

The Quantitative Importance of Technology and Demand Shocks for Unemployment Fluctuations in a Shopping Economy

Paweł Borys*

SGH Warsaw School of Economics

Paweł Doligalski

University of Bristol

Paweł Kopiec

SGH Warsaw School of Economics

May 5, 2021

Abstract

We construct and estimate a business cycle model with search and matching frictions in the labor market and in the product market. We show that the dynamic structure of the model and the endogenous job separation rate are important to accurately represent the empirical responses to the technology and the demand shocks. Our main finding is that the demand shock explains at least 58% of the unemployment fluctuations in the US, while the technology shock accounts for the residual.

Keywords: technology shocks, demand shocks, unemployment, business cycles.

JEL Codes: E24, E32.

1. Introduction

How important are technology and demand shocks in driving the unemployment fluctuations? The real business cycle literature stemming from Kydland and Prescott (1982) famously concluded that business cycles are driven by shocks to technology. The new class of macroeconomic models with shopping frictions attempts to revisit this question.

*Corresponding author: Paweł Doligalski, Department of Economics, University of Bristol, The Priory Road Complex, Priory Road, BS8 1TU Bristol, the United Kingdom, email: pawel.doligalski@bristol.ac.uk. We are grateful for useful suggestions of the associate editor and the two anonymous referees, as well as for the comments of Jakub Growiec, Krzysztof Pytka and Antoine Lepetit. Financial support from the Polish National Science Center (Narodowe Centrum Nauki) grant no. 2014/15/N/HS4/01342 is gratefully acknowledged.

In particular, Michaillat and Saez (2015) incorporate search and matching frictions in the labor and the product markets into the general-disequilibrium model of Barro and Grossman (1971). They qualitatively conclude that the labor market fluctuations mostly reflect demand shocks rather than technology shocks.

We complement these finding *quantitatively*. We extend the static framework of Michaillat and Saez (2015) into a dynamic, stochastic, general equilibrium model with frictional goods and labor markets. In addition, we endogenize the job separation rate. Both extensions play a crucial role in matching the empirical properties of the business cycle. The model is estimated on the US data using Bayesian methods. We find that the demand shock, expressed as a stochastic shifter of the marginal utility from consumption, accounts for at least 58% of the variance of unemployment and at least 74% of the variance of output. The technology shock accounts for the residual variance. These results are important for stabilization policy: it was shown in the framework similar to ours that effectiveness of various macroeconomic instruments (e.g. government spending or tax cuts) depends crucially on whether the fluctuations are supply or demand-driven (Ghassibe and Zanetti 2020).

In our model both the labor and the product markets are subject to search and matching frictions and rigidities. Just as a firm posts a costly vacancy to hire a worker, a household embarks on costly shopping visits to purchase consumption goods. The probability of a successful match with a vendor depends, through the matching function, on the aggregate number of shopping visits and the aggregate production capacity. Crucially, as the demand for goods varies in response to shocks, producers experience fluctuations in the probability of trade. Facing worsening business conditions and rigid wages, producers strongly adjust the size of the work force. This intuitive transmission mechanism shows why demand shocks are so powerful in driving output and unemployment fluctuations in our model.

In the labor market, our model features a vacancy-posting cost expressed in terms of labor: the firm directs some of its employees to the hiring process. We discover that such cost specification in the dynamic model generates a novel trade-off between employment and output which helps the model match the empirical responses to the technology shock. In the dynamic model with the cost of vacancy expressed in labor, hiring a worker can contribute *negatively* to the contemporaneous labor input devoted to production activities. That happens when the labor endowment of a new hire is lower than the cost of hiring. Hence, raising current employment can lower current output. The firm is still willing to post vacancies as long as the continuation value from hiring an employee — the expected future profit — is sufficiently high. Now, consider a transitory technological improvement. To take advantage of higher productivity, the firm directs more employees to production activities at the expense of hiring. As a result, the technology improvement leads to higher output and lower employment, as widely documented in the data (see Galí and Rabanal 2004 for the review). Thus, the vacancy-posting cost expressed

in labor offers a novel microfoundation for the observed response of unemployment to technological improvements. Furthermore, our model features endogenous separations rate: the firm can adjust its workforce by varying either the hiring or the firing intensity. In line with empirical evidence (Trigari 2009; Canova et al. 2013), we find that the firing margin plays an important role after both the technology and the demand shocks.

We calibrate the model using the Bayesian techniques on the time series of US unemployment and labor productivity. The model replicates well the empirical correlations and responses to shocks. Since we allow for a non-zero correlation between the exogenous shocks, the decomposition of the variance of model variables into the contributions of the two shocks is not unique. Rather, we obtain the intervals of possible shares. The demand shocks explains most of the variation of unemployment (58–98%) and output (74–93%), the residual being explained by the technology shock. We verify that these results are robust to assuming that the separation rate is constant. More surprisingly, we show that the demand shock remains a dominant driver of unemployment even when prices and real wages are flexible. The intuition is that in equilibrium the firm responds to the demand shock — a temporary increase of the marginal utility from consumption of the household — with greater current output. Greater current output, in the presence of the trade-off between employment and output, is provided by lowering current employment and, hence, lowering future output. Thus, unlike in the static framework of Michaillat and Saez (2015), in our dynamic framework the demand shocks can move real variables even when prices and real wages are flexible.

There is strong empirical evidence for search and matching frictions in the product markets. First, shopping is both costly and risky for consumers. In the US, individuals spend on average 42 minutes shopping daily (Petrosky-Nadeau et al. 2016) and face a non-trivial risk of 7.8% of their desired products being temporarily unavailable in store (Bils 2016). On the other side of the market, frictions manifest themselves in the idle production capacity. Michaillat and Saez (2015) report that the US workers are idle 14.8% and 17.3% of working time in the non-manufacturing sectors and in the manufacturing sector, respectively. Furthermore, Boehm and Pandalai-Nayar (2018) report that 80% of plant managers cite insufficient demand as the main reason for producing below the maximal capacity. They also find that the average capacity utilization is procyclical, which is consistent with our model.

The existing literature strongly supports incorporating shopping frictions into the business cycle models. Shopping frictions improve the propagation of the technology shocks (Petrosky-Nadeau and Wasmer 2015), reconcile puzzles in international business cycles (Bai and Rios-Rull 2015) and are useful in explaining facts about inventories over the business cycles (Den Haan 2013; Duras 2016). More related to our work, the shopping frictions may lead to self-fulfilling unemployment fluctuations (Kaplan and Menzio 2016) and amplify the non-technology shocks, allowing to model deep recessions (Huo and Rios-Rull 2013). Bai et al. (2019) conclude that the preference shock similar to

ours is a major force behind business cycle fluctuations, yet their model does not include unemployment.

The relative importance of technology and demand shocks for unemployment has been studied in the New Keynesian framework, which features oligopolistic competition and nominal rigidities. Galí et al. (2012) study the extension of the medium-scale DSGE model by Smets and Wouters (2003) with eight shocks and conclude that at the business cycle frequency demand shocks explain approximately 50% of the variance of unemployment, while the productivity shock explains less than 5%. In comparison, the approach of Michaillat and Saez (2015), on which we build on, allows to study unemployment and capacity utilization over the business cycle in a parsimonious and analytically tractable framework. As such, we view our findings as complementary to those of the New Keynesian literature.

Structure of the paper. In the next section we describe our modeling framework. We estimate the model and examine its fit with the data in Section 3. Main results are reported in Section 4. The last section concludes.

2. Framework

Our model extends the static framework of Michaillat and Saez (2015) in two respects. First, our model is dynamic, with employment as an endogenous state variable. Second, we also allow for endogenous firing, which becomes important in matching the model to the empirical evidence.

2.1. Economic environment

Time is discrete and indexed by $t = 0, 1, \dots$. Production technology uses effective labor n_t as the only input and features decreasing returns to scale. In particular, the total capacity of the supply side in period t is $k_t = A_t \cdot F(n_t)$ where A_t is the aggregate technology shock and F is a concave production function. There are two agents in the model: a representative household and a representative firm, and three markets: for goods, labor and money balances. We follow Michaillat and Saez (2015) by assuming that the markets for goods and labor are frictional.

On the goods market, the aggregate capacity of the economy k_t combined with the household's shopping visits v_t generate the following number of trades

$$M(k_t, v_t) = \left(k_t^{-\gamma} + v_t^{-\gamma} \right)^{-\frac{1}{\gamma}}, \quad (1)$$

where $\gamma > 0$ governs the elasticity of substitution of matched inputs. The goods market tightness x_t is defined as a ratio of shopping visits and aggregate capacity: $x_t = \frac{v_t}{k_t}$.

Since the matching technology M features constant returns to scale, the probability of utilizing the unit of capacity $f(x_t) = \frac{M(k_t, v_t)}{k_t} = (1 + x_t^{-\gamma})^{-\frac{1}{\gamma}}$ and the probability of a successful shopping visit $q(x_t) = \frac{M(k_t, v_t)}{v_t} = (1 + x_t^\gamma)^{-\frac{1}{\gamma}}$ can both be expressed as functions of the product market tightness alone. In particular, the capacity utilization rate $f(x_t)$ is strictly decreasing in x_t , while the shopping success rate $q(x_t)$ is strictly increasing in x_t . We define the aggregate output y_t as the aggregate number of realized trades: $y_t = M(k_t, v_t)$.

We model the labor market analogously to the goods market. Labor matching technology \hat{M} combines the unemployment rate u_t and the number of vacancies \hat{v}_t posted by the firm to generate the following number of new hires

$$\hat{M}(u_t, \hat{v}_t) = \left(u_t^{-\hat{\gamma}} + \hat{v}_t^{-\hat{\gamma}} \right)^{-\frac{1}{\hat{\gamma}}}, \quad (2)$$

where $\hat{\gamma} > 0$ controls the elasticity of substitution of unemployment and vacancies. The labor market tightness is defined as $\theta_t = \frac{\hat{v}_t}{u_t}$. Job-finding rate and vacancy-filling rate are $\hat{f}(\theta_t) = \frac{\hat{M}(u_t, \hat{v}_t)}{u_t} = \left(1 + \theta_t^{-\hat{\gamma}} \right)^{-\frac{1}{\hat{\gamma}}}$ and $\hat{q}(\theta_t) = \frac{\hat{M}(u_t, \hat{v}_t)}{\hat{v}_t} = \left(1 + \theta_t^{\hat{\gamma}} \right)^{-\frac{1}{\hat{\gamma}}}$, respectively. An increase of labor market tightness leads to higher job-finding rate $\hat{f}(\theta_t)$ and lower vacancy filling rate $\hat{q}(\theta_t)$.

2.2. Household

The representative household consists of a unit measure of workers. Each worker inelastically supplies one unit of labor. Following Merz (1995) and Andolfatto (1996), we assume that income and consumption is shared within the household, so that each worker enjoys identical utility. Consequently, we will refer to the utility of the household rather than the utility of the individual worker. The household per-period utility function over consumption c_t and real money balances $\mu_t = \frac{m_t}{p_t}$ is

$$u(c_t, \mu_t, \chi_t) = \frac{\chi_t}{1 + \chi_t} c_t^{\frac{\epsilon-1}{\epsilon}} + \frac{1}{1 + \chi_t} \mu_t^{\frac{\epsilon-1}{\epsilon}}, \quad (3)$$

where χ_t is a Markovian preference shock which increases the marginal utility from consumption relative to real money balances. The parameter $\epsilon > 1$ governs the elasticity of substitution between consumption and real money balances. The household discount factor is $\beta < 1$.

The goods market is frictional and workers need to make shopping visits v_t to purchase consumption. Each successful visit allows to buy a unit of consumption good at price p_t . However, visits are costly. Intuitively, the matching costs incurred by buyers on the goods market can take form of either a goods cost (e.g. the travel agent's fee for booking holidays with a travel agency) or a time or effort cost (e.g. the time spent in a queue or when traveling to a seller). We follow Michailat and Saez (2015) in making a simplifying

assumption that the shopping cost is expressed entirely in goods. Specifically, each visit requires spending ρ goods. Furthermore, only a share $q(x_t)$ of visits is successful. Hence, the following constraint links the number of visits and consumption

$$c_t = q(x_t) \cdot v_t - \rho \cdot v_t. \quad (4)$$

We summarize the product market frictions with a wedge $\tau(x_t) = \frac{\rho}{q(x_t) - \rho}$. This wedge tells us how many goods are spent on covering the shopping cost per each unit of consumption. The above constraint can be then expressed as $(1 + \tau(x_t)) \cdot c_t = q(x_t) \cdot v_t$.

The budget constraint equates households expenditures on consumption and money balances on the left-hand side with various sources of income on the right-hand side:

$$p_t \cdot (1 + \tau(x_t)) \cdot c_t + m_t = p_t \cdot w_t \cdot (1 + s_t \cdot \sigma_t^e) \cdot l_t + m_{t-1} + \tilde{m}_t + d_t. \quad (5)$$

The first term on the right-hand side is the labor income. The employment rate is l_t and each employed worker earns a real wage w_t . In addition, $\sigma_t^e \cdot l_t$ workers are endogenously separated from the employer and eligible for severance pay. The severance pay in real terms is $s_t \cdot w_t$ where s_t is the severance pay rate. We include severance payments in the model, because they are necessary to decentralize the efficient allocation.¹ We also allow for exogenous separations, which are implicit in the definition of the employment rate l_t and which will be introduced explicitly in the firm's problem. We assume that exogenously separated workers are not eligible for severance pay. The consecutive terms stand for money holdings chosen in the previous period m_{t-1} and exogenous growth of nominal money holdings \tilde{m}_t , which for simplicity is assumed to be deterministic. Finally, d_t is the dividend from the firm. We abstract from unemployment insurance and labor income taxation — including these elements would not affect our results.

The household takes the process for equilibrium objects, which include broadly understood prices (p_t, w_t, s_t) and tightnesses (x_t, θ_t) , as given. The intertemporal maximization problem of the household is

$$\max_{\{c_t, m_t\}} \mathbb{E}_0 \sum_{t=0}^{+\infty} \beta^t \cdot u(c_t, \mu_t, \chi_t) \text{ s.t. (4) and (5) for all } t = 0, 1, \dots \quad (6)$$

The first-order condition of the household for any $t = 0, 1, \dots$ is

$$\frac{1}{1 + \tau(x_t)} \frac{\chi_t}{1 + \chi_t} c_t^{-\frac{1}{\epsilon}} = \frac{1}{1 + \chi_t} \mu_t^{-\frac{1}{\epsilon}} + \beta \cdot \mathbb{E}_t \left[\frac{\chi_{t+1}}{1 + \chi_{t+1}} \frac{1}{1 + \tau(x_{t+1})} \frac{1}{\pi_{t+1}} c_{t+1}^{-\frac{1}{\epsilon}} \right], \quad (7)$$

where $\pi_{t+1} = \frac{p_{t+1}}{p_t}$ is the inflation rate realized in $t + 1$. The interpretation of this condition is standard: the marginal benefit from spending p_t unit of money on consumption today (the left-hand side) must equal the expected marginal benefit of increasing real

¹In Proposition 1 we show that an appropriate process for severance pay rate s_t aligns the firm's firing decision with the efficient separations chosen by the social planner.

money holdings by a unit today and spending it on consumption in the next period (the right-hand side). Note that the marginal benefit from consumption expenditure is modified by wedge $\tau(x_t)$. When the market for goods becomes tighter, this wedge increases and purchasing consumption becomes more costly.

2.3. Firm

The representative firm operates a decreasing returns to scale production technology with labor as the only input. The firm maximizes the discounted stream of profits by choosing vacancies and the firing threshold. We model firing in a way which follows Den Haan et al. (2000) and Krause and Lubik (2007).

The order of events within a period is as follows. The firm starts period t with $(1-\sigma) \cdot l_{t-1}$ employees, where l_{t-1} is the employment rate last period and σ is the rate of exogenous separations. The firm posts vacancies \hat{v}_t which are filled at rate $\hat{q}(\theta_t)$. Then, each employed worker i draws an idiosyncratic and time-independent match quality $a_{t,i} \in \mathbb{R}$ from the distribution H with mean 1. The match quality indicates the number of efficiency units of labor of a worker. The firm fires workers with match quality below a threshold \tilde{a}_t and the average match quality after firing is denoted by $\bar{a}_t = \frac{\int_{\tilde{a}_t}^{\infty} a dH(a)}{1-H(\tilde{a}_t)}$. We allow for the match quality to be negative to ensure the interior endogenous firing rate. Next, the production takes place, which we describe below. The employment at the time of production is

$$l_t = (1 - H(\tilde{a}_t)) \cdot ((1 - \sigma) \cdot l_{t-1} + \hat{q}(\theta_t) \cdot \hat{v}_t). \quad (8)$$

At the end of the period, share σ of employed workers separate exogenously. The rate of endogenous separations σ_t^e introduced in the household problem is equal to $\frac{H(\tilde{a}_t)}{1-H(\tilde{a}_t)}$.

Following Michaillat and Saez (2015), we express the vacancy-posting cost in terms of labor. Other papers making this assumptions include Farmer (2008) and Shimer (2010). Intuitively, it means that the firm directs its employees to recruitment activities rather than outsourcing them to a headhunter, in which case the vacancy-posting cost could be expressed in terms of goods. There is a strong evidence that a large share of recruitment is conducted by own employees. Among medium and large Italian firms, 49% conduct recruitment activities exclusively internally and only 9.1% outsource recruitment fully (Ordanini and Silvestri 2008). Regarding the US, about 60% of companies conduct most or all of the recruitment activities internally (Cappelli 2019).

We assume that posting a single vacancy costs the firm $\hat{\rho}$ units of effective labor. The effective labor which remains for production activities is

$$n_t = \bar{a} \cdot l_t - \hat{\rho} \cdot \hat{v}_t. \quad (9)$$

Recall that the productive capacity of the firm is $k_t = A_t \cdot F(n_t)$, where A_t is a Markovian

technology shock. A fraction $f(x_t)$ of this capacity is matched in the goods market and, hence, sold. The firm's dividend is given by the difference between the revenue and the wage bill, which includes severance payments:

$$d_t = p_t \cdot f(x_t) \cdot k_t - p_t \cdot w_t \cdot \left(1 + s_t \cdot \frac{H(\tilde{a}_t)}{1 - H(\tilde{a}_t)}\right) \cdot l_t. \quad (10)$$

The firm maximizes the expected, discounted stream of dividends, taking the process for the equilibrium objects $(p_t, w_t, s_t, x_t, \theta_t)$ as given. Define the stochastic discount factor as $\Delta_t = \beta^t \cdot \frac{u_c(c_t, \chi_t)}{p_t \cdot (1 + \tau(x_t))}$, where u_c is the household's marginal utility from consumption. The firm's intertemporal maximization problem is

$$\max_{\{l_t, \hat{v}_t, \tilde{a}_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \Delta_t \cdot d_t \text{ s.t. (8) for all } t = 0, 1, \dots, \quad (11)$$

given initial employment l_{-1} . For notational convenience, write the Lagrange multiplier on the constraint (8) as $p_t \Delta_t \lambda_t$. The first-order condition with respect to vacancies is

$$f(x_t) \cdot A_t \cdot F'(n_t) \cdot \left(\bar{a}_t - \frac{\hat{\rho}}{\hat{q}(\theta_t) \cdot (1 - H(\tilde{a}_t))}\right) + \lambda_t = w_t \cdot \left(1 + s_t \cdot \frac{H(\tilde{a}_t)}{1 - H(\tilde{a}_t)}\right). \quad (12)$$

The left-hand side captures the marginal benefit from hiring $\frac{1}{1 - H(\tilde{a}_t)}$ workers, which implies gaining exactly one worker after the endogenous firing takes place. The first term is the expected contribution of an additional worker to the firm's sales, net of the labor cost of hiring. The second term λ_t captures the continuation value of an employee to the firm. This marginal benefit is equal to the marginal cost of wages and severance packages on the right-hand side.

The firing decision is determined by the following condition

$$f(x_t) \cdot A_t \cdot F'(n_t) \cdot \tilde{a}_t + \lambda_t = (1 - s_t) \cdot w_t. \quad (13)$$

Reducing the firing threshold marginally contributes to sales and yields a positive continuation value from an employee. On the other hand, retaining an employee is costly due to the difference between the wage and the severance pay.²

Finally, the continuation value of employment is

$$\lambda_t = (1 - \sigma) \cdot \mathbb{E}_t \left[\frac{\Delta_{t+1}}{\Delta_t} \cdot \pi_{t+1} \cdot f(x_{t+1}) \cdot A_{t+1} \cdot F'(n_{t+1}) \cdot \frac{\hat{\rho}}{\hat{q}(\theta_{t+1})} \right]. \quad (14)$$

²In some contexts the firing cost is modeled as being partially a transfer to a worker (as our severance pay) and partially a deadweight loss cost, e.g. due to the administrative burden of firing. In the US the deadweight loss of firing is driven by the adoption of the Wrongful-Discharge Laws, limiting firing employees at will and creating a threat of costly litigation (Autor et al. 2006). Dertouzos and Karoly (1992) estimate the cost of Wrongful-Discharge Laws to be \$100 per termination, and only 37.5% of it can be treated as a deadweight loss. This figure amounts to less than 0.01% of quarterly output per worker in the labor force. Including such firing cost is unlikely to affect our results and we decide to ignore it for the sake of simplicity.

Intuitively, the continuation value of employment is large if the employment relationship is likely to survive to the next period and if it is costly to hire a replacement worker.

2.4. Equilibrium

The equilibrium requires that the number of successful shopping visits is equal to the number of goods sold: $q(x_t) \cdot v_t = f(x_t) \cdot k_t$, and that the number of hired workers is equal to the number of filled vacancies: $\hat{f}(\theta_t) \cdot u_t = \hat{q}(\theta_t) \cdot \hat{v}_t$. The market for money is standard and clears when the demand for money is equal to the supply: $m_t = m_{t-1} + \tilde{m}_t$.

Three variables remain to be determined: the price p_t , the real wage w_t and the severance pay rate s_t . Note that the model places no additional restrictions on these objects. Intuitively, the market clearing condition for the goods market cannot determine the equilibrium price if we already use it to determine the goods market tightness, and analogously for the labor market. Choosing additional equations to determine these broadly understood prices is akin to selecting an equilibrium. We follow Michailat and Saez (2015) and consider two types of equilibria: a fixprice equilibrium and a competitive equilibrium.

Fixprice equilibrium In a fixprice equilibrium all broadly understood prices are constant: $p_t = p$, $w_t = w$ and $s_t = s$. The frictional markets are cleared only by the adjustments of tightnesses. This equilibrium illustrates the impact of strong rigidities in setting prices and adjusting labor contracts.

Competitive equilibrium In a competitive equilibrium, inspired by Moen (1997), prices decentralize the efficient allocation. This equilibrium illustrates how the economy operates absent rigidities in setting prices and adjusting labor contracts. Michailat and Saez (2015) shows that in the static model the competitive equilibrium can also be attained with Nash bargaining. As is standard for the models with money in the utility, we define the efficient allocation as the one which maximizes the expected discounted utility of the household for a given sequence of real money balances. We set up the social planner problem and derive the efficiency conditions in Appendix A. The following proposition describes goods prices, real wages and severance pay rates in the competitive equilibrium.

Proposition 1. *Consider initial employment l_{-1} , the sequence of real money balances $\{\mu_t\}_{t=0}^{\infty}$ and the sequence of exogenous shocks $\{A_t, \chi_t\}_{t=0}^{\infty}$. Take the corresponding efficient allocation $\{c_t, v_t, l_t, \hat{v}_t, \tilde{a}_t, n_t, k_t\}_{t=0}^{\infty}$ with the associated tightnesses $\{x_t, \theta_t\}_{t=0}^{\infty}$.*

The sequence of prices, real wages and severance pay rates $\{p_t, w_t, s_t\}_{t=0}^{\infty}$ decentralizes

the efficient allocation if for all $t = 0, 1, \dots$ it satisfies

$$\frac{p_{t+1}}{p_t} = \frac{\beta u_c(c_{t+1}, \chi_{t+1})}{u_c(c_t, \chi_t) - (1 + \tau(x_t)) \cdot u_\mu(\mu_t, \chi_t)}, \quad (15)$$

$$w_t = f(x_t) \cdot A_t \cdot F'(n_t) \cdot \hat{\rho} \cdot (\delta_t + \zeta_t), \quad (16)$$

$$s_t = 1 - \frac{\zeta_t}{\delta_t + \zeta_t}, \quad (17)$$

where $p_0 > 0$, $\delta_t = \frac{1}{\hat{q}(\theta_t) + \theta_t \cdot \hat{q}'(\theta_t)} - \frac{1}{\hat{q}(\theta_t)}$, $\phi_t = u_c(c_t, \chi_t) \cdot M_k(k_t, v_t)$ and

$$\zeta_t = -(1 - \sigma) \cdot \beta \cdot \mathbb{E}_t \left[\frac{\phi_{t+1}}{\phi_t} \cdot \frac{A_{t+1}}{A_t} \cdot \frac{F'(n_{t+1})}{F'(n_t)} \cdot \left(\delta_{t+1} + \frac{\theta_{t+1}^2 \cdot \hat{q}'(\theta_{t+1})}{\hat{q}(\theta_{t+1}) + \theta_{t+1} \cdot \hat{q}'(\theta_{t+1})} \right) \right],$$

Proof. In Appendix A. □

In the competitive equilibrium the goods market tightness x_t is constant, which echoes the result of Michaillat and Saez (2015). Consequently, the goods market clears after shocks by adjusting the price rather than the tightness. The inflation rate described by (15) ensures that a constant tightness is consistent with the household's choice of shopping visits.

The real wage in the competitive equilibrium, given by (16), prevents the excessive vacancy creation of the firm. It makes the firm internalize two labor market externalities: a static and a dynamic one. The correction of the static externality, governed by δ_t , makes the firm internalize the negative impact of an additional vacancy on the vacancy-filling rate of its existing vacancies due to congestion. The correction of the dynamic externality, governed by ζ_t , makes the firm internalize that an additional vacancy today will lower the stock of unemployed next period and, hence, decrease the future vacancy-filling rate. This dynamic externality is similar to the dynamic problem of a monopolist selling a durable good (Bond and Samuelson 1984).

Finally, the competitive severance pay rate ensures that separations are efficient and is driven by two effects. First, ignore the dynamic aspects of the model by setting ζ_t to 0. Then the severance pay rate is equal to 100%: the firm pays exactly the same compensation to the workers which are fired and the workers which are retained. In this way the firm fully internalizes the impact of the firing decision on the household budget constraint, see Postel-Vinay and Turon (2014) for a similar mechanism. Second, in the presence of the dynamic effects captured by ζ_t , the severance pay rate may deviate from 100% to account for the impact of separations on the future labor market tightness.

3. Parametrization of the model

In this section we select the structural parameters of the model. As a baseline framework we choose a fixprice equilibrium model where the price, the real wage and the severance

pay rate are fully rigid. This model, as we find, approximates the data very well. We describe the calibration of the steady state and the estimation of the cyclical parameters of the model. Then we document the fit of the estimated model to the data.

3.1. Data sources

In the parametrization of the model and the validation of its fit we use the following US data series obtained from Federal Reserve Economic Data. The data period is 1967–2019 and all series are seasonally adjusted at the quarterly frequency. We measure output y_t as real output of the nonfarm business sector divided by the size of the labor force. For u_t we use the unemployment rate. The labor productivity $\frac{y_t}{l_t}$ is the real output per person of the nonfarm business sector. We measure inflation π_t by a growth rate of the consumer price index for all urban consumers. We follow Michailat and Saez (2015) in constructing the empirical series of capacity utilization and labor market tightness. The data equivalent of $f(x_t)$ is the total industry capacity utilization, which we transform to measure the utilization rate at the actual employment rather than full employment. To obtain the empirical series of labor market tightness $\theta_t = \frac{\hat{v}_t}{u_t}$ we proxy vacancies \hat{v}_t with the quarterly mean of the monthly help-wanted index, which is available for years 1967-2016 (Barnichon 2010).

We compute the cyclical component of all variables by taking logs (with an exception for u_t) and applying the one-sided Hodrick-Prescott filter with a smoothing parameter equal to 1600.

3.2. Calibration of the steady state

The time period in our model corresponds to a quarter and the targets of our calibration are moments characterizing US economy. We assume that the steady state of our model economy coincides with the competitive equilibrium, which is consistent with the intuition that rigidities are inconsequential in the long run. The first panel of Table 1 presents the parameters which determine the steady state. We explain their calibration below. The steady state values of the equilibrium variables are denoted by omitting the time index.

We set the value of discount factor β at 0.99 and assume that the production function is $F(n) = n^\alpha$ with $\alpha = 0.66$. Thus, for β and α we use standard values from the literature. We follow Michailat and Saez (2015) by assuming that parameter ϵ , governing the elasticity of substitution between consumption and real money balances, is greater than 1. We fix $\epsilon = 2$ which implies a square root utility. Note that in our model ϵ is also equal to the intertemporal elasticity of substitution. While it is common in dynamic models to assume that the intertemporal elasticity of substitution is smaller than 1, in Appendix E we show that in such case the fixprice equilibrium for a given price and real wage may

cease to be unique. While the study of equilibrium multiplicity is an interesting research avenue in the context of demand-driven fluctuations (see e.g. Kaplan and Menzio 2016), it is beyond the scope of this paper.

The vacancy-posting cost $\hat{\rho}$ is set to 0.85 to match the average unemployment rate of 6.1%. The parameter of the labor matching technology $\hat{\gamma}$ is set to 0.51 to match the elasticity of new hires with respect to unemployment of 0.5, as reported by Petrongolo and Pissarides (2001). The average quarterly job-finding rate out of unemployment is 25.4% (Krusell et al. 2017), which implies the average separation rate of 1.62%. Regarding the distribution of the match quality $H(\cdot)$, note that we do not need to specify its entire shape. To find the steady state, we only need to specify the the share of matches which are endogenously destroyed $H(\tilde{a})$ and the mean match quality after separations \bar{a} . We calibrate $H(\tilde{a})$ to match the 68% share of exogenous separations in total separations (Den Haan et al. 2000), which implies that $\sigma = 1.1\%$ and $H(\tilde{a}) = 0.52\%$. We calibrate the value of \bar{a} , i.e. the mean match quality after separations, to 1.008, which is the mean of this variable in the calibrations of Den Haan et al. (2018), Walsh (2005) and Krause and Lubik (2007).³

The average quarterly capacity utilization is 81%, which we use as a value for $f(x)$. Petrosky-Nadeau and Wasmer (2015) report that the cost of search for goods is equal to 7% of labor earnings. These two moments allow us to pin down the parameter of the goods market matching function $\gamma = 0.22$ and the unit cost of shopping $\rho = 3.69 \cdot 10^{-8}$. Although the unit cost of shopping is very low, the search frictions in the goods market are not negligible. As noted above, in the steady state 19% of capacity is not sold. Finally, we assume that nominal money balances grow at the rate of inflation, keeping the real money balances fixed. We calibrate real money balances $\mu = 5.27$ to match the average quarterly inflation rate of 1%.

3.3. Estimation of the cyclical parameters

We estimate the fixprice equilibrium model, where the price, the real wage and the severance pay rate are constant. Exogenous shocks follow the standard autoregressive process: $\log(z_t) = (1 - \rho_z) \cdot \log(z_{t-1}) + \varepsilon_{z,t}$ for $z \in \{A, \chi\}$. The innovations are Gaussian, independent over time and potentially cross-correlated:

$$\begin{bmatrix} \varepsilon_{A,t} \\ \varepsilon_{\chi,t} \end{bmatrix} \sim N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_A^2 & \sigma_{A,\chi} \\ \sigma_{A,\chi} & \sigma_\chi^2 \end{bmatrix} \right). \quad (18)$$

Parameters $\rho_A, \sigma_A, \rho_\chi, \sigma_\chi$ and $\sigma_{A,\chi}$ are estimated with Bayesian methods, using the time series of labor productivity and unemployment as observables. To simulate the model,

³The calibrations of these three papers imply very similar values of \bar{a} : 1.007, 1.009 and 1.008, respectively. They all assume that the match quality is distributed log-normally. It is not a feasible solution for us, since it would imply no endogenous separations in the steady state. We verify that our results are robust to different values of \bar{a} .

Table 1: Parameter values

Parameter	Description	Value	Target/Source
<i>Steady state parameters</i>			
β	Discount factor	0.99	Standard value
α	Curvature of the production function	0.66	Standard value
ϵ	Elasticity of substitution of c_t and μ_t	2	Michaillat and Saez (2015)
$\hat{\gamma}$	Elasticity of substitution of \hat{v}_t and u_t	0.51	Elasticity of hires w.r.t. unemployment, Petrongolo and Pissarides (2001)
σ	Rate of exogenous separations	1.1%	Job-finding rate, Krusell et al. (2017)
$H(\bar{a})$	Rate of endogenous separations	0.52%	Share of endogenous separations, Den Haan et al. (2000)
\bar{a}	Mean match quality	1.008	Den Haan et al. (2000), Walsh (2005), Krause and Lubik (2007)
$\hat{\rho}$	Vacancy posting cost	0.85	Average unemployment rate
γ	Elasticity of substitution of v_t and k_t	0.22	Average capacity utilization
ρ	Unit cost of shopping	$3.7 \cdot 10^{-8}$	Cost of goods search, Petrosky-Nadeau and Wasmer (2015)
μ	Real money balances	5.27	Average CPI inflation
<i>Cyclical parameters</i>			
ρ_A	Persistence of technology shock	0.5	Estimated
σ_A	Standard deviation of technology shock	0.012	Estimated
ρ_χ	Persistence of demand shock	0.91	Estimated
σ_χ	Standard deviation of demand shock	0.022	Estimated
$Corr(A, \chi)$	Correlation of the shocks	-0.53	Estimated
$h(\bar{a})$	Density of match quality distribution at \bar{a}	0.6	Cyclicalty of labor market tightness

we also need to specify $h(\tilde{a})$, the density of the match quality distribution at the steady state firing threshold. We choose $h(\tilde{a}) = 0.6$ such that the estimated model matches the cyclical of the labor market tightness. Our results are robust to estimating this parameter jointly with the parameters of exogenous shocks processes. The estimates are reported in the second panel of Table 1. We find that the demand shock is more volatile and more persistent than the technology shock, and that the two shocks are strongly negatively correlated. The estimation details are reported in Appendix B.

3.4. Model fit

We compare the model correlations and volatilities of key variables with their data counterparts in Table 2. The model accurately replicates the empirical correlations of output, unemployment, labor productivity, capacity utilization and labor market tightness. In particular, the model does a very good job at matching the cyclical of these variables. The remaining correlations always have a correct sign, but sometimes deviate from the data in terms of magnitude.

Regarding volatilities, the model captures well the standard deviation of labor productivity as well as the fact that the labor market tightness is an order of magnitude more volatile than labor productivity, as noted by Shimer (2005). However, it overshoots the volatility of unemployment, which also leads to an excessive volatility of output and capacity utilization. It is likely the consequence of the vacancy-posting cost which is expressed entirely in labor. Suppose instead that the vacancy-posting cost was expressed partially in labor and partially in goods or services which need to be purchased in the market. When the aggregate demand is high, it is difficult to purchase goods and services and, hence, the effective vacancy-posting cost increases. This would limit the demand-driven fall of unemployment and contribute to its lower volatility. This channel is likely to be important since, as we will discuss shortly, unemployment fluctuations are mostly demand-driven. The specification of the vacancy-posting cost which is expressed partially in labor and partially in goods or services seems realistic, but its examination is beyond the scope of this paper and we leave it for future research.

4. Results


In this section we describe the responses of our baseline economy to exogenous shocks and discuss the underlying mechanism. Then we report the variance decomposition for the key variables. We also provide comparisons with an economy with flexible prices and labor contracts (competitive equilibrium) as well as with an economy in which the separation rate is constant.

Table 2: Correlations and volatilities of key variables: baseline model & (data)

	Correlation					Standard deviation
	y_t	u_t	$\frac{y_t}{l_t}$	$f(x_t)$	θ_t	
y_t	1	-0.87 (-0.84)	0.63 (0.63)	0.78 (0.56)	0.88 (0.88)	0.028 (0.019)
u_t	.	1	-0.16 (-0.14)	-0.63 (-0.32)	-0.64 (-0.91)	0.021 (0.008)
$\frac{y_t}{l_t}$.	.	1	0.45 (0.61)	0.65 (0.33)	0.013 (0.012)
$f(x_t)$.	.	.	1	0.97 (0.51)	0.038 (0.02)
θ_t	1	0.077 (0.123)

Note: Theoretical model moments from the fixprice equilibrium of the baseline model (not in brackets) and empirical moments (in brackets). Empirical moments are computed over time period 1967– 2019 with an exception of moments of θ_t which, due to data limitations, are computed over time period 1967– 2016.

Figure 1: Impulse responses in the baseline model



Note: Impulse responses for a positive, one standard deviation shock of each type.

4.1. Impulse responses in the baseline model

In the following figures for clarity reasons we depict the response to a shock of one type (either technology or demand), keeping the other shock fixed. Thus, in the figures and the discussion which follows we abstract from the correlation between shocks. Recall that in the baseline economy p_t , w_t and s_t are fixed.

The impact of a positive technology shock is depicted on Figure 1. A positive technology shock raises the capacity brought to the goods market both mechanically and through the increased productive labor at the firm. At the household level, a greater capacity leads to a higher shopping success rate. That, in turn, generates a substitution effect, pushing for more shopping visits, and an income effect, pushing for less shopping visits. The income effect dominates and the shopping intensity falls, which means that the capacity utilization falls as well. However, the capacity increase is stronger than the fall in the capacity utilization and, hence, output — the utilized capacity — raises. The technology shock raises unemployment, which happens due to increased firing rate. At the same time the firm also creates more vacancies, although this effect is more delayed and insufficient to prevent the fall of employment.

These responses are consistent with empirical evidence. It has been widely documented that technology improvements increase unemployment (e.g. Galí 1999; Basu et al. 2006; Canova et al. 2010, see Galí and Rabanal 2004 for a review). All of these studies, with an exception of Basu et al. (2006), find that technology improvements increase output. Furthermore, Michelacci and Lopez-Salido (2007) and Canova et al. (2013) observe that the raise in unemployment is predominantly driven by intensified firing rather than lower vacancy creation.

Figure 1 also presents the impact of the demand shock. The demand shock affects the economy via the intensified shopping behavior of the household. That, in turn, raises the capacity utilization and output. The labor market impact of the demand shock is very different from the technology shock: high demand leads to lower unemployment. It happens through both lower firing and — at least on impact — higher vacancy creation. These findings are corroborated by empirical evidence on the response to demand shocks. The empirical literature has mostly studied the monetary policy shock. Although we do not study such shock explicitly, it can be represented as an innovation to the money supply and it affects the economy analogously to the innovation to the marginal utility from consumption. Christiano et al. (2020) show that the expansionary monetary policy shock raises capacity utilization and lowers unemployment. Additionally, Trigari (2009) shows that the reduction of unemployment is caused mostly by a lower separation rate and, to a lower extent, by higher job creation.

4.2. The mechanism behind the impulse responses

Notice that our dynamic model, unlike the static framework of Michaillat and Saez (2015), correctly predicts that employment falls after the positive technology shock and can account for empirical changes in the firing rate. In what follows we describe the mechanism which allows our model to match these empirical facts.

4.2.1. Response to a technology shock

It is instructive to focus first on the model in which all separations are exogenous. We can shut down the endogenous firing channel by setting $H(\tilde{a}_t) = 0$ and $\bar{a}_t = 1$. The first-order condition with respect to vacancies becomes

$$f(x_t) \cdot A_t \cdot F'(n_t) \cdot \left(1 - \frac{\hat{\rho}}{\hat{q}(\theta_t)}\right) + \lambda_t = w_t. \quad (19)$$

Notice that the continuation value of employment λ_t fully summarizes the dynamic aspects of the labor market. A high value of λ_t implies that gains from hiring a worker are spread over multiple periods. In contrast, a low value of λ_t means that gains from a new hire are mostly contemporaneous. Specifically, if $\lambda_t = 0$, which happens when the rate of exogenous separations σ is 1, then the labor market in this model is equivalent to that from the static framework of Michaillat and Saez (2015).

Consider the term $1 - \frac{\hat{\rho}}{\hat{q}(\theta_t)}$ which captures the expected contribution of a new hire to productive labor n_t . A newly hired worker brings in expectation a unit of effective labor, but reduces the effective labor of the firm by $\frac{\hat{\rho}}{\hat{q}(\theta_t)}$ due to the cost of hiring. Hence, the sign of this term tells us whether a new hire contributes positively or negatively to contemporaneous labor input n_t .

First, suppose that the continuation value of employment is small: $\lambda_t < w_t$, which is true in the static model. The condition (19) implies that $1 - \frac{\hat{\rho}}{\hat{q}(\theta_t)}$ is positive. Hence, posting a vacancy raises total employment l_t as well as productive labor n_t and output y_t . Second, suppose that the continuation value of employment is large: $\lambda_t > w_t$, which is true in our model.⁴ One can show that this inequality is likely to hold when the rates of exogenous separations σ and vacancy-filling $\hat{q}(\theta_t)$ are not very large and the household is relatively patient.⁵ Now the condition (19) can hold only if $1 - \frac{\hat{\rho}}{\hat{q}(\theta_t)}$ is negative. The contemporaneous gain from an additional worker is not sufficient to compensate for the entire hiring cost, which means that issuing vacancies increases total employment, but reduces productive labor and output. Therefore, the dynamic model, combined with the vacancy-posting cost expressed in labor, generates a novel trade-off between employment and output.

⁴In the steady state of our model $\lambda = 1.8$ and $w = 0.5$.

⁵Equations (14) and (16) imply that in the steady state $\frac{\lambda}{w} = \frac{(1-\sigma)\beta}{\hat{q}(\theta) \cdot (\delta + \zeta)}$. δ and ζ , the labor market externality correction terms, depend mainly on the curvature of the labor matching function. In our calibration they sum up roughly to 1.

Consider a positive technology shock which makes the term $f(x_t) \cdot A_t$ go up. The return to productive labor increases and the firm has incentives to increase n_t . How can it be achieved? When $\lambda_t < w_t$, the firm can increase n_t by issuing additional vacancies and raising the overall employment. Hence, unemployment falls after the technology shock, which is the prediction of the static Michaillat and Saez (2015) framework. Instead, in the dynamic model with $\lambda_t > w_t$ there is a trade-off between the employment and output. The firm can increase productive labor n_t by limiting vacancy creation and saving on the vacancy-posting cost. As a result, following the technology shock output increases but employment falls, which is consistent with the above-mentioned empirical evidence. Thus, the vacancy-posting cost expressed in terms of labor is an alternative and natural microfoundation for the observed response of employment to technology shocks. In contrast to the New Keynesian microfoundation based on nominal rigidities (Galí and Rabanal 2004), our mechanism is independent of the extent of rigidities in the economy and present even in the competitive equilibrium.

As we see, the model with exclusively exogenous separations can account for the response of unemployment to the technology shock. However, it cannot account for the fact that the raise of unemployment is mostly driven by intensified firing (Michelacci and Lopez-Salido 2007; Canova et al. 2013). To capture this feature of the data, we need to consider endogenous separations.

Let's return to our baseline model with endogenous separations to study the response of the firing margin to technology improvements. Notice that the trade-off between employment l_t and output is still present when λ_t is sufficiently high, which is evident from (12). However, now the firm has at its disposal an additional margin of response: firing threshold \tilde{a}_t . Recall that the first-order condition with respect to the firing threshold is

$$f(x_t) \cdot A_t \cdot F'(n_t) \cdot \tilde{a}_t + \lambda_t = (1 - s_t) \cdot w_t. \quad (20)$$

It is reasonable to assume that the severance pay rate is non-negative: $s_t \geq 0$. Then $\lambda_t > w_t$ implies that the firing threshold \tilde{a}_t is negative. It means that the firm is retaining some of its unproductive workers. The firm is willing to keep on the payroll some employees whose current contribution to sales is negative, because the continuation value of employment is sufficiently high. Recall that the individual match quality is identically distributed across workers and independent over time, which means that all employees have a common continuation value to the firm.

Consider a positive technology shock which lifts $f(x_t) \cdot A_t$. As a result, the negative impact of unproductive employees on sales is magnified. Supposing that the continuation value λ_t does not respond strongly, which is likely if the productivity shock is not very persistent, then the firing threshold will need to increase and firing will intensify. Thus, the firm will lay off some of its most unproductive employees and productive labor n_t will increase. By analyzing the first-order condition for the optimal vacancies (12), one

can show that this response of \tilde{a}_t implies that vacancies need to adjust less. Intuitively, if an increase of the productive labor is already achieved with separations, there is less need to adjust vacancies.

4.2.2. Response to a demand shock

A positive demand shock χ_t lifts the marginal utility from current consumption, which induces the household to embark on more shopping visits. That in turn raises the capacity utilization $f(x_t)$. Hence, the demand shock affects the firm through the rise of term $f(x_t) \cdot A_t$, similarly as the productivity shock. Why then is the firm's response to the demand shock so different than to the technology shock?


We find that the persistence of the shock is key. Our analysis of the productivity shock above assumed that the continuation value of employment λ_t did not respond much to a shock, which is true if the shock is not very persistent. The estimated persistence of the demand shock is substantially higher than that of the productivity shock. Consequently, we can expect λ_t to respond more strongly to the demand shock.

Suppose that the continuation value of employment λ_t increases strongly after the persistent demand shock. The inspection of the first-order conditions with respect to the firing threshold and vacancies reveals that the responses of these variables will be reversed in comparison to the case with low persistence. An increase of the continuation value of employment may dominate the contemporaneous impact on the marginal return to productive labor and the firm will choose to sacrifice current productive labor and output for higher total employment. That can be achieved by lowering the firing threshold and maintaining more of unproductive workers as well as issuing new vacancies.

We verify this reasoning in Figure 2, where we plot the impulse responses from the counterfactual economy in which the persistence of the demand shock is lower than estimated and equal to the estimated persistence of the technology shock. In line with the mechanism above, we find that the responses of unemployment, vacancies and separations to the two shocks are identical in sign and very similar in shape. The remaining differences stem mainly from different standard deviations of the shock innovations.

Bai et al. (2019) show that the demand shocks in economies with frictional goods markets are very similar to the technology shocks. Our model also has this feature: if the demand shock had lower persistence, the impulse responses of the labor market variables would qualitatively mimic the responses to the technology shock. Hence, we find that higher persistence of the demand shock is crucial to account for a different behavior of the economy following the two shocks.

Figure 2: Impulse responses with lower persistence of the demand shock



Note: Impulse responses for a positive, one standard deviation shock of each type. The parameters are set as in the baseline calibrated model with an exception of ρ_χ which is set to 0.5 and equal to the calibrated value of ρ_A .

4.3. Variance decomposition

Table 3 (column 1) shows the shares of the variance of key variables which are explained by the demand shock in our baseline framework. Since the shocks are correlated, we report a range of possible shares rather than a single number. The lowest (highest) value of the range correspond to a share obtained under the assumption that the covariance between shocks is attributed entirely to the technology shock (demand shock).⁶ We find that for most variables the demand shock is quantitatively more important than the technology shock. In particular, at least 74% of variance of output and at least 58% of variance of unemployment is demand-driven. The figures for other variables are similar.

There is an intuitive reason why the demand shock must be dominant in explaining the business cycle fluctuations. In the data unemployment is strongly countercyclical, with the correlation with output of -0.84 , and capacity utilization is strongly procyclical, with the correlation with output of 0.56 . A technology shock cannot account for these comovements, since it leads to a simultaneous increase of output, an increase of unemployment and a decrease of capacity utilization. The demand shock, instead, accounts well for these empirical correlations.

4.4. Comparison to the competitive equilibrium

Let's compare the baseline model with a counterfactual economy in which all structural parameters are identical, but prices, wages and severance pay rates are flexible and

⁶In other words, the lower and the upper bounds of a share correspond to the orthogonalization of the shocks based on the Cholesky decomposition with a different assumed order of shocks.

Table 3: Share of variance explained by the demand shock: lower bound & (upper bound)

	baseline model	competitive equilibrium	constant separations
output	74.3 (92.9)	18.3 (24)	76 (91.8)
unemployment	58.5 (98.1)	81.5 (96.7)	66.6 (100)
labor productivity	47.1 (85.7)	66.1 (72.4)	77.3 (87.7)
capacity utilization	55.4 (96.2)	—	50 (94.9)
labor market tightness	74.9 (99.6)	81.3 (95.6)	66.4 (100)


Note: The table shows the lower bound (not in brackets) and the upper bound (in brackets) of the % share of a given variable which is explained by the demand shock. ‘Baseline model’ refers to the model with endogenous separations and the fixprice equilibrium, ‘competitive equilibrium’ refers to the model with endogenous separations and the competitive equilibrium and ‘constant separations’ refers to the model with constant separations and the fixprice equilibrium.

determined in the competitive equilibrium, as described in Proposition 1. The variance decomposition of key variables is presented in the second column of Table 3, while the model-implied correlations between variables and the impulse responses are relegated to Appendix C.

The main difference between the competitive equilibrium and the fixprice equilibrium, which we have considered so far, is in how the goods and the labor markets clear after being disturbed by a shock. In the fixprice equilibrium these markets clear entirely by adjusting tightnesses. In the competitive equilibrium, instead, the markets clear predominantly by adjusting prices and wages. In particular, in the competitive equilibrium capacity utilization $f(x_t)$ is constant and labor market tightness θ_t is almost constant. Thus, the competitive equilibrium has no chance of fitting the data, since empirically both of these objects exhibit a substantial volatility. This point is clearly visible in Figure 3, where we plot the historical evolutions of the capacity utilization and the labor market tightnesses, as well as their predicted behavior according to the baseline model and the competitive equilibrium model. It is an evidence that price and real wage rigidities are important in explaining the data, a point raised by Shimer (2005) and Hall (2005).

We find that in the competitive equilibrium output is much less volatile than in the baseline model. In addition, as shown in the second column of Table 3, output is also much less demand-driven than in the baseline model. It is a consequence of the fact that the capacity utilization is fixed and, hence, output varies only due to changes in the aggregate capacity. Surprisingly, the situation is much different with unemployment, which is neither less volatile nor less demand-driven than in the fixprice equilibrium. That’s in contrast with the static framework of Michailat and Saez (2015), where in the competitive equilibrium the demand shock had no impact on any real variable. As it turns out, the dynamic structure of the model implies that the demand shock retains an

Figure 3: Empirical and simulated measures of tightnesses



Note: This figure present the empirical series of capacity utilization and labor market tightness, as well as their model counterparts obtained when the estimated sequence of exogenous shocks from the baseline model is fed in. 'Baseline model' refers to the model with fixprice equilibrium and endogenous separations and 'Competitive equilibrium' refers to the model with competitive equilibrium and endogenous separations.

influence on unemployment and other real variables even without any rigidities.

When the economy is hit by a positive demand shock, the household marginal utility from consumption goes up. Consequently, the household is willing to substitute future consumption for current consumption. The only way in which that can be achieved in the model is by increasing the firing threshold or reducing vacancy creation, both of which boost current productive labor and output, but reduce employment and future output. It is via this mechanism the demand shock induces unemployment fluctuations in the efficient competitive equilibrium. More generally, the demand shocks will have an impact on real allocation in the dynamic models where the household has preferences over a sequence of consumption levels and there is a trade-off between consumption levels in different periods. This observation stresses that we need to use dynamic models to study demand-driven fluctuations.

The response of the competitive equilibrium to a technology shock is qualitatively similar to its impact on the fixprice equilibrium. Since the unemployment in the competitive equilibrium is responding to both the technology and demand shocks, the strength of each shock will determine the variance decomposition of unemployment. According to the estimation of the baseline model, the demand shock has a standard deviation higher by 80% than the technology shock and a much higher persistence. Consequently, the demand shock is still dominant in driving the unemployment fluctuations even when prices and wages are flexible.

4.5. Comparison to the model with a constant separation rate

We also compare the baseline model to a fixprice equilibrium where the separation rate is fully exogenous and constant. We estimate the structural parameters of this alternative economy in a way analogous to the estimation of the baseline model. The description of the estimation as well as more detailed results are provided in Appendix D.

We find that the data favors the baseline model with an endogenous separation rate over the model where the separation rate is constant. The value of log-likelihood associated with the two series of observable variables is 1707 in the baseline model and it is significantly lower, 1477, in the model with a constant separation rate. Furthermore, the model with constant separations does not replicate the empirical volatilities nearly as well as the baseline model. Notably, the volatility of unemployment is much lower relative to other variables. For this reason, although the model with constant separations matches the volatility of unemployment equally well as the baseline model, the standard deviations of other variables are overshoot by an order of magnitude.

The mismatched volatilities imply that the model has a harder time matching the empirical correlations. Consider the correlation of labour productivity and output. In the model with constant separations this correlation is virtually 1: labor productivity is driven almost exclusively by changes in output. That is because output in this model is 5.6 more volatile in terms of standard deviation than unemployment, relative to 2.4 in the data. In contrast, the model with endogenous separations generates much more unemployment fluctuations relative to output, actually somewhat overshooting the data in this respect, such that the output is only 1.33 more volatile than unemployment. Higher relative volatility of unemployment allows the baseline model to break the perfect correlation between output and labor productivity and match the empirical correlation between these variables, equal to 0.63, exactly. These findings are line with research showing that a variable separation rate is important to explain empirical unemployment fluctuations (Elsby et al. 2009; Fujita and Ramey 2012).

The last column of Table 3 shows the share of variance of key variables which is explained by the demand shock. We find that the model with a constant separation rate implies a split of variance of the key variables which is very similar to that from the baseline model. Specifically, the demand shock is responsible for at least 76% of variance of output and 67% of variance of unemployment.

5. Conclusions

In this paper we extend the static framework of Michaillat and Saez (2015) into a dynamic, stochastic business cycle model with frictional goods and labor markets and an endogenous job separation rate. We estimate the model using the US data and find that the demand shock accounts for most (at least 58%) of the variance of unemployment in

the US. Hence, we quantitatively support the conclusions of Michaillat and Saez (2015) regarding unemployment fluctuations being mostly demand-driven. More broadly, our paper contributes to the new strand of literature incorporating shopping frictions into business cycle models.

Our analysis can be extended in a number of directions. One of them is the design of labor market institutions in an economy where demand has a productive role. Our theoretical results on the severance pay in the efficient allocation can shed light on the optimal severance pay regulation. In related works, Landais et al. (2018b,a) examine how unemployment benefits should vary over the business cycle. Second, the main force behind unemployment fluctuations in our framework is the exogenous shock to the marginal utility from consumption. A natural extension would be to provide explicit microfoundations for this shock, for instance with shocks to wealth or to financial intermediation as in Huo and Rios-Rull (2013).

Appendices

A. Derivation of the competitive equilibrium

The competitive equilibrium consists of the efficient allocation and prices (p_t, w_t, s_t) which decentralize the efficient allocation. In this section we first characterise the efficient allocation and then prove Proposition 1 which specifies the corresponding prices.

The efficient allocation solves the following planner's problem, taking as given the allocation of real money balances $\{\mu_t\}_{t=0}^{\infty}$ and initial labor supply l_{-1} :

$$\max_{\{c_t, v_t, l_t, \hat{v}_t, \bar{a}_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t, \mu_t, \chi_t) \quad (21)$$

subject to the following constraints holding in every period $t = 0, 1, \dots$

$$c_t = M(k_t, v_t) - \rho \cdot v_t, \quad (22)$$

$$l_t = (1 - H(\bar{a}_t)) \cdot ((1 - \sigma) \cdot l_{t-1} + \hat{q}(\theta_t) \cdot \hat{v}_t), \quad (23)$$

where the definitions of other variables $(k_t, n_t, \bar{a}_t, \dots)$ follow the main text. Plug in the first constraint into the objective function and denote the Lagrange multiplier on the second constraint by λ_t^e .

The first-order condition with respect to the shopping visits v_t is $M_v(k_t, v_t) = \rho$. We can rewrite the left-hand side as $q(x_t) + q'(x_t) \cdot x_t$, which is a function of x_t alone. Thus, the efficient product market tightness is constant.

Denote $\phi_t = u_c(c_t, \chi_t) \cdot M_k(k_t, v_t)$. The first-order condition with respect to vacancies is

$$\phi_t \cdot A_t \cdot F'(n_t) \cdot \left(\bar{a}_t - \frac{\hat{\rho}}{\hat{q}(\theta_t) + \theta_t \cdot \hat{q}'(\theta_t)} \cdot \frac{1}{1 - H(\bar{a}_t)} \right) + \lambda_t^e = 0 \quad (24)$$

The first-order condition with respect to \tilde{a}_t is

$$\phi_t \cdot A_t \cdot F'(n_t) \cdot \tilde{a}_t + \lambda_t^e = 0. \quad (25)$$

Finally, the Lagrange multiplier is equal to

$$\lambda_t^e = (1 - \sigma) \cdot \beta \cdot \mathbb{E}_t \left[\phi_{t+1} \cdot A_{t+1} \cdot F'(n_{t+1}) \cdot \hat{\rho} \cdot \frac{1 + \theta_{t+1}^2 \cdot \hat{q}'(\theta_{t+1})}{\hat{q}(\theta_{t+1}) + \theta_{t+1} \cdot \hat{q}'(\theta_{t+1})} \right]. \quad (26)$$

Proof of Proposition 1. Plug the prices into the first-order condition of the household (7) to verify that it is satisfied at the efficient consumption level. Thus, the household shopping behavior is efficient.

Take the equilibrium condition for the firing threshold (13) and subtract from it the efficiency condition for the firing threshold (25). Rearrange to get

$$\frac{(1 - s_t) \cdot w_t}{f(x_t) \cdot A_t \cdot F'(n_t)} = \frac{\lambda_t}{f(x_t) \cdot A_t \cdot F'(n_t)} - \frac{\lambda_t^e}{\phi_t \cdot f(x_t) \cdot A_t \cdot F'(n_t)} = \hat{\rho} \cdot \zeta_t, \quad (27)$$

where the second equality follows from (14) and (26). Consequently, we have

$$s_t = 1 - \frac{\zeta_t \cdot \hat{\rho} \cdot f(x_t) \cdot A_t \cdot F'(n_t)}{w_t}. \quad (28)$$

Take the equilibrium condition for the vacancy creation (12) and subtract from it the efficiency condition (24). After rearranging we obtain

$$\frac{w_t \cdot (1 - H(\tilde{a}_t) + s_t \cdot H(\tilde{a}_t))}{\hat{\rho} \cdot f(x_t) \cdot A_t \cdot F'(n_t)} = \delta_t + (1 - H(\tilde{a}_t)) \cdot \zeta_t. \quad (29)$$

Plugging in (28) gives us


$$\frac{w_t}{\hat{\rho} \cdot f(x_t) \cdot A_t \cdot F'(n_t)} = \delta_t + \zeta_t \quad (30)$$

and plugging that in turn into (28) yields $s_t = 1 - \frac{\zeta_t}{\delta_t + \zeta_t}$. \square

B. Estimation details

We use Bayesian methods to estimate parameters of shocks processes. We use standard approach as outlined in An and Schorfheide (2007) and Lubik (2009). The posterior distributions of the parameters are obtained by combining priors and the likelihood of

Figure 4: Observables used in estimation



Note: The trend is obtained with the one-sided Hodrick-Prescott filter with a smoothing parameter of 1600.

the observable variables evaluated with the Kalman filter. Sampling from the posterior distribution is performed with the Metropolis-Hastings algorithm involving 300,000 draws in two chains with first 50% of draws dropped. The convergence of the chains is assessed based on Brooks and Gelman (1998) charts, as implemented in Dynare package.

To identify the two shocks in the model we need to choose two observable macro variables. As our focus is on the factors driving the unemployment volatility, we select the unemployment rate as one of the observables. For the second series we choose labor productivity. Shimer (2005) points out that some search and matching models require excessively high volatility of labor productivity to match the volatility of other labor market variables. By selecting this series as an observable, we make sure that this critique does not apply to our model. The original series of observables as well as the trend which we subtract are presented in Figure 4.


In terms of the prior distributions we broadly follow Smets and Wouters (2007). The inertia parameters ρ_A and ρ_χ are beta distributed with mean 0.5. We set the standard deviation at 0.1 to prevent the excessive persistence of shocks. The standard deviations of the innovations follow the inverse-Gamma distribution with mean 0.1 and an infinite standard deviation to guarantee that the priors are relatively loose. The prior for the correlation between the shock innovation is a generalized beta distribution with support $[-1, 1]$, mean 0 and standard deviation 0.1. The results of the estimation are reported in the main text. The inertia parameter estimates are above the prior mean, while the standard deviations estimates are below the prior mean. The graphical analysis (available at the request) suggests that the posterior distributions are tightly concentrated around their modes which suggest that the shocks were strongly identified.

Table 4: Correlations and volatilities of key variables: competitive equilibrium & (data)

	Correlation				Standard deviation
	y_t	u_t	$\frac{y_t}{l_t}$	θ_t	
y_t	1	-0.09 (-0.84)	0.47 (0.63)	-0.09 (0.88)	0.018 (0.019)
u_t	.	1	0.84 (-0.14)	-0.96 (-0.91)	0.029 (0.008)
$\frac{y_t}{l_t}$.	.	1	-0.91 (0.33)	0.035 (0.012)
θ_t	.	.	.	1	0.003 (0.123)

Note: Theoretical model moments from the competitive equilibrium (not in brackets) and empirical moments (in brackets). Empirical moments are computed over time period 1967–2019 with an exception of moments of θ_t which, due to data limitations, are computed over time period 1967–2016.

Figure 5: Impulse responses in the competitive equilibrium



Note: Impulse responses for a positive, one standard deviation shock of each type.

C. Details of the model with competitive equilibrium

Table 4 presents the comparison of the model correlations and volatilities of the key variables with their empirical counterparts. Figure 5 present impulse responses.

D. Details of the model with constant separations

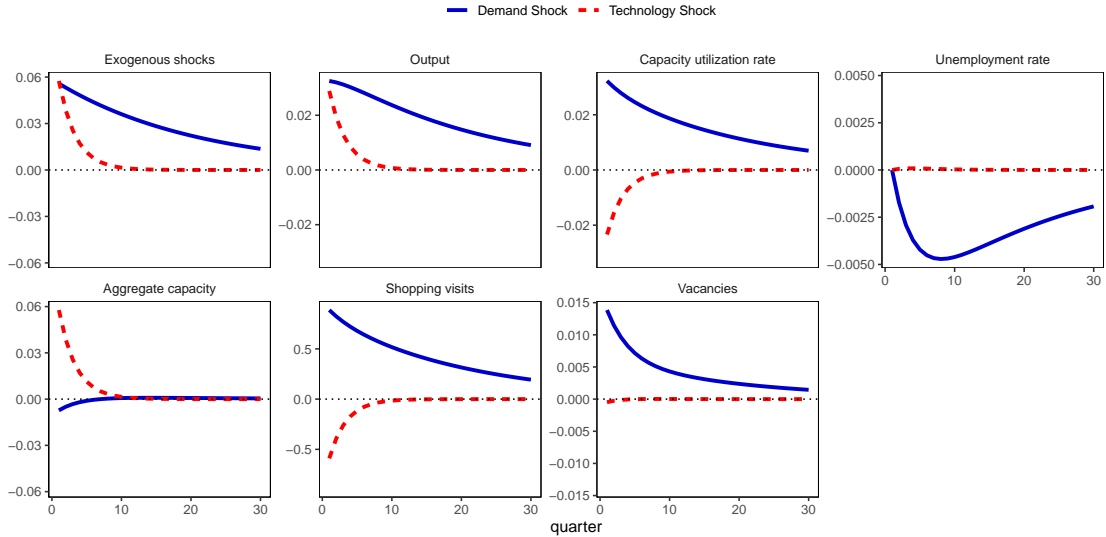
Our parametrization strategy follows closely that of the baseline model, described in Section 3. Regarding the steady state parameters, the only difference is that the exogenous separations rate σ accounts for the whole average rate of separations and is equal to 1.62%, while $H(\tilde{a}) = 0$ and $\bar{a} = 1$. The estimation method of the cyclical parameters is identical as in the baseline model and uses identical priors. The parameter estimates are

Table 5: Correlations and volatilities of key variables: constant separations & (data)

	Correlation					Standard deviation
	y_t	u_t	$\frac{y_t}{l_t}$	$f(x_t)$	θ_t	
y_t	1	-0.93 (-0.84)	1 (0.63)	0.8 (0.56)	0.93 (0.88)	0.111 (0.019)
u_t	.	1	-0.91 (-0.14)	-0.8 (-0.32)	-0.90 (-0.91)	0.020 (0.008)
$\frac{y_t}{l_t}$.	.	1	0.78 (0.61)	0.91 (0.33)	0.091 (0.012)
$f(x_t)$.	.	.	1	0.96 (0.51)	0.140 (0.02)
θ_t	1	0.747 (0.123)

Note: Theoretical model moments from the fixprice equilibrium of the model with a constant separation rate (not in brackets) and empirical moments (in brackets). Empirical moments are computed over time period 1967–2019 with an exception of moments of θ_t which, due to data limitations, are computed over time period 1967–2016.

Figure 6: Impulse responses in the model with a constant separation rate



Note: Impulse responses for a positive, one standard deviation shock of each type.

$\rho_A = 0.67$, $\sigma_A = 0.058$, $\rho_\chi = 0.95$, $\sigma_\chi = 0.056$ and the implied correlation between shock innovations is -0.57 . Both shocks are substantially more volatile than in the baseline model. The persistence parameters as well as the correlation of the innovations are in the ballpark of the baseline estimates.

Table 4 presents the comparison of the correlations and volatilities of the key variables implied by the model with constant separations with their empirical counterparts, while Figure 6 depicts impulse responses.

E. On multiplicity of fixprice equilibria when $\epsilon < 1$

In this section we demonstrate that extending the household preferences to allow for $\epsilon < 1$ can lead to the multiplicity of fixprice equilibria. We consider a more general

utility function over consumption c and real money balances μ

$$u(c, \mu, \chi) = \frac{\epsilon}{\epsilon - 1} \left(\frac{\chi}{1 + \chi} c^{\frac{\epsilon-1}{\epsilon}} + \frac{1}{1 + \chi} \mu^{\frac{\epsilon-1}{\epsilon}} \right) \quad (31)$$

with $\epsilon > 0$ (and logarithmic utility when $\epsilon = 1$). This utility function is equivalent to the one used in the main text or by Michailat and Saez (2015) when $\epsilon > 1$. In addition, it allows us to consider values ϵ in the interval $(0, 1]$.

In the following two subsections we analyze two static frameworks of Michailat and Saez (2015) with this generalized utility function. Then we discuss the significance of these results for our dynamic model.

E.1. Model with only the goods market

Consider a static model with only the goods market as in Michailat and Saez (2015), Section II. The first-order condition of the household, evaluated at the equilibrium value of real money balances, is

$$\chi \cdot \frac{c^{-\frac{1}{\epsilon}}}{1 + \tau(x)} = \mu^{-\frac{1}{\epsilon}}, \quad (32)$$

where x is the product market tightness, $\tau(x) = \frac{\rho}{q(x) - \rho}$ is the wedge due to the goods market frictions and μ are real money balances. Consumption is determined from the supply side by $c = \frac{f(x) \cdot k}{1 + \tau(x)}$, where $f(x)$ is the capacity utilization rate and k is the exogenous aggregate capacity. Combining both, we obtain the following equilibrium condition:

$$(1 + \tau(x))^{\epsilon-1} \cdot f(x) = \frac{\chi^\epsilon}{k} \cdot \mu. \quad (33)$$

Consider a fixprice equilibrium where the real money balances are fixed and the equilibrium is reached by adjusting the product market tightness x . Note that the equilibrium tightness must belong to the set $(0, \bar{x})$, where $q(\bar{x}) = \rho$, since $f(0) = 0$ and $\lim_{x \uparrow \bar{x}} \tau(x) = +\infty$. Denote the left-hand side of the equilibrium condition as $z(x) = (1 + \tau(x))^{\epsilon-1} \cdot f(x)$.

Proposition 2. *When $\epsilon \geq 1$, then there is a unique tightness which solves (33) for any $\frac{\chi^\epsilon}{k} \mu \in (0, z(\bar{x}))$. When $\epsilon < 1$, then there are at least two values of tightness which solve (33) when $\frac{\chi^\epsilon}{k} \mu \in (0, \max_{z \in (0, \bar{x})} z(x))$.*

Proof. When $\epsilon \geq 1$, $z(x)$ is strictly increasing, which implies that the equilibrium tightness is unique. When $\epsilon < 1$, $z(x)$ is a continuous function which is positive at $x \in (0, \bar{x})$ and zero at the $x = 0$ and $x = \bar{x}$. Consequently, for any value $z(x)$ which is strictly below $\max_{z \in (0, \bar{x})} z(x)$ there must be at least two values of x which attain it. \square

Intuition is as follows. As noted by Michailat and Saez (2015), the equilibrium consumption is first increasing and then decreasing with the tightness x . Consequently, the marginal utility from consumption $u_c(c, \chi)$ is first decreasing and then increasing

with tightness x . However, when $\epsilon \geq 1$, the marginal utility from *consumption spending* $\frac{u_c(c;\chi)}{1+\tau(x)}$ is decreasing for all x since changes in the denominator $1 + \tau(x)$ dominate the changes in the marginal utility from consumption. Intuitively, that's because the curvature of the utility function is limited. Since the marginal utility from consumption spending is strictly decreasing, the fixprice equilibrium is unique. When we allow for stronger curvature of the utility function, i.e. when $\epsilon < 1$, then the marginal utility from consumption spending will be first decreasing and then increasing in tightness, which leads to at least two values of tightness which satisfy the equilibrium condition.

E.2. Static model with the goods and the labor markets

Michaillat and Saez (2015) show that the equilibrium of the static model with the goods and the labor markets is described by the following equations:

$$f(x) = \frac{w}{A\alpha} \cdot \hat{f}(\theta)^{1-\alpha} \cdot (1 + \hat{\tau}(\theta))^\alpha. \quad (\text{A1})$$

$$\hat{f}(\theta) = \frac{\alpha\mu}{w} \cdot \chi^\epsilon \cdot (1 + \tau(x))^{1-\epsilon}, \quad (\text{A2})$$

where $w \geq 0$ is the real wage, θ is the labor market tightness, $\hat{\tau}(\theta) = \frac{\hat{\rho}}{\hat{q}(\theta) - \hat{\rho}}$ is the wedge due to frictions in the labor market and $\alpha \in (0, 1]$ controls the curvature of the aggregate capacity with respect to the labor in production. Consider a fixprice equilibrium with fixed and positive μ and w . (A1) always implies that x is strictly increasing with θ . When $\epsilon \geq 1$, (A2) implies that x is decreasing with θ . Hence, the equilibrium, which lays at the intersection of the two curves, is unique. When $\epsilon < 1$, (A2) implies that x is strictly increasing with θ . It means that, in principle, there may be multiple intersection points and, hence, multiple equilibria. In Figure 7 we depict an example of two equilibria for particular parameter values.


E.3. Dynamic model

Let's focus on the steady state fixprice equilibrium in the model with exogenous and constant separation rate σ . First, suppose that $\sigma = 1$, which means that the labor market operates as in the static model. Thus, the equilibrium in the labor market is described by condition (A1) with technology shock A set to its average value of 1. Regarding the goods market, it is straightforward to show that the condition analogous to (A2) is

$$\hat{f}(\theta) = \left(1 - \frac{\beta}{\pi}\right)^\epsilon \cdot \frac{\alpha\mu}{w} \cdot (1 + \tau(x))^{1-\epsilon}, \quad (34)$$

Note that this condition is simply condition (A2) with the right-hand side rescaled by a constant. Thus, the multiplicity of steady state equilibria in the dynamic model can arise for exactly the same reason as the multiplicity of equilibria in the static model.

Figure 7: Example of multiple fixprice equilibria in the static model with the goods and the labor markets



Note: The assumed parameter values are $\epsilon = 0.5, \rho = \hat{\rho} = 0.1, \gamma = \hat{\gamma} = 1, k = 1, \alpha = 0.66, w = 0.396, \mu = 0.3$.

Furthermore, the equilibrium condition regarding the labor market in the dynamic model is continuous with respect to σ and converges to the condition from the static model when $\sigma \uparrow 1$. Thus, we conjecture that the multiplicity of steady state equilibria can arise as long as σ is sufficiently large.

References

- An, S. and Schorfheide, F. (2007). Bayesian analysis of dsge models. *Econometric Reviews*, 26(2-4):113–172.
- Andolfatto, D. (1996). Business cycles and labor-market search. *American Economic Review*, 86:112–132.
- Autor, D. H., Donohue III, J. J., and Schwab, S. J. (2006). The costs of wrongful-discharge laws. *The Review of Economics and Statistics*, pages 211–231.
- Bai, Y. and Rios-Rull, J.-V. (2015). Demand Shocks and Open Economy Puzzles. *American Economic Review*, 105(5):644–649.
- Bai, Y., Rios-Rull, J.-V., and Storesletten, K. (2019). Demand shocks as productivity shocks. *Federal Reserve Board of Minneapolis*.
- Barnichon, R. (2010). Building a composite help-wanted index. *Economics Letters*, 109(3):175–178.
- Barro, R. J. and Grossman, H. I. (1971). A general disequilibrium model of income and employment. *The American Economic Review*, 61(1):82–93.
- Basu, S., Fernald, J. G., and Kimball, M. S. (2006). Are technology improvements contractionary? *American Economic Review*, 96(5):1418–1448.

- Bils, M. (2016). Deducing markups from stockout behavior. *Research in Economics*, 70(2):320–331.
- Boehm, C. E. and Pandalai-Nayar, N. (2018). Are supply curves convex? implications for state-dependent responses to shocks. Technical report.
- Bond, E. W. and Samuelson, L. (1984). Durable good monopolies with rational expectations and replacement sales. *The Rand Journal of Economics*, pages 336–345.
- Brooks, S. P. and Gelman, A. (1998). General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, 7(4):434–455.
- Canova, F., Lopez-Salido, D., and Michelacci, C. (2010). The effects of technology shocks on hours and output: a robustness analysis. *Journal of Applied Econometrics*, 25(5):755–773.
- Canova, F., Lopez-Salido, D., and Michelacci, C. (2013). The ins and outs of unemployment: an analysis conditional on technology shocks. *The Economic Journal*, 123(569):515–539.
- Cappelli, P. (2019). Your approach to hiring is all wrong. *Harvard Business Review*, 97(3):48–58.
- Christiano, L. J., Trabandt, M., and Walentin, K. (2020). Involuntary unemployment and the business cycle. *Review of Economic Dynamics*.
- Den Haan, W. J. (2013). Inventories and the role of goods-market frictions for business cycles.
- Den Haan, W. J., Rendahl, P., and Riegler, M. (2018). Unemployment (fears) and deflationary spirals. *Journal of the European Economic Association*, 16(5):1281–1349.
- Den Haan, W. J. D., Ramey, G., and Watson, J. (2000). Job destruction and propagation of shocks. *American Economic Review*, 90:482–498.
- Dertouzos, J. N. and Karoly, L. A. (1992). *Labor market responses to employer liability*. Rand Corporation.
- Duras, J. (2016). A search model of unemployment and inventories. Technical report, mimeo.
- Elsby, M. W., Michaels, R., and Solon, G. (2009). The ins and outs of cyclical unemployment. *American Economic Journal: Macroeconomics*, 1(1):84–110.
- Farmer, R. E. (2008). Aggregate demand and supply. *International Journal of Economic Theory*, 4(1):77–93.
- Fujita, S. and Ramey, G. (2012). Exogenous versus endogenous separation. *American Economic Journal: Macroeconomics*, 4(4):68–93.

- Gali, J. (1999). Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations? *American Economic Review*, pages 249–271.
- Galí, J. and Rabanal, P. (2004). Technology shocks and aggregate fluctuations: How well does the real business cycle model fit postwar us data? *NBER macroeconomics annual*, 19:225–288.
- Galí, J., Smets, F., and Wouters, R. (2012). Unemployment in an estimated new keynesian model. *NBER macroeconomics annual*, 26(1):329–360.
- Ghassibe, M. and Zanetti, F. (2020). State dependence of fiscal multipliers: the source of fluctuations matters.
- Hall, R. E. (2005). Employment fluctuations with equilibrium wage stickiness. *American Economic Review*, 95(1):50–65.
- Huo, Z. and Rios-Rull, J.-V. (2013). Paradox of thrift recessions. *NBER Working Paper Series*.
- Kaplan, G. and Menzio, G. (2016). Shopping externalities and self-fulfilling unemployment fluctuations. *Journal of Political Economy*, 124:771–825.
- Krause, M. U. and Lubik, T. A. (2007). The (ir) relevance of real wage rigidity in the new keynesian model with search frictions. *Journal of Monetary Economics*, 54(3):706–727.
- Krusell, P., Mukoyama, T., Rogerson, R., and Şahin, A. (2017). Gross worker flows over the business cycle. *American Economic Review*, 107(11):3447–76.
- Kydland, F. E. and Prescott, E. C. (1982). Time to build and aggregate fluctuations. *Econometrica*, pages 1345–1370.
- Landais, C., Michaillat, P., and Saez, E. (2018a). A macroeconomic approach to optimal unemployment insurance: Applications. *American Economic Journal: Economic Policy*, 10(2):182–216.
- Landais, C., Michaillat, P., and Saez, E. (2018b). A macroeconomic approach to optimal unemployment insurance: Theory. *American Economic Journal: Economic Policy*, 10(2):152–81.
- Lubik, T. (2009). Estimating a search and matching model of the aggregate labor market. *FRB Richmond Economic Quarterly*, 95(2):101–120.
- Merz, M. (1995). Search in the labor market and the real business cycle. *Journal of monetary Economics*, 36(2):269–300.
- Michaillat, P. and Saez, E. (2015). A model of aggregate demand, labor utilization, and unemployment. *Quarterly Journal of Economics*, 3:507–569.
- Michelacci, C. and Lopez-Salido, D. (2007). Technology shocks and job flows. *The Review of Economic Studies*, 74(4):1195–1227.

- Moen, E. (1997). Competitive search equilibrium. *Journal of Political Economy*, 105(2):385–411.
- Ordanini, A. and Silvestri, G. (2008). Recruitment and selection services: Efficiency and competitive reasons in the outsourcing of hr practices. *The International Journal of Human Resource Management*, 19(2):372–391.
- Petrongolo, B. and Pissarides, C. A. (2001). Looking into the black box: A survey of the matching function. *Journal of Economic literature*, 39(2):390–431.
- Petrosky-Nadeau, N. and Wasmer, E. (2015). Macroeconomic dynamics in a model of goods, labor, and credit market frictions. *Journal of Monetary Economics*, 72:97–113.
- Petrosky-Nadeau, N., Wasmer, E., and Zeng, S. (2016). Shopping time. *Economic Letters*, 143:52–60.
- Postel-Vinay, F. and Turon, H. (2014). The impact of firing restrictions on labour market equilibrium in the presence of on-the-job search. *The Economic Journal*, 124(575):31–61.
- Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, pages 25–49.
- Shimer, R. (2010). *Labor Markets and Business Cycles*. Princeton University Press.
- Smets, F. and Wouters, R. (2003). An estimated dynamic stochastic general equilibrium model of the euro area. *Journal of the European Economic Sssociation*, 1(5):1123–1175.
- Smets, F. and Wouters, R. (2007). Shocks and frictions in us business cycles: A bayesian dsge approach. *American Economic Review*, 97(3):586–606.
- Trigari, A. (2009). Equilibrium unemployment, job flows, and inflation dynamics. *Journal of Money, Credit and Banking*, 41(1):1–33.
- Walsh, C. E. (2005). Labor market search, sticky prices, and interest rate policies. *Review of Economic Dynamics*, 8(4):829–849.