

Layoff risk, the welfare cost of business cycles, and monetary policy *

David Berger, Ian Dew-Becker, Lawrence D.W. Schmidt, and Yuta Takahashi

August 20, 2018

First draft: September, 2015

Abstract

The single strongest predictor of changes in the Fed Funds rate in the period 1982–2008 was the level of the layoff rate (initial unemployment claims divided by total employment). That fact is puzzling from the perspective of representative-agent models of the economy, which typically imply that the welfare gains of stabilizing employment fluctuations are small. It is now well established, though, that accounting for the heterogeneous effects of business cycles can substantially increase their welfare costs. This paper augments a standard New Keynesian model with a labor market featuring countercyclical layoffs that lead to large permanent wage declines. In our benchmark calibration, welfare may be increased by 1 percent of lifetime consumption or more when the central bank’s policy rule responds to the layoff rate instead of purely targeting inflation. The theory provides a theoretical rationale for the Federal Reserve’s dual mandate and its apparent responsiveness to changes in the number of layoffs.

*Berger: Northwestern University and NBER. Dew-Becker: Northwestern University and NBER. Schmidt: University of Chicago. Takahashi: Northwestern University. We thank Adrien Auclert, Gadi Barlevy, Saroj Bhattarai, Olivier Coibion, Morten Ravn, Pontus Rendhal, Ivan Werning, Johannes Weiland, and seminar participants at the Federal Reserve Bank of Chicago, University of Texas, Austin, The University of Chicago, the Econometric Society Asian Meeting, and the SED for helpful comments.

1 Introduction

To the average person, a recession is a relatively minor event, but to those who lose their jobs, the effects are large and long-lasting. While the average person’s income might fall by a percentage point or two, a large empirical literature shows that people who are laid off in recessions can see declines in lifetime income of as much as 30 percent.¹ Idiosyncratic risk varies over the business cycle (Storesletten, Telmer, and Yaron (2001); Guvenen, Ozkan, and Song (2014)), and welfare calculations imply that such variation can be a source of large losses (Imrohroglu (1989); Krebs (2007)). So a key potential reason that stabilization of the business cycle might raise welfare is that it would stabilize idiosyncratic consumption risk.²

If time-variation in idiosyncratic risk is the source of welfare costs in recessions, then we might expect government policy to respond to it. In particular, given the evidence on the cost of job loss, a natural reading of the Federal Reserve’s mandate to promote “maximum employment” might be to attempt to minimize layoffs. This paper begins by documenting a novel fact about monetary policy in the US: in the period since 1982, the aggregate layoff rate – initial unemployment claims divided by aggregate employment – has been the single strongest predictor of changes in interest rates. That result holds even after controlling for a wide range of measures of inflation and real activity and examining both forward- and backward-looking policy rules.

In the context of many models of the business cycle and monetary policy, such an empirical finding is puzzling. A common implication of workhorse models is that policy should purely target inflation and that setting interest rates to respond to real activity in any way leads to measurable welfare losses (e.g. Yun (2005)?, Schmitt-Grohe and Uribe (2007)). Even in models where it is optimal for the central bank to put some weight on the output gap or the change in unemployment (e.g. Rotemberg and Woodford (1999), Faia (2008), Blanchard and Gali (2010), Ravenna and Walsh (2011)) the weight in the Taylor rule on measures of real activity tends to be small.

Motivated by the evidence on the welfare costs of idiosyncratic risk and the behavior of the Federal Reserve, this paper develops a tractable New Keynesian model of the business cycle featuring uninsurable risk associated with job losses to try to understand optimal monetary policy.³ While the research cited above has considered models with unemployment, our innovation is to consider optimal monetary policy when unemployment is associated with uninsurable risk.⁴

¹Notable examples are von Wachter, Song, and Manchester (2009) (finding long-term wage losses of 30 percent); Jacobson, LaLonde, and Sullivan (1993) (25 percent long-term wage loss); Davis and von Wachter (2011) (10 to 20 percent loss); Ruhm (1991) (11–15 percent); and Kletzer and Fairlie (2003) (9–13 percent losses for young workers); among many others.

²See Mukoyama and Sahin (2006), Krebs (2007), Krusell, Mukoyama, Sahin, and Smith (2009) and Krueger, Mitman and Perri (2016).

³The importance of idiosyncratic risk for policy analysis has been understood for some time. Clarida, Gali and Gertler (1999) write: “while the widely used representative agent approach may be a reasonable way to motivate behavioral relationships, it could be highly misleading as a guide to welfare analysis. If some groups suffer more in recessions than others (e.g. steel workers versus professors) and there are incomplete insurance and credit markets, then the utility of a hypothetical representative agent might not provide an accurate barometer of cyclical fluctuations in welfare.”

⁴Nakajima (2010)? is somewhat similar to this paper, but the model is centered on efficiency wages instead of

The fundamental source of heterogeneity in the model is in human capital. The key assumption is that unemployment is associated with a permanent and uninsurable reduction in human capital (which is partially an endogenous choice by the agents). Job losses are therefore associated not just with transitory periods of zero income, but in fact permanent declines in income. Those declines then pass into consumption, making separations generate large idiosyncratic tail risk in consumption. While simple, the model is able to match numerous features of income patterns following job loss.

At our baseline calibration, the welfare cost of business cycles is 0.95 percent of lifetime consumption, which is comparable to the results obtained in Krebs (2007). Of that total, though, only 0.04 percentage points is due to inflation volatility. We have the usual channel for inflation to affect welfare – if increases dispersion in prices across intermediate good producers. But as is well known, the effect of such dispersion is almost always quantitatively very small. Instead, welfare in the model is driven by variation in idiosyncratic risk, due to changes in both the probability and income consequences of job loss.

Since it is variation in idiosyncratic risk that primarily drives the cost of business cycles, the welfare maximizing policy is one that minimizes fluctuations in risk. In particular, the best policy rules focus strongly on stabilizing employment so as to minimize variation in layoffs over time. We in fact obtain nearly a corner solution: when considering Taylor rules, in the part of the parameter space where the model has a solution, welfare is almost always improved by putting more weight on the layoff rate and less on inflation. This finding runs counter to well known results on optimal monetary policy that imply that the central bank should actually ideally put no weight on output and focus purely on inflation.

Consistent with the empirical fact that the Fed Funds rate covaries more strongly with layoffs than output, we also find in the model that it is more efficient (in the sense of Ball (1999)) for the central bank to set the Taylor rule to respond to layoffs than to output. For any level of inflation volatility, a lower volatility of output or employment can be achieved by setting the Taylor rule to respond to layoffs instead of output.

Standard representative agent models imply that pure inflation targeting is optimal because realistic fluctuations in output in those settings have extremely small welfare costs (Lucas (1987)), while variation in inflation leads to relatively larger inefficiencies in production. Schmitt-Grohe and Uribe (2007) find therefore that it is optimal to stabilize inflation, which maximizes productivity and output and leads to welfare gains of approximately 0.3 percent of lifetime consumption compared to a rule that also responds to the output gap. Our results are therefore notable both for their direction – the Federal Reserve *should* pay attention to the state of the business cycle – and for their magnitude: we obtain potential welfare gains four times larger those that appear in Schmitt-Grohe and Uribe (2007). While models with search frictions, such as Ravenna and Walsh (2011), generate larger losses from output fluctuations – because variations in search intensity are costly) – the costs

human capital, and the consumption of the unemployed is set exogenously, whereas here it is endogenous. We also examine a wide range of linear policy rules.

of idiosyncratic risk here are quantitatively larger.

We focus on idiosyncratic risk because it is a well established phenomenon, whereas other proposed costs in the literature – such as price dispersion and search costs – are difficult or impossible to measure directly. An important contribution of the paper is to show how to model idiosyncratic risk and individual saving tractably in a standard linear setting. While recent work has made major advances in modeling heterogeneity in macro models (e.g. Kaplan, Moll, and Violante (2018)?), we show here that a relatively simple extension of standard models can generate sufficient heterogeneity to capture large welfare costs of business cycles and help understand optimal policy.

It is important to note that it is optimal in the model for the central bank to intervene in markets even at the flexible-price equilibrium. That is due to the simple fact that there is uninsurable risk. By stabilizing the level of employment, the central bank also stabilizes that risk. Because insurance is incomplete, the central bank’s behavior reduces the welfare cost of market incompleteness. Were markets complete, so that there were no idiosyncratic consumption risk, then our conclusions would be reversed, as the would reduce to the standard representative-agent framework that past research has focused on.

In the end, the paper is easily summarized: when a New Keynesian model is augmented so as to generate large wage losses following layoffs, it implies that optimal monetary policy will attempt to minimize fluctuations in the layoff rate. Empirically, the Federal Reserve appears to do just that.

Our work builds on a number of important and influential areas of past research. Lucas (1987, 2003) discusses the welfare cost of business cycles in settings with both representative and heterogeneous agents. Storesletten, Telmer, and Yaron (2004) and Guvenen, Ozkan, and Song (2014)), among others, provide evidence on the magnitude and countercyclicality of idiosyncratic income risk.⁵

The theoretical model builds on a large literature that examines monetary policy in micro-founded New Keynesian settings, including Rotemberg and Woodford (1999), Woodford (2003), Christiano, Eichenbaum, and Evans (2005), Schmitt-Grohe and Uribe (2007), and Coibion, Gorodnichenko, and Wieland (2012), among many others. Krause and Lubik (2007) and Braun and Nakajima (2012)?, also study optimal policy in the presence of heterogeneity. We differ from them in studying a different model of labor markets that we argue fits the data on the cost of job loss better, both on average and over the business cycle.

In considering idiosyncratic labor income risk in a model with sticky prices, our work is closely related to that of Ravn and Sterk (2015)?, Challe and Ragot (2015), Challe et al. (2015), den Haan et al. (2015)? and Werning (2015)?. This paper differs from those in two important respects. First, our primary focus is on optimal policy whereas the previous literature has primarily focused on understanding how labor market risk affects the precautionary savings motive of consumers. Second, our model is sufficiently simple that it can be easily linearized, which suggests that it will

⁵See also Lucas (2003) for a discussion of how the Lucas (1987) welfare calculation can be substantially magnified when heterogeneity is taken into account and the variance of idiosyncratic shocks is countercyclical, as in Storesletten, Telmer, and Yaron (2004) and Guvenen, Ozkan, and Song (2014).

be relatively simple to estimate or modify in future work.

In work subsequent to this paper, Challe (2018) also discusses optimal policy in the presence of uninsurable income risk. That paper focuses relatively more on the fact that uninsurable risk leads to variation in demand and hence the natural rate of interest, whereas we focus on the fact that stabilization policy reduces uninsurable risk itself. Challe (2018) also provides a more extensive analysis of the Ramsey optimal policy and different wage bargaining models, whereas we focus on robustness to preferences, unemployment insurance, and specifications of the consumption process and saving technology.

An innovation of the model is in its ability to accommodate heterogeneity, uninsurable risk, and private saving. The model uses an idea from Krebs (2003), letting agents save in terms of human capital. There is then endogenous saving when an agent is employed, and dissaving during unemployment. The dissaving both helps smooth consumption and also helps the model fit the observed persistent declines in wages following layoffs. Specifying saving in terms of human capital means that the model retains a form of linear homogeneity that makes it tractable, to the point that it can be well approximated by a simple linearization, and should be well suited to estimation.

The ability of the model to account for idiosyncratic risk in both the Euler equations and the welfare calculations represents a methodological contribution. Variation in idiosyncratic risk over time passes through to affect precautionary savings and hence consumption demand. Both the welfare and demand effects are obtained in a purely linear approximation that can easily be solved, simulated, and estimated. It is widely understood that models can generate much larger welfare costs when they account for the possibility that the pain of business cycles is focused on only a fraction of the population. But because such models are usually very difficult to work with, they are rarely used for policy analysis. So an important contribution of this paper to the literature on optimal monetary policy is to extend a standard linearizable New Keynesian model of the business cycle to account for heterogeneous effects of business cycles on workers

The remainder of the paper is organized as follows. In section 2 we briefly review the evidence on the long term earnings losses following a job displacement event. Section 3 estimates models of the interest rate rule in the U.S. since 1982 and shows that the layoff rate has been the most important driver of interest rate changes. We then proceed to build the model of the economy in section 4 and examine its basic behavior in section ???. Section 6 quantifies the welfare cost of business cycles in the model, and section 7 then studies how monetary policy can reduce those losses. Section 8 concludes.

2 Empirical estimates of the cost of job loss

There is a large literature that studies the cost of job loss using a range of data sources and methods. The studies find that wages fall by 10 to 25 percent following a job loss, with the magnitude differing depending on the population studied, the time since the layoff, and the state of the business cycle when the layoff occurred.

Jacobson, Lalonde, and Sullivan (1993) examine the income of workers in Pennsylvania between 1974 and 1986. They find large and persistent effects from job loss: the average initial drop in earnings relative to pre-displacement earnings in their data is 50 percent, and six years after the displacement event earnings are still 25 percent below their pre-displacement level. More recently von Wachter, Song, and Manchester (2009) and Davis and von Wachter (2011) use Social Security records to document economy-wide earnings consequences from displacement over a more representative sample and a longer time-span. Both studies find large and highly persistent earnings losses for both low and high tenure workers. Davis and von Wachter (2011) report that the present discounted value of earnings losses for men with three or more years of tenure after a the displacement event are 11.9 percent relative to pre-displacement earnings trends. The average effect masks significant heterogeneity in the cost of job loss across time and across different workers. The present discounted value of earnings losses is 9.9 percent of pre-displacement earnings in expansions and 19.8 percent in recessions.

Davis and von Wachter (2011) also document that while there is significant cross-sectional heterogeneity in the cost of job loss, the size and persistence of the loss is always sizable. Broadly speaking, the costs are smallest for men between the ages of 31–40 (7.7 percent) and largest for men above 50 (24 percent) and the mean cost of job loss is slightly smaller (10.9 percent) for women but still sizable. Kletzer and Fairlie (2003) examine the National Longitudinal Survey of Youth (NLSY) and find that the long-term earnings losses of male young adults with low tenure are approximately 10 percent, similar to the decline for high-tenure workers found by Davis and von Wachter (2011).

Recent work by Couch and Placzek (2010) also establishes that the cost of job loss is not isolated to a single industry. Using administrative data from Connecticut during the 1990s and 2000s, they find that earnings losses six years after a displacement event are 13-15 percent of pre-displacement earnings no matter which industry the job displacement event happens in. The smallest losses occurred in the education and health sector and the largest losses occurred in the financial sector with the losses in the manufacturing sector in the middle. Phelan (2014) also finds large effects in most sectors. These results suggest that earnings losses are not solely an artifact of depreciating human capital that is only valued in a declining industry, but rather an inherent consequence of job loss in both growing and shrinking sectors. Recent work by Guvenen, Ozkan, and Song (2014), which uses income information from Social Security records, is also consistent with that consumption. They find that both within and across industries, the probability of an extreme decline in income rises significantly in recessions, which would naturally occur if layoffs themselves cause the income declines.

3 Does the Federal Reserve react to the layoff rate?

Typical estimates of monetary policy rules do not include the layoff rate. This section shows that the Federal Reserve has in fact historically responded strongly to layoffs.

3.1 Data

The policy interest rate that we study is the target Fed Funds rate. The analysis is conducted at the quarterly frequency to allow us to include data on the output gap. All data is obtained from the Federal Reserve Bank of St. Louis’s FRED database.

We consider three measures of inflation: the personal consumption expenditures (PCE) deflator, the core PCE deflator that excludes food and energy purchases, and the gross domestic product (GDP) deflator. We measure output using real GDP either in growth rates or detrended using a Hodrick–Prescott (HP) filter with a smoothing parameter of 1600. We use the CBO’s definition of potential output to construct our baseline measure of the output gap. We also examine the unemployment rate, the HP-filtered unemployment rate (using a smoothing parameter of 12800) and the change in the unemployment rate as measures of slack in the economy. The HP filter has a number of potential drawbacks, not least that it is forward-looking, so it is included primarily to check robustness.

Layoffs are measured using weekly initial claims for unemployment benefits averaged over each quarter. Initial claims are then used either in their raw level (which has only a small trend over the sample), scaled by total employment, or smoothed using an exponentially weighted moving average with a decay rate of 0.025 (results using the HP filter are similar).

All of the variables are detrended on the 1967–2014 sample, but the analysis of the determinants of interest rates below uses the period 1982–2008 to avoid endpoint problems, major changes in monetary policy prior to the tenure of Paul Volcker, and the zero lower bound. That said, the results are highly similar in both the full (1967–2014) and post-Volcker (1982–2014) sample.

3.2 Analysis

Table 1 reports pairwise correlations between innovations in the target Fed Funds rate and innovations in the various explanatory variables (all obtained from AR(1) models). The four largest correlations are for various measures of layoffs, with detrended initial claims performing best with a correlation of -0.60. Of the other variables, the strongest is the change in the unemployment rate, with a correlation of -0.55 (and the change in the unemployment is obviously mechanically closely related to the layoff rate). The other measures of unemployment and output have substantially smaller correlations, almost all below 0.4 in absolute value. While the pairwise correlations do not represent a fully specified policy rule, they provide a simple first indication that the layoff rate is a major determinant of changes in interest rates.

Table 2 reports estimates of both backward- and forward-looking monetary policy rules. The top panel of table 2 reports results of regressions of the Fed Funds rate on its own lag and the various explanatory variables. Defining the layoff rate as detrended initial claims, in 9 of 10 specifications, layoffs have the highest t-statistic, indicating that they have the highest marginal explanatory power of any of the variables excluding the lagged fed funds rate (in the tenth case the t-stat is essentially identical to that for the change in the unemployment rate, which, again, is nearly a measure of

layoff activity). The t-statistics do not account for autocorrelation in the residuals; we use them here as measures of explanatory power – since they are monotonically related to the marginal R^2 of each variable – rather than as indicators of statistical significance (but using Newey–West (1987) standard errors does not change the conclusions). So, consistent with table 1, layoffs again appear to be most relevant for driving monetary policy out of all the variables we examine.

We next estimate forward looking policy rules of the form

$$r_t^{\$} = (1-\rho_r) \left(\alpha + \beta E_t \left[k^{-1} \sum_{m=1}^k \pi_{t+m} \right] + \gamma E_t \left[q^{-1} \sum_{m=1}^q \hat{y}_{t+m} \right] + \delta E_t \left[q^{-1} \sum_{m=1}^q \text{layoffs}_{t+m} \right] \right) + \rho_r r_{t-1}^{\$} + \varepsilon_t \quad (1)$$

where $r^{\$}$ is the Fed Funds rate, π inflation, \hat{y} the output gap, and the parameters k and q determine the forecast horizon used in setting interest rates

Forward looking rules have a strong theoretical appeal since optimal policy functions are forward-looking in standard models (Batini and Haldane (1999)). The regressions reported here closely follow the approach outlined in Clarida, Gali and Gertler (2000), and our baseline policy function is identical to theirs. The innovation here is that expected initial claims are also included as an explanatory variable. We estimate policy rules by GMM using as instruments four lags each of inflation, the output gap, the federal funds rate, the 10/1-year Treasury yield spread, and commodity price inflation.⁶

The bottom panel of table 2 reports results for five different sets of policy rules, differing in the inclusion of layoffs and the forecast horizons (k and q). For this set of models, the explanatory variables are all standardized so that the coefficients can be directly compared, in addition to the t-statistics. Across the values for the forecast horizon, whenever expected layoffs are included they again always have the highest t-statistic other than the lagged policy rate. The bottom rows show that in four of the five specifications, the coefficient on initial claims is statistically significantly higher than that on the output gap. In general, the coefficient on initial claims is around 0.8 – a unit standard deviation increase in initial claims leads to a 0.8 percentage point reduction in the Fed Funds rate. In the final column, where the forecast horizon is longest, the coefficient and power both decline, making it more difficult to distinguish from the coefficient on output.

Perhaps even more surprisingly, the coefficient on layoffs is always larger than the coefficient on inflation, and weakly significantly so in three of the five cases. Overall, across the five specifications, the estimated weight on expected layoffs is statistically and economically significantly greater than that on either inflation and the output gap. Similar to the raw correlations and the backward-looking estimates, then, in the forward-looking case the layoff rate remains dominant.

So across a range of specifications, involving three different measures of layoffs, a range of measures of inflation and the output gap, and both backward and forward looking policy rules, the layoff rate has been the dominant driver of changes in interest rates. The Federal Reserve thus

⁶This is the exact instrument set used by Clarida, Gali and Gertler (2000).

appears to follow a rule that is tightly linked to the layoff rate.

To help understand the source of the variation that explains why detrended initial claims perform so well in explaining the Fed Funds rate, figure 1 plots z-scores of detrended initial claims, the unemployment rate, and the residual of the Fed Funds rate on its own lag and PCE inflation. The scale for the Fed Funds rate innovation is reversed to help make the correlations clearer. Over the sample, the correlation of detrended initial claims with the Fed Funds residual is -0.61, compared to -0.27 for the unemployment rate. That means that the initial claims rate absorbs five times as much of the variation in the Fed Funds rate as the unemployment rate does.

Part of the difference in the correlations is due to the fact that initial claims rise faster at the beginning of recessions and fall back to their mean more quickly than the unemployment rate. The 1991 recession displays that behavior most clearly, with the Fed Funds rate and initial claims following very similar paths, and both leading the unemployment rate. Similarly, in 2008 initial claims rise much faster than the unemployment rate, consistent with the very large declines in the Fed Funds rate. But that behavior occurs more generally. The top panel of figure 5 plots initial claims and unemployment since 1967, while the middle panel plots their autocorrelations. In numerous recessions, initial claims decline faster than unemployment, and their autocorrelations are smaller at all lags in the figure than those of the unemployment rate. The model developed in the next section will be able to match that behavior, helping to rationalize the strong performance of initial claims in explaining the Fed Funds rate.

4 Model

This section develops a New Keynesian model of the business cycle augmented with heterogeneity in income and consumption driven by job loss. The fact that income risk associated with layoffs is uninsurable gives a role for the central bank in the model to play in reducing risk and thus raising welfare.

4.1 Production/employment sector

There is a competitive sector of the economy that uses labor to produce undifferentiated intermediate goods. It is the only part of the economy that uses labor.

4.1.1 Labor demand

Intermediate-good producing firms are indexed by m . The total number of people employed by firm m follows the law of motion

$$N_{m,t} = (1 - q_t) N_{m,t-1} - F_{m,t} + H_{m,t} \tag{2}$$

where $F_{m,t} \geq 0$ is firm m 's firing in period t , $H_{m,t} \geq 0$ is hiring, and q_t is the quit rate. Each firm faces the same quit rate, losing a fraction q_t of its workers at the beginning of the period. q_t will be

determined endogenously, but it is exogenous to the decisions of the individual firms. Firms that desire to expand will have $H_{m,t} > 0$ and $F_{m,t} = 0$, while those that want to contract (by more than what is induced by quits) will have $H_{m,t} = 0$ and $F_{m,t} > 0$.

The workers in each firm have average human capital of K_t^E .⁷ The firms produce an identical output which sells at price S_t . They have no intertemporal decisions and simply maximize revenue net of costs in each period,

$$\max_{N_{m,t}} S_t A_t A_{m,t} (K_t^E N_{m,t})^\gamma - W_t K_t^E N_{m,t} \quad (3)$$

where A_t is the aggregate level of technology, $A_{m,t}$ is firm m 's level of technology, and W_t is the wage paid per unit of human capital on date t . Firms have decreasing returns to scale determined by the parameter γ .

$A_{m,t}$ is a martingale with log-Normal innovations,⁸

$$\log(A_{m,t}/A_{m,t-1}) \sim N\left(-\frac{1}{2} \frac{1}{1-\gamma} \sigma_a^2, \sigma_a^2\right) \quad (4)$$

Denoting logs of variables by lower-case letters and using Δ to denote the first-difference operator, optimization by the firms implies that log labor demand follows

$$\exp(\Delta n_{m,t}) = \exp\left(\frac{1}{1-\gamma} (-\Delta w_t + \Delta s_t + \Delta a_t + \Delta a_{m,t}) - \Delta k_t^E\right) \quad (5)$$

Firms hire if the optimal change in employment is greater than the decline induced by quits. That is,

$$\exp\left(\frac{1}{1-\gamma} (-\Delta w_t + \Delta s_t + \Delta a_t + \Delta a_{m,t}) - \Delta k_t^E\right) > 1 - q_t \implies H_{m,t} > 0 \text{ and } F_{m,t} = 0 \quad (6)$$

$$\exp\left(\frac{1}{1-\gamma} (-\Delta w_t + \Delta s_t + \Delta a_t + \Delta a_{m,t}) - \Delta k_t^E\right) \leq 1 - q_t \implies H_{m,t} = 0 \text{ and } F_{m,t} \geq 0 \quad (7)$$

Firms that receive sufficiently negative shocks to their productivity fire workers, while firms that receive positive shocks hire workers. There is thus both hiring and firing in all periods. Aggregate hiring and firing in the economy are calculated by integrating over the distribution of $\Delta a_{m,t}$.

⁷The model of flows technically implies that different firms will have workers with different average levels of human capital given their history of hiring and firing. Accounting for that aspect of heterogeneity renders the model intractable. The homogenization of average human capital across firms could potentially come through unmodeled churn in employment.

⁸As a technical matter, this law of motion for productivity implies that the distribution of firm-specific productivity is undefined (or non-stationary). To account for this issue, one could allow firms to die with some fixed probability δ_F , being replaced by new firms. Our model represents the limit as $\delta_F \rightarrow 0$.

4.1.2 Quits

If the quit rate in the model were constant, then in any period when employment rose – even if it was still below its steady-state – firing would be below average. In that case, periods of high firing rates would always be followed by periods with low firing rates, meaning that their net effect on utility would be close to zero.

In the data, though, firing is countercyclical – it is high when output and employment are low, even if they are rising – and its impulse responses do not sum to zero. So periods when layoff rates are above average are not deterministically followed by periods when they are below average, which is what allows variation in the layoff rate to reduce utility in our analysis below. That behavior in the data is due to the fact that the quit rate is procyclical – when aggregate employment is below average, fewer workers quit, so firms must fire more workers just to hold employment constant. In order to accurately capture the behavior of layoffs, then, (and in order to generate welfare costs for business cycles) it is important to account for variation in the quit rate.

A parsimonious device for generating procyclical quits is to assume that a fraction λ of current employees are willing to move to a new job and that firms hire workers proportionally from the two sources of available workers, the unemployed and those willing to switch jobs (those unwilling to switch to a new job might have nonpecuniary motives, might face a large cost of switching jobs, or might be paid above the market wage in their current position (e.g. Harris and Holmstrom (1982))). Aggregate quits, Q_t , and the quit rate, q_t , are then

$$Q_t = H_t \frac{\lambda N_{t-1}}{(1 - N_{t-1}) + \lambda N_{t-1}} \quad (8)$$

$$q_t \equiv Q_t/N_{t-1} = \frac{\lambda H_t}{1 - (1 - \lambda) N_{t-1}} \quad (9)$$

where H_t denotes aggregate hiring, N_t aggregate employment, and aggregate labor supply is normalized to 1. In periods when hiring is high or employment is high, the quit rate will be high, making it procyclical.

4.1.3 Wages and disequilibrium

All people in the economy inelastically supply a unit of labor, so if wages were perfectly flexible then they would all be employed at all times. Instead, in an extreme version of Blanchard and Gali (2010) we simply assume that real wages, W_t , are constant. Employment on each date is determined by labor demand. The level of the wage is a free parameter that we use to fix the average level of employment.

4.2 Price setting

There is a set of monopolistically competitive firms that, employing no labor, buy the intermediate good and differentiate it. They then sell their output to competitive final good aggregators. The

final good aggregators have sticky prices that they reset at random intervals as in Calvo (1983). Standard results then generate a New Keynesian Phillips curve (see Gali (2015)); see appendix A.2 for details of the implementation used here.

4.3 Consumers

The model does not use a representative agent. Rather, agents are distinguished by their employment status and level of human capital. Agents can smooth effective consumption by accumulating and drawing down human capital. The model is specified with a form of linear homogeneity such that there are two representative agents, one employed and the other unemployed.

4.3.1 Human capital, earnings, and consumption

Agents have Epstein–Zin (1991) preferences over total consumption, $C_{i,t}$, where i indexes individuals:

$$U_{i,t} = \left\{ (1 - \beta) C_{i,t}^{1-\rho} + \beta E_t \left[U_{i,t+1}^{1-\alpha} \right]^{\frac{1-\rho}{1-\alpha}} \right\}^{\frac{1}{1-\rho}} \quad (10)$$

E_t is the expectation operator conditional on information available on date t . The assumption of Epstein–Zin preferences allows us to separate risk aversion and intertemporal substitution in the calibration, but the main results are not sensitive to the choice of preferences.

Total consumption, $C_{i,t}$, is the sum of market consumption and a component due to the creation or destruction of human capital. In addition to consuming market goods, a person can also choose to study or play video games. Studying is unpleasant and reduces current utility, so it contributes negatively to the $C_{i,t}$ that enters preferences, but it also increases human capital. Playing video games, on the other hand, is entertaining and contributes positively to $C_{i,t}$, but reduces human capital. Agent i 's human capital therefore follows

$$K_{i,t+1} = K_{i,t} \exp \left(- \left(1 - 1_{i,t}^E \right) d \right) (1 - \delta_K) + G_{i,t} \quad (11)$$

where δ_K is the depreciation rate of human capital, $1_{i,t}^E$ is an indicator variable for whether agent i is employed on date t , and the parameter d induces a decline in during unemployment beyond that caused by playing video games. $G_{i,t}$ is the agent's investment in human capital. Studying is represented by a positive value of $G_{i,t}$, while playing video games corresponds to a negative value.

The direct or mathematical way to think about $G_{i,t}$ is simply to say that agents have a linear technology that allows them to convert human capital into units of consumption. Agents are hand-to-mouth in the sense that they spend all wage earnings on market consumption, but they can smooth effective consumption through $G_{i,t}$. The total consumption of the employed and unemployed is

$$\text{Employed:} \quad C_{i,t+j} = (1 - \tau_t) \tilde{W}_{t+j} K_{i,t+j} - G_{i,t+j} \quad (12)$$

$$\text{Unemployed:} \quad C_{i,t+j} = b \tilde{W}_{t+j} K_{i,t+j} - G_{i,t+j} \quad (13)$$

where \tilde{W}_t is the wage per unit of capital that agents earn. The parameter b represents the replacement rate of unemployment benefits, and τ_t is the tax rate, which is set in each period to finance current unemployment benefits.⁹

Earned income \tilde{W} , is not the same as the wage faced by firms, W_t . Since W_t is not equal to the marginal product of labor, firms have profits or losses in each period. For the same of simplicity, we assume that those profits and losses are shared among workers in proportion to their human capital, so that the effective wage that workers earn – their market wage plus profits or losses – is

$$\tilde{W}_t = Y_t / (K_t^E N_t) \quad (14)$$

This ensures that the resource constraint is satisfied.

4.3.2 Optimization

The optimization problem of the consumers is

$$\max U_{i,t} + E_t \sum_{j=0}^{\infty} \lambda_{i,t+j} \left[\begin{array}{l} \left((1_{i,t+j}^E (1 - \tau_t) + (1 - 1_{i,t+j}^E) b) W_{i,t+j} K_{i,t+j} - C_{i,t+j} - G_{i,t+j} \right) \\ - \kappa_{i,t+j} \left(K_{i,t+j+1} - G_{i,t+j} + K_{i,t+j} \exp \left(- (1 - 1_{i,t+j}^E) d \right) (1 - \delta_K) \right) \end{array} \right] \quad (15)$$

where $\lambda_{i,t}$ and $\kappa_{i,t}$ are Lagrange multipliers. Appendix A.3 derives the various first-order conditions. The fact that the model is linearly homogenous implies that, rather than having to track the optimization for every individual agent, we can solve just two problems, one for the employed types and one for the unemployed. Specifically, the Euler equations for the consumption of the two types are

$$1 = E_t [M_{t+1}^{EE} P_{t+1}^{EE} ((1 - \tau_{t+1}) W_{t+1} + 1 - \delta_K) + M_{t+1}^{EU} P_{t+1}^{EU} (bW_{t+1} + \exp(-d) (1 - \delta_K))] \quad (16)$$

$$1 = E_t [M_{t+1}^{UE} P_{t+1}^{UE} ((1 - \tau_{t+1}) W_{t+1} + 1 - \delta_K) + M_{t+1}^{UU} P_{t+1}^{UU} (bW_{t+1} + \exp(-d) (1 - \delta_K))] \quad (17)$$

where M_{t+1}^{nm} represents the stochastic discount factor for an agent moving from state n to state m and P_{t+1}^{nm} is the associated probability.¹⁰ A marginal reduction in consumption on date t increases human capital by the same amount on date $t + 1$. If the person is employed at $t + 1$, that raises earnings by $(1 - \tau_{t+1}) W_{t+1}$, while it raises earnings by bW_{t+1} if they are unemployed. The human capital depreciates by the factor $(1 - \delta_K)$ in either state, with an additional $\exp(-d)$ if the person

⁹That is, $\tau_t = bK_t^U (1 - N_t) (K_t^E N_t)^{-1}$, where K_t^U is the average human capital of the unemployed on date t .

¹⁰Specifically, $P_{t+1}^{EE} \equiv \frac{N_t - F_{t+1}}{N_t}$, $P_{t+1}^{EU} \equiv \frac{F_{t+1}}{N_t}$, $P_{t+1}^{UE} \equiv \frac{H_{t+1} - Q_{t+1}}{1 - N_t}$, and $P_{t+1}^{UU} \equiv \frac{1 - N_t - H_{t+1} + Q_{t+1}}{1 - N_t}$.

is unemployed. That all represents the return to investing in human capital, which is the only method through which agents can shift (effective) consumption across dates.

The stochastic discount factor for the employed agents is

$$M_{t+1}^{E\cdot} = \beta \frac{U_{t+1}^{(\rho-\alpha)}}{(E_t^E U_{t+1}^{1-\alpha})^{\frac{\rho-\alpha}{1-\alpha}}} \left(\frac{C_{t+1}}{C_t^E} \right)^{-\rho} (1 - \delta_K + G_t^E)^{-\rho} \quad (18)$$

where variables with superscripts are scaled values – e.g. $U_t^E \equiv U_{i,t}/K_{i,t}$ – and $E_t^E U_{t+1}^{1-\alpha}$ is the expectation of scaled utility in period $t+1$ conditional on being employed on date t . The stochastic discount factor for the unemployed takes a similar form. The second and third terms represent the pure contribution of consumption growth between periods t and $t+1$. The first term is the contribution from the use of Epstein–Zin preferences (which disappears when $\rho = \alpha$). For $\alpha > \rho$, as is common, it says that states of the world in which continuation utility is above average have low marginal utility. In that case marginal utility is driven not just by current consumption but also on its expected future path and volatility.

4.3.3 Interest rates

Under the assumption that consumers face some sort of wedge or friction when borrowing, in equilibrium the agents who desire to save more will be marginal in the market for riskless bonds. At our benchmark calibration (and, one would suspect, for any reasonable calibration), the employed agents have a greater desire to save than the unemployed. We therefore use their Euler equation to price riskless bonds, which means that they are indifferent to holding those bonds (since they are in zero net supply). Specifically,

$$1 = E_t \left[\exp(\psi_t) \frac{R_t}{\Pi_{t+1}} (M_{t+1}^{EE} P_{t+1}^{EE} + M_{t+1}^{EU} P_{t+1}^{EU}) \right] \quad (19)$$

where Π_{t+1} is the gross inflation rate and R_t is the gross nominal interest rate between dates t and $t+1$. ψ_t is a demand shock or shock to the IS curve – it shifts the equilibrium interest rate for a given consumption path. Making the demand shock part of preferences, as in, e.g. Albuquerque et al. (2016), can have substantial effects on welfare. Rather than interpret ψ_t as a shock to the rate of time preference, it is simply a wedge between the intertemporal marginal rate of substitution and the interest rate.

4.3.4 Aggregate human capital

In order to ensure that human capital is stationary, at the end of each period, a fraction δ of agents die and are replaced by new agents with the same employment state but with human capital set to \bar{K} (which we assume is simply steady-state mean human capital). Average human capital of those who are employed and unemployed, K_t^E and K_t^U , respectively, therefore follows

$$N_{t+1}K_{t+1}^E = (N_t - F_{t+1})N_t((1 - \delta)K_t^E(1 - \delta_K + g_t^E) + \delta\bar{K}) \quad (20)$$

$$+ (H_{t+1} - Q_{t+1})(1 - N_t)((1 - \delta)K_t^U(\exp(-d)(1 - \delta_K) + G_t^U) + \delta\bar{K}) \quad (21)$$

$$(1 - N_{t+1})K_{t+1}^U = (1 - N_t - H_{t+1} + Q_{t+1})N_t((1 - \delta)K_t^U(\exp(-d)(1 - \delta_K) + G_t^U) + \delta\bar{K}) \quad (22)$$

$$+ (1 - N_t)F_{t+1}((1 - \delta)K_t^E(1 - \delta_K + g_t^E) + \delta\bar{K}) \quad (23)$$

4.4 Monetary policy

There is a central bank that sets the nominal interest rate following a linear rule that allows it to potentially respond to the firing rate,¹¹

$$R_t - \bar{R} = \rho_r(R_{t-1} - \bar{R}) + (1 - \rho_r)\left(\phi_\pi(\pi_t - \bar{\pi}) + \phi_y(y_t - \bar{y}) - \phi_{F/N}\left(\frac{F_t}{N_{t-1}} - \frac{\bar{F}}{\bar{N}}\right)\right) \quad (24)$$

where bars over variables represent values in the steady state with no aggregate risk. The ability of the policy rule to respond to the layoff rate is consistent with the empirical results above.

4.5 Exogenous processes

There are two exogenous shock processes in the model: the demand shock, ψ_t , and a markup shock, denoted u_t , which increases inflation for a given level of output (i.e. a shock to the Phillips curve; see section A.2). They follow

$$\psi_t = \rho_\psi\psi_{t-1} + \varepsilon_{\psi,t}, \quad \varepsilon_{\psi,t} \sim N(0, \sigma_\psi^2) \quad (25)$$

$$u_t = \rho_u u_{t-1} + \varepsilon_{u,t}, \quad \varepsilon_{u,t} \sim N(0, \sigma_u^2) \quad (26)$$

The markup shock is included in the model to generate the traditional central banking trade-off between controlling inflation and output. The demand shock drives inflation and output in the same direction and it can be fully corrected by monetary policy. It is included to help reduce the negative correlation between output and inflation in the model.

4.6 First-order approximation and precautionary saving effects

We solve the model with a standard first-order approximation around the non-stochastic steady-state. Non-stochastic here, however, only refers to aggregate risk. There is still cross-sectional

¹¹Given that layoffs are the uninsurable risk in the economy, it would be natural to manage them through a direct tax on firing. Studying the implications of such a tax is an interesting path for future work, but we focus here just on traditional monetary policy.

variation. That fact means that the first-order approximation will capture effects of cross-sectional risk on interest rates. To see why, note that the first-order approximation to the Euler equation for the nominal interest rate is

$$\frac{\hat{R}_t}{\bar{R}} - E_t \frac{\hat{\Pi}_{t+1}}{\bar{\Pi}} \approx \frac{\bar{R}}{\bar{\Pi}} \left((\bar{M}^{EE} - \bar{M}^{EU}) E_t \hat{P}_{t+1}^{EU} + E_t \left(\hat{M}_{t+1}^{EE} - \hat{M}_{t+1}^{EU} \right) \bar{P}^{EU} - E_t \hat{M}_{t+1}^{EE} \right) - \psi_t \quad (27)$$

where bars represent steady-state values and circumflexes denote deviations from steady-state.

The linearization shows that the nominal interest rate depends on variation in idiosyncratic risk through two channels. First, when the probability of job loss is above average – $E_t \hat{P}_{t+1}^{EU}$ is high – then interest rates will tend to be low, since marginal utility is relatively high when an agent is unemployed, $(\bar{M}^{EE} - \bar{M}^{EU} < 0)$. Second, when the pain of job loss is above average – as measured by the increase in marginal utility in unemployment, $\hat{M}_{t+1}^{EU} - \hat{M}_{t+1}^{EE}$, interest rates are lower. So it is the combination of the probability and pain of job loss that drive interest rates.

While the linearization separates the two effects, they are in fact multiplicative. If, as one would expect, the probability and cost of job loss are correlated, then the precautionary saving effect could be further magnified in a higher-order solution to the model. We find in unreported results that the quantitative behavior of the model is nearly identical under first- and third-order approximations. So at least for the calibration used here, the higher-order interactions are not quantitatively important.

That said, to calculate welfare, we use a fully nonlinear approximation to the utility function, discussed in section 6.1.

5 Model behavior

5.1 Calibration

Table 3 reports the benchmark calibration of the model. The majority of the parameters, e.g. those related to price stickiness and monopolistic competition, have been extensively discussed in the literature [**where do they come from, exactly?**]. We set the persistence of the markup and demand shocks to 0.9 to generate business-cycle frequency fluctuations. The relative magnitudes of the shock volatilities are chosen so that the variances of output and inflation are driven approximately equally by the two shocks (see Gali, Smets, and Wouters (2012) for estimates of a variance decomposition broadly consistent with that view). The shocks are then scaled in order to match as closely as possible the volatilities of unemployment and inflation. The parameter ϕ_π is set to 2.03 to match evidence in Stock and Watson (XXXX). ϕ_y is set to 0.25 in order to match the relative volatilities of the unemployment rate and inflation.

Table 4 reports a range of unconditional moments of the model. The standard deviation of the unemployment rate is somewhat below the empirical value, while the standard deviation of inflation is somewhat above. Those facts will tend to tilt the model in the direction of implying that there is a larger benefit from stabilizing inflation, since welfare is generally approximately linear in the

variances of unemployment and inflation.

The model implies relatively high smoothness in output, meaning that the volatility of output growth is much lower than in the data. Interest rates, on the other hand, are more volatile than what is observed empirically.

Figures 2 and 3 plot impulse responses of major aggregates to the demand shock and the markup (supply) shock. Their effects on output and inflation are what one would typically expect – the demand shock drives output and inflation in the same direction, while the supply shock drives them in opposite directions.

5.1.1 Job loss and idiosyncratic risk

The key drivers in the model of the welfare cost of business cycles and the levels and variation in the probability of job loss and the size of the income and consumption declines following job loss. The most important question about the model, therefore, is how well it matches the empirical behavior of those variables. This subsection and the next examine the model’s predictions for idiosyncratic risk and compare them to the data.

The bottom section of table 4 reports statistics from the simulations regarding job loss rates. The average probability of a job loss is 1 percent per quarter, with the rate varying between 0.7 and 1.3 percent between the 25th and 75th percentiles. These values agree closely with the calibration and evidence discussed in Krebs (2007).

The average probability of transitioning from unemployment to employment – the hiring rate – is 14.9 percent per quarter, implying that the average duration of an unemployment spell is 6.73 quarters. That value is well above the historical mean from the BLS of 1.23 quarters. At the same time, though, the proportion of workers who are laid off in each quarter is much smaller than the empirical average from the JOLTS data of 4 percent per quarter.

The calibration here is meant to capture layoffs of relatively high tenure workers that can be expected to lead to large income losses. For that reason, we assume that most of the churn observed in real-world labor markets – three quarters of the layoffs in each quarter – represents what might be thought of as costless job loss. It is the relatively small number of highly painful layoffs that we focus on. Furthermore, the 7-quarter period of unemployment in the model can be interpreted as partly proxying for the time that it takes for workers not just to find any job, but to find one that approximates the quality of the job they lost (e.g. Krolikowski (2017), Jung and Kuhn (2018), Jarosch (2018)). The model therefore should be interpreted as focusing on relatively high-tenure workers, who have jobs with lower rates of turnover (as in the previous citations) and which represent the sample typically used to measure the cost of job loss (e.g. Jacobson, Lalonde, and Sullivan (1993), etc).

An important difference between layoffs and the unemployment rate in historical data is that the layoff rate better isolates the beginning of recessions – it jumps initially, then rapidly reverts, whereas unemployment takes a longer time to recover. As discussed above, that difference is

relevant in their relative fit to interest rates. The top panel of figure 5 plots the time series of the layoff rate, measured as initial unemployment claims scaled by total employment, compared to the unemployment rate. The two series rise almost simultaneously at the beginning of recessions, but the layoff rate declines more quickly. The middle panel of figure 5 plots the first eight quarterly autocorrelations of the two series. The autocorrelations of the layoff rate decline much more rapidly. At the one-year lag, the autocorrelation of the unemployment rate is 0.70, compared to only 0.42 for the layoff rate. The bottom panel plots the corresponding autocorrelations in the model. The model matches the data qualitatively in the sense that unemployment is more persistent at all lags than layoffs (measured in the model as F_t/N_t). The spread is wider than in the data, though, since the model’s prediction for firings is that they have essentially zero autocorrelation. There are calibrations of the model that generate more realistic smoothness in firings (which was not a target moment for the calibration) but at the cost of deviating from the data in other regards.

The model also generates realistic behavior for quits. The model is calibrated to match the average 2-percent quit rate from the JOLTS data. The 25th and 75th percentiles are 1.56 and 2.44 percent in the model, compared to 1.68 and 2.1 percent in the data. Empirically, quits peak just prior to the recessions in 2001 and 2008 and have troughs two years later. In the model, quits are also procyclical, with a correlation of 0.68 with employment. The JOLTS data only begins at the end of 2000, but over the 17-year sample that is available, the model agrees well with the level, volatility, and cyclical nature of quits in the data.

5.1.2 Income and consumption losses

The top panel of figure 4 plots the average path of earnings for workers who lose jobs in a given period compared to those who do not (though note that those in the job-keeping group can potentially lose their jobs in subsequent periods). In the first period, earnings fall by 100 percent, since the job losers are unemployed. Subsequently, earnings recover as the proportion of people who have new jobs rises. In the long-run, a job loss is associated with an average decline in earnings of 15.5 percent.

The two lines split the sample based on expansions versus recessions, to see how the risk associated with job loss changes over time. In expansions – defined as periods when employment is above its 15th percentile, the long-run loss is 15.3 percent, while it is 16.5 percent in recessions. So not only is the probability of job loss higher in recessions, but the magnitude of the income decline following job loss is also larger.

The 15.5-percent average decline in income is driven by a decline in human capital of the same amount. Human capital declines during unemployment both for exogenous reasons – the $\exp(-d)$ term – and also because unemployed agents draw down human capital in order to support consumption – they play video games. At the steady-state, the relative loss in human capital during unemployment is 2.7 percent per quarter. Multiplying that value by the average unemployment duration, 6.73 quarters, approximately yields the total long-run income loss.

The middle panel of figure 4 is similar to the top panel, but instead of splitting the sample into expansions and recessions, it splits based on unemployment duration. The blue line plots the average path of earnings relative to the job keepers for people who find a new job within four quarters of being fired. For those agents, the long-run loss is only 5.5 percent – one third as large as the unconditional mean. As the duration of unemployment grows, so does the long-run loss. For those who are rehired in year 4, the total loss is double the mean, at 30.7 percent.

The middle panel of figure 4 therefore shows that idiosyncratic risk appears in the model not just because job loss itself is uninsurable, but also because the duration is uninsurable. Formally, at the nonstochastic steady-state, the long-run income loss is geometrically distributed. Income growth in the model is thus highly skewed to the left. The raw skewness of five-year income growth is -3.13, while the Kelley skewness is -13.53 percent. The latter value is highly similar to the Kelley skewness reported by Guvenen et al. (2014) of 14–30 percent.

Similar to Guvenen et al. (2017) and Storesletten, Telmer, and Yaron (2004), and consistent with figure 4, the standard deviation and skewness of five-year wage growth both rise in recessions. Kelley skewness rises from -13.50 percent to -16.30 percent.

Finally, the bottom panel of figure 4 plots the average path of consumption following a layoff. Since the agents’ optimization problem is linearly homogenous with respect to human capital, total consumption (i.e. the sum of market consumption and negative human capital investment) is also proportional to human capital in each state. During unemployment spells, human capital falls as agents consume at home, causing a decline in subsequent consumption. The long-run decline in consumption is then the same as the long-run decline in human capital and hence income.

The 15-percent decline in income and consumption that we obtain are highly similar to the calibration of Krebs (2007), and is consistent with reported long-run wage declines in the literature discussed in section 2. However, table 4 and the top panel of figure 4 show that the variation in the loss is small here relative to the calibration in Krebs (2007) and the empirical literature. Surveying the literature, Krebs (2007) argues that a reasonable range for income losses between expansions and recessions is 9–21 percent, as opposed to our 17–18 percent.

5.1.3 Precautionary saving effects

As discussed above, the model accounts for time-varying risk in the first-order approximation. One effect of variation in risk is variation in the desire of agents for riskless saving, which affects interest rates. To quantify the precautionary saving effect in the model, we examine the difference between the real interest rate that prevails under the employed agents’ Euler equation to what would be obtained if those agents did not face idiosyncratic risk. Figures 2 and 3 plot the response of the precautionary saving effect on interest rates to the two shocks. In both cases, the precautionary saving effect moves in the opposite direction of employment and layoffs – when layoffs, and hence idiosyncratic risk, rise, agents desire to save more, driving interest rates down. For both shocks, though, the magnitude of that effect is very small – at most a tenth of a basis point. The standard

deviation of the annualized real interest rate is 226 basis points, while the standard deviation of the precautionary saving effect is only 0.4 basis points. In a third-order approximation those numbers are nearly identical.

Precautionary saving effects in the calibration used here are therefore quantitatively trivial. That does not mean, though, that they cannot matter in reality. As discussed above, table 4 shows that even though the model generates a large average amount of idiosyncratic risk, the variation over time is much smaller than what has been measured empirically. One way to make precautionary saving effects larger would be to modify the model to generate a more volatile welfare cost for layoffs.

6 The welfare cost of business cycles

This section examines the model’s implications for the welfare cost of business cycles.

6.1 Calculating welfare

Calculating welfare under Epstein–Zin preferences with models of this class is not trivial. Appendix A.6 describes the details of our method. The aim is to calculate as accurately as possible the utility of an agent who receives the consumption flow implied by the first-order approximation to the model. That is, we take the approximated dynamics as though they are true (since it is the approximated dynamics that we match to the data) and then calculate utility directly. The typical method in the literature is to approximate utility using a low-order polynomial – often just a quadratic function. Here, instead, the calculation is fully nonlinear.

Utility under Epstein–Zin preferences requires calculating conditional expectations on each date, but those expectations are not available in closed form for the present model. We construct date- t expectations of date- $t+1$ outcomes by regressing the outcomes a full set of interactions of the state variables in the linear model up to the fourth order (i.e. $\{x_{1,t}, x_{2,t}, \dots, x_{1,t}^2, x_{1,t}x_{2,t}, \dots\}$). If this fourth-order polynomial regression accurately approximates the true expectation, then the method correctly calculates utility. An alternative interpretation of the calculation is that it yields utility for a boundedly rational agent who cannot calculate true statistical expectations and instead uses polynomial regressions.

The welfare calculation method does not directly account for the higher-order costs of inflation. To do that, we calculate the path of price dispersion (which is what drives the cost of inflation in much of the literature, e.g. Erceg, Henderson, and Levin (2000)) without resorting to the first-order approximation. Following Yun (1996), output in the economy is,

$$Y_t = A_t N_t^\gamma / D_t \tag{28}$$

$$\text{where } D_t = (1 - \xi) \left(\frac{1 - \xi \Pi_t^{\varepsilon-1}}{1 - \xi} \right)^{\frac{\varepsilon}{\varepsilon-1}} + \xi \Pi_t^\varepsilon D_{t-1}, \tag{29}$$

$1-\xi$ is the Calvo probability that an intermediate good producing firm is able to change its price, ε is the elasticity of substitution across intermediates, and Π_t is gross price inflation. When inflation is more volatile, D_t is higher on average, thus reducing average output. Since all output is consumed, a one percent increase in D_t is associated with a one percent decline in market consumption. Furthermore, volatility in D_t will affect the volatility of output.

D_t is calculated by iterating on equation (28) using the history of inflation from the first-order approximation. That is, it is calculated exactly conditional on the first-order simulation for inflation. We account for price dispersion in calculating welfare by replacing C_t^E and C_t^U with C_t^E/D_t and C_t^U/D_t . Denoting the average level of utility when the effects of price dispersion are ignored (i.e. with C_t^E and C_t^U) by $W^{(1)}$ and the average level of utility when the effects of price dispersion are included (using C_t^E/D_t and C_t^U/D_t) by $W^{(2)}$, the percentage cost of price distortions is then $\log W^{(1)} - \log W^{(2)}$.

Erceg, Henderson, and Levin (2000), and others, measure the welfare cost of inflation alternatively based on a second-order approximation around the zero-inflation steady-state. In their case, the cost of inflation is $(\varepsilon/2)\theta(1-\theta)^{-2}var(\Pi_t)$, where θ is the probability of a firm adjusting prices in a given period and ε is the inverse elasticity of substitution across intermediates. We compare our results to those using the second-order approximation in the next section.

The calculations here differ somewhat from those used in previous work because we are not approximating the utility function (beyond the high-order polynomial step in the expectation). Instead, we calculate utility directly taking the dynamics as given. We will show, though, that the values we obtain are similar to those obtained in past work.

6.2 Welfare costs in the calibrated model – numerical calculations

As in Krebs (2007), the measure of the welfare gain from eliminating business cycles is the increase in utility when all aggregate variables are held at their steady-state values. At that point, though, there is still idiosyncratic risk. There are then two differences that arise from stabilizing business cycles – the aggregate component of consumption risk is eliminated, and variation in idiosyncratic consumption risk is eliminated.

Table 5 reports the various components of the welfare costs of business cycles in the model. The first column reports the losses at the baseline calibration discussed above. Overall, eliminating all aggregate variation – i.e. in both total output and also in cross-sectional risk – is equivalent to a permanent increase in consumption of 1.00 percent at our baseline calibration. That value is approximately one third smaller than what is reported in Krebs (2007), with the difference caused by the fact that variation in the cost of job loss over the business cycle is much smaller. However, it remains large relative to standard calculations in representative agent models. Using the formula and volatility estimate of Lucas (2003) with our risk aversion of 3 yields a cost of business cycles of only 0.16 percent of lifetime consumption, for example.

The second and third rows of table 5 decompose the total welfare cost into components coming

from inflation volatility (in driving price dispersion) and pure consumption risk. 95 percent of the total loss is due to consumption risk. This contrasts with results from representative-agent settings where inflation is often more important. Here, inflation has effects twenty times smaller than consumption risk.

Using the formula for the cost of inflation from Erceg, Henderson, and Levin (2000) based on a second-order approximation around the zero-inflation steady-state, we would have a welfare cost of inflation dispersion of 0.05 percent of lifetime consumption, which is nearly identical to the values reported in table 5.

6.3 Welfare costs – analytic approximation

Following an analysis similar to that of Krebs (2007), it is possible to derive an analytic approximation to welfare in the model. In the case of Epstein–Zin preferences with $\rho \neq \alpha$, utility is implicitly defined by a system that is not solvable in closed form. Denote utility per unit of human capital for employed workers by U^E and unemployed workers by U^U . Using bars to denote steady-state values, we have at the steady-state

$$(\bar{U}^E)^{1-\rho} = (1-\beta)\bar{C}^{E(1-\rho)} + \beta(1-\delta_K + \bar{G}^E)^{1-\rho} \left(\begin{array}{l} \bar{E} \left[P_{t+1}^{EE} \left(\frac{U_{t+1}^E}{\bar{U}^E} \right)^{1-\alpha} \right] \bar{U}^{E(1-\alpha)} \\ + \bar{E} \left[P_{t+1}^{EU} \left(\frac{U_{t+1}^U}{\bar{U}^U} \right)^{1-\alpha} \right] \bar{U}^{U(1-\alpha)} \end{array} \right)^{\frac{1-\rho}{1-\alpha}} \quad (30)$$

$$(\bar{U}^U)^{1-\rho} = (1-\beta)\bar{C}^{U(1-\rho)} + \beta \left(\begin{array}{l} \exp(1-d)(1-\delta_K) \\ + \bar{G}^U \end{array} \right)^{1-\rho} \left(\begin{array}{l} \bar{E} \left[P_{t+1}^{UE} \left(\frac{U_{t+1}^E}{\bar{U}^E} \right)^{1-\alpha} \right] \bar{U}^{E(1-\alpha)} \\ + \bar{E} \left[P_{t+1}^{UU} \left(\frac{U_{t+1}^U}{\bar{U}^U} \right)^{1-\alpha} \right] \bar{U}^{U(1-\alpha)} \end{array} \right)^{\frac{1-\rho}{1-\alpha}} \quad (31)$$

Welfare at the steady-state, as one would expect, depends on average consumption and human capital investment – the \bar{C} and \bar{G} terms. More interestingly, though, it depends on the risk that agents face, through the transition probabilities. These equations are useful for showing how those risks matter.

Given the equations (30) and (31), \bar{U}^E and \bar{U}^U are implicit functions of the four expectation terms, $\bar{E} \left[P_{t+1}^{nm} \left(\frac{U_{t+1}^m}{\bar{U}^m} \right)^{1-\alpha} \right]$. Those expectations can be decomposed as

$$\bar{E} \left[P_{t+1}^{nm} \left(\frac{U_{t+1}^m}{\bar{U}^m} \right)^{1-\alpha} \right] = \bar{P}^{nm} \bar{E} \left[\left(\frac{U_{t+1}^m}{\bar{U}^m} \right)^{1-\alpha} \right] + cov \left(P_{t+1}^{nm}, \left(\frac{U_{t+1}^m}{\bar{U}^m} \right)^{1-\alpha} \right) \quad (32)$$

The first term represents the contribution of simple variation in utilities over time, while the second measures the contribution from the covariance of the utilities with the transition probabilities.

One interesting result is immediately obtained from the above: if U_{t+1} is constant, i.e. there is no variation in the utilities – the welfare cost of job loss is constant – then variation in the

probability of a job loss is irrelevant for utility. This result is also obtained by Krebs (2007).¹²

Table 5 uses the formulas above to calculate the parts of the total welfare losses due to business cycles coming from pure variation in utility itself – the $\bar{P}^{nm} \bar{E} \left[(U_{t+1}^m / \bar{U}^m)^{1-\alpha} \right]$ terms – versus the parts coming from covariation of the transition probabilities with the utilities.

Rows 4 and 7 of table 5 show that in our baseline calibration, both pure variation in loss sizes and also the covariation with probabilities are important contributors to the total welfare loss, with 47 percent coming from variation in the loss size and 53 percent from covariation between loss sizes and probabilities. We further decompose the covariance term into components coming from covariance in the unemployment state – i.e. $cov \left(P_{t+1}^{UE}, (U_{t+1}^E / \bar{U}^E)^{1-\alpha} \right)$ and $cov \left(P_{t+1}^{UU}, (U_{t+1}^U / \bar{U}^U)^{1-\alpha} \right)$ – versus the employment state. There is again a roughly equal split, with 53 percent driven by covariance during unemployment and 47 percent from employment. What that shows is that it is not just variation in the probability of losing one’s job or the pain of that loss that drives welfare. Just as important is variation in the probability of finding a new job and dispersion in utility between remaining unemployed or transitioning back to employment.

So it is not just the risk faced by those who are currently employed that matters, but also the risk faced by the unemployed. That fact is perhaps initially surprising since the fraction of the population that is unemployed is typically small. But since all workers have some chance of transitioning into unemployment, how painful that state of the world is ends up being an important driver of welfare overall.

6.4 Robustness

Table 5 reports the components of the welfare cost of business cycles for four perturbations of the model. First, we reduce α from 3 to 1.5. Naturally, the welfare cost of cycles falls by half, but it is still substantial and remains dominated by variation in idiosyncratic consumption risk. The cost of inflation is essentially unchanged due to the fact that inflation affects utility primarily by reducing average consumption as opposed to changing its volatility.

Second, we set the replacement rate of unemployment benefits to be half of wages (where, for simplicity, we assume that the government knows each agent’s market wage). In that case, idiosyncratic risk is smaller for two reasons. First, agents lose less income when unemployed. Second, because agents receive income during unemployment, they draw down less of their human capital – they consume market goods instead of destroying human capital by playing video games. The per-period decline in human capital during unemployment with $b = 0.5$ is only 1.82 percent, compared to 2.7 percent when $b = 0$. The long-run reduction in wages is then only 10.7 instead of 15.5 percent, and the reduction in their NPV (which takes into account the period of unemployment) is 14.4 instead of 18.4 percent. That fact causes the average utility loss following a separation to fall

¹²The result relies on the assumption that U_t^E and U_t^U are constant. If there is persistent variation in the transition probabilities, though, they will not be constant. The result in Krebs (2007) and here technically holds only when the transition probabilities are i.i.d. over time. However, in numerical experiments, we find that persistent variation in the transition probabilities alone has extremely small effects on welfare.

by more than 1/3, from 17 to 11 percent, reducing the welfare cost of fluctuations in consumption risk by 37 percent, to 0.33 percent of lifetime consumption risk. The component of welfare losses coming from consumption alone falls by half. That cost remains nearly twice as large as the cost of inflation fluctuations, though, rendering our basic results unchanged.

The third column increases the steady-state interest rate by two percentage points. A higher degree of discounting means that utility depends less on the possibility of extremely low consumption far in the future. This reduces the welfare costs by about one third to 0.71 percent of lifetime consumption.

Column 5 examines a specification of the model with power utility instead of Epstein–Zin preferences. That is done by setting the inverse IES, ρ , equal to risk aversion at 3. That change increases welfare costs very slightly to 0.98 percent of lifetime consumption, consistent with the fact that the increase in ρ makes preferences more concave than in the baseline.

Last, column 6 examines a version of the model where the cost of job loss is forced to be more strongly cyclical. As discussed above, the benchmark calibration has relatively little variation in the utility loss following a layoff – the loss has a mean of 11 percent and a standard deviation of only 1 percent. In column 5, we specify the parameter d , the human capital loss per period of unemployment, to fluctuate with the level of employment:

$$d_t = 0.01 - 1.73 \times (N_t/\bar{N}) \quad (33)$$

The coefficient 1.73 is chosen so that the utility cost of job loss varies by twice as much as in the benchmark – 2 percent (it implies a standard deviation for d of 0.0018). That degree of variability is still far less than the value used in the calibration of Krebs (2007). Nevertheless, that change more than doubles the welfare cost of business cycles to 2.16 percent of lifetime consumption. The result in column 6 therefore suggests that in a model where there was more variability in how costly the loss of a job was, welfare costs would be substantially higher. We will see in the next section, though, that even with a very small amount of variability, optimal policy still tilts strongly in favor of minimizing output and employment fluctuations.

7 Welfare across policy rules

To see the welfare effects of stabilization policies, we examine Taylor rules of the form

$$R_t = \bar{R} + \phi_\pi (\Pi_t/\bar{\Pi}) + \phi_y (Y_t/\bar{Y}) + \phi_{F/N} \left(\frac{F_t/N_{t-1}}{\bar{F}/\bar{N}} \right) \quad (34)$$

We first examine cases with ϕ_π and $\phi_{F/N}$ varying and $\phi_y = 0$, then allow ϕ_y to vary with $\phi_{F/N} = 0$. The results in this section represent the paper’s main contribution on optimal policy.

7.1 Responding to the firing rate

Figure 6 examines welfare under various choices for ϕ_π and $\phi_{F/N}$ with $\phi_y = 0$. The left-hand panel shows that welfare losses uniformly decline as $\phi_{F/N}$ increases, even to extremely high levels. Increasing ϕ_π almost always reduces welfare, except in a small region where $\phi_{F/N}$ is close to zero and ϕ_π is also small. The potential increase in welfare from shifting from the worst among the policies examined in figure 6 is 2.7 percent of lifetime consumption. That value is large relative to the numbers typically obtained in the New Keynesian literature on optimal policy, in which welfare costs from inflation and output variation are typically on the order of only 0.1 percent. Moreover, the fact that it is optimal to focus on stabilizing employment (specifically, firings) rather than inflation is also different from the literature. So in a model that can capture cross-sectional risk due to layoffs, there can be large benefits to the central bank placing high weight on the layoff rate. That is the paper's main policy result.

Looking across the policies, those that are associated with the lowest utility have low values for both ϕ_π and $\phi_{F/N}$, allowing both inflation and output to be volatile. The best policies all have high values of $\phi_{F/N}$, but there is a wide range of values of ϕ_π that produce highly similar levels of utility. Generally welfare is improved with higher values of ϕ_π when $\phi_{F/N}$ is high, but the effects are quantitatively small for ϕ_π ranging anywhere between 2 and 16, which can be seen from the fact that the contour lines are nearly vertical.

The middle panel of figure 6 plots the component of welfare losses from inflation variation. As one would expect, when policy focuses more on stabilizing layoffs, inflation volatility leads to declines in welfare. The largest declines appear when $\phi_{F/N}$ reaches the maximum values that we study. Those effects are quantitatively small compared with the overall variation in welfare. The difference in the inflation component of utility between the best and worst policies (for inflation) is only 0.33 percent of lifetime consumption, compared to 2.7 percent for total utility.

The right-hand panel plots the remaining part of the welfare losses, which come from consumption risk. The right-hand and left-hand panels look nearly identical, showing that it is consumption risk that drives the variation in welfare losses across policy rules.

To further understand how the policy rules affect welfare and the dynamics of the economy, figure 7 plots the standard deviations of inflation, unemployment, the firing rate, and the utility loss from firings. Consistent with the results in figure 6, inflation volatility rises when $\phi_{F/N}$ is relatively large compared to ϕ_π , from a low of 0.15 percent per year to peak of 4.66 percent. So across the policies, there can be massive differences in the behavior of inflation. At the same time, though, high values of $\phi_{F/N}$ reduce the standard deviations of the unemployment rate (and hence output), the firing rate, and the utility loss from firings. The standard deviation of employment falls from a maximum of 2.0 percent with low values of $\phi_{F/N}$ to a minimum of 0.05 percent with high values of $\phi_{F/N}$. The volatilities of output and the layoff rate fall by similar amounts.

Figure 6 implies that in the model welfare is maximized when the central bank focuses on stabilizing layoffs instead of inflation. To see how such a policy affects the response of the economy

to the demand and markup shocks, figures 2 and 3 plot impulse response functions when $\phi_\pi = 2.03$, $\phi_y = 0$, and $\phi_{F/N} = 40$ (marked in the figures as the “aggressive” policy. The markup shock is most important here as it is the shock that induces a trade-off between stabilizing output and inflation. The aggressive policy with large $\phi_{F/N}$ leads to an output response that is less than one fifth as large as in the baseline calibration, an employment response that is even smaller, while inflation responds by more than twice as much. Even more impressively, the firing rate response is smaller than in the baseline by nearly a factor of 10.

Table 4 also reports unconditional moments for the economy with the aggressive policy. The standard deviations of employment and output growth both fall compared to the baseline with $[\phi_\pi, \phi_{F/N}, \phi_y] = [2.03, 0, 0.125]$ by approximately a factor of 10, while the volatility of inflation rises by only a factor of 1.5, demonstrating the efficiency of the central bank responding to the layoff rate. The amount of variation in the layoff rate and utility loss due to layoffs also both fall close to zero. So the optimal policy implied by the contour plots not only leads to quantitatively large increases in welfare, it also substantially changes the dynamics of the economy, leading to much more stable output and employment and somewhat more volatile inflation.

7.2 Responding to output versus the layoff rate

It is more common to specify policy rules in terms of a response to the output gap, rather than the firing rate (even though we showed above that in fact the firing rate has more explanatory power for interest rates than any other measure of output). Figures 8 and 9 therefore replicate figures 6 and 7, but varying ϕ_y and holding $\phi_{F/N}$ at zero. Figure 8 shows that there is also a benefit to increasing ϕ_y , but that it is somewhat weaker than that for $\phi_{F/N}$. While it is possible to reduce the welfare loss from business cycles to only 0.08 percent of lifetime consumption by increasing $\phi_{F/N}$, with ϕ_y it can only be reduced to 0.16 percent. ϕ_y is also less effective for reducing the volatility of the economy. The standard deviations of unemployment and the firing rate are both substantially higher under the high- ϕ_y policy than the high- $\phi_{F/N}$ policy.

The contour plots suggest that responding to output (or unemployment, since they are over 99 percent correlated in the model as there are no productivity shocks) reduces the welfare cost of business cycles less than responding to the layoff rate. To check that more formally, the two panels of figure 10 show efficient frontiers that display two trade-offs: inflation volatility versus employment volatility, and the utility loss from inflation volatility versus the utility loss from consumption risk. Each panel plots one frontier for policies that respond to inflation and output and another for policies that respond to inflation and the layoff rate.

In both panels of figure 10, the frontier generated by responding to the layoff rate is more favorable than the frontier generated by responding to output. That is, policies that respond to layoffs can generate less volatility in output and employment for a given volatility of inflation, or less inflation volatility for a given level of volatility in output and employment than can be obtained by responding to output. For example, two policies can generate a standard deviation of

annualized inflation of 0.5 percent (i.e. allowing inflation to generally run between 1 and 3 percent, assuming a target of 2), one with $[\phi_\pi, \phi_{F/N}, \phi_y] = [16.7, 7.1, 0]$ and the other with $[\phi_\pi, \phi_{F/N}, \phi_y] = [14.6, 0, 1.18]$. The policy with $\phi_{F/N} = 7.1$ yields a standard deviation of employment of 1.34 percent, compared to 1.81 percent for the $\phi_y = 1.18$ policy. So setting policy in response to the layoff rate in this case yields output and employment volatility that is lower by a quarter for the same inflation volatility. The same result holds for the welfare cost of inflation and consumption volatility – responding to layoffs instead of output can improve welfare simultaneously along both dimensions.

The fact that the policy responding to the layoff rate is more effective than one responding to output is not surprising in light of the IRFs. The layoff rate and the utility loss due to a layoff – the two primary drivers of welfare – are 49 percent correlated with each other. So controlling layoffs, which directly affect welfare, also helps control the utility loss from layoffs, which is the other component.

8 Conclusion

Recent results on optimal monetary policy imply that central banks should not use policy rules that attempt to stabilize output. Yet statutory policy mandates and the actual behavior of central banks suggest that policymakers believe that it is important to try to stabilize the business cycle, and in particular employment. We argue in this paper that the reason that such an emphasis is placed on employment is that there are large welfare losses associated with consumption declines following job loss. Not only are these welfare losses large under standard preference specifications, but we also provide evidence that in fact it is precisely the layoff rate, as opposed to other measures of the state of the business cycle, that the Federal Reserve has historically targeted.

We then build an equilibrium model of the business cycle and use it to examine optimal policy, showing that it is in fact optimal for the central bank to respond to the layoff rate, and that the welfare gains from such a policy rule can be quantitatively large – up to 2 percent of lifetime consumption.

While our focus is on optimal policy, the modeling tools developed here are also important as a technical contribution to the literature. Building on work in the asset pricing literature (Constantinides and Ghosh (2013) and Schmidt (2015)) and on Blanchard and Gali’s (2010) model of unemployment, we develop a New Keynesian model with heterogeneous agents and endogenous layoffs that can still be solved and simulated using standard techniques. We also provide novel expressions for welfare that allow for straightforward policy analysis.

References

- Albuquerque, Rui, Martin Eichenbaum, Victor Xi Luo, and Sergio Rebelo**, “Valuation risk and asset pricing,” *The Journal of Finance*, 2016, 71 (6), 2861–2904.
- Batini, Nicoletta and Andrew Haldane**, “Forward-looking rules for monetary policy,” in “Monetary policy rules,” University of Chicago Press, 1999, pp. 157–202.
- Calvo, Guillermo A.**, “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 1983, 12 (3), 383–398.
- Challe, Edouard**, “Uninsured Unemployment Risk and Optimal Monetary Policy,” 2018. Working paper.
- and **Xavier Ragot**, “Precautionary saving over the business cycle,” *The Economic Journal*, 2015.
- , **Julien Matheron, Xavier Ragot, and Juan F. Rubio-Ramirez**, “Precautionary Saving and Aggregate Demand,” 2015. Working paper.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans**, “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 2005, 113 (1), 1–45.
- Clarida, Richard, Jordi Gal, and Mark Gertler**, “The Science of Monetary Policy: A New Keynesian Perspective,” *Journal of Economic Literature*, 1999, 37(4), 1661–1707.
- Coibion, Olivier, Yuriy Gorodnichenko, and Johannes Wieland**, “The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound?,” *The Review of Economic Studies*, 2012, 79 (4), 1371–1406.
- Couch, Kenneth A and Dana W Placzek**, “Earnings losses of displaced workers revisited,” *The American Economic Review*, 2010, 100 (1), 572–589.
- Davis, Steven J. and Till von Wachter**, “Recessions and the Costs of Job Loss,” *Brookings Papers on Economic Activity*, 2011, 2011 (2), 1–55.
- Epstein, Larry G. and Stanley E. Zin**, “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis,” *The Journal of Political Economy*, 1991, 99(2), 263–286.
- Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin**, “Optimal monetary policy with staggered wage and price contracts,” *Journal of Monetary Economics*, 2000, 46 (2), 281–313.

- Gali, Jordi**, *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and its Applications*, Princeton University Press, 2015.
- , **Frank Smets**, and **Rafael Wouters**, “Unemployment in an estimated new keynesian model,” *NBER Macroeconomics Annual*, 2012, 26 (1), 329–360.
- Guvenen, Fatih, Serdar Ozkan, and Jae Song**, “The Nature of Countercyclical Income Risk,” *Journal of Political Economy*, 2014, 122 (3), 621–660.
- Harris, Milton and Bengt Holmstrom**, “A theory of wage dynamics,” *The Review of Economic Studies*, 1982, 49 (3), 315–333.
- Imrohoroğlu, Ayşe**, “Cost of business cycles with indivisibilities and liquidity constraints,” *The Journal of Political Economy*, 1989, pp. 1364–1383.
- Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan**, “Earnings Losses of Displaced Workers,” *The American Economic Review*, 1993, 83 (4), 685–709.
- Jarosch, Gregor**, “Searching for job security and the consequences of job loss,” Technical Report, Mimeo, University of Chicago, USA 2014.
- Jr, Robert E Lucas**, “Macroeconomic priorities,” *The American Economic Review*, 2003, 93 (1), 1.
- Jung, Philip and Moritz Kuhn**, “Earnings losses and labor mobility over the life cycle,” *Journal of the European Economic Association*, 2018.
- Kletzer, Lori G. and Robert W. Fairlie**, “The Long-Term Costs of Job Displacement for Young Adult Workers,” *Industrial & Labor Relations Review*, 2003, 56 (4), 682–698.
- Krause, Michael and Thomas Lubik**, “The (ir)relevance of real wage rigidity in the New Keynesian model with search frictions,” *Journal of Monetary Economics*, 2007, 53 (3), 706–727.
- Krebs, Tom**, “Job Displacement Risk and the Cost of Business Cycles,” *American Economic Review*, 2007, 97 (3), 664–686.
- Krolikowski, Pawel**, “Job ladders and earnings of displaced workers,” *American Economic Journal: Macroeconomics*, 2017, 9 (2), 1–31.
- Krusell, Per, Toshihiko Mukoyama, Ayşegül Şahin, and Anthony A Smith**, “Revisiting the welfare effects of eliminating business cycles,” *Review of Economic Dynamics*, 2009, 12 (3), 393–404.
- Mukoyama, Toshihiko and Ayşegül Şahin**, “Costs of business cycles for unskilled workers,” *Journal of Monetary Economics*, 2006, 53 (8), 2179–2193.

- Newey, Whitney K. and Kenneth D. West**, “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 1987, *55*(3), 703–708.
- Phelan, Brian J.**, “Systematic and Idiosyncratic Heterogeneity in the Cost of Worker Displacement.” Working paper.
- Rotemberg, Julio J. and Michael Woodford**, *Interest Rate Rules in an Estimated Sticky Price Model*, University of Chicago Press,
- Ruhm, Christopher J.**, “Are Workers Permanently Scarred by Job Displacements?,” *The American Economic Review*, 1991, *81* (1), 319–324.
- Schmitt-Grohe, Stephanie and Martin Uribe**, “Optimal Simple and Implementable Monetary and Fiscal Rules,” *Journal of Monetary Economics*, 2007, *54* (6), 1702–1725.
- Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron**, “Cyclical Dynamics in Idiosyncratic Labor Market Risk,” *Journal of Political Economy*, 2004, *112* (3), 695–717.
- von Wachter, Till, Jae Song, and Joyce Manchester**, “Long-Term Earnings Losses due to Mass Layoffs During the 1982 Recession: An Analysis Using U.S. Administrative Data from 1974 to 2004,” 2009. Working paper.
- Woodford, Michael**, *Interest and Prices*, Princeton University Press, Princeton, NJ, 2003.
- Yun, Tack**, “Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles,” *Journal of Monetary Economics*, 1996, *37* (2), 345–370.

A Full model

A.1 Hiring and firing

Each firm faces a quit rate q_t . They can fire workers $F_{i,t} \geq 0$ or hire $H_{i,t} \geq 0$. For each individual firm, the number of employees follows

$$N_{m,t} = (1 - q_t) N_{m,t-1} + H_{m,t} - F_{m,t}, \quad (35)$$

In the aggregate, then, the sum across all firms is

$$N_t = (1 - q_t) N_{t-1} - F_t + H_t. \quad (36)$$

The timing assumption here is that firms first hire, which induces quits, and then fire.

A.1.1 The quit rate

Hiring is taken randomly from the pool of available workers. The workers available to be hired are the unemployed, $1 - N_t$, and those whose wage makes them indifferent to leaving the firm, λN_t (where λ is an exogenous parameter). So the fraction of hiring from currently employed workers is

$$\frac{\lambda N_{t-1}}{\underbrace{1 - N_{t-1}}_{\text{Unemployed}} + \underbrace{\lambda N_{t-1}}_{\text{Willing to accept}}} = \frac{\lambda N_{t-1}}{1 - (1 - \lambda) N_{t-1}} \quad (37)$$

Total hiring from those workers, which represents quits, is

$$Q_t = H_t \frac{\lambda N_{t-1}}{1 - (1 - \lambda) N_{t-1}}. \quad (38)$$

Quits as a fraction of last period's employment is

$$q_t = \frac{Q_t}{N_{t-1}} = \frac{\lambda H_t}{1 - (1 - \lambda) N_{t-1}}. \quad (39)$$

A.1.2 Optimization

Each firm solves the problem

$$\max_{N_{m,t}} S_t A_t A_{m,t} \left(\underbrace{K_t^E N_{m,t}}_{\text{Effective labor}} \right)^\gamma - W_t K_t^E N_{m,t} \quad (40)$$

where S_t is the real price of their output and W_t is the real wage per unit of human capital.

The first-order condition for labor yields

$$N_{m,t} = \gamma (K_t^E)^{-1} (S_t A_t A_{i,t} / W_t)^{1/(1-\gamma)} \quad (41)$$

We assume that $a_{i,t} = \log A_{i,t}$ follows a random walk with drift so that A_i is a martingale,

$$a_{i,t} = a_{i,t-1} + \varepsilon_i \quad (42)$$

$$\varepsilon_i \sim N\left(-\frac{1}{2} \frac{1}{1-\gamma} (\sigma^a)^2, \sigma^a\right) \quad (43)$$

Since $A_{i,t}$ is a martingale, its integral across firms is constant, and we have

$$N_t = (K_t^E)^{-1} (S_t A_t / W_t)^{1/(1-\gamma)}. \quad (44)$$

Taking logs and first differences,

$$\Delta n_{i,t} = \frac{1}{1-\gamma} (\Delta s_t + \Delta a_t - \Delta w_t + \Delta a_{m,t}) - \Delta k_t^E \quad (45)$$

$$\sim N\left(\underbrace{\frac{\Delta s_t + \Delta a_t - \Delta w_t}{1-\gamma} - \Delta k_t^E}_{\equiv \mu_n} - \frac{1}{2} \left(\frac{\sigma^a}{1-\gamma}\right)^2, \underbrace{\left(\frac{\sigma^a}{1-\gamma}\right)^2}_{\equiv \sigma_n^2}\right). \quad (46)$$

If $\exp(\Delta n_{i,t}) > 1 - q_t$, then firm i needs to hire, and

$$\exp(\Delta n_{i,t}) - 1 - q_t = \exp\left(\log \frac{H_{i,t}}{N_{i,t-1}}\right). \quad (47)$$

Thus

$$\frac{H_{m,t}}{N_{m,t-1}} = \begin{cases} \exp(\Delta n_{m,t}) - 1 - q_t & \Delta n_{m,t} \geq \ln(1 - q_t) \\ 0 & \text{otherwise} \end{cases}. \quad (48)$$

Similarly, firings are

$$\frac{F_{m,t}}{N_{m,t-1}} = \begin{cases} 1 - q_t - \exp(\Delta n_{m,t}) & \Delta n_{m,t} \leq \ln(1 - q_t) \\ 0 & \text{otherwise} \end{cases}. \quad (49)$$

Therefore for each m we have

$$(1 - q_t) + \frac{H_{m,t}}{N_{m,t-1}} - \frac{F_{m,t}}{N_{m,t-1}} = \exp(\Delta n_{m,t}). \quad (50)$$

To get aggregate hiring and firing, we integrate across all firms:

$$\frac{H_t}{N_{t-1}} = \int \frac{H_{m,t}}{N_{m,t-1}} di = \int 1_{\{\Delta n_{m,t} \geq \ln(1-q_t)\}} [\exp(\Delta n_{m,t}) - (1-q_t)] dm \quad (51)$$

$$= \int 1_{\{\Delta n_{m,t} \geq \ln(1-q_t)\}} \exp(\Delta n_{m,t}) dm - P(\Delta n_{m,t} \geq \ln(1-q_t)) (1-q_t) \quad (52)$$

$$= E[\exp(\Delta n_{m,t}); \Delta n_{m,t} \geq \ln(1-q_t)] - P(\Delta n_{m,t} \geq \ln(1-q_t)) (1-q_t) \quad (53)$$

$$= \exp\left(\mu_n + \frac{1}{2}\sigma_n\right) \Phi\left(\frac{\mu_n + \sigma_n^2 - \ln(1-q_t)}{\sigma_n}\right) - \left(1 - \Phi\left(\frac{\ln(1-q_t) - \mu_n}{\sigma_n}\right)\right) (1-q_t) \quad (54)$$

where the last line uses the formula for the partial expectation of a log-Normal random variable. Similar logic for the firing rate yields

$$\frac{F_t}{N_{t-1}} = \int \frac{F_{m,t}}{N_{m,t-1}} dm = (1-q_t) P(\Delta n_{m,t} \leq \ln(1-q_t)) - \int 1_{\{\Delta n_{m,t} \leq \ln(1-q_t)\}} \exp(\Delta n_{m,t}) dm \quad (55)$$

$$= (1-q_t) \Phi\left(\frac{\ln(1-q_t) - \mu_n}{\sigma_n}\right) - \exp\left(\mu_n + \frac{1}{2}\sigma_n\right) \left[1 - \Phi\left(\frac{\mu_n + \sigma_n^2 - \ln(1-q_t)}{\sigma_n}\right)\right] \quad (56)$$

A.2 Price stickiness

There are intermediate good differentiators indexed by j who take the (undifferentiated) good produced by the firms described above and transform it into a type- j output. That output is then formed into the final good with a CES aggregator. Differentiator j purchases its input for the nominal price $P_t s_t$ and sells it at price P_j . The differentiators face a demand curve such that

$$Y_{j,t+k} = \left(\frac{P_j}{P_{t+k}}\right)^{-\varepsilon} Y_{t+k} \quad (57)$$

where $Y_{j,t}$ is firm j 's output and Y_t is aggregate output. We assume Calvo pricing. When a differentiator is allowed to change their price, they solve the problem

$$\max_{P_j} E_t \sum_{k=0}^{\infty} \theta^k M_{t,t+k} \frac{P_t}{P_{t+k}} (P_j Y_{j,t+k} - P_{t+k} s_{t+k} Y_{i,t+k}) \quad (58)$$

where $M_{t,t+k}$ is the real stochastic discount factor between dates t and $t+k$ and θ is the probability of changing prices in a given period.

Optimization then yields the usual New Keynesian pricing system, with

$$\tilde{x}_t^1 = \nu \frac{\varepsilon}{\varepsilon - 1} s_t \exp(\psi_t) + \theta E_t \pi_{t+1}^\varepsilon \frac{Y_{t+1}}{Y_t} M_{t,t+1} \tilde{x}_{t+1}^1 \quad (59)$$

$$\tilde{x}_t^2 = 1 + \theta E_t \pi_{t+1}^{\varepsilon-1} \frac{Y_{t+1}}{Y_t} M_{t,t+1}^t \tilde{x}_{t+1}^2 \quad (60)$$

$$\frac{\tilde{x}_t^1}{\tilde{x}_t^2} = \left(\frac{1 - \theta \pi_t^{\varepsilon-1}}{1 - \theta} \right)^{\frac{1}{1-\varepsilon}} \quad (61)$$

where ν is a production tax, ψ a markup shock, ε is the elasticity of substitution, and π is inflation.

A.3 Households

The optimization problem from the text is

$$\max U_{i,t} + E_t \sum_{j=0}^{\infty} \lambda_{i,t+j} \left[\begin{array}{l} \left((1_{i,t+j}^E (1 - \tau_t) + (1 - 1_{i,t+j}^E) b) W_{i,t+j} K_{i,t+j} - C_{i,t+j} - G_{i,t+j} \right) \\ - \psi_{i,t+j} \left(K_{i,t+j+1} - G_{i,t+j} + K_{i,t+j} \exp \left(- (1 - 1_{i,t+j}^E) d \right) (1 - \delta_K) \right) \end{array} \right] \quad (62)$$

where λ and ψ are Lagrange multipliers. The first-order conditions are

$$C_{i,t+j} : \lambda_{i,t+j} = \frac{\partial U_{i,t}}{\partial C_{i,t+j}} \quad (63)$$

$$G_t : \psi_{i,t} = 1 \quad (64)$$

$$K_{t+1} : \psi_{i,t} = E_t \left[\frac{\lambda_{i,t+1}}{\lambda_{i,t}} \left((1_{i,t+1}^E + b (1 - 1_{i,t+1}^E)) W_{t+1} + \psi_{i,t+1} \exp \left(- (1 - 1_{i,t+1}^E) d \right) (1 - \delta_K) \right) \right] \quad (65)$$

$\psi_{i,t}$ can clearly be eliminated, so we just have the single Euler equation.

The budget constraint and dynamic equation for K can be rescaled by $K_{i,t}$, yielding

$$\frac{K_{i,t+1}}{K_{i,t}} = \exp \left(- (1 - 1_{i,t}^E) d \right) (1 - \delta_K) + g_{i,t} \quad (66)$$

$$\left((1 - \tau_t) 1_{i,t}^E + b (1 - 1_{i,t}^E) \right) W_t = c_{i,t} + g_{i,t} \quad (67)$$

where $g_{i,t} \equiv G_{i,t}/K_{i,t}$ and $c_{i,t} = C_{i,t}/K_{i,t}$. Note that the Euler equation also does not depend on any variables involving the level of $K_{i,t}$. That immediately shows that the problem is homogeneous in $K_{i,t}$ and aggregates.

We can specialize the problem for the two types, employed and unemployed.

Euler equation:

$$1 = E_t \left[M_{t+1}^{EE} P_{t+1}^{EE} \left((1 - \tau_{t+1}) W_{t+1} + 1 \right) + M_{t+1}^{EU} P_{t+1}^{EU} (b W_{t+1} + \exp(-d)) \right] \quad (68)$$

$$1 = E_t \left[M_{t+1}^{UE} P_{t+1}^{UE} \left((1 - \tau_{t+1}) W_{t+1} + 1 \right) + M_{t+1}^{UU} P_{t+1}^{UU} (b W_{t+1} + \exp(-d)) \right] \quad (69)$$

Budget constraint:

$$W_t = c_t^E + g_t^E \quad (70)$$

$$bW_t = c_t^U + g_t^U \quad (71)$$

Human capital growth:

$$\left(\frac{K_{i,t+1}}{K_{i,t}}\right)^U = (\exp(-d)(1 - \delta_K) + g_t^U) \quad (72)$$

$$\left(\frac{K_{i,t+1}}{K_{i,t}}\right)^E = (1 - \delta_K + g_t^E) \quad (73)$$

A.4 Euler equation

The general expression for the SDF under Epstein–Zin preferences is

$$M_{i,t+1} = \beta \frac{U_{i,t+1}^{\rho-\alpha}}{E_t \left[U_{i,t+1}^{1-\alpha} \right]^{\frac{\rho-\alpha}{1-\alpha}}} \left(\frac{C_{i,t+1}}{C_{i,t}} \right)^{-\rho} \quad (74)$$

Due to the linear homogeneity noted above, utility for the two types is proportional to current human capital, so that $U_{i,t} = U_t^E K_{i,t}$ for an employed worker.

$$U_t^E K_{i,t} = \left((1 - \beta) C_{i,t}^{1-\rho} + \beta E_t \left[P_{t+1}^{EE} U_{t+1}^{E(1-\alpha)} K_{i,t+1}^{1-\alpha} + P_{t+1}^{EU} U_{t+1}^{U(1-\alpha)} K_{i,t+1}^{1-\alpha} \right]^{\frac{1-\rho}{1-\alpha}} \right)^{\frac{1}{1-\rho}} \quad (75)$$

$$U_t^E = \left((1 - \beta) c_t^{E(1-\rho)} + \beta (1 - \delta_K + g_t^E)^{1-\rho} E_t \left[\left(P_{t+1}^{EE} U_{t+1}^{E(1-\alpha)} + P_{t+1}^{EU} U_{t+1}^{U(1-\alpha)} \right) \right]^{\frac{1-\rho}{1-\alpha}} \right)^{\frac{1}{1-\rho}} \quad (76)$$

where the second line uses the fact that $K_{i,t+1}/K_{i,t}$ is known as of date t . For the unemployed,

$$U_t^U = \left((1 - \beta) c_t^{U(1-\rho)} + \beta (\exp(-d)(1 - \delta_K) + g_t^U)^{1-\rho} E_t \left[\left(P_{t+1}^{UE} U_{t+1}^{U(1-\alpha)} + P_{t+1}^{UU} U_{t+1}^{U(1-\alpha)} \right) \right]^{\frac{1-\rho}{1-\alpha}} \right)^{\frac{1}{1-\rho}} \quad (77)$$

To calculate the pricing kernel, we need expectations,

$$E_t^E \left[U_{i,t+1}^{1-\alpha} \right]^{\frac{\rho-\alpha}{1-\alpha}} = E_t \left[\left(P_{t+1}^{EE} (U_{t+1}^E K_{i,t+1})^{1-\alpha} + P_{t+1}^{EU} (U_{t+1}^U K_{i,t+1})^{1-\alpha} \right) \right]^{\frac{\rho-\alpha}{1-\alpha}} \quad (78)$$

$$= K_{i,t}^{\rho-\alpha} (1 - \delta_K + g_t^E)^{\rho-\alpha} E_t \left[\left(P_{t+1}^{EE} U_{t+1}^{E(1-\alpha)} + P_{t+1}^{EU} U_{t+1}^{U(1-\alpha)} \right) \right]^{\frac{\rho-\alpha}{1-\alpha}} \quad (79)$$

If we define

$$E_t^E U \equiv E_t \left[\left(P_{t+1}^{EE} U_{t+1}^{E(1-\alpha)} + P_{t+1}^{EU} U_{t+1}^{U(1-\alpha)} \right) \right] \quad (80)$$

Then

$$E_t^E \left[U_{i,t+1}^{1-\alpha} \right]^{\frac{\rho-\alpha}{1-\alpha}} = (E_t^E U)^{\frac{\rho-\alpha}{1-\alpha}} K_{i,t}^{\rho-\alpha} (1 - \delta_K + g_t^E)^{\rho-\alpha} \quad (81)$$

Similarly,

$$E_t^U U \equiv E_t \left[\left(P_{t+1}^{UE} U_{t+1}^{E(1-\alpha)} + P_{t+1}^{UU} U_{t+1}^{U(1-\alpha)} \right) \right] \quad (82)$$

$$E_t^U \left[U_{i,t+1}^{1-\alpha} \right]^{\frac{\rho-\alpha}{1-\alpha}} = (E_t^U U)^{\frac{\rho-\alpha}{1-\alpha}} K_{i,t}^{\rho-\alpha} (\exp(-d)(1 - \delta_K) + g_t^U)^{\rho-\alpha} \quad (83)$$

$$M_{t+1}^{EE} = \beta \frac{(U_{t+1}^E K_{i,t+1})^{\rho-\alpha}}{(E_t^E U)^{\frac{\rho-\alpha}{1-\alpha}} K_{i,t}^{\rho-\alpha} (1 - \delta_K + g_t^E)^{\rho-\alpha}} \left(\frac{c_{t+1}^E K_{i,t+1}}{c_t^E K_{i,t}} \right)^{-\rho} \quad (84)$$

$$= \beta \frac{U_{t+1}^{E(\rho-\alpha)}}{(E_t^E U)^{\frac{\rho-\alpha}{1-\alpha}}} \left(\frac{c_{t+1}^E}{c_t^E} \right)^{-\rho} (1 - \delta_K + g_t^E)^{-\rho} \quad (85)$$

$$M_{t+1}^{EU} = \beta \frac{U_{t+1}^{U(\rho-\alpha)}}{(E_t^E U)^{\frac{\rho-\alpha}{1-\alpha}}} \left(\frac{c_{t+1}^U}{c_t^E} \right)^{-\rho} (1 - \delta_K + g_t^E)^{-\rho} \quad (86)$$

$$M_{t+1}^{UE} = \beta \frac{U_{t+1}^{E(\rho-\alpha)}}{(E_t^U U)^{\frac{\rho-\alpha}{1-\alpha}}} \left(\frac{c_{t+1}^E}{c_t^U} \right)^{-\rho} (\exp(-d)(1 - \delta_K) + g_t^U)^{-\rho} \quad (87)$$

$$M_{t+1}^{UU} = \beta \frac{U_{t+1}^{U(\rho-\alpha)}}{(E_t^U U)^{\frac{\rho-\alpha}{1-\alpha}}} \left(\frac{c_{t+1}^U}{c_t^U} \right)^{-\rho} (\exp(-d)(1 - \delta_K) + g_t^U)^{-\rho} \quad (88)$$

A.5 Aggregate human capital

The average human capital of the employed is K_t^E and human capital for the unemployed is K_t^U . A fraction δ of people die at the beginning of each period. They are replaced by people in the same employment state with human capital equal to \bar{K} , with

$$\bar{K} \equiv \bar{K}^E \bar{N} + \bar{K}^U (1 - \bar{N}) \quad (89)$$

where \bar{K}^E and \bar{K}^U are the steady-state values of K_t^E and K_t^U and \bar{N} is steady-state employment. So \bar{K} is the steady-state average human capital per person.

The average human capital of the unemployed is then

$$K_{t+1}^U = \frac{1 - N_t - H_{t+1} + Q_{t+1}}{1 - N_{t+1}} \left((1 - \delta) K_t^U (\exp(-d)(1 - \delta_K) + g_t^U) + \delta \bar{K} \right) \quad (90)$$

$$+ \frac{F_{t+1}}{1 - N_{t+1}} \left((1 - \delta) K_t^E (1 - \delta_K + g_t^E) + \delta \bar{K} \right) \quad (91)$$

The first term is the human capital of people who were unemployed in the previous period and remain unemployed (with probability $1 - \delta$ it is the same person, who on average has human capital of $K_t^U (\exp(-d)(1 - \delta_K) + g_t^U)$); with probability δ the old unemployed person died and was replaced by somebody with human capital of \bar{K}). The second term is the set of newly hired people. The assumption here is that death happens at the beginning of the period before anything else.

The average human capital of the employed workers is

$$K_{t+1}^E = \frac{N_t - F_{t+1}}{N_{t+1}} ((1 - \delta) K_t^E (1 - \delta_K + g_t^E) + \delta \bar{K}) \quad (92)$$

$$+ \frac{H_{t+1} - Q_{t+1}}{N_{t+1}} ((1 - \delta) K_t^U (\exp(-d)(1 - \delta_K) + g_t^U) + \delta \bar{K}) \quad (93)$$

The first term is the human capital of the previous period's workers who remain employed, while the second term is the set of people who are hired out of unemployment.

The transition probabilities are

$$P_t^{EE} = \frac{N_{t-1} - F_t}{N_{t-1}} \quad (94)$$

$$P_t^{EU} = \frac{F_t}{N_{t-1}} = 1 - P_t^{EE} \quad (95)$$

$$P_t^{UE} = \frac{H_t - Q_t}{1 - N_{t-1}} \quad (96)$$

$$P_t^{UU} = \frac{1 - N_{t-1} - H_t + Q_t}{1 - N_{t-1}} = 1 - P_t^{UE} \quad (97)$$

The condition for taxes is

$$\tau_t W_t K_t^E = b W_t K_t^U \quad (98)$$

$$\tau_t = b K_t^U / K_t^E \quad (99)$$

so taxes are countercyclical, rising when there are relatively more unemployed people.

$$\gamma^{-1} \Delta_t A_t N_t^\gamma = W_t K_t^E \quad (100)$$

$$W_t = \gamma^{-1} \Delta_t A_t N_t^\gamma / K_t^E \quad (101)$$

A.6 Calculating welfare

Utility is

$$U_t^U = \left((1 - \beta) c_t^{U(1-\rho)} + \beta (\exp(-d) + g_t^U)^{1-\rho} E_t \left[\left(P_{t+1}^{UE} U_{t+1}^{U(1-\alpha)} + P_{t+1}^{UU} U_{t+1}^{U(1-\alpha)} \right) \right]^{\frac{1-\rho}{1-\alpha}} \right)^{\frac{1}{1-\rho}} \quad (102)$$

$$U_t^E = \left((1 - \beta) c_t^{E(1-\rho)} + \beta (1 + g_t^E)^{1-\rho} E_t \left[\left(P_{t+1}^{EE} U_{t+1}^{E(1-\alpha)} + P_{t+1}^{EU} U_{t+1}^{U(1-\alpha)} \right)^{\frac{1-\rho}{1-\alpha}} \right]^{\frac{1}{1-\rho}} \right) \quad (103)$$

The difficult part of the recursion is calculating the expectation of a nonlinear function of utility. We solve that problem by constructing the expectation as the projection of $P_{t+1}^{EE} U_{t+1}^{E(1-\alpha)} + P_{t+1}^{EU} U_{t+1}^{U(1-\alpha)}$ onto a polynomial function of the underlying states on date t . Numerically, that can be done through a simple regression.

More specifically, we solve for U_t^E and U_t^U with a contraction. Suppose we have a first-order approximation to the model and have simulated a history of the transition probabilities, P_t^* , consumption, c_t , and human capital investment, g_t . Denote the current guess for the histories of U^E and U^U by ${}^n U_t^E$ and ${}^n U_t^U$. The update is then

$${}^{n+1} U_t^E = \left((1 - \beta) c_t^{E(1-\rho)} + \beta E_t^{proj} \left[P_{t+1}^{EE} ({}^n U_{t+1}^E)^{1-\alpha} + P_{t+1}^{EU} ({}^n U_{t+1}^U)^{1-\alpha} \right]^{\frac{1-\rho}{1-\alpha}} \right)^{\frac{1}{1-\rho}} \quad (104)$$

$${}^{n+1} U_t^U = \left((1 - \beta) c_t^{U(1-\rho)} + \beta E_t^{proj} \left[P_{t+1}^{UE} ({}^n U_{t+1}^E)^{1-\alpha} + P_{t+1}^{UU} ({}^n U_{t+1}^U)^{1-\alpha} \right]^{\frac{1-\rho}{1-\alpha}} \right)^{\frac{1}{1-\rho}} \quad (105)$$

where the operator E_t^{proj} is the expectation conditional on a projection onto the state variables of the first-order approximation and their products (i.e. a set of polynomials) up to some power (in practice, we use up to a fourth-order projection).

There are then two sources of approximation error here. The first is in the P_t^* , c_t , and g_t . Those errors are errors in the dynamics of the economy. The second source of error is the difference between the true expectation operator E_t and the polynomial projection, E_t^{proj} . In the absence of the second type of error, we would simply say that we are calculating welfare exactly for an agent who actually faces the consumption dynamics implied by the first-order approximation. The approximation errors say further that we are calculating welfare for an agent who is unable to calculate exact statistical expectations and instead has to use a simpler polynomial model.

To measure how accurate the polynomial approximation is, we examine changes in implied welfare as the order of the projection is increased. We find quantitatively unimportant differences across the orders – approximately one part in 10,000 between orders 3 and 4.

Table 1. Pairwise correlations with Fed Funds rate innovations

Variable	Correlation
Exponentially filtered initial claims	-0.60
Initial claims	-0.59
Change in unemployment rate	-0.55
Initial claims/total employment	-0.45
HP-filtered log output	0.42
Output growth	0.39
HP-filtered unemployment rate	-0.35
Unemployment rate	-0.28
PCE inflation	0.07
Core PCE inflation	-0.07

Table 2. Estimated Policy Rules										
Panel a. Backward looking rules										
Fed funds rate (t-1)	0.951*** (44.84)	0.975*** (45.24)	0.969*** (44.78)	0.953*** (47.15)	0.955*** (48.02)	0.941*** (36.78)	0.950*** (35.33)	0.942*** (35.52)	0.930*** (38.39)	0.937*** (39.01)
PCE inflation	0.313** (2.14)	0.376** (2.58)	0.392*** (2.64)	0.307** (2.18)	0.278** (1.99)					
IC rate	-56.100*** (-6.90)	-53.831*** (-6.56)	-50.538*** (-5.61)	-33.573*** (-5.63)	-26.961*** (-4.13)	-55.440*** (-6.65)	-54.749*** (-6.67)	-51.591*** (-5.74)	-37.259*** (-6.12)	-30.548*** (-4.49)
Unempl. rate	0.140** (2.60)					0.106* (1.71)				
HP unempl rate		0.158** (2.19)					0.121 (1.64)			
HP log output			-7.835 (-1.37)					-5.195 (-0.91)		
Output growth				0.297*** (3.78)					0.292*** (3.74)	
Δunempl. rate					-0.825*** (-4.14)					-0.791*** (-3.92)
Core PCE infl.						0.447* (1.67)	0.606** (2.58)	0.661*** (2.83)	0.567** (2.56)	0.471** (2.1)
Constant	-1.324*** (-3.76)	-0.636*** (-3.98)	-0.601*** (-3.56)	-0.551*** (-3.95)	-0.278* (-1.91)	-1.152*** (-3.09)	-0.664*** (-4.06)	-0.639*** (-3.72)	-0.623*** (-4.26)	-0.340** (-2.17)
Obs	107	107	107	107	107	107	107	107	107	107
Adjusted R2	0.965	0.964	0.963	0.967	0.968	0.964	0.964	0.963	0.967	0.968
Panel b. Forward looking rules										
Specification	k = 0, q = 0		k = 1, q = 1		k = 4, q = 1		k = 1, q = 2		k = 4, q = 2	
Constant	0.589*** (4.51)	0.215 (1.57)	0.610*** (4.58)	0.194 (1.01)	0.483*** (3.84)	0.201 (1.13)	0.605*** (4.43)	0.114 (0.70)	0.477*** (3.51)	0.214 (1.03)
Expected Inflation	0.414*** (4.62)	0.411*** (5.20)	0.377*** (3.87)	0.369*** (3.86)	0.335*** (3.01)	0.394*** (2.94)	0.520*** (5.06)	0.572*** (5.00)	0.366*** (3.09)	0.404*** (2.89)
Expected output gap	0.169*** (2.87)	-0.278** (-2.19)	0.210*** (3.40)	-0.163 (-1.06)	0.236*** (3.31)	-0.053 (-0.38)	0.348*** (4.85)	-0.184 (-1.20)	0.310*** (3.66)	0.091 (0.51)
Fed funds rate (t-1)	0.883*** (34.80)	0.925*** (38.47)	0.874*** (34.37)	0.924*** (29.12)	0.904*** (36.79)	0.932*** (29.29)	0.863*** (33.61)	0.938*** (32.33)	0.902*** (34.94)	0.931*** (26.36)
Expected IC rate		-0.862*** (-4.73)		-0.787*** (-3.13)		-0.708*** (-3.17)		-0.920*** (-4.45)		-0.569** (-2.03)
Obs	103	103	102	102	99	99	101	101	99	99
P-value: -IC coef > Output gap coef		0		0.008		0.014		0.001		0.139
P-value: -IC coef > Inflation coef		0.005		0.057		0.127		0.057		0.304
Notes: The top panel estimates a backward looking monetary policy rule by OLS. The bottom panel estimates a forward looking policy rule using GMM. The set of instruments includes four lags of inflation: output gap, the federal funds rate, the short-long spread, and commodity price inflation as in Clarida, Gali and Gertler (2000). (k,q) refer to the number of future quarters considered for expected inflation and the output gap respectively. Whenever expected initial claims are included as a regressor we always include it with the same time horizon as the output gap. All regressors are standardized so that their coefficients can be compared. Furthermore, HP-IC was multiplied by negative one so that it would have the same expected sign as the output gap in the regressions. In all panels, numbers in parentheses are t-statistics. *** indicates significance at the 1 percent level, ** the 5 percent level, and * the 10 percent level. The dependent variable is the target Fed funds rate. All data is quarterly and averaged within the period.										

Table 3. Calibrated parameters

Parameter	Value	Description	Parameter	Value	Description
β	0.9916	Time discounting	λ	0.1345	Workers willing to quit
α	3	Risk aversion	δ	0.01	Death rate
ρ	1	Inverse IES	γ	0.66	Labor's share of income
ρ_u	0.9	Persist. of demand shock	Ξ	1.869	Wage
ρ_ψ	0.9	Persist. of markup shock	d	0.01	Human capital loss
σ_u	0.0017	Volatility demand shock	θ	0.7	Price adjustment prob.
σ_ψ	0.0031	Volatility of markup shock	ε	5	Intermed. inv. elast. of subs
σ_a	0.0157	Idio. prod. shock vol.	δ_K	0.01	Human capital depreciation

Table 4. Simulated moments


	Baseline	Aggressive policy	Data	Source
$100 \times stdev(\Delta \log Y)$	0.50	0.04	0.65	GDP growth
$100 \times std(1 - N)$	1.62	0.17	1.66	Unemployment rate
$400 \times std(R)$	2.95	3.05	1.51	3-month T-bill minus linear trend
$100 \times std(\pi)$	0.43	0.66	0.43	PCE deflator
$E[P^{EU}]$	0.0100	0.0100	0.0100	Krebs (2007)
P^{EU} 25 th %tile	0.0068	0.0099	0.0075	...
P^{EU} 75 th %tile	0.0131	0.0101	0.0125	...
$E[1 - U^U/U^E]$	0.1739	0.1739	0.15	Krebs (2007)
$1 - U^U/U^E$ 25 th %tile	0.1687	0.1726	0.09	...
$1 - U^U/U^E$ 75 th %tile	0.1792	0.1752	0.21	...

Notes: Output is measured as real GDP. $1 - N$ is measured by the unemployment rate. R is the three-month T-bill rate. π is growth in the PCE deflator. The "data" values for P^{EU} and $g^{C^{EU}}$ are from Krebs (2007). The 25th and 75th percentiles are his values in the expansion and recession states, respectively.

Table 5. Components of welfare cost of business cycles (percent of lifetime consumption)

	Baseline	$\alpha = 1.5$	$b = 1/2$	$\bar{R} = 7\%$	$\rho = \alpha = 3$	Time-var. d
Total cost ($W^{(2)}$)	0.95	0.49	0.39	0.67	0.98	2.16
Inflation ($W^{(2)} - W^{(1)}$)	0.04	0.04	0.04	0.04	0.04	0.04
Consumption ($W^{(1)}$)	0.91	0.45	0.35	0.63	0.94	2.11
Frac. from $cov\left(P, (U_{t+1}/U_t)^{1-\alpha}\right)$	0.53	0.74	0.55	0.54	0.91	0.48
Unemployment state	0.53	0.51	0.51	0.53	0.06	0.52
Employment state	0.48	0.49	0.49	0.47	0.94	0.48
Frac. from dispersion in U_{t+1}/U_t	0.47	0.26	0.45	0.46	0.09	0.52

Figure 1: Z-scores of Fed Funds, unemployment, and initial claims rates



Notes: The lines are z-scores of the HP-filtered unemployment rate, the initial claims rate minus an exponentially weighted moving average, and the residual from a regression of the Fed Funds rate on its own lag and current PCE inflation. The axis for the Fed Funds rate is reversed..

Figure 2: Impulse responses to demand shock

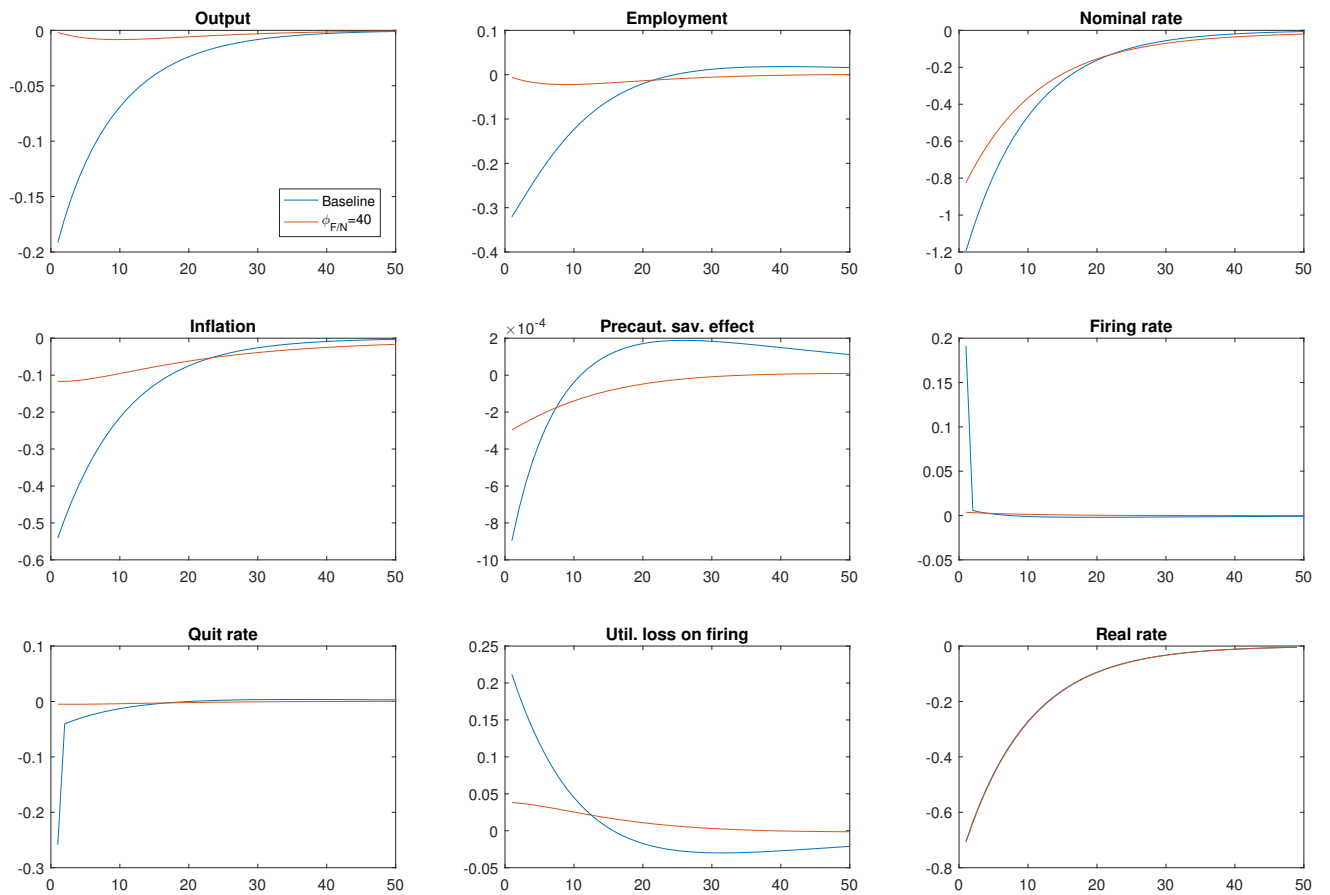


Figure 3: Impulse responses to markup shock

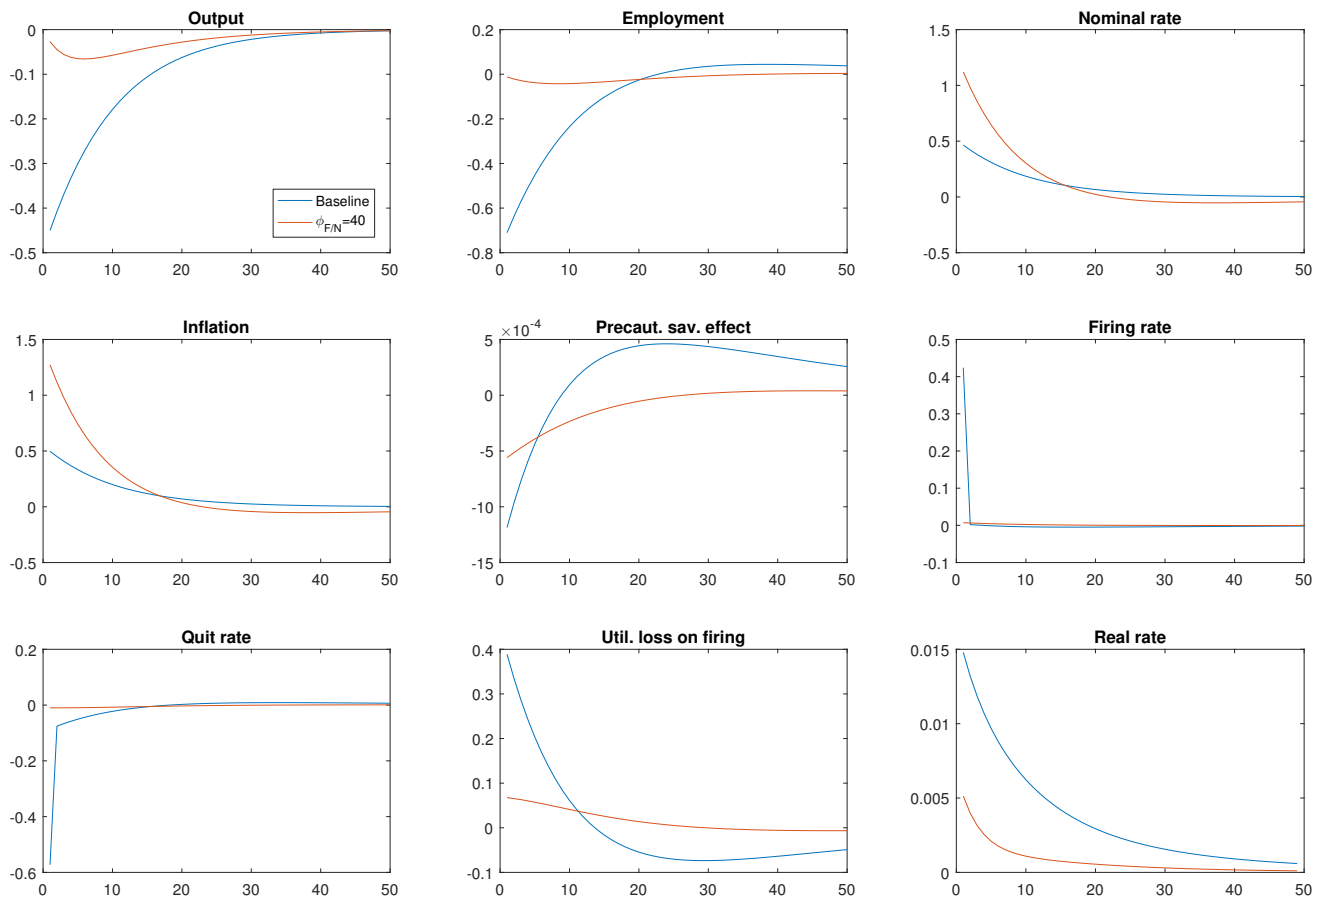
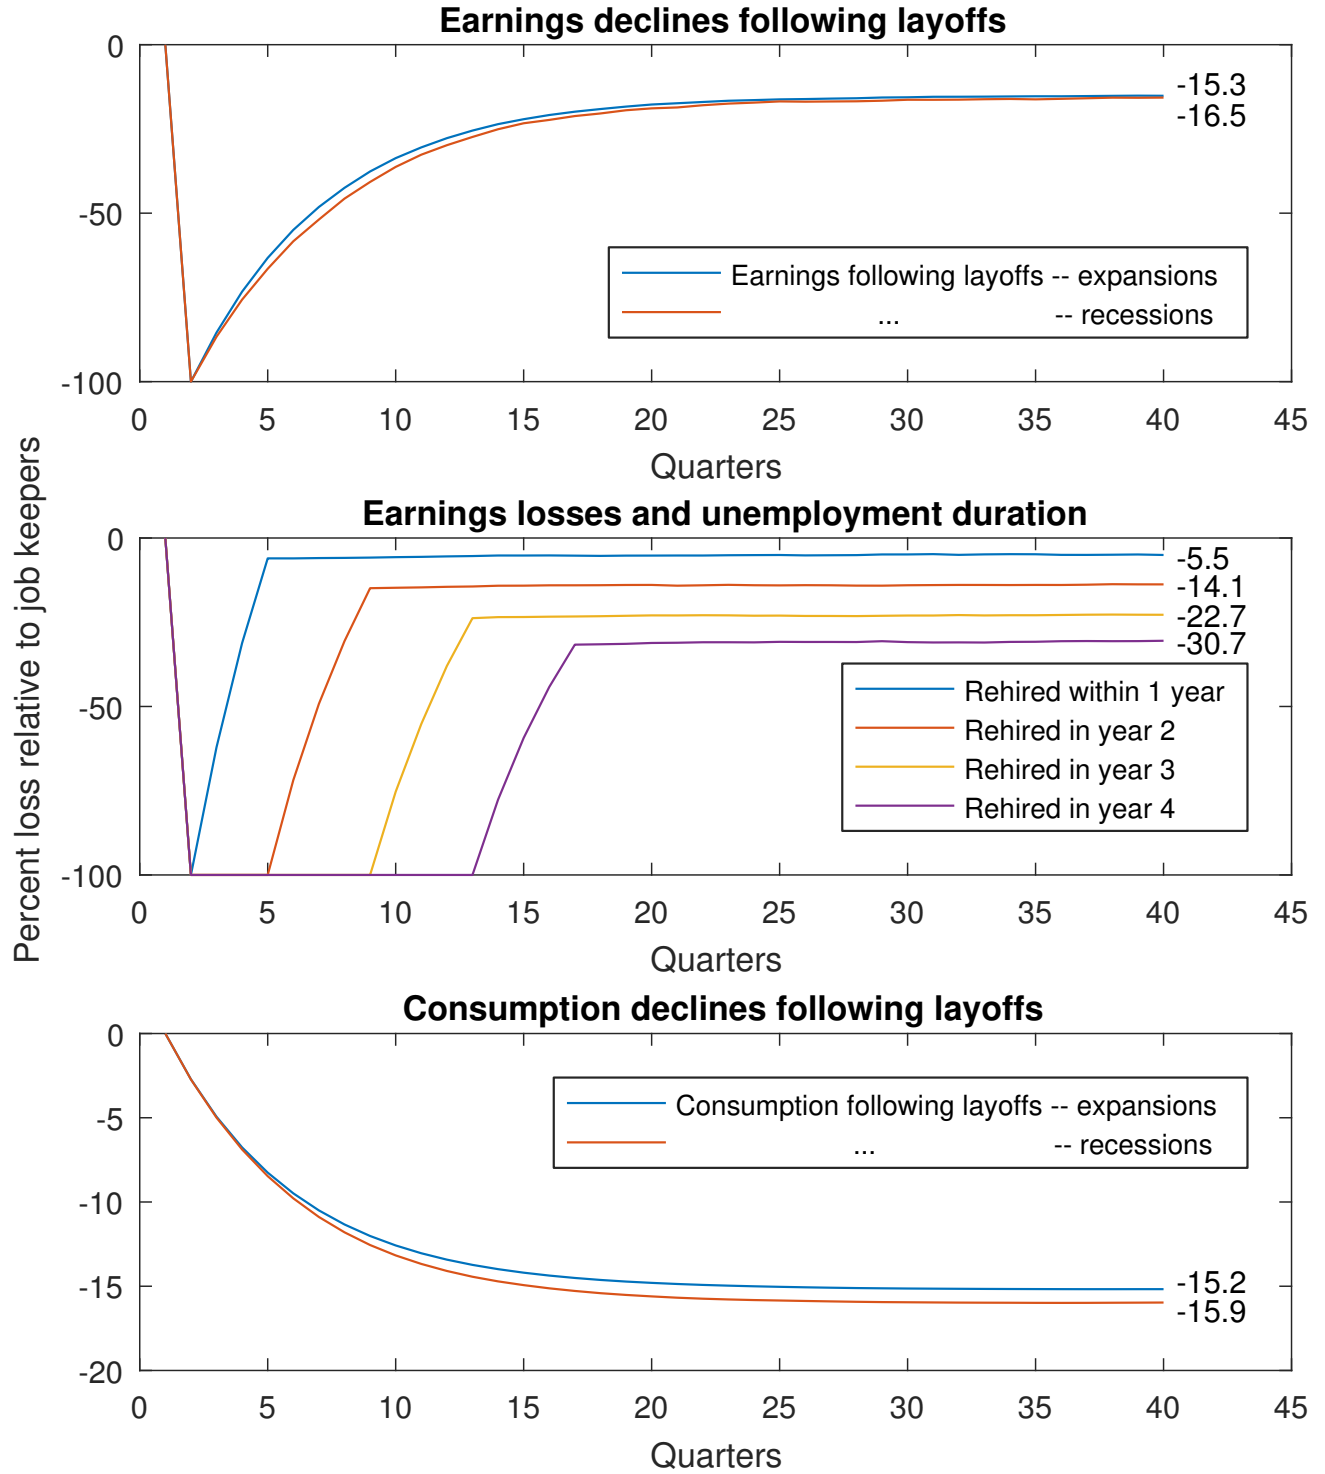
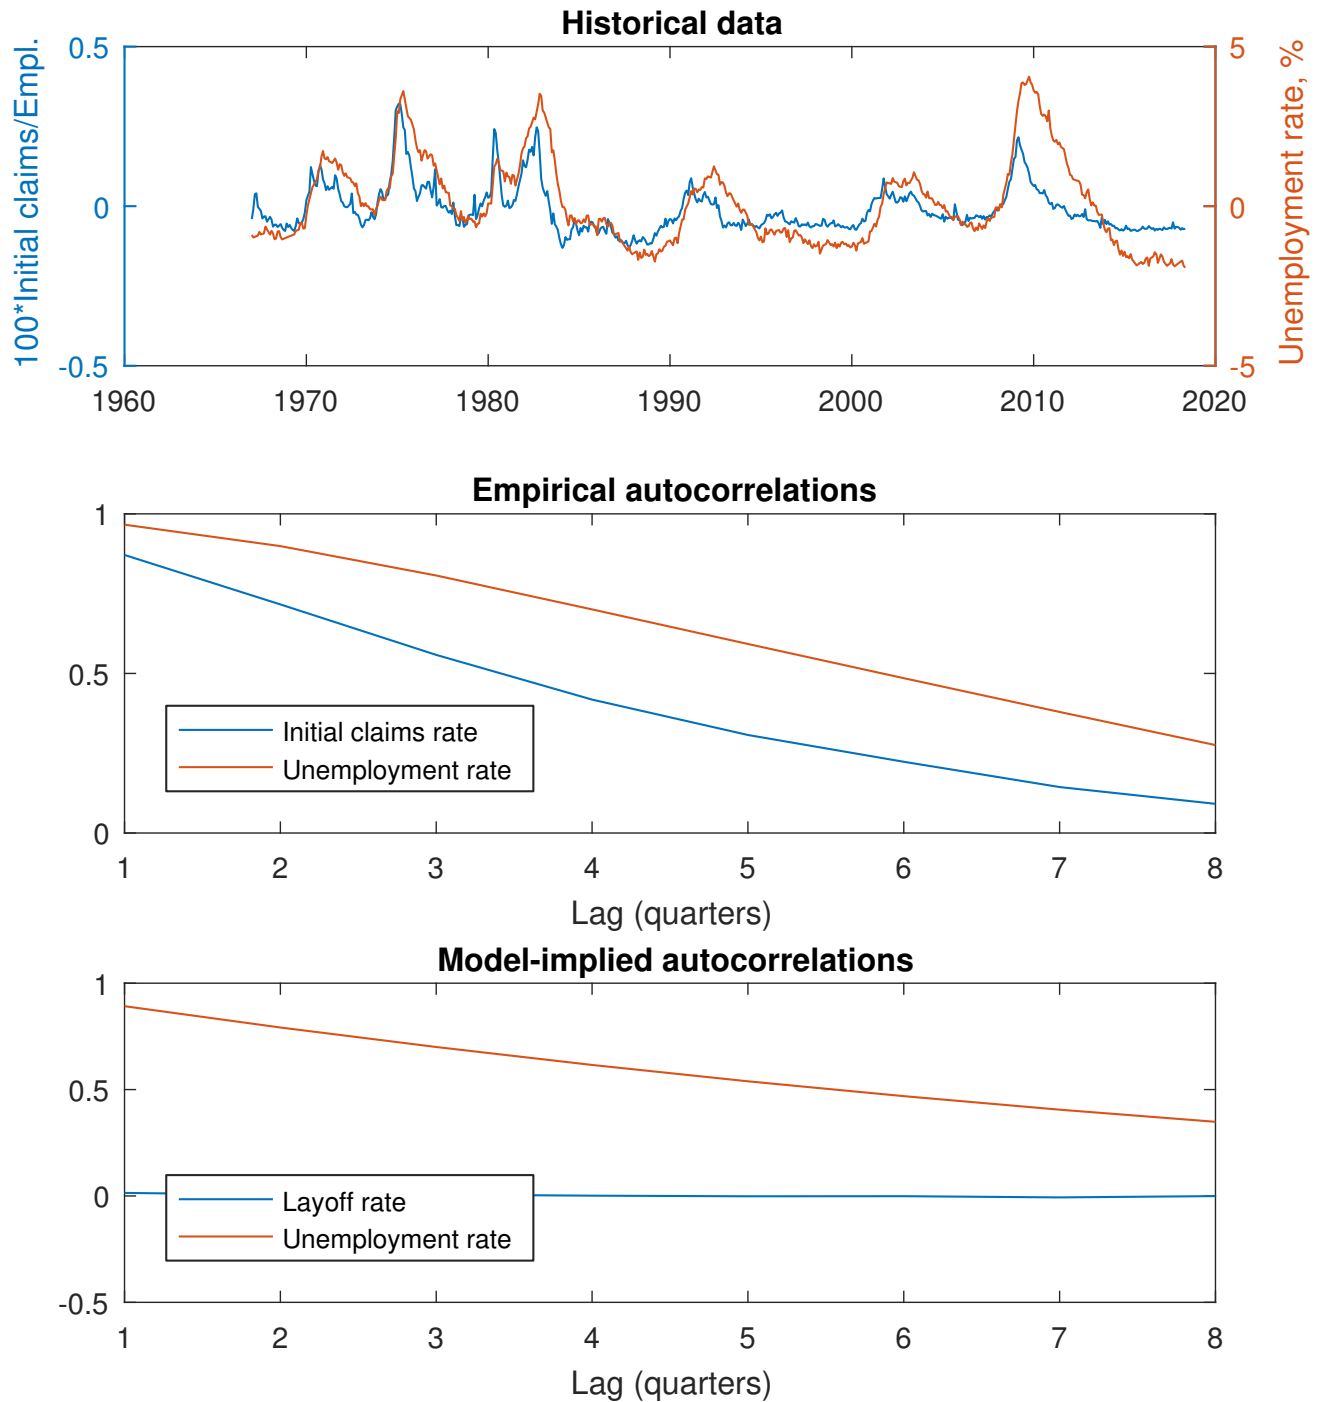


Figure 4: Dynamics following job loss



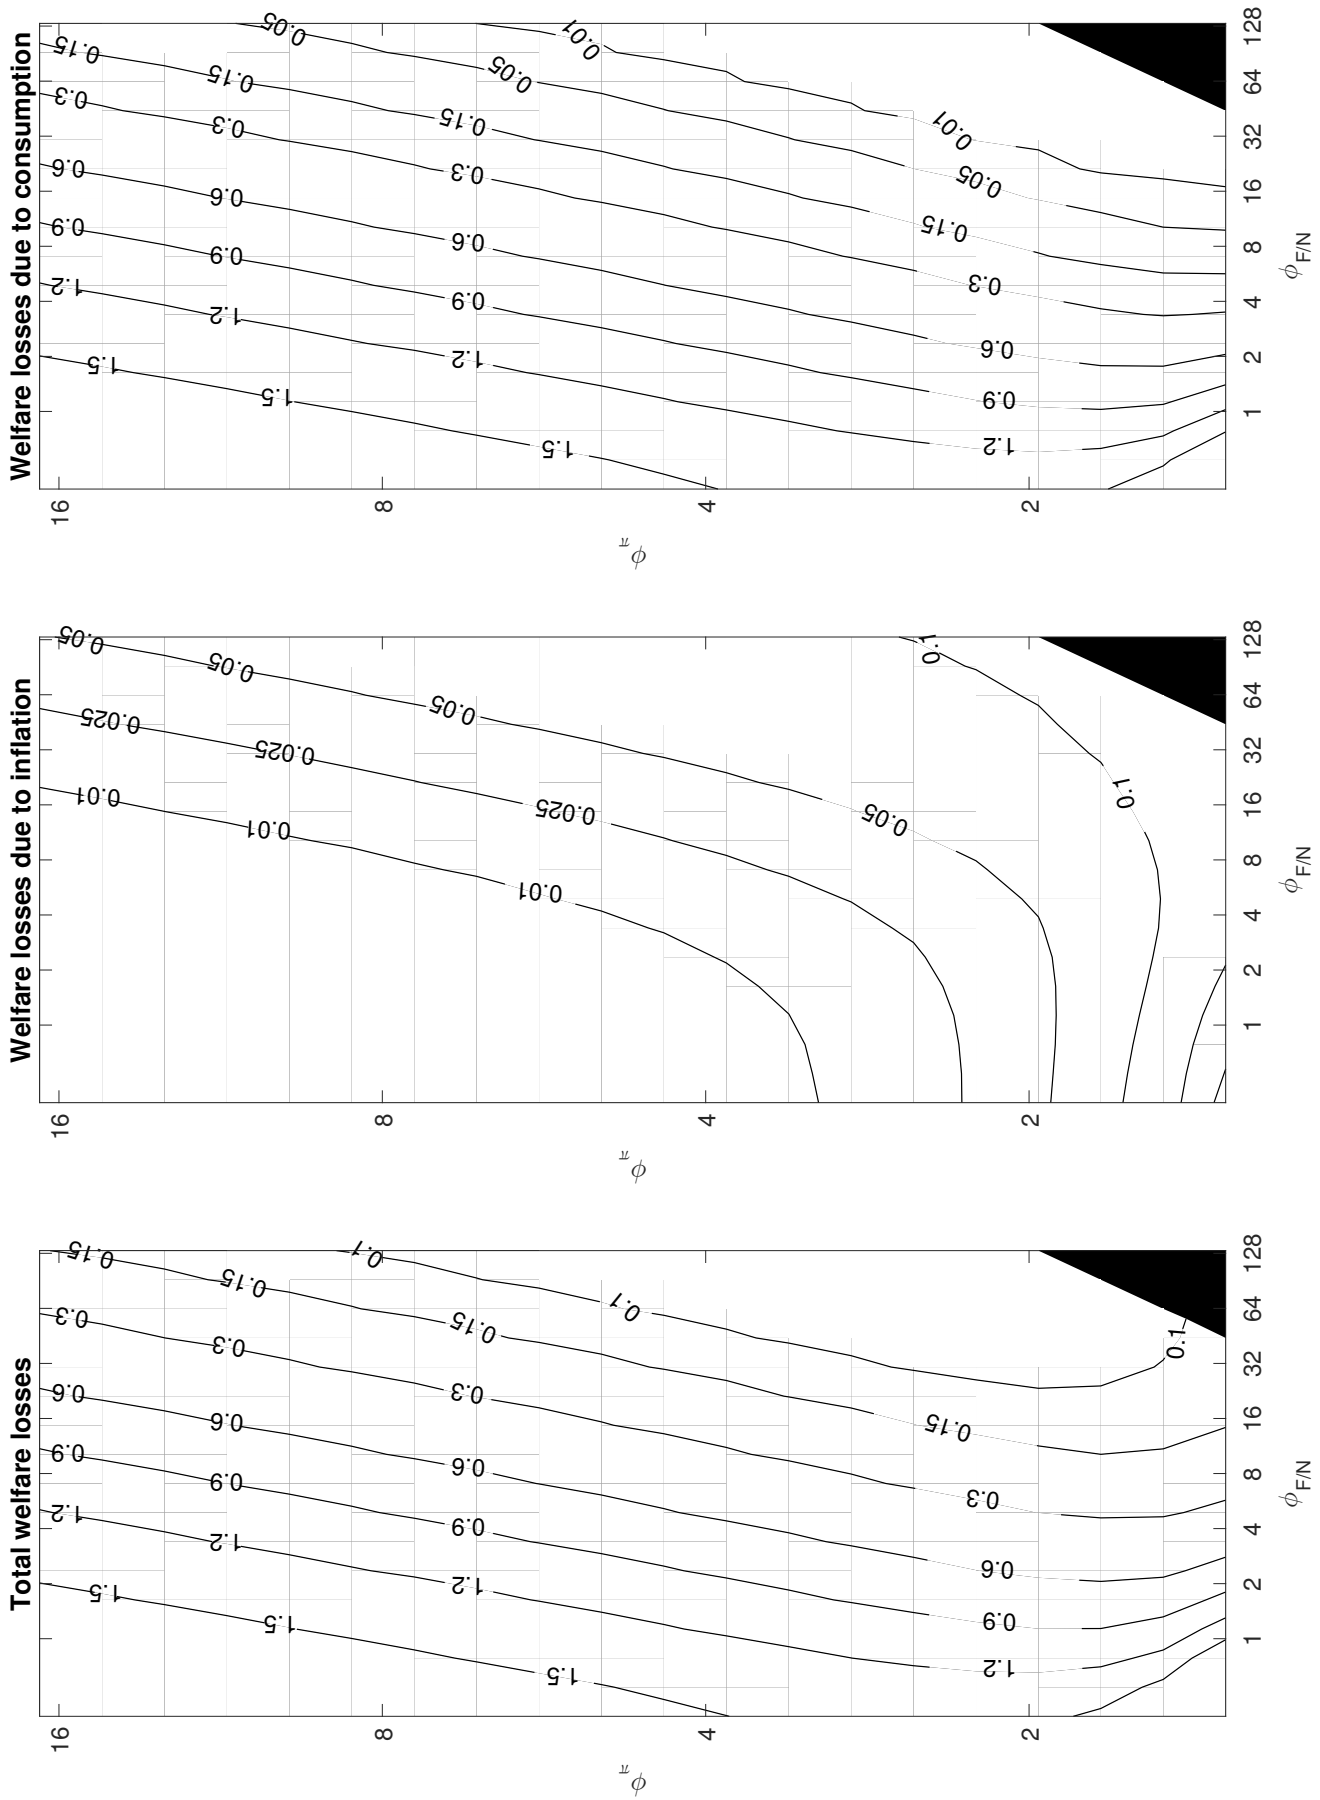
Notes: The panels plot earnings and consumption for workers who are laid off in period zero compared to workers who are not laid off in that period. Recessions are defined as periods where employment is below its fifteenth percentile. Earnings is calculated based only on wage income, not unemployment benefits. The consumption in the bottom panel is total consumption, including both the market and human capital components.

Figure 5: Behavior of unemployment and layoffs



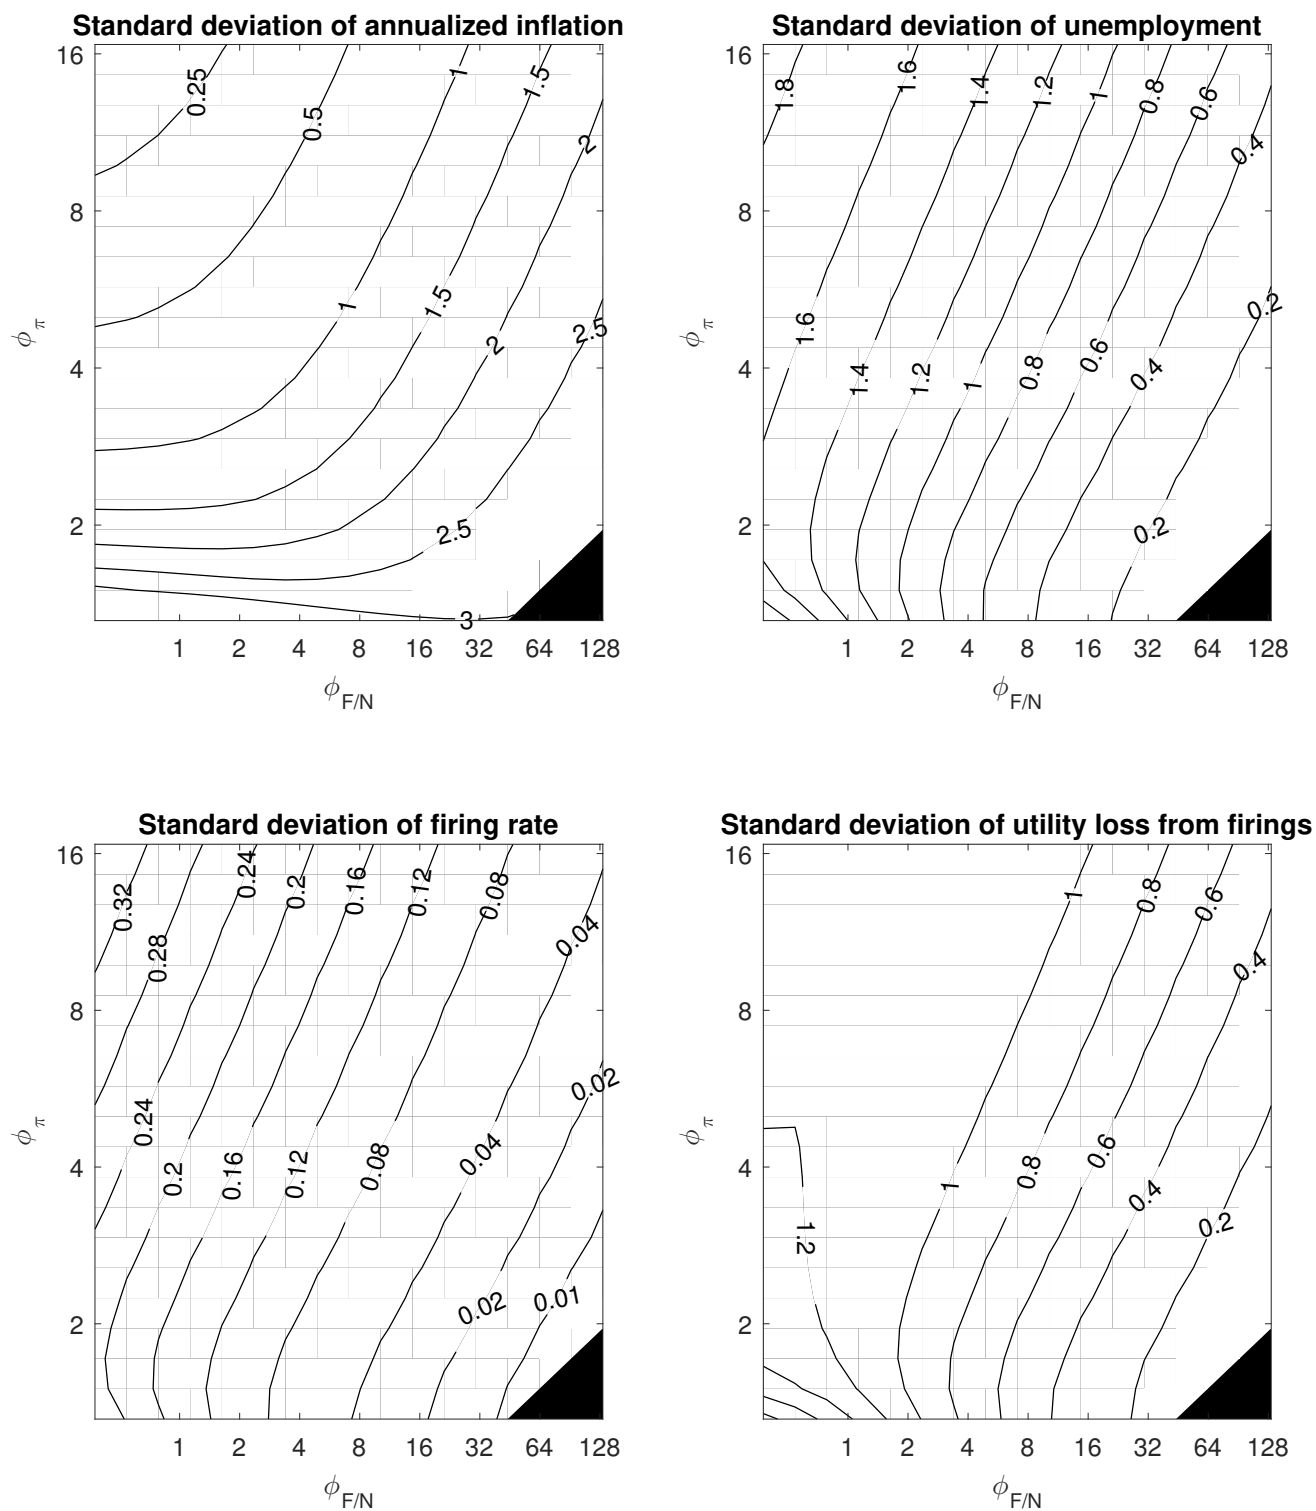
Notes: Unemployment and the initial unemployment claims rate in the data are both detrended with exponentially weighted moving averages with a persistence of 0.98 and the initial value set to minimize the mean squared error. Initial claims and total employment are for the nonfarm private economy. Unemployment in the model is $1 - N_t$ and the layoff rate is F_t/N_t .

Figure 6: Welfare losses across policy rules



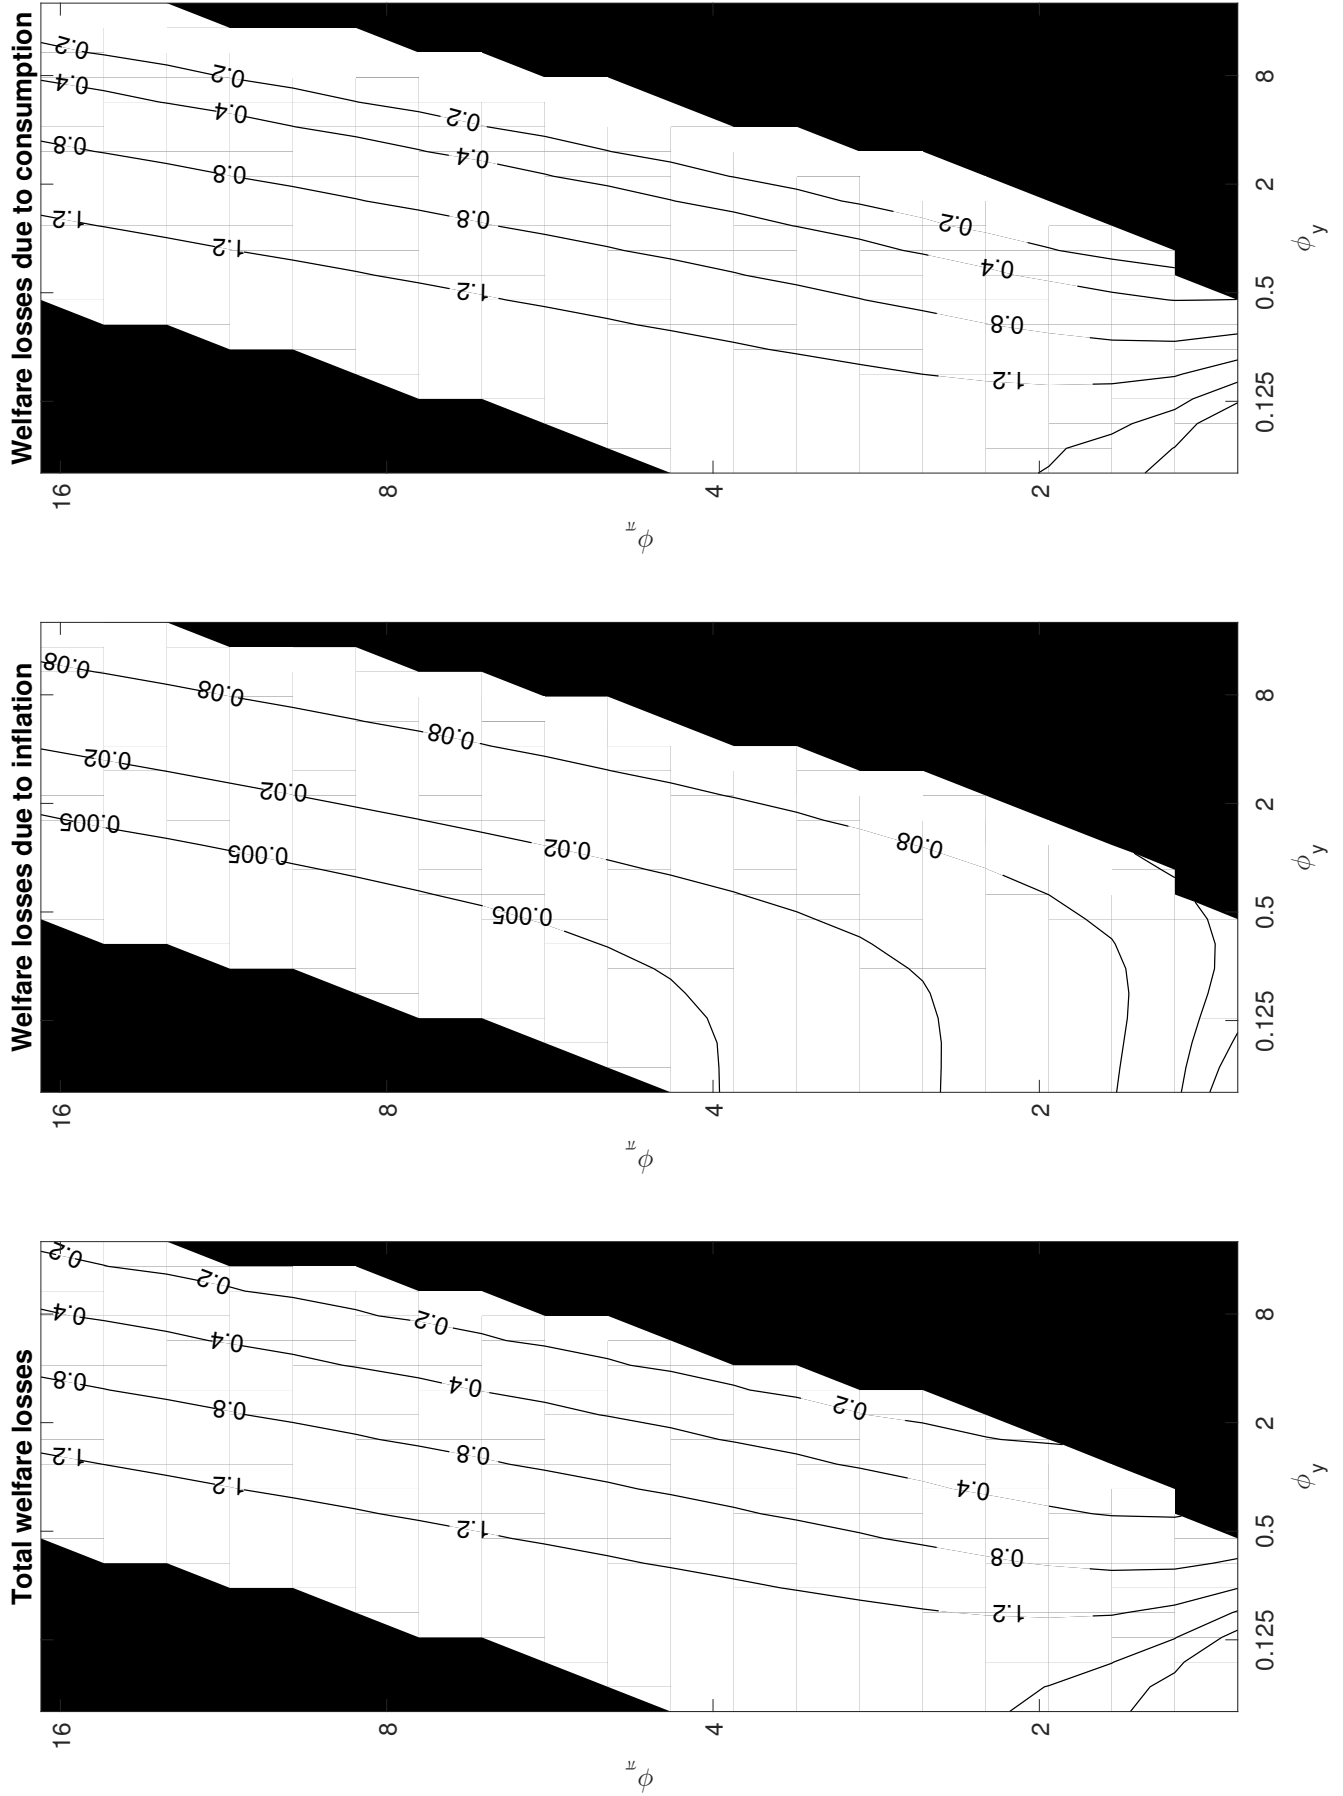
Notes: Welfare losses are in percentage points of lifetime consumption. For all simulations, $\phi_y = 0$. The white regions represent parameter combinations for which the model does not have a stable or determinate solution.

Figure 7: Policy rules and volatilities of endogenous variables



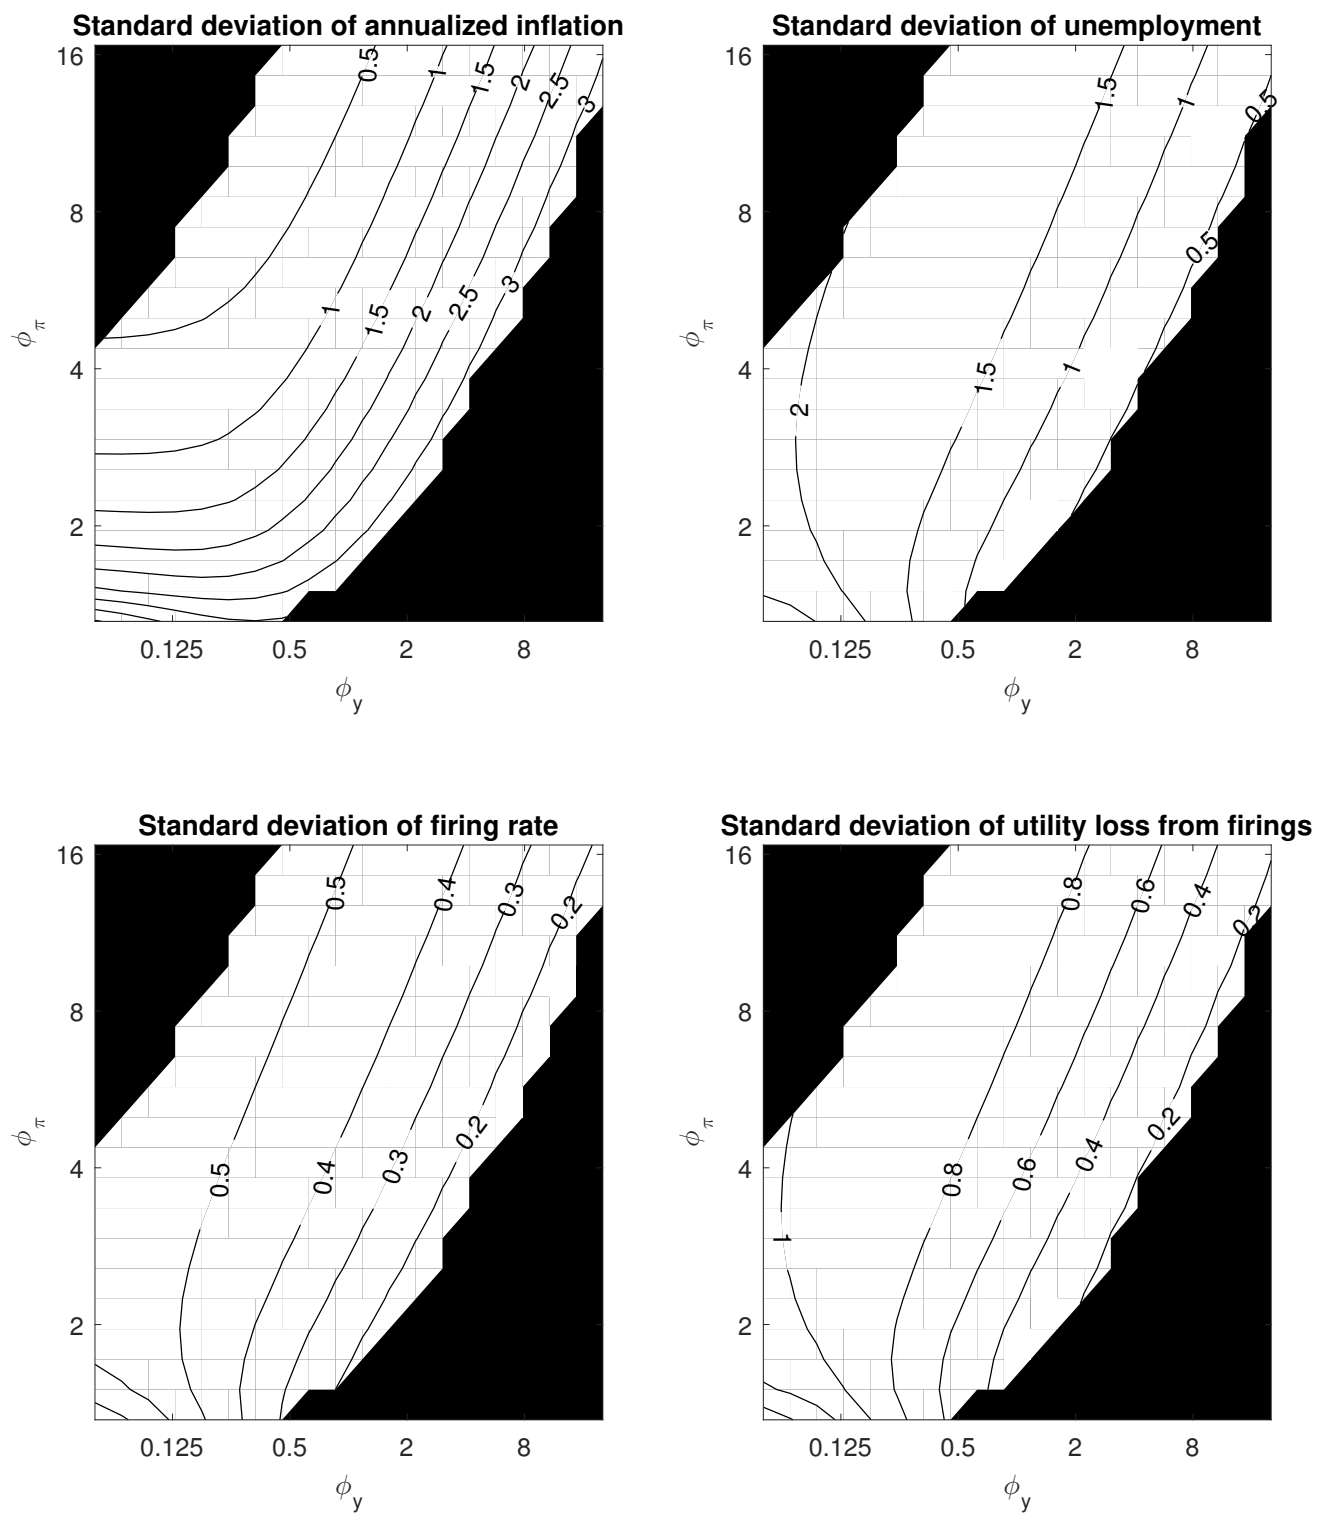
Notes: See figure 6. The utility loss from a firing in each period is equal to $1 - U_t^U / U_t^E$.

Figure 8: Welfare losses across policy rules




Notes: See figure 6. For all simulations, $\phi_{F/N} = 0$.

Figure 9: Policy rules with weight on output and volatilities of endogenous variables



Notes: See figure 7.

Figure 10: Efficient frontiers for $\phi_{F/N}$ and ϕ_y policies



Notes: The panels plot the efficient frontiers among the policies examined in figures 6 and 8. The $\phi_{F/N}$ policies set $\phi_y = 0$ and vice versa..