

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Automation, Bargaining Power, and Labor Market Fluctuations

Sylvain Leduc and Zheng Liu
Federal Reserve Bank of San Francisco

May 2022

Working Paper 2019-17

<https://www.frbsf.org/economic-research/publications/working-papers/2019/17/>

Suggested citation:

Leduc, Sylvain, Zheng Liu. 2022. “Automation, Bargaining Power, and Labor Market Fluctuations,” Federal Reserve Bank of San Francisco Working Paper 2019-17.

<https://doi.org/10.24148/wp2019-17>

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Board of Governors of the Federal Reserve System.

AUTOMATION, BARGAINING POWER, AND LABOR MARKET FLUCTUATIONS

SYLVAIN LEDUC AND ZHENG LIU

ABSTRACT. We argue that the threat of automation weakens workers' bargaining power in wage negotiations, dampening wage adjustments and amplifying unemployment fluctuations. We make this argument based on a quantitative business cycle model with labor market search frictions, generalized to incorporate automation decisions and estimated to fit U.S. time series. In the model, procyclical automation threats create real wage rigidity that amplify labor market fluctuations. We find that this automation mechanism is quantitatively important for explaining the large volatilities of unemployment and vacancies relative to that of real wages, a puzzling observation through the lens of standard business cycle models.

I. INTRODUCTION

Recent advances in robotics and artificial intelligence have raised concerns that automation is putting an increasing share of jobs at risk and reducing wages. There is a debate about whether automation reduces aggregate employment (Autor, 2015; Acemoglu and Restrepo, 2018, 2020). However, to the extent that automation is a labor-saving technology, the *threat* of automation might weaken workers' bargaining power and thus restrain wage increases, even if the technology is not actually adopted. The option to automate may become particularly attractive when firms face a tight labor market, in which hiring workers is difficult without substantial wage increases.

Date: May 18, 2022.

Key words and phrases. Automation, bargaining power, unemployment, wages, productivity, labor share, business cycles.

JEL classification: E24, J64, O33.

Leduc: Federal Reserve Bank of San Francisco. Email: Sylvain.Leduc@sf.frb.org. Liu: Federal Reserve Bank of San Francisco. Email: Zheng.Liu@sf.frb.org. The paper was previously circulated under the title "Robots or Workers? A Macro Analysis of Automation and Labor Markets." For helpful comments, we thank the Co-Editor Aysegül Sahin, three anonymous referees, Martin Eichenbaum, John Fernald, Chad Jones, Mike Keane, Emi Nakamura, Brent Neiman, Nicolas Petrosky-Nadeau, Jon Steinsson, Robert Townsend, and participants at the West Coast Search and Matching Workshop, HKUST Macro workshop, Hong Kong University, INSEAD, National University of Singapore, University of Autònoma Barcelona, University of New South Wales, Florida Macro Seminar Series, the 2020 Econometric Society World Congress, the 2020 NBER Summer Institute (Impulse and Propagation Mechanisms), and the 2021 SED meeting. We are grateful to Mollie Pepper, Lily Seitelman, and Remy Beauregard for excellent research assistance and to Anita Todd for editorial assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or of the Federal Reserve System.

In this paper, we argue that the increased threat of automation in business cycle expansions weakens workers' bargaining power and lowers the correlation between real wages and labor productivity. This creates endogenous real wage rigidity that helps explain the large fluctuations in unemployment and vacancies relative to real wages, a puzzling observation through the lens of standard business cycle models.

We formalize this argument based on a general equilibrium framework with labor market search frictions, generalized to incorporate automation decisions and estimated to fit U.S. time series. This theoretical framework allows us to study the interactions between firms' employment and automation decisions over the business cycle. A framework with search frictions in the labor market is also natural, since we are interested in the joint dynamics of unemployment, job vacancies, and real wages.

The model predicts that an increased threat of automation raises firms' reservation value in wage negotiations, strengthening their bargaining power and therefore dampening wage increases during economic expansions. In our model, the relative bargaining power is endogenous, and it fluctuates with the endogenous probability of automation over the business cycle. The automation channel creates two opposing forces on real wages. In an expansion, a tight labor market pushes real wages up. At the same time, the net benefit of automation and thus the probability of automation also rise, and the increased threat of automation restrains wage increases.

The automation channel also creates two opposing effects on employment. It has a direct job-displacing effect, because automation equipment (such as robots) can substitute for workers in production. On the other hand, automation has a job-creating effect, because the option to automate an unfilled job position raises the expected value of a job vacancy, boosting firms' incentive to create new vacancies and thus increasing the job finding rate and employment.

Our estimation suggests that the threat of automation is a quantitatively important source of real wage rigidity, and it helps account for the observed large volatilities of unemployment and vacancies. In addition, since automation raises labor productivity while depressing wages, it leads to countercyclical fluctuations in the labor share of income, as observed in the data (Ríos-Rull and Santaella-Llopis, 2010).

I.1. Model mechanism. Our theoretical framework features two significant departures from the standard Diamond-Mortensen-Pissarides (DMP) model. First, we introduce a fixed cost of vacancy creation (Fujita and Ramey, 2007; Leduc and Liu, 2020). A firm will choose to create a new vacancy if and only if the net present value of the vacancy is non-negative. Thus, in equilibrium, an unfilled vacancy has a positive present value. Second, we introduce endogenous automation decisions. A vacancy in our model can be interpreted as a continuum

bundle of tasks, which are ex ante identical, but a fraction of which can be automated depending on costs of automation. Only non-automated tasks (vacancies) are posted for hiring workers. Thus, in equilibrium, some automated tasks are performed by machines (hereafter, we refer to those machines as “robots,” but they should be thought of as labor-saving capital more generally) and some other tasks by workers (provided that a match is formed in the labor market). Since robots can substitute for workers in production, they are different from the physical capital in standard neoclassical production functions, where capital and labor are complementary inputs.¹

In line with Acemoglu and Restrepo (2018) and Zeira (1998), firms in our model first make a choice of technologies (adopting a robot or not); and then post those non-automated tasks (i.e., vacancies) for hiring workers.² To implement this approach, we assume that each firm observes an i.i.d. cost of automation in the beginning of each period and decides whether or not to automate an unfilled job vacancy carried over from the previous period. At this stage of the decisions, the tasks can be indexed by the realized automation costs. A task will be automated if and only if the cost of automation lies below a threshold determined by the net benefit of automation, in which case, the firm adopts a robot for production and takes the job vacancy offline. The probability of automation is thus the cumulative density of the realized automation costs evaluated at the automation threshold.

Our approach to modeling automation requires job vacancies to carry a positive value, which is an equilibrium outcome in our model since vacancy creation is costly. The option to automate an unfilled vacancy raises the vacancy value, providing an extra incentive for firms to create new vacancies. But an increase in the value of a vacancy would also boost the firm’s reservation value in wage negotiations, effectively weakening workers’ bargaining

¹Krusell et al. (2000) study a neoclassical model in which capital equipment complements skilled labor but substitutes for unskilled labor. He and Liu (2008) study a general equilibrium extension of the Krusell et al. (2000) model to incorporate endogenous skill accumulations. The relation between robots and workers in our model is analogous to the relation between equipment and unskilled labor in the model of Krusell et al. (2000).

²We focus on firms’ ability to use robots as a potential margin of adjustments in response to changes in labor market conditions over the business cycle. In an expansion, for example, firms face a tighter labor market in which job vacancies are harder to fill. This creates an incentive for adopting labor-saving technologies such as automation. In general, an automation process consists of an extensive margin, along which technological advancements lead to automation of some tasks previously performed by human workers; and an intensive margin where, for a given set of automatable tasks, robots can substitute for human workers depending on the relative input costs. We view automation along the extensive margin as occurring relatively infrequently over long time periods, rather than at the business cycle frequency. Since our primary interest here lies in the business cycle, we focus on the intensive margin of automation, which reacts to cyclical changes in the relative input cost.

power, depressing wage increases. Thus, the threat of automation creates a source of real wage rigidity and amplifies the fluctuations of unemployment and vacancies over the business cycle.

I.2. Model implications. To examine the quantitative importance of the automation mechanism, we estimate the model to fit quarterly U.S. time series data. These time series include unemployment, vacancies, real wage growth, and nonfarm business sector labor productivity growth, with a sample ranging from 1985:Q1 to 2018:Q4. Matching the observed fluctuations in labor productivity is an important disciplining device on the endogenous automation mechanism, especially because of the slowdown in productivity growth since the mid-2000s (Fernald, 2015).

We find that the threat of automation dampens wage increases in a business cycle boom. Since the net value of automation is procyclical, the probability of automation increases in expansions.³ By dampening wage changes, the automation threat amplifies the fluctuations in unemployment and vacancies. Increased automation in a boom also boosts aggregate productivity, further fueling the expansion. Since automation improves labor productivity while muting wage increases, it implies a countercyclical labor income share.⁴

The automation mechanism is quantitatively important. A counterfactual model with muted automation threat or one with a smaller share of intermediate goods produced by robots (relative to those produced by workers) generates a much smaller volatility of the vacancy-unemployment ratio (or the v/u ratio, a measure of labor market tightness) than that predicted by the benchmark model.

Our model implies that bargaining power is endogenous and depends on the threat of automation. Hagedorn and Manovskii (2008) argue that, in the standard DMP framework, an exogenous reduction in workers' bargaining weight or an increase in the workers' value of nonmarket activity (such as unemployment insurance) can also amplify fluctuations in unemployment and vacancies. Thus, we consider two such counterfactuals, both without the automation threat. We find that lowering workers' bargaining weight or raising the unemployment insurance (UI) can dampen wage adjustments and amplify unemployment fluctuations, in line with the findings in Hagedorn and Manovskii (2008). The Hagedorn and Manovskii (2008) solution, however, is subject to the critique by Costain and Reiter (2008),

³The procyclical automation investment implied by the model is consistent with macroeconomic evidence. For example, from 1985:Q1 to 2019:Q4, real investment in information processing equipment—a proxy for automation equipment—has a positive correlation with real GDP of 0.58.

⁴Karabarbounis and Neiman (2013) focus on the trend declines in the labor share since the mid-1970s for 59 countries. Their analysis attributes about half of the declines in the labor share to declines in the relative price of investment goods. We focus on the cyclical dynamics of the labor share, instead of its trend.

who pointed out that raising the level of UI in the standard DMP model would imply unrealistically large effects of UI policy changes on unemployment. Our model mechanism does not rely on high levels of UI benefits; instead, the automation channel in our model generates endogenous real wage rigidity that amplifies the fluctuations in unemployment and vacancies.

The automation threat is also robust to introducing heterogeneous worker skills. In a generalized version of the model, we assume that robots are substitutes for low-skilled workers but complements to high-skilled workers. Since robots and skilled workers are complementary, an increase in skilled wages in a business cycle expansion would raise the cost of using the automation technology, mitigating the incentive for firms to automate and resulting in greater fluctuations in unskilled wages. Compared to our benchmark model, the amplification effects through the automation channel are somewhat attenuated. However, the model continues to predict that the threat of automation depresses (low-skilled) wages and boosts labor productivity, leading to countercyclical labor share fluctuations.

Our framework can be generalized to study the implications of offshoring for labor market fluctuations. Elsby et al. (2013) highlight the importance of offshoring in accounting for the observed declines in the U.S. labor income share. To the extent that offshoring is a form of labor-saving technology, we conjecture that the threat of offshoring can exert similar influences on labor market outcomes as does the threat of automation. However, the relative importance of offshoring and automation may have changed over time, as we discuss in Section VI. Following the global trade collapse during the Great Recession of 2008-09, the ratio of U.S. imports to real GDP has been rising at a much slower pace than it did prior to the Great Recession. In comparison, automation investment has been increasing steadily throughout the past two decades. These time-series trends suggest that the automation channel may have become more important over time—especially after the Great Recession—for explaining the observed labor market fluctuations.

I.3. Evidence supporting the model’s mechanism. Our model’s mechanism suggests that procyclical automation probability dampens real wage fluctuations. By driving changes in bargaining powers, automation makes real wages less correlated with labor productivity. This creates endogenous real wage rigidity and amplifies fluctuations in unemployment and vacancies. These predictions are broadly in line with cross-sectional and time-series evidence, as we discuss in Section VI.

Using an unbalanced panel of industries at the two-digit level based on the North American Industry Classification System (NAICS) during the past two decades, we find that a decline in the relative price of computing equipment—a proxy for automation costs—is associated with

significantly larger increases in vacancies and the v/u ratio and declines in unemployment and real wages in industries that are more exposed to automation risks.

Our model's mechanism is also consistent with time-series data. The importance of automation (e.g., measured by the robot density or the ratio of real private investment in information processing equipment to real GDP) has increased steadily since the early 2000s. With the rising importance of automation, the model mechanism implies that the volatility of the v/u ratio should rise and the correlation between real wages and labor productivity should fall. These implications are consistent with the patterns in U.S. time-series data. In the data, the standard deviation of the v/u ratio initially declined from the 1980s to the 1990s, and then increased steadily in the 2000s and 2010s. The correlation between real wages and labor productivity initially increased from the 1980s to the 1990s and then declined substantially since the early 2000s.

II. RELATED LITERATURE

Our estimated general equilibrium model highlights the interactions between automation and labor market fluctuations at business cycle frequencies. Our study complements the empirical literature that typically focuses on longer-run implications of automation. For example, Graetz and Michaels (2018) examine the labor market impact of the cumulative changes in robot adoptions from 1993 to 2007 using a panel of industry-level data from 17 countries. They find that robot adoptions boost labor productivity and raise wages, although the positive effects on wages are much smaller than those on productivity. Arnoud (2018) also focuses on the long-run implications of automation. He examines occupation-level relations between the threat of automation and wage adjustments using data from the 2013 U.S. Current Population Survey and an index of automatability for different occupations developed by Frey and Osborne (2017). He finds that, controlling for observable characteristics, occupations that are more susceptible to automation have experienced lower wage growth. Dinlersoz and Wolf (2018) present plant-level evidence that more automated establishments in the U.S. manufacturing sector have had a smaller fraction of high-wage workers, higher labor productivity, and a smaller labor share in production. Acemoglu and Restrepo (2020) present evidence that, for U.S. commuter zones exposed to robots, the increase in the stock of industrial robots between 1990 and 2007 reduced the average employment-to-population ratio by 0.4 percentage points and average wages by 0.8 percent, relative to commuter zones with no robot exposure. Acemoglu and Restrepo (2021) document evidence that automation has reduced the relative wages of workers specialized in routine tasks and accounted for between 50 percent and 70 percent of changes in the U.S. wage structure since the early 1980s.

Automation in our model represents a labor-substituting technology, in line with Acemoglu and Restrepo (2018). There is substantial evidence that the steady progress in labor substituting technologies (such as computerization) has reduced the secular demand for workers with routine skills, contributing to job polarization in the U.S. labor market (Autor et al., 2003; Autor, 2015). Furthermore, job polarization can be linked to the jobless recoveries since the early 1990s, because most of the employment losses in routine occupations occur in recessions (Jaimovich and Siu, 2020). During the Great Recession and the subsequent recovery, employers in hard-hit areas raised the skill requirements when posting job vacancies, consistent with increased job destructions in routine occupations (Hershbein and Kahn, 2018). Although our benchmark model abstracts from skill heterogeneity, our extended model with heterogeneous worker skills captures the idea that robots are complementary with high-skill workers but are substitutes for low-skill workers. The key predictions of our benchmark model survives in this extended model. In particular, increases in automation raise labor productivity, dampen wage adjustments for low-skill workers, and reduce the labor share.

The automation threat in our paper is also related to the literature on changes in worker bargaining power. Krueger (2018) argues that declines in worker bargaining power can help explain why wage growth remained weak during periods when unemployment reached historically low levels in the United States. Stansbury and Summers (2020) also argue that forces that reduced worker power have contributed to sluggish wage growth and a declining labor share. They further argue that, while globalization and technological changes have played some part in reducing worker power, they are less important factors than declines in unionization and increases in shareholder power within firms. Our model mechanism suggests that, with the rapid rise in automation in the past two decades, labor-saving technologies might become an increasingly important factor for worker bargaining power and for labor market fluctuations.

Different from the existing literature that focuses on the secular impact of automation on labor markets, we focus on business cycle fluctuations. To our knowledge, our study provides the first quantitative general equilibrium evaluation of the interactions between automation and labor market fluctuations over the business cycle.

III. THE MODEL WITH LABOR MARKET FRICTIONS AND AUTOMATION

This section presents a dynamic stochastic general equilibrium (DSGE) model that generalizes the standard DMP model to incorporate endogenous decisions of automation.

To keep automation decisions tractable, we impose some assumptions on the timing of events. In the beginning of period t , a job separation shock δ_t is realized. Workers who

lose their jobs add to the stock of unemployment from the previous period, forming the pool of job seekers u_t . Firms carry over the stock of unfilled vacancies from the previous period, a fraction of which is automated. The stock of vacancies v_t available for hiring workers consists of the remaining vacancies after automation, the jobs separated in the beginning of the period, and newly created vacancies. Job seekers (u_t) randomly match with vacancies (v_t) in the labor market, with the number of new matches (m_t) determined by a matching technology. Final consumption goods are a composite of two types of intermediate goods—one produced with workers and the other with robots—with a constant elasticity of substitution (CES) between the two types. The unfilled vacancies and the pool of employed workers at the end of the period are carried over to the next period, and the same sequence of economic activities repeats in period $t + 1$.

Compared to the standard DMP model, our model introduces two new features. First, we replace the free-entry assumption in the DMP model with costly vacancy creation, as in Fujita and Ramey (2007) and Leduc and Liu (2020). Since creating a new vacancy incurs a fixed cost, a vacancy has a positive value even if it is not filled by a worker. The number of vacancies becomes a slow-moving state variable (instead of a jump variable as in the standard DMP framework), enabling our model to match the persistent vacancy dynamics in the data. Second, we introduce endogenous automation decisions. In the beginning of period t , each firm draws a fixed cost of automation. If the automation cost lies below a threshold value, then the firm adopts a robot and obtains the automation value, and the vacancy would be taken offline. If the automation cost exceeds the threshold, then the firm posts the vacancy for hiring a worker. Since robots can substitute for workers in production, they are different from the traditional capital input, which is typically complementary to labor input in the standard macro models.

III.1. Final goods producers. Final consumption goods are a CES composite of two types of intermediate goods: one type produced by workers and the other by robots. Specifically, aggregate output Y_t is given by

$$Y_t = \left[\alpha_n Y_{nt}^{\frac{\sigma-1}{\sigma}} + (1 - \alpha_n) Y_{at}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \quad (1)$$

where Y_{nt} and Y_{at} denote the intermediate goods produced by workers and by robots, respectively. The parameter σ measures the elasticity of substitution between the two types of intermediate inputs and the parameter α_n is a weight on worker-produced intermediate inputs in the final goods production.

The homogeneous final goods are traded in a perfectly competitive market. We use the final good as a numeraire, the price of which is normalized to one. Final goods producers take the relative prices of the intermediate goods as given and optimally choose Y_{nt} and Y_{at}

to maximize profits. With constant returns and free entry, equilibrium profits are zero. The optimal choices of intermediate inputs imply that

$$p_{nt} = \alpha_n \left(\frac{Y_t}{Y_{nt}} \right)^{\frac{1}{\sigma}}, \quad p_{at} = (1 - \alpha_n) \left(\frac{Y_t}{Y_{at}} \right)^{\frac{1}{\sigma}}, \quad (2)$$

where p_{nt} and p_{at} denote the relative price of intermediate goods produced by workers and by robots, respectively.

III.2. The Labor Market. In the beginning of period t , there are N_{t-1} existing job matches. A job separation shock displaces a fraction δ_t of those matches, such that the measure of unemployed job seekers is given by

$$u_t = 1 - (1 - \delta_t)N_{t-1}, \quad (3)$$

where we have assumed full labor force participation and normalized the size of the labor force to one.

The job separation shock δ_t follows the stationary stochastic process

$$\ln \delta_t = (1 - \rho_\delta) \ln \bar{\delta} + \rho_\delta \ln \delta_{t-1} + \varepsilon_{\delta t}, \quad (4)$$

where ρ_δ is the persistence parameter and the term $\varepsilon_{\delta t}$ is an i.i.d. normal process with a mean of zero and a standard deviation of σ_δ . The term $\bar{\delta}$ denotes the steady-state rate of job separation.

The stock of vacancies v_t consists of unfilled vacancies carried over from period $t - 1$ that are not automated, plus the separated employment matches and newly created vacancies. The law of motion for vacancies is given by

$$v_t = (1 - q_{t-1}^v)(1 - q_t^a)v_{t-1} + \delta_t N_{t-1} + \eta_t, \quad (5)$$

where q_{t-1}^v denotes the job filling rate in period $t - 1$, q_t^a denotes the automation probability in period t , and η_t denotes newly created vacancies (i.e., entry).

In the labor market, new job matches (denoted by m_t) are formed between job seekers and open vacancies based on the Cobb-Douglas matching function

$$m_t = \mu u_t^\alpha v_t^{1-\alpha}, \quad (6)$$

where μ is a scale parameter that measures matching efficiency and $\alpha \in (0, 1)$ is the elasticity of job matches with respect to the number of job seekers.

The flow of new job matches adds to the stock of employment, whereas job separations subtract from it. Aggregate employment evolves according to the law of motion

$$N_t = (1 - \delta_t)N_{t-1} + m_t. \quad (7)$$

At the end of period t , the searching workers who fail to find a match remain unemployed. Thus, unemployment is given by

$$U_t = u_t - m_t = 1 - N_t. \quad (8)$$

For convenience, we define the job finding probability q_t^u and the job filling probability q_t^v , respectively, as

$$q_t^u = \frac{m_t}{u_t}, \quad q_t^v = \frac{m_t}{v_t}.$$

III.3. The representative household. The representative household has the utility function

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t \Theta_t (\ln C_t - \chi N_t), \quad (9)$$

where $\mathbb{E}[\cdot]$ is an expectation operator and C_t denotes consumption. The parameter $\beta \in (0, 1)$ denotes the subjective discount factor, the parameter $\chi > 0$ denotes a weight on the disutility of working, and the term Θ_t denotes an exogenous shock to the subjective discount factor.

The discount factor shock $\theta_t \equiv \frac{\Theta_t}{\Theta_{t-1}}$ follows the stationary stochastic process

$$\ln \theta_t = \rho_\theta \ln \theta_{t-1} + \varepsilon_{\theta t}. \quad (10)$$

In this shock process, ρ_θ is the persistence parameter and the term $\varepsilon_{\theta t}$ is an i.i.d. normal process with a mean of zero and a standard deviation of σ_θ . Here, we have implicitly assumed that the mean value of θ is one.

The representative household chooses consumption C_t and savings B_t to maximize the utility function (9) subject to the sequence of budget constraints

$$C_t + \frac{B_t}{r_t} = B_{t-1} + w_t N_t + \phi(1 - N_t) + d_t - T_t, \quad \forall t \geq 0, \quad (11)$$

where r_t denotes the gross real interest rate, w_t denotes the real wage rate, d_t denotes the household's share of firm profits, and T_t denotes lump-sum taxes. The parameter ϕ measures the flow benefits of unemployment.

Denote by $V_t(B_{t-1}, N_{t-1})$ the value function for the representative household. The household's optimizing problem can be written in the recursive form

$$V_t(B_{t-1}, N_{t-1}) \equiv \max_{C_t, N_t, B_t} \ln C_t - \chi N_t + \beta \mathbb{E}_t \theta_{t+1} V_{t+1}(B_t, N_t), \quad (12)$$

subject to the budget constraint (11) and the employment law of motion (7).

Define the employment surplus (i.e., the value of employment relative to unemployment) as $S_t^H \equiv \frac{1}{\Lambda_t} \frac{\partial V_t(B_{t-1}, N_{t-1})}{\partial N_t}$, where Λ_t denotes the Lagrangian multiplier for the budget constraint (11). The optimizing decision for employment implies that the employment surplus

satisfies the Bellman equation

$$S_t^H = w_t - \phi - \frac{\chi}{\Lambda_t} + \mathbb{E}_t D_{t,t+1} (1 - q_{t+1}^u) (1 - \delta_{t+1}) S_{t+1}^H, \quad (13)$$

where $D_{t,t+1} \equiv \frac{\beta \theta_{t+1} \Lambda_{t+1}}{\Lambda_t}$ is the stochastic discount factor, which applies to both the household's intertemporal optimization and firms' decisions.⁵

The employment surplus has a straightforward economic interpretation. If the household adds a new worker in period t , then the current-period gains would be wage income net of the opportunity costs of working, including unemployment benefits and the disutility of working. The household also enjoys the continuation value of employment if the employment relation continues. Having an extra worker today adds to the employment pool tomorrow (provided that the employment relation survives job separation); however, adding a worker today would also reduce the pool of searching workers tomorrow, a fraction q_{t+1}^u of whom would be able to find jobs. Thus, the marginal effect of adding a new worker in period t on employment in period $t + 1$ is given by $(1 - q_{t+1}^u)(1 - \delta_{t+1})$, resulting in the effective continuation value of employment shown in the last term of Eq. (13).

It is straightforward to show that the household's optimizing consumption-savings decision leads to the intertemporal Euler equation

$$1 = \mathbb{E}_t D_{t,t+1} r_t. \quad (14)$$

III.4. The firms. A firm makes automation decisions in the beginning of the period t . Adopting a robot requires a fixed cost x in units of consumption goods. The fixed cost is drawn from the i.i.d. distribution $G(x)$. A firm chooses to adopt a robot if and only if the cost of automation is less than the benefit. For any given benefit of automation, there exists a threshold value x_t^* in the support of the distribution $G(x)$, such that automation occurs if and only if $x \leq x_t^*$.⁶ If the firm adopts a robot, then the vacancy will be taken offline and not available for hiring a worker. Thus, the automation threshold x_t^* depends on the value of automation (denoted by J_t^a) relative to the value of a vacancy (denoted by J_t^v). In particular, the threshold for the automation decision is given by

$$x_t^* = J_t^a - J_t^v. \quad (15)$$

The probability of automation is then given by the cumulative density of the automation costs evaluated at x_t^* . That is,

$$q_t^a = G(x_t^*). \quad (16)$$

⁵We provide detailed derivations in an online appendix available at https://www.frbsf.org/wp-content/uploads/sites/4/wp2019-17_appendix.pdf.

⁶Our approach to modeling robot adoption decisions here is similar in spirit to a McCall (1970) style search friction applied to the "robot labor market."

The flow of automated job positions adds to the stock of automated positions (denoted by A_t), which becomes obsolete at the rate $\rho^o \in [0, 1]$ in each period.⁷ Thus, A_t evolves according to the law of motion

$$A_t = (1 - \rho^o)A_{t-1} + q_t^a(1 - q_{t-1}^v)v_{t-1}, \quad (17)$$

where $q_t^a(1 - q_{t-1}^v)v_{t-1}$ is the number of newly automated job positions.

A robot produces $Z_t\zeta_t$ units of output, where Z_t denotes a neutral technology shock and ζ_t denotes an automation-specific shock. The neutral technology shock Z_t follows the stochastic process

$$\ln Z_t = (1 - \rho_z) \ln \bar{Z} + \rho_z \ln Z_{t-1} + \varepsilon_{zt}. \quad (18)$$

The parameter $\rho_z \in (-1, 1)$ measures the persistence of the technology shock. The term ε_{zt} is an i.i.d. normal process with a zero mean and a finite variance of σ_z^2 . The term \bar{Z} is the steady-state level of the technology shock.⁸ The automation-specific technology shock ζ_t follows a stochastic process that is independent of the neutral technology shock Z_t . In particular, ζ_t follows the stationary process

$$\ln \zeta_t = (1 - \rho_\zeta) \ln \bar{\zeta} + \rho_\zeta \ln \zeta_{t-1} + \varepsilon_{\zeta t}. \quad (19)$$

The parameter $\rho_\zeta \in (-1, 1)$ measures the persistence of the automation-specific technology shock. The term $\varepsilon_{\zeta t}$ is an i.i.d. normal process with a zero mean and a finite variance of σ_ζ^2 . The term $\bar{\zeta}$ is the steady-state level of the automation-specific technology shock.

To simplify the analysis and concentrate on the main mechanism, we assume that operating a robot incurs a flow fixed cost κ_a that captures the costs of facilities and the space for automated production.⁹ The value of automation satisfies the Bellman equation

$$J_t^a = p_{at}Z_t\zeta_t - \kappa_a + (1 - \rho^o)\mathbb{E}_t D_{t,t+1} J_{t+1}^a. \quad (20)$$

⁷If a vacancy is “filled” by a robot, it will be taken offline once and for all. Even if the robot later becomes obsolete, the vacated position does not return to the stock of vacancies. In the steady state, the number of obsolete robots equals the number of newly created vacancies (i.e., $\rho^o A = \eta$).

⁸The model can easily be extended to allow for trend growth.

⁹The flow cost of operating robots is not directly observable. A proxy for robot operating costs can be the wages of skilled labor hired to maintain robots, as we illustrate in a generalized version of our benchmark model with heterogeneous worker skills in Section VII.3. In the U.S. data, real wages for skilled workers (measured by the wages of new hires with a bachelor’s degree, deflated by the chained personal consumption expenditures price index) are roughly acyclical. For example, for the period from 2001 to 2019, the real skilled wage growth has a small positive correlation with the v/u ratio of about 0.16. An alternative proxy for robot operating costs is the quantity of energy usage. In the U.S. manufacturing sector for which we have data on energy usage, the growth rate of energy input has a correlation with the growth rate of manufacturing output of about 0.3 for the years from 1987 to 2019. Thus, our assumption of an acyclical robot operating cost is not at odds with the data.

If the automation cost exceeds the threshold x_t^* , then the firm chooses not to adopt a robot and instead chooses to post the vacancy in the labor market for hiring a worker. In addition, newly separated jobs and newly created vacancies add to the stock of vacancies for hiring workers. Following Leduc and Liu (2020), we assume that creating a new vacancy incurs an entry cost e in units of consumption goods. The entry cost is drawn from an i.i.d. distribution $F(e)$. A new vacancy is created if and only if the net value of entry is non-negative. The benefit of creating a new vacancy is the vacancy value J_t^v . Thus, the number of new vacancies η_t is given by the cumulative density of the entry costs evaluated at J_t^v . That is,

$$\eta_t = F(J_t^v). \quad (21)$$

Posting a vacancy incurs a per-period fixed cost κ (in units of final consumption goods). If the vacancy is filled (with the probability q_t^v), the firm obtains the employment value J_t^e . Otherwise, the firm carries over the unfilled vacancy to the next period, which will be automated with the probability q_{t+1}^a . If the vacancy is automated, then the firm obtains the automation value J_{t+1}^a net of the expected robot adoption costs; otherwise, the vacancy will remain open, and the firm receives the vacancy value J_{t+1}^v . Thus, the vacancy value satisfies the Bellman equation

$$J_t^v = -\kappa + q_t^v J_t^e + (1 - q_t^v) \mathbb{E}_t D_{t,t+1} \left\{ q_{t+1}^a J_{t+1}^a - \int_0^{x_t^*} x dG(x) + (1 - q_{t+1}^a) J_{t+1}^v \right\}. \quad (22)$$

If a firm successfully hires a worker, then it can produce Z_t units of intermediate goods. The value of employment satisfies the Bellman equation

$$J_t^e = p_{nt} Z_t - w_t + \mathbb{E}_t D_{t,t+1} \left\{ (1 - \delta_{t+1}) J_{t+1}^e + \delta_{t+1} J_{t+1}^v \right\}. \quad (23)$$

Hiring a worker generates a flow profit $p_{nt} Z_t - w_t$ in the current period. If the job is separated in the next period (with probability δ_{t+1}), then the firm receives the vacancy value J_{t+1}^v . Otherwise, the firm receives the continuation value of employment.

III.5. The Nash bargaining wage. When a job match is formed, the wage rate is determined through Nash bargaining. The bargaining wage optimally splits the joint surplus of a job match between the worker and the firm. The worker's employment surplus is given by S_t^H in Eq. (13). The firm's surplus is given by $J_t^e - J_t^v$. The possibility of automation affects the value of a vacancy and thus indirectly affects the firm's reservation value and its effective bargaining power.

Specifically, the Nash bargaining wage solves the problem

$$\max_{w_t} (S_t^H)^b (J_t^e - J_t^v)^{1-b}, \quad (24)$$

where $b \in (0, 1)$ represents the bargaining weight for workers.

Define the total surplus as

$$S_t \equiv J_t^e - J_t^v + S_t^H. \quad (25)$$

Then the bargaining solution implies that

$$J_t^e - J_t^v = (1 - b)S_t, \quad S_t^H = bS_t, \quad (26)$$

such that the firm's surplus is a constant fraction $1 - b$ of the total surplus S_t and the household's surplus is a fraction b of the total surplus.

The bargaining solution (26) and the expression for household surplus in equation (13) together imply that the Nash bargaining wage w_t^N satisfies the Bellman equation

$$\begin{aligned} \frac{b}{1-b}(J_t^e - J_t^v) &= w_t^N - \phi - \frac{\chi}{\Lambda_t} \\ &+ \mathbb{E}_t D_{t,t+1} (1 - q_{t+1}^u) (1 - \delta_{t+1}) \frac{b}{1-b} (J_{t+1}^e - J_{t+1}^v). \end{aligned} \quad (27)$$

We do not impose any real wage rigidity. Thus, the equilibrium real wage rate is just the Nash bargaining wage rate. That is, $w_t = w_t^N$.

III.6. Government policy. The government finances unemployment benefit payments ϕ for unemployed workers through lump-sum taxes. We assume that the government balances the budget in each period such that

$$\phi(1 - N_t) = T_t. \quad (28)$$

III.7. Search equilibrium. In a search equilibrium, the markets for bonds, final goods, and intermediate goods all clear. Since the aggregate bond supply is zero, the bond market-clearing condition implies that

$$B_t = 0. \quad (29)$$

Market clearing for intermediate goods implies that

$$Y_{nt} = Z_t N_t, \quad Y_{at} = Z_t \zeta_t A_t. \quad (30)$$

Final goods market clearing requires that consumption spending, vacancy posting costs, robot operation costs, robot adoption costs, and vacancy creation costs add up to aggregate production. This requirement yields the aggregate resource constraint

$$C_t + \kappa v_t + \kappa_a A_t + (1 - q_{t-1}^v) v_{t-1} \int_0^{x_t^*} x dG(x) + \int_0^{J_t^v} e dF(e) = Y_t. \quad (31)$$

IV. EMPIRICAL STRATEGIES

We solve the model by log-linearizing the equilibrium conditions around the deterministic steady state.¹⁰ We calibrate a subset of the parameters to match steady-state observations and the empirical literature. We estimate the remaining structural parameters and the shock processes to fit U.S. time-series data.

We focus on the parameterized distribution functions

$$F(e) = \left(\frac{e}{\bar{e}}\right)^{\eta_v}, \quad G(x) = \left(\frac{x}{\bar{x}}\right)^{\eta_a}, \quad (32)$$

where $\bar{e} > 0$ and $\bar{x} > 0$ are the scale parameters and $\eta_v > 0$ and $\eta_a > 0$ are the shape parameters of the distribution functions. We set $\eta_v = 1$ and $\eta_a = 1$, so that both the vacancy creation cost and the automation cost follow a uniform distribution.¹¹ We estimate the scale parameters \bar{e} and \bar{x} , along with the flow cost of automation κ_a and the shock processes by fitting the model to U.S. time series data.

IV.1. Steady-state equilibrium and parameter calibration. Table 1 shows the calibrated parameter values. We consider a quarterly model. We set $\beta = 0.99$, so that the model implies an annualized real interest rate of about 4 percent in the steady state. We set $\alpha = 0.5$ following the literature (Blanchard and Galí, 2010; Gertler and Trigari, 2009). In line with Hall and Milgrom (2008), we set $b = 0.5$ and $\phi = 0.25$. Based on the data from the Job Openings and Labor Turnover Survey (JOLTS), we calibrate the steady-state job separation rate to $\bar{\delta} = 0.10$ at the quarterly frequency. We set $\rho^o = 0.03$, so that robots depreciate at an average annual rate of 12 percent, in line with the estimated average life span of robots used by the International Federation of Robotics (IFR) for constructing their measure of the operation stocks of robots. We normalize the level of labor productivity to $\bar{Z} = 1$ and the average automation-specific productivity to $\bar{\zeta} = 1.5$, reflecting the relative average efficiency of the automation technology.

Based on the empirical estimates of Eden and Gaggli (2018), we calibrate the weight of worker-produced intermediate goods in final goods production to $\alpha_n = 0.535$. We calibrate the elasticity of substitution between intermediate goods produced by automation capital

¹⁰Details of the equilibrium conditions, the steady state, and the log-linearized system are presented in the online appendix.

¹¹Our assumption of the uniform distribution for the vacancy creation cost is in line with Fujita and Ramey (2007). We have estimated a version of the model in which we include the parameter η_a in the set of parameters to be estimated. We obtain a posterior estimate of η_a close to one and very similar estimates for the other parameters. For simplicity and for obtaining a closed-form solution for the steady-state equilibrium, we assume that $\eta_a = 1$ in our benchmark model.

TABLE 1. Calibrated parameters

	Parameter Description	value
β	Subjective discount factor	0.99
ϕ	Unemployment benefit	0.25
α	Elasticity of matching function	0.50
μ	Matching efficiency	0.6594
$\bar{\delta}$	Job separation rate	0.10
ρ^o	Automation obsolescence rate	0.03
κ	Vacancy posting cost	0.0409
b	Nash bargaining weight	0.50
η_v	Elasticity of vacancy creation cost	1
η_a	Elasticity of automation cost	1
χ	Disutility of working	0.3812
\bar{Z}	Mean value of neutral technology shock	1
$\bar{\zeta}$	Mean value of equipment-specific technology shock	1.5
σ	Elasticity of substitution between intermediate goods	3
α_n	Share of worker-produced intermediate goods	0.535

and by workers to $\sigma = 3$ based on the studies of Eden and Gaggl (2018) and Berg et al. (2018).¹²

We target a steady-state unemployment rate of $U = 0.0595$, corresponding to the average unemployment rate in our sample from 1985 to 2018. The steady-state employment is given by $N = 1 - U$, hiring rate by $m = \bar{\delta}N$, the number of job seekers by $u = 1 - (1 - \bar{\delta})N$, and the job finding rate by $q^u = \frac{m}{u}$. We target a steady-state job filling rate q^v of 0.71 per

¹²In the literature, there is substantial uncertainty about the value of the elasticity of substitution (EOS) between automation capital and labor. Graetz and Michaels (2018) and Acemoglu et al. (2020) assume that robots and labor are perfect substitutes in tasks that can be performed by robots, implying an infinite elasticity. Using the relative factor income shares and relative input quantities in aggregate data, Eden and Gaggl (2018) obtain an estimated EOS of about 8 between routine labor and a composite between non-routine labor and ICT capital, although they argue that an EOS between aggregate routine labor and ICT capital in the range between 2.14 and 3.27 is also plausible. Cheng et al. (2021) estimate the elasticity of substitution between labor and automation capital among automating firms using Chinese data and exploiting geographic and industry variations of government subsidies for automation under the “Made In China 2025” program. They obtain an estimated elasticity of about 3.8. Based on this literature, we choose a relatively conservative EOS value of 3. In a version of the model with perfect substitution between robots and workers, we obtained qualitatively similar results.

quarter, in line with the calibration of den Haan et al. (2000). The implied stock of vacancies is $v = \frac{m}{q^v}$. The scale of the matching efficiency is then given by $\mu = \frac{m}{u^\alpha v^{1-\alpha}}$.

Conditional on the estimated value of κ_a and the normalization of the average productivity levels (\bar{Z} and $\bar{\zeta}$), we solve for the steady-state automation value J^a from the Bellman equation (20). Then, given the estimated values of \bar{e} and \bar{x} (see below for estimation details), we use the vacancy creation condition (21), the automation adoption condition (15), and the law of motion for vacancies (5) to obtain the steady-state probability of automation, which is given by

$$q^a = \frac{J^a}{\bar{x} + \beta \bar{e}(1 - q^v)v}.$$

Given q^a and v , the law of motion for vacancies implies that the flow of new vacancies is given by $\eta = q^a(1 - q^v)v$. The vacancy value is then given by $J^v = \bar{e}\eta$. The stock of robots A can be solved from the law of motion (17), which reduces to $\rho^o A = \eta$ in the steady state. Thus, in the steady state, the newly created vacancies equal the flow of automated jobs that become obsolete. The law of motion for employment implies that, in the steady state, the flow of hiring equals the flow of separated jobs.

With A and N solved, we have $Y_n = \bar{Z}N$ and $Y_a = \bar{Z}\bar{\zeta}A$, and aggregate output is solved from Eq (1). We calibrate the vacancy posting cost to κ , such that the steady-state vacancy posting cost is 1 percent of aggregate output (i.e., $\kappa v = 0.01Y$).

Given J^v and J^a , we obtain the cutoff point for robot adoption $x^* = J^a - J^v$. The match value J^e can be solved from the Bellman equation for vacancies (22), and the equilibrium real wage rate can be obtained from the Bellman equation for employment (23). Steady-state consumption is solved from the resource constraint (31). We then infer the value of χ from the expression for bargaining surplus in Eq. (27).

IV.2. Estimation. We estimate the structural parameters \bar{e} , \bar{x} , and κ_a and the shock processes by fitting the DSGE model to quarterly U.S. time series.

IV.2.1. Data and measurement. We fit the model to four quarterly time series: the unemployment rate, the job vacancy rate, the growth rate of average labor productivity in the nonfarm business sector, and the growth rate of the real wage rate. The sample covers the period from 1985:Q1 to 2018:Q4. We provide details of the macro time-series data in Appendix A.

The unemployment rate in the data corresponds to the end-of-period unemployment rate in the model U_t . The measurement equation for unemployment is given by

$$U_t^{data} - \bar{U}^{data} = \hat{U}_t, \tag{33}$$

where U_t^{data} and \bar{U}^{data} denote the logged unemployment rate in the data and its sample mean, respectively, and \hat{U}_t denotes the log-deviations of the unemployment rate in the model from its steady-state value.

Similarly, the measurement equation for vacancies is given by

$$v_t^{data} - \bar{v}^{data} = \hat{v}_t, \quad (34)$$

where v_t^{data} and \bar{v}^{data} denote the logged vacancy rate and its sample mean in the data and \hat{v}_t denotes the log-deviations of the vacancy rate in the model from its steady-state value. Our vacancy series for the periods prior to 2001 is the vacancy rate constructed by Barnichon (2010) based on the Help Wanted Index. For the periods after 2001, we use the vacancy rate from JOLTS.

In the data, we measure labor productivity by real output per person in the nonfarm business sector. We use the demeaned quarterly log-growth rate of labor productivity and relate it to our model variable according to

$$\gamma_{p,t}^{data} - \bar{\gamma}_p^{data} = \hat{Y}_t - \hat{N}_t - (\hat{Y}_{t-1} - \hat{N}_{t-1}), \quad (35)$$

where $\gamma_{p,t}^{data}$ denotes the log growth rate of labor productivity in the data, $\bar{\gamma}_p^{data}$ denotes the sample average of labor productivity growth, and \hat{Y}_t and \hat{N}_t denote the log-deviations of aggregate output and employment from their steady-state levels in our model.

We measure the real wage rate in the data by real compensation per worker in the nonfarm business sector. We relate the observed real wage growth to the model variables by the measurement equation

$$\gamma_{w,t}^{data} - \bar{\gamma}_w^{data} = \hat{w}_t - \hat{w}_{t-1}, \quad (36)$$

where $\gamma_{w,t}^{data}$ denotes the log growth rate of the real wage rate in the data, $\bar{\gamma}_w^{data}$ denotes the sample average of real wage growth, and \hat{w}_t denotes the log-deviations of real wages from the steady-state level in the model.

IV.2.2. Prior distributions and posterior estimates. The prior and posterior distributions of the estimated parameters from our benchmark model are displayed in Table 2.

The priors for the structural parameters \bar{e} and \bar{x} are drawn from the gamma distribution. We assume that the prior mean of each of these three parameters is 5, with a standard deviation of 1. The priors of the flow cost of robot operation κ_a are drawn from the beta distribution, with a mean of 0.9 and a standard deviation of 1. The priors of the persistence parameter of each shock are drawn from the beta distribution with a mean of 0.8 and a standard deviation of 0.1. The priors of the volatility parameter of each shock are drawn from an inverse gamma distribution with a mean of 0.01 and a standard deviation of 0.1.

TABLE 2. Estimated parameters

Parameter description	Priors		Posterior		
	Type	[mean, std]	Mean	5%	95%
\bar{e} scale for vacancy creation cost	G	[5, 1]	3.0703	2.2635	3.9291
\bar{x} scale for robot adoption cost	G	[5, 1]	4.9483	3.4679	6.4300
κ_a flow cost of automation	B	[0.9, 1]	0.9788	0.9568	0.9999
ρ_z AR(1) of neutral technology shock	B	[0.8, 0.1]	0.9830	0.9696	0.9967
ρ_θ AR(1) of discount factor shock	B	[0.8, 0.1]	0.9760	0.9593	0.9930
ρ_δ AR(1) of separation shock	B	[0.8, 0.1]	0.9272	0.8838	0.9651
ρ_ζ AR(1) of automation-specific shock	B	[0.8, 0.1]	0.8924	0.8553	0.9300
σ_z std of tech shock	IG	[0.01, 0.1]	0.0168	0.0152	0.0188
σ_θ std of discount factor shock	IG	[0.01, 0.1]	0.0469	0.0366	0.0578
σ_δ std of separation shock	IG	[0.01, 0.1]	0.0502	0.0446	0.0548
σ_ζ std of automation-specific shock	IG	[0.01, 0.1]	0.0650	0.0566	0.0739
Log data density	1209.16				

Note: This table shows our benchmark estimation results. For the prior distribution types, we use G to denote the gamma distribution, B the beta distribution, and IG the inverse gamma distribution.

The posterior estimates and the 90 percent probability intervals for the posterior distributions are displayed in the last three columns of Table 2. The posterior mean estimate of the vacancy creation cost parameter is $\bar{e} = 3.07$. The posterior mean estimate of the automation cost parameter is $\bar{x} = 4.95$. The posterior mean estimate of the flow cost of robot operation is $\kappa_a = 0.98$. These parameters imply a steady-state automation probability of $q^a = 0.096$ in our quarterly model, or about 38 percent at the annual frequency, which lies in the range of the empirical estimates of automation risks (Frey and Osborne, 2017; Nedelkoska and Quintini, 2018).¹³ The 90 percent probability intervals indicate that the posterior estimates are significantly different from the priors, suggesting that the data are informative about these structural parameters.

¹³In an important study, Zolas et al. (2020) show that the automation probability (or robot adoption rate) is heterogeneous across firms, with adoptions heavily skewed toward large firms. Since our model abstracts from firm heterogeneity, it is not designed to fit the firm-level evidence. However, our model's implied automation probability lies in the range of the estimates obtained by Zolas et al. (2020) for firms with 500 or more employees (see their Figure 8).

The posterior estimation suggests that the shocks to both neutral technology and the discount factor are highly persistent, whereas the automation-specific shock is less persistent but more volatile. The 90 percent probability intervals suggest that the data are also informative for these shock processes.

Although we do not fit our quarterly model to directly match robot data (which are available at the annual frequency only), the estimated model can generate second moments of automation investment that are in line with broader measures of automation investment in the data. Specifically, in our model, automation investment corresponds to the flow of newly adopted robots, which is given by

$$I_{at} \equiv q_t^a(1 - q_{t-1}^v)v_{t-1}. \quad (37)$$

Under the estimated parameters and shocks, the model implies that the unconditional volatility of automation investment relative to that of aggregate output is about 4.71 and the correlation of automation investment with aggregate output is about 0.62. In comparison, in the quarterly U.S. data from 1985 to 2018, the year-over-year growth of real investment in information processing equipment—a proxy for automation investment—has a volatility of about 5.04 times that of real GDP growth and a correlation with real GDP growth of about 0.72.

V. ECONOMIC IMPLICATIONS

Based on the calibrated and estimated parameters, we examine the model’s transmission mechanism and its quantitative performance for explaining labor market dynamics.

V.1. Impulse responses. To illustrate the mechanism through which the threat of automation drives labor market dynamics, we calculate the impulse responses of the key variables following each shock.

Figure 1 shows the impulse responses of several key macro variables to a positive neutral technology shock in the benchmark model. The shock leads to persistent declines in unemployment and persistent increases in vacancies and hiring. The shock also raises the value of automation. Under our parameters, the value of automation rises more than does the value of a vacancy, raising the net value of robot adoption and therefore leading to an increase in the automation probability in the short run. Over time, the increase in the value of vacancies dominate that in the value of automation, leading to a decline in the automation probability about six quarters after the shock.

Increased automation following the neutral technology shock also raises labor productivity, reinforcing the initial expansionary impact of the shock. However, the increase in vacancy value also strengthens the firm’s bargaining power in wage negotiations, dampening




FIGURE 1. Impulse responses to a positive neutral technology shock in the benchmark model.

the responses of real wages. With muted wage responses and persistent increases in labor productivity, the shock leads to a persistent decline in the labor income share.

Figure 2 shows the impulse responses to a positive discount factor shock. The shock raises the present values of a job match and an open vacancy. Thus, it generates a persistent boom in employment, vacancies, and hiring. Similar to the neutral technology shock, the discount factor shock also raises the net value of automation (i.e., the difference between the value of a robot and the value of a vacancy), increasing the automation probability. Increased automation also boosts labor productivity, further fueling the boom. However, as the threat of automation rises, the workers' bargaining power weakens, leading to a modest short-run decline in the real wage. By raising productivity and reducing the real wage rate, the discount factor shock generates a persistent decline in the labor share.

Figure 3 shows the impulse responses to a positive automation-specific shock. The shock directly raises the value of automation. Since the option of automation boosts the value of vacancies, the increase in automation leads to more vacancy creation. With more job openings, the job finding rate increases, raising hiring and reducing unemployment. The increased automation in intermediate goods production also raises labor productivity. However, the shock raises the threat of automation, weakening workers' bargaining power, such

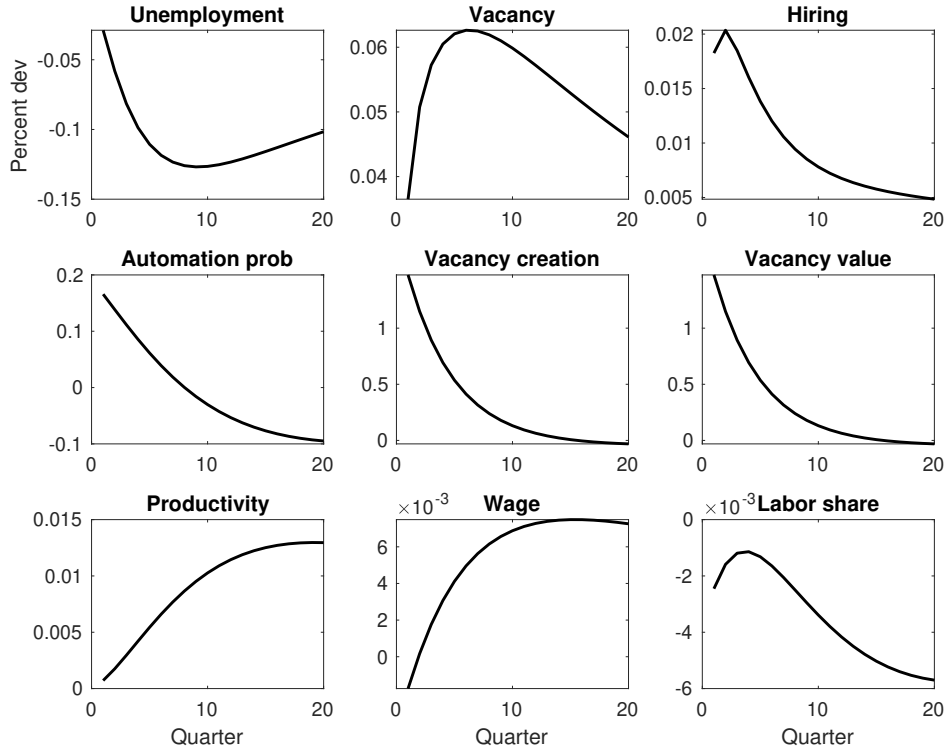


FIGURE 2. Impulse responses to a positive discount factor shock in the benchmark model.

that the rise in labor productivity does not translate fully into a rise in wages. Thus, the labor share declines persistently.¹⁴

V.2. Automation vs. other amplification mechanisms. Our model suggests that the automation threat effectively weakens workers’ bargaining power and mutes wage changes, therefore amplifying labor market fluctuations. The literature has studied other amplification mechanisms in the standard DMP framework without automation threats. For example, Hagedorn and Manovskii (2008) argue that, in the standard DMP framework, reducing workers’ bargaining weight or raising the workers’ value of nonmarket activity (such as unemployment insurance, or UI) can amplify fluctuations in unemployment and vacancies.

To evaluate the quantitative importance of the automation threat channel relative to these alternative amplification mechanisms, we study a counterfactual specification without the automation threat (labeled “no automation threat”), which is a version of our benchmark model with the automation probability q_t^a held constant at its steady-state level. We consider

¹⁴We report the impulse responses to a job separation shock in the online appendix. As we discuss there, a job separation shock raises both unemployment and vacancies. Consistent with Shimer (2005), this counterfactual positive correlation between unemployment and vacancies renders the job separation shock unimportant for driving labor market dynamics.

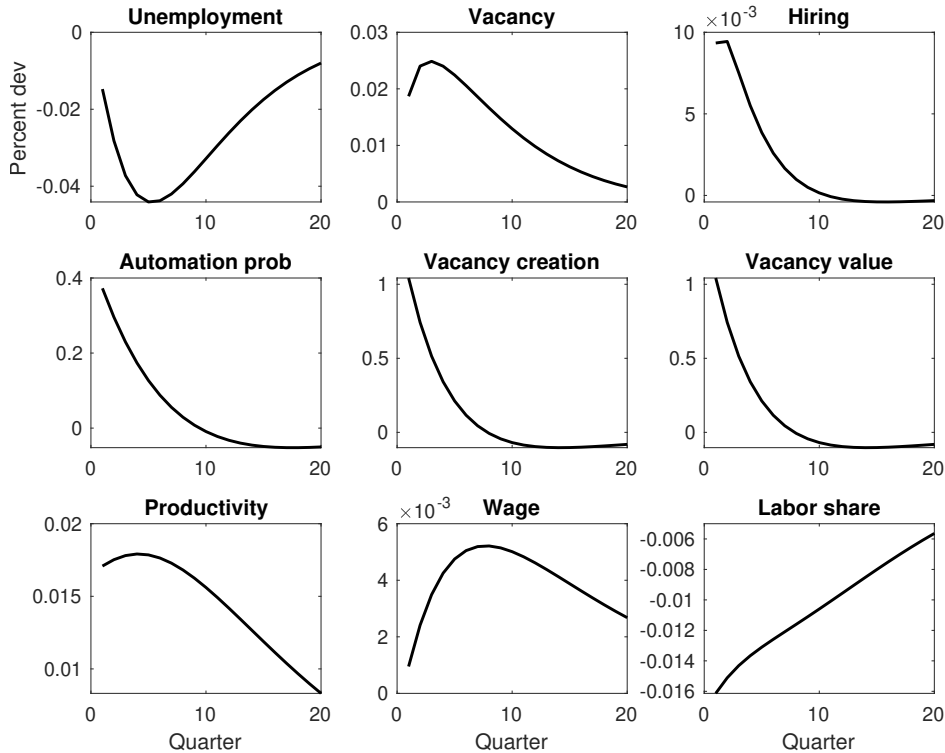


FIGURE 3. Impulse responses to a positive automation-specific shock in the benchmark model.

two variations of the “no automation threat” specification, one with a higher UI benefit (raising ϕ from 0.25 to 0.4) and the other with a lower worker bargaining weight (reducing b from 0.5 to 0.3).

Figure 4 displays the impulse responses to a positive discount factor shock in the benchmark model (black solid lines), the counterfactual with no automation threats (blue dashed lines), and the counterfactual with no automation threats and with a higher UI benefit (red dotted-dashed lines).¹⁵ Absent automation threats, the responses of unemployment and vacancies would be more muted than in the benchmark. Raising UI benefits in this case would increase workers’ outside option and amplify the responses of the labor market variables (Hagedorn and Manovskii, 2008). A key difference from the benchmark model lies in the responses of wages and the labor share. Working through the automation threat, the baseline model generates a short-run decline in wages and persistent declines in the labor share. Without the automation threat, however, the counterfactual with higher UI benefits generates less downward pressures on the real wage, leading to a short-run increase in the labor share.

¹⁵The impulse responses to a neutral technology shock in these counterfactual models display similar patterns, as we show in the online appendix.

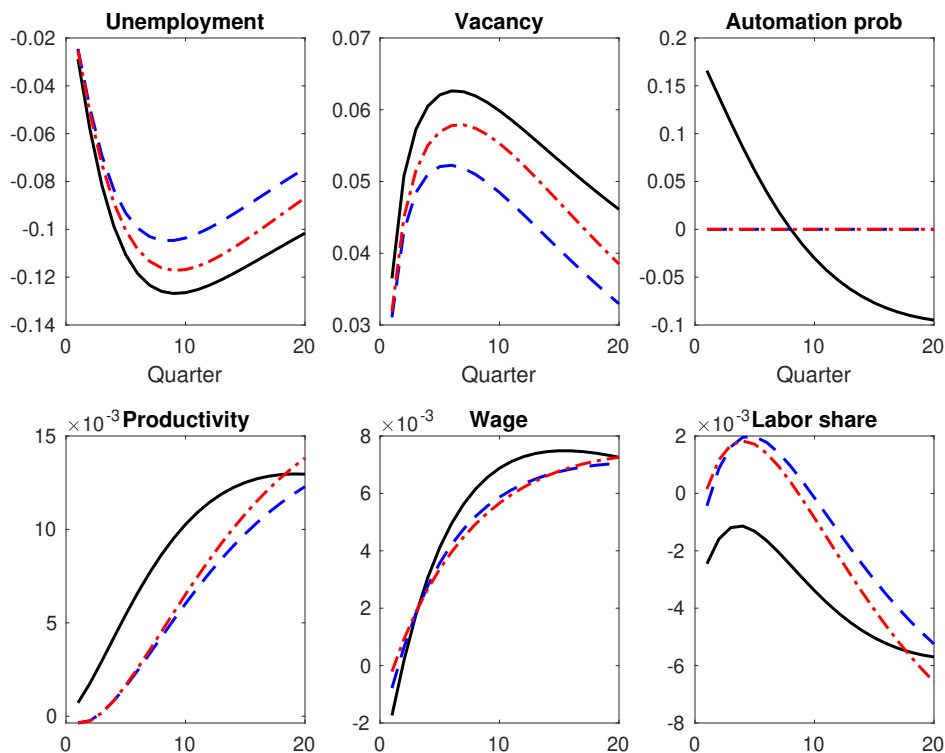


FIGURE 4. Impulse responses to a positive discount factor shock in the benchmark model (black solid lines), the counterfactual with no automation threats (blue dashed lines), and the counterfactual with no automation threats and high unemployment insurance (UI) benefits (red dot-dash lines).

In the counterfactual case without automation threats, mechanically lowering workers' bargaining weight can also dampen wage adjustments and therefore amplify the responses of unemployment and vacancies, consistent with Hagedorn and Manovskii (2008). Figure 5 shows the impulse responses to a positive discount factor shock in the benchmark model (black solid lines), the counterfactual with no automation threats (blue dashed lines), and the no-automation counterfactual with a lower worker bargaining weight (red dotted-dashed lines). Similar to the case with higher UI benefits, the real wage rate does not fall as much as in the benchmark model, such that the labor share rises in the short run, in contrast to the persistent declines under the threat of automation. These impulse responses suggest that the automation channel is an important mechanism for amplifying labor market fluctuations and generating a countercyclical labor income share.

To further illustrate the quantitative importance of the automation mechanism for labor market fluctuations, we compare the predictions from a few counterfactual models without automation threats for the volatility of labor market tightness (i.e., the v/u ratio), the correlation between real wage growth and labor productivity growth, and the volatility of

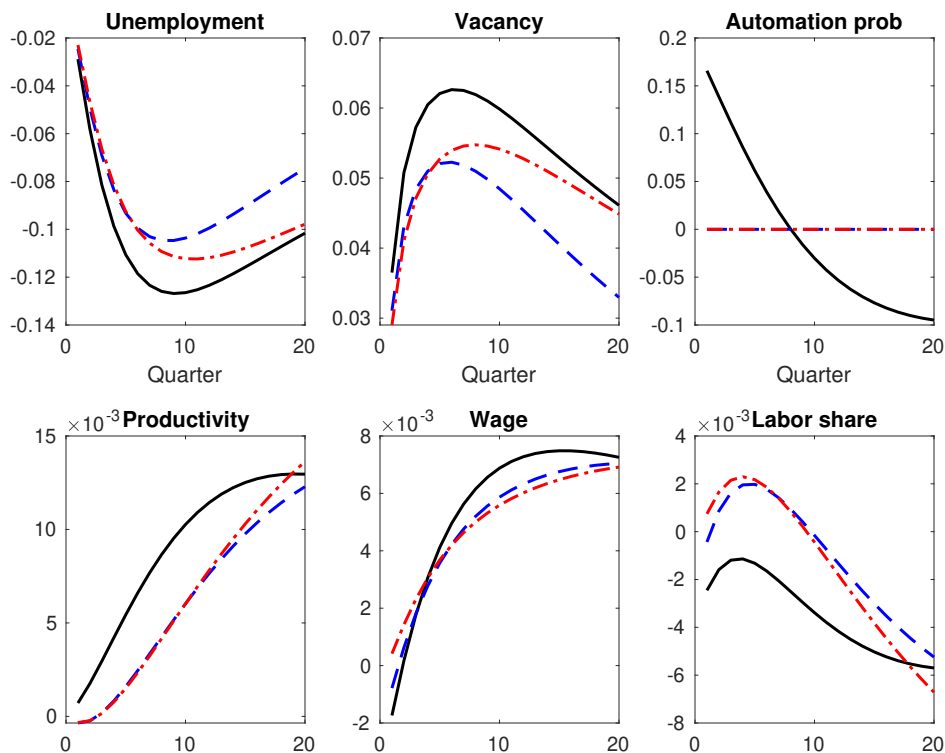


FIGURE 5. Impulse responses to a positive discount factor shock in the benchmark model (black solid lines), the counterfactual with no automation threats (blue dashed lines), and the counterfactual with no automation threats and a low worker bargaining weight (red dot-dash lines).

the v/u ratio relative to that of real wages (i.e., the volatility ratio) with the corresponding moments in the benchmark model. These moments in the estimated benchmark model are the same as in the actual data because the model is estimated to fit these time series. For ease of comparison, we normalize each of these labor market moments in the benchmark model to one and we calculate the corresponding moments in the counterfactual models relative to those in the benchmark. Table 3 displays these moments, all expressed relative to the corresponding statistics in the benchmark model.

The counterfactual model with no automation threats (row (1) of Table 3) predicts a lower volatility of the v/u ratio (about 31 percent lower), a higher correlation between real wage growth and labor productivity growth (about 10 percent higher), and a lower volatility of the v/u ratio relative to that of real wages (about 36 percent lower) than those in the benchmark model. These results reflect the Shimer (2005) volatility puzzle facing the standard DMP model.

The table also shows that, absent the automation threat, raising the UI benefits (row (2)) or reducing the worker bargaining weight (row (3)) can amplify the fluctuations in the

TABLE 3. Quantitative importance of automation threat

Counterfactual model	(a) $std(\ln(v/U))$	(b) $corr(\Delta \ln(w), \Delta \ln(p))$	(c) $\frac{std(\ln(v/U))}{std(\ln(w))}$
(1) No automation threat	0.69	1.10	0.64
(2) No automation threat and high UI benefits	0.96	0.81	0.82
(3) No automation threat and low worker bargaining weight	0.96	1.24	0.81

Note: The rows in the table correspond to the alternative models: (1) the counterfactual with the automation probability held constant at the steady-state level (“No automation threat”); (2) the model with no automation threat and a higher unemployment insurance (UI) benefit (with ϕ raised from 0.25 to 0.4); and (3) the model with no automation threat and a lower worker bargaining weight in wage negotiations (with b reduced from 0.5 to 0.3). The columns reports (a) the standard deviations of the v/u ratio, (b) the correlation between real wage growth and labor productivity growth, and (c) the ratio of the standard deviation of the $v-u$ ratio to that of real wages, all expressed relative to the corresponding statistics in the benchmark model.

v/u ratio, bringing its volatility closer to that in the benchmark model (from 69 percent to about 96 percent). The volatility ratio, however, is about 20 percent lower than that in the benchmark, because real wages fluctuate more than those in the benchmark. Overall, these counterfactual exercises illustrate that the threat of automation creates meaningful wage rigidity that amplifies labor market fluctuations.

VI. EVIDENCE FOR THE MODEL’S MECHANISM

Our model’s mechanism suggests that the procyclical automation probability dampens real wage fluctuations. By weakening workers’ bargaining power, automation makes real wages less correlated with labor productivity. This creates endogenous real wage rigidity and amplifies fluctuations in unemployment and vacancies. We now present some cross-sectional and time-series evidence that is consistent with the model’s mechanism.

VI.1. Cross-sectional evidence. We first present cross-sectional evidence that supports the model’s mechanism. We estimate the following dynamic panel data regression

$$\ln(Y_{it}) = a_0 + a_1 \ln(Y_{i,t-1}) + a_2 \ln(P_t) \times AP_i + \gamma_i + \eta_t + \varepsilon_{it}, \quad (38)$$

where the dependent variable Y_{it} includes the v/u ratio, vacancies, unemployment, and the real wage rate in industry i and quarter t . The main independent variable of interest is the interaction term $\ln(P_t) \times AP_i$, which is a proxy for time-varying automation threats. Here, P_t denotes the aggregate time series of the relative price of computer equipment and AP_i denotes the automation potential of industry i , which is a fixed characteristic of tasks in an industry, as we discuss in more detail below. By combining an aggregate relative price with physical characteristics of tasks in an industry, our proxy for the threat of automation is likely exogenous to movements in labor market variables in any given sector. We include the lagged dependent variable as a regressor to control for serial correlations in the dependent variable in our quarterly panel. The terms γ_i and η_t measure the industry fixed effects and the time fixed effects, respectively. The term ε_{it} is a regression residual.

The coefficient a_2 measures the relative sensitivity of the industry variable Y_{it} to changes in the aggregate computer prices, depending on the industry-specific automation risks. Our model suggests that, for an industry with a higher risk of automation (i.e., with a higher value of AP_i), a decline in the computer prices should lead to a larger increase in the job vacancy rate and the v/u ratio. Since the automation threat reduces workers' bargaining power and therefore dampens wage changes, the effect of a decline in the computer prices on real wages is a priori ambiguous.

We measure industry-level job vacancy rates using data from JOLTS and industry-level unemployment rates and real wages using data from the BLS. We focus on the pre-pandemic periods up to 2019. We obtain a quarterly unbalanced panel of vacancies, unemployment, and the v/u ratio that covers 15 industries at the two-digit NAICS level for the periods from 2001:Q1 to 2019:Q4. We also obtain a quarterly panel of real wages in 12 industries for the periods from 1985:Q1 to 2019:Q4. To measure the risk of automation, we use the technological potentials of automation for two-digit industries constructed by the McKinsey Global Institute (see Manyika et al. (2017)). This measure captures the (weighted-average) automation potential of tasks in an industry based on their physical characteristics. The relative price of computer equipment is the ratio of the Bureau of Economic Analysis's chain price index for private investment in computers and peripherals to the chained personal consumption expenditures price index (PCEPI). We use the time series of the computer prices to capture changes in aggregate economic conditions such as technological changes that drive automation decisions.¹⁶

¹⁶We provide detailed descriptions of these data in Appendix B.

TABLE 4. Automation threat and labor market outcomes: Cross-sectional evidence

Dep. Var.	$\ln(v/u)$	$\ln(v)$	$\ln(u)$	$\ln(w)$
$\ln(\text{ComputerPrice}) \times AP$	-0.411** (0.174)	-0.354** (0.141)	0.216* (0.115)	0.002* (0.001)
Lagged $\ln(v/u)$	0.470*** (0.028)			
Lagged $\ln(v)$		0.393*** (0.028)		
Lagged $\ln(u)$			0.407*** (0.029)	
Lagged $\ln(w)$				0.979*** (0.004)
Constant	0.868** (0.365)	1.611*** (0.302)	0.856*** (0.259)	0.061*** (0.010)
Industry fixed effect	Yes	Yes	Yes	Yes
Time fixed effect	Yes	Yes	Yes	Yes
Sargan p-value	0.224	0.409	0.056	0.019
Observations:	1125	1125	1125	1668
Time periods	2001:Q1-2019:Q4	2001:Q1-2019:Q4	2001:Q1-2019:Q4	1985:Q1-2019:Q4
No. of industries	15	15	15	12

Note: This table shows the panel estimation results using NAICS two-digit industry-level data under the empirical specification in Eq. (38). The set of dependent variables includes the labor market tightness (v/u), the vacancy rate (v), the unemployment rate (u), and the real wage rate (w), all in log units and at the quarterly frequency. In each regression, the independent variables include one lag of the dependent variable, the interaction between the relative price of computers and peripherals and the industry-specific automation potential (AP), and industry and time fixed effects. The stars denote the p-values: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

We estimate the empirical specification (38) using the Arellano-Bond estimator for dynamic panel data.¹⁷ Table 4 reports the estimation results. The columns shows the regression results for alternative industry-level dependent variables, including the v/u ratio (v/u), the vacancy rate (v), the unemployment rate (u), and the real wage rate (w), all in log units. The first row shows the estimated coefficient a_2 on the interaction term between the computer prices (also in log units) and the industry-specific automation risks.

The table shows that, in an industry more exposed to automation risks (i.e., with a higher AP_i), a decline in computer prices is associated with a greater increase in both the v/u ratio and the vacancy rate and a greater decline in the unemployment rate.¹⁸ The estimated

¹⁷The Arellano-Bond estimator first takes the differences of all variables in Eq (38) to remove the constant and the industry fixed effects. It then applies the generalized method of moments (GMM) approach to estimating the parameters of interest, using lagged dependent variables as a part of the instrumental variables. Stata implements this estimator with the command “xtabond2.” Since the estimator uses first differences of the variables, the coefficient a_2 measures the sensitivity of changes in the dependent variable for each industry to changes in the relative price of computer equipment, depending on the industry’s exposure to automation risks measured by AP_i .

¹⁸Our discussion here focuses on the effects of a decline in computer prices. If the the estimated coefficient a_2 is negative, as in the case for the v/u ratio (for instance), it means that a decline in computer prices would raise the v/u ratio.

impact of the automation threat (measured by the interaction term) on these labor market variables are statistically significant and economically important. For instance, a 1 percent increase in our measure of the automation threat (i.e., a 1 percent decline in the interaction between computer prices and AP_i) raises vacancies by 0.35 percent relative to its mean. This magnitude is in line with our model’s predicted impulse responses of vacancies to an automation-specific shock shown in Figure 3.¹⁹ A decline in computer prices is also associated with a significant reduction in real wages in industries more exposed to automation risks, consistent with the independent evidence documented by Acemoglu and Restrepo (2021). Overall, this cross-sectional evidence lends support to our model’s mechanism.²⁰

VI.2. Time-series evidence. Our model’s mechanism is also broadly consistent with time-series data. Figure 6 shows that the importance of automation has risen steadily since the early 2000s. Specifically, both the ratio of real investment in information processing equipment to real GDP and robot density (i.e., the operation stock of robots per thousand manufacturing employees) have increased during the past two decades (left panel). Since the importance of automation in the economy has been rising, our model mechanism suggests that the volatility of labor market tightness (i.e., the v/u ratio) should increase and the correlation between real wages and labor productivity should decrease over time. This is indeed the case, as shown in Figure 7. The figure shows that the standard deviation of labor market tightness initially declined from the 1980s to the 1990s, consistent with the “Great Moderation” hypothesis. Since the early 2000s, however, the volatility of the v/u ratio has increased steadily. The volatility of the v/u ratio in the 2010s (prior to the global pandemic) was about double that in the 1990s. The correlation between real wages and labor productivity initially increased from the 1980s to the 1990s, and then declined substantially since the early 2000s. The correlation in the 2010s was less than half of that in the 1990s.

¹⁹In the model, a one standard deviation positive shock to the automation-specific technology—or equivalently, a 6.5 percent increase in the shock—raises vacancies about 2 percent on impact (see Figure 3). The empirical estimate reported in Table 4 suggests that a 6.5 percent increase in the automation threat would raise vacancies by about $6.5 \times 0.35 \approx 2.3$ percent. A similar calculation suggests that the magnitude of the impulse responses of unemployment to an automation-specific shock in our model is comparable to the empirical estimate reported in Table 4.

²⁰The Sargan test suggests that, in the regressions of the v/u ratio, vacancies, and unemployment, the instrumental variables used in the GMM estimation are exogenous, although those in the real wage regression are potentially weak instruments. We have also estimated a similar empirical specification for industry-level labor shares using the BLS annual data from 1987 to 2019 (not reported in the paper). The estimation suggests that an increase in automation threat leads to a larger reduction in the labor share in an industry that is more exposed to automation, in line with our model’s implications. However, these effects (i.e., the coefficient a_2) are imprecisely estimated, partly because the sample is relatively small (at the annual frequency, instead of quarterly).




FIGURE 6. Trends in automation and offshoring in the United States. The left panel shows two measures of automation: the ratio of real investment in information processing equipment to real GDP (blue line) and the robot density measured by the operation stock of industrial robots per thousand manufacturing employees (green line). The right panel shows the ratio of real imports to real GDP. The shaded bars indicate NBER recession dates.

Source: U.S. Bureau of Economic Analysis, Bureau of Labor Statistics, Haver Analytics, International Federation of Robotics, and authors' calculations.

Our model suggests that these time-series changes in labor market dynamics might be partly driven by the rising threat of automation. To get a quantitative sense of the extent to which automation has contributed to the changes in labor market dynamics over time, we consider a counterfactual version of the model, in which we reduce the steady-state share of intermediate goods produced by automation capital (robots) relative to those produced by workers. Specifically, we raise the weight on intermediate goods produced by workers in the final goods production function from the calibrated value of $\alpha_n = 0.535$ to 0.682, such that the steady-state ratio of automation investment to aggregate output is about one-third of that in the benchmark model. This counterfactual captures the economy in earlier decades (such as the late 1980s and the early 1990s) when automation was less prevalent. For example, in the late 1980s and early 1990s, the ratio of information processing equipment investment to real GDP was about one-third of the sample median from 1985 to 2019. With a lower automation share, which translates into a lower automation probability, the volatility of labor market tightness is about half of that in the benchmark model, while the correlation between wage growth and labor productivity growth is about 30 percent higher. These patterns are consistent with the observed increases in the volatility of the v/u ratio and the




FIGURE 7. Changes in labor market volatilities and correlations over time. The left panel shows the standard deviations of labor market tightness measured by the v-u ratio for each of the four decades from 1980 to 2019. The right panel shows the correlations between year-over-year growth in real wages and in labor productivity in each of those four decades.

Source: Bureau of Labor Statistics, JOLTS, Haver Analytics, and authors' calculations.

substantial decline in the correlation between wage growth and productivity growth since the early 2000s shown in Figure 7.

Our model mechanism does not inherently rely on firms' ability to substitute robots for workers but on the possibility of adopting labor-saving technologies or offshoring of intermediate goods production. For example, Arseneau and Leduc (2014) show that the threat of offshoring can affect domestic wages through its impact on firms' outside options. Elsby et al. (2013) show that rising offshoring has played an important role in explaining the declines in the labor share from the mid-1980s up to the Great Recession in late 2000s. In the most recent decade, however, the importance of offshoring seems to have diminished, as the increase in the ratio of real imports to real GDP in the United States has slowed significantly following the trade collapse in 2008-2009 (right panel of Figure 6). The use of automation technologies, in contrast, has been steadily rising (see the left panel of the figure), suggesting that the importance of the automation channel has increased relative to the offshoring channel, at least during the recent decade.

VII. ROBUSTNESS OF THE MODEL MECHANISM

Our benchmark model shows that the threat of automation can effectively weaken workers' bargaining power, resulting in sluggish adjustments in real wages and therefore amplifying

fluctuations in unemployment and vacancies. We now consider three variations of the benchmark model and examine the robustness of the model's main transmission mechanism.²¹

VII.1. Production lags. In the benchmark model, a robot (or automation equipment in general) becomes productive without delays. In reality, however, it may take some time for a newly adopted robot to be productive, reflecting adjustment costs in using the new equipment. We now consider a variation of the benchmark model with production lags associated with the automation technology. For simplicity, we consider a one-period lag, such that a firm that operates the automation technology can produce intermediate goods using the automation equipment built in the previous period.

With production lags, the profit flow generated from using a robot built in period t can be materialized in period $t + 1$. Thus, the present value of automation would become

$$J_t^a = (1 - \rho^o) \mathbb{E}_t D_{t,t+1} [p_{a,t+1} Z_{t+1} \zeta_{t+1} - \kappa_a + J_{t+1}^a]. \quad (39)$$

The market-clearing condition for intermediate goods produced by the automation technology is now given by

$$Y_{at} = Z_t \zeta_t A_{t-1}, \quad (40)$$

where the input is the previous-period stock of automation equipment A_{t-1} instead of the current-period equipment.

The market-clearing condition (31) for final goods also needs to be modified such that the per-period fixed cost of operating the automation technology is given by $\kappa_a A_{t-1}$.

For ease of comparison, we solve this version of the model with lagged automation equipment in production using the same parameters as in the benchmark model. The implications of the model are very similar to those of the benchmark. For example, the unconditional volatility of labor market tightness and the ratio of the tightness volatility to the real wage volatility in this model with production lags are about 97 percent and 98 percent of those corresponding moments in the benchmark model. The impulse responses to each shock are also very similar to those in the benchmark model, as we show in the online appendix. Thus, our model mechanism is not sensitive to introducing production lags.

VII.2. Automating jobs. In our benchmark model, we assume that firms can automate a vacancy if that vacancy is not filled with a worker. A plausible alternative way of thinking about automation is to allow firms to automate an existing job instead of an open vacancy. We now consider that alternative setup.

²¹To conserve space, we sketch the key ingredients of each model in the text and describe the full equilibrium system in the online appendix.

In the beginning of period t , after observing all aggregate shocks, a firm can decide whether or not to replace a worker in an existing job match with a robot. The firm draws a cost x of automation from an i.i.d. distribution $F(x)$ and chooses to automate if the cost lies below the expected benefits of automation. There exists a threshold level of the automation cost—denoted by x_t^* —such that the firm automates the job position if and only if $x \leq x_t^*$. Thus, the automation probability is given by $q_t^a = F(x_t^*)$. If the firm adopts a robot, it obtains the automation value J_t^a (see Eq. (20)), but gives up the employment value J_t^e . Thus, the automation threshold is given by $x_t^* = J_t^a - J_t^e$.

The employment value takes into account the possibility of automation and is given by

$$J_t^e = p_{nt}Z_t - w_t + \mathbb{E}_t \beta \theta_{t+1} \frac{C_t}{C_{t+1}} \left\{ \delta_{t+1} J_{t+1}^v + (1 - \delta_{t+1}) \left[q_{t+1}^a J_{t+1}^a - \int_0^{x_t^*} x dG(x) + (1 - q_{t+1}^a) J_{t+1}^e \right] \right\}. \quad (41)$$

A job match yields the flow profit $p_{nt}Z_t - w_t$ in period t . In period $t + 1$, the job can be exogenously separated, in which case the firm obtains the vacancy value J_{t+1}^v . If the job is not separated, it can be automated with the probability q_{t+1}^a , in which case the firm obtains the automation value J_{t+1}^a net of the expected robot adoption costs. If the job is neither separated nor automated, then the firm obtains the continuation value of employment J_{t+1}^e .

Since a fraction of nonseparated jobs are automated, employment follows the law of motion

$$N_t = (1 - \delta_t)(1 - q_t^a)N_{t-1} + m_t. \quad (42)$$

We simulate the model based on the calibrated and estimated parameters as in the benchmark model. We recalibrate the relative productivity of automation such that the model with automated jobs implies the same steady-state probability of automation as that in the benchmark model.

Figure 8 shows the impulse responses of a few key macro and labor market variables following a positive discount factor shock. The shock raises the net present value of automation and thus increases the probability of automation. The increase in automation probability has two opposing effects on employment. Automation directly replaces workers in this model, pushing up the unemployment rate. At the same time, automation improves labor productivity, raising the employment value and boosting employment. With the calibrated parameters, the employment-boosting effect dominates the job-displaying effect, such that a positive discount factor shock leads to persistent declines in the unemployment rate. The discount factor shock boosts the value of vacancies, encouraging firms to create new vacancies and leading to large and persistent increases in the stock of vacancies. As a result, the labor market tightness (i.e., the v/u ratio) increases persistently (not shown in the figure).

Increased automation probability leads to a muted response of the real wage rate, reflecting that the threat of automation weakens workers' bargaining power. Since labor productivity

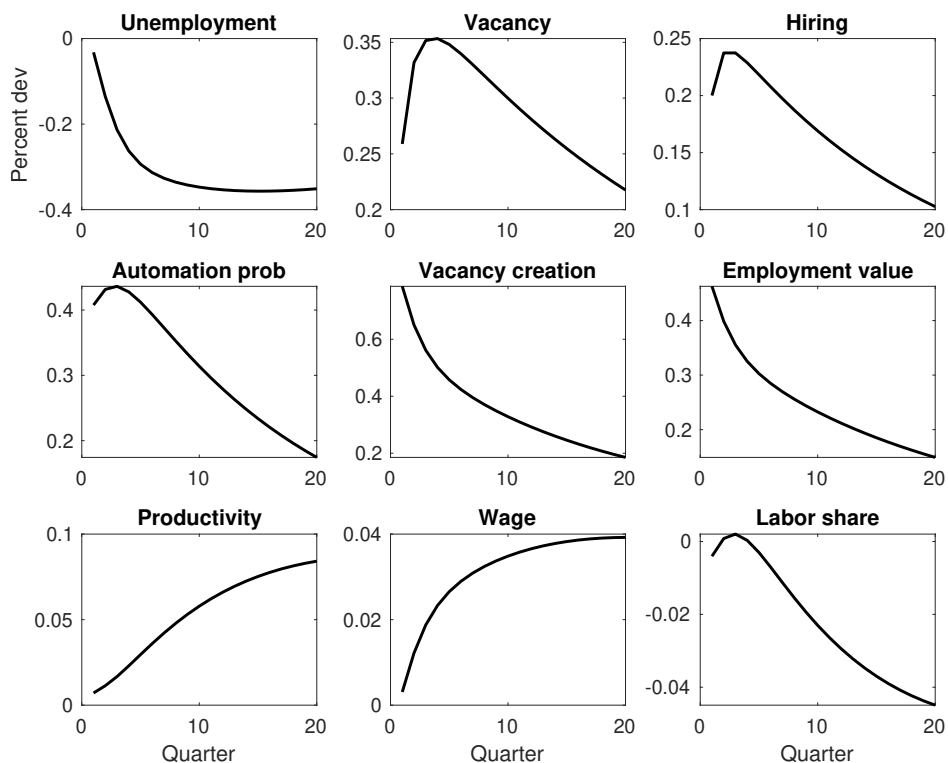


FIGURE 8. Impulse responses to a positive discount factor shock in the alternative model with automated jobs instead of vacancies.

risers and wage responses are relatively muted, the labor share declines following a positive discount factor shock, as in our benchmark model. Thus, the automation channel that we have identified in the benchmark model is robust when we consider automating jobs instead of vacancies.

VII.3. Heterogeneous worker skills. Our benchmark model features one type of workers, who compete with robots for jobs. We now generalize the model to incorporate heterogeneity in worker skills and show that the model mechanism survives this generalization.

The economy has two types of workers, skilled and unskilled, and all workers are members of the representative household family. A firm can produce a homogeneous consumption good by either hiring an unskilled worker from the frictional labor market or adopting a robot and hiring a skilled worker in a competitive spot labor market. Thus, robots are complementary inputs for skilled workers but are substitutes for unskilled workers.

To keep the model tractable, we assume that the aggregate supply of skilled workers is inelastic and fixed at \bar{s} .²² Skilled workers face a spot labor market with the competitive wage rate w_{st} . Unskilled workers face search frictions in the labor market, and they each receive

²²Given our focus on business cycles, assuming a constant supply of skilled workers seems innocuous since human capital accumulation is likely a slow-moving process. In the data, the share of skilled workers

either the unskilled wage rate w_{nt} if employed or the unemployment insurance benefit ϕ if unemployed.

The household utility function remains the same (see Eq. (9)). The budget constraint now includes the wage income from skilled workers and is given by

$$C_t + \frac{B_t}{r_t} = B_{t-1} + w_{nt}N_t + w_{st}\bar{s} + \phi(1 - N_t) + d_t - T_t. \quad (43)$$

Since the supply of skilled workers is inelastic, introducing skilled workers does not affect the household's optimizing decisions relative to the benchmark model. The only required modification in the household's problem is that, in the employment surplus expression (13), the wage rate w_t should be replaced by the unskilled wage rate w_{nt} .

An intermediate goods producer (i.e., a firm) can choose a technology at the beginning of each period: one requires an unskilled worker as the only input, and the other requires both a robot and s_t skilled workers as inputs. If the firm hires an unskilled worker for production, then it can produce $y_{nt} = Z_t$ units of output. If the firm adopts a robot, then it optimally chooses the input of skilled workers s_t , with the production function

$$y_{at} = Z_t \zeta_t^{\gamma_a} s_t^{1-\gamma_a}, \quad (44)$$

where $\alpha_a \in (0, 1)$ denotes the output elasticity of the robot input.

If a firm chooses the automation technology, then it takes the skilled wage rate w_{st} as given and chooses s_t to maximize the profit before paying robot operation costs. The value of automation is then given by

$$J_t^a = \pi_t^a + (1 - \rho^o) \mathbb{E}_t \beta \theta_{t+1} \frac{C_t}{C_{t+1}} J_{t+1}^a, \quad (45)$$

where $\pi_t^a \equiv \max_{s_t} p_{at} Z_t \zeta_t^{\gamma_a} s_t^{1-\gamma_a} - w_{st}s_t - \kappa_a$. Therefore, different from our benchmark model, automation now entails both a flow fixed cost and a variable cost through the use of skilled workers as a complementary input.

We use the calibrated and estimated parameters in the benchmark model (where appropriate), and calibrate three additional parameters in this generalized model. We set $\gamma_a = 0.32$, such that the skilled labor share is 68 percent of the revenue generated by the technology using robots and skilled workers as inputs. We normalize the supply of skilled workers and calibrate the average level of the automation-specific productivity (relative to the neutral technology) such that the model implies a steady-state skill premium of 55 percent, in line with the ratio of median weekly earnings of workers with a bachelor's degree or higher to those of workers with some college or associate degrees.

(e.g., those with a bachelor's degree or higher) has been rising steadily over time, and it shows little cyclical fluctuation.

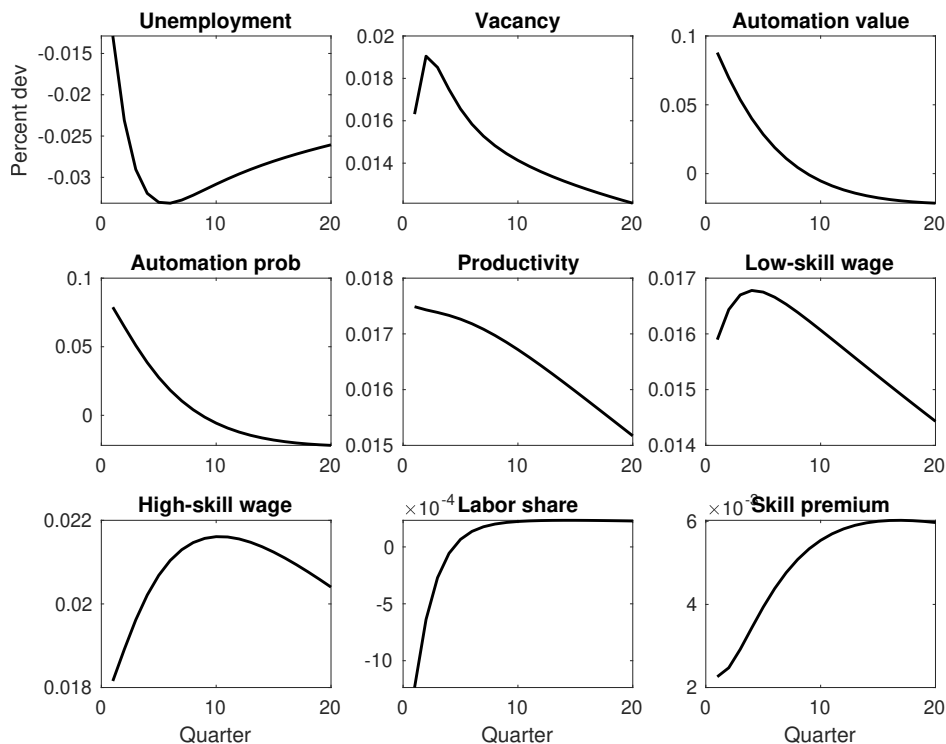


FIGURE 9. Impulse responses to a positive neutral technology shock in the model with heterogeneous worker skills.

Figure 9 shows the impulse responses following a positive neutral technology shock. The shock lowers unemployment and raises vacancies in the short run, although the responses are somewhat smaller in magnitude than those in the benchmark model. The shock boosts the present value of automation, raising the automation probability and labor productivity, which is measured by the ratio of aggregate output to aggregate employment, including both skilled and unskilled workers.

The increased threat of automation raises the firm’s reservation value in wage bargaining, restraining wage increases for unskilled workers. Since robots and skilled workers are complementary inputs in production, increases in automation raises the skilled wage rate, resulting in a persistent increase in the skill premium.

Under our calibration, despite the wage increases for both types of workers, the model predicts that the labor share (defined as the ratio of aggregate labor income—including skilled and unskilled labor—to aggregate output) falls in response to a positive neutral technology shock.²³

²³The impulse responses to a discount factor shock and an automation-specific shock are reported in the online appendix.

Since the automation technology requires both robots and skilled workers as complementary input factors, the flow cost of operating the automation technology increases with the wage rate of skilled workers. All else being equal, a business cycle boom that raises skilled wages would reduce the net value of automation, mitigating the increase in the automation probability. This feature is parallel to that for operating the manual technology, under which an increase in the wage rate of unskilled workers would reduce the present value of employment and thus discourage hiring. Nonetheless, the central mechanism of our benchmark model remains robust: increased automation raises labor productivity, dampens wage increases, boosts fluctuations in labor market tightness, and reduces the labor income share.

VIII. CONCLUSION

We have studied the role of automation in explaining the observed labor market dynamics in a tractable quantitative general equilibrium framework. The threat of automation raises the firm's reservation value in wage bargaining, reducing the worker's effective bargaining power and dampening increases in real wages in a business cycle boom. Thus, automation creates a source of real wage rigidity. At the same time, the option to automate a job position boosts the incentive for job creation, which offsets the direct job-displacing effects of automation. By muting wage increases while improving productivity, automation helps amplify fluctuations in unemployment and vacancies and also leads to a countercyclical labor share.

Our estimated general equilibrium model shows that the automation channel is quantitatively important. In particular, the automation channel helps account for the large volatility in unemployment and job vacancies relative to that of real wages, a puzzling observation through the lens of the standard DMP model without the automation channel.

A natural extension of our model framework is to incorporate firm heterogeneity in automation adoption. Recent surveys by the U.S. Census Bureau find that the use of automation technology is not widespread across firms; instead, it is highly skewed towards large and high-productivity firms (Zolas et al., 2020). Firooz et al. (2022) present a general equilibrium framework with heterogeneous firms that can match the skewed distribution of automation use in the data. Their model predicts that a rise in automation is associated with an increase in product market concentration and a decline in the labor income share. They focus on a spot labor market. In a more general framework with firm heterogeneity and labor market search frictions, one can potentially study how worker bargaining power and wages may depend on firm sizes through heterogeneous automation decisions. Research along these lines is an important and promising avenue for future work.

APPENDIX A. MACRO TIME-SERIES DATA

We fit the DSGE model to four time series of quarterly U.S. data, including the unemployment rate, the job vacancy rate, real wage growth, and labor productivity growth. The sample covers the period from 1985:Q1 to 2018:Q4.

- (1) **Unemployment:** civilian unemployment rate (age 16 years and over) from the Bureau of Labor Statistics, seasonally adjusted monthly series (LRUSECON in Haver).
- (2) **Job vacancies:** for pre-2001 periods, we use the vacancy rate constructed by Barnichon (2010) based on the Help Wanted Index. For the periods starting in 2001, we use the job openings from the Job Openings and Labor Turnover Survey (JOLTS), seasonally adjusted monthly series (LIJTLA@USECON in Haver).
- (3) **Real wages:** real compensation per worker in the nonfarm business sector. We first compute the nominal wage rate as the ratio of nonfarm business compensation for all persons (LXNFF@USECON in Haver) to nonfarm business employment (LXNFM@USECON) and then deflate it using the nonfarm business sector implicit price deflator (LXNFI@USECON).
- (4) **Labor productivity:** nonfarm business sector real output per person (LXNFS@USECON in Haver)

APPENDIX B. INDUSTRY-LEVEL DATA

The empirical regressions presented in Section VI.1 use industry-level data of the job vacancy rate, the unemployment rate, real wages, and the labor income share, along with the cross-sectional data of automation potentials and the time-series data of the relative price of computer equipment. Below, we describe the source and sample range of each variable used in the regressions.

- (1) **Job vacancy rate:** covers 15 two-digit NAICS industries in private nonfarm sectors from 2001 to 2019. Monthly values are converted to quarterly by taking the quarter-end values. Source: JOLTS.
- (2) **Unemployment rate:** covers 15 two-digit NAICS industries in private nonfarm sectors from 2001 to 2019. Monthly values are converted to quarterly by taking the quarter-end values. Source: BLS.
- (3) **Real wages:** average hourly earnings of production and nonsupervisory workers in 12 two-digit NAICS industries from 1985 to 2019, deflated by the chained personal consumption expenditures price index (PCEPI). Monthly values are converted to quarterly by taking the quarter-end values. Source: BLS.

TABLE A1. Summary statistics

	(1)				
	Mean	count	SD	Min	Max
$\ln(V)$	1.04	1140	0.40	-0.51	1.93
$\ln(U)$	1.67	1140	0.45	0.17	3.08
$\ln(V/U)$	-0.63	1140	0.72	-3.25	1.31
$\ln(w)$	3.00	1680	0.22	2.49	3.54
$\ln(\text{ComputerPrice}) \times AP$	2.72	2000	0.90	1.19	5.36

- (4) **Labor share:** percentage share of compensation of employees in value added, covering 18 two-digit NAICS industries in private nonfarm sectors, with annual data from 1987 to 2019. Source: BEA.
- (5) **Automation potentials:** technical potential for automation by industry estimated by the McKinsey Global Institute [see Manyika et al. (2017), Exhibit E4].
- (6) **Computer prices:** quarterly chain price index of private investment in computers and peripherals. This computer price series is deflated by the PCEPI to obtain the relative price of computers. Source: BEA.

Table A1 provides the summary statistics for the variables used in our industry-level regressions.

REFERENCES

- ACEMOGLU, D., C. LELARGE, AND P. RESTREPO (2020): “Competing with Robots: Firm-Level Evidence from France,” *AEA Papers and Proceedings*, 110, 383–388.
- 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, 1488–1542.
- (2020): “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy*, 128, 2188–2244.
- (2021): “Tasks, Automation, and the Rise in US Wage Inequality,” Tech. Rep. 28920, National Bureau of Economic Research.
- ARNOUD, A. (2018): “Automation Threat and Wage Bargaining,” Unpublished Manuscript, Yale University.
- ARSENEAU, D. M. AND S. LEDUC (2014): “Threatening to Offshore in a Search Model of the Labor Market,” Unpublished Manuscript.

- AUTOR, D. H. (2015): “Why Are There Still So Many Jobs? The History and Future of Workplace Automation,” *Journal of Economic Perspectives*, 29, 3–30.
- AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): “The Skill Content of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, 118, 1279–1333.
- BARNICHON, R. (2010): “Building a Composite Help-Wanted Index,” *Economics Letters*, 109, 175–178.
- BERG, A., E. F. BUFFIE, AND L.-F. ZANNA (2018): “Should We Fear the Robot Revolution? (The Correct Answer is Yes),” *Journal of Monetary Economics*, 97, 117–148.
- BLANCHARD, O. J. AND J. GALÍ (2010): “Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment,” *American Economic Journal: Macroeconomics*, 2, 1–30.
- CHENG, H., L. A. DROZD, R. GIRI, M. TASCHEREAU-DUMOUCHEL, AND J. XIA (2021): “The Future of Labor: Automation and the Labor Share in the Second Machine Age,” Working Paper WP 21-11, Federal Reserve Bank of Philadelphia.
- COSTAIN, J. S. AND M. REITER (2008): “Business Cycles, Unemployment Insurance, and the Calibration of Matching Models,” *Journal of Economic Dynamics and Control*, 32, 1120–1155.
- DEN HAAN, W. J., G. RAMEY, AND J. WATSON (2000): “Job Destruction and Propagation of Shocks,” *American Economic Review*, 90, 482–498.
- DINLERSOZ, E. AND Z. WOLF (2018): “Automation, Labor Share, and Productivity: Plant-Level Evidence from U.S. Manufacturing,” U.S. Census Bureau Center for Economic Studies Working Paper 18-39.
- EDEN, M. AND P. GAGGL (2018): “On the Welfare Implications of Automation,” *Review of Economic Dynamics*, 29, 15–43.
- ELSBY, M. W. L., B. HOBIJN, AND A. SAHIN (2013): “The Decline of the U.S. Labor Share,” *Brookings Papers on Economic Activity*, Fall, 1–63.
- FERNALD, J. G. (2015): “Productivity and Potential Output Before, During, and After the Great Recession,” *NBER Macroeconomics Annual 2014*, 29, 1–51.
- FIROOZ, H., Z. LIU, AND Y. WANG (2022): “Automation, Market Concentration, and the Labor Share,” Federal Reserve Bank of San Francisco Working Paper 2022-05.
- FREY, C. B. AND M. A. OSBORNE (2017): “The Future of Employment: How Susceptible Are Jobs to Computerisation?” *Technological Forecasting and Social Change*, 114, 254–280.
- FUJITA, S. AND G. RAMEY (2007): “Job Matching and Propagation,” *Journal of Economic Dynamics & Control*, 31, 3671–3698.

- GERTLER, M. AND A. TRIGARI (2009): “Unemployment Fluctuations with Staggered Nash Wage Bargaining,” *Journal of Political Economy*, 117, 38–86.
- GRAETZ, G. AND G. MICHAELS (2018): “Robots at Work,” *The Review of Economics and Statistics*, 107, 753–768.
- HAGEDORN, M. AND I. MANOVSKII (2008): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited,” *American Economic Review*, 98, 1692–1706.
- HALL, R. E. AND P. R. MILGROM (2008): “The Limited Influence of Unemployment on the Wage Bargain,” *American Economic Review*, 98, 1653–1674.
- HE, H. AND Z. LIU (2008): “Investment-Specific Technological Change, Skill Accumulation, and Wage Inequality,” *Review of Economic Dynamics*, 11, 314–334.
- HERSHBEIN, B. AND L. B. KAHN (2018): “Do Recessions Accelerate Routine- Biased Technological Change? Evidence from Vacancy Postings,” *American Economic Review*, 108, 1737–1772.
- JAIMOVICH, N. AND H. E. SIU (2020): “Job Polarization and Jobless Recoveries,” *The Review of Economics and Statistics*, 102, 129–147.
- KARABARBOUNIS, L. AND B. NEIMAN (2013): “The Global Decline of the Labor Share,” *Quarterly Journal of Economics*, 129, 61–103.
- KRUEGER, A. B. (2018): “Reflections on Dwindling Worker Bargaining Power and Monetary Policy,” Luncheon Address at the Jackson Hole Economic Symposium.
- KRUSELL, P., L. E. OHANIAN, J. V. RÍOS-RULL, AND G. L. VIOLANTE (2000): “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 68, 1029–1053.
- LEDUC, S. AND Z. LIU (2020): “The Weak Job Recovery in a Macro Model of Search and Recruiting Intensity,” *American Economic Journal: Macroeconomics*, 12, 310–343.
- MANYIKA, J., M. CHUI, M. MIREMADI, J. BUGHIN, K. GEORGE, P. WILLMOTT, AND M. DEWHURST (2017): “A Future That Works: Automation, Employment, and Productivity,” Tech. rep., McKinsey Global Institute, <https://www.mckinsey.com/ /media/mckinsey/featured>
- MCCALL, J. J. (1970): “Economics of Information and Job Search,” *The Quarterly Journal of Economics*, 84, 113–126.
- NEDELKOSKA, L. AND G. QUINTINI (2018): “Automation, Skills Use and Training,” OECD Social, Employment and Migration Working Papers, No. 202.
- RÍOS-RULL, J.-V. AND R. SANTAÈULÀLIA-LLOPIS (2010): “Redistributive Shocks and Productivity Shocks,” *Journal of Monetary Economics*, 57, 931–948.
- SHIMER, R. (2005): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *American Economic Review*, 95, 25–49.

- STANSBURY, A. AND L. H. SUMMERS (2020): “The Declining Worker Power Hypothesis: An Explanation for the Recent Evolution of the American Economy,” *Brookings Papers on Economic Activity*, Spring, 1–77.
- ZEIRA, J. (1998): “Workers, Machines, and Economic Growth,” *Quarterly Journal of Economics*, 113, 1091–1117.
- ZOLAS, N., Z. KROFF, E. BRYNJOLFSSON, K. MCELHERAN, D. N. BEEDE, C. BUFFINGTON, N. GOLDSCHLAG, L. FOSTER, AND E. DINLERSOZ (2020): “Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey,” Working Paper 28290, National Bureau of Economic Research.