

Labor Hoarding Policies Can Trigger Creative Destruction*

Elías Albagli[†]

Andrés Fernández[‡]

Juan Guerra-Salas[§]

Federico Huneus[¶]

September 2026

Abstract

During economic crises, firms face a trade-off between labor hoarding—retaining workers despite reduced demand—and labor shedding, which can improve efficiency through “creative destruction.” Employment protection policies, which subsidize labor costs to preserve employment relationships, are often criticized for blocking such efficiency gains. Using administrative microdata from a large-scale policy implemented in Chile during the COVID-19 pandemic—which allowed firms to furlough workers at zero hours while keeping their contracts—we show that such a policy can instead facilitate creative destruction. Firms used the subsidized suspension of contracts to selectively retain and recall higher-ability workers from furlough, while letting lower-ability workers separate. Treated firms subsequently experienced economically significant and persistent productivity gains of around 5–10% over a two-year horizon, demonstrating that employment protection policies can simultaneously preserve jobs and promote productivity-enhancing reallocation.

Keywords: Labor hoarding; furlough; selective recall; creative destruction; productivity; worker ability.

JEL classification: J65, J24, E32, D24, L25.

*For superb research assistance, we thank Mario Canales, Camila Gómez, Daniel Guzmán, Montserrat Martí, Pablo Muñoz and Manuel Taboada. We also thank Cristina Arellano, Davide Berger, Sebastian Buhai, Martin Ellison, Giulia Giupponi, Enrique Mendoza, and seminar participants at Banco de la República, Banco Central de Chile, Ecuadorian Economic Society, Fordham University, IDB, International Economic Association, IMF, LACEA, the School of Government and Economics Department at Pontificia Universidad Católica de Chile, Sociedad de Economía de Chile, Society for Computational Economics, and Universidad de Chile for useful comments. The views expressed are those of the authors and do not necessarily represent the views of the IMF, its Executive Board, or IMF management, nor those of the Central Bank of Chile or its board members.

[†]Central Bank of Chile. E-mail address: ealbagli@bcentral.cl.

[‡]International Monetary Fund. E-mail address: afernandez3@imf.org.

[§]Central Bank of Chile. E-mail address: jguerra@bcentral.cl.

[¶]Duke University and Central Bank of Chile. E-mail address: federico.huneus@duke.edu.

1 Introduction

A long-standing debate during economic crises is whether firms should hoard labor—retain workers despite reduced demand—or shed labor to facilitate reallocation and efficiency gains through “creative destruction.” While governments often subsidize job retention, a standard critique is that employment protection during crises may entrench low-quality matches and hinder productivity-enhancing reallocation. This paper shows that, under a policy design that combines subsidized contract suspension with recall flexibility, labor hoarding policies can instead facilitate reallocation and productivity improvements.

We document this finding using administrative microdata for a large employment protection program implemented in Chile during the COVID-19 pandemic—*Ley de Protección del Empleo* (LPE). LPE allowed firms to furlough workers at zero hours for up to six months while maintaining the employment relationship, prohibited remote work while on furlough, financed benefits through the unemployment insurance system, and required firms to continue paying pension and health contributions (about 20% of wages). This cost-sharing element created meaningful incentives for firms to be selective about whom to enroll and whom to recall.

The policy was deployed at broad scale: almost all formal firms were eligible, about half enrolled at least one worker, and firms that accessed the policy initially enrolled, on average, roughly 60% of their workforce. This combination of flexible design and widespread take-up makes LPE an informative setting to study whether policy-enabled suspension and selective recall can reconcile employment protection with productivity-enhancing reallocation.

We combine records of the employment protection policy with administrative tax data, credit-registry data, and employer-employee records at a monthly frequency from 2019 to 2022, covering the universe of formal firms in Chile.¹ Our empirical strategy proceeds in two stages. First, we implement a matching procedure, Coarsened Exact Matching (CEM), on rich pre-treatment firm characteristics from tax and credit records. Second, we estimate event-study difference-in-differences (DiD) using the CEM strata and weights, with fixed effects that absorb stratum-by-calendar-month shocks and event-time patterns. We restrict attention to firms that accessed the policy in May–June 2020—months of early and high take-up—and trace outcomes over a 24-month horizon. The event-study estimates show flat pre-trends for the outcomes we analyze, which supports the identifying assumptions of the empirical design.

We document two core findings. First, treated firms experience persistent productivity gains relative to matched controls. Both labor productivity and total factor productivity (TFP) rise by 5-10% over a two-year horizon, reaching about 10% two years after treatment and well after the high-take-up window had ended. These gains are reflected in higher sales without corresponding increases in labor or capital inputs, relative to controls. While employment does not change significantly two years after treatment, capital declines gradually and ends about 10% lower among treated firms.

Second, we uncover a within-firm selection mechanism on worker ability that helps account for

¹The rate of informality in Chile is around 25% of workers and lies in the low range for an emerging country.

the increase in productivity. Using worker fixed effects from Abowd–Kramarz–Margolis (AKM) regressions estimated on 2004–2023 employer–employee records as a proxy for ability, we document systematic ability-based selection in treated firms. Workers who stayed in treated firms have the highest estimated ability. Among furloughed workers, those recalled to the firm have higher ability than those not recalled. Separations without furlough are concentrated among lower-ability workers. Recall is incomplete: by late-2021, more than 20% of enrolled workers had not been recalled, consistent with policy-facilitated shedding of lower-ability matches. The ability gaps are economically large: enrolled-and-recalled workers display 0.25 log points lower ability than stayers, while enrolled-but-not-recalled workers display 0.45 log points lower ability. This selection leads to a persistent increase in firm-level average worker ability of about 2% in treated firms relative to controls.

Comparing treated and control firms reinforces this interpretation. The ability gap between stayers and separated workers is significantly smaller in treated firms than in controls, indicating that the policy raised the ability threshold below which firms separate workers. New hires at treated firms also display higher ability relative to their stayers than new hires at control firms, suggesting that the policy improved selection on the hiring margin as well. Together, these patterns indicate that the policy enabled a more comprehensive restructuring of workforce composition—not only through selective separation but also through higher-quality hiring.

If ability-based selection yields persistent improvements in productivity and sales, why is similar reorganization not implemented in the absence of the policy? We argue that three design features of LPE are central. First, the policy lowered the effective cost of maintaining employment relationships during the downturn: although furloughed workers supplied zero hours, firms continued to pay only social contributions, around 20% of wages. Second, employment contracts remained intact throughout the furlough period. Third, firms retained full discretion over which workers to enroll, which to recall, and which to separate. Together, these features relaxed adjustment frictions and enabled firms to re-optimize workforce composition through selective recall and non-recall rather than relying solely on immediate separations. By temporarily suspending a large share of employment relationships while preserving the option to recall workers at relatively low cost, the policy also created scope for firms to reassess how workers fit into their production process before making more permanent staffing decisions. While the data do not allow us to distinguish whether firms used the furlough period to learn about individual worker ability or to reassess workers’ roles in a reorganized production process, both channels are consistent with the observed selection patterns and the subsequent productivity gains.

Finally, we explore heterogeneity to assess where the policy and its mechanism are most relevant. Productivity gains are larger among smaller firms and statistically insignificant among large firms. Across industries, effects are strongest in manufacturing, wholesale and retail trade, and transport, and weaker or statistically insignificant in other services, construction, and mining.

We contribute to three strands of the literature. First, the literature that studies policies aimed at protecting employment during downturns. On one side, there are short-time work (STW) poli-

cies, which are designed to preserve employment relationships during downturns by subsidizing reduced hours rather than severing contracts (Giupponi et al., 2022; Eichhorst et al., 2022). On the other hand, there are job-retention policies that preserve employer–employee links and thereby limit “loss-of-recall” transitions from temporary separation into persistent joblessness (Gertler et al., 2025). Pre-COVID firm-level evidence is thinner and mixed across settings (Kruppe and Scholz, 2014; Cooper et al., 2017; Cravo et al., 2020; Kato and Kodama, 2019). A renewed debate, using episodes beyond COVID, stresses heterogeneity, selection, and the margin of adjustment: for France, Cahuc et al. (2024) document that STW can preserve jobs for firms facing severe revenue shocks but generate windfall take-up among less affected firms, highlighting the importance of targeting; for Germany, Brinkmann et al. (2024) find limited effects on employment retention from extensions in benefit duration but persistent wage effects, consistent with adjustment operating through bargaining and wage-setting rather than pure labor-hoarding. Our study contributes to this literature by showing that a highly flexible STW design can generate not only employment protection but also persistent firm-level productivity gains, operating through a worker selection mechanism enabled by subsidized contract suspension and selective recall.

A second strand studies the determinants of cyclical labor-market adjustment and labor hoarding. Theoretical foundations emphasize separation costs and match-specific investments that make retaining workers privately optimal even when demand is low (Malcomson, 1997), and more recent work highlights how rent-sharing and implicit contracts shape labor adjustment over the cycle (Fahn and Murooka, 2025). The LPE policy shock is informative through this lens because it effectively lowers the cost of maintaining the employment relationship while preserving firms’ discretion over whom to furlough, whom to recall, and whom to separate—creating scope for sorting on worker ability within firms. We contribute to this literature by providing direct evidence of how labor hoarding policies impact the reallocation of workers.

A third and larger branch concerns cleansing effects, jobless recoveries, and productivity dynamics in crises. Empirically, crises can trigger compositional filtering and productivity-enhancing reallocation (e.g., Ates and Saffie, 2021), though labor hoarding may dampen this margin and the strength of cleansing can vary across sectors (den Bosch and Vanormelingen, 2023). Related work links employment protection to lower worker flows, reduced entry, and lower TFP (Autor et al., 2007). Our findings relate to the classic work on creative destruction by Caballero and Hammour (1996), who argue that frictions can slow the efficient reallocation of jobs and technologies over the cycle and that creation incentives can be complemented by policies that encourage firms to keep their workers. Consistent with this logic, under design features that preserve contracts at low cost while maintaining flexibility in recall and separation, LPE can be seen as having reduced reallocation frictions and facilitated cleansing rather than entrenching low-quality matches. In standard search-and-matching frameworks (Mortensen and Pissarides, 1994), firing/hiring frictions and bargaining imply that the net effect of employment protection on productivity is theoretically ambiguous, since thresholds for separations, hiring, and match formation adjust endogenously (see also Autor et al., 2007). Complementary models highlight how worker heterogeneity, asymmetric information, and

selective firing can amplify within-firm selection over the cycle (Berger, 2018; Molavi, 2018) and how search frictions generate potentially large misallocation costs with aggregate TFP consequences (Bilal et al., 2022).²

COVID revived these questions because the shock plausibly combined large sectoral reallocation with temporary separations and recalls: early commentary framed the pandemic as a reallocation shock and warned that retention subsidies could slow efficient adjustment (Barrero et al., 2020), while descriptive evidence emphasized unusual patterns in job creation and restructuring (Haltiwanger, 2022) and the central role of temporary layoffs and recalls in the U.S. (Forsythe et al., 2022). Firm-level evidence during COVID also points to heterogeneity across the productivity distribution and suggests that support policies may offset cleansing margins (Kozeniauskas et al., 2022), while structural work incorporating recall unemployment highlights conditions under which employment losses need not be persistent when recalls are feasible and subsidies stabilize matches (Buera et al., 2021). Our contribution to this literature is to document a distinct selection margin: subsidized furlough can allow sorting—via selective recall and non-recall—without relying exclusively on firm exit or immediate separations as the cleansing channel.

The rest of the paper is organized as follows. Section 2 provides institutional background and describes the data. Section 3 outlines the empirical strategy. Section 4 documents firms’ access to the policy. Section 5 presents the main results. Section 6 reports robustness checks and additional analyses. Section 7 concludes.

2 Institutional Background and Data

This section presents the institutional background of Chile’s employment protection policy and the administrative datasets we use in the analysis.

The employment protection policy implemented in Chile during the COVID-19 pandemic, called *Ley de Protección del Empleo* (Employment Protection Law), was active from April 2020 to October 2021. It allowed firms to reduce hours of some or all workers for up to six months. Most enrolled workers were furloughed at zero hours.³ This allowed them to stay at home during some months of the pandemic to protect both their health and work. Importantly, the policy explicitly prohibited remote work during the furlough period.

Firms were free to choose which workers to enroll and were not prevented from subsequently laying them off. This highly flexible scheme was similar to those deployed in many countries during the pandemic, especially in Europe. Examples include Germany, Italy, the U.K., and France.⁴

Worker benefits under the policy were financed through Chile’s unemployment insurance system, which was temporarily strengthened with public funds. Replacement rates started at 70% of the

²A complementary macro-labor literature emphasizes that recessions can also accelerate job polarization and routine-biased technological change (see Jaimovich and Siu, 2020; Hershbein and Kahn, 2018).

³A report on the use of the policy by Chile’s *Dirección del Trabajo* shows that, by December 2020, 93% of workers were furloughed at zero hours.

⁴See Giupponi et al. (2022) for a survey of short-time work schemes during the pandemic.

worker’s previous wage and declined to 30% by the sixth month. Firms had to pay pension and health contributions of enrolled workers, which amount to roughly 20% of their salary.

For workers to be eligible, they had to be affiliated with the unemployment insurance system, having contributed for at least five months, three of which had to be in continuous months with a given employer. These requirements ensured that workers with a minimum attachment to the formal labor market could access the program, reducing the risk of opportunistic enrollment. It is useful to note that the rate of informality in Chile, at around 25% of workers, is low for an emerging country, so most workers in the economy were eligible.

We use monthly firm-level data by merging four anonymized administrative datasets covering the universe of formal firms in Chile.⁵ The first source of information is a firm production dataset, reported for tax purposes, with information on sales, expenditures on intermediate goods, and investment in machinery and equipment. The source is Form F29 collected by the Chilean tax authority (Servicio de Impuestos Internos, SII), which covers the universe of formal firms in Chile.⁶

The second source of information is a matched employer-employee dataset, which contains firm-level information on all formal labor contracts and invoices associated with professional services. These data are collected from sworn statements DJ1887 and DJ1879 that firms file with the tax authority. Although these forms are filed annually, they include information on the months of each employment relation, which allows us to create a dataset at the monthly frequency. Earnings in these datasets are not capped.

The third source is a firm-bank credit registry that contains information on credit transactions at the bank-firm-loan level, including the amount of the loan, interest rate, and indicators on firm compliance with the credit contract, such as the number of days in arrears. This credit registry also contains firm-level information on debt stocks consolidated at the banking sector. The dataset covers the universe of financial transactions between banks and firms in Chile. The sources are forms D32 (flows) and C11 (stocks) collected by the Financial Regulatory Commission in Chile (Comisión para el Mercado Financiero, CMF).

The fourth and final source of information is the employment protection policy dataset, which contains firm-level information on firms’ access to the policy. The dataset contains information on the workers that were enrolled. The source of these data is the pension regulator (Superintendencia de Pensiones).

We keep data cleaning to a minimum, requiring firms to have positive cumulative real sales in the 12-month period up to two months prior to policy access. As explained in Section 3, we restrict

⁵This paper was developed within the scope of the research agenda conducted by the Central Bank of Chile (CBC) in economic and financial affairs of its competence. The CBC has access to anonymized information from various public and private entities, by virtue of collaboration agreements signed with these institutions. To secure the privacy of households and firms, the CBC mandates that the development, extraction and publication of the results should not allow the identification, directly or indirectly, of individuals or firms. The analysis was implemented by the authors and did not involve nor compromise the institutions that provide the data (Servicio de Impuestos Internos, Comisión para el Mercado Financiero, and Superintendencia de Pensiones).

⁶The information contained in the databases of the Servicio de Impuestos Internos (SII) is of a tax nature originating in self-declarations of taxpayers presented to the Service; therefore, the veracity of the data is not the responsibility of SII.

our attention to firms that first accessed the policy in May–June 2020, months of early and high policy take-up. We present descriptive statistics for the matched sample and policy take-up after introducing the empirical design.

3 Empirical strategy

Our empirical strategy consists of two stages. First, we use a matching procedure to construct a sample of treated and control firms that are comparable along real and financial dimensions prior to treatment. Second, we implement an event-study difference-in-differences (DiD) design to estimate the dynamic effects of the treatment of the policy.

3.1 Stage 1: Matching

We define treatment at the firm level as having at least one worker enrolled in the employment protection policy. A firm’s treatment month is the first month in which any of its workers enters the program. We restrict attention to firms that first access the policy in May–June 2020 and construct a control group of firms that did not access the policy, matching treated and control firms using pre-treatment firm characteristics described below.

1. Annual sales level: We use the log of cumulative real sales in the 12 months leading up to two months before treatment (months -13 to -2) and group firms into quintiles.
2. Year-on-year sales growth before treatment (March–April 2020): To account for heterogeneous pandemic shocks, we compute the growth rate of average real sales in March–April 2020 relative to March–April 2019 and group firms into quintiles.
3. Industry: Firms are classified into one of twelve industry categories based on their industry two months prior to treatment.
4. Location: Firms are grouped into 348 geographic units based on their registered location two months prior to treatment. This controls for differences in local lockdown measures and other health policies during the pandemic that varied geographically.
5. Stock of domestic bank debt: Measured two months prior to treatment, to capture pre-pandemic credit access.
6. Days in arrears: Categorized as 0, 1–29, 30–89, and 90+ days, following international conventions, to complement the previous variable and capture credit risk.
7. Deteriorated debt (dummy): Equals 1 if the firm had any deteriorated debt two months prior to treatment, and 0 otherwise.

Missing values in any variable are treated as a separate category. The matching procedure uses Coarsened Exact Matching (CEM) (Iacus et al., 2012), which improves covariate balance by

coarsening variables into substantively meaningful bins and discarding observations outside common support. CEM weights and strata are used in subsequent regressions to estimate the average treatment effect on the treated (ATT).

We restrict the treatment cohort to firms that first accessed the policy in May–June 2020—months of early and high take-up—and trace outcomes over a 24-month horizon. Focusing on this narrow, early window serves three purposes. First, it defines a single, compact treatment cohort, which simplifies the event-time alignment and avoids well-known complications of staggered adoption designs (e.g., cohort-composition and weighting issues). Second, the policy environment in these months is relatively clean: besides the employment protection policy (our focus) and the contemporaneous credit-guarantee program (FOGAPE-COVID), which we explicitly control for in the event-study specification, major additional support programs phase in only from late-2020 onward. Limiting treatment timing to May–June therefore reduces contamination from later policy interventions and facilitates interpretation of dynamic effects. Third, take-up was both early and sizable in this window, ensuring sufficient power while preserving a long post-treatment horizon for dynamics and persistence.

While assignment to treatment is not random, the richness of the data and the matching procedure alleviate concerns about selection on unobservables. Our event-study estimates allow us to verify that treated and control firms exhibit parallel pre-treatment trends, which supports the validity of the difference-in-differences design.

3.2 Stage 2: Event-study difference-in-differences

We estimate dynamic treatment effects using a conventional event-study difference-in-differences (DiD) specification that follows treated and matched control firms over event time. Let $\tau_{it} \equiv t - t_i^*$ denote event time, i.e., the number of months between the current month t and firm i 's (first) treatment month t_i^* . For controls, we assign the treatment month of their matched treated firm so that comparisons at each τ are made within the same calendar month. Our estimating equation is:

$$Y_{it} = \sum_{\tau=-T_0}^{-2} \alpha_{\tau} \cdot \mathbf{1}[\tau_{it} = \tau] \cdot D_i + \sum_{\tau=0}^{T_f} \beta_{\tau} \cdot \mathbf{1}[\tau_{it} = \tau] \cdot D_i + \gamma_{sd} + \delta_{\tau} + \lambda \text{OtherPolicy}_i + \varepsilon_{it}, \quad (1)$$

where Y_{it} is the outcome for firm i in month t and D_i is an indicator that firm i ever receives treatment in our window. The indicators $\mathbf{1}[\tau_{it} = \tau]$ equal 1 if observation (i, t) falls τ months away from t_i^* , and zero otherwise, partitioning time into event months relative to treatment.

We include two sets of fixed effects. First, strata-month fixed effects, γ_{sd} , which absorb time-varying shocks within each CEM stratum s by calendar date t . This ensures that identification exploits only within-stratum contrasts in each calendar month, guarding against composition changes across industries, locations, size, and pre-pandemic financial conditions used in matching. It also absorbs aggregate shocks that affect all firms in the same month, regardless of treatment status or event time. Second, event-time fixed effects, δ_{τ} , capture patterns that are common across firms at

the horizon of treatment τ . Including δ_τ prevents such event-time-aligned movements from being attributed to treatment.⁷

The term OtherPolicy_i is a firm-specific, time-invariant dummy equal to one if the firm accessed the contemporaneous credit-guarantee program prior to or no later than the firm’s relevant treatment date.⁸ This dummy variable controls for selection into the other policy to mitigate the possibility of confounding the effect of the employment protection policy.

We omit $\tau = -1$ from (1), so that all pre-treatment coefficients α_τ and post-treatment coefficients β_τ are measured relative to month -1 . The pre-treatment coefficients $\{\alpha_\tau\}_{\tau < -1}$ provide a visual and statistical test of parallel trends within matched strata. The post-treatment coefficients $\{\beta_\tau\}_{\tau \geq 0}$ trace the dynamic average treatment effect on the treated (ATT) at each event month. Standard errors are clustered at the firm level. We report event-study graphs with 95% confidence intervals.

4 Firm access to the employment protection policy

This section documents the scope of take-up and the composition of firms that accessed the employment protection policy. Table 1 reports descriptive statistics for treated and control firms in the *matched* sample, measured in the pre-treatment period.⁹ These comparisons quantify baseline differences in economic size, sales dynamics around the pandemic onset, financial conditions, and productivity levels before treatment.

Our matching procedure identifies 25,111 firms for the treated group and 471,799 for the control group. Treated firms are substantially larger than controls, with average sales nearly six times higher, suggesting that larger firms were more likely to access the policy.¹⁰ The impact of the pandemic on sales, however, is similar between treated and control firms. The (symmetric) growth rate of sales in March-April 2020 with respect to March-April 2019 is, on average, -0.24 for treated firms and -0.31 for control firms, and the difference is smaller between the median treated and control firms—a decline of 0.25 and 0.23, respectively.¹¹ Treated firms have more than twice as

⁷Calendar time (t) and event time (τ) index different dimensions: t moves all firms together through the pandemic and macro environment, whereas τ aligns firms by their relative position around treatment, which differs across firms. Having both γ_{sd} and δ_τ lets us partial out shocks tied to the wall-clock (e.g., aggregate lockdowns) and shocks tied to the treatment clock (e.g., generic pre- or post-treatment adjustments) without conflating either with the treatment effect.

⁸The credit-guarantee policy, called FOGAPE-COVID, was implemented from May 2020 to April 2021, and offered government guarantees of up to 85% on loans to firms with annual sales of less than USD 40 million, approximately. These are considered mostly small and medium-sized firms in Chile.

⁹The statistics presented in Table 1 are computed two months prior to treatment, except sales and sales growth. Statistics for sales and sales growth are computed on the (pre-treatment) objects used in the matching procedure, as explained in Section 3.1.

¹⁰The difference is approximately 500% when comparing means and 440% when comparing medians. Real sales are expressed as the logarithm of the sum over the 12 months up to two months prior to treatment, i.e., between periods $t = -13$ and $t = -2$. Control firms inherit their matched treated firm’s treatment date as reference month, as mentioned in section 3. The percent difference, for the case of means, for example, is approximately equal to $(e^{(8.3639 - 6.5794)} - 1) * 100 = 495.7$.

¹¹The variable sales growth is used to gauge the impact of the pandemic on sales. It is computed as $(g_1 - g_0)/[1/2 \times (g_1 + g_0)]$, so it is bounded between -2 and 2. For small values, it is approximately equal to standard percent changes. For example, a symmetric rate of -0.24 is equivalent to -21.4%.

much domestic bank debt as control firms—110% more debt on average and 122% when comparing median firms. Treated firms have, on average, a lower number of days in arrears on their stock of bank debt—8 days relative to 14, but the median treated and control firms have zero days in arrears.¹²

The level and growth rate of sales, as well as indicators on domestic bank debt, are part of our criteria for the matching procedure. We now turn to firm characteristics not used in the matching procedure to assess baseline differences. The first are labor productivity and TFP. Treated firms have lower labor productivity than controls: -4.4% on average and -7.9% for the median firm. Treated firms also have lower TFP: -12.3% on average and -7.6% for the median firm. Regarding the number of workers, treated firms are also larger than control firms when considering this variable, which is consistent with the difference in sales. On average, treated firms have about 2.5 times more workers (15.9 versus 6.1), while the median treated firm has 4 workers, compared to 1 worker for the median control firm. Finally, treated firms have slightly lower average worker ability than controls, suggesting that selection into the policy was not driven by workforce quality. Average worker ability is 2.2% lower when averaging across firms and 2.9% lower when comparing average worker ability between the median treated and control firms.¹³

Taken together, these patterns suggest that firms accessing the employment protection policy were larger, more leveraged, and slightly less productive than controls, but not obviously different in terms of workforce quality. It is important to note that Table 1 reports unconditional means and medians for treated and control firms. Consequently, differences in these aggregate statistics do not correspond to the variation used for identification. Our empirical strategy compares treated and control firms within CEM strata, conditioning on the pre-treatment characteristics used in the matching procedure. The purpose of Table 1 is therefore descriptive: it documents selection into the policy and motivates the use of matching, rather than providing a balance test for the estimating sample.

We now describe firm-level access to the employment protection policy, which is crucial for understanding the scope of the intervention. Firm-level take-up of the policy was widespread and substantial. Figure 1 shows that about half of firms, roughly 140,000, enrolled at least one worker in the employment protection program (panel a).¹⁴ Most of these firms began using the policy in the first few months of the pandemic—from March to June 2020.

Moreover, firms that accessed the policy enrolled a substantial fraction of their workers. The blue line in panel b of figure 1 shows that treated firms enrolled, on average, up to 60% of their workforce at the time of first access. This share declines sharply six months later, consistent with the policy’s time limit on enrollment duration. The red line in panel b of figure 1 shows that, by the sixth month after treatment, 60% of treated firms had stopped using the policy. This pattern

¹²Statistics on debt stocks are reported conditional on having positive debt, so they summarize the intensive margin among indebted firms. Statistics on days in arrears include firms with no debt, capturing both incidence and intensity of delinquency in the full sample at $t = -2$.

¹³Worker ability is proxied as the worker fixed effect in AKM-style regressions, as explained in section 5.2 below.

¹⁴The share of firms that access the policy is computed with respect to the number of firms with at least one employee in February 2020.


Table 1: Descriptive statistics of matched treated and control firms

Variables	Treated		Control	
	Mean	Median	Mean	Median
Sales (log)	8.3639	8.3149	6.5794	6.6351
Sales growth (symmetric rate)	-0.2378	-0.2466	-0.3134	-0.2339
Stock of debt (log)	5.7413	6.2457	4.9948	5.4502
Days in arrears	7.7954	0	13.7739	0
Labor productivity (log)	3.3825	3.4802	3.4276	3.5623
TFP (log)	1.7750	1.7367	1.9070	1.8158
Number of workers	15.9713	4	6.0789	1
Ability	0.0009	-0.0024	0.0236	0.0271
Number of firms	25,111		471,799	

Note.— This table reports descriptive statistics for the matched sample defined in Section 3. Treated firms first accessed the employment protection policy in May–June 2020; control firms did not access the policy and are included through the matching procedure. Means and medians are reported unconditionally within the matched sample. Identification in subsequent regressions comes from comparisons within CEM strata using the matching weights described in Section 3. All variables are measured in the pre-treatment window, period $t = -2$ for most cases. Sales are real and expressed in logs of cumulative 12-month values up to t_i^* for treated firms, and the corresponding calendar months for matched controls. Sales growth describes the symmetric growth rate of average real sales, bounded between -2 and 2, between March–April 2019 and March–April 2020. Debt variables refer to domestic bank debt two months prior to treatment. Labor productivity is value added per worker; TFP is the residual from a Cobb–Douglas specification with fixed factor shares. The number of workers reflects formal employment relationships. See Section 3 for definitions, and Section 5 for outcome construction and interpretation.

suggests that firms concentrated their use of the policy in the immediate aftermath of the pandemic shock.

Figure 1: Use of the policy by firms



Note.— (a) The share of firms that access the policy is computed with respect to the number of firms with at least one employee in February 2020. (b) The blue line shows the average share of workers enrolled in the policy among the treated firms identified in the matching stage. We consider treatment ($t = 0$) in May and June 2020. This share is computed with respect to the number of workers in $t = -1$. The red line shows the share of treated firms that continue to have at least one worker enrolled, with respect to the number of treated firms in $t = 0$.

5 Results

This section presents the firm-level effects of the employment protection policy on productivity and other outcomes, and analyzes ability-based selection of workers as the mechanism behind the productivity gains.

5.1 Effects on productivity and other firm-level outcomes

Figure 2 shows that treated firms display higher productivity than controls, as proxied by two measures—labor productivity, defined as value added divided by employment, and TFP, measured as the residual from a Cobb–Douglas specification. In the baseline specification, intermediate inputs, labor, and capital are assigned constant shares of 0.5, 0.25, and 0.25, respectively, which approximate average factor cost shares in the data.¹⁵


Both measures of productivity are persistently higher in treated firms than in controls. A statistically significant effect of 5–10% higher productivity is evident throughout most of a 24-month horizon after treatment. Note that in our measures of productivity, employment excludes workers while they are enrolled in the employment protection policy. This is why labor productivity increases sharply during the first six months after treatment—the denominator (employment) declines more than the numerator (value added) due to the immediate and intense take-up of the policy, as documented in the previous section. After this initial period of high take-up, the two measures are similar, with sustained gains that reflect improvements in firm efficiency well beyond the first six months. Importantly, the coefficients in the event study are flat for pre-treatment periods, which demonstrates that productivity in treated and control firms is similar before treatment.

To assess broader firm performance, we analyze the effects of the employment protection policy on sales, labor, and capital inputs. Figure 3 shows a positive long-run effect on sales of treated firms relative to controls. This effect stabilizes around 10% roughly 15 months after treatment (panel a). During the first five months, sales of treated firms are lower than controls. This is consistent with the substantial reduction in the labor input of these firms due to the enrollment of a large fraction of workers in the policy during this period, as we previously documented. Five months after treatment, however, treated firms begin to sell more than controls.

Higher sales are driven by higher productivity, rather than higher use of labor or capital inputs. Panel b of figure 3 shows that, by the end of the 24-month horizon, there are no differences in the number of workers across treated and control firms, whereas capital is actually lower in treated firms. Regarding the number of workers, we plot two outcomes that differ in the way we classify those enrolled in the employment protection policy. The first, which we call contracts (blue dots), includes workers enrolled in the policy, whereas the second, which we call employment (red dots), excludes workers while they are enrolled. Contracts can decline due to separations, whereas employment can also decline due to enrollment in the policy. Employment is, thus, a measure of labor that is available for production in each period. These two labor measures differ during the first six months

¹⁵Using sector-specific cost shares produces similar patterns.

Figure 2: Firm-level productivity



Note.— Event study analysis comparing the productivity of firms treated in month $t = 0$ to controls. Differences in period $t = -1$ are normalized to zero. Productivity is proxied as labor productivity, defined as value added (sales minus intermediate input expenses) divided by employment, and TFP, defined as the residual from a Cobb-Douglas specification in which intermediate input expenses, employment and capital have constant shares of 0.5, 0.25 and 0.25, respectively. In both productivity measures, employment excludes workers while they are enrolled in the employment protection policy.

of high policy take-up, when contracts decline nearly 10% and employment declines up to 17%. In the following months, however, both measures are very similar and return to zero by the end of the horizon.¹⁶


Capital in treated firms declines gradually relative to controls; by the end of the horizon, it is nearly 10% lower. The employment protection policy primarily stabilized and reorganized employment relationships—through subsidized suspension and selective recall—while leaving capital accumulation largely outside the scope of the intervention, as the policy did not subsidize investment.¹⁷ Firms could therefore raise efficiency by upgrading their workforce composition and reorganizing production, and subsequently operate with *less* installed capital. Organizational streamlining following selective recall can reduce the amount of capital required for a given scale—for example, by retiring marginal or underutilized assets. In sum, the increase in the sales of treated firms relative to controls is not explained by higher labor or capital, but by higher productivity.¹⁸

¹⁶The employer–employee dataset DJ1887 is based on firm self-reports and can miss some workers who are enrolled in the policy. We therefore include such workers in contracts, since the employment relationship has not been terminated.

¹⁷In parallel, our controls for access to the credit-guarantee program mitigate the concern that capital declines reflect simple liquidity differences.

¹⁸In principle, sales could also increase through purchases of intermediate goods, but the labor productivity measure in figure 2 rules this possibility out, since it shows an increase in value added, defined as sales minus intermediate input expenses, divided by employment.

Figure 3: Firm-level sales and labor and capital inputs



Note.— Event study analysis comparing outcomes of firms treated in month $t = 0$ to controls. Differences in period $t = -1$ are normalized to zero. Contracts include workers that are enrolled in the employment protection policy; employment excludes them while they are enrolled.

5.2 Mechanism: Selection of high-ability workers

To understand the drivers of the observed productivity gains in treated firms, we analyze the role of worker selection as a potential mechanism. The employment protection policy allowed firms to temporarily suspend labor relations. Treated firms used this subsidized pause to selectively retain and recall higher-ability workers, while allowing lower-ability workers to separate. Worker ability is proxied by the estimated worker fixed effects from Abowd-Kramarz-Margolis (AKM) regressions of log wages on worker and firm fixed effects, controlling for age.¹⁹

Table 2 describes the ability of different groups of workers within treated firms. It shows results from a regression of worker ability on dummy variables for different worker categories. The reference group, for which we omit a dummy variable, is stayers, who are employed at the firm, and not enrolled in the employment protection policy, at least between periods $t = -2$ and $t = 12$. The groups of workers included in the regression are: a) those enrolled in the policy and subsequently recalled, b) those enrolled and not recalled, c) those that separated during periods $t = 0$ and $t = 1$ without going through the employment protection policy, d) those that separated beyond period $t = 1$ without going through the policy, and e) new workers, i.e., those that join the firm between

¹⁹See Abowd et al. (1999). The AKM regressions are estimated at the worker-year level on a sample from 2004 to 2023, and are given by the following equation: $\log(w_{it}) = \alpha_i + \psi_{J(i,t)} + \theta_t + \phi_1 \cdot \text{age}_{it}^2 + \phi_2 \cdot \text{age}_{it}^3 + \phi_3 \cdot (\text{age}_{it}^2 \times \theta_t) + \phi_4 \cdot (\text{age}_{it}^3 \times \theta_t) + \epsilon_{it}$, where α_i is a worker fixed effect—our proxy of ability, the $\{\psi_j\}_{j=1}^J$ are firm-specific relative pay premiums, $J(i, t)$ is a function indicating the employer of worker i in year t , θ_t is a time fixed effect, the age terms and their interactions with year capture non-linear experience profiles, and ϵ_{it} is an error term. The regression is weighted by the number of months worked in the year. Wage data come from sworn statement DJ1887, which registers formal employment contracts. We also include income reported in sworn statement DJ1879 (occasional services) only when the payer corresponds to the same employer listed in DJ1887, ensuring consistency in employer-employee links.

periods $t = -2$ and $t = 12$.^{20,21} The regression also includes industry and city fixed effects.

Table 2: Worker ability in treated firms across worker categories

Dependent variable: Ability. Reference group: Stayers.	
(a) Enrolled and recalled	-0.25*** (0.0027)
(b) Enrolled but not recalled	-0.45*** (0.0041)
(c) Separated ($t = 0, 1$)	-0.52*** (0.0050)
(d) Separated ($t > 1$)	-0.33*** (0.0033)
(e) New	-0.60*** (0.0068)
(f) Constant	0.29*** (0.0018)
Industry and city FE	✓
Observations	345,133
R ²	0.15

Note.— Cross-sectional regression, at the worker level, of ability on dummy variables indicating worker categories that takes the following form: $\text{Ability}_i = \beta_0 + \sum_g \beta_g \text{Group}_{ig} + \gamma_s + \delta_c + \epsilon_i$, where Group_{ig} are dummies for worker categories, and γ_s and δ_c are industry and city fixed effects. The reference group, for which we omit a dummy variable, is stayers, defined as workers who are employed at the firm, and not enrolled in the employment protection policy, at least between periods $t = -2$ and $t = 12$. Category a (enrolled and recalled) includes workers that were enrolled in the policy and recalled later, even if some of them subsequently separated. Category b (enrolled but not recalled) includes workers that were enrolled in the policy and did not return to the firm. Category c (separated in $t = 0, 1$) includes workers that separate during periods $t = 0$ and $t = 1$ without going through the policy. Category d (separated in $t > 1$) includes workers that separate in periods beyond $t = 1$, without going through the policy. Category e (new) includes workers who join the firm between periods $t = -2$ and $t = 12$ and had not worked at the firm since 2018. Standard errors in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% levels.

Treated firms show a clear pattern of worker selection based on ability. Stayers—defined as workers continuously employed and not enrolled in the policy—have higher estimated ability than all other groups. Workers enrolled in the employment protection policy (rows a and b) have lower ability than stayers, but higher ability than those who separated during the first two months after treatment as part of firms’ immediate crisis response (row c). Among enrolled workers, those who were subsequently recalled (row a) display higher ability than those who were not (row b). Table 2 shows that these differences are both economically and statistically significant. Stayers have AKM

²⁰Specifically, the regression takes the following form: $\text{Ability}_i = \beta_0 + \sum_g \beta_g \text{Group}_{ig} + \gamma_s + \delta_c + \epsilon_i$, where Group_{ig} are dummies for the following worker categories: a) workers enrolled and subsequently recalled, b) workers enrolled and not recalled, c) workers that separated in $t = 0$ or $t = 1$ without going through the policy, d) workers that separated in $t > 1$ without going through the policy, and e) new workers. γ_s and δ_c are industry and city fixed effects.

²¹More precisely, a worker is classified as new if she had not worked at the firm since 2018 and began a relationship between periods $t = 0$ and $t = 12$.

fixed effects that are 0.25 to 0.45 log points higher than those of enrolled workers (approximately 28% to 57%). Workers who separated in periods $t = 0$ and $t = 1$ have AKM fixed effects that are 0.52 log points lower than stayers, and 0.07 to 0.27 log points lower than enrolled workers. Among enrolled workers, those who were recalled have 0.20 log points higher ability than those who were not recalled. For completeness, we also include workers who separate after $t = 1$ as a distinct category; although they are not part of the immediate crisis response, their estimated ability lies between that of recalled workers and the early separators. New workers (row e) display the lowest ability, with a gap of 0.61 log points relative to stayers. The pattern of new workers displaying the lowest ability is typical of firms even in normal times, as we document in Section 6.2, and likely reflects lower entry wages and limited tenure.²²


This selection pattern based on worker ability indicates that the employment protection policy allowed firms to selectively retain higher-ability workers, thereby enhancing post-crisis productivity. A key feature of the policy is that it temporarily suspended labor obligations while preserving existing employment relationships and the option of recall. This arrangement may have created scope for firms to reassess workforce composition at relatively low cost before making more permanent staffing decisions. There are two possible mechanisms through which this process may have operated. First, by enrolling a worker at zero hours, firms may have improved their assessment of the worker’s ability through additional information revealed during the furlough period. Second, firms may have reassessed the worker’s contribution to production in the context of organizational changes induced by the crisis, even if individual ability was already known. In this interpretation, furlough effectively allowed firms to experiment with alternative workforce configurations and subsequently recall the workers that proved most valuable in the reorganized production process. Distinguishing these channels would require exogenous variation in information about worker ability or worker-level performance metrics (e.g., task output or evaluations), neither of which is available in our administrative data. Learning about ability and organizational reassessment are observationally similar in timing and implications for recall decisions. Both mechanisms are therefore consistent with the observed ability-based selection and the subsequent productivity gains.

The ability-based selection pattern documented in table 2 leads to higher ability in treated firms relative to controls. Figure 4 shows that treated firms display persistently higher average worker ability than controls, with a statistically significant gap of approximately 2% that stabilizes after 12 months. This result holds whether or not enrolled workers are included. During the first six months, however, when the policy is more intensively used, average ability in treated firms is much higher when excluding enrolled workers. This is a compositional effect that reflects the fact that workers enrolled in the policy had lower ability than stayers, as documented in table 2.

Table 3 presents further evidence on ability-based selection by comparing worker ability across treated and control firms. Since only treated firms enroll workers in the employment protection policy, we reclassify recalled workers as stayers and non-recalled workers as separated, so that

²²We explore recall dynamics in Section 6.3 and show that a substantial share of enrolled workers were not recalled, reinforcing the interpretation of the policy as a tool for selective retention.

Figure 4: Average ability



Note.— Event study analysis comparing the average ability of workers in firms that are treated in month $t = 0$ to controls. Differences in period $t = -1$ are normalized to zero. Ability is proxied by worker fixed effects in AKM-style regressions (Abowd et al., 1999).

worker categories are comparable across the two groups of firms.²³ Column (2), our preferred specification, includes industry and municipality fixed effects.

Two key results emerge from the interaction terms, which capture how the policy altered selection patterns relative to control firms. First, the interaction between separation and treatment status is positive and statistically significant (row c: 0.07 log points). This indicates that the ability gap between stayers and separated workers is *smaller* in treated firms than in controls. In economic terms, the policy raised the ability threshold below which firms separate workers: treated firms, enabled by the flexibility of subsidized furlough and selective recall, shed workers from a higher segment of the within-firm ability distribution. In contrast, control firms, facing standard labor market frictions, appear to retain lower-ability workers who would otherwise be closer to the separation margin—consistent with inefficient labor hoarding in the absence of the policy.

Second, the interaction between new-hire status and treatment is also positive and significant (row e: 0.06 log points), indicating that new workers hired by treated firms are drawn from a higher segment of the ability distribution, relative to their own stayers, than new hires in control firms. This reinforces the interpretation that the policy enabled a more comprehensive upgrading of workforce composition—not only through selective separation but also through higher-quality

²³As in the within-treated-firm regression discussed previously, the reference group is stayers. The regression equation is the following: $\text{Ability}_i = \beta_0 + \beta_1 \text{Separated}_i + \beta_2 \text{TreatedFirm}_i + \beta_3 (\text{Separated}_i \times \text{TreatedFirm}_i) + \beta_4 \text{New}_i + \beta_5 (\text{New}_i \times \text{TreatedFirm}_i) + \gamma_s + \delta_m + \varepsilon_i$, where γ_s and δ_m are industry and municipality fixed effects.

hiring.²⁴

The main effect of treatment status (row b) shows that stayers in treated firms have somewhat lower ability than stayers in control firms (-0.10 log points). This reflects two forces. First, treated firms had slightly lower average worker ability at baseline, as documented in Table 1. Second, in treated firms the stayer category pools workers who were never enrolled with those who were enrolled and subsequently recalled; since recalled workers have lower ability than never-enrolled stayers (Table 2, row a), this compositional effect mechanically reduces the average ability of the stayer group. Importantly, this level difference does not affect the interpretation of the interaction terms, which capture differential selection patterns induced by the policy. Moreover, Figure 4 shows that average worker ability in treated firms *increases* relative to controls over time—consistent with the ability-upgrading mechanism documented throughout this section.

Table 3: Worker ability in treated firms relative to controls

Dependent variable: Ability. Reference group: Stayers.		
	(1)	(2)
(a) Separated ¹	-0.40*** (0.0013)	-0.36*** (0.0014)
(b) Firm treated	-0.13*** (0.0015)	-0.10*** (0.0016)
(c) Separated $\times$ treated	0.10*** (0.0027)	0.07*** (0.0027)
(d) New	-0.62*** (0.0029)	-0.55*** (0.0029)
(e) New $\times$ treated	0.09*** (0.0072)	0.06*** (0.0071)
(f) Constant	0.31*** (0.0008)	0.28*** (0.0007)
Industry and municipality FE	–	✓
Observations	1,647,370	1,647,370
R ²	0.09	0.17

Note.— Cross-sectional regressions at the worker level of ability on dummy variables for worker categories interacted with treatment status. The sample includes workers in both treated and control firms. Stayers are the reference group. In treated firms, stayers include workers who were enrolled in the policy and subsequently recalled, even if they eventually separated. Separated workers in treated firms include those enrolled in the policy but not recalled. Column (1) includes no fixed effects beyond the intercept; column (2) adds industry and municipality fixed effects. Standard errors in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% levels.

²⁴Using column (2) estimates, the stayer–separated ability gap in control firms is $|\beta_1| = 0.36$ log points, whereas in treated firms it is $|\beta_1 + \beta_3| = 0.29$ log points. Similarly, the stayer–new-hire gap is 0.55 in controls and 0.49 in treated firms. Both gaps are significantly smaller in treated firms, consistent with the policy raising selection thresholds for both separations and new hires.

While the documented pattern of ability-based worker selection is economically and statistically significant, we do not interpret it as the sole explanation for the observed productivity gains. Productivity improvements following the policy likely reflect a combination of ability upgrading, organizational restructuring, and other within-firm adjustments that are difficult to disentangle empirically.

6 Robustness and additional results

This section offers additional results on four dimensions. First, we characterize the heterogeneity in the productivity effects of the employment protection policy across firms of different size and industry. Second, we provide further evidence on worker ability in treated firms. Third, we describe the extent of worker recall from the policy. Finally, we discuss the behavior of labor income.

6.1 Productivity effects by firm size and industry

To analyze the heterogeneity in our results, figures 5 and 6 show the evolution of total factor productivity (TFP) by firm size and industry, respectively. Results for labor productivity follow similar patterns and are available upon request.

Regarding firm size, figure 5 compares the aggregate TFP effect of the policy, which we discussed previously, with the effect across firms of four sizes—micro, small, medium, and large, classified by sales.²⁵ Micro, small and medium-sized firms display positive and significant TFP gains from the policy, with micro firms showing an effect about twice as large as the aggregate. Large firms do not display significant TFP effects. The stronger effects among micro firms may reflect greater marginal returns to labor optimization or higher sensitivity to labor cost subsidies that make the employment protection policy particularly valuable. The lack of productivity gains among large firms may reflect lower marginal benefits from labor cost subsidies or less flexibility in workforce restructuring.²⁶


Regarding heterogeneity across industries, figure 6 compares the aggregate TFP effect of the policy with the effect across four selected industries—manufacturing, commerce, transport, and personal services.²⁷ Firms in manufacturing, commerce and transport display economically and statistically significant productivity gains, often higher than the aggregate, though wider standard errors indicate these effects are less precisely estimated. The personal services sector shows no signs of productivity gains. Similarly, there is no evidence of productivity gains in mining, construction, and financial services, industries which we do not plot for brevity. Taken together, these

²⁵The classification of firms by size considers annual real sales in year 2019 and is held fixed throughout the analysis to avoid endogenous reclassification. The categories follow National Accounts practices, with thresholds based on sales expressed in a unit of account indexed to inflation, called *unidad de fomento* (UF). Micro firms: less than 2,400 UF. Small firms: 2,400–25,000 UF. Medium firms: 25,000–100,000 UF. Large firms: more than 100,000 UF. Based on inflation and market exchange rates at the time of writing, 25,000 UF are approximately equivalent to 1 million USD.

²⁶Recall that our administrative data refers to the universe of *formal* firms.

²⁷Firms are classified across industries based on pre-treatment characteristics in year 2019. This classification is held fixed throughout the analysis to avoid endogenous reclassification.

Figure 5: Firm-level TFP effects by firm size



Note.— Event study analysis comparing the productivity of firms treated in month $t = 0$ to controls. Differences in period $t = -1$ are normalized to zero. TFP is measured as the residual from a Cobb-Douglas specification in which intermediate input expenses, employment and capital have constant shares of 0.5, 0.25 and 0.25, respectively. Employment excludes workers while they are enrolled in the policy. The classification of firms by size considers annual real sales in year 2019. The categories follow National Accounts practices, with thresholds based on sales expressed in a unit of account indexed to inflation, called *unidad de fomento* (UF). Micro firms: less than 2,400 UF. Small firms: 2,400–25,000 UF. Medium firms: 25,000–100,000 UF. Large firms: more than 100,000 UF. Based on inflation and market exchange rates at the time of writing, 25,000 UF are approximately equivalent to 1 million USD.


results suggest that the productivity gains from the policy are concentrated in industries such as manufacturing, commerce, and transport, where person-to-person contact is particularly intense.

6.2 Ability-based worker selection

We offer two pieces of additional evidence on the documented pattern of ability-based worker selection. First, in addition to the regression-based results, the selection pattern among treated firms is evident in raw comparisons of mean and median ability across worker groups. Second, new workers have the lowest ability also in normal times.

Section 5.2 showed that treated firms display a clear pattern of ability-based selection of workers. Specifically, a regression of ability on dummy variables for different worker categories that controls for industry and location (table 2) showed stayers have higher ability than all other groups, workers enrolled in the policy have higher ability than those who separated, and among those enrolled,

Figure 6: Firm-level TFP effects by firm industry



Note.— Event study analysis comparing the productivity of firms treated in month $t = 0$ to controls. Differences in period $t = -1$ are normalized to zero. TFP is measured as the residual from a Cobb-Douglas specification in which intermediate input expenses, employment and capital have constant shares of 0.5, 0.25 and 0.25, respectively. Employment excludes workers while they are enrolled in the policy. Firms are classified across industries based on pre-treatment characteristics in year 2019.

recalled workers have higher ability than those that were not recalled. Table 4 shows that the same selection pattern is evident in raw comparisons of mean and median ability across worker groups. As in the main regression, new workers display the lowest ability.

Table 5 shows that new workers also display the lowest ability during normal times. It computes mean and median ability for stayers, separated and new workers for 2018, a year that can be considered as normal for the Chilean economy.²⁸ For this exercise, May 2018 is considered $t = 0$, and stayers, separated and new workers are defined according to their status in periods $t = -2$ to $t = 12$.²⁹ We use a 14-month window around May 2018 to ensure comparability with the crisis-period analysis. During this normal period, new workers also have the lowest mean and median ability. This confirms that lower AKM fixed effects among new workers are a persistent feature of labor markets, likely reflecting entry-level wages and limited tenure. Consistent with the crisis-

²⁸We use 2018 as a benchmark year, given its stable macroeconomic conditions—real GDP growth of 4%, unemployment rate of 7.3%, close to the 2010-2024 average of 7.6%, and no major labor market disruptions.

²⁹Stayers are employed at the firm between periods $t = -2$ to $t = 12$, separated cease to be employed at a firm between $t = 0$ to $t = 12$ after being employed in $t = -2$ and $t = -1$, and new workers join the firm between $t = -2$ to $t = 12$.

period findings, stayers in 2018 also exhibit higher ability than separated workers.

Table 4: Worker ability in treated firms

	Mean	Median	SD	N° of Workers
Stayers	0.29	0.19	0.65	126,795
Short-time, recalled	0.04	0.00	0.57	101,794
Short-time, not recalled	-0.18	-0.18	0.59	28,499
Separated	-0.25	-0.25	0.67	20,742
New	-0.35	-0.33	0.75	13,172

Note.— Ability is proxied by worker fixed effects in AKM regressions of log wages on worker and firm fixed effects, controlling for age. Stayers are employed at the firm, and not enrolled in the policy, at least between periods $t = -2$ and $t = 12$. Short-time recalled workers were enrolled in the policy and recalled later, even if some of them subsequently separated. Short-time not recalled workers were enrolled in the policy and did not return to the firm. Separated workers were employed in periods $t = -2$ and $t = -1$, and separated in periods $t = 0$ or $t = 1$. New workers join the firm between $t = -2$ and $t = 12$.

Table 5: Worker ability in 2018

	Mean	Median	SD	N° of Workers
Stayers	0.235	0.144	0.621	194,527
Separated	-0.165	-0.178	0.660	118,829
New	-0.332	-0.314	0.737	32,989


Note.— Ability is proxied by worker fixed effects in AKM regressions. $t = 0$ is set in May 2018. Stayers are employed between periods $t = -2$ and $t = 12$, separated cease to be employed at a firm between $t = 0$ and $t = 12$ after being employed in $t = -2$ and $t = -1$, and new workers join the firm between $t = -2$ and $t = 12$.

6.3 Worker recall from the employment protection policy

To characterize the extent of worker recall from the policy, figure 7 shows the share of workers that were recalled at some point and are working at the firm in month t , with respect to the number of enrolled workers up to that month (panel a), and the share of workers that were recalled at some point, regardless of whether they are working at the firm in month t , also with respect to the number of enrolled workers up to that month (panel b). Panel a captures the fact that recalled workers may subsequently separate. Each panel displays two series, which differ in the way these average rates are computed. For the blue lines, we compute the recall rates in each firm, and then average across firms, in practice giving equal weight to all firms. For the red lines, the recall rates are computed using aggregates.

Panel b shows that by late-2021, more than 20% of workers enrolled in the employment protection policy were not recalled. This implies that one in five workers enrolled in the policy were permanently separated, consistent with firms using the policy to shed lower-ability workers.

Figure 7: Recall rate



Note.— (a) At each month t , workers that were enrolled in the policy and are back at the firm, as a share of the total number of enrolled workers up to t . (b) At each month t , workers that were enrolled in the policy and were recalled at some point, regardless of whether they are working at the firm in t , as a share of the total number of enrolled workers up to t . For the blue lines (mean firm ratio), we compute the recall rates in each firm and then average across firms. For the red lines (total ratio), the recall rates are computed using aggregates.

Panel a shows that about half of the workers that had gone through the policy were working at the firm by late-2021. The difference between panels a and b is due to the fact that some recalled workers subsequently separated. These separations may be due to firms firing workers, or to voluntary quits, potentially reflecting improved outside options during the furlough period.

6.4 Labor income

The final additional result is about labor income. Figure 8 shows the effect of the employment protection policy on worker earnings of treated firms relative to controls. We consider workers that are at the firm at least from January 2016 to December 2023. By focusing on stayers, the evolution of earnings is free from compositional effects.³⁰


In the two years following treatment, earnings increase about 2.5%. However, the observed post-treatment earnings gap mirrors pre-treatment differences, suggesting that the policy did not materially affect wage trajectories. Therefore, higher productivity does not seem to have resulted in higher worker earnings. This disconnect may reflect wage rigidity or limited rent-sharing in the Chilean labor market.

7 Conclusion

This paper provides causal evidence that employment-protection policies can simultaneously preserve jobs and enable productivity-enhancing reallocation when they relax contractual frictions in ways that allow selective furlough and recall. Leveraging Chile's short-time work scheme implemented during COVID-19 and rich administrative microdata, we show that treated firms exhibit

³⁰Monthly worker earnings are obtained from administrative records on unemployment insurance contributions. These data are compiled by the Pensions Authority (Superintendencia de Pensiones).

Figure 8: Average worker earnings



Note.— Event study analysis comparing average worker earnings of firms treated in month $t = 0$ to controls. Differences in period $t = -1$ are normalized to zero. We consider only stayers, i.e., those who work at their firm at least from January 2016 to December 2023, and are not enrolled in the policy, so the evolution of earnings is free from compositional effects.

sustained 5–10% productivity gains, higher sales without higher inputs, and ability upgrading via selective retention and recall of higher-ability workers and non-recall/separation of lower-ability workers. Consistent with this interpretation, the comparison with control firms suggests that treated firms separated workers from a higher segment of the ability distribution and hired higher-ability new workers.

The design elements that likely underpin these results include easy access, zero-hours furlough with no remote work, time-limited eligibility, and continued employer contributions of about 20% of wages that sharpen the incentive to enroll and recall selectively.

The evidence suggests that employment protection policies need not impede creative destruction; instead, they can mediate it by shifting selection from outright separations to policy-enabled recall decisions, improving match quality and TFP from within firms under designs that combine flexibility with employer cost-sharing.

Although our results provide strong evidence that treated firms engage in ability-based selection, we do not attempt to quantify the fraction of the productivity gains that can be directly attributed to changes in workforce composition. Productivity improvements following the policy likely reflect a combination of ability upgrading, organizational restructuring, and other within-firm adjustments that are difficult to disentangle empirically. We therefore view ability-based worker selection as a central mechanism enabled by the policy’s design, but not necessarily the sole driver of the observed

productivity gains.

References

- ABOWD, J. M., F. KRAMARZ, AND D. N. MARGOLIS (1999): “High Wage Workers and High Wage Firms,” *Econometrica*, 67, 251–333.
- ATES, S. T. AND F. E. SAFFIE (2021): “Fewer but Better: Sudden Stops, Firm Entry, and Financial Selection,” *American Economic Journal: Macroeconomics*, 13, 304–56.
- AUTOR, D. H., W. R. KERR, AND A. D. KUGLER (2007): “Does Employment Protection Reduce Productivity? Evidence From US States,” *Economic Journal*, 117, 189–217.
- BARRERO, J. M., N. BLOOM, AND S. J. DAVIS (2020): “COVID-19 Is Also a Reallocation Shock,” *Brookings Papers on Economic Activity*, 2020, 329–383.
- BERGER, D. (2018): “Countercyclical Restructuring and Jobless Recoveries,” Mimeo.
- BILAL, A., N. ENGBOM, S. MONGEY, AND G. L. VIOLANTE (2022): “Firm and Worker Dynamics in a Frictional Labor Market,” *Econometrica*, 90, 1425–1462.
- BRINKMANN, C., S. JÄGER, M. KUHN, F. SAIDI, AND S. WOLTER (2024): “Short-Time Work Extensions,” Discussion Paper 606, CRC TR 224 (Collaborative Research Center Transregio 224), University of Bonn and University of Mannheim.
- BUERA, F. J., R. N. FATTAL-JAEF, H. HOPENHAYN, P. A. NEUMEYER, AND Y. SHIN (2021): “The Economic Ripple Effects of COVID-19,” NBER Working Papers 28704, National Bureau of Economic Research, Inc.
- CABALLERO, R. J. AND M. L. HAMMOUR (1996): “On the Timing and Efficiency of Creative Destruction,” *The Quarterly Journal of Economics*, 111, 805–852.
- CAHUC, P., F. KRAMARZ, S. NEVOUX, AND N. VIEIRA (2024): “The Heterogeneous Impact of Short-Time Work: From Saved Jobs to Windfall Effects,” Manuscript dated August 16, 2024.
- COOPER, R., M. MEYER, AND I. SCHOTT (2017): “The Employment and Output Effects of Short-Time Work in Germany,” Working Paper 23688, National Bureau of Economic Research.
- CRAVO, T., P. D. A. JACINTO, AND R. QUINTANA (2020): “Can work sharing sustain employment during economic downturn? Evidence from Brazil,” Tech. rep.
- DEN BOSCH, J. V. AND S. VANORMELINGEN (2023): “Productivity growth over the business cycle: cleansing effects of recessions,” *Small Business Economics*, 60, 639–657.
- EICHHORST, W., P. MARX, U. RINNE, AND J. BRUNNER (2022): “Job Retention Schemes during COVID-19: A Review of Policy Responses,” IZA Policy Papers 187, Institute of Labor Economics (IZA).

- FAHN, M. AND T. MUROOKA (2025): “Informal Incentives and Labor Markets,” *The Economic Journal*, 135, 144–179.
- FORSYTHE, E., L. B. KAHN, F. LANGE, AND D. WICZER (2022): “Where have all the workers gone? Recalls, retirements, and reallocation in the COVID recovery,” *Labour Economics*, 78.
- GERTLER, M., C. HUCKFELDT, AND A. TRIGARI (2025): “Temporary Layoffs, Loss-of-Recall, and Cyclical Unemployment Dynamics,” *American Economic Review*, forthcoming. Manuscript dated April 13, 2025.
- GIUPPONI, G., C. LANDAIS, AND A. LAPEYRE (2022): “Should We Insure Workers or Jobs during Recessions?” *Journal of Economic Perspectives*, 36, 29–54.
- HALTIWANGER, J. (2022): “Spatial and Sectoral Reallocation of Firms, Workers and Jobs in the Pandemic and Recovery.” Mimeo.
- HERSHBEIN, B. AND L. B. KAHN (2018): “Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings,” *American Economic Review*, 108, 1737–1772.
- IACUS, S., G. KING, AND G. PORRO (2012): “Causal Inference without Balance Checking: Coarsened Exact Matching,” *Political Analysis*, 20, 1–24.
- JAIMOVICH, N. AND H. E. SIU (2020): “Job Polarization and Jobless Recoveries,” *Review of Economics and Statistics*, 102, 129–147.
- KATO, T. AND N. KODAMA (2019): “The Consequences of Short-Time Compensation: Evidence from Japan,” Discussion papers 19056, Research Institute of Economy, Trade and Industry (RIETI).
- KOZENIAUSKAS, N., P. MOREIRA, AND C. SANTOS (2022): “On the cleansing effect of recessions and government policy: Evidence from Covid-19,” *European Economic Review*, 144, None.
- KRUPPE, T. AND T. SCHOLZ (2014): “Labour hoarding in Germany : employment effects of short-time work during the crises,” IAB-Discussion Paper 201417.
- MALCOMSON, J. M. (1997): “Contracts, Hold-Up, and Labor Markets,” *Journal of Economic Literature*, 35, 1916–1957.
- MOLAVI, P. (2018): “A Theory of Dynamic Selection in the LaborMarket,” Mimeo.
- MORTENSEN, D. T. AND C. A. PISSARIDES (1994): “Job Creation and Job Destruction in the Theory of Unemployment,” *The Review of Economic Studies*, 61, 397–415.