

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

April 2021

Working Paper 2019-17

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

**Suggested citation:**

Leduc, Sylvain, Zheng Liu. 2021. “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>

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

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

SYLVAIN LEDUC AND ZHENG LIU

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

## I. INTRODUCTION

Recent advances in robotics and artificial intelligence have raised concerns that automation is putting an increasing share of jobs at risk and reducing wages. There is an on-going debate about whether automation reduces aggregate employment (Autor, 2015; Acemoglu and Restrepo, 2018, 2020). However, to the extent that automation is a labor-saving technology, the *threat* of automation might weaken workers' bargaining power and thus restrain wage increases, even if robots are not actually adopted. The option to automate may become particularly attractive when firms face a tight labor market, in which hiring workers is difficult without substantial wage increases. In this paper, we argue that, by lowering workers' effective bargaining power, the increased threat of automation in expansions acts as a form of endogenous wage rigidity and helps explain the large fluctuations in unemployment and

---

*Date:* April 12, 2021.

*Key words and phrases.* Robots, automation, unemployment, wages, productivity, Shimer puzzle, 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, HKUST Macro workshop, Hong Kong University, INSEAD, National University of Singapore, University of Autonomia Barcelona, University of New South Wales, and the 2020 Econometric Society World Congress. We are grateful to Lily Seitelman and Remy Beauregard for excellent research assistance and to Anita Todd for editorial assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or of the Federal Reserve System.

vacancies relative to real wages, a puzzling observation through the lens of standard business cycle models (Shimer, 2005).

We formalize this argument based on a general equilibrium framework with labor market search frictions, generalized to incorporate automation decisions and estimated to fit U.S. time series. This theoretical framework allows us to study the interactions between firms' employment and automation decisions over the business cycle.<sup>1</sup> A framework with search frictions in the labor market is also natural, since we are interested in the joint dynamics of unemployment, job vacancies, and wages.

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

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

Our estimation suggests that the automation mechanism is a quantitatively important source of real wage rigidities, and it helps account for the observed large volatilities of unemployment and vacancies. In addition, since automation raises labor productivity while depressing wages, it leads to countercyclical fluctuations in the labor share of income, as observed in the data.

**I.1. Model mechanism.** We generalize the standard Diamond-Mortensen-Pissarides (DMP) model with labor market search frictions to build a tractable framework that incorporates automation decisions.

Our theoretical framework features two important departures from the standard DMP model. First, we introduce a fixed cost of vacancy creation (Fujita and Ramey, 2007; Leduc

---

<sup>1</sup>Our general equilibrium approach is complementary to a large body of empirical literature that focuses on the causal impact of automation on different outcome variables (e.g., employment and wage growth) using instrumental variables to control for the possibility of reverse causation (Autor and Salomons, 2018; Acemoglu and Restrepo, 2020; Acemoglu et al., 2020). For instance, low employment growth resulting from hiring difficulties in expansions may lead to more automation.

and Liu, 2020). A firm will choose to create a new vacancy if and only if the net present value of the vacancy is non-negative. Thus, in an equilibrium, an unfilled vacancy has a positive present value. Second, we introduce endogenous automation decisions. A vacancy in our model can be interpreted as a continuum bundle of tasks, which are ex ante identical, but a fraction of which can be automated depending on costs of automation. Only non-automated tasks (vacancies) are posted for hiring workers. Thus, in equilibrium, some tasks are performed by robots (i.e., automated tasks) and some other tasks by workers (provided that a match is formed in the labor market). Since robots can substitute for workers in production, they are different from the physical capital in standard neoclassical production functions, where capital and labor are complementary inputs.<sup>2</sup>

In line with Acemoglu and Restrepo (2018) and Zeira (1998), firms in our model first make a choice of technologies (adopting a robot or not); and then post those non-automated tasks (i.e., vacancies) for hiring workers.<sup>3</sup>

To implement this approach, we assume that each firm observes an i.i.d. cost of automation in the beginning of each period and decides whether or not to automate an unfilled job vacancy carried over from the previous period. At this stage of the decisions, the tasks can be indexed by the realized automation costs. A task will be automated if and only if the cost of automation lies below a threshold determined by the net benefit of automation; in which case, the firm adopts a robot for production and takes the job vacancy offline. The probability of automation is thus the cumulative density of the automation cost draws evaluated at the automation threshold.

After the automation decisions, the firm pools the non-automated tasks (which are again identical because the automation costs are irrelevant at the stage for hiring decisions) and posts those vacancies in the labor market to search for a potential match with job seekers. If a match is successful, a vacancy will be filled with a worker and both the firm and the worker obtain their respective share of the 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 it in future periods.

---

<sup>2</sup>Krusell et al. (2000) study a neoclassical model in which capital equipment complements skilled labor but substitutes for unskilled labor. He and Liu (2008) study a general equilibrium extension of the Krusell et al. (2000) model to incorporate endogenous skill accumulations. The relation between robots and workers in our model is analogous to the relation between equipment and unskilled labor in the model of Krusell et al. (2000).

<sup>3</sup>We focus on firms' ability to use automation as a potential margin of adjustments in response to changes in labor market conditions over the business cycles. In an expansion, for example, firms face a tighter labor market, in which job vacancies are harder to fill. This creates an incentive for adopting labor-saving technologies such as automation.

Our approach to modeling automation requires job vacancies to carry a positive value, which is an equilibrium outcome in our model since vacancy creation is costly. The option to automate an unfilled vacancy raises the vacancy value, providing an extra incentive for firms to create new vacancies. But an increase in the value of a vacancy would also boost the firm’s reservation value in wage negotiations, effectively weakening workers’ bargaining power, depressing wage increases. Thus, the threat of automation (i.e., the endogenous probability of robot adoptions) creates a source of real wage rigidity that contains wage changes and amplify the fluctuations of unemployment and vacancies over the business cycle.

**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).

We find that the threat of automation dampens wage increases in a business cycle boom. Since the net value of automation is procyclical, the probability of automation increases in expansions. The model’s implied procyclical automation investment is consistent with macroeconomic evidence. For example, Figure 1 shows that, in the U.S. data, aggregate real private investment in computers and peripheral equipment and in information processing equipment—both categories closely related to automation equipment—rise in expansions and fall in recessions, as does aggregate equipment investment. All three categories of equipment investment are positively correlated with real GDP growth in the sample period from 1985:Q1 to 2019:Q4, with correlation coefficients of 0.46, 0.58, and 0.71, respectively. Furthermore, the shipment of industrial robots is also procyclical, as reported by Kopytov et al. (2018) based on data from the International Federal of Robotics (IFR).

An increase in automation probability in a business cycle boom raises the firm’s reservation value (i.e., the value of a vacancy) and effectively weakens the worker’s bargaining power in wage negotiations, and therefore muting wage increases. By dampening wage changes, the automation channel amplifies the fluctuations in unemployment and vacancies. Increased automation in a boom also boosts aggregate productivity, further fueling the expansion. Since automation improves labor productivity while muting wage increases, it implies a countercyclical labor income share, as observed in the data.<sup>4</sup>

---

<sup>4</sup>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




FIGURE 1. Year-over-year percentage changes in real private investment in computers and peripherals (blue line), information processing equipment (green line), and aggregate equipment (yellow line). The gray shaded bars indicate recessions based on the NBER recession dates. The sample period covers 1985:Q1-2019:Q4. Source: U.S. Bureau of Economic Analysis and Haver Analytics.

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 is important for amplifying labor market fluctuations (Christiano et al., 2020). In addition, automation raises aggregate productivity in a business cycle boom,

---

relative price of investment goods. Elsby et al. (2013) also study the trend declines in the U.S. labor share since the 1980s. They argue that offshoring of labor-intensive component of the U.S. supply chain is an important factor that may explain the declines in the labor share. We focus on the cyclical dynamics of the labor share, instead of its trend.

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 37 times that of the real wage rate.<sup>5</sup> 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 from the threat of automation, making unemployment and vacancies less volatile. In this sense, automation and labor market search frictions are both important for understanding the observed labor market fluctuations.

The threat of automation effectively weakens workers' bargaining power and therefore mutes wage changes and amplifies unemployment fluctuations. Hagedorn and Manovskii (2008) argue that, in the standard DMP framework, reducing workers' bargaining weight or raising the workers' value of non-market activity (such as unemployment insurance) can amplify fluctuations in unemployment and vacancies. Thus, we consider two such counterfactuals, both without the automation mechanism. We find that lowering workers' bargaining weight or raising the unemployment insurance (UI) can dampen wage adjustments and amplify unemployment fluctuations, in line with the findings in Hagedorn and Manovskii (2008). However, the magnitude of amplification is substantially smaller than that arising from the threat of automation. In addition, the impulse responses of wages and the labor share in the counterfactual models are qualitatively different from those in our benchmark model. For example, following a positive discount factor shock, the benchmark model predicts that the real wage and the labor share both decline, whereas the counterfactuals without automation predict that they both rise. These differences reflect the importance of the automation threat for wage bargaining and the endogenous productivity changes through the automation channel.

The automation mechanism is also robust to introducing heterogeneous worker skills. In a generalized version of the model, we assume that robots are substitutes for low-skilled workers but complements to high-skilled workers. In this case, we continue to find that the threat of automation depresses wages and boosts productivity, leading to countercyclical labor share fluctuations. We also find that the volatility of the vacancy-unemployment ratio relative to that of real wages remains high in this environment, though somewhat lower than that obtained in our benchmark model. This attenuation effect reflects the business-cycle movements in skilled wages and the resulting changes in the cost of using robots. Since robots and skilled workers are complementary, an increase in skilled wages in a business

---

<sup>5</sup>Since we fit our model to these time series, the actual volatility ratio in the data is the same.

cycle expansion would raise the cost of using the automation technology, mitigating the incentive for firms to automate and resulting in greater fluctuations in unskilled wages.

Our framework can be generalized to study the implications of offshoring for labor market fluctuations. Elsbey et al. (2013) highlight the importance offshoring in accounting for the observed declines in the U.S. labor income share. To the extent that offshoring is a form of labor-saving technology, we conjecture that the threat of offshoring can exert similar influences on labor market outcomes as does the threat of automation.

**I.3. Related literature.** Our estimated general equilibrium model highlights the interactions between automation and labor market fluctuations at the business cycle frequencies. Our findings complements the empirical literature that typically focuses on longer-run implications of automation. For example, Graetz and Michaels (2018) examine the labor-market impact of cumulative changes in robot adoptions from 1993 to 2007 using a panel of industry-level data from 17 countries. They find that robot adoptions boost labor productivity and raise wages, although the positive effects on wages are much smaller than those on productivity. Arnoud (2018) also focuses on the long-run implications of automation. He examines occupation-level relations between the threat of automation and wage adjustments using data from the 2013 U.S. Current Population Survey and an index of automatability for different occupations developed by Frey and Osborne (2017). He finds that, controlling for observable characteristics, occupations that are more susceptible to automation have experienced lower wage growth. Dinlersoz and Wolf (2018) present plant-level evidence that more automated establishments in the U.S. manufacturing sector have had a smaller fraction of high-wage workers, higher labor productivity, and a smaller labor share in production. Acemoglu and Restrepo (2020) present evidence that, for U.S. commuter zones exposed to robots, the increase in the stock of industrial robots between 1990 and 2007 reduced the average employment-to-population ratio by 0.4 percentage points and average wages by 0.8 percent, relative to commuter zones with no robot exposure.

Automation in our model represents a labor-substituting technology, in line with Acemoglu and Restrepo (2018). There is substantial evidence that the steady progress in labor substituting technologies (such as computerization) has reduced the secular demand for workers with routine skills, contributing to job polarization in the U.S. labor market (Autor et al., 2003; Autor, 2015). Furthermore, job polarization can be linked to the jobless recoveries since the early 1990s, because most of the employment losses in routine occupations occur in recessions (Jaimovich and Siu, 2020). During the Great Recession and the subsequent recovery, employers in hard-hit areas raised the skill requirements when posting job vacancies, consistent with increased job destructions in routine occupations (Hershbein and Kahn, 2018). Although our benchmark model abstracts from skill heterogeneity, our extended



model with heterogeneous worker skills captures the idea that robots are complementary with high-skill workers but are substitutes for low-skill workers. The key predictions of our benchmark model survives in this extended model. In particular, increases in automation raises labor productivity, reduces wages and the labor share.

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

## 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  $\delta_t$  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, as in Fujita and Ramey (2007) and Leduc and Liu (2020). Since creating a new vacancy incurs a fixed cost, a vacancy has a positive value even if it is not filled by a worker. The number of vacancies becomes a slow-moving state variable (instead of a jump variable as in the standard DMP framework), enabling our model to match the persistent vacancy dynamics in the data. Second, we introduce endogenous automation decisions. We focus on firms' ability to use automation as a labor-saving device when labor markets tighten in expansions. In the beginning of period  $t$ , each firm draws a fixed automation cost, 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, then the firm adopts a robot and obtains the automation value, and the vacancy would be taken offline. If the automation cost exceeds the

threshold, then the firm posts the vacancy for hiring a worker. Since robots can substitute for workers 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  $\delta_t$  of those matches, so that the measure of unemployed job seekers is given by

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

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

The job separation rate shock  $\delta_t$  follows the stationary stochastic process

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

where  $\rho_\delta$  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  $\sigma_\delta$ . The term  $\bar{\delta}$  denotes the steady-state rate of job separation.

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

$$v_t = (1 - q_{t-1}^v)(1 - q_t^a)v_{t-1} + \delta_t N_{t-1} + \eta_t, \quad (3)$$

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

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

$$m_t = \mu u_t^\alpha v_t^{1-\alpha}, \quad (4)$$

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

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

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

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

$$U_t = u_t - m_t = 1 - N_t. \quad (6)$$

For convenience, we define the job finding probability  $q_t^u$  as

$$q_t^u = \frac{m_t}{u_t}. \quad (7)$$

Similarly, we define the job filling probability  $q_t^v$  as

$$q_t^v = \frac{m_t}{v_t}. \quad (8)$$

**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

$$x_t^* = J_t^a - J_t^v. \quad (9)$$

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

$$q_t^a = G(x_t^*). \quad (10)$$

The flow of automated job positions adds to the stock of automated positions (denoted by  $A_t$ ), which becomes obsolete at the rate  $\rho^o \in [0, 1]$  in each period.<sup>6</sup> Thus,  $A_t$  evolves according to the law of motion

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

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

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

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

The parameter  $\rho_z \in (-1, 1)$  measures the persistence of the technology shock. The term  $\varepsilon_{zt}$  is an i.i.d. normal process with a zero mean and a finite variance of  $\sigma_z^2$ . The term  $\bar{Z}$  is the steady-state level of the technology shock.<sup>7</sup> The automation-specific technology shock

---

<sup>6</sup>If a vacancy is “filled” by a robot, it will be taken offline once and for all. Even if the robot later becomes obsolete, the vacated position does not return to the stock of vacancies.

<sup>7</sup>The model can easily be extended to allow for trend growth.

$\zeta_t$  follows a stochastic process that is independent of the neutral technology shock  $Z_t$ . In particular,  $\zeta_t$  follows the stationary process

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

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  $\sigma_\zeta^2$ . The term  $\bar{\zeta}$  is the steady-state level of the automation-specific technology shock.

To simplify the analysis and concentrate on the main mechanism, we assume that operating the robot incurs a flow fixed cost  $\kappa_a$  that captures the costs of facilities and the space for automated production.<sup>8</sup> The value of automation satisfies the Bellman equation

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

where  $D_{t,t+1}$  denotes the stochastic discount factor determined by the marginal utility of the households.

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

$$\eta_t = F(J_t^v). \quad (15)$$

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

$$J_t^v = -\kappa + q_t^v J_t^e + (1 - q_t^v) \mathbb{E}_t D_{t,t+1} \left\{ q_{t+1}^a \left[ J_{t+1}^a - \int_0^{x_t^*} x dG(x) \right] + (1 - q_{t+1}^a) J_{t+1}^v \right\}. \quad (16)$$

---

<sup>8</sup>This simplification, however, does not drive our main results, as we show in Section V.2, where we generalize the benchmark framework to allow for worker skill heterogeneity. In that generalized framework, operating a robot requires skilled workers. Thus the flow cost of using the automation technology becomes endogenous, and in particular, it depends on the wage rate of skilled workers.

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

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

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  $\delta_{t+1}$ ), 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

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

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  $\Theta_t$  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

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

In this shock process,  $\rho_\theta$  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  $\sigma_\theta$ . Here, we have implicitly assumed that the mean value of  $\theta$  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

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

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  $\phi$  measures the flow benefits of unemployment.

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

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

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

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

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  $\Lambda_t$  denotes the Lagrangian multiplier for the budget constraint (20). The optimizing decision for employment implies that the employment surplus satisfies the Bellman equation

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

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.<sup>9</sup>

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 (provided that the employment relation survives job separation); however, adding a worker today would also reduce the pool of searching workers tomorrow, a fraction  $q_{t+1}^u$  of whom would be able to find jobs. Thus, the marginal effect of adding a new worker in period  $t$  on employment in period  $t + 1$  is given by  $(1 - q_{t+1}^u)(1 - \delta_{t+1})$ , resulting in the effective continuation value of employment shown in the last term of Eq. (23).

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

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

**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

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

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

<sup>9</sup>We provide detail derivations in an online appendix available at [https://www.frbsf.org/economic-research/files/wp2019-17\\_appendix.pdf](https://www.frbsf.org/economic-research/files/wp2019-17_appendix.pdf).

Define the total surplus as

$$S_t \equiv J_t^e - J_t^v + S_t^H. \quad (26)$$

Then the bargaining solution is given by

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

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

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

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  $\phi$  for unemployed workers through lump-sum taxes. We assume that the government balances the budget in each period such that

$$\phi(1 - N_t) = T_t. \quad (29)$$

**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

$$B_t = 0. \quad (30)$$

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

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

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

$$Y_t = Z_t N_t + Z_t \zeta_t A_t. \quad (32)$$

## III. EMPIRICAL STRATEGIES

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

We focus on the parameterized distribution functions

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

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.<sup>11</sup> 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  $\rho^o = 0.02$ , so that robots depreciate at an average annual rate of 8 percent, in line with the estimated average life span of robots used by the International Federation of Robotics (IFR) for constructing their measure of the operation stocks of robots. We normalize the level of labor productivity to  $\bar{Z} = 1$  and 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

---

<sup>10</sup>Details of the equilibrium conditions, the steady state, and the log-linearized system are presented in the appendix.

<sup>11</sup>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  $\eta_a$  in the set of parameters to be estimated. We obtain a posterior estimate of  $\eta_a$  close to one and very similar estimates for the other parameters. For simplicity and for obtaining a closed-form solution for the steady-state equilibrium, we assume that  $\eta_a = 1$  in our benchmark model.



TABLE 1. Calibrated parameters

Parameter	Description	value
$\beta$	Subjective discount factor	0.99
$\phi$	Unemployment benefit	0.25
$\alpha$	Elasticity of matching function	0.50
$\mu$	Matching efficiency	0.6594
$\bar{\delta}$	Job separation rate	0.10
$\rho^o$	Automation obsolescence rate	0.02
$\kappa$	Vacancy posting cost	0.1041
$b$	Nash bargaining weight	0.50
$\eta_v$	Elasticity of vacancy creation cost	1
$\eta_a$	Elasticity of automation cost	1
$\kappa_a$	Flow cost of automated production	0.98
$\chi$	Disutility of working	0.7237
$\bar{Z}$	Mean value of neutral technology shock	1
$\bar{\zeta}$	Mean value of equipment-specific technology shock	1

is  $v = \frac{m}{q^v}$ . The scale of the matching efficiency is then given by  $\mu = \frac{m}{u\alpha v^{1-\alpha}}$ . 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  $\kappa$ , 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  $\chi$  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.<sup>12</sup>

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

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

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

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

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

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

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.

<sup>12</sup>We provide more details of the macro time-series data in Appendix A.

TABLE 2. Estimated parameters

Parameter description	Priors		Posterior		
	Type	[mean, std]	Mean	5%	95%
$\bar{e}$ scale for vacancy creation cost	G	[5, 1]	7.7091	5.1363	9.7973
$\bar{x}$ scale for robot adoption cost	G	[5, 1]	2.6433	1.7908	3.5023
$\rho_z$ AR(1) of neutral technology shock	B	[0.8, 0.1]	0.9394	0.9195	0.9572
$\rho_\theta$ AR(1) of discount factor shock	B	[0.8, 0.1]	0.9701	0.9527	0.9887
$\rho_\delta$ AR(1) of separation shock	B	[0.8, 0.1]	0.9445	0.9109	0.9814
$\rho_\zeta$ AR(1) of automation-specific shock	B	[0.8, 0.1]	0.8595	0.8293	0.8889
$\sigma_z$ std of tech shock	IG	[0.01, 0.1]	0.0111	0.0098	0.0124
$\sigma_\theta$ std of discount factor shock	IG	[0.01, 0.1]	0.0076	0.0050	0.0107
$\sigma_\delta$ std of separation shock	IG	[0.01, 0.1]	0.0467	0.0422	0.0519
$\sigma_\zeta$ std of automation-specific shock	IG	[0.01, 0.1]	0.0285	0.0197	0.0370
Log data density	1258.27				

*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.

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

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

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  $\bar{e} = 7.71$ . The posterior mean estimates of the automation cost parameter is  $\bar{x} = 2.64$ . These parameters imply a steady-state share of output produced by automation of  $A/Y = 0.32$ . Thus, our model implies that, in the long run, about 32 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 posterior estimates are significantly different from the priors, suggesting that the data are informative about these 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 2 shows the impulse responses of several key macro variables to a positive neutral technology shock in the benchmark model. The shock leads to persistent declines in unemployment and persistent increases in vacancies and hiring. The shock also raises the value of automation. Under our parameters, the value of automation rises more than does the value of a vacancy, raising the net value of robot adoption and therefore leading to an increase in the automation probability. 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 3 shows the impulse responses to a positive discount factor shock. The shock raises the present values of a job match and an open vacancy. Thus, it generates a persistent boom

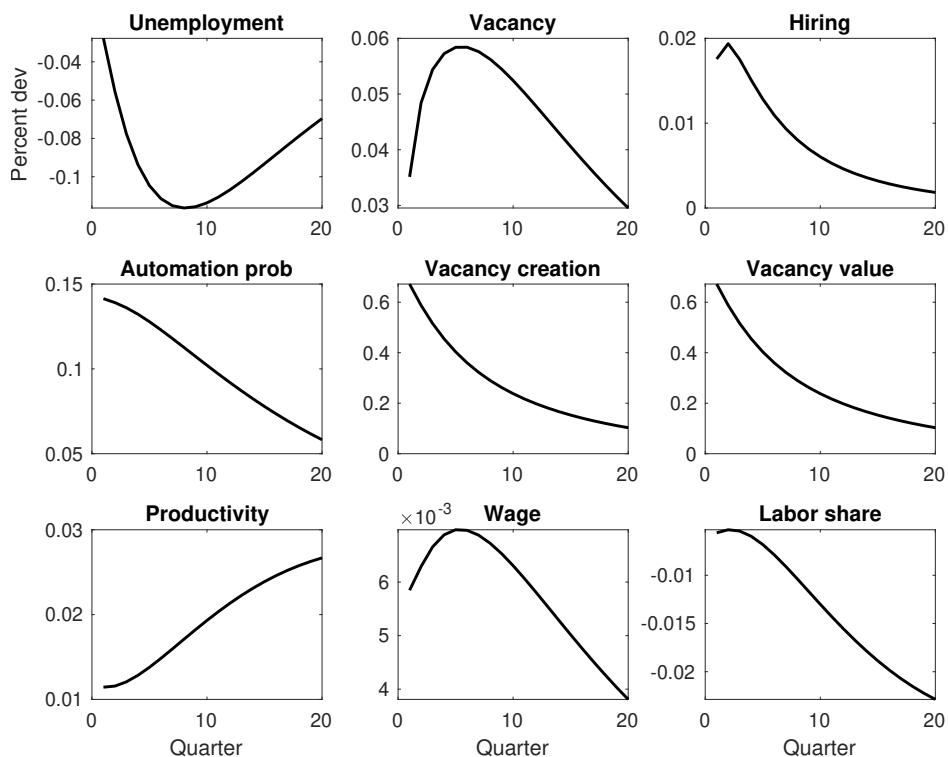


FIGURE 2. Impulse responses to a positive neutral technology shock in the benchmark model.

in employment, vacancies, and hiring. The shock also raises the net value of automation (i.e., the difference between the value of a robot and the value of a vacancy), increasing the probability of robot adoption. The increase in robot adoptions boosts labor productivity, further fueling the boom. However, as the threat of automation rises, the workers' bargaining power weakens, leading to a modest short-run decline in the real wage. By increasing productivity and reducing the real wage rate, the discount factor shock generates a persistent decline in the labor share.

Figure 4 shows the impulse responses to a positive automation-specific shock. The shock directly raises the value of automation. Since the option of automation boosts the value of vacancies, the increase in automation leads to more vacancy creation. With more job openings, the job finding rate increases, raising hiring and reducing unemployment. 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.<sup>13</sup>

<sup>13</sup>We report the impulse responses to a job separation shock in the online appendix at [https://www.frbsf.org/economic-research/files/wp2019-17\\_appendix.pdf](https://www.frbsf.org/economic-research/files/wp2019-17_appendix.pdf). As we discuss there,

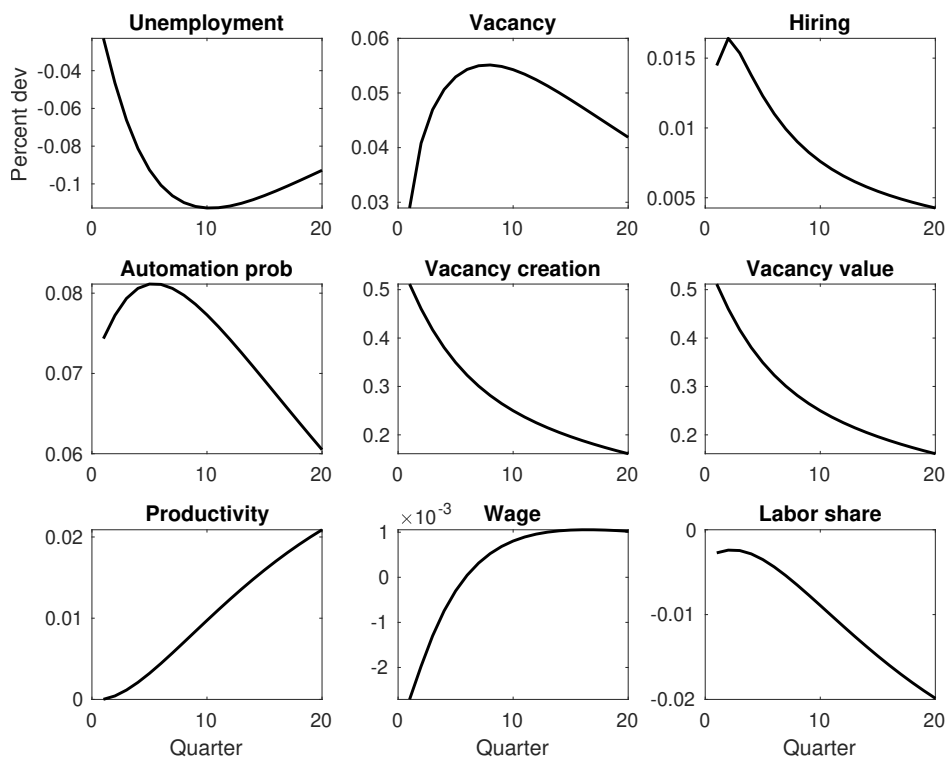


FIGURE 3. Impulse responses to a positive discount factor shock in the benchmark model.

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. Table 3 displays the results.<sup>14</sup>

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 a little under 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 a job separation shock raises both unemployment and vacancies. Consistent with Shimer (2005), this counterfactual positive correlation between unemployment and vacancies renders the job separation shock unimportant for driving labor market dynamics.

<sup>14</sup>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.

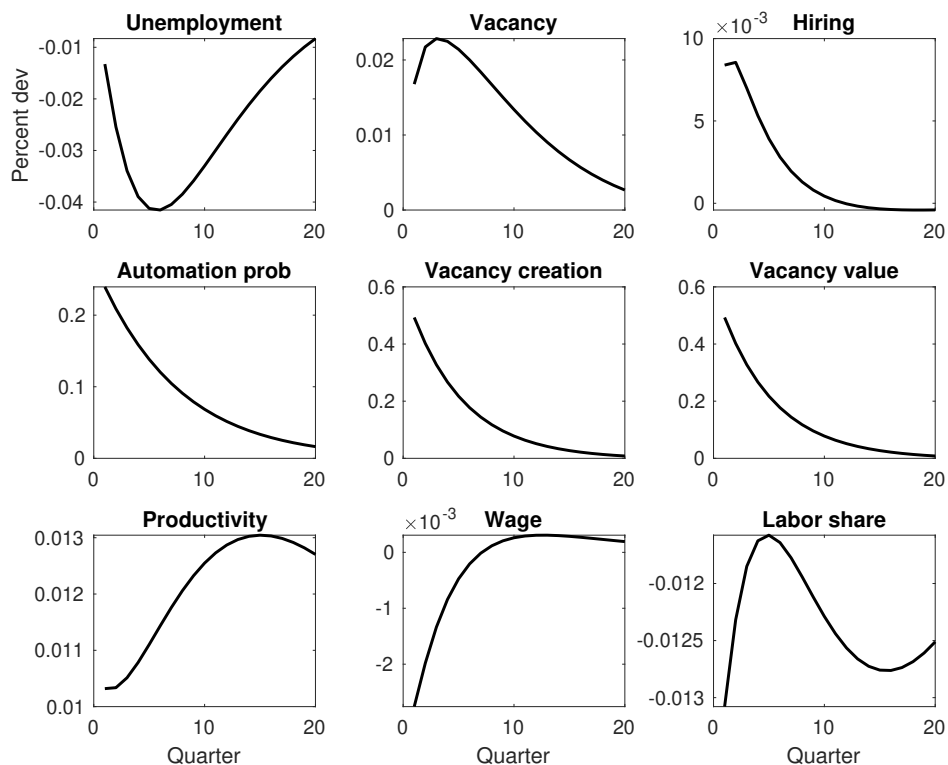


FIGURE 4. Impulse responses to a positive automation-specific shock in the benchmark model.

TABLE 3. Forecasting Error Variance Decomposition

Variables	Neutral technology shock	Discount factor shock	Job separation shock	Automation specific shock
Unemployment	36.62	59.38	1.18	2.82
Vacancy	33.40	53.46	10.42	2.73
Productivity growth	46.21	14.81	0.29	38.69
Real wage growth	61.12	17.98	0.52	20.38

*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.

the probability of automation. These two shocks explain about 62 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, accounting for about one-third of their variances. Perhaps not surprisingly, the neutral technology shock is also important for explaining the fluctuations in labor productivity, explaining about 46 percent of its variance.<sup>15</sup> In addition, about 80 percent of the real wage fluctuations are accounted for by shocks to the neutral technology and the discount factor.

**IV.2. Automation vs. other amplification mechanisms.** Our model suggests that automation effectively weakens workers’ bargaining power and mutes wage changes, and therefore amplifying labor market fluctuations. Absent the automation mechanism, however, it is difficult to reproduce the observed labor market dynamics. The literature has studied other amplification mechanisms in the standard DMP framework without automation. For example, Hagedorn and Manovskii (2008) argue that, in the standard DMP framework, reducing workers’ bargaining weight or raising the workers’ value of non-market activity (such as unemployment insurance, or UI) can amplify fluctuations in unemployment and vacancies.

To evaluate the quantitative importance of the automation channel relative to these alternative amplification channels, we study a counterfactual specification without the automation mechanism (labelled “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. We consider two variations of the “no automation” specification, one with a lower worker bargaining weight (reducing  $b$  to 0.25 from 0.5) and the other with a higher UI benefit (raising  $\phi$  to 0.5 from 0.25).

Figure 5 displays the impulse responses to a positive discount factor shock in the benchmark model with automation (the black solid lines), the counterfactual with no automation (the blue dashed lines), and the no-automation counterfactual with a lower worker bargaining weight (the red dotted-dashed lines). Without the automation channel, reducing workers’ bargaining weight dampens wage increases and amplifies unemployment responses (the red dotted-dashed lines vs. the blue dashed lines), in line with the findings of Hagedorn and Manovskii (2008). However, the magnitude of amplification is small relative to that from the automation mechanism (the black solid lines vs. the blue dashed lines). Furthermore, in the benchmark model, the shock raises the value of automation, raising the threat of automation and the actual use of robots for production. The increased use of robots

---

<sup>15</sup>In the standard DMP model without automation, labor productivity fluctuations would be entirely driven by the neutral technology shock.



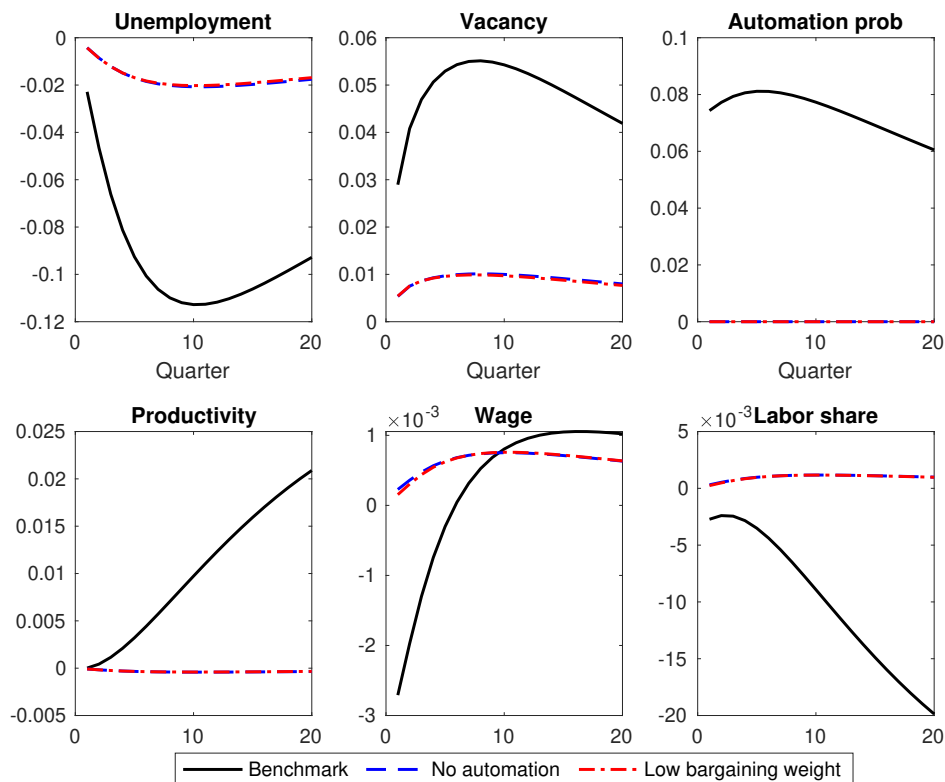


FIGURE 5. Impulse responses to a positive neutral technology shock in the benchmark model (black solid lines), the counterfactual with no automation (blue dashed lines), and the counterfactual with no automation and low worker bargaining power (red dot-dash lines).

boosts labor productivity, while the elevated threat of automation reduces wages. Working together, these forces lead to a persistent decline in the labor share following the positive technology shock. In contrast, the no-automation counterfactuals—including the case with a lower worker bargaining weight—predict that the shock raises the real wage but does not affect labor productivity, leading to an increase in the labor share.

Figure 6 shows the impulse responses to a positive discount factor shock in the benchmark model (the black solid lines), the counterfactual with no automation (the blue dashed lines), and the no-automation counterfactual with a higher UI benefit (the red dotted-dashed lines). Absent the automation channel, raising the UI benefit dampens wage adjustments and amplifies the responses of unemployment and vacancies, consistent with Hagedorn and Manovskii (2008). However, similar to the case with a lower worker bargaining weight, the magnitude of amplification is small relative to that from the automation mechanism.<sup>16</sup> These impulse

<sup>16</sup>The impulse responses to a neutral technology shock in these counterfactual models display similar patterns, as we show in the online appendix at [https://www.frbsf.org/economic-research/files/wp2019-17\\_appendix.pdf](https://www.frbsf.org/economic-research/files/wp2019-17_appendix.pdf).

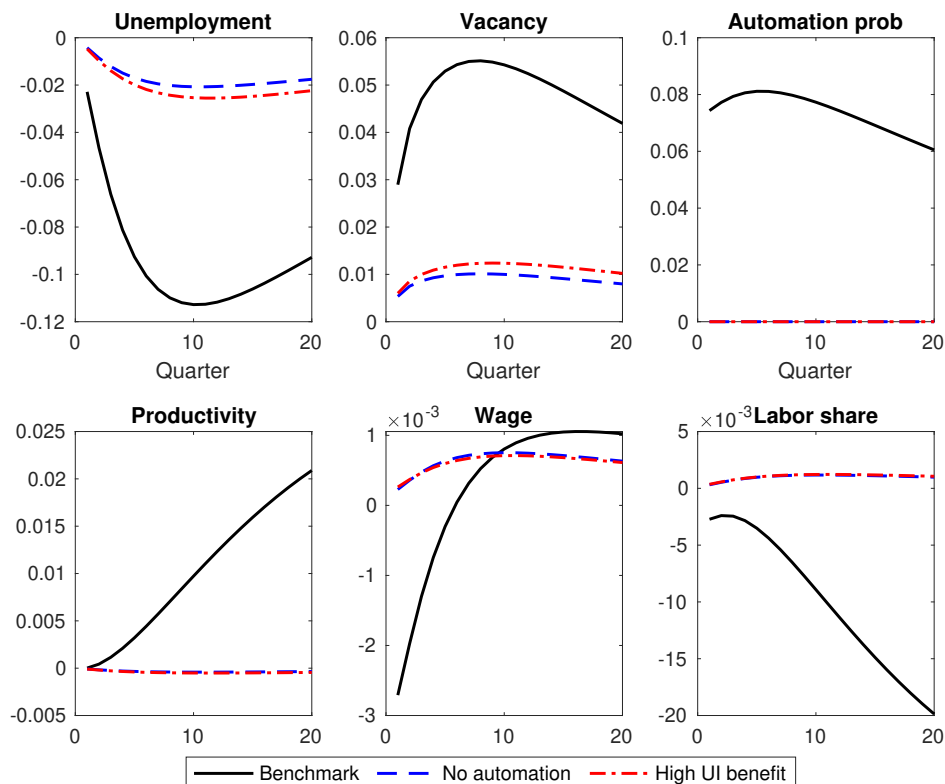


FIGURE 6. Impulse responses to a positive discount factor shock in the benchmark model (black solid lines), the counterfactual with no automation (blue dashed lines), and the counterfactual with no automation and high unemployment insurance (UI) benefits (red dot-dash lines).

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. In the standard real business cycle (RBC) model with a spot labor market, productivity fluctuations completely pass through to real wage adjustments, rendering it difficult to generate large employment fluctuations. The presence of search frictions in the labor market introduces a wedge between the equilibrium wage rate and labor productivity. However, absent the automation channel, the DMP model has also difficulties in generating large fluctuations in unemployment and job vacancies relative to those in real wages (i.e., the Shimmer puzzle). In this sense, search frictions and automation are complementary mechanisms for resolving the Shimer puzzle, as we discuss in this section.

To illustrate the importance of search frictions, we consider a counterfactual version of the model which features low levels of labor search frictions. In particular, that counterfactual

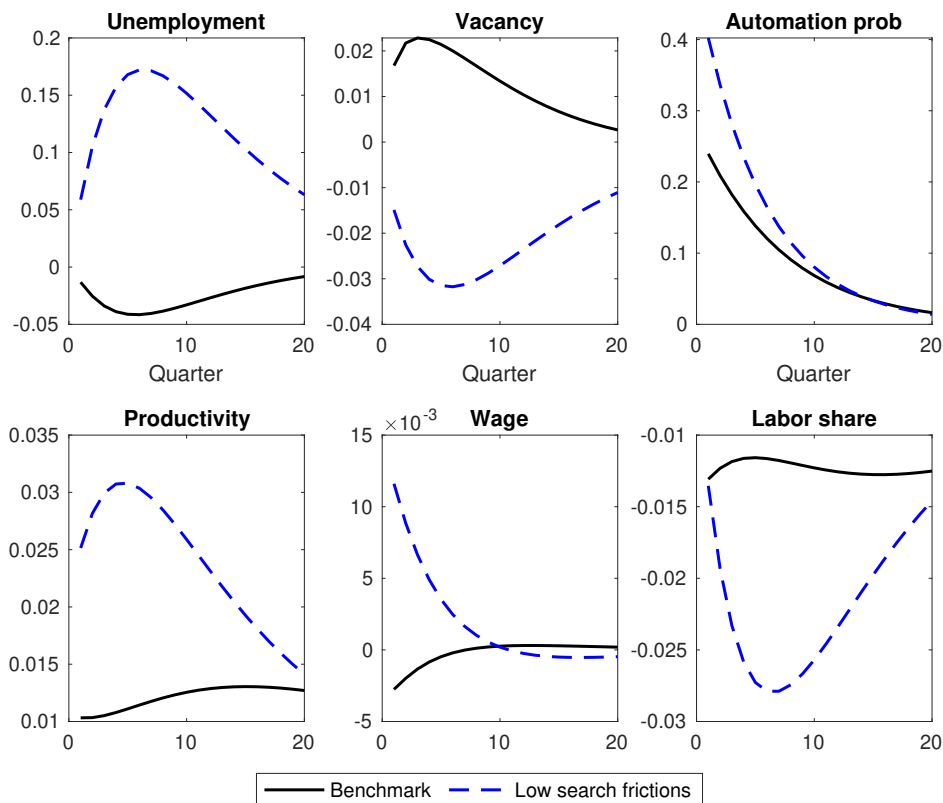


FIGURE 7. Impulse responses to a positive automation-specific shock in the benchmark model (black solid lines) and the counterfactual with low search frictions (blue dashed lines).

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).<sup>17</sup>

Figure 7 shows the impulse responses of the macro variables following a positive automation-specific 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, they produce starkly different responses of the labor market variables. In the benchmark model, the shock reduces unemployment and increases vacancies, suggesting that the job-creating effect of automation dominates the job-displacing effect. In contrast, in the model with low search frictions, the shock raises unemployment and reduces vacancies. In an economy with a spot labor market (i.e., the extreme case without search frictions), an employment relation would cease to be a

<sup>17</sup>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.

long-term relation, and the option of automation in the future would not affect current hiring decisions. In an economy with low search frictions, the option of automation would have a weak effect on job-creating incentives, and thus the job-displacing effect would become dominant.

In addition, in the case with low search frictions, the present value of a vacancy responds less to the automation-specific shock (because the model becomes closer to a spot labor market). Since the shock directly raises the value of automation, the automation threshold (i.e.,  $x_t^* = J_t^a - J_t^v$ ) and the automation probability rises more sharply than in the benchmark model, as shown in Figure 7, leading to sharper 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.<sup>18</sup>

The impulse responses shown in Figure 7 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 automation channel and the labor search frictions, working together, helps generate the observed large fluctuations in unemployment and vacancies relative to those in real wages.

**IV.4. Automation threat and labor market dynamics.** Our benchmark model implies that the automation channel 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 (2005) puzzle.

To further illustrate the quantitative importance of the automation mechanism for resolving the Shimer puzzle, we compare our benchmark model's predictions for the volatilities of labor market tightness and real wages with those from two counterfactuals: one without the automation channel and the other with low search frictions. In the no-automation case, we also consider two variations: one with a lower worker bargaining weight and the other with an increase in UI benefit, as we have done in Section IV.2.

Table 4 displays the standard deviations of the labor market tightness (measured by the v-u ratio  $v/U$ ), 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 37 times as volatile

---

<sup>18</sup>We have also compared the impulse responses to a neutral technology shock between the benchmark model and the counterfactual with low search frictions. We find that the benchmark model produces much stronger amplifications for the labor market variables than does the counterfactual with low search frictions. For details, see the online appendix at [https://www.frbsf.org/economic-research/files/wp2019-17\\_appendix.pdf](https://www.frbsf.org/economic-research/files/wp2019-17_appendix.pdf).

TABLE 4. Labor market volatilities implied by alternative models

Model	Market tightness	Real wage	Relative volatility
	(a)	(b)	(c)
(1) Benchmark model	1.0856	0.0294	36.9298
(2) No automation	0.2640	0.0303	8.6999
(2A) Low bargaining weight	0.2881	0.0301	9.5820
(2B) High UI benefit	0.4613	0.0276	16.6869
(3) Low search friction	1.1441	0.0410	27.8909
(4) Automating jobs	0.5074	0.0319	15.9095
(5) Skill heterogeneity	1.6857	0.0638	26.4417

*Note:* The rows in the table correspond to the alternative models: (1) the benchmark model, (2) the no-automation counterfactual, (3) the low-search-friction counterfactual, (4) the model with automation of jobs (instead of vacancies), and (5) the model with heterogeneous worker skills. We consider two additional variations of the no-automation case: (2A) low worker bargaining weight and (2B) high UI benefits. For each model, the numbers in the three columns are (a) the standard deviation of labor market tightness (measured by  $v/U$ , the ratio of vacancies to unemployment), (b) the standard deviation of the real wage rate, and (c) the ratio of the first two columns. In the model with skill heterogeneity, column (b) reports the real wage rate of unskilled workers.

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 (0.264 vs. 1.086) and slightly larger volatility of the real wage rate (0.030 vs. 0.029) than does the benchmark model, implying a much smaller relative volatility (8.70 vs. 36.93). This no-automation case essentially reproduces the Shimer (2005) volatility puzzle. Reducing the worker bargaining weight or raising the UI benefit boosts the relative volatility modestly (from 8.70 to 9.58 and 16.69, respectively), but the relative volatility remains smaller than in the data (36.93). The counterfactual with low search frictions (but with automation) also generates a smaller relative volatility than in the benchmark model (27.89 vs. 36.93). Thus, both automation and labor search frictions are important for the model's transmission mechanism.

## V. ROBUSTNESS OF THE MODEL MECHANISM

Our benchmark model shows that the threat of automation can effectively weaken workers' bargaining power, resulting in sluggish adjustments in real wages and therefore amplifying fluctuations in unemployment and vacancies. We now consider two variations of the benchmark model and examine the robustness of the model's main transmission mechanism.<sup>19</sup>

**V.1. Automating jobs.** In our benchmark model, we assume that firms can automate a vacancy if that vacancy is not filled with a worker. A plausible alternative way of thinking about automation is to allow firms to automate an existing job instead of an open vacancy. We now consider that alternative setup.

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

$$J_t^e = Z_t - w_t + \mathbb{E}_t \beta \theta_{t+1} \frac{C_t}{C_{t+1}} \left\{ \delta_{t+1} J_{t+1}^v + (1 - \delta_{t+1}) \left[ q_{t+1}^a \left( J_{t+1}^a - \int_0^{x_t^*} x dG(x) \right) + (1 - q_{t+1}^a) J_{t+1}^e \right] \right\}. \quad (38)$$

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  $J_{t+1}^a$  net of the expected robot adoption costs. If the job is neither separated nor automated, then the firm obtains the continuation value of employment  $J_{t+1}^e$ .

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

$$N_t = (1 - \delta_t)(1 - q_t^a)N_{t-1} + m_t. \quad (39)$$

We simulate the model based on the same calibrated and estimated parameters as in the benchmark model. Figure 8 shows the impulse responses of a few key macro and labor market variables following a positive discount factor shock. The shock raises the net present value of automation and thus increases the probability of automation. The increase in automation

<sup>19</sup>To conserve space, we sketch the key ingredients of each model in the text and describe the full equilibrium system in the online appendix at [https://www.frbsf.org/economic-research/files/wp2019-17\\_appendix.pdf](https://www.frbsf.org/economic-research/files/wp2019-17_appendix.pdf).

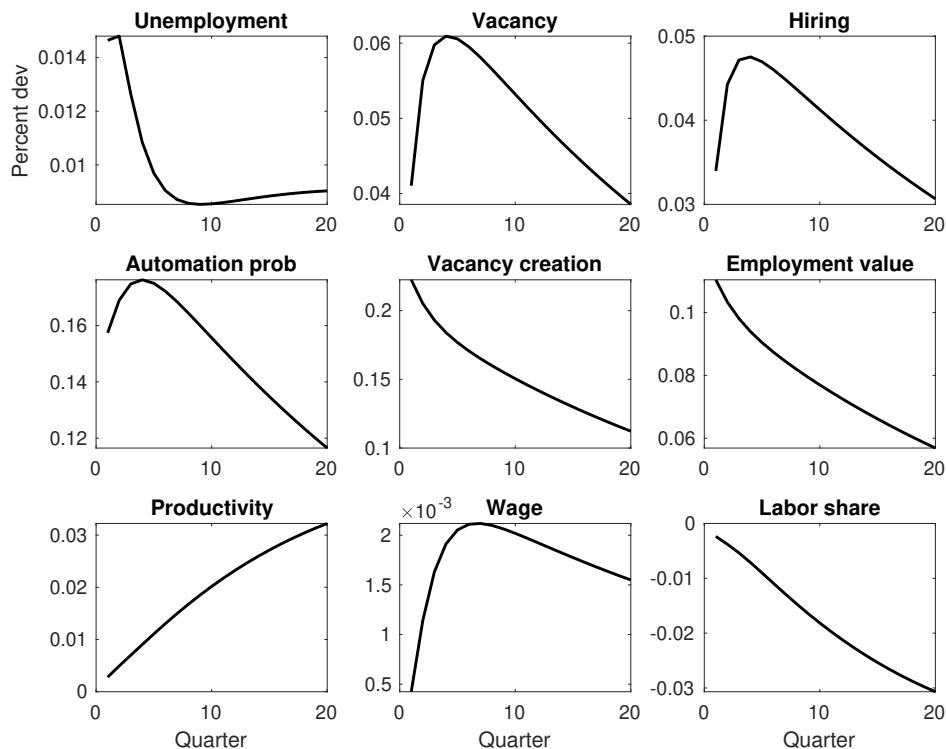


FIGURE 8. Impulse responses to a positive discount factor shock in the alternative model with automated jobs instead of vacancies.

probability has two opposing effects on employment. Automation directly replaces workers in this model, pushing up the unemployment rate. At the same time, automation improves labor productivity, raising the employment value and boosting employment. With the calibrated parameters, the job-displacing effect modestly dominates the employment boosting effect, such that a positive discount factor shock leads to a small increase in the unemployment rate. In contrast, the discount factor shock boosts the value of vacancies, encouraging firms to create new vacancies and leading to large and persistent increases in the stock of vacancies. As a result, the labor market tightness (i.e., the v-u ratio) increases persistently (not shown in the figure).

Increased automation probability leads to a muted response of the real wage rate, reflecting that the threat of automation weakens workers' bargaining power. Since labor productivity rises and wage responses are relatively muted, the labor share declines following a positive discount factor shock, as in our benchmark model. Thus, the automation channel that we have identified in the benchmark model is robust when we consider automating jobs instead of vacancies.

By muting wage adjustments, the automation mechanism creates a source of real wage rigidity that helps amplify the fluctuations of the labor market tightness. As shown in

Table 4, the model with automating jobs generates a relative volatility of the labor market tightness (relative to that of real wages) of about 16. This volatility ratio is smaller than the benchmark, although similar in magnitude to that of doubling the unemployment benefits in a model without automation. This relatively weaker amplification effect through the automation channel partly reflects that an increase in automation in the current model acts like an endogenous job separation shock, which raises both unemployment and vacancies.

**V.2. Heterogeneous worker skills.** Our benchmark model features one type of workers, who compete with robots for jobs. We now generalize the model to incorporate heterogeneity in worker skills and show that the model mechanism survives this generalization.

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

To keep the model tractable, we assume that the aggregate supply of skilled workers is inelastic and fixed at  $\bar{s}$ .<sup>20</sup> Skilled workers face a spot labor market with the competitive wage rate  $w_{st}$ . Unskilled workers face search frictions in the labor market, and they each receives the low-skilled wage rate  $w_{nt}$  if employed or the unemployment insurance benefit  $\phi$  if unemployed.

The household utility function remains the same (see Eq. (18)). The budget constraint now includes the wage income from skilled workers and is given by

$$C_t + \frac{B_t}{r_t} = B_{t-1} + w_{nt}N_t + w_{st}\bar{s} + \phi(1 - N_t) + d_t - T_t, \quad (40)$$

Since the supply of skilled workers is inelastic, introducing skilled workers does not affect the household's optimizing decisions relative to the benchmark model. The only required modification in the household's problem is that, in the employment surplus expression (23), the wage rate  $w_t$  should be replaced by the low-skilled wage rate  $w_{nt}$ .

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

---

<sup>20</sup>Given our focus on business cycles, assuming a constant supply of skilled workers seems innocuous since human capital accumulation is likely a slow-moving process. In the data, the share of skilled workers (e.g., those with a bachelor's degree or higher) has been rising steadily over time, although it shows little cyclical fluctuations.



with the production function

$$y_{at} = Z_t \zeta_t^{\alpha_a} s_t^{1-\alpha_a}, \quad (41)$$

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

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

$$J_t^a = \pi_t^a + (1 - \rho^o) \mathbb{E}_t \beta \theta_{t+1} \frac{C_t}{C_{t+1}} J_{t+1}^a, \quad (42)$$

where  $\pi_t^a \equiv \max_{s_t} Z_t \zeta_t^{\alpha_a} s_t^{1-\alpha_a} - w_{st} s_t$ . Therefore, in contrast to our benchmark model, automation now entails a variable cost through the use of skilled workers as a complimentary input.

Aggregate output is the sum of goods produced by low-skill workers and those produced by robots operated by skilled workers. In particular, aggregate output is given by

$$Y_t = Z_t N_t + Z_t (\zeta_t A_t)^{\alpha_a} \bar{s}^{1-\alpha_a}. \quad (43)$$

We use the calibrated and estimated parameters in the benchmark model (where appropriate), and calibrate three additional parameters in this generalized model. We set  $\alpha_a = 0.3$ , such that the skilled labor share is 70% of the revenue generated by the technology using robots and skilled workers as inputs. We normalize the supply of skilled workers and calibrate the average value of the automation-specific productivity  $\bar{\zeta}$  such that the model implies a steady-state skill premium of 58%, in line with the ratio of median weekly earnings of workers with a bachelor's degree or higher to those of workers with some college or associate degrees.

Figure 9 shows the impulse responses following a positive automation-specific productivity shock. The shock lowers unemployment and raises vacancies in the short run. Overtime, however, unemployment overshoots its steady state. These patterns reflect the two opposing forces created by the automation mechanism that we have discussed in the benchmark model: the job displacing effect and the job creating effect. The shock boosts the present value of automation, raising the automation probability and labor productivity, which is measured by the ratio of aggregate output to aggregate employment, including both skilled and unskilled workers.

The increased threat of automation raises the firm's reservation value in wage bargaining, reducing the wage rate for unskilled workers, as shown in the figure. Since robots and skilled workers are complementary in production, increases in automation raises the wage rate for skilled workers, leading to an increase in the skill premium and income inequality.

Since the skilled wage rises while the unskilled wage falls, it is *a priori* unclear how the shock affects aggregate labor income and the labor share in aggregate output. Under our calibration, the model predicts that the labor share (defined as the ratio of aggregate labor

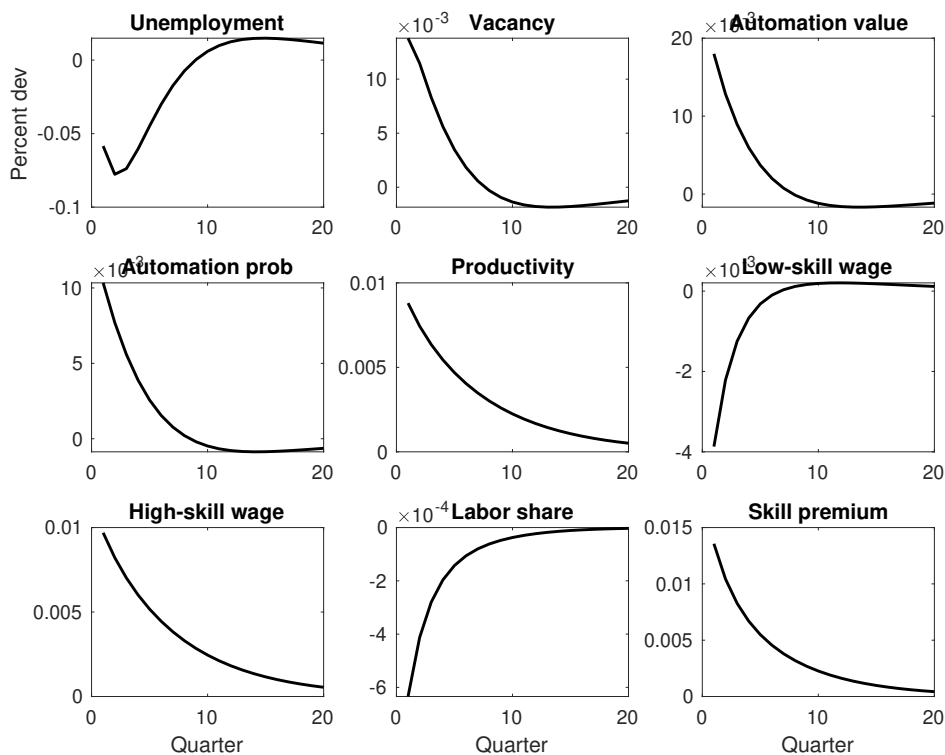


FIGURE 9. Impulse responses to a positive automation-specific shock in the alternative model with both skilled and unskilled workers.

income—including skilled and unskilled labor—to aggregate output) falls in response to a positive automation-specific shock. Thus, the decline in unskilled wage income stemming from the threat of automation dominates the rise in skilled wage income.<sup>21</sup>

Since the automation technology requires both robots and skilled workers as complementary input factors, the flow cost of operating the automation technology increases with the wage rate of skilled workers. All else being equal, a business cycle boom that raises skilled wages would reduce net value of automation, mitigating the increase in the automation probability. This feature is parallel to that for operating the manual technology, under which an increase in the wage rate of unskilled workers would reduce the present value of employment and thus discourage hiring. Nonetheless, the central mechanism of our benchmark model remains robust: increased automation raises labor productivity, dampens wage increases, boosts fluctuations in labor market tightness, and reduces the labor income share. As shown in Table 4, the model with heterogeneous worker skills also generates a large volatility of the v-u ratio relative to that of real wages, although the magnitude of the relative volatility

<sup>21</sup>The labor share also declines following a positive TFP shock and a positive discount factor shock. Details are available in the online appendix.

is modestly smaller than that in the benchmark economy (26.44 vs. 36.93), reflecting the procyclical costs (i.e., skilled wages) of operating the automation technology.

## VI. CONCLUSION

We have studied the role of automation in explaining the observed labor market dynamics in a tractable quantitative general equilibrium framework. The threat of automation raises the firm's reservation value in wage bargaining, 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 and also leads to countercyclical labor share.

Our estimated general equilibrium model shows that the automation channel is quantitatively important. In particular, the automation channel helps account for the large volatility in unemployment and job vacancies relative to that of real wages, a puzzling observation through the lens of the standard DMP model with labor search frictions. Our analysis also suggests that search frictions and automation are both important for generating the observed large Shimer volatility ratio.

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 tractable model with automation provides a useful framework that can be enriched to study the relative importance of these alternative channels and to improve our understanding of labor market policies. Future research along these lines should be important and promising. Our work represents a useful step in that direction.

## 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)

## REFERENCES

- ACEMOGLU, D., C. LELARGE, AND P. RESTREPO (2020): “Competing with Robots: Firm-Level Evidence from France,” *AEA Papers and Proceedings*, 110, 383–388.
- ACEMOGLU, D. AND P. RESTREPO (2018): “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,” *American Economic Review*, 108, 1488–1542.
- (2020): “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy*, 128, 2188–2244.
- 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., F. LEVY, AND R. J. MURNANE (2003): “The Skill Content of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, 118, 1279–1333.
- 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.
- 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.
- 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.
- HAGEDORN, M. AND I. MANOVSKII (2008): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited,” *American Economic Review*, 98, 1692–1706.
- HALL, R. E. AND P. R. MILGROM (2008): “The Limited Influence of Unemployment on the Wage Bargain,” *American Economic Review*, 98, 1653–1674.
- HE, H. AND Z. LIU (2008): “Investment-Specific Technological Change, Skill Accumulation, and Wage Inequality,” *Review of Economic Dynamics*, 11, 314–334.
- HERSHBEIN, B. AND L. B. KAHN (2018): “Do Recessions Accelerate Routine- Biased Technological Change? Evidence from Vacancy Postings,” *American Economic Review*, 108, 1737–1772.
- JAIMOVICH, N. AND H. E. SIU (2020): “Job Polarization and Jobless Recoveries,” *The Review of Economics and Statistics*, 102, 129–147.
- KARABARBOUNIS, L. AND B. NEIMAN (2013): “The Global Decline of the Labor Share,” *Quarterly Journal of Economics*, 129, 61–103.

- KOPYTOV, A., N. ROUSSANOV, AND M. TASCHEREAU-DUMOUCHEL (2018): “Short-run pain, long-run gain? Recessions and technological transformation,” *Journal of Monetary Economics*, 97, 29–44.
- KRUSELL, P., L. E. OHANIAN, J. V. RÍOS-RULL, AND G. L. VIOLANTE (2000): “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 68, 1029–1053.
- LEDUC, S. AND Z. LIU (2020): “The Weak Job Recovery in a Macro Model of Search and Recruiting Intensity,” *American Economic Journal: Macroeconomics*, 12, 310–343.
- 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.