ROBOT ADOPTION AND INFLATION DYNAMICS

2025

BANCO DE ESPAÑA
Eurosistema

Documentos de Trabajo
N.º 2536

Henrique S. Basso and Omar Rachedi

ROBOT ADOPTION AND INFLATION DYNAMICS (*)

Henrique S. Basso (**)

BANCO DE ESPAÑA AND CEMFI

Omar Rachedi (***)

ESADE, UNIVERSITAT RAMON LLULL

(*) We thank Martin Beraja, Saroj Bhattarai, Federico Cacciatore, Antoine Camous, Carlos Carvalho, Fabrizio Colella, Cristiano Cantore, Davide Debortoli, Harris Dellas, Federico di Pace, Luca Fornaro, Francesco Furlanetto, Pedro Gomes, Luis Guimarães, Andreas Gulyas, Marcus Hagedorn, Juan Francisco Jimeno, Chad Jones, Anastasios Karantounias, Jennifer La'O, Zheng Liu, Matthias Meier, Ricardo Nunes, Alessandro Peri, Federico Ravenna, Dominik Thaler, Antonella Trigari, Ludo Visschers, and presentation participants at the Bank of England, Birkbeck University of London, Queen Mary University of London, Trinity College Dublin, Universidad Carlos III de Madrid, Università della Svizzera Italiana, University of Kent, University of Mannheim, University of Surrey, the LACEA LAMES 2022 Annual Meeting, the Leuven Summer Event, the Bank of Finland and CEPR Joint Conference on Monetary Policy in Times of Large Shocks, the EEA-ESEM in Barcelona, the EABCN Conference on New Challenges in Monetary Economics and Macro-Finance in Mannheim, the Sailing the Macro Workshop in Siracusa, the PSE Macro Days 2023, the Dynare Conference in Valletta, the LuBraMacro in Porto, the X Amsterdam Macroeconomic Workshop in Amsterdam, the Economics Winter Workshop in Dublin, and the Annual Meeting Society for Economic Dynamics in Barcelona. Omar Rachedi acknowledges financial support from the Spanish Ministry of Science and Innovation with the project PID2023-153073NB-I00.
(**) henrique.basso@bde.es
(***) omar.rachedi@esade.edu

Documentos de Trabajo. N.º 2536

September 2025

https://doi.org/10.53479/40845

ISSN: 1579-8666 (online edition)

Abstract

Leveraging variation in robot adoption across U.S. metropolitan areas, we document that automation reduces the sensitivity of inflation to unemployment. To rationalize this finding, we build a New Keynesian model with search frictions in the labor market where robot adoption flattens the Phillips curve. The key channel is the option value of automation: the threat of automating labor tasks alters workers' effective bargaining power, muting the wage sensitivity to unemployment. We validate the relevance of this channel in the data by showing that robot adoption reduces the sensitivity of inflation to unemployment relatively more in highly unionized metropolitan areas.

Keywords: automation, Phillips curve, unemployment, bargaining power.

JEL classification: E24, E31, J31, O33.

Resumen

Aprovechando que la adopción de robots en las áreas metropolitanas de Estados Unidos muestra variación, en este trabajo documentamos que la automatización reduce la sensibilidad de la inflación al desempleo. Para explicar este resultado, construimos un modelo neokeynesiano con fricciones en el mercado laboral, donde la adopción de robots aplana la curva de Phillips. El canal clave es la opción de la empresas a automatizar: la amenaza de automatizar tareas laborales altera el poder de negociación efectivo de los trabajadores, atenuando la sensibilidad de los salarios al desempleo. Validamos la relevancia de este canal con los datos, al mostrar que la adopción de robots reduce la sensibilidad de la inflación al desempleo relativamente más en las áreas metropolitanas altamente sindicalizadas.

Palabras clave: automatización, curva de Phillips, desempleo, poder de negociación de los trabajadores.

Códigos JEL: E24, E31, J31, O33.

1 Introduction

The rise of automation has reshaped labor markets in advanced economies (Acemoglu and Restrepo, 2018, 2020; Graetz and Michaels, 2018). However, notwithstanding the key role of labor market conditions on firms' price-setting decisions, little is known on the effect of labor-displacing technologies on inflation dynamics. This paper shows empirically, theoretically, and quantitatively that automation can help explain the subdued inflation dynamics until the COVID-19 pandemic.

Specifically, automation weakens workers' bargaining power, making labor demand more elastic. As a result, it flattens the wage and price Phillips curves. We first provide empirical evidence across U.S. cities showing that automation reduces the inflation sensitivity to unemployment—especially in areas where workers have strong bargaining power. Second, we build a New Keynesian model with search frictions in the labor market and endogenous robot adoption. These two features are essential to capturing the interaction between automation and wage bargaining in shaping wage and price dynamics. In this setting, the option to automate limits workers' ability to extract rents in tight labor markets, muting the inflationary impact of low unemployment. As in the data, the effect of the automation threat on wage setting hinges on the combination of the stock of robots and workers' bargaining power. Finally, the model attributes 9% of the decline in inflation sensitivity over the last three decades to automation.

We begin by leveraging geographical variation in automation across U.S. metropolitan areas (MSAs), and document that robot adoption reduces the sensitivity of local wage and price inflation to unemployment. To do so, we build an annual panel of inflation, unemployment, and robot adoption over 2008-2018. As in Acemoglu and Restrepo (2020), we measure robots per employees combining MSA industry employment shares in the year 2000 with national industry-year robot installations. We then regress different measures of inflation on unemployment and its interaction with robot adoption. We also include year fixed effects, which control for aggregate supply shocks and common inflation expectations across areas, and absorb the response of monetary policy to aggregate demand shocks (Beraja et al., 2019; McLeay and Tenreyro, 2020; Fitzgerald et al., 2024).

This approach may fail to uncover the local Phillips curve if idiosyncratic supply shocks correlate with local labor market conditions. To address this concern, we follow Hazell et al. (2022) and instrument unemployment with local tradable demand spillovers, using a shift-share variable that weights national value-added growth in tradable industries by their city-level shares. Then, to identify the causal effect of automation, we instrument U.S. automation using robot installations in five major European economies, as in Acemoglu and Restrepo (2020).

The interaction between unemployment and robot adoption is positive, statistically significant, and economically relevant: a one standard deviation increase in automation reduces the sensitivity of price and wage inflation to unemployment by 15-17% and 10%, respectively. These magnitudes crucially depend on workers' bargaining power. Indeed, robot adoption alters relatively more inflation dynamics in areas with high union membership rates. Thus, insofar as workers' bargaining power is sufficiently high, robot adoption weakens the relationship between price inflation and unemployment via a muted wage responsiveness. In order words, the effect of automation on inflation depends on the combination of the stock of installed robots and the degree of workers' bargaining power.

To rationalize this evidence, we extend a standard New Keynesian model with robot adoption and search frictions in the labor market. The key model block consists of producers—interpretable as tasks—that decide to either use robots or post job vacancies, in the spirit of Acemoglu and Restrepo (2020), Basso and Jimeno (2021) and Beraja and Zorzi (2025). Specifically, producers draw an idiosyncratic worker efficiency and decide to use either a labor technology (labor firms) or a machine technology (robot firms). This setting implies an automation threshold, defined as the worker efficiency that equalizes the value of being a labor firm to that of robot firms. Producers with efficiencies above the threshold post job vacancies, whereas the rest install machines.

To isolate the effect of automation, we vary only its steady-state level. In our setting, final goods are transformed into machines with a linear technology, subject to exogenous robot-specific technical change. An increase in technical change reduces the steady-state price of robots, thus raising the automation threshold (and the fraction of robot firms), in line with the evidence of Graetz and Michaels (2018). While the automation cut-off varies across steady states with technical change, it may also fluctuate around the steady state in response to endogenous price movements. Following a shock, a rise in wages relative to the price of robots raises the incentives to substitute workers with machines.

When characterizing the Phillips curve, we show that higher automation in the steady state does flatten the slope only if automation can adjust upon a shock. In other words, what matters for inflation dynamics is not just the existing stock of robots, but the capacity for further adjustment: the stock matters only insofar as it can respond at the margin. Crucially, the magnitude of this effect is modulated by workers' bargaining power, in line with our evidence: higher bargaining power amplifies the effect of robot installations on inflation dynamics. Intuitively, automation alters workers' expected payout, as firms could replace them if wages rise substantially. Thus, the outside option of automation—which gains traction as the stock of robots increases—alters workers' effective bargaining power in wage setting, leading to a flat Phillips curve.

Next, we compare the quantitative implication of our model with those estimated in the data. This analysis builds on a key premise: Hazell et al. (2022)

demonstrate there is a direct mapping between the slopes of the local and aggregate Phillips curves insofar as one abstracts from wealth effects on labor supply. This restriction circumvents the missing intercept problem, allowing the use of our model as a laboratory to interpret and dissect the empirical evidence.

To run this exercise, we simulate a transition in which robot adoption evolves to match the change observed in our data, and study its impact on upon preference shocks.[1] We then regress price inflation on unemployment and its interaction with automation, replicating our empirical specification at the aggregate level. Strikingly, the model delivers estimates that closely match those in the data, giving credence to the joint role of the stock of installed robots and workers' bargaining power in shaping the effect of automation on inflation dynamics. The model implies that a surge in automation like that of the past 30 years reduces the inflation sensitivity to unemployment by 9%. For context, Inoue et al. (2025) estimate a 68% decline over the past 50 years, suggesting that automation accounts for a non-negligible part of the recent changes in inflation dynamics.

Our work builds on the literature suggestive of a flat Phillips curve post 1980s (Blanchard, 2016; Del Negro et al., 2020; Stock and Watson, 2020). This may be due to policy improvements and expectations anchoring (Ball and Mazumder, 2011; McLeay and Tenreyro, 2020; Hazell et al., 2022; Bergholt et al., 2024), labor market changes (Stansbury and Summers, 2020; Lombardi et al., 2023; Siena and Zago, 2024; Faccini and Melosi, 2025), globalization (Forbes, 2019; Heise et al., 2022), changes in shock composition (Coibion and Gorodnichenko, 2015), changes in firm inter-linkages (Galesi and Rachedi, 2019; Rubbo, 2023), financial frictions (Gilchrist et al., 2017), and non-linear price setting (Harding et al., 2022). We emphasize that labor-displacing technology may be contributing to the muted inflation sensitivity to the unemployment rate in the pre-Covid period.

[1] We start from a steady state with the robot-per-employee ratio of the early 2000s, and increase robot-specific technical change to reach a second steady state with an automation as in the late 2010s.

The two closest papers are Fornaro and Wolf (2022) and Leduc and Liu (2024). Fornaro and Wolf (2022) build a sticky-price model with robot adoption to show that monetary accommodations can reduce the long-term effects of automation on employment. Our complementary approach emphasizes that robot adoption decouples inflation and labor market dynamics in the short run, taking as given the stance of monetary policy. Leduc and Liu (2024) build a real model with robot adoption and search frictions to explain unemployment fluctuations. While both contributions focus on the threat of robots to workers' bargaining power, we look at how automation alters the slope of the Phillips curve.

2 Empirical Evidence

This section provides novel empirical evidence on how robot adoption leads to a decoupling between inflation and unemployment. To establish this result, we leverage geographical variation in robot adoption across U.S. metropolitan areas.

2.1 Data

We build an annual panel of inflation, unemployment rates, and robot adoption across 384 MSAs from 2008-2018. We use the regional price parities of the U.S. Bureau of Economic Analysis (BEA), which provides data on total prices, the price of goods, as well as distinct series for the price of rents, utilities, and other services at the MSA level. We complement it with wages (defined as the average compensation per job) from the BEA, the unemployment rate from the Local Area Unemployment Statistics of the U.S. Bureau of Labor Statistics (BLS), and employment and value added at the industry-MSA level from the BEA. The information on robots installed at the industry level for the U.S. and the five largest European countries comes from the International Federation of Robotics. Our measure of robot adoption is the ratio of installed robots per 1000 employees at the MSA-year level, and is derived in the two-step procedure of Acemoglu and

Restrepo (2020): we compute robots per employee for each industry-year at the national level, and combine it with employment shares at the industry-MSA level as of year 2000. Appendix A reports some summary statistics on key variables.

2.2 Econometric Specification

We estimate the causal effect of robot adoption on the sensitivity of price inflation to unemployment using the following panel regression:

$$\pi_{i,t} = \beta\, u_{i,t-1} + \gamma\, u_{i,t-1}\,(m_{i,t-1} - \bar{m}) + \mathbf{X}'_{i,t-1}\theta + \alpha_i + \delta_t + \epsilon_{i,t}, \tag{1}$$

where $\pi_{i,t}$ is a measure of inflation for MSA i at year t. We consider three alternative dependent variables: (i) the inflation rate of non-tradable goods, defined as the log-difference of the price of services excluding rents and utilities; (ii) total price inflation, defined as the log-difference of total CPI; and (iii) wage inflation, defined as the log-difference of the average compensation per job. Then, $u_{i,t}$ is the unemployment rate, $m_{i,t}$ denotes robot adoption (i.e., installed robots per 1000 employees), and $\bar{m} = \sum_i \sum_t \frac{m_{i,t}}{n_i n_t}$ is its average value across all MSA-year observations, where n_i is the number of MSAs in the sample and n_t is the number of years. The regression includes a set of covariates, $\mathbf{X}_{i,t-1}$, which consist of the lagged values of robot adoption and the relative price of non-tradable goods, as well as MSA fixed effects, α_i, and year fixed effects, δ_t. Errors, $\epsilon_{i,t}$, are clustered at the two-way MSA-year level to account for any cross-regional residual correlation.

In this setting, the coefficient β denotes the local (reduced-form) sensitivity of inflation to the unemployment rate for a MSA with an average robot adoption. The parameter γ associated with our regressor of interest—the interaction between the unemployment rate and the (demeaned) robot-per-employee ratio—captures how the inflation sensitivity to unemployment varies with automation.[2]

[2] As shown in Basso and Rachedi (2021), demeaning the robot-per-employee ratio, $m_{i,t-1} - \bar{m}$, in the interaction does not alter the estimation of the interaction term γ. This normalization allows us to directly interpret the parameter β as the sensitivity for a MSA with the average degree of automation, that is, when $m_{i,t} = \bar{m}$. Appendix A.2 reports the estimation results of a version of the empirical model that does not

2.3 Identification Strategy

Regression (1) incorporates the role of automation into the specification of Hazell et al. (2022), which leverages cross-sectional variation to estimate the slope of the local Phillips curve. Our strategy to identify the causal effect of automation hinges on shift-share instruments for the unemployment rate and robot adoption.

As far as the unemployment rate is concerned, McLeay and Tenreyro (2020) shows that there are two main challenges in estimating the true inflation sensitivity to labor market conditions.[3] To address these concerns, McLeay and Tenreyro (2020) and Fitzgerald et al. (2024) suggest to estimate the Phillips curve at the regional level. In such a setting, one can absorb any common demand and supply shocks via the presence of time fixed effects. In addition, the presence of city fixed effects absorbs any fixed heterogeneity across areas, such as fixed differences in inflation expectations. However, the estimate of the slope of the local Phillips curve would still be biased in the presence of idiosyncratic local supply shocks. To purge local supply shocks from the variation in local unemployment rates, McLeay and Tenreyro (2020) and Hazell et al. (2022) propose the use of instruments that capture shifts in local demand which are heterogeneous across areas.

In our main specification, we instrument the unemployment rate using tradable demand spillovers, as in Hazell et al. (2022). Specifically, local tradable demand spillovers in area i at year t are determined as a shift-share variable:

$$\text{Tradable Demand}_{i,t} = \sum_{x} \bar{s}_{x,i} \times \Delta \log s_{-i,x,t}, \tag{2}$$

where $\bar{s}_{x,i}$ is the average value-added share of tradable industry x in city i, and

feature any demeaning for the robot-per-employee ratio.

[3] First, the observed inflation response to unemployment incorporates the endogenous response of monetary policy. For instance, if monetary policy is carried out optimally so that there is no variation in inflation, one would find no relationship whatsoever between price changes and labor market conditions and wrongly attribute it to a flat Phillips curve. Second, while the Phillips curve is well identified upon a demand shock, this is not the case upon supply shocks, which would reverse the direction of the relationship between inflation and unemployment.

$\Delta \log s_{-i,x,t}$ is the log change in national real value added of sector x in year t, excluding the contribution of MSA i.[4] In this shift-share instrument, the identifying assumption is that neither the cross-sectional shares $\bar{s}_{x,i}$ nor the time-series shifts $\Delta \log s_{-i,x,t}$ are correlated with any supply factor.[5] Since the instrument relies on national tradable demand, we use non-tradable inflation as our baseline dependent variable to avoid direct effects of national demand on local pricing.

To corroborate the robustness to different identification strategies, Appendix A.2 instruments the unemployment rate with local government demand shifts, as in McLeay and Tenreyro (2020). In this case, we derive an alternative shift-share instrument, using the share of defense spending in a city as a fraction of GDP, and the shifts are related to the national real defense spending.[6]

Regarding automation, the main concern is that it could be due to local economic conditions, either to demand shifts, as a local surge in labor costs could put upward pressures on automation, or supply shifts, as a local surge in energy costs could curtail robot adoption. To rule this out, we follow Acemoglu and Restrepo (2020) and instrument the robot-to-employee ratio with a measure that replaces U.S. industries' robot installations with the average robot installation per industry in Denmark, Finland, France, Italy, and Sweden. This approach yields a shift-share variable in which the share captures the employment share of each industry in each MSA, and the shift is the average foreign robot installation per industry.

The first identifying assumption is that the industry shares in each MSA are

[4] The tradable industries are agriculture, mining, manufacturing, and utilities.

[5] For example, consider the surge in shale gas in the U.S. The Marcellus shale play—spanning Ohio, Pennsylvania, and West Virginia—is the largest source of shale-based natural gas. This boom disproportionately increased local demand for non-tradables in those states, raising labor demand in sectors like accommodation and food services. As a result, wages rose more in these areas than elsewhere. However, because natural gas is highly tradable, the increase in supply does not generate differential changes in energy marginal costs.

[6] This case weakens the concern on the correlation with supply factors because Nakamura and Steinsson (2014) show that the allocation of defense spending across areas is very persistent over time, as it is determined by the presence of military bases (i.e., the shares can be considered orthogonal to current supply shocks), and defense spending is the component of GDP which is least dependent on domestic economic conditions.

not related to the future trends in robot adoption. This condition is verified in two ways. On the one hand, we use employment shares in the year 2000, which is way before the increase in automation in the U.S. On the other hand, Appendix A.4 shows that our results hold above and beyond a wide set of potential confounding factors that could lead to a correlation between the shares and the shifts.

The second identifying assumption is that the shift does not reflect variation in U.S. industries' demand for automation. Insofar as the demand for robots weakly correlates across countries, our instrument isolates the supply-side component which caused the surge in the efficiency—and widespread usage—of robots. In this way, our instrument captures the fact that the production of robots has become more productive, and thus cheaper (Graetz and Michaels, 2018).[7,8]

2.4 Results

Columns (1)-(2) of Table 1 report the OLS and IV results, respectively, on how automation alters the sensitivity of non-tradable inflation to unemployment. In either case, the inflation sensitivity β is highly statistically significant: the OLS estimate is -0.1884, while the IV yields a steeper Phillips curve, with an estimate of -0.5248. We follow the approach of Hazell et al. (2022) to map our IV estimate into the slope of the aggregate Phillips curve at the quarterly frequency, and find a value of -0.0114,[9] in the ballpark of the values of the literature.

In either column, the role of automation is statistically significant, with a magnitude that rises substantially in the IV regression. The positive coefficients

[7] A potential concern is that both tradable demand spillovers and robot adoption are shift-share variables, which leverage MSA industry composition. We ensure that these two variables barely comove by using industry value-added shares to build the tradable demand spillovers, and industry employment shares for robot adoption. In this way, the correlation is 0.2. In addition, this concern does not hold when we instrument the unemployment rate with changes in defense spending, as in McLeay and Tenreyro (2020).

[8] To further corroborate the identification of the effect of robots on inflation dynamics, Appendix A.2 considers a case in which we fix automation to its 2008 levels. In this way, we abstract from the contribution of the shift, and rely only on the first identifying assumption.

[9] The slope is $[-0.5248 \times (1 - 0.995 \times 0.918)]/4$, where 0.995 is the quarterly discount factor, 0.918 is the implied quarterly persistence of unemployment in our data (estimated as an AR(1) process with MSA and year fixed effects), and the factor of 4 accounts for time aggregation from quarterly into annual inflation.

Table 1: Robot Adoption and Inflation across MSAs

Dependent Variable:	Non-tradable Inflation		CPI Inflation		Wage Inflation	
	OLS	IV	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)
$u_{i,t-1}$	-0.1884***	-0.5248***	-0.1268***	-0.2882***	-0.3855***	-0.9815***
	(0.0221)	(0.1469)	(0.0157)	(0.0856)	(0.0330)	(0.2483)
$u_{i,t-1} \times (m_{i,t-1} - \bar{m})$	0.0010**	0.0058**	0.0007***	0.0028*	0.0016**	0.0060**
	(0.0004)	(0.0028)	(0.0003)	(0.0016)	(0.0007)	(0.0028)
Year & MSA Fixed Effects	✓	✓	✓	✓	✓	✓
N. Observations	3,205	3,205	3,205	3,205	3,205	3,205
Cragg-Donald F stat.	-	20.59	-	22.33	-	64.14

Note: The table reports regression estimates on annual data over 2008-2018 across U.S. MSAs. The dependent variable is non-tradable inflation in Columns (1)-(2), CPI inflation in Columns (3)-(4), and wage inflation in Columns (5)-(6). The key independent variables are the lagged unemployment rate, $u_{i,t-1}$, and its interaction with robot adoption, $u_{i,t-1} \times (m_{i,t-1} - \bar{m})$. All regressions include lagged values of robot adoption and the relative price of non-tradables, as well as year and MSA fixed effects. In Columns (2)-(4)-(6), unemployment is instrumented with a shift-share variable that captures tradable demand spillovers, and robot adoption is instrumented with the robot penetration in a pool of European countries. Two-way clustered standard errors are reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10%, respectively.

imply that price inflation is less reactive to changes in the local labor market in cities with more robots. In other words, automation decouples inflation from unemployment. This effect is highly economically significant: a one standard deviation rise in robot adoption reduces the sensitivity of non-tradable inflation to unemployment by 17%. Columns (3)-(4) provide similar results—with a slightly lower magnitude—when looking at total inflation: a one standard deviation rise in robots reduces the inflation sensitivity to unemployment by 15%.

Columns (5)-(6) look at wage inflation: the interaction term is again statistically significant, with a magnitude implying that a one standard deviation rise in robots reduces the wage inflation sensitivity to unemployment by 10%. Thus, the flattening of the local wage Phillips curve accounts for 59%-67% of the drop in the sensitivity of price inflation to unemployment.

This evidence uncovers two novel findings relating automation to inflation dynamics: (*i*) robot adoption reduces the sensitivity of price inflation to unemployment, (*ii*) this effect hinges on a muted wage inflation sensitivity. These results are validated in an extensive battery of robustness checks in Appendix A.

2.5 The Role of Workers' Bargaining Power

What drives the dampening effect of automation on the inflation sensitivity to unemployment? We show that this is modulated by workers' bargaining power. To do so, we modify regression (1) as follows:

$$\pi_{i,t} = \beta_1\, u_{i,t-1} + \beta_2\, u_{i,t-1}\mathbb{I}\{\text{HBP}\}_{i,t-1} + \gamma_1\, u_{i,t-1}\,(m_{i,t-1} - \bar{m}) + \dots \qquad (3)$$
$$\dots + \gamma_2\, u_{i,t-1}\,(m_{i,t-1} - \bar{m})\,\mathbb{I}\{\text{HBP}\}_{i,t-1} + \mathbf{X}'_{\mathbf{i,t-1}}\theta + \alpha_i + \delta_t + \epsilon_{i,t},$$

where $\mathbb{I}\{\text{HBP}\}_{i,t-1}$ is an indicator function that equals 1 for those MSAs with high workers' bargaining power in year $t-1$. We consider two proxies for workers' bargaining power. First, we use information from the Current Population Survey (CPS) and compute the union membership rate for each MSA-year. The indicator equals 1 for those MSAs-year in which the membership rate is above the median in the sample.[10] Second, we follow a similar approach based on the union coverage rate, which also includes non-member workers covered by a union. We then estimate the regression using the same IV method as in the baseline regression (1).

In this setting, we can test whether the flattening effect of automation positively depends on workers' bargaining power by evaluating whether the estimated value of γ_2 is greater than that of γ_1. Insofar as $\gamma_2 > \gamma_1$, the flattening effect of automation is relatively larger in cities with high workers' bargaining power.

The results in Table 2 show that the unionization rate does shape how automation reduces the inflation sensitivity to unemployment. The fact that γ_2, the coefficient of the triple interaction, is positive and statistically significant indicates that the effect of robot penetration in dampening the inflation sensitivity to unemployment is stronger in cities with sufficiently high workers' bargaining power.

[10]Our approach of interacting robot adoption and unionization rates is justified by the fact that the correlation is virtually zero.

Table 2: Robot Adoption and Inflation across MSAs

Dependent Variable:	Non-tradable Inflation	
	Union Membership	Union Coverage
	IV	IV
	(1)	(2)
$u_{i,t-1}$	-0.5446***	-0.5354***
	(0.1465)	(0.1475)
$u_{i,t-1} \times \mathbb{I}\{\text{HBP}\}_{i,t-1}$	0.0291***	0.0218**
	(0.0091)	(0.0091)
$u_{i,t-1} \times (m_{i,t-1} - \bar{m})$	0.0020	0.0020
	(0.0025)	(0.0025)
$u_{i,t-1} \times (m_{i,t-1} - \bar{m}) \times \mathbb{I}\{\text{HBP}\}_{i,t-1}$	0.0029**	0.0031**
	(0.0013)	(0.0012)
Year & MSA Fixed Effects	✓	✓
N. Observations	3,205	3,205

Note: The table reports regression estimates as in Table 1 also including the double interaction between the unemployment rate with the indicator function $\mathbb{I}\{\text{HBP}\}_{i,t-1}$, that equals 1 for MSA-year with high workers' bargaining power, and the triple interaction that multiplies the double interaction with the demeaned amount of robot penetration. In Column (1), the indicator function equals 1 for those MSA-year with a fraction of workers that are union members which is above the sample median. In Column (2), the indicator function considers both union members and workers covered by union. *** and ** indicate statistical significance at the 1% and 5%, respectively.

This holds for union members and non-member workers covered by unions.[11]

In Appendix A.5, we also evaluate the implications of potential alternative explanations for the main channel through which automation alters inflation dynamics. Specifically, we consider the role of geographical differences in job flows (hires, separations, job-to-job), utilization rates (proxied by hours worked per full-time production workers), and sectoral composition (shares of manufacturing and tradable industries in total local GDP). To do so, we estimate a regression similar to that in Equation (3), with the only difference that we replace the dummy variable that captures high workers' bargaining power, $\mathbb{I}\{\text{HBP}\}_{i,t-1}$, with a similar indicator based on the alternative conditioning variables. We find no

[11] Appendix A.5 reports the results of a similar exercise in which we use the union membership and coverage rates as of 2008—rather than using its time-varying values—to address any potential endogeneity of unionization with respect to the surge of automation throughout the sample period.

evidence that the flattening effect of automation on local Phillips curves hinges on these mechanisms. This result corroborates that workers' bargaining power is the key channel that modulates the way in which robot installations alter inflation dynamics, and this effect holds above and beyond potential alternative explanations.

All in all, our evidence points out that automation reduces the inflation sensitivity to unemployment, and this effect depends on the combination of the stock of installed robots and the degree of workers' bargaining power.

3 Model

To rationalize the evidence in the previous section, we extend a standard New Keynesian model with robot adoption and labor market search frictions. This framework allows wage setting—and thus the responsiveness of wages to labor market conditions—to be jointly shaped by workers' bargaining power and the stock of robots. Bargaining power determines workers' ability to extract rents, while the stock of robots captures how easily firms can further substitute labor, limiting that ability. Their interaction explains how the threat of automation weakens the wage-unemployment link, in line with our empirical findings.

The economy is populated by a household consisting of a continuum of workers, who look for a job, and pool income with perfect consumption insurance. The production side has three layers: (*i*) a varying measure of producers that either use machines or post job vacancies and look for workers; (*ii*) monopolistically competitive wholesalers that convert producers' goods into different varieties, subject to a price setting friction; and (*iii*) a retailer that assembles the varieties into the final good. Final goods are sold to the household and the machine manufacturer, which transforms them into machines. We describe the main aspects of the model below and refer to Appendix B for all details, derivations, and equilibrium conditions.

3.1 Producers

Every period, there is a measure Ξ_t of producers that pay a fixed nominal monetary cost κ to enter the market.[12] We index each producer by $j \in [0, \Xi_t]$. Upon entry, producers draw an idiosyncratic efficiency in operating with a labor technology, γ_j, from a distribution $f(\gamma;\alpha)$ with support $[1, \gamma_H]$, and shape parameter α.

After drawing the labor efficiency, producers decide to complete tasks with either machines (i.e., robot firms) or workers (i.e., labor firms). In case a producer uses machines, they produce with certainty using a linear technology with unit efficiency—at the lower bound of producers' labor efficiency—and sell their output to wholesalers at price $P_{P,t}$. The nominal value of robot firms equals the value of sales net of the cost of purchasing a machine and the entry cost,

$$V_{M,j,t} = P_{P,t} - P_{M,t} - \kappa. \tag{4}$$

Since all robot firms operate at the same efficiency, $V_{M,j,t} = V_{M,t}$, for all j.

In case a producer decides to use labor, it opens a vacancy, which it fills with probability $q_t(\theta_t)$, which we define below. In such a case, the labor firm produces using a linear technology with efficiency γ_j, and sells its output to wholesalers at price $P_{P,t}$. The nominal value of a labor firm equals the value of sales net of the wage, multiplied by the probability of filling the vacancy, minus the entry cost,

$$V_{L,t}(\gamma_j) = q_t(\theta_t)\left[P_{P,t}\gamma_j - W_t\right] - \kappa. \tag{5}$$

How do producers sort into labor firms and robot firms? A producer j opens a vacancy and operates the labor technology when the value of being a labor firm is greater than the value of a robot firm, that is, $V_{L,t}(\gamma_j) > V_{M,t}$. Since the value of being a labor firm increases with the labor efficiency level γ_j, there exists a cut-off point for the labor efficiency, $\gamma_t^\star$, such that

$$V_{L,t}(\gamma_t^\star) = V_{M,t}, \tag{6}$$

[12] Entry costs are transferred to households. Results are equivalent if entry costs were to be paid in goods.

and firms are indifferent between the two technologies. The cut-off shapes the automation choices: producers with a labor efficiency above $\gamma_t^\star$ turn into labor firms, whereas the rest become robot firms.[13]

Given the automation cut-off point, we can characterize the measure of labor firms and robot firms in the economy. The measure of labor firms integrates across all the producers with efficiency above $\gamma_t^\star$, and the measure of robot firms captures low-efficiency producers:

$$\Xi_{L,t} = \Xi_t \int_{\gamma_t^\star}^{\gamma_H} f\left(\gamma;\alpha\right)\, \mathrm{d}\gamma, \tag{7}$$

and

$$\Xi_{M,t} = \Xi_t \int_{1}^{\gamma_t^\star} f\left(\gamma;\alpha\right)\, \mathrm{d}\gamma. \tag{8}$$

In equilibrium, the sum of the measures of labor firms and robot firms equals the mass of producers that have entered the market: $\Xi_{L,t} + \Xi_{M,t} = \Xi_t$. These measures are set such that the expected value of entering the market is zero, that is $V_{E,t} = V_{M,t} \int_1^{\gamma_t^\star} f\left(\gamma;\alpha\right) \mathrm{d}\gamma + \int_{\gamma_t^\star}^{\gamma_H} V_{L,t}\left(\gamma\right) f\left(\gamma;\alpha\right) \mathrm{d}\gamma = 0$.

Given the measure of labor firms and robot firms, we can define the total amount of goods assembled by producers, Z_t, as

$$Z_t = \Xi_t \left[\int_{\gamma_t^\star}^{\gamma_H} q_t(\theta_t)\, \gamma\, f\left(\gamma;\alpha\right) \mathrm{d}\gamma \right] + \Xi_{M,t}. \tag{9}$$

3.2 Labor Market

As in Pissarides (2000), the labor market is characterized by N_t individuals who search for a job and firms posting v_t vacancies. Since firms are heterogenous in labor productivity and workers cannot observe it, firms would have the incentive to claim they have the lowest productivity level, $\gamma_t^\star$, to maximize their surplus. However, workers have information on prices and the distribution of productivity levels across firms. Therefore they can obtain the expected value of the matched

[13] While workers and machines are perfect substitutes at the producer level, Appendix C.1 shows that they are imperfect substitutes at the aggregate level.

firm's surplus, by inferring the set of firms who posted vacancies, that is, those firms whose $\gamma_j > \gamma_t^\star$, which implies that total vacancies equal $v_t = \Xi_{L,t}$. The expected value of the surplus of firms posting vacancies is

$$\mathbb{E}_t\left[S_t\right] = \int_{\gamma_t^\star}^{\gamma_H} \left[q_{P,t}\gamma - w_t\right] f(\gamma;\alpha)d\gamma. \tag{10}$$

Firms and workers match randomly, with a flow, $x_t(v_t, s_t)$, determined by:

$$x_t(v_t, s_t) = \xi v_t^\eta s_t^{1-\eta}, \tag{11}$$

where η is the elasticity of the matching function with respect to vacancies, and ξ denotes the matching efficiency. Matches last for one period.

Given the matching function (11) and the tightness $\theta_t \equiv v_t/s_t$, which describes the ratio of vacancies to job searchers, the probability that a person finds a job and the probability that a firm fills a vacancy equal, respectively,

$$p_t\left(\theta_t\right) = \frac{x_t(v_t, s_t)}{s_t} = \xi\theta_t^\eta, \tag{12}$$

and

$$q_t\left(\theta_t\right) = \frac{x_t(v_t, s_t)}{v_t} = \xi\theta^{\eta-1}. \tag{13}$$

Upon a match, workers and firms bargain on the real wage by splitting the total surplus of the match. The wage bargaining problem is

$$\arg\max_{w_t} w_t^\tau \; \mathbb{E}_t\left[S_t\right]^{1-\tau}, \tag{14}$$

where τ denotes workers' bargaining power. The optimal wage then equals

$$w_t = \tau q_{P,t}\frac{\int_{\gamma_t^\star}^{\gamma_H} \gamma f\left(\gamma;\alpha\right) \mathrm{d}\gamma}{\int_{\gamma_t^\star}^{\gamma_H} f\left(\gamma;\alpha\right) \mathrm{d}\gamma}, \tag{15}$$

which depends on both workers' bargaining power and the automation threshold.

The unemployment rate equals the ratio between the measures of individuals that have not matched with a producer and those actively looking for a job:

$$u_t = \left[N_t - \Xi_t q_t(\theta_t)\left(\int_{\gamma_t^\star}^{\gamma_H} \gamma f(\gamma;\alpha)d\gamma\right)\right]/N_t. \tag{16}$$

3.3 Wholesalers

A unit measure of monopolistically competitive wholesalers, indexed by $i \in [0,1]$, purchase goods $Z_{i,t}$ from producers at price $P_{P,t}$, and transform them into different varieties $Y_{i,t}$ with the linear technology:

$$Y_{i,t} = Z_{i,t}. \tag{17}$$

Varieties are sold to the retailer at price $P_{i,t}$, so that profits equal $P_{i,t}Y_{i,t} - P_{P,t}Z_{i,t}$.

Wholesalers face a price-setting friction in the form of a Rotemberg adjustment cost, denoted by the parameter ϕ. Wholesalers optimally set their price $P_{i,t}$ by maximizing expected profits net of Rotemberg costs. In equilibrium, wholesalers set the same price: $P_{i,t} = P_t$ for all i. We denote by $\pi_t = \frac{P_t}{P_{t-1}}$ the inflation rate.

Market clearing implies that the goods produced by wholesalers—net of the Rotemberg cost—equal those produced by both labor firms and robot firms,

$$\int_0^1 \left[1 - \frac{\phi}{2}\left(\frac{P_{i,t}}{P_{i,t-1}} - 1\right)^2\right] Y_{i,t}\,\mathrm{d}i = \int_0^1 Z_{i,t}\,\mathrm{d}i = Z_t. \tag{18}$$

3.4 Retailer

A perfectly competitive representative retailer assembles wholesalers' varieties, $Y_{i,t}$, into the final good of the economy, Y_t, with a CES technology:

$$Y_t = \left[\int_0^1 Y_{i,t}^{\frac{\epsilon-1}{\epsilon}}\,\mathrm{d}i\right]^{\frac{\epsilon}{\epsilon-1}}, \tag{19}$$

where ϵ is the elasticity of substitution across varieties. Final goods are sold at price P_t to household, in form of consumption goods C_t, and to machine manufacturers, in form of investment goods I_t, such that

$$Y_t = C_t + I_t. \tag{20}$$

3.5 Machine Manufacturer

A perfectly competitive representative manufacturer purchases goods I_t from the retailer at price P_t, and transforms them into machines M_t with the technology

$$M_t = \zeta I_t, \tag{21}$$

where ζ is the exogenous level of robot-specific technical change. The manufacturer sells the machines to robot firms at price $P_{M,t}$. This price inversely relates to the level of technical change: $P_{M,t} = P_t/\zeta$. Robot-specific technical change raises the efficiency of the production of machines, thus curtailing their price.

In equilibrium, the machines sold by the manufacturer equals those demanded by robot firms (i.e., the measure of robot firms): $M_t = \Xi_{M,t}$.

3.6 Household

The household consists of a unit measure of individuals, indexed by $x \in [0, 1]$, with perfect consumption insurance. Given total nominal labor earnings X_t, which we describe below, firms' nominal profits D_t, and entry costs T_t that are rebated to the household, the household decides the optimal levels of consumption goods, C_t, purchased from retailers at price P_t, and savings in one-period nominal bonds, B_t. Specifically, the household maximizes its lifetime utility

$$\max_{C_t, B_{t+1}} \mathbb{E}_t \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\sigma}}{1-\sigma} \tag{22}$$

$$\text{s.t. } P_t C_t + B_t = B_{t-1} R_{t-1} + X_t + D_t + T_t, \tag{23}$$

where R_t is the nominal interest rate, β is the discount factor, and σ captures risk aversion. We then alter households' inter-temporal Euler equation by adding a preference shifts, Ω_t, that follows the AR(1) process

$$\log \Omega_t = \rho_\Omega \log \Omega_{t-1} + \varepsilon_{\Omega,t}, \tag{24}$$

with auto-regressive parameter ρ_Ω and preference shock $\varepsilon_{\Omega,t}$.

Individuals decide whether to look for a job: each individual draws a searching cost, χ, from a uniform distribution, $U(0, \chi_H)$, and look for a job when the real value of doing so (w_t) exceeds the searching cost, that is when $w_t \geq \lambda$. This specification rules out any wealth effect on labor supply, as they prevent a direct mapping between the slopes of the regional and aggregate Phillips curves (see Hazell et al., 2022). As a result, a measure of $N_t = w_t/\chi_H$ individuals search for jobs.

Once the matches are realized, all wages are pooled together within the household, so that total nominal labor earnings is $X_t = \Xi_t q_t(\theta_t) W_t \int_{\gamma_t^\star}^{\gamma_H} f(\gamma;\alpha)\,\mathrm{d}\gamma$.

3.7 Monetary Authority

The monetary authority sets the nominal interest rate with a Taylor rule that reacts to inflation and the unemployment gap, u_t/u_t^F, where u_t^F is the unemployment rate in the flexible-price economy, such that

$$R_t/\bar{R} = \left[R_{t-1}/\bar{R}\right]^{\psi_R}\left[(1+\pi_t)^{\psi_\pi}\left(u_t/u_t^F\right)^{\psi_u}\right]^{1-\psi_R}, \tag{25}$$

where ψ_R captures the degree of interest-rate smoothing, and ψ_π and ψ_u denote the responsiveness of interest rates to the inflation rate and the unemployment gap (u_t^F is define as the level of unemployment under flexible prices), respectively, and $\bar{R}$ is the steady-state nominal interest rate.[14]

3.8 Model Assumptions

Let us detail several key features of the model. First, we assume that matches last for only one period, as all firms are short-lived and exit after a single period of operation. While this assumption is admittedly stark, Appendix D.6 shows that allowing for longer-lasting job matches does not alter the model's core implications. Moreover, our approach aligns with the empirical evidence in Appendix A.5, which finds that automation's dampening effect on inflation sensitivity to unemployment is unrelated to local labor market characteristics such as hiring, separation, and job-to-job transition rates.

Second, labor supply in the model operates exclusively along the extensive margin: households cannot adjust hours worked, so total hours vary only through

[14] Throughout the paper, we denote by $\bar{A}$ the steady-state value of variable A_t.

participation decisions. This assumption is consistent with the empirical findings of Blundell et al. (2011), who show that the extensive margin accounts for most of the variation in working hours in the U.S. economy. To ensure an interior participation solution, we introduce idiosyncratic search costs drawn from a uniform distribution, which prevent the labor force from being trivially full or empty.

Third, we assume that firms and workers are atomistic in wage bargaining and take the automation threshold as given. While the equilibrium automation cutoff still depends on wages, we rule out the additional feedback between wages and automation decisions during bargaining.

Fourth, during the matching process, workers cannot observe firms' productivity. This choice is necessary for maintaining tractability: if workers were to observe the firm-specific productivity, the equilibrium would feature a distribution of wages, requiring us to track the entire distribution of firms when solving the model—effectively introducing heterogeneous producers. In an alternative version of the model with directed search (available upon request), we relax this assumption and allow workers to observe firm productivity. Importantly, this alternative setup yields quantitative results that are nearly identical to those of the baseline model.

Finally, we focus on preference shocks as the sole source of uncertainty. Appendix C.5 shows that our results remain robust when introducing other disturbances, such as monetary policy and productivity shocks.

4 Characterization of the Phillips Curve

This section analytically derives how the slope of the Phillips curve depends on automation. Within a simplified, tractable framework, we highlight three results: (i) a higher stock of robots flattens the Phillips curve; (ii) this effect is amplified when workers' bargaining power is high; and (iii) the channel operates only if automation responds to shocks. This setup clarifies our core mechanism: the impact of the automation threat on inflation dynamics depends on the joint influence of the stock of robots and workers' bargaining power. To gain tractability while preserving this mechanism, we introduce two simplifying assumptions. We later verify that these assumptions do not materially affect our results by validating them within the full baseline model.

Our two key simplifications are as follows. First, we assume producers draw one of two productivity levels: either $\gamma = 0$—and operate using machines—with probability $\Phi_t^\star$, or $\gamma = \gamma_H > 1$ with probability $1 - \Phi_t^\star$. The probability $\Phi_t^\star$ corresponds to the share of robot firms, that is, $\Phi_t^\star = \Xi_{M,t}/\Xi_t$. Second, we posit that this share is a reduced-form function of wages:

$$\Phi_t^\star = \bar{\Phi}^\star \left[1 + \varrho(w_t - \bar{w})\right], \tag{26}$$

where $\bar{\Phi}^\star$ and $\bar{w}$ are the steady-state share of robot firms and real wages, and $0 < \varrho < 1$ pins down the elasticity of automation to real wages, denoted by $\xi_{\Phi^\star,w}$.[15]

In this setting, we can explicitly express the relationship between inflation and unemployment. If we denote by $\hat{A}_t$ the log deviations of variable A_t from its steady state $\bar{A}$, the Phillips curve reads:

$$\hat{\pi}_t = \Psi(\bar{\Phi}^\star;\tau)\hat{u}_t + \beta \mathbb{E}_t\left[\hat{\pi}_{t+1}\right], \tag{27}$$

where the slope $\Psi(\bar{\Phi}^\star;\tau)$—which describes the sensitivity of inflation to the unemployment rate—explicitly depends on the degree of automation at the steady state, $\bar{\Phi}^\star$, and workers' bargaining power, τ, as follows:

$$\Psi(\bar{\Phi}^\star;\tau) = -\frac{\epsilon - 1}{\phi}\left[\frac{\bar{u}}{1-\bar{u}}\frac{1-\eta}{\eta}\left(\varrho\,\tau\gamma_H\frac{1}{1-\bar{\Phi}^\star}\frac{\epsilon-1}{\epsilon} + \frac{\epsilon}{(\epsilon-1)\zeta - \epsilon}\right)^{-1}\right]. \tag{28}$$

The slope is always negative, so that inflation rises when unemployment is low.

We then derive three key analytical results. We start with Proposition 1,

[15] In this setting, robot-specific technical change ζ is inconsequential for the level of automation as long as all firms that draw $\gamma = \gamma_H$ select to produce with labor. Without loss of generality, we assume that $\zeta > \left(2 - \bar{\Phi}^\star\right)\epsilon/(\epsilon - 1)$ to ensure firms find robot adoption profitable.

which states that the slope varies as a function of automation: a higher share of robot firms at steady state flattens the Phillips curve.

Proposition 1. *A higher degree of automation in steady states increases the slope of the Phillips curve towards zero:* $\frac{\partial \Psi(\bar{\Phi}^{\star};\tau)}{\partial \bar{\Phi}^{\star}} > 0$.

Proof. See Appendix B.2. □

Since the slope is negative, the positive derivative of Proposition 1 implies that higher steady-state robots pushes the value of the slope towards zero, thus muting the relationship between inflation and unemployment. How does automation alters inflation dynamics? Proposition 2 uncovers that firms' bargaining power is key in modulating the effect of automation on the slope.

Proposition 2. *A surge in steady-state automation flattens relatively more the Phillips curve at higher values of workers' bargaining power:* $\frac{\partial^2 \Psi(\bar{\Phi}^{\star};\tau)}{\partial \bar{\Phi}^{\star} \partial \tau} > 0$.

Proof. See Appendix B.2. □

Intuitively, automation alters workers' and firms' expected payout, as firms have the possibility to replace workers if wages rise substantially. Thus, the outside option of automation alters workers' effective bargaining power in wage setting, leading to a flat Phillips curve.[16]

Importantly, Proposition 3 establishes that higher automation in steady state does reduce the slope only insofar as automation can adjust upon a shock.

Proposition 3. *Steady-state automation flattens the Phillips curve if and only if automation varies with wages upon shocks, i.e., insofar as* $\varrho \neq 0$*:* $\left.\frac{\partial \Psi(\bar{\Phi}^{\star};\tau)}{\partial \bar{\Phi}^{\star}}\right|_{\varrho=0} = 0$.

Proof. See Appendix B.2. □

In other words, what matters for inflation dynamics is not just the existing stock of robots, but rather firms' ability to further automate.

[16]This result complements the evidence of Lombardi et al., 2023, in which a reduction in explicit workers' bargaining power reduces the inflation sensitivity to unemployment. In our case, changes in automation affects inflation dynamics even absent formal changes in workers' protection, as the option of installing robots curtails workers' effective bargaining power.

5 Quantitative Analysis

Having characterized the effects of automation on inflation dynamics in a simplified setting, we turn to the quantitative predictions of the full model. We begin by validating the two simplifying assumptions behind the analytical results on the role of automation. We then confirm those findings numerically by studying the impact of robots on the Phillips curve. Next, we show that the model delivers estimates closely aligned with the data, lending credibility to its core mechanism: the joint influence of the stock of robots and the degree of workers' bargaining power in modulating the effect of the automation threat on inflation dynamics. Finally, we use the model to quantify how much automation has contributed to the decline in the sensitivity of inflation to unemployment over the past three decades.

5.1 Calibration

We calibrate the model to the U.S. economy by considering that one period corresponds to a quarter. We consider a zero net inflation rate in the steady state. We set households' risk aversion to the standard value of $\sigma = 2$, and the discount factor to $\beta = 0.995$, which implies a 2% annual steady-state real interest rate.

The labor efficiency γ_j is drawn from a Truncated Pareto Distribution $f(\gamma) = \frac{\alpha\gamma_j^{-\alpha-1}}{1-\gamma_H{}^{-\alpha}}$. The location parameter is $\gamma_H = 1.1$, so the most productive firms have a 10% efficiency gain in using the labor technology than the least productive ones.[17] The scale parameter, $\alpha = 5$, and the matching efficiency, $\xi = 0.9$, ensure that the vacancy filling and job finding probabilities are within zero and one. The elas-

[17] Empirically disciplining this parameter is particularly challenging, as it reflects the productivity differential between tasks performed by robots and those carried out by workers. Nonetheless, it can be treated as a degree of freedom in the model, since it jointly determines the automation threshold alongside the rate of robot-specific technical change, ζ. For any given value of γ_H, the model can be recalibrated to replicate the same steady-state level of automation. Recent evidence by Bonhomme et al. (2023) points to a limited role for firm-specific productivity—estimating that only 5% of the variance in workers' earnings is attributable to firm fixed effects—which supports the view that this productivity differential is relatively small. This, in turn, reinforces the interpretation of the parameter as a modeling degree of freedom, rather than a key structural margin that requires tight empirical discipline.

ticity of labor matches with respect to vacancies equals $\eta = 0.5$ as in Petrongolo and Pissarides (2001). We set the entry fixed cost, $\kappa = 0.42$, and the searching cost, $\lambda = 1.8$, such that the unemployment rate is 5.7%, and the participation rate is 63%, in line with the average rates observed in the 2010's in the U.S.

For any given parametrization of the labor efficiency distribution, robot-specific technical change pins down automation in the steady state. We discipline it by targeting a 0.2% robot to (full-time equivalent) employee ratio, as documented for the U.S. in the 2000s by Acemoglu and Restrepo (2020). To match this value, we set $\zeta = 2.26$. The elasticity of substitution across varieties is $\epsilon = 4$, so that the markup is around 30%, in line with the recent evidence of De Loecker et al. (2020). The Rotemberg cost is $\phi = 46$ targeting a slope of the Phillips curve of -0.0114, as implied by our evidence. This yields a 13-month price duration.

We parameterize the Taylor rule following the estimates of Clarida et al. (2000): an inertia of $\psi_R = 0.8$ and responsiveness to inflation and the unemployment gap of $\psi_\pi = 1.5$ and $\psi_u = -0.125$. Finally, the persistence of the preference shock process is $\rho_\Omega = 0.93$, in line with Justiniano and Primiceri (2008).

5.2 Validation of the Tractable Framework

To derive analytical results on how robot adoption affects the slope of the Phillips curve in Section 4, we simplified the model into a tractable framework where automation responds to wages via a reduced-form relationship: higher labor costs lead to greater robot installation, as captured by Equation (26). To validate this assumption, we now show that, in a similar fashion, automation endogenously increases with labor costs in the full baseline model. In other words, while the analytical framework imposes this relationship by assumption, it endogenously arises from the underlying economic structure of the full model.

In the full model, wages and robots are linked by the log-linear relation:

$$\hat{\gamma}_t^\star = \Omega\left(\bar{\gamma}^\star, \tau\right)\hat{w}_t, \tag{29}$$

Figure 1: The automation elasticity to the real wage

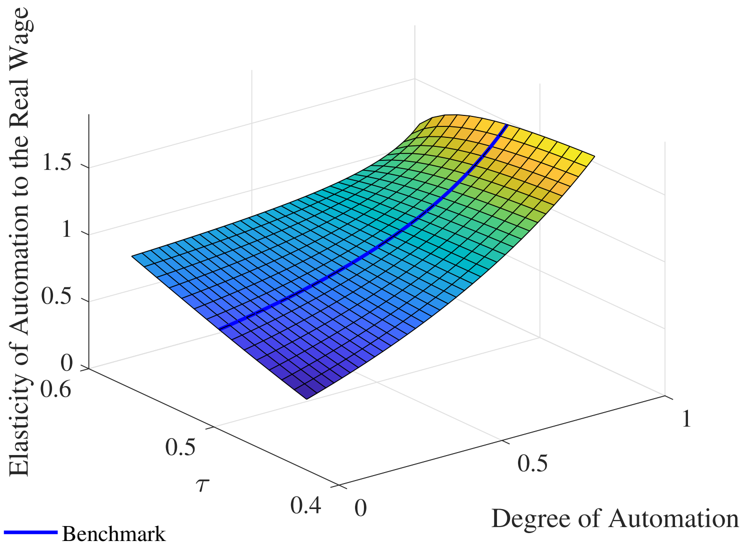

Note: The figures show how the elasticity of automation to the real wage varies with robot-specific technical change ζ (i.e., the degree of automation at the steady state) and workers' bargaining power τ.

where the loading factor, $\Omega\left(\bar{\gamma}^{\star}, \tau\right)$, depends on the steady-state level of automation, $\bar{\gamma}^{\star}$, and workers' bargaining power, τ.[18] To illustrate that the model endogenously delivers the same relationship between automation and real wages as in the tractable framework, Figure 1 shows that the elasticity of automation to wages is indeed positive and increasing in both the level of automation and the degree of workers' bargaining power.

Insofar the automation threat to wages is pinned down by the level of automation threshold—and its changes over time— Equation (29) clarifies that the stock of installed robots and workers' bargaining power are the two key variables that jointly characterize its behavior. This result supports our empirical strategy: we proxy the threat of automation with the stock of robots—and its variations across U.S. cities—and we show that its effect in reducing the sensitivity of inflation to unemployment is amplified in areas with higher workers' bargaining power.

[18] Appendix B.5 provides the full functional form of the loading factor, $\Omega\left(\bar{\gamma}^{\star}, \tau\right)$.

5.3 Automation and the Phillips Curve

Does the analytical characterization of the Phillips curve slope extend to the full baseline model, where firms endogenously choose whether to automate? This section addresses this question by numerically examining how the slope varies with robot installations when automation is an equilibrium outcome.

We proceed by computing the slope of the Phillips curve across economies that differ only in the level of robot-specific technical change, ζ. A higher ζ implies greater efficiency in producing machines, lowering their relative price and raising the steady-state automation cut-off, $\bar{\gamma}^{\star}$. As a result, the share of robot firms increases in the steady state.[19] By varying ζ while holding all other parameters fixed, we isolate how long-run changes in automation influence inflation dynamics outside the steady state.

Panel (a) of Figure 2 reports how the slope of the Phillips curve varies as a function of both the level of steady-state automation and workers' bargaining power.[20] The slope is consistently negative, reflecting the usual inverse relationship between inflation and unemployment. Crucially, we find two results: (i) automation flattens the Phillips curve in the full model, and (ii) this effect is stronger when workers' bargaining power is high. In other words, robot adoption alters inflation dynamics only when workers have leverage in wage negotiations—since the outside option of automation constrains their ability to extract rents. Bargaining power is thus a key force modulating the impact of robot installations on wages and, consequently, on inflation.

Panel (b) replicates this exercise in a counterfactual version of the economy

[19] Graetz and Michaels (2018) document that the surge in robot adoption has coincided with a marked decline in robot prices.

[20] In this exercise, we vary the price of robots and bargaining power while holding unemployment and participation rates constant across steady states, by adjusting κ and λ_H. All other parameters are set to their calibrated values from Section 5.1. The solid line represents the slope as automation varies at the calibrated value $\tau = 0.5$.

Figure 2: Automation, workers' bargaining power, and the slope of the Phillips curve

(a) Baseline (b) Automation fixed at steady-state level

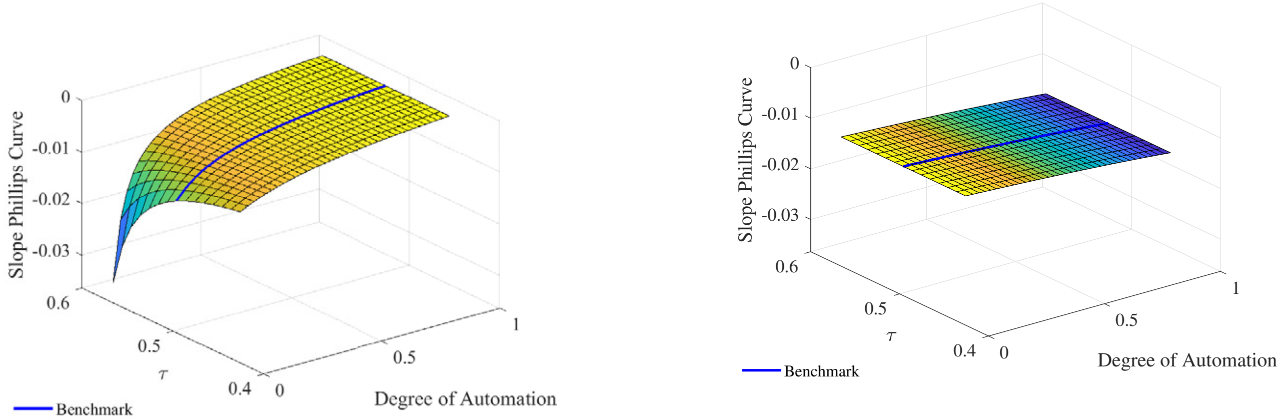

Note: The figures show how the slope of the Phillips curve of the baseline model (Panel a) and a counterfactual economy with automation fixed at its steady-state value (Panel b) varies with robot-specific technical change ζ (i.e., the degree of automation) and workers' bargaining power τ.

where automation varies across steady states but remains fixed in response to shocks. In this case, changes in steady-state automation have little effect on the Phillips curve slope. This finding echoes the analytical results from Section 4: the existing stock of robots is only relevant to understand inflation dynamics when automation responds to shocks.

The results reinforce the finding that automation flattens the Phillips curve—especially when workers' bargaining power is high—but this mechanism is operative only insofar as firms can adjust automation in response to shocks. It is not merely the stock of robots that matters, but the ability of firms to alter automation when wage pressures arise. This underscores the joint importance of installed robots and bargaining power in shaping how the threat of automation affects wage—and therefore price—dynamics. These insights corroborate our empirical approach, which leverages variation in robot adoption and workers' bargaining power to identify the causal effect of automation on inflation.

Appendix C.2 and C.3 provide additional support by dissecting the wage dynamics channel. There, we show that automation dampens wage pressures primarily through its effect on wage setting. In our model, wages respond to labor market tightness—rising when vacancies are harder to fill. However, as automation increases, workers' effective bargaining power declines, weakening the link between tightness and wage growth. This decoupling effect is amplified at higher levels of bargaining power, again underscoring its central role in shaping the inflation response to automation.

Appendix C.2 quantifies the contribution of the wage-setting channel of automation. Again, we compare the baseline model, where automation responds endogenously to wages, with a counterfactual in which robot adoption is fixed. In the baseline economy, the wage-setting effect of automation becomes increasingly important as steady-state automation rises: it accounts for a growing share of the total wage response to a preference shock. In contrast, when automation cannot adjust, the wage-setting contribution remains flat and small, regardless of the stock of robots. These findings, together with the earlier result that the Phillips curve slope does not vary with automation or bargaining power when automation is held fixed, underscore the necessary role of adjustment at the margin for the mechanism to operate.

An additional exercise in Appendix C.4 further emphasizes the central role of the robot stock. We study how much steady-state automation must increase to generate a given reduction in the inflation sensitivity to unemployment across different levels of workers' bargaining power. The results show that when workers have high bargaining power, the model requires a much smaller increase in the stock of robots to match the same dampening in inflation responsiveness observed under weak bargaining power. In other words, the inflation effect of automation depends not only on the size of the stock, but also on how strongly workers can resist wage compression. As we highlight in the introduction, the threat of automation shapes wage-setting dynamics through the joint influence of the stock of robots and workers' bargaining power.

5.4 Quantification of the Flattening due to Automation

After characterizing how automation alters the Phillips curve, we quantify the extent to which the observed surge in robot installations in recent decades has reduced inflation's sensitivity to unemployment. We do so through a transitional dynamics exercise.

Specifically, we simulate the impact of a surge in robot adoption—mirroring the U.S. trend over recent decades—on inflation sensitivity to unemployment upon preference shocks. Starting from a steady state with a 0.2% robot-per-employee ratio (early 2000s level), we gradually increase robot-specific technical change (ζ) until reaching a new steady state with a 0.8% ratio. This value ensures that the rise in automation across the two steady states roughly mimics that of the U.S. in our sample period. Then, at each step of the transition, we feed the model with a set of preference shocks. We discipline the automation changes and preference shocks by matching the relative standard deviation of the robot-to-employee ratio and unemployment rate in our data.[21]

This exercise builds on the key premise that we can closely tie our empirical evidence with the model predictions. This is possible because our theory abstracts from wealth effects on labor supply, which is the key restriction for directly mapping the slopes of the local and aggregate Phillips curves (Hazell et al., 2022). Intuitively, the inflation sensitivity to changes in labor market conditions is the by-product of national and local general-equilibrium forces, but our empirical evidence uncovers only the latter. Thus, the presence of national general-equilibrium adjustments would impede to trace back the slope of the national Phillips curve in a regional setting. However, it turns out that the

[21]Throughout the transition, we also change the entry cost κ and searching cost λ_H to preserve the same unemployment and participation rates across steady states.

national general-equilibrium adjustment is captured by the presence of a consumption term, which disappears if we assume no wealth effects on labor supply. The lack of wealth effects on labor supply rules out the additional term due to national general-equilibrium forces, and allows us to directly map the slopes of the local and aggregate Phillips curves.[22]

To run this exercise, we take the simulated data on inflation, the unemployment rate, and robot adoption, and estimate a regression which is the aggregate national-level counterpart of the empirical specification of Equation (1):

$$\pi_t = \beta u_t + \gamma u_t \left(m_t - \bar{m}\right) + \varsigma m_t + \epsilon_t, \tag{30}$$

where β is the inflation sensitivity to unemployment when automation is at its average value, and γ informs on the inflation sensitivity varies with robot adoption. To estimate regression (30) in a way that can be compared to our empirical analysis, we convert the simulated data from quarterly into annual frequency.

We report the results in Table 3. The model replicates the estimated inflation sensitivity to unemployment and automation's effect at the city level. In the data, we find $\hat{\beta} = -0.5248$ and $\hat{\gamma} = 0.0058$. The model implies estimates of $\hat{\beta} = -0.5036$ and $\hat{\gamma} = 0.0058$. Thus, our economy almost perfectly matches both the inflation sensitivity to unemployment and the effect of automation of our empirical evidence. In contrast, if firms have close to full bargaining power ($\tau = 0.01$), the inflation sensitivity to unemployment shrinks, more than halving from -0.5059 to -0.1848. In this case, automation is not relevant for inflation dynamics: the estimate of the interaction term is reduced fivefold. This confirms the effect of automation on inflation dynamics hinges on the way in which robots interact with workers' bargaining power in altering wage setting.

[22] Appendix D.2 studies a version of the model with wealth effects of labor supply, and shows that this specification understates both the average inflation sensitivity to unemployment, as well as the effect of automation on this sensitivity. Yet, even in this case the effect of automation on inflation dynamics is empirically relevant.

Table 3: Robot Adoption and Inflation: Transitional Dynamics Exercise

	Dependent Variable: π_t		
	Baseline Model	Random Search $\tau = 0.01$	Data MSA-level
	(1)	(2)	(3)
u_t	-0.5036	-0.1819	-0.5248
	(0.0005)	(0.0001)	(0.1469)
$u_t \times (m_t - \bar{m})$	0.0058	0.0011	0.0058
	(0.0003)	(0.0001)	(0.0028)

Note: The table reports the estimates of the regression in Equation (30) based on simulated data of the baseline model (in Column 1) and the alternative model specification with random search and full bargaining power to firms (in Column 2). Column 3 reports the estimate based on a panel regression across U.S. MSAs which correspond to the case with non-tradable inflation as dependent variable and IV method of Table 1. The simulated data of the baseline model are based on a transitional dynamics exercise between a low-automation steady state, in which the robot-to-employee ratio is 0.2%, and a high-automation steady state, in which the robot-to-employee ratio is 0.8%.

Importantly, the model's ability to replicate the empirical effect of automation on the inflation–unemployment trade-off validates our modeling approach. It shows that capturing the joint influence of the stock of installed robots and workers' bargaining power is not only quantitatively relevant but essential to explain how automation shapes inflation dynamics. This interaction is the key channel through which the automation threat flattens the Phillips curve—highlighting the value of modeling both forces together, rather than in isolation.

What is the magnitude of automation's impact on the decoupling of inflation from labor market conditions? Between 1993 (the first year with U.S. robot data) and 2019, robot installations per employee rose by a factor of 7.7. According to Column (1) of Table 3, the model predicts a 9% ($\frac{7.7 \times 0.0058}{|-0.5059|}$) drop in the inflation sensitivity to unemployment due to automation. For context, Inoue et al. (2025) estimate a 68% decline in the price inflation sensitivity to unemployment over the five decades. Thus, the model suggests that the surge in automation accounts for a non-negligible part of the recent changes in inflation dynamics.

We replicate this exercise in a number of extensions of our baseline model:

Appendix D.4 adds physical capital, Appendix D.5 considers a two-agent setting, with hand-to-mouth and Ricardian households, and Appendix D.6 incorporates long-lasting job matches by allowing firms to operate for more than one period. In all cases, we confirm that automation does significantly flatten the inflation sensitivity to labor market conditions.

6 Conclusion

Leveraging geographical variation across U.S. MSAs, we uncover that automation decouples price inflation from unemployment, and this effect mainly hinges on a reduction in the wage sensitivity to changes in labor market conditions. We rationalize these facts through a New Keynesian model with search frictions in the labor market and robot adoption. In this economy, increasing automation to an amount that replicates the variation in robot penetration across MSAs reduces the slope of the price and wage Phillips curve by 15%. The key channel is the option of automatizing labor tasks, which alters workers' effective bargaining power, leading to a muted wage sensitivity to unemployment.

References

Acemoglu, Daron, and Pascual Restrepo. (2018). "The race between man and machine: Implications of technology for growth, factor shares, and employment". *American Economic Review*, 108(6), pp. 1488-1542. https://doi.org/10.2139/ssrn.2781320

Acemoglu, Daron, and Pascual Restrepo. (2020). "Robots and jobs: Evidence from US labor markets". *Journal of Political Economy*, 128(6), pp. 2188-2244. https://doi.org/10.2139/ssrn.2940245

Ball, Laurence, and Sandeep Mazumder. (2011). "Inflation dynamics and the Great Recession". *Brookings Papers on Economic Activity*, pp. 337-381. https://doi.org/10.2139/ssrn.1833648

Basso, Henrique S., and Juan F. Jimeno. (2021). "From secular stagnation to robocalypse? Implications of demographic and technological changes". *Journal of Monetary Economics*, 117, pp. 833-847. https://doi.org/10.1016/j.jmoneco.2020.06.004

Basso, Henrique S., and Omar Rachedi. (2021). "The young, the old, and the government: Demographics and fiscal multipliers". *American Economic Journal: Macroeconomics*, 13(4), pp. 110-141. https://doi.org/10.1257/mac.20190174

Beraja, Martin, Erik Hurst and Juan Ospina. (2019). "The aggregate implications of regional business cycles". *Econometrica*, 87(6), pp. 1789-1833. https://doi.org/10.3982/ecta14243

Beraja, Martin, and Nathan Zorzi. (2025). "Inefficient automation". *Review of Economic Studies*, 92, pp. 69-96. https://doi.org/10.1093/restud/rdae019

Bergholt, Drago, Francesco Furlanetto and Etienne Vaccaro-Grange. (2024). "Did monetary policy kill the Phillips curve? Some simple arithmetics". *The Review of Economics and Statistics*, pp. 1-45. https://doi.org/10.1162/rest_a_01524

Blanchard, Olivier. (2016). "The Phillips Curve: Back to the '60s?". *American Economic Review*, 106(5), pp. 31-34. https://doi.org/10.1257/aer.p20161003

Blundell, Richard, Antoine Bozio and Guy Laroque. (2011). "Labor supply and the extensive margin". *American Economic Review*, 101(3), pp. 482-486. https://doi.org/10.1257/aer.101.3.482

Bonhomme, Stéphane, Kerstin Holzheu, Thibaut Lamadon, Elena Manresa, Magne Mogstad and Bradley Setzler. (2023). "How much should we trust estimates of firm effects and worker sorting?". *Journal of Labor Economics*, 41(2), pp. 291-322. https://doi.org/10.1086/720009

Clarida, Richard, Jordi Gali and Mark Gertler. (2000). "Monetary policy rules and macroeconomic stability: Evidence and some theory". *The Quarterly Journal of Economics*, 115(1), pp. 147-180. https://doi.org/10.1162/003355300554692

Coibion, Olivier, and Yuriy Gorodnichenko. (2015). "Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation". *American Economic Journal: Macroeconomics*, 7(1), pp. 197-232. https://doi.org/10.1257/mac.20130306

De Loecker, Jan, Jan Eeckhout and Gabriel Unger. (2020). "The rise of market power and the macroeconomic implications". *The Quarterly Journal of Economics*, 135(2), pp. 561-644. https://doi.org/10.1093/qje/qjz041

Del Negro, Marco, Michele Lenza, Giorgio E. Primiceri and Andrea Tambalotti. (2020). "What's up with the Phillips Curve?". *Brookings Papers on Economic Activity*, Spring, pp. 301-357. https://doi.org/10.1353/eca.2020.0004

Faccini, Renato, and Leonardo Melosi. (2025). "Job-to-job mobility and inflation". *Review of Economics and Statistics*, 107(4), pp. 1027-1041. https://doi.org/10.1162/rest_a_01312

Fitzgerald, Terry, Callum Jones, Mariano Kulish and Juan Pablo Nicolini. (2024). "Is there a stable relationship between unemployment and future inflation?". *American Economic Journal: Macroeconomics*, 16(4), pp. 114-142. https://doi.org/10.1257/mac.20220273

Forbes, Kristin. (2019). "Inflation dynamics: Dead, dormant, or determined abroad?". *Brookings Papers on Economic Activity*, 2019(2), pp. 257-338. https://doi.org/10.1353/eca.2019.0015

Fornaro, Luca, and Martin Wolf. (2022). "Monetary policy in the age of automation". Mimeo.

Galesi, Alessandro, and Omar Rachedi. (2019). "Services deepening and the transmission of monetary policy". *Journal of the European Economic Association*, 17(4), pp. 1261-1293. https://doi.org/10.1093/jeea/jvy041

Gilchrist, Simon, Raphael Schoenle, Jae Sim and Egon Zakrajšek. (2017). "Inflation dynamics during the financial crisis". *American Economic Review*, 107(3), pp. 785-823. https://doi.org/10.1257/aer.20150248

Graetz, Georg, and Guy Michaels. (2018). "Robots at work". *Review of Economics and Statistics*, 100(5), pp. 753-768. https://doi.org/10.1162/rest_a_00754

Harding, Martin, Jesper Lindé and Mathias Trabandt. (2022). "Resolving the missing deflation puzzle". *Journal of Monetary Economics*, 126, pp. 15-34. https://doi.org/10.1016/j.jmoneco.2021.09.003

Hazell, Jonathon, Juan Herreño, Emi Nakamura and Jón Steinsson. (2022). "The slope of the Phillips curve: Evidence from US states". *The Quarterly Journal of Economics*, 137(3), pp. 1299-1344. https://doi.org/10.1093/qje/qjac010

Heise, Sebastian, Fatih Karahan and Ayşegül Şahin. (2022). "The missing inflation puzzle: The role of the wage-price pass-through". *Journal of Money, Credit and Banking*, 54(S1), pp. 7-51. https://doi.org/10.1111/jmcb.12896

Inoue, Atsushi, Barbara Rossi and Yiru Wang. (2025). "Has the Phillips curve flattened?". Mimeo.

Justiniano, Alejandro, and Giorgio E. Primiceri. (2008). "The time-varying volatility of macroeconomic fluctuations". *American Economic Review*, 98(3), pp. 604-641. https://doi.org/10.1257/aer.98.3.604

Leduc, Sylvain, and Zheng Liu. (2024). “Automation, bargaining power, and labor market fluctuations”. *American Economic Journal: Macroeconomics*, 16(4), pp. 311-349. https://doi.org/10.1257/mac.20220181

Lombardi, Marco J., Marianna Riggi and Eliana Viviano. (2023). “Workers’ bargaining power and the Phillips curve: A micro-macro analysis”. *Journal of the European Economic Association*, 21(5), pp. 1905-1943. https://doi.org/10.1093/jeea/jvad016

McLeay, Michael, and Silvana Tenreyro. (2020). “Optimal inflation and the identification of the Phillips curve”. *NBER Macroeconomics Annual*, 34(1), pp. 199-255. https://doi.org/10.1086/707181

Nakamura, Emi, and Jón Steinsson. (2014). “Fiscal stimulus in a monetary union: Evidence from US regions”. *American Economic Review*, 104(3), pp. 753-792. https://doi.org/10.1257/aer.104.3.753

Petrongolo, Barbara, and Christopher A. Pissarides. (2001). “Looking into the black box: A survey of the matching function”. *Journal of Economic Literature*, 39(2), pp. 390-431. https://doi.org/10.1257/jel.39.2.390

Pissarides, Christopher A. (2000). *Equilibrium Unemployment Theory*. MIT Press.

Rubbo, Elisa. (2023). “Networks, Phillips curves, and monetary policy”. *Econometrica*, 91(4), pp. 1417-1455. https://doi.org/10.3982/ecta18654

Siena, Daniele, and Riccardo Zago. (2024). “Job polarization, labor market fluidity and the flattening of the Phillips curve”. *Economic Journal*, 134(661), pp. 2141-2174. https://doi.org/10.1093/ej/ueae006

Stansbury, Anna, and Lawrence H. Summers. (2020). “The declining worker power hypothesis”. *Brookings Papers on Economic Activity*, Spring, pp. 1-77. https://doi.org/10.1353/eca.2020.0000

Stock, James H., and Mark W. Watson. (2020). “Slack and cyclically sensitive inflation”. *Journal of Money, Credit and Banking*, 52(S2), pp. 393-428. https://doi.org/10.1111/jmcb.12757

A Empirical Evidence: Robustness

This section provides additional information on our empirical evidence by presenting the results associated with alternative specifications. Moreover, we evaluate the robustness of our empirical findings as well as corroborate the validity of our identification strategy by reporting a comprehensive battery of checks.

A.1 Descriptive Statistics and First-Stage Regressions

We report the descriptive statistics of the key set of variables of our empirical evidence: tradable price inflation, total CPI inflation, wage inflation, the unemployment rate, and robot adoption. We show the average value and standard deviation in Table A.1.

To study the relevance of our instruments, we report in Table A.2 (see at the end of the section) the first-stage regressions for each of the three endogenous variables—unemployment rate, robot adoption, and their interaction—separately, using either only the excluded instruments motivated by the endogenous variable of interest, or also the full set of instruments. The results establish the explanatory power of the instruments, both when considered in isolation and when featured all together.

Table A.1: Descriptive Statistics

Variable	Mean	Standard Deviation	N. Observations
Tradable Price Inflation	-0.18%	10.05%	3,205
Price Inflation	-0.14%	7.79%	3,205
Wage Inflation	0.15%	8.55%	3,205
Unemployment Rate	6.85%	2.99%	3,205
Robot Adoption	0.68%	15.54%	3,205

Note: The table reports the mean and standard deviation of the key variables of our empirical evidence: tradable price inflation, total CPI inflation, wage inflation, the unemployment rate, and robot adoption.

A.2 Additional Results

This section shows that our findings do not change if we alter either the model specifications or the instrumenting strategy. We start by estimating the main regression without demeaning robot adoption. Thus, regression (1) becomes

$$\pi_{i,t} = \beta\, u_{i,t-1} + \gamma\, u_{i,t-1} m_{i,t-1} + \mathbf{X'_{i,t-1}}\theta + \alpha_i + \delta_t + \epsilon_{i,t}, \tag{A.1}$$

so that $m_{i,t-1}$ enters the regression in levels, and not in differences from its sample mean. As in the baseline regression, we consider as the independent variable (i) non-tradable price inflation, (ii) total CPI inflation, and (iii) wage inflation. We report the results in Table A.3. We confirm that the estimate of the interaction term does not change with respect to the baseline regression.[1] In other words, the demeaning does not alter at all the estimate of the interaction term. The only implication of the demeaning is the value of the estimated coefficient associated with the unemployment rate, β, as it changes its interpretation. While in the baseline regression, the coefficient β describes the inflation sensitivity to the unemployment rate for a MSA with an average degree of automation, that is when $m_{i,t-1} = \bar{m}$, in the no demeaning regression β captures the inflation sensitivity to the unemployment rate for a MSA with no robots, that is when $m_{i,t-1} = 0$.

Then, we consider a version of regression (1) in which we instrument the unemployment rate with the tradable demand spillovers variable, but do not instrument robot adoption. We report the results in Columns (1) and (2) of Table A.4. Since Column (1) is based on the OLS estimates, its results coincide with Column (1) of Table 1. When looking at Column (2), we find that, as in the baseline, the IV estimates of both the average inflation sensitivity and its interaction with robot adoption are substantially larger than the OLS estimates. This result corroborates the difference in the magnitude between the IV and OLS estimates is not due to the instrumenting of robots. Instead, as explained by McLeay and Tenreyro (2020), our

[1] As shown in Basso and Rachedi (2021), demeaning the robot-per-employee ratio, $m_{i,t-1} - \bar{m}$, in the interaction does not alter the estimation of the interaction term γ. This normalization allows us to directly interpret the parameter β as the sensitivity for a MSA with the average degree of automation, that is, when $m_{i,t} = \bar{m}$.

instrumenting strategy absorbs the role of the monetary policy targeting rule, which blurs the relevance of the Phillips curve.

Columns (3) and (4) of Table A.4 address the endogeneity of robot adoption. Rather than instrumenting local robot installations with those predicted by the robot adoption of foreign countries, we consider a regression in which robot adoption is kept fixed at its initial 2008 levels. Specifically, we modify regression (1) as follows:

$$\pi_{i,t} = \beta\, u_{i,t-1} + \gamma\, u_{i,t-1}\,(m_i - \tilde{\bar{m}}) + \mathbf{X}'_{\mathbf{i,t-1}}\theta + \alpha_i + \delta_t + \epsilon_{i,t}, \tag{A.2}$$

where m_i is robot adoption as of 2008, and $\tilde{\bar{m}} = \frac{1}{n}\sum_i m_i$ is its average value computed across MSAs. In this setting, the regression does not include the robot adoption variable in isolation in the covariates $\mathbf{X}_{\mathbf{i,t-1}}$, since it is absorbed by the MSA fixed effect as the values of robot adoption are kept constant for each city at its 2008 levels. We find that the interaction between unemployment and robot adoption is still negative and highly statistically significant, implying that the effect of automation on the inflation sensitivity to unemployment is driven by the heterogeneity across MSAs in robot adoption rather than from the cross-section of the year-to-year variation in robot installations.

We address the concern that both tradable demand spillovers and robot adoption are shift-share variables that leverage the industry composition at the city level. We then consider a different identification strategy and instrument the unemployment rate with changes in defense government spending, in the spirit of McLeay and Tenreyro (2020). Specifically, we follow Nakamura and Steinsson (2014) and define changes in defense government spending at the city level as the product of the share of defense spending in the GDP of a city (computed between 2000 and 2005) and the change in national real defense spending from 2008 to 2018. The results in Table A.5 confirm that automation reduces the price and wage inflation sensitivity to unemployment even under this identification strategy.

We modify regression (1) in three additional ways. First, we add the lagged value of the non-tradable inflation rate as an additional control. Second, we abstract from the lags in the independent variables and consider their contemporaneous values, leading to the following regression:

$$\pi_{i,t} = \beta\, u_{i,t} + \gamma\, u_{i,t}\,(m_{i,t} - \bar{m}) + \mathbf{X'_{i,t}}\theta + \alpha_i + \delta_t + \epsilon_{i,t}. \tag{A.3}$$

Then, we capture the forward-looking component of the local Phillips curve following Hazell et al. (2022) as follows:

$$\pi_{i,t} = \beta \sum_{j=0}^{3} u_{i,t+j} + \gamma \sum_{j=0}^{3} u_{i,t+j}\,(m_{i,t-1} - \bar{m}) + \mathbf{X'_{i,t-1}}\theta + \alpha_i + \delta_t + \epsilon_{i,t}, \tag{A.4}$$

which replaces the lagged value of the unemployment rate with the cumulative discounted sum of its future realizations. We set the discount parameter to $\beta = 0.995$ in line with the calibration of our model, and choose to account for realized unemployment up to three years to strike a where $\mathbb{I}\{\text{High Automation}\}_{i,t}$ is an indicator function that equals one if the change in the robot-to-employee ratio in a given city in a given year, $\Delta m_{i,t} \equiv m_{i,t} - m_{i,t-1}$, is above the median value of the robot adoption change in our sample. We report the results of this exercise in Table A.7, which confirm that while the interaction term between unemployment and robot adoption on average is still positive and statistically significant, there is an additional decoupling effect from automation into the relationship between inflation and unemployment that is concentrated in cities-year with high surges in automation: the estimates of both β_2 and γ_2 are positive and statistically significant.

Finally, we test whether the effect of automation in decoupling inflation from unemployment is pronounced in the cities-year associated with a relatively large surge in robot adoption. To do so, we run the following regression

$$\begin{aligned} \pi_{i,t} = \;& \beta_1\, u_{i,t-1} + \beta_1\, u_{i,t-1} \times \mathbb{I}\{\text{High Automation}\}_{i,t-1} + \dots \\ & \dots + \gamma_1\, u_{i,t-1}\,(m_{i,t-1} - \bar{m}) + \gamma_2\, u_{i,t-1}\,(m_{i,t-1} - \bar{m}) \times \mathbb{I}\{\text{High Automation}\}_{i,t-1} + \dots \\ & \dots + \mathbf{X'_{i,t-1}}\theta + \alpha_i + \delta_t + \epsilon_{i,t}, \end{aligned} \tag{A.5}$$

where $\mathbb{I}\{\text{High Automation}\}_{i,t}$ is an indicator function that equals one if the change in the robot-to-employee ratio in a given city in a given year, $\Delta m_{i,t} \equiv m_{i,t} - m_{i,t-1}$, is above the median value of the robot adoption change in our sample. We report the results of this exercise in Table A.7, which confirm that while the interaction term between unemployment and robot adoption on average is still positive and statistically significant, there is an additional decoupling effect from automation into the relationship between inflation and unemployment that is concentrated in city-year with high surges in automation: the estimates of both β_2 and γ_2 are positive and statistically significant.

A.3 Fixed Effects and Error Correlation

We consider exercises which explicitly study the role of any cross-section correlation across metropolitan areas that is not removed by the combination of the MSA and year fixed effects, as well as the clustering of errors at the two-way MSA-year level.

We start by substituting the year fixed effects with the more granular state-year fixed effects. This approach allows us to account for any comovement across MSAs within the same state during the same year. We then consider two additional cases that combine the year fixed effects with either deciles of the share of manufacturing GDP over total GDP or deciles of the share of tradable industries GDP over total GDP. In this way, we can absorb any role of the time variation in the sectoral composition of the economy in driving the inflation sensitivity to unemployment (Galesi and Rachedi, 2019; Rubbo, 2023). Finally, we use information from the Statistics of U.S. Businesses of the U.S. Census Bureau to compute the employment share of large firms, defined as firms with more than 500 employees, and again compute deciles of this share. This case allows us to account for the possibility that automation is entirely driven by large firms, and thus what we are capturing with the interaction between unemployment and robot adoption is something that has to do with the relevance of large firms on aggregate price setting. The results in Table A.8 indicate that the interaction term between automation and unemployment is always statistically different from zero.

Table A.9 considers five additional specifications. First, we consider a less restricting setting in which the clustering of the standard errors is at the MSA level, rather than the two-way MSA-year clustering. Second, we introduce two-way state-year clustering of the standard errors, rather than a two-way MSA-year clustering, capturing any correlation across MSAs within the same state. Third, we estimate the regression with Driscoll and Kraay (1998) standard errors, which are robust to general forms of spatial and temporal dependence. Fourth, we estimate the standard errors by imposing a common factor structure as in Norkutė et al. (2021), which absorbs the common time-variation across MSAs. Again, we find that the role of robot adoption in dampening the local sensitivity of inflation to unemployment is always statistically significant. Finally, we compute standard errors as in Adao et al. (2019) and Borusyak et al. (2022), which are robust to the presence of unobserved shocks that may affect a sub-group of industries. Specifically, the standard errors are computed by taking into account a matrix of regional industry shares that drives the correlation structure of the error terms. While this procedure widens the standard errors with respect to the baseline specification, we still find a statistically significant role for automation in muting the inflation sensitivity to the unemployment rate.

A.4 Alternative Channels

This section examines to what extent our findings keep holding when accounting for the role of potential alternative explanations for the decoupling of inflation and unemployment dynamics, including variables that could be highly correlated with the surge of automation. To do so, we estimate a sequence of additional regressions in which we introduce each time a new key potential confounding factor and we explicitly control for both its local lagged level and its interaction with the local unemployment rate. In this way, we can evaluate whether the effect of automation on inflation dynamics keeps holding above and beyond the interaction that unemployment may have with other MSA-level characteristics.

Our first set of potential alternative explanations relates to heterogeneity in demographic characteristics across metropolitan areas. We merge our data with information from the Current Population Survey (CPS) of the U.S. Census Bureau, and compute for each metropolitan area the following characteristics: (*i*) the share of young people in total population, defined as the share of individuals whose age is below 30 years, (*ii*) the share of old people in total population, defined as the share of individuals whose age is above 60 years, (*iii*) the female labor market participation, (*iv*) the black people labor market participation, (*v*) the Asian people labor market participation, (*vi*) the share of individuals with low educational attainments, defined as those people who have attended at most until the tenth grade, (*vii*) the overall labor market participation, and (*viii*) the average marginal propensity to consume (MPC). We proxy the latter in two ways. First, we follow **?** and combine the estimate of the MPC by demographic characteristics derived by Patterson (2023) with the share of each of these characteristics in each metropolitan area in each year of our sample. Overall, merging our initial data with the CPS information slightly reduces the total number of observations in our panel, from 3,205 to 2,270. Second, we proxy high MPCs with high debt-to-income ratios, as provided at the MSA level by the Consumer Credit Panel/Equifax data of the Federal Reserve Bank of New York.

We report the results of extending our baseline regression to include the lagged value of each of the above demographic characteristics—one at a time—both as its lagged values and its interaction with the unemployment rate in Tables A.10 and A.11. Overall, we find that the role of automation is always highly statistically significant and rather constant across the different specifications. These results also suggest that our baseline setting does not capture the relationship that automation has with the aging labor force (Basso and Jimeno, 2021; Acemoglu and Restrepo, 2022), and in turn the effect of the aging population on long-run inflation dynamics (Aksoy et al., 2019). The effect of automation holds also above and beyond the way in which differences in the MPC across metropolitan areas modulate the transmission of monetary policy, as documented by **?**.

The second set of confounding factors we consider is related to the heterogeneous variations in the content of occupations across metropolitan areas. Indeed, robot adoption has led to a decline in both routine and manual occupations (Acemoglu and Restrepo, 2018, 2020), a phenomenon which is intrinsically related to the job polarization emphasized by Autor et al. (2013). We evaluate the role of changes in the occupational structure as Siena and Zago (2024) show that the disappearance of routine and manual occupations is a potential explanation for the flattening of the price Phillips curve in the early 2000s. To document that the effect of automation on inflation dynamics holds above and beyond that of job polarization, we take the information on occupations provided by the CPS, and assign them to either manual, routine, abstract, and offshorable occupations based on the matching procedure of Autor et al. (2013). We report the results of extending our baseline regression to include the lagged value of each of the above occupational characteristics—one at a time—both as its lagged values and its interaction with the unemployment rate in Table A.12. Again, we find the effect of automation on inflation dynamics holds even when explicitly controlling for the time-variation in the occupational structure across metropolitan areas.

The third set of potential alternative explanations relates to the key role that foreign import competition has had on recent changes in inflation dynamics (Forbes, 2019; Heise et al., 2022; Amiti et al., 2023). Specifically, we consider to what extent the effect of automation on inflation could hold when including in our regressions the role of imports from China and Mexico, the two countries that have been providing the greatest competition threats to U.S. products. To do so, we closely follow the steps of Autor et al. (2013): we get import data from the United Nations Commodity Trade Statistics (UN Comtrade) on imports from China and Mexico at the 6 digit Harmonized System product level, we convert this information into 1987 four-digit SIC codes, and finally transform the information at the 1997 six-digit NAICS codes. We use the employment structure of each metropolitan area at the industry level to compute a time-varying measure of Chinese and Mexican import competition over the entire sample period, and merge it with our original data. We then report the results of extending our baseline regression to include the lagged value of each of the above import variables—either the imports from China, or the imports from Mexico, or the sum of imports from the two countries—both as its lagged values and its interaction with the unemployment rate in

Table A.13. We find that although imports did flatten the price Phillips curve, the effect of automation on inflation dynamics holds above and beyond the time-variation in import competition across metropolitan areas, and the magnitude of this effect is remarkably similar to the flattening implied by foreign import competition. For instance, the estimates in Column (2) of Table A.13 imply that an increase in robot adoption by one standard deviation reduces the sensitivity of prince inflation by 17%, while an increase in import competition by one standard deviation reduces the sensitivity of price inflation by 16%.

Finally, we consider a last group of key potential confounding factors. We start by considering the share of government value added in total GDP, the ratio of unemployment benefits over total personal income, and the ratio of income from dividends, interests, and rents over total personal income. Then, we also consider the employment share of large firms, computed as in Table A.8. The results are reported in Table A.14. Again, we find that the coefficient associated with the interaction between the unemployment rate and robot adoption is always statistically different from zero.

A.5 More on the Validation of Model Mechanism

Section 2.5 validates the main model mechanism in the data. Specifically, in the model, automation alters inflation dynamics through a wage-setting effect, through which robot adoption dampens workers' effective bargaining power. We validate this model mechanism in the data by extending the baseline empirical setting to test whether the effect of automation in decoupling inflation and unemployment depends on workers' bargaining power. We proxy workers' bargaining power with a measure of either the union membership rate or the union coverage rate for each MSA-year.

One potential concern is that the unionization rate could be endogenously related to the automation patterns. To address this concern, we consider an alternative specification of the validation regression (3) as follows:

$$\pi_{i,t} = \beta_1 \, u_{i,t-1} + \beta_2 \, u_{i,t-1} \mathbb{I}\{\text{HBP}\}_i + \gamma_1 \, u_{i,t-1} \, (m_{i,t-1} - \bar{m}) + \ldots \qquad \text{(A.6)}$$
$$\cdots + \gamma_2 \, u_{i,t-1} \, (m_{i,t-1} - \bar{m}) \, \mathbb{I}\{\text{HBP}\}_i + \mathbf{X'_{i,t-1}} \theta + \alpha_i + \delta_t + \epsilon_{i,t},$$

where $\mathbb{I}\{\text{HBP}\}_i$ is an indicator function that equals 1 for those MSAs with high workers' bargaining power in the year 2008, defined as those MSA in which either the union membership rate or the union coverage rate in the year 2008 is above the median in the sample. By fixing the value of the unionization rate to that observed in the first year of the sample, we rule out any potential endogenous reaction of the unionization rate to the installation of robots.

We report the results of this exercise in Table A.15, which confirm that the effect of automation on decoupling inflation from unemployment is concentrated in those MSAs with a high workers' bargaining power, independently of which of the two proxies we use. In this case, the statistical significance is slightly lower than in the baseline, but this is just the product of the fact that we keep the unionization rate fixed as of the value in the year 2008, and thus disregard and reduce substantially the variation in this key variable in the data. However, the main conclusion of this exercise coincides with that reported in Section 2.5: automation alters inflation dynamics by reducing workers' bargaining power.

We then consider potential alternative explanations to workers' bargaining power that could be behind the influence of automation on inflation dynamics. Specifically, we consider three alternative hypotheses. First, we consider the role of geographical differences in job flows. We take information on hires, separations, and job-to-job transition from the Job-to-Job Flows (J2J) of the U.S. Census Bureau, and take the ratio of all these flows relative to total private employment for each MSA-year pair. Second, we consider the role of capacity utilization, and proxy it with hours worked per employee as in Basu et al. (2006). We do so using information from the U.S. Bureau of Labor Statistics. Third, we consider the role of differences in the sectoral composition of the economy and consider the shares of either manufacturing industries or tradable industries in total local GDP. We get the information on these shares from the U.S. Bureau of Economic Analysis (BEA).

In all these three cases, we run exactly the same specification as in Section 2.5:

$$\pi_{i,t} = \beta_1\, u_{i,t-1} + \beta_2\, u_{i,t-1}\mathbb{I}\{\text{VAR}\}_{i,t-1} + \gamma_1\, u_{i,t-1}\,(m_{i,t-1} - \bar{m}) + \dots \tag{A.7}$$

$$\dots + \gamma_2\, u_{i,t-1}\,(m_{i,t-1} - \bar{m})\,\mathbb{I}\{\text{VAR}\}_{i,t-1} + \mathbf{X'_{i,t-1}}\theta + \alpha_i + \delta_t + \epsilon_{i,t},$$

with the only difference that we substitute the role of the bargaining power in regression (3) each time with a different variable, $\mathbb{I}\{\text{VAR}\}_{i,t-1}$, each one proxying one of the three alternative channels described above.

More precisely, consider the following cases. In the first group of alternative hypotheses, we define $\mathbb{I}\{\text{VAR}\}_{i,t-1}$ as an indicator that equals 1 for those MSA-year in which either the hiring rate, or the separation rate, or the job-to-job flow rate is above the median in the sample. That is, we have three cases that capture differences across areas in the dynamism of the labor market. In the second group of hypotheses, we define $\mathbb{I}\{\text{VAR}\}_{i,t-1}$ as an indicator that equals 1 for those MSA-year in which hours worked per employee is above the median in the sample. This case captures the role of differences in the utilization rate across cities. Finally, we define $\mathbb{I}\{\text{VAR}\}_{i,t-1}$ as an indicator that equals 1 for those MSA-year in which either the manufacturing share of the share of tradable industries is above the median in the sample. These two cases allow us to look into the role of differences in the geographical composition of the economy at the city level as a potential confounding factor behind the effect of automation on inflation dynamics. Then, we estimate all regressions using the same IV method as in the regression (3).

We report the results of these six regressions in Table A.16. We find that although all these channels matter for inflation dynamics, as their interaction with the unemployment rate is statistically different from zero (but for the case of the manufacturing share), they do shape how automation influences the local Phillips curve. Indeed, while the double interaction of the unemployment rate and the robot adoption variable is always positive and statistically significant, the triple interaction with each of these additional variables is not.

Thus, this exercise showcases that while automation does interact with the degree of unionization in shaping the local inflation sensitivity to the unemployment rate, there is no evidence suggesting that the variables that capture geographical differences in labor market flows, capacity utilization, and the sectoral composition of the economy may be behind the role of automation. From this perspective, we confirm the key role of workers' bargaining power in modulating how robot installations impact inflation dynamics.

Table A.2: Robot Adoption and Inflation across MSAs - First-Stage Regressions

Dependent Variable:	Unemployment Rate		Robot Adoption		Unemployment Rate × Robot Adoption	
	(1)	(2)	(3)	(4)	(5)	(6)
Tradable Demand Spillovers	16.8903***	15.9390***		8.6557		16.9700***
	(2.0800)	(2.1139)		(6.5892)		(8.9246)
EU Robot Penetration		-0.0938	6.8690***	2.7433**		51.4628**
		(0.2130)	(1.2886)	(1.2730)		(22.7917)
Tradable Demand Spillovers × EU Robot Penetration		0.0538		-0.5497	10.0544***	16.3056***
		(0.0360)		(0.3456)	(2.7199)	(5.1356)
Year Fixed Effects	✓	✓	✓	✓	✓	✓
MSA Fixed Effects	✓	✓	✓	✓	✓	✓
N. Observations	3,205	3,205	3,205	3,205	3,205	3,205

Note: The table reports the first-stage regressions associated with the empirical setting of Table 1. Columns (1), (3), and (5) report the results the first-stage regressions for each of the three endogenous variables—unemployment rate, robot adoption, and their interaction—separately, using only the excluded instruments motivated by the endogenous variable of interest. Columns (2), (4), and (6) report analogous results focusing on the full set of instruments. The instrument of unemployment is a shift-share variable that captures tradable demand spillovers. The instrument of robot adoption is a measure of robot penetration in a pool of European countries. All regressions include year and MSA fixed effects. Columns (2), (4), and (6) also feature the lagged value of the relative price of tradable goods. Two-way clustered standard errors are reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10%, respectively.

Table A.3: Robot Adoption and Inflation across MSAs - No Robot Adoption Demeaning

Dependent Variable:	Non-tradable Inflation		CPI Inflation		Wage Inflation	
	OLS (1)	IV (2)	OLS (3)	IV (4)	OLS (5)	IV (6)
$u_{i,t-1}$	-0.1970***	-0.5743***	-0.1331***	-0.3118***	-0.3986***	-1.0296***
	(0.0227)	(0.1388)	(0.0162)	(0.0815)	(0.0340)	(0.2542)
$u_{i,t-1} \times m_{i,t-1}$	0.0010**	0.0058**	0.0007***	0.0028*	0.0016**	0.0060**
	(0.0004)	(0.0028)	(0.0003)	(0.0016)	(0.0007)	(0.0028)
Year Fixed Effects	✓	✓	✓	✓	✓	✓
MSA Fixed Effects	✓	✓	✓	✓	✓	✓
N. Observations	3,205	3,205	3,205	3,205	3,205	3,205

Note: The table reports the estimates of panel regressions similar to that of Table 1 with the only difference that we do not demean the robot adoption variable with its sample average. Two-way clustered standard errors are reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10%, respectively.

Table A.4: Robot Adoption and Inflation across MSAs - The Role of Robot Adoption

	Dependent Variable: Non-tradable Inflation			
	No Instrumenting for Robots		Fixed Robots as of 2008	
	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)
$u_{i,t-1}$	-0.1884***	-0.6740***	-0.1946***	-0.6958**
	(0.0221)	(0.1356)	(0.0222)	(0.1329)
$u_{i,t-1} \times (m_i - \bar{m})$	0.0010**	0.0050**	0.0016***	0.0050***
	(0.0004)	(0.0021)	(0.0041)	(0.0012)
Year Fixed Effects	✓	✓	✓	✓
MSA Fixed Effects	✓	✓	✓	✓
N. Observations	3,205	3,205	3,200	3,200

Note: The table reports the estimates of panel regressions similar to that of Table 1 with the only difference that Column (2) instruments only the unemployment rate with the shift-share variable that captures tradable demand spillovers but does not instrument the robot adoption variable. Then, Column (3) and (4) fix the values of robot adoption as the robot-per-employee ratio as of 2008, m_i. In all columns, the dependent variable is the non-tradable good inflation rate, and the regressions feature the lagged value of the relative price of non-tradable goods, as well as year and MSA fixed effects. Columns (1) and (3) reports the OLS results, and Columns (2) and (4) report the IV results in which the unemployment rate is instrumented with a shift-share variable that captures tradable demand spillovers. Two-way clustered standard errors are reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10%, respectively.

Table A.5: Robot Adoption and Inflation across MSAs - Government Spending as Instrument

Dependent Variable:	Non-tradable Inflation	CPI Inflation	Wage Inflation
	IV	IV	IV
	(1)	(2)	(3)
$u_{i,t-1}$	-0.7403**	-0.4453**	-0.5571***
	(0.3383)	(0.2017)	(0.2496)
$u_{i,t-1} \times (m_{i,t-1} - \bar{m})$	0.0059**	0.0024*	0.0010**
	(0.0026)	(0.0017)	(0.0006)
Year Fixed Effects	✓	✓	✓
MSA Fixed Effects	✓	✓	✓
N. Observations	2,795	2,795	2,795

Note: The table reports the estimates of panel regressions similar to that of Table 1 with the difference that the unemployment rate is instrumented by the change in real defense spending, computed as the product of the share of defense spending in GDP and the change in national real defense spending. The dependent variable is the inflation rate of non-tradable goods in Column (1), CPI inflation in Column (2), and wage inflation in Column (3). All cases are estimated using instrumental variables. Two-way clustered standard errors are reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10%, respectively.

Table A.6: Robot Adoption and Inflation across MSAs - Alternative Specifications

	Dependent Variable: Non-tradable Inflation					
	Lagged Inflation Control		Contemporaneous Values		Forward Looking	
	OLS (1)	IV (2)	OLS (3)	IV (4)	OLS (5)	IV (6)
$u_{i,t-1}$	-0.1599*** (0.0222)	-0.5255*** (0.1431)				
$u_{i,t-1} \times (m_{i,t-1} - \bar{m})$	0.0006 (0.0005)	0.0046** (0.0026)				
$u_{i,t}$			-0.1788*** (0.0209)	-0.5594*** (0.1585)		
$u_{i,t} \times (m_{i,t} - \bar{m})$			0.0002 (0.0006)	0.0064** (0.0022)		
$\sum_{j=0}^{3} \beta^j u_{i,t+j}$					-0.0258*** (0.0079)	-0.0651 (0.0546)
$\sum_{j=0}^{3} \beta^j u_{i,t+j} \times (m_{i,t-1} - \bar{m})$					0.0006** (0.0003)	0.0035** (0.0016)
Lagged Inflation Control	✓	✓	✗	✗	✗	✗
Year Fixed Effects	✓	✓	✓	✓	✓	✓
MSA Fixed Effects	✓	✓	✓	✓	✓	✓
N. Observations	2,884	2,884	3,205	3,205	2,563	2,563

Note: The table reports the estimates of panel regressions similar to that of Table 1 in which either we add the lagged value of the non-tradable inflation rate as an additional control, in Columns (1) and (2), or we substitute all lagged independent variables with their contemporaneous values, in Columns (3) and (4), or we substitute the lagged value of unemployment with the cumulative discounted sum of realized unemployment from the year contemporaneous to that of the dependent variable up to three years in the future, in Columns (5) and (6). In all columns, the dependent variable is the non-tradable good inflation rate, and the regressions feature the (lagged) value of the relative price of non-tradable goods, as well as year and MSA fixed effects. Columns (1), (3), and (5) report the OLS results, and Columns (2), (4), and (6) report the IV results in which the unemployment rate term is instrumented with a shift-share variable that captures tradable demand spillovers. Two-way clustered standard errors are reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10%, respectively.

Table A.7: Robot Adoption and Inflation across MSAs - Effect of High Automation

	Dependent Variable: Non-tradable Inflation	
	OLS	IV
	(1)	(2)
$u_{i,t-1}$	-0.2033***	-0.6705***
	(0.0237)	(0.1553)
$u_{i,t-1} \times \mathbb{I}\{\text{High Automation}\}_{i,t-1}$	0.0129*	0.0345***
	(0.0072)	(0.0125)
$u_{i,t-1} \times (m_{i,t-1} - \bar{m})$	0.0007**	0.0030
	(0.0003)	(0.0025)
$u_{i,t-1} \times (m_{i,t-1} - \bar{m}) \times \mathbb{I}\{\text{High Automation}\}_{i,t-1}$	0.0007**	0.0060*
	(0.0003)	(0.0036)
Year Fixed Effects	✓	✓
MSA Fixed Effects	✓	✓
N. Observations	3,205	3,205

Note: The table reports the estimates of panel regressions similar to that of Table 1 with the difference that it adds the interaction with those cities-years in which the growth rate in robot adoption has been above the median value in our sample. Column (1) is estimated using OLD methods and Column (2) is estimated using instrumental variables. In both columns, the dependent variable is the non-tradable good inflation rate. All regressions also include the lagged values of robot adoption, the relative price of non-tradable goods, year and MSA fixed effects. Two-way clustered standard errors are reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10%, respectively.

Table A.8: Robot Adoption and Inflation across MSAs - Alternative Fixed Effects

	Dependent Variable: Non-tradable Inflation			
	State-Year Fixed Effects	Deciles Manufacturing GDP Share-Year Fixed Effects	Deciles Tradable GDP Share-Year Fixed Effects	Deciles Large Firms Employment Share-Year Fixed Effects
	IV	IV	IV	IV
	(1)	(2)	(3)	(4)
$u_{i,t-1}$	-0.5127**	-0.7659***	-0.3206**	-0.6887***
	(0.2176)	(0.1668)	(0.1626)	(0.1646)
$u_{i,t-1} \times (m_{i,t-1} - \bar{m})$	0.0074**	0.0072**	0.0064**	0.0047**
	(0.0029)	(0.0032)	(0.0032)	(0.0024)
MSA Fixed Effects	✓	✓	✓	✓
N. Observations	3,185	3,205	3,205	3,205

Note: The table reports the estimates of panel regressions similar to that of Table 1 with the difference that Column (1) replaces the year fixed effects with state-year fixed effects, Column (2) replaces the year fixed effects with the interaction of year fixed effects with deciles of the share of manufacturing GDP in total GDP, Column (3) replaces the year fixed effects with the interaction of year fixed effects with deciles of the share of tradable industries GDP in total GDP, and Column (4) replaces the year fixed effects with the interaction of year fixed effects with deciles of the employment share of large firms in total employment. In all columns, the dependent variable is the non-tradable good inflation rate, and all cases are estimated with IV methods, in which the unemployment rate is instrumented with a shift-share variable that captures tradable demand spillovers, and the robot-adoption variable is instrumented with the industry-level robot penetration in a pool of European countries. All regressions also include the lagged values of robot adoption, the relative price of non-tradable goods, and MSA fixed effects. Two-way clustered standard errors at the MSA-year level are reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10%, respectively.

Table A.9: Robot Adoption and Inflation across MSAs - Alternative Clustering Choices

	Dependent Variable: Non-tradable Inflation				
	MSA Clustering	State-Year Clustering	Driscoll-Kraay Errors	Factor-Structure Errors	Cluster-Robust Errors
	IV	IV	IV	IV	IV
	(1)	(2)	(3)	(4)	(5)
$u_{i,t-1}$	-0.5248***	-0.5248***	-0.5248*	-0.5382***	-0.5382***
	(0.1190)	(0.1878)	(0.2761)	(0.1770)	(0.1865)
$u_{i,t-1} \times (m_{i,t-1} - \bar{m})$	0.0058***	0.0058*	0.0058**	0.0200*	0.0058**
	(0.0021)	(0.0033)	(0.0021)	(0.0115)	(0.0029)
Year Fixed Effects	✓	✓	✓		✓
MSA Fixed Effects	✓	✓	✓	✓	✓
N. Observations	3,205	3,205	3,205	3,205	3,205

Note: The table reports the estimates of panel regressions similar to that of Table 1 with the difference that Column (1) replaces the clustering of standard errors at the two-way MSA-year level with a clustering at only the MSA level, Column (2) replaces the clustering of standard errors at the two-way MSA-year level with a clustering at the two-way state-year level, Column (3) considers Driscoll-Kraay standard errors, Column (4) allows for a factor-structure in the errors, and Column (5) considers cluster-robust errors as in Adao et al. (2019) and Borusyak et al. (2022). In all columns, the dependent variable is the non-tradable good inflation rate, and all cases are estimated with IV methods, in which the unemployment rate is instrumented with a shift-share variable that captures tradable demand spillovers, and the robot-adoption variable is instrumented with the industry-level robot penetration in a pool of European countries. All regressions also include the lagged values of robot adoption, the relative price of non-tradable goods, as well as year and MSA fixed effects. Two-way clustered standard errors at the MSA-year level are reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10%, respectively.

Table A.10: Robot Adoption and Inflation across MSAs - The Role of Demographics

	Dependent Variable: Non-tradable Inflation			
	Young People	Old People	MPC	Debt-to-Income
	IV	IV	IV	IV
	(1)	(2)	(3)	(4)
$u_{i,t-1}$	-0.5942***	-0.6016***	-0.6001***	-0.6534***
	(0.1511)	(0.1500)	(0.1498)	(0.2087)
$u_{i,t-1}\times$ $(m_{i,t-1} - \bar{m})$	0.0140***	0.0140***	0.0140***	0.0039*
	(0.0051)	(0.0051)	(0.0051)	(0.0022)
$u_{i,t-1}\times$ $\left(VAR_{i,t-1} - V\bar{A}R\right)$	-0.0402	-0.0326	-0.1412	0.1446**
	(0.0804)	(0.0685)	(0.2669)	(0.0700)
Year Fixed Effects	✓	✓	✓	✓
MSA Fixed Effects	✓	✓	✓	✓
N. Observations	2,270	2,270	2,270	2,899

Note: The table reports the estimates of panel regressions similar to that of Table 1 with the difference that we also include the interaction of the unemployment rate with a set of potential confounding factors one at a time, a term we refer to as $u_{i,t-1} \times \left(VAR_{i,t-1} - V\bar{A}R\right)$, where $VAR_{i,t-1}$ is the value that each of this additional confounding factors take in metropolitan area i at year t, and $V\bar{A}R$ is the associated average value in the sample. In all columns, the dependent variable is the non-tradable good inflation rate, and all cases are estimated with IV methods, in which the unemployment rate is instrumented with a shift-share variable that captures tradable demand spillovers, and the robot-adoption variable is instrumented with the industry-level robot penetration in a pool of European countries. All regressions also include the lagged values of robot adoption variable, the relative price of non-tradable goods, and the confounding variable used in the interaction term, as well as year and MSA fixed effects. Column (1) considers the role of the share of young people in total population, defined as those below 30 years old, Column (2) considers the role of the share of old people in total population, defined as those above 60 years old, Column (3) considers the role of workers marginal propensity to consume, and Column (4) considers the role of the debt-to-income ratio. Two-way clustered standard errors are reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10%, respectively.

Table A.11: Robot Adoption and Inflation across MSAs - The Role of Demographics

	Dependent Variable: Non-tradable Inflation				
	Female Labor Particip.	Black Labor Particip.	Asian Labor Particip.	Low Education	Labor Force Particip.
	IV	IV	IV	IV	IV
	(1)	(2)	(3)	(4)	(5)
$u_{i,t-1}$	-0.5997***	-0.5927***	-0.5889***	-0.6009***	-0.6039***
	(0.1381)	(0.1483)	(0.1463)	(0.1491)	(0.1488)
$u_{i,t-1} \times$ $(m_{i,t-1} - \bar{m})$	0.0143***	0.0138***	0.0140***	0.0140***	0.0136***
	(0.0051)	(0.0050)	(0.0051)	(0.0051)	(0.0050)
$u_{i,t-1} \times$ $\left(VAR_{i,t-1} - V\bar{A}R\right)$	-0.1496	-0.0672	0.1181	-0.0699	0.1975**
	(0.1127)	(0.0585)	(0.1485)	(0.0858)	(0.0806)
Year Fixed Effects	✓	✓	✓	✓	✓
MSA Fixed Effects	✓	✓	✓	✓	✓
N. Observations	2,270	2,270	2,270	2,270	2,270

Note: The table reports the estimates of panel regressions similar to that of Table 1 with the difference that we also include the interaction of the unemployment rate with a set of potential confounding factors one at a time, a term we refer to as $u_{i,t-1} \times \left(VAR_{i,t-1} - V\bar{A}R\right)$, where $VAR_{i,t-1}$ is the value that each of this additional confounding factors take in metropolitan area i at year t, and $V\bar{A}R$ is the associated average value in the sample. In all columns, the dependent variable is the non-tradable good inflation rate, and all cases are estimated with IV methods, in which the unemployment rate is instrumented with a shift-share variable that captures tradable demand spillovers, and the robot-adoption variable is instrumented with the industry-level robot penetration in a pool of European countries. All regressions also include the lagged values of robot adoption variable, the relative price of non-tradable goods, and the confounding variable used in the interaction term, as well as year and MSA fixed effects. Column (1) considers the role of the female labor participation, Column (2) considers the role of the labor participation of black workers, Column (3) considers the role of the labor participation of Asian workers, Column (4) considers the share of workers with low educational attainments, defined as those workers who have attended school up to grade 10, and Column (5) considers the role of the labor participation of all workers. Two-way clustered standard errors are reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10%, respectively.

Table A.12: Robot Adoption and Inflation across MSAs - The Role of Occupations

	Dependent Variable: Non-tradable Inflation			
	Abstract Occupations	Routine Occupations	Manual Occupations	Offshorable Occupations
	IV	IV	IV	IV
	(1)	(2)	(3)	(4)
$u_{i,t-1}$	-0.5897***	-0.5916***	-0.5964***	-0.5959***
	(0.1658)	(0.1660)	(0.1674)	(0.1664)
$u_{i,t-1}\times$ $(m_{i,t-1}-\bar{m})$	0.0139***	0.0140***	0.0143***	0.0139***
	(0.0051)	(0.0051)	(0.0052)	(0.0051)
$u_{i,t-1}\times$ $\big(VAR_{i,t-1}-V\bar{A}R\big)$	-0.0200*	0.0173	0.0029	0.0392
	(0.0119)	(0.0118)	(0.0220)	(0.0252)
Year Fixed Effects	✓	✓	✓	✓
MSA Fixed Effects	✓	✓	✓	✓
N. Observations	2,268	2,268	2,268	2,268

Note: The table reports the estimates of panel regressions similar to that of Table 1 with the difference that we also include the interaction of the unemployment rate with a set of potential confounding factors one at a time, a term we refer to as $u_{i,t-1}\times\big(VAR_{i,t-1}-V\bar{A}R\big)$, where $VAR_{i,t-1}$ is the value that each of this additional confounding factors take in metropolitan area i at year t, and $V\bar{A}R$ is the associated average value in the sample. In all columns, the dependent variable is the non-tradable good inflation rate, and all cases are estimated with IV methods, in which the unemployment rate is instrumented with a shift-share variable that captures tradable demand spillovers, and the robot-adoption variable is instrumented with the industry-level robot penetration in a pool of European countries. All regressions also include the lagged values of robot adoption, the relative price of non-tradable goods, and the confounding variable used in the interaction term, as well as year and MSA fixed effects. Column (1) considers the share of abstract occupations in total occupations, Column (2) considers the share of routine occupations in total occupations, Column (3) considers the share of manual occupations in total occupations, and Column (3) considers the share of offshorable occupations in total occupations. Two-way clustered standard errors are reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10%, respectively.

Table A.13: Robot Adoption and Inflation across MSAs - The Role of Import Competition

	Dependent Variable: Non-tradable Inflation		
	Chinese Imports	Mexican Imports	Chinese & Mexican Imports
	IV	IV	IV
	(1)	(2)	(3)
$u_{i,t-1}$	-0.5239***	-0.7713***	-0.6486***
	(0.1477)	(0.2223)	(0.1832)
$u_{i,t-1}\times$ $(m_{i,t-1}-\bar{m})$	0.0058**	0.0086***	0.0076**
	(0.0028)	(0.0031)	(0.0030)
$u_{i,t-1}\times$ $\left(VAR_{i,t-1}-V\bar{A}R\right)$	0.0110	0.4417***	0.1921**
	(0.0579)	(0.1364)	(0.0784)
Year Fixed Effects	✓	✓	✓
MSA Fixed Effects	✓	✓	✓
N. Observations	3,205	3,205	3,205

Note: The table reports the estimates of panel regressions similar to that of Table 1 with the difference that we also include the interaction of the unemployment rate with a set of potential confounding factors one at a time, a term we refer to as $u_{i,t-1} \times \left(VAR_{i,t-1} - V\bar{A}R\right)$, where $VAR_{i,t-1}$ is the value that each of this additional confounding factors take in metropolitan area i at year t, and $V\bar{A}R$ is the associated average value in the sample. In all columns, the dependent variable is the non-tradable good inflation rate, and all cases are estimated with IV methods, in which the unemployment rate is instrumented with a shift-share variable that captures tradable demand spillovers, and the robot-adoption variable is instrumented with the industry-level robot penetration in a pool of European countries. All regressions also include the lagged values of robot adoption, the relative price of non-tradable goods, and the confounding variable used in the interaction term, as well as year and MSA fixed effects. Column (1) considers the share of abstract occupations in total occupations, Column (2) considers the share of routine occupations in total occupations, Column (3) considers the share of manual occupations in total occupations, and Column (3) considers the share of offshorable occupations in total occupations. Two-way clustered standard errors are reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10%, respectively.

Table A.14: Robot Adoption and Inflation across MSAs - Additional Key Covariates

	Dependent Variable: Non-tradable Inflation			
	Government GDP Share	Unemployment Benefits to Personal Income Ratio	Income from Rents to Personal Income Ratio	Employment Share of Large Firms
	IV	IV	IV	IV
	(1)	(2)	(3)	(4)
$u_{i,t-1}$	-0.7203***	-0.5454**	-0.5221***	-0.5708***
	(0.2246)	(0.2185)	(0.1423)	(0.1505)
$u_{i,t-1}\times$ $(m_{i,t-1} - \bar{m})$	0.0095**	0.0046*	0.0057**	0.0046*
	(0.0041)	(0.0025)	(0.0029)	(0.0027)
$u_{i,t-1}\times$ $\left(VAR_{i,t-1} - V\bar{A}R\right)$	3.1062**	12.1585***	-0.1731	-1.1483***
	(1.2078)	(4.4928)	(0.7115)	(0.3920)
Year Fixed Effects	✓	✓	✓	✓
MSA Fixed Effects	✓	✓	✓	✓
N. Observations	3,205	3,185	3,205	3,205

Note: The table reports the estimates of panel regressions similar to that of Table 1 with the difference that we also include the interaction of the unemployment rate with a set of potential confounding factors one at a time, a term we refer to as $u_{i,t-1} \times \left(VAR_{i,t-1} - V\bar{A}R\right)$, where $VAR_{i,t-1}$ is the value that each of this additional confounding factors take in metropolitan area i at year t, and $V\bar{A}R$ is the associated average value in the sample. In all columns, the dependent variable is the non-tradable good inflation rate, and all cases are estimated with IV methods, in which the unemployment rate is instrumented with a shift-share variable that captures tradable demand spillovers, and the robot-adoption variable is instrumented with the industry-level robot penetration in a pool of European countries. All regressions also include the lagged values of robot adoption, the relative price of non-tradable goods, and the confounding variable used in the interaction term, as well as year and MSA fixed effects. Column (1) considers the government GDP share, Column (2) considers the ratio of unemployment benefits to total personal income, Column (3) considers the ratio of income from dividends, interests, and rents to total personal income, and Column (4) considers the employment share of large firms, defined as those with more of 500 employees. Two-way clustered standard errors are reported in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10%, respectively.

Table A.15: Robot Adoption and Inflation across MSAs - Unionization as of 2008

Dependent Variable:	Non-tradable Inflation	
	Union Membership	Union Coverage
	IV	IV
	(1)	(2)
$u_{i,t-1}$	-0.7059***	-0.7037***
	(0.1655)	(0.1755)
$u_{i,t-1} \times \mathbb{I}\{\text{HBP}\}_i$	0.2250	0.2108
	(0.1413)	(0.0091)
$u_{i,t-1} \times (m_{i,t-1} - \bar{m})$	-0.0098	-0.0109
	(0.0089)	(0.0096)
$u_{i,t-1} \times (m_{i,t-1} - \bar{m}) \times \mathbb{I}\{\text{HBP}\}_i$	0.0155*	0.0167*
	(0.0092)	(0.0096)
Year Fixed Effects	✓	✓
MSA Fixed Effects	✓	✓
N. Observations	3,205	3,205

Note: The table reports the estimates of panel regressions similar to that of Table 2 with the difference that we consider an indicator function $\mathbb{I}\{\text{HBP}\}_i$ which does not vary over time, and depends on the value of the union membership and coverage rate of MSA i in the year 2008. In Column (1), the indicator function equals 1 for those MSA with a fraction of workers that are union members which is above the sample median. In Column (2), the indicator function is similarly defined but considering workers that are either union members or covered by union. *** and ** indicate statistical significance at the 1% and 5%, respectively.

Table A.16: Robot Adoption and Inflation across MSAs - Alternative Channels

Dependent Variable:	Non-tradable Inflation					
	Hires	Separations	Job-to-Job Flows	Hours per Employee	Manufacturing Share	Tradable Share
	IV	IV	IV	IV	IV	IV
	(1)	(2)	(3)	(4)	(5)	(6)
$u_{i,t-1}$	-0.4810**	-0.5321***	-0.4110**	-0.3768**	-0.5055***	-0.4891***
	(0.2896)	(0.1370)	(0.1475)	(0.1664)	(0.1371)	(0.1431)
$u_{i,t-1} \times \mathbb{I}\{\text{VAR}\}_{i,t-1}$	-0.0586***	-0.0620***	-0.0886***	-0.0179*	-0.0087	-0.0279*
	(0.0209)	(0.0149)	(0.0091)	(0.0106)	(0.0215)	(0.0169)
$u_{i,t-1} \times (m_{i,t-1} - \bar{m})$	0.0142*	0.0068*	0.0122*	0.0141***	0.0050*	0.0090**
	(0.0085)	(0.0040)	(0.0072)	(0.0052)	(0.0025)	(0.0038)
$u_{i,t-1} \times (m_{i,t-1} - \bar{m}) \times \mathbb{I}\{\text{VAR}\}_{i,t-1}$	-0.0079	-0.0005	-0.0086	-0.0050	0.0665	-0.0020
	(0.0076)	(0.0050)	(0.0065)	(0.0036)	(0.0289)	(0.0027)
Year Fixed Effects	✓	✓	✓	✓	✓	✓
MSA Fixed Effects	✓	✓	✓	✓	✓	✓
N. Observations	3,205	3,205	3,205	3,205	3,205	3,205

Note: The table reports regression estimates as in Table 2 with the only difference that it substitutes the indicator function $\mathbb{I}\{\text{HBP}\}_{i,t-1}$ with another one $\mathbb{I}\{\text{VAR}\}_{i,t-1}$, which varies across specifications. In Column (1), the indicator function equals 1 for those MSA-year with a hiring rate above the sample median. In Column (2), the indicator function equals 1 for those MSA-year with a seperation rate above the sample median. In Column (3), the indicator function equals 1 for those MSA-year with a job-to-job flow rate above the sample median. In Column (4), the indicator function equals 1 for those MSA-year with an amount of hours worked per production worker above the sample median. In Column (5), the indicator function equals 1 for those MSA-year with a share of manufacturing GDP above the sample median. In Column (1), the indicator function equals 1 for those MSA-year with a share of tradable industries GDP above the sample median. ***, **, and * indicate statistical significance at the 1%, 5%, and 10%, respectively.

B More on the Model

This section provides additional details on the model. We start by showing a graphical representation of the structure of the baseline economy and the interplay between the different agents in Figure B.1. We then provide the definition of equilibrium of the baseline model, as well as the set of equilibrium conditions, in Section B.1.

Then, we discuss the simplified economy we have used in Section 4 for deriving the analytical results connecting automation to inflation dynamics. We do so in Section B.2, which provides the proof of our three theoretical insights on the way in which robot installations alter the price Phillips curve, and the role of workers' bargaining power.

Next, Section B.3 describes the derivation of the price Phillips curve in the baseline model with endogenous automation, and Section B.4 replicates the same but for the model specification in which automation is fixed at its steady-state level. Finally, Section B.5 illustrates the relationship between robot adoption and wages in the baseline economy.

Figure B.1: The Structure of the Model

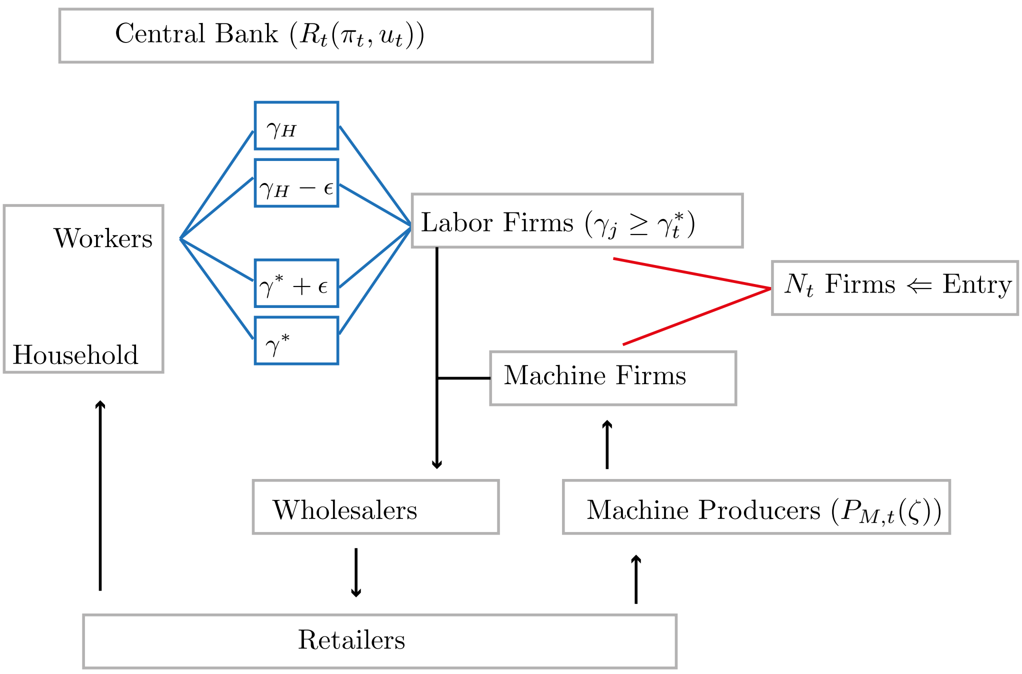

Note: This figure gives a graphical representation of the structure of the model economy.

B.1 Equilibrium

Before defining the equilibrium of the model, let us provide some additional details to it. First, we formally state the optimal price decision of wholesalers, which equals to

$$\max_{P_{i,t}} \mathbb{E}_t \left\{ \sum_{k=t}^{\infty} Q_{k,t} \left(P_{i,k} Y_{i,k} - P_{P,k} Z_{i,k} - \frac{\phi}{2} \left[\frac{P_{i,k}}{P_{i,k-1}} - 1 \right]^2 Y_{i,k} \right) \right\}, \tag{B.1}$$

where $Q_{k,t}$ is households' stochastic discount factor.

Then, notice that since in equilibrium the expected value of the marginal firm entering the market, $V_{E,t}$, is zero, the measure of producers entering the market is set such that

$$V_{E,t} = V_{M,t} \int_1^{\gamma_t^\star} f(\gamma; \alpha) \mathrm{d}\gamma + \int_{\gamma_t^\star}^{\gamma_H} V_{L,t}(\gamma) f(\gamma; \alpha) \mathrm{d}\gamma = 0. \tag{B.2}$$

Now, we can provide the definition of equilibrium in the baseline model. Specifically, the equilibrium is defined as a set of prices $\{\pi_t, q_{P,t}, w_t, R_t\}$, allocations $\{Y_t, Y_{i,t},\ Z_t, C_t, I_t, M_t, N_t, \Xi_t, \gamma_t^\star\}$ and labor market outcomes $\{q_t(\theta_t), u_t\}$ such that producers, the wholesaler, retailer and machine producers maximize profits, the household maximize utility, individuals optimally set labor participation, firms and workers set wages optimally and machine, labor and goods markets clear.

The equilibrium is described by the solutions to the household problem (22) and wholesalers' price setting problem (B.1), the free entry condition (B.2), the automation cut-off condition (6), the definition of total amount of producers' goods (9), the definition of the probability to fill a vacancy (13), the conditions defining the optimal wage in (15) and the equilibrium unemployment rate in (16), wholesalers' market clearing condition (18), retailer's technology (19), retailer's market clearing condition (20), machine manufacturer's technology (21), machine manufacturer's market clearing condition $M_t = \Xi_{M,t}$, the condition determining the numbers of workers searching for jobs, $N_t = w_t/\chi_H$, and the Taylor Rule in Equation (25).

If we define the real cost of entering the market for producers as $\tilde{\kappa} \equiv \kappa/P_t$, and the relative price of machines as $q_{M,t} \equiv P_{M,t}/P_t$, we can rewrite the above set of conditions determining the equilibrium of the model as:

$$U_{C,t} = \beta \mathbb{E}_t \left[\frac{R_t U_{C,t+1}}{\pi_{t+1}} \right], \tag{B.3}$$

$$Y_t = C_t + I_t + \frac{\phi}{2}(\pi_t - 1)^2 (C_t + I_t), \tag{B.4}$$

$$(1-\psi)(C_t + I_t) + \psi q_{P,t}(C_t + I_t) - \phi(\pi_t - 1)\pi_t (C_t + I_t) + \dots$$

$$\dots + \beta \mathbb{E}_t \left[\frac{U_{C,t+1}}{U_{C,t}} \phi(\pi_{t+1} - 1)\pi_{t+1}(C_{t+1} + I_{t+1}) \right] = 0, \tag{B.5}$$

$$q_{M,t} \Xi_t \int_1^{\gamma_t^\star} f(\gamma;\alpha) d\gamma = I_t, \tag{B.6}$$

$$q_{P,t} - q_{M,t} = q_t(\theta_t) \left(q_{P,t}\gamma_t^\star - w_t \right), \tag{B.7}$$

$$\int_1^{\gamma_t^\star} \left(q_{P,t} - q_{M,t} - \kappa \right) f(\gamma;\alpha) d\gamma + \dots$$

$$\dots + \int_{\gamma_t^\star}^{\gamma_H} \left[q_t(\theta_t) \left(q_{P,t}\gamma - w_t \right) - \kappa \right] f(\gamma;\alpha) d\gamma = 0, \tag{B.8}$$

$$q_t(\theta_t) = \xi \left(\frac{\Xi_t \int_{\gamma_t^\star}^{\gamma_H} f(\gamma;\alpha) d\gamma}{N_t} \right)^{\eta - 1}. \tag{B.9}$$

$$N_t = \frac{w_t}{\lambda_H}, \tag{B.10}$$

$$w_t = \frac{\int_{\gamma_t^\star}^{\gamma_H} q_{P,t}\gamma\tau f(\gamma;\alpha) d\gamma}{\int_{\gamma_t^\star}^{\gamma_H} f(\gamma;\alpha) d\gamma}, \tag{B.11}$$

$$N_t u_t = N_t - \Xi_t q_t(\theta_t) \left(\int_{\gamma_t^\star}^{\gamma_H} \gamma f(\gamma;\alpha) d\gamma \right), \tag{B.12}$$

$$\Xi_t \left[\int_1^{\gamma_t^\star} f(\gamma;\alpha) d\gamma + \int_{\gamma_t^\star}^{\infty} q_t(\theta_t)\gamma f(\gamma;\alpha) d\gamma \right] = Y_t. \tag{B.13}$$

Equation (B.3) comes from the solution to the household problem (22), Equation (B.4) combines the clearing conditions in (18), (19) and (20). Equation (B.5) is the solution to the pricing problem in (B.1), whereas Equation (B.6) combines the condition in (21) and the fact that in equilibrium machines manufacturing provide as many robots as the amount of robots firms in the economy, that is, $M_t = \Xi_{M,t}$. Then, Equations (B.7)-(B.12) determine the automation cut-off $\gamma^\star$, the firm entry condition, the probability firms fill a vacancy, the labor market participation condition, the equilibrium wage, and the unemployment rate, respectively. Equation (B.13) is the intermediate market clearing condition defined in (9). Finally, we close the model with the monetary policy rule of Equation (25).

We then provide the full set of log-linear equilibrium conditions of the model:

$$-\sigma\hat{C}_t = \beta\hat{R}_t\mathbb{E}_t\left[-\sigma\hat{C}_{t+1} - \hat{\pi}_{t+1}\right], \tag{B.14}$$

$$\hat{Y}_t = \frac{\bar{C}}{\bar{Y}}\hat{C}_t + \frac{\bar{I}}{\bar{Y}}\hat{I}_t, \tag{B.15}$$

$$\hat{\pi}_t = \frac{\epsilon-1}{\phi}\widehat{q_{P,t}} + \beta\mathbb{E}_t\left[\hat{\pi}_{t+1}\right], \tag{B.16}$$

$$\hat{\Xi}_t + \frac{\frac{\alpha}{\bar{\gamma}^{\star\,\alpha}}}{1-\frac{1}{\bar{\gamma}^{\star\,\alpha}}}\widehat{\gamma_t^{\star}} = \hat{I}_t, \tag{B.17}$$

$$\bar{q_P}\widehat{q_{P,t}} = \left[\bar{q_P} - \bar{q}_M\right]\widehat{q_t(\theta_t)} + q(\bar{\theta})\bar{q_P}\bar{\gamma}^{\star}\left[\widehat{\gamma_t^{\star}} + \widehat{q_{P,t}}\right] - q(\bar{\theta})\bar{w}\hat{w}_t, \tag{B.18}$$

$$\begin{aligned}\left[1-\left(\frac{1}{\bar{\gamma}^{\star}}\right)^{\alpha}\right]\bar{q_P}\widehat{q_{P,t}} + \frac{\alpha}{\alpha-1}\left[\bar{\gamma}^{\star\,1-\alpha} - \gamma_H^{1-\alpha}\right]q(\bar{\theta})\bar{q_P}\left[\widehat{q_{P,t}} + \widehat{q_t(\theta_t)}\right] - \dots \\ \dots - \left(\bar{\gamma}^{\star\,-\alpha} - \gamma_H^{-\alpha}\right)q(\bar{\theta})\bar{w}\left[\widehat{q_t(\theta_t)} + \hat{w}_t\right],\end{aligned} \tag{B.19}$$

$$\begin{aligned}(\hat{Y}_t - \hat{\Xi}_t)\frac{\bar{Y}\left(1-\gamma_H^{-\alpha}\right)}{\bar{\Xi}} = \alpha\left(\frac{1}{\bar{\gamma}^{\star}}\right)^{\alpha}\left[1-q(\bar{\theta})\bar{\gamma}^{\star}\right]\widehat{\gamma_t^{\star}} + \dots \\ \dots + \frac{\alpha}{\alpha-1}\left[\bar{\gamma}^{\star\,1-\alpha} - \gamma_H^{1-\alpha}\right]q(\bar{\theta})\widehat{q_t(\theta_t)},\end{aligned} \tag{B.20}$$

$$\widehat{q_t(\theta_t)} = (\eta-1)\left[\hat{\Xi}_t - \hat{N}_t - \frac{\alpha\bar{\gamma}^{\star\,-\alpha}}{\bar{\gamma}^{\star\,-\alpha} - \gamma_H^{-\alpha}}\widehat{\gamma_t^{\star}}\right], \tag{B.21}$$

$$\hat{w}_t = \widehat{q_{P,t}} - (\alpha-1)\frac{\bar{\gamma}^{\star\,1-\alpha}}{\bar{\gamma}^{\star\,1-\alpha} - \gamma_H^{1-\alpha}}\widehat{\gamma_t^{\star}} + \alpha\frac{\bar{\gamma}^{\star\,-\alpha}}{\bar{\gamma}^{\star\,-\alpha} - \gamma_H^{-\alpha}}\widehat{\gamma_t^{\star}}, \tag{B.22}$$

$$\hat{N}_t = \hat{w}_t, \tag{B.23}$$

$$\frac{\bar{u}}{1-\bar{u}}\hat{u}_t = \hat{\Xi}_t - \hat{N}_t + \widehat{q_t(\theta_t)} - \frac{\alpha\bar{\gamma}^{\star\,-\alpha}}{\bar{\gamma}^{\star\,-\alpha} - \gamma_H^{-\alpha}}\widehat{\gamma_t^{\star}}. \tag{B.24}$$

B.2 Analytical Results in the Simplified Model

The model version that we present in Section 4 simplifies the baseline economy in two ways. First, we consider two values for labor productivity: producers either draw $\gamma = 0$—and thus operate using machines—with probability $\Phi_t^{\star}$, or get $\gamma = \gamma_H > 1$ with the remaining probability $1 - \Phi_t^{\star}$. The probability $\Phi_t^{\star}$ is then the ratio between the measure of robot firms and the total measure of firms, that is, $\Phi_t^\star = \Xi_{M,t}/\Xi_t$. Second, we posit that the fraction of robot firms, $\Phi_t^\star$, is determined by a reduced-form function that depends on wages. Specifically, we set

$$\Phi_t^\star = \bar{\Phi}^\star \left[1 + \varrho(w_t - \bar{w})\right], \tag{B.25}$$

where $\bar{\Phi}^\star$ and $\bar{w}$ are the steady-state share of robot firms and real wages, and $0 < \varrho < 1$ pins down the elasticity of automation to the real wage, denoted by $\xi_{\Phi^\star,w}$

In this setting, the free entry condition for producers implies the following:

$$V_{E,t} = \Phi_t^\star(P_{P,t} - P_{M,t}) + (1 - \Phi_t^\star)q_t(\theta_t)\left[P_{P,t} - W_t\right] = 0. \tag{B.26}$$

We can also determine the expected real value of the surplus of firms who match with a worker to be

$$\mathbb{E}_t\left[S_t\right] = \frac{P_{P,t}}{P_t}\gamma_H - \frac{W_t}{P_t} = q_{p,t}\gamma_H - w_t. \tag{B.27}$$

As all firms that draw $\gamma_j = \gamma_H$ find it optimal to open a vacancy, $v_t = \Xi_t(1 - \Phi_t^\star)$. We can define the labor market tightness and the probability that a firm fills a vacancy as

$$\theta_t = \frac{\Xi_t(1 - \Phi_t^\star)}{N_t}, \tag{B.28}$$

and

$$q_t(\theta_t) = \xi\left[\frac{\Xi_t(1 - \Phi_t^\star)}{N_t}\right]^{\eta-1}. \tag{B.29}$$

Upon a match, workers and firms bargain on the real wage by splitting the total surplus of the match. The wage bargaining problem is

$$\arg\max_{w_t} = w_t^\tau\ \mathbb{E}_t\left[S_t\right]^{1-\tau}, \tag{B.30}$$

where τ denotes the firms' bargaining power. Using the definition of the surplus of a match of Equation (B.27), we can defined the optimal wage as

$$w_t = \tau q_{P,t}\gamma_H. \tag{B.31}$$

The unemployment rate equals the ratio between the measures of individuals that have not matched with a producer and those actively looking for a job:

$$u_t = \left[N_t - \Xi_t(1 - \Phi_t^\star)q_t(\theta_t)\right]/N_t. \tag{B.32}$$

Then, leveraging the wholesalers' price setting problem of Equation (B.1) and log-linearizing the optimal pricing decision around the steady state yields the dynamics of inflation:

$$\hat{\pi}_t = \frac{\epsilon - 1}{\phi}\widehat{q_{P,t}} + \beta\mathbb{E}_t\left[\hat{\pi}_{t+1}\right]. \tag{B.33}$$

From the labor market block of the model we establish a link between unemployment and the probability of filling a vacancy, combining the log-linearized versions of (B.29) and (B.32), obtaining

$$\frac{\bar{u}}{1-\bar{u}}\hat{u}_t = \frac{\eta}{\eta - 1}\widehat{q_t(\theta_t)}. \tag{B.34}$$

From the producer conditions, combining the log-linearized versions of the free entry conditions in (B.26), the wage equation in (B.31), and the exogenous assumption on automation in (B.25), we establish a link between the relative price of goods, $q_{P,t}$, and the probability of filling a vacancy, that is

$$\widehat{q_t(\theta_t)} = \frac{1}{1-\bar{\Phi}^\star}\widehat{\Phi_t^\star} + \frac{q_M}{\bar{q}_P - \bar{q}_M}\widehat{q_{P,t}} = \tag{B.35}$$

$$= \left[\frac{1}{1-\bar{\Phi}^\star}\varrho\frac{\tau(\epsilon-1)\gamma_H}{\epsilon} + \frac{\epsilon}{(\epsilon-1)\zeta - \epsilon}\right]\widehat{q_{P,t}}. \tag{B.36}$$

The last equality of Equation (B.35) comes from the fact that $\bar{w} = \tau\bar{q}_P\gamma_H$, $\bar{q}_P = (\epsilon-1)/\epsilon$, and $q_M = 1/\zeta$.

Finally, we can express the New Keynesian Phillips curve as the relationship between inflation and unemployment, that is

$$\hat{\pi}_t = \Psi(\bar{\Phi}^\star;\tau)\hat{u}_t + \beta\mathbb{E}_t\left[\hat{\pi}_{t+1}\right], \tag{B.37}$$

where $\Psi(\bar{\Phi}^\star;\tau)$ denotes the slope describing the sensitivity of inflation to changes in unemployment. The slope explicitly depends on the degree of automation at the steady state, $\bar{\Phi}^\star$, and workers' bargaining power, τ, as follows:

$$\Psi(\bar{\Phi}^{\star};\tau) = -\frac{\epsilon-1}{\phi}\left[\frac{\bar{u}}{1-\bar{u}}\frac{1-\eta}{\eta}\left(\varrho\frac{1}{1-\bar{\Phi}^{\star}}\frac{\tau(\epsilon-1)\gamma_H}{\epsilon}+\frac{\epsilon}{(\epsilon-1)\zeta-\epsilon}\right)^{-1}\right]. \tag{B.38}$$

Importantly, the slope is always negative, since lower inflation rates are associated with times of relatively high unemployment.

In this context, we can derive the following three key analytical insights on the way in which robot installations shape inflation dynamics. Specifically, we look at the way in which the slope of the Phillips curve, $\Psi(\bar{\Phi}^{\star};\tau)$, vary with automation. The first result characterizes the derivative of the slope with respect to the value of robot installations at the steady state.

Proposition 1. *A higher degree of automation at the steady state flattens the slope of the Phillips curve:*

$$\frac{\partial\Psi(\bar{\Phi}^{\star};\tau)}{\partial\bar{\Phi}^{\star}} > 0.$$

Proof.

$$\frac{\partial\Psi(\bar{\Phi}^{\star};\tau)}{\partial\bar{\Phi}^{\star}} = -\Psi(\bar{\Phi}^{\star};\tau)\frac{\frac{\varrho}{(1-\bar{\Phi}^{\star})^2}\frac{\tau(\epsilon-1)\gamma_H}{\epsilon}}{\left[\varrho\frac{1}{1-\bar{\Phi}^{\star}}\frac{\tau(\epsilon-1)\gamma_H}{\epsilon}+\frac{\epsilon}{(\epsilon-1)\zeta-\epsilon}\right]}.$$

Since the slope of the Phillips curve, $\Psi(\bar{\Phi}^{\star};\tau)$, is negative and the last term is always positive, the derivative is negative. Finally, note that as robot installations crowd out completely labor at the steady state, then the slope converges to zero: $\lim_{\bar{\Phi}^{\star}\to 1}\Psi(\bar{\Phi}^{\star};\tau) = 0$. □

The second result is on the way in which automation and workers' bargaining power interact in altering inflation dynamics.

Proposition 2. *A surge in steady-state automation flattens relatively more the Phillips curve at higher values of workers' bargaining power:*

$$\frac{\partial^2\Psi(\bar{\Phi}^{\star};\tau)}{\partial\bar{\Phi}^{\star}\partial\tau} > 0.$$

Proof. The derivative of the slope of the Phillips curve relative to the steady-state share of robot firms and workers' bargaining power is the following:

$$\frac{\partial^2\Psi(\bar{\Phi}^\star;\tau)}{\partial\bar{\Phi}^\star\partial\tau} = \frac{\frac{\epsilon-1}{\phi}\frac{\bar{u}}{1-\bar{u}}\frac{1-\eta}{\eta}\frac{1}{(1-\bar{\Phi}^\star)^2}}{\left[\varrho\frac{1}{1-\bar{\Phi}^\star}\frac{\tau(\epsilon-1)\gamma_H}{\epsilon}+\frac{\epsilon}{(\epsilon-1)\zeta-\epsilon}\right]^2} \times \dots$$

$$\dots \times \left\{1-2\frac{\varrho\frac{1}{1-\bar{\Phi}^\star}\frac{\tau(\epsilon-1)\gamma_H}{\epsilon}}{\left[\varrho\frac{1}{1-\bar{\Phi}^\star}\frac{\tau(\epsilon-1)\gamma_H}{\epsilon}+\frac{\epsilon}{(\epsilon-1)\zeta-\epsilon}\right]}\right\} \times \left[\frac{\varrho(\epsilon-1)\gamma_H}{\epsilon}\right].$$

The first and third term are always positive, whereas the second term is positive in the empirically relevant case in which the elasticity of robots relative to the real wage is between 0 and 1, that is, $0 < \xi_{\Phi^\star,w} < 1$.[2] $\square$

Thus, the flattening effect of automation is amplified at high levels of workers' bargaining power.

Finally, the third key result highlights the key role of changes in automation upon a shock. That is, higher automation in steady state does flatten the slope only insofar automation can adjust upon a shock.

Proposition 3. *Changes in steady-state automation alter the slope of the Phillips curve if and only if automation does vary as a function of wages upon a shock, that is, insofar $\varrho \neq 0$:*

$$\left.\frac{\partial\Psi(\bar{\Phi}^\star;\tau)}{\partial\bar{\Phi}^\star}\right|_{\varrho=0} = 0.$$

Proof. Recall that the derivative of the slope of the Phillips curve relative to the steady-state share of robot firms is $\frac{\partial\Psi(\bar{\Phi}^\star;\tau)}{\partial\bar{\Phi}^\star} = -\Psi(\bar{\Phi}^\star;\tau)\frac{\frac{\varrho}{(1-\bar{\Phi}^\star)^2}\frac{\tau(\epsilon-1)\gamma_H}{\epsilon}}{\left[\varrho\frac{1}{1-\bar{\Phi}^\star}\frac{\tau(\epsilon-1)\gamma_H}{\epsilon}+\frac{\epsilon}{(\epsilon-1)\zeta-\epsilon}\right]}$. Thus, when $\varrho = 0$ and automation does not react to changes in wages above steady state, then $\frac{\partial\Psi(\bar{\Phi}^\star;\tau)}{\partial\bar{\Phi}^\star} = 0$. $\square$

In other words, what matters for inflation dynamics is not just the existing stock of robots, but rather firms' ability to further automate.

[2]These boundaries arise because in this setting automation is not optimally set, but rather depends on a reduced-form function. If automation were to respond too aggressively to changes in the real wage, then it may lead to a "suboptimal" fall in output and a further increase in prices, offsetting the bargaining mechanism that ultimately dampens the price response as firms threaten to automate. In the full version of the model, automation changes optimally and thus only the bargaining mechanism is in place.

B.3 Derivation of the Price Phillips Curve

After having characterized the price Phillips curve in the simplified version of our model, this section derives the same relationship in our baseline economy. In doing that, we show that automation keeps influencing inflation dynamics, as the slope of the Phillips curve depends on the amount of robot installations at the steady state.

Combining Equations (B.21) and (B.24), we establish a link between unemployment and the probability of filling a vacancy, thus obtaining

$$\frac{\bar{u}}{1-\bar{u}}\hat{u}_t = \frac{\eta}{\eta-1}\widehat{q_t(\theta_t)}. \tag{B.39}$$

Then, rearranging Equation (B.19) yields

$$\hat{w}_t = \frac{\omega_1}{\tau}\widehat{q_{P,t}} + \frac{1-\tau}{\tau}\widehat{q_t(\theta_t)}, \tag{B.40}$$

where ω_1 is a convolution of parameters and steady-state values, and equals to

$$\omega_1 \equiv \frac{\left[1-\left(\frac{1}{\bar{\gamma}^\star}\right)^{\alpha}\right] + \xi^{\frac{1}{\eta}}(1-\bar{u})^{\frac{\eta-1}{\eta}}\frac{\alpha}{\alpha-1}\left(\bar{\gamma}^{\star\,1-\alpha} - \gamma_H^{1-\alpha}\right)}{\xi^{\frac{1}{\eta}}(1-\bar{u})^{\frac{\eta-1}{\eta}}\frac{\alpha}{\alpha-1}\left(\bar{\gamma}^{\star\,1-\alpha} - \gamma_H^{1-\alpha}\right)}. \tag{B.41}$$

We can relate changes in the automation cut-off to changes in the labor market tightness and producers' price by combining Equation (B.40) and (B.22):

$$\widehat{\gamma_t^\star} \;=\; \omega_2\frac{1-\tau}{\tau}\widehat{q_t(\theta_t)} + \omega_2\left[\frac{\omega_1}{\tau} - 1\right]\widehat{q_{P,t}}, \tag{B.42}$$

where ω_2 is a convolution of parameters and steady-state values, and equals to

$$\omega_2 \equiv \left[(1-\alpha)\frac{\bar{\gamma}^{\star\,1-\alpha}}{\bar{\gamma}^{\star\,1-\alpha} - \gamma_H^{1-\alpha}} + \alpha\frac{\bar{\gamma}^{\star\,-\alpha}}{\bar{\gamma}^{\star\,-\alpha} - \gamma_H^{-\alpha}}\right]^{-1}. \tag{B.43}$$

Nexy, we can characterize the changes in producers' price as a function of the automation cut-off and workers' bargaining power by combining Equation (B.40) and (B.42) into Equation (B.18) and (B.39):

$$\widehat{q_{P,t}} = \frac{\bar{u}}{1-\bar{u}}\frac{\bar{\gamma}^\star\left(1+\omega_2(1-\tau)/\tau\right) - \omega_3/\tau}{-\frac{\eta}{1-\eta}\left[1/(\xi^{1/\eta}(1-\bar{u})^{(\eta-1)/\eta}) + \bar{\gamma}^\star(\omega_2-1) + \frac{\omega_1(\omega_3-\bar{\gamma}^\star\omega_2)}{\tau}\right]}\hat{u}_t, \tag{B.44}$$

where ω_3 is a convolution of parameters and steady-state values, and equals to

$$\omega_3 \equiv \tau \frac{\alpha}{\alpha-1} \frac{\bar{\gamma}^{\star\,1-\alpha} - \gamma_H^{1-\alpha}}{\bar{\gamma}^{\star\,-\alpha} - \gamma_H^{-\alpha}}. \tag{B.45}$$

We can define the price Phillips curve as

$$\hat{\pi}_t = \Psi(\bar{\Phi}^\star;\tau)\hat{u}_t + \mathbb{E}_t\left[\beta\hat{\pi}_{t+1}\right], \tag{B.46}$$

where the slope, $\Psi(\bar{\Phi}^\star;\tau)$, equals to

$$\Psi(\bar{\Phi}^\star;\tau) = -\frac{\epsilon-1}{\phi}\frac{\bar{u}}{1-\bar{u}}\frac{\bar{\gamma}^\star\left(1+\omega_2(1-\tau)/\tau\right)-\omega_3/\tau}{\frac{\eta}{1-\eta}\left[1/(\xi^{1/\eta}(1-\bar{u})^{(\eta-1)/\eta}) + \bar{\gamma}^\star(\omega_2-1) + \frac{\omega_1(\omega_3-\bar{\gamma}^\star\omega_2)}{\tau}\right]}. \tag{B.47}$$

B.4 Price Phillips Curve with Fixed Automation

To uncover that automation influences inflation dynamics only when it is allowed to change upon a shock, we consider a version of the model in which automation is kept fixed at its steady-state level, that is, $\gamma_t^\star = \bar{\gamma}^\star$ or $\widehat{\gamma_t^\star} = 0$.

In this case, we can use Equations (B.18) and (B.22) to define what is the log-linear relationship between the probability to fill a vacancy and producers' price:

$$\widehat{q_t(\theta_t)} = \frac{\bar{q}_p - \bar{q}_M}{\bar{q}_M}\widehat{q_{P,t}}. \tag{B.48}$$

If we plug in this relationship into the Phillips cuve, we have the following

$$\hat{\pi}_t = \Psi(\bar{\Phi}^\star;\tau)\hat{u}_t + \beta\mathbb{E}_t\left[\hat{\pi}_{t+1}\right], \tag{B.49}$$

where the slope becomes

$$\Psi(\bar{\Phi}^\star;\tau)\mid_{\widehat{\gamma_t^\star}=0} = -\frac{\epsilon-1}{\phi}\left[\frac{\bar{u}}{1-\bar{u}}\frac{1-\eta}{\eta}\left(\frac{\epsilon}{(\epsilon-1)\zeta-\epsilon}\right)^{-1}\right], \tag{B.50}$$

which does not depend anymore on steady-state automation. Thus, once more this result uncovers that changes in steady-state robot installations alter inflation dynamics only insofar as automation can adjust upon a shock, as shown in Figure 2 of the paper.

B.5 Automation Elasticity to Wages

This section studies the relationship between automation and wage dynamics. Specifically, we represent automation as a function of wages, confirming the reduced-form representation we used in simplified version of the model in Section B.2.

We start by combining Equation (B.42), (B.18) and (B.40) to derive the relationship between changes in the probability to fill a vacancy and changes in producers' price

$$\widehat{q_t(\theta_t)} = -\left[\vartheta_1\left(\bar{\gamma}^{\star}, \tau\right)\right]^{-1} \widehat{q_{P,t}}, \tag{B.51}$$

where ϑ_1 is a convolution of parameters and steady-state values, and equals to

$$\vartheta_1\left(\bar{\gamma}^{\star}\right) = \frac{\bar{\gamma}^{\star}\left(1+\omega_2\tau/(1-\tau)\right) - \omega_3/(1-\tau)}{1/(\xi^{1/\eta}(1-\bar{u})^{(\eta-1)/\eta}) + \bar{\gamma}^{\star}(\omega_2 - 1) + \frac{\omega_1(\omega_3 - \bar{\gamma}^{\star}\omega_2)}{(1-\tau)}}. \tag{B.52}$$

Then, we can connect changes in the real wage to changes in producers' price by plugging the previous condition that back into Equation (B.40), which gives

$$\hat{w}_t = \left[\frac{\omega_1}{\tau} - \frac{1-\tau}{\tau}\vartheta_1\left(\bar{\gamma}^{\star}, \tau\right)^{-1}\right] \widehat{q_{P,t}}. \tag{B.53}$$

Finally, if we use Equation (B.22), we have that wages relate to changes in the automation cut-off as follows

$$\hat{\gamma}_t^{\star} = \omega_2 \left[1 - \frac{\tau}{\omega_1(\bar{\gamma}^{\star}) - (1-\tau)\vartheta_1\left(\bar{\gamma}^{\star}, \tau\right)^{-1}}\right] \hat{w}_t = \Omega\left(\bar{\gamma}^{\star}, \tau\right)\hat{w}_t. \tag{B.54}$$

Thus, the elasticity of the automation to the real wage varies itself with the level of steady-state automation, $\bar{\gamma}^{\star}$, and with workers' bargaining power, τ. To highlight that the baseline model implies the same relationship between automation and real wages of the simplified economy, in which robot installations increase with labor costs via a reduced-form function, we show in Figure B.2 how the elasticity of automation to the real wage varies with workers' bargaining power and steady-state automation. The figure confirms that the elasticity is positive and rises with both automation and the degree of bargaining power.

Figure B.2: Automation, workers' bargaining power, and the automation elasticity to the real wage

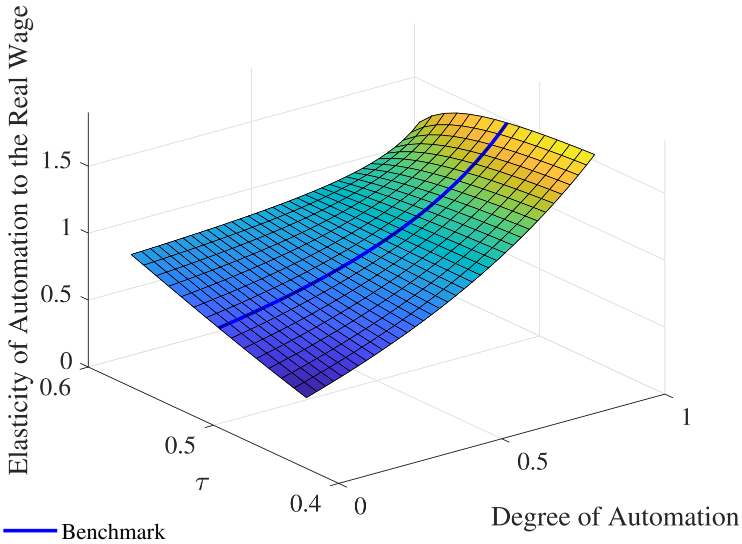

Note: The figures show how the elasticity of automation to the real wage varies with robot-specific technical change ζ (i.e., the degree of automation at the steady state) and workers' bargaining power τ.

C Additional Results of the Model

This section provides an extensive set of additional results of the model. We start with Section C.1 by deriving the implied aggregate elasticity of substitution between workers and machines in the model, then dissect the way in which automation alters the wage dynamic in Section C.2, and inspect the key model mechanisms in an impulse-response exercise in Section C.3. Section C.5 reports the responses of the key set of variables of the model to both monetary policy and productivity shocks. Finally, we study a version of the model in which automation is kept fixed at its steady-state level in Section C.6.

C.1 Elasticity of Substitution between Workers and Machines

Section 3.1 provides the details of producers in our model and their decisions to use either machines or post job vacancies and look for workers. The structure of our economy implies that at the producer level machines and workers are perfectly substitutes: each producer can choose one or the other. Notwithstanding this condition, workers and machines are imperfect substitutes at the aggregate level.

To uncover this result, we back out from the model simulated values of the total number of workers, $\mathcal{N}_t$, the wage rate, W_t, the total number of machines, M_t, and the price of machines, $P_{M,t}$, after feeding the economy with 10,000 realizations of the preference shock. Let us conjecture that the relationship between machines and workers at the aggregate level can be represented by the following CES function with an elasticity of substitution between workers and machines ϱ_e,

$$\left[\varphi_e \mathcal{N}_t^{\frac{\varrho_e-1}{\varrho_e}} + (1-\varphi_e)\, M_t^{\frac{\varrho_e-1}{\varrho_e}}\right]^{\frac{\varrho_e}{\varrho_e-1}}, \tag{C.1}$$

whose first-order condition reads

$$\frac{\mathcal{N}_t}{M_t} = \frac{\varphi_e}{1-\varphi_e}\left(\frac{P_{M,t}}{W_t}\right)^{\varrho_e}. \tag{C.2}$$

Then, we can estimate the implied elasticity of substitution, by estimating the following regression using simulated data from the model:

$$\log\left(\frac{\mathcal{N}_t}{M_t}\right) = \alpha + \beta \log\left(\frac{P_{M,t}}{W_t}\right) + \epsilon_t, \tag{C.3}$$

where $\alpha = \log\left(\frac{\varphi_e}{1-\varphi_e}\right)$, and $\beta = \varrho_e$. This exercise leads to an estimate of the elasticity of substitution that equals $\hat{\varrho}_e = 2.31$, with a p-value below 0.001. This means that the model features a degree of imperfect substitutability between workers and machines at the aggregate level.

What drives the imperfect degree of substitutability between robots and workers at the aggregate level? There are two forces behind this result. First, as we emphasize in Section 4, automation alters prices through the threat that the option value of automation exerts on wage negotiations. Consequently, in the model robot adoption alters price and wage dynamics even absent changes in machine installations at the business-cycle frequency. Second, the incentives to automate rise upon expansionary preference shocks, also leading to an increase in producers' incentives to enter the economy. This mechanism mutes the fluctuations in the relative share of robot firms for any given change in prices.

C.2 Automation and Wage Dynamics

To further dissect how automation affects price inflation in the baseline model, we focus on its effect on the dynamics of the real wage, as any change in wage dynamics ends up exerting pressure on prices. Combining Equation (B.18), (B.40), (B.42) into Equation (B.22), we get a relationship between the real wage, automation cut-off, and the probability firms fill a vacancy:

$$\hat{w}_t = -\underbrace{\vartheta_1\left(\bar{\gamma}^\star\right)\widehat{q_t(\theta_t)}}_{\text{Wage Setting Effect}} - \underbrace{\vartheta_2\left(\bar{\gamma}^\star\right)\hat{\gamma}_t^\star}_{\text{Labor Displacing Effect}} + \underbrace{\vartheta_3\left(\bar{\gamma}^\star\right)\hat{\gamma}_t^\star}_{\text{Selection Effect}}, \tag{C.4}$$

where

$$\vartheta_1\left(\bar{\gamma}^\star\right) = \frac{\bar{\gamma}^\star\left(1+\omega_2\tau/(1-\tau)\right) - \omega_3/(1-\tau)}{1/(\xi^{1/\eta}(1-\bar{u})^{(\eta-1)/\eta}) + \bar{\gamma}^\star(\omega_2 - 1) + \frac{\omega_1(\omega_3 - \bar{\gamma}^\star\omega_2)}{(1-\tau)}}, \tag{C.5}$$

$$\vartheta_2(\bar{\gamma}^\star) = (\alpha - 1)\frac{\bar{\gamma}^{\star\,1-\alpha}}{\bar{\gamma}^{\star\,1-\alpha} - \gamma_H^{1-\alpha}}, \tag{C.6}$$

$$\vartheta_3(\bar{\gamma}^\star) = \alpha\frac{\bar{\gamma}^{\star\,-\alpha}}{\bar{\gamma}^{\star\,-\alpha} - \gamma_H^{-\alpha}}. \tag{C.7}$$

Equation (C.4) relates the log deviations of the real wage from its steady state, $\hat{w}_t$, to changes in the probability of filling a vacancy, $\widehat{q_t(\theta_t)}$, and the automation cut-off, $\hat{\gamma}_t^\star$, with positive loading factors, $\vartheta_1(\bar{\gamma}^\star)$, $\vartheta_2(\bar{\gamma}^\star)$, $\vartheta_3(\bar{\gamma}^\star) > 0$, that depend on the automation cut-off at the steady state, $\bar{\gamma}^\star$. This condition posits that real wage dynamics depend on three components.

The first term is the *Wage Setting Effect*, and it amplifies wage dynamics via the movement in the probability of filling a vacancy, $\widehat{q_t(\theta_t)}$. Intuitively, when the probability of filling a vacancy decreases, workers can bargain for a higher wage. This component reflects that the wage bargained responds to changes in the tightness of the labor market, θ. Hence, the greater the relevance of this term, the larger the response of wages to changes in labor market conditions.

The second term describes the *Labor Displacing Effect*, which subdues wage dynamics. This comes from the fact that the automation cut-off rises in times of low unemployment, as producers replace workers with machines to suppress labor demand.

The third term captures the *Selection Effect* of automation, which magnifies wage dynamics. When the automation cut-off rises to suppress labor demand, the average efficiency of the remaining labor firms increases, which in turn pushes upward the real wage.

How does an increase in automation at the steady state alter these three effects? We show that a surge in $\bar{\gamma}^\star$ dampens the relevance of all three channels. This result mostly comes from the fact that the three loading factors decrease with higher steady-state automation, that is, $\partial\vartheta_1(\bar{\gamma}^\star)/\partial\bar{\gamma}^\star < 0$, $\partial\vartheta_2(\bar{\gamma}^\star)/\partial\bar{\gamma}^\star < 0$, $\partial\vartheta_3(\bar{\gamma}^\star)/\partial\bar{\gamma}^\star < 0$. Consequently, a higher $\bar{\gamma}^\star$ directly reduces the relevance of the three channels.

However, higher steady-state automation also alters the probability of filling a vacancy, and through that further reduces the influence of the wage setting effect, above and beyond how automation impacts the loading factor. This happens because, during wage bargaining, workers and firms internalize that if wages were to rise substantially, some firms would switch to machines, especially so when robots are relatively cheaper at the steady state. As a result, the number of firms posting vacancies—and thus the probability of filling them—is influenced by automation. In other words, the outside option of automation alters workers' effective bargaining power, reducing their influence on the equilibrium wage for a given level of labor market tightness.

Also, the labor displacing and selection effects shrink with automation above and beyond the influence of robot installations on the loading factor. This happens via general equilibrium forces. Indeed, higher automation at the steady state dampens the *changes upon a shock* in the automation cut-off, $\hat{\gamma}_t^{\star}$: the drop in workers' effective bargaining power implies that firms, upon a shock, no longer need to switch into machines as much as before to curb the increase in production costs. This reduces the contribution of labor displacing and selection effects.

We then proceed to quantify the impact of automation on wage dynamics by dissecting the relevance of the wage setting channel vis-à-vis the joint role of the selection effect and the labor displacing effect. To do so, we consider two steady states that only differ in the amount of automation. We use the two steady states that bound the transitional dynamics of Section 5.4: a first steady state with a 0.2% robot-per-employee ratio which describes the U.S. economy in the early 2000s, and a second steady state with a 0.8% robot-per-employee ratio to capture the amount of automation in the U.S. the late 2010s. To target these two economies, we change the level of robot-specific technical change, ζ.[3] We then refer to the first economy as the "low automation steady state", and the second economy as the "high automation steady state".

We consider an expansionary preference shock, $\varepsilon_{\Omega,t}$, in either economy, compute the response of the wage, and disentangle the contribution of each of the three effects of Equation (C.4). In presenting the results, we bundle the selection and labor displacing effects together

[3] We also change the entry cost κ and searching cost λ_H preserving the same unemployment and participation rates across steady states.

Figure C.3: Wage Setting, Selection and Labor Displacing Effects

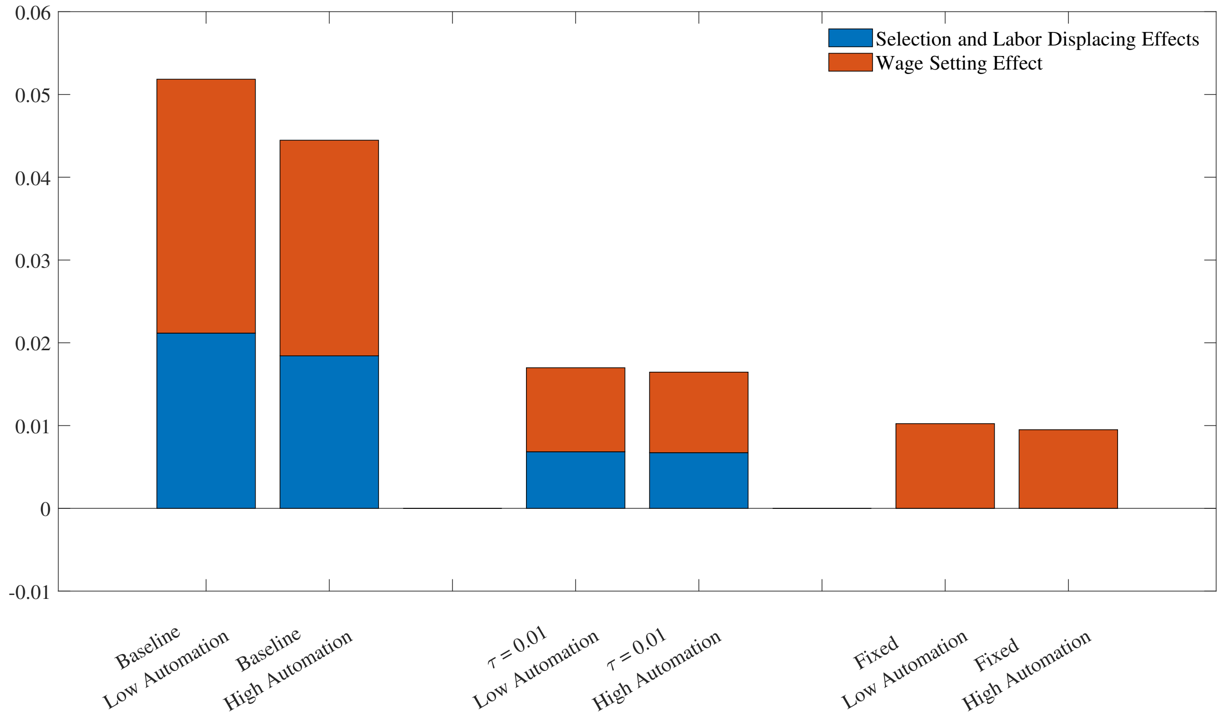

Note: The bars depict the contribution of the wage setting effects (the red bars), and the sum of the selection and labor displacing effects (the blue bars) in the first period after the realization of an expansionary preference shock for the baseline economy, the model version in which firms have almost full bargaining power, that is, $\tau = 0.01$, and the case where automation is fixed at its steady state value (*Fixed*), in both the low automation and high automation steady states.

since they both depend on the changes in the automation cut-off upon the shock. Thus, this decomposition highlights the role of changes in workers' expected payoff—the wage setting effect—and changes in the automation cut-off—the sum of the selection and labor displacing effects. Figure C.3 reports the results of this exercise.

The figure shows that while both the selection and labor displacing effects are muted when moving from the low automation to the high automation steady state, their joint contribution barely changes across the two equilibria. This is because the contribution of each of the two effects are almost mirror images of one another: the selection effect is positive, and becomes relatively less positive in the high automation steady state, whereas the labor displacing effect is negative, and becomes relatively less negative in the high automation steady state. Thus, when bundling together these two channels, their joint contribution to the dampening of the real wage response is limited. Instead, the lower responsiveness of the wage setting effect in the high automation steady state accounts for the lion's share of the muted dynamics of the real wage.

When we consider the model version in which the bargaining power is tilted towards firms, so that $\tau = 0.01$, and replicate the same exercise as before. In this case, we find that neither the wage setting channel nor the sum of the selection and labor displacing effects vary across the low and high automation steady states. These results once more corroborate the key role of the wage setting effect of automation in driving the reduction in the price and wage inflation sensitivity to unemployment.

Finally, we consider a third case, in which automation cannot adjust upon a shock. In this case, while robot adoption varies across steady states, it then cannot change endogenously in response to preference shocks. In other words, automation is kept fixed at its steady state level. We refer to this case as the (*Fixed*) economy. In this instance, we find that the selection and labor displacing effects are equal to zero and the wage setting effects is largely unchanged in the two steady states. Once again this results provides further evidence that what matters for inflation dynamics is that the existing stock of robots delimits firms' ability to further automate and this effect is modulated by workers' bargaining power. When automation is higher at steady state, firms' ability to further automate is greater, which dampens wage movements particularly when workers' bargaining power is high.

C.3 Inspecting the Mechanism

To isolate the relevance of the wage setting, selection, and labor displacing effects, we compare again the baseline economy to the alternative specification in which we set $\tau = 0.01$ assuming wage bargaining is done almost entirely to maximize the firms' surplus from the match. Intuitively, this case shuts down the wage setting effect: the use of automation in curtailing workers' wage negotiation requests is not operative when firms already have full bargaining power. We devise an exercise in which—for both the baseline model ($\tau = \eta = 0.5$) and the alternative specification with $\tau = 0.01$—we evaluate the response of price and wage inflation

Figure C.4: Impulse Responses in the Baseline Model

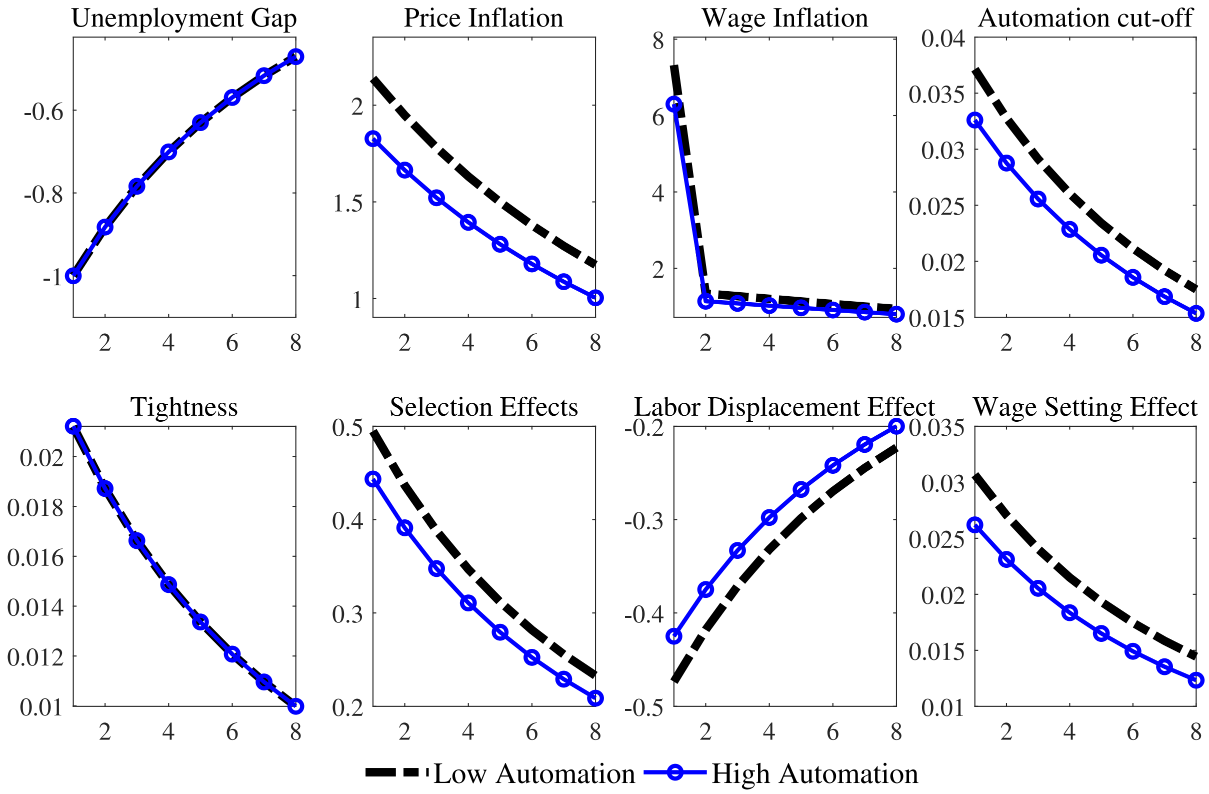

Note: Impulse responses in the first eight quarters after an expansionary preference shock of the unemployment gap, price inflation, wage inflation, the automation cut-off, workers' expected gains from searching for a job, the selection effect, the labor displacing effect, and the wage setting effect in the low-automation and high-automation steady states of the baseline economy. Responses are in percentage-point deviations from steady state for unemployment, price and wage inflation, and percentage deviations for the other variables.

to an expansionary preference shock, $\varepsilon_{\Omega,t}$, around the two steady states introduced in the previous section. These two steady states only differ in the level of robot-specific technical change: the *low-automation* and the *high-automation* economy.

Figure C.4 depicts the results of this exercise for the baseline economy. Importantly, we set it up such that the unemployment gap drops by 1 percentage point (p.p.) on impact, and the entire response coincides across steady states. In this way, we can directly identify the effects of automation on the decoupling of unemployment and inflation by looking at the response of the latter.[4]

In the baseline model, the surge in automation across the two steady states reduces

[4]Since we consider preference shocks, the variation in unemployment coincides with that of the unemployment gap.

the impact response of price inflation by 14%, from 2.2 to 1.8 p.p. Also, wage inflation is less responsive, with a similar reduction of 14% across steady states, from 7.3 to 6.3 p.p. The responses of tightness and the automation cut-off to a preference shock shrink in the high-automation economy. That is consistent with the characterization of wage dynamics in Equation (C.4) discussed in the previous section.

Next, we tilt the bargaining power entirely towards firms and consider the economy with $\tau = 0.01$. Figure C.5 shows that there is virtually no longer any difference across the low-automation and high-automation steady states in the responses of both price inflation and wage inflation to preference shocks.[5] The same applies to changes in the automation cut-off. As a result, the relevance of the three channels virtually coincides across the low-automation and high-automation steady states. In other words, when shutting down the wage setting effect, also the variation in the selection and labor displacing effects become immaterial for the comparison of the wage dynamics across steady states; the low-frequency movements of automation do not alter inflation dynamics anymore. This analysis establishes that the effect of automation on inflation dynamics is mainly modulated through a wage setting effect: what dampens wage—and thus price—inflation is the option of automating labor tasks.

From this perspective, robot adoption alters inflation through a mechanism that resembles the one emphasized by Leduc and Liu (2024): the pro-cyclical threat of automation acts as a real wage rigidity that dampens wage and price pressures in booms while mitigating wage and price declines in recessions.[6,7] The implications of this mechanism are twofold. First, the effects of automation on inflation dynamics come from the combination of potential robot adoption (i.e., the threat of robot adoption) and actual robot installations. Second,

[5] Although the response of price inflation does not change across steady states, its magnitude is substantially lower than in the baseline economy. This is because tilting all the bargaining power to firms shifts most variation due to preference shocks into firms' surplus while dampening wage dynamics. As a result, the responsiveness of price inflation to shocks shrinks.

[6] As wages fall, the threat of automation becomes relatively less important, allowing workers to bargain for lower wage cuts.

[7] Appendix B.4 studies a setting in which automation is kept constant at its steady-state level, in the spirit of the exercise in Leduc and Liu (2024). In this case, changes in automation at the steady state have no impact on the selection, labor displacing and wage setting effects, since switching to machines is no longer possible, and thus does not pose any threat to workers.

Figure C.5: Impulse Responses in the Model with $\tau = 0.01$

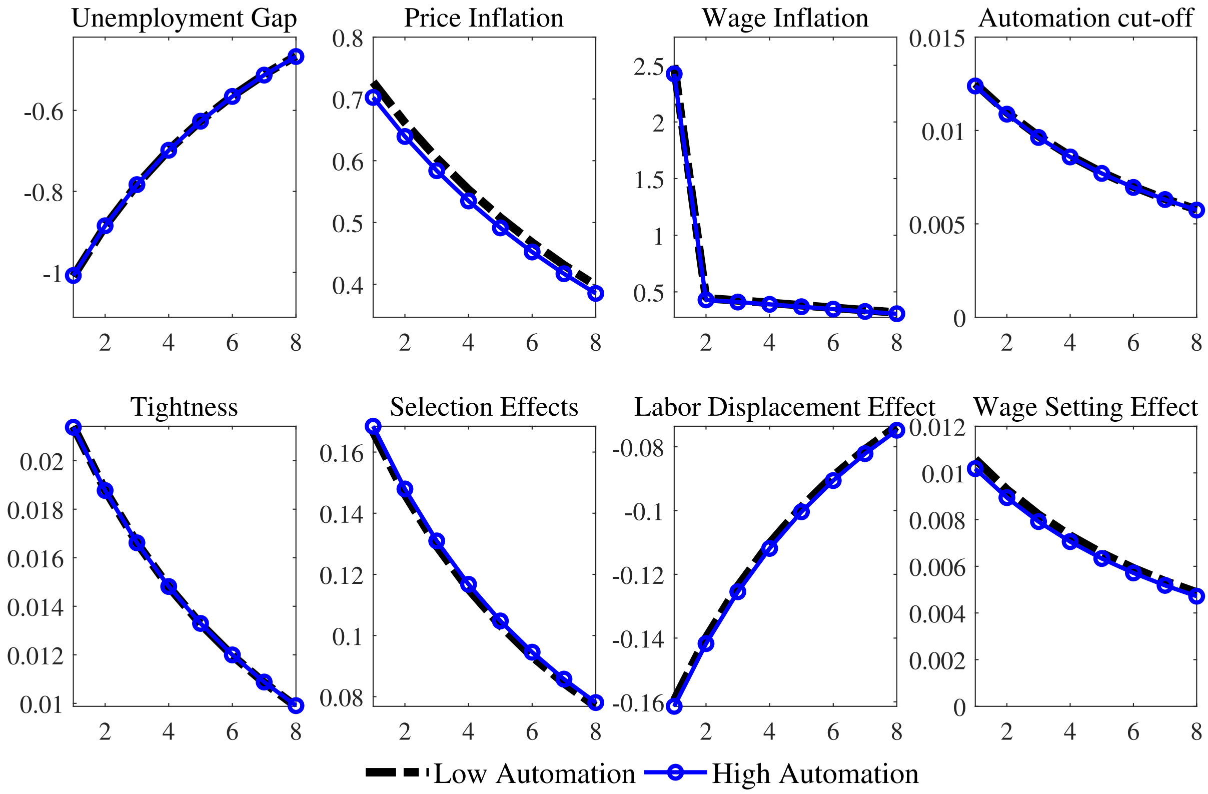

Note: Impulse responses in the first eight quarters after an expansionary preference shock of the unemployment gap, price inflation, wage inflation, the automation cut-off, workers' expected gains from searching for a job, the selection effect, the labor displacing effect, and the wage setting effect in the low-automation and high-automation steady states of the model economy in bargaining power is almost entirely tilted towards firms, that is, $\tau = 0.01$. Responses are in percentage-point deviations from steady state for unemployment, price and wage inflation and percentage deviations for the other variables.

the dampening effect of robot adoption on the relationship between inflation and unemployment can be generalized to any labor-displacing technology that poses a threat to workers' bargaining power.

C.4 The Role of the Robot Stock

First, the effect of automation depends critically on workers' bargaining power: when firms have full bargaining power, the outside option of automation is irrelevant, as firms already capture the entire surplus from the match. Second, this interaction matters only if automation can adjust in response to shocks. If automation varies across steady states but not around a steady state in response to a shock, the slope of the Phillips curve does not vary as a function of the steady-state stock of robots.

In this section, we provide additional evidence underscoring the importance of the stock of robots. While all previous exercises already assume that automation adjusts across steady states—establishing it as a necessary condition for affecting inflation dynamics—this exercise serves two complementary goals: (i) it further illustrates the role of the robot stock in shaping the Phillips curve, and (ii) it highlights the joint influence of automation and workers' bargaining power in determining the strength of this effect.

To do so, we devise an exercise as follows. We take the baseline economy in which the workers' bargaining power is $\tau = 0.5$ and start from the low automation steady state in which the the robot-to-employee ratio is 0.2%. Then, we compute much the degree of automation has to increase in percentage terms in a second new steady state such that the inflation responds is 10% less for a 1 percentage point change in unemployment. Then, we replicate the same analysis for two distinct values of workers' bargaining power: $\tau = 0.45$ and $\tau = 0.55$.

We report the results in Table C.17. In the baseline economy, with workers' bargaining power set at $\tau = 0.5$, the degree of automation must increase by 132% across steady states to achieve a 10% flattening of the Phillips curve—implying a new robot-to-employee ratio of 0.046%. This required change varies substantially with bargaining power. When $\tau = 0.45$, automation must rise by 172%, leading to a new robot stock of 0.054%. In contrast, when $\tau = 0.55$, the same flattening can be achieved with just a 73% increase, yielding a lower robot stock of 0.035%.

Thus, this exercise shows how different values of stock of robots are required to match the same flattening as we change the bargaining power. Thus, the interaction between these two dimensions is key to understand the effect of the threat of automation on wage and price dynamics: the effectiveness of automation in moderating inflation dynamics depends on how much leverage workers have in wage bargaining. The higher the bargaining power, the stronger the disciplining effect of the automation threat—requiring a smaller change in the stock of robots to generate the same outcome. In addition, these results highlight the

Table C.17: Change in Degree of Automation across Steady States

	Target: 10% Flattening		
	$\tau = 0.5$	$\tau = 0.45$	$\tau = 0.55$
$\Delta\Xi_{M,SS}$	132%	172%	73%

Note: Starting from a low-automation steady state, in which the robot-to-employee ratio is 0.2%, we report how much the degree of automation has to increase in percentage terms in the new steady state such that the inflation responds is 10% less for a 1pp change in unemployment in the new steady state for different values of the workers' bargaining power.

critical importance of the steady-state level of automation in determining the strength of the threat of automation on price dynamics.

C.5 Model with Monetary Policy and Productivity Shocks

In our analysis of Section C.3, we study the response of price inflation, wage inflation to expansionary preference shocks, $\varepsilon_{\Omega,t}$, and disentangle the relevance of the wage setting, selection, and labor displacing effects of automation. In this section, we show that our results are robust to studying additional sources of exogenous variation. Specifically, we study the changes in the responsiveness of price and wage inflation and the changes in relevance of the three effects across the low- and high-automation economies produced by two alternative sources of exogenous uncertainty.

In the first case, we consider a monetary policy shock. Specifically, we alter the Taylor rule of Equation (25) by adding the monetary policy shock $\varepsilon_{R,t}$ as follows

$$R_t/\bar{R} = \left[R_{t-1}/\bar{R}\right]^{\psi_R} \left[(1+\pi_t)^{\psi_\pi} \left(u_t/u_t^F\right)^{\psi_u}\right]^{1-\psi_R} + \varepsilon_{R,t}. \tag{C.8}$$

As we have done for the case of the preference shocks, we compare the responses to an expansionary monetary policy shock of the four key variables by equalizing the response of unemployment in the two low and high automation economies, such that on impact,

Figure C.6: Impulse Responses to Monetary Policy Shocks

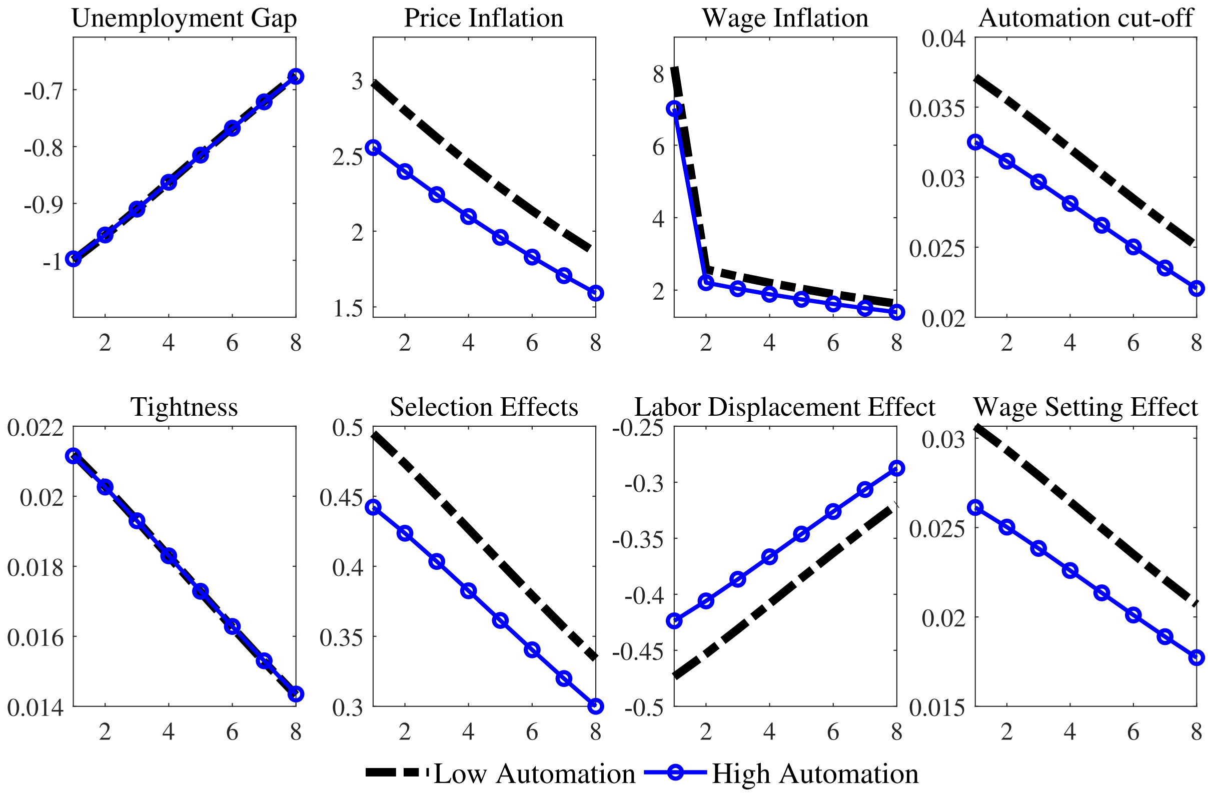

Note: Impulse responses in the first eight quarters after a monetary policy shock of the unemployment gap, price inflation, wage inflation, the automation cut-off, workers' expected gains from searching for a job, the selection effect, the labor displacing effect, and the wage setting effect in the low-automation and high-automation steady states of the baseline economy. Responses are in percentage-point deviations from steady state for unemployment, price and wage inflation, and percentage deviations for the other variables.

unemployment drops by 1 p.p. We report the results of this exercise in Figure C.6, which confirms the patterns derived by the responses associated with the preference shock: low-frequency changes in automation across the two steady states lead to a muted sensitivity of price inflation for any given change in the unemployment gap: the impact response of price inflation goes from 3 p.p. in the low-automation economy to 2.6 p.p. in the high-automation economy, a drop by 14%. This effect hinges on the same 14% reduction in the responsiveness of wage inflation: the impact response of wage inflation shrinks from 8.1 p.p. to 7 p.p.

In the second case, we consider the effect of a productivity shock. Specifically, we alter the wholesalers' technology of Equation (17) as follows

$$Y_{i,t} = A_t Z_{i,t}, \tag{C.9}$$

Figure C.7: Impulse Responses to Productivity Shocks

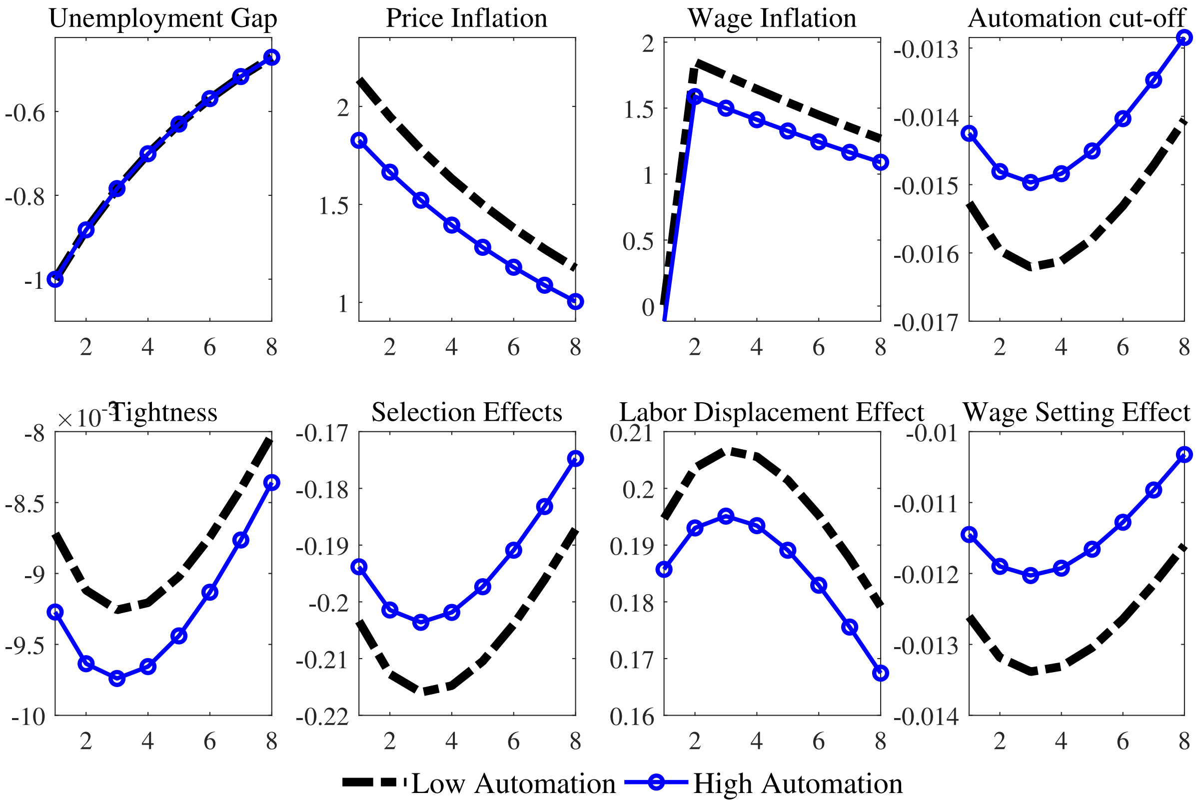

Note: Impulse responses in the first eight quarters after a productivity shock of the unemployment gap, price inflation, wage inflation, the automation cut-off, workers' expected gains from searching for a job, the selection effect, the labor displacing effect, and the wage setting effect in the low-automation and high-automation steady states of the baseline economy. Responses are in percentage-point deviations from steady state for unemployment, price and wage inflation, and percentage deviations for the other variables.

where A_t is the productivity level that is common across all wholesalers. The logarithm of productivity follows the first-order auto-regressive process

$$\log A_t = \rho_A \log A_{t-1} + \varepsilon_{A,t}, \tag{C.10}$$

in which ρ_A captures the persistence of the process, and $\varepsilon_{A,t}$ is the productivity shock. We set the persistence to $\rho_A = 0.9$ so to generate exactly the same auto-correlation patterns for the shock process as those associated with the preference process of Equation (24).

As in the previous case, we compare the responses to a contractionary productivity shock of the four key variables again by equalizing the response of the unemployment gap in the two low- and high-automation steady states, such that on impact the unemployment gap drops by one percentage point. Note that in this setting, a contractionary productivity shock raises the unemployment rate but it does so to a relatively lower extent than in the model specification with flexible prices. As a result, the unemployment gap goes down. We report the results in Figure C.7. Also in this case, we find that the response of price inflation shrinks in the high-automation steady state compared to the low-automation one: the impact response of price inflation drops by 14%, from 2.2 p.p. to 1.8 p.p. Also the responsiveness of wage inflation at the peak shrinks by the same amount, from 1.9 p.p. to 1.6 p.p.

Importantly, independently of whether we consider the monetary policy shock or the productivity shock, we confirm two key findings of the analysis based on preference shocks: (*i*) the responsiveness of both the automation cut-off and workers' expected gains from searching for a job decreases when moving from the low automation to the high automation steady state; and (*ii*) an increase in robot adoption at the steady state mutes the relevance of the selection, labor displacing, and wage setting effects, thus dampening the responsiveness of the real wage to shocks.

C.6 Model with Fixed Automation

In this section, we run an exercise in the spirit of Leduc and Liu (2024), in which while we let robot adoption to vary between the low automation and high automation steady state, we impose that producers cannot change robot installations upon a shock. In other words, automation is kept constant at its steady-state level. We then run the same set of impulse responses of Figure C.4.

The results in Figure C.8 indicate that all variables display virtually the same response to an expansionary preference shock in both the low automation and high automation steady state. This means that changes in the steady state robot adoption are immaterial for price and wage dynamics. Hence, this result numerically confirms the analytical result of Proposition 3 in Section 4, in which we have transparently derived how the possibility of adjusting robots upon a shock is a necessary condition for changes in steady-state automation to im-

Figure C.8: Impulse Responses in the Baseline Model with Fixed Short-run Automation

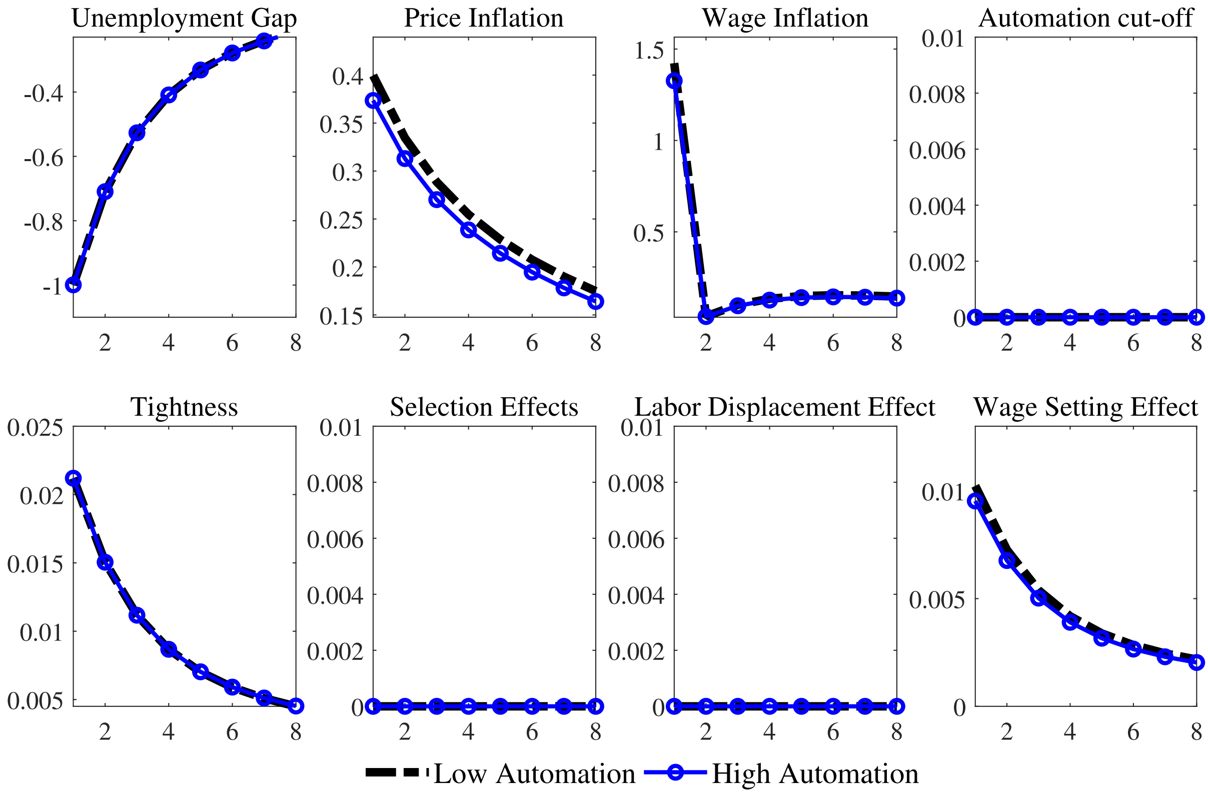

Note: Impulse responses in the first eight quarters after an expansionary preference shock of the unemployment gap, price inflation, wage inflation, the automation cut-off, workers' expected gains from searching for a job, the selection effect, the labor displacing effect, and the wage setting effect in the low-automation and high-automation steady states of the baseline economy holding automation fixed in the short-run. Responses are in percentage-point deviations from steady state for unemployment, price and wage inflation, and percentage deviations for the other variables.

pact inflation dynamics. This is also consistent with our numerical analysis in Section 5.3, showing that the slope of the Phillips curve hardly varies with steady-state automation and workers' bargaining power if robots cannot adjust upon a shock.

We can understand this same result also through a different perspective, by looking at the characterization of the dynamics of the real wage of Equation (C.4). From this condition, it follows immediately that fixing automation at its steady state shuts down both the selection and labor displacement effects, as these two channels require that producers can switch to machines upon a shock. As a result, robot adoption does not pose anymore a threat to wage negotiations. This explains why also the responsiveness of workers' expected gains from searching for a job to the preference shock coincides across the two steady states.

Once more, this analysis highlights that allowing robots to pose a threat to workers, and thus curtail their effective bargaining power, is the key channel that modulates the effect of changes in steady-state levels of automation on inflation dynamics. From this perspective, robot adoption alters inflation through a mechanism that coincides with that emphasized by Leduc and Liu (2024): the pro-cyclical threat of automation acts as a real wage rigidity that dampens wage and price pressures in booms, while mitigating wage and price declines in recessions.

D Model Extensions

This section evaluates the implications of different settings, in which we either change the structure of the economy relative to the baseline model, or consider a different specification of the transition dynamics exercise. Section D.1 studies the shock-size-dependency of the inflation response to expansionary shocks in a model version featuring adjustment cost in automation, Section D.2 uncovers the implied inflation sensitivity to unemployment when we add wealth effects on labor supply, Section D.3 modifies the specifics of the transition exercise by allowing the unemployment and participation rates to vary with the level of robot technical change, Section D.4 introduces physical capital, and Section D.5 recovers the inflation sensitivity to unemployment in an economy with two agents: Ricardian and hand-to-mouth households. Section D.6 considers a setting in which labor matches last longer than one period, and finally **??** relaxes the assumption that workers and firms take the automation cut-off as given while bargaining the wage.

D.1 Model with Costly Automation

Automation flattens the slope of the Phillips curve by curbing workers' bargaining power. In this appendix, we show that the relevance of this mechanism may depend on the size of the shock hitting the economy. The aim of this exercise is to uncover that while robot adoption reduces the sensitivity of inflation to labor market conditions, this mechanism is not inconsistent with the possibility of a sudden resurgence of a steep Phillips curve. To do so, we provide the conditions under which our economy can generate an inflation spike notwithstanding a higher degree of automation. Specifically, when ramping up robot adoption in the short term is costly, the effect of automation on the slope of the Phillips curve is shock-size-dependent: robot adoption flattens the Phillips curve upon small shocks but not amidst large shocks.

To uncover this result, we introduce asymmetric convex adjustment costs as in Varian

(1975) into the production function of machine manufacturers:

$$M_t = \left[\zeta - \Delta\left(I_t/\bar{I}\right)\right] I_t, \tag{D.11}$$

where $\Delta\left(I_t/\bar{I}\right)$ denotes the asymmetric adjustment costs, such that

$$\Delta\left(I_t/\bar{I}\right) = \frac{\delta}{\varrho^2}\left\{\exp\left[\varrho\left(\frac{I_t}{\bar{I}} - 1.05\right)\right] + \varrho\left(\frac{I_t}{\bar{I}} - 1.05\right) - \vartheta\right\}, \tag{D.12}$$

in which δ controls the magnitude of the cost, ϱ is the degree of asymmetry (i.e., adjustment costs are symmetric when $\varrho \to 0$), and ϑ is a residual that ensures the steady-state cost is zero. In this specification, the adjustment costs kick in when investment is 5 p.p. above its steady state.

The adjustment cost function implies that the price of machines rises substantially when robot production ramps up. A higher investment demand implies a surge in the adjustment costs that boosts the price of machines. The outside option of automation ceases to be profitable for some producers, which then opt to operate with the labor technology. Workers exploit this situation by negotiating higher wages, which are passed into prices. In this setting, large expansionary preference shocks that substantially reduce unemployment are accompanied by a spike in wages and prices, leading to the resurgence of a steep Phillips curve. Indeed, when the price of robots hikes up, also the low-efficiency firms that keep operating using robots after the shock face an increase in marginal costs, putting further upward pressure on prices.

The quantitative analysis of this mechanism is challenged by the uncertainty on how to discipline the cost function and the size of the shock. We thus interpret it as a proof-of-concept illustration of how costly robot adoption can account for a steep Phillips curve upon large shocks. Our calibration choices capture the idea that while the adjustment costs are negligible for limited investment changes, any variation in investment from its steady state that is above 5% leads to a convexly increasing cost. We set the cost parameters to $\delta = 0.0015$ and $\varrho = 100$ to imply that a 7 p.p. increase in investment above its steady-state level implies a cost which is three times as large as that associated to a 6 p.p. increase.

Figure D.9: Costly Automation

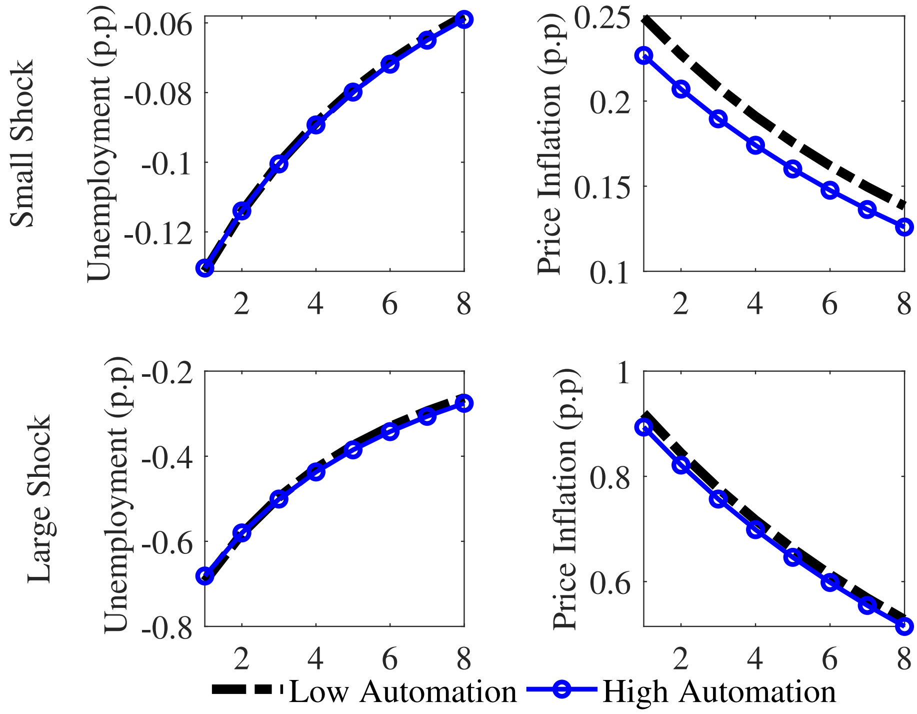

Note: Impulse responses in the first eight quarters of the unemployment gap (in percentage-point deviations from steady state) and price inflation (in annualized percentage-point deviations from steady state) to a small and large preference shocks designed to generate similar responses in unemployment in economies with high and low degrees of automation.

We then consider the response of the low-automation and high-automation economies with costly robot adoption to two expansionary preference shocks that only differ in their size: a small shock decreasing unemployment by 0.13 p.p., and a large shock reducing unemployment by 0.7 p.p..

The results of the exercise are reported in Figure D.9. The graphs at the top show the response of unemployment and inflation under the small shock realization, whereas the two graphs at the bottom indicate the behavior of the same two variables following the large shock. The responses of unemployment in the low- and high-automation economies are equalized. Upon a small shock, inflation surges relatively less in the high-automation economy, confirming that robot adoption flattens the price Phillips curve. However, when the shock is large, there is no difference in the response of inflation in the low-automation and high-automation economies. Thus, while an increase in automation flattens the price Phillips curve when the size of the shock is small, automation does not influence at all the inflation sensitivity to unemployment amidst large shocks, leading to the resurgence of a steep Phillips curve.

D.2 Model with Wealth Effects on Labor Supply

In our baseline model, households' participation decision to search for a job in the labor market is not influenced by their marginal utility of consumption. We abstract from wealth effects on labor supply since Hazell et al. (2022) show that they prevent a direct mapping between the slopes of the regional and aggregate Phillips curves. This restriction circumvents the missing intercept problem, allowing us to use our model as a laboratory to directly interpret and dissect the empirical evidence.

In this section, we evaluate the quantitative implications of a model specification that does feature wealth effects on labor supply and show that, although this specification does not match the reduced form estimated on the effects of automation on the price inflation sensitivity to unemployment, it also delivers a flattening of price Phillips curve implied by the baseline model. In other words, the wealth effects on labor supply do not significantly alter the key implications of the model.

Specifically, we modify the equilibrium condition $N_t = w_t/\chi_H$ of the baseline model that governs households' labor market participation decision as follows

$$N_t = w_t \frac{U_{C,t}}{\chi_H}, \tag{D.13}$$

where $U_{C,t}$ is households' marginal utility of consumption. We then study the quantitative implications of this model specification in two dimensions: we replicate the impulse-response functions of Figure C.4, and derive a transitional dynamics exercise as in Table 3.

Figure D.10 reports the response of price inflation and wage inflation to an expansionary preference shock that reduces unemployment by 1 p.p. in both the low automation and high

Figure D.10: Impulse Responses in a Model with Wealth Effects on Labor Supply

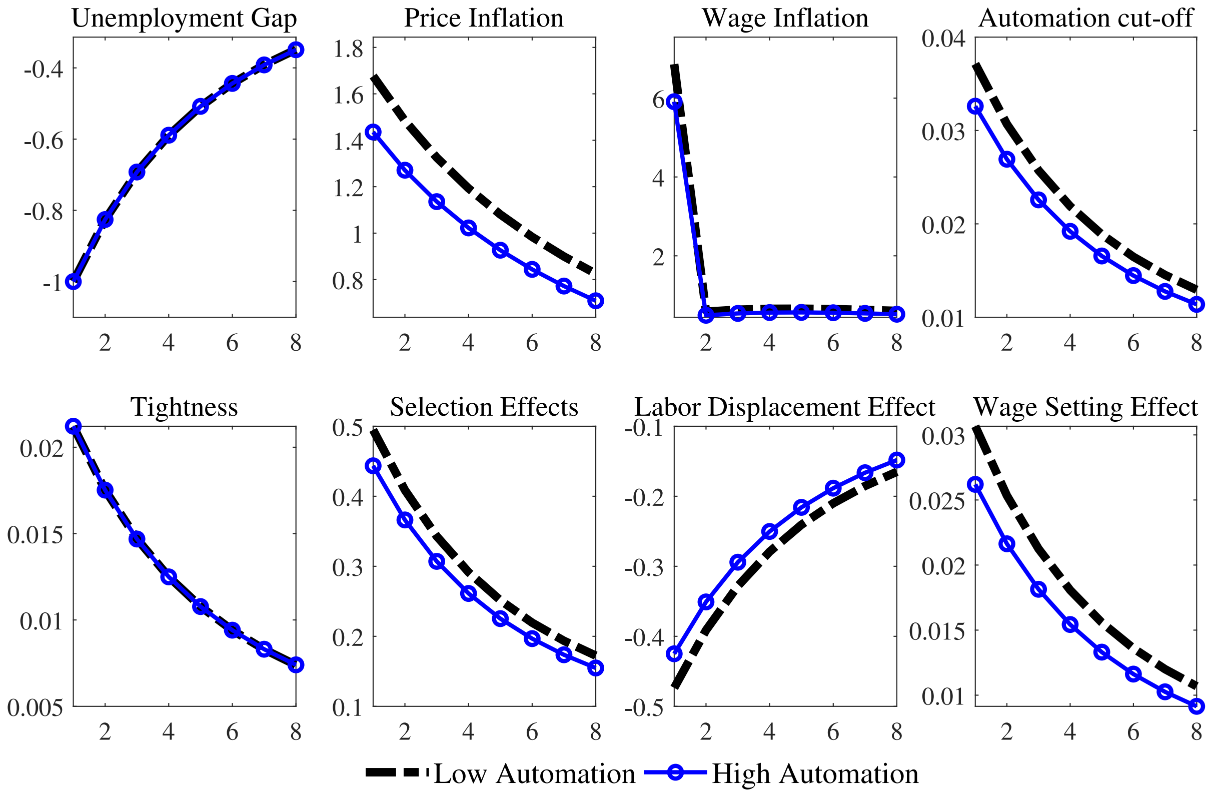

Note: Impulse responses in the first eight quarters after an expansionary preference shock of the unemployment gap, price inflation, wage inflation, the automation cut-off, workers' expected gains from searching for a job, the selection effect, the labor displacing effect, and the wage setting effect in the low-automation and high-automation steady states of the model economy featuring wealth effects on labor supply. Responses are in percentage-point deviations from steady state for unemployment, price and wage inflation, and percentage deviations for the other variables.

automation steady states. We find that higher automation at the steady state reduces the responsiveness of price inflation by 14%, from 1.7 p.p. to 1.5 p.p. Again, the reduction in the responsiveness of wage inflation coincides with that of price inflation, with a drop from 6.8 p.p to 5.9 p.p. Although the magnitude of the responses is slightly lower than in the baseline model which abstracts from wealth effects on labor supply, the reduction in the responsiveness of price and wage inflation across the two steady states matches exactly that generated by the baseline model.

We then turn into the transitional dynamics exercise, estimate regression (30) using data simulated from the model specification featuring wealth effects on labor supply, and compare the estimates with those produced with the data of the baseline economy. We report the

Table D.18: Transitional Dynamics Exercise - Wealth Effects on Labor Supply

	Dependent Variable: π_t		
	Baseline Model	Wealth Effects Labor Supply	Data MSA-level
	(1)	(2)	(3)
u_t	-0.5036	-0.4362	-0.5248
	(0.0005)	(0.0008)	(0.1469)
$u_t \times (m_t - \bar{m})$	0.0058	0.0033	0.0058
	(0.0003)	(0.0005)	(0.0028)

Note: The table reports the estimates of the regression in Equation (30) based on simulated data of the baseline model (in Column 1) and the alternative model specification with wealth effects on labor supply (in Column 2). Column 3 reports the estimate based on a panel regression across U.S. MSAs which correspond to the case with non-tradable inflation as dependent variable and IV method of Table 1. The simulated data of the baseline model are based on a transitional dynamics exercise between a low-automation steady state, in which the robot-to-employee ratio is 0.2%, and a high-automation steady state, in which the robot-to-employee ratio is 0.8%.

results in Table D.18. We find that the estimates of the alternative model specification are relatively lower than those of the baseline model: the estimate of the average price inflation sensitivity to unemployment is -0.5036 in the baseline model and -0.4362 with the wealth effect, and the estimate of the interaction term of unemployment with robot adoption equals 0.0058 is in the baseline model and 0.0033 in this new specification. Thus, relative to the baseline economy and the data, the model with wealth effects understates both the average inflation sensitivity to unemployment and the effect of automation on this sensitivity. However, also in this alternative specification robot adoption does mute the relationship between inflation dynamics and labor market conditions.

From this perspective, these results confirm the insights of Hazell et al. (2022): an aggregate-level economy that abstracts from wealth effects on labor supply can correctly and directly map the slopes of the regional and aggregate Phillips curves. However, we add by showing that although the presence of the wealth effects impedes a correct mapping with the data, it does not alter the implications on the way in which automation alters the elasticity of price inflation to changes in the unemployment rate.

D.3 Baseline Model - Transition with Varying Unemployment

In the quantification of the effect of robot adoption in flattening the Phillips curve of Section 5.4, we consider a transition dynamics in which we start from a low automation steady state, which proxies the U.S. economy in the early 2000s, and progressively increase the level of robot-specific technical change, ζ, so that the economy reaches a high automation steady state, which describes the U.S. economy in the late 2010s. Throughout the transition, we also change the entry cost κ and searching cost λ_H to preserve the same unemployment and participation rates across steady states. This choice is consistent with the lack of trends in both the unemployment rate and the participation rate over the first two decades of the 2000s.

In this section, we replicate the exercise in Section 5.4 with the only difference that this time we only vary the level of robot-specific technical change, ζ, but keep fixed at their initial values both the entry cost κ and the searching cost λ_H. In this way, the unemployment rate increases throughout the transition. We also find that the participation slightly increases as well.

Again, we take the simulated data on inflation, the unemployment rate, and robot adoption, and estimate exactly the same regression of Equation (1), that is:

$$\pi_t = \beta u_t + \gamma u_t \left(m_t - \bar{m}\right) + \varsigma m_t + \epsilon_t, \tag{D.14}$$

We report the results of this exercise in Table D.19.

We find that varying the unemployment as limited effects of the way in which automation reduces the inflation sensitivity to labor market conditions. While the baseline model implies that the increase in robot installations in the U.S. over the last three decades predicts a 9% drop in the inflation sensitivity to unemployment, the alternative specification of this section yields a 8% drop in the inflation sensitivity.

Table D.19: Transitional Dynamics Exercise - Varying Unemployment

	Dependent Variable: π_t		
	Baseline Model	Varying Unemployment	Data MSA-level
	(1)	(2)	(3)
u_t	-0.5036	-0.5087	-0.5248
	(0.0005)	(0.0005)	(0.1469)
$u_t \times (m_t - \bar{m})$	0.0058	0.0049	0.0058
	(0.0003)	(0.0003)	(0.0028)

Note: The table reports the estimates of the regression in Equation (30) based on simulated data of the baseline model (in Column 1) and the alternative exercise in which we not keep fixed the unemployment and participation rates throughout the transition (in Column 2). Column 3 reports the estimate based on a panel regression across U.S. MSAs which correspond to the case with non-tradable inflation as dependent variable and IV method of Table 1. The simulated data of the baseline model are based on a transitional dynamics exercise between a low-automation steady state, in which the robot-to-employee ratio is 0.2%, and a high-automation steady state, in which the robot-to-employee ratio is 0.8%.

D.4 Model with Physical Capital

We then consider a version of the model that incorporates physical capital as Christiano et al. (2005), so that investment is subject to convex adjustment costs and capital features variable capacity utilization. We change the economy by considering that the household (i) owns the stock of capital, K_t, (ii) chooses the level of utilization, u_t, and (ii) rents out capital services, $\tilde{K}_t = u_t K_t$, to wholesalers at a rental rate $R_{K,t}$.

We start by positing that the monopolistically competitive wholesalers produce different varieties $Y_{i,t}$ by assembling together the goods $Z_{i,t}$ purchased from producers at price $P_{P,t}$, and capital services $\tilde{K}_{i,t}$ rented from the household at the nominal rental rate $R_{K,t}$. To do so, the wholesalers use a Cobb-Douglas technology:

$$Y_{i,t} = Z_{i,t}^{1-v} \tilde{K}_{i,t}^{v}. \tag{D.15}$$

where v denotes the share of capital in wholesalers' output. As in the baseline model, varieties are sold to the retailer at price $P_{i,t}$, so that profits equal $P_{i,t}Y_{i,t} - P_{P,t}Z_{i,t} - R_{K,t}\tilde{K}_{i,t}$.

Then, the retailers assembles wholesalers' varieties into the final good, Y_t, with the same CES technology of the baseline model in Equation (19). Final goods are sold at price P_t to household, in form of consumption C_t, investment I_t, and to machine manufacturers, in form of robot intermediate inputs A_t, such that

$$Y_t = C_t + I_t + A_t. \tag{D.16}$$

The machine manufacturers purchase the robot intermediate inputs, A_t, and transform them into robots with the linear technology

$$M_t = \zeta A_t, \tag{D.17}$$

where ζ is the level of robot-specific technical change.

Regarding the households' problem of Equation (22), households now maximizes utility also by optimally choosing the aggregate stock of capital, K_{t+1}, investment, I_t, and the capacity utilization, u_t, such that:

$$\max_{C_t,B_{t+1},K_{t+1},I_t,u_t} \mathbb{E}_t \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\sigma}}{1-\sigma} \tag{D.18}$$

$$\text{s.t. } P_tC_t + B_t + P_tI_t = B_{t-1}R_{t-1} + K_tR_{K,t}u_t + X_t + D_t + T_t, \tag{D.19}$$

$$K_{t+1} = \left[1 - \delta\left(u_t\right)\right]K_t + I_t\left[1 - \frac{\varkappa}{2}\left(\frac{I_t}{I_{t-1}} - 1\right)^2\right], \tag{D.20}$$

$$\delta\left(u_t\right) = \delta_0 + \delta_1\left(u_t - 1\right) + \delta_1\frac{\delta_2}{2}\left(u_t - 1\right)^2, \tag{D.21}$$

where Equation (D.20) is the law of motion of physical capital, which depreciates at a rate, $\delta\left(u_t\right)$, that varies with the degree of utilization, following a functional form defined in Equation (D.21). The parameter $\varkappa$ captures the adjustment costs, δ_0 denotes depreciation in steady state (i.e., when $\bar{u} = 1$), and δ_1 and δ_2 captures the extent to which higher capital utilization leads to relatively faster depreciation rates.

Finally, the market clearing of physical capital implies that the total capital services demanded by wholesalers equals that rented out by the household, that is:

$$\int_0^1 \tilde{K}_{i,t} \mathrm{d}i = \tilde{K}_t = u_t K_t. \tag{D.22}$$

To discipline this economy, we have five new parameters relative to the baseline model that we need to calibrate: the capital share in production υ, the steady-state depreciation rate δ_0, the parameters that determine the relationship between utilization and depreciation, δ_1 and δ_2, and the adjustment cost parameter $\varkappa$. We set the capital share to the standard value of $\upsilon = 0.4$ and the depreciation rate in steady state to $\delta = 0.015$ to match the implied depreciation rate of the BEA as computed by Ramey (2021). The choice that capacity utilization in the steady state equals to $\bar{u} = 1$ implies that the following holds: $\bar{R}_K = \frac{1}{\beta} - (1 - \delta_0)$. If we evaluate the first order condition of utilization in the steady state, we have that $\bar{R}_K = \delta_1 + \delta_1 \delta_2 (\bar{u} - 1)$. Combining the last two conditions yields that $\delta_1 = \frac{1}{\beta} - (1 - \delta_0)$. We then set the remaining utilization parameter, $\delta_2 = 0.75$, so that the model matches the ratio between the standard deviation of capacity utilization, as estimated by Fernald (2014), and the standard deviation of GDP in the data. Finally, we set the adjustment cost to $\varkappa = 3$ to match the ratio between the standard deviation of investment and the standard deviation of GDP.

We run the exercise similarly to Section 5.4, take the simulated values of the model, and estimate regression (30). We report the results in Table D.20. In this case, the inflation sensitivity to unemployment increases substantially, more than doubling with respect to the baseline. To understand this result, note that we calibrate this alternative model using exactly the same value of the Rotemberg cost of the baseline economy. Although in the baseline economy this choice allowed us to match the slope of the Phillips curve implied by our empirical evidence, in this case we find a slope which is four times larger. From this perspective, a doubling of the inflation sensitivity notwithstanding a slope four times as large implies that physical capital is actually muting the inflation responsiveness to labor market conditions.

Table D.20: Transitional Dynamics Exercise - Physical Capital

	Dependent Variable: π_t		
	Baseline Model	Physical Capital	Data MSA-level
	(1)	(2)	(3)
u_t	-0.5036	-1.3231	-0.5248
	(0.0005)	(0.0237)	(0.1469)
$u_t \times (m_t - \bar{m})$	0.0058	0.0165	0.0058
	(0.0003)	(0.0093)	(0.0028)

Note: The table reports the estimates of the regression in Equation (30) based on simulated data of the baseline model (in Column 1) and the alternative economy with physical capital (in Column 2). Column 3 reports the estimate based on a panel regression across U.S. MSAs which correspond to the case with non-tradable inflation as dependent variable and IV method of Table 1. The simulated data of the baseline model are based on a transitional dynamics exercise between a low-automation steady state, in which the robot-to-employee ratio is 0.2%, and a high-automation steady state, in which the robot-to-employee ratio is 0.8%.

In this context, we also find that the magnitude of the dampening effect of automation on the sensitivity more than doubles relative to the baseline. To quantify this magnitude, we replicate the same comparison of the baseline exercise: if we consider the increase in robot installations in the U.S. over the last three decades, the model predicts a 12.5% drop in the inflation sensitivity to unemployment, slightly higher than in the baseline.

All in all, this exercise showcases that the presence of physical capital does not alter our main finding on the way in which automation flattens the Phillips curve.

D.5 Model with Hand to Mouth

In this section, we extend the baseline economy such that it features a two-agent structure, with Ricardian and hand-to-mouth households, in the spirit of Galí et al. (2007). Specifically, the economy is populated by a fraction ς of hand-to-mouth individuals (H), who are indexed by $x \in [0, \varsigma]$ and form the hand-to-mouth household, and a fraction $1 - \varsigma$ of Ricardian individuals (R), who are indexed by $x \in [\varsigma, 1]$ and form the Ricardian household.

All individuals of either type supply labor, and, as in baseline model, they decide whether to look for a job after drawing a searching cost, χ, from a uniform distribution, $U(0, \chi_H)$. Individuals decide to work if the real wage exceeds the searching cost. Thus, in equilibrium, the participation rate of hand-to-mouth individuals is

$$N_{H,t} = \varsigma w_t / \chi_H, \tag{D.23}$$

and that of Ricardian households only differs in the terms of the distinct measure of individuals, and reads

$$N_{R,t} = (1 - \varsigma)\, w_t / \chi_H, \tag{D.24}$$

Since both types of individuals have the same unit efficiency of working hours and face the same searching cost, they also share the same participation constraint. As a result, the total measure of participating households is

$$N_t = N_{H,t} + N_{R,t}. \tag{D.25}$$

The Ricardian household sets its consumption, $C_{R,t}$, optimally following exactly the same problem as in the baseline model. Instead, the hand-to-mouth households consume their entire labor income at every period, such that their consumption, $C_{H,t}$, can be written as

$$C_{H,t} = \frac{X_t}{P_t} \tag{D.26}$$

where X_t denotes total nominal labor earnings. Consequently, total consumption in this economy equals:

$$C_t = \varsigma C_{H,t} + (1 - \varsigma)\, C_{R,t}, \tag{D.27}$$

where the weight of each household is given by the its share of individuals in the economy.

To discipline this economy, we set the fraction of hand-to-mouth households to 30%, in line with the evidence of Weidner et al. (2014) on the U.S. regarding the sum of poor and wealthy hand-to-mouth households as a fraction of total households. We then replicate the

Table D.21: Transitional Dynamics Exercise - Hand-to-Mouth and Ricardian Households

	Dependent Variable: π_t		
	Baseline Model	TANK Economy	Data MSA-level
	(1)	(2)	(3)
u_t	-0.5036	-0.5082	-0.5248
	(0.0005)	(0.0005)	(0.1469)
$u_t \times (m_t - \bar{m})$	0.0058	0.0065	0.0058
	(0.0003)	(0.0003)	(0.0028)

Note: The table reports the estimates of the regression in Equation (30) based on simulated data of the baseline model (in Column 1) and the alternative economy with two type of agents, hand-to-mouth and Ricardian households (in Column 2). Column 3 reports the estimate based on a panel regression across U.S. MSAs which correspond to the case with non-tradable inflation as dependent variable and IV method of Table 1. The simulated data of the baseline model are based on a transitional dynamics exercise between a low-automation steady state, in which the robot-to-employee ratio is 0.2%, and a high-automation steady state, in which the robot-to-employee ratio is 0.8%.

exercise of Section 5.4: we simulate the model, recover the values of inflation, unemployment, and automation, and estimate regression (30). The results in Table D.21 indicate that both the sensitivity of inflation to unemployment and the way in which automation alters it are very much in line with the estimates based on the baseline model. If anything, in this model specification, the magnitude of the effect of automation in flattening the Phillips curve is slightly amplified. In this case, the increase in robot installations in the U.S. over the last three decades predicts a 10% drop in the inflation sensitivity to unemployment, an effect one percentage point larger than in the baseline exercise.

D.6 Model with Long-Lasting Job Matches

In our baseline economy, job matches last for one period. This happens because entering firms operate just for one period, and then exit the economy. While this assumption is stark, the results in Section A.5 indicate that the effect of automation on reducing the inflation sensitivity to unemployment does not vary with local labor market characteristics, such as hiring rates, separation rates, and job-to-job transition rates. Thus, our baseline modeling of the labor market is consistent with the evidence ruling out a role for labor flows in modulating the influence of robot adoption on inflation dynamics.

However, this section shows that the main conclusions of our economy do not change if we extend it by incorporating long-lasting job matches by allowing firms to operate for more than one period. In this version, job separations emerge both exogenously as well as endogenously, in case firms decide to terminate an employment relationship to install a robot. Specifically, the economy features three types of incumbent producers:

(i) labor firms, i.e. the producers that post a vacancy and successfully fill it;

(ii) idle firms, i.e. the producers that do not manage to match with a worker, and thus do not produce;

(iii) robot firms, i.e. the producers that install a robot.

In addition, every period a new vintage of firms enters the economy.

The decisions of incumbent firms vary with their type. Labor firms can keep operating with the same workers, or substitute them with robots. In addition, the labor firms that keep operating with the same workers face an exogenous job separation probability, $1 - \Theta_t$. Thus, labor firms can either decide to substitute their workers with robots, triggering an endogenous job separation (displacement effect) or can be forced into an exogenous job separation. Instead, idle and robot firms can post a vacancy and look for a worker, or operate with a machine. Then, at the end of every period, producers may exogenously exit the economy with probability δ, which provides a second source of exogenous job separations.

On the other side, workers decide whether to participate in the labor market and look for a job. In each period, the total pool of individuals searching for a job consists of previously unemployed people, the individuals that have been laid off, either exogenously (due to firm exit) or endogenously (displacement due to automation), and the new fraction of individuals that decide to participate in the labor market.

The timing of the model is as follows. First, at the beginning of the period, entrant firms join the economy and incumbent firms decide which option to pursue in terms of their production structure. Consequently, some labor incumbents may decide to endogenously separate from their workers, and opt to install a machine. In this case, the endogenously displaced workers join the current pool of unemployed workers. Second, the labor market takes place, and individuals search for jobs posted by the producers. With some probability firms and individuals match, and thus the firm becomes a labor firm and the individual is employed. Otherwise, the firm becomes an idle producer, and the individual enters the pool of unemployed carrying over to the following period. Third, after the labor market has taken place, a fraction of incumbent labor firms experience an exogenous job separation. As a result, the firm becomes an idle producer, and the worker joins the pool of unemployed carrying over to the following period. Fourth, production takes place, using workers and machines. Fifth, as we transition to a new period a fraction of producers exogenously exit the economy. In this case, the workers join the pool of unemployed and a new period begins.

Let us denote by $\Xi_{L,t}$, $\Xi_{I,t}$, and $\Xi_{M,t}$ the measures of labor firms, idle firms, and machine firms at the end of each period. Upon entry producers draw a level of labor efficiency, γ, from a distribution $f(\gamma;\alpha)$ with support $[1,\gamma_H]$ and shape parameter α. The level of labor efficiency remains fixed for all subsequent periods.

Labor firms face two options: either keep the previously hired workers, or substitute them with a machine. The value of a labor firm that keeps the previous employment relationship, $V_{L,t}(\gamma)$, is:

$$V_{L,t}(\gamma) = \Theta_t \left[P_{P,t}\gamma - W_t\right] + \beta (1-\delta) V_{C,t+1}(\gamma), \tag{D.28}$$

where $V_{C,t+1}(\gamma)$ is the continuation value. Equation (D.28) posits that the value of keeping a previous employment relationship equals the product between the probability a previous match keeps being active in the current period, Θ_t, and the difference between the value of production, $P_{P,t}\gamma$, and the wage bill, W_t. In addition, there is the continuation value discounted by the time discount factor of the household, β, and the probability of not exiting the economy at the end of the period, $1-\delta$. In setting Equation (D.28), we make two assumptions: (i) upon the separation of a previous job match, the firm turns into an idle producer and can operate only in the next period; (ii) the wage rate is the one determined by the Nash bargaining problem of the producers that have opened a vacancy and the matched workers in each period.

The value of a labor firm that substitute its worker with a machine, $V_{M,t}(\gamma)$, is:

$$V_{M,t}(\gamma) = P_{P,t} - P_{M,t} + \beta(1-\delta) V_{C,t+1}(\gamma), \tag{D.29}$$

which consists in the difference between the value of production, $P_{P,t}$, and the price of purchasing a robot, $P_{M,t}$, and then adds the discounted continuation value. Note that the continuation value is identical independently of whether the labor firm decides for either of the two options. We show why this is the case below.

Thus, a labor firm decides to keep its previous job match if $V_{L,t}(\gamma) > V_{M,t}(\gamma)$. Since the value of being a labor firm increases with the labor efficiency level, γ, there exists a cut-off point for the labor efficiency, $\gamma^{\star}_{L,t}$, such that

$$V_{L,t}\left(\gamma^{\star}_{L,t}\right) = V_{M,t}\left(\gamma^{\star}_{L,t}\right). \tag{D.30}$$

The two options of idle firms are to either post a vacancy and look for a job, or install a machine. In the former case, the value, $V_{I,t}(\gamma)$, is:

$$V_{I,t}(\gamma) = q(\theta_t)\left[P_{P,t}\gamma - W_t\right] + \beta(1-\delta) V_{C,t+1}(\gamma), \tag{D.31}$$

which multiply the difference between the production value and the wage bill by the probability of fill the vacancy, $q(\theta_t)$. Instead, the value of installing a robot, $V_{M,t}(\gamma)$, is the one described in Equation (D.29). Thus, idle firms post a vacancy if and only if $V_{I,t}(\gamma) > V_{M,t}(\gamma)$, which implies the automation cut-off, $\gamma^{\star}_{I,t}$, such that

$$V_{I,t}\left(\gamma^{\star}_{I,t}\right) = V_{M,t}\left(\gamma^{\star}_{I,t}\right). \tag{D.32}$$

Finally, machine firms face the same two options of idle producers, and thus face the same automation threshold, $\gamma^{\star}_{I,t}$.

We then proceed into a crucial assumption that allows us to solve for the equilibrium of the model without tracking how the distribution of producers evolves across the three types over time. We do so by positing that the probability to fill a job equals to the probability that a previous job match keeps being operative in the current period, that is

$$q(\theta_t) = \Theta_t. \tag{D.33}$$

Importantly, this assumption is very much in line with the data, and thus we can consider more a feature of the data rather than a parameter restriction.[8] This condition is crucial because it makes the current-period value of idle firms that post a vacancy to be equal to that of labor firms, that is

$$q(\theta_t)\left[P_{P,t}\gamma - W_t\right] = \Theta_t\left[P_{P,t}\gamma - W_t\right]. \tag{D.34}$$

As a result, although the three incumbents face two options with different values, these values are common across the three types. This implies that (i) the continuation value is exactly the same for the three types, and (ii) all incumbents face the same automation threshold, $\gamma^{\star}_t = \gamma^{\star}_{L,t} = \gamma^{\star}_{I,t}$.

In addition to the incumbents, every period, there is a measure $\Xi_{e,t}$ of producers that pay a fixed nominal monetary cost κ_t and enter the market. As in the baseline model, entrant producers decide to operate their task by either employing workers or installing a machine. Then, the value of a labor entrant producer, $V_{E,L,t}(\gamma)$ is:

$$V_{E,L,t}(\gamma) = q_t(\theta_t)\left[P_{P,t}\gamma - W_t\right] - \kappa_t + \beta(1-\delta)V_{C,t+1}(\gamma), \tag{D.35}$$

where the continuation value, $V_{C,t+1}(\gamma)$, is the same than for incumbents, as this entrant producer can either become a labor incumbent producer upon a successful match, or an idle incumbent producer, if it does not manage to fill the vacancy. Instead, the value of a machine

[8] The job filling probability estimated by Davis et al. (2013) and Mongey and Violante (2025) implied an average job tenure of about 5.5 years, just one year above the average observed in the data over the decade of 2010.

entrant producer, $V_{E,M,t}$, is:

$$V_{E,M,t}(\gamma) = P_{P,t} - P_{M,t} - \kappa_t + \beta(1-\delta)V_{C,t+1}(\gamma). \tag{D.36}$$

An entrant decides to post a vacancy if the value to do so is larger than the value of installing a machine, that is, if $V_{E,L,t}(\gamma) > V_{E,M,t}(\gamma)$. Since also in this case the value of being an entrant producer increases with the labor efficiency level, γ, there exists a cut-off point for the labor efficiency, $\gamma^{\star}_{E,L,t}$, such that

$$V_{E,L,t}\left(\gamma^{\star}_{E,L,t}\right) = V_{E,M,t}\left(\gamma^{\star}_{E,L,t}\right). \tag{D.37}$$

Importantly, since the two values faced by the entrant are identical to the two values faced by the incumbents, with the only difference that in both of the values of the entrant we need to deduce the entry cost κ_t, it follows that the automation cut-off of entrant producers is exactly the same as for incumbents, that is, $\gamma^{\star}_t = \gamma^{\star}_{E,L,t}$.

To characterize the law of motion of the measures of each type of incumbent, let us first determine the measure of entrants, $\Xi_{E,t}$. This measure is pinned down by the free entry condition, which posits that producers enter the economy until the ex-ante value of operating in the economy net of the entry cost, $V_{E,t}$, equals to zero. Since entrants have to decide whether to enter the economy before drawing their labor efficiency level, the free entry condition reads:

$$V_{E,t} = \int_1^{\gamma^{\star}_t} V_{E,M,t}(\gamma)\, f(\gamma;\alpha)\,\mathrm{d}\gamma + \int_{\gamma^{\star}_t}^{\gamma_H} V_{E,L,t}(\gamma)\, f(\gamma;\alpha)\,\mathrm{d}\gamma = 0 \tag{D.38}$$

The measure of labor entrants integrates over all entrants that have sufficiently high labor efficiency, that is, a labor efficiency above the automation cut-off:

$$\Xi_{E,L,t} = \Xi_{E,t} \int_{\gamma^{\star}_t}^{\gamma_H} f(\gamma;\alpha)\,\mathrm{d}\gamma. \tag{D.39}$$

Instead, the measure of machine entrants integrates over the entrants whose labor efficiency is below the automation threshold:

$$\Xi_{E,M,t} = \Xi_{E,t} \int_1^{\gamma^{\star}_t} f(\gamma;\alpha)\,\mathrm{d}\gamma. \tag{D.40}$$

To then characterize the evolution of the measure of incumbent, we consider two different cases: one in which the current automation cut-off is equal to or larger than the previous one, $\gamma_t^\star \geq \gamma_{t-1}^\star$, and one in which he current automation cut-off is lower than in the previous period, $\gamma_t^\star < \gamma_{t-1}^\star$.

Let us start with the case in which $\gamma_t^\star < \gamma_{t-1}^\star$, so that labor incumbents decide not to substitute their workers with machines. Since the current automation threshold is lower than in the previous period, it follows that labor incumbents have a labor efficiency above the current automation threshold, and thus they do not install a machine. In this case, the measure of labor incumbents evolves as:

$$\Xi_{L,t} = \Theta_t \left(1-\delta\right) \Xi_{L,t-1} + q\left(\theta_t\right) \left\{ \left(1-\delta\right) \left[\Xi_{I,t-1} + \Xi_{M,t-1} \frac{\int_{\gamma_t^\star}^{\gamma_{t-1}^\star} f\left(\gamma;\alpha\right) \mathrm{d}\gamma}{\int_1^{\gamma_{t-1}^\star} f\left(\gamma;\alpha\right) \mathrm{d}\gamma} \right] + \Xi_{E,L,t} \right\}. \quad \text{(D.41)}$$

Thus, the measure of labor incumbents consists of four components: (i) the measure of previous labor incumbents that have neither exited the economy not experienced an exogenous job separation, (ii) the measure of previous idle firms that did not exit, have posted a vacancy in the current period, and successfully filled it; (iii) the measure of previous machine firms that did not exit, have a labor efficiency between the current and previous cut-offs, have posted a vacancy, and successfully filled it, and (iv) the measure of labor entrants that have matched with a worker.

Then, the measure of idle incumbents evolves as:

$$\Xi_{I,t} = \left(1-\delta\right) \left\{ \left[1 - q\left(\theta_t\right)\right] \Xi_{I,t-1} + \left(1-\Theta_{t-1}\right) \Xi_{L,t-1} \right\} + \Xi_{E,L,t}, \quad \text{(D.42)}$$

which sums over (i) the measure of previous idle firms that have not exited in the previous period and have posted a vacancy in the current period without being able to fill it, (ii) the measure of previous labor firms which have experienced the exogenous job separation and have unsuccessfully posted a vacancy, and (iii) the measure of labor entrants that have not filled their vacancy.

Finally, the law of motion of the measure of machine incumbents is:

$$\Xi_{M,t} = (1-\delta)\,\Xi_{M,t-1}\frac{\int_{1}^{\gamma_t^\star} f(\gamma;\alpha)\,\mathrm{d}\gamma}{\int_{1}^{\gamma_{t-1}^\star} f(\gamma;\alpha)\,\mathrm{d}\gamma} + \Xi_{E,M,t}, \tag{D.43}$$

which sums of over the measure of previous machine firms that have not exited and have a labor efficiency below the automation cut-off, and the measure of machine entrants.

Instead, in case in which $\gamma_t^\star \geq \gamma_{t-1}^\star$, all previous machine firms decide to keep operating with robots, whereas there is a measure of previous labor firms with a labor efficiency that was above the previous automation cut-off, but now is below the current one, that decide to endogenously separate from its workers and install a machine.

Thus, the law of motion of labor incumbents becomes:

$$\Xi_{L,t} = \left\{ (1-\delta)\left[\Theta_t\Xi_{L,t-1} + q(\theta_t)\,\Xi_{I,t-1}\right]\frac{\int_{\gamma_t^\star}^{\gamma_H} f(\gamma;\alpha)\,\mathrm{d}\gamma}{\int_{\gamma_{t-1}^\star}^{\gamma_H} f(\gamma;\alpha)\,\mathrm{d}\gamma}\right\} + q(\theta_t)\,\Xi_{E,L,t} \tag{D.44}$$

in which the first component is the measure of previous labor firms that have neither exited the economy nor exogenously separated from their workers and have a labor efficiency above the new current automation cut-off, the second component is the measure of non-exited idle firms with sufficiently high labor efficiency and have successfully matched with a worker, and the third component is the measure of labor entrants that fill their job vacancies.

Then, the measure of idle incumbents evolves over time as:

$$\begin{aligned}\Xi_{I,t} = \Bigg\{ (1-\delta)\Big[(1-\Theta_t)\,\Xi_{L,t-1} + [1-q(\theta_t)]\,\Xi_{I,t-1}\Big] \times \dots \\ \dots \times \frac{\int_{\gamma_t^\star}^{\gamma_H} f(\gamma;\alpha)\,\mathrm{d}\gamma}{\int_{\gamma_{t-1}^\star}^{\gamma_H} f(\gamma;\alpha)\,\mathrm{d}\gamma}\Bigg\} + [1-q(\theta_t)]\,\Xi_{E,L,t}\end{aligned} \tag{D.45}$$

which posits that the measure of current idle incumbents are the sum of the previous labor firms with sufficiently high labor efficiency, that did not exit, and experienced the exogenous job separation, the previous idle firms with sufficiently high labor efficiency, that did not exit, posted a vacancy in the current period but did not fill it, and the measure of labor entrants that did not filled the job vacancy.

We can characterize the law of motion of the measure of machine incumbents as:

$$\Xi_{M,t} = (1-\delta)\left\{\Xi_{M,t-1} + \left[(\Xi_{I,t-1} + \Xi_{L,t-1})\frac{\int_{\gamma_{t-1}^{\star}}^{\gamma_t^{\star}} f(\gamma;\alpha)\,\mathrm{d}\gamma}{\int_{\gamma_{t-1}^{\star}}^{\gamma_H} f(\gamma;\alpha)\,\mathrm{d}\gamma}\right]\right\} + \Xi_{E,M,t}, \tag{D.46}$$

which sums of over (i) the measure of previous machine firms that have not exited, (ii) the measure of non-exited previous idle firms with labor efficiency below the current automation cut-off, even though this was not the case in the previous period, (iii) the measure of non-exited previous labor firms whose labor efficiency is currently below the automation cut-off, even though this was not the case in the previous period, and (iv) the measure of machine entrants.

In either case, given that the total measure of producers in a given period is Ξ_t, its law of motion equals to

$$\Xi_t = (1-\delta)\,\Xi_{t-1} + \Xi_{E,t}, \tag{D.47}$$

that is, the current measure of producers equals the measure of producers in the previous period that have not exited the economy, plus the new measure of entrants.

We can now define the amount of vacancies, v_t, that are opened in every period, which equals to:

$$\begin{aligned} v_t = (1-\delta)\Xi_{I,t-1}\min\left\{1, \frac{\int_{\gamma_t^{\star}}^{\gamma_H} f(\gamma;\alpha)\ \mathrm{d}\gamma}{\int_{\gamma_{t-1}^{\star}}^{\gamma_H} f(\gamma;\alpha)\ \mathrm{d}\gamma}\right\} + \dots \\ \dots + (1-\delta)\,\Xi_{M,t-1}\max\left\{0, \frac{\int_{\gamma_t^{\star}}^{\gamma_{t-1}^{\star}} f(\gamma;\alpha)\ \mathrm{d}\gamma}{\int_{1}^{\gamma_{t-1}^{\star}} f(\gamma;\alpha)\ \mathrm{d}\gamma}\right\} + \Xi_{E,L,t}, \end{aligned} \tag{D.48}$$

which shows that we need to keep track of whether the automation cut-off increases or decreases. In other words, in this model extension the structure of the labor market is non-convex, and the number of workers displaced when existing labor firms with matched workers decide to automate.

After determining the number of vacancies, we proceed into computing the measure of individuals searching for a job, s_t. As we have mentioned above, the total pool of individuals searching for a job consists of previously unemployed people, U_{t-1}, the new fraction of individuals that decide to participate in the labor market, $N_t - N_{t-1}$, the individuals that have been exogenously laid off in the previous period, $\delta\Xi_{L,t-1}$, and those that have been endogenously displaced by machines in the current period, D_t, so that:

$$s_t = U_{t-1} + (N_t - N_{t-1}) + \delta\Xi_{L,t-1} + D_t \tag{D.49}$$

where the fraction of displaced workers, D_t, equals to:

$$D_t = (1-\delta)\Xi_{L,t-1} \max\left\{0, \frac{\int_{\gamma_t^\star}^{\gamma_{t-1}^\star} f(\gamma;\alpha)\, \mathrm{d}\gamma}{\int_{\gamma_{t-1}^\star}^{\gamma_H} f(\gamma;\alpha)\, \mathrm{d}\gamma}\right\}. \tag{D.50}$$

Thus, in every period the total number of matches, $x(v_t, s_t)$, is pinned down by the same matching function of the baseline economy.

$$x(v_t, s_t) = \xi v_t^\eta s_t^{1-\eta}, \tag{D.51}$$

where η is the elasticity of the matching function with respect to vacancies, and xi denotes the matching efficiency. Given the matching function and the fact that the tightness of the labor market is the ratio between vacancies and searching individuals, $\theta_t = v_t/s_t$, we can then derive the probability that a person finds a job as

$$p(\theta_t) = \frac{x(v_t, s_t)}{s_t} = \xi\theta_t^\eta \tag{D.52}$$

and the probability that a firm fills a vacancy as

$$q(\theta_t) = \frac{x(v_t, s_t)}{v_t} = \xi\theta_t^{\eta-1}. \tag{D.53}$$

As in the baseline model, upon a match, workers and firms bargain on the real wage by splitting the total surplus of the match. However, the wage bargaining problem changes with respect the baseline model, and in this novel model version it is given by

$$\begin{aligned} &\arg\max_{w_t} \mathbb{E}_t\left\{\int_{\gamma_t^\star}^{\gamma_H}\left(q_{P,t}\gamma - w_t + \beta(1-\delta)\left[V_{L,t+1}(\gamma) - V_{I,t+1}(\gamma)\right]\right) f(\gamma;\alpha) d\gamma\right\}^{1-\tau} \times \dots \\ &\qquad \cdots \times \left[w_t + \beta(V_{W,t+1} - V_{U,t+1})\right]^\tau, \end{aligned} \tag{D.54}$$

where τ denotes workers' bargaining power. The first part of the argument of the maximization refers to firms' surplus, and consists of the current gains from employing a worker, which equals to the value of production minus the wage bill, $q_{P,t}\gamma - w_t$, and it adds the difference between discounted continuation values of being a labor firm and being an idle firm. These two components are then integrated over the measure of labor firms, that is, the producers whose labor efficiency is above the automation cut-off.

The second part of the argument of the maximization refers to the surplus of the workers, which equals to the wage rate they earn upon being hired, plus the discounted continuation value of being unemployment. The latter consists in the difference between the value of entering the following period as a worker matched to a firm, $V_{W,t}$, and the value of doing so unemployed, $V_{U,t}$. The value of being a worker attached to a firm after production takes place equals to

$$V_{W,t} = \delta V_{U,t} + (1-\delta)\left(\Theta_t\left(w_t + \beta V_{W,t+1}\right) + \beta\left(1-\Theta_t\right)V_{U,t+1}\right), \tag{D.55}$$

in which the first component indicates the probability that the worker is still attached to the same firm (i.e., the firm does not exit and does not experience the exogenous separation shock), the second component indicates that in case the firm exits, then the worker will enter into the unemployment pool in the next period, and the third component refers to the worker whose firm has not exited, but have experienced the exogenous job separation, in which case the worker enters the pool of unemployed with a period of delay.

Then, the value of being unemployment is

$$V_{U,t} = p(\theta_t)\left(w_t + \beta V_{W,t+1}\right) + \beta\left[1-p\left(\theta_t\right)\right]V_{U,t}, \tag{D.56}$$

in which the first component refers to the case in which, with the job finding probability $p(\theta_t)$, the unemployed individual is matched to a firm, and the second component indicates the case in which the match does not happen, so that the individual remains unemployed over the following period.

The Nash bargaining problem in Equation (D.54) implies that the optimal wage equals to

$$w_t = \tau q_{P,t} \frac{\int_{\gamma_t^\star}^{\gamma_H} \gamma f\left(\gamma;\alpha\right) \mathrm{d}\gamma}{\int_{\gamma_t^\star}^{\gamma_H} f\left(\gamma;\alpha\right) \mathrm{d}\gamma} + \beta\tau(V_{M,t+1} - V_{U,t+1}), \tag{D.57}$$

which extends the optimal wage of the baseline model in Equation (15) by adding the second term linked to workers' discounted continuation value, $\beta\tau(V_{W,t+1} - V_{U,t+1})$.

Finally, we describe labor participation, N_t, and the unemployment rate, u_t. The former is determined as

$$N_t = (1-\delta)\Xi_{L,t-1} + \left[1-(1-\delta)\Xi_{L,t-1}\right] \frac{w_t + \beta\left(V_{W,t} - V_{U,t}\right)}{\chi_H}. \tag{D.58}$$

The term $(1-\delta)\Xi_{L,t-1}$ refers to the mass of workers which were previously attached to a firm and still are, as the firm has not exited the economy. Accordingly, the complementary mass of individuals, $1-(1-\delta)\Xi_{L,t-1}$ is the one that can potentially decide to work. They do so only if the first order condition associated to a problem similar to that of Section 3.6 in the baseline model is satisfied. Consequently, the mass of individuals that decide to look for a job is: $\left[w_t + \beta\left(V_{W,t} - V_{U,t}\right)\right]/\chi_H$. Instead, the unemployment rate equals to

$$u_t = \frac{N_t - \Xi_{L,t}}{N_t}, \tag{D.59}$$

which uses the fact that the number of unemployed individuals equals the difference between the total amount of participating households, N_t, and those employed by firms, $\Xi_{L,t}$.

We then proceed in calibrating the model. We follow exactly the same parametrization of the baseline economy, as presented in Section 5.1. However, we have two differences. First, we set the exit rate to $\delta = 0.05$. This choice allows us to match an average tenure at a job of about 4 years. Second, we now consider the entry cost as given by $\kappa_t = \kappa + \varpi\left(\Xi_{E,t} - \bar{\Xi}_E\right)/\bar{\Xi}_E$, so that the cost increases when the measure of entrant firms rises substantially relative to steady state with a sensitivity captured by the parameter ϖ. We do so because otherwise the cyclicality of entry is way above what we observe in the data. In other words, a constant entry cost would imply a counterfactual high degree of pro-cyclicality of entry. Indeed, as indicated by Campbell (1998), Lee and Mukoyama (2015), and Clementi and Palazzo (2016),

Table D.22: Transitional Dynamics Exercise - Long-Lasting Job Matches

	Dependent Variable: π_t		
	Baseline Model	Long-Lasting Job Matches Economy	Data MSA-level
	(1)	(2)	(3)
u_t	-0.5036	-0.1473	-0.5248
	(0.0005)	(0.0031)	(0.1469)
$u_t \times (m_t - \bar{m})$	0.0058	0.0042	0.0058
	(0.0003)	(0.0005)	(0.0028)

Note: The table reports the estimates of the regression in Equation (30) based on simulated data of the baseline model (in Column 1) and the alternative economy with long-lasting job matches (in Column 2). Column 3 reports the estimate based on a panel regression across U.S. MSAs which correspond to the case with non-tradable inflation as dependent variable and IV method of Table 1. The simulated data of the baseline model are based on a transitional dynamics exercise between a low-automation steady state, in which the robot-to-employee ratio is 0.2%, and a high-automation steady state, in which the robot-to-employee ratio is 0.8%.

while entry is procyclical, its correlation with output is somewhat limited, as it equals to 0.4. We then calibrate the sensitivity parameter ϖ to match this moment.

We then replicate the baseline transition dynamics exercise of Section5.4: we simulate the model, recover the values of inflation, unemployment, and automation, and estimate regression (30). The results in Table D.22 indicate that, given our baseline calibration of the Rotemberg cost, the Phillips Curve is much flatter than in our baseline economy: the average inflation sensitivity to unemployment is roughly 30% of that of the baseline model. On the other hand, while the magnitude of the effect of automation on inflation dynamics is lower than in the baseline, the difference this time is much more muted: the new model implies an effect of automation which is 70% that of the baseline economy. Thus, again we have found that altering the model specification does not modify its quantitative implications on the empirically relevant way in which robot installations reduces the inflation sensitivity to labor market conditions.

References

Acemoglu, Daron, and Pascual Restrepo. (2018). "The race between man and machine: Implications of technology for growth, factor shares, and employment". *American Economic Review*, 108(6), pp. 1488-1542. https://doi.org/10.2139/ssrn.2781320

Acemoglu, Daron, and Pascual Restrepo. (2020). "Robots and jobs: Evidence from US labor markets". *Journal of Political Economy*, 128(6), pp. 2188-2244. https://doi.org/10.2139/ssrn.2940245

Acemoglu, Daron, and Pascual Restrepo. (2022). "Demographics and automation". *Review of Economic Studies*, 89(1), pp. 1-44. https://doi.org/10.1093/restud/rdab031

Adao, Rodrigo, Michal Kolesár and Eduardo Morales. (2019). "Shift-share designs: Theory and inference". *The Quarterly Journal of Economics*, 134(4), pp. 1949-2010. https://doi.org/10.1093/qje/qjz025

Aksoy, Yunus, Henrique S. Basso, Ron Smith and Tobias Grasl. (2019). "Demographic structure and macroeconomic trends". *American Economic Journal: Macroeconomics*, 11(1), pp. 193-222. https://doi.org/10.1257/mac.20170114

Amiti, Mary, Sebastian Heise, Fatih Karahan and Ayşegül Şahin. (2023). *Inflation Strikes Back: The Role of Import Competition and the Labor Market*. University of Chicago Press. https://doi.org/10.1086/729195

Autor, David, David Dorn and Gordon Hanson. (2013). "The China syndrome: Local labor market effects of import competition in the United States". *American Economic Review*, 103(6), pp. 2121-2168. https://doi.org/10.1257/aer.103.6.2121

Basso, Henrique S., and Juan F. Jimeno. (2021). "From secular stagnation to robocalypse? Implications of demographic and technological changes". *Journal of Monetary Economics*, 117, pp. 833-847. https://doi.org/10.1016/j.jmoneco.2020.06.004

Basso, Henrique S., and Omar Rachedi. (2021). "The young, the old, and the government: Demographics and fiscal multipliers". *American Economic Journal: Macroeconomics*, 13(4), pp. 110-141. https://doi.org/10.1257/mac.20190174

Basu, Susanto, John G. Fernald and Miles S. Kimball. (2006). "Are technology improvements contractionary?". *American Economic Review*, 96(5), pp. 1418-1448. https://doi.org/10.1257/aer.96.5.1418

Borusyak, Kirill, Peter Hull and Xavier Jaravel. (2022). "Quasi-experimental shift-share research designs". *Review of Economic Studies*, 89(1), pp. 181-213. https://doi.org/10.1093/restud/rdab030

Campbell, Jeffrey R. (1998). "Entry, exit, embodied technology, and business cycles". *Review of Economic Dynamics*, 1(2), pp. 371-408. https://doi.org/10.1006/redy.1998.0009

Christiano, Lawrence J., Martin Eichenbaum and Charles L. Evans. (2005). "Nominal rigidities and the dynamic effects of a shock to monetary policy". *Journal of Political Economy*, 113(1), pp. 1-45. https://doi.org/10.1086/426038

Clementi, Gian Luca, and Berardino Palazzo. (2016). "Entry, exit, firm dynamics, and aggregate fluctuations". *American Economic Journal: Macroeconomics*, 8(3), pp. 1-41. https://doi.org/10.1257/mac.20150017

Davis, Steven J., R. Jason Faberman and John C. Haltiwanger. (2013). "The establishmentlevel behavior of vacancies and hiring". *The Quarterly Journal of Economics*, 128(2), pp. 581-622. https://doi.org/10.1093/qje/qjt002

Driscoll, John C., and Aart C. Kraay. (1998). "Consistent covariance matrix estimation with spatially dependent panel data". *Review of Economics and Statistics*, 80(4), pp. 549-560. https://doi.org/10.1162/003465398557825

Fernald, John. (2014). "A quarterly, utilization-adjusted series on total factor productivity". Mimeo.

Forbes, Kristin. (2019). "Inflation dynamics: Dead, dormant, or determined abroad?". *Brookings Papers on Economic Activity*, 2019(2), pp. 257-338. https://doi.org/10.1353/eca.2019.0015

Galesi, Alessandro, and Omar Rachedi. (2019). "Services deepening and the transmission of monetary policy". *Journal of the European Economic Association*, 17(4), pp. 1261-1293. https://doi.org/10.1093/jeea/jvy041

Galí, Jordi, J. David López-Salido and Javier Vallés. (2007). "Understanding the effects of government spending on consumption". *Journal of the European Economic Association*, 5(1), pp. 227-270. https://doi.org/10.1162/jeea.2007.5.1.227

Hazell, Jonathon, Juan Herreño, Emi Nakamura and Jón Steinsson. (2022). "The slope of the Phillips curve: Evidence from US states". *The Quarterly Journal of Economics*, 137(3), pp. 1299-1344. https://doi.org/10.1093/qje/qjac010

Heise, Sebastian, Fatih Karahan and Ayşegül Şahin. (2022). "The missing inflation puzzle: The role of the wage-price pass-through". *Journal of Money, Credit and Banking*, 54(S1), pp. 7-51. https://doi.org/10.1111/jmcb.12896

Leduc, Sylvain, and Zheng Liu. (2024). "Automation, bargaining power, and labor market fluctuations". *American Economic Journal: Macroeconomics*, 16(4), pp. 311-349. https://doi.org/10.1257/mac.20220181

Lee, Yoonsoo, and Toshihiko Mukoyama. (2015). "Entry and exit of manufacturing plants over the business cycle". *European Economic Review*, 77, pp. 20-27. https://doi.org/10.1016/j.euroecorev.2015.03.011

McLeay, Michael, and Silvana Tenreyro. (2020). "Optimal inflation and the identification of the Phillips curve". *NBER Macroeconomics Annual*, 34(1), pp. 199-255. https://doi.org/10.1086/707181

Mongey, Simon, and Giovanni L. Violante. (2025). "Macro recruiting intensity from micro data". *American Economic Journal: Macroeconomics*. Forthcoming. https://doi.org/10.3386/w26231

Nakamura, Emi, and Jón Steinsson. (2014). "Fiscal stimulus in a monetary union: Evidence from US regions". *American Economic Review*, 104(3), pp. 753-792. https://doi.org/10.1257/aer.104.3.753

Norkutė, Milda, Vasilis Sarafidis, Takashi Yamagata and Guowei Cui. (2021). "Instrumental variable estimation of dynamic linear panel data models with defactored regressors and a multifactor error structure". *Journal of Econometrics*, 220(2), pp. 416-446. https://doi.org/10.1016/j.jeconom.2020.04.008

Patterson, Christina. (2023). "The matching multiplier and the amplification of recessions". *American Economic Review*, 113(4), pp. 982-1012. https://doi.org/10.1257/aer.20210254

Ramey, Valerie A. (2021). "The macroeconomic consequences of infrastructure investment". In Edward L. Glaeser and James M. Poterva (eds.), *Economic Analysis and Infrastructure Investment.* University of Chicago Press, p. 219. https://doi.org/10.7208/chicago/9780226800615.003.0005

Rubbo, Elisa. (2023). "Networks, Phillips curves, and monetary policy". *Econometrica,* 91(4), pp. 1417-1455. https://doi.org/10.3982/ecta18654

Siena, Daniele, and Riccardo Zago. (2024). "Job polarization, labor market fluidity and the flattening of the Phillips curve". *Economic Journal,* 134(661), pp. 2141-2174. https://doi.org/10.1093/ej/ueae006

Varian, Hal R. (1975). "A Bayesian approach to real estate assessment". *Studies in Bayesian Econometric and Statistics in Honor of Leonard J. Savage,* pp. 195-208.

Weidner, Justin, Greg Kaplan and Gianluca Violante. (2014). "The wealthy hand-to-mouth". *Brookings Papers on Economic Activity,* 1, pp. 77-153. https://doi.org/10.1353/eca.2014.0002

BANCO DE ESPAÑA PUBLICATIONS

WORKING PAPERS

2420 FRANCISCO GONZÁLEZ, JOSÉ E. GUTIÉRREZ and JOSÉ MARÍA SERENA: Shadow seniority? Lending relationships and borrowers' selective default.

2421 ROBERTO BLANCO, MIGUEL GARCÍA-POSADA, SERGIO MAYORDOMO and MARÍA RODRÍGUEZ-MORENO: Access to credit and firm survival during a crisis: the case of zero-bank-debt firms.

2422 FERNANDO CEREZO, PABLO GIRÓN, MARÍA T. GONZÁLEZ-PÉREZ and ROBERTO PASCUAL: The impact of sovereign debt purchase programmes. A case study: the Spanish-to-Portuguese bond yield spread.

2423 EDGAR SILGADO-GÓMEZ: Sovereign uncertainty.

2424 CLODOMIRO FERREIRA, JULIO GÁLVEZ and MYROSLAV PIDKUYKO: Housing tenure, consumption and household debt: life-cycle dynamics during a housing bust in Spain.

2425 RUBÉN DOMÍNGUEZ-DÍAZ and SAMUEL HURTADO: Green energy transition and vulnerability to external shocks.

2426 JOSEP GISBERT and JOSÉ E. GUTIÉRREZ: Bridging the gap? Fintech and financial inclusion.

2427 RODOLFO G. CAMPOS, MARIO LARCH, JACOPO TIMINI, ELENA VIDAL and YOTO V. YOTOV: Does the WTO Promote Trade? A Meta-analysis.

2428 SONER BASKAYA, JOSÉ E. GUTIÉRREZ, JOSÉ MARÍA SERENA and SERAFEIM TSOUKAS: Bank supervision and non-performing loan cleansing.

2429 TODD E. CLARK, GERGELY GANICS, and ELMAR MERTENS: Constructing fan charts from the ragged edge of SPF forecasts.

2430 MIGUEL GARCÍA-POSADA and PETER PAZ: The transmission of monetary policy to credit supply in the euro area.

2431 KLODIANA ISTREFI, FLORENS ODENDAHL and GIULIA SESTIERI: ECB communication and its impact on financial markets.

2432 FRUCTUOSO BORRALLO, LUCÍA CUADRO-SÁEZ, CORINNA GHIRELLI and JAVIER J. PÉREZ: “El Niño” and “La Niña”: Revisiting the impact on food commodity prices and euro area consumer prices.

2433 VÍCTOR CABALLERO, CORINNA GHIRELLI, ÁNGEL LUIS GÓMEZ and JAVIER J. PÉREZ: The public-private wage GAP in the euro area a decade after the sovereign debt crisis.

2434 LIDIA CRUCES, ISABEL MICÓ-MILLÁN and SUSANA PÁRRAGA: Female financial portfolio choices and marital property regimes.

2435 RODOLFO G. CAMPOS, ANA-SIMONA MANU, LUIS MOLINA and MARTA SUÁREZ-VARELA: China's financial spillovers to emerging markets.

2436 LUDOVIC PANON, LAURA LEBASTARD, MICHELE MANCINI, ALESSANDRO BORIN, PEONARE CAKA, GIANMARCO CARIOLA, DENNIS ESSERS, ELENA GENTILI, ANDREA LINARELLO, TULLIA PADELLINI, FRANCISCO REQUENA and JACOPO TIMINI: Inputs in Distress: Geoeconomic Fragmentation and Firms' Sourcing.

2437 DANIEL DEJUAN-BITRIA, WAYNE R. LANDSMAN, SERGIO MAYORDOMO and IRENE ROIBÁS: How do changes in financial reporting standards affect relationship lending?

2438 ALICIA AGUILAR and RICARDO GIMENO: Discrete Probability Forecasts: What to expect when you are expecting a monetary policy decision.

2439 RODOLFO G. CAMPOS, JESÚS FERNÁNDEZ-VILLAVERDE, GALO NUÑO and PETER PAZ: Navigating by Falling Stars: Monetary Policy with Fiscally Driven Natural Rates.

2440 ALEJANDRO CASADO and DAVID MARTÍNEZ-MIERA: Local lending specialization and monetary policy.

2441 JORGE ABAD, DAVID MARTÍNEZ-MIERA and JAVIER SUÁREZ: A macroeconomic model of banks' systemic risk taking.

2442 JOSEP PIJOAN-MAS and PAU ROLDAN-BLANCO: Dual labor markets and the equilibrium distribution of firms.

2443 OLYMPIA BOVER, LAURA HOSPIDO and ANA LAMO: Gender and Career Progression: Evidence from the Banco de España.

2444 JESÚS FERNÁNDEZ-VILLAVERDE, GALO NUÑO and JESSE PERLA: Taming the curse of dimensionality: quantitative economics with deep learning.

2445 CLODOMIRO FERREIRA and STEFANO PICA: Households' subjective expectations: disagreement, common drivers and reaction to monetary policy.

2446 ISABEL MICÓ-MILLÁN: Inheritance Tax Avoidance Through the Family Firm.

2447 MIKEL BEDAYO, EVA VALDEOLIVAS and CARLOS PÉREZ: The stabilizing role of local claims in local currency on the variation of foreign claims.

2501 HENRIQUE S. BASSO, MYROSLAV PIDKUYKO and OMAR RACHEDI: Opening the black box: aggregate implications of public investment heterogeneity.

2502 MARCO BARDOSCIA, ADRIAN CARRO, MARC HINTERSCHWEIGER, MAURO NAPOLETANO, LILIT POPOYAN, ANDREA ROVENTINI and ARZU ULUC: The impact of prudential regulations on the UK housing market and economy: insights from an agent-based model.

2503 IRINA BALTEANU, KATJA SCHMIDT and FRANCESCA VIANI: Sourcing all the eggs from one basket: trade dependencies and import prices.

2504 RUBÉN VEIGA DUARTE, SAMUEL HURTADO, PABLO A. AGUILAR GARCÍA, JAVIER QUINTANA GONZÁLEZ and CAROLINA MENÉNDEZ ÁLVAREZ: CATALIST: A new, bigger, better model for evaluating climate change transition risks at Banco de España.

2505 PILAR GARCÍA and DIEGO TORRES: Perceiving central bank communications through press coverage.

2506 MAR DELGADO-TÉLLEZ, JAVIER QUINTANA and DANIEL SANTABÁRBARA: Carbon pricing, border adjustment and renewable energy investment: a network approach.

2507 MARTA GARCÍA RODRÍGUEZ: The role of wage expectations in the labor market.

2508 REBECA ANGUREN, GABRIEL JIMÉNEZ and JOSÉ-LUIS PEYDRÓ: Bank capital requirements and risk-taking: evidence from Basel III.

2509 JORGE E. GALÁN: Macroprudential policy and the tail risk of credit growth.

2510 PETER KARADI, ANTON NAKOV, GALO NUÑO, ERNESTO PASTÉN and DOMINIK THALER: Strike while the Iron is Hot: Optimal Monetary Policy with a Nonlinear Phillips Curve.

2511 MATTEO MOGLIANI and FLORENS ODENDAHL: Density forecast transformations.

2512 LUCÍA LÓPEZ, FLORENS ODENDAHL, SUSANA PÁRRAGA and EDGAR SILGADO-GÓMEZ: The pass-through to inflation of gas price shocks.

2513 CARMEN BROTO and OLIVIER HUBERT: Desertification in Spain: Is there any impact on credit to firms?

2514 ANDRÉS ALONSO-ROBISCO, JOSÉ MANUEL CARBÓ, PEDRO JESÚS CUADROS-SOLAS and JARA QUINTANERO: The effects of open banking on fintech providers: evidence using microdata from Spain.

2515 RODOLFO G. CAMPOS and JACOPO TIMINI: Trade bloc enlargement when many countries join at once.

2516 CORINNA GHIRELLI, JAVIER J. PÉREZ and DANIEL SANTABÁRBARA: Inflation and growth forecast errors and the sacrifice ratio of monetary policy in the euro area.

2517 KOSUKE AOKI, ENRIC MARTORELL and KALIN NIKOLOV: Monetary policy, bank leverage and systemic risk-taking.

2518 RICARDO BARAHONA: Index fund flows and fund distribution channels.

2519 ALVARO FERNÁNDEZ-GALLARDO, SIMON LLOYD and ED MANUEL: The Transmission of Macroprudential Policy in the Tails: Evidence from a Narrative Approach.

2520 ALICIA AGUILAR: Beyond fragmentation: unraveling the drivers of yield divergence in the euro area.

2521 RUBÉN DOMÍNGUEZ-DÍAZ and DONGHAI ZHANG: The macroeconomic effects of unemployment insurance extensions: A policy rule-based identification approach.

2522 IRMA ALONSO-ALVAREZ, MARINA DIAKONOVA and JAVIER J. PÉREZ: Rethinking GPR: The sources of geopolitical risk.

2523 ALBERTO MARTÍN, SERGIO MAYORDOMO and VICTORIA VANASCO: Banks vs. Firms: Who Benefits from Credit Guarantees?

2524 SUMIT AGARWAL, SERGIO MAYORDOMO, MARÍA RODRÍGUEZ-MORENO and EMANUELE TARANTINO: Household Heterogeneity and the Lending Channel of Monetary Policy.

2525 DIEGO BONELLI, BERARDINO PALAZZO, and RAM YAMARTHY: Good inflation, bad inflation: implications for risky asset prices.

2526 STÉPHANE BONHOMME and ANGELA DENIS: Fixed Effects and Beyond. Bias Reduction, Groups, Shrinkage and Factors in Panel Data.

2527 ÁLVARO FERNÁNDEZ-GALLARDO and IVÁN PAYÁ: Public debt burden and crisis severity.

2528 GALO NUÑO: Three Theories of Natural Rate Dynamics.

2529 GALO NUÑO, PHILIPP RENNER and SIMON SCHEIDEGGER: Monetary policy with persistent supply shocks.

2530 MIGUEL ACOSTA-HENAO, MARÍA ALEJANDRA AMADO, MONTSERRAT MARTÍ and DAVID PÉREZ-REYNA: Heterogeneous UIPDs across Firms: Spillovers from U.S. Monetary Policy Shocks.

2531 LUIS HERRERA and JESÚS VÁZQUEZ: Learning from news.

2532 MORTEZA GHOMI, JOCHEN MANKART, RIGAS OIKONOMOU and ROMANOS PRIFTIS: Debt maturity and government spending multipliers.

2533 MARINA DIAKONOVA, CORINNA GHIRELLI and JAVIER J. PÉREZ: Political polarization in Europe.

2534 NICOLÁS FORTEZA and SERGIO PUENTE: Measuring non-workers' labor market attachment with machine learning.

2535 GERGELY GANICS and LLUC PUIG CODINA: Simple Tests for the Correct Specification of Conditional Predictive Densities.

2536 HENRIQUE S. BASSO and OMAR RACHEDI: Robot adoption and inflation dynamics.