Bad Jobs and Low Inflation[*]

Renato Faccini

Queen Mary, University of London

Centre for Macroeconomics (LSE)

Leonardo Melosi

FRB Chicago

European University Institute

CEPR

April 30, 2019

Abstract

In a dynamic general equilibrium model with a job ladder, inflation rises when most workers are employed in high-productivity jobs because in this case, poaching leads to wage increases that are not backed by changes in productivity. The model predicts that the drop in the job-to-job flow rate, which started before the onset of the Great Recession, has significantly slowed the pace at which the U.S. labor market turns low-productivity jobs into high-productivity ones. As a result, inflation has fallen below trend in the last decade, despite the marked decline in the unemployment rate and the dismal labor productivity growth. The model explains well the cyclical dynamics of inflation over the last 25 years. The declining trend in the employment-to-employment flow rate is crucial for the model to explain why the last three recoveries were not characterized by significant spikes in inflation.

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

JEL codes: E31, E24, C78

1 Introduction

After experiencing the most severe contraction in the postwar history, the U.S. economy has been expanding for almost 10 years. While this long spell of recovery has brought about the lowest unemployment rate ever recorded in the last fifty years, inflation has remained subdued. As shown in Figure 1, since March 2017 the civilian unemployment rate has hovered consistently below its average level measured over the last twelve months of the previous expansion, suggesting that the labor market has been very tight. At the same time, PCE core inflation has remained persistently below its long-run trend. Moreover, the gap of labor productivity relative to its pre-Great Recession level widened to negative six percent in 2016 and has remained stable at that level ever since. Understanding why such tight labor market conditions coupled with low productivity growth have not sparked inflation yet, proves to be very challenging for standard macroeconomic models.

We show that a dynamic general equilibrium model with a job ladder can reconcile these three empirical facts. In the model the quality of jobs is match specific and can be either good or bad. Workers searching on the job climb the ladder by moving to more productive jobs. In recessions, workers are shaken off the ladder and the average job quality worsens as workers who have lost their jobs start climbing up again from the lowest rung. In the aftermaths of a recession, when the labor market is flooded with low quality jobs, wage rises are largely accompanied by changes in employers that bring about increases in productivity. As a result, wage pressures are not inflationary. But after years of economic expansion, most workers will have climbed up the ladder and the average job quality will be high. In this case, wage rises obtained from firms poaching workers in good matches do not come along with any rise in productivity. As a result, the real marginal cost rises, and wage renegotiations are inflationary.

Figure 2 provides some evidence supporting the view that an anomalous malfunctioning of the job ladder may have played a role in persistently slowing down price dynamics during the post-Great Recession recovery. Indeed, the employment-to-employment (EE) flow rate and the quit rate have remained persistently below their pre-recession levels during the recovery. The lowering of these flows may have hindered the reallocation of labor by slowing the transition of workers from low- to high-productivity jobs. Following the arguments above, a general equilibrium model with a job ladder could potentially account for the behavior of inflation over the last ten years if the legacy of the great recession was a large mass of low productivity jobs, which failed to be fully reabsorbed for as long as a decade. Could the sclerosis of job-to-job flows possibly account for such a persistent rise in bad jobs? The right panel of figure 2 provides some suggestive evidence in this direction. The workers employed part time for economic reasons (i.e., part timers who want a full time job but could not find it) as a fraction of civilian employment more than doubled during the Great Recession, highlighting the massive disruptive effects that
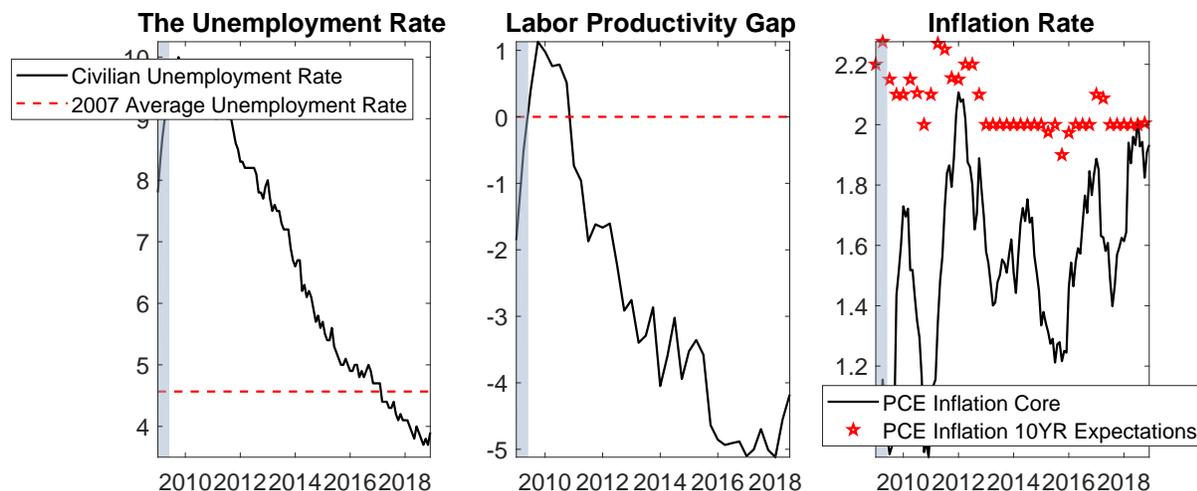
Figure 1: Labor market and inflation dynamics during the post-Great Recession recovery. The left graph: the civilian unemployment rate (solid line) and its average computed over the twelve months that preceded the Great Recession (red dashed), computed over the period December 2006 through November 2007. Frequency: Monthly. Source: Bureau of Labor Statistics (BLS). The center graph: labor productivity gap computed in deviation from their eight-year moving-average trend and normalized to be zero at the beginning of the Great Recession in 2007Q4. Frequency: Quarterly. Sources: BLS (NFBUS, nonfarm business sector, all employed person, index). The 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 shaded areas denote NBER recessions.

the recession had on the allocation of labor. Quite strikingly, this mass of workers trapped in unwanted jobs has reached its pre-recession average level (February 1996-November 2007, the red dashed line) only after nine years of economic recovery. This indicator suggests that a considerable fraction of workers may be still stuck in low-productivity jobs.

We calibrate the model and we obtain its in-sample predictions about inflation using only labor market variables: (i) the civilian unemployment rate and (ii) the quit rate from the Job Openings and Labor Turnover Survey (JOLTS) database administered by the Bureau of Labor Statistics (BLS). These series are observed in the U.S. from January 1996 through September 2018. We show that the persistent fall in the quit rate can explain in the model why the inflation rate has been so low during the post-Great recession recovery, despite record-low unemployment rates. The drop in this rate considerably slows down the creation of good jobs during the recovery, leading to a persistent surge in bad matches, which lasts for an entire decade, and is remindful of the increase in part-time workers for economic reasons observed in the data. Replacing the quit rate with the employment-to-employment rate from the U.S. Census leads to even more striking results because this rate exhibits a slower reversion to its pre-Great Recession level relative to the quit rate. As a result the model predicts an even more persistent response of bad jobs and weaker pickup of inflation during the post-Great Recession recovery.

The relative size of good and bad matches in the model, which captures the state of the job ladder and reflects the degree of cyclical labor misallocation, is closely tied to the dynamics

2

of inflation. An important property of the model is that the unemployment rate and the EE transition rate are sufficient statistics for the state of the job ladder and labor productivity. Thus, our model allows us to derive implied series for good and bad matches, which are not directly available in the data. We also show that the persistent fall in EE flows helps explain, to some extent, the slowdown in the U.S. labor productivity growth observed during the recovery. A conservative calibration implies a peak fall in productivity of one-percentage-point, relative to trend.

The reason why the model is able to generate such a persistent rise in bad jobs and fall in inflation is because EE flow rates are very small, in the order of 2% per month. As a result, it takes a very long time for an increase in bad matches to be reabsorbed, and for the economy to revert to steady-state. The persistence of this propagation mechanism is striking, particularly when compared to the lack of persistence that characterizes the standard search and matching model without on the job search. In that framework the speed at which the economy reverts to steady state is governed by the UE rate, which is higher than the EE rate by an order of magnitude. So in that baseline model, an increase in unemployment is reabsorbed very quickly.

This paper deliberately focuses on the post-Great Recession inflation dynamics, for two main reasons. First, the failure of the job-to-job flow rate to recover after this recession is an interesting anomaly that makes the most recent years a natural laboratory to test the theory of the job ladder. Second, we show that in this period the traditional New Keynesian models most spectacularly fail to explain why inflation has remained so moderate and for so long in spite of tight labor market conditions and low productivity growth. We argue that this prediction arises precisely because these theories only focus on conventional measures of labor market slack and thus fail to take into account the extensive and persistent labor market misallocation generated by the Great Recession.

When we extend the empirical evaluation of the model to the period that precedes the Great Recession, we find that the model explains fairly well the dynamics of inflation in the previous two cycles too. At the beginning of last three recoveries, the number of unemployed workers is relatively large. Our model with the job ladder predicts that the large majority of the unemployed will be first offered a bad job. Since poaching workers in bad matches is not inflationary, inflation does not increase at the beginning of a recovery. This is strikingly different from the prediction of standard New Keynesian DSGE models in which inflation is driven by the flows from unemployment to employment, which raise hours worked. This prediction of standard models is at odds with the data that suggest a fairly persistent fall in the inflation gap at the beginning of the last three economic recoveries. As the recoveries consolidate, the job ladder transforms the bad matches into good matches and inflation rises because poaching workers away from good jobs is costly for employers.

While in traditional New Keynesian models unemployment to employment flows are the
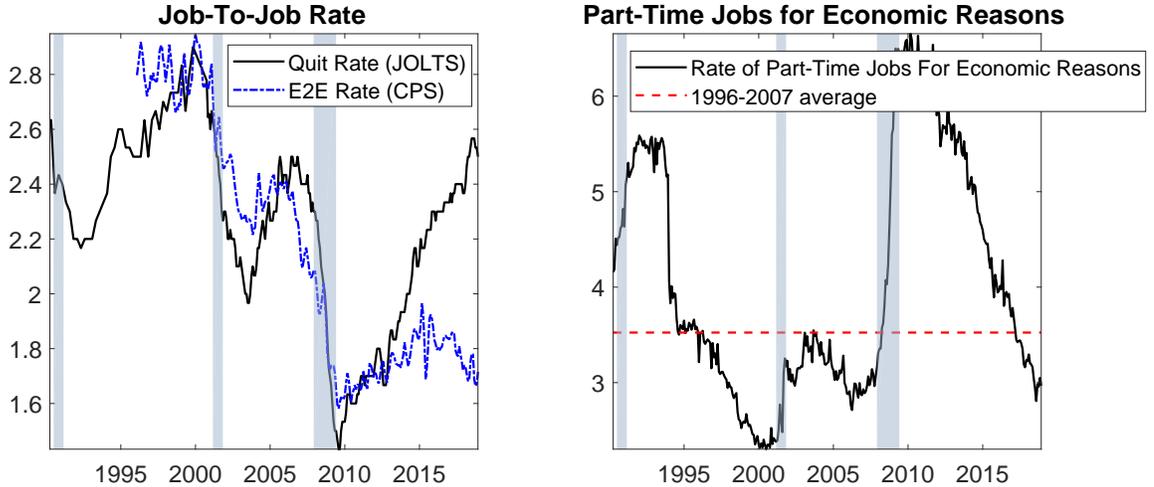
Figure 2: The job ladder. The left plot: The quit rate and employment-to employment (EE) flow rate. Frequency: Monthly. Sources: BLS (JOLTS dataset) and the Census Bureau (Current Population Survey), respectively. The right graph: Part-time workers for economic reasons as a fraction of civilian employment (over sixteen years old), black line, and its average value from February 1996 through November 2007. Frequency: Monthly. Sources: BLS (JOLTS dataset). The shaded areas denote NBER recessions.

leading driving factor of inflation, in our model with the job ladder the emphasis is on the EE flows. Indeed, our model predicts that the EE rate is a very reliable predictor of future inflationary pressures. This result rationalizes why central banks look at the health of the labor market as a multidimensional object and supplement the data on employment and unemployment with measures of gross job flows, such as the quit rate (Yellen 2013).

Furthermore, the model explains the lack of major inflationary spikes in the last 25 years with the marked decline of the job-to-job flow rate since the end of the 2001 recession. This decline is evident in the EE rate measured by the CPS, but it can also be noticed in the quit rate from JOLTS. The average quit rate has fallen by roughly 25 basis points in each of the last three NBER cycles (peak to peak): from July 1990 through February 2001 the quit rate averaged at 2.53%; from March 2001 through November 2007, this rate was 2.23% on average; and from December 2007 through December 2018, the average quit rate was 2.00%. This rate controls the pace at which the job ladder in our model transforms bad matches into good ones, which requires poaching firms to rise wages to attract workers. Since these pay rises are not backed by any gain in labor productivity, inflation rises. Therefore, the lower the job-to-job rate, the longer it takes for the economy to create good jobs, and the slower the recovery of inflation.

A convenient feature of the model, as shown by Moscarini and Postel-Vinay (2019), is that it can be solved – and inflation dynamics can be derived – by characterizing expected surpluses of good and bad matches, without explicitly characterizing the wage distribution. This property stems from the assumption that wages are bargained following the sequential auction protocol developed by Postel-Vinay and Robin (2002). So while discussing the dynamics of wages is

helpful to convey intuition for the mechanism at work, it is not necessary to solve for inflation dynamics. Another interesting feature of the model is that it allows for exogenous shocks to the probability that employed workers search on the job. This is a simple assumption, which allows the model to capture the joint dynamics of flow rates from employment to employment and from unemployment to employment (UE). Our model takes the paths for these two flow rates as given, and subject to those, it asks whether a general equilibrium model with the job ladder can account for the recent dynamics of key U.S. macroeconomic aggregates. Note indeed that while UE flows are not directly observed, they are in fact implied by the unemployment rate, in a way that tracks the data pretty closely.

The paper builds on the empirical literature that has investigated the role of search and matching frictions in New Keynesian models. Key contributions include Gertler, Sala, and Trigari (2008), Krause, Lopez-Salido, and Lubik (2008), Ravenna and Walsh (2008), Sala, Soderstrom, and Trigari (2013) and Christiano, Eichenbaum, and Trabandt (2016). We deviate from these studies by considering the role of on the job search, inspired by 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, wages 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 (2019) introduce on-the-job search and heterogenous jobs in a New-Keynesian framework to investigate the propagation of shocks and the role of EE flows as a predictor of inflation dynamics. Unlike that paper, our emphasis is on the empirical evaluation of this class of general equilibrium models. As in Moscarini and Postel-Vinay (2019), 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 variation in the probability that employed workers receive outside wage offers, as a means of accounting for the joint dynamics of UE and EE rates, 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 to the

probability that employed workers search on the job in order to explain the data.

Our contribution relative to the works by Moscarini and Postel-Vinay is to evaluate the properties of New Keynesian models with a job ladder using state-of-the-art time series techniques. We believe that this is the first paper that carries out this empirical analysis, which leads us to highlight 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. This can be thought of as an exceptionally protracted sullying effect of recessions, in the spirit of Barlevy (2002), one that extends to the ensuing recovery.

Key to these two findings is the disruption of the ladder-based mechanism of labor reallocation that has characterized the post-Great Recession 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. These findings reveal an anomaly in the cyclical patter of the job ladder in the U.S., whereby in economic recoveries employment typically tends to flow out of small, less productive firms and into larger, more productive ones (see Moscarini and Postel-Vinay 2009, 2012). This normal pattern did not carry over to the post-Great Recession recovery. with important implications for the dynamics of productivity and inflation.

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

$$\pi_t = \kappa \varphi_t + \beta E \pi_{t+1}, \tag{1}$$

where φ_t denotes the real marginal cost, π_t is inflation and β is the discount factor. 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
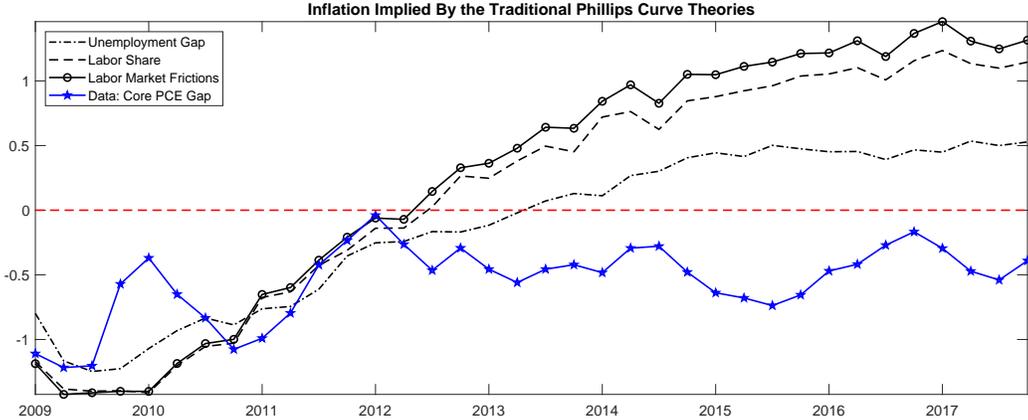
Figure 3: Core PCE inflation gap from 2009Q1 through 2017Q3 and inflation dynamics predicted using three traditional theories of inflation.

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 as well as UE flow rates (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 to facilitate comparability with the model presented in the next Section. We discuss the extension to the case of price indexation in Appendix D and show that it does not matter for the results.

Solving eq.(1) forward, expected inflation can be expressed as the sum of the current and 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 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, and the federal funds rate (FFR).[1] The data sample covers the period 1958q4 through 2017Q4. The VAR model is estimated using these nine observables, and is then used to forecast the future expected path of real marginal costs; that is, the 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. We assume in estimation a discount factor β to equal 0.99 and a slope of the Phillips curve κ equal to 0.005, so as to fit inflation at the beginning of the post-Great Recession recovery. While the slope of the Phillips curve affects the magnitude of inflation predicted by

---

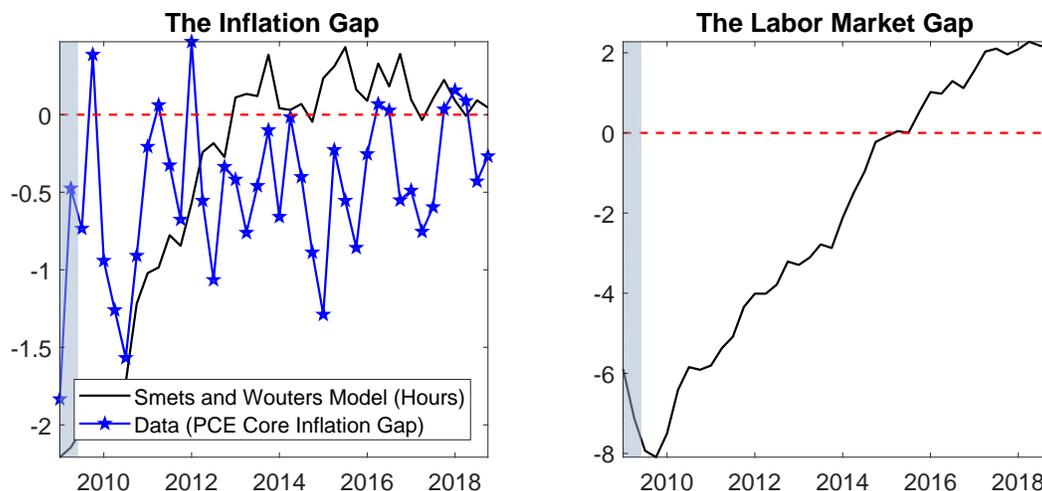[1]Details on how these series are constructed is in Appendix B.

7

Figure 4: Left plot: the blue solid line: Inflation predicted by the model of Smets and Wouters (2007) conditional on observing the series of hours plotted on the right graph. The black starred line: the difference between the annualized quarter-to-quarter PCE core inflation rate and the ten-year-ahead PCE core inflation expectations based on the Survey of the Professional Forecasters. Right plot: Hours worked detrended using their eight-year moving average.

the three theories, it does not affect significantly the point in time when inflation rises above target. See Appendix C.

Using a VAR model to approximate the private sector's forecasts is desirable for three 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. Model misspecifications 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 assuming that agents form their expectations according to a VAR model does not necessarily imply a deviation from rationality. Third, large Bayesian VAR models are well-known to provide relatively accurate macroeconomic forecasts.

Figure 3 shows that all the traditional series suggest that inflation should have been above steady state starting around 2013 or 2014. None of these theories is able to account for why inflation has been so low for so many years after the Great Recession because all the three proxies for marginal costs improved quickly in the first years of the economic recovery. Consequently, the VAR model's forecasts of future marginal costs go up at a relatively early stage of the recovery, which leads the three Phillips curve to predict inflation above target.

2.1 A State-of-the-Art Dynamic General Equilibrium Model

We also evaluate the ability of a leading empirical general equilibrium model to reconcile labor market and inflation dynamics in the post-Great Recession recovery. We use the popular model introduced by Smets and Wouters (2007a) to perform this exercise. This is a model with many real and nominal frictions and a large array of shocks and is well known to fit the U.S. macro series well. Smets and Wouters conduct Bayesian estimation of the parameters of their model using seven observables: consumption growth, investment growth, GDP growth, hours (detrended for the labor force participation), inflation, real wage, and the federal funds rate. Their sample period goes from 1966Q1 through 2004Q4. We extend their data set to 2018Q4 and detrend the series of hours using a eight-year moving average. We make the latter change because the series of hours exhibited a significant downward shift after the onset of the Great Recession and has never attained its pre-recession level again.

We use the extended data set to estimate the model. Then the same data set is used to filter the state variables of the estimated model from the first quarter of 1966 through the fourth quarter of 2008. For the subsequent periods (2009Q1-2018Q4), we filter the state variables of the estimated model using only the series of hours in order to obtain the model's predictions about inflation giving the model only labor market data. Recall the emphasis of this paper is on the apparently waning link between the labor market and inflation.

Based on the series of hours, the Smets and Wouters' model predicts that inflation is above target already in 2012. See the black solid line in the left Figure 4. The plot also reports the inflation gap in the data and computed by taking the difference between the annualized quarter-to-quarter PCE core inflation rate and the ten-year-ahead PCE core inflation expectations based on the Survey of the Professional Forecasters (blue starred line). The inflation gap based on PCE core inflation rate remains persistently below zero whereas the Smets and Wouters' model predicts that inflation move above its long-run expected level as early as 2012. The solid blue line in the right plot shows the series of hours detrended using a eight-year moving average, which we use to simulate the Smets and Wouters model. This series suggests that the labor market becomes tight (positive labor market gap) in 2015 after the labor market boom in 2014.

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 of cyclical labor misallocation implied by the state of 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.

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 any employed worker can search in a given period with probability s_t, which is assumed to follow an AR(1) stochastic process.

We distinguish three types of firms: service producers, price setters and packers. 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. The output of the match is sold to price setting firms in a competitive market at the real price φ_t, and transformed into a differentiated product. Specifically, one unit of the service is transformed by firm i into one unit of a differentiated good $y_t(i)$. These firms set the price of their goods subject to Calvo price rigidities. Finally, the last layer of firms packages all differentiated 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:

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

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 unemployed. 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:

$$n_t = 1 - u_t. \tag{2}$$

Aggregate unemployment at the beginning of a period is given by

$$u_{0,t} = u_{t-1} + \delta n_{t-1}, \tag{3}$$

while aggregate unemployment at the end of a period is

$$u_t = u_{0,t} \left[1 - \phi_t\right]. \tag{4}$$

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 government bond holdings B_t, profits of firms that produce differentiated goods, D_t^P, profits of service firms D_t^S, wages from the workers who are employed and transfers from the government T_t. 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

11

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:

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

subject to the budget constraint,

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

the stochastic process for the employment status,

$$
\begin{aligned}
prob \{e_{t+1}(j) = 1 \mid e_t(j)\} &= e_t(j)(1 - \delta) + [1 - e_t(j)] \phi(\theta_{t+1}) & (5) \\
prob \{e_{t+1}(j) = 0 \mid e_t(j)\} &= 1 - prob \{e_{t+1}(j) = 1 \mid e_t(j)\},
\end{aligned}
$$

and the stochastic process for equilibrium wages $w_t(j)$, to be determined below. Equation (5) implies 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.(5) 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:

$$1 = (1 + R_t) \beta E_t \frac{\lambda_{t+1}}{\lambda_t} \Pi_{t+1}, \tag{6}$$

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:

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

12

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:

$$\lambda_t V_{e,t}^j \left(w_t\left(j\right), y_t\left(j\right)\right) = \lambda_t \frac{w_t\left(j\right)}{P_t}$$
$$+ \beta E_t \lambda_{t+1} \left[\delta V_{u,t+1}^j + (1-\delta) V_{e,t+1}^j \left(w_{t+1}\left(j\right), y_{t+1}\left(j\right)\right)\right],$$

where $E_t V_{e,t+1}\left[w_{t+1}\left(j\right), y_{t+1}\left(j\right)\right]$ 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 $y_t\left(j\right)$ at the promised wage $w_t\left(j\right)$. 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. $\lambda_t V_{e,t}\left(w_t\left(j\right), y\left(j\right)\right) = b + \beta E_t \lambda_{t+1} V_{u,t+1}^j$ independently of productivity. It follows that

$$V_{u,t}^j = \frac{b}{\lambda_t} + \beta E_t \frac{\lambda_{t+1}}{\lambda_t} V_{u,t+1}^j = V_{u,t}. \tag{7}$$

Let $V_{e,t}^*\left(y\right)$ 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 $\{w_s^*\left(y\right)\}_{s=t}^\infty$ denote the state-contingent contract that delivers $V_{e,t}^*(y) \equiv V_{e,t}\left(w_t^*, y\right)$. By promising to pay the contract $\{w_s^*\left(y\right)\}_{s=t}^\infty$, 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 $\{w_s\left(y\right)\}_{s=t}^\infty$. 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. $V_{e,t}^*\left(y'\right) < V_{e,t}\left(w_t, y\right)$; 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 $V_{e,t}\left(w_t, y\right)$.

2. $V_{e,t}\left(w_t, y\right) \leq V_{e,t}^*\left(y'\right) < V_{e,t}^*\left(y\right)$; 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 $V_{e,t}^*\left(y'\right)$.

3. $V_{e,t}^*\left(y\right) \leq V_{e,t}^*\left(y'\right)$; in this case the poaching firm has a willingness to pay that is no less than the incumbent's. If this condition holds with strict inequality, the current match is terminated and the worker is poached at the maximum value of the contract that

13

the incumbent is willing to pay. If instead the worker is poached by a firm with equal productivity, it is assumed that job switching takes place with probability υ. In either case, the continuation value of the contract obtained by the worker is $V_{e,t}^{*}(y)$.

3.4 Service sector firms: 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 sunk, fixed cost of hiring c^{f}. 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 surplus after the match is formed. The free entry conditions is:

$$
\begin{aligned}
c^{f} + \frac{c}{\varpi_t} &= \frac{u_{0,t}}{u_{0,t} + s_t(1 - u_{0,t})} \left[\xi_b S_t(y_b) + \xi_g S_t(y_g) \right] \\
&\quad + \frac{s_t(1 - u_{0,t})}{u_{0,t} + s_t(1 - u_{0,t})} \left\{ \xi_g \frac{l_{b,t}^0}{1 - u_{0,t}} \left[S_t(y_g) - S_t(y_b) \right] \right\},
\end{aligned}
\tag{8}
$$

where $l_{b,t}^0$ denotes the measure of workers who, at the beginning of a period, are employed in low quality matches. The expected return from forming a match, on the right hand side of the equation, is expressed as a weighted average of the surplus expected from meeting with an unemployed worker or with a worker employed and searching on the job, with weights given by their relative measure in the pool of job seekers. The first term on the R.H.S. implies that if a vacancy meets with an unemployed worker it will extract the entire surplus of the match, whose value will depend on the realized quality of the match. The second term instead implies that when meeting with an employed worker, a vacant job can extract a strictly positive surplus only if the prospective match quality is good, and match quality with the incumbent is bad. In this case, the firm will poach the worker offering a wage contract with present value $S_t(y_b)$, and keeping the remaining surplus $S_t(y_g) - S_t(y_b)$ for itself.

Substituting out for the surplus functions in the above equations requires some steps. Start by considering 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) in the previous subsection, receives a contract from the new firm which is also worth $V_{e,t}(w', y') = V_{e,t}^*(y)$. 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 price setters. Subtracting (7) from the above equation yields:

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

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:

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

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

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

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 derived from selling one unit of the service until separation. From the point of view of a price setting 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.

Using eq.(9) we can now substitute for the surplus functions and rearrange to rewrite the free entry condition (8) as:

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

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 equi-

librium, 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 nominal prices of the service, $\varphi_t P_t$, fall. With perfectly flexible prices, all of the adjustment in nominal prices takes place through P_t. With sticky prices instead the current *real* prices of the service φ_t also 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:

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

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

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, $\left[1 - s_t\phi\left(\theta_t\right)\xi_g\right] l_{b,t}^0$, plus the new hires from the unemployment pool who turned out to draw a low quality match, $\phi\left(\theta_t\right)\xi_b u_{0,t}$. 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 $\phi\left(\theta_t\right)$ 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 $l_{g,t}^0$, plus the job-to-job inflows from low quality matches $s_t\phi\left(\theta_t\right)\xi_g l_{b,t}^0$, and the new hires from unemployment $\phi\left(\theta_t\right)\xi_g u_{0,t}$. Using

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

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

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

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

16

3.6 Price Setters and Packers

The final layer of firms packages differentiated products into a homogeneous good, which is sold in perfect competition. As in Smets and Wouters (2007b), the consumption bundle is given by the general Kimball (1995) aggregator:

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

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 $G\left(q_t\left(i\right)/Q_t\right)$:

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

where $\omega^k = \frac{\chi(1+\varkappa)}{1+\varkappa\chi}$, $\varkappa \leq 0$ is a parameter that governs the degree of curvature of the demand curve for the differentiated goods, and χ captures the gross markup.

The maximization problem reads:

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

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

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

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).

Price setters buy the (homogeneous) output produced by the service firms in a competitive market at the real price φ_t and sell it to the packers 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 price setting 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 price setting firms 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, aggregate output from price setters must equal aggregate demand from the packers:

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

where we have made use of the demand function in eq.(17). Substituting the profits of all firms into the household's budget constraint and assuming that government bonds are in zero net supply yields the aggregate resource constraint:

$$C_t + v_t c + c^f \phi_t \left[u_{0,t} + s_t (1 - u_{0,t}) \right] = Q_t.$$

3.8 Taylor rule

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

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

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-

Calibration

Parameters	Description	Value	Target/source
Parameters that affect the steady-state			
β	Discount factor	0.9980	Real interest rate 2% p.a.
ϕ_0	Scale parameter matching fn	0.3200	Job finding rate 32%
δ	Job separation rate	0.0200	Unemployment rate 5.5%
y_b	Productivity bad matches	1.0000	Normalization
y_g	Productivity good matches	1.1250	
υ	Prob. of job switching if indifferent	0.5000	
b	Utility of leisure	0.8427	Calibrated
c	Flow cost of vacancy	0.0134	Calibrated
c^f	Fixed cost of hiring	0.5376	Calibrated
s	On the job search rate	0.2393	Calibrated
ξ_g	Probability draw good match	0.2800	Calibrated
Parameters that do not affect the steady-state			
χ	Mark-up parameter	1.2000	20% mark-up
$\varkappa$	Scale param. Kimball aggregator	10.0000	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	0.6500	Conventional
ϕ_π	Taylor rule response to inflation	1.8000	Conventional
ϕ_y	Taylor rule response to output	0.2500	Conventional
ψ	Elasticity of matching function	0.5000	Moscarini and Postel-Vinay (2018)
ρ_λ	Autocorrel. preference shock	0.9932	MLE estimation
$100\sigma_\lambda$	St. dev. preference shock	0.0218	MLE estimation
ρ_S	Autocorrel. job search rate	0.8818	MLE estimation
$100\sigma_S$	St. dev. job search rate	2.5581	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	0.0244	Pre-Great Recession average quit rate
θ	Labor market tightness	1.000	Normalization
$\frac{l_g^0}{l_g^0+l_b^0}$	Employment share in good jobs	0.6700	Employment share at top 10% firms
$\frac{(Q-C)/N}{\varphi}$	Hiring costs over wages	1.500	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	
$\frac{Q-C}{Q}$	Hiring costs over output	0.5000	
u_0	unemployment rate	0.0550	

Table 1: Model parameters and implied steady-state values

Douglas matching function $M_t = \phi_0 \left[u_{0,t} + s_t \left(1 - u_{0,t}\right)\right]^{1-\psi} v_t^{\psi}$, 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 $\theta_t = v_t / \left[u_{0,t} + s_t \left(1 - u_{0,t}\right)\right]$.

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 regard this productivity differential as conservative, in the light of values that have been assigned in the calibration of other comparable models with on-the-job search. For instance, in Gertler, Huckfeldt, and Trigari (2016) the productivity of bad matches is 35% below that of good matches. We have checked that the dynamics of inflation predicted by the model are robust to the precise parameterization of y_g. Moreover, we noticed 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 when solving the model in the empirical exercise. Finally, we set the probability that workers will accept an equally valuable outside offer to be $\upsilon = 0.5$.

This leaves us with five parameters to calibrate, the parameter governing the utility of leisure 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 value of expected hiring costs, including both the variable and the fixed cost component, equal to six weeks of wages;[2] *(ii)* a fraction of good jobs in steady state equal to 67%, which is the share of employment for the top 10% US firms by employment size in year 2000 (the reasons why this is a reasonable target for calibration of this parameter are provided below) ; *(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;[3] *(v)* a monthly job-to-job transition rate of 2.44%, which is the average quit rate measured in the pre-Great Recession sample (April 1990 through November 2007). We have checked that the value of b implied by

---

[2]The wage is measured as the price of the labor service φ.

[3]Silva and Toledo (2009) indicate in Table 1, p.80, that the average pre-match recruiting cost costs is 105.1\$, while the average post-match training cost amounts to 1,359.4\$.

the calibration is consistent with a positive surplus for low quality matches both in steady state and when solving the model in the empirical exercise.

The calibration of the probability of a good match ξ_g (conditional on receiving a job offer) relies on the empirical strategy in Moscarini and Postel-Vinay (2016a), who exploit the notorious correlation between firm size and productivity by assuming that employed workers climb the ladder when moving to larger firms. Other evidence supporting this target includes the share of the U.S. full-time private industry workers who are offered access to medical care benefits. This share was 69% in March 2018 according to the BLS *Employee Benefits Survey*. This number is strikingly close to our target of 67%. It is conceivable that employers are willing to bear the higher costs of offering medical care benefits only for their most productive employees. Indeed, this share correlates with the wage distribution: these benefits are offered to 86 percent of top 25 percent of U.S. wages earners, 75 percent of the third 25 percentile, 61 percent of the second 25 percentile, and only 26 percent of the lowest 25 percentile. It is worthy emphasizing that in the model workers in a good match that have received a poaching offer from a firm offering a match of good quality are the top earners. Furthermore, to the extent that wage differentials reflect productivity differentials, this statistic can be appropriately used as a target for the fraction of high-productivity jobs in steady state. The share of the U.S. full-time private industry workers who are offered access to retirements benefits (defined benefits plus defined contributions) was 68% in March 2018 according to the BLS *Employee Benefits Survey*. This statistic has very similar properties and magnitude to the one based on employer's size and medical care benefits.

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 (2007b). 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 (2018), 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

---

[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

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.

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

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

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 EE 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 also includes workers who quit their job to become unemployed or to leave the labor force, the vast majority of the workers who quit switch jobs. Elsby, Hobijn, and Sahin (2010) document that 86% of workers 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 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 filter can be applied. Log-linearized variables are denoted with $\hat{\,}$. Rates and shares are linearized and denoted by $\tilde{\,}$.

---

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).

Identification of Key Labor Market Variables in the Data It is important to notice that for a given initial value of 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 for price setting firms, $\mathcal{W}_t$.

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 (2) and (3) and linearizing.

Endowed with the end of period unemployment rate $\tilde{u}_t$, we can linearize equation (4) 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 $\tilde{s}_t$:

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

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 $\tilde{s}_t$. 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 matches in the next period $t+1$, using the linearized laws of motion for low and high quality matches in (14) and (15), which read

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

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

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

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

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 use our model with the job ladder to assess the effects of cyclical 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 have been highly and persistently misallocated, as suggested by Figure 1. To this end, we use the Kalman filter to estimate the model's variables.[5] The filter returns the real-time model's predicted value of any variable conditional on the observations we will use in the exercise. We will use two sets of observables. Both include the civilian unemployment rate as a measure of the rate of unemployment in the model. But while one data set includes the quit 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. While the latter measure is available from February 1996, the former measure goes back to April 1990 using the methodology developed by Davis, Faberman, and Haltiwanger (2012). From January 2000 and on the series of quit rate is based on JOLTS. When we run the filter with the EE rate as observable variable, we initialize the states of the model using the filtered estimate of the states implied by the quit rate.

As shown in the previous section, observing the unemployment rate and one series for EE 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 property of the model 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 estimated dynamics of bad jobs.[7]

Figure 5 show the estimated dynamics of a set of model's variables from April 1990 (February 1996 when the EE rate from CPS is used) through December 2018. The blue solid lines mark the simulation based on the quit rate in JOLTS and the black dashed-dotted lines are the estimates based on the E2E rate measured by the CPS. The upper plots show the dynamics of the job search intensity of the employed s_t and two traditional labor market variables -the unemployment rate and the job finding rate- over the sample period. All these measures suggest that the U.S. labor market has become tight starting from 2015 and on. The job finding rate is not observed but is exactly identified by the observed unemployment rate as we already noticed when we introduced equation (21). This feature implies an almost perfect negative correlation

[5]We could have used the Kalman smoother (two-sided filter) for this exercise. We can show that the model of the job ladder would not materially change its predictions about inflation dynamics over the sample period considered in this study. We work with the filter estimates, because their construction is more intuitive. Indeed, for the model's labor market variables (e.g., bad jobs) its outcome is very similar to what one would obtain from simulating the model using the series of unemployment rate and the job-to-job flow rate. Furthermore, the filter returns the real-time estimates of the states of the ladder which are actually what policymakers can observe to inform their choices.

[6]The other side of the coin of this implication is that the econometrician's uncertainty about the model's predictions on 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.
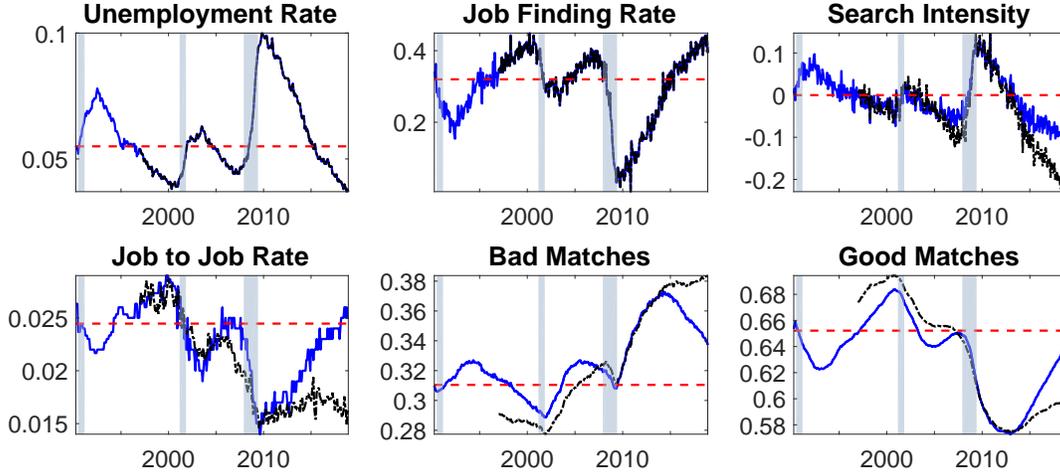
Figure 5: Filtered estimates of labor market variables obtained conditional on the calibrated model, the observed unemployment rate, and a proxy for E2E flows. We consider two proxies: the quit rate (blue solid line), which is available from April 1990 and the EE rate from the CPS (the black dashed-dotted line), which is availble from February 1996.

of these two series in Figure 5.[8]

In the bottom panels of this figure, we report the observed job to job flow rate. It should be noticed that up to 2005 the quit rate (the blue solid line) is a very good proxy of the EE rate (the black dashed-dotted line) as most of workers quit their job just to start a new one. In the most recent years we observe a fairly different dynamics of these two rates, with the quit rate slowly reverting back to its precession level

Since the fractions of bad and good matches are fairly slow moving, equation (22), the swings in the job search intensity of the employed, s_t, appears to be mainly driven by the joint dynamics of the observed job-to-job flow rate and the job finding rate. For given job finding rate, a lower job to job rate implies less job search intensity. As a result, the lack of recovery in the EE rate measured by the CPS pushes the search intensity to very low value. As we will show, such a low value severely impairs the working of the job ladder, delaying the rate at which bad jobs are transformed into good jobs.

While these traditional labor market measures in the upper panel of Figure 5 suggest that the U.S. labor market has become quite tight in recent years, our model predicts that the cross-sectional slack has been much more persistent: bad matches have increased since the end of the last recession, remaining at a high level throughout the recovery. As the job finding rate plummeted to historically low levels at the end of the recession, matches that were destroyed in every period failed to be reabsorbed, which led to a sharp increase in the unemployment rate. Consequently, an increasing fraction of displaced workers started climbing the ladder

---

[8]After the Great Recession, the model's implied series for the job finding rate drops by less and recovers more quickly than in the data. Hence, using the time series of job finding rates as an observable in lieu of unemployment rates, would slow down the working of the ladder even more.
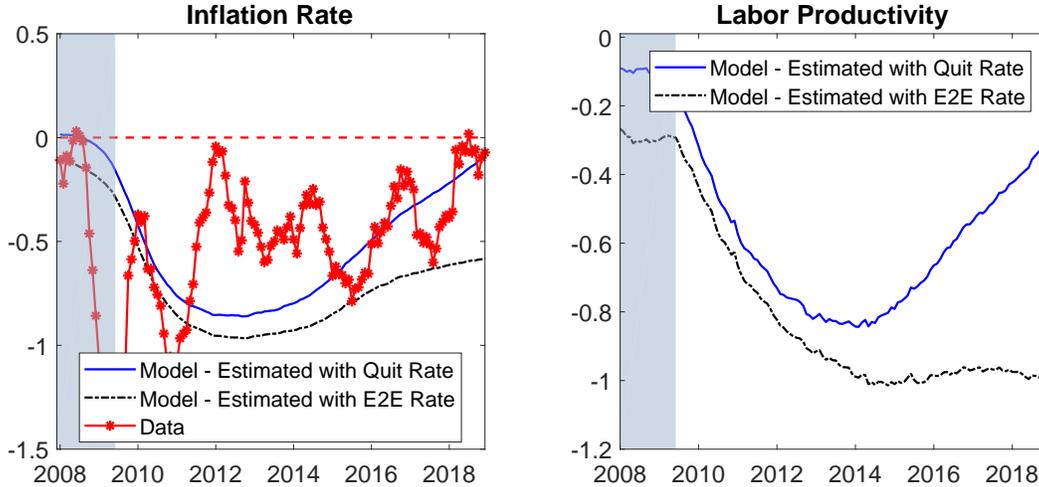
Figure 6: Right Graph: PCE core inflation gap (red line with dot markers) and model generated paths for inflation (in deviation from its steady state value) obtained when observing the JOLTS quit rate (solid blue line) and the CPS EE rate (dashed black line), respectively. The core PCE inflation gap is obtained by subtracting the ten-year PCE inflation expectations measured by the Survey of Professional Forecasters from the year-over-year core PCE inflation rate. All rates are 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.

anew, raising the fraction of bad matches, $l_{b,t}^0$. Because the job to job rate was low by historical standards already at the onset of the financial crisis, job matches failed to be reabsorbed quickly, inducing hump shaped dynamics in bad jobs and depressing the fraction of good jobs for an entire decade. Unlike the number of bad jobs, the fraction of good jobs declined in the first few years of the post-Great Recession recovery and started to rise only in 2014. 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 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 EE flow rate remains subdued and never fully recovers, unlike in the simulation based on the quit rate. This missed recovery has the striking effect of blocking entirely the working of the job ladder. As the EE 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.

The behavior of bad matches is reflected in a mirror image response of the present discounted value of the labor service 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. Because price setters face nominal rigidities, 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 6 shows the effects of labor misallocation accumulated during and after the Great Recession on inflation. The blue line (dashed-dot black line) represents the dynamics of inflation implied by the JOLTS quit rate (CPS EE rate). The red line with dotted markers denotes the observed core PCE inflation gap, which is obtained by subtracting the ten-year PCE inflation expectations measured by the Survey of Professional Forecasters from the year-over-year core PCE inflation rate.[9] In both cases, the labor misallocation produced by the Great Recession has contributed to keep inflation below its long-term level (red dashed line) for several years. Even though the model is stylized and endowed with only two shocks, it seems to capture well the cyclical dynamics of PCE core inflation rate during the post-Great Recession recovery. This finding suggests that the cross-section slack captured by the job ladder is important to understand the dynamics of inflation in the post-Great Recession recovery. Crucially, the model predicts subdued inflation even though the rate of unemployment and the job finding rate suggest that the U.S. labor market has been very tight, as shown in Figure 5. This result suggests that focusing only on aggregate measures of slack, such the unemployment rate or the job finding rate, to understand inflation may be misleading.

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 6 shows the model's predicted percentage deviations of average labor productivity from its long-run value during the Great Recession and the following recovery. Labor misallocations persistently lower average labor productivity at the beginning of the recovery, as bad matches increase and good matches fall. Depending on the job-to-job flow rate we observe, i.e. the quit rate or the EE rate from Census, we obtain fairly different patterns for labor productivity in the second half of the recovery. Because the EE rate implies a more persistent rise in bad matches, productivity remains stuck at about one percentage point below steady state, whereas in the case where the quit rate is used, labor productivity shows some reversion to the steady state in the most recent years. Given our conservative calibration of productivity differentials $(y_g - y_b)$, we regard these results as lower bound estimates of the contribution of cyclical misallocation to the dynamics of productivity over the post-Great Recession recovery.

---

[9]The PCE inflation rate expected by the professional forecasters ten year from now is extremely stable around two percent. The series is shown in the right plot of Figure 1.
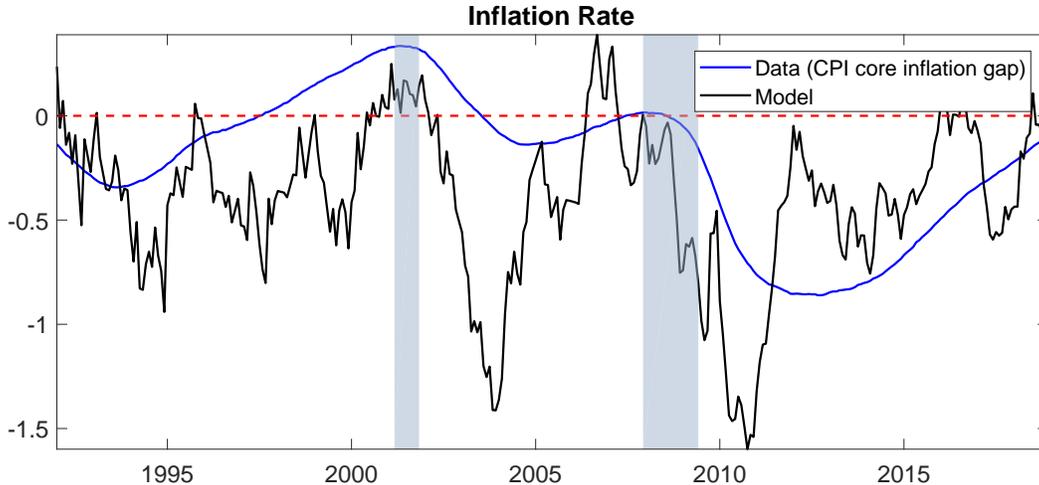
Figure 7: The observed CPI core inflation gap (the black solid line) and model generated paths for inflation (in deviation from its steady state value) obtained when observing the JOLTS quit rate (the blue dashed-dotted line). The core CPI inflation gap is obtained by subtracting the ten-year CPI inflation expectations measured by the Survey of Professional Forecasters from the year-over-year core CPI inflation rate. All rates are in percent.

4.4 The Theory of the Job Ladder Before the Great Recession

This paper primarily focuses on inflation dynamics after the Great Recession because of two reasons. First, it is during the post-Great Recession recovery that inflation dynamics become more spectacularly decoupled from standard measures of marginal costs. As shown in Section 2, the dynamics of inflation in that period represent a conundrum for traditional models based on the New Keynesian Phillips curve. Second, both the quit rate and the EE rate start falling even before the Great Recession with the former recovering slowly and the latter remaining stuck at a lower level. Our model with the job ladder allows us to understand the implications of lower job-to-job flows for price dynamics.

That said, it is interesting to evaluate the empirical performance of the theory of the job ladder before the Great Recession. While the CPS data on the EE rate are available only from February 1996, Davis, Faberman, and Haltiwanger (2012) extend the series of the U.S. quit rate back to April 1990.[10] Thus we can use the series of the quit rate for this exercise. We compare the year-over-year core CPI inflation gap with the model's year-over-year inflation rate, which is estimated using the Kalman filter. We compute the core CPI inflation gap by taking the difference between the monthly CPI core inflation gap and the ten-year CPI inflation expectations measured by the Survey of Professional Forecasters. These expectations are available starting from the first quarter of 1992. We do not use PCE core inflation gap, as done in Section 4.3, because data from the Survey of Professional Forecasters for this definition of inflation is only available since 2007.

---

[10]This rate is computed at quarterly frequency in the pre-2001. We make it monthly by dividing the rate by three. From December 2001 and on, we use the quit rate based on JOLTS.
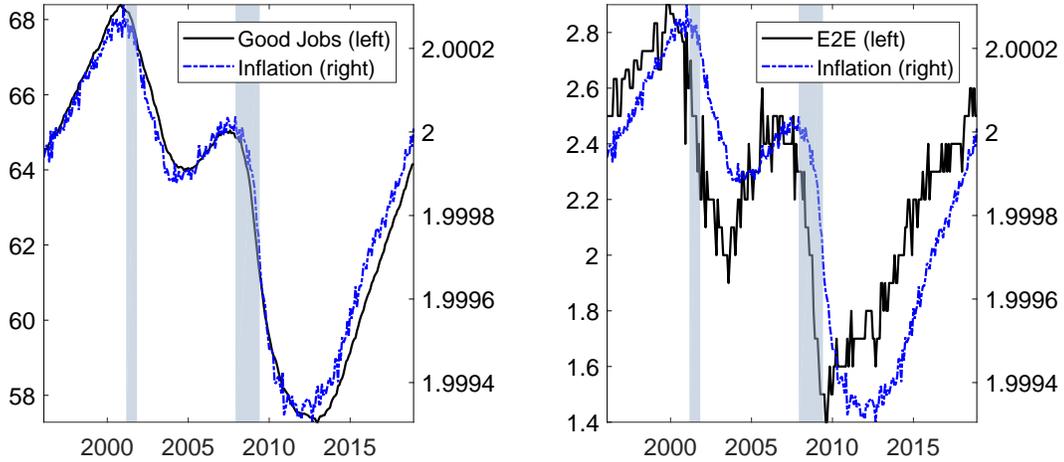
Figure 8: Left Graph: Model-implied inflation and fraction of good matches. Right Graph: Model-implied inflation and the job-to-job flow rate. All series are expressed in percent and the inflation rate is annualized.

The observed and the model-implied inflation gap are shown in Figure 7. In the data, the inflation gap falls deep into the negative territory at the beginning of the three economic expansions and then reaches a trough, after which it gradually increases. The model can account for these cyclical patterns of inflation fairly accurately. Our stylized model with only two shocks cannot explain the dramatic drop in inflation observed at the end of 2003 as well as the erratic high-frequency swings in inflation.

Our model explains the fall in the inflation gap at the beginning of each expansion with a large number of unemployed workers, most of whom get a bad matches at first. Poaching workers in bad matches is not inflationary and as shown in Figure 5 the fraction of good matches falls at the beginning of economic expansions. This is strikingly different from the prediction of standard New Keynesian DSGE models with search and matching frictions in which inflation is driven by the flow from unemployment to employment. This prediction of standard models is at odds with the data that suggest a fairly persistent fall in the inflation gap at the beginning of an economic boom (see the black line in Figure 7). As the economic boom consolidates, the job ladder transforms these bad matches into good matches. When more workers are in good matches, the expected profits for service firms that are searching a worker squeeze. This is because firms can extract positive surpluses only if they meet an unemployed or a bad match (conditional on the original match being good). Lower expected returns to searching for worker causes service firms to leave the industry, which lowers the supply of the homogenous good and raises its price. When the price of the homogeneous goods rises so do the marginal costs of price setters and hence inflation. The right plot of Figure 8 gives us a graphical confirmation of the working of the ladder. This plot compares the model's estimates of inflation and good matches. The two series are almost perfectly correlated, confirming that in our model inflation

rises in tandem with the fraction of good matches.

Another interesting feature of Figure 7 is the asymmetric dynamics of the CPI core inflation gap: this gap lives mostly on the negative territory in the last 25 years. Our model also predicts asymmetric dynamics of good jobs and inflation during the three expansions of our sample. This gap drops considerably at the beginning of the three economic expansions in our sample and recovers only very slowly later. The model with the ladder points to a main culprit for this asymmetric pattern of inflation: the step wise decline of the job-to-job flow rate over our sample period. This decline is evident when one uses the EE rate measured by the CPS, but it is also visible if one looks at the quit rate. The average quit rate has fallen by roughly 25 basis points in each of the last three NBER cycles (peak to peak): from July 1990 through February 2001 the quit rate averaged at 2.53%; from March 2001 through November 2007, this rate was 2.23% on average; and from December 2007 through December 2018, the average quit rate was 2.00%. As we already pointed out, this rate controls the pace at which the job ladder in our model transforms bad matches into good ones. As more and more matches slowly turn into good matches, more poaching firms have to rise their wage bids. However, these wage rises are not matched by any increase in labor productivity, and thus bring about inflationary pressures. Therefore, the lower job-to-job rate, the longer it takes for the economy to create good jobs, and the slower the recovery of inflation.

While in traditional New Keynesian models unemployment to employment flows are the leading driving factor of inflation, in our model with the job ladder the emphasis is on the employment-to-employment flows. The left plot of Figure 8 compares the filtered estimates of the job-to-job flow rate and the model's inflation rate. Note that the model's job-to-job rate is identical to the observed quit rate. The quit rate systematically leads model's inflation. Taking into account the model's good fit of inflation shown in Figure 7, the quit rate can be regarded as a reliable predictor of inflation. This result may explain why central banks look at the health of the labor market as a multidimensional object and supplement the data on employment and unemployment with measures of gross job flows, such as job loss and hiring (Yellen 2013).

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 stylized 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 emphasizes the central role of the inflationary pressures arising from the renegotiation of the wages of the 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 bears some resemblance to that of part-time workers for economic reasons. Seen through the lenses of our model, this time series is informative about inflationary pressure if considered as a proxy for bad jobs. This interpretation differs from the traditional view that sees part-time workers for economic reasons simply as a measure of underemployment or an extended measure of unemployment slack. All in all, 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 the dynamics of job-to-job flows over the cycle, and its implications for the wages of the employed workers.

This paper has taken a first step in evaluating the empirical bite of cyclical misallocation in explaining inflation dynamics. The focus on the post-Great Recession recovery was motivated by the inability of alternative theories to account for the recent behavior of inflation. Moreover, the evidence reviewed in the introduction suggests that a malfunctioning of the job ladder may have made the role of cyclical misallocation particularly relevant over the last decade. The next step in this research agenda will be to embed the job ladder into a richer general equilibrium model to quantitatively assess the relative merit of alternative propagation mechanisms over a wider period of time, where each different channel has a chance to succeed.

References

BARLEVY, G. (2002): "The Sullying Effect of Recessions," *Review of Economic Studies*, 69(1), 65–96.

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.

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.

GERTLER, M., L. SALA, AND A. TRIGARI (2008): "An Estimated Monetary DSGE Model with Unemployment and Staggered Nominal Wage Bargaining," *Journal of Money, Credit and Banking*, 40(8), 1713–1764.

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 (2009): *The Timing of Labor Market Expansions: New Facts and a New Hypothesis*pp. 1–51. University of Chicago Press.

——— (2012): "The Contribution of Large and Small Employers to Job Creation in Times of High and Low Unemployment," *American Economic Review*, 102(6), 2509–2539.

——— (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.

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

——— (2019): "The Job Ladder: Inflation vs Reallocation," 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.

RAVENNA, F., AND C. E. WALSH (2008): "Vacancies, unemployment, and the Phillips curve," *European Economic Review*, 52(8), 1494–1521.

SALA, L., U. SODERSTROM, AND A. TRIGARI (2013): "Structural and Cyclical Forces in the Labor Market during the Great Recession: Cross-Country Evidence," *NBER International Seminar on Macroeconomics*, 9(1), 345–404.

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 (2007a): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *The American Economic Review*, 97(3), 586–606.

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

YELLEN, J. (2013): "Challenges Confronting Monetary Policy," Discussion paper, 2013 National Association for Business Economics Policy Conference, Washington, D.C.

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 $\phi_t = 1 - \left(u_{t+1} - u_{t+1}^s\right)/u_t$, where u_{t+1}^s 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 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 Traditional New Keynesian Theories:Robustness

Figure 9 shows the inflation rate predicted by the three traditional theories introduced in Section 2. In the left graph, we set a slope of the Phillips curve to be equal to 0.001 and in the right graph the slope is set to 0.01. In the main text, the slope is set equal to 0.005. These results suggest that perturbing the slope of Phillips would not change the predictions of the three popular theories of inflation about the time in which the inflation gap becomes positive.
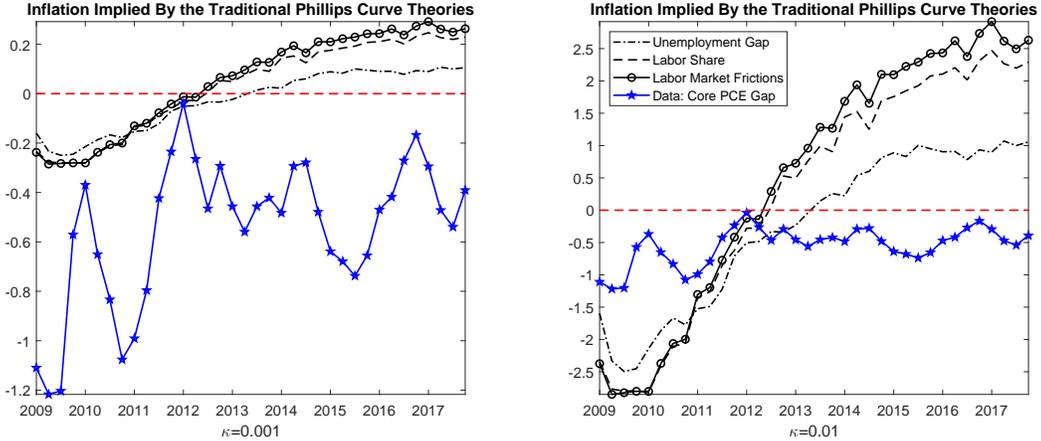
35

Figure 9: Shows the inflation dynamics from 2009Q1 through 2017Q3 using three traditional theories of inflation using a slope of the phillips curve of 0.001 (left graph) and 0.01 (right graph) and the core PCE inflation gap.

D   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. The time when the inflation gap is predicted to become positive would not significantly change if one modifies this parameter within the range of values used by 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 10 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.
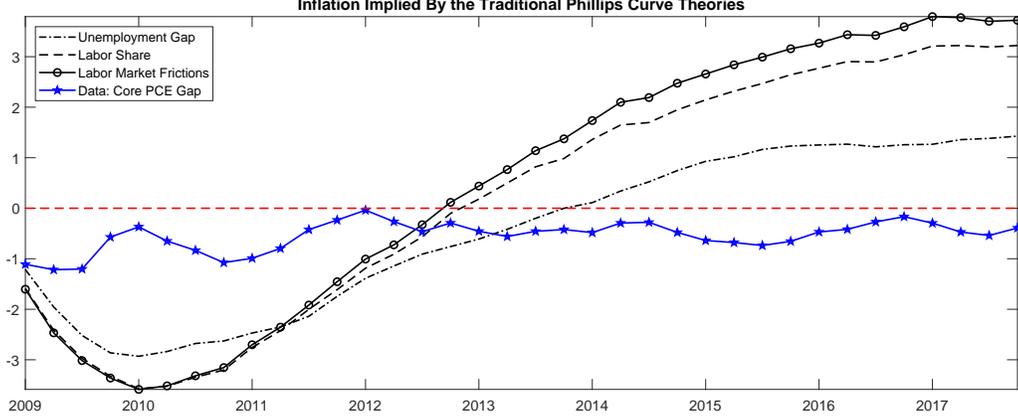
36

Figure 10: 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.

E 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:

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

where W_t denotes the real hourly wage, y_t/n_t is the average product of labor, $c'(v_t)$ is the derivative of the vacancy cost function with respect to vacancies, $q(\theta_t)$ 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:

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

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 $\hat{\beta}_t = 0$.