ROBOTS OR WORKERS? A MACRO ANALYSIS OF AUTOMATION AND LABOR MARKETS

SYLVAIN LEDUC AND ZHENG LIU

Abstract. We argue that automation has contributed to sluggish wage growth despite historically low unemployment rates during the recovery from the Great Recession. This argument is based on a quantitative general equilibrium framework that incorporates automation decisions. A firm can choose to hire a worker or adopt a robot to produce goods. The *threat* of automation strengthens the firm's bargaining power against job seekers in wage negotiations, suppressing wages in economic expansions. The option to automate also raises the value of a vacancy, boosting the incentive for job creation, and thereby amplifying fluctuations in vacancies and unemployment. Our mechanism helps account for the large fluctuations in unemployment and vacancies relative to that in real wages, a puzzling observation through the lens of the standard labor search models.

I. Introduction

Despite the longest economic expansion and the lowest unemployment rate in recent U.S. history, wage growth has been subdued since the Great Recession (Krueger, 2018). Accompanying these developments in the labor market, recent advances in robotics and artificial intelligence have raised concerns that automation might put an increasing share of workers at risk of losing their jobs. There is an on-going debate about whether automation reduces overall employment (Autor, 2015; Acemoglu and Restrepo, 2017). But workers may have become more reluctant to ask for pay increases in a tight labor market, if they perceive their jobs as being more at risk of automation.

In this paper, we develop and estimate a general equilibrium framework to assess the quantitative impact of automation on wages and employment. Our analysis highlights a

---

Date: July 4, 2020.

Key words and phrases. Robots, automation, unemployment, wage bargaining, productivity, 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. For helpful comments, we thank 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, the HKUST Macro workshop, the Hong Kong University, the INSEAD, and the University of New South Wales. 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.

novel channel through which automation drives labor market fluctuations: the ability to automate, or the "threat of automation," provides an outside option to firms in wage negotiations, effectively raising their bargaining power. The automation threat creates a source of real wage rigidity that mutes wage increases during economic expansions, while boosting employment. The predictions of our estimated model are consistent with the strong labor markets and muted wage growth observed during the later part of the recovery from the Great Recession. More broadly, our model helps explain the large fluctuations in unemployment and vacancies relative to that in real wages in the data, a puzzling observation through the lens of the standard labor search models (Shimer, 2005).

I.1. **Model mechanism.** Our framework builds on the standard Diamond-Mortensen-Pissarides (DMP) model with labor market search frictions and generalizes it to incorporate automation decisions. In line with Acemoglu and Restrepo (2018) and Zeira (1998), firms in our model first make a choice of technologies (adopting or not adopting a robot); and only non-automated tasks (or vacancies) are available for hiring workers. Firms can produce consumption goods using either workers or robots. Since robots are perfect substitutes for workers in production, they are different from the physical capital in the standard neoclassical production functions, where capital and labor are complementary inputs.[1]

In the beginning of each period, a firm observes an i.i.d. cost of automation and decides whether or not to automate an unfilled job position that is carried over from the previous period. If the cost of automation lies below a threshold determined by the net benefit of automation, then the firm adopts a robot for production and takes the job vacancy offline. The probability of automation is thus the cumulative density of the automation cost draws evaluated at the automation threshold.

If the job position is not automated, then the firm posts the vacancy in the labor market to search for a potential match with a job seeker. If the match is successful, the vacancy will be filled with a worker and both the firm and the worker obtain their respective employment surplus from bargaining over the wage rate. If no match is formed, then the vacancy remains open and the firm obtains the continuation value of the vacancy, including the option to automate the position in future periods.[2]

[1]Krusell et al. (2000) study a neoclassical model in which capital equipment complements skilled labor but substitutes for unskilled labor. The relation between robots and workers in our model is analogous to the relation between equipment and unskilled labor in their model. For simplicity, we abstract from labor heterogeneity (skilled vs. unskilled). See He and Liu (2008) for a general equilibrium extension of the Krusell et al. (2000) model to incorporate endogenous skill accumulations.

[2]We interpret a job position broadly as consisting of a bundle of tasks, which are ex ante identical, but a fraction of which will be automated depending on the realization of the idiosyncratic costs of automation.

Our approach to modeling automation decisions requires a job vacancy to carry a positive value in equilibrium. Unlike the standard DMP model with free entry, we introduce a fixed cost of vacancy creation. A firm will choose to create a new vacancy if the vacancy-creation cost (drawn from an i.i.d. distribution) is below the value of the vacancy. Since vacancy creation is costly, an unfilled vacancy carries a positive value, allowing the firm to choose whether or not to automate an unfilled vacancy. Furthermore, different from the standard DMP model where the number of vacancies is a jump variable, it becomes a slow-moving state variable in our setup, enabling the model to match the persistent vacancy dynamics in the data (Leduc and Liu, 2019). Importantly, an increase in the value of a vacancy raises the firm's reservation value in wage negotiations, putting downward pressures on equilibrium wages.

I.2. **Model implications.** 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. To fit these four time series, we assume four shocks in our model, including a discount factor shock, a neutral technology shock, an automation-specific shock, and a job separation shock. We find that 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).

In our estimated model, we find that automation amplifies employment fluctuations. Automation has a direct job-displacing effect since goods produced by robots are perfect substitutes for goods produced by workers. But automation has also a job-creation effect: the option of automating an unfilled job vacancy boosts the present value of a vacancy, raising the incentive for firms to create vacancies. The net effect of automation on employment can be ambiguous, depending on the relative strength of the two opposing effects. Under our estimated parameters, automation amplifies employment fluctuations.

We also 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 good economic times, raising the firm's reservation value (i.e., the value of a vacancy) in wage bargaining, and therefore muting wage increases.

This approach simplifies our analysis significantly. We have considered an alternative timing of the automation decisions, under which firms first post the job vacancies for hiring workers, and then decide whether or not to automate the unfilled vacancies. The results are similar.

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, as observed in the data.[3]

Overall, automation helps generate large fluctuations in unemployment and vacancies relative to that in real wages. The threat of automation gives rise to a source of endogenous real wage rigidities, which are important for amplifying labor market fluctuations (Christiano et al., 2020). In addition, automation raises aggregate productivity in a business cycle boom, further fueling the boom. This mechanism is quantitatively important. In our estimated model and the data, the volatility of the vacancy-unemployment ratio (i.e., the v-u ratio), which is a measure of labor market tightness, is about 40 times that of the real wage rate.[4] In contrast, a counterfactual model without the automation mechanism produces a much smaller volatility ratio of about 9, less than 25 percent of that predicted by our estimated model. Furthermore, we show that search frictions are also important: a more competitive labor market tends to mitigate the real wage rigidity stemming form the threat of automation, making unemployment and vacancies less volatile. Thus, examining automation decisions within a search framework is important.

Although automation effectively weakens workers' bargaining power, its quantitative general equilibrium impact on the labor market cannot be reproduced in a counterfactual model with a lower worker bargaining power but without the automation mechanism. Absent the automation channel, reducing workers' bargaining power can amplify unemployment and vacancy fluctuations and dampen wage adjustments, yet these effects are substantially weaker than those in our benchmark model. In addition, the implied dynamics of wages are qualitatively different from those in our model with automation. For instance, in response to a positive discount factor shock, our model predicts a decline in the real wage, while it increases in the counterfactual model without automation and with a lower workers' bargaining weight. In addition, labor productivity rises in our benchmark model, which, coupled with the fall in the real wage, leads to a fall in the labor share. In contrast, the labor share increases in the counterfactual model. These differences reflect the threat of automation on wage bargaining and the endogenous productivity changes through the automation channel.

I.3. **Evidence for the model's mechanism.** Independent micro-level studies support our model's prediction that the threat of automation suppresses wage growth. For example,

[3]Karabarbounis 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. Our model has implications for the cyclical dynamics of the labor share, instead of its trends.

[4]Since we fit our model to these time series, the actual volatility ratio in the data is the same.

Arnoud (2018) 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 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, in line with our model's mechanism. 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. Graetz and Michaels (2018) use a panel of robot adoptions within industries in 17 countries from 1993 to 2007 and find that adoption of industrial robots boosts labor productivity and also raises wages, although the positive effects on wages are much smaller than those on productivity. Acemoglu and Restrepo (2017) study U.S. county-level data and report a negative effect of robot adoptions on local wages.

While the empirical literature using disaggregated data is well suited to examine the impact of automation on different types of industries, jobs, and tasks, it is difficult to aggregate the micro-level effects into a macroeconomic impact. Our dynamic stochastic general equilibrium (DSGE) model instead embeds the general equilibrium effects, although it does not directly speak to heterogeneous effects across different types of jobs or workers. In an environment with heterogeneous agents, automation can have important distributional consequences because it affects different occupations differently. For example, the recent work by Jaimovich et al. (2020) shows that increased automation has accounted for a large fraction of the declines in jobs in routine occupations. To the best of our knowledge, our model is the first quantitative general equilibrium model that incorporates automation decisions into a framework with labor search frictions to study the interactions between automation, bargaining power, and labor market fluctuations over the business cycle.

II. The model with labor market frictions and automation

This section presents a 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 by adopting robots. 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. The job seekers (u_t) randomly match with the vacancies (v_t) in the labor market, with the number of new matches (m_t)

determined by a matching technology. Production then takes place, with a homogeneous consumption good produced using either workers or robots. 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. Creating a new vacancy incurs a fixed cost e, which is drawn from an i.i.d. distribution $F(e)$, as in Leduc and Liu (2019). Thus, 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 automation cost x from an i.i.d. distribution $G(x)$, the realization of which determines whether the firm will adopt a robot or post the vacancy for hiring a worker. If the automation cost lies below a threshold value x_t^*, 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.[5] Since robots can substitute for workers for production, they are different from the traditional capital input, which is typically complementary to labor input in the standard macro models.

II.1. **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, so that the measure of unemployed job seekers is given by

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

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

The job separation rate shock δ_t follows the stationary stochastic process

$$\ln \delta_t = (1 - \rho_\delta) \ln \bar{\delta} + \rho_\delta \ln \delta_{t-1} + \varepsilon_{\delta t}, \tag{2}$$

[5] A plausible alternative way of thinking about automation is to allow firms to automate an existing job instead of an open vacancy. We have considered such an alternative setup. We find that automation implies counterfactual comovements between unemployment and vacancies, since it acts as an endogenous job destruction. See Appendix D for details. Furthermore, the alternative model with automated jobs underperforms our baseline model with automated vacancies in fitting the time-series data, with much lower log data density (1026.18 vs. 1258.54). Thus, the data strongly prefer our baseline model.

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 the 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

$$v_t = (1 - q^v_{t-1})(1 - q^a_t))v_{t-1} + \delta_t N_{t-1} + \eta_t, \tag{3}$$

where q^v_{t-1} denotes the job filling rate in period $t-1$, q^a_t denotes the automation rate in period t, and η_t denotes the newly created vacancies (i.e., entry).

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

$$m_t = \mu u_t^\alpha v_t^{1-\alpha}, \tag{4}$$

where μ is a scale parameter that measures match 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 employment pool, and job separations subtract from it. Aggregate employment evolves according to the law of motion

$$N_t = (1 - \delta_t)N_{t-1} + m_t. \tag{5}$$

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

$$U_t = u_t - m_t = 1 - N_t. \tag{6}$$

For convenience, we define the job finding probability q^u_t as

$$q^u_t = \frac{m_t}{u_t}. \tag{7}$$

Similarly, we define the job filling probability q^v_t as

$$q^v_t = \frac{m_t}{v_t}. \tag{8}$$

II.2. **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 $x \leq x^*_t$. If the firm adopts a robot to replace the job position, 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 automation decision is given by

$$x_t^* = J_t^a - J_t^v. \tag{9}$$

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

$$q_t^a = G(x_t^*). \tag{10}$$

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

$$A_t = (1 - \rho^o)A_{t-1} + q_t^a(1 - q_{t-1}^v)v_{t-1}, \tag{11}$$

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

Once adopted, 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

$$\ln Z_t = (1 - \rho_z)\ln \bar{Z} + \rho_z \ln Z_{t-1} + \varepsilon_{zt}. \tag{12}$$

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.[6] 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

$$\ln \zeta_t = (1 - \rho_\zeta)\ln \bar{\zeta} + \rho_\zeta \ln \zeta_{t-1} + \varepsilon_{\zeta t}. \tag{13}$$

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.

Operating the robot incurs a flow fixed cost of κ_a. The value of automation satisfies the Bellman equation

$$J_t^a = Z_t\zeta_t - \kappa_a + (1 - \rho^o)\mathbb{E}_t D_{t,t+1} J_{t+1}^a, \tag{14}$$

where the term κ_a captures the costs of energy, facilities, and space for automated production, and $D_{t,t+1}$ is a stochastic discount factor of the households.

If the automation cost exceeds the threshold x_t^*, then the vacancy will be posted 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 (2019),

[6] The model can easily be extended to allow for trend growth. We do not present that version of the model to simplify presentation.

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,

$$\eta_t = F(J_t^v). \tag{15}$$

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; 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

$$J_t^v = -\kappa + q_t^v J_t^e + (1 - q_t^v)\mathbb{E}_t D_{t,t+1}\left[q_{t+1}^a J_{t+1}^a + (1 - q_{t+1}^a) J_{t+1}^v\right]. \tag{16}$$

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

$$J_t^e = Z_t - w_t + \mathbb{E}_t D_{t,t+1}\left\{(1 - \delta_{t+1}) J_{t+1}^e + \delta_{t+1} J_{t+1}^v\right\}, \tag{17}$$

where w_t denotes the real wage rate. Hiring a worker generates a flow profit $Z_t - w_t$ 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.

II.3. **The representative household.** The representative household has the utility function

$$\mathbb{E}\sum_{t=0}^{\infty}\beta^t\Theta_t\left(\ln C_t - \chi N_t\right), \tag{18}$$

where $\mathbb{E}[\cdot]$ is an expectation operator, C_t denotes consumption, and N_t denotes the fraction of household members who are employed. The parameter $\beta \in (0, 1)$ denotes the subjective discount factor, and the term Θ_t denotes an exogenous shifter to the subjective discount factor.

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

$$\ln\theta_t = \rho_\theta \ln\theta_{t-1} + \varepsilon_{\theta t}. \tag{19}$$

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 (18) subject to the sequence of budget constraints

$$C_t + \frac{B_t}{r_t} = B_{t-1} + w_t N_t + \phi(1 - N_t) + d_t - T_t, \quad \forall t \geq 0, \tag{20}$$

where r_t denotes the gross real interest 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

$$V_t(B_{t-1}, N_{t-1}) \equiv \max \ \ln C_t - \chi N_t + \beta \mathbb{E}_t \theta_{t+1} V_{t+1}(B_t, N_t), \tag{21}$$

subject to the budget constraint (20) and the employment law of motion (5), the latter of which can be written as

$$N_t = (1 - \delta_t) N_{t-1} + q^u u_t, \tag{22}$$

where we have used the definition of the job finding probability $q_t^u = \frac{m_t}{u_t}$, with the measure of job seekers u_t given by Eq. (1). In the optimizing decisions, the household takes the economy-wide job finding rate q_t^u as given.

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 (20). We show in the Appendix that the employment surplus satisfies the Bellman equation

$$S_t^H = w_t - \phi - \frac{\chi}{\Lambda_t} + \mathbb{E}_t D_{t,t+1} (1 - q_{t+1}^u)(1 - \delta_{t+1}) S_{t+1}^H, \tag{23}$$

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.

The employment surplus has a straightforward economic interpretation. If the household adds a new worker in period t, then the current-period gain 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 (if 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. (23).

We also show in the appendix that the household's optimizing consumption-savings decision implies the intertemporal Euler equation

$$1 = \mathbb{E}_t D_{t,t+1} r_t. \tag{24}$$

II.4. **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. (23). 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 bargaining decisions.

The Nash bargaining problem is given by

$$\max_{w_t} \quad \left(S_t^H\right)^b \left(J_t^e - J_t^v\right)^{1-b}, \tag{25}$$

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

Define the total surplus as

$$S_t \equiv J_t^e - J_t^v + S_t^H. \tag{26}$$

Then the bargaining solution is given by

$$J_t^e - J_t^v = (1-b) S_t, \quad S_t^H = b S_t. \tag{27}$$

The bargaining outcome implies 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 (27) and the expression for household surplus in equation (23) together imply that the Nash bargaining wage w_t^N satisfies the Bellman equation

$$\begin{aligned} \frac{b}{1-b}(J_t^e - J_t^v) \quad &= \quad w_t^N - \phi - \frac{\chi}{\Lambda_t} \\ &+ \mathbb{E}_t D_{t,t+1}(1 - q_{t+1}^u)(1 - \delta_{t+1}) \frac{b}{1-b}(J_{t+1}^e - J_{t+1}^v). \end{aligned} \tag{28}$$

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

II.5. **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 so that

$$\phi(1 - N_t) = T_t. \tag{29}$$

II.6. **Search equilibrium.** In a search equilibrium, the markets for bonds and goods both clear. Since the aggregate bond supply is zero, the bond market-clearing condition implies that

$$B_t = 0. \tag{30}$$

Goods market clearing requires that consumption spending, vacancy posting costs, automation costs, and vacancy creation costs add up to aggregate production. This requirement yields the aggregate resource constraint

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

where Y_t denotes aggregate output, which equals the sum of goods produced by workers and by robots and is given by

$$Y_t = Z_t N_t + Z_t \zeta_t A_t. \tag{32}$$

III. Empirical Strategies

We solve the model by log-linearizing the equilibrium conditions around the deterministic steady state.[7] 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

$$F(e) = \left(\frac{e}{\bar{e}}\right)^{\eta_v}, \quad G(x) = \left(\frac{x}{\bar{x}}\right)^{\eta_a}, \tag{33}$$

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.[8] We estimate the scale parameters $\bar{e}$ and $\bar{x}$ and the shock processes by fitting the model to U.S. time series data.

[7]Details of the equilibrium conditions, the steady state, and the log-linearized system are presented in the appendix.

[8]Our 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.

III.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. We normalize the level of labor productivity to $\bar{Z} = 1$ and automation-specific productivity to $\bar{\zeta} = 1$.

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 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 $\mu = \frac{m}{u^\alpha v^{1-\alpha}} = 0.6594$. We set the flow cost of operating robots to $\kappa_a = 0.98$. Given the average productivities $\bar{Z} = \bar{\zeta} = 1$, this implies a quarterly profit of 2 percent of the revenue by using a robot for production. The steady-state automation value J^a can then be solved from the Bellman equation (14).

Conditional on J^a and the estimated values of $\bar{e}$ and $\bar{x}$ (see below for estimation details), we use the vacancy creation condition (15), the automation adoption condition (9), and law of motion for vacancies (3) 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 automatons A can be solved from the law of motion (11), which reduces to $\rho^o A = q^a(1 - q^v)v = \eta$ 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 employment relations.

With A and N solved, we obtain the aggregate output $Y = \bar{Z}(N + \bar{\zeta}A)$. We calibrate the vacancy posting cost to $\kappa = 0.0939$, so 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 $x^* = J^a - \beta J^v$. The match value J^e can be solved from the Bellman equation for vacancies (16), and the equilibrium real wage rate can be obtained from the Bellman equation for employment (17). Steady-state consumption is solved from the resource constraint (31). We then infer the value of $\chi = 0.7292$ from the expression for bargaining surplus in Eq. (28).

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

III.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.

The unemployment rate in the data (denoted by U_t^{data}) corresponds to the end-of-period unemployment rate in the model U_t. We demean the unemployment rate data (in log units) and relate it to our model variable according to the measurement equation

$$\ln(U_t^{data}) - \ln(\bar{U}^{data}) = \hat{U}_t, \tag{34}$$

where $\bar{U}^{data}$ denotes the sample average of the unemployment rate in the data and $\hat{U}_t$ denotes the log-deviations of the unemployment rate in the model from its steady-state value.

Similarly, we use demeaned vacancy rate data (also in log units) and relate it to the model variable according to

$$\ln(v_t^{data}) - \ln(\bar{v}^{data}) = \hat{v}_t, \tag{35}$$

where $\bar{v}^{data}$ denotes the sample average of the vacancy rate 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 the 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 (denoted by $\Delta \ln p_t^{data}$) and relate it to our model variable according to

$$\Delta \ln(p_t^{data}) - \Delta \ln(p^{data}) = \hat{Y}_t - \hat{N}_t - (\widehat{Y}_{t-1} - \widehat{N}_{t-1}), \tag{36}$$

where $\Delta \ln(p^{data})$ denotes the sample average of 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 compensations per worker in the nonfarm business sector. We relate the observed real wage growth (denoted by $\Delta \ln(w_t^{data})$) to the model variables by the measurement equation

$$\Delta \ln(w_t^{data}) - \Delta \ln(w^{data}) = \hat{w}_t - \hat{w}_{t-1}, \tag{37}$$

where $\Delta \ln(w^{data})$ denotes the sample average of wage growth in the data and $\hat{w}_t$ denotes the log-deviations of real wages from its steady-state level in the model.

III.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 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.

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 $\bar{e} = 8.60$. The posterior mean estimates of the automation cost parameter is $\bar{x} = 1.80$. These parameters imply a steady-state share of output produced by automation of $A/Y = 0.24$. Thus, our model implies that, in the long run, about 24 percent of the jobs will be performed by robots, which lies in the range of the estimates in the empirical literature (Nedelkoska and Quintini, 2018). The 90 percent probability intervals indicate that the data are informative about the structural parameters.

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.

IV. Economic implications

Based on the calibrated and estimated parameters, we examine the model's transmission mechanism and its quantitative performance for explaining the labor market dynamics. We also present some counterfactuals to illustrate the quantitative importance of both automation and labor market search frictions.

IV.1. **The model's transmission mechanism.** The equilibrium dynamics in our model are driven by both the exogenous shocks and the model's internal propagation mechanism. To help understand the contributions of the shocks and the model's mechanism, we examine impulse response functions and forecast error variance decompositions.

IV.1.1. *Impulse responses.* 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, leading to an increased probability of robot adoption, which raises the value of a vacancy and boosts the incentive for vacancy creation. The increase in vacancy value also strengthens the firm's bargaining power in wage negotiations, dampening the responses of real wages. Increased automation also raises labor productivity, reinforcing the initial expansionary impact of the technology shock. The increase in labor productivity, coupled with muted wage responses, implies persistent declines 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, an open vacancy, and a worker's employment surplus. Thus, it generates a persistent boom in employment, vacancies, and hiring. The shock also raises the value of automation and therefore increases the probability of robot adoption. The increased automation probability raises the vacancy value, incentivizing vacancy creation. The increase in actual robot adoption raises 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 boosting productivity and reducing the real wage rate, the discount factor shock generates a persistent decline in the labor share.

A job separation shock raises both unemployment and vacancies and mechanically boosts hiring through the matching function, as shown in Figure 3. This finding is consistent with Shimer (2005), who argues that the counterfactual implication of the job separation shock for the correlation between unemployment and vacancies renders the shock unimportant for explaining observed labor market dynamics. The shock reduces the automation probability. Labor productivity increases slightly, since the decline in employment outpaces the decline in aggregate output. The shock also leads to small declines in real wages and the labor income share.

Figure 4 shows the impulse responses to a positive automation-specific shock. The shock directly raises the value of automation. In turn, the increased probability of automation raises the vacancy value and boosts the incentive for vacancy creation. With more job openings, the job finding rate increases, raising hiring and reducing unemployment. Since a greater fraction of output is produced with robots, labor productivity improves. The increased threat of automation weakens the worker's bargaining power, leading to a decline in the real wage rate. The improvement in labor productivity and the reduction in the real wage rate result in a persistent decline in the labor income share.

IV.1.2. *Forecast error variance decompositions.* We now examine the unconditional forecast error variance decompositions for the four observable labor market variables used for our estimation.[9] Table 3 displays the results.

The variance decompositions suggest that fluctuations of unemployment and vacancies are mostly driven by the neutral technology shock and the discount factor shock. The neutral technology shock accounts for about one-third of the variances of unemployment and vacancies, and the discount factor shock accounts for about 60 percent. The job separation shock is not important for these labor market variables, consistent with Shimer (2005).

The automation-specific shock does not directly contribute to the fluctuations in unemployment and vacancies; instead, the threat of automation works to amplify the effects of the other shocks, particularly the neutral technology and the discount factor shocks, by raising the probability of automation. These two shocks explain about 70 percent of the fluctuations in the automation probability (not shown in the table). As discussed in the previous section, the resulting procyclical threat of automation dampens real wage adjustments and thus magnifies the impact of the neutral technology and the discount factor shocks on labor market variables.

While the threat of automation dampens wage adjustments, the actual adoption of robots raises labor productivity. Through these channels, the automation-specific shock plays a quantitatively important role in driving fluctuations of the growth rates of both labor productivity and real wages. This shock accounts for about 32 percent of the variance of productivity growth and 37 percent of that of real wage growth. Perhaps not surprisingly, the neutral technology shock is also important for explaining the fluctuations in labor productivity, explaining about half of its variance.[10] In addition, about 62 percent of the real wage fluctuations are accounted for by shocks to the neutral technology and the discount factor.

IV.2. **The role of automation in the propagation mechanism.** To isolate the role of automation in driving labor market dynamics, we consider a counterfactual specification of "no automation," which is a version of our benchmark model with all automation-related variables held constant at their steady-state levels and with no automation-specific shocks. To highlight the effect of the threat of automation on the worker's bargaining power in wage negotiations, we also compare our benchmark model's impulse responses to a version of the

[9]We have also computed the conditional forecast error variance decompositions with forecasting horizons between 4 quarters and 16 quarters and found that they deliver the same message as the unconditional variance decomposition.

[10]In the standard DMP model without automation, labor productivity fluctuations would be entirely driven by the neutral technology shock.

"no automation" case, in which we also reduce the bargaining weight for workers by a half (i.e., setting $b = 0.25$). We refer to this version of the model as the "low bargaining power" case.

Figure 5 displays the impulse responses to a discount factor shock in the three models: the benchmark model (the black solid lines), the counterfactual with no automation (the blue dashed lines), and the counterfactual with no automation and a lower bargaining power for workers (the red dot-dash lines). These impulse responses suggest that the automation channel is a powerful amplification mechanism for labor market dynamics. Without automation, the counterfactual model implies much more muted responses of unemployment, vacancies, and hiring to the discount factor shock than those in the benchmark model. The responses of the real wage rate are also different: the counterfactual model implies a small increase in the real wage rate, whereas the benchmark model implies a modest decrease. This pattern suggests that the threat of automation is important for suppressing wage adjustments. Without automation, labor productivity is solely driven by the neutral technology shock, so that productivity does not respond to the discount factor shock. With automation, as in our benchmark model, labor productivity rises following a discount factor shock, because the shock raises the value of automation and thus leads to increased adoption of robots. As a consequence, the labor income share rises in the counterfactual but falls in our benchmark model.

In the no-automation model, reducing workers' bargaining weight mechanically dampens real wage adjustments and thus should help amplify the responses of the unemployment and vacancy rates. This can be seen from the steady-state version of the Nash bargaining wage solution in Eq. (28):

$$w^N = \phi + \frac{\chi}{\Lambda} + \frac{b}{1-b}[1 - \beta(1-q^u)(1-\delta)](J^e - J^v). \tag{38}$$

Clearly, *ceteris paribus*, the Nash bargaining wage w^N increases with the worker's bargaining weight b and decreases with the firm's reservation value J^v. A reduction in b reduces the equilibrium wage, as does an increase in J^v when firms have the option to automate.

However, the impulse responses shown in Figure 5 reveal that reducing the worker's bargaining weight is much less effective in amplifying labor market fluctuations than granting firms the ability to automate job positions. In addition, even if the worker's bargaining weight is reduced, the real wage rate still rises following the discount factor shock, leading to an increase in the labor share. In contrast, our benchmark model with automation implies a fall in the real wage and a persistent decline in the labor share.

The impulse responses to a neutral technology shock in these counterfactual models display similar patterns, as shown in Figure 6. These impulse responses suggest that the automation channel is an important mechanism for amplifying labor market fluctuations and generating a countercyclical labor income share.

IV.3. **The role of labor market search frictions in the propagation mechanism.** The model's amplification mechanism depends not only on automation, but also on labor market search frictions. To illustrate the importance of the search frictions, we consider a counterfactual version of the model which features low levels of labor search frictions. In particular, that counterfactual model has a smaller vacancy posting cost (of 0.5 percent of aggregate output in the steady state instead of 1 percent) and a higher average job separation rate (with $\bar{\delta} = 0.5$ instead of 0.1).[11]

Figure 7 shows the impulse responses of the macro variables following a positive neutral technology shock, and compares the impulse responses from the benchmark model (the black solid lines) with those from the counterfactual with low search frictions (the blue dashed lines). Although both models have the automation channel operating, the benchmark model produces much stronger amplification effects of the shock on unemployment and vacancies than does the counterfactual with low search frictions. As we have discussed, automation displaces jobs because robots and workers are substitutable inputs of production; meanwhile, automation also boosts job creation because the option of automating increases the present value of a job vacancy. In an economy with a spot labor market without search frictions, an employment relation would cease to be a long-term relation, and the option of automation in the future would not directly affect current hiring decisions. In that case, robots would replace workers, and increased automation in response to a positive technology shock would raise unemployment and reduce vacancies. The counterfactual model with lower search frictions lies between our benchmark model and the spot labor market, with mitigated job-creating effect of automation, and thus more muted responses in unemployment and vacancies than those in the benchmark economy.[12]

With low search frictions, the present value of a vacancy responds less to the technology shock (because the model becomes closer to a spot labor market). Since the neutral technology shock directly raises the productivity of both workers and robots, the value of automation rises on impact. Thus, the automation threshold (i.e., $x_t^* = J_t^a - J_t^v$) and the

[11]In the limit with $\kappa = 0$ and $\delta = 1$, there is no vacancy posting cost and employment becomes a jump variable, approximating a spot labor market. We do not consider that extreme case to minimize deviations from our benchmark framework.

[12]We have also considered the case with a higher average job separation rate of $\bar{\delta} = 0.8$ (not reported in the paper). In that case, we find that a positive neutral technology shock raises unemployment and lowers vacancies, because the forward-looking job-creating effect becomes further mitigated.

automation probability rises more sharply than in the benchmark model, as shown in Figure 7, leading to stronger increases in labor productivity. Although the real wage rate also increase more than that implied by the benchmark model, the productivity effects dominate, leading to a more pronounced and persistent decline in the labor share.

The importance of labor search frictions can be further illustrated by comparing the impulse responses following a positive automation-specific shock, as shown in Figure 8. In the counterfactual model with low search frictions, the automation-specific shock raises unemployment and reduces vacancies because the direct job-displacing effect dominates the (mitigated) job-creating effect. In contrast, in the benchmark model, the same shock reduces unemployment and increases vacancies. As in the case with a neutral technology shock, the automation-specific shock directly raises the value of automation, while the counterfactual with low search frictions implies a more muted response of the value of vacancies. Thus, the automation probability, labor productivity and real wages all increase more sharply than those in the benchmark model.

The impulse responses shown in Figures 7 and 8 suggest that search frictions are important because they give rise to forward-looking hiring decisions, generating a job-creating effect of automation that would be otherwise absent in a spot labor market.

The results discussed in this section and those presented in Section IV.2, taken together, suggest that the automation channel interacts with labor search frictions, enabling our model to confront the time-series data of the U.S. labor market.

IV.4. **Automation threat and labor market dynamics.** Our model predicts that the threat of automation dampens wage adjustments and amplifies labor market fluctuations. Is the automation mechanism quantitatively important? To examine the empirical importance of the automation mechanism, we compare our model's predictions for labor productivity and real wages with those from two counterfactuals: one without the automation channel and the other with low search friction. As discussed in Section IV.2, the no-automation counterfactual is a version of our benchmark model with the automation-related variables held constant at their steady-state values and the automation-specific shock shut off. The low search frictions counterfactual is a version of the benchmark model with a lower vacancy posting cost and a higher average job separation rate, as in Section IV.3.

Our benchmark model implies that the probability of automation is procylical, rising in business cycle booms and falling in recessions. Thus, the increased threat of automation mutes wage growth in a business cycle boom, allowing the model to generate large volatilities of the labor market tightness (the v-u ratio) relative to that of the real wage rate. In this sense, the automation channel helps resolve the Shimer puzzle.

Indeed, both automation and search frictions are important for generating the observed large relative volatility of the labor market tightness, as illustrated in Table 4. The table shows the standard deviations of the labor market tightness (i.e., the v-u ratio), the real wage rate (w), and the relative volatility of the tightness (relative to that of real wages). In the benchmark model, the v-u ratio is about 39 times as volatile as the real wage rate. This relative volatility is the same as in the actual data, because the model is estimated to fit these time series.

The counterfactual with no automation generates a much smaller volatility of the v-u ratio (0.298 vs. 1.159) and modestly larger volatility of the real wage rate (0.031 vs. 0.029) than does the benchmark model, implying a much smaller relative volatility (9.56 vs. 39.47). This no-automation case essentially reproduces the Shimer (2005) volatility puzzle. The counterfactual with low search frictions also generates less volatility of the v-u ratio (0.99) and more volatility of the real wage rate (0.034), implying a smaller relative volatility than in the benchmark model (29.48). Thus, both automation and labor search frictions are important for the model's transmission mechanism.

Our model's mechanism also sheds light on the driving factors for the muted wage growth over the past few years despite an increasingly tightened labor market. To assess the contribution of the automation mechanism to the observed dynamics of real wages and unemployment, we simulate a counterfactual model in which we shut off the automation channel, and into which we feed in the same parameters and shocks as those in our benchmark model. We then compare the smoothed time series generated from the counterfactual model with those from the benchmark model, the latter of which replicates the actual time series data under our Bayesian estimation. We find that, between 2013 and 2018, the automation mechanism has reduced the cumulative real wage growth by roughly 10 percent; it has also reduced the unemployment rate by an average of about 2.5 percentage points per year during the same period. Thus, absent automation, both the real wage rate and the unemployment rate would have been substantially higher during this period. In this sense, the automation mechanism has contributed significantly to the persistent declines in unemployment accompanied by sluggish wage growth in the recovery from the Great Recession.

V. Conclusion

We have studied the role of automation in explaining the observed labor market dynamics in a quantitative general equilibrium framework. The threat of automation raises the firm's reservation value in wage bargaining, 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 growth while improving productivity, automation helps amplify fluctuations in unemployment and vacancies.

Our estimated general equilibrium model shows that the automation channel is quantitatively important. The automation mechanism has contributed significantly to the observed sluggish wage growth despite strong labor markets during the long expansion following the Great Recession. More broadly, automation 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 with labor search frictions.

Similar effects could also arise from other labor-saving mechanism, such as offshoring. When firms have the option of importing intermediate goods instead of producing them domestically, the threat of offshoring could also weaken domestic workers' bargaining power in wage negotiations, similar to the threat of automation in our model. Other factors such as increases in product market concentration and declines in union powers may have also contributed to the observed labor market dynamics in the past few decades. Assessing the quantitative importance of these alternative contributing factors requires a coherent general equilibrium framework that can be used to fit time series data. Our framework with automation provides a useful step in that promising direction for future research.

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	0.0939
b	Nash bargaining weight	0.50
η_v	Elasticity of vacancy creation cost	1
η_a	Elasticity of automation cost	1
κ_a	Flow cost of automated production	0.98
χ	Disutility of working	0.7292
$\bar{Z}$	Mean value of neutral technology shock	1
$\bar{\zeta}$	Mean value of equipment-specific technology shock	1

Table 2. Estimated parameters

		Priors		Posterior		
	Parameter description	Type	[mean, std]	Mean	5%	95%
$\bar{e}$	scale for vacancy creation cost	G	[5, 1]	8.6001	7.7510	9.4607
$\bar{x}$	scale for robot adoption cost	G	[5, 1]	1.8020	1.3155	2.2557
ρ_z	AR(1) of neutral technology shock	B	[0.8, 0.1]	0.9432	0.9235	0.9598
ρ_θ	AR(1) of discount factor shock	B	[0.8, 0.1]	0.9702	0.9535	0.9838
ρ_δ	AR(1) of separation shock	B	[0.8, 0.1]	0.9354	0.9019	0.9662
ρ_ζ	AR(1) of automation-specific shock	B	[0.8, 0.1]	0.8205	0.7991	0.8471
σ_z	std of tech shock	IG	[0.01, 0.1]	0.0109	0.0095	0.0128
σ_θ	std of discount factor shock	IG	[0.01, 0.1]	0.0124	0.0090	0.0170
σ_δ	std of separation shock	IG	[0.01, 0.1]	0.0489	0.0452	0.0515
σ_ζ	std of automation-specific shock	IG	[0.01, 0.1]	0.0320	0.0244	0.0406
	Log data density		1258.54			

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.

Table 3. Forecasting Error Variance Decomposition

Variables	Neutral technology shock	Discount factor shock	Job separation shock	Automation specific shock
Unemployment	34.48	63.30	1.22	0.99
Vacancy	31.69	57.27	10.03	1.01
Productivity growth	50.16	17.60	0.25	32.00
Real wage growth	43.64	18.48	0.48	37.40

Note: The numbers reported are the posterior mean contributions (in percentage terms) of each of the four shocks in the benchmark estimation to the forecast error variances of the variables listed in each row.

Table 4. Labor market volatilities implied by alternative models

Model	Labor market tightness	Real wage	Relative volatility
Benchmark model	1.1591	0.0294	39.4676
No automation	0.2982	0.0312	9.5605
Low search friction	0.9904	0.0336	29.4758

Note: The three rows in the table correspond to three alternativ models: the benchmark model, the no-automation counterfactual, and the low-search-friction counterfactual (see the text for more detailed explanations of these models). For each model, the numbers in the three columns are (1) the standard deviation of labor market tightness measured by the ratio of vacancies to unemployment, (2) the standard deviation of the real wage rate, and (3) the ratio of the first two columns.

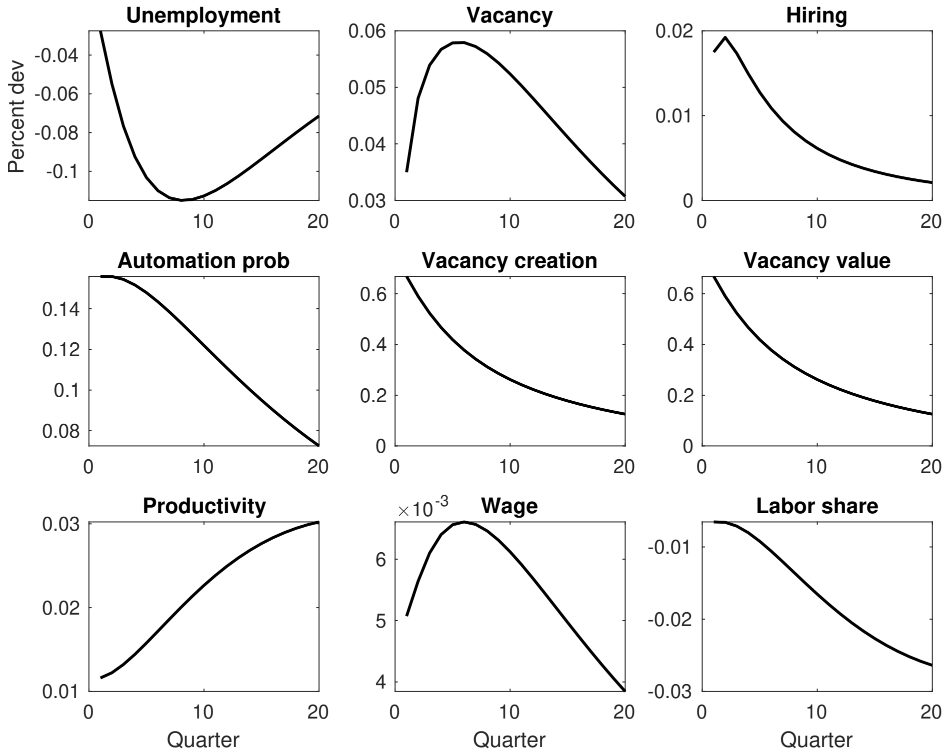

FIGURE 1. Impulse responses to a positive neutral technology shock in the benchmark model.

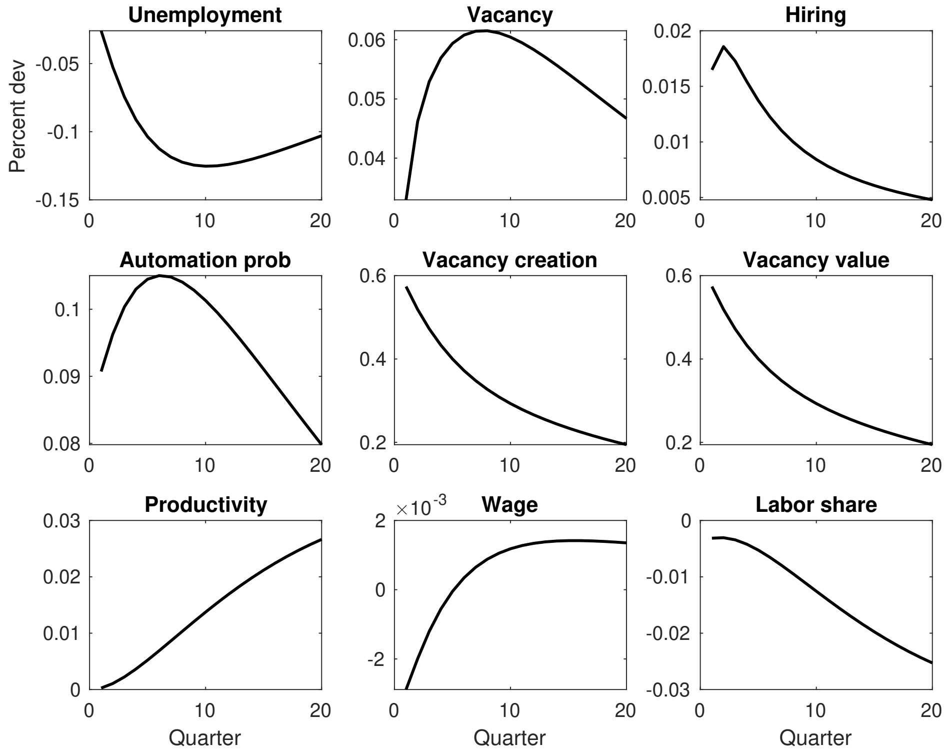

FIGURE 2. Impulse responses to a positive discount factor shock in the benchmark model.

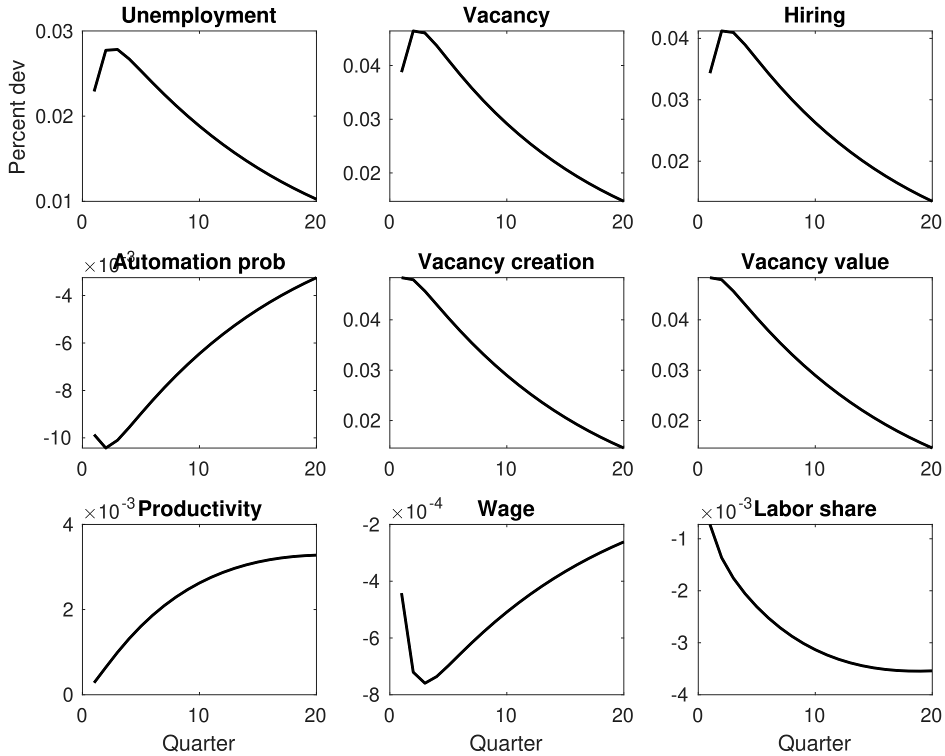

FIGURE 3. Impulse responses to a job separation shock in the benchmark model.

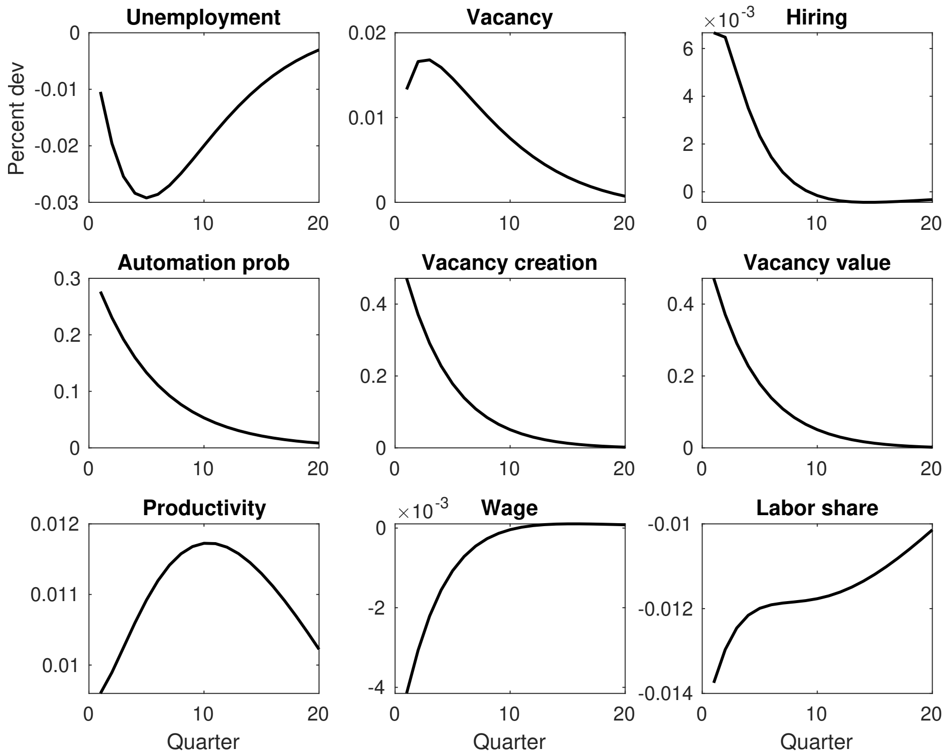

FIGURE 4. Impulse responses to a positive automation-specific shock in the benchmark model.

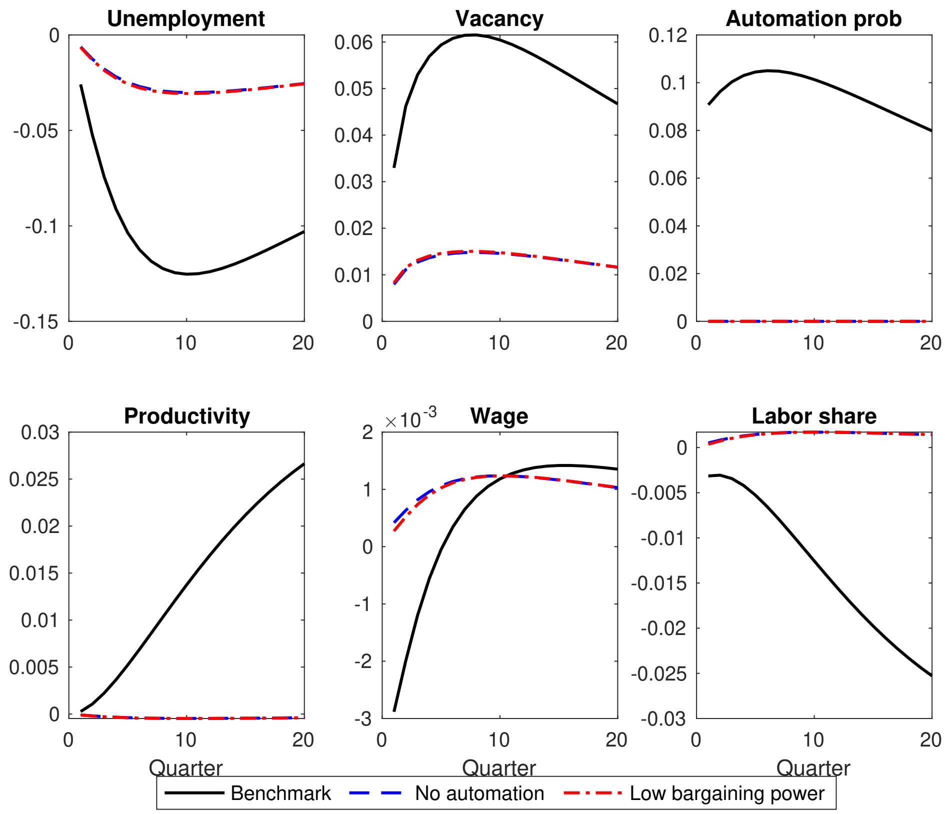

Figure 5. Impulse responses to a positive discount factor shock in the benchmark model (black solid lines), the counterfactual with no automation (blue dashed lines), and the counterfactual with no automation and low worker bargaining power (red dot-dash lines).

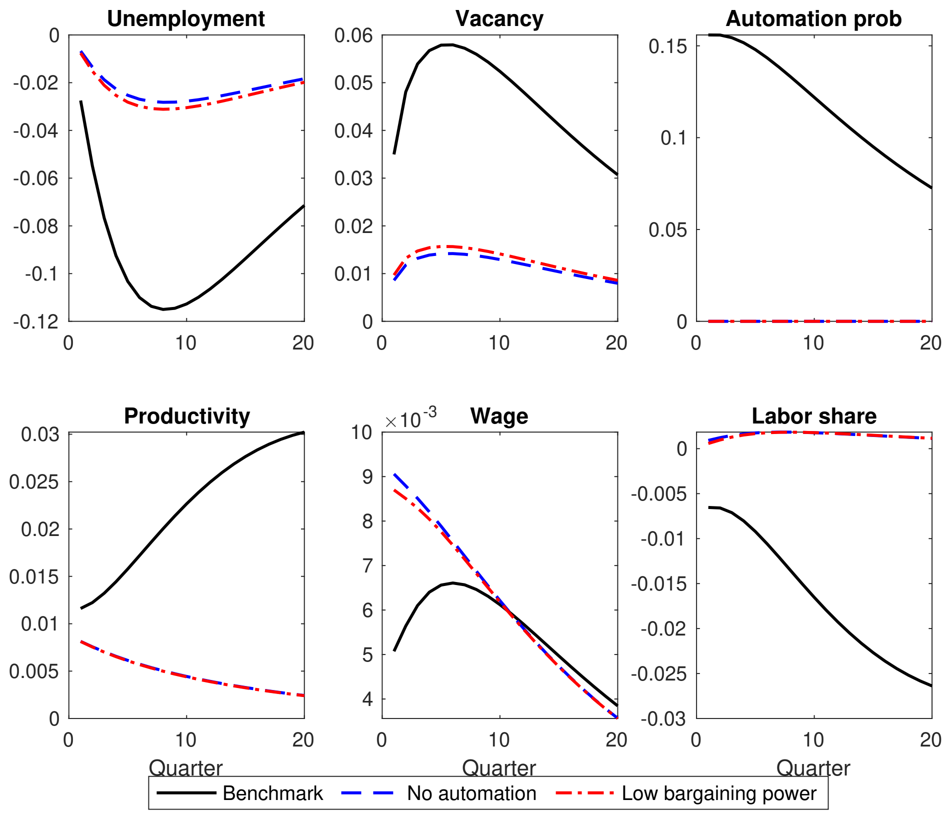

FIGURE 6. Impulse responses to a positive neutral technology shock in the benchmark model (black solid lines), the counterfactual with no automation (blue dashed lines), and the counterfactual with no automation and low worker bargaining power (red dot-dash lines).

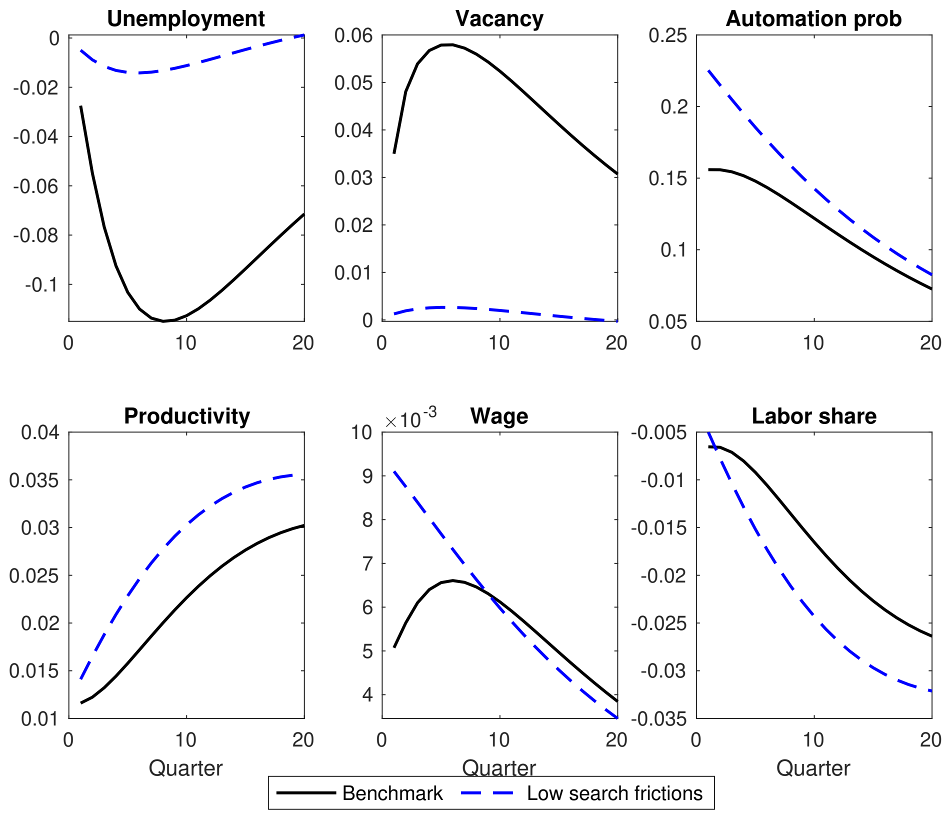

FIGURE 7. Impulse responses to a positive neutral technology shock in the benchmark model (black solid lines) and the counterfactual with low search frictions (blue dashed lines).

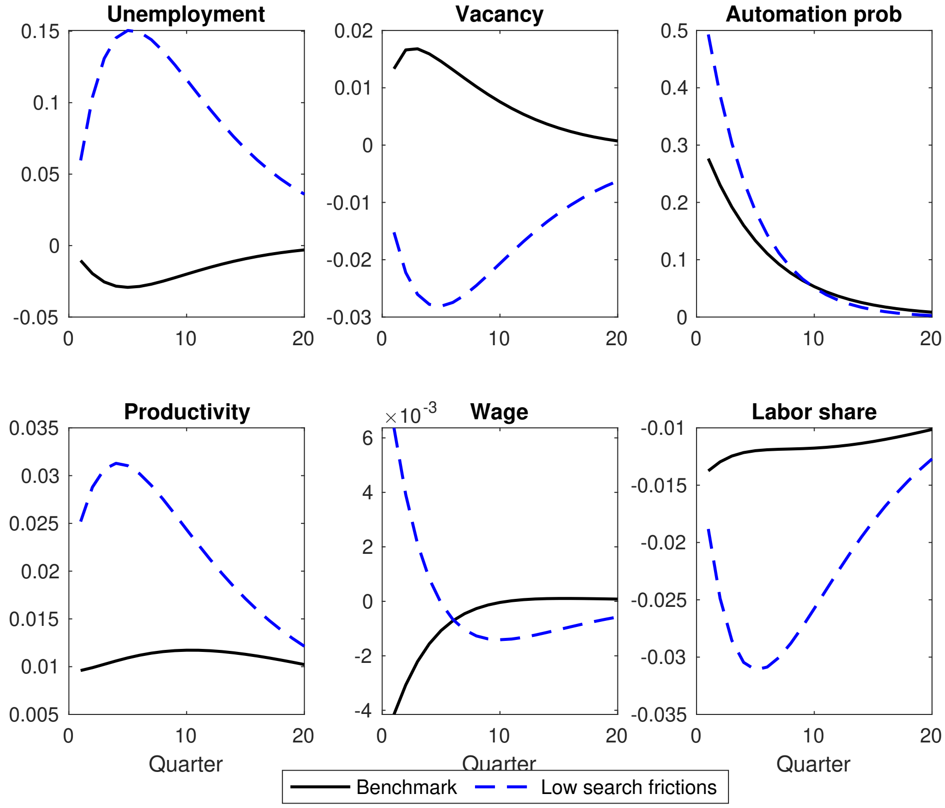

FIGURE 8. Impulse responses to a positive automation-specific shock in the benchmark model (black solid lines) and the counterfactual with low search frictions (blue dashed lines).

References

Acemoglu, D. and P. Restrepo (2017): "Robots and Jobs: Evidence from US Labor Markets," NBER Working Paper No. 23285.

——— (2018): "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, 108, 1488–1542.

Arnoud, A. (2018): "Automation Threat and Wage Bargaining," Unpublished Manuscript, Yale University.

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.

Barnichon, R. (2010): "Building a Composite Help-Wanted Index," *Economics Letters*, 109, 175–178.

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.

Christiano, L. J., M. S. Eichenbaum, and M. Trabandt (2020): "Why is Unemployment so Countercyclical?," Unpublished Manuscript, Northwestern University.

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.

Fernald, J. G. (2015): "Productivity and Potential Output Before, During, and After the Great Recession," *NBER Macroeconomics Annual 2014*, 29, 1–51.

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.

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.

Jaimovich, N., I. Saporta-Eksten, H. Siu, and Y. Yedid-Levi (2020): "The Macroeconomics of Automation: Data, Theory, and Policy Analysis," Unfinished Manuscript.

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 (2019): "Robots or Workers? A Macro Analysis of Automation and Labor Markets," Federal Reserve Bank of San Francisco Working Paper No. 2019-17.

Nedelkoska, L. and G. Quintini (2018): "Automation, Skills Use and Training," OECD Social, Employment and Migration Working Papers, No. 202.

Shimer, R. (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95, 25—49.

Zeira, J. (1998): "Workers, Machines, and Economic Growth," *Quarterly Journal of Economics*, 113, 1091–1117.

Appendix A. Data

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

(1) **Unemployment:** Civilian unemployment rate (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. Derivations of household's optimizing conditions

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

$$V_t(B_{t-1}, N_{t-1}) \equiv \max \ \ln C_t - \chi N_t + \beta \mathbb{E}_t \theta_{t+1} V_{t+1}(B_t, N_t), \tag{A1}$$

subject to the budget constraint

$$C_t + \frac{B_t}{r_t} = B_{t-1} + w_t N_t + \phi(1 - N_t) + d_t - T_t, \tag{A2}$$

and the law of motion for employment

$$N_t = (1 - \delta_t) N_{t-1} + q_t^u u_t, \tag{A3}$$

where the measure of job seekers is given by

$$u_t = 1 - (1 - \delta_t) N_{t-1}. \tag{A4}$$

The household chooses C_t, B_t, and N_t, taking prices and the average job finding rate as given.

Denote by Λ_t the Lagrangian multiplier for the budget constraint (A2). The first-order condition with respect to consumption implies that

$$\Lambda_t = \frac{1}{C_t}. \tag{A5}$$

The optimizing decision for B_t implies that

$$\frac{\Lambda_t}{r_t} = \beta \mathrm{E}_t \theta_{t+1} \frac{\partial V_{t+1}(B_t, N_t)}{\partial B_t}. \tag{A6}$$

The envelope condition with respect to B_{t-1} implies that

$$\frac{\partial V_t(B_{t-1}, N_{t-1})}{\partial B_{t-1}} = \Lambda_t. \tag{A7}$$

We thus obtain the intertemporal Euler equation

$$1 = \mathrm{E}_t \frac{\beta \theta_{t+1} \Lambda_{t+1}}{\Lambda_t} r_t, \tag{A8}$$

which is equation (24) in the text.

The envelope condition with respect to N_{t-1} implies that

$$\frac{\partial V_t(B_{t-1}, N_{t-1})}{\partial N_{t-1}} = \left[\Lambda_t(w_t - \phi) - \chi + \beta \mathrm{E}_t \theta_{t+1} \frac{\partial V_{t+1}(B_t, N_t)}{\partial N_t}\right] \frac{\partial N_t}{\partial N_{t-1}}. \tag{A9}$$

Equations (A3) and (A4) imply that

$$\frac{\partial N_t}{\partial N_{t-1}} = (1 - \delta_t)(1 - q_t^u) \tag{A10}$$

and that

$$\frac{\partial u_t}{\partial N_{t-1}} = -(1 - \delta_t). \tag{A11}$$

Define the employment surplus (i.e., the value of employment relative to unemployment) as

$$S_t^H = \frac{1}{\Lambda_t} \frac{\partial V_t(B_{t-1}, N_{t-1})}{\partial N_t} = \frac{1}{\Lambda_t} \frac{\partial V_t(B_{t-1}, N_{t-1})}{\partial N_{t-1}} \frac{\partial N_{t-1}}{\partial N_t} = \frac{1}{\Lambda_t} \frac{\partial V_t(B_{t-1}, N_{t-1})}{\partial N_{t-1}} \frac{1}{(1 - \delta_t)(1 - q^u(s_t))}. \tag{A12}$$

Thus, S_t^H is the value for the household to send an additional worker to work in period t. Then the envelope condition (A9) implies that

$$S_t^H = w_t - \phi - \frac{\chi}{\Lambda_t} + \mathrm{E}_t \frac{\beta \theta_{t+1} \Lambda_{t+1}}{\Lambda_t} (1 - \delta_{t+1})(1 - q_{t+1}^u) S_{t+1}^H. \tag{A13}$$

The employment surplus S_t^H derived here corresponds to equation (23) in the text, and it is the relevant surplus for the household in the Nash bargaining problem.

Appendix C. Summary of equilibrium conditions

A search equilibrium is a system of 18 equations for 18 variables summarized in the vector

$$[C_t, r_t, Y_t, m_t, u_t, v_t, q_t^u, q_t^v, q_t^a, N_t, U_t, \eta_t, J_t^e, J_t^v, J_t^a, A_t, x_t^*, w_t]\,.$$

We write the equations in the same order as in the dynare code.

(1) Household's bond Euler equation:

$$1 = \mathrm{E}_t \beta \theta_{t+1} \frac{C_t}{C_{t+1}} r_t, \tag{A14}$$

(2) Matching function

$$m_t = \mu u_t^{\alpha} v_t^{1-\alpha}, \tag{A15}$$

(3) Job finding rate

$$q_t^u = \frac{m_t}{u_t}, \tag{A16}$$

(4) Vacancy filling rate

$$q_t^v = \frac{m_t}{v_t}, \tag{A17}$$

(5) Employment dynamics

$$N_t = (1-\delta_t) N_{t-1} + m_t, \tag{A18}$$

(6) Number of searching workers

$$u_t = 1 - (1-\delta_t) N_{t-1}, \tag{A19}$$

(7) Unemployment

$$U_t = 1 - N_t, \tag{A20}$$

(8) Vacancy dynamics

$$v_t = (1 - q_{t-1}^v)(1 - q_t^a) v_{t-1} + \delta_t N_{t-1} + \eta_t, \tag{A21}$$

(9) Automation dynamics

$$A_t = (1-\rho^o) A_{t-1} + q_t^a (1 - q_{t-1}^v) v_{t-1}, \tag{A22}$$

(10) Employment value

$$J_t^e = Z_t - w_t + \mathbb{E}_t \beta \theta_{t+1} \frac{C_t}{C_{t+1}} \left[\delta_{t+1} J_{t+1}^v + (1-\delta_{t+1}) J_{t+1}^e\right], \tag{A23}$$

(11) Vacancy value

$$J_t^v = -\kappa + q_t^v J_t^e + (1 - q_t^v) \mathbb{E}_t \beta \theta_{t+1} \frac{C_t}{C_{t+1}} \left[(1 - q_{t+1}^a) J_{t+1}^v + q_{t+1}^a J_{t+1}^a\right]. \tag{A24}$$

(12) Automation value

$$J_t^a = Z_t\zeta_t - \kappa_a + (1-\rho^o)\mathbb{E}_t\beta\theta_{t+1}\frac{C_t}{C_{t+1}}J_{t+1}^a, \tag{A25}$$

(13) Automation threshold

$$x_t^* = J_t^a - J_t^v, \tag{A26}$$

(14) Robot adoption

$$q_t^a = \left(\frac{x_t^*}{\bar{x}}\right)^{\eta_a}, \tag{A27}$$

(15) Vacancy creation

$$\eta_t = \left(\frac{J_t^v}{\bar{e}}\right)^{\eta_e}, \tag{A28}$$

(16) Aggregate output

$$Y_t = Z_t N_t + Z_t\zeta_t A_t. \tag{A29}$$

(17) Resource constraint

$$C_t + \kappa v_t + \kappa_a A_t + \frac{\eta_a}{1+\eta_a}q_t^a x_t^*(1-q_{t-1}^v)v_{t-1} + \frac{\eta_e}{1+\eta_e}\eta_t J_t^v = Y_t, \tag{A30}$$

(18) Nash bargaining wage

$$\frac{b}{1-b}(J_t^e - J_t^v) = w_t - \phi - \chi C_t + \mathbb{E}_t\frac{\beta\theta_{t+1}C_t}{C_{t+1}}(1-q_{t+1}^u)(1-\delta_{t+1})\frac{b}{1-b}(J_{t+1}^e - J_{t+1}^v), \tag{A31}$$

Appendix D. Automating jobs instead of vacancies

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.

D.1. **Main ingredients in the model.** 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 by 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. (14)), 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

$$J_t^e = Z_t - w_t + \mathbb{E}_t\beta\theta_{t+1}\frac{C_t}{C_{t+1}}\left\{\delta_{t+1}J_{t+1}^v + (1-\delta_{t+1})\left[q_{t+1}^a J_{t+1}^a + (1-q_{t+1}^a)J_{t+1}^e\right]\right\}, \tag{A32}$$

A job match yields the flow profit $Z_t - w_t$ in period t. In period $t+1$, the job can be exogenously separated, in which case the firm obtains the vacancy value J^v_{t+1}. If the job is not separated, it can be automated with the probability q^a_{t+1}, in which case the firm obtains the automation value J^a_{t+1}. If the job is neither separated nor automated, then the firm obtains the continuation value of employment J^e_{t+1}.

Since a fraction of non-separated jobs are automated, the employment stock follows the law of motion

$$N_t = (1-\delta_t)(1-q^a_t)N_{t-1} + m_t. \tag{A33}$$

The system of equilibrium conditions is summarized in Appendix D.2.

The law of motion for employment (A33) reveals that, in this model setup, automation acts like a job separation shock. This intuition is confirmed by the impulse responses to a discount factor shock in Figure A1. The figure shows that a positive discount factor shock raises the net present value of automation and thus increases the probability of automation. Since automation directly replaces workers, the unemployment rate rises following the shock. At the same time, automation improves labor productivity and boosts employment and vacancy creation, offsetting its direct job-displacing effect. With estimated parameters and shocks in the model (using the same time-series data as in our benchmark case), the job-displacing effect dominates the employment boosting effect in the short run, so that a positive discount factor shock raises unemployment and vacancies, similar to the effects of an exogenous job separation shock.

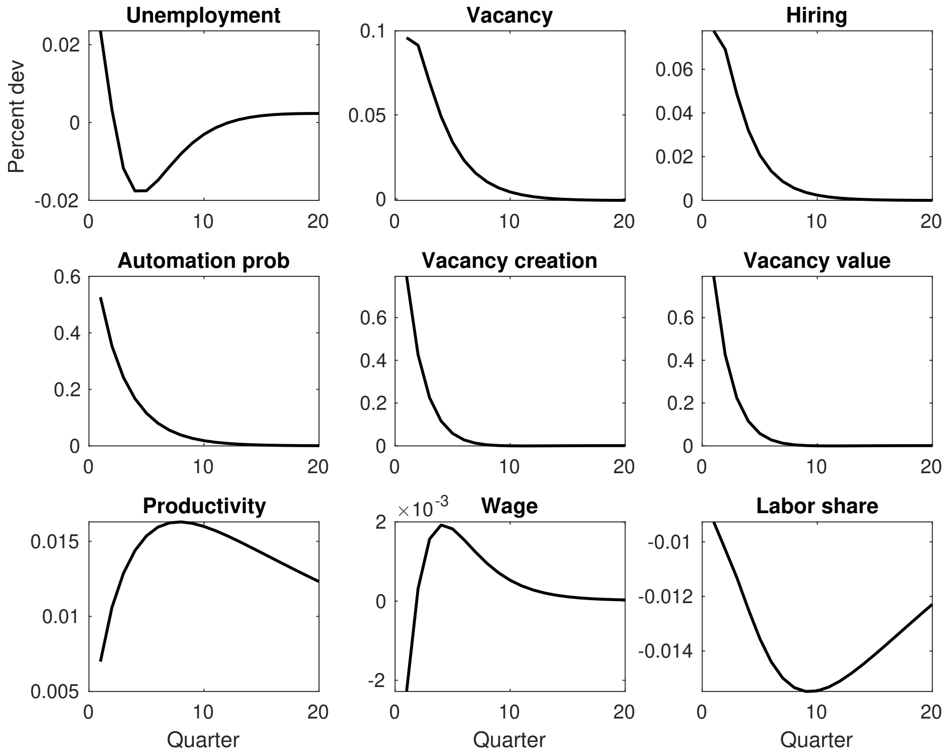

FIGURE A1. Impulse responses to a positive discount factor shock in the alternative model with automated jobs instead of vacancies.

D.2. **Equilibrium conditions in the model with automated jobs.** A search equilibrium is a system of 18 equations for 18 variables summarized in the vector

$$[C_t, r_t, Y_t, m_t, u_t, v_t, q_t^u, q_t^v, q_t^a, N_t, U_t, \eta_t, J_t^e, J_t^v, J_t^a, A_t, x_t^*, w_t] \, .$$

(1) Household's bond Euler equation:

$$1 = \mathrm{E}_t \beta \theta_{t+1} \frac{C_t}{C_{t+1}} r_t, \tag{A34}$$

(2) Matching function

$$m_t = \mu u_t^{\alpha} v_t^{1-\alpha}, \tag{A35}$$

(3) Job finding rate

$$q_t^u = \frac{m_t}{u_t}, \tag{A36}$$

(4) Vacancy filling rate

$$q_t^v = \frac{m_t}{v_t}, \tag{A37}$$

(5) Employment dynamics

$$N_t = (1 - \delta_t)(1 - q_t^a) N_{t-1} + m_t, \tag{A38}$$

(6) Number of searching workers

$$u_t = 1 - (1 - \delta_t)(1 - q_t^a)N_{t-1}, \tag{A39}$$

(7) Unemployment

$$U_t = 1 - N_t, \tag{A40}$$

(8) Vacancy dynamics

$$v_t = (1 - q_{t-1}^v)v_{t-1} + \delta_t N_{t-1} + \eta_t, \tag{A41}$$

(9) Automation dynamics

$$A_t = (1 - \rho^o)A_{t-1} + q_t^a(1 - \delta_t)N_{t-1}, \tag{A42}$$

(10) Employment value

$$J_t^e = Z_t - w_t + \mathbb{E}_t\beta\theta_{t+1}\frac{C_t}{C_{t+1}}\left\{\delta_{t+1}J_{t+1}^v + (1 - \delta_{t+1})\left[q_{t+1}^a J_{t+1}^a + (1 - q_{t+1}^a)J_{t+1}^e\right]\right\}, \tag{A43}$$

(11) Vacancy value

$$J_t^v = -\kappa + q_t^v J_t^e + (1 - q_t^v)\mathbb{E}_t\beta\theta_{t+1}\frac{C_t}{C_{t+1}}J_{t+1}^v, \tag{A44}$$

(12) Automation value

$$J_t^a = Z_t\zeta_t - \kappa_a + (1 - \rho^o)\mathbb{E}_t\beta\theta_{t+1}\frac{C_t}{C_{t+1}}J_{t+1}^a, \tag{A45}$$

(13) Automation threshold

$$x_t^* = J_t^a - J_t^e, \tag{A46}$$

(14) Robot adoption

$$q_t^a = \left(\frac{x_t^*}{\bar{x}}\right)^{\eta_a}, \tag{A47}$$

(15) Vacancy creation

$$\eta_t = \left(\frac{J_t^v}{\bar{e}}\right)^{\eta_e}, \tag{A48}$$

(16) Aggregate output

$$Y_t = Z_t N_t + Z_t\zeta_t A_t. \tag{A49}$$

(17) Resource constraint

$$C_t + \kappa v_t + \kappa_a A_t + \frac{\eta_a}{1 + \eta_a}q_t^a x_t^*(1 - \delta_t)N_{t-1} + \frac{\eta_e}{1 + \eta_e}\eta_t J_t^v = Y_t, \tag{A50}$$

(18) Nash bargaining wage

$$\frac{b}{1 - b}(J_t^e - J_t^v) = w_t - \phi - \chi C_t + \mathbb{E}_t\frac{\beta\theta_{t+1}C_t}{C_{t+1}}(1 - q_{t+1}^u)(1 - \delta_{t+1})\frac{b}{1 - b}(J_{t+1}^e - J_{t+1}^v), \tag{A51}$$