FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Robots or Workers?
A Macro Analysis of Automation and Labor Markets

Sylvain Leduc and Zheng Liu
Federal Reserve Bank of San Francisco

August 2019

Working Paper 2019-17

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

Suggested citation:

Leduc, Sylvain, Zheng Liu. 2019. "Robots or Workers? A Macro Analysis of Automation and Labor Markets," Federal Reserve Bank of San Francisco Working Paper 2019-17. https://doi.org/10.24148/wp2019-17

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

SYLVAIN LEDUC AND ZHENG LIU

Abstract. We study the implications of automation for labor market fluctuations in a Diamond-Mortensen-Pissarides (DMP) framework that is generalized to incorporate automation decisions. If a job opening is not filled with a worker, a firm can choose to automate that position and use a robot instead of a worker to produce output. The threat of automation strengthens the firm's bargaining power against job seekers in wage negotiations, depressing equilibrium real wages in a business cycle boom. The option of automation also increases the value of a vacancy, raising the incentive for job creation, and thereby amplifying fluctuations in vacancies and unemployment relative to the standard DMP framework. Since automation improves labor productivity while muting wage increases, it implies a countercyclical labor income share, as observed in the data.

I. Introduction

Recent developments in robotics and artificial intelligence have renewed concerns that robots could render workers redundant, possibly leading to what Keynes (1930) called "technological unemployment." Stagnant wage growth and recent declines in the labor share of income have also heightened fears that automation's growing importance may have weakened workers' bargaining power. However, the notion of machines replacing workers is an oversimplification of the macroeconomic impact of automation (Autor, 2015). While robots are increasingly performing standardized tasks that workers previously performed, new tasks for which workers have a comparative advantage are being created (Acemoglu and Restrepo, 2018). As a result, the impact of automation on employment, wages, and the labor share can be ambiguous, and assessing such impacts requires a coherent quantitative general equilibrium framework.

Date: August 29, 2019.

Key words and phrases. Robots, threat of automation, unemployment, wage bargaining, productivity, labor share, business cycles.

JEL classification: E24, E25, O33, J64.

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. We thank Nicolas Petrosky-Nadeau, participants at the West Coast Search and Matching Workshop and the HKUST Macro workshop, as well as seminar participants at the Hong Kong University of Science and Technology and 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.

In this paper, we develop such a framework to examine the interactions between automation and the labor market over the business cycle. Our general equilibrium model features a labor market with search frictions and endogenous automation decisions. We study the mechanism through which automation—either adopting a robot to perform a job or the threat of doing so—can drive changes in unemployment, vacancies, real wages, and labor productivity. To assess the quantitative importance of the automation mechanism, we estimate the model to fit quarterly U.S. time series. We then compare the labor market implications of the estimated model to those of a counterfactual version of the model with no changes in automation decisions.

We build on the standard Diamond-Mortensen-Pissarides (DMP) model with labor market search frictions and generalize it to incorporate automation decisions. In our model, consumption goods can be produced using either a worker or a robot. If a job position is filled by a worker, the firm obtains the employment value. If not, the firm can choose to automate the open position subject to a fixed cost drawn from an i.i.d. distribution. If the firm draws an automation cost that lies below a threshold determined by the expected benefit of automation, then the position is automated, in which case the firm uses a robot instead of a worker to produce output. Thus, the probability of automation is the cumulative density of the automation cost draws evaluated at the automation threshold. If the unfilled position is not automated because the cost exceeds the threshold, the firm keeps the position open and receives the continuation value of the vacancy, including the option to automate the position in future periods.[1]

Our approach to modeling automation decisions requires a job vacancy to carry a positive value in equilibrium. In the standard DMP model, however, free entry implies a zero value of an unfilled vacancy. We thus introduce a fixed cost of vacancy creation in our model. 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, unlike 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).

[1]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. This approach simplifies our analysis significantly. Acemoglu and Restrepo (2018) use an alternative framework to study automation. Building on the earlier work of Zeira (1998), they consider a job consisting of a continuum of tasks, a fraction of which are technologically automatable, while the other tasks need to be performed by human workers (see also Autor and Salomons (2018)). In our model, automation is an endogenous decision.

Our model yields a few theoretical insights. First, 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 job creation. 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 unemployment fluctuations.

Second, 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. In contrast, in a standard model with a spot labor market, automation would boost wages by raising productivity.

Third, increased automation in a boom raises aggregate productivity, further fueling the expansion. By dampening wage increases while boosting labor productivity, automation contributes to a decline in the labor income share in a business cycle expansion.

To assess the quantitative importance of our 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. 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. Since we fit the model to match the growth rates of real wages and labor productivity, the model also fits the observed labor share by construction.

We find that automation plays an important role behind the decline in the labor share, particularly during the recovery from the Great Recession. In a counterfactual version of our model with the probability of automation kept constant at its steady-state value, the predicted path of the labor share is relatively flat, in contrast to the observed sharp decline since the early 2000s. In addition, absent automation and its threat, wages would have been higher during this period, while productivity growth would have been even slower.

Our model's prediction that increased automation has contributed significantly to the declines in the labor share in recent years is consistent with the evidence presented by Autor and Salomons (2018). Using data on 28 industries across 18 OECD countries between 1970 and 2007, Autor and Salomons (2018) find that automation has substantially reduced the labor share since the early 2000s, although its labor-share displacing effect was much more muted in earlier periods (see also Acemoglu and Restrepo (2019)).[2]

Automation also helps generate large fluctuations in unemployment and vacancies. The threat of automation dampens wage adjustments relative to unemployment and vacancies, giving rise to a source of endogenous real wage rigidities, which are important for amplifying labor market fluctuations (Shimer, 2005). In addition, automation raises aggregate productivity in a business cycle boom, further fueling the boom. This mechanism is quantitatively important. In our estimated model, 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.[3] In contrast, a counterfactual model without automation produces a volatility ratio of about 9, less than 25 percent of that predicted by our estimated model.

We further show that the effects of automation on the labor market cannot simply be replicated in a counterfactual exercise without automation but with a lower bargaining power for the workers. 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 in our framework. In addition, the implied dynamics of the labor share 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, an increase in labor productivity, and thus a decline in the labor share. The counterfactual model without automation and with a lower workers' bargaining weight predicts a small increase in the real wage, no change in labor productivity, and thus an increase in the labor share. This difference reflects the threat of automation on wage bargaining and the endogenous productivity enhancement from automation.

Our model predicts that automation dampens wage growth, boosts labor productivity, and reduces the labor share. These predictions are supported by independent micro-level evidence in recent studies. For example, Dinlersoz and Wolf (2018) use plant-level data from the 1991 U.S. Census Bureau's Survey of Manufacturing Technology and document evidence that more-automated establishments have a smaller fraction of high-wage workers, higher labor productivity, and a smaller labor share in production. Furthermore, plants with higher investment in automation have experienced larger declines in labor shares. Arnoud (2018) uses data from the U.S. Current Population Survey in 2013 and an index of automatability developed in the engineering literature (Frey and Osborne, 2017) to examine occupation-level relations between the threat of automation and wage adjustments. He finds that, controlling

[2]Because their sample ends in 2007, Autor and Salomons (2018) are silent on the impact of automation on the labor share following the Great Recession.

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

for observable characteristics, occupations that are more susceptible to automation have experienced lower wage growth, in line with our model's mechanism. Acemoglu and Restrepo (2017) study U.S. county-level data and report a negative effect of robot adoptions on local employment and wages. 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.

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 all the general equilibrium effect, although it does not directly speak to heterogeneous effects across different types of jobs or workers.

Our work adds to a growing strand of literature that attempts to explain the declines in the labor share. Karabarbounis and Neiman (2013) document a general decline in the labor share across 59 countries between 1975 and 2012. They argue that the secular falls in the relative price of investment goods induced firms to substitute capital for labor, contributing to about half of the observed declines in the labor share since the 1980s. Elsby et al. (2013) examine the role of offshoring for the decline in the U.S. labor share. They find that industries facing greater import competition, particularly from China, experienced a larger decline in the labor share. Autor et al. (2017) attribute the declines in the labor share to the rise of superstar firms, based on a panel of firm-level data from the U.S. Census since 1982. Krueger (2018) emphasizes the growing importance of firms' monopsony power in the labor market (in the forms of implicit collusions, noncompete restrictions in labor contracts, and outsourcing) as a factor that diminishes workers' bargaining power. Automation can also lead to reallocation of workers across occupations and can thus contribute to job polarization (Cortes et al., 2017).

Our model highlights a different mechanism for explaining the declines in the labor share. An increase in the threat of automation in a business cycle boom suppresses wage increases while boosting labor productivity, reducing the labor share. To our knowledge, our model is the first quantitative general equilibrium model that incorporates automation into a DMP framework to study interactions between automation 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 post vacancies v_t at a fixed cost κ. The stock of vacancies v_t includes the unfilled vacancies that were not automated at the end of period $t-1$, the jobs separated in the beginning of period t, and new vacancies created in the beginning of period t. Creating a new vacancy incurs a fixed cost, which is drawn from an i.i.d. distribution $G(\cdot)$, as in Leduc and Liu (2019). In the labor market, a matching technology transforms job seekers and vacancies into an employment relation, with a wage rate determined through Nash bargaining between the employer and the job seeker. Once an employment relation is formed, production takes place, and the firm receives the employment value. An unfilled vacancy can be either carried forward to the next period or automated at a fixed cost. Similar to the vacancy creation cost, the automation cost x is drawn from an i.i.d. distribution $F(x)$. If a firm draws an automation cost that is below a threshold value x_t^*, then the firm adopts a robot and closes the job opening. In that case, the firm obtains the automation value. Otherwise, the vacancy remains open and the firm receives the continuation value of the vacancy. Newly adopted robots add to the stock of automation, which becomes obsolete over time at a constant rate ρ^o. Final goods output is the sum of the goods produced by workers and by robots. The final good is used for household consumption and also for paying the costs of vacancy posting, vacancy creation, and robot adoption.

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

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

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 in the beginning of period t consists of the vacancies in period $t-1$ that were not filled with workers and not automated, plus the separated employment matches and newly created vacancies. The law of motion for vacancies is given by

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

where q^v_{t-1} denotes the job filling rate in period $t-1$, q^a_{t-1} denotes the automation rate in period $t-1$, 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.** If a firm successfully hires a worker, then it can produce Z_t units of intermediate goods. The technology shock Z_t follows the stochastic process

$$\ln Z_t = (1 - \rho_z) \ln \bar{Z} + \rho_z \ln Z_{t-1} + \varepsilon_{zt}. \tag{9}$$

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.[4]

The value of employment satisfies the Bellman equation

$$J^e_t = Z_t - w_t + \mathbb{E}_t D_{t,t+1} \left\{ (1 - \delta_{t+1}) J^e_{t+1} + \delta_{t+1} J^v_{t+1} \right\}, \tag{10}$$

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

where w_t denotes the real wage rate and $D_{t,t+1}$ is a stochastic discount factor of the households. 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^v_{t+1}. Otherwise, the firm receives the continuation value of employment.

Following Leduc and Liu (2019), 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^v_t. Thus, the number of new vacancies η_t is given by the cumulative density of the entry costs evaluated at J^v_t. That is,

$$\eta_t = F(J^v_t). \tag{11}$$

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^v_t), the firm obtains the employment value J^e_t. If the vacancy is not filled, then the firm can choose to automate the position and close the vacancy (with the probability q^a_t), in which case the firm obtains the automation value J^a_t. If the firm does not automate the unfilled position, then it receives the continuation value of the vacancy. Thus, the vacancy value satisfies the Bellman equation

$$J^v_t = -\kappa + q^v_t J^e_t + (1 - q^v_t) q^a_t J^a_t + (1 - q^v_t)(1 - q^a_t)\mathbb{E}_t D_{t,t+1} J^v_{t+1}. \tag{12}$$

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

$$A_t = (1 - \rho^o)A_{t-1} + q^a_t(1 - q^v_t)v_t, \tag{13}$$

where $q^a_t(1-q^v_t)v_t$ is the number of unfilled job positions that are newly automated in period t.

Once adopted, a robot produces $Z_t\zeta_t$ units of output, where ζ_t denotes an automation-specific technology shock, which follows a stochastic process that is independent of the neutral technology shock Z_t. In particular, ζ_t follows the stationary process

$$\ln \zeta_t = (1 - \rho_\zeta) \ln \bar{\zeta} + \rho_\zeta \ln \zeta_{t-1} + \varepsilon_{\zeta t}. \tag{14}$$

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^2_ζ. 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

$$J^a_t = Z_t\zeta_t - \kappa_a + (1 - \rho^o)\mathbb{E}_t D_{t,t+1} J^a_{t+1}, \tag{15}$$

where the term κ_a captures the costs of energy, facilities, and space for automated production.

Automating a vacancy 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 $x \leq x_t^*$. The threshold value x_t^* depends on the value of automation J_t^a relative to the continuation value of a vacancy. In particular, the threshold for automation decision is given by

$$x_t^* = J_t^a - \mathbb{E}_t D_{t,t+1} J_{t+1}^v. \tag{16}$$

Thus, the probability of automation is the cumulative density of the automation costs evaluated at x_t^*. More formally, the automation probability is determined by

$$q_t^a = G(x_t^*). \tag{17}$$

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

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

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

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

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

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

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.[5] 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.[6] We estimate the scale parameters $\bar{e}$ and $\bar{x}$ and the shock processes by fitting the model to U.S. time series data.

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 $U = 0.058$, 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 $\mu = \frac{m}{u^\alpha v^{1-\alpha}} = 0.6629$. We set the flow cost of operating robots to $\kappa_a = 0.98$. Given the average productivities $\bar{Z} = \bar{\zeta} = 1$,

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

[6] 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.

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 (15).

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 (11), the automation adoption condition (16), 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 automation A can be solved from the law of motion (13), 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 $\kappa = 0.0819$, 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 (12), and the equilibrium real wage rate can be obtained from the Bellman equation for employment (10). Steady-state consumption is solved from the resource constraint (31). We then infer the value of $\chi = 0.7261$ 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 $\bar{e} = 8.1044$. The posterior mean estimates of the automation cost parameter is $\bar{x} = 4.1579$. These parameters imply a steady-state share of output produced by automation of $A/Y = 0.13$. Thus, our model implies that, in the long run, about 13 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.

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.* We use the impulse responses of several key labor market variables to each shock to illustrate the model's transmission mechanism. To highlight the role of automation in the model's transmission, we compare the impulse responses from our estimated benchmark model to those from a counterfactual model without automation.

Figure 1 shows the impulse responses 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.[7] 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 55 percent of the variances of unemployment and vacancies, while the discount factor shock accounts for about 40 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 much 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 nearly 90 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

[7] 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.

and thus magnifies the impact of the neutral and 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 42 percent of the variance of productivity growth and 37 percent of that of real wage growth. Perhaps unsurprisingly, the neutral technology shock is also important for fluctuations in labor productivity, explaining about half of its variance.[8] In addition, about 60 percent of the real wage fluctuations are accounted for by the neutral and discount factor shocks.

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 "no automation" case, in which we reduce the bargaining weight for workers by a half (i.e., setting $b = 0.25$).

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 small 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.

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

In the no-automation model, reducing workers' bargaining weight mechanically dampens real wage adjustments and thus helps 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 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.

To assess the quantitative magnitude of amplification stemming from the automation channel, we compute the unconditional volatility of labor market tightness (measured by the v-u ratio) relative to the unconditional volatility of the real wage rate. This volatility ratio is about 40 in our benchmark model. In contrast, the counterfactual model without automation implies a relative volatility of about 9. Thus, without the automation channel, the standard DMP model has difficulties in generating large volatilities in the labor market, in line with the well-known Shimer puzzle. In this sense, the automation channel provides a quantitatively important amplification mechanism that helps explain the observed labor market dynamics.

IV.3. **Automation threat and labor market dynamics.** Our model predicts that the threat of automation dampens wage adjustments and amplifies labor market fluctuations. By raising labor productivity while muting wage responses, automation also implies a countercyclical labor income share. Is the automation mechanism quantitatively important? To examine the empirical importance of the automation mechanism, we compare our model's predictions for labor productivity, real wages, and the labor share with those from counterfactuals without the automation channel. As in the previous section, this no-automation counterfactual is a version of our benchmark model with all the automation-related variables held constant at their steady-state values and without automation-specific shocks.

Our benchmark model implies that the probability of automation is procylical, rising in business cycle booms and falling in recessions. This procyclical behavior can be seen in Figure 7, which shows the smoothed series of the automation probability obtained from the estimated benchmark model.

Procyclical automation boosts labor productivity and mutes wage increases in a business cycle boom. Figure 8 shows the labor productivity implied by our estimated model (the blue solid line) and that from a counterfactual model without automation (the red dashed line). Since the mid-2000s, increased automation has boosted labor productivity, with the effects becoming more pronounced after the Great Recession.[9] At the same time, increased automation probability in a business cycle boom raises the firm's reservation value in wage bargaining and dampens wage increases. This wage-suppressing effect of automation is quantitatively significant, especially after the Great Recession, as shown in Figure 9.

The automation channel boosts labor productivity and depresses real wages, leading to countercyclical labor share fluctuations. Figure 10 shows the paths of the labor income share implied by our benchmark model (the blue solid line) and that from the counterfactual model with no automation (the red dashed line). Since we estimate the benchmark model to fit both labor productivity growth and real wage growth, we can recover the labor income share from the smoothed series of real wages and labor productivity. We follow the same procedure to calculate the labor share in the counterfactual model with no automation. The figure shows that automation played a quantitatively important role in explaining the observed labor share fluctuations. The actual labor share has declined sharply since the early 2000s, and the pace of declines accelerated during the Great Recession. After the recession, the labor share has remained at historically low levels. Without automation, the counterfactual path of the labor share trended down together with the actual labor share until 2009, although it missed the short-run fluctuations.[10] After 2009, however, the two paths diverged: while

[9]The actual labor productivity growth in the post-recession period has been significantly slower than that observed in the late 1990s and 2000s (Fernald, 2015). From the lens of our model, this slowdown in labor productivity growth is driven by persistent negative shocks to the neutral technology, partly offset by positive automation shocks. Had there been no increases in automation, labor productivity growth would have been even slower since the mid-2000s.

[10]Since we demean the growth rates in the data, we need to add back the mean growth rates of productivity and wages and then convert the growth rates into levels (using the actual initial observations in the data). The labor share corresponds to the ratio of the real wage level to the labor productivity level. Since the productivity series and the real wage series are both indexed in our data (with 2012 as the base year), we convert the index number of the labor share to percentage terms. To do this, we first compute the growth

the actual labor share fell sharply, the counterfactual model predicts that the labor share in 2018 would have stayed flat, not very different from the pre-recession level. The divergence between the two paths of the labor share after the Great Recession highlights the importance of the automation channel in explaining the observed labor share dynamics.

V. Robustness

We now examine the robustness of our quantitative results to alternative measurements and model specifications.

V.1. **Measurement of labor income.** We have estimated our model using data on real wage and labor productivity in the nonfarm business sector (NFBS). However, the NFBS measure of the labor income share can be potentially biased because of imputation problems related to self-employment income.

To address this issue, we follow Karabarbounis and Neiman (2013) and and measure the real wage rate and labor productivity using data in the non-financial corporate sector (see also Elsby et al. (2013)). That sector consists of large corporations and thus does not have any imputation issues for self-employment income. The corporate sector is also large, accounting for about 52 percent of U.S. total value added in our sample period.

We re-estimate our benchmark model to fit the quarterly time series of unemployment, vacancies, and the growth rates of real wages and labor productivity in the non-financial corporate sector. As in our benchmark estimation, we focus on the sample period from 1985:Q1 to 2018:Q4.

Figure 11 shows the smoothed series of the labor share from the estimated benchmark model (the blue solid line), which coincides with the actual data (by construction). Compared to the NFBS labor share shown in Figure 10, the corporate-sector labor share has declined more sharply since the early 2000s, from about 63 percent in 2001 to about 53 percent in 2014. Unlike the NFBS labor share, which stayed roughly flat after the Great Recession, the corporate-sector labor share has rebounded since 2014 to a level of about 55.5 percent by the end of 2018.

To assess the importance of the automation channel in explaining the observed dynamics in the corporate-sector labor share, Figure 11 also plots the path of the labor share implied by the counterfactual no-automation model (the red dashed line). The no-automation model predicts a much smaller decline in the corporate-sector labor share in the early 2000s, and a modest dip of the labor share during the Great Recession, but a rapid recovery after the recession. Thus, without the automation channel, the model fails to generate the large

rates of the labor share index, and then use the observed labor share in percentage terms at the beginning of our sample (1985:Q1) as the initial value to convert the labor share growth rates into levels.

fluctuations in the labor share observed in the data, consistent with the results obtained in our benchmark model estimated with NFBS data.

V.2. **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. We describe the main ingredients in the alternative model here and relegate the details to Appendix D.

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. (15)), 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

$$J_t^e = Z_t - w_t + \mathbb{E}_t \beta \theta_{t+1} \frac{C_t}{C_{t+1}} \left\{ \delta_{t+1} J_{t+1}^v + (1 - \delta_{t+1}) \left[q_{t+1}^a J_{t+1}^a + (1 - q_{t+1}^a) J_{t+1}^e \right] \right\}, \tag{39}$$

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_{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. 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 non-separated jobs are automated, the employment stock follows the law of motion

$$N_t = (1 - \delta_t)(1 - q_t^a) N_{t-1} + m_t. \tag{40}$$

The remaining equilibrium conditions are straightforward to derive (see Appendix D).

The law of motion for employment (40) 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 12. 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.

Since automation implies counterfactual comovements between unemployment and vacancies, the model with automated jobs underperforms our baseline model with automated vacancies in fitting the time-series data. The marginal log data density of the model with automated jobs is about 1026.17, much smaller than that of our baseline model (1226.57). Thus, the data strongly prefer our baseline model.

VI. Conclusion

We have studied how automation interacts with labor market dynamics in a DMP framework that incorporates endogenous automation decisions, emphasizing the linkages between automation and bargaining power. The option to automate a job position induces a job-creation incentive, which offsets the direct job-displacing effects of automation, and can thus amplify unemployment fluctuations. The procyclical threat of automation raises the firm's reservation value in wage bargaining, dampening increases in real wages in a business cycle boom, while also amplifying fluctuations in unemployment and vacancies. Furthermore, actual adoptions of robots raise labor productivity. With muted wage increases and amplified labor productivity, automation leads to a countercyclical labor share, as observed in the data.

To assess the quantitative importance of the automation channel, we estimate our DMP model with automation to fit the U.S. time series data. We find that automation contributed significantly to the observed sluggish wage growth and the rapid declines in the labor share since the Great Recession.

Our focus on automation does not necessarily preclude other factors that may have contributed to the observed changes in wage and labor-share dynamics in the past two decades. Examples include increases in market concentration, declines in union power, and greater use of offshoring, particular since China's entry into the World Trade Organization. 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 first 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	0.6629
$\bar{\delta}$	Job separation rate	0.10
ρ^o	Automation obsolescence rate	0.03
κ	Vacancy posting cost	0.0819
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	0.7261
$\bar{Z}$	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]	8.1044	6.1781	9.9046
$\bar{x}$	scale for robot adoption cost	G	[5, 1]	4.1579	2.8365	5.4505
ρ_z	AR(1) of neutral technology shock	B	[0.8, 0.1]	0.9768	0.9620	0.9915
ρ_θ	AR(1) of discount factor shock	B	[0.8, 0.1]	0.9776	0.9620	0.9946
ρ_δ	AR(1) of separation shock	B	[0.8, 0.1]	0.9400	0.9101	0.9733
ρ_ζ	AR(1) of automation-specific shock	B	[0.8, 0.1]	0.7207	0.6857	0.7620
σ_z	std of tech shock	IG	[0.01, 0.1]	0.0116	0.0104	0.0128
σ_θ	std of discount factor shock	IG	[0.01, 0.1]	0.0190	0.0112	0.0266
σ_δ	std of separation shock	IG	[0.01, 0.1]	0.0472	0.0417	0.0515
σ_ζ	std of automation-specific shock	IG	[0.01, 0.1]	0.0803	0.0501	0.1064

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	57.10	41.96	0.51	0.44
Vacancy	53.58	39.29	6.54	0.59
Productivity growth	48.60	9.08	0.10	42.23
Real wage growth	39.72	22.64	0.61	37.03

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.

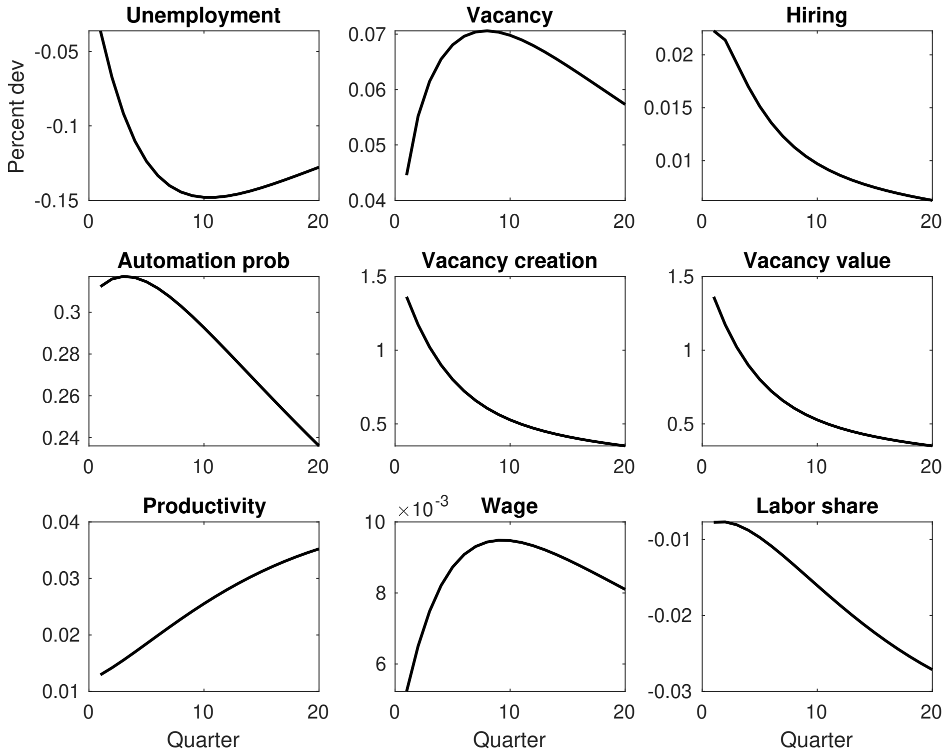

Figure 1. Impulse responses to a positive neutral technology shock in the benchmark model.

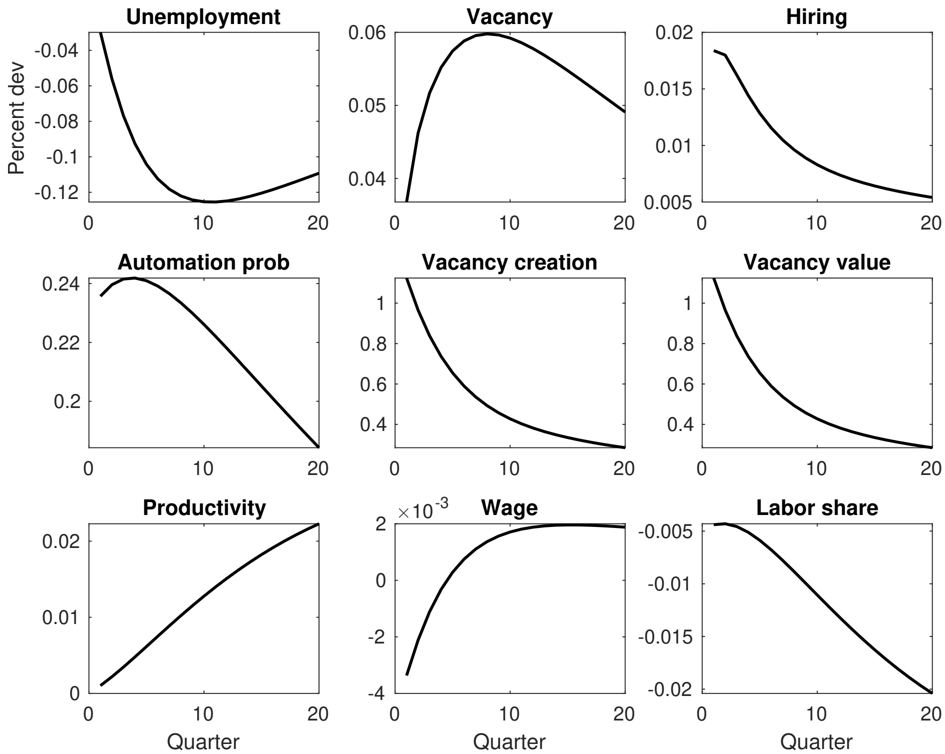

FIGURE 2. Impulse responses to a positive discount factor shock in the benchmark model.

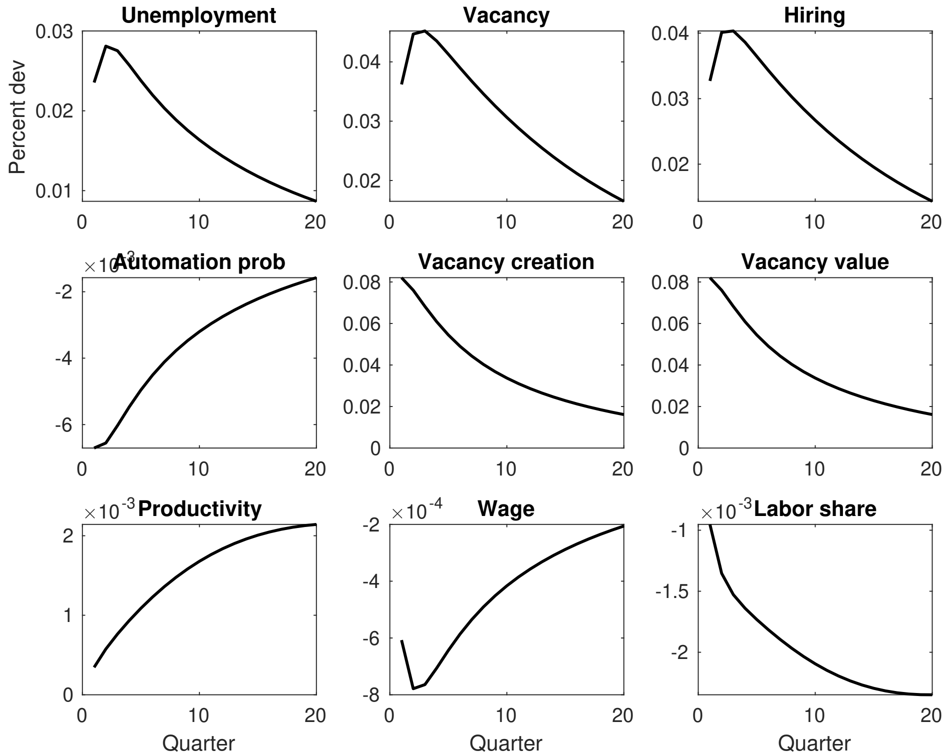

FIGURE 3. Impulse responses to a job separation shock in the benchmark model.

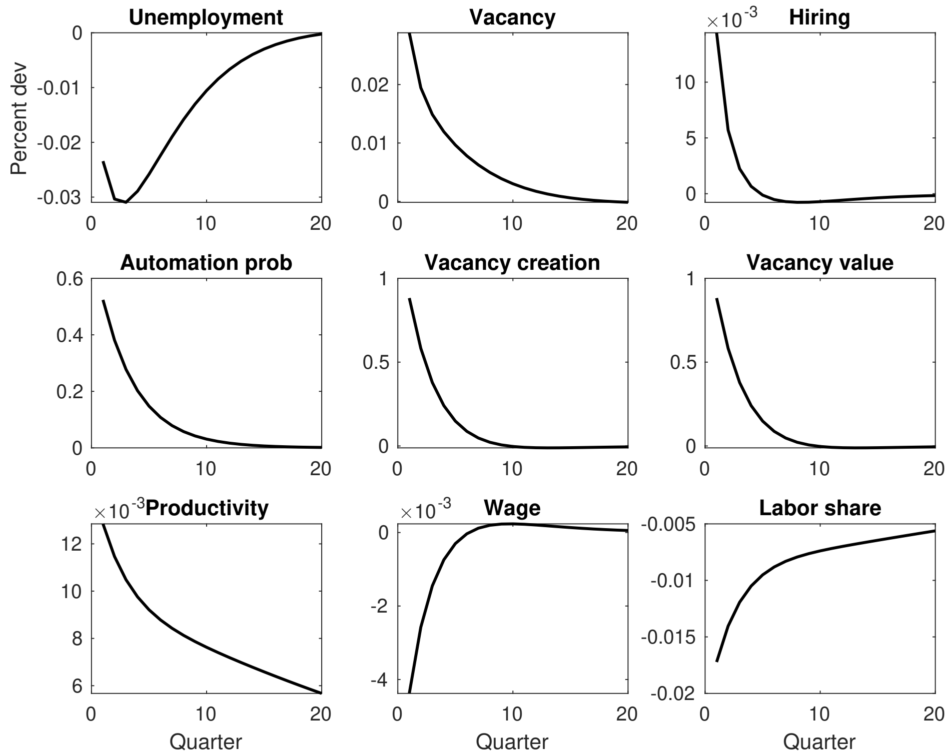

FIGURE 4. Impulse responses to a positive automation-specific shock in the benchmark model.

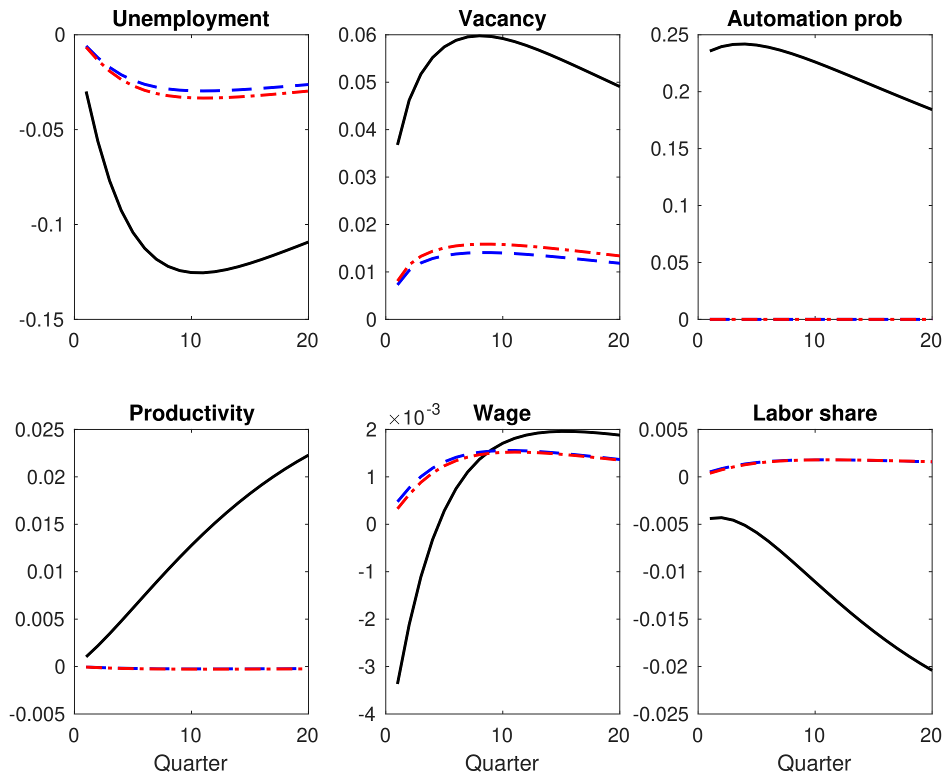

FIGURE 5. Impulse responses to a positive discount factor shock in the benchmark model (black solid lines), the counterfactual with no automation (blue dashed lines), and the counterfactual with no automation and low worker bargaining power (red dot-dash lines).

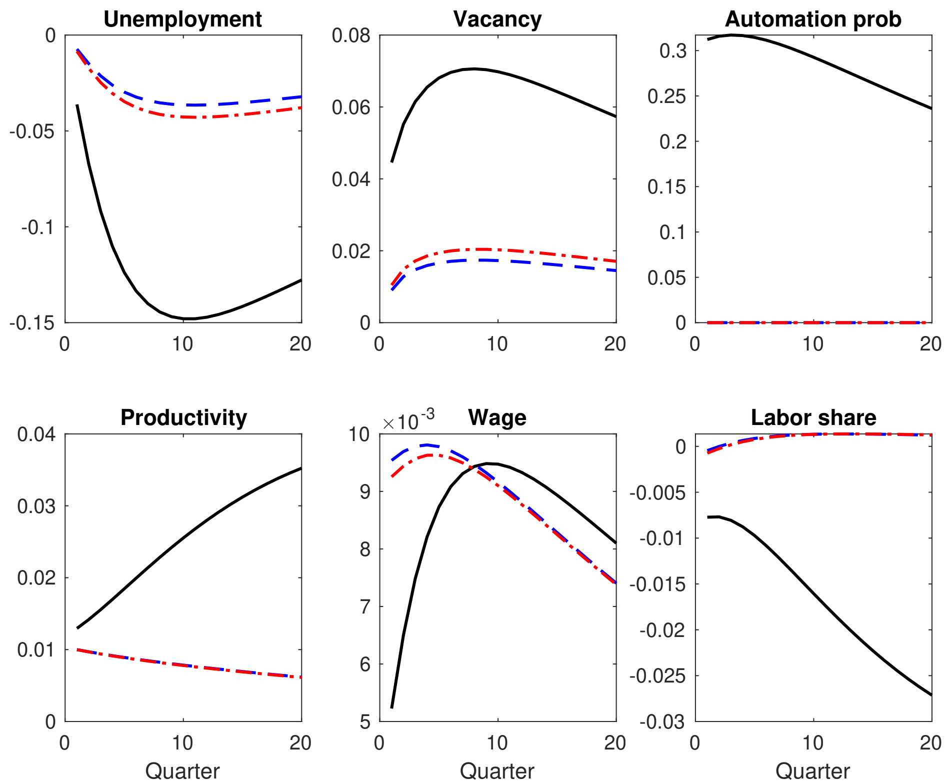

FIGURE 6. Impulse responses to a positive neutral technology shock in the benchmark model (black solid lines), the counterfactual with no automation (blue dashed lines), and the counterfactual with no automation and low worker bargaining power (red dot-dash lines).

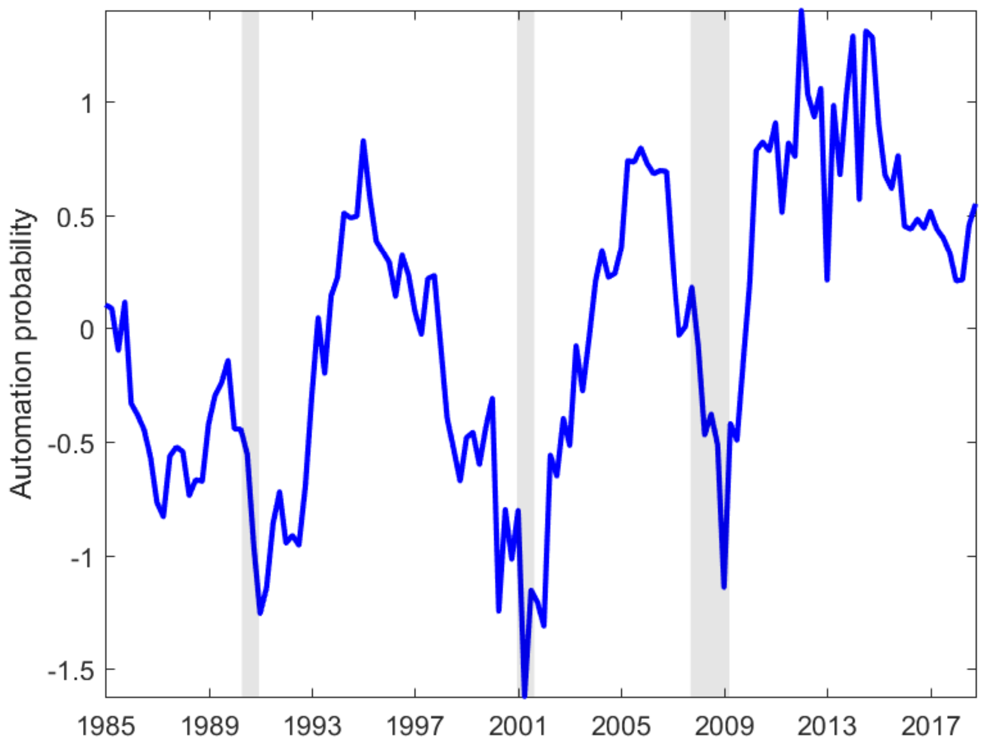

FIGURE 7. The smoothed time series of the automation probability.

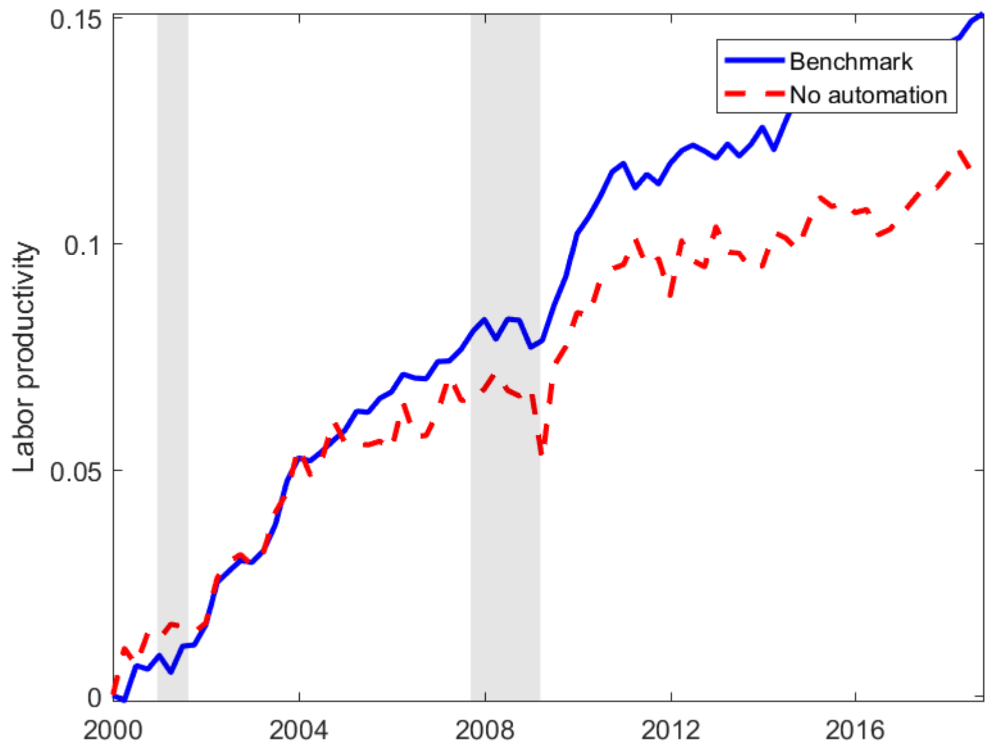

Figure 8. The smoothed time series of labor productivity in the benchmark model (the blue solid line) versus the counterfactual with no automation (the red dashed line). The counterfactual is a version of our benchmark model with all automation related variables (including automation specific shocks) kept at their steady-state values. The labor productivity levels are computed based on the smoothed labor productivity growth series in each model (with the initial value fixed at the actual productivity level in the beginning of the sample).

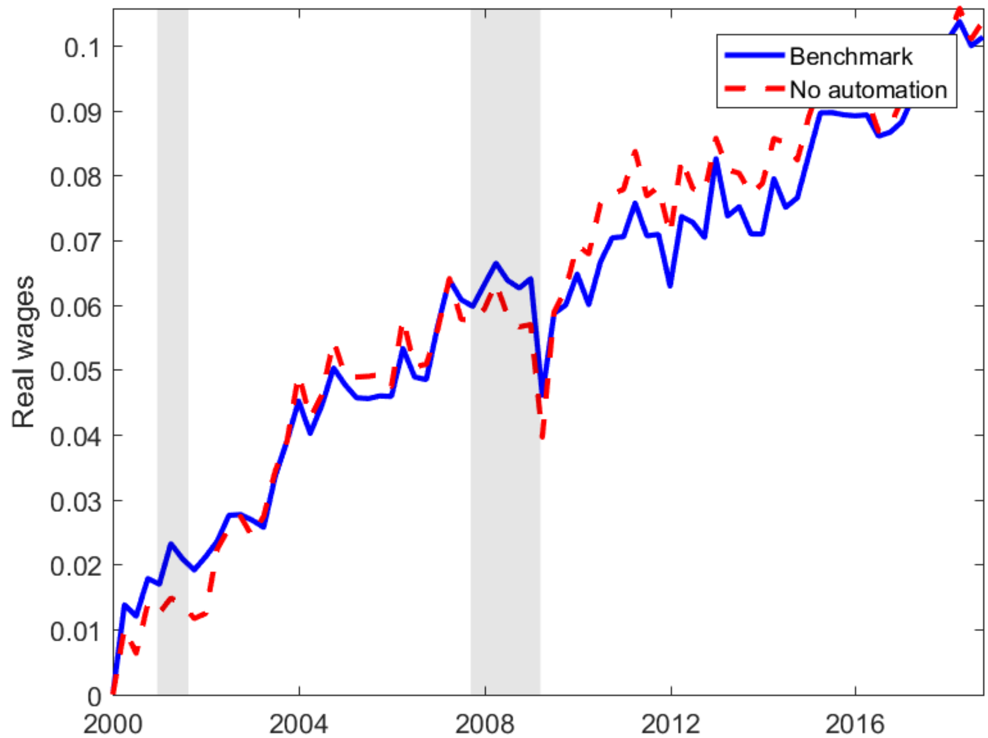

FIGURE 9. The smoothed time series of real wages in the benchmark model (blue solid line) versus the counterfactual with no automation (red dashed line). The counterfactual is a version of our benchmark model with all automation-related variables (including automation-specific shocks) kept at their steady-state values. The real wage levels are computed based on the smoothed real wage growth series in each model (with the initial value fixed at the actual real wage level in the beginning of the sample).

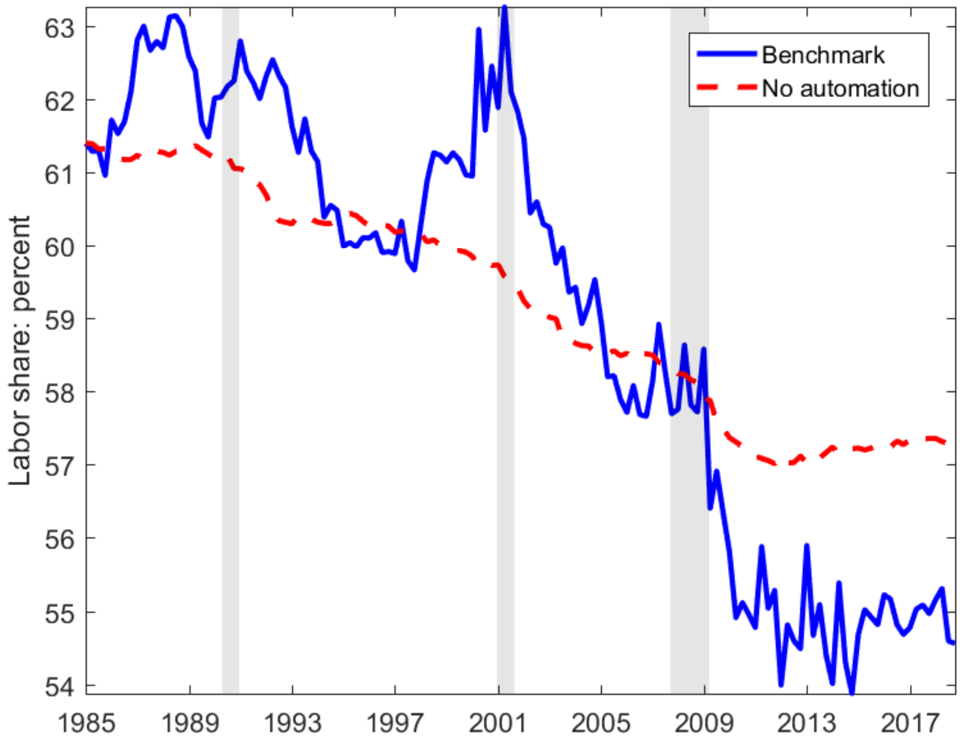

FIGURE 10. The smoothed time series of the labor share in the benchmark model (blue solid line) versus the counterfactual with no automation (red dashed line). The counterfactual is a version of our benchmark model with all automation-related variables (including automation-specific shocks) kept at their steady-state values. The benchmark model series is identical to the actual data because we estimated the model to fit the real wage and labor productivity series, and the labor share can be directly inferred from these two variables.

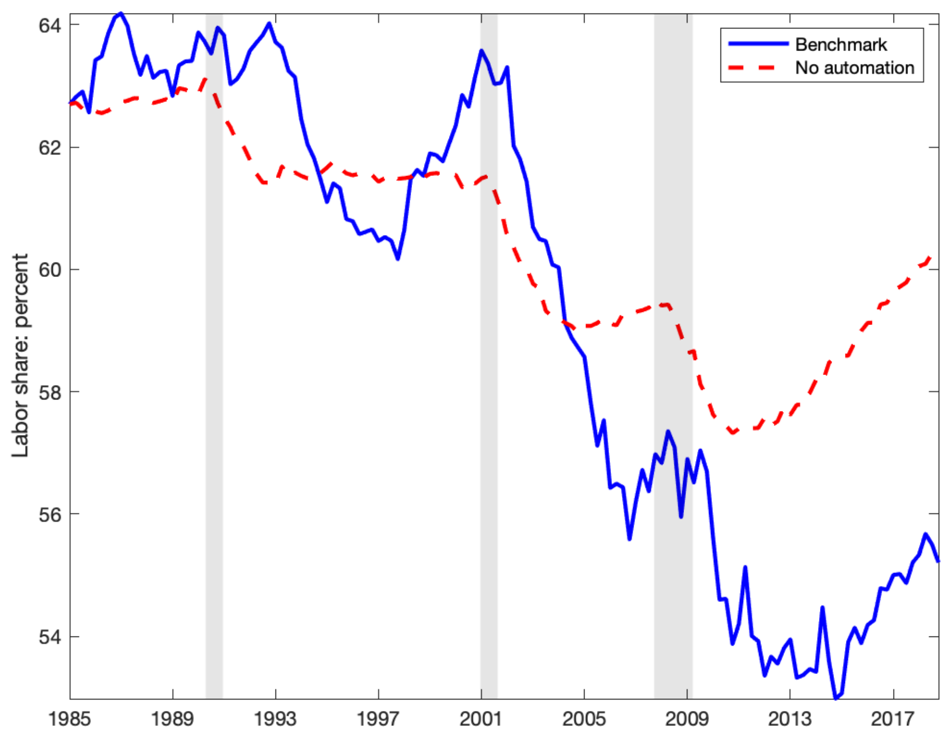

FIGURE 11. The smoothed time series of the labor share in the benchmark model estimated using the non-financial corporate sector data (blue solid line) versus the counterfactual with no automation (red dashed line). The counterfactual is a version of our benchmark model with all automation-related variables (including automation-specific shocks) kept at their steady-state values. The benchmark model series is identical to the actual labor share data for the non-financial corporate sector because we estimated the model to fit the real wage and labor productivity series, and the labor share can be directly inferred from these two variables.

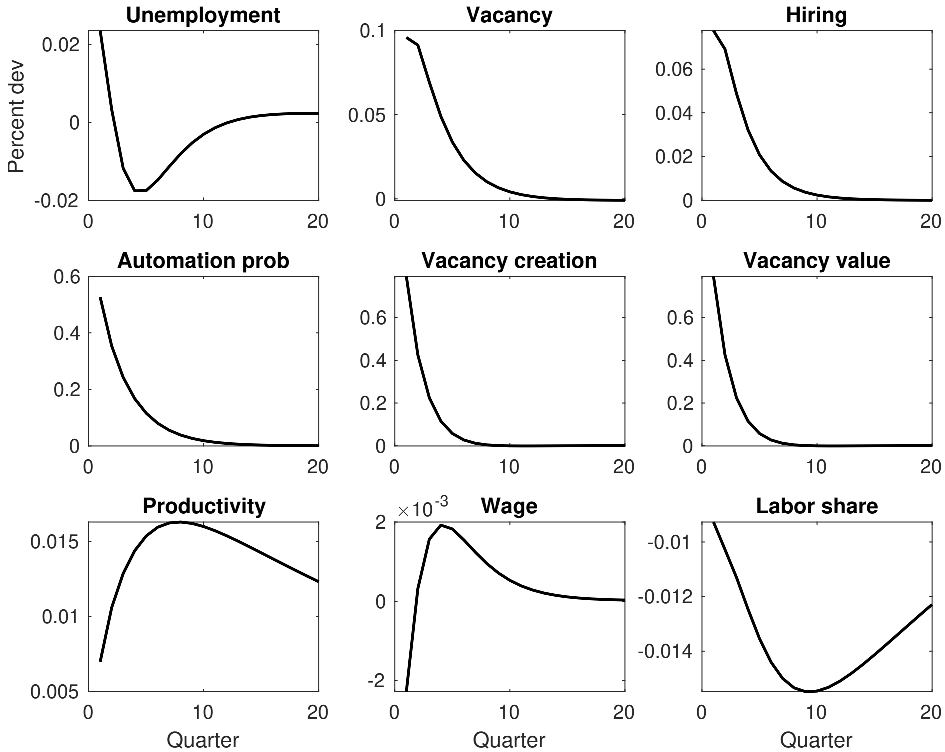

FIGURE 12. Impulse responses to a positive discount factor shock in the alternative model with automation of jobs instead of vacancies.

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.

——— (2019): "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 33, 3–30.

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.

Autor, D. H., D. Dorn, L. F. Katz, C. Patterson, and J. V. Reenen (2017): "The Fall of the Labor Share and the Rise of Superstar Firms," NBER Working Paper No. 23396.

Autor, D. H. and A. Salomons (2018): "Is Automation Labor-Displacing? Productivity Growth, Employment, and the Labor Share," *Brookings Papers on Economic Activity*, 1, 1–87.

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.

Cortes, G. M., N. Jaimovich, and H. E. Siu (2017): "Disappearing Routine Jobs: Who, How, and Why?" *Journal of Monetary Economics*, 91, 69–87.

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.

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.

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.

Karabarbounis, L. and B. Neiman (2013): "The Global Decline of the Labor Share," *Quarterly Journal of Economics*, 129, 61–103.

Keynes, J. M. (1930): "Economic Possibilities for Our Grandchildren," in *Essays in Persuasion*, New York: Norton Co.

Krueger, A. B. (2018): "Reflections on Dwindling Worker Bargaining Power and Monetary Policy," Luncheon Address at the Jackson Hole Economic Symposium.

Leduc, S. and Z. Liu (2019): "The Weak Job Recovery in a Macro Model of Search and Recruiting Intensity," *American Economic Journal: Macroeconomics*, (forthcoming).

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

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

(9) Automation dynamics

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

(10) Employment value

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

(11) Vacancy value

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

(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

$$x_t^* = J_t^a - \mathbb{E}_t\beta\theta_{t+1}\frac{C_t}{C_{t+1}}J_{t+1}^v, \tag{A26}$$

(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

$$Y_t = Z_tN_t + Z_t\zeta_tA_t. \tag{A29}$$

(17) Resource constraint

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

(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 a job instead of a vacancy

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:

$$1 = \mathrm{E}_t\beta\theta_{t+1}\frac{C_t}{C_{t+1}}r_t, \tag{A32}$$

(2) Matching function

$$m_t = \mu u_t^{\alpha} v_t^{1-\alpha}, \tag{A33}$$

(3) Job finding rate

$$q_t^u = \frac{m_t}{u_t}, \tag{A34}$$

(4) Vacancy filling rate

$$q_t^v = \frac{m_t}{v_t}, \tag{A35}$$

(5) Employment dynamics

$$N_t = (1-\delta_t)(1-q_t^a)N_{t-1} + m_t, \tag{A36}$$

(6) Number of searching workers

$$u_t = 1-(1-\delta_t)(1-q_t^a)N_{t-1}, \tag{A37}$$

(7) Unemployment

$$U_t = 1 - N_t, \tag{A38}$$

(8) Vacancy dynamics

$$v_t = (1-q_{t-1}^v)v_{t-1} + \delta_t N_{t-1} + \eta_t, \tag{A39}$$

(9) Automation dynamics

$$A_t = (1-\rho^o)A_{t-1} + q_t^a(1-\delta_t)N_{t-1}, \tag{A40}$$

(10) Employment value

$$J_t^e = Z_t - w_t + \mathbb{E}_t\beta\theta_{t+1}\frac{C_t}{C_{t+1}}\left\{\delta_{t+1}J_{t+1}^v + (1-\delta_{t+1})\left[q_{t+1}^a J_{t+1}^a + (1-q_{t+1}^a)J_{t+1}^e\right]\right\}, \tag{A41}$$

(11) Vacancy value

$$J_t^v = -\kappa + q_t^v J_t^e + (1-q_t^v)\mathbb{E}_t\beta\theta_{t+1}\frac{C_t}{C_{t+1}}J_{t+1}^v, \tag{A42}$$

(12) Automation value

$$J_t^a = Z_t\zeta_t - \kappa_a + (1-\rho^o)\mathbb{E}_t\beta\theta_{t+1}\frac{C_t}{C_{t+1}}J_{t+1}^a, \tag{A43}$$

(13) Automation threshold

$$x_t^* = J_t^a - J_t^e, \tag{A44}$$

(14) Robot adoption

$$q_t^a = \left(\frac{x_t^*}{\bar{x}}\right)^{\eta_a}, \tag{A45}$$

(15) Vacancy creation

$$\eta_t = \left(\frac{J_t^v}{\bar{e}}\right)^{\eta_e}, \tag{A46}$$

(16) Aggregate output

$$Y_t = Z_t N_t + Z_t\zeta_t A_t. \tag{A47}$$

(17) Resource constraint

$$C_t + \kappa v_t + \kappa_a A_t + \frac{\eta_a}{1+\eta_a}q_t^a x_t^*(1-\delta_t)N_{t-1} + \frac{\eta_e}{1+\eta_e}\eta_t J_t^v = Y_t, \tag{A48}$$

(18) Nash bargaining wage

$$\frac{b}{1-b}(J_t^e - J_t^v) = w_t - \phi - \chi C_t + \mathbb{E}_t \frac{\beta \theta_{t+1} C_t}{C_{t+1}} (1 - q_{t+1}^u)(1 - \delta_{t+1}) \frac{b}{1-b}(J_{t+1}^e - J_{t+1}^v), \quad \text{(A49)}$$