Some Like It Hot: Monetary Policy Under Okun's Hypothesis*

Felipe Alves Giovanni L. Violante

December 22, 2023

Abstract

Okun (1973) argued that running a 'high-pressure economy' could persistently improve labor market outcomes of low-wage workers. This hypothesis is more than ever pertinent today since the ZLB, by exacerbating downturns, is especially costly for the bottom of the distribution. By Okun's premise, the recent reformulation of the Fed's framework—which intends to run the economy hot for longer during economic recoveries—has the potential to be more inclusive toward low-skill workers, possibly at the cost of higher inflation. To evaluate this conjecture, we develop a Heterogeneous-Agent New-Keynesian framework with a three-state frictional model of the labor market where employment trajectories of low-skilled workers are more exposed to the business cycle, and where recessions have long-lasting effects on labor force participation and earnings, in line with U.S. data. We find that the new monetary policy strategy of the Fed gives rise to a meaningful 'inflation-inclusion trade-off', which we quantify. While the two novel components of the reformed framework, average inflation targeting and employment shortfall targeting, are ineffective (and sometimes deleterious) in isolation, jointly they succeed in inducing significant gains for low-wage workers at only a moderate cost in terms of inflation.

JEL Codes: E21, E24, E31, E32, E52, J24, J64

Keywords: Distribution, Earnings Losses Upon Displacement, Inflation, Labor Market Frictions, Monetary Policy, Participation, Okun, Unemployment, Upward Mobility, Zero Lower Bound.

*Alves: Bank of Canada, felipe.a.alves0@gmail.com. Violante: Princeton University, CEPR, IFS, IZA and NBER, violante@princeton.edu. We thank Edouard Challe, Marco Del Negro, Jordi Gali, Mark Gertler, Greg Kaplan, Ben Moll and many workshop participants at the Bank of Canada, Bank of Chile, Bank of Portugal, 2023 Barcelona Summer Forum, Canadian Economic Association, Danmarks Nationalbank, Dartmouth, ECB, 2023 IARIW–Bank of Italy conference (Naples), Maryland, National Bank of Belgium, 2023 NBER Summer Institute, OFCE-Sciences Po-PSE Workshop, Penn, Richmond Fed, 2023 SED (Cartagena), St. Louis Fed, Wharton, Workshop on Micro Data Meet Macro Models (Lund), Workshop on Heterogeneity in Macroeconomics (Rennes), Yale.

1 Introduction

The new monetary policy framework unveiled by Federal Reserve Board in 2020 is motivated by the concern that, in a permanently low interest rate environment, there might be less scope to support the economy during a downturn. When the zero lower bound (ZLB) constrains the conduct of monetary policy, *the result is worse economic outcomes in terms of both employment and price stability, with the costs of such outcomes likely falling hardest on those least able to bear them* (Powell, 2020).

To address this challenge, the new framework introduced two major novelties. The first one is the shift from a strict 2 percent inflation targeting to a more flexible interpretation of price stability whereby the Fed aims to achieve inflation that averages 2 percent over a period of time. As a result, following periods where inflation runs persistently below 2 percent, the Fed would be willing to keep rates low and accept an inflation rate that overshoots the target for some time during the recovery. This strategy is, however, intended to be asymmetric, i.e., inflation would not be pushed persistently below 2% following unexpected inflationary episodes (Clarida, 2022).

The second novelty is a reinterpretation of the maximum employment goal –the other mandate of the Fed's mission– as *broad-based and inclusive*. This policy shift originated mainly from FedListens, a series of events held around the country that engaged a wide range of organizations to hear about how monetary policy affects peoples' livelihoods. In the foreword of the final report on this project (Federal Reserve System, 2020), Chair Powell writes that: *One clear takeaway from these events was the importance of sustaining a strong job market, particularly for people from low- and moderate-income communities. Everyone deserves the opportunity to participate fully in our society and in our economy.*[1] Coherently, the new framework specifies that monetary policy is no longer informed by deviations, but only by shortfalls of employment from its maximum level, suggesting again an element of asymmetry in the response to aggregate disturbances. The Fed would remain equally aggressive as in the past in cutting rates to face a surge in unemployment, but it will wait longer to raise them during tight labor markets periods, as long as this loose policy stance does not pose an inflationary danger.

Jointly, these two shifts in the conduct of monetary policy suggest that, during recoveries, rates might be kept *lower for longer* in order to (i) make up for past disinflation and

[1] Interestingly, also from the latest monetary policy strategy review of the European Central Bank (ECB), it emerges that even though the ECB's objective remains only price stability, monetary policy now recognizes more explicitly the existence of indirect effects of policy choices on certain dimensions of social progress such as wealth and income inequality (Ioannidis et al., 2021).

(ii) protract a strong job market. This strategy, which would reduce the likelihood of recessions being aggravated by the ZLB, would mostly serve low-wage workers, who are more severely hurt by downturns.

The idea that workers at the bottom of the distribution especially benefit from a hot economy is not new. In his article entitled *Upward Mobility in a High Pressure Economy* (Okun, 1973), Arthur Okun already conjectured that a high-pressure economy, i.e. an economy that keeps slack to a minimum, can persistently improve the economic circumstances of more disadvantaged workers by creating opportunities for them to remain attached to the labor force, find steady employment, strengthen their skills, and gradually climb the job ladder. As he put it: *the sacrifice of upward mobility must be carefully reckoned as one part of the high cost of accepting slack as an insurance policy against inflation* (page 244).

The post-pandemic experience of the U.S. offers some support to Okun's high-pressure hypothesis and to the idea that the Fed is operating under a new framework. First, according to several indicators, the post-Covid labor market has been extremely tight and, simultaneously, for the first time in decades the economy has witnessed a substantial compression in the earnings distribution as well as unusually high employment rates for low-educated workers (Autor et al., 2023). Second, the Fed seems to have abandoned its strategy of 'preemptive restraint'. In the recovery from the Great Recession, the Fed began raising rates as soon as inflation started picking up, thus aggressively cooling down a labor market that was just beginning to heat. In the post-pandemic recovery, instead, the Fed kept rates lower for longer: it did not hike rates after inflation took off in early 2021, but it waited an additional year during which unemployment kept falling steadily, and inflation kept rising.

Taken together, these observations point to a potential stark *inflation-inclusion* trade-off, intrinsic in the new framework, along the lines envisioned by (Okun, 1973). On the one hand, the new regime, by reducing the negative effects associated with the ZLB, should persistently improve labor market outcomes of low-wage workers. On the other hand, running the economy hot for longer can open the door to more inflation. The objective of our paper is to develop a quantitative macroeconomic model to assess the effectiveness of the new framework, ask how steep this trade-off is, investigate its key determinants, and assess whether alternative rules might present a better trade-off.

The first challenge we face is to give content to Okun's hypothesis and formalize it within a modern macroeconomic framework. At the time of his writing, Okun offered only suggestive empirical evidence in support of his conjecture. Nearly half a century later, much more is understood about the labor market trajectories of workers across the entire distribution, and a number of facts support the insight that a hot labor market especially favors low earners. First, employment outcomes of the low-income groups are more cyclically sensitive than for the rest of the population (Aaronson et al., 2019; Cajner et al., 2017). When the economy weakens, these groups suffer disproportionately through higher unemployment; when it recovers, their employment increases disproportionately through higher job finding rates and lower job separation rates.[2] Second, the key factor driving fluctuations of labor force participation over the business cycle is the size of the unemployment pool because a large share of unemployed workers leaves the labor force (Hobijn and Şahin, 2021). This 'participation cycle' is more pronounced for low-skill groups who have higher unemployment incidence and duration and are close to indifference between being in or out the labor force. Keeping these workers away from unemployment thus persistently improves their attachment to the labor force. Third, earnings losses upon displacement are sizable, long-lasting, and countercyclical (Davis and Von Wachter, 2011).[3] Combining these three mechanisms implies that a protracted strong labor market curtails the frequency layoffs and, as a result, exit from the labor force and, by keeping workers employed, prevents the erosion of human capital. Especially so for the low-income groups. These three channels—differential exposure to aggregate fluctuations, labor force attachment wedge, and persistent effects of displacement—is how we concretely think about Okun's high-pressure hypothesis.

We then incorporate these mechanisms into a heterogeneous-agent incomplete-markets framework where workers move across three labor market states: employment, unemployment, and nonparticipation. It is labor market frictions which prevent full employment: some workers who would retain their job are laid off; others who search for a job, and would work at the ongoing equilibrium wage, can't find it. As in the data, both the level and the cyclicality of separation and job finding rates depend on individual skill levels, with low-wage workers suffering more exposure to aggregate fluctuations. A large systematic exit rate from unemployment into nonparticipation gives rise to the participation cycle in the model, and makes recessions times where many workers, especially low-skilled ones, become unattached from the labor force. In the model, individual skill levels evolve stochastically during a career depending on labor market status: skills grow during employment through returns to experience, and gradually depreciate when the

[2] Recent work by Graves et al. (2023) looking at the response of labor market flows to identified monetary policy shocks also finds similar patterns for the employment outcomes of low-educated workers.

[3] In addition, there is evidence suggesting that earnings losses from job loss are worse at the bottom of the distribution precisely because they lead to disattachment from the labor force, i.e. through the extensive margin (Guvenen et al., 2017; Athey et al., 2023).

worker is not employed. These three Okun's channels reinforce each other in generating heterogeneous labor market trajectories across workers and over the cycle. Aggregate shocks on both demand and supply side drive business cycles in the model. Nominal wage rigidity gives monetary policy the ability to impact the real economy, as in the standard representative agent New Keynesian model (Erceg et al., 2000), but here its consequences are uneven across the distribution.

The model is calibrated to match match several dimensions of the US labor market and is consistent with (i) the size and cyclicality of flows across the three labor market states (Krusell et al., 2017; Cairó et al., 2022; Graves et al., 2023), and (ii) the dependence of job-finding and separation rates on skill level (which we document). Our calibration also captures key facts of individual earnings dynamics such as average life-cycle earnings growth and persistent earnings losses upon displacement (Davis and Von Wachter, 2011). We solve for the linearized equilibrium dynamics using the sequence space Jacobian method developed by Auclert et al. (2021), which we extend to incorporate an occasionally binding zero lower bound (ZLB) and asymmetric policy rules by adapting insights from Holden (2016), and Hebden and Winkler (2021) to our setup.

We use the model to study the inflation-inclusion trade-off induced by the new monetary policy framework (Lower for Longer), carefully separating the role played by its two novel components, average inflation targeting (AIT), and tracking only the employment shortfall instead of upward and downward deviations (Shortfall).

We investigate the differential dynamics of the economy under the traditional strict inflation targeting rule (IT) and the new Lower for Longer (LfL) strategy in the short- and the long-run. Our short-run exercise illustrates how the rules modify the dynamics of a single business cycle. Namely, it asks how the dynamics of the US economy around the Great Recession would have been had the Fed followed a Lower for Longer strategy. Our long-run analysis compares the ergodic distribution of the model economy subject to aggregate disturbances under the new and old frameworks. In all these exercises, our assessment of the inflation-inclusion trade-off is centered on a credible estimation of the *gains* of inclusive stabilization policy by carefully modeling labor market dynamics across the distribution. We cast the costs in terms of higher inflation which, through the wage Phillips curve of our model, arises directly from a widher labor wedge that reflects deviations from aggregate productive efficiency.[4]

[4] There exist several other potential uneven costs of inflation across households, such as heterogeneous expenditure baskets, nominal net positions, and degrees of nominal wage rigidity. We purposefully neutralize all these additional channels by designing an economy with a unique final good, real assets, and

Taken together, our set of experiments offers four main findings. The first message pertains to the two components of the LfL strategy taken in isolation. The AIT component is effective at reducing the likelihood and severity of ZLB episodes. As a result, it acts to lessen the adverse effects of recessions on labor market outcomes of low-wage workers but, quantitatively, the gains it generates for this group are small. The Shortfall component is instead deleterious for the economy because it produces additional inflation and aggregate volatility, with little or no average labor market gains. There are two main reasons behind this result. The first reason is that this rule doesn't address the negative consequences of the ZLB. The second is that the persistent and sizable productivity and participation gains generated at the end of long expansions under this rule can actually turn into a liability once the economy falls into a recession. This is because these forces tend to exacerbate the positive output gap opened up by a contraction in aggregate demand, forcing inflation to decline by more in response to negative demand shocks which, in turn, exacerbates the ZLB and the depth of recessions. Interestingly, these perverse effects are not apparent in our short-run analysis, a feature that highlights the importance of both the short-run and long-run perspectives.

Second, we find that the combined LfL strategy (AIT + Shortfall) is remarkably successful in producing sizable gains at the bottom of the distribution at only moderate costs in terms of additional inflation. In particular, by decreasing the likelihood of the ZLB binding, the AIT component almost entirely eliminates the perverse dynamic effects of the Shortfall component during contractions, while preserving its positive implications for low-wage workers along expansions.

Third, we assess quantitatively the slope of the inflation-inclusion trade-off implicit in the LfL rule, i.e., how costly it is, in terms of additional inflation, to achieve some given gains at the bottom of the skill distribution. To answer this question we produce what we call *Okun's cones*. These plots show, relative to the non-stochastic steady state, the improvement in a number of distributional outcomes of interest (e.g., lower unemployment, higher participation, more robust earnings at various points in the skill distribution) as a function of additional percentage points (ppts) of inflation required to achieve them. The cones thus map out the entire menu of outcomes available to a policymaker considering running the new LfL strategy. Our third message is that this long-run trade-off appears quite favorable. Just by eliminating the contractionary inflation bias caused by the ZLB would produce permanent gains at the bottom quartile equal to 0.5 ppts of lower unem-

constant real wage. Thus, we accurately describe the distribution of gains but, for simplicity, reduce costs to a single number, i.e., the extra inflation generated by the policy.

ployment and 3 ppts of higher labor earnings. Going even further, an inequality-averse central bank willing to tolerate 25 basis points of additional average inflation in excess of its target could permanently boost participation by 1 ppt and earnings by 10 ppts for workers in the bottom quartile of the wage distribution. As the central bank keeps pushing for bigger gains, the systematic positive inflation bias becomes sizable and, according to recent evidence on expectation formation, this trade-off is less and less likely to be sustained in the long run.

Finally, we argue that, through the lens of our model, there exist alternative monetary policy rules that present a more favorable inflation-inclusion trade off.

Related Literature. Fueled by the redefinition of the maximum employment mandate of the Fed as broad-based and inclusive, a number of recent papers have started to investigate the effect of monetary policy on racial income and wealth gaps (Bartscher et al., 2021; Bergman et al., 2020; Lee et al., 2021; Nakajima, 2022; Cairó and Lipton, 2023). Our approach differs from this literature along several dimensions. First, we focus on a more comprehensive notion of skills and inequality whose sources can be both unobservable (e.g., innate abilities, specialized knowledge) and observable demographic characteristics (e.g., race, gender, education). Second, these papers contain at most one of the three Okun's channels we model (the uneven exposure), but abstract from human capital losses from non-employment and from the labor force attachment margin. We show that it is the last two mechanisms that mostly matter for our findings. Third, most of these papers concentrate on the impact of expansionary monetary policy *shocks* to standard Taylor rules on racial inequality, while our emphasis is on alternative monetary policy *rules* in reaction to aggregate shocks that can generate long-lasting gains for low-income groups. Focusing on rules allows us to quantify the inflation-inclusion trade-off.[5]

Two recent papers explore asymmetric monetary policy rules within representative agent models. Bianchi et al. (2021) study how an asymmetric rule with respect to inflation gaps can correct the deflationary bias caused by the ZLB. Bundick and Petrosky-Nadeau (2021) analyze the impact of employment shortfall-based rules on business cycle dynamics. Relative to this work, our interest lies in assessing how inclusive these rules are with respect to the bottom of the wage distribution. This requires moving beyond a standard representative agent model and modelling new transmission mechanisms. Fernández-Villaverde et al. (2023) model an incomplete-market economy with heterogeneous agents

[5] An exception in this list is contemporaneous work by Cairó and Lipton (2023), who also explores the impact of the Shortfall rule on the unemployment gap between different racial groups.

where the ZLB binds occasionally. Their main insight is that, by increasing the frequency of ZLB episodes, a lower inflation target amplifies aggregate precautionary saving which, in turn, lowers the natural real rate and further reduces the effectiveness of monetary policy. We relate to this work by identifying a novel channel, based on the persistent effects of the business cycle on labor force participation and human capital, through which a monetary policy rule (the Shortfall), can worsen the severity and frequency of ZLB episodes.

More broadly, our study also relates to a small, but growing, literature on the insight that monetary policy can have long-run repercussions on the economy (Jordà et al., 2020; Fornaro and Wolf, 2020; Ma and Zimmermann, 2023). In our framework, persistent effects of monetary policy on aggregate productivity operate through the human capital channel—workers build skills when employed and suffer long-lasting scars upon displacement, in line with micro evidence. We show that, depending on the rule adopted, this mechanism can either improve or perversely worsen the inclusion-inflation trade-off.

The rest of the paper is organized as follows. Section 2 outlines the model and Section 3 its parameterization. Section 4 contains all our counterfactuals. Section 5 concludes.

2 Model

It is useful to start with an overview of the model. Individuals consume, save and can be employed, unemployed, or out of the labor force. One way to depict the structure of the labor market is to envision three separate "islands"(Lucas and Prescott, 1978; Alvarez and Shimer, 2011). Some transitions across islands are exogenous and some are endogenous. Workers endogenously choose whether to participate to the labor market. Job finding opportunities for unemployed, and at a lower rate for non-participants, arise exogenously, but workers choose whether to accept them or not. Layoffs are also exogenous. Individuals are endowed with efficiency units of labor that are subject to persistent uninsurable idiosyncratic shocks whose distribution depends on labor market status. Both job finding and job separation rates are indexed by skill level.

Every worker on the employment island adheres to labor unions which set nominal wages. Monopolistically competitive intermediate good producers with flexible prices take wages as given and demand the profit-maximizing amount of labor. A competitive final good sector packages the intermediate goods into a final good, the numeraire of the economy, and sells it to households.

Households hold and trade shares of a mutual fund which owns claims to the firms' profits and holds government bonds. The government (fiscal authority) finances expenditures and transfers by levying taxes on households and issuing debt. The central bank (monetary authority) sets the nominal interest rate. Both authorities follow exogenous policy rules.

2.1 Households

Time is continuous and indexed by t. The economy is populated by a continuum of households (or individuals) with measure 1 who discount the future at rate $\tilde{\rho} > 0$ and face mortality rate θ. Let $\rho = \tilde{\rho} + \theta$ be the effective discount rate.

At any date t, individuals can be in one of three mutually exclusive labor market states s_t: employed and earning labor income, ($s_t = e$), unemployed and searching for a job ($s_t = u$), and non-participant, or outside the labor force, ($s_t = n$). Among the unemployed, we distinguish between those who are eligible ($u = u_1$) and not eligible ($u = u_0$) for unemployment insurance (UI) benefits. Workers gain eligibility when they enter the unemployment pool due to an exogenous separation, and they lose it at some constant rate which reflects benefit duration. Among those out of the labor force, we distinguish between "active" non-participants ($n = n_1$) and "passive" non-participants ($n = n_0$). The former can, at a lower rate than the unemployed, transition back into employment, while the latter cannot. This differentiation is meant to capture the fact that the pool of non-participants is heterogeneous (Hall and Kudlyak, 2019). Some individuals in this pool are able and willing to work. Others, instead, are unable to accept any job offer (e.g., because they are sick, or heavily involved in household care) or not searching at all (e.g., because discouraged by the failure of previous job search).

Households derive utility from consumption c_t, and suffer disutility from the effort cost κ^s associated to being in labor market status s (the extensive margin) and from the effort cost of working h_t hours (the intensive margin). We specify the following functional form for period utility

$$u^s(c_t, h_t) = \frac{c_t^{1-\gamma}}{1-\gamma} - \psi \frac{h_t^{1+\frac{1}{\sigma}}}{1+\frac{1}{\sigma}} - \kappa^s \tag{1}$$

where $\gamma > 0$ is the coefficient of relative risk aversion as well as the inverse of the IES, and $\sigma > 0$ is the Frisch elasticity of labor supply. We assume that $\kappa^e > \kappa^u > \kappa^n \geq 0$.

Each individual is endowed with efficiency units of labor (or skills) z evolving according to a Ornstein-Uhlenbeck diffusion process which depends on labor market status s_t:

$$d\log z_t = \Big\{ -\rho_z \log z_t + \mathcal{I}_{\{s_t=e\}}\, \delta_z^+ - \mathcal{I}_{\{s_t \neq e\}}\, \delta_z^- \Big\} dt + \sigma_z d\mathcal{W}_t \tag{2}$$

	e	u_1	u_0	n_1	n_0
e	$\ddots$	λ^{eu}_{zt}	$\times$	▶	η^{en_0}
u_1	$\lambda^{ue}_{zt} \cdot \triangleright$	$\ddots$	$\eta^{u_1 u_0}$	▶	η^{un_0}
u_0	$\lambda^{ue}_{zt} \cdot \triangleright$	$\times$	$\ddots$	▶	η^{un_0}
n_1	$\lambda^{ne}_{zt} \cdot \triangleright$	$\times$	▶	$\ddots$	$\eta^{n_1 n_0}$
n_0	$\times$	$\times$	$\times$	$\eta^{n_0 n_1}$	$\ddots$

Table 1: Transition matrix across the 5 employment states. The $\times$ symbol means that transition cannot happen. The ▶ symbol means that an endogenous participation decision moves the individual in that state. The $\triangleright$ symbol means that an endogenous job acceptance decision moves the individual into employment. $\lambda^{ss'}_{zt}$ and $\eta^{ss'}$ are exogenous Poisson rates. The diagonal dots stand for the negative of the sum of all the other entries on that line.

When workers are employed ($s_t = e$), skills drift up at rate $\delta^+_z > 0$, and when they are not employed ($s_t = u, n$) they drift down at rate $\delta^-_z < 0$. The parameter $\rho_z > 0$ measures the degree of mean reversion in skill dynamics, the standard deviation σ_z determines uncertainty about future realizations, and $\mathcal{W}_t$ is a Wiener process. Upon death, workers are replaced by an offspring with log skill drawn from a Normal distribution with mean $\bar{z}_0$ and variance σ^2_{0z}.

Every period individuals can transition between employment states through a combination of exogenous Poisson rates and optimal mobility decisions. Table 1 describes all the possible transitions and their endogenous/exogenous nature.

At any date t, every employed and unemployed worker can choose to quit the labor force and enter active non-participation (rows 1, 2, 3 of Table 1). Similarly, an active non-participant can choose to re-enter the labor force as unemployed ineligible for UI (row 4). Employed workers who decide to remain attached can still be laid off, and thus move from e to u at an exogenous rate λ^{eu}_{zt} which depends on the worker's skill level z (row 2). Eligible and ineligible unemployed who choose to remain attached to the labor force draw an employment opportunity at an exogenous rate λ^{ue}_{zt}. Eligible ones decide whether to accept it or not (rows 2 and 3).[6] In addition, UI benefits can expire at rate η^{u1u_0} and an eligible unemployed becomes ineligible (row 2). Also active participants receive job

[6]The unemployed ineliglible for UI always accept job offers because in equilibrium there is a unique wage per effective hours, and if they did not want to work, they would choose non-participation where the fixed cost κ^s is lower. Eligible unemployed instead may turn down job opportunities if UI benefits are generous enough.

opportunities at rate λ_{zt}^{ne} and decide whether to accept them or not (row 3). All workers can exogenously transition into passive nonparticipation at rate η^{s,n_0} (rows 1, 2, 3, 4). At rate $\eta^{n_0 n_1}$, passive nonparticipants become active again (row 5).

Employed individuals earn labor income $w_t h_t z_t$, where w_t is the real wage per effective hour, and eligible unemployed receive benefits $b(z_t)$. We let UI benefits be a function of current worker productivity z_t, as a proxy for actual replacement rates. Both types of income are taxed at a proportional rate $\mathfrak{t}_t$. Every household is entitled to a lump-sum transfer ϕ_t. Households can save through a financial asset a_t with rate of return r_t, and can borrow up to the exogenous credit limit $-\bar{a}$. Newborn workers are start with zero wealth holdings. Perfect annuity markets insure workers against survival risk, so that the wealth holdings of the deceased are redistributed to surviving workers in proportion to their position in the mutual fund.[7]

Household problem The triplet (s, a, z) fully characterizes the household state vector. The dynamic problem solved by the household at time t is a mix of an optimal control problem, the choice of $c_t > 0$, and two optimal stopping problems: a continuous one, the participation decision $\mathfrak{p}_t^s \in \{0, 1\}$, and one arising at random Poisson jump times, the job acceptance decision $\mathfrak{f}_t^s \in \{0, 1\}$. The stochastic nature of the problem is due to both the exogenous Poisson arrival rates that determine transitions across labor market states, and the exogenous diffusion that determines the evolution of skills z_t. Conditional on these realizations, wealth evolves deterministically. Let $v_t^s(a, z)$ be the value at date t of an individual with employment state s, wealth a, and productivity z.

Consider first the problem of the passive non-participant (n_0):

$$v_0^{n_0}(a_0, z_0) = \max_{\{c_t\}_{t\geq 0}} \quad \mathbb{E}_0 \int_0^{\tau^{n_1}} e^{-\rho t} \mathfrak{u}^n(c_t, h_t)\, dt + e^{-\rho \tau^{n_1}} v_{\tau^{n_1}}^{n_1}(a_{\tau^{n_1}}, z_{\tau^{n_1}}) \tag{3}$$

$$\begin{aligned} & s.t. \\ c_t + \dot{a}_t &= r_t a_t + \phi_t \\ a_t &\geq -\bar{a} \end{aligned}$$

Passive non-participants do not receive any job opportunity. At rate η_1, with τ^{n_1} being the first arrival rate of this event, they become active non-participants and enter employment status n_1. The conditional expectation reflects the uncertainty in transition rates and in the evolution of skill dynamics. In addition to the participation decision $\mathfrak{p}_t^{n_0}$, at every instant

[7] To ease notation, we fold this adjustment directly into the rate of return of the fund r_t, which should be interpreted as including the return from insurance contracts.

the worker chooses its consumption flow c_t. The last two lines of this problem state the budget constraint (in real terms) and the borrowing limit.

The problem of the active non-participant (n_1) is:

$$\begin{aligned} v_0^{n_1}(a_0,z_0) = \max_{\{c_t\}_{t\geq 0},\tau^*} \quad & \mathbb{E}_0\Bigg[\int_0^{\tau^{\min}} e^{-\rho t}\mathfrak{u}^n(c_t,h_t)\,dt + \mathbb{I}_{\{\tau^{\min}=\tau^e\}} e^{-\rho\tau^e}\max\left\{v_{\tau^e}^e(a_{\tau^e},z_{\tau^e}), v_{\tau^e}^{n_1}(a_{\tau^e},z_{\tau^e})\right\} \\ & +\mathbb{I}_{\{\tau^{\min}=\tau^*\}}e^{-\rho\tau^*}\left(v_{\tau^*}^{u_0}(a_{\tau^*},z_{\tau^*})-\xi\right) + \mathbb{I}_{\{\tau^{\min}=\tau^{n_0}\}}e^{-\rho\tau^{n_0}}v_{\tau^{n_0}}^{n_0}(a_{\tau^{n_0}},z_{\tau^{n_0}})\Bigg] \\ s.t. & \\ c_t + \dot{a}_t &= r_t a_t + \phi_t \\ a_t &\geq -\bar{a} \end{aligned} \tag{4}$$

Active non participants receive job opportunities at rate λ_{zt}^{ne}, with τ^e being the first arrival time of this event. Conditional on receiving this job offer, they decide whether to accept it or not. At every instant, the non-participant chooses whether to remain unattached $\left(\mathfrak{p}_t^{n_1}=0\right)$ or re-enter the labor force $\left(\mathfrak{p}_t^{n_1}=1\right)$, in which case they become unemployed, but are not eligible for UI benefits ($u=u_0$). We assume that re-entering the labor force involves a small fixed switching cost ξ.[8] The optimal stopping time τ^* represents the first instant in which the choice $\mathfrak{p}_t^{n_1}$ switches from 0 to 1. Finally, at rate $\eta^{n_1 n_0}$ (with τ^{n_0} being the first arrival rate of this shock) active non-participants become passive non-participants.

The problem of an unemployed household who is not eligible for UI benefits is:

$$\begin{aligned} v_0^{u_0}(a_0,z_0) = \max_{\{c_t\}_{t\geq 0},\tau^*} \quad & \mathbb{E}_0\Bigg[\int_0^{\tau^{\min}} e^{-\rho t}\mathfrak{u}^u(c_t,h_t)\,dt + \mathbb{I}_{\{\tau^{\min}=\tau^e\}} e^{-\rho\tau^e} v_{\tau^e}^e(a_{\tau^e},z_{\tau^e}) \\ & +\mathbb{I}_{\{\tau^{\min}=\tau^*\}}e^{-\rho\tau^*}v_{\tau^*}^{n_1}(a_{\tau^*},z_{\tau^*}) + \mathbb{I}_{\{\tau^{\min}=\tau^{n_0}\}}e^{-\rho\tau^{n_0}}v_{\tau^{n_0}}^{n_0}(a_{\tau^{n_0}},z_{\tau^{n_0}})\Bigg] \\ s.t. & \\ c_t + \dot{a}_t &= r_t a_t + \phi_t \\ a_t &\geq -\bar{a} \end{aligned} \tag{5}$$

Ineligible unemployed workers receive a job opportunity at rate λ_{zt}^{ue} (with τ^e being the first arrival time of this event) and always take it. At any time τ^* during the unemployment spell, the individual can quit the labor force $\left(\mathfrak{p}_t^u=0\right)$. Finally, at rate η_0 (with τ^{n_0}

[8] The presence of a small switching cost is mostly a technical assumption to avoid "chattering", i.e. infinitely fast switching between n_1 and u_0, in the optimal solution of the problem. For all other participation decisions, this problem does not arise because switching back can only occur upon the realization of Poisson shocks.

being the first arrival rate of this shock) they can become passive non-participants.

The problem of an unemployed household who is eligible for UI benefits is:

$$\begin{aligned}
v_0^{u_1}(a_0,z_0) = \max_{\{c_t\}_{t\geq 0},\tau^*} \quad & \mathbb{E}_0\Bigg[\int_0^{\tau^{\min}} e^{-\rho t}\mathfrak{u}^u(c_t,h_t)\,dt + \mathbb{I}_{\{\tau^{\min}=\tau^e\}}e^{-\rho\tau^e}\max\left\{v^e_{\tau^e}(a_{\tau^e},z_{\tau^e}),v^{u_1}_{\tau^e}(a_{\tau^e},z_{\tau^e})\right\} \\
& +\mathbb{I}_{\{\tau^{\min}=\tau^*\}}e^{-\rho\tau^*}v^{n_1}_{\tau^*}(a_{\tau^*},z_{\tau^*}) + \mathbb{I}_{\{\tau^{\min}=\tau^{u_0}\}}e^{-\rho\tau^{u_1}}v^{u_0}_{\tau^{u_0}}(a_{\tau^{u_0}},z_{\tau^{u_0}}) \\
& +\mathbb{I}_{\{\tau^{\min}=\tau^{n_0}\}}e^{-\rho\tau^{n_0}}v^{n_0}_{\tau^{n_0}}(a_{\tau^{n_0}},z_{\tau^{n_0}})\Bigg] \\
s.t. \quad & \\
c_t+\dot{a}_t \quad = \quad & r_t a_t + (1-\mathfrak{t}_t)b_t(z_t)+\phi_t \qquad (6) \\
a_t \quad \geq \quad & -\bar{a}
\end{aligned}$$

Besides receiving job opportunities and choosing whether to take them, choosing to drop out of the labor force, and exogenously switching to passive non-participant status, the eligible unemployed could lose its entitlement to UI benefit at rate $\eta^{u_1u_0}$, with τ^{u_0} being the first arrival time of this event.

Finally, the problem of the employed household is:

$$\begin{aligned}
v_0^{e}(a,z) = \max_{\{c_t\}_{t\geq 0},\tau^*} \quad & \mathbb{E}_0\Bigg[\int_0^{\tau^{\min}} e^{-\rho t}\mathfrak{u}^e(c_t,h_t)\,dt + \mathbb{I}_{\{\tau^{\min}=\tau^u\}}e^{-\rho\tau^u}v^{u_1}_{\tau^u}(a_{\tau^u},z_{\tau^u}) \qquad (7) \\
& +\mathbb{I}_{\{\tau^{\min}=\tau^*\}}e^{-\rho\tau^*}v^{n_1}_{\tau^*}(a_{\tau^*},z_{\tau^*}) + \mathbb{I}_{\{\tau^{\min}=\tau^{n_0}\}}e^{-\rho\tau^{n_0}}v^{n_0}_{\tau^{n_0}}(a_{\tau^{n_0}},z_{\tau^{n_0}})\Bigg] \\
s.t. \quad & \\
c_t+\dot{a}_t \quad = \quad & r_t a_t + (1-\mathfrak{t}_t)\,w_t z_t h_t+\phi_t \\
a_t \quad \geq \quad & -\bar{a}
\end{aligned}$$

Employed workers (e) can be laid-off at rate λ^{eu}_{zt}, in which case they become eligible for UI benefits ($u=u_1$). Let τ^u be the first arrival time of this Poisson shock. At every instant τ^*, the employed worker can choose to quit the labor force $\left(\mathfrak{p}^e_t=0\right)$.[9] In addition, an employed worker can exogenously switch to passive non-participant status at rate η^{en_0}, with τ^{n_0} being the first arrival time of this event.

Each of these five problems can be expressed recursively as a Hamilton-Jacobi-Bellman Quasi-Variational Inequality (HJBQVI) which can, in turn, be appropriately discretized to numerically solve the household problem. See Appendix A for details on the computa-

[9]Quitting into unemployment is never optimal, because the worker would not receive UI benefits, and would pay a higher disutility cost κ for the opportunity to be re-employed at the same wage.

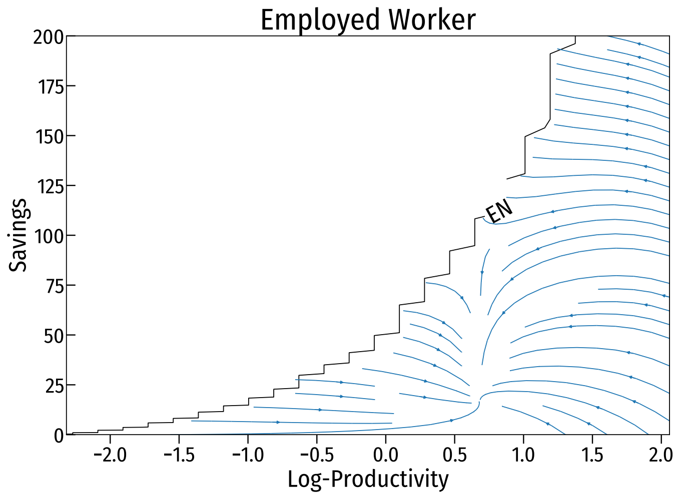

Figure 1: Participation decision for an employed worker represented in the state space of productivity and savings. The white area is the non-participation region. The arrows indicate productivity and wealth dynamics as in a phase diagram.

tion.

Figure 1 plots the participation threshold for an employed worker, as a function of the individual states (a, z). As expected, wealth-poor and highly productive individuals are more likely to remain attached to the labor force.

2.2 Firms

Final-goods producers. A competitive representative final-good producer aggregates a continuum of intermediate inputs indexed by $j \in [0, 1]$ with technology

$$Y_t = \left(\int_0^1 y_{jt}^{\frac{\nu-1}{\nu}} dj \right)^{\frac{\nu}{\nu-1}} \tag{8}$$

where $\nu > 0$ is the elasticity of substitution across inputs. This firm takes prices as given and solves

$$\max_{\{y_{jt}\}} P_t Y_t - \int_0^1 p_{jt} y_{jt} dj \tag{9}$$

subject to (8). Cost minimization implies that demand for intermediate good j at price p_{jt} is

$$y_{jt}(p_{jt}) = \left(\frac{p_{jt}}{P_t} \right)^{-\nu} Y_t, \quad \text{where} \quad P_t = \left(\int_0^1 p_{jt}^{1-\nu} dj \right)^{\frac{1}{1-\nu}} \tag{10}$$

is the price of the final good and the numeraire of the economy.

Intermediate-goods producers. A continuum of measure one of monopolistically competitive firms produce the intermediate goods using labor. Production requires hiring labor on a continuum of tasks indexed by $k \in [0,1]$. Each firm j hires labor services (efficiency-weighted hours) ℓ_{jkt} on every task k, combines them into a final labor input ℓ_{jt} using a Dixit-Stiglitz aggregator with elasticity of substitution ε, and produces the intermediate good according to the linear technology $y_{jt} = \alpha \ell_{jt}$.[10] Every period firms face a fixed operating cost χ expressed in terms of final good. At every date t, these firms take the task-specific wage as given, and maximize profits by solving

$$\begin{aligned} \max_{p_{jt},\{\ell_{jkt}\}_k} \quad & \left(\frac{p_{jt}}{P_t}\right) y_{jt} - \int_0^1 w_{kt}\ell_{jkt}dk - \chi \\ s.t. \quad & \\ y_{jt} &= \alpha \ell_{jt} \\ \ell_{jt} &= \left[\int_0^1 \ell_{jkt}^{\frac{\varepsilon-1}{\varepsilon}} dk\right]^{\frac{\varepsilon}{\varepsilon-1}} \\ y_{jt} &= \left(\frac{p_{jt}}{P_t}\right)^{-\nu} Y_t \end{aligned} \tag{11}$$

where w_{kt} is the real wage on task k. Cost minimization yields the relative demand of labor for task k

$$\ell_{jkt} = \left(\frac{w_{kt}}{w_t}\right)^{-\varepsilon} \ell_{jt}, \tag{12}$$

where w_t is the Dixit-Stiglitz real aggregate wage index $w_t = \left[\int_0^1 w_{kt}^{1-\varepsilon} dk\right]^{\frac{1}{1-\varepsilon}}$ that satisfies $\int_0^1 w_{kt}\ell_{jkt}dk = w_t \ell_{jt}$. The profit-maximizing price setting decision yields the standard expression whereby the relative price equals a markup over the marginal cost of production

$$\frac{p_{jt}}{P_t} = \frac{\nu}{\nu - 1}\left(\frac{w_t}{\alpha}\right) \tag{13}$$

This price corresponds an optimal quantity sold and, in turn, to a total amount of labor demanded to produce it which, in the symmetric equilibrium with $p_{jt} = P_t$, satisfies $\ell_{jt} =$

[10] An alternative structure, which gives rise to the same allocations, would be to assume that there is a competitive labor intermediary that hires task-specific labor services for all tasks $k \in [0,1]$ and packages into a CES aggregate labor input sold to firms.

ℓ_t with

$$\ell_t = \left(\frac{1}{\alpha}\right) Y_t.$$

Imposing $p_{jt} = P_t$ in (13) implies that the equilibrium aggregate real wage per effective hour is constant over time.[11] Thus, price inflation equals wage inflation. Finally, the real aggregate profits of the production sector are

$$\Pi_t = Y_t - w_t \ell_t - \chi. \tag{14}$$

Every period, profits are paid as dividends to the mutual fund that owns all intermediate producers.

2.3 Wage Setting

This block of the model adapts the wage setting mechanism of Erceg et al. (2000) —i.e., the standard New Keynesian *sticky wage* model— to an heterogeneous-agent economy.[12]

Every worker i at date t supplies hours on each task k. The *nominal* wage ω_{kt} per effective hour worked on task k is set by a union that represents all workers on that particular task. By adhering to the union, each employed worker agrees to supply, at that wage, the same number of hours h_{kt} to producers. The problem of each union is:

$$\begin{aligned} \max_{\{\omega_{kt}\}_{t\geq 0}} & \quad \int_0^\infty e^{-\rho t} \left[\int_{s_{it}=e} \mathrm{u}^e\left(c_{it}, h_{it}\right) di - \frac{\Theta}{2} \left(\frac{\dot{\omega}_{kt}}{\omega_{kt}} - \pi^* \right)^2 \right] dt \\ s.t. & \\ h_{it} &= \int_0^1 h_{kt} dk \\ c_{it} + \dot{a}_{it} &= r_t a_{it} + (1 - \mathrm{t}_t) \frac{1}{P_t} z_{it} \int_0^1 \omega_{kt} h_{kt} dk + \phi_t \\ h_{kt} \int_{s_{it}=e} z_{it} di &= \left(\frac{\omega_{kt}}{\omega_t} \right)^{-\varepsilon} \ell_t \end{aligned} \tag{15}$$

At every date t, the union sets the nominal wage ω_{kt} in order to maximize welfare of its current members (all individuals employed at date t) subject to a Rotemberg-style quadratic costs of adjusting the nominal wage, in utility terms, with scaling parameter Θ. Let inflation be denoted by $\pi_t = \dot{P}_t / P_t$. This cost is expressed in terms of deviations of

[11] This property descends from constant returns to scale in production.

[12] We follow closely the approach of Auclert et al. (2018, 2020), with the necessary modifications due to our continuous time formulation and the presence of the extensive margin in labor supply.

nominal wage growth from the central bank's inflation target, the deterministic steady-state trend inflation rate π^*.[13] The first constraint faced by the union states that total hours worked by an employed worker equal the sum of hour worked on each task. The second constraint is the budget constraint of employed workers. The third one states that contractual hours worked required by the union from all workers must satisfy firm's demand for effective labor on task k, ℓ_k.[14]

Because each task-specific union is "small" (there is a continuum of tasks) the impact of a union's wage on individual income and firm's employment is negligible. As a result, the union takes as given all individual decisions embedded in the budget constraint, and the firm's labor demand curves for each task. In a symmetric equilibrium $\omega_{kt} = \omega_t$ and $h_{kt} = h_t$. Since $\ell_t = h_t \int_{s_{it}=e} z_{it} di$, once we know the set of individuals who are employed at t, we can residually derive hours of work h_t that each employed individual is contractually required to supply.[15]

In Appendix B we show that the solution to this problem yields the wage Phillips curve

$$\rho(\pi_t - \pi^*) = \frac{\epsilon}{\Theta} H_t \left[\psi h_t^{\frac{1}{\sigma}} - \left(\frac{\epsilon - 1}{\epsilon} \right) (1 - \mathfrak{t}) w_t Z_t^e \int_{s_{it}=e} c_{it}^{-\gamma} \left(\frac{z_{it}}{\int_{s_{it}=e} z_{it} di} \right) di \right] + \dot{\pi}_t \tag{16}$$

where π_t is aggregate (wage and price) inflation rate, H_t aggregate hours worked, and Z_t^e average productivity of the employed. The square bracket of equation (16) indicates that whenever the marginal disutility of an extra hour exceeds the marginal utility of an extra unit of after-tax real wage income averaged across all employed workers and weighted by their productivity, the union will push up nominal wages to reduce labor demand and close the gap between these two margins. Note the role of average productivity Z_t^e. Ceteris paribus, an increase in labor productivity reduces wage inflation pressures.

[13] Our interpretation of this adjustment cost technology is therefore that wage setters can freely index nominal wage growth to π^*. They understand that inflation in the long-run always converges to that value, and they can take advantage of this information in making their wage setting plans. It is only costly for them to set a value for nominal wage growth which deviates from it.

[14] Note that the right hand side of this latter constraint equals (12).

[15] Huo and Ríos-Rull (2020) raise a valid criticism to the RANK model featuring nominal wage rigidity. In that model, along the equilibrium path workers may end up being forced to supply hours against their will, violating the principle of voluntary exchange. They propose a resolution based on a different equilibrium concept. Here, we have a different solution: in our model, unions propose all workers an employment contract that specifies a non-negotiable pair of wages and hours, but workers can always voluntarily choose not to participate in it and, in fact, in equilibrium some do and quit employment.

2.4 Mutual Fund

A competitive risk-neutral mutual fund owns all intermediate good firms and holds all debt issued by the government.[16] Let X_t^m denote the shares of the intermediate good producers held by the mutual fund, q_t the unit share price, Π_t per-share dividends (or profits), B_t^m the amount of government bonds held by the fund, and r_t^b the real interest rate on government bonds. Let A_t be the value of the fund. The problem of the mutual fund, which takes prices as given, entails choosing the optimal portfolio composition between bonds and equity:

$$\begin{aligned} r_t A_t\left(X^m, B^m\right) = \max_{\dot{X}_t^m, \dot{B}_t^m} \quad & \Pi_t X_t^m - q_t \dot{X}_t^m + r_t^b B_t^m - \dot{B}_t^m \\ & + \partial_X A_t\left(X^m, B^m\right)\dot{X}^m + \partial_B A_t\left(X^m, B^m\right)\dot{B}_t^m + \partial_t A_t\left(X^m, B^m\right) \end{aligned} \tag{17}$$

with first-order conditions with respect to $\dot{X}_t^m$ and $\dot{B}_t^m$

$$\begin{aligned} q_t &= \partial_X A_t\left(X^m, B^m\right) \\ 1 &= \partial_B A_t\left(X^m, B^m\right) \end{aligned}$$

Substituting these first order conditions into (17) and exploiting the linear homogeneity of the problem which implies that $A_t = q_t X_t^m + B_t^m$, we arrive at

$$r_t\left(q_t X_t^m + B_t^m\right) = \Pi_t X_t^m + r_t^b B_t^m + \dot{q}_t X^m.$$

By matching coefficients on equity and bonds, we obtain

$$r_t = \frac{\Pi_t + \dot{q}_t}{q_t} = r_t^b \tag{18}$$

which determines the real return on the household financial asset a_t (wealth invested in the mutual fund), and establishes a no-arbitrage condition between government bonds and firm equity which holds at every t, except when a shock hits the economy, in which case the price q_t features a jump.[17]

[16] The set up in this section follows closely Alves et al. (2020).

[17] Note that, because of the absence of trading frictions, the mutual fund is willing to absorb any amount of each asset.

2.5 Fiscal Authority

Let G_t be the units of the final goods purchased by the government (fiscal authority), ϕ_t lump-sum transfers, b_t UI benefits, $\mathfrak{t}_t$ the labor income tax and $B_t^g > 0$ outstanding real government debt. The government faces the following intertemporal budget constraint:

$$G_t + \phi_t + (1 - \mathfrak{t}_t) \int_{s_{it}=u^1} b(z_{it})di + r_t^b B_t^g = \mathfrak{t}_t w_t h_t \int_{s_{it}=e} z_{it}di + \dot{B}_t^g \tag{19}$$

Outside of steady-state, we assume that the government follows the passive fiscal policy rule:

$$G_t = G^* - \beta_B(B_t^g - B^*), \quad \beta_B > 0 \tag{20}$$

where the superscript * denotes steady-state values. Thus, following an aggregate shock debt adjusts to satisfy the government budget constraint, and government expenditures respond to deviations of debt from its steady-state level to keep debt from growing too quickly.

2.6 Monetary Authority

In our baseline, the monetary authority sets the nominal interest rate i_t according to an Inflation Targeting (IT) rule that reacts to deviations of inflation and unemployment rate from their targets with some inertia. If we let $\imath_t$ denote the shadow policy instrument not subject to the ZLB, then the IT rule is defined as

$$\frac{d\imath_t}{dt} = -\beta_i\Big(i_t - i^* - \beta_\pi(\pi_t - \pi^*) - \beta_u(u_t - u^*)\Big) \tag{21}$$

$$i_t = \max\{0, \imath_t\} \tag{22}$$

where we let i^* denote the steady-state nominal rate, $\pi_t = \dot{P}_t/P_t$ the aggregate inflation rate at date t, and u_t is aggregate unemployment rate at date t. The coefficients $\beta_\pi > 1$ and $\beta_u \leq 0$ capture the strength of the policy response to deviations of inflation from target π^* and of unemployment from its steady-state value u^*. The coefficient β_i captures the degree of interest rate smoothing. The monetary authority is constrained by a zero lower bound (ZLB) on nominal rates which forces realized rates i_t weakly above zero at all times.

The nominal interest rate and the real interest rate on government bonds r_t^b are linked through the Fisher equation $r_t^b = i_t - \pi_t$.

2.7 Equilibrium

An equilibrium for this economy is defined as time paths for household consumption decisions $\{c_t^s(a,z)\}_{t\geq 0}$ for $s \in \{e, u_0, u_1, n_0, n_1\}$, participation and job offer acceptance decisions $\{\mathfrak{p}_t^s(a,z), \mathfrak{f}_t^s(a,z)\}_{t\geq 0}$ for all s, unions' nominal wage setting $\{\omega_{kt}\}_{t\geq 0}$ for all labor types k, intermediate producers' hiring decisions $\{\ell_{kt}\}_{t\geq 0}$ for all k, mutual fund allocations between equity and government bonds $\{X_t^m, B_t^m\}_{t\geq 0}$, real rates of return on the mutual fund and on government bonds $\{r_t, r_t^b\}_{t\geq 0}$, firms' share price $\{q_t\}_{t\geq 0}$, fiscal variables (taxes, transfers, UI benefits, expenditures and debt) $\{\mathfrak{t}_t, \phi_t, b_t(z), G_t, B_t^g\}_{t\geq 0}$, nominal interest rates $\{i_t\}_{t\geq 0}$, aggregate output, consumption, profits, contractual hours worked, and inflation $\{Y_t, C_t, \Pi_t, h_t, \pi_t\}_{t\geq 0}$, and measures of households $\{\mu_t^s(a,z)\}_{t\geq 0}$ for all s such that at every t: (i) households solve problems (3)-(7); (ii) final good and intermediate good producers solve (9) and (11), respectively; (iii) unions solve (15) and inflation satisfies the Phillips curve in (16) (iv) the mutual fund solves (17); (v) the government budget constraint (19) holds; (vi) the fiscal and monetary authorities follow their policy rules (20) and (22); (vii) the sequence of distributions satisfies aggregate consistency conditions, and (viii) all good and asset markets clear.[18]

Besides the continuum of intermediate goods' markets with equilibrium condition $y_{jt} = Y_t$, there are five other markets in our economy: the mutual fund shares market, the intermediate firms' shares market, the government bond markets, the final good market, and the labor market. The first three markets clear when, respectively

$$\begin{aligned} X_t^m &= 1 \\ A_t := \sum_{s\in\{e,u,n\}} \int a_t d\mu_t^s &= q_t + B_t^g \\ B_t^m &= B_t^g \end{aligned}$$

where, without loss of generality, we normalized the measure of firms' shares to 1. These market clearing conditions, together with the no-arbitrage condition (18) and the definition of firm profits (14), determine firm share prices, real interest rates, and aggregate profits. The final goods market clears when

$$Y_t = C_t + G_t + \chi$$

The labor market is frictional and at any t some workers are 'involuntarily' unemployed.

[18] We report the KFE for the distributions in Appendix C.

Steady State Parameters

Parameter		Value	Target
Preferences			
Risk aversion	γ	1.00	—
Labor supply elasticity	σ	1.00	—
Credit limit	$\bar{a}$	0.00	—
Discount rate	ρ	0.0087	Liquid wealth to annual earnings (0.56)
Utility weight on hours	ψ	0.85	No wage inflationary pressures at SS
Disutility of working	κ^e	0.90	Average labor market flows
Disutility of searching	κ^u	0.0527	
Disutility of nonparticipation	κ^n	0	—
Demographics			
Death rate	θ	$1/312$	Average worker lifespan of 36 years
Variance of initial skill distribution	σ^2_{0z}	0.25	P90-P50 hourly wage ratio for 23-27 age group
Productivity process			
Skill drift while employed	δ^+	0.0024	Average log earnings growth from 25 to 55
Skill drift while non-employed	δ^-	-0.0171	10-Year earnings losses from displacement
Skill mean reversion	ρ_z	0.0017	
Skill diffusion	σ_z	0.0465	P90-P50 hourly wage ratio for all workers
Labor market frictions			
Job-separation rate out of E		–	Average labor market flows
Job-finding rate out of U		–	Average labor market flows
Job-finding rate out of N		–	Average labor market flows
Passive nonparticipation rate during E	η^{en_0}	0.007	$e \to n$ for high-z workers
Passive nonparticipation rate during U/N	$\eta^{un_0}, \eta^{n_1 n_0}$	0.070	$u \to n$ for high-z workers
Passive nonparticipation exit rate	$\eta^{n_0 n_1}$	0.130	Average labor market flows
Taxes and transfers			
UI replacement rate	$\bar{b}$	0.50	—
UI expiration rate	$\eta^{u_1 u_0}$	0.167	Average duration of UI (6 months)
Lump-sum transfer	ϕ	0.055	6% of annual average earnings
Labor tax rate	$\mathfrak{t}$	0.2	—
Technology and Price/Wage Setting			
Firm productivity	α	1.38	—
Firm fixed cost	χ	0.12	Steady-state real rate of 2%
Price/Wage markups	ν, ε	10	—
Slope of the wage Phillips curve (quarterly)	–	0.015	—
Fiscal and Monetary Policy			
Trend inflation	π^*	2%	—
Taylor rule persistence	β_i	0.07	—
Taylor rule reaction to inflation	β_π	2.25	—
Taylor rule reaction to unemployment rate	β_u	-0.15	—
Government expenditures response to debt	β_B	0.10	—

Table 2: Parameter values needed to determine steady state. The corresponding targeted moments are listed in the main text. The model period is one month.

A stationary equilibrium is a particular case of our definition where –absent aggregate shocks– all decisions, prices, aggregate variables, and distributions are constant over time, i.e., nothing in that definition is indexed by t.

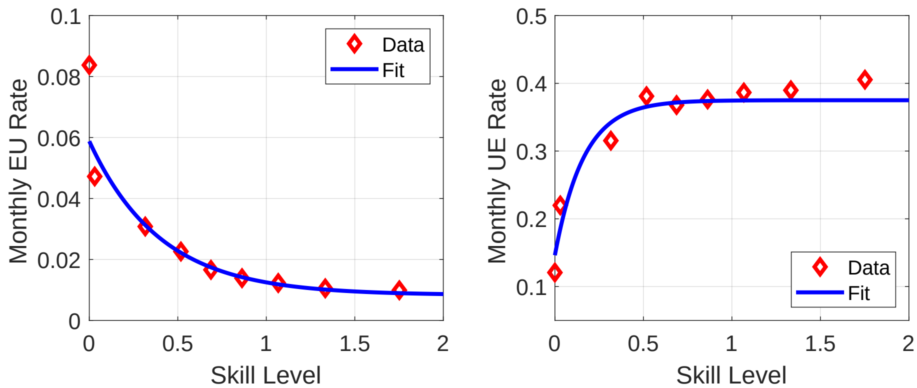

Figure 2: Empirical transition rates as a function of skill levels (weekly earnings) measured from the Current Population Survey (CPS) and model fit. See Appendix D for details.

3 Parameterization

Preferences. We set $\gamma = 1$ (log-utility over consumption expenditures) and $\sigma = 1$ (quadratic disutility of hours worked). The discount rate ρ is set to target a ratio of mean wealth to annual earnings of 0.56, corresponding to the amount of liquid wealth immediately available for consumption smoothing among US households (Kaplan and Violante, 2022). This choice allows the model to match a sizable aggregate marginal propensity to consume of 0.10 without adding illiquid assets or preference heterogeneity. The credit limit $\bar{a}$ is set to zero, so workers cannot hold short positions on the mutual fund.

Working entails a variable and a fixed cost. The variable disutility parameter ψ is set so that in steady state there is no inflationary pressure, beyond trend inflation. The fixed disutility of work κ^e is set to match the average flows as discussed below. The disutility cost of searching κ^u is set to match the observation that job-seekers spend less than 30 minutes per day searching (Faberman et al., 2017). The flow utility of non-participation κ^n is normalized to zero.[19]

Demographics. We set the monthly mortality rate θ so that workers are on an average active for 36 years (ages 25 to 60). The initial skill dispersion σ^2_{0z} is set to match the P90-P50 hourly wage ratio of 2.00 for workers aged 23-27 in the 2019 CPS (Heathcote et al.,

[19]The switching cost ξ is set to a very small number to make the optimal stopping problem well behaved.

2023). The average initial skill level $\bar{z}_0$ serves only as a normalization.

Labor market frictions. Three labor market transition rates depend on worker's skill level z: the separation rate λ_z^{eu}, and the job finding rates for unemployed λ_z^{ue} and the job-finding rate for active non-participants λ_z^{ne}. In the steady state, we model the dependence of each one of these labor market frictions on worker's skill level z as

$$\lambda^{ss'}(z) = \lambda_0^{ss'} + \lambda_1^{ss'} \exp(\lambda_2^{ss'} z). \tag{23}$$

We choose the coefficients in (23) in two steps. In the first step, we use data on transition rates across the workers' wage distribution to get an estimate of $\lambda_0^{ss'}, \lambda_1^{ss'}, \lambda_2^{ss'}$ for *eu*, *ue* and *ne*.[20] Figure 2 plots the measured *eu* and *ue* rates in the data as a function of our measure of skills together with the fit that comes out from the first stage. The two flows show a clear dependence on the skill level: the separation rate *eu* is four-five times larger, and job finding rates *ue* three-four times smaller, for workers at the bottom of the distribution compared to a worker with average skills. The second step takes place at the calibration stage. At that stage, we keep the "shape" from the first stage fixed, but allow the coefficients to be rescaled to target average worker flows EU, UE and NE estimated from the CPS.

The exogenous transition rates to and from the passive nonparticipant state don't depend on skill z and are set as follows. Transitions η^{en_0} and η^{un_0} out of participation are set to match the measured $e, u \rightarrow n$ flows at the very top of the skill distribution. In the data, outflows from participation are decreasing in our measure of skill, but are still positive even for worker in the top deciles. In the model, these transitions occur primarily through the forced transitions to passive nonparticipant state, as strong substitution effects among high-skill workers generally implies that these workers would prefer to remain in the labor force. The exit rate of passive nonparticipation $\eta^{n_0 n_1}$ adjusts to match the flows out of participation.

Table 3 shows that the model can reproduce all average flows well. In particular, because the *un* flow is an order of magnitude larger than the *eu* flow, the model can correctly replicate the participation cycle.

[20] We don't use *ne* transitions directly in our estimation because the job acceptance decisions from nonparticipants creates an wedge between job-finding rates out of non-participation $\lambda^{ne}(z)$ (our object of interest) and the observed *ne* flows (our empirical measure) . Instead, we impose that the job-finding rate out of nonparticipation shares the same shape as the job-finding rate out of unemployment. Appendix D provides more details on the measurement and estimation of coefficients.

Productivity dynamics. The mean reversion parameter ρ_z is set to -0.0017, corresponding to an annual autocorrelation of $\exp(-12 \times \rho_z) = 0.98$. The negative drift δ^- is set to match the evidence on earnings losses upon displacement from Davis and Von Wachter (2011). Specifically, we target the estimate that laid-off workers still earn on average 15% less than their control group 10 years after separation. We set the positive drift δ^+ to match the average worker log earnings growth between ages 25 and 55 of 0.70 log points, consistent with US data download from the Global Repository of Income Dynamics (GRID).[21] Finally, we choose the standard deviation σ_z to match a 90-50 wage ratio of 3, the value for the 2019 CPS (Heathcote et al., 2023).[22]

Taxes and transfers. We assume that unemployment benefits are given by $b(z_{it}) = \bar{b}\ w_t h_t z_{it}$, and set the UI replacement rate to 0.5 of individual earnings. We set rate $\eta^{u_1 u_0}$ to 0.167 to reflect an average UI benefits duration of 6 months. The proportional tax rate t is set to 0.2 and the lump-sum transfer ϕ is set to match 6% of average earnings in steady-state.[23] The amount of government debt is set to equal 1/4 of total equity (2019 Flow of Funds, Table B.101.h Balance Sheet of Households). Government expenditures are set residually to satisfy the budget constraint in steady state.

Technology and Phillips curve. Firm productivity α is set so that the net hourly wage per efficiency unit in steady state is normalized to 1. The fixed operating cost χ affects the value of equity and, therefore, the size of the aggregate supply of liquid wealth. We set χ so that, given the household demand curve, the annual real interest rate that clears the asset market is 2%.

Both elasticities of substitution across labor types (ε) and across intermediate goods (ν) are set to 10 which implies wage and price markups around 10 percent. The nominal wage adjustment cost Θ is set to match a slope of the structural wage Phillips curve (the semi-elasticity of inflation to deviations of marginal rate of substitution from the real

[21] The data is available through the website https://www.grid-database.org/. See Guvenen et al. (2022) for a description of the database.

[22] We target the 90-50 ratio because earnings variation at the top of the distribution is more directly associated with productivity variation, which is what we aim to measure, compared to the rest of the distribution where the extensive margin of labor supply plays a bigger role.

[23] This number is obtained by dividing Government Social Benefits by Wages and Salaries. Transfers are computed as: Workers' compensation, SNAP, Supplemental security income, Refundable tax credits, Temporary disability insurance, Workers' compensation, Family assistance, General assistance, Energy assistance, Employment and training, Other benefits, and 0.4*Medicaid (Table 3.12 of NIPA). Wages and salaries are taken from Table 2.1 of NIPA for 2019. The share of Medicaid expenditures that are effective transfers to households (0.4) is obtained from Finkelstein et al. (2019).

	Data			Model		
	$\mathbb{E}(x)$	$std(x)$	$cor(x, Y)$	$\mathbb{E}(x)$	$std(x)$	$cor(x, Y)$
Unemployment rate	0.0552	12.6632	-0.8655	0.0546	12.6600	-0.8258
Labor Force Participation	0.8083	0.3793	0.1462	0.7668	0.4585	0.8857
(-) EU	0.0127	8.5263	-0.7713	0.0132	7.5188	-0.8175
(+) UE	0.2479	8.5669	0.7697	0.3020	7.6027	0.7789
(+) NE	0.0696	3.8214	0.4351	0.0253	6.3583	0.7711
(+) EN	0.0170	3.9220	0.3120	0.0110	1.8292	0.3367
(+) UN	0.1326	8.6439	0.6660	0.0915	2.9095	0.7167
(-) NU	0.0267	8.3435	-0.6507	0.0239	6.8588	-0.6954

Table 3: Cyclical properties of labor market and stocks and gross worker flows (monthly rates) in the data and in the calibrated model.

wage) of 0.015 quarterly as recently estimated by Del Negro et al. (2020).[24]

Monetary and fiscal policy. We set steady-state (trend) inflation rate π^* at 2%. In our baseline Inflation Targeting (IT) rule (22), we set the interest rate smoothing to $\beta_i = 0.07$, the reaction coefficient on deviations of inflation from its trend to $\beta_\pi = 2.25$, and the coefficient on the unemployment gap to $\beta_u = -0.15$. In the fiscal rule (20), we set $\beta_B = 0.1$.

3.1 Calibration of Aggregate Shocks

Aggregate shocks. We posit that, outside of steady-state, our economy is subject to two aggregate shocks, which we label (with a slight abuse of language) demand and supply shocks. The demand shock ζ^d shows up as a wedge between the rate of return on equity and the return on government bonds

$$\frac{\Pi_t + \dot{q}_t}{q_t} = r_t^b + \zeta^d \tag{24}$$

[24] In log-linearized models, there exists a mapping between this parameter and the Poisson adjustment rate parameter of a model with Calvo-style nominal rigidities. Our value of Θ would correspond to a monthly Calvo parameter of 10%, which is somewhat above the average monthly frequency of wage changes of 5% estimated by Grigsby et al. (2021). A different strategy for the calibration of Θ which replicates the frequency of wage adjustment would therefore lead to a flatter curve. We prefer to err on the side of a steeper Phillips curve because in our model wage inflation equals price inflation and price Phillips curves are estimated to be steeper that wage ones, in general. In fact, the monthly frequency of price adjustment is higher than that of wage adjustment, i.e. around 15% (Nakamura and Steinsson, 2008). Our 10% implied frequency of adjustment is therefore in between wage and price adjustment frequencies.

which can be interpreted as a 'risk-premium shock'. Workers experience the shock through changes to the rate of return on wealth, as the price of the fund adjusts to the shock. The supply shock ζ^s appears as a wedge in the wage Phillips curve (16), which can be interpreted as a a cost-push shock caused by time-varying wage mark-up. Both shocks are common in the representative-agent New Keynesian literature and contribute significantly to the overall fluctuations of the US economy (Smets and Wouters, 2007). We assume that the shocks follow Ornstein-Uhlenbeck diffusion processes

$$d\zeta_t^k = -\rho_k \zeta_t^k dt + \sigma_k dW_t, \quad k = d, s$$

with annualized persistence of 0.90 for both shocks. In our simulations, we choose the volatility of the two shocks to match the standard deviation of the unemployment rate while being consistent with the fact that 50% of unemployment rate variations at business cycle frequencies (6 and 32 quarters) are driven by demand shocks.

Cyclicality of frictions. The behavior of labor market frictions (i.e., separation and job finding rates) out of steady state is modeled in a simple way. Specifically, we make the job-finding and separation rates functions $\lambda_t^{eu}(z)$, $\lambda_t^{ue}(z)$ and $\lambda_t^{ne}(z)$ vary in proportion to changes in the average hours per worker h_t. This approach allows us to capture the main facts of how job finding and separation rates fluctuate over the business cycle without further complicating the model. We discipline the elasticities of frictions to average hours by targeting (i) the contribution of extensive and intensive margins to aggregate hours fluctuations (the ratio of the covariance of log average hours worked with log of total hours to the variance of log of total hours worked is 1/3), and (ii) the relative volatility of EU and UE transitions (the ratio of the variance of log *ue* rate to the variance of log *eu* rate is approximately 1).[25]

Fluctuations in stocks and gross worker flows. Table 3 reports the key business cycle statistics of the labor market stocks and gross worker flows in the U.S. economy and in our model simulation. Our calibration strategy targets the volatility of unemployment, EU and UE flows, so it is no surprise that their values are close to our empirical estimates. More interesting is that the model also closely matches the volatility of the labor force participation rate, which displays relatively modest fluctuations compared to fluctuations in

[25] In keeping with the calibration strategy for the steady state, we set the elasticity of the *ne* rate with respect to hours equal to the elasticity for the *ue* rate.

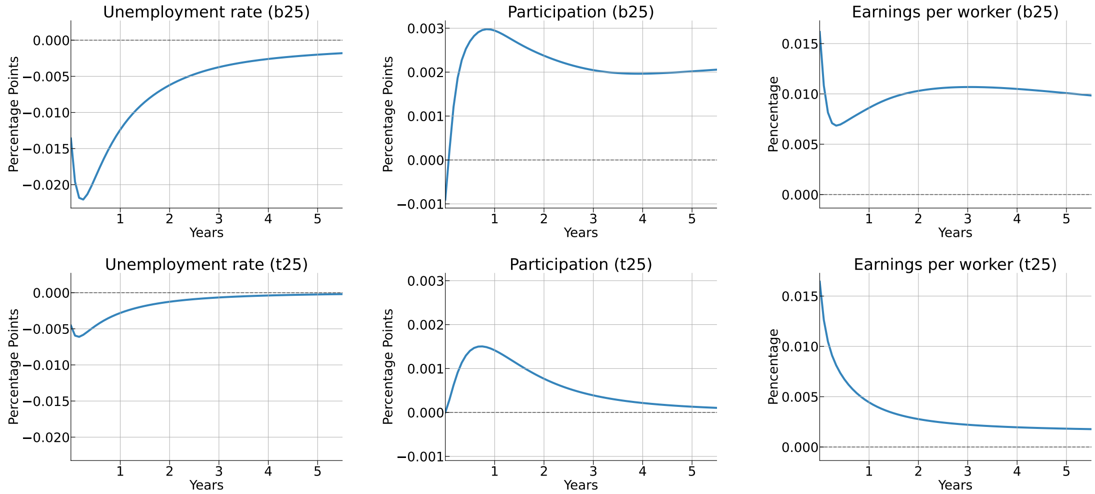

Figure 3: Model's IRFs of unemployment, participation, and earnings per worker in the bottom and top 25% of the skill distribution to a demand shock.

unemployment or employment rates. The calibration also does well in terms of the dynamics of labor market flows, capturing the right cyclicality of all six flows. Importantly for the dynamics of participation, the model matches the cyclicality of EN and UN flows, which act to offset the procyclical pressures in participation stemming from the strong movements in EU and UE flows, i.e., the "participation cycle" discussed in Hobijn and Şahin (2021).

3.2 Okun's Hypothesis Through the Lens of the Model

How do the three mechanisms of Okun's Hypothesis shape the labor market outcomes of workers at different points of the skill distribution in response to aggregate fluctuations? Figure 3 answers this question. The two rows feature the response of unemployment rate, participation, and earnings per worker to a positive demand shock for the bottom and top 25% of the skill distribution. The left panels show the exposure channel: workers at the bottom of the skill distribution are disproportionately sensitive to the increase in labor demand: their unemployment rate falls three times as much compared to the top quartile. The central panels show the attachment channel: because of the larger drop in unemployment, participation rates rise across the distribution, especially at the bottom, where marginally attached workers are concentrated. The right panels show the persistence channel. Earnings per worker are determined by hours worked and skills (recall that the real wage is constant). Human capital takes time to adjust, so the initial rise in labor earnings is mostly due to higher hours worked, which are common among workers.

Baseline Inflation Target	$\frac{d\iota_t}{dt} = -\beta_i \Big(\iota_t - i^* - \beta_\pi (\pi_t - \pi^*) - \beta_u (u_t - u^*) \Big)$
Average Inflation Target	$\frac{d\iota_t}{dt} = -\beta_i \Big(\iota_t - i^* - \beta_\pi (\pi_t - \pi^*) - \beta_{AIT} (\pi_t^{MA} - \pi^*)^- - \beta_u (u_t - u^*) \Big)$
Shortfall Rule	$\frac{d\iota_t}{dt} = -\beta_i \Big(\iota_t - i^* - \beta_\pi (\pi_t - \pi^*) - \beta_u^+ (u_t - u^*)^+ \Big)$
Lower for Longer Rule	$\frac{d\iota_t}{dt} = -\beta_i \Big(\iota_t - i^* - \beta_\pi (\pi_t - \pi^*) - \beta_{AIT} (\pi_t^{MA} - \pi^*)^- - \beta_u^+ (u_t - u^*)^+ \Big)$
and for all rules	$\text{i}_t = \max\{\iota_t, 0\}$

Table 4: Monetary policy rules used in the counterfactual experiments. π_t^{MA} is the exponential moving average of past inflation $\int_0^\infty (1/48)\ e^{-(1/48)\tau} (\pi_{t-\tau} - \pi^*) d\tau$ with a smoothing factor $1/48$ to target a window of 4 years. X^+ is the shorthand for $\max\{X, 0\}$, while X^- stands for $\min\{X, 0\}$. Our parameterization of the rules is as follows. Baseline Inflation Target: $\beta_i = 0.07$, $\beta_\pi = 2.25$ and $\beta_u = -0.15$. Asymmetric Average Inflation Target: $\beta_i = 0.07$, $\beta_\pi = 2.25$, $\beta_{AIT} = 5.00$, and $\beta_u = -0.15$. Shortfall: $\beta_i = 0.07$, $\beta_\pi = 2.25$ and $\beta_u^+ = -0.15$. Lower for Longer: $\beta_i = 0.07$, $\beta_\pi = 2.25$, $\beta_{AIT} = 5.00$, and $\beta_u^+ = -0.15$.

Over time, the larger job-finding rates and smaller separation rates lengthen the employment spells of workers, leading them to accumulate more skills on the job and to avoid skill deterioration due to nonemployment. This force is extremely persistent at the bottom of the distribution because higher skills help sustaining a higher labor force participation for longer, as seen from the central panels. These channels reinforce one another with human capital growth begetting participation, and viceversa.

The next step in our analysis is asking whether, in an economy where the labor market operates under Okun's hypothesis, a lower for longer strategy can be more effective in dealing with the ZLB and be more inclusive with respect to low-wage workers compared to strict IT. All this without, at the same time, generating systematic inflation in excess of the target. This was the intent of the new framework and we'll assess if, through the lens of the model, this strategy can be successful.

4 Results

We divide this section in four parts. In the first one, we focus on short-run dynamics and ask how the fall into and the recovery from the Great Recession would have looked like had the Fed already been operating under its new framework. This counterfactual exercise is helpful in comparing how the different rules operate to modify the aggregate

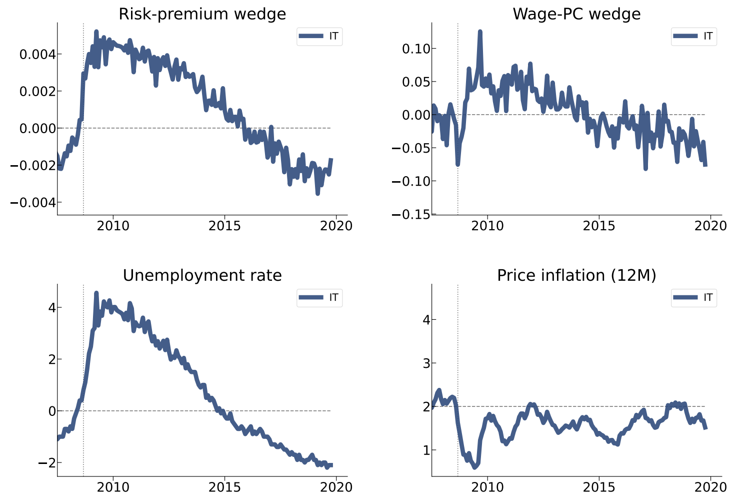

Figure 4: Bottom panels: Unemployment rate (Fred Series: UNRATE) and Inflation (12months MA, Fred Series: CPILFESL). Top panels: estimated and demand and supply shocks.

and distributional dynamics of the economy over a single business cycle.

In the second part, we simulate our economy under the new framework for a long period of time, as it is repeatedly hit by aggregate demand and supply shocks. We then use this ergodic distribution to quantify the inflation-inclusion trade-off generated by pursuing a lower for longer strategy. As clear from the illustration of Okun's channels, following such policy rule over a cycle can have persistent effects on participation and productivity that spillover into the next cycle, and influence the effectiveness of the rule. This mechanism is missing from the short-run analysis.

In the third part, we show that by varying the size of the coefficients on the LfL rule, one can obtain a menu for the policymaker. This menu describes the long-run inflation-inclusion trade-off implicit in the ergodic distributions corresponding to each parameter configuration of the rule.

In the last part, we argue that, based on our findings, one can design more effective rules that lead to a better inflation-inclusion trade-off, and we study such rules.

Alternative monetary policy rules Table 4 lists the set of policy rules we consider in our counterfactuals. The Inflation Targeting (IT) rule, described in Section 3, serves as our baseline representation of the Fed reaction function prior to the change in the policy framework. The Average Inflation Targeting (AIT) rule is meant to capture the first component of the new monetary policy framework of the Federal Reserve, whereby it aims to achieve an inflation that averages 2 percent over time instead of targeting current

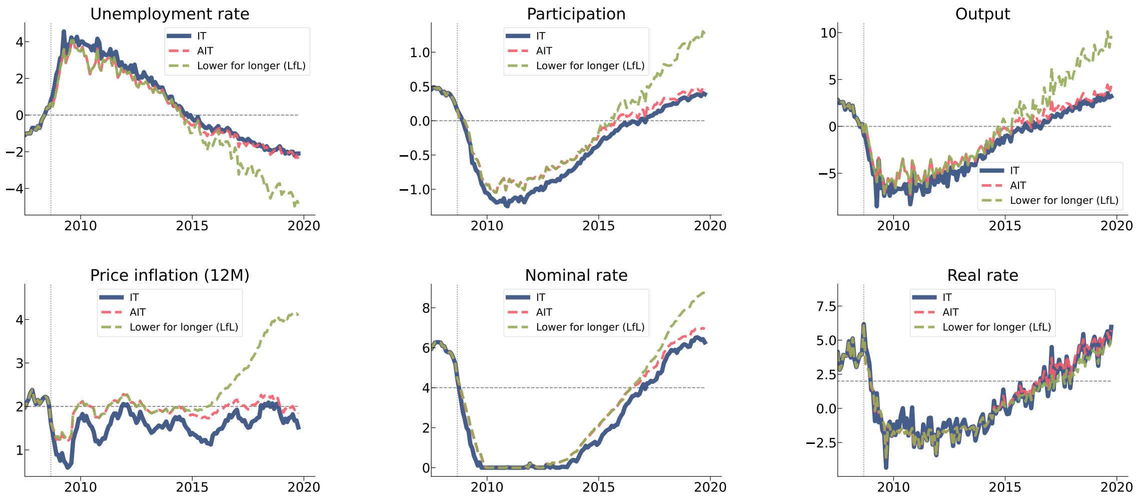

Figure 5: Model simulation of the economy under the baseline Inflation Target, asymmetric AIT, and Lower for Longer rule. See Table 4 for the specification of the different monetary policy rules.

inflation only. To reflect this shift in focus, we include a measure of average of past inflation π^{MA} to the set of variables that the Fed systematically reacts to. As long as average past inflation falls short of the 2 percent target, the rule prescribes that rates should be kept lower than under IT to make up for past downward (but not upward) deviations of inflation from the target.

The Shortfall rule (SR) introduces a different asymmetry to the monetary policy. We substitute the reaction to unemployment fluctuations in the baseline IT for a rule that reacts only to positive unemployment gaps, reflecting the reinterpretation that rates should be informed by shortfalls instead of deviations of employment from its maximum level. Accordingly, the Shortfall rule still promises to aggressively lower rates whenever the economy contracts and the unemployment rate increases. Relative to the baseline, however, the new rule is more tolerant during periods of low unemployment as long as they are not accompanied by higher inflation. This captures the idea that the Fed might want to let the economy run hot following recessions as a way to benefit workers at the bottom of the distribution.

Finally, the Lower for longer (LfL) rule combines both aspects of the new policy framework, incorporating the idea that, following recessions, the Fed might want to keep rates lower for longer to run the economy hot and make-up for past shortfalls in inflation.

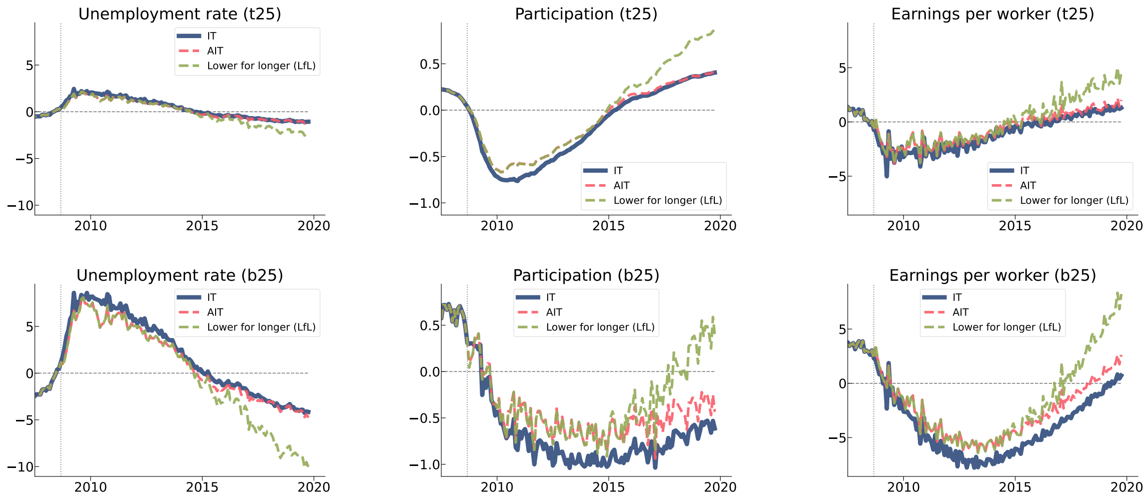

Figure 6: Model simulation of the economy under the baseline Inflation Target, asymmetric AIT, and Lower for Longer rule. See Table 4 for the specification of the different monetary policy rules.

4.1 Short-run

In our first counterfactual exercise, we ask how would the US labor market and inflation dynamics have looked like, around the Great Recession, had the Fed followed a LfL strategy.

To generate this counterfactual, we first estimate the sequence of demand and supply innovations that replicate, under the baseline IT rule, the U.S. unemployment and inflation dynamics from 1990:1 to 2019:12.[26] Figure 4 reports the resulting filtered shock series along with the model implied unemployment rate and inflation paths, which, by construction, replicate the data exactly. The filtered path for the risk-premium wedge closely follows the movements in the unemployment rate. In contrast, the cost-push shock tends to peak in periods where inflation is, through the lens of the model, too high relative to the slack of the labor market.[27] We use this series of filtered shocks to simulate the Great Recession and its recovery dynamics (2008:1 through 2019:12) under a LfL monetary policy rule, and contrast its predicted outcomes to the baseline economy. Figures 5 and 6 summarize our results.

At the onset of the recession and upon reaching the ZLB, the new strategy promises to keep rates lower for longer through its reaction to past average inflation, helping to

[26] We follow the filtering algorithm described in McKay and Wieland (2021) and compute the sequence of shocks recursively. See Appendix E for more details on this step.

[27] As seen in the top-right panel of Figure 4, the shock peaks around 2010 when inflation remained in the proximity of 2% in spite of a very strong labor market.

alleviate the downfall relative to the baseline IT. Around 2015, as unemployment starts to fall below trend, the asymmetric reaction to unemployment deviations embedded in the rule kicks in and amplifies the effects of expansionary shocks on labor market outcomes. Between 2015 and 2019, the simulation shows a marked decoupling of the counterfactual economy under the LfL rule from the data: at the end of the sample period, for example, unemployment is 2 ppts lower, participation 1 ppt higher, and output 5 ppt higher. Relative to IT, the new rule effectively "lets the economy run hot" during the recovery.

Figure 6 shows that, as expected, these better labor market outcomes relative to the baseline IT rule are especially pronounced at the bottom of the skill distribution: compared to the top quartile, the gains in participation in the bottom quartile are twice as large, and the gains on earnings per worker —through additional skill accumulation and foregone skill depreciation— three times as large.

Figure 5 illustrates that the these more inclusive outcomes are obtained by keeping real rates lower for longer compared to the strict IT rule. The cost of the new strategy is higher inflation: while the data show that inflation, under strict IT, remained 0.5 ppts below target throughout the recovery, under the LfL rule, inflation is on average 0.75 ppts above it. Thus, the new framework transforms the negative inflation bias caused by the ZLB into a positive bias.

The figures also report simulations under the AIT only rule, which serves to highlight the role of this component of the LfL rule. The average inflation block of the new framework is especially important in preventing excessively depressed participation rates and severe earnings losses for low-wage workers at the height of the Great Recession (bottom row of Figure 6).

4.2 Ergodic Distribution

In this section we ask what would happen if the Fed were to pursue this alternative LfL strategy for a prolonged stretch of time, covering multiple business cycles episodes. For this purpose, we perform a series of long simulations where we continuously subject the model economy to demand and supply shocks that, occasionally, throw the economy into the ZLB. Table 5 reports the standard deviation and averages, expressed as deviations from the non-stochastic steady state, for some variables of interest under the baseline and alternative monetary policy rules discussed in Table 4. A visual representation of the same exercise is shown in Figure E1 in Appendix E. In what follows, we discuss the results of the ergodic simulation for each monetary policy rule.

	Inflation Targeting		AIT		Shortfall		Lower for Longer	
	std	mean	std	mean	std	mean	std	mean
Price inflation	0.92	-0.10	0.46	0.12	1.35	0.01	0.57	0.32
Output	2.77	-0.90	1.52	0.21	4.98	-0.45	2.25	1.22
Hours	1.07	-0.19	0.59	0.03	1.84	-0.22	0.89	0.18
Unemployment rate	1.26	0.25	0.69	-0.04	2.19	0.27	1.04	-0.25
Participation	0.48	-0.19	0.25	0.03	0.86	-0.09	0.37	0.24
Total Labor Earnings (b25)	9.71	-4.34	5.28	0.76	17.72	-1.40	7.82	5.88
Earnings per worker (b25)	4.10	-1.69	2.36	0.38	7.53	-0.39	3.46	2.52
Unemployment rate (b25)	2.68	0.67	1.46	-0.12	4.73	0.57	2.16	-0.75
Participation (b25)	0.95	-0.50	0.55	0.06	1.68	-0.09	0.81	0.66
Recessions (% simulation)	—	0.2415	—	0.2009	—	0.2630	—	0.2179
Expansions (% simulation)	—	0.2259	—	0.2655	—	0.3487	—	0.3605
ZLB frequency	—	0.0760	—	0.0270	—	0.0866	—	0.0313

Table 5: Standard deviatons and averages of aggregate and distributional series in ergodic simulation rule expressed as deviations from their non-stochastic steady state values under the baseline Inflation Target, Asymmetric AIT, Shortfall, and Lower for Longer rule. The labels 'b25' and 't25' indicate outcomes for the bottom and top quartile of the skill distribution. Deviations for inflation, unemployment and participation rates are in ppts, whereas for output, labor productivity, and earnings they are in percentages. Recessions (expansions) are times when the aggregate unemployment rate is more than 0.5 ppts above (below) its deterministic steady-state.

Strict inflation targeting. We start by analyzing the results under the baseline IT rule, reported in the first two columns of Table 5. The average frequency of ZLB episodes is 7.6%. The occasionally binding ZLB produces a negative inflation bias, with average inflation falling 10 basis points short of the 2 percent target. The bias also shows up in labor market outcomes, with employment and participation systematically running below their steady-state values. As a result, the negative bias on output is nearly 1 ppt. In line with our discussion of Section 3.2, the negative employment effects of the ZLB are much more pronounced at the bottom quartile of the skill distribution. For this subset of workers, the constraint on nominal rates amplifies unemployment by almost 0.7 ppts, shrinks participation by 0.5 ppts, and lowers total labor earnings by over 4 ppts relative to the non-stochastic steady state.

Average inflation targeting. The second two columns of Table 5 illustrate the performance of the economy under the first of the two components of the LfL strategy, the asymmetric AIT. This rule is successful in significantly reducing the dispersion of inflation and unemployment rates, which has the consequence of significantly decreasing the frequency and the intensity of ZLB episodes. As a result, the counterfactual economy spends less time in recessions and more time in expansions compared to the data. Overall, the volatility of other aggregates and distributional variables is also much reduced.

A less frequent binding ZLB and the asymmetric reaction with respect to past inflation is enough to fully reverse the negative bias on both inflation and real allocations. For example, the total earnings of low-wage workers rise by over 5 ppts compared to the IT rule.

Shortfall rule. Columns five and six of Table 5 summarize the dynamics of the economy under the Shortfall rule alone, the second component of the LfL strategy. This rule, in isolation, performs very poorly. Even though it corrects the negative inflation bias induced by the ZLB, the performance of the real economy in various dimensions is even inferior to that under IT: the economy is much more volatile, the ZLB binds even more often, and some of the real aggregate (like unemployment) are even worsened under this rule.

These results are surprising and, at a first pass, counterintuitive in light of what we had observed in our Great Recession counterfactual. To understand these findings, one must return to our short-run analysis and note two features of this rule. First, because this component gets activated only once unemployment has fully recovered, the Shortfall rule is not effective at improving the state of the labor market during recessions. Second, by protracting a strong labor market, the Shortfall rule generates significant and persistent participation and labor productivity gains during expansions.

It is the combination of these two properties that has perverse effects in an economy that goes repeatedly through multiple business cycles. When a sequence of negative demand shocks hit after a long expansion, labor productivity and participation are still high. As a result, hours worked have to fall substantially to accommodate the decline in aggregate demand, which amplifies the recession and the severity of ZLB episodes. Put differently, the gains of the Shortfall rule at the end of an expansion can cast a dark shadow on the outcomes over the next recession.

Lower for Longer. The perverse dynamic effects of the Shortfall rule nearly vanish when this rule is implemented jointly with the asymmetric AIT strategy, as we can see from the last two columns of Table 5 that report the results for the Lower for Longer strategy. Adopting both components together causes output, participation and earnings to average above their steady state values. This is because the two strategies operate to boost real

	Total		Top 25		Bottom 25	
	$\mathbb{E}[X_t^{LfL}] - \mathbb{E}[X_t^{IT}]$	% Total	$\mathbb{E}[X_t^{LfL}] - \mathbb{E}[X_t^{IT}]$	% Total	$\mathbb{E}[X_t^{LfL}] - \mathbb{E}[X_t^{IT}]$	% Total
Total Labor Income	2.1172	1.0000	1.5402	1.0000	10.2281	1.0000
Unemployment rate	-0.5014	0.2505	-0.2373	0.1578	-1.4173	0.1661
Participation	0.4312	0.2660	0.0929	0.0639	1.1651	0.4219
Earnings per worker	1.0240	0.4836	1.1986	0.7782	4.2073	0.4113

Table 6: Earnings gain decomposition into the three Okun's channels. Numbers under the $\mathbb{E}[X_t^{LfL}] - \mathbb{E}[X_t^{IT}]$ column denote average differences between outcomes under the LfL regime and the baseline IT rule. Number under % Total column indicate the contribution of each one of the channels to the changes in average total labor income.

outcomes at different points of the business cycle.

By helping with the ZLB, the asymmetric AIT component makes recessions less frequent (they fall from 24% to 22% in our simulation) and less severe. As a consequence, it lessens the scarring effects of recessions on participation and earnings. The Shortfall component still leads to stronger expansions, during which the economy experiences large gains in unemployment, participation and earnings. Crucially, the negative effects over future recessions that we discussed above are weakened by the AIT component because the AIT is very effective at preventing the ZLB from binding: under the LfL rule the ZLB binds almost 1/3 of the times than under the Shortfall rule alone.[28]

The LfL strategy leads to large gains at the bottom quartile of the skill distribution, with total earnings of low-wage workers rising by over 10 ppts compared to the IT rule. Since effects are much stronger at the bottom than at the top quartile, both earnings and consumption inequality are reduced. The cost of running this strategy is that it pushes average inflation roughly 30 basis points above the target, more than the combined bias of each rule implemented individually.

4.2.1 The Role of the Three Okun's Channels

Lower unemployment, stronger participation and higher labor earnings per worker all contribute to the increase in total labor earnings we observe under the LfL rule. But how important is each one of these Okun's channels? To answer this question, note that total

[28] The negative long-run effects of stronger expansions under the Shortfall rule is still present in the simulation as the ZLB binds more frequently under the LfL strategy than under the asymmetric AIT rule alone. The magnitude of these effects is, however, much reduced.

earnings of group g (e.g., the bottom quartile) at time t can be written as:

$$\text{Total earnings}_{gt} = (1 - u_{gt}) \times P_{gt} \times w Z^e_{gt} h_t$$

where u_{gt} is the unemployment rate for that group, P_{gt} its number of labor force participants, w the real wage, Z^e_{gt} average labor productivity, and h_t hours per worker. The last three terms, combined, equal earnings per worker.

Table 6 decomposes the average difference in total earnings between the baseline IT rule and the LfL strategy into the contributions of changes to the unemployment, participation and earnings per worker for the bottom and top quartiles of the skill distribution.

First, we note that the LfL effect on total earnings is much more pronounced at the bottom quartile of the skill distribution. This group experiences earnings that are almost 10% higher under the LfL strategy relative to the baseline IT rule. Out of this effect, over 40% comes from the higher participation (1.2 ppts above IT) and a similar share comes from higher human capital (4.2 ppts above IT). This pattern stands in stark contrast with the behavior at the top quartile of the wage distribution, where stronger participation plays virtually no role, and nearly 80% of the observed labor income gains is due to higher earnings conditional on employment.[29]

4.3 Okun's Cones: A Menu for Policymakers

Up to now, we have analyzed and quantified the inflation-inclusion trade-off exclusively for a specific parameterization of the LfL rule. We now assess the entire menu of outcomes available to the policymaker interested in exploring some version of the Lower for Longer strategy. Ultimately, we measure the cost, in terms of average long-run inflation, required to achieve a certain average long-run gain (e.g., 1 ppt) in a particular outcome variable (e.g., participation) for a group of workers (e.g., bottom quartile of the skill distribution).

For this purpose, we repeat the simulations from Section 4.2 varying the reaction coefficients β_{AIT} and β_u^- between the values taken under the LfL rule and the IT rule. For each simulation, we track the average biases for unemployment, participation and average earnings per worker at the aggregate, bottom and top 25% of the skill distribution. The resulting mapping which, due to their shape we denote by *Okun's cones*, is plotted in Figure 7.

The cones' lower and upper limits reproduce the outcomes under baseline IT and

[29] In addition, for the bottom quartile higher earnings per worker are almost entirely explained by higher skills, while for the top quartile 1/3 of the rise in earnings per worker is accounted for by longer hours

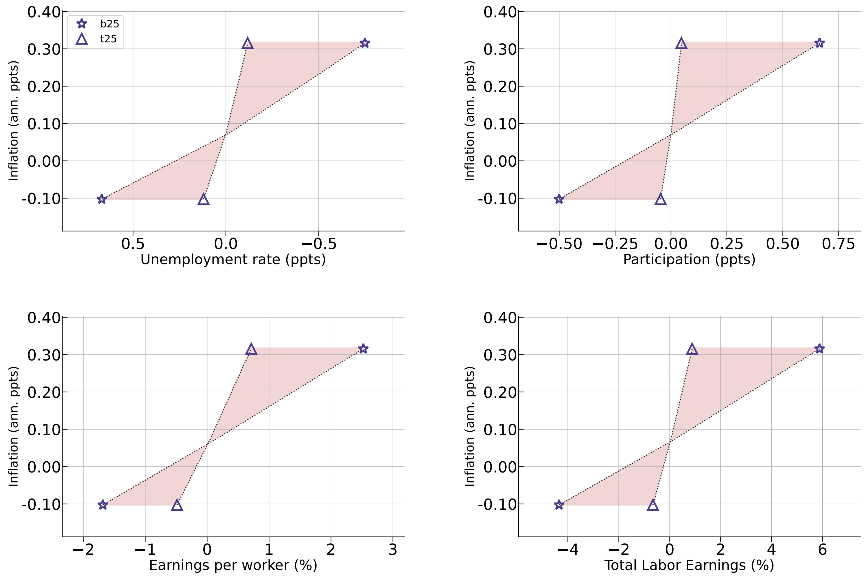

Figure 7: Okun's cones. The outcomes for different points of the distribution are indicated as follows: star denotes outcomes for the bottom 25% of the skill distribution, triangle denotes outcomes for the top 25% of the skill distribution, circle denotes the average/aggregate outcomes. The pink area represent outcomes for the rest of the worker skill distribution.

the parameterized Lower for Longer rule. Starting from the bottom, we observe that the 10bp deflationary bias in inflation caused by the ZLB under IT is accompanied by higher unemployment, lower participation and labor earnings, especially so at the bottom quartile of the skill distribution (denoted by the star in the plot). A central bank that aims at fully closing these contractionary biases relative to the non-stochastic steady-state, can do so with a LfL type of rule at the cost of pushing inflation 5bp above the target.

An especially inequality-averse policymaker can go beyond this point and push for even stronger "inclusion". For example, to obtain an additional 5 ppts gain in average earnings in the bottom quartile, the policymaker should be willing to accept inflation to rises, on average, by 25 basis points.

Negative and positive inflation bias in the long-run. The model simulations of the Lower for Longer rule show that running a more "inclusive" monetary policy rule might result in a persistent gap between the economy's average inflation and the central bank's target π^*. This relation between average labor market outcomes and inflation over our simulation exercise amounts to a long-run trade-off. After any shock, the central bank commits to bring inflation back to its long-run target π^*. But as the economy is continuously buffered by shocks, the asymmetric response of the Lower for Longer strategy results in stronger labor market outcomes and a slightly positive inflation bias, as we showed in Table 5.

In the same way that the deflationary bias arising from the ZLB can tarnish central bank's credibility and, in more extreme circumstances, push the economy into deflationary spirals, a monetary framework that runs inflation systematically above target also risks de-anchoring long-run inflation expectations from the target. In the spirit of Lucas Jr (1972), one would think that over time this inflationary bias should become ingrained in agents' expectations and wage setters behavior in a way that weakens the trade-off.

Whether the trade-off depicted by Okun's cones can be even partially exploited beyond simply eliminating the contractionary bias induced by the ZLB, and if so for how long, is a question that pertains to how agents form expectations. This question goes beyond the scope of the current paper. We do, however, note that in all our simulations, solid gains for the bottom of the wage distribution can be achieved for very moderate deviations of inflation from target, e.g. around 25 basis points. The literature on expectation formation concludes that expectations are sticky, and that agents become more attentive and informed only when inflation rises or fall very significantly relative to the status quo (Weber et al., 2023).

4.4 More Effective Monetary Policy Rules.

In progress.

5 Conclusions

In the next version of the paper, we plan to incorporate three additional exercises into the analysis. First, we plan to integrate our cross-sectional analysis with a longitudinal simulation to report gains in terms of upward mobility. Second, we plan to explore alternative monetary policy rules that can achieve a more favorable inflation-inclusion trade-off. Third, we will analyze the extent to which this inflation-inclusion trade-off can also be generated by an asymmetric fiscal policy rule, e.g. a rule that spends aggressively in recessions and is not too conservative in expansions. This policy will be associated with a higher level of debt and interest payments: the additional taxes necessary to balance the government budget might partially moderate the labor market gains.

Going forward, the analysis in this paper can be extended in two natural directions. One could analyze the Covid episode in isolation. This episode is characterized by high inflation, but also by extraordinary employment and wage gains for the low-wage workers. This last recession occurred under the new monetary policy regime. One could filter demand and supply shocks that explain the data from 2020:1 until today under the current policy rule, and ask the model to what extent, had the Fed instead followed a traditional IT rule, inflation would have been lower and labor market outcomes for the low-wage workers worse.

As explained in the Introduction, our analysis of the inflation-inclusion trade-off is focused on a credible estimation of the *gains* of inclusive stabilization policy by carefully modelling labor market dynamics across the distribution. Throughout the paper, we have casted inflation as a cost for the economy as a whole, but we have not yet explicitly quantified such cost. A natural next step in this research agenda should be a theoretical framework that integrates the gains-side of the inflation-inclusion trade-off developed in this paper with the many channels through which inflation can be costly for households, and in an uneven way (e.g., heterogeneous consumption baskets, nominal net positions, and nominal wage rigidity). Cardoso et al. (2022), Del Canto et al. (2023), and Pallotti et al. (2023) are recent examples of empirical studies that try to quantify the relative role of these channels.

Appendix

A Recursive Formulation of the Household Problem

In this section we show how the individual problems can be expressed recursively as a Hamilton-Jacobi-Bellman Quasi-Variational Inequality (HJBQVI). We follow Laibson et al. (2021) in using the notation with the max operator which we find more intuitive. We use the symbol $\mathcal{Z}^s$ to denote the infinitesimal generator of the stochastic process defined in (2), transformed in levels, in employment state s. Then, $\mathcal{Z}^s v = \mu^s(z)\partial_z v + \sigma(z)\partial_{zz} v$, with $\mu^s(z) = \left(-\rho_z \log z + \mathbb{I}_{\{s_t=e\}}\,\delta_z^+ - \mathbb{I}_{\{s_t\neq e\}}\,\delta_z^- + \frac{\sigma_z^2}{2}\right) z$ and $\sigma(z) = \sigma_z z$.

The problem of the employed households in (7) at time t can be written as:

$$\rho v_t^e = \max\ \Big\{ \max_{c_t} \left\{ \mathfrak{u}^e\left(c_t, h_t\right) + \partial_a v_t^e \left(r_t a_t + (1 - \mathfrak{t}_t)\, w_t z_t h_t + \phi_t - c_t \right) \right\} + \lambda_{zt}^{eu}\left(v_t^{u_1} - v_t^e\right) \\ + \eta^{en}(v^{n_0} - v^e) + \mathcal{Z}^e v^e + \partial_t v_t^e, \rho v_t^{n_1} \Big\}$$

The problem of the passive non-participant in (3) is:

$$\rho v_t^{n_0} = \max_{c_t}\ \left\{ \mathfrak{u}^n\left(c_t, h_t\right) + \partial_a v_t^{n_0}\left(r_t a_t + \phi_t - c_t\right)\right\} + \eta^{n_0 n_1}\left(v_t^{n_1} - v_t^{n_0}\right) \\ + \mathcal{Z}^{n_0} v_t^{n_0} + \partial_t v_t^{n_0}$$

The problem of the active non-participant in (4) is:

$$\rho v_t^{n_1} = \max\ \Big\{ \max_{c_t} \left\{ \mathfrak{u}^n\left(c_t, h_t\right) + \partial_a v_t^{n_1}\left(r_t a_t + \phi_t - c_t\right)\right\} + \lambda_{zt}^{ne} \max\left\{v_t^e - v_t^{n_1}, 0\right\} + \eta^{n_1 n_0}\left(v_t^{n_0} - v_t^{n_1}\right) \\ + \mathcal{Z}^{n_1} v_t^{n_1} + \partial_t v_t^{n_1}, \rho v_t^{u_0} - \xi \Big\}$$

The problem of the non-eligible unemployed in (5) becomes:

$$\rho v_t^{u_0} = \max\ \Big\{ \max_{c_t} \left\{ \mathfrak{u}^u\left(c_t, h_t\right) + \partial_a v_t^{u_0}\left(r_t a_t + \phi_t - c_t\right)\right\} + \lambda_{zt}^{ue} v_t^e + \eta^{u n_0}\left(v_t^{n_0} - v_t^{u_0}\right) \\ + \mathcal{Z}^{n_0} v_t^{u_0} + \partial_t v_t^{u_0}, \rho v_t^{n_1} \Big\}$$

The problem of the eligible unemployed in (5) becomes:

$$\rho v_t^{u_1} = \max \left\{ \max_{c_t} \left\{ \mathfrak{u}^u \left(c_t, h_t\right) + \partial_a v_t^{u_1} \left(r_t a_t + \phi_t + b\left(z_t\right) - c_t\right) \right\} + \lambda_{zt}^{ue} \max \left\{ v_t^e - v_t^{u_1}, 0 \right\} \right.$$
$$\left. + \lambda^{u_1 u_0} \left(v_t^{u_0} - v_t^{u_1}\right) + \eta^{u n_0} \left(v_t^{n_0} - v_t^{u_1}\right) + \mathcal{Z}^{u_1} v_t^{u_1} + \partial_t v_t^{u_1}, \rho v_t^{n_1} \right\}$$

The four HJBQVI's above, jointly with the five first-order conditions at every t

$$\partial_c \mathfrak{u}^s \left(c_t, h_t\right) = \partial_a v_t^s \quad s \in \{e, u_0, u_1, n_0, n_1\}$$

can be solved for value functions $\left\{v_t^s\left(a_t, z_t\right)\right\}_{t\geq 0}$, consumption decision rules $\left\{c_t^s\left(a_t, z_t\right)\right\}_{t\geq 0}$ for $s \in \{e, u_0, u_1, n_0, n_1\}$, for participation rules $\left\{\mathfrak{p}_t^s\left(a_t, z_t\right)\right\}_{t\geq 0}$ for $s \in \{e, u_0, u_1, n_1\}$, and job acceptance rules $\left\{\mathfrak{f}_t^s\left(a_t, z_t\right)\right\}_{t\geq 0}$ for $s \in \{u_1, n_1\}$.

B Derivation of the Wage Phillips Curve

We reproduce the union problem here for convenience:

$$\max_{\{\omega_{kt}\}_{t\geq 0}} \int_0^\infty e^{-\rho t} \left[\int_{s_{it}=e} \mathfrak{u}^e \left(c_{it}, h_{it}\right) di - \frac{\Theta}{2} \left(\frac{\dot{\omega}_{kt}}{\omega_{kt}} - \pi^* \right)^2 \right] dt$$

$$\begin{aligned}
& s.t. \\
h_{it} &= \int_0^1 h_{kt} dk \\
c_{it} + \dot{a}_{it} &= r_t a_{it} + \left(1 - \mathfrak{t}_t\right) y_{it} + \phi_t \\
y_{it} &= \frac{1}{P_t} z_{it} \int_0^1 \omega_{kt} h_{kt} dk \\
h_{kt} &= \left(\frac{\omega_{kt}}{\omega_t} \right)^{-\varepsilon} h_t
\end{aligned}$$

We can write this problem recursively as

$$
\begin{aligned}
\rho J_t(\omega_{kt}) &= \max_{\pi_{kt}} \int_{s_{it}=e} \mathfrak{u}^e(c_{it}, h_{it})\, di - \frac{\Theta}{2}\left(\pi_{kt} - \pi^*\right)^2 + \partial_\omega J_t(\omega_{kt})\, \omega_{kt}\pi_{kt} + \partial_t J_t(\omega_{kt}) \\
& \quad s.t. \\
c_{it} &= r_t a_{it} + (1 - \mathfrak{t}_t)\, y_{it} + \phi_t - \dot{a}_{it} \\
y_{it} &= \frac{1}{p_t} z_{it} \int_0^1 \omega_{kt} \left[\left(\frac{\omega_{kt}}{\omega_t}\right)^{-\varepsilon} h_t\right] dk \\
h_{it} &= \int_0^1 \left[\left(\frac{\omega_{kt}}{\omega_t}\right)^{-\varepsilon} h_t\right] dk
\end{aligned}
$$

The first order condition with respect to π_{kt} and the envelope condition are, respectively:

$$
\begin{aligned}
\partial_\omega J_t(\omega_{kt})\, \omega_{kt} &= \Theta\left(\pi_{kt} - \pi^*\right) \\
(\rho - \pi_{kt})\, \partial_\omega J_t(\omega_{kt}) &= \int_{s_{it}=e} \left\{\partial_c \mathfrak{u}^e(c_{it}, h_{it}) \left[(1 - \mathfrak{t}_t)\, \partial_\omega y_{it}\right] di + \partial_h \mathfrak{u}^e(c_{it}, h_{it})\, \partial_\omega h_{it}\right\} di \\
& \quad + \partial_{\omega\omega} J_t(\omega_{kt})\, \omega_{kt}\pi_{kt} + \partial_{\omega t} J_t(\omega_{kt})
\end{aligned}
$$

where

$$
\begin{aligned}
\partial_\omega y_{it} &= \frac{1}{P_t} z_{it} \omega_t^{\varepsilon} h_t (1 - \varepsilon) \int_0^1 \omega_{kt}^{-\varepsilon} dk \\
\partial_\omega h_{it} &= -\varepsilon h_t \omega_t^{\varepsilon} \int_0^1 \omega_{kt}^{-(1+\varepsilon)} dk
\end{aligned}
$$

and note that we did not differentiate neither consumption nor participation/job acceptance decisions with respect to ω_{kt} because of the assumption of a 'small union' which cannot affect any individual decisions.

Imposing symmetry $(\omega_{tk} = \omega_t)$, we obtain

$$
\begin{aligned}
\partial_\omega y_{it} &= \frac{1}{P_t} z_{it} h_t (1 - \varepsilon) \qquad \text{(B1)} \\
\partial_\omega h_{it} &= -\varepsilon h_t \omega_t^{-1}
\end{aligned}
$$

and

$$
\begin{aligned}
\partial_\omega J_t(\omega_t) &= \frac{\Theta\left(\pi_t - \pi^*\right)}{\omega_t} \\
\left(\rho - \pi_t\right)\partial_\omega J_t(\omega_t) &= \int_{s_{it}=e} \left\{\partial_c \mathfrak{u}^e(c_{it}, h_{it})\left[(1-\mathfrak{t}_t)\partial_\omega y_{it}\right] di + \partial_h \mathfrak{u}^e(c_{it}, h_{it})\partial_\omega h_{it}\right\} di \\
&\quad + \partial_{\omega\omega} J_t(\omega_t)\,\omega_t \pi_t + \partial_\omega J_t(\omega_t)
\end{aligned}
$$

Substituting (B1) into the envelope condition

$$
\begin{aligned}
\left(\rho - \pi_t\right)\frac{\Theta\left(\pi_t - \pi^*\right)}{\omega_t} &= (1-\mathfrak{t}_t)\frac{1}{P_t}h_t(1-\varepsilon)\int_{s_{it}=e}\partial_c \mathfrak{u}^e(c_{it}, h_{it})\, z_{it} di \\
&\quad - \varepsilon\frac{h_t}{\omega_t}\int_{s_{it}=e}\partial_h \mathfrak{u}^e(c_{it}, h_{it})\, di + \partial_{\omega\omega} J_t(\omega_t)\,\omega_t\pi_t + \partial_{\omega t} J_t(\omega_t).
\end{aligned}
$$

Differentiating the first-order condition (B1) with respect to time yields

$$
\partial_{\omega t} J_t(\omega_t) + \partial_{\omega\omega} J_t(\omega_t)\,\dot{\omega}_t = \frac{\Theta\dot{\pi}_t}{\omega_t} - \frac{\Theta\left(\pi_t - \pi^*\right)}{\omega_t}\left(\frac{\dot{\omega}_t}{\omega_t}\right) \tag{B2}
$$

Substituting (B2) into the envelope condition:

$$
\begin{aligned}
\left(\rho - \pi_t\right)\frac{\Theta\left(\pi_t - \pi^*\right)}{\omega_t} &= (1-\mathfrak{t}_t)\frac{1}{P_t}h_t(1-\varepsilon)\int_{s_{it}=e}\partial_c \mathfrak{u}^e(c_{it}, h_{it})\, z_{it} di - \varepsilon\frac{h_t}{\omega_t}\int_{s_{it}=e}\partial_h \mathfrak{u}^e(c_{it}, h_{it})\, di \\
&\quad + \frac{\Theta\dot{\pi}_t}{\omega_t} - \frac{\Theta\left(\pi_t - \pi^*\right)}{\omega_t}\pi_t
\end{aligned}
$$

Multiply both sides by ω_t/Θ to arrive at:

$$
\begin{aligned}
\left(\rho - \pi_t\right)\left(\pi_t - \pi^*\right) &= \frac{\omega_t}{\Theta}(1-\mathfrak{t}_t)\frac{1}{P_t}h_t(1-\varepsilon)\int_{s_{it}=e}\partial_c \mathfrak{u}^e(c_{it}, h_{it})\, z_{it} di - \frac{\varepsilon}{\Theta}h_t\int_{s_{it}=e}\partial_h \mathfrak{u}^e(c_{it}, h_{it})\, di \\
&\quad + \dot{\pi}_t - \left(\pi_t - \pi^*\right)\pi_t
\end{aligned}
$$

Simplifying this expression, we obtain

$$
\rho\left(\pi_t - \pi^*\right) = \frac{\varepsilon}{\Theta}h_t\left[-\int_{s_{it}=e}\partial_h \mathfrak{u}^e(c_{it}, h_t)\, di - \left(\frac{\varepsilon-1}{\varepsilon}\right)(1-\mathfrak{t}_t)\, w_t\int_{s_{it}=e}\partial_c \mathfrak{u}^e(c_{it}, h_{it})\, z_{it} di\right] + \dot{\pi}_t
$$

Using the fact that, in our functional form for period utility in (1), $\partial_h \mathfrak{u}^e(c_{it}, h_t)$ is independent of i, letting Z_t^e denote average skills of employed workers, and letting H_t denote aggregate hours worked, i.e. h_t times the number of employed at t, we arrive at

$$\rho\left(\pi_t-\pi^*\right)=\frac{\varepsilon}{\Theta}H_t\left[-\partial_h u^e\left(c_{it},h_t\right)-\left(\frac{\varepsilon-1}{\varepsilon}\right)\left(1-\mathfrak{t}_t\right)w_t Z_t^e\int_{s_{it}=e}\partial_c u^e\left(c_{it},h_{it}\right)\left(\frac{z_{it}}{\int_{s_{it}=e}z_{it}di}\right)di\right]+\dot{\pi}_t$$

which equals the expression for the Phillips curve in the main text.

C Kolmogorov-Forward Equations

TBC

D Measurement of Worker Flows by Skills

To estimate how the separation and job finding rates vary with the skill level z_{it}, we use the Basic Monthly Current Population Survey (CPS) from 1989 to 2019, merged with the Annual Social and Economic Supplement (ASEC). The Basic Monthly CPS reports the employment status [EMPSTAT] of each individual interviewed. The ASEC asks each individual additional questions about past earnings and weeks worked in the CPS in March. We use the ASEC supplement because there is no measure of individual labor earnings in the Basic Monthly CPS. A well known concern in the literature measuring worker flows in three-state models is misclassification in labor market status between unemployment and nonparticipation. We follow the correction approach of Elsby et al. (2015) where some sequences of reported labor market states are recoded to eliminate high-frequency reversals of transitions between unemployment and nonparticipation. Examples of the latter are recorded sequences 'uunu' reclassified as 'uuuu' or 'nnun' reclassified as 'nnnn', and so on. We are interested in month-to-month changes in employment status, and therefore treat each observation as a pair of months, $(t, t+1)$. We clean the sample so as to keep only workers observed in at least two consecutive months.

We keep in the sample individuals who appear in the ASEC and are between 25 and 59 years old and exclude self-employed and government employees and unpaid family workers.[30] We further drop all individuals who report positive weeks worked but no earnings and all entries that are assigned zero weight. This sample selection leaves us with a sample of $942,546$ workers, or $5,095$ individuals per month on average.

[30] Because of the rotating nature of the CPS, in which households are interviewed for four months, then left aside for eight months, and finally interviewed again for four months, this leaves many individuals in the Basic Monthly CPS out of our sample because they are not part of any ASEC supplement. More specifically, workers interviewed from July to November are not in our sample.

In the ASEC, we measure individuals' earnings from the annual pre-tax labor income measure available in the ASEC supplement [INCWAGE]. To obtain a proxy for skills, we compute weekly real earnings. We multiply the reported nominal labor earnings by the Consumer Price Index (CPI-U) for that year, and then divide it by the number of annual weeks worked [WKSWORK1].

Every individual in our sample is interviewed at most twice in the ASEC supplement, 12 months apart, in month (always March) t and $t + 12$. As a proxy for skills, we assign the average of weekly earnings in the two supplements to all months in which the individual is present in the Basic Monthly CPS. After this earnings imputation, we have a sample of $3, 933, 752$ entries, or $21, 264$ entries per month on average.

We express skills in relative terms, as the ratio of weekly individual earnings to the average weekly earnings of that year. We then compute, on the data pooled across the whole time period, 10 deciles of the relative skill distribution. We measure, for example, the separation rate for quantile q in month t as the share of workers employed at time t with weekly earnings in quantile q who are no longer employed at $t + 1$. And similarly for the job finding rates.

In the estimation step, we use an equally weighted minimum quadratic distance estimator that minimizes the difference between the transition probabilities predicted by the statistical model in (23) and their empirical counterparts estimated in the data. To estimate the steady-state parameters of our transition functions $(\lambda_0^{ss'}, \lambda_1^{ss'}, \lambda_2^{ss'})$, we pooled data across all months together.

To compute earnings growth by skill level, we select the observations in our sample which are in the outgoing rotation group (CPS-ORG) and have strictly positive values for the variable EARNWEEK (weekly earnings) in both their 4th and their 8th month in the survey (two observations which are 1 year apart). We then compute the growth rate of EARNWEEK across these two points.

E Computation

Impulse response functions We solve for the aggregate dynamics of the economy using the sequence-space approach of Auclert et al. (2021).[31]

[31] In the tradition of other linear methods, the approach linearizes with respects to aggregates and thus does not take into account the ZLB on nominal interest rates. As we discuss in the filtering step, we impose the ZLB through monetary news shocks, following the method described by Holden (2016). The method substitutes the constraint on nominal rates with a sequence of anticipated news shocks, which are chosen to ensure that the nominal rate never violates the constraint (i.e., $i_t \geq 0$) whenever negative shocks hitting the economy would otherwise call for negative rates.

Zero Lower Bound and counterfactuals [TBC]

Filtering For the analysis of alternative monetary policy rules, we estimate the sequence of innovations to demand and supply shocks that replicate the U.S. unemployment and inflation dynamics from 1990:1 to 2019:12. We follow the filtering algorithm described in McKay and Wieland (2021), which does not require information on the state transition matrix, but relies only on the impulse response functions. We compute the sequence of shocks recursively. For each date t, we solve for innovations that explain the surprise movement in inflation and unemployment, that is, the difference between their $t-1$ forecast and their time t observed values. Due to linearity, this step involves solving a linear system where the time t innovations are weighted by the unemployment and inflation impact responses with respect to demand and supply shocks. The innovations and the previous projection then determine the updated projected paths for t + 1 onward.

Before moving to $t+1$, however, we check whether the ZLB is violated at time t or at any other time period in the future projected path. If the ZLB is not violated, we find $t+1$ innovations in the same way described above. If instead the constraint is binding, we solve for the path of anticipated monetary news shocks that incorporate the ZLB along the projected path of nominal interest rates following the methodology of Holden (2016). Because the added news shocks also affect unemployment and inflation at date t, we must update the previously computed innovations so that they are in line with observed data. We thus face a simultaneity problem where innovations and the news shocks must be set jointly to be consistent with the observed data at the t and the ZLB along the entire projected path. An iterative procedure works well in practice, and we are able to recover the time t demand and supply innovations after a few iterations.

E Extra Figures and Tables

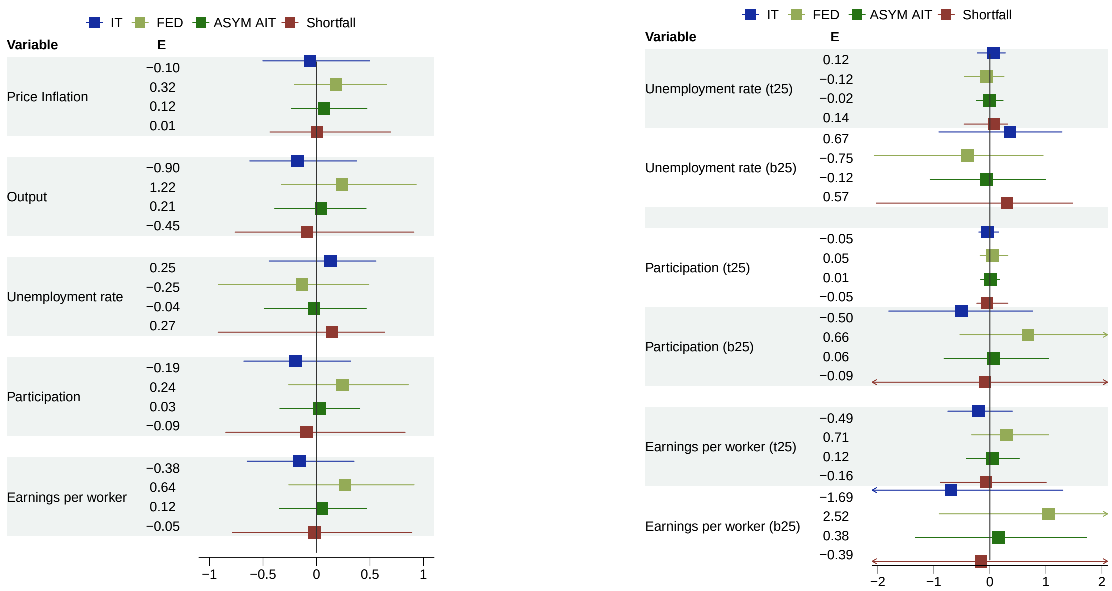

Figure E1: Summary statistics of the simulated economy under alternative monetary policy rules. The middle column reports the averages (E), 10^{th} ($p10$) and 90^{th} ($p90$) percentile of each variable expressed in deviations from their non-stochastic steady-state. The lines show the same information, but are standardized so that the length of each line for an aggregate variable equals 1.0 under the Inflation Target rule.

References

Aaronson, Stephanie R, Mary C Daly, William L Wascher, and David W Wilcox (2019), "Okun revisited: Who benefits most from a strong economy?" *Brookings Papers on Economic Activity*, 2019(1), 333–404.

Alvarez, Fernando, and Robert Shimer (2011), "Search and Rest Unemployment," *Econometrica*, 79(1), 75–122.

Alves, Felipe, Greg Kaplan, Benjamin Moll, and Giovanni L Violante (2020), "A Further Look at the propagation of Monetary Policy Shocks in HANK," *Journal of Money, Credit and Banking*, 52(S2), 521–559.

Athey, Susan, Lisa K. Simon, Oskar N. Skans, Johan Vikstrom, and Yaroslav Yakymovych (2023), "The Heterogeneous Earnings Impact of Job Loss Across Workers, Establishments, and Markets," Working paper, Stanford University.

Auclert, Adrien, Bence Bardóczy, and Matthew Rognlie (2020), "MPCs, MPEs and Multipliers: A Trilemma for New Keynesian Models," National Bureau of Economic Research.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub (2021), "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models," *Econometrica*, 89(5), 2375–2408.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub (2018), "The Intertemporal Keynesian Cross," Working paper, Stanford University.

Autor, David, Arindrajit Dube, and Annie McGrew (2023), "The unexpected compression: Competition at work in the low wage labor market," Working paper, MIT.

Bartscher, Alina K, Moritz Kuhn, Moritz Schularick, and Paul Wachtel (2021), "Monetary policy and racial inequality," CEPR Discussion Paper No. DP15734.

Bergman, Nittai, David A Matsa, and Michael Weber (2020), "Heterogeneous labor market effects of monetary policy," *Chicago Booth Research Paper*, (21-02).

Bianchi, Francesco, Leonardo Melosi, and Matthias Rottner (2021), "Hitting the elusive inflation target," *Journal of Monetary Economics*, 124, 107–122.

Bundick, Brent, and Nicolas Petrosky-Nadeau (2021), "From Deviations to Shortfalls: The Effects of the FOMC's New Employment Objective," Working paper, Federal Reserve Bank of San Francisco.

Cairó, Isabel, Shigeru Fujita, and Camilo Morales-Jiménez (2022), "The cyclicality of labor force participation flows: The role of labor supply elasticities and wage rigidity," *Review of Economic Dynamics*, 43, 197–216.

Cairó, Isabel, and Avi Lipton (2023), "Labor Market Discrimination and the Racial Unemployment Gap: Can Monetary Policy Make a Difference?" Finance and Economics Discussion Series, Federal Reserve Board 2023-065.

Cajner, Tomaz, Tyler Radler, David Ratner, and Ivan Vidangos (2017), "Racial gaps in labor market outcomes in the last four decades and over the business cycle," FED Working Paper.

Cardoso, Miguel, Clodomiro Ferreira, José Miguel Leiva, Galo Nuño, Álvaro Ortiz, Tomasa Rodrigo, and Sirenia Vazquez (2022), "The heterogeneous impact of inflation on households balance sheets," Red Nacional de Investigadores en Economıa Working Paper.

Clarida, Richard (2022), "The Federal Reserve's New Framework: Context and Consequences," FED Working Paper.

Davis, Steven J., and Till Von Wachter (2011), "Recessions and the Costs of Job Loss," *Brookings Papers on Economic Activity*, 2011(2), 1–72.

Del Canto, Felipe N, John R Grigsby, Eric Qian, and Conor Walsh (2023), "Are Inflationary Shocks Regressive? A Feasible Set Approach," National Bureau of Economic Research.

Del Negro, Marco, Michele Lenza, Giorgio E Primiceri, and Andrea Tambalotti (2020), "What's up with the Phillips Curve?" *Brookings Papers on Economic Activity*, 301–357.

Elsby, Michael WL, Bart Hobijn, and Ayşegül Şahin (2015), "On the importance of the participation margin for labor market fluctuations," *Journal of Monetary Economics*, 72, 64–82.

Erceg, Christopher J, Dale W Henderson, and Andrew T Levin (2000), "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics*, 46(2), 281–313.

Faberman, R Jason, Andreas I Mueller, Ayşegül Şahin, and Giorgio Topa (2017), "Job search behavior among the employed and non-employed," National Bureau of Economic Research.

Federal Reserve System (2020), "FedListens: Perspectives from the Public," Federal Reserve Board.

Fernández-Villaverde, Jesús, Joël Marbet, Galo Nuño, and Omar Rachedi (2023), "Inequality and the Zero Lower Bound," Working Paper 31282, National Bureau of Economic Research.

Finkelstein, Amy, Nathaniel Hendren, and Erzo FP Luttmer (2019), "The value of medicaid: Interpreting results from the oregon health insurance experiment," *Journal of Political Economy*, 127(6), 2836–2874.

Fornaro, Luca, and Martin Wolf (2020), "The scars of supply shocks," CEPR Discussion Paper No. DP15423.

Graves, Sebastian, Christopher Huckfeldt, and Eric T Swanson (2023), "The Labor Demand an Labor Supply Channels of Monetary Policy," Working Paper, Federal Reserve Board.

Grigsby, John, Erik Hurst, and Ahu Yildirmaz (2021), "Aggregate nominal wage adjustments: New evidence from administrative payroll data," *American Economic Review*, 111(2), 428–71.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song (2017), "Heterogeneous scarring effects of full-year nonemployment," *American Economic Review*, 107(5), 369–373.

Guvenen, Fatih, Luigi Pistaferri, and Giovanni L Violante (2022), "Global trends in income inequality and income dynamics: New insights from GRID," *Quantitative Economics*, 13(4), 1321–1360.

Hall, Robert E., and Marianna Kudlyak (2019), "Job-finding and job-losing: A comprehensive model of heterogeneous individual labor-market dynamics," National Bureau of Economic Research.

Heathcote, Jonathan, Fabrizio Perri, Giovanni L. Violante, and Lichen Zhang (2023), "More unequal we stand? Inequality in the United States from the Great Recession to the Covid Pandemic," *Review of Economic Dynamics*, 50, 235–266.

Hebden, James, and Fabian Winkler (2021), "Impulse-Based Computation of Policy Counterfactuals," Working paper, Federal Reserve Board.

Hobijn, Bart, and Ayşegül Şahin (2021), "Maximum Employment and the Participation Cycle," National Bureau of Economic Research.

Holden, Tom D (2016), "Computation of solutions to dynamic models with occasionally binding constraints," Working paper, University of Surrey.

Huo, Zhen, and José-Víctor Ríos-Rull (2020), "Sticky Wage Models and Labor Supply Constraints," *American Economic Journal: Macroeconomics*, 12(3), 284–318.

Ioannidis, Michael, Sarah Jane Hlásková, and Chiara Zilioli (2021), "The mandate of the ECB: Legal considerations in the ECB's monetary policy strategy review," *ECB Occasional Paper*, (2021-276).

Jordà, Òscar, Sanjay R Singh, and Alan M Taylor (2020), "The long-run effects of monetary policy," National Bureau of Economic Research.

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

Krusell, Per, Toshihiko Mukoyama, Richard Rogerson, and Ayşegül Şahin (2017), "Gross Worker Flows over the Business Cycle," *American Economic Review*, 107(11), 3447–76.

Laibson, David, Peter Maxted, and Benjamin Moll (2021), "Present bias amplifies the household balance-sheet channels of macroeconomic policy," National Bureau of Economic Research.

Lee, Munseob, Claudia Macaluso, and Felipe Schwartzman (2021), "Minority Unemployment, Inflation, and Monetary Policy," Federal Reserve Bank of Richmond.

Lucas, Robert E, and Edward C Prescott (1978), "Equilibrium Search and Unemployment," in *Uncertainty in Economics*, 515–540, Elsevier.

Lucas Jr, Robert E (1972), "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 4(2), 103–124.

Ma, Yueran, and Kaspar Zimmermann (2023), "Monetary Policy and Innovation," National Bureau of Economic Research.

McKay, Alisdair, and Johannes F Wieland (2021), "Lumpy durable consumption demand and the limited ammunition of monetary policy," *Econometrica*, 89(6), 2717–2749.

Nakajima, Makoto (2022), "Monetary Policy with Racial Inequality," Federal Reserve Bank of Philadelphia.

Nakamura, Emi, and Jón Steinsson (2008), "Five facts about prices: A reevaluation of menu cost models," *The Quarterly Journal of Economics*, 123(4), 1415–1464.

Okun, Arthur M (1973), "Upward mobility in a high-pressure economy," *Brookings Papers on Economic Activity*, 1973(1), 207–261.

Pallotti, Filippo, Gonzalo Paz-Pardo, Jiri Slacalek, Oreste Tristani, and Giovanni L. Violante (2023), "Who Bears The Costs of Inflation? Euro Area Households and the 2021-22 Shock," ECB Working Paper.

Powell, Jerome H. (2020), "New Economic Challenges and the Fed's Monetary Policy Review," Technical report, Remarks at the jackson Hole Symposium sponsored by the Federal Reserve Bank of Kansas City.

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

Weber, Michael, Bernardo Candia, Tiziano Ropele, Rodrigo Lluberas, Serafin Frache, Brent H Meyer, Saten Kumar, Yuriy Gorodnichenko, Dimitris Georgarakos, Olivier Coibion, Geoff Kenny, and Jorge Ponce (2023), "Tell Me Something I Don't Already Know: Learning in Low and High-Inflation Settings," Working Paper 31485, National Bureau of Economic Research.