of being observed *only* in monotributo rises steadily from essentially zero to +0.49 percentage points by event year +5 (s.e. 0.19), a quantitatively small but precisely estimated and monotonically growing margin.¹⁶ At the original employer, months fall by 0.33 per year by $t + 5$ (s.e. 0.20) and the probability of having left the original employer rises to +2.5 percentage points (s.e. 1.9); the separation point estimates grow monotonically across the post-period but are imprecise, with confidence intervals that include both zero and substantially larger effects. The number of distinct formal employers, by contrast, shows no systematic movement in the baseline specification. Since exit indicators are mechanically zero for every incumbent before the event, the pre-period cannot probe parallel trends for those outcomes; the months and earnings outcomes, whose pre-event coefficients are centered on zero, carry that burden.

To put these magnitudes in perspective, consider 1,000 incumbent workers at an automating firm, relative to comparable incumbents at firms that adopt only later. By event year +5, the point estimates imply 25 additional departures from the original employer: 15 of these workers move to another registered wage job, while 10 leave formal wage employment altogether. Of those 10, 5 appear in the simplified self-employment regime and 5 are no longer observed in any of our records, consistent with informal work or non-employment.¹⁷ Since the average adopter employs roughly 150 such incumbents, this amounts to about four additional separations from the firm, and between one and two additional departures from formal wage employment, per adopting firm. The months outcomes put the same adjustment on a time scale: in event year +5 alone, the 1,000-worker cohort spends 144 fewer months in registered wage employment and 329 fewer months at the original employer. These are persistent shifts in employment states for a small but non-negligible share of incumbents, not mass displacement—and they are fully consistent with the flat firm-level separations of Section 5.3. Most separations at any firm are of short-tenured workers, so as adopters cut hiring, the inflow of such workers and their later separations fall, holding the total separation count flat even as the stable incumbents tracked here grow more likely to leave. The flat firm total, in other words, is largely mechanical—fewer hires mean fewer workers who can later separate—and does not preclude the genuine, if modest, separation of attached incumbents that the worker panel isolates.

These margins matter in the Argentine context. The formal registry captures registered wage employment; the simplified regime captures a less formal status that typically comes with lower and more volatile income and weaker social protection than a registered salaried job, and workers who leave the registry without appearing in monotributo may be unemployed, out of the labor force, or working in a more informal environment that is not even captured in the simplified regime data. The evidence therefore points less to an immediate reported-wage loss than to a gradual loosening of ties to stable registered wage employment: fewer months, more exits, and a small but growing monotributo-only margin. Bessen et al. (2025) finds a similar pattern: the quantity of employment attachment adjusts more clearly than the wage rate. The difference is that in our data part of the displacement lands on a more precarious rung rather than in outright non-employment: the simplified

¹⁶The monotributo (simplified-regime) records cover 2011–2018 only, so this outcome is defined for the worker-years that fall within that window—8.4 million, against 11.5 million for the other panels. Later event times therefore rest on the earlier cohorts, whose post-event years still lie inside the coverage window.

¹⁷This decomposition cumulates three margins of differing precision—the exit from the original employer is imprecisely estimated, exit from the formal registry is significant only at the 10 percent level, and the monotributo margin is the most precise—so we read it as illustrative rather than as exact counts.

self-employment regime, where workers still earn an income but with weaker protection and less stable income than in registered wage employment—a status closer to informality, though still within the tax system. This is consistent with recent empirical evidence for Latin America, where the informal sector and self-employment absorb shocks that would otherwise raise unemployment more sharply (Brambilla et al., 2023; Dix-Carneiro and Kovak, 2019; Dix-Carneiro et al., 2021).



FIGURE 10 Worker-level attachment and reallocation outcomes

Notes: Each panel plots the treated $\times$ event-time interaction coefficients from the stacked event study, with 95 percent confidence intervals. Exit and simplified-regime outcomes are in percentage points and are mechanically zero for all incumbents before the event, so their pre-event coefficients are plotted at zero without confidence intervals; the simplified-regime (monotributo) outcome is observed for 2011–2018 only. Event time $t = -1$ is the omitted period. The specification uses worker weights, all sub-experiments, base fixed effects, and the full incumbent-worker sample; standard errors are clustered by anchor firm. *Source:* Own computations based on Argentine customs, SIPA, and monotributo records.

5.5 | Worker-Level Heterogeneity

It is interesting to study which incumbents bear the attachment margins documented above. To keep the splits comparable, the heterogeneity figures use the same six outcomes throughout: total earnings across registered wage jobs, exit from the formal registry, displacement to the simplified-regime-only status, total earnings at the original employer, exit from the original employer, and the number of formal employers. Corrective weights are recomputed within every group's effective estimation sample. Three splits organize the analysis: age, position in the pre-event wage distribution, and the size of the original employer. Together, these splits show that the near-zero average earnings response in Figure 9 is not an absence of effects but a cancellation of them: some incumbents gain and others lose, depending on their position in the wage distribution and the size of their employer.

Figure 11 splits workers by age at $t = -1$ (18–35 versus 36 and older). These estimates are hard to interpret: most margins move similarly across the two groups and are imprecisely estimated, so we interpret them with caution. Separation from the original employer is comparable (+2.3 and +2.7 percentage points at $t + 5$), and is different from zero only for the older group. The clearest difference is in the destination. Older incumbents appear somewhat more likely to end up only in the simplified self-employment regime (+0.60 percentage points at $t + 5$, s.e. 0.22, against +0.31 for the young), while younger incumbents

are slightly more likely to move to another registered employer—their count of distinct formal employers rises by 0.02—though these coefficients are very small. Earnings differ little between the groups. We do not draw a strong conclusion from this. One possibility is that older workers find it easier to shift into self-employment—because they are more skilled, better networked, or more used to operating on their own—but the data cannot separate this from many other explanations, and the magnitudes are modest throughout. At most, the split suggests that displacement takes a slightly different form by age: some registered-employer churn among the young, a small move toward self-employment among the old.




FIGURE 11 Worker heterogeneity by age at $t = -1$

Notes: Each panel overlays the treated $\times$ event-time interaction coefficients, with 95 percent confidence intervals, for the groups shown (18–35 versus 36 and older); groups are defined at $t = -1$. Corrective weights are recomputed within every group’s effective estimation sample, and sub-experiments without both treated and control workers in a group are dropped. Exit and simplified-regime outcomes are in percentage points and are mechanically zero before the event, so their pre-event coefficients are plotted at zero without confidence intervals; the simplified-regime (monotributo) outcome is observed for 2011–2018 only. Event time $t = -1$ is the omitted period. The specification uses worker weights, all sub-experiments, base fixed effects, and the full incumbent-worker sample; standard errors are clustered by anchor firm. *Source:* Own computations based on Argentine customs, SIPA, and monotributo records.

Figure 12 splits workers by the quartile of their observed pre-event wage, and the two ends of the distribution behave very differently. Workers in the bottom quartile are no more likely to leave the formal registry or their original employer after the event, and their earnings rise (+14 percent at $t + 5$); on this margin, automation does not appear to displace the lowest-paid incumbents—plausibly the least skilled, though we cannot verify that—and, among those who remain in registered work, may even leave them somewhat better off. The upper quartiles show the opposite: workers in the third and fourth quartiles are the most likely to separate from the original employer (+5 and +6 percentage points at $t + 5$) and migrate disproportionately into the simplified self-employment regime, while the earnings of those who stay in registered wage jobs do not fall. Whether these separations are involuntary we cannot say. Two features make pure forced displacement less than obvious: the workers who separate are the higher-paid, who tend to have more skills and outside options, and they land in monotributo—a structured, observed status—rather than

vanishing from the records into informality or non-employment; this is consistent with transitions that are at least partly chosen, into independent or self-employed work, though we cannot rule out involuntary separations that these workers simply absorb at lower cost. More broadly, the pattern cuts against a mechanical reading in which automation simply displaces the cheapest, least-skilled workers: the bottom quartile is, if anything, spared, and exposure is greater higher up the wage distribution. In that respect it echoes evidence that recent technological shocks concentrate away from the bottom of the distribution rather than on it, in the spirit of the job-polarization literature (Autor, 2015)—though not a clean case of it, since here exposure reaches the top quartile rather than hollowing out only the middle.




FIGURE 12 Worker heterogeneity by observed pre-event wage quartile

Notes: Each panel overlays the treated $\times$ event-time interaction coefficients, with 95 percent confidence intervals, for the groups shown; groups are defined at $t = -1$. Corrective weights are recomputed within every group's effective estimation sample, and sub-experiments without both treated and control workers in a group are dropped. Exit and simplified-regime outcomes are in percentage points and are mechanically zero before the event, so their pre-event coefficients are plotted at zero without confidence intervals; the simplified-regime (monotributo) outcome is observed for 2011–2018 only. Event time $t = -1$ is the omitted period. The specification uses worker weights, all sub-experiments, base fixed effects, and the full incumbent-worker sample; standard errors are clustered by anchor firm. *Source:* Own computations based on Argentine customs, SIPA, and monotributo records.

Figure 13 splits workers by the size of their original employer at $t = -1$ and reveals that firm size determines not the intensity of adjustment but its *margin*. Incumbents of larger adopters bear the quantity margin: separation from the original employer rises immediately and monotonically, reaching +6.6 percentage points by $t + 5$ (s.e. 2.7), exit from the formal registry is elevated from the event year onward (+0.8 to +1.9 percentage points over the window), and the simplified-regime-only margin climbs to +0.9 percentage points (s.e. 0.3), while their earnings drift upward imprecisely. Incumbents of smaller adopters show the mirror image: they are *less* likely to leave the original employer in the first post-event years, show no simplified-regime response, and instead show an earnings response—total earnings decline steadily to –4.0 percent by $t + 5$ (s.e. 1.2) across all jobs and –3.4 percent at the original employer. They are the only subgroup in our data with a significant negative earnings response. The imprecision on the large-firm side has a simple explanation. We

split incumbent workers into two halves of equal size, so each group contains the same number of workers but not the same number of firms: because large firms employ many workers each, the workers in the large-firm group come from only 208 firms, against 4,953 in the small-firm group. Since standard errors are clustered by firm, precision is governed by the number of firms rather than the number of workers, and the large-firm estimates are correspondingly less precise despite the many workers behind them.¹⁸

Argentina's collective-bargaining and wage-floor institutions do not constrain all firms equally: coverage and enforcement are weaker among small employers, which therefore have more room to adjust pay. Consistent with that, it is precisely at small adopters that an earnings response appears—incumbent annual earnings fall by 4 percent—while large adopters, where wage-setting constraints bite hardest, adjust through separations and transitions into self-employment instead. We read the magnitude cautiously, however, since annual earnings combine the monthly wage with months worked, and our worker-level outcomes cannot separate a lower wage rate from fewer months within the year.




FIGURE 13 Worker heterogeneity by initial firm size

Notes: Groups are the bottom and top halves of the worker distribution by the employment of the original employer at $t = -1$ (the median worker, not the median firm), so the large group covers far fewer firms. Each panel overlays the treated $\times$ event-time interaction coefficients, with 95 percent confidence intervals, for the groups shown; groups are defined at $t = -1$. Corrective weights are recomputed within every group's effective estimation sample, and sub-experiments without both treated and control workers in a group are dropped. Exit and simplified-regime outcomes are in percentage points and are mechanically zero before the event, so their pre-event coefficients are plotted at zero without confidence intervals; the simplified-regime (monotributo) outcome is observed for 2011–2018 only. Event time $t = -1$ is the omitted period. The specification uses worker weights, all sub-experiments, base fixed effects, and the full incumbent-worker sample; standard errors are clustered by anchor firm. *Source:* Own computations based on Argentine customs, SIPA, and monotributo records.

5.6 | Robustness to the Definition of an Automation Event

Our treatment is defined by two choices: the spike threshold γ and the basket of HS products classified as automation goods. This subsection varies each in turn. The two exercises are conceptually different, and we present them differently. Varying γ perturbs the sensitivity

¹⁸With 208 clusters, cluster-robust inference remains reliable; the cost of the split is precision, not validity.

of the event detector while holding the technology concept fixed; it is a robustness check in the narrow sense, and the message is that nothing depends on it. Varying the HS basket changes *which technology* defines the event—and with it the set of adopting firms, their event dates, and their incumbent workforces—so we read it as a technology-specificity analysis rather than a symmetric robustness check.

5.6.1 | The spike threshold γ

Figure 14 re-estimates the firm-level employment and median-wage event studies for $\gamma \in \{2, 2.5, 3, 3.5, 4\}$, rebuilding the treated set, the event dates, the clean-control pools, and the weights from scratch for each threshold. The five employment series are nearly indistinguishable: the event-year-+5 estimates range from -0.183 at $\gamma = 2$ to -0.216 at $\gamma = 4$, all within a standard error of the baseline -0.207 , and the pre-event coefficients remain centered on zero throughout. Median wages stay within one percentage point of zero for every threshold. Figure 15 repeats the exercise for four core worker outcomes. The series again overlap: at event year +5, exit from the original employer ranges from $+2.0$ to $+3.4$ percentage points across thresholds, months in the registry from -0.12 to -0.14 , the monotributo-only margin from $+0.37$ to $+0.53$ percentage points, and all-jobs earnings remain virtually flat. The gradient across thresholds is easy to read: a higher γ isolates larger, cleaner import spikes, so the treatment is sharper and the estimated effect slightly larger, while a lower γ admits smaller episodes that dilute it. Precision is essentially unchanged, and no conclusion in the paper depends on the choice of γ .




FIGURE 14 Robustness of firm-level results to the spike threshold γ

Notes: Each panel overlays the treated $\times$ event-time interaction coefficients, with 95 percent confidence intervals, for five spike thresholds $\gamma \in \{2, 2.5, 3, 3.5, 4\}$; outcomes are in logs. Each series rebuilds the stacked design—treated firms, event dates, clean controls, and the weights—from scratch under its own threshold, holding the HS basket at the baseline definition, so the baseline ($\gamma = 3$) reproduces Figures 5 and 6. Event time $t = -1$ is the omitted period. Estimates use corrective weights, and standard errors are clustered by firm. *Source:* Own computations based on Argentine customs and SIPA records.



FIGURE 15 Robustness of worker-level results to the spike threshold γ

Notes: Each panel overlays the treated $\times$ event-time interaction coefficients, with 95 percent confidence intervals, for five spike thresholds $\gamma \in \{2, 2.5, 3, 3.5, 4\}$. Each series rebuilds the worker-level stacked design from scratch under its own threshold, including the anchor-employer match and the incumbency condition, so the baseline ($\gamma = 3$) reproduces Figures 9 and 10. Exit and simplified-regime outcomes are in percentage points and are mechanically zero before the event; the simplified-regime (monotributo) outcome is observed for 2011–2018 only. Event time $t = -1$ is the omitted period. The specification uses worker weights, all sub-experiments, base fixed effects, and the full incumbent-worker sample; standard errors are clustered by anchor firm. *Source:* Own computations based on Argentine customs, SIPA, and monotributo records.

5.6.2 | The technology basket: from robots only to broad automation

The baseline treatment counts a broad basket of automation capital. To ask whether the results hinge on that breadth, we re-estimate them under progressively narrower baskets, down to industrial robots alone. These form five nested definitions: Definition 1 is industrial robots only—the strict Acemoglu–Restrepo notion; Definition 2 adds machining centres; Definition 3 adds numerically controlled (NC/CNC) machinery; Definition 4 adds automatic data-processing and computing equipment; and Definition 5, the baseline, adds control instruments and panels. Each basket contains the previous one. Narrowing the basket does not simply drop firms: because the spike detector runs on each firm’s redefined import series, events are re-dated as well—the first-spike year coincides with the baseline year for only 64 percent of Definition-1 adopters, 77 percent under Definition 2, 90 percent under Definition 3, and 84 percent under Definition 4. Each definition therefore delivers a different experiment, with its own adopters, cohorts, controls, and incumbent workers—not the same sample with fewer firms—and support shrinks sharply at the narrow end: the estimation sample contains 178 treated firm-by-cohort units under Definition 1 and 487 under Definition 2, against 2,883 in the baseline. Both features discipline the comparisons below.

The firm-level employment result survives every basket. Figure 16 shows negative, monotonically deepening employment paths under all five definitions, with event-year+5 estimates of -0.18 under the robots-only definition (s.e. 0.07), -0.15 under Definition 2, and -0.11 to -0.21 across the broader baskets. The robots-only estimate is noisy, as its sample size dictates, but its confidence interval excludes zero at the end of the window: firms whose automation events consist specifically of industrial-robot imports also contract registered employment. Median wages tell the same story as the baseline: they stay small under every basket—within about one to three percent of zero, and, apart from Definition 2 (-2.7 percent, s.e. 1.4), statistically indistinguishable from zero—so the firm-side response is a quantity adjustment, not a wage adjustment, whatever the definition.




FIGURE 16 Firm-level results under nested HS baskets: robots only to broad automation

Notes: Each panel overlays the treated $\times$ event-time interaction coefficients, with 95 percent confidence intervals, for five nested HS baskets; outcomes are in logs. Definitions are nested: Def. 1 = industrial robots; Def. 2 adds machining centres; Def. 3 adds NC/CNC machinery; Def. 4 adds automatic data-processing equipment; Def. 5 (baseline) adds control instruments. Because the spike detector runs on each redefined import series, definitions re-date events and re-select firms; the five series are distinct experiments on samples of very different size (178 treated firm-by-cohort units under Def. 1 vs. 2,883 in the baseline), not subsamples of one experiment. Event time $t = -1$ is the omitted period. Estimates use corrective weights, and standard errors are clustered by firm. *Source:* Own computations based on Argentine customs and SIPA records.

The worker-level patterns, in contrast, differ across baskets in an economically interesting way, and Figure 17 presents them as heterogeneity by technology rather than as a robustness check. Under the two narrow definitions, incumbent earnings across all jobs drift *upward* after the event: under the robots-only definition the path rises steadily to $+0.12$ log points at event years +4 and +5 (s.e. 0.06 and 0.05), statistically distinguishable from zero at the end of the window, with a similar but weaker profile under Definition 2 ($+0.08$, s.e. 0.05, not significant at the 5 percent level); under Definitions 3–5 the same coefficient is a tight zero. The exit margin is the mirror image: point estimates for leaving the original employer are *negative* under the narrow definitions (-2.6 and -0.3 percentage points at +5 under Definitions 1 and 2, neither significant), while the positive separation margin that characterizes the baseline is concentrated in the broad baskets ($+1.9$ to $+2.7$ percentage points). The remaining margins further reinforce this distinction. Under the robots-only definition, point estimates for months in the formal registry increase following the event, but they are not statistically distinguishable from zero. The probability of being observed only in monotributo also increases in most post-event periods, but these estimates are noisy; the exception is event year +5, when the effect turns negative. In contrast, the broader automation definitions produce the opposite pattern: formal-registry months tend

to fall, while monotributo-only participation rises, pointing to displacement from formal employment toward monotributo-only status.

One reading is that imports of robots specifically—the narrowest and most distinct technology—touch incumbents differently: their earnings tend to rise rather than fall, and they are no more likely to leave the firm, in contrast to the displacement that broad automation produces. This is consistent with robots complementing the workers who remain. We read it cautiously, however. The narrow definitions change more than the technology: because the spike detector runs on each firm's redefined import series, they also re-date events and re-select firms, cohorts, and incumbents, and the robots-only estimates are imprecise (332 anchor-firm clusters, with most attachment margins indistinguishable from the baseline). We therefore treat the pattern as suggestive—a sign contrast that recurs across the earnings and attachment margins, motivating a view of automation as a heterogeneous bundle rather than a single treatment—and leave it for future work with sharper within-firm technology variation to settle. What the exercise does establish is the asymmetry that matters for our headline results: the firm-level employment contraction holds under every basket, while the worker-level displacement margins are a property of broad automation adoption, not of robot adoption narrowly defined.




FIGURE 17 Worker-level results under nested HS baskets: heterogeneity by technology

Notes: Each panel overlays the treated $\times$ event-time interaction coefficients, with 95 percent confidence intervals, for five nested HS baskets (Def. 1 = robots only through Def. 5 = baseline). Each series rebuilds the worker-level stacked design under its basket, including event dates, anchors, and incumbency; definitions re-date events and re-select firms and workers (first-spike years coincide with the baseline for 64/77/90/84 percent of adopters under Definitions 1–4), so the series are distinct experiments, and the Definition-1 estimates rest on 332 anchor-firm clusters and carry wide confidence intervals. Exit and simplified-regime outcomes are in percentage points and are mechanically zero before the event; the simplified-regime (monotributo) outcome is observed for 2011–2018 only. Event time $t = -1$ is the omitted period. The specification uses worker weights, all sub-experiments, base fixed effects, and the full incumbent-worker sample; standard errors are clustered by anchor firm. *Source:* Own computations based on Argentine customs, SIPA, and monotributo records.

6 | CONCLUSION

This paper combines three administrative sources in Argentina—customs records on firm-level imports of automation-related capital goods, matched employer–employee records from the formal social-security registry (SIPA), and records from monotributo, the simplified self-employment tax regime—to trace how directly observed automation events reshape firm labor demand and the subsequent trajectories of incumbent workers over 2007–2022. At the firm level, an automation event is followed by a substantial contraction in registered labor demand: by event year +5, employment, positions, and average positions at adopters are 19 percent below comparable non-adopters, and the wage bill is 20 percent lower, while firm-level wages drift by less than one percent. That contraction runs through hiring rather than through dismissals: hires fall by 21 points of pre-event employment by event year +5, while separations never rise significantly above the control path at any horizon. Adopters stop replacing the workers who leave rather than shedding the ones they have, so a substantial part of the incidence falls on the workers the firm never hires, who are invisible in any analysis based on incumbents alone. Pre-event coefficients are centered

on zero, and the Argentine institutional setting—non-automatic import-licensing regimes (DJAI and SIMI), exchange-rate volatility, and binding credit constraints on imported capital goods—places a wedge between firms' desired adoption and the date at which an import spike is recorded, weakening the demand-driven adoption mechanism that would otherwise generate positive pre-trends.

The worker-level evidence reshapes the interpretation of this firm-level contraction. Reported earnings among incumbents who remain in registered wage employment do not move robustly in either direction, either across all employers or at the original employer. The clearest adjustment margin is attachment. In event year +5 alone, a cohort of 1,000 pre-event incumbents spends 144 fewer months—12 worker-years—in registered wage employment, and 329 fewer months, 27 worker-years, at the original employer. Exit from the formal registry rises by 1.0 percentage point and months in the registry fall by 0.14 per year, both distinguishable from zero only at the 10 percent level, while separation from the original employer rises by 2.5 percentage points, a sizeable point estimate but an imprecise one, with confidence intervals that include zero as well as substantially larger effects. The sharpest and most precisely estimated margin is the move into the simplified self-employment regime: the probability of being observed only in *monotributo* grows monotonically to +0.49 percentage points (s.e. 0.19) by the end of the window. These patterns are insensitive to the spike threshold γ that defines an automation event, and the firm-level employment contraction holds under all five nested baskets of automation goods. Heterogeneity points in the same direction: workers attached to larger adopters show the clearest detachment, and displacement is concentrated at the top of the pre-event wage distribution rather than at the bottom, with the probability of leaving the original employer rising monotonically across pre-event wage quartiles to 6 percentage points in the highest. Earnings are nowhere the principal margin of adjustment; the one group with a clear decline is incumbents of small adopters, whose annual earnings fall by 4 percent. The technology that defines the event matters more than the threshold that detects it: when the basket is narrowed to industrial robots alone, the worker-level displacement margins disappear and incumbent earnings drift upward instead. Those estimates rest on far fewer adopters and are correspondingly imprecise, but the contrast suggests that the displacement we document is a property of broad automation adoption rather than of robot adoption narrowly defined.

Read jointly, the evidence traces an adjustment that operates on the margin where institutions most plausibly matter. Automation weakens incumbents' attachment to formal wage employment—fewer months in the registry, more exits, and a growing simplified-self-employment margin—but reported wage earnings among incumbents who remain in registered employment do not fall in step. Adjustment in Argentina proceeds primarily through transitions across employment states—from registered wage work into simplified self-employment or out of the formal registry altogether, and, for those who remain in registered work, across employers—consistent with a labor market in which collective bargaining, wage floors, and payroll-tax institutions constrain within-job wage adjustment while informality and simplified self-employment provide a visible extensive margin. The reported-wage response should nonetheless be read with a measurement caveat: the administrative records we observe do not capture off-payroll compensation, which evidence from similar regulatory environments suggests is a non-trivial margin.

The broader implication is that, in developing economies with regulated formal wages and pervasive informality, the costs of automation need not materialize through observed payroll wage changes inside formal contracts. They can instead appear as a gradual erosion of formal labor-market attachment: fewer months in registered wage employment, higher exit probabilities, and migration into non-wage formal statuses. Wage-centered estimates that extrapolate from advanced economies to these contexts may therefore understate the

overall incidence of adjustment. Part of that incidence, moreover, never reaches incumbents at all: because the firm-level contraction operates through hiring rather than dismissals, those who bear it include the workers who are never hired, and whom no incumbent-based design can observe. Understanding which workers leave standard registered wage employment, where they go, how persistent those transitions are, and how the design of social-insurance and simplified-tax regimes shapes those destinations is a natural direction for future research on automation in developing countries.

REFERENCES

- Acemoglu, D. and Autor, D. (2011) Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics*, vol. 4, 1043–1171. Elsevier.
- Acemoglu, D. and Restrepo, P. (2018) The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, **108**, 1488–1542.
- (2020) Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, **128**, 2188–2244.
- (2022) Demographics and automation. *The Review of Economic Studies*, **89**, 1–44.
- Aghion, P., Antonin, C., Bunel, S. and Jaravel, X. (2020) What are the labor and product market effects of automation? new evidence from france. *Discussion Paper 14443*, CEPR.
- Autor, D. H. (2015) Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives*, **29**, 3–30.
- Autor, D. H., Levy, F. and Murnane, R. J. (2003) The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, **118**, 1279–1333.
- Bessen, J., Goos, M., Salomons, A. and Van den Berge, W. (2025) What happens to workers at firms that automate? *Review of Economics and Statistics*, **107**, 125–141.
- Bonfiglioli, A., Crino, R., Fadinger, H. and Gancia, G. (2024) Robot imports and firm-level outcomes. *The Economic Journal*, **134**, 3428–3444.
- Brambilla, I., César, A., Falcone, G. and Gasparini, L. (2023) The impact of robots in latin america: Evidence from local labor markets. *World Development*, **170**, 106271.
- Callaway, B. and Sant’Anna, P. H. C. (2021) Difference-in-differences with multiple time periods. *Journal of Econometrics*, **225**, 200–230.
- Carbonero, F., Ernst, E. and Weber, E. (2023) Robots worldwide: The impact of automation on employment and trade. *Applied Economics*.
- 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**, 1405–1454.
- de Chaisemartin, C. and d’Haultfœuille, X. (2020) Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, **110**, 2964–2996.
- Cheng, H., Jia, R., Li, D. and Li, H. (2019) The rise of robots in china. *Journal of Economic Perspectives*, **33**, 71–88.
- Dauth, W., Findeisen, S., Suedekum, J. and Woessner, N. (2021) The adjustment of labor markets to robots. *Journal of the European Economic Association*, **19**, 3104–3153.
- Davis, S. J. and Haltiwanger, J. (1999) Gross job flows. In *Handbook of Labor Economics*, vol. 3, 2711–2805. Elsevier.
- Dix-Carneiro, R., Goldberg, P. K., Meghir, C. and Ulyssea, G. (2021) Trade and informality in the presence of labor market frictions and regulations. *NBER Working Paper*.
- Dix-Carneiro, R. and Kovak, B. K. (2019) Margins of labor market adjustment to trade. *Journal of International Economics*, **117**, 125–142.
- Faber, M. (2020) Robots and reshoring: Evidence from mexican labor markets. *Journal of International Economics*, **127**, 103384.
- Feinmann, J. (2024) Payments under the table.
- Goodman-Bacon, A. (2021) Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, **225**, 254–277.

- Graetz, G. and Michaels, G. (2018) Robots at work. *Review of Economics and Statistics*, **100**, 753–768.
- Humlum, A. (2022) Robot adoption and labor market dynamics. *Rockwool Foundation Research Unit*.
- Koch, M., Manuylov, I. and Smolka, M. (2021) Robots and firms. *The Economic Journal*, **131**, 2553–2584.
- Sun, L. and Abraham, S. (2021) Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, **225**, 175–199.
- Wing, C., Freedman, S. M. and Hollingsworth, A. (2024) Stacked difference-in-differences. *NBER Working Paper 32054*, National Bureau of Economic Research.

A | ADDITIONAL FIGURES

A.1 | Automation-Importer Profiles




FIGURE A.1 Share of Automation-Importer Firms and of Automation-Import Value, by Import Pattern

Notes: The figure reports, for each import pattern, the share of automation-importer firms and the share of total automation-import value. “Spike” firms experience a discrete automation-import spike; “large regular” firms never spike but import automation goods in at least three years with an average annual value above the median of regular non-spiking importers; “small or occasional” firms are the remaining non-spiking importers. Import values are nominal customs values in U.S. dollars, summed over the period. *Source:* Own computations based on Argentine customs records.

A.2 | Sectoral and Geographic Distribution of Automation-Spike Firms




FIGURE A.2 Sectoral Distribution of Automation-Spike Firms

Notes: The figure reports the number of automation-spike firms by broad sector and two-digit activity. Sector classifications follow the firm's main activity in SIPA. *Source:* Own computations based on Argentine customs and SIPA records.




FIGURE A.3 Geographic Distribution of Automation-Spike Firms

Notes: The figure reports the number of automation-spike firms by province. To assign exposure precisely, the map uses firms that operate in a single province during the relevant period. Firms observed in multiple provinces are excluded from the map. *Source:* Own computations based on Argentine customs and SIPA records.

A.3 | Labor-Market Informality




FIGURE A.4 Composition of Employment by Labor-Market State, 2007–2022

Notes: The gap corresponds to 2015H2–2016H1, when the survey was not fielded. *Source:* EPH (INDEC), harmonized by SEDLAC (CEDLAS–World Bank).

A.4 | Composition and Scale of Automation Events




FIGURE A.5 Share of Automation-Spike Events by Dominant Good

Notes: The figure reports, for the twelve most common dominant codes, the share of baseline automation-spike events assigned to each. An event's dominant code is the automation code with the largest import value in the firm's spike year. Because less common codes are omitted, the displayed shares do not sum to 100 percent. *Source:* Own computations based on Argentine customs records.




FIGURE A.6 Share of Events and of Spike-Year Automation Value by Robot-Proximity Tier

Notes: The figure reports the share of events and of spike-year automation value by robot-proximity tier. Broader automation capital goods account for more than three quarters of both; strict robots for about two and a half percent. *Source:* Own computations based on Argentine customs records.




FIGURE A.7 The Size of the Automation Purchase at the Adoption Event

Notes: The figure reports the median value of the entire automation purchase made at the adoption event, by assigned product family. Panel (a) expresses it in nominal FOB US dollars; panel (b) expresses it as a multiple of the firm's average annual gross wage per worker in the year before the event, with the dashed line marking one worker-year. Physical quantities are not reported for programmable controllers because they are almost never declared in units. *Source:* Own computations based on Argentine customs and SIPA records.




FIGURE A.8 International Benchmarks and the Argentine Robot Stock

Notes: Panel (a) compares the task-level engineering anchor with net equilibrium estimates from the literature: G-M = [Graetz and Michaels \(2018\)](#); Dauth = [Dauth et al. \(2021\)](#); A-R = [Acemoglu and Restrepo \(2020\)](#). Panel (b) reports the operational stock of industrial robots in Argentina. *Source:* IFR World Robotics and own computations.

B | ADDITIONAL TABLES

B.1 | Composition and Scale of Automation Events

TABLE B.1 Codebook of the baseline automation measure and robot-proximity tiers

Code	Customs description	Description	Proximity
<i>Robots</i>			
847950	Industrial robots, n.e.s.	Free-standing industrial robots	Strict robot
<i>Machining centres</i>			
845710	Machining centres, metal	Multi-tool CNC cutting cells	CNC / automatic
<i>CNC machine tools</i>			
845811	Horizontal lathes, NC	CNC horizontal lathes	CNC / automatic
845819	Horizontal lathes, other	Other turning centres	CNC / automatic
845921	Drilling machines, NC	CNC drilling machines	CNC / automatic
845931	Boring-milling machines, NC	CNC boring/milling	CNC / automatic
845951	Milling machines, knee-type, NC	CNC knee-type mills	CNC / automatic
845961	Milling machines, other, NC	Other CNC mills	CNC / automatic
846011	Flat-surface grinders, NC	CNC surface grinders	CNC / automatic
846021	Other grinding machines, NC	Other CNC grinders	CNC / automatic
846031	Sharpening (tool/cutter) machines, NC	CNC tool/cutter grinders	CNC / automatic
846221	Bending/folding/straightening, NC	CNC bending/folding presses	CNC / automatic
846231	Shearing machines, NC	CNC shears	CNC / automatic
846239	Shearing machines, other	Other CNC shears/presses	CNC / automatic
846241	Punching/notching machines, NC	CNC punching machines	CNC / automatic
<i>Computing (automatic data processing)</i>			
847141	ADP machines, CPU with I/O units	Complete desktop computers	Broader automation
847149	ADP machines presented as systems	Computer systems	Broader automation
847150	ADP processing units	Servers / processing units	Broader automation
847160	ADP input or output units	Monitors, keyboards, printers	Broader automation
847170	ADP storage units	Disk / storage units	Broader automation
847180	Other ADP units	Other peripherals	Broader automation
<i>Automatic control instruments and controllers</i>			
903210	Thermostats	Temperature controllers	Broader automation
903220	Manostats	Pressure controllers	Broader automation
903281	Regulating instr., hydraulic/pneumatic	Process controllers	Broader automation
903289	Automatic regulating instruments, other	Other process controllers	Broader automation
903290	Parts for automatic control instruments	Control-instrument parts	Broader automation
8537.10.20	Programmable controllers	PLC controllers	Broader automation
8537.10.11	Numerical controllers (CNC), high-spec	CNC machine controllers	Broader automation
8537.10.19	Numerical controllers (CNC), other	Other CNC controllers	Broader automation

Notes: Codes are six-digit tariff headings except the programmable and numerical controllers, which are eight-digit NCM sublines. “NC” / “ADP” abbreviate numerically controlled and automatic data processing. Codes 846011 and 846021 enter only through 2016 owing to a tariff-nomenclature revision. Proximity tiers group the codes by closeness to a true industrial robot.

Source: Own classification based on Argentine customs records and [Acemoglu and Restrepo \(2022\)](#).

TABLE B.2 Automation-spike events by product family and robot-proximity tier

Product family	% events	% value
<i>Strict robots</i>		
Industrial robots	2.5	3.7
<i>CNC / automatic machinery</i>		
Machining centres	5.9	6.5
CNC machine tools	13.9	13.0
<i>Broader automation capital goods</i>		
Computing (ADP)	30.6	60.1
Control instr. & controllers	47.2	16.7
<i>Strict robot</i>	2.5	3.7
<i>CNC / automatic machinery</i>	19.7	19.5
<i>Broader automation</i>	77.7	76.8
Total	100.0	100.0

Notes: Each event is assigned to the product family of its dominant code, defined as the code with the largest automation-import value in the firm's spike year. The value column reports the share of total spike-year automation value (nominal FOB, US dollars) accounted for by the events assigned to each family. Shares are calculated across all baseline events. The lower panel aggregates the five families into the three proximity tiers. Tier totals are calculated from unrounded family shares, so displayed sums may differ slightly because of rounding.

Source: Own computations based on Argentine customs records.

TABLE B.3 Imported automation goods by product family

Product family	Transactions	% of trans.	Importing firms	Value (US\$m)	% of value	Median unit price (US\$)
Computing and data processing	183054	34.8	7153	8691.1	59.8	147
Automatic control instruments	297426	56.6	8590	4007.5	27.6	67
CNC machine tools	10217	1.9	2749	884.2	6.1	23803
Machining centres	2679	0.5	1194	387.7	2.7	65500
Programmable and numerical controllers	30181	5.7	3767	376.1	2.6	–
Industrial robots	1818	0.3	564	196.1	1.3	29100
Total	525375	100.0	16016	14542.7	100.0	

Notes: Argentine customs records, 2007–2022. The table covers all transactions involving automation goods and all firms importing them, regardless of whether they enter the estimation sample. Values are nominal FOB US dollars. Transaction counts refer to customs records rather than physical units. The unit price is the median transaction value divided by the declared physical quantity for product lines whose statistical unit is a count of units. All transactions involving industrial robots, machining centres, and CNC machine tools report quantities in units, so their median prices use the full set of transactions. Programmable and numerical controllers are almost never declared in units, so no median unit price is reported. Firm counts across families do not sum to the total because firms may import goods from more than one family. Percentages may not sum exactly because of rounding.

Source: Own computations based on Argentine customs records.

TABLE B.4 Spikes by product family

Family generating the spike	% of events	Median spike (US\$)	Median units	Worker-years	Payback (months)	Displ. needed (%)
Control instruments	34.9	2717	5	0.21	2.5	0.9
Computing/data processing	30.0	14855	15	1.10	13.2	6.1
CNC machine tools	14.0	61500	1	6.36	76.4	29.0
Programmable controllers	12.6	5692	–	0.44	5.3	1.3
Machining centres	6.0	78877	1	7.79	93.4	48.2
Industrial robots	2.6	52636	1	4.09	49.1	11.7
All spikes	100.0	11860		0.95	11.4	4.1

Notes: The spike value is the firm's whole automation-import purchase in the event year, not only the dominant family. Units are the physical quantity of the dominant family declared to customs; programmable and numerical controllers are almost never declared in units. A worker-year is the annual gross wage of one worker at the firm in the year before the event, in nominal dollars of that year; the median is 12,486 dollars. Payback is the time the purchase takes to repay itself if it allows the firm to shed exactly one worker. Displacement needed is the share of the firm's own pre-event workforce that would have to go for the purchase to repay itself within one year, at the median firm in each family.

Source: Own computations based on Argentine customs and SIPA records.

C | OTHER DESCRIPTIVE STATISTICS

C.1 | Anatomy of the Automation Spikes

The baseline treatment combines a broad set of automation-related imports. We therefore document both the goods included in the measure and their contribution to the resulting events. The measure aggregates twenty-nine tariff codes into five product families: industrial robots; metal-working machining centres; numerically controlled (CNC) machine tools such as lathes, drills, mills, grinders, shears and presses; automatic data-processing (computing) equipment; and automatic regulating instruments together with the programmable and numerical controllers that steer automated machinery. Table B.1 lists every code, its customs description and a plain-language gloss, and groups the codes into three tiers according to their technological proximity to industrial robots: *strict robots* (the single robot heading), *CNC and automatic machinery* (machining centres and computer-controlled machine tools, which are programmable but task-specific), and *broader automation capital goods* (computing equipment, automatic regulating instruments, and programmable or numerical controllers).

An automation event is defined at the firm level from these imports, following the spike rule in equation 1: a firm-year is flagged as a spike when automation imports scaled by the firm's average wage bill exceed three times the firm's own average automation intensity across its other years, and the firm's first spike is its event. Applied to the full measure this rule yields the set of baseline events over 2008–2022, and everything that follows decomposes those same events.

To attribute each event to a product, we rank by value the automation codes imported by the firm in its spike year. We define the code with the largest value as the event's dominant code and classify events containing more than one automation code as multi-code. Roughly 28 percent of events are multi-code. Yet concentration within events is high: on average, the dominant code accounts for 94 percent of spike-year automation value. Table B.2 reports the resulting distribution across the five product families, while Figure A.5 shows the twelve most common dominant codes.

The decomposition is lopsided, and not in the direction the term “automation” might suggest. The most common dominant code is automatic regulating instruments, accounting for 23.6 percent of all events, followed by programmable logic controllers at 10.9 percent. Taken together, automatic regulating instruments and programmable and numerical controllers dominate 47.2 percent of events; events assigned to this family account for 16.7 percent of aggregate spike-year value. Computing equipment shows the opposite pattern: it dominates 30.6 percent of events, while computing-dominated events account for 60.1 percent of aggregate value. Events dominated by machining centres or CNC machine tools account for 19.7 percent of events and 19.5 percent of value. Industrial robots rank last among the twelve codes shown in Figure A.5, dominating only 2.5 percent of events; robot-dominated events account for 3.7 percent of value.

Grouping the codes by their proximity to industrial robots clarifies the composition of the baseline treatment (Figure A.6). Events assigned to broader automation capital goods—computing equipment, control instruments and controllers—account for 77.7 percent of events and 76.8 percent of aggregate spike-year automation-import value. Events assigned to CNC and automatic machinery account for 19.7 percent of events and 19.5 percent of value, while robot-dominated events account for just 2.5 percent and 3.7 percent, respectively. The baseline treatment therefore captures threshold crossings in broadly defined automation-related capital much more often than robotization in the narrow sense. The estimated employment and wage-bill responses should be interpreted accordingly: as responses to this broad measure rather than to industrial robots alone.

C.2 | Economic Magnitudes

The anatomy of the spikes raises a natural question of scale: how much labor could the purchase behind a typical adoption episode plausibly replace or complement? Customs records report import values and, for most product lines, physical quantities, allowing us to characterize both total expenditure and the number and cost of the units imported. Before turning to event-level purchases, Table B.3 describes all automation imports across six product categories. For this magnitude analysis, we separate programmable and numerical controllers from the automatic control instruments with which they were grouped in Section C.1, because controllers are almost never declared in physical units. Computing equipment and automatic control instruments jointly account for 87.3 percent of total value and 91.5 percent of customs transactions. Their individual units are relatively inexpensive: the median imported unit costs 147 US dollars for computing equipment and 67 dollars for automatic control instruments. These categories include keyboards, monitors, storage units, thermostats, pressure regulators and process-control devices, often purchased in bulk. Industrial robots, by contrast, account for 1.3 percent of value and 0.3 percent of transactions, but the median imported robot costs 29,100 dollars and the median machining centre 65,500 dollars. The measure therefore combines a small number of expensive machines with a much larger number of inexpensive components and devices.

To connect product composition to adoption events, the relevant unit is the entire spike-year purchase rather than the price of an individual item. Unit prices alone can understate the scale of an event when a purchase bundles many inexpensive devices. We therefore characterize each event by its total automation-import value in the spike year, the number of physical units where reported, and the product family accounting for the largest share of the purchase. To compare this expenditure with the firm's wage scale, we express the whole spike purchase in *worker-years*: its value divided by the annual wage of one worker at the same firm, with both quantities measured in nominal US dollars for the corresponding year.¹⁹

Table B.4 and Figure A.7 report the result. The two families that generate nearly two thirds of all adoption events—control instruments at 34.9 percent and computing equipment at 30.0 percent—also generate spikes that are far smaller than those associated with the machine families. The median control-instrument spike is 2,717 dollars and five physical units, or 0.21 worker-years, while the median computing spike is 14,855 dollars and fifteen units, or 1.10 worker-years. Machine-family events are less common but substantially larger: the median machining-centre spike is 78,877 dollars and 7.79 worker-years, while the median CNC machine-tool spike is 61,500 dollars and 6.36 worker-years. Robot spikes sit between these categories, at 52,636 dollars and 4.09 worker-years. For each of the three machine families, the median reported quantity of the dominant product is one unit.²⁰

The small size of the typical spike does not by itself determine how much labor the technology may replace or complement. Purchase prices need not map one-for-one into

¹⁹The wage is the firm's nominal wage bill in the year before the event divided by its employment in that year and converted into US dollars at that year's average official exchange rate. The median annual wage per worker at treated firms is 12,486 dollars, ranging from about 7,800 dollars for events early in the sample to 18,600 in 2013 and back to about 8,200 in 2020. Employer contributions, which would raise the firm's labor cost by roughly one quarter, are left out of the figures below, which are therefore conservative in the sense of understating the value of the labor a purchase could replace.

²⁰The medians conceal substantial heterogeneity within the control-instrument family. Almost 60 percent of control-instrument spikes involve fewer than ten instruments, with a median purchase of 1,500 dollars. At the other extreme, 8.4 percent involve a thousand units or more, with medians of 4,598 units and 24,510 dollars, consistent with plant-scale instrumentation programs rather than purchases of a few gauges. Manufacturing accounts for 54.8 percent of the firms in this group, the highest manufacturing share across purchase-size classes.

employment effects, since equipment can reorganize tasks and generate savings or productivity gains over several years. The last two columns of Table B.4 therefore provide accounting benchmarks rather than estimates of actual displacement. The median control-instrument purchase is equivalent to 2.5 months of one worker's gross wage, and the median programmable-controller purchase to 5.3 months. Relative to the buyer's pre-event wage bill, the median purchase represents 0.9 percent for control instruments and 6.1 percent for computing equipment. The corresponding shares are substantially larger for the machine families: 48.2 percent for machining centres, 29.0 percent for CNC machine tools, and 11.7 percent for robots. Overall, just over half of all spikes cost no more than the annual gross wage of one worker at the importing firm.

Two cautions frame this accounting exercise. First, estimates from the international literature provide useful context but are not direct benchmarks. Graetz and Michaels (2018) find no significant effect of robot adoption on total employment across countries, despite substantial productivity gains. Dauth et al. (2021) estimate that each additional robot displaces two manufacturing jobs in Germany, but find that these losses are fully offset by employment gains in services. For the United States, Acemoglu and Restrepo (2020) estimate a reduction of about 6.6 jobs within exposed local labor markets and 3.3 jobs nationally per additional robot. These studies estimate net equilibrium effects in high-income economies with factor prices, plant scales, and capital vintages that differ sharply from Argentina's. They also concern industrial robots, which account for only 2.6 percent of the events in our data; Argentina's operational stock reached about 3,600 units by 2022 (Figure A.8, panel (b)). The worker-year and payback measures therefore complement rather than supersede these estimates: they place purchase values on the wage scale of firms in our sample, but do not measure actual employment displacement.

Second, descriptive employment paths around the event cannot identify an employment effect: they lack a counterfactual and are affected by substantial attrition. Among firms observed in the year before the event, 89.8 percent remain observed one year after it, 71.9 percent three years after, and 57.0 percent five years after. This attrition also limits the interpretation of the stacked estimates. Because the baseline stack requires units to be observed in every period from $e = -3$ to $e = 5$, the estimates apply to firms that remain in the registry throughout the event window; the comparison design does not by itself address selection into this balanced sample. The product data play a narrower role: they establish the composition and scale of the treatment and clarify the population to which the estimates apply. Valued at the median imported-robot price, continuous importers that the design does not treat purchase the equivalent of about four robots per year, compared with 0.4 robot-equivalents at the median treatment event. The estimates therefore pertain to moderate, lumpy automation purchases dominated by instrumentation and computing rather than to sustained investment in industrial robotics.

D | ROBUSTNESS

D.1 | Gross Flows: Alternative Denominator and Levels among Adopters

Figure D.1 re-estimates the hiring and separation event studies dividing each flow by previous-year employment instead of pre-event employment. Because the lag is unavailable at $e = -3$, these series start at $e = -2$. The message of Figure 8 is unchanged: hiring falls significantly after the event (-5 points of previous-year employment at $e + 5$), while separations never rise. Note that with a shrinking denominator the lag-based rates understate the decline in flow levels; the pre-event denominator in the main text avoids exactly that mechanical effect.




FIGURE D.1 Hiring and separations divided by previous-year employment

Notes: Each panel plots the treated $\times$ event-time interaction coefficients, with 95 percent confidence intervals; each flow is divided by previous-year employment, so the series begin at $t = -2$ (the lag is unavailable at $t = -3$). Hires and separations are adjacent-year worker flows (a hire is present in t but not $t - 1$; a separation present in t but not $t + 1$), undefined at the panel borders, so flows in a firm's first and last observed year are dropped. Event time $t = -1$ is the omitted period. Estimates use corrective weights, and standard errors are clustered by firm. *Source:* Own computations based on Argentine customs and SIPA records.

Figure D.2 reports the descriptive counterpart requested in the CAF review: raw hiring and separation levels among adopting firms only, with one-standard-error bands around each mean and no treatment-control comparison. These panels are descriptive and carry no causal interpretation; they show the scale on which the event-study rates operate, and they are discussed in Section 5.3.




FIGURE D.2 Worker flows among adopting firms: levels (descriptive)

Notes: Treated firms only; error bars show $\pm$ one standard error of the mean. Panels (c)–(d) plot small (left axis) and large (right axis) adopters on separate scales. Descriptive evidence; no causal interpretation. Reference year is $t = -1$. Hires and separations are adjacent-year worker flows (a hire is present in t but not $t - 1$; a separation present in t but not $t + 1$), undefined at the panel borders, so flows in a firm's first and last observed year are dropped. *Source:* Own computations based on Argentine customs and SIPA records.