ECONOMIC RESEARCH
FEDERAL RESERVE BANK OF ST. LOUIS

WORKING PAPER SERIES

Labor Market Shocks and Monetary Policy

Authors	Serdar Birinci, Fatih Karahan, Yusuf Mercan, and Kurt See
Working Paper Number	2022-016J
Revision Date	November 2025
Citable Link	https://doi.org/10.20955/wp.2022.016
Suggested Citation	Birinci, S., Karahan, F., Mercan, Y., See, K., 2025; Labor Market Shocks and Monetary Policy, Federal Reserve Bank of St. Louis Working Paper 2022-016. URL https://doi.org/10.20955/wp.2022.016

Federal Reserve Bank of St. Louis, Research Division, P.O. Box 442, St. Louis, MO 63166

Labor Market Shocks and Monetary Policy*

Serdar Birinci — St. Louis Fed
Fatih Karahan — CBRT
Yusuf Mercan — Kansas City Fed
Kurt See — NUS

November 2025

Abstract

We develop a heterogeneous-agent New Keynesian model with a frictional labor market and on-the-job search to study how employer-to-employer (EE) transitions affect macroeconomic outcomes and monetary policy. We find that EE dynamics significantly shaped inflation during the Great Recession and COVID-19 recoveries. Despite similar unemployment paths, the former experienced weaker EE transitions and lower inflation. Optimal monetary policy prescribes a strong positive response to EE fluctuations, implying central banks should distinguish between episodes with similar unemployment but different EE patterns. We show that accounting for market incompleteness alters macroeconomic outcomes and optimal monetary policy prescriptions upon changes in EE transitions.

Keywords: Job mobility, monetary policy, HANK, job search
JEL Codes: E12, E24, E52, J31, J62, J64

*First version: August 2022. We thank our discussants Masao Fukui, Chris Huckfeldt, Andreas Mueller, Leena Rudanko, Ettore Savoia, as well as Sushant Acharya, Adrien Auclert, Bence Bardóczy, Guido Menzio, Giuseppe Moscarini, Chris Moser, Fabien Postel-Vinay, Bruce Preston, Matthew Rognlie, Benjamin Schoefer, Ludwig Straub, Ayşegül Şahin, Gianluca Violante for helpful comments. We also thank participants at NBER SI sessions (The Micro and Macro Perspectives of the Aggregate Labor Market; Micro Data and Macro Models; Workshop on Methods and Applications for Dynamic Equilibrium Models), The Minnesota Workshop in Macroeconomic Theory, San Francisco Fed Macroeconomics and Monetary Policy Conference and participants at various other seminars and conferences. This research was supported through computational resources provided by the Big-Tex High Performance Computing Group at the Federal Reserve Bank of Dallas. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Central Bank of the Republic of Türkiye (CBRT), the Federal Reserve Bank of St. Louis, the Federal Reserve Bank of Kansas City, or the Federal Reserve System. Birinci: birincise@gmail.com, Karahan: yfkarahan@gmail.com, Mercan: aymmercan9@gmail.com, See: kurtgerrardsee@gmail.com.

1 Introduction

The Federal Reserve has the dual mandate of fostering price stability and maximum employment. Its main tool for this goal is the federal funds rate, which it sets based on inflation and measures of economic slack. Labor market slack measures focus on the *quantity* of employment (e.g. the unemployment rate) and underemphasize its *quality* dimension. Starting with the seminal work by Moscarini and Postel-Vinay (2022), there has been a growing interest in understanding the role of employer-to-employer (EE) flows in determining the quality of employment and its implications for inflation. EE flows affect production costs via firm competition for workers and the economy's productive capacity through labor reallocation. Job mobility is also a determinant of wages, and thus impacts aggregate demand—all important determinants of inflation.

In this paper, we quantitatively analyze the positive and normative implications of EE flows for inflation and monetary policy, and ask two questions. First, how much and through which channels do EE flows affect macroeconomic outcomes, particularly inflation? Second, how should monetary policy respond to changes in EE flows?

To this end, our paper makes four contributions. First, we develop a model that combines the Heterogeneous Agent New Keynesian (HANK) framework with a frictional labor market and on-the-job search (OJS). Second, we use the model to quantify the impact of EE fluctuations on inflation. We show that muted job mobility caused 0.42 percentage points (pp) lower inflation during the 2016–2019 recovery episode, which saw "missing inflation" despite a large decline in the unemployment rate. We also find that the elevated EE rate during the "Great Resignation" of 2021–2022 generated an additional 0.77 pp of inflation. Third, we study optimal monetary policy within a class of Taylor rules under a dual-mandate central bank objective function. The optimal policy prescribes a less aggressive response to the unemployment gap than the conventional value used in the literature, and a strong positive response to the EE gap. This result suggests that the central bank should distinguish episodes with different EE dynamics, even under similar unemployment dynamics. Finally, we show that an alternative model approximating perfect self-insurance generates largely different outcomes in response to EE fluctuations and prescribes an optimal policy with a much smaller response to the EE gap. Hence, accounting for market incompleteness matters for the positive and normative implications of EE fluctuations.

We start by developing a Two Agent New Keynesian (TANK) model to demonstrate our key insights, which serves as a segue to our quantitative (HANK) model. In this model, a hand-to-mouth household (HtM) consumes its entire labor income, while a permanent-income household (PIH) collects labor income and firm profits, and can save in a risk-free asset. Both unemployed and employed members of these households search for jobs in a frictional labor market and work in firms that produce labor services. We assume that all firm-worker pairs are homogeneous in productivity and that workers hired out of unemployment receive a fraction of the flow surplus as wage, while poached workers extract the entire surplus from the new match and thus always switch jobs when possible. The rest of the model follows the New Keynesian tradition. Monopolistically competitive firms buy labor services in a competitive market to produce intermediate goods, which are then sold to a final-good producer. The monetary authority follows a Taylor rule to control the short-term nominal rate.

We use this model to illustrate how EE dynamics affect macroeconomic outcomes. Starting from the steady state, we raise the EE rate by exogenously increasing employed workers' contact probability with external firms, modeled as a shock to the relative job search efficiency of employed workers, i.e., an "OJS efficiency shock." Holding the real interest rate fixed, a higher OJS efficiency increases output, the real marginal cost for intermediate firms, and inflation. This rise in output is driven by an increase in labor income due to wage increases from more-frequent job switches. Furthermore, a higher OJS efficiency (i) increases the probability that a vacancy meets with an employed worker who in turn extracts the entire surplus and (ii) reduces the expected duration of a match as workers are more likely to be poached. Both these composition and match duration channels lead to a decline in the expected match value for labor services firms, generating a rise in the marginal cost for intermediate firms and thus inflation. Allowing the real rate to adjust, the rise in inflation induces a more than one-for-one increase in the nominal rate. Despite the rise in labor income, the rise in the real rate together with lower firm profits lead to lower demand and output. Therefore, in equilibrium, the shift in EE behaves like a cost-push shock.

Importantly, to understand whether accounting for imperfect self-insurance changes the effects of EE fluctuations on aggregate outcomes, we compare the results in the TANK model with a nested Representative Agent New Keynesian (RANK) version. We uncover two key differences. First, the declines in aggregate demand, output, and labor market tightness in response to an OJS efficiency shock are around *half* as large in the TANK model. Second, the real interest rate increase is around *twice* as large in the TANK model.[1] Hence, accounting for incomplete markets influences both aggregate outcomes and the conduct of monetary policy in response to EE fluctuations.

We continue by extending the stylized TANK model into a fully-fledged quantitative HANK model that can be more tightly linked to the data. We highlight three reasons for building this HANK version. First, to correctly quantify the effects of job mobility, it is crucial to capture the extent of wealth heterogeneity and allow for an endogenous interaction between heterogeneity in the marginal propensities to consume (MPC) and idiosyncratic income risk as emphasized in the broader HANK literature. Second, capturing the supply-side effects of job switches requires productivity differences across matches in the labor services sector. Third, to empirically discipline the wage dynamics of job switchers, job stayers, and job losers, we need to account for the evolution of productivity both on and off the job, and wage bargaining.

Thus, our HANK model has the following additional features. Individuals can save in shares of a mutual fund to self insure against idiosyncratic income fluctuations due to unemployment risk, human capital depreciation in unemployment, and stochastic retirement events. They are heterogeneous in their discount factors, allowing the model to generate an empirically plausible average MPC. Worker-firm pairs in the labor services sector are heterogeneous in their match productivity and workers' human capital. The wage is an endogenous piece-rate of the pair's output, which is determined through Bertrand competition. The heterogeneity in match productivity together with this bargaining protocol allows the model to capture the effect of job switches on productivity and wage gains. Finally, an individual's human capital stochastically appreciates when employed and depreciates when unemployed, allowing to generate wage growth among job stayers and scarring from job loss. The economy is subject to aggregate shocks to demand, productivity, and the OJS efficiency.

Estimating shocks and studying optimal monetary policy in this environment requires solving and simulating the model under aggregate shocks efficiently. We overcome the well-known challenges in solving HA models under aggregate uncertainty by implementing the sequence-space Jacobian (SSJ) method of Auclert, Bardóczy, Rognlie, and Straub (2021). One complication in adapting this method to our context is that endogenous worker distributions *directly* enter equilibrium conditions and optimization problems, in contrast to simpler settings where *only* scalars enter such conditions. In sum, we implement the SSJ method in a multi-stage model with a

[1] We explain the reasons behind these results in our discussions below and in Section 2.

frictional labor market, where one needs to track the discretized worker distributions.

We calibrate the model's steady state to match the average MPC and relevant labor market moments in the U.S. prior to the Great Recession, as well as targets for the New Keynesian components, including the average level of markups, the slope of the Phillips curve, and the responsiveness of the nominal rate to inflation and unemployment gaps. We then jointly estimate aggregate shocks to demand, productivity, and the OJS efficiency by targeting the empirical correlations of the unemployment rate, EE rate, and inflation with output, as well as their standard deviations. We find that OJS efficiency shocks account for 16 percent of inflation fluctuations.

We then use the model to quantify (i) the drag on inflation from suppressed EE flows during the expansion after the Great Recession and (ii) the rise in inflation from elevated job mobility after the COVID-19 recession. In simulating our model to emulate the first episode, we first back out the sequence of positive demand and negative OJS efficiency shocks that replicate the declining path of unemployment and stable EE rates observed in this period. We then compare the outcomes to a counterfactual economy matching *only* the decline in unemployment by estimating demand shocks alone. We find that the OJS efficiency shocks we infer lower annual inflation by 0.42 pp. In a similar exercise, we simulate the recovery from the COVID-19 recession and show that the sizable rise in the EE rate lead to 0.77 pp higher inflation. The model's prediction that higher job mobility generates a rise in inflation is consistent with the empirical observation on the differential growth of unit labor costs (ULC) between these two episodes that feature similar unemployment but differential EE dynamics.

Next, we study the normative implications of job mobility for monetary policy. We consider an augmented Taylor rule that responds to both unemployment and EE rate deviations from their steady states as well as the inflation gap. For a dual-mandate central bank, we find that the optimal policy in response to OJS efficiency shocks prescribes a less aggressive response to the unemployment gap than what is commonly assigned in the literature, and a strong positive response to the EE gap. In practice, this optimal policy implies that during the last two recoveries with very similar declines in the unemployment rate, monetary policy should have been more aggressive between 2021–2022 when job mobility was much higher.

Critically, in addition to comparing outcomes between the TANK and RANK models, we present two more results using the HANK model that reveal why imperfect self-insurance affects the positive and normative implications of EE fluctuations.

First, in equilibrium, higher OJS efficiency leads to an increase in the real interest rate, suppressing aggregate demand. Notably, we show that wealth-rich (low-MPC) agents reduce their consumption more sharply due to a drop in their wealth. This decline stems from falling real dividends and share prices, as higher labor costs driven by elevated OJS efficiency erode firm profits. As wealthier agents rely relatively more on financial income, the reduction in their wealth triggers a larger consumption cutback. Overall, a key implication of this result is that using a model approximating perfect self-insurance (i.e., an economy populated by low-MPC agents) would have led to a much larger decline in demand, output, and market tightness, changing the magnitude of the monetary policy response through real rate adjustments—a result that is in line with our findings comparing the TANK and RANK models.

Second, we show that accounting for imperfect self-insurance also matters for optimal monetary policy. In particular, we recalibrate our HANK model without discount factor heterogeneity to match the same moments as in our baseline calibration, except for the empirically realistic high average MPC. In this alternative calibration, the optimal policy in response to OJS efficiency shocks prescribes an EE gap coefficient that is around *half* of that in the baseline optimal policy. Since a model with perfect self-insurance implies much larger declines in demand, output, and tightness due to a higher demand sensitivity to the real rate among agents with low MPCs, the monetary authority responds to inflationary pressures from a higher EE rate by raising the nominal rate less aggressively relative to a model with imperfect self-insurance.

Related literature. Our work is most closely related to Moscarini and Postel-Vinay (2022), who are the first to introduce search frictions with OJS into a RANK model to study the inflationary effects of job mobility. Relative to their work, we have two key differences. First, our HANK model features imperfect self-insurance, and therefore changes in job mobility are an important determinant of aggregate demand as well as supply. Crucially, we show that accounting for imperfect insurance largely alters aggregate outcomes and optimal monetary policy prescriptions upon changes in EE flows. Second, Moscarini and Postel-Vinay (2022) emphasize how the effects of EE flows depend on the distribution of workers across match productivity. In particular, if the employed are well-matched, a rise in outside offers is inflationary as counteroffers build inflationary pressure. If, however, the employment distribution is concentrated at the bottom rungs of the ladder, a rise in outside offers leads to a higher EE rate and lower inflation due to the productivity improvements dominating wage increases. While our model allows for either effect to dominate, when we calibrate it to the microdata on wage dynamics among job switchers and stayers, we find that a positive OJS efficiency shock leads to a higher EE rate, marginal cost, and inflation. This result is in line with the data that ULC (and wage) growth was muted between 2016 and 2019 when inflation and the EE rate remained low, and that ULC growth was strong between 2021 and 2022 when inflation and the EE rate were elevated.

The HANK literature emphasizes empirically realistic distributions of wealth and MPCs to quantify the effects of monetary policy (Kaplan, Moll, and Violante, 2018; Auclert, Rognlie, and Straub, 2020; Auclert, Rognlie, and Straub, 2023). This literature often assumes a stylized labor supply block using an exogenous income process. Our contribution lies in a model with an endogenous income process via an explicit job ladder and human capital evolution. This allows us to endogenize income risk, its variation with aggregate shocks and correlation with MPCs, which are key elements for the amplification properties of HANK models (Acharya and Dogra, 2020). On the computational side, we implement the SSJ method of Auclert et al. (2021) to handle discretized worker distributions that enter *directly* into optimization problems and equilibrium conditions, which they flag as a limitation of the method.[2]

Several papers synthesize elements from labor search models with those of the NK literature (Ravn and Sterk, 2021; Gornemann, Kuester, and Nakajima, 2021). Relative to these, our model features endogenous EE transitions, which we use to study how job mobility affects aggregate dynamics. Our paper also complements recent studies that use labor search models to investigate how wage rigidity affects unemployment fluctuations through its impact on job-finding rates (Fukui, 2020) and job-separation rates (Blanco, Drenik, Moser, and Zaratiegui, 2023). We highlight how wage changes due to job switching and rebargaining affect the macro economy.

Finally, Alves (2020) also embeds the insights from Moscarini and Postel-Vinay (2022) in a HANK model. Our work differs from his in three ways. First, our model features richer labor-market heterogeneity by human capital and match productivity. Second, we not only quantify the effect of job mobility on inflation but also decompose its underlying drivers, including through the changes in the worker distribution. Third, we study the normative implications of job mobility for monetary policy.[3]

[2] The paper states that "An important limitation [of the method], however, is that ... the value function $\mathbf{v_t}$ is not allowed to depend on [the distribution of agents in the economy] $\mathbf{D_t}$."

[3] In very recent work, Darougheh et al. (2025) adopt our model—where they omit human capital dynamics and discount factor heterogeneity, but add a discrete OJS choice—to study the impact of

2 A TANK model with a frictional labor market

We start by studying a TANK model, which is a simplified version of the quantitative model in Section 3, for two purposes.[4] First, it clarifies the channels through which job mobility impacts the economy with simple equations. Second, it allows us to compare two economies with identical steady states but different levels of market incompleteness and show how imperfect insurance affects outcomes upon EE changes. To this end, we make simplifying assumptions to highlight the main model mechanisms.[5]

2.1 Environment

Time is discrete and runs forever. There is a unit mass of individuals belonging to one of two types of households, firms in three vertically integrated sectors (labor services, intermediate goods, and final goods), and a monetary authority.

Households. Individuals are infinitely lived and discount the future with $\beta \in (0, 1)$. A fraction μ belongs to a hand-to-mouth (HtM) household and consumes their real labor income $c_t^{HtM} = W_t$. The remaining $1-\mu$ belong to a permanent-income household (PIH) that saves in a risk-free asset a at rate r_t, whose consumption c_t^{PIH} solves:

$$\max_{c_t^{PIH}} \sum_{t=0}^{\infty} \beta^t u(c_t^{PIH}) \quad \text{s.t.} \quad c_t^{PIH} + a_t = (1 + r_t)a_{t-1} + Z_t,$$

where Z_t is total real household income consisting of labor income and firm profits, described below.[6] Assuming a CRRA utility with parameter σ, c_t^{PIH} is given by:[7]

$$c_t^{PIH} = \frac{1}{\sum_{s=0}^{\infty} \beta^{s/\sigma} q_{t+s}^{1-1/\sigma}} \Big((1 + r_t)a_{t-1} + \sum_{s=0}^{\infty} q_{t+s} Z_{t+s} \Big), \tag{1}$$

with $q_{t+s} = \frac{1}{1+r_{t+1}} \times \cdots \times \frac{1}{1+r_{t+s}}$, $q_t = 1$, and a no-Ponzi condition $\lim_{s\to\infty} q_{t+s}a_{t+s} = 0$. There is perfect risk-sharing across the members within each household, and aggregate consumption is $C_t = (1-\mu)\, c_t^{PIH} + \mu c_t^{HtM}$.

---

a tax change in Denmark on the incentives to climb the job ladder.

[4]The model presented here is inspired by Masao Fukui's insightful discussion of our paper at the San Francisco Fed Macroeconomics and Monetary Policy Conference on March 2023.

[5]Namely, we assume that self-insurance is independent of income risk, all matches are homogeneous in productivity, i.e., job switches do not have supply-side effects, all workers are homogeneous in productivity, and wages follow an exogenous rule described below. We also abstract from fiscal policy. Section 2.4.2 discusses how we relax all of these assumptions in our HANK model.

[6]We assume that the PIH household owns the firms and collect their profits, while the HtM household does not hold assets. This assumption is motivated by the fact that some households do not possess any liquid assets, including shares of publicly traded firms.

[7]See Appendix A.1 for the derivation.

Firms. Service firms hire workers to produce labor services, sold in a competitive market to intermediate firms. These monopolistically-competitive firms produce a differentiated variety, facing a downward-sloping demand from a final-goods producer.

Labor market. The labor market features random search. Both unemployed and employed agents search for jobs, with their contact probabilities determined by their respective job search efficiencies and the market tightness, θ_t. In particular, a worker's contact rate per unit of search efficiency is given by $f(\theta_t)$ defined later. The job search efficiency of the unemployed is normalized to 1, and the relative job search efficiency of the employed (or, OJS efficiency) is given by ν_t. We assume that all job offers are accepted and real wages are paid according to a rule described below. In the beginning of each period, matches dissolve at an exogenous rate δ. All matches are homogeneous in productivity and produce one unit of labor service each period.

Labor services firms pay a fixed cost κ to post vacancies and sell their output to intermediate firms at nominal price P_t^l ($p_t^l = P_t^l/P_t$ in units of the final good).[8] Thus, the output is valued at p_t^l in real terms. We assume that when workers find a job while unemployed, they receive a real wage equal to a fraction $\alpha \in [0, 1]$ of match revenue p_t^l. When employed workers are poached, we assume that they extract the entire flow surplus and receive a real wage of p_t^l. Hence, there are two types of employed: those who found a job out of unemployment with wages equal to αp_t^l and those who are poached with wages equal to p_t^l. We denote the masses of these groups by $e_t(\alpha)$ and $e_t(1)$, respectively, and the mass of unemployed is u_t. Aggregate labor income is:

$$W_t = \alpha p_t^l e_t(\alpha) + p_t^l e_t(1), \tag{2}$$

while aggregate real profits in the economy are $\Gamma_t = \Gamma_t^I + \Gamma_t^S$, where Γ_t^I and Γ_t^S are the aggregate per-period real profits of the intermediate and labor services firms, described in Section 2.2. Per-capita real income of the PIH household is then:[9]

$$Z_t = W_t + \Gamma_t/(1-\mu). \tag{3}$$

Finally, the laws of motion for workers' end-of-period labor market status are:

$$\begin{aligned} u_t &= (1 - f(\theta_t))u_{t-1} + \delta(1 - f(\theta_t))(1 - u_{t-1}) \\ e_t(\alpha) &= f(\theta_t)u_{t-1} + (1 - \delta)(1 - \nu_t f(\theta_t))e_{t-1}(\alpha) + \delta f(\theta_t)(1 - u_{t-1}) \\ e_t(1) &= (1 - \delta)e_{t-1}(1) + (1 - \delta)\nu_t f(\theta_t)e_{t-1}(\alpha), \end{aligned} \tag{4}$$

where workers who lose their jobs can search for a new one within the same period.

[8] Unless otherwise stated, we use uppercases (lowercases) for nominal (real) variables.

[9] Since all profits accrue to the PIH household, total profits $\Gamma_t/(1-\mu)$ are scaled by their measure.

Monetary authority. The monetary authority controls the short-term nominal interest rate i_t according to the following reaction function:

$$i_t = i^* + \Phi_\pi(\pi_t - \pi^*) + \Phi_u(u_t - u^*). \tag{5}$$

Here, i^* denotes the steady-state nominal interest rate, Φ_π governs the responsiveness of the central bank to deviations of inflation from its target π^*, and Φ_u controls how much the central bank responds to deviations of the unemployment rate from its steady-state value u^*. Finally, the real interest rate r_t satisfies the Fisher equation:[10]

$$1 + i_t = (1 + \pi_{t+1})(1 + r_{t+1}). \tag{6}$$

The timing convention for these variables is as follows: The nominal interest rate i_t is indexed to the period in which it is set and applies between periods t and $t+1$. The inflation rate is denoted by the period in which it is measured; i.e., π_{t+1} is the realized inflation between periods t and $t+1$. The real rate has the same timing convention as inflation: r_{t+1} is the ex-post realized real interest rate from t to $t+1$.

Timing of events. At the start of each period, aggregate and job destruction shocks realize (details in Section 2.4). Next, the job search stage opens: firms post vacancies, workers search for jobs, and new matches are formed. Then, in the production stage, each worker-firm pair produces labor services, intermediate firms produce differentiated goods using labor services and set their prices, and final goods firms produce the numeraire using intermediate goods. Given the unemployment rate and inflation, the monetary authority sets the nominal rate. Next, service firms pay wages to their workers and intermediate and service firms realize their profits. Finally, households pool income across their members and decide on how much to consume.

2.2 Production

Final goods. The final-good producer purchases intermediate goods $y_t(j)$ at relative price $p_t(j) = P_t(j)/P_t$ and produces the final good Y_t using the technology:

$$Y_t = \left(\int_0^1 y_t(j)^{\frac{\eta-1}{\eta}} dj \right)^{\frac{\eta}{\eta-1}},$$

where η is the elasticity of substitution between varieties, solving the following:

$$\max_{\{y_t(j)\}} \; Y_t - \int p_t(j) y_t(j) dj.$$

[10]In solving the TANK model and later the HANK model, we use a first-order perturbation method around the non-stochastic steady state. Thus, we express the equations in a way that agents take expectations with respect to the idiosyncratic shocks, but not with respect to the aggregate shocks.

This problem determines the demand for each variety, $y_t(j) = p_t(j)^{-\eta} Y_t$, as a function of its relative price $p_t(j)$ and aggregate demand conditions Y_t with an ideal price index satisfying $1 = \left(\int p_t(j)^{1-\eta} dj\right)^{\frac{1}{1-\eta}}$ that the intermediate firms take as given.

Intermediate goods. Intermediate firms produce $y_t(j)$ using a linear technology: $y_t(j) = l_t(j)$, where $l_t(j)$ denotes the amount labor services bought from labor services firms. They set the price for their variety, taking into account demand from the final-good producer and price adjustment costs à la Rotemberg (1982). Pricing frictions render last period's *relative* price $p_{t-1}(j)$ a state variable. These firms solve the following profit maximization problem:

$$\Theta(p_{t-1}(j)) = \max_{p_t(j)} \; p_t(j) y_t(j) - p_t^l y_t(j) - Q(p_{t-1}(j), p_t(j)) Y_t + \frac{1}{1 + r_{t+1}} \Theta(p_t(j)), \quad (7)$$

where adjustment costs are given by $Q(p_{t-1}(j), p_t(j)) = \frac{\eta}{2\vartheta} \log\left(\frac{p_t(j)}{p_{t-1}(j)}(1+\pi_t) - \pi^*\right)^2$. Appendix A.2 shows that the New Keynesian Phillips curve (NKPC) is given by:

$$\Pi_t = \vartheta\left(mc_t - \frac{\eta - 1}{\eta}\right) + \frac{1}{1 + r_{t+1}} \Pi_{t+1} \frac{Y_{t+1}}{Y_t}, \quad (8)$$

where $\Pi_t = \frac{\log(1+\pi_t-\pi^*)(1+\pi_t)}{1+\pi_t-\pi^*}$ and $mc_t = p_t^l$ is the real marginal cost of production. Finally, the aggregate per-period real profits of intermediate firms are given by

$$\Gamma_t^I = \left(1 - p_t^l - \frac{\eta}{2\vartheta} \log(1 + \pi_t - \pi^*)^2\right) Y_t, \quad (9)$$

where we use the fact that equilibrium relative prices are one, $p_t(j) = 1$ by symmetry.

Labor services. A continuum of labor services firms post vacancies, incurring cost κ per vacancy. Labor market tightness, θ_t, is defined as the ratio of vacancies v_t to the aggregate measure of job searchers, including both unemployed and employed workers weighted by their job search efficiencies, $S_t = u_{t-1} + \delta(1 - u_{t-1}) + \nu_t(1-\delta)(1-u_{t-1})$. Let $M(v, S)$ be a constant-returns-to-scale (CRS) matching function that determines the number of matches as a function of vacancies and effective job seekers. We define $q(\theta) = \frac{M(v,S)}{v} = M\left(1, \frac{1}{\theta}\right)$ as the firm's contact rate and $f(\theta) = \frac{M(v,S)}{S} = M(\theta, 1)$ as the worker's contact rate per unit of search efficiency.

We now turn to the problem of labor services firms. Consider a firm that hires an unemployed worker. The worker-firm pair produces one unit of labor services which is then sold to intermediate goods producers at real price p_t^l, and the worker is paid real wages equal to αp_t^l. The real value of this firm is then given by:

$$J_t = (1 - \alpha) p_t^l + \frac{1}{1 + r_{t+1}} (1 - \delta)\left(1 - \nu_{t+1} f(\theta_{t+1})\right) J_{t+1}. \quad (10)$$

We assume that labor services firms are managed by risk-neutral entrepreneurs and thus they discount the future at the real interest rate r_{t+1}.[11] Note that the value of a poaching firm is zero. Here, the worker extracts the entire match revenue as wage by assumption. Therefore, the per-period real profits of service firms are given by[12]:

$$\Gamma_t^S = (1-\alpha)p_t^l e_t(\alpha). \tag{11}$$

The real value of a labor services firm posting a vacancy is as follows:

$$V_t = -\kappa + q(\theta_t)\frac{u_{t-1} + \delta(1-u_{t-1})}{u_{t-1} + \delta(1-u_{t-1}) + \nu_t(1-\delta)(1-u_{t-1})} J_t = -\kappa + q(\theta_t)\mathbb{E}J_t, \tag{12}$$

where $\mathbb{E}J_t$ is the ex-ante expected value of posting a vacancy conditional on meeting a worker. Further, a free-entry condition implies that $V_t = 0$.

2.3 Equilibrium

Goods market clearing requires that output Y_t equals aggregate demand C_t, and labor market clearing requires labor demand to be equal to employment $1-u_t$:

$$C_t = Y_t = 1 - u_t. \tag{13}$$

Appendix A.3.1 defines the stationary equilibrium and Appendices A.3.2 and A.3.3 describe the solution of this model. Further computational details for solving the transition path are provided in Section 4.2 and in Appendices A.3.4 and B.3.

2.4 Results

In this section, we use our model to study how job mobility dynamics affect inflation and other macroeconomic outcomes. We use the same parameter values from the HANK model described in Section 3 and calibrated in Section 4.[13]

2.4.1 Understanding the mechanisms: Market incompleteness

Our main experiment involves introducing exogenous variation to job mobility dynamics. We implement this by simulating a series of shocks to the relative job search

[11] This assumption is similar to the HANK model where a risk-neutral mutual fund owns firms.

[12] Here, we assume that vacancy creation cost κ does not exhaust real resources.

[13] There are two exceptions. First, while in the HANK model we set the OJS efficiency ν to match the average EE rate in the data, we set a different value for this parameter here. This is because the TANK model does not feature any job ladder with respect to productivity and workers with $\alpha < 1$ are forced to move to a new job upon an outside offer. This implies that the average EE rate in the TANK model would be a function of the mass $e(\alpha)$, and that the empirical EE rate is not very informative. We simply set $\nu = 0.5$ in the TANK model to capture the findings by Faberman, Mueller, Şahin, and Topa (2022). They show that hours spent searching and mean applications sent among employed job-searchers are around half of those among unemployed workers. Second, in both the TANK and HANK models, we set the cost of vacancy creation κ to match the empirical unemployment rate. In addition, there are two parameters specific to the TANK model: the share of HtM households μ and the piece rate for jobs out of unemployment α. We set $\mu = 0.3$ and $\alpha = 0.5$.

Figure 1: Impulse responses to a positive OJS efficiency shock: Constant real rate

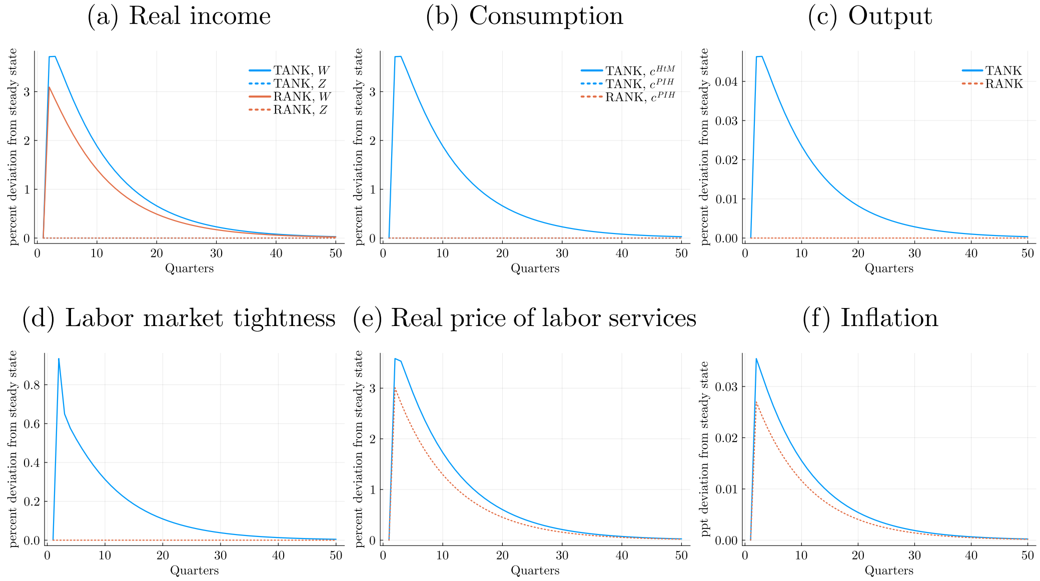

Notes: This figure presents the impulse responses of aggregate real labor income W, total income of the PIH household Z, consumption, output, labor market tightness, the real price of labor services p^l, and inflation in an economy subject to a 1 percent persistent increase in the OJS efficiency. Blue lines represent the TANK model and red lines plot represent the RANK model. In this exercise, we keep the real interest rate r constant.

efficiency of the employed ν, which we label as OJS efficiency shocks.[14] To analyze the role of market incompleteness, we also run the same experiments in a nested RANK model with the same steady state, which we obtain by setting the fraction of the HtM household μ to zero. To keep the discussion concise, we defer the details of estimating OJS shocks and their potential microfoundations to Section 4.3.

Constant real rate. We first study the effects of an OJS shock under a real interest rate fixed at its steady-state value. Here, the goal is to build intuition from a partial-equilibrium setting, focusing on the first-order effects of the shock.[15] Figure 1 presents impulse responses from a 1 percent increase in ν that gradually recovers.

We first focus on the RANK model. Aggregate real labor income W increases upon a rise in the OJS efficiency (Panel (a)). A higher OJS efficiency raises the probability that a worker with piece-rate $\alpha < 1$ is poached and experiences a real wage increase as she extracts the full flow surplus in her next job. Higher wages paid by labor services firms lead to an increase in the real price of labor services p^l charged to intermediate firms. Because of nominal rigidities, intermediate firms

[14] This shock is first introduced by Faccini and Melosi (2023) who build on Moscarini and Postel-Vinay (2022) and quantify the effects of job mobility on inflation over the last two decades.

[15] We use the linearized system obtained under a monetary policy reaction function that satisfies the Taylor principle, but we simply feed a constant real interest rate path to the system.

cannot immediately adjust their prices, leading to a decline in their profits Γ^I. As shown in Figure A.1, total real profits Γ decline. Because the PIH household collects these profits, the total income of the PIH household Z does not change, as the rise in labor income is offset by the decline in profits. Thus, aggregate demand c^{PIH} (Panel (b)), output (Panel (c)), and market tightness (Panel (d)) do not change.

A higher OJS efficiency reduces the ex-ante expected value from posting a vacancy conditional on meeting a worker, $\mathbb{E}J$, for two reasons.[16] First, a higher ν makes it more likely for a vacant firm to meet an employed worker, who in turn extracts the entire surplus. Second, the match is more likely to dissolve in the future as the worker enjoys more outside opportunities, reducing expected match duration. Given a fixed real rate and absent a change in tightness, the decline in $\mathbb{E}J$ requires a rise in p^l for equilibrium conditions to hold (Panel (e)). That is, to meet the same amount of demand, the intermediate firm needs to pay a higher p^l to convince the service firm to produce the same amount of labor services. The NKPC implies that inflation—to a first-order approximation—is driven by the real price of labor services p^l, which determines the real marginal cost for intermediate firms. As a result, inflation increases (Panel (f)).[17]

Next, we discuss results from the TANK model. Labor income increases more in TANK since a higher ν now has positive real effects: aggregate demand and tightness increase. Unlike in RANK, higher labor income arising from increased job-to-job moves directly translates into higher demand by the HtM household (Panel (b)), given that their income is unaffected by firm profits. This increase in demand leads to an increase in output. Further, an increase in tightness amplifies the contact probability of employed workers, $\nu f(\theta)$. As a result of higher demand and higher poaching probability, p^l (and thus inflation) also increases by more in TANK.

Overall, according to the RANK model under a constant real rate, the rise in job mobility acts like a pure cost-push shock for intermediate firms without any effect on demand or productivity, as we assumed all worker-firm matches are homogeneous in productivity. Thus, the OJS efficiency shock is completely absorbed by the rise in

[16]In this model, it is easier to see why an increase in ν leads to a higher real price of labor services as $\mathbb{E}J$ in Equation (12) is less complicated than its HANK counterpart in Equation (20). This is because the TANK model abstracts from heterogeneity in match-specific productivity, human capital, wealth, and discount factor, as well as from bargaining, all of which are featured in HANK.

[17]The reason behind the rise in inflation upon an increase in ν is the decline in $\mathbb{E}J$, driven by *both* the change in the composition of job seekers and a shorter match duration. Figure A.4 in Appendix A.4 shows that the shorter match duration channel is the main reason behind the rise in inflation. Thus, cost-push arising from wage increases due to poaching is not necessary to generate inflation.

Figure 2: Impulse responses to a positive OJS efficiency shock: Real rate response

(a) Real income (b) Consumption (c) Output

(d) Labor market tightness (e) Real interest rate (f) Inflation

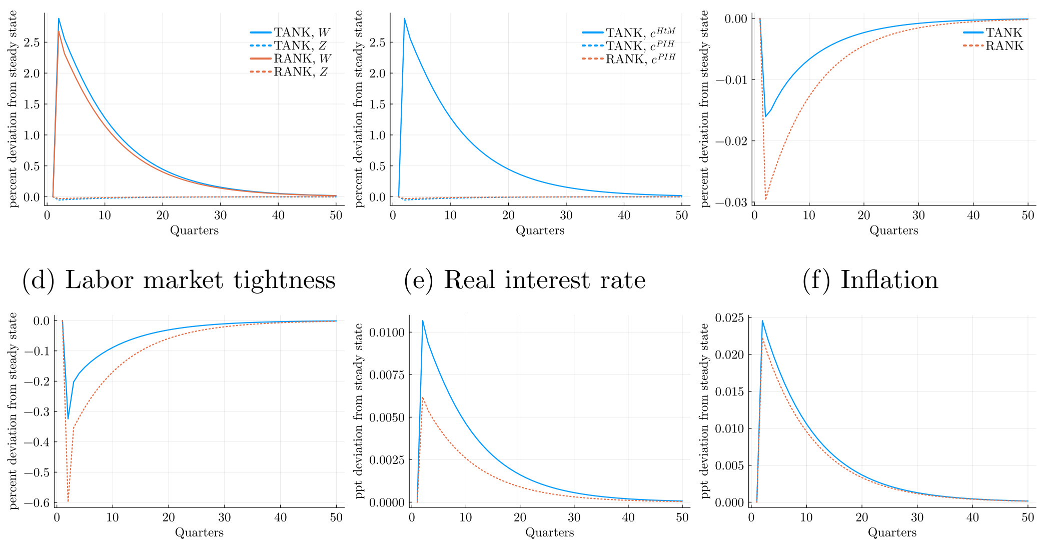

Notes: This figure presents the impulse responses of aggregate labor income W, total income of the PIH household Z, consumption, output, market tightness, real interest rate, and inflation in an economy subject to a 1 percent persistent increase in the OJS efficiency. Blue lines represent the TANK model and red lines represent the RANK model.

the marginal cost of production for intermediate firms and inflation. In the TANK model, the shock also increases demand. Hence, under a constant real interest rate, inflation in the TANK model rises by much more than it does in RANK. However, as we discuss below, the net effect of an increase in the OJS efficiency on the real outcomes depends on how much the real rate changes in equilibrium.

Real rate response. Next, we study the same experiment but allow the monetary authority to respond to the shock according to the reaction function in Equation (5). Figure 2 summarizes the results from this exercise.

We again start by focusing on the RANK case. Recall that a higher OJS efficiency only leads to inflation in the RANK model under a constant real rate. Here, through the response of the monetary authority, higher inflation induces an increase in the nominal rate and therefore an increase in the real rate (Panel (e)).[18] Despite the rise in labor income (Panel (a)), the rise in the real rate together with the decline in total profits (Figure A.2) lead to a decline in aggregate demand (Panel (b)). Lower aggregate demand implies lower output (Panel (c)) and market tightness (Panel (d)).[19]

In equilibrium, a higher OJS efficiency leads to a decline in the expected value of

[18] A rise in inflation induces a more than one-for-one increase in the nominal rate since $\Phi_\pi > 1$. This effect dominates the opposing effect due to higher unemployment (Figure A.2).

[19] In both models, changes in *aggregate* consumption and output are identical (see Figure A.3).

a match $\mathbb{E}J$, again through the compositional shift towards higher-paid employed job seekers and shorter match durations.[20] Thus, $\mathbb{E}J$ declines and p^l and inflation rise. Since demand now declines, the increase in inflation is smaller (Figure 2 Panel (f)) than the increase under a constant real rate (Figure 1 Panel (f)).

The TANK model outcomes differ from RANK in two important ways. First, the decline in aggregate demand is much smaller in the TANK model. This is because demand of the HtM household increases as its labor income increases. The resulting declines in output and tightness are around *half* as large in the TANK model. Second, the nominal rate (and real rate) increases by much more in the TANK model to curb what would otherwise have been a stronger inflation response. Unlike the PIH household, demand of the HtM household is impacted only indirectly by the real rate through its effect on tightness and income. As such, the HtM household's demand is less elastic to real rate changes. This leads to a much higher real rate, achieving a rise in inflation in the TANK model close to that in the RANK model. Put differently, to achieve a similar rise in inflation, the rise in the real rate should be around *twice* as large in the TANK model.[21] Importantly, these significant differences in real variables such as output and tightness as well as in monetary policy responses through real rate adjustments highlight the importance of modeling market incompleteness.

Overall, a positive OJS shock leads to a rise in the marginal cost and inflation, and a decline in demand, output, and tightness in both the RANK and TANK models, but there are two key differences in their magnitudes. First, the HtM household in the TANK model mitigates the decline in demand, output, and tightness. Second, the equilibrium real rate response to the shock is substantially stronger in the TANK model. These indicate that both the real outcomes and monetary policy in response to job mobility shocks differ greatly when the model features incomplete markets.

2.4.2 Taking stock and motivations for a HANK framework

This section shows that a positive OJS efficiency shock leads to a decline in demand and tightness and to an increase in the real marginal cost and inflation in both the TANK and RANK models.[22] We also find that the size of these changes differ

[20] In Section 5.4, we use the HANK model to provide a full decomposition of inflation.

[21] While RANK and TANK inflation dynamics are similar in this calibration, we find that the gap between inflation responses becomes larger as the fraction of HtM households rises. This is because the HtM household experiences a labor income increase upon a positive OJS shock.

[22] We reiterate that whether an increase in the OJS efficiency leads to an increase or a decrease in demand and tightness depends on, e.g., the response of monetary policy, as in Figures 1 and 2.

significantly between the two models. This finding suggests that incorporating imperfect self-insurance is relevant for both quantifying how job mobility dynamics affect macroeconomic outcomes and the conduct of monetary policy.

While our TANK model is useful for understanding the critical channels through which job mobility dynamics affect the macroeconomy, we have three motivations for studying a quantitative HANK model with a richer labor market. First, despite the TANK model highlights the key role of self-insurance in shaping economic outcomes and monetary policy, it is limited by the assumption that agents are either fully-insured or hand-to-mouth; i.e, insurance is independent of income fluctuations. For proper quantification and to establish an explicit interaction between MPC heterogeneity and idiosyncratic income risk, the model must feature a more realistic wealth distribution. Second, the TANK model captures only the inflationary effects of job mobility and ignores that some job switches are productivity enhancing. In order to capture this margin, the labor market must feature heterogeneity in match productivity and an explicit job ladder. Third, the stylized TANK model is not suitable for disciplining wage dynamics. Specifically, the model assumes that an employed worker is forced to switch jobs upon contact and the associated wage gains are substantial. While we find that this wage increase is not the main driver of inflation when job mobility increases, response magnitudes are still affected by the wage dynamics.[23] In addition, some external offers may not necessarily increase the worker's wage (for example, when the offer is much worse than current terms). This stylized model also counterfactually assumes that job stayers do not experience any wage increase over time. Finally, in terms of wage losses upon displacement, the model cannot account for the scarring effects of unemployment, as it assumes that all workers hired from unemployment receive the same wage and that unemployment does not affect a worker's productivity. For these reasons, we need to explicitly model changes in worker productivity on- and off-the-job together with endogenously determined wages to discipline income dynamics among job losers, switchers, and stayers with microdata.

As we build towards a substantially richer HANK model in the next section, two main insights will remain the same. Specifically, in the HANK model, higher job mobility generates higher inflation and the presence of high-MPC agents (imperfect insurance) shapes the responses of aggregate outcomes to job mobility shocks.

[23] For example, the decline in output in the TANK model without wage increases upon external offers (Figure A.4 (a)) is smaller than the decline in output in the TANK model (Figure 2 (c)).

3 A HANK model with a frictional job ladder

We now describe our quantitative model that combines a New Keynesian framework with heterogeneous agents and a frictional labor market with OJS. In our presentation of this model, we focus on components that are new relative to the TANK model.

3.1 Environment

The economy is populated by a measure one of ex-ante identical individuals, firms in three sectors, a mutual fund, a fiscal authority, and a monetary authority.

Firms. While this model incorporates the same three vertically integrated firms as in the TANK model, there are two differences. First, worker-firm matches in the labor services sector are heterogeneous in productivity. Second, the production function of the intermediate goods firms depends on aggregate productivity z as well.

Individuals. An individual's life consists of a working stage and a retirement stage. Individuals are heterogeneous in their discount factors $\beta \in \mathcal{B} = \left\{\underline{\beta}, \ldots, \overline{\beta}\right\}$, holdings of mutual fund shares $s \geq 0$, employment status e (employed E, unemployed U, or retired R), human capital (skill) $h \in \mathcal{H} = \left\{\underline{h}, \ldots, \overline{h}\right\}$,—among the employed—match productivity $x \in \mathcal{X} \equiv \{\underline{x}, \ldots, \overline{x}\}$, and endogenous and history-dependent piece-rate $\alpha \in (0, 1]$ governing the share of output they receive as wages.

Individuals are born with skill h and discount factor β drawn from distributions Γ^h and Γ^β, respectively. The heterogeneity in discount factors enables us to generate an empirically plausible value for the average MPC. While β is permanent once drawn, skill h evolves stochastically depending on an agent's employment status, as in Ljungqvist and Sargent (1998). An employed individual's skill increases by Δh percent with probability π^E, while an unemployed individual's skill declines by Δh percent with probability π^U in each period. Formally,

$$h' = \begin{cases} h \times (1 + \Delta h) & \text{with probability } \pi^E \\ h & \text{with probability } 1 - \pi^E \end{cases}$$

when employed, and,

$$h' = \begin{cases} h \times (1 - \Delta h) & \text{with probability } \pi^U \\ h & \text{with probability } 1 - \pi^U \end{cases}$$

when unemployed. We model human capital evolution to better capture wage dynamics in the data. Skill depreciation helps the model capture scarring from job loss, while skill appreciation enables the model to generate wage growth among job stayers.

Individuals trade shares of the mutual fund and make consumption decisions in the face of income risk due to the stochastic evolution of their skill and labor market frictions. Each period, working-age individuals retire with probability ψ^R. Retirees ($e = R$) finance consumption through their savings and from pension income ϕ^R. They die at rate ψ^D, upon which they are replaced with unemployed agents.[24]

Labor market. Unlike in the TANK model, worker-firm pairs are heterogeneous in match-specific productivity and worker skill. Upon contact, the pair draws a productivity x from distribution Γ^x, which remains constant throughout the match. The match operates a production technology given by $F(h, x) = hx$. The worker is paid real wages according to a predetermined rule $w(h, x, \alpha)$ (detailed below) every period until the match dissolves. The match can dissolve through exogenous job separation, which occurs at rate δ, retirement, or an endogenous EE transition by the worker. Unemployed individuals receive unemployment insurance (UI) benefits according to the function $UI(h) = \phi^U F(h, \underline{x})$ (denoted in consumption units), where we assume that UI payments are designed as a replacement rate ϕ^U of output the worker would have produced working at a job with the lowest match productivity $\underline{x}$.

Wage determination. Wage is an endogenous piece-rate α of the the match's flow output. We follow a static bargaining protocol—a simplified version of Postel-Vinay and Robin (2002)—for determining α, where firms Bertrand compete based on flow output (instead of present values).[25] Figure 3 summarizes this bargaining protocol.

Consider a worker with skill h, match productivity x, and piece rate α, whose wage is given by $w(h, x, \alpha) = \alpha\phi^E F(h, x)$, where $\phi^E \in (0, 1)$ represents the maximum share of output that a worker with maximum piece rate $\alpha = 1$ can capture as wage.[26] If this worker meets with a new firm with a higher productivity $x' > x$, she switches jobs. This is because the most the incumbent firm can offer to the worker is $w(h, x, 1) = \phi^E F(h, x)$. We assume that the new firm is willing to match this wage, i.e., $w(h, x', \alpha') = w(h, x, 1)$, implying a new piece rate $\alpha' = x/x'$. The worker is

[24] When an individual dies, she is replaced by an offspring who inherits her savings and enters the working stage as unemployed with skill h and discount factor β drawn from Γ^h and Γ^β, respectively.

[25] Bargaining over the entire present discounted value of the surplus would introduce an additional loop into the solution, where for a given wage function, we would have to check whether it is consistent with the worker and firm value functions, as in Krusell et al. (2010). To avoid this added computational burden, we assume that bargaining takes place over the flow surplus.

[26] When $\phi^E < p^l$, the firm's profit is greater than zero and there are no firms with negative surplus.

Figure 3: Wage determination

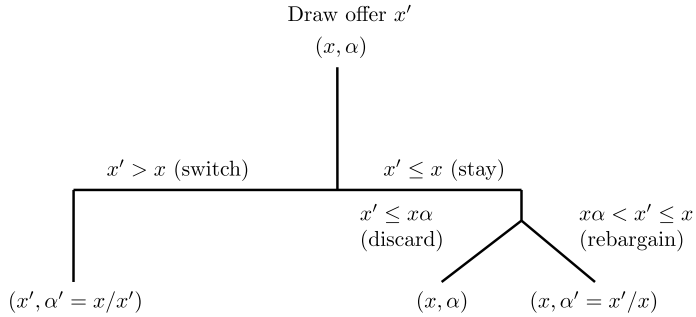

better off switching to the more-productive firm as the piece rate can only become (weakly) larger in the future upon new contacts, as discussed below.

Now suppose the worker receives a lower-productivity offer $x' < x$, resulting in the worker staying with the incumbent firm. This case induces two scenarios. First, x' could be so low that even the maximum potential wage from the new job cannot match the worker's wage, i.e., $w(h, x', 1) < w(h, x, \alpha)$, which happens when $x' < \alpha x$. In this case, the worker discards the offer and continues with the same piece rate. Second, x' could be sufficiently high to serve as a credible threat for the worker to bid up her wage. This happens when $w(h, x, 1) > w(h, x', 1) > w(h, x, \alpha)$, i.e., $x > x' > \alpha x$, in which case the incumbent firm matches the maximum potential wage from the outside offer, $w(h, x, \alpha') = w(h, x', 1)$, implying an updated piece-rate $\alpha' = x'/x$.

The heterogeneity in match-specific productivity together with this bargaining protocol allows the model capture the productivity effects of job switches, which we ignored in Section 2. In addition, workers are not forced to accept all offers—they now decline outside offers with lower match productivity; i.e., job switches are rational.

There are three sources of wage growth in this model. First, output increases as the worker's human capital appreciates over time, leading to higher wages given a piece rate. Second, poaching yields a wage increase as the worker extracts the entire surplus from the incumbent firm. Third, external offers not good enough to poach the worker but with sufficiently high match productivity lead to an increase in the piece rate via rebargaining with the incumbent firm. The latter two channels capture the potential inflationary effects of EE flows that stem from wage increases in excess of productivity gains upon job switch and from counteroffers by incumbents.[27]

[27] We explicitly model these sources of wage growth for three reasons. First, a higher EE rate is empirically associated with higher wage growth for both stayers and switchers (Moscarini and Postel-Vinay, 2016). Second, in the absence of wage increases due to external offers, wage growth among stayers would counterfactually be attributed to human capital appreciation alone. Third, while

We assume unemployment is akin to being employed at the lowest match productivity $\underline{x}$. Thus, for a match with productivity x', the piece rate out of unemployment is $\alpha' = \underline{x}/x'$. We further assume that all offers out of unemployment are accepted.[28]

Mutual fund. A risk-neutral mutual fund owns all the firms in the economy, holds all nominal bonds B_t issued by the government, and sells shares in return. The fund pays a nominal dividend D_t per share and can be traded by individuals at price P_t^s.

Fiscal and monetary authorities. The government implements a linear consumption tax τ_c and a progressive income tax. For any gross income level ω, net income is given by $\tau_t \omega^{1-\Upsilon}$, where τ_t captures the level of taxation and $\Upsilon \geq 0$ captures the rate of progressivity built into the tax system, as in Benabou (2002) and Heathcote, Storesletten, and Violante (2014).[29] Together with the tax revenue, the government issues nominal bonds B_t to finance UI benefits, retirement pensions, and an exogenous stream of nominal expenditures G_t. The central bank sets the short-term nominal interest rate using the reaction function given in Equation (5).

Timing of events. First, aggregate shocks (detailed in Section 4.3) are realized. Then, the government sets taxes. Exogenous retirement, mortality, and job destruction shocks are realized. Next, workers' skill evolve and new workers replenish the dead. In the search stage, firms post vacancies and individuals search for jobs. Once new contacts are made, match productivities are observed, and endogenous job switches occur. Then, each match produces labor services; intermediate firms produce differentiated goods and set their prices; and final goods are produced. Given the unemployment rate and inflation, the monetary authority sets the nominal rate. Next, intermediate and service firms realize their profits; service firms pay wages to workers; the mutual fund pays out dividends; and the government collects taxes, issues new bonds, pays out benefits, and spends an exogenous amount. Finally, individuals decide how much to consume and how many shares to buy.

---

wage-inflation induced by EE is not indispensable for generating a positive relationship between the level of EE rate and inflation (see Figure A.4 in Appendix A.4), it has effects on how EE dynamics shape demand through consumption responses triggered by future income changes.

[28]Under our calibration, we verify that all matches out of unemployment have positive surplus, as dynamic gains of employment dominate the value of waiting for a match with higher productivity.

[29]Parameter τ is inversely related to the tax rate. Under $\Upsilon = 0$, the tax rate is $1 - \tau$.

3.2 Individuals

Individuals decide on accepting a job offer while employed, as well as consumption and share holdings subject to budget and short-selling constraints. We cast the problems recursively, where time subscripts encode all the relevant aggregate state variables, as of the consumption stage within a period.

Unemployment. Let $V_t^U(s,h,\beta)$ denote the value of an unemployed with shares s, skill h, and discount factor β in period t. The problem of this agent is given by

$$\begin{aligned} V_t^U(s,h,\beta) &= \max_{s',\, c\geq 0} \; u(c) + \beta(1-\psi^R)\mathbb{E}_{h'|h}\left[\Omega_{t+1}^U(s',h',\beta)\right] + \beta\psi^R V_{t+1}^R(s',\beta) \\ \text{s.t.} \quad & P_t c(1+\tau_c) + P_t^s s' = P_t \tau_t UI(h)^{1-\Upsilon} + (P_t^s + D_t)s, \end{aligned} \tag{14}$$

where $\Omega_{t+1}^U(s',h',\beta)$ is the value of job search for unemployed workers at the beginning of the next period described below. Unemployed workers receive dividends D_t from the mutual fund in proportion to their share holdings s, receive real after-tax UI benefits, and make a consumption-saving decision.

Employment. Let $V_t^E(s,h,x,\alpha,\beta)$ denote the value of a worker with shares s, skill h, match productivity x, piece rate α, and discount factor β, solving the problem:

$$\begin{aligned} V_t^E(s,h,x,\alpha,\beta) &= \max_{s',\, c\geq 0} u(c) + \beta(1-\psi^R)\mathbb{E}_{h'|h}\left\{(1-\delta)\Omega_{t+1}^E(s',h',x,\alpha,\beta) + \delta\Omega_{t+1}^U(s',h',\beta)\right\} \\ &\quad + \beta\psi^R V_{t+1}^R(s',\beta) \\ \text{s.t.} \quad & P_t c(1+\tau_c) + P_t^s s' = P_t \tau_t w(h,x,\alpha)^{1-\Upsilon} + (P_t^s + D_t)s. \end{aligned} \tag{15}$$

Employed agents collect dividends, receive a real after-tax wage, and choose consumption and share holdings. At the beginning of the next period, the job might dissolve exogenously, in which case the worker becomes unemployed and searches for a new job. If not, the worker can engage in OJS, whose value is given by $\Omega_{t+1}^E(\cdot)$.

Retirement. Finally, the value of retirement is given by

$$\begin{aligned} V_t^R(s,\beta) &= \max_{s',\, c\geq 0} u(c) + \beta(1-\psi^D)V_{t+1}^R(s',\beta) \\ \text{s.t.} \quad & P_t c(1+\tau_c) + P_t^s s' = P_t \tau_t (\phi^R)^{1-\Upsilon} + (P_t^s + D_t)s. \end{aligned} \tag{16}$$

Retirees face mortality risk and make consumption decisions given pension income.

Job search problems. Let $f(\theta_t)$ be the worker contact rate per unit of search efficiency. The value of job search for an unemployed worker is given by:

$$\Omega_t^U(s,h,\beta) = f(\theta_t)\,\mathbb{E}_x V_t^E\left(s,h,x,\underline{x}/x,\beta\right) + \left(1-f(\theta_t)\right)V_t^U(s,h,\beta), \tag{17}$$

where job search efficiency of the unemployed is set to one. The value of OJS is

$$\Omega_t^E(s,h,x,\alpha,\beta) = \nu_t f(\theta_t)\,\mathbb{E}_{\widetilde{x}}\left[\max\left\{V_t^E(s,h,\widetilde{x},x/\widetilde{x},\beta), V_t^E(s,h,x,\max\{\alpha,\widetilde{x}/x\},\beta)\right\}\right]$$
$$+ (1-\nu_t f(\theta_t))\,V_t^E(s,h,x,\alpha,\beta), \tag{18}$$

where ν_t is the OJS efficiency. Upon contact, the worker-firm pair takes expectations with respect to the sampling distribution Γ^x. The first term inside the expectation is the worker's value when she switches to a new job with productivity $\widetilde{x}$ and new piece rate $\alpha' = x/\widetilde{x}$. The second term is her value of staying with the incumbent firm, with either the current piece rate α (if $\widetilde{x} < \alpha x$) or a higher piece rate $\widetilde{x}/x$ (if $\widetilde{x} > \alpha x$).

3.3 Production

The final-good firm and intermediate firms are the same as in Section 2 except that the production function of intermediate firms is given by $y_t(j) = z_t l_t(j)$. The NKPC is identical to Equation (8) except that the real marginal cost is now $mc_t = p_t^l/z_t$.

Labor services. Tightness θ_t is the ratio of vacancies v_t to the measure of job search $S_t = \int d\mu_t^U(s,h,\beta) + \nu_t \int d\mu_t^E(s,h,x,\alpha,\beta)$, where μ^U and μ^E are the distributions of unemployed and employed workers at the search stage within the period.

The real value of a labor services firm that employs a worker with human capital h and piece rate α in a match with productivity x is given by

$$J_t(h,x,\alpha) = p_t^l F(h,x) - w(h,x,\alpha) + \frac{1}{1+r_{t+1}}(1-\psi^R)(1-\delta) \tag{19}$$
$$\times\, \mathbb{E}_{h'|h}\Big\{ (1-\nu_t f(\theta_{t+1}))\, J_{t+1}(h',x,\alpha) + \nu_t f(\theta_{t+1}) \int_{\underline{x}}^{x} J(h',x,\max\{\alpha,\widetilde{x}/x\})\, d\Gamma^x(\widetilde{x})\Big\},$$

where the match survives if the worker does not retire, does not lose her job, and does not find a new job. The worker accepts a new job if $\widetilde{x} > x$, in which case the firm's value is 0. If $\widetilde{x} < x$, then the firm keeps the worker either at a higher piece rate $\widetilde{x}/x$ (if $\widetilde{x} > \alpha x$) or at the current piece rate α (if $\widetilde{x} < \alpha x$). As firms are owned by the risk-neutral mutual fund, they discount the future at the real rate r_{t+1}.

The real value of posting a vacancy is

$$V_t = -\kappa + q(\theta_t)\frac{1}{S_t}\left[\int_{s,h,\beta}\int_{\widetilde{x}} J_t(h,\widetilde{x},\underline{x}/\widetilde{x})\, d\Gamma^x(\widetilde{x})\, d\mu_t^U(\cdot)\right. \tag{20}$$
$$\left. + \nu_t \int_{s,h,x,\alpha,\beta}\int_{x}^{\overline{x}} J_t(h,\widetilde{x},x/\widetilde{x})\, d\Gamma^x(\widetilde{x})\, d\mu_t^E(\cdot)\right] = -\kappa + q(\theta_t)\,\mathbb{E}J_t,$$

where the first and second terms capture the value of filling a vacancy out of unemployment and employment, respectively. A free-entry condition implies that $V_t = 0$.

Mutual fund. The mutual fund issues shares at price P^s, owns all firms and holds all government debt. A no-arbitrage condition equalizes the share and bond returns:

$$\frac{P^s_{t+1} + D_{t+1}}{P^s_t} = 1 + i_t. \tag{21}$$

The mutual fund is not allowed to retain any funds. All balances (positive or negative) are distributed to share owners in the form of dividends given by

$$D_t = B_{t-1} - \frac{B_t}{1+i_t} + P_t\Gamma^I_t + P_t\Gamma^S_t, \tag{22}$$

where the total real profits of intermediate and labor services firms are given by:[30]

$$\Gamma^I_t = \left(1 - \frac{p^l_t}{z_t} - \frac{\eta}{2\vartheta}\log(1 + \pi_t - \pi^*)^2\right) Y_t, \tag{23}$$

and

$$\Gamma^S_t = \int \left(p^l_t F(h,x) - w(h,x,\alpha)\right) d\lambda^E_t(s,h,x,\alpha,\beta), \tag{24}$$

where $\lambda^E_t(\cdot)$ is the distribution of employed at the end of the period. Equation (22) clarifies that the mutual fund collects payments for the existing debt obligations B_{t-1}, collects profits of intermediate and labor services firms Γ^I_t and Γ^S_t, and finances all new debt issuance B_t. The remaining balance accrues to the individuals as dividends.

Fiscal authority. The fiscal authority taxes individuals and issues bonds to finance G_t as well as UI and retirement benefits. The government budget constraint is

$$\begin{aligned} B_{t-1} + G_t + P_t\int UI(h)d\lambda^U_t(\cdot) + P_t\int \phi^R d\lambda^R_t(\cdot) = &\frac{B_t}{1+i_t} + P_t\tau_c\int c(e,s,h,x,\alpha,\beta)d\lambda_t(\cdot) \\ &+P_t\int\left(UI(h) - \tau_t UI(h)^{1-\Upsilon}\right)d\lambda^U_t(\cdot) \\ &+P_t\int\left(w(h,x,\alpha) - \tau_t w(h,x,\alpha)^{1-\Upsilon}\right)d\lambda^E_t(\cdot) \\ &+P_t\int\left(\phi^R - \tau_t(\phi^R)^{1-\Upsilon}\right)d\lambda^R_t(\cdot), \end{aligned} \tag{25}$$

where the left-hand side is total government expenses and the right-hand side is total government revenues generated from collecting taxes and issuing bonds. Here, $\lambda_t(\cdot)$, $\lambda^E_t(\cdot)$, $\lambda^U_t(\cdot)$, and $\lambda^R_t(\cdot)$ are the distributions of all, employed, unemployed, and retired individuals over relevant state variables at the end of the period, respectively.

Monetary authority. A monetary authority controls the short-term nominal rate following the Taylor rule (Equation 5), and the real rate satisfies the Fisher equation (Equation 6). The timing convention for these variables are the same as in Section 2.

[30]We assume that vacancy cost κ is psychic in that it does not affect profits.

3.4 Equilibrium

Market clearing requires that labor services demanded by intermediate firms Y_t/z_t is equal to the aggregate supply of labor services and that mutual fund shares demanded by all individuals aggregate to one. Formally, these conditions are:

$$Y_t/z_t = \int F\left(h, x\right) d\lambda_t^E(\cdot), \tag{26}$$

$$1 = \int g_t^{Us}(s, h, \beta) d\lambda_t^U(\cdot) + \int g_t^{Es}(s, h, x, \alpha, \beta) d\lambda_t^E(\cdot) + \int g_t^{Rs}(s, \beta) d\lambda_t^R(\cdot), \tag{27}$$

where g_t^{es} denotes the saving decision of workers with employment status $e \in \{E, U, R\}$.

In equilibrium, tightness θ and real price of labor services p^l are jointly determined to satisfy the free-entry condition in Equation (20) and the market-clearing condition for labor services in Equation (26). This is because the supply of labor services is determined by the distribution of workers, which is governed by contact rate $f(\theta)$. Raising the supply of labor services requires an increase in θ, which is induced by a rise in p^l through the free-entry condition.

In steady state, we assume that tax parameter $\tau_t = \tau^*$ clears the government budget constraint and that outstanding debt and government expenditures are a fraction of the steady-state level of output $B^* = x_B Y^*$ and $G^* = x_G Y^*$, respectively. We define the model's stationary equilibrium in Appendix B.1.1, describe the solution of the steady state in Appendices B.1.2 and B.1.3, and detail the computation of the economy's transitional dynamics in Section 4.2 and Appendix B.3.

4 Calibration and estimation

We now calibrate the steady state of the model to match moments of the U.S. over 2004–2006. Then, we describe the estimation of the shock processes for discount factor β, aggregate productivity z, and OJS efficiency ν. A model period is a quarter.[31]

4.1 Calibration of parameters

Functional forms and external calibration. Table B.1 in Appendix B.2 summarizes the externally calibrated parameters. The utility function over consumption is of the CRRA form with $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$ where we set $\sigma = 2$. As for the life cycle, workers spend on average 40 years in the labor force and 20 years in retirement, which require setting $\psi^R = 0.625$ percent and $\psi^D = 1.25$ percent on a quarterly basis.

[31] We pick a quarterly frequency because some data (e.g., output and ULC) are available quarterly.

We use three equally-spaced grid points for the discount factor such that $\beta \in \left\{\bar{\beta} - \Delta\beta, \bar{\beta}, \bar{\beta} + \Delta\beta\right\}$ and assume that Γ^{β} is a uniform distribution. For worker skills, we use five equally-spaced (in log space) grid points between the lowest value $\underline{h} = 1$ and the highest value $\overline{h} = 3$. These choices imply that skills change by a proportion $\Delta h = (ln(3) - ln(1))/4 = 0.275$ between grid points. We discipline the probability of skill appreciation for the employed π^E by the annual wage growth of job stayers. Karahan et al. (2022) document that this is around 2 percent, which implies that expected quarterly wage growth of job stayers is around 0.5 percent, which requires setting $\pi^E = 0.005/0.275 \approx 0.018$. We assume that Γ^h is such that individuals start their life at the lowest skill $\underline{h}$. We further assume that the match productivity x is drawn from a log-normal distribution Γ^x with standard deviation σ_x (to be discussed below). We discretize this process with 7 equally-spaced grid points (again in log space) between the 1st and 99th percentiles of the log-normal distribution.

We pick a CES matching function so that the worker and firm contact rates are bounded between 0 and 1, and given by $f(\theta) = \theta(1 + \theta^{\xi})^{-1/\xi}$ and $q(\theta) = (1 + \theta^{\xi})^{-1/\xi}$, respectively. We set the elasticity parameter to $\xi = 1.6$ following Schaal (2017).

The elasticity of substitution across intermediate goods η controls the markup of prices over the marginal cost and therefore the profit share. We set $\eta = 6$ to obtain a profit share of $\eta/(\eta - 1) - 1 = 20$ percent (Auclert et al., 2021). We normalize the productivity of the intermediate sector to $z = 1$ at the steady state. Finally, as Equation (8) shows, the price adjustment cost parameter ϑ directly dictates the slope of the Phillips curve, which we set to 0.021 as estimated by Galí and Gertler (1999).

As Ricardian equivalence does not hold in our model, fiscal policy matters for how the economy responds to shocks. Over our calibration period of 2004–2006, the ratio of government spending to GDP was around 19 percent, so we set $x_G = 0.19$. We also target a realistic amount of government debt. In the data, the ratio of debt stock to annual GDP averages to 60.8 percent over the same period. The quarterly calibration of the model dictates that we set $x_B = 4 \times 0.608 = 2.43$. We set the consumption tax to $\tau_c = 3.02$ percent, which we obtain as the ratio of state and local sales tax receipts to personal consumption expenditures for 2006. We set the parameter governing the progressivity of income taxes to $\Upsilon = 0.151$, as in Heathcote et al. (2014).

The central bank targets 2 percent annual inflation. Our quarterly calibration requires to set $\pi^* = 1.02^{1/4} - 1 \approx 0.496$ percent. For the Taylor rule, we follow Taylor (1993) and Galí (2015), and set the coefficient on inflation to $\Phi_\pi = 1.5$. Galí (2015)

Table 1: Internally calibrated parameters

Parameter	Explanation	Value	Target	Data	Model
$\bar{\beta}$	Middle value of discount factor	0.939	Income-weighted average MPC	0.20	0.20
$\Delta\beta$	Gap between discount factor values	0.050	Average wealth/annual income	4.10	4.89
κ	Vacancy creation cost	0.636	Unemployment rate	0.051	0.049
δ	Job separation probability	0.102	EU rate	0.038	0.033
ν	Search efficiency of employed	0.097	EE rate	0.020	0.019
π^U	Skill depreciation probability	0.107	Earnings drop upon job loss	-0.35	-0.35
σ_x	Standard deviation of productivity	0.068	Wage growth of job switchers	0.09	0.10
ϕ^E	Maximum share of output as wages	0.822	Labor share	0.67	0.73
ϕ^U	UI replacement rate	0.406	UI replacement rate	0.40	0.46
ϕ^R	Retirement benefit amount	0.412	Retirement income/labor income	0.34	0.38

Notes: This table summarizes the internally calibrated model parameters. See the main text for a detailed discussion.

sets the the coefficient on the output gap to 0.125. To map this to the unemployment gap, we use an Okun's law coefficient of -2 (Okun, 1962) yielding $\Phi_u = -0.25$.

Internal calibration. The remaining ten parameters to be jointly determined are the middle value of the discount factor $\bar{\beta}$, the distance between adjacent β values $\Delta\beta$, vacancy cost κ, job separation probability δ, OJS efficiency ν, skill depreciation probability when unemployed π^U, the standard deviation of the match productivity distribution σ_x, maximum share of output potentially paid to workers as wages ϕ^E, UI replacement rate ϕ^U, and retirement benefit amount ϕ^R. These parameters are calibrated jointly by matching a set of data moments. Table 1 summarizes these calibrated parameter values next to their most informative target moment.[32]

To discipline $\bar{\beta}$, we target an income-weighted average quarterly MPC value of 0.2, as in Auclert et al. (2020) and Auclert et al. (2024).[33] We use $\Delta\beta$ to target an average wealth to annual income ratio of 4.1, as in Kaplan and Violante (2022).

On the labor market side, we target a steady-state unemployment rate of 5.1 percent, as well as labor market transition rates. Using data from the Current Population Survey (CPS), we compute the average monthly employment-to-unemployment rate over 2004–2006. We convert this monthly rate to a quarterly frequency and obtain our target of 3.8 percent. Using data from the Longitudinal Employer-Household Dynamics (LEHD), we find that the quarterly EE rate, measured as the job-switching rate of workers who do not have any intervening nonemployment spell, is 2 percent over the same period. These three moments are informative about the vacancy creation cost κ, job separation rate δ, and employed search efficiency ν, respectively.

[32] We use the simulated method of moments, where we minimize the sum of squared percentage deviations of the model moments from their empirical counterparts.

[33] We follow the approach in Kaplan and Violante (2022) where we calculate individual level MPCs out of the model equivalent of a \$500 windfall to their wealth.

The probability of skill depreciation when unemployed π^U governs the magnitude of earnings drop upon job loss. A large literature has estimated earnings losses upon displacement using a variety of datasets (see, Jacobson, LaLonde, and Sullivan, 1993; Stevens, 1997; Davis and von Wachter, 2011; Birinci, 2021; Jarosch, 2023, among others). Across these studies, the median estimate of earnings loss in the year of displacement is 35 percent.[34] To this end, we simulate a panel of agents, aggregate quarterly observations to an annual frequency, and run a distributed-lag regression on the simulated data analogously with the empirical studies. Hence, our simulation captures both employment and wage losses upon unemployment, as in the data.

Wage changes upon job switching affect the role of EE flows in aggregate demand. Using publicly-available state-level data from the LEHD, we calculate the rise in earnings from a job switch for continuously employed workers to be 9 percent. This moment informs the dispersion parameter for match productivity σ_x, which governs the wage rise upon a job switch. Given σ_x, we set the location parameter to $\mu_x = -\sigma_x^2/2$ so that its mean is normalized to one. Finally, we choose the maximum output share paid to workers ϕ^E to target an average labor income share of 0.67.

Turning to the fiscal side, we calibrate the UI replacement rate ϕ^U to match an average replacement rate of 40 percent. To discipline retirement pensions ϕ^R, we calculate the average retirement income to average labor income ratio in the Survey of Income and Program Participation (SIPP). Specifically, we add up Social Security Income and pension incomes for each retiree in our sample and calculate the average per-person retirement income. We then divide this measure by the average labor income among non-retirees to obtain a target ratio of 0.34.

4.2 Solving for transitional dynamics

We now summarize how we solve the model under aggregate uncertainty, relegating further details to Appendix B.3. We use the SSJ method developed by Auclert et al. (2021), which allows us to efficiently compute impulse responses. Importantly, we apply the SSJ method in a context where model blocks *directly* interact not only via aggregate variables but also through the discretized worker distributions, as shown by the DAG representation of the model in Figure B.1. We do this by treating the

[34] We follow the literature in targeting short-term earnings losses upon displacement to discipline the probability of human capital depreciation, e.g., as in Huckfeldt (2022). To generate persistent income losses upon job loss, the literature turns to long and slippery job ladders, e.g., as in Krolikowski (2017) and Jarosch (2023) or forces like congestion in hiring, e.g., Mercan et al. (2024).

worker distributions as histograms and calculate Jacobians for each mass point, which significantly increases the computation time given the large state space, but still remains feasible. This modification is crucial because outcomes of the HA block include distributions of agents, which are required inputs for labor service firms and equilibrium conditions. First, the distribution of employed across human capital and match productivity and the distribution of unemployed across human capital at the job search stage, i.e., $\mu^E(h,x)$ and $\mu^U(h)$, affect the expected match value $\mathbb{E}J$ for firms deciding on vacancy creation. This is because (i) human capital affects match output and (ii) employed workers' match productivity affects their job-acceptance decisions and the piece rate that a poaching firm would offer to the worker. Second, the distribution of employed across human capital, match productivity, and piece rate at the production stage within a period, $\lambda^E(h,x,\alpha)$, affects service firms' profits Γ^S by determining the output and wage levels in a match, which in turn affect dividends.[35]

In sum, we apply the SSJ method incorporating discretized worker distributions along the DAG, combining the SSJ method of Auclert et al. (2021) with Reiter (2009).

4.3 Estimation of aggregate shocks

We now use our calibrated model to estimate aggregate shock processes that we then use in our positive and normative analysis in Sections 5 and 6. The economy is subject to demand, supply, and labor market shocks, which are modeled as innovations to the discount rate β, aggregate labor productivity z, and OJS efficiency ν.[36] We assume AR(1) processes for all the shocks given by:

$$\beta_t = (1-\rho_\beta)\beta^* + \rho_\beta\beta_{t-1} + \sigma_\beta\epsilon_{\beta,t},\ z_t = (1-\rho_z)z^* + \rho_z z_{t-1} + \sigma_z\epsilon_{z,t},\ \nu_t = (1-\rho_\nu)\nu^* + \rho_\nu\nu_{t-1} + \sigma_\nu\epsilon_{\nu,t},$$

where ρ_j denotes the persistence of the AR(1) process, $\epsilon_j \sim N(0,1)$ is i.i.d., and $\sigma_j > 0$ denotes the standard deviation of innovations for $j \in \{\beta, z, \nu\}$.

Off the steady state, the government adjusts debt to satisfy its budget constraint. Thus, along a transition, debt can deviate from its steady-state level $b^* = x_B Y^*$. In such cases, the fiscal authority follows an exogenous rule to adjust τ to gradually bring the real debt level back to steady state b^*.[37] This response is governed by

$$\tau_t = \tau^* - \rho_\tau \left(b_{t-1} - b^*\right)/Y^*, \tag{28}$$

[35] If we were to assume directed search in the labor market, block-recursivity would allow us to drop $\mu^E(\cdot)$ and $\mu^U(\cdot)$. However, we would still need to keep track of $\lambda^E(\cdot)$ to calculate firm profits.

[36] Specifically, we introduce shocks to the middle discount factor $\bar{\beta}$ without changing $\Delta\beta$ such that the low and high discount factors scale with the middle value.

[37] Absent this countervailing force, the model would have a unit root resulting in debt exploding.

Table 2: Estimation of aggregate shocks: Data vs model moments

	Data			Model		
	Std. dev	Autocorr.	Corr. w/ Y	Std. dev	Autocorr.	Corr. w/ Y
Y	0.017	0.926	1	0.008	0.942	1
u	0.130	0.942	-0.626	0.136	0.889	-0.675
EE	0.046	0.700	0.202	0.047	0.567	0.204
θ	0.268	0.923	0.661	0.080	0.317	0.535
π	0.304	0.749	0.328	0.351	0.777	0.331

Notes: This table compares outcomes in the model and the data using the estimated shock processes. Y u, EE, θ, and π denote output, unemployment rate, EE rate, labor market tightness, and inflation.

where $\rho_\tau = 0.1$, as in Auclert et al. (2020), controlling tax sensitivity to debt.[38]

We estimate parameters of the AR(1) processes for β, z and ν (in total six parameters for the persistence and standard deviations for each) by targeting the autocorrelation of output, correlations of the unemployment rate, EE rate, and inflation with output, as well as the standard deviations of output, unemployment rate, EE rate, and inflation in the data.[39] This estimation yields $\rho_\beta = 0.912$, $\rho_z = 0.834$, $\rho_\nu = 0.975$ and $\sigma_\beta = 0.0015$, $\sigma_z = 0.0027$, $\sigma_\nu = 0.0009$.

Table 2 compares model and data moments. The model generates a nearly perfect match for the correlations of unemployment rate, EE rate, and inflation with output, as well as the autocorrelations of output, unemployment rate, and inflation. Standard deviations of the unemployment rate, EE rate, and inflation are also close to their empirical values. However, the model is less successful in matching the volatilities of output and tightness as well as the autocorrelation of tightness.

To gauge the contribution of each shock to the cyclicality of these moments, Table B.2 in Appendix B.4 presents a variance decomposition. We find that shocks to β explain 73.5 percent of unemployment rate fluctuations, while shocks to z and ν jointly account for the remaining 26.5 percent. We also find that shocks to ν are

[38] As mentioned in Footnote 29, τ_t is inversely related to the average tax rate. The negative sign in front of ρ_τ means that whenever $b_{t-1} > b^*$, we have $\tau_t < \tau^*$, i.e., taxes increase.

[39] We obtain monthly data on the unemployment rate and PCE price index from the BLS. We take quarterly averages for the unemployment rate. We calculate 12-month inflation and convert it to a quarterly rate, and take quarterly averages. We take quarterly data on real GDP from the U.S. Bureau of Economic Analysis; monthly data on the EE rate from Fujita, Moscarini, and Postel-Vinay (2023), which we convert to a quarterly frequency; and monthly data on the number of vacancies from JOLTS and divide that by unemployment, and take quarterly averages to obtain tightness. All data cover the period between 1995:Q3 and 2008:Q4. To calculate data moments, we take logs and detrend the time series of output (real GDP), unemployment rate, EE rate, and tightness using the HP filter with a smoothing parameter of 10^5 and calculate correlations and standard deviations of their cyclical components. Because inflation is negative in some periods, we detrend the level of inflation and then divide the cyclical component by the trend level to obtain percent deviations.

an important driver of the EE rate and inflation; they explain 46.3 percent of the fluctuations in the EE rate and 15.8 percent of the fluctuations in inflation.

What are OJS efficiency shocks? We showed that fluctuations in OJS efficiency have sizable effects on the EE rate and inflation. A natural question to ask is: What are possible micro-foundations for these shocks? For instance, Bilal et al. (2022) show that worsening financial frictions led to a decline in firm entry around the Great Recession. Importantly, they argue that the job ladder collapse during this episode was due to a decline in vacancy creation among firms whose net poaching rate is typically high. Along these lines, to microfound OJS efficiency shocks, our model can be extended to feature ex-ante heterogeneous labor-services firms, where the entry decision is subject to financial frictions. In this model, firms can choose to hire senior workers through poaching or junior workers out of unemployment who earn lower wages but are also less productive. Firms that are close to the borrowing constraint may opt to hire junior workers instead of poaching seniors. This would generate a depressed EE rate in a recovery despite a declining unemployment rate, a pattern consistent with the data during the episode that we analyze in Section 5.1.

While such examples provide guidance on how to think about changes in the OJS efficiency, we do not attempt to microfound OJS shocks given the complexity of our model and leave this for future research. We simply treat these shocks as a reduced-form time-varying wedge that occasionally breaks the negative correlation between the unemployment rate and the EE rate as in the data that we study in Section 5. We estimate such shocks using the model, and study their macroeconomic implications. We show that an upward shift in OJS efficiency causes higher inflation and unemployment, and hence—unlike demand or supply shocks—they break divine coincidence and result in a trade-off for monetary policy.[40]

5 Positive implications of job mobility on inflation

Now, we use our model to quantitatively study the economy's response to exogenous shifts in job mobility captured by OJS efficiency shocks. These shocks are disciplined by two episodes we document in Section 5.1: (i) a stable EE rate despite a large decline

[40]Demand and productivity shocks typically move inflation and unemployment gaps in the opposite directions—absent time varying real rigidities, which our model features—leading to divine coincidence. OJS efficiency shocks move inflation and unemployment gaps in the same direction, breaking divine coincidence, and introducing a trade-off for monetary policy. In this sense, OJS efficiency shocks act like cost-push shocks.

Figure 4: Unemployment and EE rates, unit labor cost, and inflation over time

(a) Unemployment and EE rates (b) Unit labor cost, inflation, unemployment

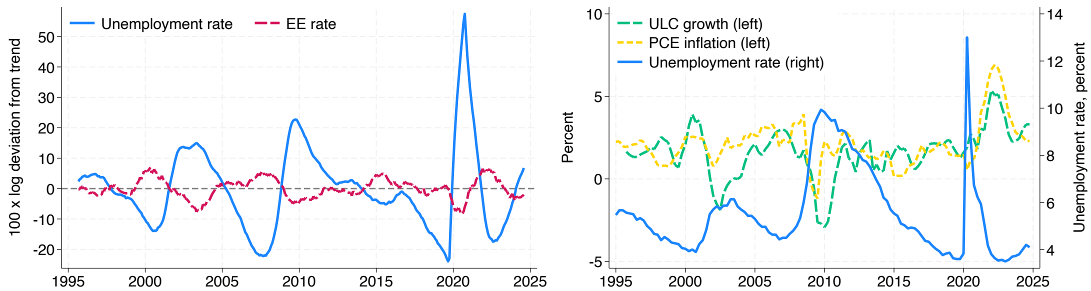

Notes: Panel (a) plots the cyclical components of the log unemployment rate and EE rate (Fujita, Moscarini, and Postel-Vinay, 2023). Both series are detrended using the HP filter with a smoothing parameter of 10^5 and we take 12-month centered moving average for visual clarity. Panel (b) plots the annual growth rate of the quarterly ULC index and the quarterly average of the 12-month PCE inflation rate (left axis), and the quarterly average of the monthly unemployment rate (right axis). We obtain the quarterly ULC index for nonfarm business sector from the BLS. For visual clarity, we take a four-quarter moving averages for all three series.

in unemployment during the post-Great Recession expansion and (ii) a sizable rise in the EE rate post COVID-19. In Sections 5.2 and 5.3, we quantify (i) the drag on inflation due to muted job mobility, relating this to the "missing inflation" puzzle in the first episode, and (ii) the rise in inflation driven by elevated job mobility during the latter. In Section 5.4, we decompose the channels through which OJS efficiency shocks affect inflation. Finally, in Section 5.5, we reiterate the importance of modeling market incompleteness and discuss how it affects our quantitative results.

5.1 EE and labor cost dynamics in the last two recoveries

We start by an investigation of the empirical relationship between the unemployment and EE rates, with a particular focus on the recovery following the last two recessions. Figure 4 Panel (a) plots the cyclical components of these two labor market variables over time. Typically, periods of low unemployment and a tight labor market are associated with high EE rates. The 2016–2019 period, however, displays a breakdown in this historical correlation, which can be traced to a flat EE rate despite a roughly 30 percent decline in the unemployment rate from trend. This is in contrast to the recovery from the COVID-19 recession, when the unemployment rate declined by almost the same amount while the EE rate increased by around 8 percent above its trend. Figure B.2 in Appendix B.5 provides complementary charts, showing that the change in the correlation between unemployment and EE rates during the 2016–2019 episode is neither an artifact of detrending nor picking a specific time window.

The goal of our quantitative exercises is to study how OJS efficiency shocks change the marginal cost of production p^l/z and, ultimately, inflation. The closest data counterpart to p^l/z is the unit labor cost (ULC), an index adjusted for worker composition and productivity. Panel (b) presents the time-series of ULC growth. During the recovery from the COVID-19 recession, when the unemployment rate was below 4 percent, ULC growth reached around 6 percent and the average annual PCE inflation was 7 percent. However, despite a similar decline in unemployment between 2016 and 2019, ULC growth was around 2 percent and inflation was 1.5 percent. These differences in ULC growth and inflation during the last two recoveries cannot be accounted for by unemployment dynamics alone. Instead, differential EE dynamics may have played a crucial role. Figure 4 illustrates that a key distinction between the two episodes was in EE dynamics. We now use our model to show that the differences in ULC growth and inflation over these episodes can be attributed to differential EE dynamics.

5.2 Missing inflation and worker mobility during 2016–2019

We first simulate the Great Recession recovery and quantify the inflation that is "missing" due to muted job mobility. We compare two economies that mimic the path of the unemployment rate but *differ* in their EE rates. The first economy features an endogenously rising EE rate due to a tightening labor market driven by positive demand shocks, consistent with its negative historical correlation with the unemployment rate but inconsistent with observation. The second economy—further subject to negative OJS shocks—not only features the same unemployment path but also replicates the stable EE rate in this episode.

Starting from steady state, we let two types of shocks to hit the economy starting in 2016.[41] We model demand and OJS efficiency shocks as innovations to the discount factor β and OJS efficiency ν, respectively, following AR(1) processes given by:[42]

$$\beta_t = (1-\rho_\beta)\beta^* + \rho_\beta \beta_{t-1} + \varepsilon_{\beta,t}, \qquad \nu_t = (1-\rho_\nu)\nu^* + \rho_\nu \nu_{t-1} + \varepsilon_{\nu,t}. \tag{29}$$

To simulate the first economy, we turn off OJS shocks ($\varepsilon_{\nu,t} = 0$) and back out the path of $\varepsilon_{\beta,t}$ to generate a desired decline and a subsequent mean reversion in unemployment. Specifically, we assume that the unemployment rate declines by 23 percent relative to its steady state level of 4.9 percent linearly over $\overline{T} = 16$ quarters (corresponding to 2016–2019).[43] Upon reaching its trough, it mean reverts geometrically

[41] That the economy is in steady state in 2016 is plausible from the labor market perspective. The unemployment rate in 2016 was around 5 percent, close to the estimates of its natural rate.

[42] In Equation (29), we use the estimated persistence parameters ρ_β and ρ_ν in Section 4.3.

[43] This is consistent with the decline in unemployment explained by the increase in the job-finding rate alone between 2016 and 2019, holding the separation rate fixed at its January 2016 level.

Figure 5: Effects of muted worker mobility on aggregates: Post-Great Recession

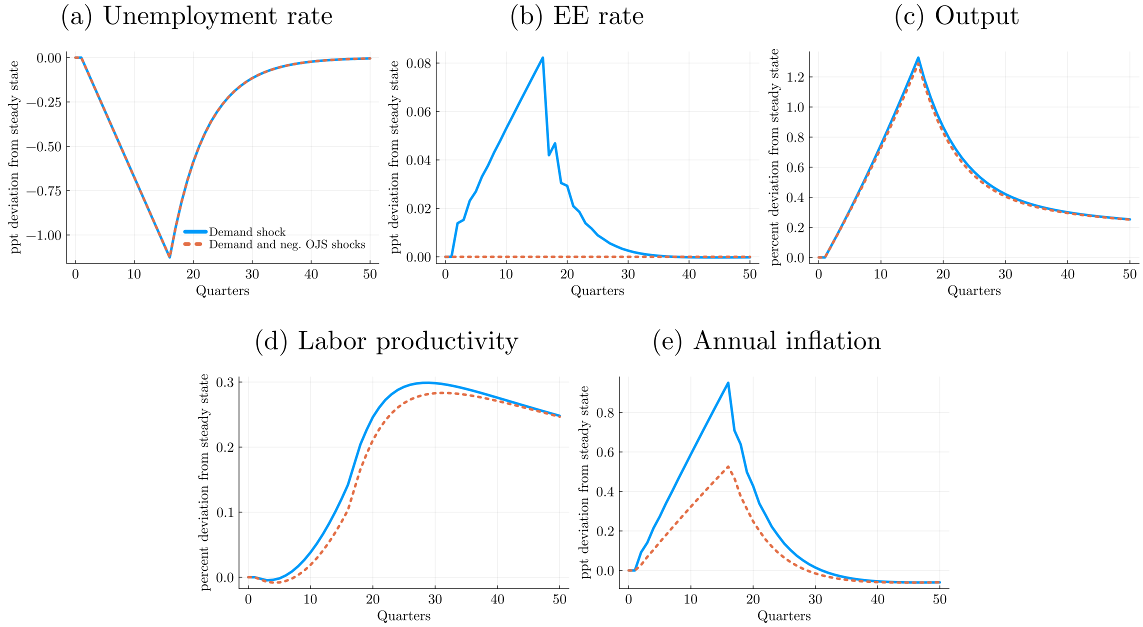

Notes: This figure presents the dynamics of unemployment rate, EE rate, output, average labor productivity, and annual inflation in an economy subject to (i) a series of positive demand shocks (solid-blue lines) and (ii) series of positive demand shocks and negative OJS efficiency shocks (dashed-orange lines). The shocks in the two economies are estimated to generate the same path of unemployment. The EE rate is untargeted in the first economy whereas the OJS efficiency shocks are such that the EE rate remains unchanged in the second economy as in the data.

at rate ρ_u. In this scenario, movements in the EE rate arise solely from endogenous responses dictated by the model. The second scenario turns on OJS efficiency shocks to mimic the actual expansionary episode during 2016–2019. We jointly estimate the paths of $\varepsilon_{\beta,t}$ and $\varepsilon_{\nu,t}$, such that the economy generates the same unemployment rate path as in the first economy, but additionally has the EE rate unchanged throughout.

We compare these two economies with identical unemployment rates but differing EE rates to quantify how much of the inflation gap arise solely due to differences in EE dynamics. This exercise directly addresses the question we are after: What would have been inflation after the Great Recession if the EE rate had increased following its historically negative correlation with the unemployment rate?[44]

Figure 5 presents the results.[45] The two economies have identical unemployment paths (Panel (a)) and differ in their EE rates (Panel (b)) by construction: In the first

[44]An alternative exercise is to first estimate paths of $\varepsilon_{\beta,t}$, $\varepsilon_{\nu,t}$, and $\varepsilon_{z,t}$ to match (smoothed) time series of several data moments and then to turn off these shocks one at a time to understand how the model matches the data. This exercise would quantify the change in inflation in the absence of, say, OJS efficiency shocks. However, this departs from the comparison we seek to make—turning off this shock would also lead to changes in unemployment, output, productivity, etc.

[45]Figure B.3 Panel (a) plots the paths of the estimated shocks for each economy. Figures B.4 and B.5 further show the impulse responses of model outcomes to shocks to β and ν, respectively.

economy, positive demand shocks increase vacancies and, consequently, the EE rate. In the second economy, negative OJS efficiency shocks keep the EE rate suppressed as in the data. The first economy produces slightly more output (Panel (c)), reflected in average labor productivity (ALP) since the path of (un)employment is identical across the two economies. ALP initially declines slightly as the rise in the job-finding rate brings more of the unemployed into employment. These workers often have lower human capital and accept jobs at any level of productivity. ALP eventually rises as human capital and match productivity improve (Panel (d)). While the increase is slightly larger in the first economy due to the rise in productivity-improving job switches, the gap is small as the productivity distribution evolves slowly (Figure B.6).

If we were to solely focus on the joint dynamics of unemployment and output, and ignore job mobility, we would have inferred that the two economies were hit by very similar shocks. However, looking at inflation alters this conclusion. In particular, annual inflation is 0.42 percentage points (pp) smaller in the economy with a constant EE rate (Panel (e)).[46] This is not a small quantitative effect, as this gap implies that inflation would have been around 1.8 percent instead of 1.4 percent in 2019 had the EE rate increased commensurately with the decline in unemployment.[47,48]

5.3 High inflation and worker mobility during 2021–2022

We now study the role of the high EE rate in the rise in inflation during the recovery from the COVID-19 recession. We proceed similarly to the previous exercise.

Here, the first economy features positive demand shocks and positive OJS efficiency shocks to emulate the path of unemployment and EE rates in this episode.[49]

[46] To obtain this number, we calculate the annual inflation rate in each economy and report their maximum difference, which materializes 16 quarters after, when unemployment is at its lowest level.

[47] Bostanci, Koru, and Villalvazo (2022) and Pilossoph and Ryngaert (2022) document that higher inflation expectations might lead to higher job-to-job transitions as workers try to escape inflation by moving to better-paying jobs. It is feasible to incorporate this channel into our framework by introducing (i) nominal wage rigidity, (ii) endogenous job search effort, and (iii) by preventing offer matching for a fraction of external offers. In such a model, an employed worker would increase her search effort in episodes of high inflation under rigid nominal wages to get a better-paying outside offer. If such an offer is not matched by the incumbent, the worker would switch jobs. In this case, the missing inflation explained by our model would be even larger because the increase in EE rate in the first economy would be larger. Therefore, we view our estimate as a lower bound.

[48] Inflation in 2019 was expected to be above 2 percent as the unemployment rate was historically low. Thus, the muted EE rate explains a sizable portion but not all of the missing inflation.

[49] Since the unemployment dynamics in this episode are almost the same as the post-Great Recession, we target the same unemployment path as in Section 5.2. We also find that positive demand shocks alone generate a smaller rise in the EE rate than that was observed during the COVID-19

Figure 6: Effects of elevated worker mobility on aggregates: COVID-19 recovery

(a) EE rate (b) Annual inflation

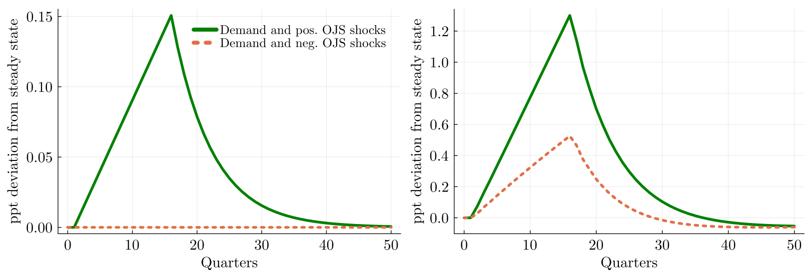

Notes: This figure presents the dynamics of the EE rate and annual inflation in an economy subject to (i) a series of positive demand shocks and positive OJS efficiency shocks (solid-green lines) and (ii) a series of positive demand shocks and negative OJS efficiency shocks (dashed-orange lines). The shocks in the two economies are estimated to generate the same path of unemployment as in Figure 5 Panel (a). The EE rate in the first economy targets an increase by around 8 percent as in the data, while the second economy targets a constant EE rate.

We inherit the second economy from the previous exercise with a flat EE rate. The rest of the details are identical to the previous section.[50]

Figure 6 presents the dynamics of the EE rate and inflation.[51] Panel (a) shows that the first economy (solid-green line) generates a 0.15 pp increase in the EE rate as in the data. Panel (b) shows that a higher EE rate leads to higher inflation. We find that annual inflation is 0.77 pp in the first economy, i.e., had the EE rate remained flat during the recovery following the COVID-19 recession (as it did during the 2016–2019 episode), inflation would have been 0.77 pp lower. Since annual inflation increased by 5.5 pp during this episode in the data, our result indicates that strong worker mobility accounts for 14 percent of the total increase in annual inflation.

Taking stock. Sections 5.2 and 5.3 show that a recovery episode with an increasing EE rate generates a larger increase in inflation compared to an episode with identical unemployment dynamics but a low EE rate. This implies that the empirical Phillips curve is flatter when the EE rate is sluggish—a result that we show in Figure B.7.[52]

recovery. As Figure 5 Panel (b) shows, under such shocks, the EE rate increases by 0.08 ppts, while the EE rate increased by 0.16 ppts in the data during 2021–2022.

[50]Figure B.3 Panel (b) plots the estimated path of shocks for these two economies.

[51]We only present the EE rate and inflation as the intuitions carry over from the previous section.

[52]We do not claim that the differential inflation dynamics between the two recovery episodes are entirely driven by their differential EE dynamics. For example, the underlying match distribution being different at the beginning of each episode might lead to the same shocks having different effects on inflation, which is mainly emphasized by Moscarini and Postel-Vinay (2022). Another contributing factor could be the transition to the so-called "flexible average inflation targeting" announced in August 2020, whereby the Federal Reserve became willing to temporarily accommodate an inflation above target while maintaining a longer-term average inflation around 2 percent.

5.4 Decomposing the inflationary effects of OJS shocks

The preceding case studies show that changes in job mobility have sizable effects on inflation, which are an outcome of several opposing forces that depend on the changes in match values and demand responses. We now decompose the channels through which a rise in OJS efficiency translates to higher inflation utilizing the model's DAG representation (see Figure B.8) and the system of Jacobians we compute to solve the model. Here, we summarize our results and relegate the details to Appendix B.6.

A higher OJS efficiency reduces the expected match value for service firms due to the composition and match duration channels, as discussed in the TANK model. In the HANK model, the match value additionally decreases through more frequent wage re-bargaining, where outside offers can raise wages without any rise in productivity. This decline in the match value requires an increase in the price of labor services (i.e., the marginal cost for intermediate firms) to satisfy the free-entry condition. This direct effect explains 164 percent of the increase in the marginal cost. However, a higher OJS efficiency also reduces market tightness because of (i) an increased supply of labor services as the match-productivity distribution improves and (ii) lower aggregate demand due to a higher real rate and unemployment, both of which dictate a lower tightness to clear the labor services market. All else the same, a lower tightness implies that firms fill vacancies faster and workers find it harder to find jobs, generating an increase in the expected match value. Therefore, the price of labor services declines to respect the free-entry condition. This decline in tightness explains -66 percent of the rise in the marginal cost. Thus, labor market effects explain 98 percent of the overall effect. The remaining 2 percent is due to changes in the real rate in response to inflation and unemployment. Figure B.9 shows the contribution of these channels to the total rise in p^l upon a positive unit shift in OJS efficiency ν.

5.5 Revisiting the importance of incomplete markets

Section 2 showed that, upon a positive OJS efficiency shock, a TANK model with imperfect self-insurance generates much lower declines in demand, output, and tightness, and a much larger rise in the real rate relative to a RANK model with perfect self-insurance. This is because demand is less sensitive to the real rate for the HtM household with a high MPC than the PIH household with a low MPC. Despite the differences between the HANK and TANK models, the same intuition holds here. We now explain the underlying mechanisms in the HANK model by analyzing how

Figure 7: Impulse responses of consumption to a positive OJS efficiency shock

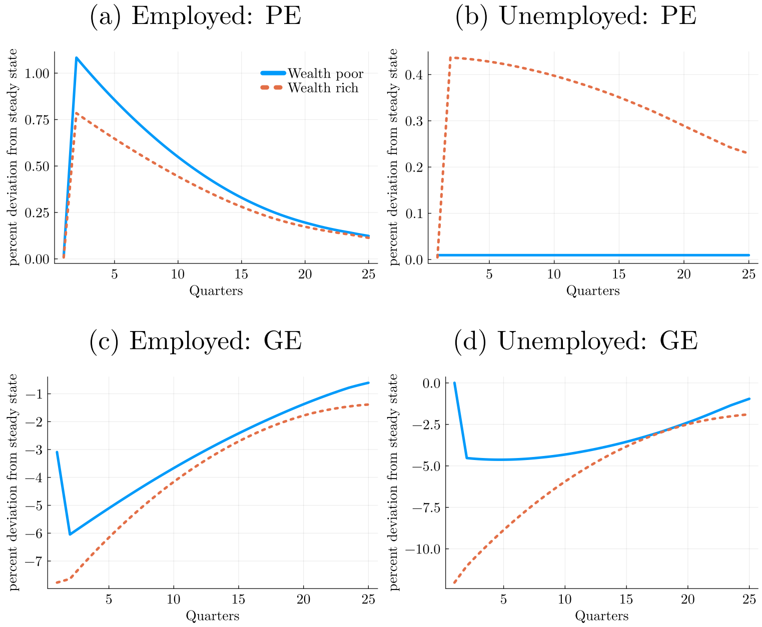

Notes: This figure presents how the consumption responses to a positive OJS efficiency shock differ across wealth and employment status in partial (PE) and general equilibrium (GE) settings. We compare the consumption responses of two individuals who have the same the human capital, and when employed, the same match-specific productivity and the piece-rate but differ only in their wealth. The wealth-poor individual has no wealth, and the wealth-rich individual holds shares close to the median of the wealth distribution.

the consumption response to a 1 percent increase in OJS efficiency that gradually recovers back to the mean differs across agents with the same state variables except their wealth. We study these IRFs for agents with different employment states, separately in partial equilibrium (PE) and in general equilibrium (GE).[53]

In PE, a persistent positive OJS efficiency shock affects consumption through several mechanisms. Employed workers anticipate faster wage growth from more frequent job offers and thus raise consumption. Additionally, with improved job mobility post-unemployment, the wage scar from job loss diminishes, lowering the precautionary savings motive—a force that is stronger for those with less wealth. Panel (a) of Figure 7 shows that consumption rises more among low-wealth employed agents. Among the unemployed (Panel b), however, it's the *wealth-rich* who increase consumption. Wealth-poor agents are constrained by limited access to borrowing and cannot smooth consumption despite higher expected income. This contrasts with standard HANK models, where poorer households with higher MPCs respond more

[53] PE is where we fix all equilibrium objects relevant for the budget constraint (e.g., prices, dividends, taxes, etc.) at their steady-state values and only allow for the individual consumption choices to respond to higher OJS efficiency. GE responses are the optimal consumption decisions for individuals that expect all budget-relevant variables to evolve endogenously.

to shocks. Here, since the shock mainly affects *future* income, current consumption rises only for those with wealth or sufficient income.

In GE, the response is markedly different. Panels (c) and (d) show that consumption falls across all groups, reducing aggregate demand and output (Figure B.5). A higher ν raises inflation and unemployment while tightening monetary policy—leading to higher real interest rates and lower demand. Notably, the wealth-rich cut consumption more. This is not due to taxes because both workers pay the same taxes as they have the same income. Moreover, if the decline in tightness was the primary driver, consumption would have fallen by more for the wealth-poor due to the precautionary motive. Instead, the larger decline in consumption for the wealth-rich stems from falling financial wealth, caused by declining real dividends and share prices, as higher labor costs reduce firm profits. Since wealthier agents rely more on financial assets to fund consumption, the decline in wealth triggers a larger cutback. The wealth-poor, by contrast, rely on labor income and are less affected by asset price changes.

Critically, this result implies that an alternative model approximating perfect self-insurance (i.e., a model populated only by low-MPC agents) would have generated a much larger decline in demand, output, tightness, a much larger rise in the real rate—a result that is consistent with our findings in Section 2.4. As such, our findings reveal novel mechanisms behind why accounting for market incompleteness matters for the positive implications of EE fluctuations. In the next section, we discuss the role of EE fluctuations and market incompleteness in the conduct of monetary policy.

6 Monetary policy under job mobility shocks

We have established that shifts in EE flows relative to the unemployment rate alter the relationship between unemployment and inflation. We now study whether explicitly accounting for EE flows as an additional proxy of economic slack matters for the conduct of monetary policy when the economy is subject to job mobility shocks.

To this end, we assume the economy is subject to OJS efficiency shocks estimated in Section 4.3. Recall that positive OJS efficiency shocks lead to higher unemployment and higher inflation, breaking divine coincidence and introducing a trade-off for monetary policy.[54] To quantify the gains from reacting to job mobility, we solve

[54]On the flip side, demand and productivity shocks typically lead to divine coincidence. To keep the discussion here focused, we consider an economy with OJS efficiency shocks only. However, the main conclusions remain similar if we incorporate demand and productivity shocks as well.

for optimal monetary policy within generalized Taylor rules—where the nominal rate reacts to inflation, unemployment, and the EE rate—under a dual-mandate central bank loss function. Then, to contextualize our results, we compare the implied nominal rates under the baseline vs optimal policy rules during the COVID-19 recovery episode. Finally, to gauge whether imperfect self-insurance matters for policy, we repeat this exercise in a calibration that approximates perfect self-insurance.

Central bank loss function. We start by positing that the central bank sets monetary policy to minimize the following loss function:

$$\mathcal{W} = \text{var}(\pi_t - \pi^*) + \Psi \text{var}(Y_t - Y^*), \tag{30}$$

which penalizes the quarterly inflation and output gap variances, for which we have two motivations.[55,56] First, our solution method relies on a first-order approximation, implying that it is not conducive to capturing the second-order effects of inflation on consumption and hence welfare. Therefore, we assume an objective function that explicitly involves inflation volatility. Second, this function approximates the dual mandate of the Federal Reserve and is commonly used in the literature. We set the relative weight of the output gap to $\Psi = 0.25$—a conventional value in the literature (see, for example, Jensen 2002 and Walsh 2003).[57]

Central bank reaction function. We study monetary policy within a class of Taylor-rule type monetary policy reaction functions in the following form:

$$i_t = i^* + \Phi_\pi (\pi_t - \pi^*) + \Phi_u (u_t - u^*) + \Phi_{EE} (EE_t - EE^*), \tag{31}$$

where Φ_{EE} governs the response of the central bank to the EE rate. This is a generalized version of Equation (5). We emphasize that this is a *feasible* policy, as both the EE and unemployment rates are measured at the same frequency using the CPS.[58]

Optimal policy in the baseline model. The optimal monetary policy consists of a combination of coefficients for the generalized Taylor rule in Equation (31) that

[55] Ideally, these gaps should be measured relative to their natural levels, which need not be constant with frictions (or productivity shocks). We use steady-state values for computational feasibility.

[56] Results remain similar if we replace the output gap with the unemployment gap and set $\Psi = 1$.

[57] Okun's law implies $u_t - u_t^* = \frac{Y_t - Y_t^*}{2}$, which would reduce the loss function to $\mathcal{W} = \text{var}(\pi_t - \pi^*) + \text{var}(u_t - u^*)$. An equal weight on inflation and unemployment is also consistent with how the Federal Reserve trades off inflation and unemployment (Yellen, 2012; Debortoli et al., 2018).

[58] To the extent that OJS efficiency shocks are persistent, monetary policy might be a blunt tool to manage the inflationary effects of job mobility. However, the largest effects of these shocks are upon impact (Figure B.5). Therefore, monetary policy is useful in managing these short-term effects.

Table 3: Optimal monetary policy and implied nominal interest rates

Taylor rule coefficients	Baseline	Optimal	Optimal (low MPC)
Inflation π	1.5	1.5	1.5
Unemployment rate u	-0.25	-0.125	-0.125
EE rate EE	0	1.125	0.625
Nominal interest rate i^*	2.13	2.22	2.11

Notes: This table presents Taylor rule coefficients under the baseline rule, the optimal rule, and the optimal rule in a calibration that approximates perfect self-insurance (i.e., a low average MPC). The bottom row compares the nominal interest rate implied by these rules when we feed data on π, u, and EE from 2022:Q2 when inflation peaked.

minimizes the loss function in Equation (30).[59] As our focus is on the response of monetary policy to the labor market, we keep the coefficient on inflation at its baseline value of $\Phi_\pi = 1.5$.[60] We find that optimal policy prescribes $\Phi_u^* = -0.125$ and $\Phi_{EE}^* = 1.125$, implying a less aggressive response to the unemployment gap than the calibrated value of $\Phi_u = -0.25$ (and $\Phi_{EE} = 0$), which is oft-used in the literature, and a strong positive response to the EE gap. In this case, the central bank loss shrinks by 95 percent relative to the baseline Taylor rule.

In practice, the optimal policy implies that the central bank should separate recovery episodes where job mobility is high from those where job mobility is low. For instance, despite the decline in the unemployment rate being similar between the last two recovery episodes, the rise in the nominal rate should have been more aggressive under the optimal policy between 2021 and 2022 when job mobility was much higher.

We further illustrate the implications of the optimal policy for the aftermath of the COVID-19 recession that saw high inflation. We focus on 2022:Q2 when inflation peaked and use the baseline and optimal Taylor rule coefficients—summarized in Table 3—together with inflation, unemployment rate, and EE rate gaps in that quarter to calculate the implied nominal interest rates i^*.[61] The bottom row of Table 3 shows that the prescribed quarterly nominal interest rate is 2.22 percent under the optimal policy and 2.13 percent under the baseline policy, translating to an around 0.4 pp higher annual nominal interest rate under the optimal policy.

Optimal policy in a model with low MPC. To gauge whether accounting for market incompleteness matters for monetary policy, we consider an alternative version of our model without discount factor heterogeneity that is recalibrated to match the

[59] Appendix B.7 provides details on evaluating the objective under alternative rules.

[60] This restriction allows us to search over a broader range of coefficients. We also considered a reasonable range for *all* three coefficients, including the one on inflation. We found similar qualitative results in that the optimal policy responds to both u and EE but with opposite signs.

[61] Inflation, unemployment, and EE rate gaps were 1.19, -1.26, and 0.22 pp., respectively.

same set of moments in Table 1, except for the moments disciplining $\bar{\beta}$ and $\Delta\beta$.[62] Importantly, the income-weighted average quarterly MPC in this model is only 0.02 instead of the baseline 0.20, approximating a model with perfect self-insurance.

The last column of Table 3 shows that the optimal policy now prescribes $\Phi_u^* = -0.125$ and $\Phi_{EE}^* = 0.625$, i.e., the monetary authority responds to the EE gap much less aggressively relative to the optimal policy in the baseline calibration. We find that the implied annual optimal nominal rate in 2022:Q2 is around 0.5 pp lower in the low-MPC model compared to the optimal nominal rate in the baseline model. This result is consistent with our earlier findings in Section 2 and in Section 5.5. When the EE rate increases, relative to a model with perfect self-insurance, a model with imperfect self-insurance generates smaller declines in demand, output, and tightness due to the lower elasticity of demand to the real rate among high-MPC agents. As such, monetary policy is able to respond to inflation due to an elevated EE rate more aggressively in a model with imperfect self-insurance.

Taking stock. We show that there are large stabilization gains from accounting for EE dynamics in monetary policy when the economy is subject to job mobility shocks. Crucially, we find that the optimal coefficient on the EE gap is around twice as large with imperfect self-insurance relative to an economy with perfect self-insurance, implying that accounting for market incompleteness also matters for monetary policy.

7 Conclusions

We show that fluctuations in the EE rate have key macroeconomic and policy-relevant implications. Using a HANK model, we find that muted job mobility during the recovery from the Great Recession caused 0.42 pp lower inflation, while elevated job mobility during the recovery from COVID-19 generated 0.77 pp of additional inflation. On the normative side, we show that optimal monetary policy in response to EE fluctuations prescribes a less aggressive nominal rate response to unemployment than what is commonly prescribed by the literature, and a strong positive response to the EE gap. This implies that central banks should behave differently between cases where EE dynamics differ, even when unemployment dynamics are identical. Crucially, we show that accounting for market incompleteness matters for the positive and normative implications of changes in EE transitions.

[62] Table B.3 in Appendix B.8 summaries the calibrated parameter values of this model.

Our model features a rich set of fiscal instruments, such as a consumption tax, progressive labor income tax, unemployment and retirement benefits, and government debt. Thus, it provides a framework to study fiscal and monetary policy interactions, accounting for rich labor market dynamics. In addition, it is straightforward to introduce other shocks (e.g., shocks to monetary policy, markups, and other labor-market-related parameters) into our model. Our solution method enables us to estimate such shocks jointly to evaluate the model's performance in matching richer time-series and cross-sectional moments. We leave these considerations for future research.

References

Acharya, S. and K. Dogra (2020): "Understanding HANK: Insights from a PRANK," *Econometrica*, 88, 1113–1158.

Alves, F. (2020): "Job ladder and business cycles," Tech. rep., Technical report, New York University.

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2021): "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models," *Econometrica*, 89, 2375–2408.

Auclert, A., M. Rognlie, and L. Straub (2020): "Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model," Tech. rep., NBER.

——— (2023): "The intertemporal keynesian cross," Tech. rep., Working paper.

——— (2024): "Fiscal and monetary policy with heterogeneous agents," Tech. rep., National Bureau of Economic Research.

Benabou, R. (2002): "Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency?" *Econometrica*, 70, 481–517.

Bilal, A., N. Engbom, S. Mongey, and G. L. Violante (2022): "Firm and worker dynamics in a frictional labor market," *Econometrica*, 90, 1425–1462.

Birinci, S. (2021): "Spousal labor supply response to job displacement and implications for optimal transfers," *FRB St. Louis Working Paper*.

Blanco, J., A. Drenik, C. Moser, and E. Zaratiegui (2023): "A theory of non-coasean labor markets," Tech. rep., IZA Discussion Paper No. 16121.

Bostanci, G., O. Koru, and S. Villalvazo (2022): "Changing jobs to fight inflation: Labor market reactions to inflationary shocks," Tech. rep., Penn State University.

Darougheh, S., R. Faccini, L. Melosi, and A. Villa (2025): "On-the-job search and inflation under the microscope," Tech. rep., CEPR Discussion Paper No. 19795.

Davis, S. J. and T. von Wachter (2011): "Recessions and the costs of job loss," *Brookings Papers on Economic Activity*, 1–73.

Debortoli, D., J. Kim, J. Lindé, and R. Nunes (2018): "Designing a simple loss function for central banks: Does a dual mandate make sense?" *The Economic Journal*, 129, 2010–2038.

Faberman, R. J., A. I. Mueller, A. Şahin, and G. Topa (2022): "Job search behavior among the employed and non-employed," *Econometrica*, 90, 1743–1779.

Faccini, R. and L. Melosi (2023): "Job-to-job mobility and inflation," *Review of Economics and Statistics*.

Fujita, S., G. Moscarini, and F. Postel-Vinay (2023): "Measuring employer-to-employer reallocation," Tech. rep., National Bureau of Economic Research.

Fukui, M. (2020): "A theory of wage rigidity and unemployment fluctuations with on-the-job search," Tech. rep., Massachusetts Institute of Technology.

Galí, J. (2015): *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*, Princeton University Press.

Galí, J. and M. Gertler (1999): "Inflation dynamics: A structural econometric analysis," *Journal of Monetary Economics*, 44, 195–222.

Gornemann, N., K. Kuester, and M. Nakajima (2021): "Doves for the rich, hawks for the poor? Distributional consequences of systematic monetary policy," Tech. rep., University of Bonn.

Heathcote, J., K. Storesletten, and G. L. Violante (2014): "Consumption and labor supply with partial insurance: An analytical framework," *American Economic Review*, 104, 2075–2126.

Huckfeldt, C. (2022): "Understanding the scarring effect of recessions," *American Economic Review*, 112, 1273–1310.

Jacobson, L. S., R. J. LaLonde, and D. G. Sullivan (1993): "Earnings losses of displaced workers," *American Economic Review*, 685–709.

Jarosch, G. (2023): "Searching for job security and the consequences of job loss," *Econometrica*, 91, 903–942.

Jensen, H. (2002): "Targeting nominal income growth or inflation?" *American Economic Review*, 92, 928–956.

Kaplan, G., B. Moll, and G. L. Violante (2018): "Monetary policy according to HANK," *American Economic Review*, 108, 697–743.

Kaplan, G. and G. L. Violante (2022): "The marginal propensity to consume in heterogeneous agent models," *Annual Review of Economics*, 14, 747–775.

Karahan, F., S. Ozkan, and J. Song (2022): "Anatomy of lifetime earnings inequality: Heterogeneity in job ladder risk vs. human capital," *FRB St. Louis Working Paper*.

Krolikowski, P. (2017): "Job ladders and earnings of displaced workers," *American*

Economic Journal: Macroeconomics, 9, 1–31.

Krusell, P., T. Mukoyama, and A. Şahin (2010): "Labour-market matching with precautionary savings and aggregate fluctuations," *The Review of Economic Studies*, 77, 1477–1507.

Ljungqvist, L. and T. J. Sargent (1998): "The European unemployment dilemma," *Journal of Political Economy*, 106, 514–550.

McKay, A. and C. Wolf (2022): "What can time-series regressions tell us about policy counterfactuals?" Tech. rep., Federal Reserve Bank of Minneapolis.

Mercan, Y., B. Schoefer, and P. Sedláček (2024): "A congestion theory of unemployment fluctuations," *American Economic Journal: Macroeconomics, forthcoming.*

Moscarini, G. and F. Postel-Vinay (2016): "Wage posting and business cycles," *American Economic Review: Papers & Proceedings*, 106, 208–213.

——— (2022): "The job ladder: Inflation vs. reallocation," *Yale University.*

Okun, A. (1962): "Potential 1962 GNP: Its measurement and its significance," in *Proceedings of the Business and Economic Statistics: Section American Statistical Association*, 98–103.

Pilossoph, L. and J. M. Ryngaert (2022): "Job search, wages, and inflation," Tech. rep., Duke University.

Postel-Vinay, F. and J.-M. Robin (2002): "Equilibrium wage dispersion with worker and employer heterogeneity," *Econometrica*, 70, 2295–2350.

Ravn, M. O. and V. Sterk (2021): "Macroeconomic fluctuations with HANK & SAM: An analytical approach," *Journal of the European Economic Association*, 19, 1162–1202.

Reiter, M. (2009): "Solving heterogeneous-agent models by projection and perturbation," *Journal of Economic Dynamics and Control*, 33, 649–665.

Rotemberg, J. (1982): "Sticky prices in the United States," *Journal of Political Economy*, 1187–1211.

Schaal, E. (2017): "Uncertainty and unemployment," *Econometrica*, 85, 1675–1721.

Stevens, A. H. (1997): "Persistent effects of job displacement: The importance of multiple job losses," *Journal of Labor Economics*, 15, 165–188.

Taylor, J. B. (1993): "Discretion versus policy rules in practice," in *Carnegie-Rochester conference series on public policy*, Elsevier, vol. 39, 195–214.

Walsh, C. E. (2003): "Speed limit policies: the output gap and optimal monetary policy," *American Economic Review*, 93, 265–278.

Yellen, J. L. (2012): "The Economic Outlook and Monetary Policy: A speech at the Money Marketeers of New York University, New York, New York, April 11, 2012," Speech 638, Board of Governors of the Federal Reserve System (U.S.).

Supplemental Appendix

A TANK model

In this section, we provide derivations for the TANK model, discuss its solution, and present additional results.

A.1 Consumption rule of the PIH household

The intertemporal budget constraint (IBC) for the PIH household is:

$$
\begin{aligned}
a_{t-1} &= \frac{1}{(1+r_t)}\left(c_t^{PIH} - Z_t\right) + \frac{1}{(1+r_t)} a_t \\
&= \frac{1}{(1+r_t)}\left(c_t^{PIH} - Z_t\right) + \frac{1}{(1+r_t)}\frac{1}{(1+r_{t+1})}\left(c_{t+1}^{PIH} - Z_{t+1}\right) + \frac{1}{(1+r_t)}\frac{1}{(1+r_{t+1})} a_{t+1} \\
&= \frac{1}{(1+r_t)}\left(c_t^{PIH} - Z_t\right) + \frac{1}{(1+r_t)}\frac{1}{(1+r_{t+1})}\left(c_{t+1}^{PIH} - Z_{t+1}\right) \\
&+ \frac{1}{(1+r_t)}\frac{1}{(1+r_{t+1})}\frac{1}{(1+r_{t+2})}\left(c_{t+2}^{PIH} - Z_{t+2}\right) + \frac{1}{(1+r_t)}\frac{1}{(1+r_{t+1})}\frac{1}{(1+r_{t+2})} a_{t+2} \\
&\vdots \\
a_{t-1} &= \frac{1}{(1+r_t)} \sum_{s=0}^{\infty} q_{t+s}(c_{t+s}^{PIH} - Z_{t+s}),
\end{aligned}
$$

where we define $q_{t+s} = \frac{1}{1+r_{t+1}} \cdots \times \frac{1}{1+r_{t+s}}$ with $q_t = 1$ and impose a no-Ponzi condition $\lim_{s\to\infty} q_{t+s}a_{t+s} = 0$. We can equivalently express the IBC as follows:

$$
\sum_{s=0}^{\infty} q_{t+s} c_{t+s}^{PIH} = (1+r_t)a_{t-1} + \sum_{s=0}^{\infty} q_{t+s} Z_{t+s}.
$$

We express the Euler equation (assuming CRRA utility) to tie current period to s-period ahead consumption as:

$$
\begin{aligned}
c_t^{PIH} &= \beta^{-1/\sigma}(1+r_{t+1})^{-1/\sigma} c_{t+1}^{PIH} \\
&= \beta^{-1/\sigma}(1+r_{t+1})^{-1/\sigma}\beta^{-1/\sigma}(1+r_{t+2})^{-1/\sigma} c_{t+2}^{PIH} \\
&= \beta^{-2/\sigma}(1+r_{t+1})^{-1/\sigma}(1+r_{t+2})^{-1/\sigma} c_{t+2}^{PIH} \\
&= \beta^{-2/\sigma}\left[(1+r_{t+1})(1+r_{t+2})\right]^{-1/\sigma} c_{t+2}^{PIH} \\
&\vdots \\
c_t^{PIH} &= \beta^{-s/\sigma} q_{t+s}^{1/\sigma} c_{t+s}^{PIH} \\
\Rightarrow c_{t+s}^{PIH} &= \beta^{s/\sigma} q_{t+s}^{-1/\sigma} c_t^{PIH}.
\end{aligned}
$$

For $t = 0$, we have the following:

$$c_s^{PIH} = \beta^{s/\sigma} q_s^{-1/\sigma} c_0^{PIH}.$$

Plugging this into the IBC, we obtain Equation (1):

$$\sum_{s=0}^{\infty} q_{t+s} c_{t+s}^{PIH} = (1+r_t)a_{t-1} + \sum_{s=0}^{\infty} q_{t+s} Z_{t+s}$$

$$\sum_{s=0}^{\infty} q_{t+s} \beta^{s/\sigma} q_{t+s}^{-1/\sigma} c_t^{PIH} = (1+r_t)a_{t-1} + \sum_{s=0}^{\infty} q_{t+s} Z_{t+s}$$

$$c_t^{PIH} = \frac{1}{\sum_{s=0}^{\infty} \beta^{s/\sigma} q_{t+s}^{1-1/\sigma}} \Big((1+r_t)a_{t-1} + \sum_{s=0}^{\infty} q_{t+s} Z_{t+s} \Big).$$

A.2 Intermediate firm's problem

The problem of an intermediate firm j with last period *relative* price $p_{t-1}(j)$ is

$$\begin{aligned} \Theta\left(p_{t-1}\left(j\right)\right) = \max_{p_t(j)} & \, p_t\left(j\right) y_t\left(j\right) - p_t^l y_t\left(j\right) \\ & - \frac{\eta}{2\vartheta} \log \left(\frac{p_t\left(j\right)}{p_{t-1}\left(j\right)} \left(1+\pi_t\right) - \pi^* \right)^2 Y_t + \frac{1}{1+r_{t+1}} \Theta\left(p_t\left(j\right)\right). \end{aligned}$$

Substituting in the CES demand for each variety, $y_t(j) = p_t(j)^{-\eta} Y_t$, we have

$$\begin{aligned} \Theta\left(p_{t-1}\left(j\right)\right) = \max_{p_t(j)} & \, p_t\left(j\right)^{1-\eta} Y_t - p_t^l p_t\left(j\right)^{-\eta} Y_t \\ & - \frac{\eta}{2\vartheta} \log \left(\frac{p_t\left(j\right)}{p_{t-1}\left(j\right)} \left(1+\pi_t\right) - \pi^* \right)^2 Y_t + \frac{1}{1+r_{t+1}} \Theta\left(p_t\left(j\right)\right). \end{aligned}$$

The first-order condition with respect to relative price $p_t(j)$ is given by

$$\begin{aligned} 0 = & \left(1-\eta\right) p_t\left(j\right)^{-\eta} Y_t + \eta p_t^l p_t\left(j\right)^{-\eta-1} Y_t \\ & - \frac{\eta}{\vartheta} \log \left(\frac{p_t\left(j\right)}{p_{t-1}\left(j\right)} \left(1+\pi_t\right) - \pi^* \right) \frac{1}{\frac{p_t(j)}{p_{t-1}(j)} \left(1+\pi_t\right) - \pi^*} \frac{1+\pi_t}{p_{t-1}\left(j\right)} Y_t + \frac{1}{1+r_{t+1}} \Theta'\left(p_t\left(j\right)\right), \end{aligned}$$

and the envelope condition is

$$\Theta'\left(p_{t-1}\left(j\right)\right) = \frac{\eta}{\vartheta} \log \left(\frac{p_t\left(j\right)}{p_{t-1}\left(j\right)} \left(1+\pi_t\right) - \pi^* \right) \frac{1}{\frac{p_t(j)}{p_{t-1}(j)} \left(1+\pi_t\right) - \pi^*} \frac{p_t\left(j\right)\left(1+\pi_t\right)}{p_{t-1}\left(j\right)^2} Y_t.$$

Iterating the envelope condition forward by one period yields

$$\Theta'\left(p_t\left(j\right)\right) = \frac{\eta}{\vartheta} \log \left(\frac{p_{t+1}\left(j\right)}{p_t\left(j\right)} \left(1+\pi_{t+1}\right) - \pi^* \right) \frac{1}{\frac{p_{t+1}(j)}{p_t(j)} \left(1+\pi_{t+1}\right) - \pi^*} \frac{p_{t+1}\left(j\right)\left(1+\pi_{t+1}\right)}{p_t\left(j\right)^2} Y_{t+1}.$$

Consolidating the envelope and the first-order conditions, we obtain:

$$
\begin{aligned}
0 = {} & (1-\eta)\, p_t(j)^{-\eta}\, Y_t + \eta p_t^l p_t(j)^{-\eta-1}\, Y_t \\
& - \frac{\eta}{\vartheta} \log\left(\frac{p_t(j)}{p_{t-1}(j)}(1+\pi_t) - \pi^*\right) \frac{1}{\frac{p_t(j)}{p_{t-1}(j)}(1+\pi_t) - \pi^*} \frac{1+\pi_t}{p_{t-1}(j)} Y_t \\
& + \frac{1}{1+r_{t+1}} \underbrace{\frac{\eta}{\vartheta} \log\left(\frac{p_{t+1}(j)}{p_t(j)}(1+\pi_{t+1}) - \pi^*\right) \frac{1}{\frac{p_{t+1}(j)}{p_t(j)}(1+\pi_{t+1}) - \pi^*} \frac{p_{t+1}(j)(1+\pi_{t+1})}{p_t(j)^2} Y_{t+1}}_{\Theta'(p_t(j))}.
\end{aligned}
$$

All firms set the same price due to symmetry, $p_t(j) = 1 \; \forall\, t,\, j$. Thus, we have

$$
\begin{aligned}
0 = {} & (1-\eta)\, Y_t + \eta p_t^l Y_t - \frac{\eta}{\vartheta} \frac{\log(1+\pi_t-\pi^*)(1+\pi_t)}{1+\pi_t-\pi^*} Y_t \\
& + \frac{1}{1+r_{t+1}} \frac{\eta}{\vartheta} \frac{\log(1+\pi_{t+1}-\pi^*)(1+\pi_{t+1})}{1+\pi_{t+1}-\pi^*} Y_{t+1}.
\end{aligned}
$$

Rearranging terms and using the definition of π_t, we have the NKPC in Equation (8):

$$
\frac{\log(1+\pi_t-\pi^*)(1+\pi_t)}{1+\pi_t-\pi^*} = \vartheta\left(p_t^l - \frac{\eta-1}{\eta}\right) + \frac{1}{1+r_{t+1}} \frac{\log(1+\pi_{t+1}-\pi^*)(1+\pi_{t+1})}{1+\pi_{t+1}-\pi^*} \frac{Y_{t+1}}{Y_t}.
$$

A.3 Steady state

A.3.1 Equilibrium definition

Given the monetary policy rule in Equation (5) and a path of OJS efficiency ν_t, an equilibrium is a sequence of consumption levels c_t^{PIH} and c_t^{HtM}, real price of labor services p_t^l, market tightness θ_t, interest rates r_t and i_t, inflation π_t, and worker masses u_t, $e_t(\alpha)$, and $e_t(1)$ such that

- Household consumption decisions are given by Equation (1) and $c_t^{HtM} = W_t$, where the aggregate real labor income W_t and the per capita income of the PIH household Z_t are given by Equations (2) and (3), and firm profits satisfy Equations (9) and (11).
- Aggregate demand is $C_t = (1-\mu)\, c_t^{PIH} + \mu c_t^{HtM}$.
- Tightness θ_t satisfies the law of motion for unemployment in Equation (4).
- Employed masses $e_t(\alpha)$ and $e_t(1)$ evolve according to Equation (4).
- The free-entry condition $V_t = 0$ and Equation (12) yield the value of a matched firm J_t. Equation (10) then gives the real price of labor services p_t^l.
- The real interest rate r_t satisfies the Fisher equation (6).
- Inflation π_t satisfies the NKPC in Equation (8).

A.3.2 Steady state in real terms

Nominal frictions are not relevant in the steady state, where prices rise by the rate of long-run inflation π^*; hence, intermediate firms do not incur price adjustment costs. To facilitate the solution, we express the model in relative prices and real allocations.

- Evaluating Equation (8) at the steady state yields the real marginal cost:
$$mc = p^l = \frac{\eta - 1}{\eta}. \tag{A.1}$$
- Using Equation (9), per-period real profits of the intermediate firms are
$$\Gamma^I = (1 - p^l)Y. \tag{A.2}$$
- Evaluating Equation (4) at the steady state and solving for the worker's contact rate per unit of search efficiency $f(\theta)$, we obtain:
$$f(\theta) = \delta(1-u) / [u + \delta(1-u)]. \tag{A.3}$$
- Let $\widehat{u}$ and $\widehat{e}$ denote the masses of unemployed and employed job searchers at the search stage, respectively. They are given by:
$$\begin{aligned} \widehat{u} &= u + \delta(1-u), \\ \widehat{e} &= (1-\delta)(1-u). \end{aligned} \tag{A.4}$$

  We then back out the real value of a matched service firm J from Equation (12) imposing the free-entry condition $V = 0$ as follows:
$$J = \kappa / [q(\theta)(\widehat{u}/(\widehat{u} + \nu\widehat{e}))]. \tag{A.5}$$
- Finally, solving for the relative price of labor services p^l in Equation (10) yields:
$$p^l = J[1 - (1/(1+r))(1-\delta)(1 - \nu f(\theta))] / (1-\alpha). \tag{A.6}$$

A.3.3 Solution algorithm

Next, we provide the algorithm to solve for the steady state of the TANK model.

1. Guess output Y and real interest rate r.
2. Solve for labor market outcomes:
   - Given Y, calculate the mass of unemployed at the steady state using the equilibrium condition $u = 1 - Y$.
   - Given u, use Equation (A.3) to obtain the worker's contact rate per unit of search efficiency $f(\theta)$. Given $f(\theta)$, calculate market tightness θ by inverting the CES matching function $\theta = (f^\xi/(1-f^\xi))^{1/\xi}$. Given $f(\theta)$ and θ, calculate the firm's contact rate $q(\theta) = f(\theta)/\theta$.

- Given $f(\theta)$, $q(\theta)$, and the definitions of $\widehat{u}$ and $\widehat{e}$ in Equation (A.4), calculate the value of a matched firm J using Equation (A.5). Then, use Equation (A.6) to obtain the relative price of labor services p^l.
- Compute the stationary worker masses $e(\alpha)$ and $e(1)$ using Equation (4).

3. Solve for aggregate demand:
   - Given the worker distribution, obtain aggregate real labor income W using Equation (2), aggregate real service firm profits Γ^S using Equation (11), and aggregate real intermediate firm profits using Equation (A.2).
   - Given aggregate profits $\Gamma = \Gamma^S + \Gamma^I$ and labor income, obtain the total income of the PIH household $Z = W + \Gamma/(1-\mu)$.
   - Given W and Z, solve for steady-state demand of the PIH household c^{PIH} using Equation (1) and of the HtM household using $c^{HtM} = W$. Then, obtain aggregate demand using $C = (1-\mu)c^{PIH} + \mu c^{HtM}$.
4. Check whether goods market clearing $Y = C$ and the Phillips curve $p^l = (\eta - 1)/\eta$ hold. If not, return to Step 1 and update the guesses for Y and r until convergence. In practice, we use a non-linear solver over Y and r.

A.3.4 Transitional dynamics

Online Appendix Section 1 represents the model as a DAG and discusses how we solve for the dynamics of the TANK model.[63]

A.4 Additional results

We present additional results from the TANK model that supplement Section 2.4.

Impulse responses. Figure A.1 plots the dynamics of total real profits ($\Gamma = \Gamma^S + \Gamma^I$), real service firm profits Γ^S, real intermediate firm profits Γ^I, and the unemployment rate following a 1 percent persistent increase in the OJS efficiency, as described in Section 2.4 for both the RANK and TANK models under a constant real rate r. The rise in OJS efficiency leads to a decline in Γ in both models (Panel (a)). Panels (b) and (c) show that the decline in total profits is driven by the decline in the profits of intermediate firms, while the profits of service firms increase in both models. Because the OJS shock does not have any real effects in the RANK model under a constant real rate—it merely redistributes income from workers to firms,

[63] A separate Online Appendix for this paper is available at: `https://tinyurl.com/4hrkfyxs`.

Figure A.1: Impulse responses to a positive OJS efficiency shock: Constant real rate

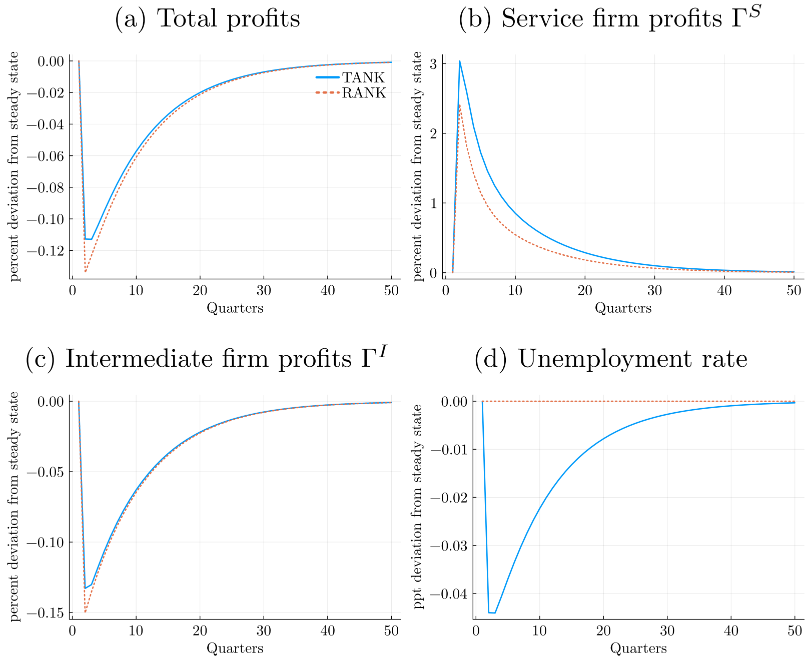

Notes: This figure presents the impulse responses of total real profits ($\Gamma = \Gamma^S + \Gamma^I$), real service firm profits Γ^S, real intermediate firm profits Γ^I, and unemployment rate in an economy subject to a 1 percent persistent increase in the OJS efficiency. The blue lines represent the TANK model and the red lines represent the RANK model. Here, we keep the real interest rate r constant.

which are then remitted back to workers as dividends—the unemployment rate does not change. However, the increase in demand in the TANK model leads to a rise in vacancy creation and a decline in the unemployment rate, as shown in Panel (d).

Figure A.2 presents the same exercise when the monetary authority reacts to shocks according to the Taylor rule. Different from Figure A.1, the unemployment rate now increases. This is driven by a decline in demand and vacancy creation. The smaller increase in the unemployment rate in the TANK model is due to a smaller decline in demand and market tightness as the HtM households in the TANK model, who are less sensitive to the real rate, mitigate fluctuations in demand and tightness.

Figure A.3 presents the *level* deviations of aggregate consumption and output from their respective steady states. Aggregate consumption and output responses are identical, confirming the equilibrium condition.

Understanding the mechanisms: Composition vs shorter duration. Section 2.4 argues that the main reason behind the rise in p^l upon a positive OJS efficiency shock is the decline in the expected match value $\mathbb{E}J$, driven by *both* the change in the composition of job seekers towards the employed and a shorter match duration.

Figure A.2: Impulse responses to a positive OJS efficiency shock: Real rate response

(a) Total profits (b) Firm profits

(c) Unemployment rate (d) Real price of labor services

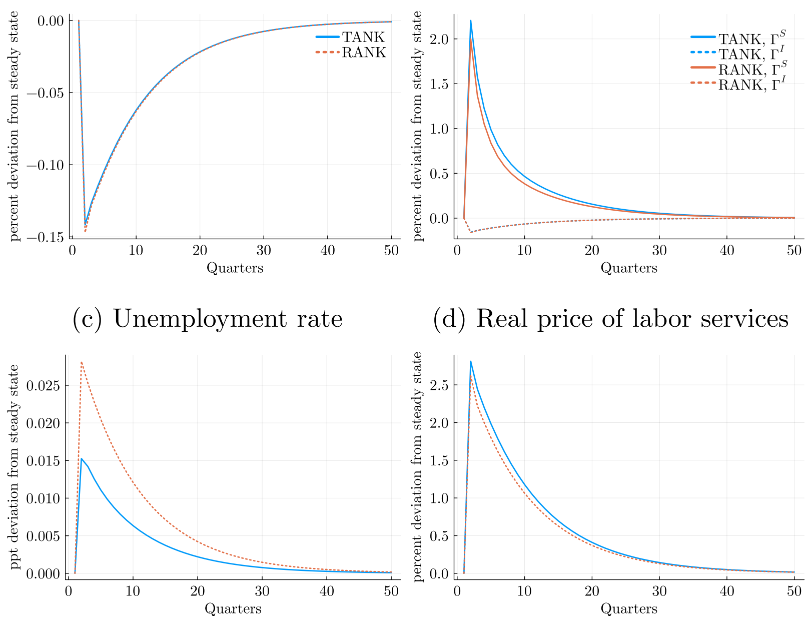

Notes: This figure presents the impulse responses of total profits ($\Gamma = \Gamma^S + \Gamma^I$), service firm profits Γ^S, intermediate firm profits Γ^I, unemployment rate, and real price of labor services p^l in an economy subject a 1 percent persistent increase in the OJS efficiency. Blue lines represent the TANK model and red lines represent the RANK model.

Figure A.3: IRFs to a positive OJS efficiency shock: Demand and output levels

(a) Consumption (b) Output

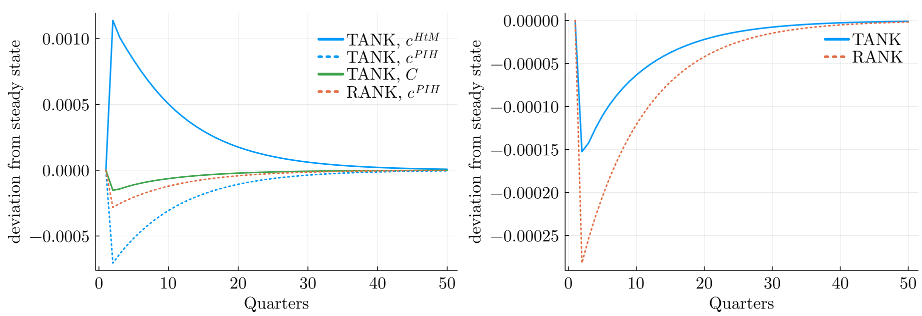

Notes: This figure presents the impulse responses (as level deviations from steady state) of aggregate consumption and output in an economy subject a 1 percent persistent increase in the OJS efficiency. The blue and green lines represent the TANK model, and the red lines represent the RANK model.

Here, we implement an exercise to isolate the role of the latter channel.

The composition channel is active to the extent that the unemployed and employed job seekers are paid different wages. To this end, we set the piece rate of workers hired out of unemployment to $\alpha = 0.95$ instead of 0.5, while recalibrating κ to match the same steady-state unemployment rate as in the baseline model, keeping all other

Figure A.4: IRFs to a positive OJS efficiency shock: Without wage increases

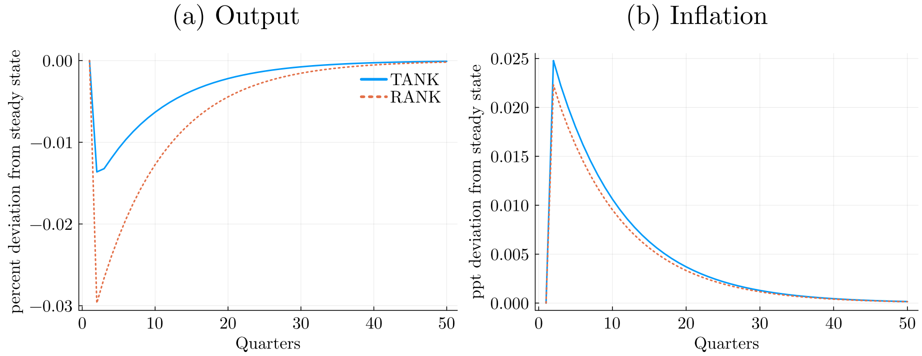

Notes: This figure presents the impulse responses of output and inflation in an economy subject to a 1 percent persistent increase in the OJS efficiency. Blue lines represent the TANK model and red lines represent the RANK model. Here, to isolate the role of shorter match duration channel on the rise in inflation, we set the piece rate of workers who are hired out of unemployment $\alpha = 0.95$ instead of 0.5, while recalibrating κ to match the same steady-state unemployment rate as in the baseline model.

parameters the same.[64] This way, outside offers result in small wage gains, and the effect of worker composition on $\mathbb{E}J$ becomes negligible. Therefore, this exercise isolates the role of the (shorter) match duration channel in inflation.

Figure A.4 shows that a positive OJS efficiency shock still leads to a decline in output and to a rise in inflation in both versions of the model.[65] As we turn off the composition channel, intermediate firms still need to pay a higher price for labor services, given that labor services firms expect shorter-lived matches. Overall, a similar inflationary response between Figures 2 and A.4 imply that the shorter match duration effect is the main reason behind the rise in inflation.

B HANK model

Here, we discuss the solution of the steady state and transition dynamics in the HANK model, present additional results, and provide details on the normative exercise.

B.1 Steady state

B.1.1 Equilibrium definition

Given fiscal and monetary policies, as well as paths of aggregate shocks, an equilibrium is a sequence of decision rules for consumption $g_t^{Ec}, g_t^{Uc}, g_t^{Rc}$ and mutual fund share demand $g_t^{Es}, g_t^{Us}, g_t^{Rs}$, intermediate and service firms' profits Γ_t^I and Γ_t^S, dividends D_t,

[64] Given a common job separation rate δ and unemployment rate, this recalibration implies that the job finding rate and hence tightness and vacancies are the same across the two exercises.

[65] The change in the value of α does not change the results in the RANK model as the total income of the PIH household does not depend on how firm revenue is shared as profits vs wages.

real price of labor services p_t^l, share price P_t^s, tightness θ_t, interest rates r_t, i_t, inflation π_t, bond holdings B_t, and worker distributions $\left\{\lambda_t^E, \lambda_t^U, \lambda_t^R\right\}$ such that

- The nominal and real interest rates satisfy Equations (5) and (6).
- Intermediate and service firms' profits satisfy Equations (23) and (24).
- Share prices satisfy Equation (21) and dividends are given by Equation (22).
- Bonds are such that the government budget constraint in Equation (25) holds.
- Individual decisions g_t^{Ec}, g_t^{Uc}, g_t^{Rc}, g_t^{Es}, g_t^{Us} and g_t^{Rs} are optimal.
- Tightness θ_t and real price of labor services p_t^l are such that the value of posting a vacancy expressed in Equation (20) is zero and Equation (26) holds.
- Inflation π_t satisfies the NKPC in Equation (8).
- The share market clears as specified in Equation (27).[66]
- Worker distributions follow the laws of motion in Online Appendix Section 2.

B.1.2 Casting the model in relative prices and real variables

As in the TANK model, we start by deriving the equations governing relative prices, real dividends, and real profits of intermediate firms in steady state.

- Evaluating the NKPC in steady state, we obtain the real marginal cost:
$$mc = \frac{\eta - 1}{\eta}.$$
- The price of labor services is then given by
$$p^l = mc \times z = \frac{\eta - 1}{\eta} z. \tag{B.1}$$
- Per-period real profits of the intermediate firms are given by
$$\Gamma^I = (1 - mc)Y = \frac{Y}{\eta}. \tag{B.2}$$
- Real dividends are given by
$$\begin{aligned} d &= x_B Y - \frac{x_B Y(1 + \pi^*)}{(1 + i)} + \Gamma^I + \Gamma^S \\ &= x_B Y \frac{r}{1 + r} + \Gamma^I + \Gamma^S. \end{aligned} \tag{B.3}$$
- Dividing the no-arbitrage condition by the aggregate price P, we have
$$\frac{(p^s + d)(1 + \pi^*)}{p^s} = 1 + i$$
$$p^s = \frac{d}{r}. \tag{B.4}$$

[66] We do not check for goods market clearing due to Walras's Law.

- We rewrite the government budget constraint in real terms. Let $b_t = B_t/P_{t+1}$. Then, dividing both sides by P_t, multiplying the first term on the right-hand side by $\frac{P_{t+1}}{P_{t+1}}$, and recognizing that $1 + i_t = (1 + r_{t+1})(1 + \pi_{t+1})$, we get

$$
\begin{aligned}
b_{t-1} + g_t + \int UI\left(h\right) d\lambda_t^U\left(\cdot\right) + \int \phi^R d\lambda_t^R(\cdot) = \frac{b_t}{1 + r_{t+1}} &+ \tau_c \int c\left(e, s, h, x, \alpha, \beta\right) d\lambda_t\left(\cdot\right) \\
&+ \int \left(UI\left(h\right) - \tau_t\left(UI\left(h\right)\right)^{1-\Upsilon}\right) d\lambda_t^U\left(\cdot\right) \\
&+ \int \left(w(h, x, \alpha) - \tau_t w(h, x, \alpha)^{1-\Upsilon}\right) d\lambda_t^E\left(\cdot\right) \\
&+ \int \left(\phi^R - \tau_t\left(\phi^R\right)^{1-\Upsilon}\right) d\lambda_t^R(\cdot).
\end{aligned}
$$

It is useful to define the real net revenue of government (tax proceeds minus outlays for pensions and unemployment insurance), $\mathcal{R}_t$, as

$$
\begin{aligned}
\mathcal{R}_t = &- \int UI\left(h\right) d\lambda_t^U\left(\cdot\right) - \int \phi^R d\lambda_t^R\left(\cdot\right) \\
&+ \tau_c \int c\left(e, s, h, x, \alpha, \beta\right) d\lambda_t\left(\cdot\right) + \int \left(UI\left(h\right) - \tau\left(UI\left(h\right)\right)^{1-\Upsilon}\right) d\lambda_t^U\left(\cdot\right) \\
&+ \int \left(w(h, x, \alpha) - \tau w(h, x, \alpha)^{1-\Upsilon}\right) d\lambda_t^E\left(\cdot\right) + \int \left(\phi^R - \tau\left(\phi^R\right)^{1-\Upsilon}\right) d\lambda_t^R\left(\cdot\right).
\end{aligned} \tag{B.5}
$$

With these definitions, the government budget constraint in real terms is

$$
\begin{aligned}
b_{t-1} + g_t &= \frac{b_t}{1 + r_{t+1}} + \mathcal{R}_t \\
\Rightarrow 0 &= (1 + r_{t+1})(b_{t-1} + g_t - \mathcal{R}_t) - b_t.
\end{aligned} \tag{B.6}
$$

B.1.3 Solution algorithm

We solve for the steady state using the following algorithm by bisecting over a nominal interest rate i that clears the share market given by Equation (27).

1. For a given nominal rate i and π^*, obtain r from the Fisher equation (6).
2. Outer loop: Guess a tax parameter τ, output Y, and service firm profits Γ^S.
   - Calculate the relative price of labor services using Equation (B.1), real bond holdings $b = x_B Y$, real government expenditures $g = x_G Y$, and real intermediate firm profits Γ^I using Equation (B.2).
   - Calculate real dividends d using Equation (B.3).
   - Calculate real share price p^s using Equation (B.4).
3. Inner loop: Guess a market tightness θ.

Table B.1: Externally calibrated parameters of the baseline HANK model

Parameter	Explanation	Value	Reason
σ	Curvature in utility function	2	Standard
ψ^R	Retirement probability	0.00625	40 years of work stage
ψ^D	Death probability	0.0125	20 years of retirement stage
Δh	Skill appreciation/depreciation amount	0.275	Set
π^E	Skill appreciation probability	0.018	Wage growth for job stayers
ξ	Matching function elasticity	1.6	Set
η	Elasticity of substitution	6	20 percent markup
ϑ	Price adjustment cost parameter	0.021	Slope of Phillips curve, Galí and Gertler (1999)
x_G	Government spending/GDP ratio	0.19	Total net federal outlay/GDP
x_B	Debt/GDP ratio	2.43	Total public debt/GDP
τ_c	Consumption tax rate	0.0302	Sales tax receipt/consumption expenditure
Υ	Progressivity of income tax	0.151	Heathcote, Storesletten, and Violante (2014)
ρ_τ	Responsiveness of income tax to debt	0.10	Auclert, Rognlie, and Straub (2020)
π^*	Steady-state inflation rate	0.00496	2 percent annual inflation rate
Φ_π	Responsiveness of i to inflation gap	1.5	Taylor (1993) and Galí and Gertler (1999)
Φ_u	Responsiveness of i to unemployment gap	-0.25	Taylor (1993) and Galí and Gertler (1999)

Notes: This table summarizes the externally calibrated parameters of the model. See the main text for details.

- Calculate worker contact rate $f(\theta)$.
- Solve the workers' problems given by Equations (14), (15), (16).
- Compute the stationary worker distributions over state variables μ^E, μ^U, μ^R, λ, λ^E, λ^U, and λ^R.
- Solve the matched firm problem given by Equation (19).
- Given the solution to the firm problem and worker distributions, calculate the implied market tightness $\widetilde{\theta}$ consistent with the free-entry condition $V = 0$, where V satisfies Equation (20).
- Iterate over the inner loop until $\widetilde{\theta}$ agrees with the guessed tightness θ.

4. Using the worker distributions, calculate the implied output $\widetilde{Y}$ using Equation (26) and real service firm profits $\widetilde{\Gamma}^S$ using Equation (24).
5. Calculate the implied tax parameter $\widetilde{\tau}$ that clears the government budget constraint, which can be obtained from Equations (B.5) and (B.6) as:

$$\widetilde{\tau} = \frac{-\frac{r}{1+r}x_B Y - x_g Y + \tau_c \int c d\lambda + \int w d\lambda^E + \int UI d\lambda^U + \int \phi^R d\lambda^R}{\int w^{1-\Upsilon} d\lambda^E + \int UI^{1-\Upsilon} d\lambda^U + \int \left(\phi^R\right)^{1-\Upsilon} d\lambda^R}.$$

6. Iterate until $\widetilde{\tau}$, $\widetilde{Y}$, and $\widetilde{\Gamma}^S$ agree with the guesses for τ, Y, and Γ^S.

B.2 Externally calibrated parameters

Table B.1 provides a list of parameters calibrated outside the model.

B.3 Solving for the transition path using the SSJ method

We now detail the implementation of the SSJ method of Auclert et al. (2021). Relative to a standard shooting algorithm to obtain impulse responses, the SSJ method provides major efficiency gains along two dimensions. The first is the computation of policy function responses by a single backward iteration. The second improvement offers an efficient method for the forward iteration of equilibrium distributions, both of which we implement following the fake news algorithm in Auclert et al. (2021).

We assume that the economy is in steady state at time $t = 0$. At $t = 1$, there is an unexpected and transitory shock to the economy (e.g. to productivity, discount rate, and OJS efficiency), and the economy returns to the same real allocations but potentially at different nominal prices by $t = T$ for some large T.

We first rewrite the model in real terms and relative prices so that the initial and terminal steady states coincide. We then cast the model as a DAG, depicted in Figure B.1. Different from the TANK model, the intermediate nodes in the DAG for the HANK model can be categorized into simple blocks and a heterogeneous agent block. An example of the former would be model components that relate various aggregate variables, such as the fiscal policy rule (Equation 28), the Taylor rule (Equation 6), or the expression for dividends and no-arbitrage that relate to the mutual fund (Equations 21 and 22). The latter is the most-complex model component, where heterogeneous agents make consumption-saving choices and labor market transitions, which play a key role in the dynamics of aggregates and distributions. Importantly, as we discuss in Section 4.2, model blocks *directly* interact not only via aggregate variables but also through the discretized worker distributions—namely, the distribution of employed individuals across human capital and match productivity and the distribution of unemployed individuals across human capital at the job search stage, and the distribution of employed workers across human capital, match productivity, and piece rate levels at the consumption/production stage. We apply the SSJ method in a context where discretized worker distributions are direct inputs and outputs along the DAG, as opposed to model blocks interacting only via aggregate variables.

There are three exogenous variables (shocks), eight endogenous variables (unknowns) whose dynamics we are interested in, and thus eight target sequences that must equal zero in equilibrium (market clearing and consistency conditions) in the sequence space.[67] Formally, let $\zeta = \left(\{\pi_t, Y_t, p^l_t, b_t, u_t, \theta_t, \Gamma^S_t, e2e_t\}_{t=0}^{T-1}\right)$ represent the path of the unknown endogenous variables and $\Theta = \left(\{z, \beta, \nu\}_{t=0}^{T-1}\right)$ represent the path of the exogenous variables. The system of equations, labeled as "targets" in the

[67]These equations in order capture the NKPC, market clearing for labor services, market clearing for mutual fund shares, the government budget balance, consistency of the unemployment rate, the free-entry condition, consistency of services firms' profits, and consistency of EE transitions.

rightmost node, that govern the transition path is:

$$H(\zeta;\Theta) = \begin{pmatrix} \Pi_t - \vartheta\left(\frac{p^l_t}{z_t} - \frac{\eta-1}{\eta}\right) - \frac{1}{1+r_{t+1}}\Pi_{t+1}\frac{Y_{t+1}}{Y_t} \\ \mathcal{L}_t - L_t \\ \mathcal{S}_t - 1 \\ (1+r_{t+1})(b_{t-1}+g_t-\mathcal{R}_t) - b_t \\ \mathcal{U}_t - u_t \\ \theta_t - q^{-1}\left(\kappa/\mathbb{E}J_t\right) \\ \mathit{\Gamma}^S_t - \Gamma^S_t \\ \mathcal{E}2\mathcal{E}_t - e2e_t \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}. \tag{B.7}$$

The main purpose of the DAG is to systematically solve for Jacobians that summarize the partial equilibrium responses of each node's each output (including targets in the rightmost node) with respect to each direct input to that node. We are then able to forward accumulate—that is, apply the chain rule in a systematic fashion—the partial Jacobians along a topological sort of the DAG to obtain the total Jacobians of any output (again including targets) with respect to any exogenous variable or unknown endogenous variable. A total Jacobian combines the direct and indirect responses of an output to an input. For example, the response of the expected match value conditional on contact $\mathbb{E}J$ (service firm block output) is affected directly by the real rate r, but also indirectly through how the real rate affects demand and output, which ultimately affect tightness, the unemployment rate, and the worker distributions. These Jacobians are what we also use in our decomposition exercise.

Having obtained the total Jacobians of targets $H(\zeta;\Theta)$ with respect to endogenous unknowns ζ and to exogenous variables Θ, we can apply the implicit function theorem to compute the response of any endogenous unknown $d\zeta$ to a change in the exogenous variables $d\Theta$. Formally, let H_ζ and H_Θ denote the total Jacobians of the targets with respect to endogenous unknowns and exogenous variables; then, the impulse responses of unknowns is given by:

$$d\zeta = \underbrace{-H_\zeta^{-1}H_\Theta}_{G_\zeta} d\Theta,$$

where G_ζ denotes the GE Jacobians of the endogenous variables.

Using the partial Jacobians of the intermediate variables and GE Jacobians of the unknown variables, we compute the GE Jacobians of the intermediate variables too. We then use the equivalence of the impulse response function with the moving-average process representation of a time series to simulate time-paths of aggregate variables.

Figure B.1: DAG representation of the HANK model with a frictional labor market and on-the-job search

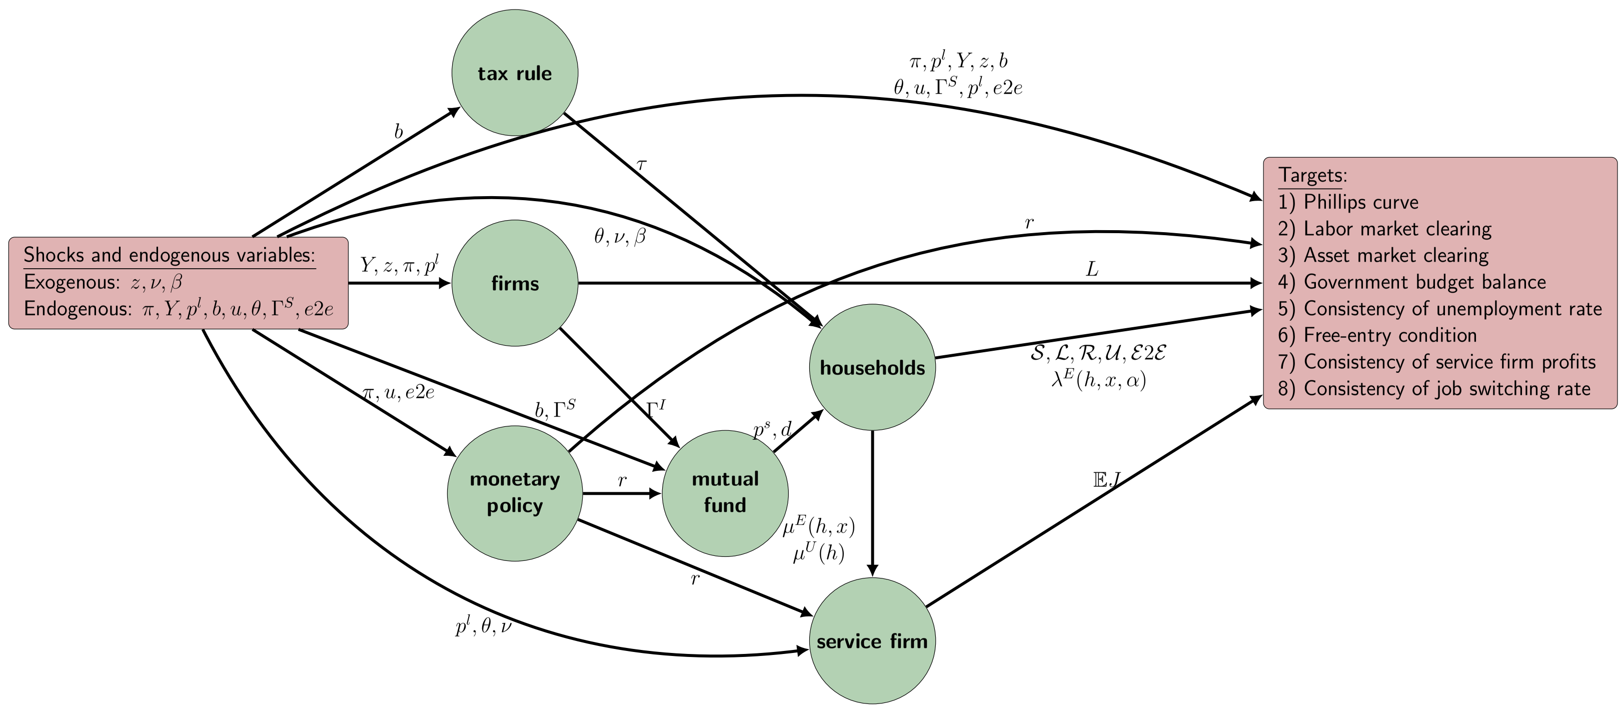

Table B.2: Variance decomposition of moments

	Share of variance explained by		
	z	ν	β
Y	0.051	0.056	0.893
u	0.232	0.033	0.735
EE	0.162	0.463	0.375
θ	0.280	0.021	0.699
π	0.338	0.158	0.504

Notes: This table presents a variance decomposition of output, unemployment rate, EE rate, tightness, and inflation. The columns represent the fraction of each moment's variance explained by shocks to productivity z, OJS efficiency ν, and discount factor β alone.

B.4 Variance decomposition

Table B.2 summarizes the variance decomposition exercise results in Section 4.3.

B.5 Additional results

This appendix provides additional results to complement Section 5.

Additional empirical results. Figure B.2 provides complementary figures to Section 5.1. Panel (a) presents a scatter plot of the monthly EE rate and the unemployment rate (i) prior to the Great Recession (1995–2007), (ii) during the Great Recession and the subsequent recovery (2008–2015), (iii) post-Great Recession (2016–2019), and (iv) the COVID-19 period and the subsequent recovery episode (2020–2024). The correlation between the two series is negative and significant, except during 2016–2019, when it turned slightly positive (but insignificant). That is, a tightening labor market characterized by a declining unemployment rate is typically also when the EE rate picks up, except for the 2016–2019 period. To present a more continuous view and to separate trend from the cycle, Panel (b) plots the rolling correlation between the cyclical components of log unemployment and EE rates over time using a five-year window. There is historically a strong negative comovement among the two series, which disappeared during 2016–2019 and reappeared following the COVID-19 recession. We conclude that the EE rate and the unemployment rate contain independent information about the labor market, a case in point provided by the 2016–2019 period.

Estimated shocks. In Sections 5.2 and 5.3, we conduct two exercises to quantify the magnitude of the missing inflation due to a low EE rate during the post-Great Recession episode and the additional rise in inflation due to a high EE rate during the recovery from the COVID-19 recession. In these exercises, we estimate demand and

Figure B.2: Unemployment rate and EE rate dynamics over time

(a) Unemployment vs EE rate (b) Rolling correlation: unemployment and EE

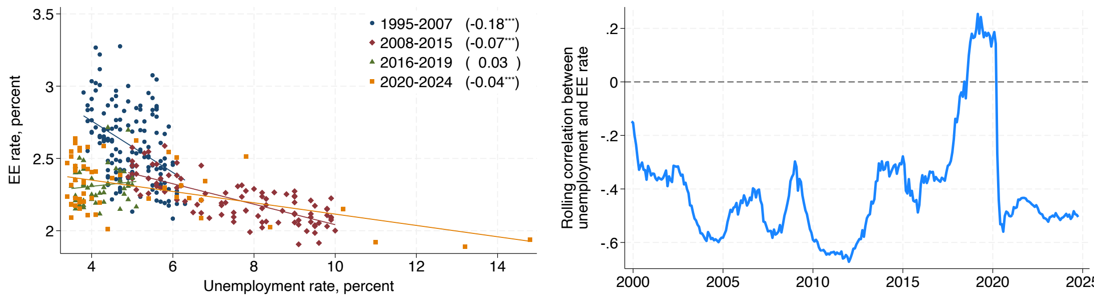

Notes: Panel (a) scatter plots the monthly EE rate (Fujita, Moscarini, and Postel-Vinay, 2023) and unemployment rate across different time periods. Values in parenthesis report the coefficient from regressing the EE rate on the unemployment rate and *** denotes significance at the 1 percent level. Panel (b) presents the rolling correlation between the cyclical components of the logs of unemployment and EE rates using a five-year window. Both series are logged and HP-filtered using a smoothing parameter of 10^5.

Figure B.3: Estimated paths of innovations to demand and OJS efficiency

(a) Shocks for post-Great Recession (b) Shocks for COVID-19

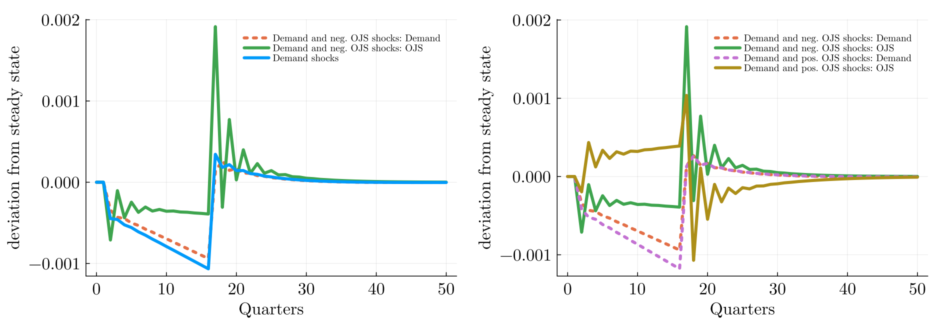

Notes: This figure plots the estimated paths of innovations to discount factor $\varepsilon_{\beta,t}$ and OJS efficiency $\varepsilon_{\nu,t}$ processes in Equation (29) for the post-Great Recession exercise in Section 5.2 and for the COVID-19 recovery exercise in Section 5.3. In Panel (a), the solid-blue line is the sequence of innovations to the discount factor to match the path of unemployment in the post-Great Recession episode without targeting the path of the EE rate. The dashed-orange and solid-green lines are the paths of innovations to the discount factor and OJS efficiency to match the same unemployment rate as in the first economy and the flat EE rate in the post-Great Recession period. In Panel (b), the dashed-orange and solid-green lines are the same as those in Panel (a) that generate the same unemployment rate and flat EE rate, while the dashed-purple and solid-dark-gold lines are the paths of innovations to the discount factor and OJS efficiency to match the same unemployment rate and the rise in EE rate during the COVID-19 recovery.

OJS efficiency shocks to match the unemployment rate and EE rate dynamics. Panels (a) and (b) in Figure B.3 show the estimated paths of innovations to the discount factor and OJS efficiency processes, $\varepsilon_{\beta,t}$ and $\varepsilon_{\nu,t}$, in Equation (29).

Impulse responses to a positive demand shock. Figure B.4 plots impulse responses to a negative discount factor (i.e., a positive demand) shock. A positive demand shock leads to an increase in tightness, output, total real profits, consumption, the real price of labor services, the real interest rate, inflation, and EE rate, and

Figure B.4: Impulse responses to a positive demand shock

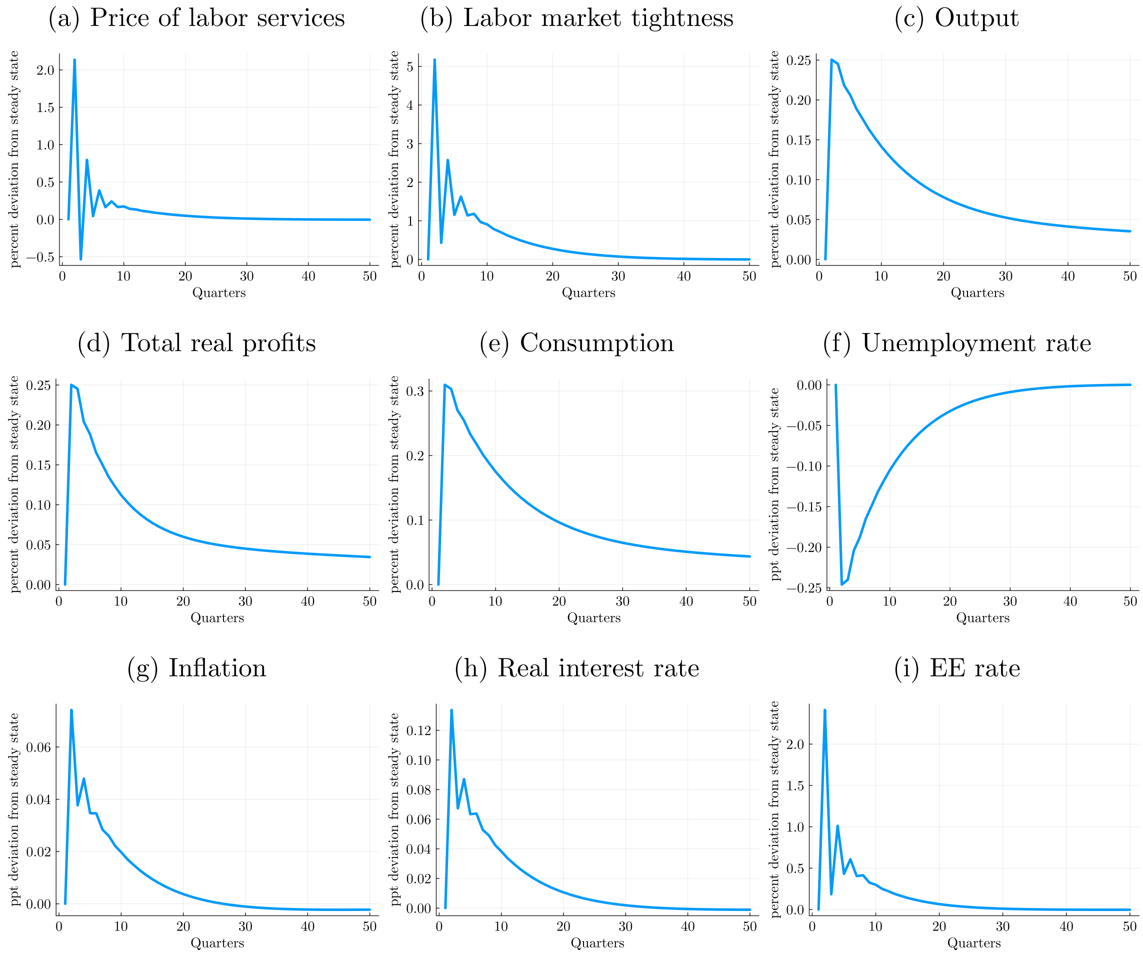

Notes: This figure presents impulse responses to a negative one standard deviation shock to discount factor β.

to a decrease in the unemployment rate.[68]

Impulse responses to a positive OJS efficiency shock. Figure B.5 plots impulse responses to a positive OJS efficiency shock. An increase in ν leads to a decline in tightness, output, total profits, and consumption, and to an increase in the real price of labor services, unemployment rate, inflation, the real rate, and EE rate.

Effects of OJS efficiency shocks on ALP and piece rate. Section 5.2 studies the macroeconomic implications of negative OJS efficiency shocks by comparing outcomes between two different transitions starting from the same steady state. There,

[68]The presence of oscillations in some IRFs is a well-known issue of the SSJ method when solving single-asset models with rich heterogeneity. When there is a single interest rate to clear the market, these oscillations occasionally appear in IRFs, as discussed by the authors: `https://tinyurl.com/46vk5tpu`.

Figure B.5: Impulse responses to a positive OJS efficiency shock

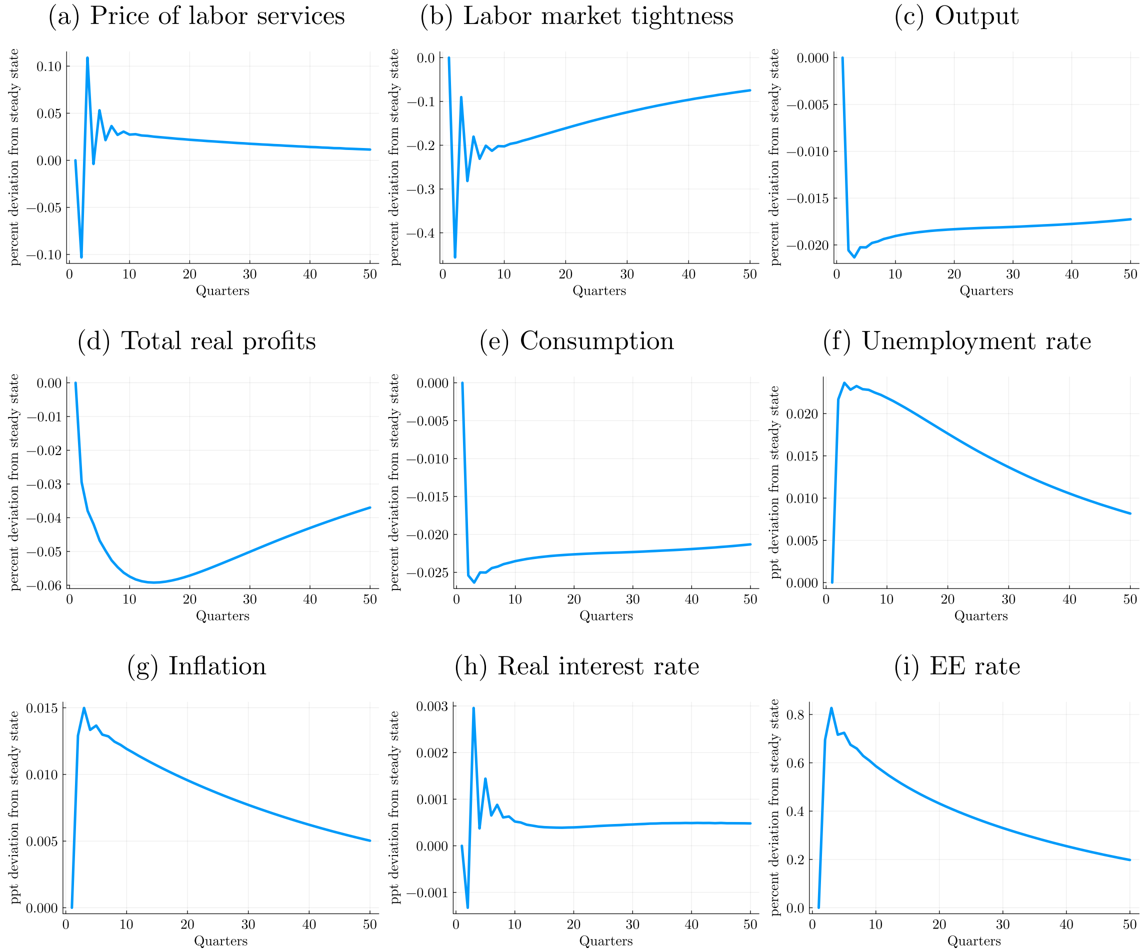

Notes: This figure presents impulse responses to a positive one standard deviation shock to OJS efficiency ν.

we also present ALP dynamics in Figure 5 to discuss the shocks' effects on output. In Figure B.6 Panel (a), we present the dynamics of average match-specific productivity. We show that because negative OJS efficiency shocks limit employed workers' ability to change jobs, they temper the increase in average match productivity during the labor market recovery. In fact, because unemployed workers accept their first job offer, they initially lower the average match quality, which only recovers after 20 quarters as they climb up the ladder slowly. Importantly, Panel (b) shows that the average piece rate increases slightly more when there are only positive demand shocks. This implies that wages in this economy are higher than in the economy with negative OJS efficiency shocks, despite similar ALP dynamics, as shown in Panel (d) of Figure 5. Therefore, inflationary wage pressures are stronger when we observe an EE rate rise coincidently with an unemployment decline. Quantitatively, this effect is small, as

Figure B.6: Effects of negative OJS efficiency shocks on productivity and piece rate

(a) Average match productivity (b) Average piece rate

(c) Productivity distribution (d) Piece-rate distribution

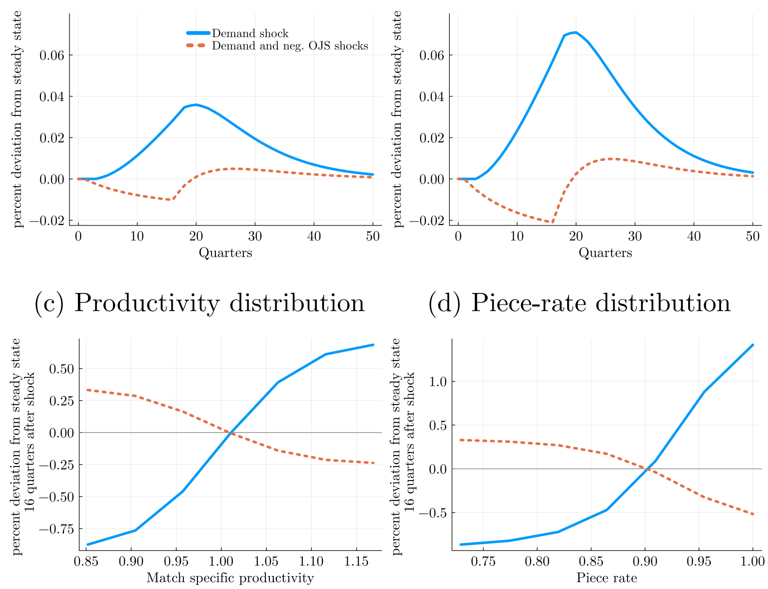

Notes: This figure presents the dynamics of average match-specific productivity over time (Panel (a)), average piece rate over time (Panel (b)), the distribution of match-specific productivity (Panel (c)), and the distribution of piece rate (Panel (d)) 16 quarters after the shock relative to their respective steady states in an economy with (i) only a series of positive demand shocks (solid-blue lines) and (ii) series of positive demand shocks and negative OJS efficiency shocks (dashed-orange lines), i.e., economies studied in Section 5.2. The shocks in the two economies are estimated such that they lead to the same path for the unemployment rate. The additional negative OJS shocks in the second economy are estimated to keep the EE rate unchanged.

evidenced by a very small gap between the average piece rates in the two economies. Therefore, the model generates a rise in inflation even without wage re-bargaining for incumbent workers with external offers.

We also compare the evolution of match-specific productivity and piece rate in the cross-section between the two economies, as looking at averages alone may mask interesting cross-sectional patterns. The next two panels in Figure B.6 plot *changes* in the distributions of match-specific productivity (Panel (c)) and of the piece rate (Panel (d)) 16 quarters after the shock relative to their respective steady states. Under positive demand shocks alone, the match productivity and piece-rate distributions exhibit a rightward shift. In contrast, when there are negative OJS efficiency shocks as well, both distributions shift leftward as these shocks decelerate the job ladder.

Implications for the slope of the empirical Phillips Curve. Sections 5.2 and 5.3 show that two recovery episodes with identical unemployment rates but different EE rate dynamics generate different inflation dynamics. As a result, the slope of the Phillips Curve would be different across these two recoveries. Figure B.7 presents

Figure B.7: Effects of OJS efficiency shocks on the slope of NKPC

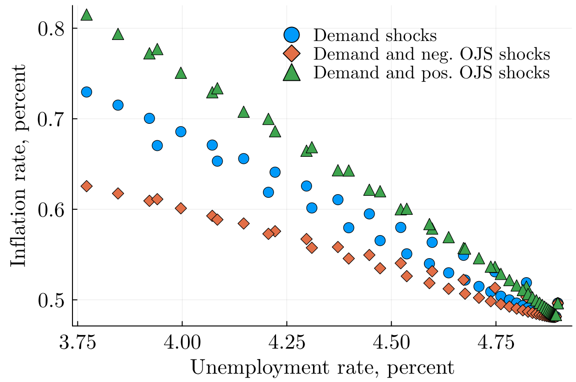

Notes: This figure plots the quarterly inflation rate against the quarterly unemployment rate for three different recovery episodes with the same unemployment dynamics: Blue circles represent a recovery episode with positive demand shocks alone; orange diamonds represent a recovery episode with positive demand shocks and negative OJS efficiency shocks generating a flat EE rate as in the post-Great Recession episode; green triangles represent a recovery episode with positive demand shocks and positive OJS efficiency shocks generating a rise in the EE rate as in the recovery from the COVID-19 recession.

simulations from the exercises in Sections 5.2 (Great Recession) and 5.3 (COVID-19).

The blue circles plot the quarterly inflation rate against the quarterly unemployment rate for the economy with positive demand shocks alone, while the orange diamonds represent the same for the economy with positive demand shocks and negative OJS efficiency shocks. Because the latter economy generates a smaller rise in inflation, the Phillips curve is flatter during the post-Great Recession episode. The green triangles plot inflation against unemployment for the post-COVID economy, which experiences a larger increase in inflation given the same unemployment rate path. Hence, the Phillips Curve is steeper during the recovery from COVID-19.

B.6 Decomposing the inflationary effects of OJS shocks

In Section 5.4, we discuss the decomposition of channels through which a rise in OJS efficiency increases inflation. Here, we provide further details.

We quantify the channels through which a unit increase in OJS efficiency ν translates to higher inflation by providing a decomposition of its impact response that utilizes the model's DAG representation in Figure B.1 and the system of Jacobians we compute to solve the model.[69]

The NKPC in Equation (8) reveals that—to a first-order approximation around the non-stochastic steady state—inflation is driven by the real price of labor services

[69]This exercise can be extended to the inflation IRF period-by-period; for brevity, we focus on the impact effect given that it is the largest. We acknowledge that if the channels considered here affect future marginal costs with varying degrees of persistence, their relative contributions to inflation may change over the response horizon.

Figure B.8: Decomposing the effect of OJS efficiency shocks on inflation

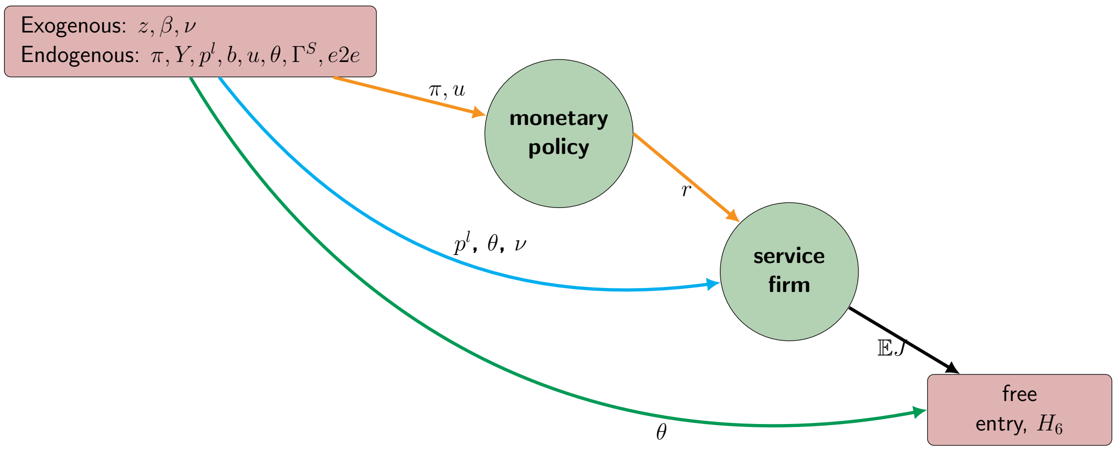

Notes: This figure summarizes model channels through which an OJS efficiency shock ν affects the real price of labor services p^l. The light-blue line captures the effect of ν on p^l through the shock's direct effect on expected match value for service firms. The green line captures the GE effect of ν on p^l through the shock's indirect effect on equilibrium tightness that dictates a change in p^l via the free-entry condition. Finally, the orange line refers to the GE effect of ν on p^l through the shock's indirect effects on equilibrium inflation and unemployment that necessitates a change in p^l via the real rate that affects the continuation value of service firms and thus the free-entry condition.

p^l, which, absent aggregate productivity z shocks, alone determines the real marginal cost p^l/z. Therefore, we provide a decomposition of the effects of an increase in ν on p^l to characterize inflation's behavior.

The decomposition applies the implicit function theorem to equilibrium conditions of choice in a particular order, to express an outcome variable as a (linear) function of the shocks and other endogenous variables. As discussed in Section 3.4, the free-entry condition is key to pinning down p^l. Hence, we focus on variables that enter this condition, whose effects can be grouped into labor market effects via ν directly and through market tightness θ, and discount rate effects via real rate r. Because r is determined in the monetary policy block reacting to inflation and unemployment, the contribution of r to p^l can be further broken down to the effects of π and u on p^l via r, denoted by $r(\pi)$ and $r(u)$, respectively. Similarly, the contribution of θ to inflation can be further dissected by recognizing that it is pinned down by Equation (26) equating the demand and supply of aggregate labor services, which we denote by $\theta(Y)$ and $\theta(\mathcal{L})$. Therefore, we end up with (i) the effect of ν on p^l through the shock's direct effect on expected value from a match for labor services firms (light-blue line in Figure B.8), (ii) the GE effect of ν on p^l through the shock's indirect effect on equilibrium tightness that dictates a change in p^l due to the change in the free-entry condition (green line), and (iii) the GE effect of ν on p^l through the shock's indirect effects on equilibrium inflation and unemployment that necessitates a change in p^l

Figure B.9: Decomposition of the inflationary effects of OJS efficency shocks

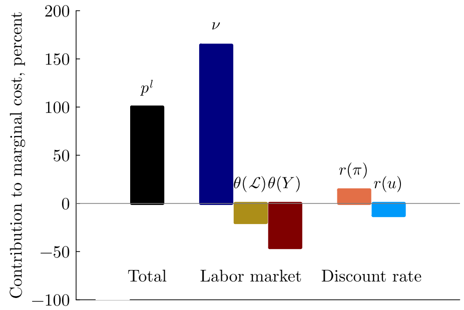

Notes: This figure presents a decomposition of the overall increase in the marginal cost, i.e., the real price of labor services p^l, explained by various channels in response to an increase in the OJS efficiency parameter ν—in particular, the fraction of the total change in p^l is accounted for by labor market effects and discount rate effects. ν refers to the direct effect of OJS efficiency on p^l; $\theta(\mathcal{L})$ refers to the effect of ν on market tightness θ through its effect on the total supply of labor services $\mathcal{L}$; $\theta(Y)$ denotes the effect of ν on θ through its effect on output Y; $r(\pi)$ denotes the effect of ν on real rate r through inflation π; and $r(u)$ refers to the effect of ν on real rate r through unemployment u.

due to the change in the real rate that affects the continuation value of labor services firms and thus the free-entry condition (orange lines).

Figure B.9 shows the contribution of each channel to the total rise in p^l in response to a positive unit shift in OJS efficiency ν. Now, we describe these results.

Labor market effects. An increase in ν leads to a decline in the expected match value, $\mathbb{E}J$, due to the composition and match duration channels, as discussed in Section 2.4. In this model, an additional channel that lowers $\mathbb{E}J$ is more frequent wage re-bargaining where outside offers can bid up wages without changing productivity.[70] All else the same, this decline in $\mathbb{E}J$ necessitates an increase in p^l for the free-entry condition to hold. We find that the direct effect of ν on p^l, labeled as ν in Figure B.9, explains 164 percent of the total rise in p^l.

There are further general equilibrium (GE) effects. In the labor market, an increase in ν leads to a decline in θ (Figure B.5 Panel (b)). According to the DAG in Figure B.8, θ enters the free-entry condition through its effects on the service firm. For an unmatched service firm, a lower θ increases the probability of filling its vacancy $q(\theta)$. In addition, a lower θ reduces the worker's probability of contacting other firms in the future, implying less-frequent wage re-bargaining and longer match durations. Thus, a matched firm's value J and hence $\mathbb{E}J$ increase. All else constant, this higher expected match value requires a decline in p^l for the free-entry condition to hold. The effect of ν on p^l through θ overall accounts for -66 percent of the increase in p^l.

[70] We note that the decline in firm value is especially driven by a shorter expected match duration, as we also showed in Appendix A.4. The composition and wage re-bargaining channels have relatively small effects on $\mathbb{E}J$.

Because the GE effect through θ is large with the opposite sign to the direct effect, we use the labor services market clearing condition in Equation (26) to further decompose the response of θ. The direct effect of ν on θ is through the supply of labor services in the HA block, $\mathcal{L} = \int F(h,x)\, d\lambda_t^E(s,h,x,\alpha,\beta)$. All else the same, a higher ν implies an improved match distribution and, therefore, a greater supply of labor services. For the labor services market to clear, the productivity gains that raise $\mathcal{L}$ should be offset by a decline in θ. This effect of ν on θ through the supply of labor services, $\theta(\mathcal{L})$, explains -20 percent of the rise in p^l. This is a result of match heterogeneity and productivity-enhancing job switches in the quantitative model.

The ν shock has a separate effect on θ through output Y, which declines in GE (Figure B.5 Panel (c)). This is driven by a lower demand (Panel (e)), which itself is a result of higher unemployment (Panel (f)) and a higher real rate. Lower output implies less demand for labor services $L = Y/z$. All else the same, a commensurate decline in θ is required for the labor services market to clear. This effect, $\theta(Y)$, accounts for -46 percent of the total increase in p^l. Overall, the GE effects of ν on θ through $\mathcal{L}$ and Y mitigate the direct effect of ν on p^l.

To summarize, a higher ν lowers expected match values due to the composition, match duration, and wage re-bargaining channels. This direct effect entails a compensatory rise in p^l to maintain the free-entry condition. However, a higher ν raises the supply of labor services and lowers demand, both of which require a decrease in tightness θ to clear the market for labor services. Lower θ leads to higher match values as firms expect to fill vacancies faster and face less quits. This requires a decline in p^l to reduce firm entry, partially mitigating the direct effect.

Discount rate effects. We now turn to the GE effects of ν on p^l through the monetary authority's reaction to inflation and unemployment, which we label as the discount rate effects. When ν rises, the unemployment rate and inflation both increase (Figure B.5 Panels (f) and (g)). Higher inflation induces a more than one-for-one increase in the nominal rate as we assume $\Phi_\pi > 1$ and hence an increase in the real rate. Higher real rate reduces the valuation of service firms (the right-hand side of Equation (19)), which in turn puts downward pressure on $\mathbb{E}J$. All else constant, this requires a rise in p^l for the free-entry condition to hold. We find that the inflation channel $r(\pi)$ accounts for 14 percent of the rise in p^l, a much smaller effect compared to the labor market effects. Similar reasoning implies that the rise in unemployment induces a decline in the real rate; $r(u)$ explains -12 percent of the rise in p^l.

Table B.3: HANK model without discount factor heterogeneity

Parameter	Explanation	Value	Target	Data	Model
κ	Vacancy creation cost	0.701	Unemployment rate	0.051	0.048
δ	Job separation probability	0.111	EU rate	0.038	0.033
ν	Search efficiency of employed	0.110	EE rate	0.020	0.022
π^U	Skill depreciation probability	0.134	Earnings drop upon job loss	-0.35	-0.34
σ_x	Standard deviation of productivity	0.071	Wage growth of job switchers	0.09	0.10
ϕ^E	Maximum share of output as wages	0.795	Labor share	0.67	0.70
ϕ^U	UI replacement rate	0.383	UI replacement rate	0.40	0.45
ϕ^R	Retirement benefit amount	0.406	Retirement income/labor income	0.34	0.40

Notes: This table summarizes the recalibrated parameters of the HANK model without discount factor heterogeneity.

B.7 Evaluating alternative monetary policy rule coefficients

Here, we provide computational details on how we evaluate the objective function in Equation (30) under alternative Taylor-rule coefficients as described in Section 6.

Naively, solving and simulating the model, and calculating the variances of inflation and output to evaluate the objective function in Equation (30) under alternative Taylor-rule coefficients in Equation (31) requires the repeated computation of the entire SSJ system. As discussed in Section 4.2, worker distributions directly enter the equilibrium conditions and agents' problems, as opposed to model blocks depending on one another only through aggregate variables. This added layer of complication in our context makes the computation of derivatives costly. To overcome this challenge, we use the approach in McKay and Wolf (2022). Details are provided in Section 3 of Online Appendix.

B.8 The HANK model without discount factor heterogeneity

In Section 6, we use an alternative calibration of our HANK model without discount factor heterogeneity to study how optimal Taylor-rule coefficients change in a model that approximates an environment with perfect self-insurance. In this model, we assume a common discount factor, which we set to the highest value of the discount factor $\bar{\beta} + \Delta\beta = 0.989$ in the baseline calibration *with* discount factor heterogeneity. Therefore, in this version, we eliminate two parameters ($\bar{\beta}$ and $\Delta\beta$) as well as their associated empirical targets. Then, we recalibrate the remaining set of parameters to match the same data moments as before. The resulting parameter values, and a comparison of data and model moments are presented in Table B.3.