Automation and Employment Recovery After the Great Recession*

Lalit Contractor† Ronit Mukherji‡ Abhishek Rai§

September 11, 2025

Abstract

We document a novel empirical finding: U.S. regions with higher pre-Great Recession exposure to automation experienced significantly faster employment recoveries after the recession. To explain this result, we develop a simple model of firm heterogeneity with capital accumulation. The model predicts that, following a transitory negative TFP shock, economies with higher automation intensity recover more quickly than those with lower automation intensity. In less automation-intensive regions, firms start with relatively lower levels of automation capital, and this gap is exacerbated during the shock. As a result, these regions benefit less from the complementarity between automation capital and labour, leading to a slower employment recovery.

Keywords: automation, hysteresis, firm heterogeneity, business cycles, Great Recession

JEL Codes: E23, E24, E32, J24

*For comments and suggestions, we thank Anand Chopra, Kunal Dasgupta, Aruni Mitra and participants at the 2025 SERI workshop, CAFRAL seminar and Macro Seminar Series at Ashoka University. This research uses proprietary data made available to Ronit Mukherji through The University of British Columbia, CIDER Grant. We received no other grant from public, commercial, or not-for-profit funding agencies. We alone are responsible for all errors and interpretations.

†Ashoka University. Email: lalit.contractor@ashoka.edu.in.

‡Ashoka University. Email: ronit.mukherji@ashoka.edu.in.

§Ashoka University. Email: abhishek.rai@ashoka.edu.in.

1 Introduction

Automation is widely recognised as a disruptive force in the labour market. By reducing labour demand and reshaping the occupational structure of employment, it has led to both employment and wage losses. However, much less is known about how automation influences labour market adjustment over the business cycle. Beyond its steady-state effects, an open question is how automation has shaped the capacity of labour markets to absorb and recover from severe aggregate shocks. This paper examines how automation intensity across labour markets lead to systematically different employment adjustment paths following the global financial crisis (GFC). The existing literature on automation shows that robots affect employment and wages differently across tasks and regions (e.g. Acemoglu and Restrepo (2020); Graetz and Michaels (2018)). In parallel, studies of the Great Recession document persistent local employment losses and long-run scarring (e.g. Yagan (2019); Fernald, Stock, Hall, and Watson (2017); Hershbein and Stuart (2024); Jaimovich and Siu (2020)). What remains unknown is whether pre-crisis automation intensity conditioned the resilience of local economies in the recovery. This paper addresses that gap by combining novel evidence from U.S. commuting zones with a quantitative model to show how automation shaped the trajectory of post-crisis labour market adjustment.

Empirically, we measure automation exposure as the CZ-level employment share weighted growth of industry robots prior to 2007. We proxy the shock of the local recession by the change in unemployment between 2007 and 2009, following Yagan (2019). Areas that adopt robots may differ systematically in industry composition, exposure to import competition, demographics, or pre-trends that also affect recoveries. To address this, besides having a rich set of controls at the CZ-level to control for pre-trends, we implement a Bartik style strategy that instruments CZ automation exposure with the same employment weights interacted with industry level robot adoption in a reference set of European economies, thereby isolating variation driven by global robotics trends rather than local U.S. shocks. Our empirical contributions are twofold. First, we document a new fact: conditional on the severity of the local Great Recession shock, commuting zones (CZs) with greater pre-crisis robot adoption experienced smaller long-run employment losses and faster recoveries. Second, we show that this faster recovery in employment rates is not simply compositional and it holds for both high- and low-skilled workers.

The empirical results establish a robust positive interaction between automation and employment recovery but do not identify the mechanisms driving the pattern by themselves. Since automation is typically associated with reduced labour demand, its apparent role in mitigating recession scarring poses a conceptual puzzle that requires formal analysis. We therefore propose a structural model which serves two purposes. First, the model disciplines the interpretation of the empirical results by embedding firms in an environment with explicit technological and financial heterogeneity. Second, it isolates the channels through which automation intensity shapes recovery dynamics, separating compositional forces from complementarities and investment behaviour over the cycle. Our heterogeneous-firm model with automation capital and labour can therefore reconcile the two key features of the data: the adverse long-run employment effects of automation, and the faster post-crisis recovery in more automation-intensive economies.

Firms are heterogeneous in productivity and produce using labour and automation capital, which they accumulate subject to the constraint that investment cannot exceed revenues net of labour costs. We study stationary equilibria that differ in automation intensity—interpreted as regions or commuting zones (CZs)—captured by the share of automation capital in a constant elasticity of substitution (CES) production function. In line with Acemoglu and Restrepo (2020), the model first shows that economies with higher automation intensity feature lower equilibrium wages and employment. We then analyse employment dynamics along the transition path following a transitory negative aggregate productivity (TFP) shock. Consistent with our empirical results, economies with greater automation intensity experience faster recoveries. This arises because firms in more automation-intensive economies hold larger stocks of automation capital and cut investment by less in response to the shock. The higher installed capital base also buffers adverse firm-level productivity shocks, sustaining production and labour demand through complementarities between labour and automation capital. These results are robust across parameter choices and to alternative assumptions on the size and persistence of the considered TFP shock.

Firm heterogeneity is central to reconciling the model with our main empirical finding. Differences in automation capital across firms, driven by idiosyncratic shocks, shape their capacity to recover once the TFP shock dissipates. Firms that exit the recession with relatively high automation capital are better able to absorb adverse shocks, sustain investment, and expand labour hiring given the partial complementarity between labour and automation capital. Because more automation-intensive economies have higher average capital stocks, firms in these regions reduce investment by less than those in less automated economies, leading to faster aggregate employment recovery. Finally, relaxing the investment constraint—allowing firms to invest beyond revenues net of labour costs—also strengthens recovery in less automation-intensive economies, as greater capital accumulation activates the same complementarity and resilience mechanisms that operate in more automation-intensive settings.

We establish both a novel empirical fact and a theoretical mechanism that together reframe how automation is understood in the context of macroeconomic adjustment. Empirically, pre-crisis automation exposure is associated with more resilient local employment dynamics: the adverse effect of the Great Recession on long-run employment is attenuated in regions with higher robot adoption. Theoretically, when firms differ in productivity and accumulate automation capital under financing frictions, greater automation intensity accelerates recovery because a larger installed capital base dampens the contraction in investment at the trough and supports rehiring as productivity normalises. Thus, while automation lowers steady-state employment in more automated regions, its dynamic role during recoveries operates through an "insurance-cum-complementarity" channel that mitigates hysteresis.

This paper sits at the intersection of three literatures. First, we advance the automation and labour markets literature by moving beyond average employment effects to examine cyclical adjustment, highlighting the role of installed capital and financing frictions in driving rehiring dynamics. Second, we contribute to research on regional hysteresis after macroeconomic shocks by showing that pre-crisis technology intensity systematically shaped recovery. Third, we relate to quantitative models with heterogeneous firms and investment frictions, introducing a novel extension—automation capital embedded in a CES production structure with GHH preferences (Greenwood, Hercowitz, & Huffman, 1988)—that generates sharp recovery predictions consistent with the data.

A large literature studies the labour-market effects of automation. Acemoglu and Restrepo (2018b) and Acemoglu and Restrepo (2020) show that U.S. commuting-zone exposure to robots reduces employment and wages, while D. Autor and Salomons (2018) link automation to labour displacement and falling labour shares. Cross-country evidence from Graetz and Michaels (2018) shows productivity gains with mixed employment effects, and Dauth, Findeisen, Suedekum, and Woessner (2021) find German job losses in manufacturing due to automation are offset by service-sector gains. Fewer papers examine cyclical dynamics: Jaimovich and Siu (2020) show routine jobs contract disproportionately in recessions and fail to rebound, whereas Graetz and

Michaels (2017) argue new technologies are unlikely to cause jobless recoveries outside the U.S. Our findings extend Jaimovich and Siu (2020) by showing that regions more automated before the GFC—already less reliant on the routine margin—exhibited greater resilience in recovery. Thus, while automation intensifies routine displacement, it also operates as a form of "insurance" that speeds post-crisis adjustment.

A burgeoning literature has documented the severity and persistence of recessions, particularly following the protracted recovery from the Global Financial Crisis (GFC) (see Cerra, Fatás, and Saxena (2023) for a recent survey). Yagan (2019) shows that a one–percentage point higher local unemployment shock in 2007–2009 reduced subsequent employment probabilities in 2015 by 0.4 percentage points, while Hershbein and Stuart (2024) find that local employment shocks from 1973–2009 failed to mean revert even after a decade. Explanations for slow recoveries include composition effects, duration dependence, and non-participation driving long-term unemployment (Kroft, Lange, Notowidigdo, and Katz (2016)); sharp declines in matching efficiency when worker characteristics deteriorate or regional conditions diverge (Barnichon and Figura (2015)); and weak aggregate demand, with only a modest role for increased labour-market frictions (Elsby et al. (2011)). Hershbein and Kahn (2018) add a technological channel, showing that hard-hit MSAs raised skill requirements in job postings, with effects persisting through 2015 and correlating with capital investment. Our finding that pre-GFC automation predicts stronger recoveries provides a place-based complement to this cyclical technology mechanism, corroborating the evidence on the persistence of downturns while showing that recoveries are systematically stronger in regions with higher pre-crisis automation intensity.

Recent work develops quantitative models of automation (e.g. Kopytov et al. (2018); Acemoglu and Restrepo (2018b); Jaimovich et al. (2021); Eden and Gaggl (2018); Hémous and Olsen (2022); Hubmer and Restrepo (2021); Firooz et al. (2025)). Most emphasise worker heterogeneity by skill or task, while abstracting from firm heterogeneity driven by idiosyncratic productivity shocks. Two exceptions—Hubmer and Restrepo (2021) and Firooz et al. (2025)—model costly adoption decisions by heterogeneous firms, showing that larger and more productive firms are more likely to automate, a margin that would also arise in our framework if adoption were endogenised. This literature has largely focused on long-run trends in wages, employment, the skill premium, the labour share, job polarisation, and the rise of superstar firms. Our emphasis is different: we analyse how automation interacts with the business cycle. While our model is consistent with evidence that wages and employment are lower in more automated regions Acemoglu and Restrepo (2020), our contribution is to show that prior automation exposure conditions the dynamics of employment recovery after a major downturn. Closest to our approach, Kopytov et al. (2018) study technology adoption during recessions, highlighting its role in accelerating job polarisation. Unlike their framework, however, we do not endogenise adoption. Instead, we provide evidence that employment recoveries differ systematically across regions with heterogeneous automation intensities, and offer a model in which complementarities between automation and labour play a central role in shaping post-crisis adjustment.

The remainder of the paper is structured as follows. Section 2 describes the data used to estimate our main empirical specification, presented in Section 3. Section 4 reports our baseline results and robustness checks. Section 5 develops a model of firm heterogeneity to interpret the empirical findings, and Section 6 presents the calibration and quantitative analysis. Section 7 concludes the analysis.

2 Data and Empirical Framework

2.1 Data on stock of industrial robots

Our primary measure of robot adoption is drawn from the International Federation of Robotics (IFR), which has conducted annual surveys of robot suppliers in more than 60 countries since 1993. The IFR dataset is the most widely used cross-country source on industrial robots (see Acemoglu and Restrepo (2020); Graetz and Michaels (2018)), and provides information on robot stocks at the industry–year–country level. Its consistent methodology and long time horizon make it uniquely suited for analysing how variation in automation intensity interacts with business cycle dynamics. The coverage spans six broad sectors—agriculture, mining, utilities, construction, education, and services—and thirteen detailed manufacturing industries. To complement these data, we use industry-level employment and output growth rates from the EU KLEMS Growth and Productivity Accounts Board (2023). Appendix Table 5 documents the evolution of robot adoption across U.S. industries between 1993 and 2015: automation expanded most rapidly in automotive, chemicals, and electronics, while growth was modest in construction and services.

2.2 Outcomes and robot exposure at the commuting zone level

To measure long-term changes in local labour markets, we draw on the IPUMS public-use census samples for 1990 and 2000 and the 2013–2017 American Community Survey (ACS) (Ruggles et al. (2023)).[1] The sample consists of non-institutionalised individuals aged 16–64. Our analysis is conducted at the commuting zone (CZ) level, using 722 CZs that cover the contiguous U.S. (excluding Alaska and Hawaii).[2]

Following Acemoglu and Restrepo (2020), we measure robot exposure in CZ i at time t as a shift-share index, where national changes in robot adoption at the industry level are weighted by local industry employment shares:

$$\Delta R_{i,t}^{US} = \sum_{j} \left[\frac{L_{i,j,1970}}{L_{i,1970}} \times \Delta R_{j,t} \right]. \tag{1}$$

Here, $\Delta R_{j,t}$ denotes the change in robot use in industry j, while $\frac{L_{i,j,1970}}{L_{i,1970}}$ is the share of industry j in CZ i's total employment in 1970. Using pre-determined 1970 weights avoids mechanical correlation between robot adoption and contemporaneous industry shares (cf. Chopra and Mukherji (2024)).

A concern with this measure is that unobserved local demand shocks could simultaneously affect employment and firms' incentives to adopt robots. To address this, we follow Acemoglu and Restrepo (2020) and instrument U.S. robot exposure with robot adoption patterns in five European countries (Denmark, Finland, France, Italy, and Sweden—"EURO5"). The IV replaces U.S. industry-level robot growth with its European counterpart in equation (1):

$$\Delta R_{i,t}^{EURO5} = \sum_{j} \left[\frac{L_{i,j,1970}}{L_{i,1970}} \times \Delta R_{j,t}^{EURO5} \right], \tag{2}$$

where $\Delta R_{j,t}^{EURO5}$ is the average growth in robot adoption across the five countries. Since technological advances in robotics are global, European adoption provides a powerful predictor of U.S. adoption while being plausibly orthogonal to local U.S. demand shocks. Appendix Figure 9 shows broadly parallel trends in robot use across North America, Germany, and the EURO5.

Figure 1a documents the strong correlation between U.S. and EURO5 robot adoption at the

[1] Following D. H. Autor et al. (2013), we use the pooled 2013–2017 ACS to measure outcomes in 2015, thereby increasing sample size. The census samples are 5%.

[2] CZs are groupings of counties linked by strong labour market and commuting ties (Tolbert & Sizer, 1996), and are a standard unit of analysis in the local labour markets literature (cf. D. H. Autor and Dorn (2013)).

industry level, while Figure 1b shows that the EURO5 exposure measure is a strong predictor of U.S. robot penetration at the CZ level. The instrument explains 87% of the cross-CZ variation in robot exposure, with a statistically significant first-stage coefficient. Our identification assumption is that national trends in robot adoption are exogenous to local economic conditions (Borusyak, Hull, & Jaravel, 2022). Under this assumption, cross-CZ differences in exposure can be interpreted as quasi-random, conditional on controls, allowing us to isolate the causal effect of automation on local employment.

Figure 1: Relationship between US and EURO5 robot exposure

(a) Industry

(b) Commuting Zone

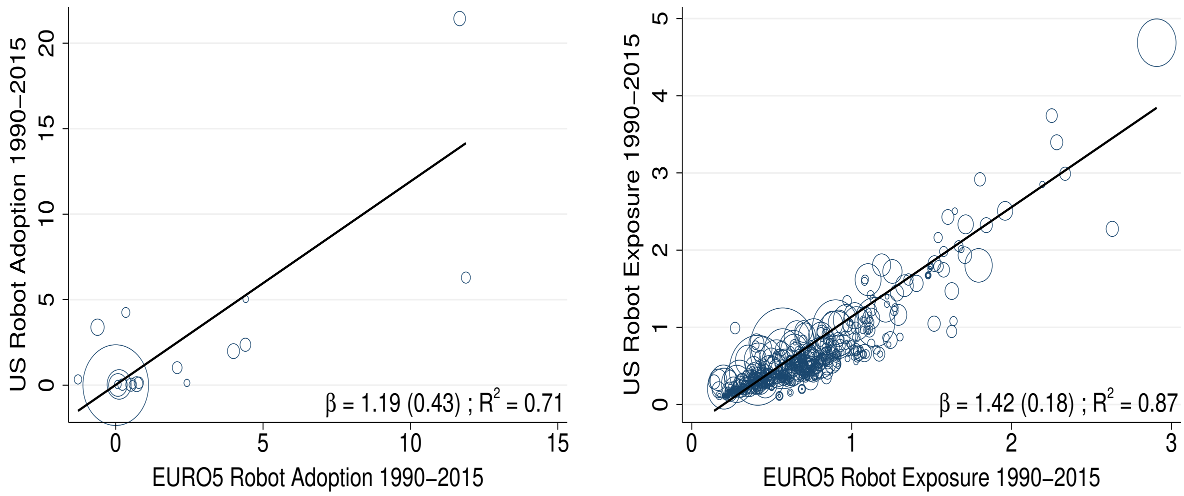

Note: Panel (a) plots the growth in robots per thousand workers at the industry level in US and EURO5 countries. The marker size indicates the US industry employment shares in 1990. Robust standard errors are displayed parentheses. Panel (b) shows the relationship between US and EURO5 robot exposure at the CZ level. The marker size indicates the 1990 population in the CZ. Clustered standard errors at the state level are displayed in parentheses.

2.3 GFC Shock

We measure the severity of the Great Recession shock following Yagan (2019). Using county-level unemployment rates from the Bureau of Labor Statistics Local Area Unemployment Statistics (BLS-LAUS), aggregated to the CZ, we compute the change in unemployment between 2007 and 2009:

$$\text{GFC Shock} = \text{Unemployment Rate}2009 - \text{Unemployment Rate}2007. \tag{3}$$

This measure captures the intensity of the labour market shock across local economies during the GFC and provides a basis for examining how recovery dynamics differ by pre-crisis automation exposure.

2.4 Other data

We incorporate data on trade exposure from D. Autor et al. (2019) to account for the employment effects of rising Chinese import competition.[3] We also control for demographic ageing, which has been shown to increase automation intensity across industries (Acemoglu & Restrepo, 2022).

To mitigate concerns that CZs more exposed to robots may also have experienced differential local population growth, we follow Lewis (2011) and include low-skilled population growth between 1970 and 1990. In addition, we absorb broad regional trends by including division–time fixed effects, following Faber, Sarto, and Tabellini (2022), and we control for a comprehensive set of pre-determined demographic and industrial characteristics measured in 1990. Demographic controls include log population and population shares by gender (male), educational attainment (no college, some college, college or higher), and race/ethnicity (White, Black, Hispanic). Industry controls include employment shares in manufacturing, light manufacturing, agriculture, construction, and mining. Because Acemoglu and Restrepo (2020) document that employment declines in light manufacturing are negatively correlated with robot penetration, we explicitly control for employment in textiles and printing. Finally, to account for nationwide compositional shifts in task content, we control for the employment shares of routine and offshorable jobs in 1990.[4] Taken together, these controls ensure that our estimates isolate the effect of robot adoption from confounding local shocks and pre-existing structural differences.

3 Empirical Specification

To estimate how pre-GFC crisis robot adoption influenced the long-term employment effects of the 2007–2009 Global Financial Crisis (GFC), we employ the following two-stage least squares (2SLS) regression model:

[3]Trade exposure at the CZ level is constructed as the weighted sum of industry-level growth in Chinese import penetration, where weights are industry employment shares. To address endogeneity between U.S. import demand and Chinese exports, the growth in U.S. imports from China is instrumented with Chinese export growth to a set of eight other developed economies: Australia, Denmark, Finland, Germany, Japan, New Zealand, Spain, and Switzerland.

[4]Following D. H. Autor et al. (2013), we compute the share of workers performing routine, manual, and abstract tasks, and construct industry-level measures of task offshorability.

$$\Delta\text{Employment}_{cz,2007-15} = \beta_1 \cdot \text{GFC Shock}_{cz} + \beta_2 \cdot \Delta\text{Robots}_{cz} + \beta_3 \cdot (\text{GFC Shock}_{cz} \times \Delta\text{Robots}_{cz}) + \mathbf{X}'_{cz}\gamma + \delta_r + \varepsilon_{cz} \tag{4}$$

where:

- $\Delta\text{Employment}_{cz,2007-15}$ is the change in the employment rate in commuting zone cz from 2007 to 2015.
- GFC Shock_{cz} is the change in the unemployment rate from 2007 to 2009, which captures local exposure to the GFC. (which is the same shock measure used by Yagan (2019))
- ΔRobots_{cz} measures the increase in robot exposure from 2000 to 2007.
- $\mathbf{X}_{cz}$ is a vector of control variables including demographics, industry shares, trade exposure, and offshoring risk.
- δ_r represents Census division fixed effects.
- ε_{cz} is the error term.

To address potential endogeneity in robot adoption, we instrument ΔRobots_{cz} using a shift-share instrument based on 1990 industry employment shares and national trends in robot adoption in EURO5 countries as explained in Section2.2. All specifications include robust standard errors clustered at the CZ level, and the first-stage F-statistics exceed 100 in all models, indicating strong instrument relevance. The key coefficient of interest, β_3, captures whether regions with greater adoption of robots were more insulated from persistent employment losses following the GFC shock. The regressions are weighted by the CZ population of 1990 to give more influence to larger labour markets. We include division-time dummies and a number of CZ-level controls $\mathbf{X}_{cz}$ that are essential to reduce the omitted variable bias. They ensure that the estimated interaction effect between the GFC shock and the adoption of robots (β_3) is not confounded by differences related to demographic, industrial, or trade between the CZs.

4 Results

In this section, we begin by presenting raw descriptive evidence on how local employment responded to the GFC in regions with differing levels of pre-crisis automation. We then turn to our IV regression framework, where we incorporate a full set of controls, and finally examine the marginal effect of pre-GFC robot adoption on long-run employment recovery.

Figure 2: Employment Recovery and Pre-GFC Robot Adoption

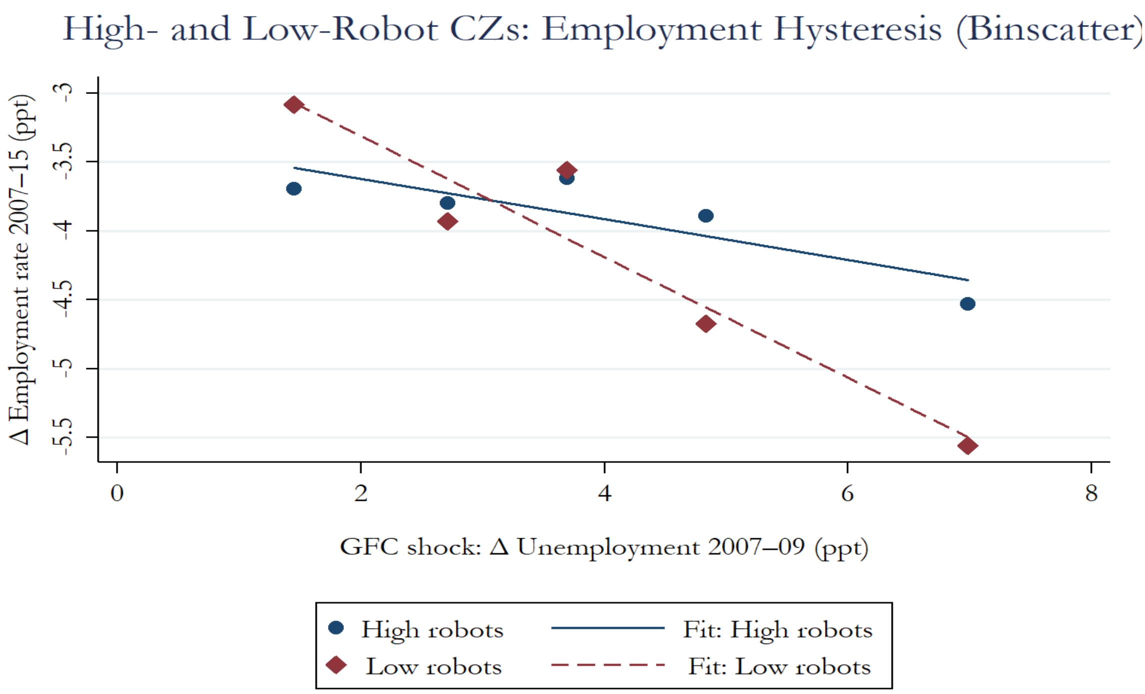

Figure 2 compares post-crisis employment dynamics in commuting zones (CZs) with high versus low levels of robot adoption. The horizontal axis measures the severity of the GFC shock, defined as the change in unemployment between 2007 and 2009, while the vertical axis plots the change in the employment rate from 2007 to 2015. High-robot CZs (blue circles, solid line) are those with greater pre-crisis robot adoption; low-robot CZs (red diamonds, dashed line) are those with less.

Both groups exhibit negative slopes: areas hit harder by the GFC experienced larger long-term employment losses. Yet the slope is markedly steeper for low-robot CZs, implying more persistent disemployment. By contrast, high-robot CZs display a flatter slope, indicating that regions with greater pre-crisis automation were more resilient, with employment recovering more quickly for a given initial shock.

These raw patterns suggest that automation cushioned the long-run employment impact of the GFC. However, the figure reflects unconditional correlations: the binscatter does not control for differences in demographics, industry mix, or trade exposure across CZs, nor for broader macroeconomic trends that may affect both robot adoption and crisis severity. In the next step, we turn to a regression analysis in Table 1, where we explicitly account for these covariates to assess whether the resilience of high-robot regions survives when confounding factors are addressed.

Table 1: 2SLS: Instrument = ΔRobot Shift–Share, 1990 Baseline

	(1)	(2)	(3)	(4)	(5)
Shock (β_1)	-0.53^{***}	-0.45^{***}	-0.43^{***}	-0.42^{***}	-0.43^{***}
	(0.09)	(0.10)	(0.10)	(0.10)	(0.10)
Δ Robots (β_2)	-2.91^{***}	-0.86	-0.72	-0.70	-0.66
	(0.92)	(0.87)	(0.95)	(0.94)	(0.90)
Shock $\times$ ΔRobots (β_3)	0.46^{***}	0.24^{**}	0.23^{**}	0.23^{**}	0.23^{**}
	(0.13)	(0.11)	(0.11)	(0.11)	(0.10)
Observations	722	722	722	722	722
R^2	0.25	0.46	0.47	0.47	0.47
Census divisions	Yes	Yes	Yes	Yes	Yes
Demographics		Yes	Yes	Yes	Yes
Industry shares			Yes	Yes	Yes
Trade, routine jobs				Yes	Yes
Offshore index					Yes

Notes: 2SLS regression instrumenting ΔRobots with the 2000–2007 robot adoption shift–share (ΔRA_{iv}). First-stage F-statistics for excluded instruments >100 in all columns. See text for control definitions. ***/**/* denote p<0.01/0.05/0.1.

Table 1 estimates how the local employment effects of the GFC shock vary with pre-crisis exposure to robots, instrumenting robot adoption with the EURO5 measure to address endogeneity. The coefficient on the GFC shock, β_1, is negative and statistically significant across all specifications, confirming that regions hit harder by the crisis experienced larger and more persistent employment losses between 2007 and 2015. The coefficient on robot adoption alone, β_2, is negative but generally imprecise once controls are included, suggesting that robot exposure by itself does not robustly predict long-run employment changes. Our central coefficient of interest, β_3, is positive and statistically significant in every specification. This interaction term shows that regions with greater pre-crisis robot adoption were systematically less affected by GFC shocks, consistent with automation cushioning the long-term employment impact of the downturn. In other words, prior automation exposure appears to have shielded local labour markets from persistent hysteresis effects of the Great Recession.

The specification includes a rich set of CZ-level controls measured in 1990. Robustness exercises (available in the Appendix) confirm the stability of the result: Table 6 uses 2000 covariates in place of 1990 covariates, while Tables 7 and 8 employ an alternative instrument—the log of robot stock in CZs in 2000—combined with both 1990 and 2000 covariates. This confirms that it is not the EURO5 measure that is leading to our results. Across all these variations, the interaction term β_3 remains positive and statistically significant, reinforcing the conclusion that employment recovered more strongly in regions with greater prior exposure to robots. These regression results confirm the raw patterns shown in Figure 2, where high-robot regions displayed systematically flatter recovery slopes.

Figure 3: Relationship between GFC shock, robot exposure, and employment

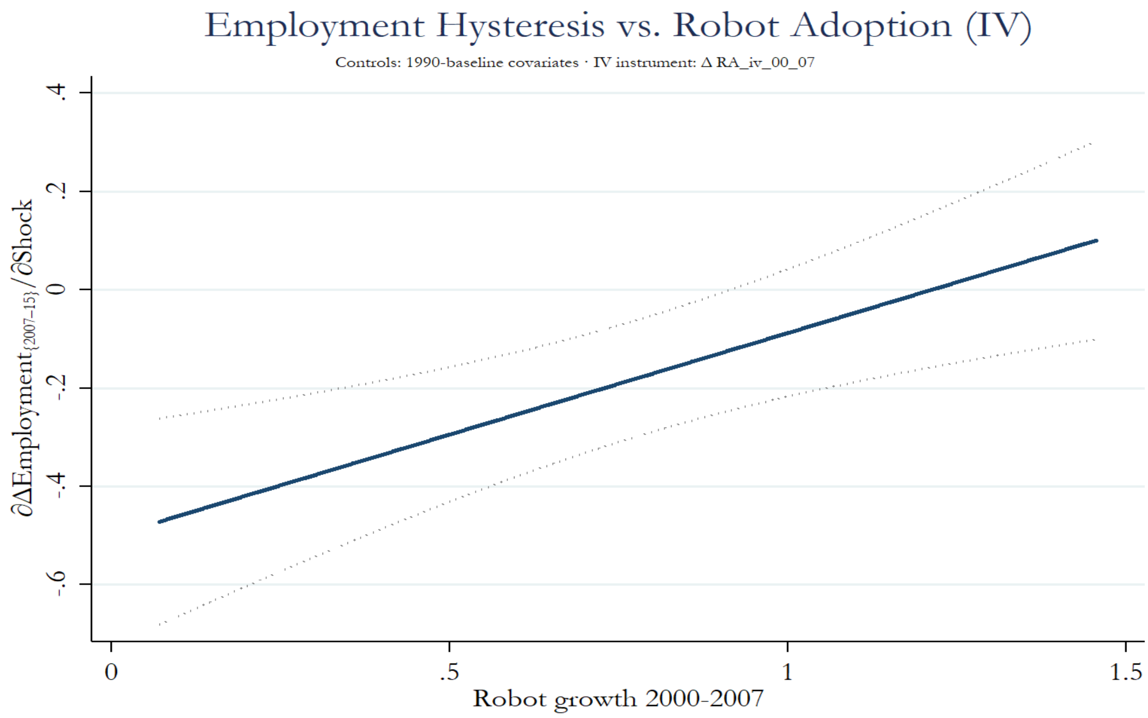

We next examine how the marginal effect of the GFC shock on long-run employment varies with pre-crisis robot adoption using our regression estimates. Figure 3 plots this relationship. The vertical axis measures $\partial\Delta\text{Employment}_{2007\text{\textasciitilde}15}/\partial\text{Shock}$, that is, the marginal effect of a one–percentage point increase in the 2007–2009 unemployment shock on employment between 2007 and 2015. The horizontal axis shows the growth of robots from 2000 to 2007. The solid blue line traces the estimated relationship, while the dotted lines display 95% confidence intervals. The slope is positive and statistically significant: as robot adoption rises, the long-run employment loss associated with a given GFC shock becomes less severe. At low levels of robot growth, the effect of the GFC shock is strongly negative (around –0.5 or worse), whereas at higher levels of robot adoption the effect approaches zero or even turns slightly positive, indicating a substantial attenuation of hysteresis.

Because automation may affect different types of labour unequally—displacing routine and low-skill workers more readily—we also plot analogous figures disaggregated by skill group. We classify high-skilled individuals as those with at least a high school diploma, and low-skilled individuals as those with less than a high school education. This allows us to assess whether the cushioning effect of automation differs systematically across skill categories. As we show below, the resilience effect of robot adoption is evident for both groups, though somewhat stronger among high-skilled workers.

Figures 4a and 4b plot regional robot adoption during 2000–2007 against the estimated employment response from 2007 to 2015 per unit increase in GFC exposure, $(\partial\Delta\text{Employment}/\partial\text{shock})$. In both panels, the slope of the estimated relationship is positive: regions with greater pre-crisis robot adoption experienced less severe long-run employment losses following the GFC. Strikingly, this pattern holds for both low- and high-skilled workers. Although robots are generally thought to complement high-skill tasks and substitute for low-skill tasks, the figures show that recovery dynamics are broadly similar across groups: employment losses are attenuated in high-adoption regions regardless of skill composition. This suggests a common recovery mechanism rather than a widening divergence in adjustment across skill groups. To probe heterogeneity further, we next examine recoveries by occupational task content, since certain tasks are more directly exposed to automation Acemoglu and Restrepo (2020). Table 2 presents results from the same specification as Table 1, disaggregated by task category.

Table 2: Recovery by Task

	(1) Routine	(2) Manual	(3) Abstract
Shock (β_1)	-0.04	-0.14^{*}	-0.30^{***}
	(0.07)	(0.07)	(0.11)
$\Delta\text{RA}_{00\text{–}07}$ (β_2)	-0.60	0.34	-1.03
	(0.55)	(0.58)	(0.81)
Shock $\times$ ΔRA (β_3)	0.11	0.04	0.18^{*}
	(0.09)	(0.09)	(0.11)
Observations	722	722	722
R^2	0.13	0.48	0.30
Census divisions	Yes	Yes	Yes
Demographics	Yes	Yes	Yes
Trade	Yes	Yes	Yes
Offshore index	Yes	Yes	Yes

Figure 4: Relationship between GFC shock, robot exposure, and employment by skill

(a) Low-skilled

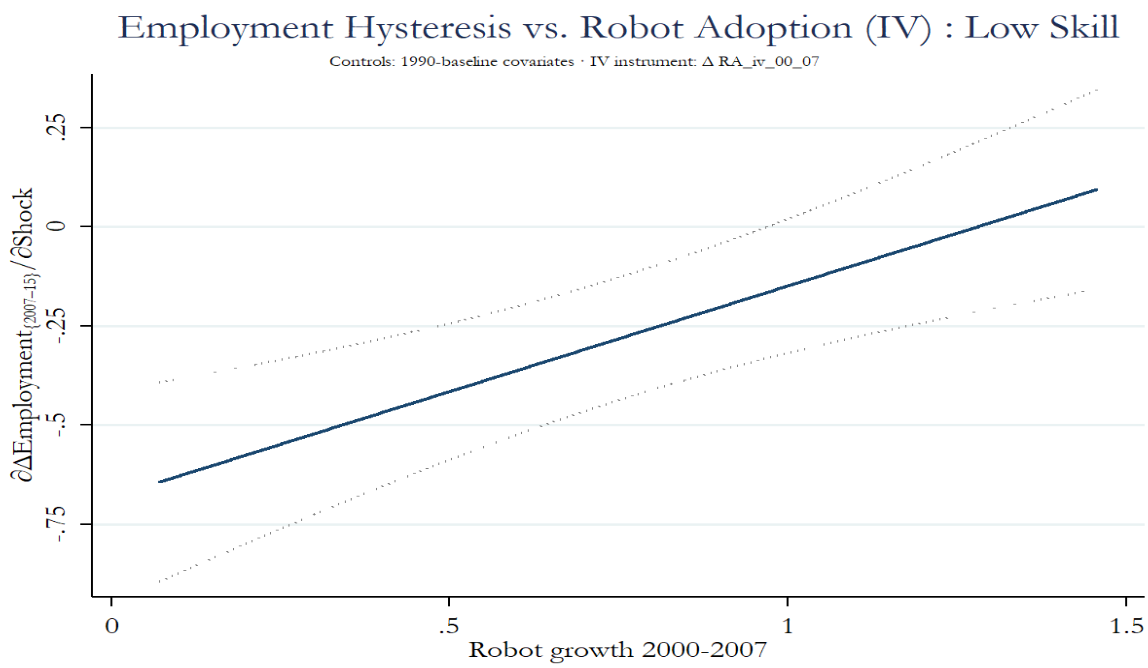

(b) High-skilled

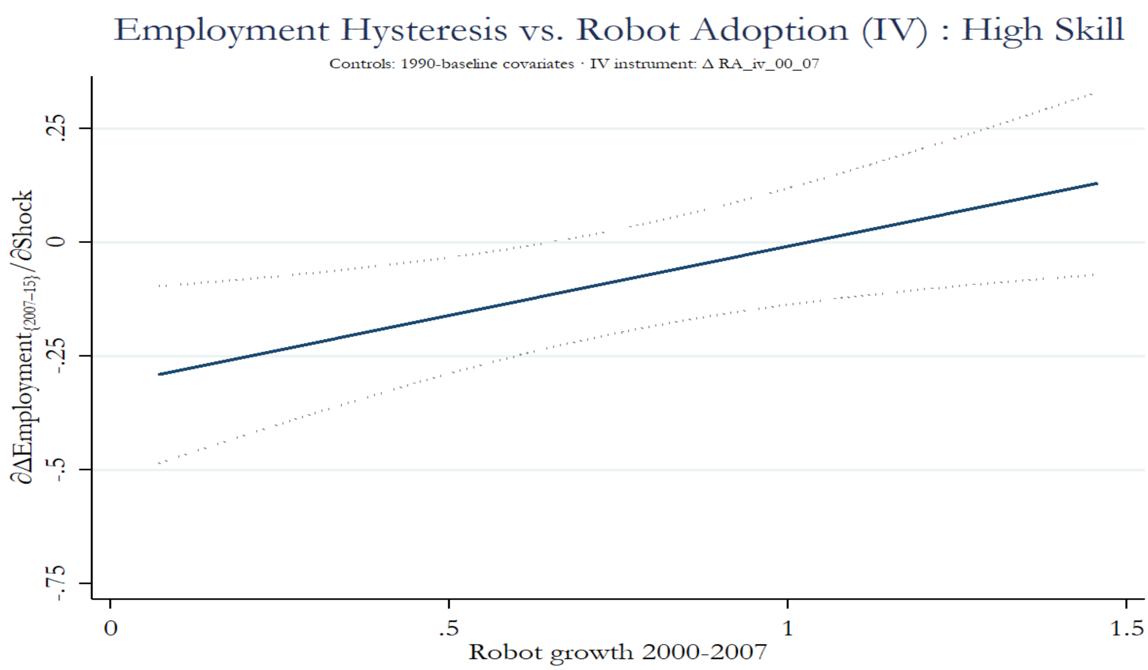

Table 2 disaggregates the effects of the GFC-related automation shock across routine, manual, and abstract tasks. The direct impact of the shock (β_1) is negative for all three categories, underscoring that the crisis generated broad-based employment losses. We find the interaction with robot adoption (β_3) is positive across tasks and statistically significant for abstract work. This pattern indicates that regions with greater pre-2007 exposure to robots experienced a cushioning effect: employment losses were systematically smaller, and recoveries faster, particularly in abstract occupations where complementarities with automation are strongest. The fact that the interaction is positive for routine and manual tasks as well, though less precisely estimated, suggests that the insurance effect of automation may be more general, with limited statistical power rather than an absence of adjustment driving the weaker significance. Taken together, the results point that while automation does displace certain tasks in the short run, prior adoption also appears to stabilize labour markets during downturns by accelerating recovery, especially in higher-skill abstract employment.

Table 3 reports a placebo exercise that examines employment changes in the pre-period (2000–2007), before the onset of the GFC. The specification mirrors that used in our baseline regressions, but the dependent variable is the change in overall employment during the pre-period. Across both OLS and IV, the estimated coefficients are small and statistically indistinguishable from zero. In particular, the interaction term between the automation shock and prior robot adoption is negative but economically negligible and not significant. These results confirm that our main findings are not driven by pre-existing differential trends in employment between high- and low-adoption regions. Instead, the cushioning effect of prior automation exposure emerges only in the post-crisis recovery period, strengthening the interpretation that the interaction captures a genuine insurance mechanism rather than spurious pre-trends.

Table 3: Pre-Period Change in Overall Employment Rate (2000–2007)

	OLS	IV
Shock (β_1)	-0.062	-0.006
	(0.089)	(0.091)
ΔRA$_{00\text{-}07}$ (β_2)	-0.315	0.053
	(0.780)	(0.916)
Shock $\times$ ΔRA (β_3)	-0.061	-0.134
	(0.141)	(0.145)
Observations	722	722
R^2	0.605	0.604
Census divisions	Yes	Yes
Demographics	Yes	Yes
Industry shares	Yes	Yes
Trade, routine jobs	Yes	Yes
Offshore index	Yes	Yes

5 Model

We now present a model that is consistent with the main empirical finding reported above, viz. that regions (CZs) with greater exposure to robot usage prior to the great recession witnessed faster employment recoveries following the great recession.

The model features heterogeneous firms that receive idiosyncratic productivity shocks following H. A. Hopenhayn (1992), and can invest in automation capital ('robots') subject to a non-negativity requirement on firm cash-in-hand (revenue net of labour cost). Firms produce the output good using robots and labour, and the share of robots utilised in the production process is governed by the automation threshold that will be varied in numerical simulations in section 6.

5.1 Model environment

Time is discrete and continues forever. There are two types of agents: (i) a representative consumer, who is risk-averse, supplies labour, receives firm profits and consumes the output good; and (ii) a unit continuum of heterogeneous firms, which are risk-neutral, face idiosyncratic productivity shocks as in H. A. Hopenhayn (1992), and produce output by hiring labour and accumulating automation capital (robots). Hence, we consider a heterogeneous firm model with (automation) investment in the spirit of Gomes (2001), Cooley and Quadrini (2001) and

Clementi and Palazzo (2016); although we abstract from the entry-exit margin. Aggregate TFP (z) is normalised to 1 in steady state. We now describe the behaviour of these agents in greater detail.

5.1.1 Firms

Firms receive idiosyncratic productivity shocks y that follow a Markov process independent across firms with conditional distribution $P(y'|y)$, where primes are used to represent variables in the following period. Firms discount profits using the discount factor β. Firms are owned by the representative worker-consumers, hence profits accrue to those agents.

Firms behave competitively and produce the numeraire homogeneous output good using labour (l) hired at unit wage w and automation capital[5] (k) that they accumulate. Automation capital depreciates at rate δ. The combination of idiosyncratic productivity shocks and automation capital investment result in firm heterogeneity. The idiosyncratic state vector for firms is (y, k).

Labour and automation capital produce output as per the following constant elasticity of substitution (CES) production function:

$$F(k, l; y, z, \theta) = zy\big(\theta k^{\epsilon} + (1-\theta)l^{\epsilon}\big)^{\frac{1}{\epsilon}} \tag{5}$$

The parameter θ represents the automation share that is common to all firms in an economy, and shall play an important role in our analysis. We shall vary θ in order to represent economies or regions with different intensities of robot usage. The parameter $\epsilon \equiv \frac{\gamma-1}{\gamma}$ is related to the elasticity of substitution (γ) between automation capital and labour.

A firm with state (y, k) chooses labour (l) and automation capital for the following period, (k'), in order to maximize its value:

$$J(y, k) = \max_{k', l} F(k, l; y, z, \theta) - wl - \big(k' - (1-\delta)k\big) + \beta \mathbb{E}_{y'|y}\, J(y', k') \tag{6}$$

subject to the following constraint:

[5] We discuss the nature of capital in section 5.3.

$$0 \leq k'(y,k) \leq F\big(k, l(y,k); y, z, \theta\big) - wl(y,k) + (1-\delta)k \tag{7}$$

The constraint restricts firms from investing an amount that results in negative profits in any period. In addition, firms cannot have negative capital, implying that firms can reduce their capital stock subject to it remaining non-negative. Consequently, the policy function for labour, $l(y,k)$, follows from a static constrained profit maximization problem. The unconstrained value of next period's automation capital, $k^u(y,k)$, is obtained from the Euler equation:

$$1 = \beta \mathbb{E}_{y'|y}\Bigg[F_k\big(k', l(y',k'); y', z, \theta\big) + (1-\delta)\Bigg] \tag{8}$$

Incorporating the constraint in equation (7), the policy function for automation capital next period is:

$$k'(y,k) = \max\Big\{0, \min\{k^u(y,k), F\big(k, l(y,k); y, z, \theta\big) - wl(y,k)(1-\delta)k\}\Big\} \tag{9}$$

The policy functions in turn yield firm output and profits:

$$Y(y,k) = F\big(k, l(y,k); y, z, \theta\big) \tag{10}$$

$$\pi(y,k) = F\big(k, l(y,k); y, z, \theta\big) - wl(y,k) - \big(k'(y,k) - (1-\delta)k\big) \tag{11}$$

The endogenous distribution of firms is denoted by $\lambda(y,k)$. One can then use the distribution and policy functions in order to obtain aggregate labour demanded by firms, the aggregate profits rebated to consumers and aggregate automation capital:

$$L^d = \int l(y,k)\; d\lambda(y,k) \tag{12}$$

$$\Pi = \int \pi(y,k)\; d\lambda(y,k) \tag{13}$$

$$K = \int k\; d\lambda(y,k) \tag{14}$$

5.1.2 Consumers

The representative consumer owns firms and hence receives firms profits, supplies labour to these firms and consumes the output good. We assume that the representative consumer has preferences over consumption (C) and labour (L) that take the following form based on Greenwood et al. (1988):

$$U(C,L) = \frac{\left(C - \frac{L^{1+\psi}}{1+\psi}\right)^{1-\sigma}}{1-\sigma} \tag{15}$$

This functional form ('GHH' preferences) eliminates the wealth effect on labour supply, and is often employed in the real business cycle literature. As we will note below, the use of GHH preferences allows us to be consistent with the empirical finding that high robot usage economies have lower wages and employment relative to low robot usage economies. Furthermore, the heterogeneous firms literature (e.g. Gomes (2001)) often makes alternative assumptions about preferences, specifically that they are quasi-linear in leisure following Hansen (1985), that also eliminate the wealth effect on labour supply. We now show that GHH preferences will also yield a labour supply function that is consistent with the specifications considered in the heterogeneous firms literature.

As in the heterogeneous firm literature following H. Hopenhayn and Rogerson (1993), the representative consumer does not save and consumes her resources in each period, which are the sum of wage income and aggregate profits rebated, $\Pi + wL$.

With GHH preferences, the intra-temporal labour supply condition is:

$$L = w^{\frac{1}{\psi}} \tag{16}$$

We observe that the labour supply condition derived in equation (16) is similar to the exogenous labour supply function assumed in heterogeneous firm models like Clementi and Palazzo (2016).

5.2 Stationary equilibrium

We consider a stationary equilibrium of the model, wherein aggregate variables are time invariant. A stationary equilibrium consists of the wage (w), policy functions $\{k'(y,k),\ l(y,k)\}$ and labour supply choice (L) by the representative agent such that: (i) Firms choose labour and automation capital to maximize value as in equation (6); (ii) The representative worker chooses labour supply based on equation (16); (iii) The labour market clears: $L = \int l(y,k)\ d\lambda(y,k)$; (iv) The stationary distribution λ^* evolves as per:

$$\lambda^*(y',k') = \int \mathbb{I}_{\{k'(y,k)=k'\}}\ P(y'|y)\ d\lambda^*(y,k) \tag{17}$$

In the expression above, $\mathbb{I}_{\{k'(y,k)=k'\}}$ is an indicator function for when the automation capital policy function for a firm with state (y,k) equals k'; while $P(y'|y)$ is the Markov transition probability.

5.3 Discussion

The model of firm heterogeneity presented above abstracts from certain modeling assumptions made by the related literature on automation, which we now address. Firstly, we do not consider worker heterogeneity by skill, as in e.g. Kopytov et al. (2018) and Jaimovich et al. (2021). This is because we find empirically (cf. fig-4a) that in any region, regardless of its automation intensity, the speed of employment recovery is similar by worker skill levels. This is also why we do not include physical capital in addition to automation capital in our baseline model, as it is often assumed to be complementary to both types of workers, and this complementarity is captured in reduced form in our simpler setup.

However, we do consider a more general specification incorporating both physical and automation capital in appendix D. It could be argued that regions might have relatively fewer robots, but firms in those regions might still be able to insure against bad draws of the idiosyncratic productivity shock through the accumulation of physical capital. We also note that robots are conventionally thought of as substituting for labour and being more complementary with machinery or equipment, which the more general specification can account for. We find that our broad result is unchanged even after incorporating physical capital.

The production function used in equation (5) bears similarities to the one used in the task-based framework of Acemoglu and Restrepo (2018a) and Acemoglu and Restrepo (2018b).

Specifically, we distinguish between less and more automated economies based on the automation share parameter θ, while Acemoglu and Restrepo (2018b) vary the task threshold to model the process of automation[6]. While Acemoglu and Restrepo (2018b) assume that automated tasks are performed by capital that is accumulated, they do not distinguish between physical and automated capital. However, an increase in the automation task share also increases capital demand in their setup. To emphasise the fact that an increase in the automation share (θ) leads to greater adoption of robots, we choose to refer to k as 'automation capital'.

We do not allow firms to choose their automation share (θ), unlike Hubmer and Restrepo (2021) and Firooz et al. (2025). While this is conceptually straightforward, we note that it is not necessary to explain our main empirical finding. Furthermore, an advantage of endogenising the choice of automation shares by firms would be the ability to assert whether automation intensity rises following a recession[7]. However, the empirical evidence for this is inconclusive[8].

6 Numerical results

We first discuss how some of the key parameters of the model are calibrated. We then compare the stationary equilibria of economies (i.e. regions) with different degrees of automation intensity, examining in particular the behaviour of wages and employment. We find that economies with a higher automation threshold, resulting in greater robot usage in these economies, have lower wages and employment. This is consistent with the findings of Acemoglu and Restrepo (2020). Our main numerical exercise simulates a recession using a transitory negative TFP shock and examines the recovery of employment along the transition path back to stationary equilibrium. We shall demonstrate that high robot usage economies hit by a negative TFP shock recover faster than low robot usage economies facing the same shock. Finally, we examine the role of constraints on firm investment in the baseline model in enabling employment

[6]One can show that if tasks are symmetric and the threshold separating robot and labour use is exogenous, then the resulting CES production function in Hubmer and Restrepo (2021) is similar to the production function employed in equation 5.

[7]It would also not be possible in the present setup to solve for the transition path to a stationary equilibrium that would result in a different degree of automation adoption in stationary equilibrium simply following a transitory negative TFP shock, unless one also considers a permanent change in the economy wide technological limit on automation, $\bar{\theta}$.

[8]Prior research shows that investments in research and development (R&D) tend to be less cyclical or even countercyclical compared to aggregate investment or physical capital investment (e.g. Bloom (2007), Benhabib et al. (2014)). At the same time, recent research in labour economics suggests that technological change often accelerates its impact on workers and jobs during downturns (e.g. Hershbein and Kahn (2018), Jaimovich and Siu (2020) and Kopytov et al. (2018)). Compared to R&D and investment, evidence on the cyclicality of frontier technology adoption like robots remains scarce (Cerra et al. (2023)), with some evidence of firm reducing robot adoption during the COVID-19 pandemic in Germany (Arntz et al. (2024))

recoveries following a recession.

6.1 Calibration

The baseline model has five production function and preference parameters that need to be calibrated: $\{\theta, \epsilon, \beta, \sigma, \psi\}$. In addition, there are two more parameters related to the persistence (ρ) and volatility (σ_ϵ) of the productivity shock that are used to obtain the Markov transition probability matrix. We use $\{\rho = 0.9, \sigma_\epsilon = 0.1\}$, which is slightly higher than corresponding estimates in Khan and Thomas (2008), and we shall vary these parameters in sensitivity analyses in appendix B.

We choose standard values for the time preference factor $(\beta = 0.95)$ and the coefficient of relative risk aversion $(\sigma = 2)$. The parameter ψ, the inverse of the Frisch elasticity, is chosen to equal 1. This is above the estimate of 1.33 recommended by Chetty et al. (2011), which we consider in the sensitivity analysis reported in appendix B.

The parameter ϵ, which is positively related to the elasticity of substitution between automation capital and labour, is chosen to be 0.5. This implies an elasticity of substitution equal to 2. As noted above, we do not incorporate skill heterogeneity (as in, e.g., Kopytov et al. (2018)), and it has been observed that low skill workers are more substitutable with automation capital, while high skill workers are more complementary with automation capital. One could interpret our calibrated parameter value as reflecting the higher substitutability between automation and especially low skill labour.

This parameter value is also in the range of values reported in the literature. For example, DeCanio (2016) finds the elasticity of substitution between human and robotic labour is greater than about 1.9. Berg et al. (2025) find in their econometric estimations with data up to 2020 that the elasticity of substitution between 'robots' (including AI) and low skill labour is between 2.2 and 2.5. Eden and Gaggl (2018) finds that the elasticity of substitution has increased rapidly since the late 1990s, rising from 2.5 to 3.27.

Finally, we consider a range of values[9] for the automation capital share threshold, θ, that exceed 0.1. We shall typically assume that economies that are less automation intensive have a value of $\theta \in (0.2, 0.3]$, while economies that are more automation intensive have values of θ exceeding 0.3. We have considered alternative thresholds for classifying economies based

[9]Values of θ at or below 0.1 are inconsistent with the empirical finding that wages and employment are lower in more automation intensive economies.

Table 4: Variation of wages and employment with automation usage

	(1) Log wages	(2) Log employment
Log automation capital	-0.052***	-0.052***
	(0.00542)	(0.00536)
N	500	500
adj. R^2	0.391	0.390

Robust standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

on automation intensity (conditional on $\theta > 0.1$) and observe that our main findings are unaffected.

6.2 Wages and employment in stationary equilibrium: variation with automation intensity

We now show that the model is consistent with the empirical finding that wages and employment are lower in more automated regions (see e.g. Acemoglu and Restrepo (2020)). In order to do so, we solve for the stationary equilibria of a range of model economies differentiated by values of the automation capital share parameter, θ. We then regress the log of equilibrium wages or employment on log automation capital usage, and find a significant negative relationship[10], as listed in table 4.

We note that the assumption of GHH preferences is important in yielding this result. As noted above, workers earn wage income and receive profits from firms. When comparing outcomes between less and more automation intensive economies, it is the latter which tend to have higher aggregate profits. With preferences that are separable between consumption and leisure, and therefore feature a wealth effect on labour supply, the higher profits in more automation intensive economies diminish labour supply adequately so as to offset the dampened labour demand by firms and consequently *raise* wages. When the wealth effect on labour supply is switched off with GHH preferences, then the reduced labour demanded by firms in automation intensive economies tends to lower both wages and employment.

6.3 Employment recoveries following a recession

We now demonstrate that the model is capable of replicating our main empirical finding, concerning the relatively slower recovery of employment in regions that are less automation in-

[10]The coefficients on both dependent variables are identical because of the intra-temporal labour supply condition, equation (16), and the assumption that $\psi = 1$. Setting $\psi = 1.33$, following Chetty et al. (2011), does not affect the negative relationship between wages or employment and automation capital.

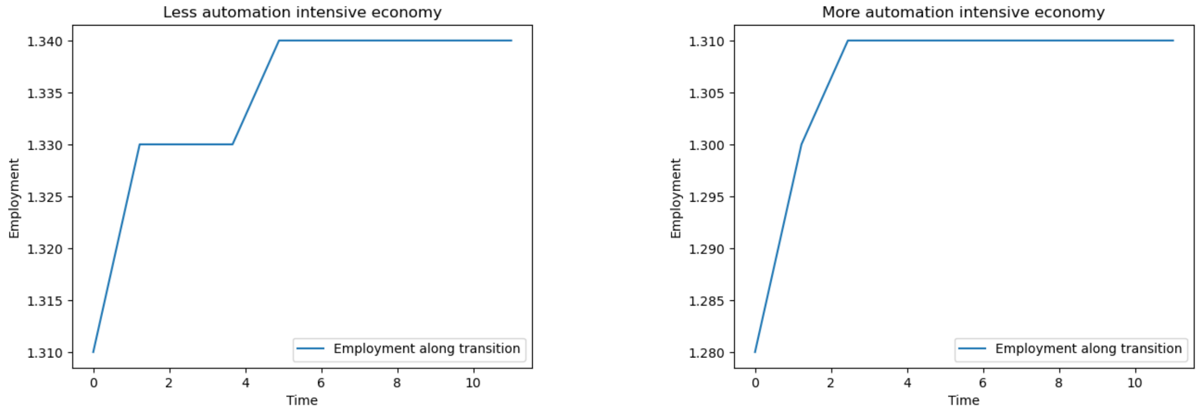

(a) Low θ economy

(b) High θ economy

Figure 5: **Employment recovery following a transitory negative TFP shock**. In (a), the employment recovery in a less automation intensive economy is slower than the corresponding recovery in a more automation intensive economy depicted in (b).

tensive.

We simulate a recession by lowering the value of aggregate TFP from its benchmark value of 1 in stationary equilibrium to 0.97 for a single period. We assume that this shock is not anticipated by agents in an economy, and is transitory. Hence, the shock is a 'MIT shock', defined by Boppart et al. (2018) as an unexpected shock that hits an economy at its steady state, leading to a transition path back towards the economy's steady state. In appendix B, we show that our result is robust to larger and more persistent shocks to TFP.

We compute the transition path back to the initial stationary equilibrium of two economies or regions facing a negative TFP shock of the same magnitude but which are differentiated by their automation intensities, parameterised by θ. Here, we choose values of $\theta_l = 0.3$ and $\theta_h = 0.4$ to demonstrate our result, although appendix B shows that this result is robust to alternative combinations of $\{\theta_l, \theta_h\}$ which exceed 0.1.

Figure 5 contrasts the employment recovery following a single period negative TFP shock in economies with low and high automation intensities (panels (a) and (b) respectively). Clearly, employment attains its stationary equilibrium value faster (in 2 periods) in the more automation intensive economy as compared to the less automation intensive economy, where it takes around 5 periods to reach the corresponding stationary equilibrium level of employment.

6.3.1 Understanding the variation in employment recovery rates by automation intensity

The relatively faster recovery in more automation intensive regions is driven by the fact that firms in such regions tend to have more automation capital in any time period. Firms that are

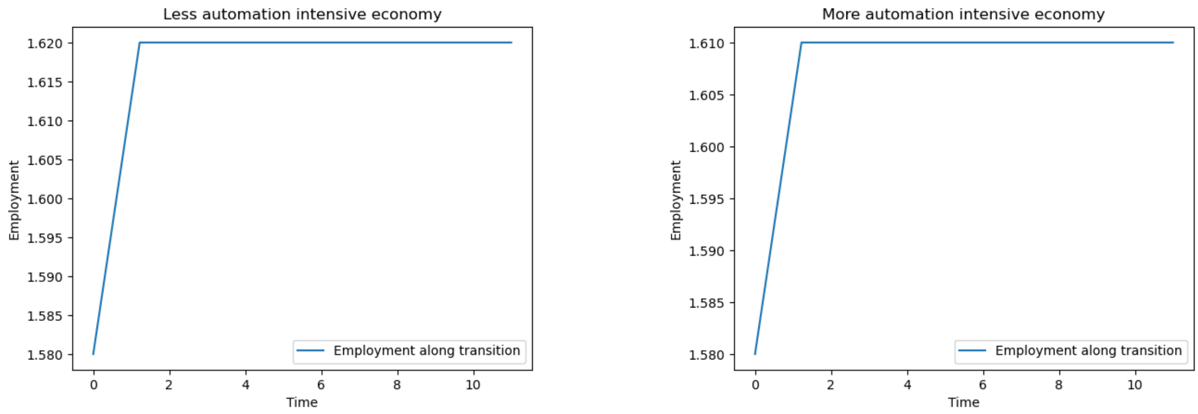

(a) Low θ economy

(b) High θ economy

Figure 6: **Employment recovery following a transitory negative TFP shock: higher elasticity of substitution**. In (a), the employment recovery in a less automation intensive economy is similar to the corresponding recovery in a more automation intensive economy depicted in (b).

more reliant on automation capital also tend to reduce investment in such capital to a lesser degree than firms in economies that are more reliant on labour when hit by a negative TFP shock. Higher automation capital also allows firms to offset low idiosyncratic productivities in the production process. Collectively, this strengthens the ability of firms based in more automation intensive economies to pick up production following the dissipation of the negative TFP shock. One could think of this as the 'insurance' role of automation capital.

Consequently, there is a relatively greater mass of firms in the automation intensive economy with average or higher levels of capital stock conditional on receiving a low draw of the idiosyncratic productivity shock. Such firms could offset the low idiosyncratic productivity if they have a high stock of capital, which boosts revenue and investment. As labour is (partially) complementary to automation capital, with the strength of the complementarity being driven by the elasticity of substitution (ϵ) between the two factors, a higher automation capital stock raises labour demand and leads to a faster employment recovery. Note that while automation capital and labour are partially complementary, the *level* of employment in more automation intensive economies is still lower than the corresponding level in less automation intensive economies[11].

Interestingly, we find that raising the elasticity of substitution parameter (ϵ) from a benchmark value of 0.5 to 0.7 results in employment recoveries at similar rates following the nega-

[11]As noted by Acemoglu and Restrepo (2018a), an increase in the automation share has two opposing effects: a greater dependence on cheaper capital boosts productivity and wages ('productivity effect'), while also displacing labour ('displacement effect'). In our setup, the negative effect on labour demand and wages dominates.

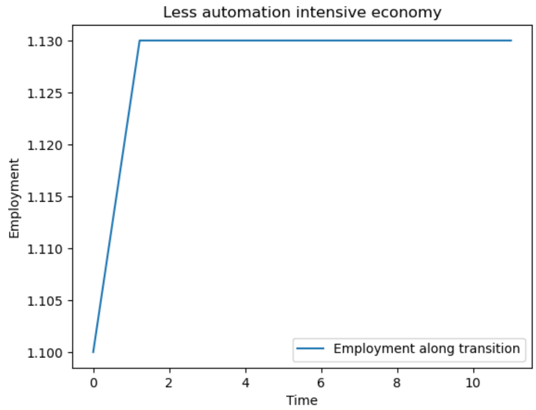

(a) Low θ economy

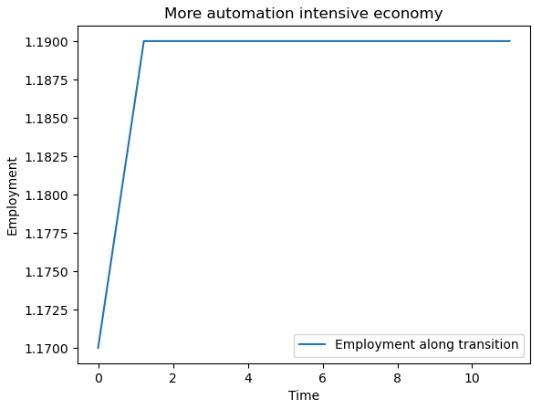

(b) High θ economy

Figure 7: **Employment recovery following a transitory negative TFP shock when firm heterogeneity is effectively removed**. In (a), the employment recovery in a less automation intensive economy is similar to the corresponding recovery in a more automation intensive economy depicted in (b).

tive TFP shock in more and less automation intensive economies, as seen in figure 6. This is because a higher elasticity of substitution raises firm demand for automation capital, which reinforces the demand for labour along the transition path. When automation capital is above a certain threshold, the demand for labour rises sufficiently to aid the recovery of employment regardless of the automation intensity of an economy.

Firm heterogeneity also plays an important role in driving the employment recovery. We consider the role of heterogeneity by lowering the variance of the productivity shocks to 0.001, such that firms receive approximately the average level of the idiosyncratic productivity each period. We find that the employment recovery rates are independent of automation intensity (θ). This is driven by two factors. *First*, there is a greater mass of firms with a high stock of automation capital when all firms receive the average level of the idiosyncratic productivity shock. Firms steadily accumulate automation capital when there is no dispersion in the idiosyncratic productivity shock. *Second*, when there is firm heterogeneity by productivity levels, the relatively lower automation capital stock in less automation intensive economies interacts with the idiosyncratic productivity shocks that firms face so as to affect employment recovery. Firms that receive a low idiosyncratic productivity shock following the negative TFP shock would also tend to have lower automation capital, and hence benefit less from the complementarity between automation capital and labour than do firms facing similar idiosyncratic shocks in more automation intensive economies. In other words, the higher automation capital stock in the more automation intensive economy offsets the low idiosyncratic productivity shock that a firm might receive (referred to as the 'insurance channel' above); and this channel is diluted in an economy which is less automation intensive.

6.4 Relaxing the constraint on investment

The model in section 5 imposed a constraint (equation (7)) on the investment choices of firms, which prevented firms from investing an amount that would exceeded their revenues net of their wage bills, which we have referred to as their 'cash in hand'. Hence, firms in the baseline model might not be able to choose their optimal level of investment if that investment choice exceeds their cash in hand. This is, therefore, a stringent version of a financing constraint on investment.

We now show that relaxing this constraint, by allowing firms to choose investment that allows them to move closer to their optimal choices results in faster employment recoveries even in less automation intensive economies. The intuition behind this is straightforward: by permitting greater automation capital accumulation in an economy, the looser constraint facilitates the channels leading to a faster employment recovery that we discussed in section 6.3.1.

While we do not model a borrowing constraint explicitly (see e.g. Gomes (2001), Albuquerque and Hopenhayn (2004) and Clementi and Hopenhayn (2006)), we now permit firms to choose investment that could exceed their cash in hand, but which cannot result in a negative firm value. We denote by $\bar{k}'(y,k)$ the upper bound on future automation capital choice, defined implicitly by:

$$0 = \max_{l} \Big[F(k,l;y,z,\theta) - wl \Big] - \big(\bar{k}'(y,k) - (1-\delta)k \big) + \beta \mathbb{E}_{y'|y}\, J(y', \bar{k}'(y,k)) \tag{18}$$

Firms would not wish to choose a level of investment resulting in negative firm value, as they would rather choose not to produce in that case[12].

When firms are less constrained in their choice of automation capital investment, the average level of automation capital should be higher regardless of the automation capital share (θ). This intuition is confirmed in the stationary equilibria of both less and more automation intensive economies, using the same parameter values as in the baseline model. We then simulate a transitory recession by lowering the value of aggregate TFP from its benchmark value

[12]This resembles the margin of firm entry and exit in models of firm dynamics, e.g. Gomes (2001).

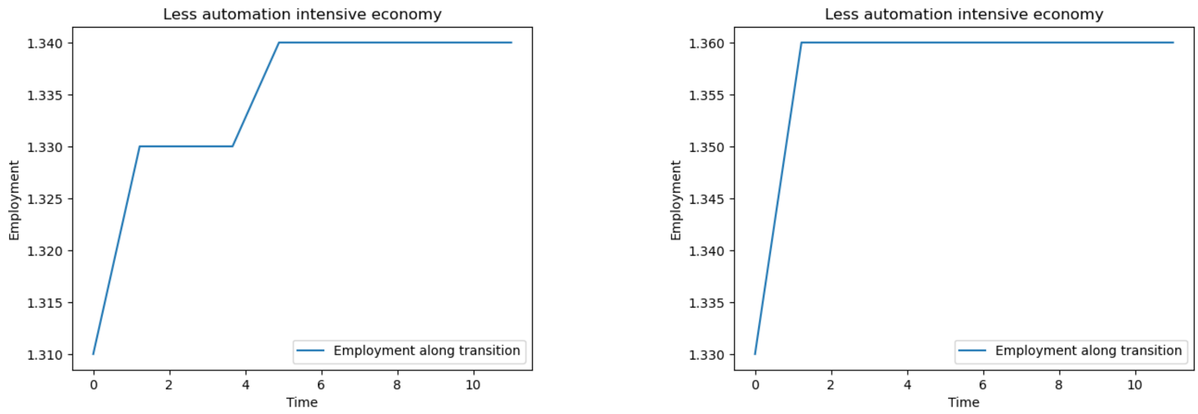

(a) Model with tighter financing constraint

(b) Model with looser financing constraint

Figure 8: **Employment recovery in a less automation intensive economy following a transitory negative TFP shock: financing constraints**. In the baseline model (a), the employment recovery is slower than the corresponding recovery in the economy with looser financing constraints depicted in (b).

of 1 in stationary equilibrium to 0.97 for a single period, as in section 6.3.

Figure 8 depicts the employment recovery along the transition path to stationary equilibrium in the less automation intensive economy[13]. We observe that the employment recovery is slightly quicker (in 1 as opposed to 5 periods) when firm investment choices are less constrained, as they can now accumulate a level of automation capital that is closer to its optimal level. The resulting greater level of automation capital in turn promotes the recovery of employment following the recession, due to the complementarity and 'insurance' channels discussed in section 6.3.1.

An implication is that policies that seek to loosen constraints on firm investment, either through direct assistance or indirectly, by building the supporting financial infrastructure, would enable productive but constrained firms to expand, reduce misallocation (see e.g. H. A. Hopenhayn (2014)) and also facilitate faster employment recoveries from negative aggregate shocks.

7 Conclusion

We document a novel empirical finding: U.S. regions more exposed to automation before the Great Recession experienced faster employment recoveries afterward. Using an instrumental variable approach and detailed commuting zone (CZ) covariates, we show that higher pre-crisis automation is associated with stronger post-crisis employment rebounds.

[13] The employment recovery rates in more automation intensive economies with varying degrees of financing constraints are broadly similar.

We show that this finding can be understood using a simple model of firm heterogeneity. Comparing economies with different degrees of automation intensity, we find that less automation-intensive economies recover slower relative to more automation-intensive economies following a transitory negative TFP shock, as they have a relatively lower level of automation capital and tend to reduce automation capital investment more after the shock. Consequently, they benefit to a lesser degree from the complementarity between automation capital and labour, resulting in a slower recovery of employment. Finally, we find that loosening the constraints on firm investment choices speeds up employment recovery even in less automation-intensive economies, which points to another beneficial consequence of policies aimed at relaxing firm financing constraints.

Our model could be extended along certain dimensions that have been considered in the recent literature on automation. For instance, one could introduce skill heterogeneity alongside task heterogeneity, wherein low and high-skilled workers perform different tasks that might be more or less substitutable with automated capital. It would also be interesting to endogenise firms' choices of automation intensities in a model of business cycles, in order to understand when firms are likely to choose to automate their production processes, and how that choice varies with the business cycle.

References

Acemoglu, D., & Restrepo, P. (2018a). Modeling automation. In *Aea papers and proceedings* (Vol. 108, pp. 48–53).

Acemoglu, D., & Restrepo, P. (2018b). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American economic review*, *108*(6), 1488–1542.

Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from us labor markets. *Journal of political economy*, *128*(6), 2188–2244.

Acemoglu, D., & Restrepo, P. (2022). Demographics and automation. *The Review of Economic Studies*, *89*(1), 1–44.

Albuquerque, R., & Hopenhayn, H. A. (2004). Optimal lending contracts and firm dynamics. *The Review of Economic Studies*, *71*(2), 285–315.

Arntz, M., Böhm, M., Graetz, G., Gregory, T., Lehmer, F., & Lipowski, C. (2024). *Firm-level technology adoption in times of crisis* (Tech. Rep.). ZEW Discussion Papers.

Autor, D., Dorn, D., & Hanson, G. (2019). Replication data for: When Work Disappears: Manufacturing Decline and the Falling Marriage Market Value of Young Men.

Autor, D., & Salomons, A. (2018). *Is automation labor-displacing? productivity growth, employment, and the labor share* (Tech. Rep.). National Bureau of Economic Research.

Autor, D. H., & Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review*, *103*(5), 1553–1597.

Autor, D. H., Dorn, D., & Hanson, G. H. (2013). The china syndrome: Local labor market effects of import competition in the united states. *American economic review*, *103*(6), 2121–2168.

Barnichon, R., & Figura, A. (2015). Labor market heterogeneity and the aggregate matching function. *American Economic Journal: Macroeconomics*, *7*(4), 222–249.

Benhabib, J., Perla, J., & Tonetti, C. (2014). Catch-up and fall-back through innovation and imitation. *Journal of Economic Growth*, *19*, 1–35.

Berg, A., Buffie, E. F., Comunale, M., Papageorgiou, C., & Zanna, L.-F. (2025). Searching for wage growth: Policy responses to the "new machine age". *Review of Economic Dynamics*, 101286.

Bloom, N. (2007). Uncertainty and the dynamics of r&d. *American Economic Review*, *97*(2), 250–255.

Board, T. C. (2023). EU KLEMS July 2018 Release. DataverseNL. Retrieved from https://doi.org/10.34894/6GDD7Q doi: 10.34894/6GDD7Q

Boppart, T., Krusell, P., & Mitman, K. (2018). Exploiting mit shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. *Journal of Economic Dynamics and Control*, *89*, 68–92.

Borusyak, K., Hull, P., & Jaravel, X. (2022). Quasi-experimental shift-share research designs. *The Review of economic studies*, *89*(1), 181–213.

Cerra, V., Fatás, A., & Saxena, S. C. (2023). Hysteresis and business cycles. *Journal of Economic Literature*, *61*(1), 181–225.

Chetty, R., Guren, A., Manoli, D., & Weber, A. (2011). Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. *American Economic Review*, *101*(3), 471–475.

Chopra, A., & Mukherji, R. (2024). Automation and local labour markets: Impact of immigrant mobility.

Clementi, G. L., & Hopenhayn, H. A. (2006). A theory of financing constraints and firm dynamics. *The Quarterly Journal of Economics, 121*(1), 229–265.

Clementi, G. L., & Palazzo, B. (2016). Entry, exit, firm dynamics, and aggregate fluctuations. *American Economic Journal: Macroeconomics, 8*(3), 1–41.

Cooley, T. F., & Quadrini, V. (2001). Financial markets and firm dynamics. *American economic review, 91*(5), 1286–1310.

Dauth, W., Findeisen, S., Suedekum, J., & Woessner, N. (2021). The adjustment of labor markets to robots. *Journal of the European Economic Association, 19*(6), 3104–3153.

DeCanio, S. J. (2016). Robots and humans–complements or substitutes? *Journal of Macroeconomics, 49*, 280–291.

Eden, M., & Gaggl, P. (2018). On the welfare implications of automation. *Review of Economic Dynamics, 29*, 15–43.

Elsby, M. W., Hobijn, B., Şahin, A., Valletta, R. G., Stevenson, B., & Langan, A. (2011). The labor market in the great recession—an update to september 2011 [with comment and discussion]. *Brookings Papers on Economic Activity*, 353–384.

Faber, M., Sarto, A. P., & Tabellini, M. (2022). *Local shocks and internal migration: The disparate effects of robots and chinese imports in the us* (Tech. Rep.). National Bureau of Economic Research.

Fernald, J., Stock, J., Hall, R., & Watson, M. (2017). The disappointing recovery of output after 2009. *Brookings Papers on Economic Activity*.

Firooz, H., Liu, Z., & Wang, Y. (2025). Automation and the rise of superstar firms. *Journal of Monetary Economics*, 103733.

Gomes, J. F. (2001). Financing investment. *American Economic Review, 91*(5), 1263–1285.

Graetz, G., & Michaels, G. (2017). Is modern technology responsible for jobless recoveries? *American Economic Review, 107*(5), 168–173.

Graetz, G., & Michaels, G. (2018). Robots at work. *Review of economics and statistics, 100*(5), 753–768.

Greenwood, J., Hercowitz, Z., & Huffman, G. W. (1988). Investment, capacity utilization, and the real business cycle. *The American Economic Review*, 402–417.

Hansen, G. D. (1985). Indivisible labor and the business cycle. *Journal of monetary Economics, 16*(3), 309–327.

Hémous, D., & Olsen, M. (2022). The rise of the machines: Automation, horizontal innovation, and income inequality. *American Economic Journal: Macroeconomics*, *14*(1), 179–223.

Hershbein, B., & Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, *108*(7), 1737–1772.

Hershbein, B., & Stuart, B. A. (2024). The evolution of local labor markets after recessions. *American Economic Journal: Applied Economics*, *16*(3), 399–435.

Hopenhayn, H., & Rogerson, R. (1993). Job turnover and policy evaluation: A general equilibrium analysis. *Journal of political Economy*, *101*(5), 915–938.

Hopenhayn, H. A. (1992). Entry, exit, and firm dynamics in long run equilibrium. *Econometrica: Journal of the Econometric Society*, 1127–1150.

Hopenhayn, H. A. (2014). Firms, misallocation, and aggregate productivity: A review. *Annu. Rev. Econ.*, *6*(1), 735–770.

Hubmer, J., & Restrepo, P. (2021). *Not a typical firm: The joint dynamics of firms, labor shares, and capital–labor substitution* (Tech. Rep.). National Bureau of Economic Research.

Jaimovich, N., Saporta-Eksten, I., Siu, H., & Yedid-Levi, Y. (2021). The macroeconomics of automation: Data, theory, and policy analysis. *Journal of Monetary Economics*, *122*, 1–16.

Jaimovich, N., & Siu, H. E. (2020). Job polarization and jobless recoveries. *Review of Economics and Statistics*, *102*(1), 129–147.

Khan, A., & Thomas, J. K. (2008). Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. *Econometrica*, *76*(2), 395–436.

Kopytov, A., Roussanov, N., & Taschereau-Dumouchel, M. (2018). Short-run pain, long-run gain? recessions and technological transformation. *Journal of Monetary Economics*, *97*, 29–44.

Kroft, K., Lange, F., Notowidigdo, M. J., & Katz, L. F. (2016). Long-term unemployment and the great recession: the role of composition, duration dependence, and nonparticipation. *Journal of Labor Economics*, *34*(S1), S7–S54.

Lewis, E. (2011). Immigration, skill mix, and capital skill complementarity. *The Quarterly Journal of Economics*, *126*(2), 1029–1069.

Ruggles, S., Flood, S., Sobek, M., Brockman, D., Cooper, G., Richards, S., & Schouweiler, M. (2023). Ipums usa: Version 13.0 [dataset]. *Minneapolis, MN: IPUMS*. doi: https://doi.org/10.18128/D010.V13.0

Tolbert, C. M., & Sizer, M. (1996). *Us commuting zones and labor market areas: A 1990 update*

(Tech. Rep.). ERS Staff Paper 9614, Rural Economic Division, US Department of Agriculture.

Yagan, D. (2019). Employment hysteresis from the great recession. *Journal of Political Economy, 127*(5), 2505–2558.

Appendix

A Automation: data, trends and regional dispersion

A.1 IFR Robot data

The IFR has collected data on the stock of industrial robots at the country-industry level since 1993. Industrial robots are defined as an "automatically controlled, reprogrammable multipurpose manipulator, programmable in three or more axes, which can be either fixed in place or fixed to a mobile platform for use in automation applications in an industrial environment (ISO 8373:2021)."[14]

The IFR data has a few limitations. Industry-specific data is available for North America from 2004 onwards. For the years before 2004, we classify the data into industries using the distribution from 2010. Not all data can be categorised by sectors; for example, around 11% of total robots remained unclassified in 2015. We allocated these unclassified robots proportionally to the classified data. Additionally, the stock of robots for the US includes data from Canada and Mexico before 2011. To ensure consistency, we use the data for North America. This is not problematic as our instrumental variable strategy will account for any measurement errors.

A.2 Robots per thousand workers in industrialised economies

Figure 9: Robots per thousand workers in US and selected countries

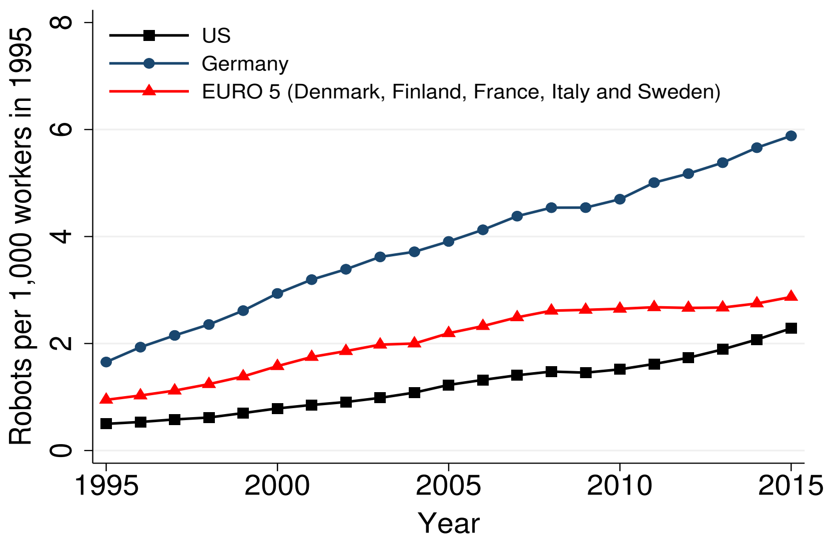

[14]The definition can be viewed at the IFR website https://ifr.org/industrial-robots.

Figure 9 shows the trend of robots per thousand workers in North America, Germany, and EURO5 (Denmark, Finland, France, Italy and Sweden) countries. The average growth in robot adoption for the EURO5 countries is a simple average across all the countries. The number of industrial robots per thousand workers has steadily increased in all the aforementioned countries. In North America, the stock of robots increased from 0.5 per thousand workers in 1995 to 2.28 per thousand workers in 2015.

A.3 Robot per thousand workers by industry in US

Table 5 shows that the automotive industry experienced the strongest growth in North America between 1993 and 2015, while the service industry saw the smallest increase in robot usage.

Table 5: Robot per thousand workers by industry

	Robot per $1,000$ workers in 1990		
Industry	1993	2015	Difference
All Industries	0.404	2.424	2.02
Automotive	11.033	65.117	54.083
Metal products	1.777	6.411	4.633
Plastics and chemicals	3.298	17.757	14.459
Electronics	2.611	14.869	12.259
Food and beverages	1.227	6.678	5.451
Textiles	0.003	0.062	0.06
Wood and furniture	0.009	0.294	0.285
Paper and printing	0.002	0.131	0.129
Minerals	0.028	0.342	0.314
Basic metals	0.046	11.123	11.078
Industrial machinery	0.052	2.317	2.265
Shipbuilding and aerospace	0.047	0.815	0.768
Manufacturing Miscellaneous	0.387	9.825	9.437
Agriculture	0.004	0.074	0.07
Mining	0.001	0.056	0.054
Utilities	0	0.085	0.085
Construction	0.004	0.027	0.023
Education and Research	0.008	0.105	0.098
Services	0	0.005	0.004

A.4 Exposure at local labour market level

Following Acemoglu and Restrepo (2020), the growth in the stock of industrial robots in industry j over time is expressed as follows:

$$\Delta R_{j,t} = \frac{R_{j,t_1} - (1 + g_{j,(t,t_1)}) \cdot R_{j,t}}{L_{j,t}} \tag{19}$$

where $R_{j,t}$ is the number of robots in industry j at year t, $L_{j,t}$ is the employment count (in thousands) in industry j in year t and $g_{j,(t,t_1)}$ is the rate of growth of output over the period from t to t_1 in industry j. t_1 is 2000 and 2015 when t equals 1990 and 2000, respectively. Equation (19) captures the additional acquisition of robot capital while considering the overall growth of the industry and keeping employment fixed at year t. Similarly, the EURO5 industry-level robot growth is calculated as:

$$\Delta R_{j,t}^{EURO5} = \frac{1}{5} \sum_{c} \frac{R_{j,t_1}^{c} - (1 + g_{j,(t,t_1)}^{c}) \cdot R_{j,t}^{c}}{L_{j,t}^{c}} \tag{20}$$

where $R_{j,t}^{c}$ is the stock of robots in country c and industry j at year t, $g_{j,(t,t_1)}^{c}$ is the growth rate of output in country c and industry j between time t and t_1, and $L_{j,t}^{c}$ denotes the number of employed workers in country c and industry j at time t.

B Robustness Checks

Table 6: 2SLS: Instrument = ΔRobot Shift–Share, 2000 Baseline

	(1)	(2)	(3)	(4)	(5)
Shock (β_1)	-0.53^{***}	-0.46^{***}	-0.46^{***}	-0.46^{***}	-0.48^{***}
	(0.10)	(0.10)	(0.10)	(0.11)	(0.11)
Δ Robots (β_2)	-2.91^{***}	-0.77	-0.75	-0.84	-0.78
	(0.92)	(0.90)	(1.00)	(1.02)	(0.99)
Shock × ΔRobots (β_3)	0.46^{***}	0.23^{**}	0.23^{**}	0.22^{*}	0.21^{*}
	(0.13)	(0.11)	(0.11)	(0.11)	(0.11)
Observations	722	722	722	722	722
R^2	0.25	0.46	0.47	0.49	0.49
Census divisions	Yes	Yes	Yes	Yes	Yes
Demographics		Yes	Yes	Yes	Yes
Industry shares			Yes	Yes	Yes
Trade, routine jobs				Yes	Yes
Offshore index					Yes

Notes: 2SLS with 2000-baseline controls. Instrument as in previous table. Shock–ΔRobots interaction remains significant at 5–10% across full specifications despite inclusion of additional controls (offshore index not applicable to 2000 baseline).

Table 7: 2SLS: Instrument = $\ln$(Robot Stock 2000), 1990 Baseline

	(1)	(2)	(3)	(4)	(5)
Shock (β_1)	-0.53^{***}	-0.47^{***}	-0.44^{***}	-0.44^{***}	-0.45^{***}
	(0.09)	(0.09)	(0.09)	(0.09)	(0.10)
Δ Robots (β_2)	-2.44^{**}	-1.41^{*}	-1.32^{*}	-1.32^{*}	-1.38^{*}
	(0.99)	(0.74)	(0.77)	(0.77)	(0.78)
Shock × ΔRobots (β_3)	0.41^{***}	0.30^{***}	0.28^{***}	0.28^{***}	0.29^{***}
	(0.13)	(0.09)	(0.09)	(0.09)	(0.09)
Observations	722	722	722	722	722
R^2	0.26	0.45	0.46	0.46	0.47
Census divisions	Yes	Yes	Yes	Yes	Yes
Demographics		Yes	Yes	Yes	Yes
Industry shares			Yes	Yes	Yes
Trade, routine jobs				Yes	Yes
Offshore index					Yes

Notes: 2SLS regression instrumenting ΔRobots with $\ln(\text{robots}_{2000})$. The robot stock instrument yields a somewhat smaller estimated impact of ΔRobots than the shift–share instrument (see Table above), but the interaction effect remains positive and significant. Standard errors in parentheses; ***/**/* p<0.01/0.05/0.1.

C Numerical sensitivity analyses

We first show that the main numerical result, regarding the relatively faster employment recovery in a more automation intensive economy following a negative TFP shock, is robust to considering larger and more persistent negative TFP shocks. We then consider the sensitivity of our main result to choices of certain parameters, particularly the automation share parame-

Table 8: 2SLS: Instrument = $\ln$(Robot Stock 2000), 2000 Baseline

	(1)	(2)	(3)	(4)	(5)
Shock (β_1)	-0.53^{***}	-0.48^{***}	-0.49^{***}	-0.48^{***}	-0.52^{***}
	(0.09)	(0.09)	(0.10)	(0.10)	(0.11)
Δ Robots (β_2)	-2.44^{**}	-1.37^{*}	-1.95^{**}	-1.68^{**}	-1.81^{**}
	(0.99)	(0.76)	(0.84)	(0.80)	(0.80)
Shock × ΔRobots (β_3)	0.41^{***}	0.29^{***}	0.33^{***}	0.29^{***}	0.31^{***}
	(0.13)	(0.10)	(0.11)	(0.10)	(0.10)
Observations	722	722	722	722	722
R^2	0.26	0.45	0.46	0.48	0.49
Census divisions	Yes	Yes	Yes	Yes	Yes
Demographics		Yes	Yes	Yes	Yes
Industry shares			Yes	Yes	Yes
Trade, routine jobs				Yes	Yes
Offshore index					Yes

Notes: 2SLS with 2000-baseline controls. Robot stock (2000) used as instrument for ΔRobots. Results are broadly similar to the shift–share IV (Tables above), with a significant negative shock effect and positive interaction term in all specifications.

ter (θ), the Frisch elasticity parameter (ψ) and the persistence of the idiosyncratic productivity shock.

C.1 Varying the magnitude and persistence of the negative TFP shock

Figure 10 shows that employment recovery continues to be faster in more as opposed to less automation intensive economies when one considers larger negative TFP shocks. Specifically, we consider an unanticipated reduction in TFP (z) from its baseline value of 1 to 0.9 for one period. In contrast, figure 5 in the paper considered a one period reduction in z to a value of 0.97. We note that the more automation intensive economy does take considerably longer to reach steady state employment as compared to the scenario depicted in figure 5, owing to the larger disruption faced by firms. Hence, the main numerical result is robust to varying the size of the negative TFP shock.

It is also robust to considering more persistent TFP shocks. Figure 11 plots an unanticipated negative TFP shock $\{z_t\}$ with $z_0 = 0.95$, $z_1 = 0.98$ and $z_s = 1\ \forall s > 1$. Again, the recovery of employment is faster in more as opposed to less automation intensive economies. We note again that the recovery is more protracted in more automation intensive economies compared to the scenario depicted in figure 5, as firms face disruptive shocks that dissipate gradually.

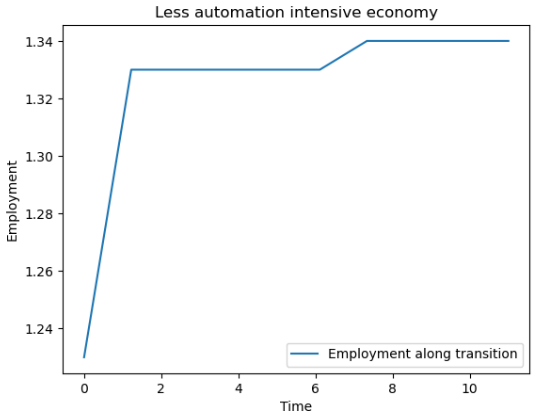

(a) Low θ economy

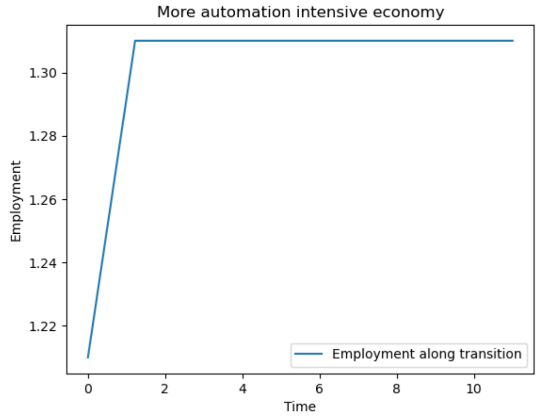

(b) High θ economy

Figure 10: **Employment recovery following a larger transitory negative TFP shock**. In (a), the employment recovery in a less automation intensive economy is slower than the corresponding recovery in a more automation intensive economy depicted in (b).

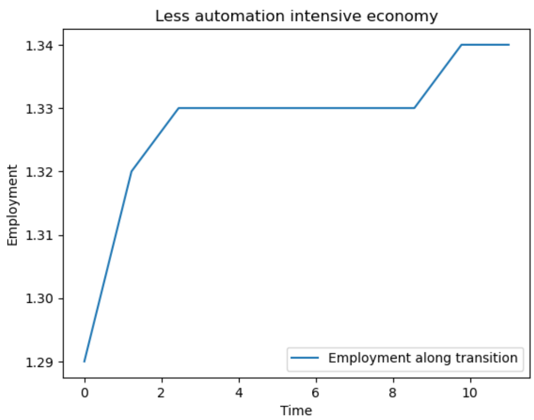

(a) Low θ economy

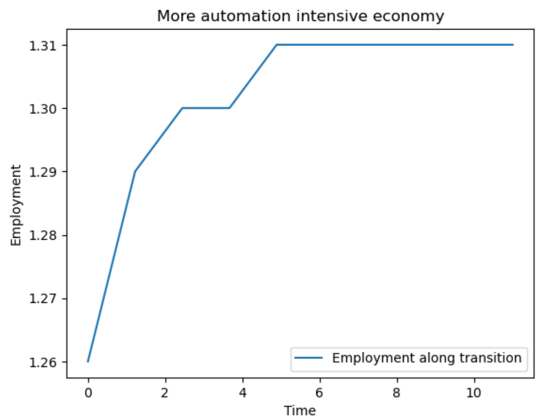

(b) High θ economy

Figure 11: **Employment recovery following a persistent negative TFP shock**. In (a), the employment recovery in a less automation intensive economy is slower than the corresponding recovery in a more automation intensive economy depicted in (b).

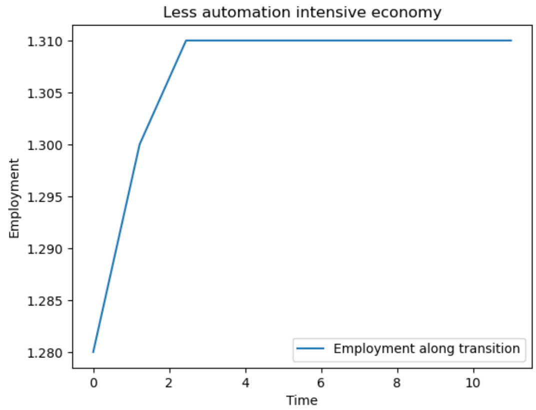

(a) Low θ economy

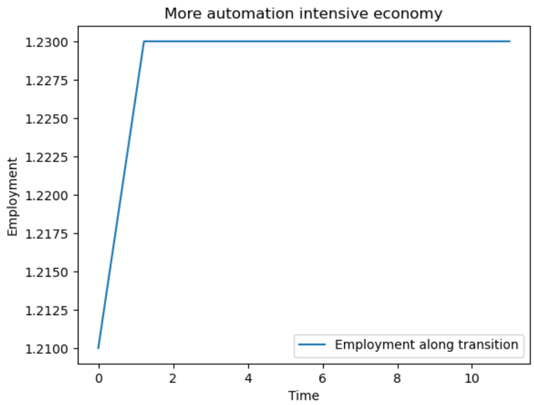

(b) High θ economy

Figure 12: **Employment recovery following a transitory negative TFP shock: robustness to higher automation shares**. In (a), the employment recovery in a less automation intensive economy is slightly slower than the corresponding recovery in a more automation intensive economy depicted in (b).

C.2 Robustness to varying the automation share parameter

We consider alternative choices for the automation share parameter in low and high automation intensity economies. Recall that the parameter values used for low and high automation intensity economies in figure 5 are $\{\theta_l, \theta_h\} = \{0.3, 0.4\}$. Figure 12 is plotted using $\{\theta_l, \theta_h\} = \{0.4, 0.5\}$. While the higher automation capital speeds up employment recovery in both economies, we note that the recovery remains faster in the more automation intensive economy (in a single period), although the recoveries in both economies are greatly accelerated.

Alternatively, figure 13 is plotted using lower values of automation shares in both economies $\{\theta_l, \theta_h\} = \{0.25, 0.35\}$ relative to the benchmark. We see that employment recoveries in low and high automation intensity economies follow the pattern first displayed in figure 5, with the recovery being even more protracted in the less automation intensive economy. As discussed in the paper, this pattern is driven by lower automation capital in the less automation intensive economy.

C.3 Robustness to varying the Frisch elasticity parameter

As noted in the paper, Chetty et al. (2011) suggest that Macro models should aim to match a Frisch elasticity of 0.75, implying a value of $\psi = 1.33$. Figure 14 compares employment recoveries in economies that are less or more automation intensive, using $\{\theta_l, \theta_h\} = \{0.3, 0.4\}$, but with a lower Frisch elasticity of 0.75. We observe that the pattern documented in the paper continues to hold: employment recoveries are relatively slower in less automation intensive

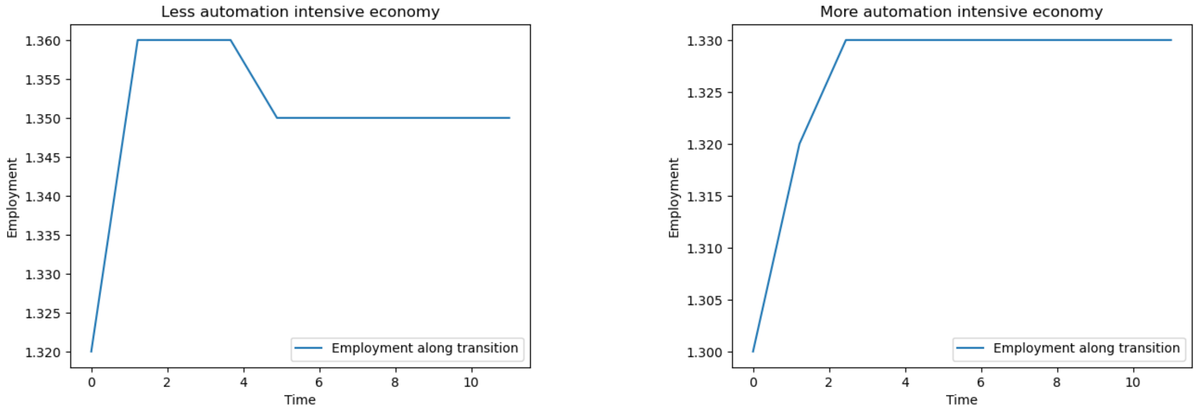

(a) Low θ economy

(b) High θ economy

Figure 13: **Employment recovery following a transitory negative TFP shock: robustness to lower automation shares**. In (a), the employment recovery in a less automation intensive economy is much slower than the corresponding recovery in a more automation intensive economy depicted in (b).

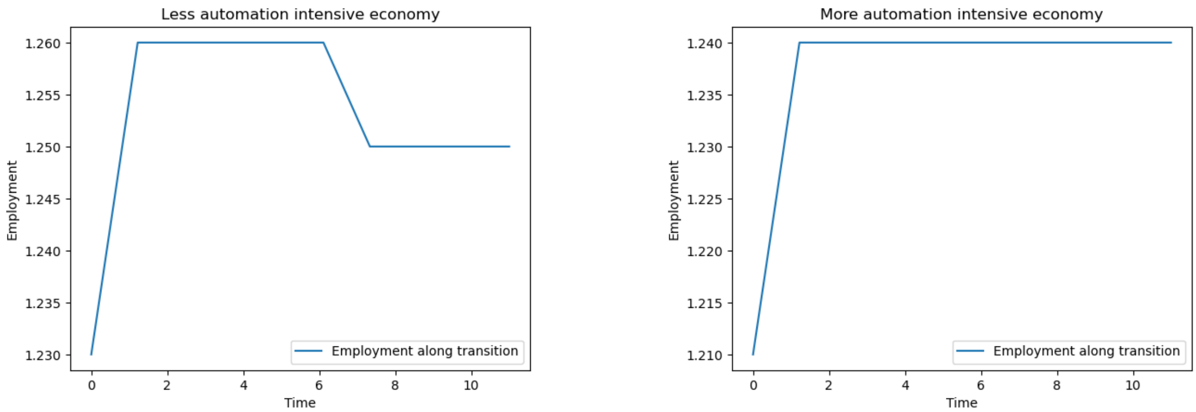

(a) Low θ economy

(b) High θ economy

Figure 14: **Employment recovery following a transitory negative TFP shock: robustness to lower Frisch elasticity**. In (a), the employment recovery in a less automation intensive economy is slower than the corresponding recovery in a more automation intensive economy depicted in (b).

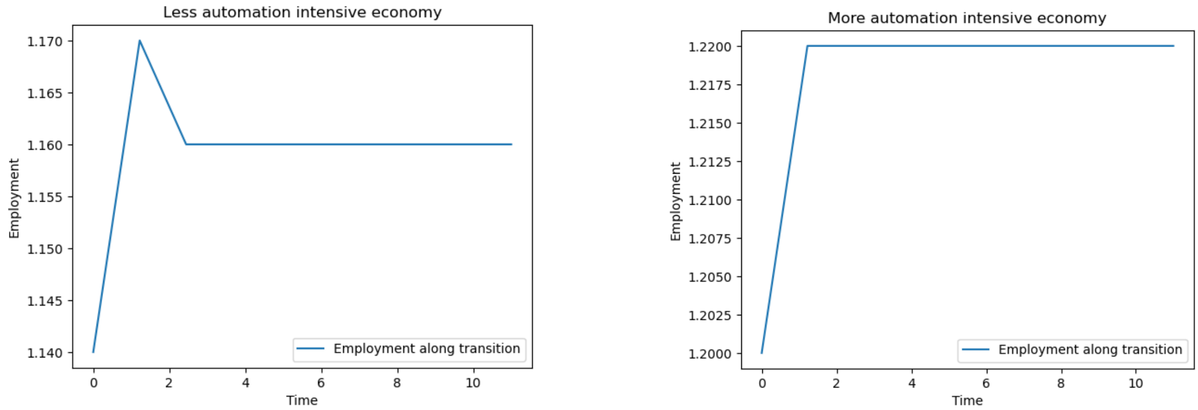

(a) Low θ economy

(b) High θ economy

Figure 15: **Employment recovery following a transitory negative TFP shock: robustness to lower shock persistence**. In (a), the employment recovery in a less automation intensive economy is slower than the corresponding recovery in a more automation intensive economy depicted in (b).

economies. Wages are higher relative to the benchmark calibration used in the paper when the Frisch elasticity is lower, which dampens the complementarity channel associated with greater automation capital in more automation intensive economies, although it doesn't affect the employment recovery rates seen in panel (b) of figure 14.

C.4 Robustness to varying the persistence of the idiosyncratic shock

Figure 15 compares employment recoveries in economies that are less or more automation intensive, using $\{\theta_l, \theta_h\} = \{0.3, 0.4\}$, but with a lower shock persistence of 0.5. Once again, the pattern documented in the paper continues to hold: employment recoveries are relatively slower in less automation intensive economies.

D Incorporating physical capital

In this section, we add physical capital as an input that firms can accumulate. As noted in section 5.3, this allows us to generalise our results to an environment where firms have an additional instrument available for insurance against idiosyncratic productivity shocks. In addition, Acemoglu and Restrepo (2020) find some evidence for a positive correlation between industries that witnessed greater automation and industries that witnessed rapid growth in total capital, although they also note that their findings about the negative effect of automation on employment and wages are not driven by physical capital.

We assume a firm's production function is Cobb-Douglas in physical capital (k) and a CES composite of robots (k_a) and labour (l). We assume that the share parameter for physical capital is ζ, while the elasticity of substitution between robots and labour is ϵ. The production function is:

$$F(k, k_a, l; y, z, \theta) = zy\, k^{\zeta}\Big(\theta k_a^{\epsilon} + (1-\theta)l^{\epsilon}\Big)^{\frac{1-\zeta}{\epsilon}} \tag{21}$$

A firm with state (y, k, k_a) chooses labour (l), physical capital (k') and automation capital for the following period, (k_a'), in order to maximize its value:

$$J(y, k, k_a) = \max_{k', k_a', l} F(k, k_a, l; y, z, \theta) - wl - \big(k' - (1-\delta)k\big) - \big(k_a' - (1-\delta)k_a\big) + \beta \mathbb{E}_{y'|y}\, J(y', k', k_a') \tag{22}$$

subject to the following constraints:

$$0 \leq k'(y, k, k_a) \leq F\big(k, k_a, l(y, k, k_a); y, z, \theta\big) - wl(y, k, k_a) + (1-\delta)k + (1-\delta)k_a - k_a'(y, k, k_a)$$
$$0 \leq k_a'(y, k, k_a) \leq F\big(k, k_a, l(y, k, k_a); y, z, \theta\big) - wl(y, k, k_a) + (1-\delta)k + (1-\delta)k_a - k'(y, k, k_a)$$

The constraints restrict firms from investing an amount that results in negative profits in any period. In addition, firms cannot have negative capital of either type, implying that firms can reduce their physical or automation capital stock subject to it remaining non-negative. Consequently, the policy function for labour, $l(y, k, k_a)$, follows from a static constrained profit maximization problem. The unconstrained value of next period's physical capital, $k^u(y, k, k_a)$, is obtained from the Euler equation:

$$1 = \beta \mathbb{E}_{y'|y}\left[F_k\big(k', k_a', l(y', k', k_a'); y', z, \theta\big) + (1-\delta)\right] \tag{23}$$

Similarly, the unconstrained value of next period's automation capital, $k_a^u(y, k, k_a)$, is obtained from the Euler equation:

$$1 = \beta \mathbb{E}_{y'|y} \left[F_{k_a}\big(k', k_a', l(y', k', k_a'); y', z, \theta\big) + (1-\delta) \right] \tag{24}$$

Incorporating the constraints, the policy function for physical capital next period is:

$$\begin{aligned} k'(y,k,k_a) = \max\Big\{0, \min\{k^u(y,k,k_a), F\big(k, k_a, l(y,k,k_a); y, z, \theta\big) - wl(y,k,k_a) \\ +(1-\delta)k + (1-\delta)k_a - k_a'(y,k,k_a)\}\Big\} \end{aligned} \tag{25}$$

Similarly, the policy function for automation capital next period is:

$$\begin{aligned} k_a'(y,k,k_a) = \max\Big\{0, \min\{k_a^u(y,k,k_a), F\big(k, k_a, l(y,k,k_a); y, z, \theta\big) - wl(y,k,k_a) \\ +(1-\delta)k + (1-\delta)k_a - k'(y,k,k_a)\}\Big\} \end{aligned} \tag{26}$$

The policy functions in turn yield firm output and profits:

$$\begin{aligned} Y(y,k,k_a) &= F\big(k, k_a, l(y,k,k_a); y, z, \theta\big) \\ \pi(y,k,k_a) &= F\big(k, k_a, l(y,k,k_a); y, z, \theta\big) - wl(y,k,k_a) - \big(k'(y,k,k_a) - (1-\delta)k\big) \\ &\quad +\big(k_a'(y,k,k_a) - (1-\delta)k_a\big) \end{aligned} \tag{27}$$

The endogenous distribution of firms is denoted by $\lambda(y,k,k_a)$. One can then use the distribution and policy functions in order to obtain aggregate labour demanded by firms, the aggregate profits rebated to consumers and aggregate automation capital:

$$\begin{aligned} L^d &= \int l(y,k,k_a)\ d\lambda(y,k,k_a) \\ \Pi &= \int \pi(y,k,k_a)\ d\lambda(y,k,k_a) \\ K &= \int k\ d\lambda(y,k,k_a) \\ Ka &= \int k_a\ d\lambda(y,k,k_a) \end{aligned}$$

We note that the consumer's problem is unchanged, while the definition of a stationary equilibrium is modified to incorporate the choice of physical capital, and in the equation of motion for the stationary distribution λ^*:

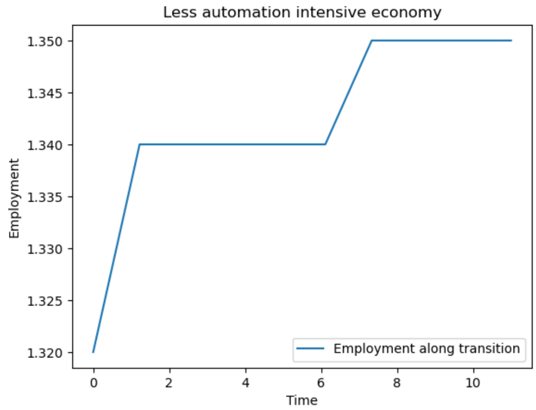

(a) Low θ economy

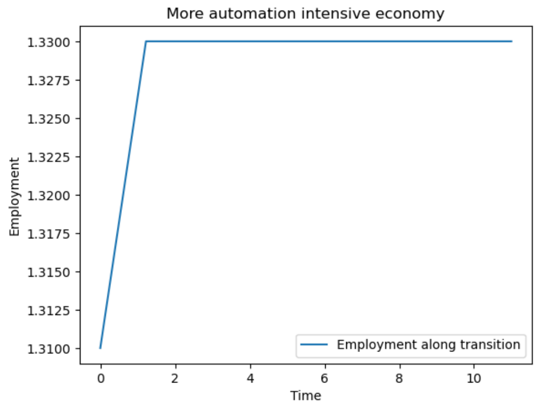

(b) High θ economy

Figure 16: **Employment recovery following a transitory negative TFP shock**. In (a), the employment recovery in a less automation intensive economy is slower than the corresponding recovery in a more automation intensive economy depicted in (b).

$$\lambda^*(y', k', k_a') = \int \mathbb{I}_{\{k'(y,k)=k'\}} \, \mathbb{I}_{\{k_a'(y,k,k_a)=k_a'\}} \, P(y'|y) \, d\lambda^*(y, k, k_a) \tag{28}$$

In the expression above, $\mathbb{I}_{\{k'(y,k,k_a)=k'\}}$ is an indicator function for when the physical capital policy function for a firm with state (y, k, k_a) equals k'; $\mathbb{I}_{\{k_a'(y,k,k_a)=k_a'\}}$ is an indicator function for when the automation capital policy function for a firm with state (y, k, k_a) equals k_a'; while $P(y'|y)$ is the Markov transition probability.

D.1 Employment recoveries following a recession

We now replicate the exercise in section 6.3 of the paper, i.e. we simulate a recession by lowering the value of aggregate TFP from its benchmark value of 1 in stationary equilibrium to 0.97 for a single period. We choose a physical capital share parameter $\zeta = 0.25$, while the elasticity of substitution between robots and labour, $\epsilon = 0.8$, in order to be consistent with Acemoglu and Restrepo (2020)'s finding that wages and employment are lower in more automated regions. All other parameter values are identical to those reported in 6 above.

We compute the transition path back to the initial stationary equilibrium of two economies or regions facing a negative TFP shock of the same magnitude but which are differentiated by their automation intensities, parameterised by θ. As above, we choose values of $\theta_l = 0.3$ and $\theta_h = 0.4$ to demonstrate our result.

Figure 16 contrasts the employment recovery following a single period negative TFP shock in economies with low and high automation intensities (panels (a) and (b) respectively). Clearly, employment attains its stationary equilibrium value faster (in 1 period) in the more automation intensive economy as compared to the less automation intensive economy, where it takes around 7 periods to reach the corresponding stationary equilibrium level of employment.

The reason for the relatively slower employment recovery in the less automated economy is, once again, that firms tend to reduce investment in capital (of either type) to a greater degree than firms in economies that are less reliant on labour when hit by a negative TFP shock. Consequently, there is a relatively greater mass of firms in the automation intensive economy with average or higher levels of both types of capital stock conditional on receiving a low draw of the idiosyncratic productivity shock. Such firms could offset the low idiosyncratic productivity if they have a high stock of physical and/or automation capital, which boosts revenue and investment. As labour is (partially) complementary to automation capital and highly complementary to physical capital, a higher capital stock raises labour demand and leads to a faster employment recovery.