Bad Jobs and Low Inflation*

Renato Faccini
Queen Mary, University of London
Centre for Macroeconomics (LSE)

Leonardo Melosi
FRB Chicago
European University Institute
CEPR

February 15, 2019

Abstract

Since 2014 the U.S. economy has been characterized by (i) a tight labor market with a record-low unemployment rate and very high job finding rates, (ii) disappointing labor productivity growth, and (iii) low inflation. We propose a model with the job ladder that can reconcile these three facts. In the model inflation picks up only when most jobs are concentrated at the high rung of the ladder: as firms compete for efficiently allocated employed workers, outside offers are declined and matched, triggering an increase in production costs that is not backed by an increase in productivity. The model is estimated using unemployment and quit rates, which allow the model to precisely identify the distribution of the quality of jobs. After the Great Recession, the observed structural drop in the job-to-job rate has slowed down the pace at which the U.S. labor market turns bad jobs into good jobs. As a result, inflation has not escalated even though the labor market appears to be very tight. Furthermore, the model predicts that labor productivity persistently fell by up to 70 bps in the post-Great Recession recovery owing to this protracted misallocation in the labor market.

Keywords: Job Ladder, Labor Productivity, Misallocation, Phillips curve

JEL codes: E31, E24, C78

*Correspondence to: r.faccini@qmul.ac.uk and lmelosi@frbchi.org. We thank Giuseppe Moscarini and Fabien Postel-Vinay for their precious comments and suggestions. The views in this paper are solely those of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of Chicago or any other person associated with the Federal Reserve System.

1 Introduction

In the span of just six quarters between 2007 and 2009, the U.S. economy lost 8.1 million jobs in the worst economic contraction since the Great Depression. The fraction of workers employed part time, but who also wanted a full time job and could not find it, doubled during that period, which is suggestive of the massive disruptive effects of the Great Recession on the allocation of labor. While the unemployment rate has improved relatively quickly during the subsequent recovery, there are facts suggesting that the labor misallocation inherited from the Great Recession might have been more long-lasting. For instance, Figure 1 shows that the job-to-job flow rate and the quit rate have remained persistently below their pre-recession level during the recovery. The lowering of these flows may have hindered the reallocation of labor by slowing the transition of workers from bad jobs to good jobs. Indeed, the rate of part-time jobs for economic reasons to employment reaches its pre-recession level only after nine years of economic recovery. Furthermore, the U.S. labor productivity has been growing at a considerably slower pace relative to its pre-Great Recession average growth rate, which may reflect, at least to some extent, enduring misallocation in the labor market. Another peculiarity of the current expansion is that even though the unemployment rate reached its 50-years low in 2018, price and wage dynamics have remained tepid.

We develop a dynamic general equilibrium model with a "job ladder" that can reconcile all these empirical facts; that is, low job to job transitions, low productivity growth, tight labor market conditions, and low inflation. In the model workers search on the job to climb up the job ladder. The quality of jobs is match specific and can be either good or bad. We assume that firms Bertrand-compete for employed workers, following the sequential auction protocol of Postel-Vinay and Robin (2002). In this setup, wages are insulated from unemployment and are renegotiated upwards only when firms attempt to poach workers from incumbents. Inflationary pressures stemming from this wage protocol are tied to the distribution of jobs on the ladder. In recession, workers are shaken off the ladder and the average match quality worsens as workers start climbing up the ladder again. At this stage, wage rises are mostly accompanied by changes in jobs that bring about increases in productivity. As a result, wage pressures are not inflationary. When instead most workers have already climbed up the job ladder, the degree of misallocation is low and the average job quality is high. Wage rises obtained from poaching firms do not come along with any increase in productivity. As a result, the real marginal cost rises, and wage renegotiations are inflationary.

To quantitatively assess the role of labor misallocation on the recent inflation dynamics, we calibrate the model and we filter it using only labor market variables: the civilian unemployment rate and the quit rate from employment from the Bureau of Labor Statistics (BLS) database "JOLTS". These series are observed in the U.S. from January 1996 through September 2018.

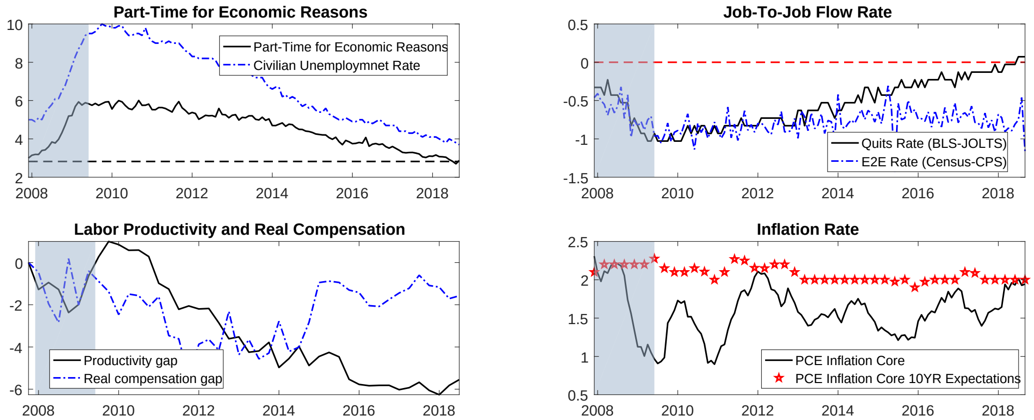

Figure 1: Stylized facts about the labor market and inflation during the current cycle (December 2008-September 2018). The upper left graph: part time workers for economic reasons as a fraction of labor force participation (black line), the average of this ratio computed over the period February 1996 through November 2007 (black dashed-line) and the civilian unemployment rate (dashed-dot blue line). Frequency: Monthly. Source: Bureau of Labor Statistics (BLS). The upper right graph: deviation of the quits rate and employment-to employment (E2E) flow rate from their pre-Great Recession trend over the period from February 1996 through November 2007 (black line and dashed-dot blue line, respectively). Frequency: Monthly. Frequency: Monthly. Sources: BLS (JOLTS dataset) and the Census Bureau (Current Population Survey), respectively. The lower left graph: labor productivity gap and real hourly compensation gap computed in deviation from their ten-year moving-average trend and normalized to be zero in 2017Q4 (black line and dashed-dot blue line, respectively). Frequency: Quarterly. Sources: BLS. The lower right graph: core PCE inflation (annualized in percentage, black line). Frequency: Monthly. Source: BEA. The red stars denote the 10-year ahead expectations about the PCE inflation rate. Frequency: Quarterly. Source: The Federal Reserve Bank of Philadelphia - the Survey of Professional Forecasters.

The model can explain why inflation has been so low while the unemployment rate and the job finding rate suggest that the labor market is extremely tight. We show that the persistent fall in the quit rate has critically contributed to dampening the rise in inflation. The drop in this rate considerably slowed down the creation of good jobs during the post-Great Recession recovery, leading to a persistent surge in bad matches, which is remindful of the increase in part-time workers for economic reasons. Replacing the quit rate with the employment-to-employment rate from the U.S. Census leads to even more striking results because this rate does not show a quick reversion back to its pre-Great Recession level as the quit rate does. Hence, the job-to-job flow rate remains persistently below its long-run value in the model impairing the working of the job ladder until late 2018. As a consequence the model predicts an even weaker pickup of inflation during the post-Great Recession recovery.

While the employment-to-employment (E2E) flow rate is a good predictor for inflation in sample in line with findings in Faberman and Justiniano (2015), the best predictor in population is the fraction of good jobs. This property stems from the job ladder. The only way for firms to hire workers away from a good match is to offer a higher wage. However, this salary increase is not accompanied by any increase in productivity, if the worker is already in a good match. Therefore, when the large majority of workers is efficiently matched, rational firms expect future marginal costs to increase, which leads to a rise in inflation. An important property of the model is that the unemployment rate and the job-to-job rate are sufficient statistics for the state of the job ladder (i.e., the fraction of bad jobs and good jobs) and labor productivity. Thus, our model allows us to measure the series of good matches, which is not available in the data. We also show that the disruptive effects of the recession on the job ladder helps explaining to some extent the slowdown in the U.S. labor productivity growth, which we observe during the recovery. See the lower left graph of Figure 1.

The standard New Keynesian models fail to explain the recent U.S. inflation dynamics because they do not take into account the extensive and persistent labor market misallocation generated by the Great Recession. In the paper we review the traditional theories of the Phillips curve and show that they cannot explain why inflation has remained so moderate for so long while labor market conditions have been so tight.

This paper builds on a series of important contributions that have been made by Moscarini and Postel-Vinay in recent years. Moscarini and Postel-Vinay (2016, 2017) provide empirical evidence that any correlation between real wage growth and unemployment, the core of the Phillips curve, is purely spurious. When unemployment is low, the job finding rate is high and employed workers also benefit from more job opportunities; as employers compete for employed workers, wage will rise, either because workers change job, moving to better paid employers, or because offers are declined and matched. Wage growth seems to be more closely associated with job-to-job reallocations than with the dynamics of flows from unemployment to employment, suggesting that monetary authorities concerned with the sources of wage inflation should pay more attention to job-to-job flows, rather than focusing on unemployment rates as a measure of slack.

Moscarini and Postel-Vinay (2018a) introduce on-the-job search and heterogenous jobs in a New-Keynesian framework to investigate the transmission mechanism of monetary and technology shocks in a New Keynesian model with a job ladder. Unlike that paper, our emphasis is on the empirical evaluation of this class of general equilibrium models. As in Moscarini and Postel-Vinay (2018a), we retain the key assumptions that prices are sticky, that workers search on the job, and that wages are negotiated following the bargaining protocol of Postel-Vinay and Robin (2002). On the other hand, our model also differs in a number of ways that make our model amenable to empirical analysis. For instance, we introduce time-varying on-the-job search intensity as a means of accounting for the joint dynamics of the UE rate and the EE rate, playing a key role in the analysis. We also assume a two steps ladder with only good and bad matches, rather than a continuum of job types, and a Kimball aggregator of consumer preferences, as a way of reconciling marginal costs dynamics with an empirically reasonable degree of price rigidity. Finally, the model relies on estimated shocks to preferences and on-the-job search intensity to explain the data.

Our contribution relative to the works by Moscarini and Postel-Vinay is to embed the job ladder into a simple New Keynesian model and evaluate its properties using state-of-the-art

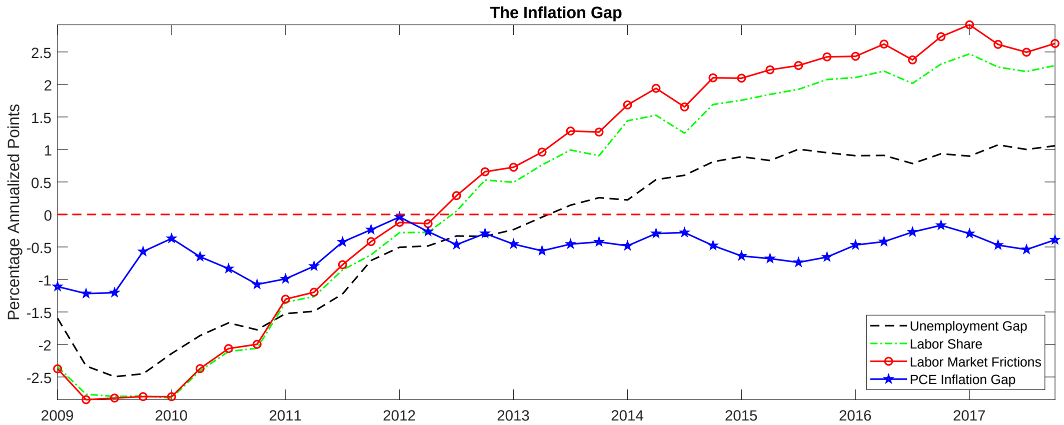

Figure 2: Shows the inflation dynamics from 2009Q1 through 2017Q3 using three traditional theories of inflation and the core PCE inflation gap.

time series techniques. We believe that this is the first paper that carries out this empirical analysis. This analysis highlights two main findings. First, the dynamics of inflation implied by the theory of the job ladder can successfully account for the missing inflation in the post-Great Recession recovery. Second, this theory partially explains the slowdown in the U.S. labor productivity observed over the last decade.

Key to these two findings is the disruption of the ladder-based mechanism of labor reallocation that has characterized the Great Recession and the ensuing recovery. Moscarini and Postel-Vinay (2016a) provide cross-sectional evidence in this direction. Specifically, they show that job-to-job quit rates declined sharply, not only in the aggregate, but especially from smaller, less productive employers, which are considered as a proxy for jobs at the bottom rungs of the ladder. Because these are always the main source of job reallocation, Moscarini and Postel-Vinay (2016a) conclude that workers almost stopped climbing the ladder during the Great recession and the recovery was almost absent.

The paper is organized as follows. Section 2 presents a formal analysis of the missing inflation puzzle. Section 3 presents the general equilibrium model with the job ladder, and Section 4 the empirical analysis. Section 5 contains our concluding remarks.

2 The Existing Theories of Inflation

The New Keynesian model is the most popular workhorse to study inflation dynamics. In this set-up, inflation dynamics hinge upon the New Keynesian Phillips curve, which in its simpler format reads as follows

$$\pi_t = \kappa\varphi_t + E\pi_{t+1}, \quad (1)$$

where φ_t denotes the real marginal cost, and π_t is inflation. In its empirical applications, φ_t is typically proxied by alternative theory based measures. We consider marginal costs time series related to the following three traditional theories of the Phillips curve: (1) Old-fashioned theories, recently revived by Galí, Smets, and Wouters (2011), which link inflation to the current and expected unemployment gap; (2) the standard NK theory, derived from models with no labor frictions as in Galí (2008), suggesting that the labor share alone is the key determinant of the inflation rate; (3) a variant of the standard NK theory, based on models that account for search and matching frictions, which explains inflation using current and expected measures of the labor share, the job finding rate, and the real wage (Krause, Lopez-Salido, and Lubik, 2008). While there are more sophisticated versions of the New Keynesian Phillips curve, which, for instance, feature price indexation, we focus here on the simpler version of this curve. We will discuss the extension to the case of price indexation in the appendix of the paper and show that assuming price indexation does not change the main result.

Solving eq.(1) forward, expected inflation can be expressed as the sum of future expected real marginal costs. To estimate measures of inflation expectations, we estimate a Bayesian Vector autoregressions (VAR) model to predict different measures of real marginal costs and launch forecasts of marginal costs starting in 2017Q4. We construct a gap measure for nine quarterly macro observables by using their 8-year past moving average trend. The observables are: the labor share, the job finding rate, real wages, the civilian unemployment rate, real GDP, real consumption, real investment, CPI inflation (log differences in deviations from its 8-year moving average plus three other concepts of inflation gaps based on pce core, cpi core), and the federal funds rate (FFR). The data sample covers the period 1958q4 through 2017Q4. The VAR model is estimated using these nine observables. The estimated VAR model is used to forecast the future expected path of real marginal costs; that is, unemployment gap, the unit labor cost, and a modified measure of unit labor costs augmented to account for labor market frictions depending on which of the three theories we are considering.

Using a VAR model to approximate the private sector's forecasts is desirable for two reasons. First, this approach does not require us to take a stand on what model agents use to form their expectations about future marginal costs. Misspecifications in the type of model would likely affect our results in an undesirable way. Our task is to evaluate how well alternative definitions of marginal costs that correspond to three well-known theories of the Phillips curve explain inflation dynamics over the current business cycle. We do not want our findings to be biased by the choice of the model agents use to form their forecasts. Since VAR models can be regarded as reduced form representations for the data that are less prone to misspecifications than structural models and are very suitable to construct forecasts, VAR-based forecasts help us obtain more robust results. Second, many structural economic models, such as linearized dynamic stochastic general equilibrium (DSGE) models, are approximated by VAR models. So

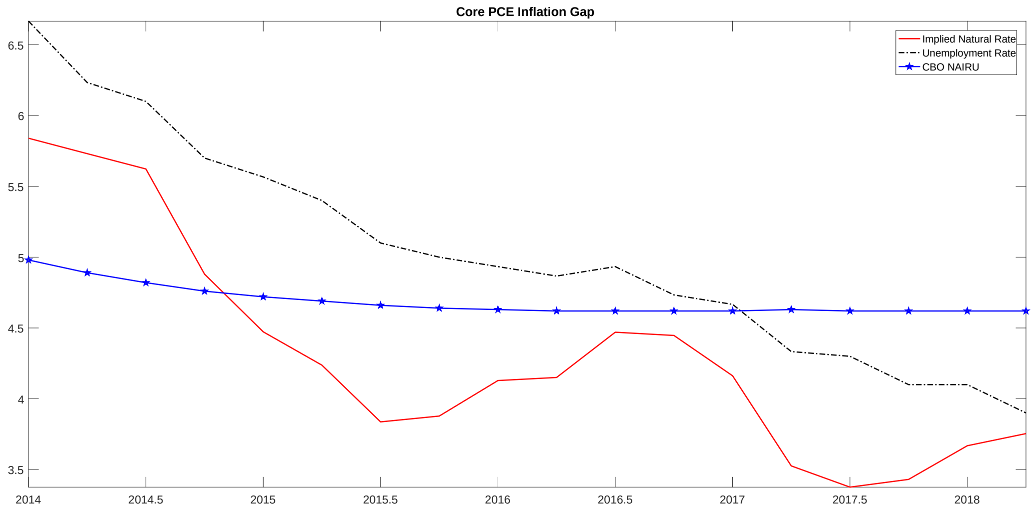

Figure 3: The natural rate of unemployment required to keep the inflation gap consistent with the estimated Phillips curve.

assuming that agents form their expectations according to a VAR model does not necessarily imply a deviation from rationality.

Figure 2 shows that all the traditional series suggest that inflation should be above steady state since 2013 or 2014. So none of these theories is able to account for why inflation has been so low in recent years.

Another way to evaluate whether the traditional New Keynesian models fail to account for the dynamics of inflation in the Post-Great Recession recovery is to estimate the Phillips curve $\pi_t - \beta E_t \pi_{t+1} = -\kappa \left(u_t - u_t^N\right) + \varepsilon_t$ using the OLS estimator over the period ranging from the first quarter of 1996 through 2018q2. We construct a series for the left-hand-side variable by taking the difference between the rate of core PCE inflation and the Survey of Professional Forecasters (SPF) median one-quarter-ahead expectations about PCE core inflation. The unemployment gap, $u_t - u_t^N$, is measured by taking the difference between the civilian unemployment rate computed by the BLS, u_t, and a measure of the natural rate of unemployment, u_t^N.[1] The residual ε_t is assumed to be Gaussian. Then we compute the series of the natural rate of unemployment, u_t^*, that would have allowed the estimated Phillips curve to perfectly fit the data. This counterfactual natural rate of unemployment is hence given by

$$\hat{u}_t^* \equiv \hat{\kappa}^{-1} \left(\pi_t - \beta E_t \pi_{t+1}\right) + u_t, \tag{2}$$

where $\hat{\beta}$ denotes the least squares estimate of the slope of the Phillips curve.

[1]We use the CBO's estimate of the non accelerating inflation rate of unemployment (NAIRU) as a measure of the natural rate. Using the long-term or the short-term NAIRU will not materially affect the results. Prior to 2007Q1 SPF inflation expectations are spliced from SPF expectations about CPI inflation.

The results are shown in Figure 3. The red line marks the natural rate of unemployment u_t^* that would have allowed the estimated New Keynesian Phillips curve to explain PCE core inflation perfectly. This counterfactual natural rate has been almost always below the civilian unemployment rate (dashed-dotted black line) during the post-Great Recession recovery. This finding suggests that in order for the New Keynesian Phillips curve to explain the persistent low inflation, the natural rate of unemployment must be lower than the record-low unemployment rate observed in the period 2014-2018. In addition, the natural rate that explains the recent period of low inflation has been persistently below the non accelerating inflation rate of unemployment (NAIRU) computed by the Congressional Budget Office (CBO). This finding suggests that the magnitude of labor market slack needed for the New Keynesian Phillips curve to explain the recent inflation dynamics has little empirical support.

3 A General Equilibrium Model with the Job Ladder

The limitation of the traditional theories of inflation discussed above motivates the need of an alternative theory, which will be introduced in this section. This theory is nested in the traditional New Keynesian model of inflation and focuses on the role for labor misallocation along the job ladder. We will use this model as a quantitative tool to assess the macroeconomic implications of the misallocation generated by the Great Recession, as suggested by Figure 1.

3.1 The Economy

The economy is populated by a representative, infinitely lived household, whose members' labor market status is either unemployed or employed. All members of the households are assumed to pool their income at the end of each period and thereby fully share their labor income risk. The labor market is frictional and workers search for jobs both whether they are unemployed or employed. While all unemployed workers are also job seekers, it is assumed that only a random fraction s_t of the employed workers look for vacancies in each period and that s_t follows an AR(1) stochastic process.

On the firms' side we distinguish three sectors: Service, Intermediate and Final goods. The service sector comprises an endogenous measure of worker-firm pairs who match in a frictional labor market and produce a homogeneous non-storable good. Productivity $y \in \{y_g, y_b\}$ is match-specific and can be either good or bad, with $y_g > y_b > 0$. We let ξ_g denote the probability that upon matching the productivity draw is good and $\xi_b = 1 - \xi_g$ the probability that the draw is bad.[2] The output of the match is sold to intermediate firms in a competitive market at the real price φ_t, and transformed into a differentiated product. Specifically, one unit of

[2]For a two-steps model with a job ladder see also Gertler, Huckfeldt, and Trigari (2016).

the service is transformed by firm i into one unit of a differentiated intermediate good $y_t(i)$. Intermediate producers set the price of their goods subject to quadratic costs of adjustment á la Rotemberg. Finally, the final good sector packages all different goods into a homogenous consumption product Q_t, which is sold to the households at unit price P_t.

3.2 The Labor Market

The labor market is frictional and governed by a meeting function which brings together vacancies and job seekers. The pool of workers looking for jobs at each period of time t is given by the measure of workers who are unemployed at the beginning of a period, $u_{0,t}$ plus a fraction s_t of the workers who are employed, $1 - u_{0,t}$. Denoting the aggregate mass of vacancies by v_t, we can define labor market tightness as:

$$\theta_t = \frac{v_t}{u_{0,t} + s_t\left(1 - u_{0,t}\right)}.$$

We assume that the meeting function is homotetic, which implies that the rate at which searching workers locate a vacancy, $\phi(\theta) \in [0,1]$, and the rate at which vacancies locate job seekers, $\phi(\theta)/\theta \in [0,1]$, depend exclusively on θ and are such that $d\phi(\theta)/d\theta > 0$ and $d[\phi(\theta)/\theta]/d\theta < 0$.

Because of frictions in the labor market, wages deviate from the competitive solution. It is assumed that wage bargaining follows the sequential auction protocol of Postel-Vinay and Robin (2002). Namely, the outcome of the bargaining is a wage contract, i.e. a sequence of state contingent wages, which promises to pay a given utility payoff in expected present value terms, accounting also for expected utility from future spells of unemployment and wages paid by future employers. The commitment of the worker-firm pair to the contract is limited, in the sense that either party can unilaterally break-up the match if either the present value of firm profits becomes negative, or the present value utility from being employed falls below the value of being employed. The contract can be renegotiated only by mutual consent: if an employed worker meets a vacancy, the current and the prospective employer observe first the productivity associated with both matches, and then engage in Bertrand competition over contracts. The worker chooses the contract that delivers the larger value.

The within-period timing of actions is as follows: all the unemployed workers and a fraction s_t of the employed search for a job at the beginning of the period. Next, some workers move out of the unemployment pool, while successful on-the-job seekers have their wage renegotiated and possibly move up the ladder. Then production takes place and wages are paid. Finally, a fraction δ of the existing matches is destroyed. This timing implies the following dynamics for the aggregate state of unemployment. Denote the stock of end-of-period employed workers as:

$$n_t = 1 - u_t. \tag{3}$$

Aggregate unemployment at the beginning of a period is given by

$$u_{0,t} = u_{t-1} + \delta n_{t-1}, \tag{4}$$

while aggregate unemployment at the end of a period is

$$u_t = u_{0,t} [1 - \phi_t]. \tag{5}$$

3.3 Households

3.3.1 The intertemporal maximization problem

The representative household enjoys utility from the consumption basket C_t and from the fraction of its members who are not working and are therefore free to enjoy leisure. The utility function is linear in consumption and subject to preference shocks, λ_t, which are assumed to follow an AR(1) stochastic process in logs. We have assumed risk neutrality as it reduces problems of indeterminacy with the solution of the model. The resources available to consume at a given point in time t, include bond holdings B_t, profits Π_t and wages from the workers who are employed. We assume that all unemployed workers look for jobs, and restrict attention to equilibria where the value of being employed for any worker is no less than the value of being unemployed. In this set-up, the measure of workers who are employed is not a choice variable of the household, but is driven by aggregate labor market conditions through the job finding probability $\phi(\theta_t)$. Let $e_t(j) \in \{0, 1\}$ be an indicator function which takes the value of one if a worker j is employed at the end of the period, and zero otherwise. The intertemporal maximization problem reads:

$$\max_{\{C_t, B_{t+1}\}} E_0 \sum_{t=0}^{\infty} \beta^t \lambda_t \left[C_t + b \int_0^1 (1 - e_t(j)) \, dj \right],$$

subject to the budget constraint,

$$P_t C_t + \frac{B_{t+1}}{1 + R_t} \leq B_t + \int_0^1 e_t(j) \, w_t(j) + \Pi_t,$$

the stochastic process for the employment status,

$$\begin{aligned} prob\{e_{t+1}(j) = 1 \mid e_t(j)\} &= e_t(j)(1 - \delta) + [1 - e_t(j)] \phi(\theta_{t+1}) \\ prob\{e_{t+1}(j) = 0 \mid e_t(j)\} &= 1 - prob\{e_{t+1}(j) = 1 \mid e_t(j)\}, \end{aligned} \tag{6}$$

and the stochastic process for equilibrium wages $w_t(j)$, to be determined below. The above equation writes that a worker who is registered as unemployed at the end of a period, i.e.

$e_t(j) = 0$, will only have a chance to look for jobs at the beginning of next period, and get one with probability $\phi(\theta_{t+1})$. Moreover, a worker employed at the end of a period, i.e. $e_t(j) = 1$, will also be in employment next period unless separation occurs at the exogenous rate δ. So while in principle separation could occur endogenously if firms profits fall below zero, or if the value of an employed worker falls below the value of being unemployed, eq.(6) implicitly restricts attention to those equilibria where all separations are exogenous.

The first order conditions with respect to B_{t+1} and C_t yield a standard Euler equation:

$$1 + R_t = \beta^{-1} E_t \frac{\lambda_t}{\lambda_{t+1}} \Pi_{t+1}, \tag{7}$$

where Π_t denotes the (gross) inflation rate at time t.

3.3.2 Job values and sequential auctions

Here we characterize the value functions for the states of employment and unemployment. The value of unemployment to a worker j, measured at the end of a period and expressed in utility units reads:

$$\lambda_t V_{u,t}^j = \lambda_t b + \beta E_t \phi(\theta_{t+1}) \lambda_{t+1} \left[V_{e,t+1}^j \mid e_t(j) = 0 \right] + \beta E_t (1 - \phi(\theta_{t+1})) \lambda_{t+1} V_{u,t+1}^j.$$

The value to a worker j of being employed at the end of a period in a job of productivity y_t at wage w_t, after reallocation has taken place reads:

$$\begin{aligned} \lambda_t V_{e,t}^j (w_t(j), y_t(j)) \quad &= \quad \lambda_t \frac{w_t(j)}{P_t} + \beta E_t \lambda_{t+1} \left\{ \delta V_{u,t+1}^j \right. \\ &\quad \left. + (1-\delta) V_{e,t+1}^j (w_{t+1}(j), y_{t+1}(j)) \mid e_t(j) = 1, w_\tau(j), y_\tau(j) \right\}, \end{aligned}$$

where $E_t V_{e,t+1} [w_{t+1}(j), y_{t+1}(j) \mid e_t(j) = 1, w_t(j), y_t(j)]$ indicates the value in units of the numeraire good of being employed at the end of the next period in a match with productivity y_{t+1} at the wage w_{t+1}, conditional on being currently employed in a match with productivity $y_t(j)$ at the promised wage $w_t(j)$. We assume that firms have all the bargaining power, and hence the unemployed workers who take up a new offer are indifferent between being employed or unemployed, i.e. $\lambda_t V_{e,t}(w_t(j), y(j)) = \lambda_t b + \beta E_t \lambda_{t+1} V_{u,t+1}^j$ independently of productivity. It follows that

$$V_{u,t}^j = b + \beta E_t \frac{\lambda_{t+1}}{\lambda_t} V_{u,t+1}^j = V_{u,t}. \tag{8}$$

Let $V_{e,t}^*(y)$ denote the value to the worker of being employed under full extraction of a firm's willingness to pay at the end of time t. In this case a worker of productivity y receives the maximum value that the firm is willing to promise in period t, including the payment of the current period wage. Let $\{w_s^*(y)\}_{s=t}^{\infty}$ denote the state-contingent contract that delivers $V_{e,t}^*(y) \equiv V_{e,t}(w_t^*, y)$. By promising to pay the contract $\{w_s^*(y)\}_{s=t}^{\infty}$, the firm breaks even in expectation, that is, the expected present value of future profits is zero.

Now consider a firm that is currently employing a worker with productivity y under any promised contract $\{w_s(y)\}_{s=t}^{\infty}$. Assume that the worker is poached by a firm with match productivity y'. The outcome of the auction must be one of the following three:

1. $V_{e,t}^*(y') < V_{e,t}(w_t, y)$; in this case the willingness to pay of the poaching firm is less than the value of the contract that the worker is currently receiving. As a result, the incumbent firm retains the worker with the same wage contract with value $V_{e,t}(w_t, y)$.

2. $V_{e,t}(w_t, y) \leq V_{e,t}^*(y') < V_{e.t}^*(y)$; In this case the willingness to pay of the poaching firm is greater or equal to the value of the contract the worker is receiving in his current job, but lower than the willingness to pay of the incumbent firm. The two firms engage in Bertrand competition and as a result, the incumbent firm retains the worker offering the new contract $V_{e,t}^*(y')$.

3. $V_{e,t}^*(y) \leq V_{e,t}^*(y')$; in this case the poaching firm has a willingness to pay that is no less than the incumbent's. As a result, the current match is terminated and the worker is poached at the maximum value of the contract that the incumbent is willing to pay. The continuation value of the contract obtained by the worker is thus $V_{e,t}^*(y)$. It is assumed that if a worker is poached by a firm with equal productivity, job switching takes place with probability υ.

Consider the case of a firm that has promised to pay the contract $\{w_s^*(y)\}_{s=t}^{\infty}$, which implies that the firm breaks even in expectation and is not able to promise higher wage payments in case it enters an auction with a poaching firm. In this case, if no outside offers arrive the worker receives a continuation value of $V_{e,t}^*(y)$ from the incumbent firm. Otherwise the worker is poached and, in accordance with point (3) above, receives a contract from the new firm which is also worth $V_{e,t}^*(y)$. So either way, the worker receives a contract of value $V_{e,t}^*(y)$. The value to a worker of being employed under the contract $\{w_s^*(y)\}_{s=t}^{\infty}$ can therefore be written as:

$$V_{e,t}^*(y) = \varphi_t y + \beta E_t \frac{\lambda_{t+1}}{\lambda_t}\left[\delta V_{u,t} + (1-\delta) V_{e,t+1}^*(y)\right]$$

where $\varphi_t y$ is the marginal revenue product of selling y units of the service to the intermediate producers. Subtracting (8) from the above equation yields:

$$V_{e,t}^*(y) - V_{u,t} = \varphi_t y - b + (1-\delta)\beta E_t \frac{\lambda_{t+1}}{\lambda_t}\left[V_{e,t+1}^*(y) - V_{u,t+1}\right].$$

Notice that the value to the worker of extracting all the rents associated with a type-y match, $V_{e,t}^{*}(y) - V_{u,t}$, is in fact simply the surplus $S_t(y)$. Iterating forward on the above expression, we can define the surplus of a match with productivity y as:

$$S_t(y) = E_t\left[\sum_{\tau=0}^{\infty}(1-\delta)^{\tau}\frac{\lambda_{t+\tau}}{\lambda_t}\left(\varphi_{t+\tau}y - b\right)\right].$$

Notice that the surplus function above is affine increasing in y, which implies that firms with higher productivity win the auction, and therefore workers cannot move to jobs with lower productivity. For convenience, we can rearrange the above expression as

$$S_t(y) = y\mathcal{W}_t - \frac{b}{1-\beta(1-\delta)}, \tag{9}$$

where

$$\mathcal{W}_t = \varphi_t + \beta(1-\delta)E_t\frac{\lambda_{t+1}}{\lambda_t}\mathcal{W}_{t+1}. \tag{10}$$

Seen from the point of view of a Service sector firm, $\mathcal{W}_t$ can be interpreted as the expected present discounted value of the entire stream of current and future real marginal revenues expected by the worker until separation. From the point of view of an intermediate firm, who purchases labor services, $\mathcal{W}_t$ can be interpreted as the expected present discounted value of the cost of purchasing one unit of the labor service by a firm until separation.

3.4 Service sector firms: match values and free entry

In order to advertise a vacant job in the Service sector, firms need to pay an advertising cost c per period. In addition, to form a match, they also have to pay a fixed cost c^f, which can be interpreted as a sunk cost of training. The expected cost of creating a job equals $c^f + \frac{c}{\varpi_t}$, where ϖ_t is the vacancy filling rate and ϖ^{-1} measures the expected number of periods that are required to meet a worker. Free entry in the Service sector implies that entrant firms will make zero profits on average, i.e. expected costs are equal to the expected profits. In turn, expected profits will be a weighted average of the profits expected from meeting with an unemployed worker, and the profits expected from meeting with a worker employed and searching on the job, with weights given by their relative measure in the pool of job seekers. The free entry conditions is:

$$\begin{aligned} c^f + \frac{c}{\varpi_t} &= \frac{u_{0,t}}{u_{0,t} + s_t(1-u_{0,t})}\left[\xi_b S_t(y_b) + \xi_g S_t(y_g)\right] \\ &\quad + \frac{s_t(1-u_{0,t})}{u_{0,t} + s_t(1-u_{0,t})}\left\{\xi_g \frac{l_{b,t}^{0}}{1-u_{0,t}}\left[S_t(y_g) - S_t(y_b)\right]\right\}, \end{aligned}$$

where $l_{b,t}^{0}$ denotes the measure of workers who, at the beginning of a period, are employed in low quality matches. Hence, the expression on the second row in curly brackets implies that a vacant job can extract a strictly positive surplus from a worker who is searching on the job, only if his match quality is good with the poaching firm, and bad with the incumbent. In this case, the firm will offer the worker a wage contract with present value $S_t(y_b)$, keeping the remaining surplus $S_t(y_g) - S_t(y_b)$ for itself.

Using eq.(9) to substitute for the surplus functions and rearranging, we can rewrite the above expression as:

$$\begin{aligned} c^f + \frac{c}{\varpi_t} &= \frac{u_{0,t}}{u_{0,t} + s_t(1-u_{0,t})}\left[\mathcal{W}_t\left(\xi_b y_b + \xi_g y_g\right) - \frac{b}{1-\beta(1-\delta)}\right] \\ &\quad + \frac{s_t}{u_{0,t} + s_t(1-u_{0,t})}\xi_g l_{b,t}^{0}\mathcal{W}_t\left(y_g - y_b\right). \end{aligned} \tag{11}$$

Consider a shock to the above free entry condition, that causes the expected profits to become larger than the expected costs. There are two key forces that help restore the equilibrium, making sure that this condition is always satisfied. First, vacancies increase, lowering the vacancy filling rate and raising the expected costs of entry, as in any standard search and matching model. Second, as the measure of jobs in the service sector rise, aggregate output increases and both the current and the future expected competitive prices of the service falls. Because of sticky prices in the intermediate goods sector, it will take a while for this reduction in the prices of the labor service to pass-through to the aggregate price level P_t. As a result, the current *real* price of the service φ_t falls and so does the entire expected discounted stream of real prices, $\mathcal{W}_t$. A lower $\mathcal{W}_t$ decreases the expected revenue from entering the service sector, contributing to restore the equilibrium in this market.

3.5 The dynamic distribution of match types

We characterize the distribution of match types and its dynamics with the laws of motion for bad and good matches:

$$l_{b,t} = \left[1 - s_t\phi(\theta_t)\xi_g\right] l_{b,t}^{0} + \phi(\theta_t)\xi_b u_{0,t}, \tag{12}$$

$$l_{g,t} = l_{g,t}^{0} + s_t\phi(\theta_t)\xi_g l_{b,t}^{0} + \phi(\theta_t)\xi_g u_{0,t}. \tag{13}$$

In the above equations we let $l_{b,t}$ and $l_{g,t}$ denote the end-of-period measure of bad and good matches respectively. We let $l_{b,t}^{0}$ and $l_{g,t}^{0}$ denote beginning-of-period values, instead. In turn, $l_{b,t}$ is equal to the sum of the bad matches at the beginning of a period who did not move up the ladder by finding a good quality match within the period, $\left[1 - s_t\phi(\theta_t)\xi_g\right] l_{b,t}^{0}$, plus the new hires from the unemployment pool who turned out to draw a low quality match, $\phi(\theta_t)\xi_b u_{0,t}$.

Indeed, job-to-job flows from bad to good quality matches are given by the fraction of badly matched employed workers who search on the job with exogenous probability s_t, meet a vacancy with probability $\phi(\theta_t)$ and draw a good quality match with probability ξ_g. The end-of period measure of good matches is instead given by the beginning of period measure of good matches $l^0_{g,t}$, plus the job-to-job inflows from low quality matches $s_t \phi(\theta_t) \xi_g l^0_{b,t}$, and the new hires from unemployment $\phi(\theta_t) \xi_g u_{0,t}$. Using

$$l^0_{i,t+1}(y) = (1-\delta) l_{i,t}(y) \text{ for } i = \{b, g\},$$

we can rewrite the dynamic equations (12) and (13) to express laws of motion for bad and good jobs at their beginning-of-the period values:

$$l^0_{b,t+1} = (1-\delta)\left\{\left[1 - s_t \phi(\theta_t) \xi_g\right] l^0_{b,t} + \phi(\theta_t) \xi_b u_{0,t}\right\}, \tag{14}$$

$$l^0_{g,t+1} = (1-\delta)\left\{l^0_{g,t} + s_t \phi(\theta_t) \xi_g l^0_{b,t} + \phi(\theta_t) \xi_g u_{0,t}\right\}. \tag{15}$$

3.6 Intermediate and Final Good Producers

Final firms package differentiated intermediate products into a homogeneous good, which is sold in perfect competition. As in Smets and Wouters (2007), the consumption bundle is given by the general Kimball (1995) aggregator:

$$\int_0^1 G(q_t(i)/Q_t)\, di = 1, \tag{16}$$

which nests Dixit-Stiglitz as a special case. Relative to the latter, the Kimball aggregator introduces more strategic complementarity in price setting, which causes firms to adjust prices by less to a given change in marginal costs. As in Dotsey and King (2005), Levin, Lopez-Salido, and Yun (2007) and Lindé and Trabandt (2018), we assume the following strictly concave and increasing function for $G(q_t(i)/Q_t)$:

$$G(q_t(i)/Q_t) = \frac{\omega^k}{1+\varkappa}\left[(1+\varkappa)\frac{q_t(i)}{Q_t} - \varkappa\right]^{\frac{1}{\omega^k}} + 1 - \frac{\omega^k}{1+\varkappa},$$

where $\omega^k = \frac{\chi(1+\kappa)}{1+\varkappa\chi}$, $\varkappa \leq 0$ is a parameter that characterizes the Kimball aggregator and χ captures the gross markup.

The maximization problem reads:

$$\Pi^F_t = \max_{q_t(i), i\in[0,1]} P_t Q_t - \int_0^1 p_t(i)\, q_t(i)\, di.$$

The solution of this maximization problem is a demand function for intermediate good (i):

$$\frac{q_t(i)}{Q_t} = \frac{1}{1+\varkappa}\left(\frac{P_t(i)}{P_t\Xi_t}\right)^{\iota} + \frac{\varkappa}{1+\varkappa}, \tag{17}$$

where $\varkappa \leq 0$ is a parameter, $\iota = \frac{\chi(1+\varkappa)}{1-\chi}$, and Ξ is the Lagrange multiplier associated with the constraint (16).

Intermediate firms buy the (homogeneous) output produced by the service sector firms in a competitive market at the real price φ_t and sell it to the final goods firms in a monopolistic competitive market. They can re-optimize their price $P_t(i)$ with probability $1-\zeta$. If they cannot reoptimize, they adjust their price at the steady state inflation rate Π. Therefore, the problem of the intermediate firm is:

$$\max_{P_{t+s}(i)} E_t \sum_{s=0}^{\infty} \beta^{t+s}\zeta^s \frac{\lambda_{t+s}}{\lambda_t}\left(P_t(i)\Pi^s - P_{t+s}\varphi_{t+s}\right) q_{t+s}(i) \tag{18}$$

subject to the demand function (17).

3.7 Market clearing

Market clearing in the market of intermediate good producers implies that the quantity sold summing over all producers i, must be equal to the production in the Service sector:

$$y_g l_{g,t} + y_b l_{b,t} = \int_0^1 q_t(i)\, di,$$

In turn, the aggregate production in the intermediate sector must equal the aggregate demand from final firms:

$$\int_0^1 q_t(i)\, di = Q_t \int_0^1 \left(\frac{1}{1+\varkappa}\left(\frac{P_t(i)}{P_t\Xi_t}\right)^{\iota} + \frac{\varkappa}{1+\varkappa}\right) di,$$

where we have made use of the demand function in eq.(17). Because in equilibrium all firms set the same price, it follows that

$$\int_0^1 q_t(i)\, di = Q_t.$$

3.8 Taylor rule

We assume that the conduct of monetary policy is described by a standard Taylor rule:

$$\frac{1+R_t}{1+R^*} = \left(\frac{1+R_{t-1}}{1+R^*}\right)^{\rho_r}\left[\left(\frac{1+\pi_t}{1+\pi^*}\right)^{\phi_\pi}\left(\frac{\tilde{Q}_t}{Q^*}\right)^{\phi_y}\right]^{1-\rho_r}, \tag{19}$$

where $\rho_R \in [0,1)$ captures the degree of interest rate smoothing and the parameters $\phi_\pi > 1$ and $\phi_y > 0$ capture the sensitivity of the interest rate set by the monetary authority to the deviations of the inflation rate and output from their steady state values π^* and Q^*, respectively.

4 Empirical Strategy

The model is log-linearized around its steady state equilibrium. In section 4.1, we discuss the calibration strategy. In Section 4.2, we explain how we implement our empirical strategy, which is to use the unemployment rate and the quit rate to obtain an implied series of price inflation.

4.1 Calibration

We calibrate the steady state of the model economy presented in Section (3) to the US economy at monthly frequency over the period 1996Q1-2018Q2. In order to do so, we assume a Cobb-Douglas matching function $M_t = \phi_0 \left[u_{0,t} + s_t (1 - u_{0,t})\right]^{1-\psi} v_t^{\psi}$, where $\psi \in (0,1)$ is an elasticity parameter and $\phi_0 > 0$ is a scale factor. The functional form above implies the job finding rate $\phi(\theta_t) = \phi_0 \theta_t^{\psi}$, where $\theta_t = v_t / \left[u_{0,t} + s_t (1 - u_{0,t})\right]$.

The calibration of the steady-state requires assigning values to the following eleven parameter values: β, ϕ_0, δ, y_b, y_g, υ, b, ξ_g, c, c^f and s. We set the discount factor β to 0.998, in order to match an annual real interest rate of 2%. We normalize θ to unity, which allows us to pin down the scale factor ϕ_0, so as to match a job finding rate of 0.32, the sample mean. The job separation rate δ is implied by the Beveridge curve, under the assumption of a steady state rate of unemployment of 5.5%. Namely, solving the Beveridge curve for $\delta = \frac{\phi_0 u_0}{1-u_0+\phi_0 u_0}$ yields a separation rate of 0.02. The productivity of a bad match is normalized to one and the productivity in a good match is set to be 12.5% higher. We show that results are robust to the precise parameterization of y_g and that assigning higher values would violate the incentive compatibility constraint, which requires that the surplus of bad matches should be positive both in steady state and in the simulation. Finally, we set the probability that workers will accept an equally valuable outside offer to be $\upsilon = 0.5$. We will show how the results change if we modify this parameter value.

This leaves us with five parameters to calibrate, the parameter governing unemployment benefits b, the probability of drawing a good match ξ_g, the flow cost of advertising a vacancy c, the fixed cost c^f, and the parameter governing search intensity s. These are calibrated in order to match: *(i)* a share of labor adjustment costs over consumption of 5%, which is implied by the assumption that the total costs of hiring, which include vacancy posting, screening and training amount to six weeks of wages, measured as the price of the labor service φ, in line with

Calibration			
Parameters	Description	Value	Target/source
	Parameters that affect the steady-state		
β	Discount factor	0.998	Real interest rate 2% p.a.
ϕ_0	Scale parameter matching fn	0.32	Job finding rate 32%
δ	Job separation rate	0.02	Unemployment rate 5.5%
y_b	Productivity bad matches	1	Normalization
y_g	Productivity good matches	1.125	Normalization
υ	Prob. of job switching if indifferent	0.5	Normalization
b	Unemployment benefits	0.845	Calibrated
c	Flow cost of vacancy	0.0067	Calibrated
c^f	Fixed cost of vacancy	0.3248	Calibrated
s	On the job search rate	0.372	Calibrated
ξ_g	Probability draw good match	0.28	Calibrated
	Parameters that do not affect the steady-state		
χ	Mark-up parameter	1.2	20% mark-up
$\varkappa$	Scale param. Kimball aggregator	10	Smets and Wouters (2007)
ζ	Calvo price parameter	0.9250	Quarterly probability is 80%
Π	Steady-state gross inflation rate	1.0016	Net inflation rate of 2% p.a.
ρ_r	Taylor rule smoothing parameter	0.65	Conventional
ϕ_π	Taylor rule response to inflation	1.8	Conventional
ϕ_y	Taylor rule response to output	0.25	Conventional
ψ	Elasticity of matching function	0.5	Moscarini and Postel-Vinay (2018)
ρ_S	Autocorrel. job search rate	0.7074	MLE estimation
$100\sigma_S$	St. dev. job search rate	3.3012	MLE estimation
ρ_λ	Autocorrel. preference shock	0.9894	MLE estimation
$100\sigma_D$	St. dev. preference shock	0.0165	MLE estimation
Variable	Description	Value	Target/source
	Steady state calibration targets		
$\frac{c}{\varpi}/c^f$	Ratio of variable to fixed cost	0.078	Silva and Toledo (2009)
$EE_t = \frac{\frac{1}{\upsilon}s\phi\left[l_b^0\left(\xi_b+\upsilon^{-1}\xi_g\right)+l_g^0\xi_g\right]}{l_b^0+l_g^0}$	EE transition rate	0.0233	Presample estimation using quit rate
θ	Labor market tightness	1	Normalization
$\frac{l_g^0}{l_g^0+l_b^0}$	Employment share in good jobs	0.67	Employment share at top 10% firms
$\frac{Q-C}{Q}$	Labor adjustment costs to output	0.05	Hiring costs totalling 6 weeks of wages
	Other implied steady state values		
$l_{b,0}$	measure of bad matches	0.3165	
$l_{g,0}$	measure of good matches	0.6285	
W	Present value of unit service	41.1579	
u_0	unemployment rate	0.055	

Table 1: Model parameters and implied steady-state values

the evidence reviewed in Faccini and Yashiv (2017);[3] *(ii)* a fraction of good jobs in steady state equal to 67%. The latter is the share of employment for the top 10% US firms by employment size in year 2000. The calibration of this parameter makes use of the empirical strategy in Moscarini and Postel-Vinay (2016a), who exploit the well-known correlation between firm size and productivity by assuming that employed workers climb the ladder when moving to larger firms; *(iii)* a normalized value of labor market tightness equal to one; (iv) a ratio of total variable costs of hiring to fixed costs $\frac{c}{\varpi}/c^f$ equal to 0.078. This value is the ratio of pre-match recruiting, screening and interviewing costs to post-match training costs in the US, following the analysis of Silva and Toledo (2009), which is based on the 1982 Employer Opportunity Pilot Project (EOPP), a cross-sectional firm-level survey that contains detailed information on both pre-match and post-match labor turnover costs in the United States; *(v)* a monthly job-to-job transition rate of 2.33%, which is the average quit rate measured in the pre-Great Recession sample. We have checked that the value of b implied by the calibration is consistent with a positive surplus for low quality matches both in steady state and in the simulation.

Turning now to the parameters that do not affect the steady-state, we set the smoothing coefficient of the Taylor rule to the value of 0.65, and the response parameters to inflation and output to the values of 1.8 and 0.25, respectively. The mark-up parameter χ is set to equal 1.2, which implies a 20% price mark-up. The Calvo parameter governing price stickiness is set to 0.925, which in quarterly frequency implies a probability of not readjusting prices equal to 0.8. The scale parameter of the Kimball aggregator is set to 10, the value used by Smets and Wouters (2007). The steady state gross rate of inflation is set to equal 1.0016, which implies a 2% annualized rate of inflation. Finally, we set the elasticity of vacancies in the matching function ψ to equal 0.5 to be consistent with estimates by Moscarini and Postel-Vinay (2018b), which account for workers searching on the job.

The autocorrelation parameters and the standard deviations of the exogenous process for the demand shock, λ_t, and the fraction of employed workers looking for a new job, s_t, are estimated with maximum likelihood using the model and time series data for both the civilian unemployment rate and the quit rate from February 1996 through September 2018.[4] These data are available at monthly frequency. As shown in Table 4.1, the preference shocks λ_t follows an almost unit root process with small variance whereas the fraction of employed workers looking for a new job s_t is substantially less persistent and more volatile.

[3] This follows from the observation that the measure of matches in the calibration is around 0.042. Assuming that hiring costs are 1.5 times the average cost of the labor service, which is 5/6, makes total hiring costs 0.052. Given that total output equals 1.082, total hiring costs are 4.8% of output.

[4] The civilian unemployment rate (LNS14000000) comes from comes from the 'Current Population Survey (Household Survey)' and is released by the U.S. Bureau of Labor Statistics. The quit rate is the number of quits in a given month as a percent of total employment. Total employment is total nonfarm payroll employment from the Establishment Survey (Haver: LANAGRA@USECON). Source: Bureau of Labor Statistics. Before 2000, we use the series of quit rate constructed by Davis, Faberman, and Haltiwanger (2012).

The quit rate is assumed to be informative about the job-to-job flow rate EE_t, which is defined in the model as

$$EE_t \equiv \frac{\upsilon s_t \phi_t \left[l^0_{b,t} \left(\xi_b + \upsilon^{-1} \xi_g \right) + l^0_{g,t} \xi_g \right]}{l^0_{b,t} + l^0_{g,t}}. \tag{20}$$

This measurement equation reflects the assumption that all bad workers who find a good match change job, whereas workers employed in low and high quality matches who find an equally valuable match only change job with probability υ. We show that using the employment-to-employment flow rate measured by the CPS would deliver very similar results. While this rate is in principle a better choice to inform job to job transitions, it is known to be affected by important measurement errors. We are going to show that using this measure would strengthen our results. In fact, the E2E rate computed by the CPS has not reverted back yet to its pre-Great Recession level. This pattern turns out to totally freeze the working of the ladder rather than simply slowing it down as done by the quit rate. Since we deem this result as a bit extreme, we prefer estimating the parameters of the processes λ_t and s_t using the quit rate. While the quit rate includes also people who quit their job to either become unemployed or to leave the labor force, people generally switch jobs by quitting (rather than losing) their previous job. Elsby, Hobijn, and Sahin (2010) document that the vast majority (86%) of people observed quitting their job tend to move directly to a new job, rather than becoming unemployed or exiting the labor force.

4.2 Taking the Model to the Data

We set the parameters to the values reported in Table 4.1 and then we filter the model variables using the monthly series of the unemployment rate and the quit rate. These series are the same as those we used to estimate the parameters of the two exogenous processes in the model. The filter returns the model's predicted dynamics of inflation and the other variables conditional on observing the series of the unemployment rate and the quit rate from February 1996 through September 2018. Before filtering, the model is log-linearized around its deterministic steady-state equilibrium so that the Kalman smoother can be applied. Log-linearized variables are denoted with $\hat{}$. Rates and shares are linearized and denoted by $\tilde{}$.

Identification of Key Labor Market Variables in the Data It is important to notice that for a given initial value of the bad matches, observing the unemployment rate and the quit rate pins down uniquely the entire time series of the two variables that characterize the state of the job ladder; that is, the fraction of good and bad matches $l^0_{g,t+1}$ and $l^0_{b,t+1}$. It can be shown that these two observable time series also precisely determine time sequences for the rate of on-the-job search s_t, the job finding rate ϕ_t, the vacancy filling rate ϖ_t, and the entire expected discounted stream of real marginal costs of intermediate goods, $\mathcal{W}_t$.

To see this, note that the observed series of unemployment rates informs $u_{0,t+1}$ and hence the aggregate unemployment at the end of the period, u_t, through the following equation

$$\tilde{u}_t = \frac{\tilde{u}_{0,t+1}}{1-\delta},$$

which is is obtained by combining equations (3) and (4) and linearizing.

Endowed with the end of period unemployment rate $\tilde{u}_t$, we can linearize equation (5) to pin down the job finding rate $\tilde{\phi}_t$:

$$\tilde{\phi}_t = \frac{(1-\phi)\,\tilde{u}_{0,t} - \tilde{u}_t}{u_0}, \tag{21}$$

where u_0 denotes the unemployment rate at the beginning of the period in steady state.

Using the observed job-to-job flow rate, we can linearize the measurement equation (20), and obtain the following equation that allows us to pin down the on-the-job search rate $\tilde{s}_t$:

$$\begin{aligned} \tilde{s}_t \quad = \quad & \frac{s}{EE}\widetilde{EE}_t - \frac{s}{\phi}\tilde{\phi}_t - \frac{s}{\upsilon\left(l_b^0 + l_g^0\right)}\left[\frac{s\phi\left[\left(\xi_b + \upsilon\xi_g\right)\right]}{EE} - 1\right]\tilde{l}^0_{b,t} \\ & - \frac{s}{\upsilon\left(l_b^0 + l_g^0\right)}\left[\frac{s\phi\xi_g}{EE} - 1\right]\tilde{l}^0_{g,t}. \end{aligned} \tag{22}$$

This equation shows that given the observed values of the employment to employment transition rate, $\widetilde{EE}_t$, and of the unemployment rate, which in turn implies the job finding rate $\tilde{\phi}_t$ via equation (21), equation (22) pins down the on-the-job search rate s_t. Note that $\tilde{l}^0_{b,t}$ and $\tilde{l}^0_{g,t}$ are predetermined at time t.

With the rates $\tilde{\phi}_t$ and $\tilde{s}_t$ at hand, we can use the observed unemployment rate $\tilde{u}_{0,t}$ to pin down the fraction of bad and good match in the next period $t+1$, using the linearized laws of motion for low and high quality matches in (14) and (15), which read

$$\begin{aligned} \tilde{l}^0_{b,t+1} \quad = \quad & -(1-\delta)\,\phi\xi_g l_b^0 \tilde{s}_t - \left[(1-\delta)\,s\xi_g l_b^0 + \xi_b u_0\right]\tilde{\phi}_t \\ & + (1-\delta)\left[1 - s\phi\xi_g\right]\tilde{l}^0_{b,t} + (1-\delta)\,\phi\xi_b\tilde{u}_{0,t} \end{aligned} \tag{23}$$

$$\tilde{l}^0_{g,t+1} = (1-\delta)\left[\tilde{l}^0_{g,t} + \phi\xi_g l_b^0 \tilde{s}_t + s\phi\xi_g\tilde{l}^0_{b,t} + \phi\xi_g\tilde{u}_{0,t} + \left[s\xi_g l_b^0 + \xi_g u_0\right]\tilde{\phi}_t\right]. \tag{24}$$

In addition, the assumption of a Cobb-Douglas matching function implies that given the job finding rate $\tilde{\phi}_t$ we can solve for labor market tightness

$$\hat{\theta}_t = -\frac{\psi\tilde{\phi}^{-\psi-1}}{\theta_*}\tilde{\phi}_t, \tag{25}$$

and therefore recover the vacancy filling rate

$$\tilde{\varpi} = \phi_0 \left(\psi - 1\right) \tilde{\theta}^{\psi-1} \hat{\theta}_t. \tag{26}$$

It should be noted that the observed unemployment rate uniquely pins down the series of the labor market tightness $\hat{\theta}_t$ and the vacancy filling rate via equations (25) and (26).

Finally, we can linearize the free entry condition in eq.(11) and then solve the resulting equation for $\widehat{\mathcal{W}}_t$. Since we know the time$-t$ value of all the variables appearing in this equation, we can obtain the implied expected present discounted value of the entire stream of current and future expected real marginal revenues from selling a unit of service until separation, $\widehat{\mathcal{W}}_t$. Furthermore, note that the linearized equation (10) becomes

$$\widehat{\mathcal{W}}_t = \frac{\varphi}{\mathcal{W}} \widehat{\varphi}_t - (1-\delta)\beta \left[\widehat{\lambda}_t - E_t[\widehat{\lambda}_{t+1}] - E_t[\widehat{\mathcal{W}}_{t+1}] \right]. \tag{27}$$

Inflation in the model is given by the New Keynesian Phillips Curve:

$$\widehat{\pi}_t = \frac{(1-\zeta)(1-\zeta\beta)}{\zeta(1-\chi\varkappa)} \widehat{\varphi}_t + \beta E_t \widehat{\pi}_{t+1}. \tag{28}$$

Comparing equations (27) and (28) reveals that the variable $\widehat{\mathcal{W}}_t$ is governed by an equation that is very akin to the New Keynesian Phillips curve that determines inflation in the model. In fact, the difference between these two equations is given by an exogenous factor $(\widehat{\lambda}_t - E_t[\widehat{\lambda}_{t+1}])$, the discount rate, and a scaling factor. While the discounting of future real marginal costs in the Phillips curve is given by β, to get $\widehat{\mathcal{W}}_t$ one needs to discount at a slightly lower rate $(1-\delta)\beta$. Since the monthly job separation rate δ is a small number, the two discounting rates are not very different in practice. Therefore, the variable $\widehat{\mathcal{W}}_t$ can be regarded as a salient proxy for the inflationary pressures stemming from the labor market in our model. Importantly, our model pins down the evolution of this crucial labor market variable by using only the series of the unemployment rate and the employment-to employment flow rate.

4.3 The Effects of Labor Market Misallocation on Inflation

We want to use our stylized model with the job ladder to assess the effects of labor market misallocation on inflation. We are particularly interested in evaluating these effects during the Great Recession and its aftermath since this is a period where labor resources were highly and persistently misallocated, as Figure 1 suggests. To this end, we use the Kalman smoother to simulate the model's variables.[5] The smoother returns the model's predicted value of any variable conditional on all the observations used in the exercise. We use two data sets. Both

[5]Using the Kalman filter (one-sided filter) would produce very similar pattern for all the model's variables.

data sets include the civilian unemployment rate as a measure of the rate of unemployment in the model. But while one data set includes the quits rate as a proxy for the employment-to-employment transition rate, the other one relies on the job-to-job flow rate measured by the CPS.

As shown in the previous section, observing the unemployment rate and one series for E2E transitions is sufficient to identify the degree of labor misallocation (i.e., the dynamics of the bad and good matches) in the model (for a given initial value of bad jobs). This model's property has two important implications. First, adding other observables will not change the model's assessment of the degree of misallocation.[6] Second, introducing additional business cycle shocks, such as shocks to TFP would not change the model's evaluation of the degree of misallocation in our sample period.[7] This consideration as well as our maintained objective of assessing the effects of bad jobs on inflation rather than providing a universal theory of inflation determination prompted us to keep the number of shocks to the minimum. We consider this goal as quite ambitious already as this is the first paper that conducts a time-series investigation of a general equilibrium model with the job ladder.

The upper plots of Figure 4 show the traditional labor market variables: the unemployment rate, the job finding rate, and the vacancy filling rate. All these measures suggest that the U.S. labor market has become tight starting from 2015 and on. Note that the unemployment rate in the model is identical to the civilian unemployment rate in the data. The job finding rate is not observed but is effectively pinned down by the civilian unemployment rate and is highly correlated with the series computed by Shimer (2005).[8]

While the labor market has become quite tight in recent years, according to the measures in the upper panel of Figure 4, the variables plotted in the lower panel suggest that the labor market went through a prolonged period of misallocation: bad jobs persist for a very long while after the end of the Great Recession. Bad matches have risen dramatically since the end of the last recession: as the job finding rate plummets to historically low levels, matches that are destroyed every period fail to be reabsorbed, generating a sharp increase in the rate of unemployment. In turn, this implies that an increasing fraction of displaced workers start climbing the ladder anew, raising the fraction of bad matches, $l_{b,t}^{0}$. Because the job to job rate is low by historical standards already at the onset of the financial crisis, job matches fail to be reabsorbed quickly, inducing hump shaped dynamics that lasts for an entire decade.

Unlike the number of bad jobs, the fraction of good jobs keeps declining for the first few

[6] The other side of the coin of this implication is that the econometrician's uncertainty about the model's predictions about the fraction of bad and good matches is always zero.

[7] Adding some shocks to the labor market, such as shocks to the job separation rate would affect the model's assessment of the fraction of bad jobs.

[8] After the Great Recession, our implied series for the job finding rate drops less and recover more quickly than that in Shimer. Hence, using the series constructed by Shimer would have slowed down the working of the ladder even more.

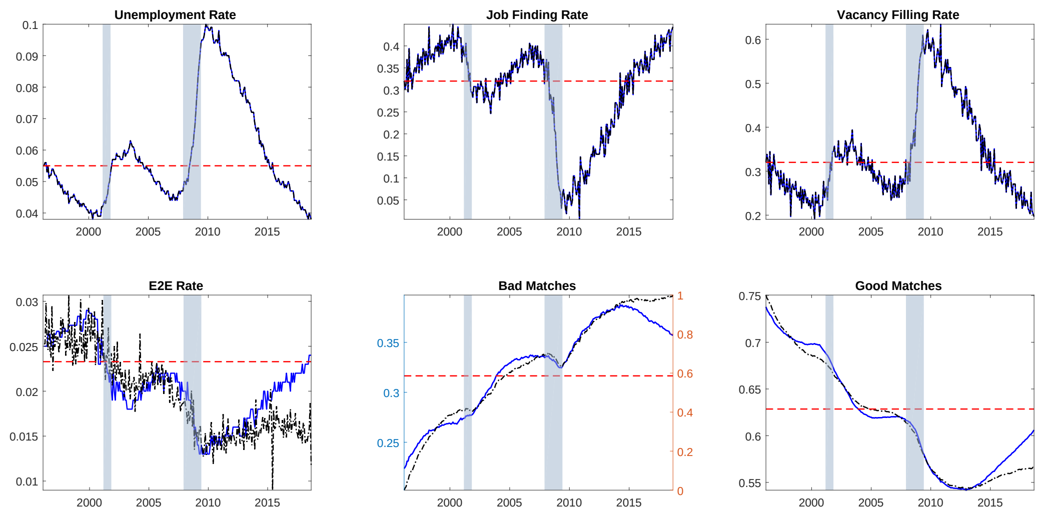

Figure 4: Labor market variables obtained by feeding the calibrated model with historical realizations of the unemployment rate and the quit rate.

years of the post-Great Recession recovery. Only in 2014, the model predicts that the fraction of good jobs starts recovering. This delayed creation of good jobs is an implication of the ladder. Given that the conditional probability of finding a good match in every period $(1-\xi_b)$ is only 28 percent, most workers will need to go through employment spells in bad matches before finding a good one.

Using the employment-to-employment flow rate from the CPS (the black dashed-dot lines) leads to very similar implications until 2013. Afterwards, the employment-to-employment rate remains subdued and never fully recovers contrary to the JOLTS quit rate. This missed recovery has the striking effect of blocking entirely the working of the job ladder. As the E2E rate fails to recover, the fraction of bad jobs remains permanently high. This freezing of the job ladder will have important implications for inflation as we will show next.

Crucially, the behavior of bad matches is reflected in a mirror image response of the present discounted value of the labor service $\mathcal{W}_t$ (not shown). The intuition is as follows. At times when labor is highly misallocated, entrant firms are more likely to poach employed workers away from their current bad match and into a good one. This implies a higher expected return from posting vacancies, given that under Bertrand competition workers switch from bad to good jobs without any gain in present value terms. As more vacancies enter the labor market, the aggregate output of the service sector rises, and its price falls. Price stickiness in the intermediate sector implies that the current and future real prices of the service falls, leading to a fall in $\mathcal{W}_t$ and in price inflation. So ultimately, the behavior of bad matches is key to explain the dynamics of real marginal costs and inflation.

The left graph of Figure 5 shows the effects of labor misallocation accumulated during and

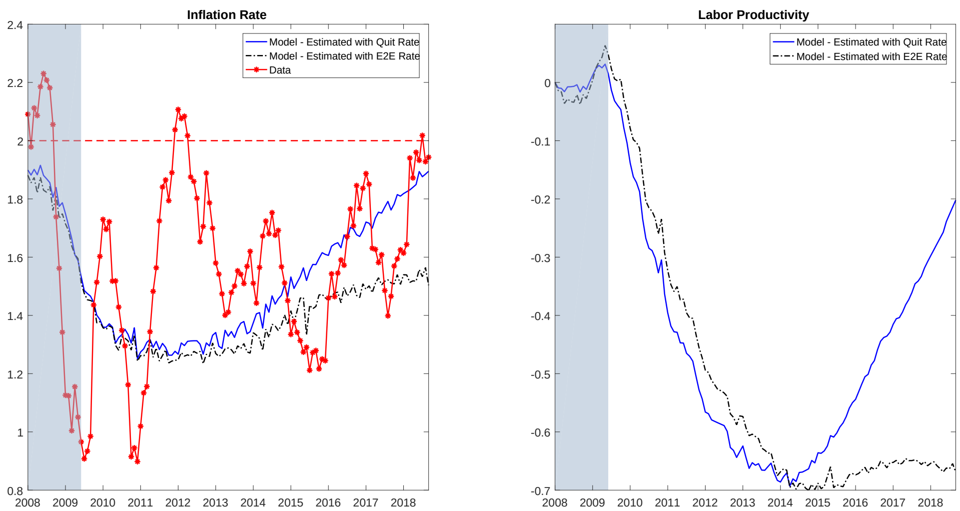

Figure 5: Right Graph: Observed core PCE inflation rate (red line with dot markers) and model generated paths for inflation obtained when observing the JOLTS quit rate (solid blue line) and the CPS EE rate (dashed black line), respectively. All rates are annualized and in percent. Left Graph: Labor productivity in percentage deviation from its long-run level. We normalize the series of labor productivity so that the deviation in December 2017 to zero.

after the Great Recession on inflation. The blue line (dashed-dot blue line) represents the dynamics of inflation implied by the JOLTS quits rate (CPS E2E rate). The red line with dots denotes the observed core PCE inflation rate. In both cases, the labor misallocation produced by the Great Recession has contributed to keeping inflation below its long-term level (red dashed line) for several years. The model seems to deliver predictions that captures fairly well the low frequency dynamics of PCE core inflation rate in the data. This finding suggests that labor market misallocation seems to have played an important role in shaping the dynamics of inflation in the post-Great Recession recovery. The fit does not have to be perfect given the stylized nature of the model and the fact that only two series from the labor market are used as observables for the exercise. The model delivers this prediction even though the rate of unemployment and the job finding rate suggest that the U.S. labor market is very hot, as shown in Figure 4. When the employment-to-employment flow rate of CPS is observed, inflation does not revert to its long-run level, not even in late 2018. This is because the E2E rate computed by the CPS does not revert back to its long-run trend, which causes the job ladder to stop working and the fraction of bad jobs accumulated up to 2014 to plateau in the subsequent years.

The labor market misallocation owing to the Great Recession does not only have an effect on price dynamics but it also works as a drag for labor productivity. The right graph of Figure 5 shows the model's predicted percentage deviations of labor productivity from its long-run value during the Great Recession and the following recovery.

Labor misallocations persistently lower labor productivity at the beginning of the recovery. The rising number of bad jobs keeps labor productivity down and away from its long-run value.

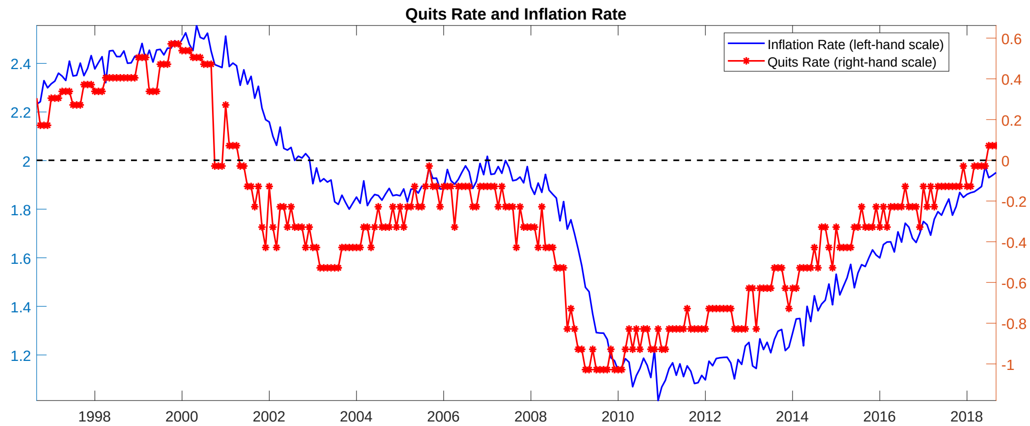

Figure 6: JOLTS quit rate and core PCE inflation.

Depending on the job-to-job flow rate we observe (i.e., the quits rate of the E2E rate of Census), we obtain fairly different patterns for labor productivity in the second half of the recovery. The reason is that these two rates imply very different dynamics for the bad matches in the model, as shown in the lower graphs of Figure 4. The E2E rate of Census fell and has remained below its long-run value so far. Thus, the fraction of bad jobs $l_{b,t}$ remains high until the end of the sample and, as a consequence, labor productivity gets stuck at 0.65 percent below its long-run value since the beginning of the recovery. In this scenario the growth rate of labor productivity is particularly subdued. If the quits rate from JOLTS is observed instead, then the E2E flow rate in the model converges to its long-run value in 2018. As a result, the fraction of bad matches $l_{b.t}$ also slowly converges towards its steady-state value and so does labor productivity. Yet, the model predicts that the gap in labor productivity is not fully reabsorbed at the end of 2018 even when we use the quits rate. Either ways, labor productivity remains below its long-run level by the end of 2018 because persistent labor misallocation.

What Predictor for Inflation? Figure 6 shows that the quit rate effectively anticipates the changes in the rate of inflation in sample. Therefore, this rate can be considered a good predictor of the component of inflation that is affected by labor market misallocation. This result echoes the findings in the reduced-form literature, such as Faberman and Justiniano (2015).

Nevertheless, this result seems less striking when one looks at these correlations in the model's population. The blue bars in Figure 7 show the cross-correlation between inflation and the E2E rate at monthly horizon zero through twelve. The 12-month-ahead correlation between inflation and the E2E rate is roughly 20 percent. In population, a much better predictor for inflation seems to be the fraction of good jobs, whose cross-correlation with inflation in the model is shown by the yellow bars in Figure 7. While the fraction of good matches is hard to

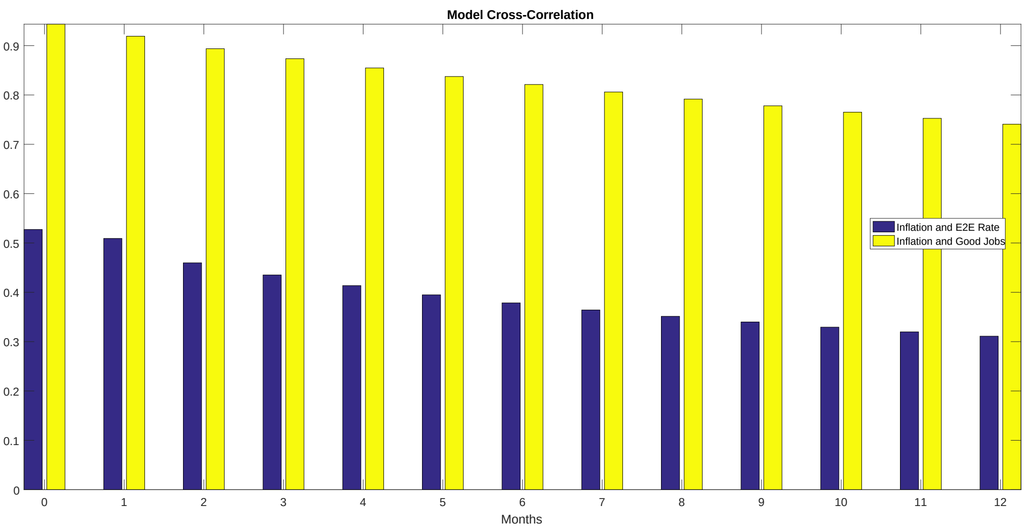

Figure 7: Model generated cross correlations in population of the EE rate (purple bars) and the measure of good matches (yellow bars) with the core CPE rate of inflation.

measure directly in the data, we showed that our model with the job ladder precisely pins down the dynamics of the good matches given the rate of unemployment and a measure of the E2E flow rate.

5 Concluding Remarks

We have shown that standard New Keynesian models cannot explain why, during the current economic expansion, U.S. price and wage inflation have remained subdued in spite of a record tight labor market. We have presented a general equilibrium model with the job ladder that successfully reconciles these facts and partially accounts for the observed slow down in labor productivity growth. The model emphasized the central role of the inflationary pressures arising from the renegotiation of the wages of employed workers, who attract outside offers when searching on the job. In this model, inflationary pressures are tied to the distribution of match quality across jobs and are lower the greater the amount of misallocation. Simulating the model using only labor market data reveals that inflation has been low in recent years because of the persistent amount of misallocation induced by the Great Recession. The behavior of bad matches predicted by the model resembles those of part-time workers for economic reasons. Our empirical evidence provides further support to the view that monetary authorities concerned with the sources of wage and price inflation should pay more attention to fluctuations of labor misallocation over the cycle, and its implications for the wage dynamics of the employed workers.

References

Christiano, L. J., M. S. Eichenbaum, and M. Trabandt (2016): "Unemployment and Business Cycles," *Econometrica*, 84, 1523–1569.

Davis, S. J., R. J. Faberman, and J. Haltiwanger (2012): "Labor market flows in the cross section and over time," *Journal of Monetary Economics*, 59(1), 1 – 18.

Dotsey, M., and R. G. King (2005): "Implications of state-dependent pricing for dynamic macroeconomic models," *Journal of Monetary Economics*, 52(1), 213–242.

Elsby, M. W. L., B. Hobijn, and A. Sahin (2010): "The Labor Market in the Great Recession," *Brookings Papers on Economic Activity*, 41(1 (Spring), 1–69.

Faberman, R. J., and A. Justiniano (2015): "Job Switching and Wage Growth," *Chicago Fed Letter*.

Faccini, R., and E. Yashiv (2017): "The Importance of Hiring Frictions in Business Cycles," Queen Mary University, mimeo.

Galí, J. (2008): "Monetary Policy and the Open Economy," in *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Introductory Chapters. Princeton University Press.

Galí, J., F. Smets, and R. Wouters (2011): "Unemployment in an Estimated New Keynesian Model," in *NBER Macroeconomics Annual 2011, Volume 26*, NBER Chapters, pp. 329–360. National Bureau of Economic Research, Inc.

Gertler, M., C. Huckfeldt, and A. Trigari (2016): "Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires," NBER Working Papers 22341, National Bureau of Economic Research, Inc.

Kimball, M. S. (1995): "The Quantitative Analytics of the Basic Neomonetarist Model," *Journal of Money, Credit and Banking*, 27(4), 1241–1277.

Krause, M. U., D. Lopez-Salido, and T. A. Lubik (2008): "Inflation dynamics with search frictions: A structural econometric analysis," *Journal of Monetary Economics*, 55(5), 892–916.

Levin, A., J. D. Lopez-Salido, and T. Yun (2007): "Strategic Complementarities and Optimal Monetary Policy," CEPR Discussion Papers 6423, C.E.P.R. Discussion Papers.

Lindé, J., and M. Trabandt (2018): "Resolving the Missing Deflation Puzzle," mimeo, Freie Universitat Berlin.

Moscarini, G., and F. Postel-Vinay (2016a): "Did the Job Ladder Fail after the Great Recession?," *Journal of Labor Economics*, 34(S1), 55–93.

——— (2016b): "Wage Posting and Business Cycles," *American Economic Review*, 106(5), 208–213.

——— (2017): "The Relative Power of Employment-to-Employment Reallocation and Unemployment Exits in Predicting Wage Growth," *American Economic Review*, 107(5), 364–368.

——— (2018a): "The Job Ladder: Inlfation vs Reallocation," Mimeo Yale U. and U. College London.

——— (2018b): "On the Job Search and Business Cycles," Mimeo Yale U. and U. College London.

Postel-Vinay, F., and J.-M. Robin (2002): "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica*, 70(6), 2295–2350.

Shimer, R. (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *The American Economic Review*, 95(1), 25–49.

Silva, J. I., and M. Toledo (2009): "Labor Turnover Costs and the Cyclical Behavior of Vacancies and Unemployment," *Macroeconomic Dynamics*, 13(S1), 76–96.

Smets, F., and R. Wouters (2007): "Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach," *American Economic review*, 97(3), 586–606.

A Appendix

B Construction of the time series and their sources

The time series used for the VAR analysis have been constructed from the following data downloaded from the Federal Reserve Economic Data (FRED). The labor share of income is computed as the ratio of total compensation in the non-farm business sector divided by nominal non-farm GDP. In turn, total compensation is computed as the product of compensation per hour (COMPNFB) times total hours (HOANBS) and nominal GDP is the product of real output (OUTNFB) times the appropriate deflator (IPDNBS). All series are quarterly and seasonally adjusted. We compute the deviations of the labor share from its trend by computing log deviations from an eight year moving average.

We follow Shimer (2005) and compute the job finding rate as $\phi_t = 1 - \left(u_{t+1} - u^s_{t+1}\right)/u_t$, where u^s_{t+1} denotes the number of workers employed for less than five weeks in month $t+1$ (UEMPLT5). The total number of workers unemployed in each month is computed as the sum of the number of civilians unemployed less than five weeks (UEMPLT5), for 5 to 14 (UEMP5TO14), 15 to 26 weeks (UEMP15T26), and 27 weeks and over (UEMP27OV). Primary data is constructed by the U.S. Bureau of Labor Statistics from the CPS and seasonally adjusted. To obtain quarterly percentage point deviations of the job finding rate from its trend we average monthly data over each quarter and then subtract the actual job finding rate from its eight year moving average.

We also use data on real gross domestic product (GDPC1), real gross private domestic

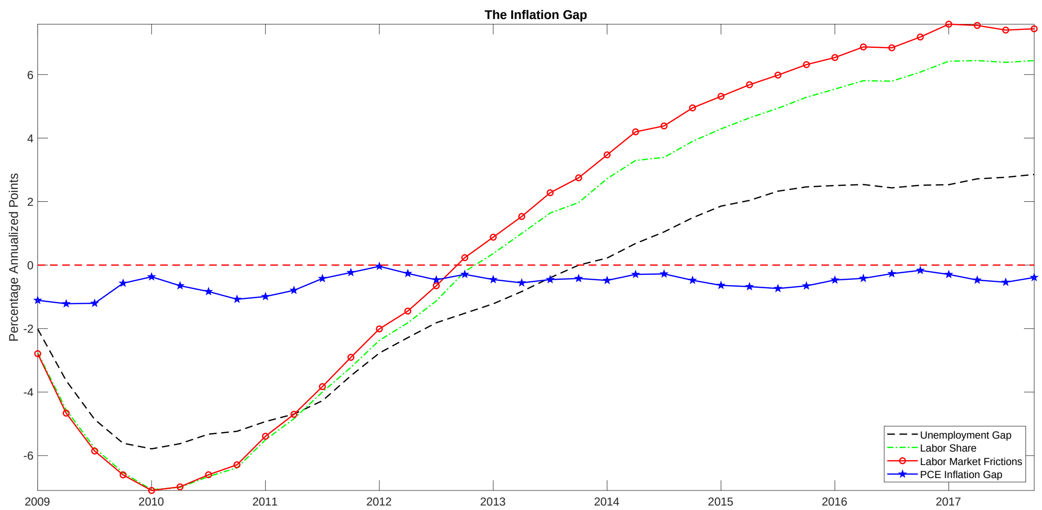

Figure 8: Shows the inflation dynamics from 2009Q1 through 2017Q3 using three traditional theories of inflation augmented with the assumption of price indexation and the core PCE inflation gap.

investment (GDPIC1) and real personal consumption expenditures (PCECC96). All data are quarterly and seasonally adjusted. When computing percentage deviations of these time series from their trend we first remove a quadratic trend from the variables in logs,and then take the difference from their eight year moving averages. To compute percentage deviations of real wages from the trend we first remove a linear trend to the log of compensation per hour (COMPNFB) and then take the difference with respect to its eight year moving average.

We measure aggregate price inflation by taking log differences on the previous quarter of the seasonally adjusted consumer price index for all urban consumers (CPIAUCSL). We also use quarterly data on the effective Federal Funds rate (FFR) and on the short-term Natural Rate of Unemployment (NROUST). We compute percentage point deviations of inflation, the Federal Funds rate and the natural rate of unemployment from their trend as the difference from their eight year moving average.

C Testing the Traditional New Keynesian Theories with Price Indexation

With price indexation, the New Keynesian Phillips curve becomes:

$$\pi_t = \iota \pi_{t-1} + \kappa \varphi_t + E\pi_{t+1}, \tag{29}$$

where the parameter ι controls the degree of price indexation, which affects the relative importance of the backward component of the New Keynesian Phillips curve. We can redo the same VAR-based exercise made in Section 2 in order to assess the results of that section, which were based on assuming no indexation. We set the degree of price ι_p to 0.65. Higher values are rarely used in the empirical literature. The lagged inflation value in the first quarter of 2009 is taken from the data (i.e., it is core PCE inflation in the fourth quarter of 2018).

Figure 8 confirms the main result in the text: Price indexation just makes the drop in inflation in 2009 more pronounced and delays the period after which inflation goes above its long-run level by just three quarters. The New Keynesian Phillips curve cannot explain why we have not observed high inflation lately even if we assume price indexation.

D Computation of real marginal costs in a standard NK model with search and matching frictions

We follow the work by Krause, Lopez-Salido, and Lubik (2008), who study the behavior of real marginal costs in a simple New-Keynesian model with search and matching frictions in the labor market. Eq.(32) in page 898 defines the real marginal cost as:

$$mc_t = \frac{W_t}{\alpha\left(\frac{y_t}{n_t}\right)} + \frac{c'\left(v_t\right)/q\left(\theta_t\right) - \left(1-\rho\right)E_t\beta_{t+1}c'\left(v_{t+1}\right)/q\left(\theta_{t+1}\right)}{\alpha\left(\frac{y_t}{n_t}\right)}, \tag{30}$$

where W_t denotes the real hourly wage, y_t/n_t is the average product of labor, $c'\left(v_t\right)$ is the derivative of the vacancy cost function with respect to vacancies, $q\left(\theta_t\right)$ is the vacancy filling rate β_{t+1} is the discount factor and α is the elasticity of output to employment in the production function. The first component on the RHS of eq.(30) is the unit labor cost, i.e. the ratio of the labor cost and the marginal product of labor. The second component is stems from the existence of search and matching frictions and can be interpreted as cost savings from not having to hire in the following period.

Let $s_t \equiv W_t/\alpha\left(\frac{y_t}{n_t}\right)$ denote the unit labor cost, which equals the labor share of income divided by the elasticity of output to employment. Krause, Lopez-Salido, and Lubik (2008) show that linearizing eq.(30) and rearranging, leads to the following expression:

$$\widehat{mc_t} = \hat{s}_t + \frac{1-\phi}{1-\tilde{\beta}}\left[\frac{\xi}{1-\xi}\left(\hat{h}_t - \tilde{\beta}E_t\hat{h}_{t+1}\right) + \left(\varepsilon_c - 1\right)\left(\hat{v}_t - \tilde{\beta}E_t\hat{v}_{t+1}\right) - \tilde{\beta}E_t\hat{\beta}_{t+1} - \left(1-\tilde{\beta}\right)\hat{w}_t\right], \tag{31}$$

where a hat variable is used to denote log deviations from the steady-state, h_t denotes the job finding rate, $\tilde{\beta}$ is a discount factor adjusted for the rate of job separation,ε_c is the elasticity of vacancy costs to vacancies, ξ is the elasticity of the matching function with respect to unemployment and $\phi = s/mc$ is the share of unit labor cost over total marginal costs. We follow the calibration in Krause, Lopez-Salido, and Lubik (2008) and assume that $\xi = 0.5$, $1 - \phi = 0.05$ and $\tilde{\beta} = 0.943$. In line with the model specified in Section (3), we assume a linear vacancy cost function, which implies $\varepsilon_c = 1$, and risk neutrality, which implies that $\hat{\beta}_t = 0$.