Robot Adoption and Inflation Dynamics*

Henrique S. Basso
Banco de España

Omar Rachedi
ESADE Business School
Universitat Ramon Llull

March 7, 2023

Abstract

We leverage variation in robot adoption across U.S. metropolitan areas to document that automation reduces the sensitivity of inflation to unemployment. A New Keynesian model with search frictions and automation rationalizes our empirical findings through two mechanisms. First, automation shrinks workers' bargaining power, dampening the sensitivity of wages to unemployment. Second, automation reduces the labor share, decoupling output changes from unemployment variation. Both channels flatten the price Phillips curve. However, when boosting automation is costly, the threat of robot adoption is no longer effective in curtailing workers' bargaining power amidst large expansionary shocks, leading to a steeper Phillips curve.

Key Words: Automation, robots, inflation, Phillips curve, unemployment.

JEL Classification Codes: E24, E31, J31, O33.

*Addresses: henrique.basso@bde.es, omar.rachedi@esade.edu. We thank presentation participants at the Bank of England, the University of Mannheim, and the LACEA LAMES 2022 Annual Meeting for useful comments and suggestions. The views expressed in this paper are those of the authors and do not necessarily represent the views of the Banco de España or the Eurosystem.

1 Introduction

Over the past few decades, advanced economies have witnessed a substantial increase in the use of robots and other forms of automation in production processes. This phenomenon has generated comprehensive implications on the labor market, contributing to the polarization of employment opportunities and the decline of middle-skilled jobs, and compressing wages at the lower end of the earnings distribution (Acemoglu and Restrepo, 2018, 2020a, 2020b, 2022; Graetz and Michaels, 2018; Acemoglu et al., 2020). However, notwithstanding the key role that labor market conditions have on wage and price setting, little is known about how robot adoption may influence inflation dynamics. In this paper, we show empirically, theoretically, and quantitatively that the surge in automation could explain the muted sensitivity of inflation to unemployment observed in advanced economies until the Covid pandemic.

We start by providing novel empirical evidence showing that robot adoption alters both price inflation and wage inflation. To do so, we build a panel of non-tradable goods inflation, wage inflation, unemployment rate and robot adoption at the U.S. metropolitan area (MSA) level. To measure automation, we follow Acemoglu and Restrepo (2020a) and combine the robot installation for each industry at the U.S. national level with the employment share of each industry at the MSA level. In this way, we measure the robot installed per employee for each metropolitan area. We end up with a panel across 384 MSAs at the annual frequency from 2008 and 2018. While 2008 is the first year for which the U.S. Bureau of Economic Analysis provides price information across MSAs, our sample period tracks the years in which the surge of automation took place.

Our empirical approach closely follows that of Hazell et al. (2022), which we generalize to incorporate the role of automation on inflation dynamics. Specifically, we regress both non-tradable goods inflation and wage inflation on the lagged values of the unemployment rate and its interaction with robot adoption, while controlling in isolation for the role of robot adoption and the non-tradable goods relative price. Hazell et al. (2022) show that the estimated sensitivity of inflation to unemployment maps into the slope of the aggregate price Phillips curve implied by a multi-region model. This setting allows to saturate the regression with year fixed effects, which not only control for supply shocks and inflation expectations that are common across areas, but most importantly permit to absorb the endogenous response of monetary policy to common demand shocks (Beraja et al., 2019; McLeay and Tenreyro, 2020; Fitzgerald et al., 2023).

However, this cross-sectional analysis of inflation dynamics would not suffice to uncover the Phillips curve if local idiosyncratic supply shocks are correlated with the changes in the local labor market. To purge the variation of unemployment from idiosyncratic supply shocks, we follow Hazell et al. (2022) and instrument the unemployment rate with local tradable demand spillovers. In particular, for each metropolitan area we build a shift-share instrument the weights the log difference of value added of tradable industries at the national level with the value-added share of each tradable industry across areas. Then, to uncover the causal effect of automation, we instrument robot adoption with a variable that replaces the robot installation across industries observed in the U.S. with those of the five largest European economy, as in Acemoglu and Restrepo (2020a). Under the identifying restriction that robot demand shocks are weakly correlated across advanced countries, our instrumenting strategy isolates the supply-side component which caused the surge in the efficiency and widespread usage of robots.

In our baseline results, the interaction of unemployment and robot adoption is positive and highly statistically significant, indicating a significant role of automation in decoupling inflation and unemployment. This effect is also economically relevant: an increase in robot adoption by one standard deviation reduces the sensitivity of prince inflation and wage inflation to unemployment by 17% and 9%, respectively. In other words, the dampening of the wage sensitivity to unemployment due to automation accounts for 42% of the associated reduction in the price inflation responsiveness to unemployment. This differential magnitude suggests that that robot adoption also diminishes the influence of wage changes onto price changes. Overall, our empirical analysis uncovers three novel findings relating automation to inflation dynamics: robot adoption reduces (*i*) the sensitivity of price inflation to unemployment, (*ii*) the sensitivity of wage inflation to unemployment, and (*iii*) the pass-through from wages to prices.

Our empirical findings keeps holding in a comprehensive battery of robustness checks that validate the extent to which the role of automation in decoupling the movements of inflation and unemployment holds above and beyond both potential alternative explanations and confounding factors related to the surge in automation. For instance, we establish that the role of robot adoption in dampening the sensitivity of non-tradable goods inflation to the unemployment rate is always highly statistically significant even when controlling the role of the time-varying differences across MSAs in the age structure of the population (Aksoy et al., 2019; Acemoglu and Restrepo, 2022; Basso and Jimeno, 2021), the labor market participation of workers with different gender, race, and education, differences in the average marginal propensity to consume (Herreno and Pedemonte, 2022), the relevance of abstract, routine, manual, and offshorable occupations (Autor et al., 2013; Siena and Zago, 2021), as well as the exposure of the metropolitan areas to foreign import competition (Forbes, 2019; Heise et al., 2022, 2023).

To rationalize our empirical evidence on how automation alters inflation dynamics, we extend an otherwise standard New Keynesian model with two key features: search frictions in the labor market and the possibility of robot adoption, in the spirit of Acemoglu and Restrepo (2020a). The economy features a representative household consisting of a continuum of workers with perfect consumption insurance, who directly search for a job. The production sector is split over three layers: (*i*) a varying measure of producers that operate a linear technology using either robots or workers and post vacancies in the labor market, (*ii*) a continuum of monopolistically competitive wholesalers that purchase the goods of the producers and transform them into different varieties, and face a price setting friction in the form of Rotemberg costs, and (*iii*) a representative retailer that aggregates the different varieties with a CES technology into the final good. The economy is closed by a Taylor rule that sets the nominal interest rate responding to changes in inflation and the unemployment gap, subject to inertia.

Automation is modulated by producers' decision to use either workers or machines. Producers trade off the certainty of installing and operate with a robot with the uncertainty of possibly hiring a worker but – conditional on that – operate at a relatively higher efficiency. Specifically, upon entry – and after paying a fixed operating cost – producers draw an idiosyncratic efficiency in employing workers, and then decide to use either a labor technology (i.e., labor firms) or a machine technology (i.e., robot firms). Labor firms open costly vacancies at given posted wage, which are filled with a probability that depends on the labor market tightness. Machine firms purchase a robot from machine manufacturers, and produce with certainty. Machine manufacturers transform final goods into machines with a linear technology featuring robot-specific technological change. Accordingly, the relative price of robots declines with the level of technological change.

This setting defines an automation threshold, that is, a level of the efficiency in operating the labor technology that defines whether firms opt to either post a vacancy and look for workers or install a machine. This threshold crucially depends on the job filling probability and the levels of both wages and the price of robots. An increase in wages relative to the price of robots leads to more automation, as firms can replace workers with machines. In the model, the automation cut-off varies across steady states, as a function of the exogenous level of robot-specific technological change, and around the steady state upon the occurrence of a shock, as a function of the endogenous response of prices.

We then characterize the price Phillips curve and show that automation – interpreted as a positive robot-specific technological change – reduces its slope. The dampening effect of automation on the relationship between inflation and the unemployment gap is due to two main mechanisms. First, automation raises the fraction of firms operating with machines, reducing the labor share in value added. As a result, part of the aggregated demand adjustment is unrelated to changes in wage and unemployment dynamics. Second, the outside option of automating production negatively affects workers' bargaining power, dampening the responsiveness of wages to changes in the unemployment gap.

In the quantitative analysis, we consider two steady states that differ uniquely in the level of robot-specific technological change. These two steady states are carefully calibrated to replicate the standard deviation of robot penetration across MSAs in the data, which implies a 200% rise in the ratio of robots per employee. We find that a positive demand shock – which reduces the unemployment gap by the same amount across the two steady states – reduces the responsiveness of price inflation and wage inflation in the high-automation economy by 18% and 14%, respectively. These changes are remarkably in line with the magnitude of the effects of automation on the price and wage Phillips curve estimated in our empirical evidence, that suggest a flattening of 17% and 9%, respectively. Thus, our model implies a relatively larger flattening of the wage Phillips curve and a relatively more muted drop in the wage-to-price pass-through than that uncovered in the data.

Our model can rationalize not only the flattening of the price and wage Phillips curve in the pre-Covid period, but also the sudden resurgence of a steep Phillips curve. When ramping up automation is costly and machine manufacturers face adjustment costs, the threat that robots pose to workers' bargaining power crucially depend on the size of the shock realizations. When facing a small expansionary shock, firms can purchase additional machines without facing a sharp increase in robot prices, and thus gain an upper hand on wage negotiations. In this case, both the wage and price Phillips curves are flat. However, when the size of an expansionary shock is substantial, installing all the required robots to meet demand would be increasingly costly, forcing producers to continue to operate using labor. Consequently, the threat of robot adoption is no longer effective in curtailing workers' bargaining power, and wages highly react to changes in the unemployment gap. In other words, robot adoption alters the price Phillips curve such that its slope is relatively flat when the size of shocks is small, but can quickly become steep amidst large shock realizations.

Our work relates to the literature on the subdued dynamics of inflation in the post 1980s, suggestive of a flat Phillips curve (Blanchard, 2016; Stock and Watson, 2020). This evidence may be due to policy improvements and better anchoring of inflation expectations (Ball and Mazumder, 2011; McLeay and Tenreyro, 2020; Hazell et al., 2022; Bergholt et al., 2023), labor market changes muting the responsiveness of wages (Stansbury and Summers, 2020; Del Negro et al., 2020; Siena and Zago, 2021), globalization (Forbes, 2019; Heise et al., 2022), changes in the shocks composition (Gordon, 2013; Coibion and Gorodnichenko, 2015), changes in firm inter-linkages (Galesi and Rachedi, 2019; Höynck, 2020; Rubbo, 2023), financial frictions (Gilchrist et al., 2017), and a non-linear Phillips curve (Harding et al., 2022). We emphasize that automation can account for the flattening of price and wage inflation observed in the pre-Covid period, while also rationalizing a steep Phillips curve amidst large expansionary shocks.

The two closest papers to ours are Fornaro and Wolf (2021) and Leduc and Liu (2023). Fornaro and Wolf (2021) build a New Keynesian model with robot adoption to show that monetary policy accommodations can reconcile firms' intense usage of automation with limited effect on employment and inflation in medium and long run. We take a complementary approach by emphasizing that robot adoption decouples inflation and labor market dynamics in the short run, taking as given the stance of monetary policy. Leduc and Liu (2023) build a real model with robot adoption and search frictions to account for the business cycle fluctuations of unemployment. While our work share with theirs the focus on the threat that robots pose to workers' bargaining power, we look at how automation alters the slope of the price and wage Phillips curve.

2 Empirical Evidence

This section provides novel empirical evidence on how robot adoption leads to a decoupling between inflation and unemployment. Specifically, we study a panel of price inflation, wage inflation, unemployment, and robot adoption across U.S. metropolitan areas. To estimate the effect of automation on the relationship between inflation and unemployment, we use the variation across U.S. metropolitan areas in both tradable demand spillovers and robot adoption.

2.1 Data

We build a data set of non-tradable goods inflation, wage inflation, the unemployment rate, and robot adoption across 384 U.S. metropolitan areas at the annual frequency from 2008 to 2018. The frequency and the time period of our panel differ from those of Hazell et al. (2022) and Fitzgerald et al. (2023), as we start much later in time, from the early 2000s on, to capture the period in which automation took place.[1]

We use the information on the regional price parities of the U.S. Bureau of Economic Analysis (BEA), which gives a breakdown of prices at the MSA level by providing data on total prices, the price of goods, as well as distinct series for the price of rents, utilities, and other services. We complement it with information on 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), robot installed at the industry level for the U.S. and the five largest European countries from the International Federation of Robotics, employment at the industry-MSA level from the Quarterly Census of Employment and Wages of the BLS. To derive a measure of robot adoption at the MSA-year level, we follow the two-step procedure of Acemoglu and Restrepo (2020a): we compute the robot per employee for each industry at the U.S. national level, and combine it with the employment share of each industry at the MSA level. In this way, we end up with a ratio of installed robots per employee for each MSA-year pair.

Finally, we also consider value added at the industry-MSA level from the BEA, and employment at the industry-country level for the five largest European countries from EUKLEMS.

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:

$$\pi_{N,i,t} = \beta\, u_{i,t-1} + \gamma\, u_{i,t-1}\,(m_{i,t-1} - \bar{m}) + \zeta\, m_{i,t-1} + \chi\, p_{N,i,t} + \alpha_i + \delta_t + \epsilon_{i,t}, \quad (1)$$

[1] The data on prices at the annual frequency across 384 MSAs start in 2008. Although prices at the metropolitan areas are available also at the quarterly and semi-annual frequency well before than 2008, they only track around 20 MSAs. Consequently, we opt for a panel at the annual frequency from 2008 on to focus on the period of robot adoption while maximizing the cross-sectional dimension of our data.

where $\pi_{N,i,t}$ is the inflation rate of non-tradable goods of MSA i at year t, defined as the log-difference of the price of services excluding rents and utilities, $u_{i,t}$ is the lagged unemployment rate, $m_{i,t}$ denotes robot adoption, $\bar{m} = \sum_i \sum_n \frac{m_{i,t}}{n_i n_t}$ is its average value across all MSA-year observations, where n_i is the number of MSA in the sample and n_t is the number of years, and $p_{N,i,t}$ is the relative price of non-tradable goods. As in Ball and Mazumder (2011), Hazell et al. (2022) and Fitzgerald et al. (2023), we consider the unemployment rate as lagged by one year. Similarly, we also lag by one year the robot adoption variable. The regression also includes state fixed effects, α_i, and year fixed effects, δ_t.

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

We estimate the coefficients β and γ leveraging cross-sectional differences in unemployment rate, inflation, and robot adoption across metropolitan areas. For instance, the average value of the the unemployment rate at the MSA level in our sample equals 6.8%, but it is highly heterogeneously distributed, as it ranges from a value of 3% in Bismarck, ND up to 23.1% in Barnstable Town, MA. Metropolitan areas also differ substantially in the time variation of unemployment over time: the area with the smallest fluctuations is Anchorage, AK, in which the unemployment rate ranged between 5.4% and 7.4%, whereas Elkhart-Goshen, IN experienced swings between 2.5% and 18.1%. If anything, the variation in robot adoption across MSAs is even larger, since the metropolitan-level standard deviation of robot per employee is twice as large as its average value.

Importantly, our specification of regression (1) extends the approach of Hazell et al. (2022) in leveraging cross-sectional information to identify the slope of the Phillips curve to incorporate the role of automation. In a setting which abstracts from robot adoption (i.e., imposing $\gamma = \zeta = 0$), Hazell et al. (2022) show that the estimate of the coefficient β in regression (1) can be mapped into the aggregate slope of the Phillips curve implied by a multi-region model. This result hinges on

[2] As shown in Basso and Rachedi (2021), considering the interaction term of the unemployment rate with the demeaned robot-per-employee ratio, $m_{i,t-1} - \bar{m}$, does not alter the estimation of how robot adoption affects the relationship between inflation and unemployment. Rather, this normalization allows us to directly interpret the parameter β as the sensitivity of non-tradable goods inflation to the unemployment rate for a MSA with the average degree of robot adoption, that is, when $m_{i,t} = \bar{m}$.

the following conditions. First, the cross-sectional setting allows to saturate the regression with year fixed effects, which absorb the endogenous response of monetary policy to common demand shocks as well as capture the time-variation in common inflation expectations and supply shock realizations across metropolitan areas. Second, the presence of MSA fixed effects permits to control for any fixed heterogeneity across areas, such as time-invariant differences in inflation expectations, and also ameliorates the negative omitted variable bias from estimating the regression using actual unemployment rates and not the unemployment gaps.

Notwithstanding, this setting would not suffice to identify the slope of the Phillips curve because the presence of local idiosyncratic supply shocks which may be correlated with local unemployment rate could bias the estimate of β, as discussed by McLeay and Tenreyro (2020). To purge the variation in local unemployment rate from idiosyncratic supply shocks, we follow Hazell et al. (2022) and instrument the unemployment rate with local tradable demand spillovers. Specifically, the local tradable demand spillovers in area i at year t equals

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

where $s_{x,i}$ denotes the average value-added share of industry x in the metropolitan area i, and $\Delta \log s_{-i,x,t}$ is the log change in the national real value added of sector x excluding the contribution of the MSA i at year t. In other words, local tradable demand spillovers are defined as a shift-share variable in the spirit of Bartik (1991). As long as supply disturbances that may drive the time variation in national industry value added are not correlated with the heterogeneous relevance of industry value added across areas, the tradable demand spillovers provide a valid instrument.[3] As in Mian and Sufi (2014), the tradable industries are agriculture, mining, and manufacturing.

Since automation could be driven by local demand factors related to the dynamics of wages, prices, and the conditions of the labor market in each metropolitan area, we sharpen our identification of the effect of robot adoption on the relationship between inflation and unemployment following Acemoglu and Restrepo (2020a). In particular, we instrument the robot-to-employee ratio at the MSA-year pair with an alternative measure which replaces the robot installations

[3] Although the tradable demand spillovers are defined as a shift-share variable as for the case of automation, we use industry value-shares for the former and industry employment shares for the latter. In this way, we make sure that the two variables do not strongly comove. In our sample, the correlation between the tradable demand instrument and the robot adoption variable is around 0.2.

for each industry at the U.S. nation level with the average robot installation per industry in the largest five European economies. Under the identifying restriction that robot demand shocks are weakly correlated across advanced countries, our instrumenting strategy isolates the supply-side component which caused the surge in efficiency of robots, and thus boosted their widespread usage.

We also study the effect of robots on the sensitivity of wage inflation to unemployment by considering a setting identical to regression (1), with the only difference that the dependent variable is $\pi_{W,i,t}$, defined as the log-difference in the average compensation per job of MSA i at year t. This case allows us to study whether automation alter the relationship between wage changes and unemployment, and to what extent robot adoption imply a differential sensitivity of unemployment for wage and price inflation.

2.3 Results

Panel A of Table 1 reports the results on how automation alters the sensitivity of non-tradable goods inflation to unemployment. Columns (1) and (2) focus on a case of regression (1) which abstracts from the interaction between robot adoption and the unemployment rate, with the only difference that Column (1) uses OLS methods whereas Column (2) instruments unemployment with tradable demand spillovers. The OLS estimate of the sensitivity of price inflation to unemployment equals -0.1884, is highly statistically significant, and its magnitude is in line with previous estimates of Hazell et al. (2022), while being substantially lower than those of McLeay and Tenreyro (2020). However, the results of Column (2) provide a much steeper relationship between unemployment and inflation, with an estimate of β that equals -0.7031, slightly above the IV estimate of McLeay and Tenreyro (2020) that leverages variation in government spending across metropolitan areas. Our results is consistent also with the evidence of Hazell et al. (2022) and Fitzgerald et al. (2023), that point out how using variation across regional areas leads to a much steeper relationship between inflation and unemployment than when focusing on aggregate data at the national level.

Columns (3) and (4) report the results of the baseline regression that includes the interaction of robot adoption and unemployment, estimated with OLS and IV methods, respectively. In either case, the role of automation is statistically significant at the 5% confidence level, and the magnitude of the coefficient rises substantially when instrumenting both unemployment with tradable spillovers and robot adoption with that implied by the automation patterns of European countries. The fact that the estimated coefficient displays a negative sign implies

Table 1: Robot Adoption and Inflation across MSAs

	No Interaction Term		Baseline	
	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)
	Panel A — Dependent Variable: $\pi_{N,i,t}$			
$u_{i,t-1}$	-0.1884***	-0.7031***	-0.1884***	-0.5069***
	(0.0226)	(0.1364)	(0.0221)	(0.1381)
$u_{i,t-1} \times (m_{i,t-1} - \bar{m})$			0.0010**	0.0066**
			(0.0004)	(0.0030)
Year Fixed Effects	✓	✓	✓	✓
MSA Fixed Effects	✓	✓	✓	✓
N. Observations	3,205	3,205	3,205	3,205
	Panel B — Dependent Variable: $\pi_{W,i,t}$			
$u_{i,t-1}$	-0.3848***	-1.0341***	-0.3855***	-0.9580***
	(0.0330)	(0.1503)	(0.0330)	(0.2450)
$u_{i,t-1} \times (m_{i,t-1} - \bar{m})$			0.0016**	0.0049**
			(0.0007)	(0.0024)
Year Fixed Effects	✓	✓	✓	✓
MSA Fixed Effects	✓	✓	✓	✓
N. Observations	3,205	3,205	3,205	3,205

Note: The table reports the estimates of panel regressions across U.S. MSAs on annual data from 2008 to 2018. In Panel A, the dependent variable is the inflation rate of non-tradable goods, $\pi_{N,i,t}$. In Panel B, the dependent variable is wage inflation, $\pi_{W,i,t}$. In all regressions, the key independent variables are the lagged value of the unemployment rate, $u_{i,t-1}$, the interaction between the lagged value of the unemployment rate and the lag value of the demeaned robot-adoption variable, $u_{i,t-1} \times (m_{i,t-1} - \bar{m})$. In the IV regressions, the unemployment rate is instrumented with a shift-share variable that captures tradeable 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 value of the robot-adoption variable, $m_{i,t-1}$, the relative price of non-tradable goods, $p_{N,i,t-1}$, as well as year and MSA fixed effects. Columns (1) and (2) report the results of a regression which abstracts from the interaction between the lagged value of the unemployment rate and the lag value of the demeaned robot-adoption variable, while Columns (3) and (4) report the results of the baseline regression which explicitly incorporates the role of the interaction term. Columns (1) and (3) are estimated using OLS methods, and Columns (2) and (4) are estimated using instrumental variables. Double-clustered standard errors are reported in brackets. *** and ** indicate statistical significance at the 1% and 5%, respectively.

that metropolitan areas with relatively more robots feature a price inflation that is relatively less reactive to changes in the local labor market. In other words, automation decouples inflation from unemployment. Importantly, the estimation of the role of automation is also highly economically significant: a one standard deviation in robot adoption reduces the sensitivity of inflation to unemployment by 19% with respect to the sensitivity of the metropolitan area featuring the average value of robots per employee.

Similarly to the different cases presented by Panel A of Table 1, Panel B reports the results on how automation alters the relationship between unemployment and wage inflation. Also in this case the coefficient associated to the interaction term between unemployment and automation is statistically significant at the 5% confidence level for both the OLS and IV regressions. Interestingly, while the effect of automation on the implied wage Phillips curve at the MSA level is economically relevant, its magnitude falls short of the magnitude of the effects of robot adoption on the price Phillips curve: a one standard deviation in robot adoption reduces the sensitivity of wage inflation to unemployment by 8% with respect to the sensitivity of the metropolitan area featuring the average value of robots per employee. In other words, the dampening of the wage sensitivity to unemployment due to automation accounts for 42% of the associated reduction in the price inflation responsiveness to unemployment, suggesting that robot adoption could also blunt the influence of wage inflation into price inflation.

Overall, this analysis has established three main results: automation reduces (*i*) the sensitivity of price inflation to unemployment, (*ii*) the sensitivity of wage inflation to unemployment, and (*iii*) the pass-through from wages to prices.

2.4 Robustness Check

Our results on the relationship between robot adoption and inflation dynamics is validated in an extensive battery of robustness checks. We use this analysis to evaluate the extent to which the effect of automation in decoupling inflation and unemployment holds above and beyond alternative explanations. In particular, we consider three groups of potential confounding factors with differences across metropolitan areas in demographic characteristics, occupational structure, and exposure to international trade. We report the results of these exercises in Appendix A.

First, we show that robot adoption dampens the sensitivity of inflation to unemployment even when including the interaction of the unemployment rate with differences in the age structure of the population across MSA, proxied with either the share of individuals below 30 years old, or the share of individuals above 60 years old, indicating that the effect of automation on price changes is not related to its relationship with an aging labor force (Acemoglu and Restrepo, 2022; Basso and Jimeno, 2021), and the way in which population aging affects the long-run dynamics of inflation (Aksoy et al., 2019). Our evidence holds even when interacting unemployment with measures capturing differences across MSAs in the labor force participation of women, black people, and asians, as well as in differences in educational attainments and overall labor force participation. We also show that the role of robot adoption keeps being statistically significant even when including differences in the marginal propensity to consume across areas (Herreno and Pedemonte, 2022).

Second, our results hold above and beyond on the interaction of unemployment with differences across MSAs in occupations. In particular, we consider variations in the presence of either abstract, routine, manual, as well as the extent to which occupations are offshorable. These characteristics are relevant as Siena and Zago (2021) document that the flattening of the price Phillips curve is related to the phenomenon of job polarization away from routine occupations, which is also directly related to the offshoring of routine activities toward low labor-cost countries (Autor et al., 2013).

Third, the automation dampening of the inflation sensitivity to unemployment is also robust to explicitly incorporating the role of import competition, measured in terms of MSA exposure to either Chinese imports, or Mexican imports, or both. Thus, our findings holds above and beyond the way in which variations in import competition alter wage and price inflation dynamics (Forbes, 2019; Heise et al., 2022, 2023).

3 Model

The model extends a standard New Keynesian economy to incorporate search frictions in the labor market and robot adoption, in the spirit of Acemoglu and Restrepo (2020a). The production side is split over three layers: (*i*) a varying measure of producers that can post vacancies in the labor market and opt to operate with a linear technology using either labor or machines, (*ii*) a continuum of monopolistically competitive wholesalers, that purchase the goods of producers, convert them into different varieties, and face price setting frictions, and (*iii*) a representative retailer, that purchases the different varieties and assemble them into the final good. Final goods are sold to the household and machine manufacturers, that transform them into machines using a linear technology subject to robot-specific technological change. The household consists of a continuum of workers, who directly look for a job. Income is pooled at the household level, who collectively decides consumption and asset holdings. The monetary authority sets the nominal interest rate according to a Taylor rule.[4]

3.1 Labour Market

Labour markets consist of a set of sub-markets with unit measure, which are indexed by $\omega \in [0,1]$. At each point in time, there is a time-varying measure $\Xi_{L,t}$ of producers posting vacancies at a given wage, which we refer to as labor firms. We denote with $v_{\omega,t}$ as the number of vacancies in each sub-market, such that $\int_0^1 v_{\omega,t}\,\mathrm{d}\omega = \Xi_{L,t}$, and $W_{\omega,t}$ is the associated nominal wage upon a successful match. On the other side there are workers, who decide in which sub-market to search for a job.[5] We denote by s_ω as the measure of workers searching in each sub-market. If workers match with a producer, they earn the nominal wage, and otherwise they receive no income.[6] Given the number of vacancies and searching workers in each sub-market, the successful flow of matches, $x_{\omega,t}(v_{\omega,t}, s_{\omega,t})$, is pinned down by the matching function

$$x_{\omega,t}(v_{\omega,t}, s_{\omega,t}) = \xi v_{\omega,t}^{\eta} s_{\omega,t}^{1-\eta}, \tag{3}$$

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

Given the matching function (3) and the labor market tightness in the sub-market ω, $\theta_{\omega,t} = v_{\omega,t}/s_{\omega,t}$, which describes the ratio between number of vacancies and number of searching workers, the probability that worker finds a job equals

$$p_{\omega,t}\left(\theta_{\omega,t}\right) = \frac{x_{\omega,t}(v_{\omega,t}, s_{\omega,t})}{s_{\omega,t}} = \xi \theta_{\omega,t}^{\eta} \tag{4}$$

and the probability of filling a vacancy is

$$q_{\omega,t}\left(\theta_{\omega,t}\right) = \frac{x_{\omega,t}(v_{\omega,t}, s_{\omega,t})}{v_{\omega,t}} = \xi \theta_{\omega,t}^{\eta-1}. \tag{5}$$

[4] Appendix B provides a graphical description of the structure of the model.

[5] In the baseline model, we assume that all workers search for a job in each period. In Appendix, we consider an extension in which labor market participation is allowed to vary across workers.

[6] We abstract from the presence of unemployment benefits as we assume perfect consumption insurance across workers within the household.

The payoff of workers searching in the sub-market ω equals the product between the nominal wage rate in case of a successful match and the probability of finding a job, that is,

$$J_{s,\omega,t} = p_{\omega,t}(\theta_{\omega,t})\, W_{\omega,t}. \tag{6}$$

Workers decide in which sub-market to search for a job trading off the wage rate and the probability to find a job. In a symmetric equilibrium, workers' payoff should be equalized across all active sub-markets, such that $J_{s,\omega,t} = J_{s,t}$ for all ω. Consequently, sub-markets offering higher wage rates feature lower probabilities to find a job.

The equilibrium in the labor market implies the sum of workers searching in all sub-markets equals the overall unit measure of workers in the household,

$$1 = \int_0^1 s_{\omega,t}\, \mathrm{d}\omega. \tag{7}$$

Accordingly, at the end of the period the unemployment rate equals the difference between the total number of workers and the measure of workers that have matched with a producer,

$$u_t = 1 - \int_0^1 p_{\omega,t}(\theta_{\omega,t}) s_{\omega,t}\, \mathrm{d}\omega. \tag{8}$$

3.2 Producers

At each point of time, there is a total measure Ξ_t of producers that decide to pay a per-period fixed nominal operating cost κ to enter the market. We index each producer with $j \in [0, \Xi_t]$. Upon entry, producers draw an idiosyncratic efficiency in operating with a labor technology, $\gamma_{L,j}$, from a distribution $f(\gamma)$ with support $[\underline{\gamma}, \bar{\gamma}]$, where $\underline{\gamma}$ and $\bar{\gamma}$ denote the minimum and maximum level of producers' labor efficiency.

After drawing the labor efficiency level, producers decide to operate employing either machines (i.e., robot firms) or workers (i.e., labor firms). In case a producer decides to operate using machines, it purchases a robot from machine manufacturers at price $P_{M,t}$, and produces with certainty using a linear technology that is characterized by an ex-ante known efficiency level γ_M. The efficiency level of robot firms lies below the upper bound of producers' labor efficiency, such that $\gamma_M < \bar{\gamma}$. Robot firms then sell their output to to wholesalers at price $P_{P,t}$, such that their nominal value equals the nominal value of sales net of the cost of purchasing a robot and the entry cost,

$$V_{M,j,t} = P_{P,t}\gamma_M - P_{M,t} - \kappa. \tag{9}$$

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

In case a producer decides to operate using labor, then it opens a vacancy in a given sub-market at the nominal wage rate $W_{\omega,t}$. Posting a vacancy comes at a cost χ. Upon filling the vacancy, the labor firm produces using a linear technology at the labor efficiency rate γ_j, and sells its output to wholesalers at price $P_{P,t}$. Consequently, the nominal value of a labor firm equals the nominal value of sales net of the wage rate, multiplied by the probability of filling the vacancy, minus both the entry cost and the vacancy posting cost,

$$V_{L,j,t} = q_{\omega,t}(\theta_{\omega,t})\left(P_{P,t}\gamma_{L,j} - W_{\omega,t}\right) - \kappa - \chi. \tag{10}$$

Labor firms decide the nominal wage rate associated to their vacancies to maximize their value given the labor market tightness and subject to preserving a positive payoff for workers in each sub-market. Optimality then implies that the nominal wage rate equals

$$W_{\omega,t} = P_{P,t}\gamma_{L,j}(1-\eta). \tag{11}$$

In other words, the variation in wages across labor firms is uniquely pinned by the dispersion in the labor efficiency values. This result implies that in equilibrium firms with different efficiency levels, $\gamma_{L,j}$ sort themselves into different sub-markets, ω. Since the labor efficiency is assigned randomly, hereafter we use firms' labor efficiency levels to denote the sub-markets. For instance, we refer to wage $W_{\gamma_{L,j},t}$ as the rate offered by firms posting a vacancy in the sub-market populated by labor firms with efficiency level $\gamma_{L,j}$.

How do producers sort into labor firms and robot firms? A producer j opts to open a vacancy and operate the labor technology if and only if the value of being a labor firms is greater than the value of being a robot firm, that is, $V_{L,j,t} > V_{M,t}$. Since the value of being a labor firm increases with the labor efficiency level $\gamma_{L,j}$,[7],

[7]See the Appendix for a proof of this property.

there exists a cut-off point for the labor efficiency level, $\gamma_t^\star$, such that

$$V_{L,j,t}\left(\gamma_t^\star\right) = V_{M,t}, \tag{12}$$

and firms are indifferent between operating the labor technology or the machine technology. The cut-off point crucially defines the automation choices, since all producers with a labor efficiency level above $\gamma_t^\star$ become labor firms, whereas all the rest become robot firms.

Given the 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 an efficiency above $\gamma_t^\star$,

$$\Xi_{L,t} = \Xi_t \int_{\gamma_t^\star}^{\bar{\gamma}} f\left(\gamma\right)\, \mathrm{d}\gamma, \tag{13}$$

whereas the measure of robot firms captures all producers with sufficiently low labor efficiency:

$$\Xi_{M,t} = \Xi_t \int_{\underline{\gamma}}^{\gamma_t^\star} f\left(\gamma\right)\, \mathrm{d}\gamma. \tag{14}$$

In equilibrium, the sum of the measures of labor firms and robot firms equals the total amount of producers that have entered the market, that is, $\Xi_{L,t} + \Xi_{M,t} = \Xi_t$.

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

$$Z_t = \Xi_t \int_{\gamma_t^\star}^{\bar{\gamma}} q_{\gamma_{L,j},t}(\theta_{\gamma_{L,j},t})\gamma_{L,j}\, \mathrm{d}j + \Xi_{M,t}\gamma_M. \tag{15}$$

Finally, we can characterize what is the total measure of producers entering the market: a producer enters the market as long as its expected value, $V_{e,t}$, equals zero:

$$V_{e,t} = \int_{\underline{\gamma}}^{\gamma_t^*} V_{M,t} f(\gamma_{L,j}) d\gamma + \int_{\gamma_t^*}^{\bar{\gamma}} V_{L,j}(\gamma_{L,j}) f(\gamma) d\gamma = 0. \tag{16}$$

We can also obtain the average wage:

$$\overline{W}_t = \frac{\int_{\gamma_t^*}^{\bar{\gamma}} P_{P,t}\gamma_{L,j}(1-\eta) f(\gamma) d\gamma}{\int_{\gamma_t^*}^{\bar{\gamma}} f(\gamma) d\gamma}. \tag{17}$$

3.3 Wholesalers

There is a unit measure of monopolistically competitive wholesalers, indexed by $i \in [0,1]$. Each wholesaler purchases goods $Z_{i,t}$ from the producers at price $P_{P,t}$, and transforms them into a different variety $Y_{i,t}$ with the linear technology:

$$Y_{i,t} = Z_{i,t}. \tag{18}$$

The varieties are sold to retailers at the price $P_{i,t}$. Then, wholesalers' profits equal to $P_{i,t}Y_{i,t} - P_{P,t}Z_{i,t}$.

Wholesalers face price-setting friction in the form of Rotemberg adjustment cost, denoted by the parameter ϕ. Thus, wholesalers optimally set their price $P_{i,t}$ by maximizing expected profits net of the Rotemberg costs

$$\max_{P_{i,t}} \mathbb{E}_t \left\{ \sum_{k=t}^{\infty} Q_{k,t} \left(P_{i,k}Y_{i,k} - P_{P,k}Z_{i,k} - \frac{\phi}{2} \left[\frac{P_{i,k}}{P_{i,k-1}} - 1 \right]^2 Y_{i,k} \right) \right\}, \tag{19}$$

where $Q_{s,t}$ is households' stochastic discount factor. In a symmetric equilibrium, all wholesalers set the same price, such that $P_{i,t} = P_t$ for all i. We denote by $\pi_t = \frac{P_t}{P_{t-1}}$ the inflation rate.

The market clearing condition implies that the total amount of goods produced by the wholesalers – net of the Rotemberg adjustment cost – equals those produced by both labor firms and machine firms,

$$\int_0^1 \left[1 - \frac{\phi}{2} \left(\frac{P_{i,t}}{P_{i,t-1}} - 1 \right)^2 \right] Y_{i,t} \,\mathrm{d}i = \int_0^1 Z_{i,t} \,\mathrm{d}i = Z_t. \tag{20}$$

3.4 Retailers

There is a perfectly competitive representative retailer that purchases all the varieties from the wholesalers, $Y_{i,t}$, and assembles them into the final good of the economy, Y_t, with a CES technology:

$$Y_t = \left[\int_0^1 Y_{i,t}^{\frac{\epsilon-1}{\epsilon}} \,\mathrm{d}i \right]^{\frac{\epsilon}{\epsilon-1}}, \tag{21}$$

where ϵ is the elasticity of substitution across varieties. The retailer then sells the final goods at price P_t to households and machine manufacturers. Accordingly, retailers' optimal demand of each variety is

$$Y_{i,t} = \left(\frac{P_{i,t}}{P_t}\right)^{-\epsilon} Y_t, \tag{22}$$

where the price of final goods is given by

$$P_t = \left[\int_0^1 P_{i,t}^{1-\epsilon}\,\mathrm{d}i\right]^{\frac{1}{1-\epsilon}}. \tag{23}$$

The total amount of final goods is then sold to household, in form of consumption goods C_t, and to machine manufacturers, in form of investment goods I_t, such that the following clearing condition applies

$$Y_t = C_t + I_t. \tag{24}$$

3.5 Machine Manufacturers

There is a perfectly competitive representative machine manufacturer that purchases final goods from the retailer I_t at price P_t, and transform them into machines M_t with the linear technology

$$M_t = \zeta I_t, \tag{25}$$

where ζ is the level of robot-specific technological change. The manufacturers sell the machines to the robot firms at price $P_{M,t}$.This price inversely relates to the level of technological change, such that

$$P_{M,t} = \frac{1}{\zeta} P_t. \tag{26}$$

A higher value of robot-specific technological change implies that the production of machines is becoming relatively more efficient. Consequently, the price of machines goes down.

In equilibrium, the total amount of machines sold by the manufacturers equals the total amount of machines demanded by the robot firms (i.e., the measure of robot firms), that is

$$M_t = \Xi_{M,t}. \tag{27}$$

3.6 Households

The household consists of a unit measure of workers featuring perfect consumption insurance. Consequently, after the workers have been searching for the job and the matches are realized, all the nominal labor earnings X_t are pooled together within the household, such that

$$X_t = \Xi_t \int_{\gamma_t^\star}^{\bar{\gamma}} q_{\gamma_{L,j},t} \left(\theta_{\gamma_{L,j},t}\right) W_{\gamma_{L,j},t} f\left(\gamma\right) \mathrm{d}\gamma. \tag{28}$$

In other words, taking the nominal wage rate of all the sub-markets/efficiency levels which are not automated and multiplying for the associated probability to find a job yields the aggregate labor earnings of the household.

The household then decides the optimal levels of consumption, C_t to purchase from retailers at price P_t, and savings in one-period nominal bonds, B_t. Specifically, the household maximizes its lifetime utility

$$\max_{C_t, B_{t+1}} \mathbb{E}_t \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\sigma}}{1-\sigma} \tag{29}$$

$$\text{s.t. } P_t C_t + B_t = B_{t-1} R_{t-1} + X_t \tag{30}$$

where R_t denotes the return of bonds.

3.7 Monetary Authority

The monetary authority sets the nominal interest rate R_t following a standard Taylor rule that reacts to the inflation rate, π_t, and the unemployment gap, u_t/u_t^F, where u_t^F is the unemployment rate in a version of the economy featuring flexible prices, such that

$$\frac{R_t}{\bar{R}} = \left[\frac{R_{t-1}}{\bar{R}}\right]^{\psi_R} \left[(1+\pi_t)^{\psi_\pi} \left(u_t/u_t^F\right)^{\psi_u}\right]^{1-\psi_R}, \tag{31}$$

where $\bar{R}$ is the steady-state nominal interest rate, 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, respectively.

4 Quantitative Analysis

4.1 Calibration

We assume the distribution of productivity γ is a Truncated Pareto Distribution with location parameters, $\underline{\gamma}$ and $\bar{\gamma}$, and shape parameter, α. Consequently, $f(\gamma) = \frac{\alpha\underline{\gamma}^{\alpha}\gamma^{-\alpha-1}}{1-\underline{\gamma}^{\alpha}\bar{\gamma}^{-\alpha}}$ and set the lowest value in the support $\underline{\gamma} = 1$. To ensure the probability of finding a job in the low wage sub-markets (where few workers search) and that the probability of filling a vacancy in the high wage sub-markets (where many workers search) are within zero and one, we set $\alpha = 7$, $\bar{\gamma} = 1.1$ (thus labour productivity of the most productive firms is 10% higher than the worst firm) and the efficiency of the matching function $\xi = 0.9$. Households utility is given by $U(C) = \frac{C^{1-\sigma}}{1-\sigma}$, where we set $\sigma = 2$ and the discount factor is set to 0.995. We set the elasticity of substitution across varieties to $\epsilon = 9$, which implies a markup of 12.5%, in the ball park of the estimates used in the literature of New Keynesian models. The price rigidity parameter (ϕ) is set such that on average firms adjust their prices every 12 months. We set the cost of opening vacancies (χ) and the fixed cost of entry (κ) such that the unemployment rate is 5.7%, matching the month average in the US, and the the relative price of robots $P_{M,t}/P_t$ is below one, thus operating profits of low labour productivity firms, which decide to use machines, is positive. We follow Petrongolo and Pissarides (2001) and set the elasticity of matches to unemployment $(1-\eta)$ to 0.5. Finally, the Taylor rule parameters are: the inertia parameter equals $\psi_R = 0.8$, the degree of response to the inflation rate is $\psi_\pi = 1.5$, and the degree of response to the unemployment gap is $\psi_u = -0.2$.

Parameter	Values	Criteria
$\underline{\gamma}$	1	Low Productivity normalize to 1
α	7	Set to ensure
$\bar{\gamma}$	1.1	$p_{\omega=0}\left(\theta_{\omega=0}\right) < 1$
ξ	0.9	$q_{\omega=1}\left(\theta_{\omega=1}\right) < 1$
σ	2	Standard Parameter Calibration
β	0.995	Standard Parameter Calibration
ϵ	9	Mark-up = 12.5 %
ϕ	94.6	Average firms adjust their prices every 12 months
χ	0.2	
κ	0.44	Given χ, unemployment = 5.7%
η	0.5	Petrongolo and Pissarides (2001)
$P_{M,t}/P_t$	< 1	Given $\underline{\gamma}$, robot production is feasible
ψ_R	0.8	Standard Parameter Calibration
ψ_π	1.5	Standard Parameter Calibration
ψ_u	-0.2	Standard Parameter Calibration

4.2 Degree of Automation and the Slope of the Phillips Curve

Our model can be described as a standard New Keynesian model augmented to incorporate an automation choice in production that by replacing workers with machines also affects labour markets. As such the equilibrium conditions can then be segregated into four groups, the standard three equation blocks from the NK model (the policy rule (31), the IS curve coming from the household Euler equation (solution of (29)) and goods market clearing conditions, and the pricing equation (solution to (19)) and a fourth set of equations that determines the labour market outcomes, the automation level and ultimately the price of intermediate goods that serve as input into the pricing decision (combining (7), (8), (12), (16), and (27)).

After log-linearisation around the steady state, we incorporate the equilibrium conditions in the input production and labour markets into the pricing equation to determine the relationship between inflation (prices) and unemployment, or the Phillips curve. Details of the derivation are found in the Appendix. Let $\Theta \equiv \{\eta, \gamma_m, \gamma_H, \alpha, \psi\}$ represent the set of key structural parameters, and γ_S^* the cut-off point that determines the share of production that is automated versus the one that is labour-intensive at the steady state, then the Phillips Curve in our model economy is given by

$$\hat{\pi}_t = \frac{\psi - 1}{\phi}\Xi(\gamma_S^*; \Theta)(\hat{u}_t - \hat{u^F}_t) + E_t\beta\frac{\psi - 1}{\phi}\hat{\pi}_{t+1} \tag{32}$$

where

$$\Xi(\gamma_S^*; \Theta) = \frac{u_S}{1 - u_S}\frac{1}{\left(-\frac{\eta}{1-\eta} - \frac{\eta((\psi-1)/\psi)\gamma_m}{\omega_{1,S}}\left(\frac{1}{1-\eta}\varpi_{2,S} - \varpi_{3,S}\left(1 + \varpi_{2,S}\right)\right)\right)} \tag{33}$$

$$\varpi_{1,S}(\gamma_S^*; \Theta) = \frac{\xi^{1/\eta}\eta((\psi-1)/\psi)\left(\gamma_S^*\right)^{(1/\eta)}}{\left(\frac{(1-u_S)\left((\gamma_S^*)^{-\alpha+1/\eta} - (\gamma_H)^{-\alpha+1/\eta}\right)}{\left((\gamma_S^*)^{-\alpha+(1-\eta)/\eta} - (\gamma_H)^{-\alpha+(1-\eta)/\eta}\right)}\frac{\alpha-(1-\eta)/\eta}{\alpha-1/\eta}\right)^{(1-\eta)/\eta}}$$

$$\varpi_{2,S}(\gamma_S^*; \Theta) = \frac{1 - \frac{\gamma_m^\alpha}{(\gamma_S^*)^\alpha}}{\frac{\alpha}{\alpha-1/\eta}\left(\frac{\gamma_m^\alpha}{(\gamma_S^*)^\alpha} - \frac{\gamma_m^\alpha}{(\gamma_H)^\alpha}\frac{(\gamma_H)^{1/\eta}}{(\gamma_S^*)^{1/\eta}}\right)}$$

$$\varpi_{3,S}(\gamma_S^*; \Theta) = \left(\frac{\left(\alpha - \frac{1-\eta}{\eta}\right)\left(\gamma_S^*\right)^{-\alpha+(1-\eta)/\eta}}{\left(\gamma_S^*\right)^{-\alpha+(1-\eta)/\eta} - \left(\gamma_H\right)^{-\alpha+(1-\eta)/\eta}} - \frac{\left(\alpha - \frac{1}{\eta}\right)\left(\gamma_S^*\right)^{-\alpha+1/\eta}}{\left(\gamma_S^*\right)^{-\alpha+1/\eta} - \left(\gamma_H\right)^{-\alpha+1/\eta}}\right)$$

Thus, the slope of the Phillips curve in our setting, as in the standard NK model, is a function of the firms mark-up (controlled by parameter ψ), the degree of nominal rigidity (controlled by parameter ϕ), but is also a function of the degree of automation in the economy (controlled by parameter γ_S^*).[8] Holding the initial level of unemployment, u_S (which is pinned down by the cost of entry κ_f, affecting only the steady state and bearing no direct implications for the dynamics around the steady state) and all the key structural parameters (Θ) constant, the degree of automation is determined by the relative price of robots (q). Graetz and Michaels (2018) show that the price of robots have fallen during the last decades while automation has increased, interpreting this as the result of robot-specific technological change. What is the effect of this robot-specific technological change that lead the economy towards a new steady state with lower prices of robots and higher degree of automation on the Phillips Curve?

Figure 1 shows the effect of increasing the degree of automation on the slope of the Phillips curve for the benchmark calibration and also the effect of the key structural parameters in this relationship.

Firstly, we see that increasing automation decreases the slope of the Phillips Curve for all the relevant parameter ranges. Second, the effect of automation

[8]Formally the degree of automation is defined as $\int_{\gamma_m}^{\gamma_S^*} f(\gamma)d\gamma = \frac{(1-\gamma_m^\alpha(\gamma_t^*)^{-\alpha})}{(1-\gamma_m^\alpha(\gamma_H)^{-\alpha})}$.

Figure 1: Degree of automation and the slope of the Phillips curve

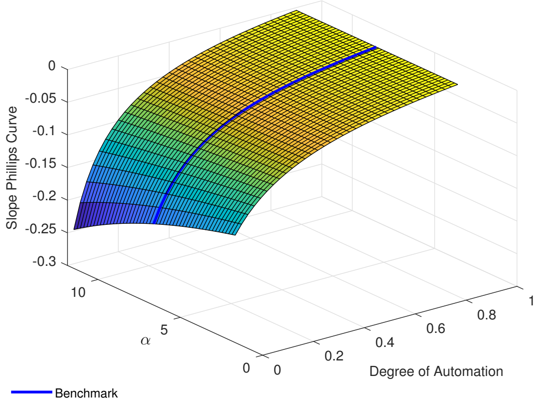

(a) Effect of α - Shape of productivity distribution

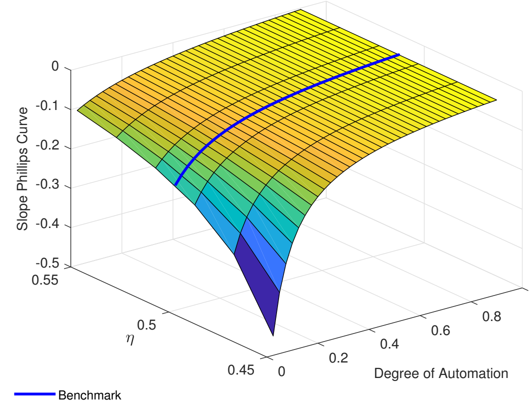

(b) Effect of η - Elasticity of vacancies to matches

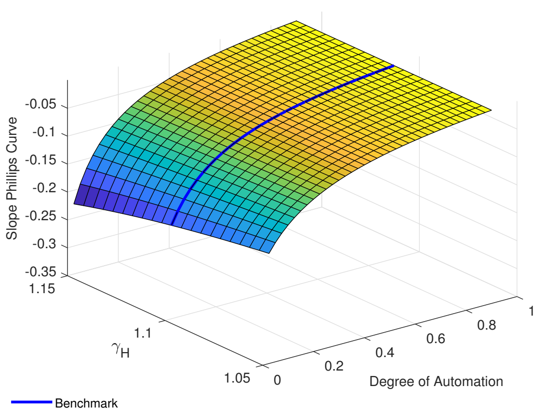

(c) Effect of γ_H - Upper bound on labour productivity

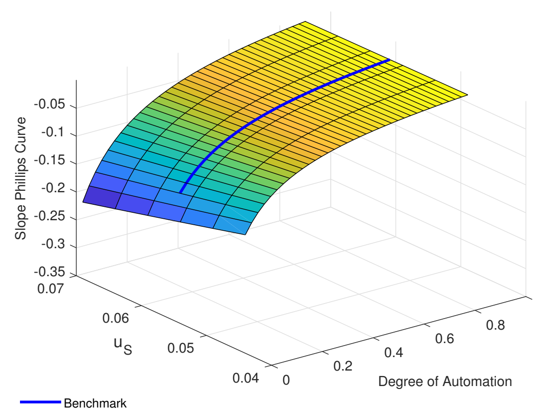

(d) Effect of u_S - Unemployment at steady state on the Phillips Curve is weaker, the bigger the elasticity of unemployment to matches (η) and the smaller the initial level of unemployment (u_S), thus the potential effect of automation is stronger when unemployment is relatively less important (whereas the effect of firms' decision gain weight) in determining the equilibrium in the labour markets. Third, the higher the density of firms in the lower end of the productivity scale (α is high), the larger the change in the slope when automation increases from initially low levels (below 30%). Fourth, in economies with higher upper bound in labour productivity imply more automation has only a slightly stronger effect on the slope of the Phillips Curve. Finally, at high levels of automation changes in unemployment have hardly any effect on inflation, independent of the parametrisation.

We can also obtain a close form solution for the first order approximation of the Phillips Curve in an extension of the model in which labour market participation is endogenously set (the details of the derivation are shown in the Appendix). In this case labour supply is set such that $L_t = \frac{(\frac{J_{E,t}}{P_t} U'(C_t))^{\frac{1}{\zeta}}}{\lambda_H}$, depending on the expected gain from joining the labour markets and the marginal utility of consumption. The main qualitative results remain unchanged, a higher degree of automation is associated with a flatter Phillips curve for all parameter ranges (see figure C.2 in the Appendix).

The level of automation and how it changes in response to shocks influence the dynamics of wages and prices. As automation decreases the labour share, part of the adjustment of aggregate demand will be done by firms using robots and thus follows different dynamics than the one determined at the labour market. Furthermore, in our model firms have the outside option of using robots instead of opening vacancies. As a result, automation will also influence the bargaining power of workers and firms that result in different wage responses for a given aggregate demand conditions. In order to analyse the relationship between wage and prices and how automation affect their responses, we also report the effect of automation on the wage to price pass-through. We use the condition that determines the average wage ((17)) and following a similar procedure done for the Phillips Curve (see the Appendix for detail) we can determine the relationship between the changes in the price of inputs P_t^I (effectively the marginal cost, which direct impact prices through the firms optimal pricing condition) and the changes in the average wage, obtaining

$$\widehat{p_{I,t}} = \Upsilon(\gamma_S^*;\Theta)\hat{w}_t \tag{34}$$

$$\Upsilon(\gamma_S^*;\Theta) = \frac{1}{1+\varpi_{pw,S}\frac{\eta p_{I,S}\gamma_m}{\omega_{1,S}}\left(1+\varpi_{2,S}\right)}$$

$$\varpi_{pw,S}(\gamma_S^*;\Theta) = \frac{(\gamma_S^*)^{-\alpha}\left((\gamma_S^*)^{-\alpha}-(\gamma_H)^{-\alpha}\right)+\alpha(\gamma_S^*)^{-\alpha}\gamma_H^{-\alpha}(\gamma_S^*-\gamma_H)}{\left((\gamma_S^*)^{-\alpha}-\gamma_H^{-\alpha}\right)\left((\gamma_S^*)^{-\alpha+1}-\gamma_H^{-\alpha+1}\right)}$$

$\Upsilon(\gamma_S^*;\Theta)$, therefore provides a measures the wage to inflation pass-through. We can then verify how $\Upsilon(\gamma_S^*;\Theta)$ changes with the degree of automation and the structural parameters as done for the slope of the Phillips Curve. Pass-through decreases with automation for all the relevant parameter ranges (Figure 2 show the results for the effect of automation when α and u_S also change, see the appendix for the additional parameters in Θ). Although lower wage inflation responses due to automation is part of the mechanism delivering a flatter Phillips Curves, automation also weakens the relationship between wages and prices. Thus, our model provides a rationale for the findings of Del Negro et al. (2020) and Heise et al. (2022), who uncover evidence that the flattening of the Phillips curve is related to the decoupling of wage and price dynamics.[9]

Figure 2: Degree of automation and the wage to price pass-through

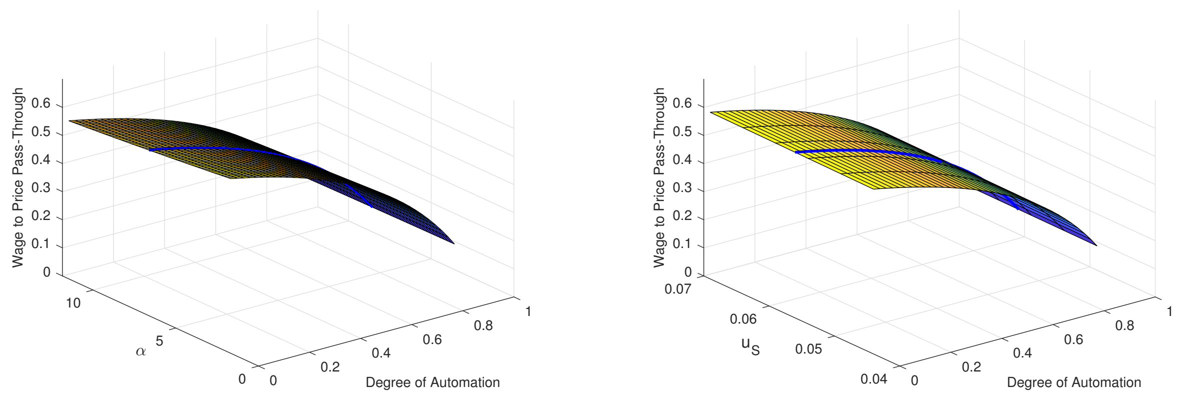

(a) Effect of α - Shape of productivity distribution

(b) Effect of u_S - Unemployment at steady state

4.3 Automation and Inflation Responses

We first look at the prices and wages responses after a monetary policy shock to gain understanding of the implications of higher level of automation for short-

[9]For more on the difference between in price and wage Phillips Curves in the past decades, see also Fitzgerald et al. (2023).

term dynamics (results for a consumption demand shock and a productivity shock generate similar implications and thus are shown in the Appendix). In this section we employ the richer model with endogenous labour participation detailed in the appendix (similar responses are obtained under the simpler model with fixed labour participation). We study two otherwise identical economies that differ only in the degree of automation resulted from robot-specific technological change. A one standard deviation of robots penetration across MSAs in the US implies a 200% increase in the ratio of robots to employees. We therefore set the initial share of automation to be 2% (an economy where robots are used in a significantly small share of production, denoted the *Low Automation* economy) and increase it by 200%, denoting it the *High Automation* economy, to compare the model results with our empirical estimates.

Figure 3 shows the results. We subject both economies to shocks that generate the same response on unemployment and verify the responses of price inflation, average wage, price to wage mark-up, number of firms and the automation cut-off point (γ^*) after the shock. With a similar exercise, Del Negro et al. (2020) shows that inflation responses to unemployment shocks have decreased in the past decades.

After a monetary policy shock, demand for final goods fall, pushing the demand for inputs down. Lower demand for inputs bring the value of operating either a robot intensive or a labour intensive firm down, reducing entry. Lower entry implies less vacancies are opened in the labour market, depressing wages. As the degree of automation is endogenous, there are two added elements altering the dynamics (these comprise the direct and the indirect *Automation Effects*). First, as wage decreases, some firms decide to use labour instead of robots, sustaining labour demand and wages. This direct channel of variation in automation moderates the wage and consequently the price responses after a monetary shock. We called this the *Moderation Effect of Automation.* Second, as using robots is always a choice firms can fall back on, robot adoption increases firms' bargaining power, dampening the responsiveness of wages to changes in the unemployment gap. We call this the *Wage Setting Effect of Automation.* Finally, at the initial steady state, a share of the production uses (and will continue to do so) robots. Therefore, a part of the adjustment process as a response to the shock occurs independently of the equilibrium changes in labour markets, decoupling the variation of unemployment to changes in aggregate demand. We denote this as the *Labour Share Effect*; note that as it will become clear below, this does not depend

Figure 3: Impulse Responses - Monetary Policy Shock

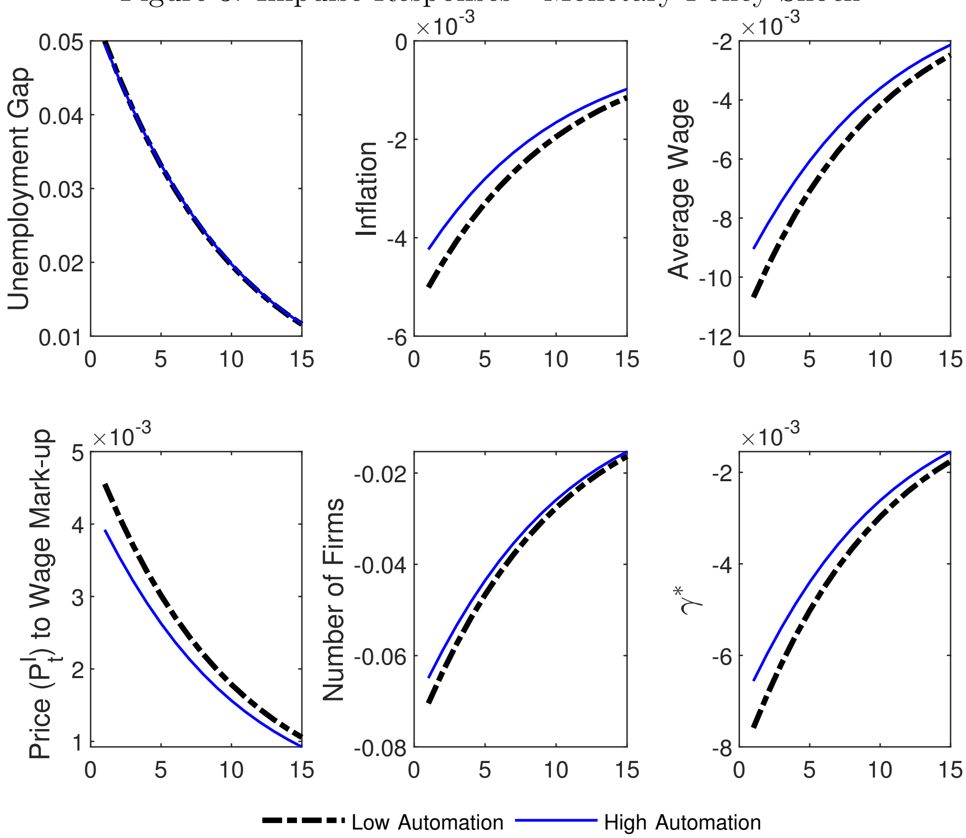

Note: Impulse responses of unemployment gap, price inflation, average wage, price to wage mark-up, number of firms and the automation cut-off point (γ^*) to monetary policy shocks designed to generate similar responses in unemployment in economies with high and low degrees of automation. For a variable x, impulse responses are $(x_t - x_S)/x_S$, where x_S is the value of x at the steady state.

on the endogenous changes in automation after the shock.

In the *low automation* economy, goods markets is almost entirely labour dominated and thus changes in firm entry lead to more volatile wages and prices given the unemployment level relative to the *high automation* economy in which a bigger share of the adjustment in goods' markets does not depend on unemployment movements. Thus, the *Labour Share Effect* implies greater price and wage response in the *low automation* economy. Second, automation is more responsive in a *low automation* economy, strengthening the *Moderation Effect*, potentially reducing the wage and price responses in the *low automation* economy. However, the *Wage Setting Effect*, which is stronger in the *high automation* economy, dominates (we look at this in more detail below). Consequently, we find that in the *low automation* economy, wages and consequently prices, are more sensitive and drop more significantly after a monetary shock. Therefore, confirming the results from the relationship between the degree of automation and the slope of the Phillips curve, high automation at steady state leads to lower inflation responses for a given change in unemployment.

In order to isolate these different channels through which automation affects inflation responses, we consider three alternative model specifications, aside from our core economy, which we denote as the *Directed Search* model. In the first, we continue using directed search in labour markets as in the core model but do not allow automation to endogenously change after the shock, thus the degree of automation (γ^*) is fixed in the short-term being unresponsive to monetary policy shocks. We denote this specification as *Directed Search - Fixed*. The second we assume firms and workers are randomly matched in labour markets and share the surplus of the match under Nash bargaining (see Pissarides (2000)). We name this model version, the *Random Search* model. Finally, we also consider the model under random search in which the automation level is fixed in the short-run, denoting it *Random Search - Fixed* model.

First we focus on the two cases where we fixed automation in the short-run (dotted lines with and without circles in Figure 4). Under these scenarios, both the moderation via directly changing automation and the wage setting effect boosted by the treat of robot adoption are no longer operational. Therefore, the only difference between the *low automation* and *high automation* economies is the size of the labour share in equilibrium. A lower labour share (*high automation* economy) implies wages and prices are less responsive to variation in unemployment and thus the Phillips curve is flatter and inflation responds less for a given unemployment response. Under fixed short-term automation, the labour markets have similar responses under random and directed search for each respective steady state (note that the key parameters that determine matching and surplus sharing are the same in both specifications).

If we allow automation to endogenously change after the shock then the labour market specification becomes crucial to determine the difference in wages and prices between the *low* and *high automation* economies. On the one hand, under directed search the firm that is at the margin between using robots or labour in production is setting the wage based on its own productivity level (γ_j), thus the probability of filling a vacancy, or the labour market conditions the firm faces, is a function of its own productivity level (or $q(\theta_{\gamma_j,t})$). The firm therefore internalizes the dependency between γ_j and the labour market it faces while measuring its value under robot adoption or opening a vacancy. The value while opening a vacancy of the firm at the margin under directed search (DS) is given by

$$V_{j,t}^{W,DS}(\gamma_j) = q(\theta_{\gamma_j,t})\left(P_{P,t}\gamma_j - W_{\gamma_j,t}\right) - \kappa_v - \kappa_f.$$

Wage setting (or selecting the labour sub-market and consequently the probability of filling a vacancy) and the automation decision interact, increasing the bargaining power of firms.

On the other hand, under random search the labour market conditions are the same for all firms; the probability of filling a vacancy for an specific firm j depends on the number of firms searching (so γ^*, via general equilibrium), but not γ_j. The automation decision of the firm at a margin does not alter its labour market outcome, and thus wage setting only indirectly affects the automation decisions via general equilibrium effects, decreasing the bargaining power of each firm. Under random search (RS) the value of opening a vacancy for the firm at the margin is given by

$$V_{j,t}^{W,RS}(\gamma_j) = q(\theta_{\gamma^*,t})\left(P_{P,t}\gamma_j - W_{\gamma^*,t}\right) - \kappa_v - \kappa_f.$$

When automation is endogenous and firms wage setting and automation decisions interact, the *Wage Setting Effect* is stronger, leading to a greater difference between the inflation response under the *low* and *high automation* economies (dashed dark line in Figure 4, depicting the inflation difference for our core economy). In contrast, when automation and wage setting only interact via indirect equilibrium conditions with labour markets characterized by random search and endogenous change in automation, the *Wage Setting Effect* weakens significantly. In this case the *Moderation Effect of Automation*, or the direct effect of automation changes in increasing labour demand after a negative monetary shock as the degree of automation falls in the short-run, dominates. The difference in inflation response for a given change in unemployment between the *low* and *high automation* economies in this case is the smallest (continuous line in Figure 4); the *Moderation Effect of Automation*, which is stronger in *low automation* economy, offsets the *Labour Share Effect*, which is higher in the *high automation* economy, bringing the inflation response for a given change in unemployment in these two economies closer together.

Figure 4: Difference of Inflation Responses after a Monetary Policy Shock for Low and High Automation Economies under Alternative Models

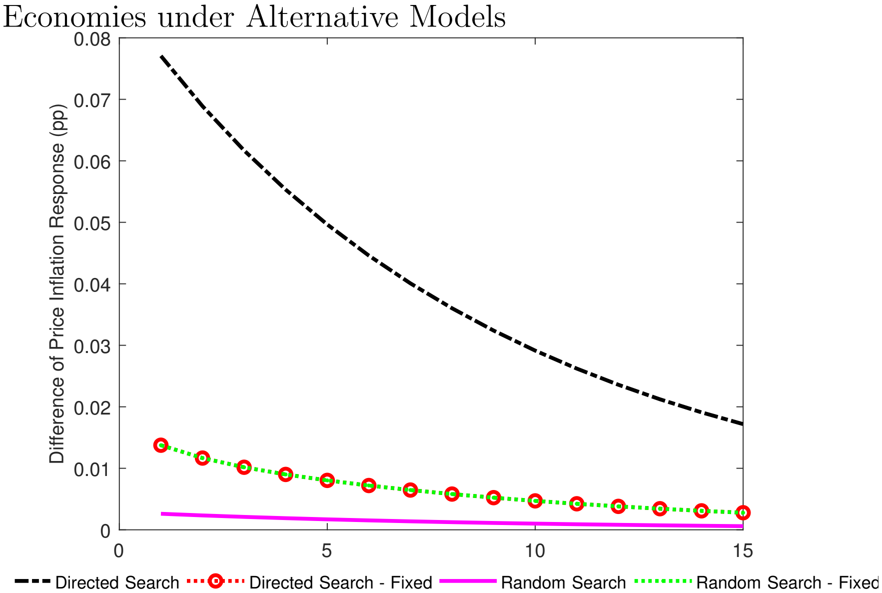

Note: Figure shows the absolute difference in the price inflation in the Low and High Automation economies for each model version, the *Directed Search* is our core model with endogenous labour participation, the *Random Search* model assumes random search instead of directed search, the *Directed Search - Fixed* is the core model assuming automation is fixed in the short-term and *Random Search - Fixed* is the random search model assuming automation is fixed in the short-term. The Inflation difference is defined as $100 * |$Impulse Response of *Low Automation* Economy - Impulse Response of *High Automation* Economy$|$.

So far we have kept q, the relative price of robots, constant in the short-term. As such, when automation is endogenously moving, it does so without restrictions. Next we consider the case in which ramping up automation becomes costly. As such we assume that the relative price of robots increases when the investment in robots (I_t) increases substantially from its steady state point (see Appendix for details). We then measure the responses of a small (one standard deviation) and a large (five standard deviations) positive demand shocks. Results are shown in figure 5. For a small positive demand shock inflation responds less in the *high automation* economy for a given change in unemployment. The treat of robot adoption dampens the wage responses leading to lower inflation increases for a given fall in unemployment. However, for larger demand shocks, the potential increase in robot investment would lead to rises in the price of robots making the treat of robot adoption no longer as effective in dampening wage responses. Moreover, the actual increase in the price of robots also lead to an increase in the cost of production of firms who select robots (low γ firms), adding an additional channel pushing inflation up, particular in the *high automation* economy, in which a greater share of output is produced by robots. As a result, after a large demand shock we no longer observe a difference in inflation responses for a given level of unemployment, the Phillips curve is no longer flatter in the *high automation* economy.

Figure 5: Costly Automation Changes

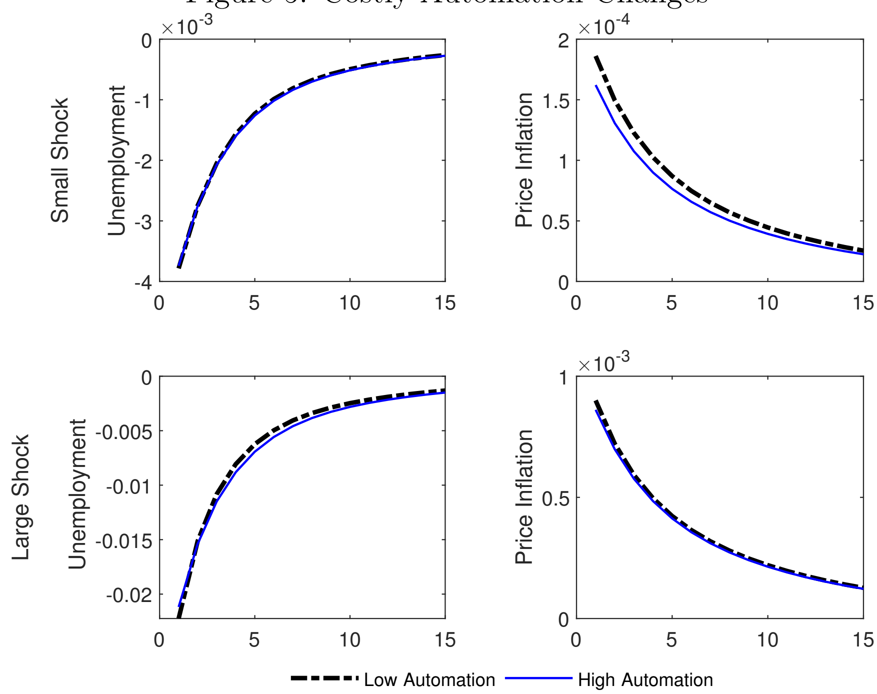

Note: Impulse responses of unemployment gap and price inflation to a small and large demand shocks designed to generate similar responses in unemployment in economies with high and low degrees of automation. For a variable x, impulse responses are $(x_t - x_S)/x_S$, where x_S is the value of x at the steady state.

In our final model exercise we simulate a series of demand shocks under the *low* and *high automation* economies and plot the series of unemployment and inflation deviations for each period, and their resulting regression lines (a crude measure of the Phillips curve). Details of the simulation exercise are provided in the Appendix. We find that the price inflation Phillips curve is 18% flatter in the *high automation*, relative to the *low automation* economy, while the wage inflation Phillips curve is 14% flatter in the *high automation*. Our results are in line with the empirical estimates discussed above. An one standard deviation increase in robot penetration across MSAs, which implies a 200% increase in the ratio of robots to employees, delivers a 17% flattening of the price inflation Phillips curve and a 9% flattening in the wage inflation Phillips curve. Starting from a low level of automation (2%), a 200% increase in the degree of automation in our model economy delivers a similar flattening in the price Phillips curve but a more substantial flattening of the wage Phillips curve.

Figure 6: Simulated Phillips Curves

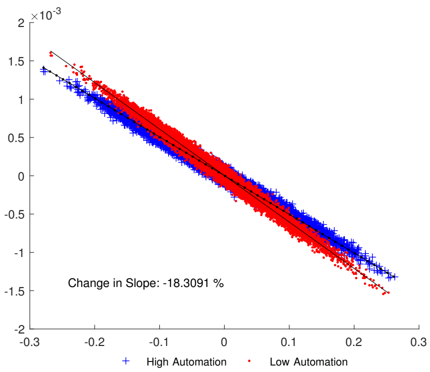

(a) Price Inflation Phillips Curve

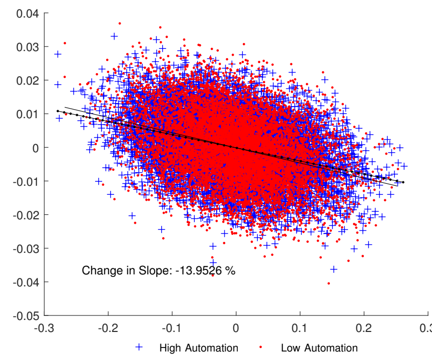

(b) Wage Inflation Phillips Curve

5 Conclusion

How does robot adoption influence inflation dynamics? We show empirically and theoretically that economies characterized by a higher degree of automation experience a lower sensitivity of inflation to movements in unemployment. As such, the substantial increase in the use of robots and other forms of automation in production processes experienced in most advanced economies in the last decades may be associated with the missing inflation observed during the same period, when inflationary pressures did not materialize despite the fluctuations observed in unemployment rates.

Employing a panel of nontradable goods inflation, wage inflation, unemployment rate and robot adoption at the U.S. metropolitan area (MSA) level, we uncover the causal relationship between automation and inflation by instrumenting the unemployment rate with local tradable demand spillovers and robot adoption with the robot installation across industries observed in the the five largest European economies. In our baseline results, the relationship between unemployment and inflation is affected by the degree of robot adoption in each MSA, indicating a significant role of automation in altering the relationship between inflation and unemployment. This effect is also economically relevant: an increase in robot adoption by one standard deviation reduces the sensitivity of prince inflation and wage inflation to unemployment by 17% and 9%, respectively. Overall, our empirical analysis uncovers three novel findings relating automation to inflation dynamics: robot adoption reduces (*i*) the sensitivity of price inflation to unemployment, (*ii*) the sensitivity of wage inflation to unemployment, and (*iii*) the pass-through from wages to prices.

We rationalize the empirical findings building a standard New Keynesian model with two key augmented features: search frictions in the labor market and the possibility of robot adoption. Firms, upon entry, draws an idiosyncratic efficiency in employing workers, and then decide to use either a labor technology (i.e., labor firms) or a machine technology (i.e., robot firms). Labor firms open costly vacancies at given posted wage, which are filled with a probability that depends on the labor market tightness. Workers observe posted wages and directly search for jobs of a given wage. Machine firms purchase a robot, at a given price determined by the level of robot-specific technological progress, from machine manufacturers, and produce with certainty. This setting defines an automation threshold, that is, a level of the efficiency in operating the labor technology that defines whether firms opt to produce using labor or robots. This threshold depends on the job filling probability and the levels of both wages and the price of robots. Therefore, in the model, the automation cut-off varies across steady states, as a function of the exogenous level of robot-specific technological change, and around the steady state upon the occurrence of a shock, as a function of the endogenous response of prices.

We then characterize the price Phillips curve and show that the degree of automation reduces its slope. The dampening effect of automation on the relationship between inflation and the unemployment gap is due to two main mechanisms. First, automation reduces the labor share in value added and thus, part of the aggregated demand adjustment is unrelated to changes in wage and unemployment dynamics. Second, the outside option of robot adoption negatively affects workers' bargaining power, dampening the responsiveness of wages to changes in the unemployment gap. Finally, we find that in an economy with a higher degree of automation, designed to reflect the changes observed in robot penetration across MSAs, the responsiveness of price inflation and wage inflation to unemployment changes after a demand shock is reduced by 18% and 14%, respectively. These changes are in line with the magnitude of the effects of automation on the price and wage Phillips curve estimated in our empirical evidence, that suggest a flattening of 17% and 9%, respectively.

In the aftermath of the Covid pandemic inflation and low rates of unemployment have been observed, indicating a sudden strengthening of the relationship between inflation and unemployment, or a reversal of the trend observe in the past decades. The key element in our model that explains the flattening of the Phillips curve in our setting is the effect of automation on the wage setting and the bargaining power of workers. If ramping up automation is costly and machine manufacturers face adjustment, then the threat that robots pose to workers' bargaining power crucially depend on the size of the shock realizations. When facing a small expansionary shock, firms can purchase additional machines without facing a sharp increase in robot prices, and thus gain an upper hand on wage negotiations. In this case, both the wage and price Phillips curves are flat. However, when the size of an expansionary shock is substantial, installing all the required robots to meet demand would be increasingly costly, forcing producers to continue to operate using labor. Consequently, the threat of robot adoption is no longer effective in curtailing workers' bargaining power, and wages highly react to changes in the unemployment gap. As such, the model would be consistent with an more substantial response of wages and prices after a large demand shock observed after Covid. Indeed, Autor et al. (2023) shows that the increases in wage post-Covid have been stronger for low educated and low income workers, whose wages have been compressed during the past decades. That is consistent with the implications of our analysis; workers who got displaced or faced the risk of displacement during the last decades, and experienced lower wage gains despite the unemployment rate, have recently observed an increase in their bargaining power and are finding new jobs at greater wages.

References

Acemoglu, D., C. Lelarge, and P. Restrepo (2020). Competing with robots: Firm-level evidence from France. *AEA Papers and Proceedings 110*, 383–88.

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

Acemoglu, D. and P. Restrepo (2020a). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy 128*(6), 2188–2244.

Acemoglu, D. and P. Restrepo (2020b). Unpacking skill bias: Automation and new tasks. *AEA Papers and Proceedings 110*, 356–61.

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

Aksoy, Y., H. S. Basso, R. Smith, and T. Grasl (2019). Demographic structure and macroeconomic trends. *American Economic Journal: Macroeconomics 11*(1), 193–222.

Autor, D., A. Dube, and A. McGrew (2023, March). The unexpected compression: Competition at work in the low wage labor market. Working Paper 31010, National Bureau of Economic Research.

Autor, D. H., D. Dorn, and G. H. Hanson (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review 103*(6), 2121–2168.

Ball, L. and S. Mazumder (2011). Inflation dynamics and the Great Recession. *Brookings Papers on Economic Activity*, 337–381.

Bartik, T. J. (1991). Who benefits from state and local economic development policies? *Mimeo*.

Basso, H. S. and J. F. Jimeno (2021). From secular stagnation to robocalypse? Implications of demographic and technological changes. *Journal of Monetary Economics 117*, 833–847.

Basso, H. S. and O. Rachedi (2021). The young, the old, and the government: Demographics and fiscal multipliers. *American Economic Journal: Macroeconomics 13*(4), 110–141.

Beraja, M., E. Hurst, and J. Ospina (2019). The aggregate implications of regional business cycles. *Econometrica 87*(6), 1789–1833.

Bergholt, D., F. Furlanetto, and E. Vaccaro-Grange (2023). Did monetary policy kill the Phillips curve? Some simple arithmetics. *Mimeo*.

Blanchard, O. (2016). The Phillips Curve: Back to the '60s? *American Economic Review 106*(5), 31–34.

Coibion, O. and Y. Gorodnichenko (2015). Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics 7*(1), 197–232.

Del Negro, M., M. Lenza, G. E. Primiceri, and A. Tambalotti (2020). What's up with the Phillips Curve? *Brookings Papers on Economic Activity Spring*, 301–357.

Fitzgerald, T., C. Jones, M. Kulish, and J. P. Nicolini (2023). Is there a stable relationship between unemployment and future inflation? *American Economic Journal: Macroeconomics*, forthcoming.

Forbes, K. (2019). Inflation dynamics: Dead, dormant, or determined abroad? *Brookings Papers on Economic Activity 2019*(2), 257–338.

Fornaro, L. and M. Wolf (2021). Monetary policy in the age of automation. *Mimeo*.

Galesi, A. and O. Rachedi (2019). Services deepening and the transmission of monetary policy. *Journal of the European Economic Association 17*(4), 1261–1293.

Gilchrist, S., R. Schoenle, J. Sim, and E. Zakrajšek (2017). Inflation dynamics during the financial crisis. *American Economic Review 107*(3), 785–823.

Gordon, R. J. (2013). The Phillips curve is alive and well: Inflation and the NAIRU during the slow recovery. *Mimeo*.

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

Harding, M., J. Lindé, and M. Trabandt (2022). Resolving the missing deflation puzzle. *Journal of Monetary Economics 126*, 15–34.

Hazell, J., J. Herreno, E. Nakamura, and J. Steinsson (2022). The slope of the Phillips Curve: Evidence from US states. *The Quarterly Journal of Economics 137*(3), 1299–1344.

Heise, S., F. Karahan, and A. Şahin (2022). The missing inflation puzzle: The role of the wage-price pass-through. *Journal of Money, Credit and Banking 54*(S1), 7–51.

Heise, S., F. Karahan, and A. Şahin (2023). Inflation strikes back: The return of wage to price pass-through. *Mimeo*.

Herreno, J. and M. Pedemonte (2022). The geographic effects of monetary policy shocks. *Mimeo*.

Höynck, C. (2020). Production networks and the flattening of the Phillips curve. *Mimeo*.

Leduc, S. and Z. Liu (2023). Automation, bargaining power, and labor market fluctuations. *Mimeo*.

McLeay, M. and S. Tenreyro (2020). Optimal inflation and the identification of the Phillips curve. *NBER Macroeconomics Annual 34*(1), 199–255.

Mian, A. and A. Sufi (2014). What explains the 2007–2009 drop in employment? *Econometrica 82*(6), 2197–2223.

Patterson, C. (2023). The matching multiplier and the amplification of recessions. *American Economic Review*, forthcoming.

Petrongolo, B. and C. A. Pissarides (2001). Looking into the black box: A survey of the matching function. *Journal of Economic literature 39*(2), 390–431.

Pissarides, C. A. (2000, December). *Equilibrium Unemployment Theory* (2nd Edition ed.), Volume 1 of *MIT Press Books*. The MIT Press.

Rubbo, E. (2023). Networks, Phillips curves, and monetary policy. *Econometrica*, forthcoming.

Siena, D. and R. Zago (2021). Job polarization and the flattening of the price Phillips curve. *Mimeo*.

Stansbury, A. and L. H. Summers (2020). The declining worker power hypothesis. *Brookings Papers on Economic Activity Spring*, 1–77.

Stock, J. H. and M. W. Watson (2020). Slack and cyclically sensitive inflation. *Journal of Money, Credit and Banking 52*(S2), 393–428.

A Empirical Evidence: Robustness

This section evaluates the robustness of our empirical findings as well as corroborates the validity of our identification strategy by reporting a comprehensive battery of checks. Specifically, we consider to what extent our findings keep holding when accounting for the role of potential alternative explanations for the decoupling of inflation and unemployment dynamics, and when including variables which could be highly correlated (across states and over time) 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 relate to heterogeneity in demographic characteristics across metropolitan areas. To address this set of variables, we merge our data with information from the Current Population Survey (CPS) of the U.S. Census Bureau, and we 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). To compute the latter, we follow Herreno and Pedemonte (2022) and combine the estimate of the MPC by demographic characteristics derived by Patterson (2023) with the share of each of this characteristic 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.

We then 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 Table A. Overall, we find that the role of automation is always highly statistically significant and rather constant across the different specifications. These results also suggest our baseline setting does not capture the relationship that

Table A.1: Robot Adoption and Inflation across MSAs - The Role of Demographics

	Dependent Variable: $\pi_{N,i,t}$							
	Young People	Old People	Female Labor Particip.	Black Labor Particip.	Asian Labor Particip.	Low Education	Labor Force Particip.	MPC
	IV (1)	IV (2)	IV (3)	IV (4)	IV (5)	IV (6)	IV (7)	IV (8)
$u_{i,t-1}$	-0.5942***	-0.6016***	-0.5997***	-0.5927***	-0.5889***	-0.6009***	-0.6039***	-0.6001***
	(0.1511)	(0.1500)	(0.1381)	(0.1483)	(0.1463)	(0.1491)	(0.1488)	(0.1498)
$u_{i,t-1} \times (m_{i,t-1} - \bar{m})$	0.0140***	0.0140***	0.0143***	0.0138***	0.0140***	0.0140***	0.0136***	0.0140***
	(0.0051)	(0.0051)	(0.0051)	(0.0050)	(0.0051)	(0.0051)	(0.0050)	(0.0051)
$u_{i,t-1} \times \big(VAR_{i,t-1} - V\bar{A}R\big)$	-0.0402	-0.0326	-0.1496	-0.0672	0.1181	-0.0699	0.1975**	-0.1412
	(0.0804)	(0.0685)	(0.1127)	(0.0585)	(0.1485)	(0.0858)	(0.0806)	(0.2669)
Year Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓
MSA Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓
N. Observations	2,270	2,270	2,270	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 $u_{i,t-1} \times \big(VAR_{i,t-1} - V\bar{A}R\big)$, 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-tradables good inflation rate, $\pi_{N,i,t}$, and all cases are estimated with IV methods, in which the unemployment rate is instrumented with a shift-share variable that captures tradeable 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 value of the robot-adoption variable, $m_{i,t-1}$, the lagged value of the confounding variable used in the interaction term, $VAR_{i,t-1}$, the relative price of non-tradable goods, $p_{N,i,t-1}$, 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 the female labor participation, Column (4) considers the role of the labor participation of black workers, Column (5) considers the role of the labor participation of asian workers, Column (6) considers the share of workers with low educational attainments, defined as those workers who have attended school up to grade 10, Column (7) considers the role of the labor participation of all workers, and Column (8) considers the role of workers marginal propensity to consume. Double-clustered standard errors are reported in brackets. *** and ** indicate statistical significance at the 1% and 5%, respectively.

automation has with the aging labor force (Acemoglu and Restrepo, 2022; Basso and Jimeno, 2021), 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 Herreno and Pedemonte (2022).

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 lead to a decline in both routine and manual occupations (Acemoglu and Restrepo, 2018, 2020a, 2020b), 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 (2021) shows 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 show that the effect of automation on inflation dynamics holds above and beyond that of job polarization, we merge our data with the information on occupations provided by the CPS, and the assignment of these occupations to manual, routine, and abstract, as well as their offshorable content, all of which come from Autor et al. (2013). We then 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. Again, we find that although the occupation offshorability also leads to a flattening of the price Phillips curve, the effect of automation on inflation dynamics holds even when explicitly controlling for the time-variation in the occupational structure across metropolitan areas.

Finally, the third set of potential alternative explanations relates to the key role that foreign import competition has had on the changes in inflation dynamics in the pre-Covid and the post-Covid periods (Forbes, 2019; Heise et al., 2022, 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, which are the two countries which have been providing the largest 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 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

Table A.2: Robot Adoption and Inflation across MSAs - The Role of Occupations

	Dependent Variable: $\pi_{N,i,t}$			
	Abstract Occupations	Routine Occupations	Manual Occupations	Offshorable Occupations
	IV (1)	IV (2)	IV (3)	IV (4)
$u_{i,t-1}$	-0.5888*** (0.1364)	-0.5842*** (0.1358)	-0.5921*** (0.1372)	-0.5928*** (0.1365)
$u_{i,t-1}\times$ $(m_{i,t-1}-\bar{m})$	0.0114*** (0.0044)	0.0125*** (0.0044)	0.0127*** (0.0044)	0.0124*** (0.0044)
$u_{i,t-1}\times$ $\left(VAR_{i,t-1}-V\bar{A}R\right)$	-0.0175 (0.0109)	0.0170 (0.0170)	0.0051 (0.0202)	0.0429* (0.0242)
Year Fixed Effects	✓	✓	✓	✓
MSA Fixed Effects	✓	✓	✓	✓
N. Observations	2,489	2,489	2,489	2,489

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-tradables good inflation rate, $\pi_{N,i,t}$, and all cases are estimated with IV methods, in which the unemployment rate is instrumented with a shift-share variable that captures tradeable 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 value of the robot-adoption variable, $m_{i,t-1}$, the lagged value of the confounding variable used in the interaction term, $VAR_{i,t-1}$, the relative price of non-tradable goods, $p_{N,i,t-1}$, 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. Double-clustered standard errors are reported in brackets. *** and ** indicate statistical significance at the 1% and 5%, respectively.

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 imports variable – either the imports from China, or the imports from Mexico, or the sum imports from the two countries – both as its lagged values and its interaction with the unemployment rate in Table A. We find that although the total imports did flatten the price Phillips curve, the effect of automation on inflation dynamics hold above and beyond the time-variation in import competition across metropolitan areas.

Table A.3: Robot Adoption and Inflation across MSAs - The Role of Import Competition

	Dependent Variable: $\pi_{N,i,t}$		
	Chinese Imports	Mexican Imports	Chinese & Mexican Imports
	IV	IV	IV
	(1)	(2)	(3)
$u_{i,t-1}$	-0.5687***	-0.7265***	-0.6056***
	(0.1399)	(0.2033)	(0.1549)
$u_{i,t-1}\times$ $(m_{i,t-1}-\bar{m})$	0.0063**	0.0105***	0.0077**
	(0.0032)	(0.0040)	(0.0044)
$u_{i,t-1}\times$ $\left(VAR_{i,t-1}-V\bar{A}R\right)$	0.0141	-0.8281	0.1812*
	(0.0675)	(0.5082)	(0.1011)
Year Fixed Effects	✓	✓	✓
MSA Fixed Effects	✓	✓	✓
N. Observations	3,526	3,526	3,526

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-tradables good inflation rate, $\pi_{N,i,t}$, and all cases are estimated with IV methods, in which the unemployment rate is instrumented with a shift-share variable that captures tradeable 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 value of the robot-adoption variable, $m_{i,t-1}$, the lagged value of the confounding variable used in the interaction term, $VAR_{i,t-1}$, the relative price of non-tradable goods, $p_{N,i,t-1}$, 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. Double-clustered standard errors are reported in brackets. *** and ** indicate statistical significance at the 1% and 5%, respectively.

B Further Details on the Model

Figure B.1: The Structure of the Model

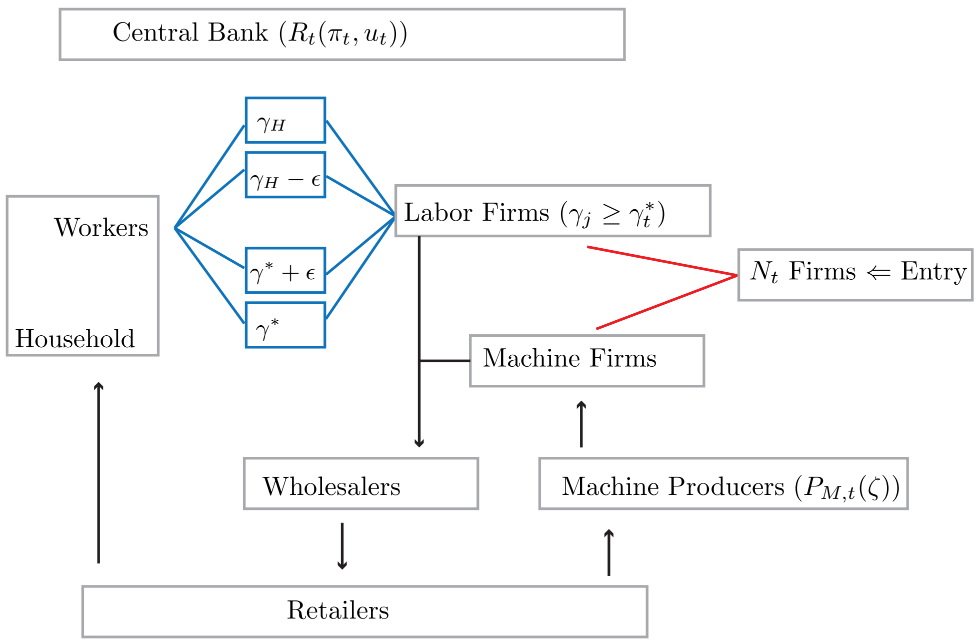

Note: This figure gives a graphical representation of the structure of the model economy.

C Additional Results

Figure C.2 shows the effect of increasing the degree of automation on the slope of the Phillips curve for this extension, for the benchmark calibration with $\zeta = 2.5$ and also the effect of the key structural parameters in this relationship. The main results of the model are unchanged under this extension.

Figure C.3 shows how the Wages to Price Pass-through changes as automation changes for additional structural parameters.

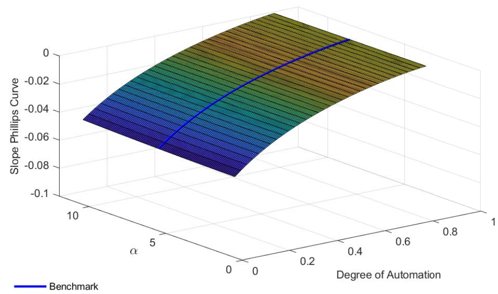

(a) Effect of α - Shape of productivity distribution

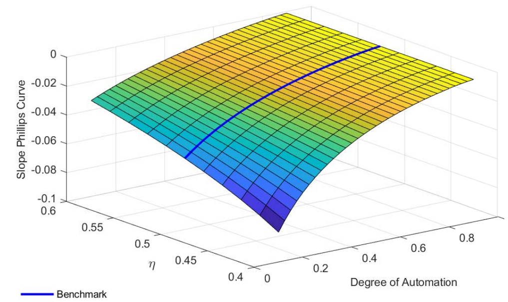

(b) Effect of η - Elasticity of vacancies to matches

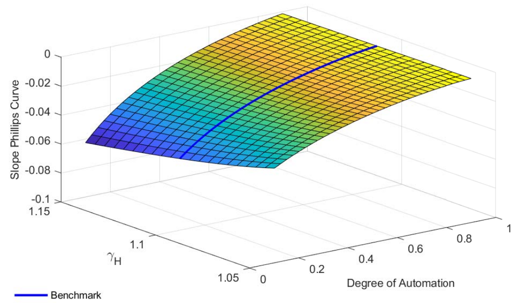

(c) Effect of γ_H - Upper bound on labour productivity

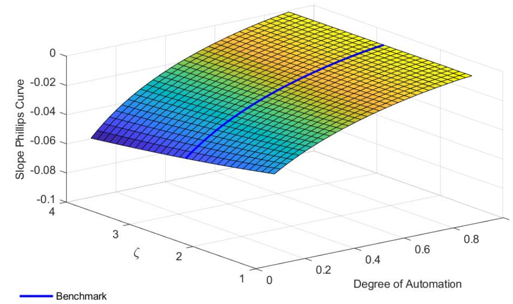

(d) Effect of ζ - Elasticity of Labour Participation

Figure C.2: Degree of automation and the slope of the Phillips curve with Endogenous Labour Participation

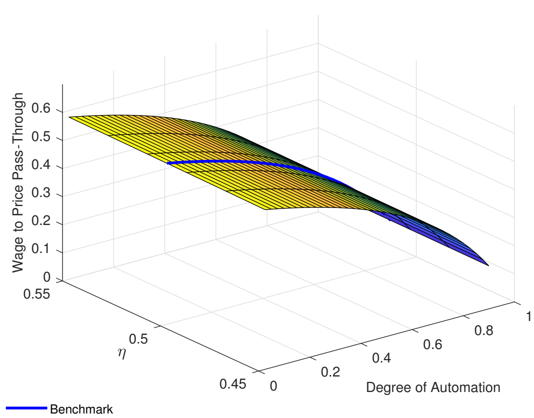

(a) Effect of η - Elasticity of vacancies to matches

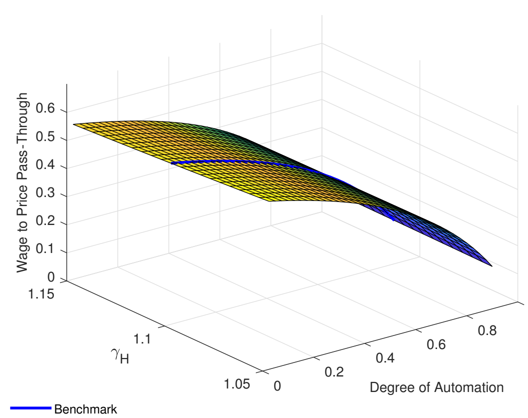

(b) Effect of γ_H - Upper bound on labour productivity

Figure C.3: Degree of automation and the wage to price pass-through