

A Plucking Model of Business Cycles

Stéphane Dupraz
Banque de France

Emi Nakamura
UC Berkeley

Jón Steinsson*
UC Berkeley

May 18, 2022

Abstract

In standard models, economic activity fluctuates symmetrically around a “natural rate” and stabilization policies can dampen these fluctuations but do not affect the average level of activity. An alternative view—labeled the “plucking model” by Milton Friedman—is that economic fluctuations are drops below the economy’s full potential ceiling. We show that the dynamics of the unemployment rate in the US display a striking asymmetry that strongly favors the plucking model: increases in unemployment are followed by decreases of similar amplitude, while the amplitude of the increase is not related to the amplitude of the previous decrease. We develop a microfounded plucking model of the business cycle. In our model, the plucking asymmetry results from downward nominal wage rigidity, which we rationalize using a search model. Our model also matches the fact that unemployment rises much faster during recessions than it falls during expansions. In our plucking model, stabilization policy lowers average unemployment and thereby yields sizable welfare gains.

Keywords: Downward Nominal Rigidity, Stabilization Policy, Labor Search.

JEL Classification: E24, E30, E52

*We would like to thank Gadi Barlevy, Marco Bassetto, Pierpaolo Benigno, Fabrice Collard, Jordi Gali, François Gourio, Robert Hall, John Leahy, Julien Matheron, Andreas Mueller, Ricardo Reis, Carlos Thomas, Antonella Trigari, and Johannes Wieland for valuable discussions. Emi thanks the National Science Foundation (grant SES-1056107) and the George S. Eccles Research Award in Finance and Economics at Columbia Business School for financial support. Emi and Jon thank the Smith Richardson Foundation for financial support. The views expressed herein are those of the authors and do not necessarily reflect those of the Banque de France or the Eurosystem.

1 Introduction

In the workhorse models currently used for most business cycle analysis, economic activity fluctuates symmetrically around a “natural rate” and stabilization policy does not appreciably affect the average level of output or unemployment. At best, stabilization policy can reduce inefficient fluctuations around the natural rate. As a consequence, in these models the welfare gains of stabilization policy are trivial (Lucas, 1987, 2003).

An alternative view is that economic contractions involve drops below the economy’s full-potential ceiling or maximum level. Milton Friedman proposed a “plucking model” analogy for this view of business cycles: “In this analogy, [...] output is viewed as bumping along the ceiling of maximum feasible output except that every now and then it is plucked down by a cyclical contraction” (Friedman, 1964, 1993).¹ In the plucking model view of the world, improved stabilization policy that eliminates or dampens the “plucks”—i.e., contractions—increases the average level of output and decreases the average unemployment rate. Stabilization policy can therefore potentially raise welfare by substantial amounts (De Long and Summers, 1988; Benigno and Ricci, 2011; Schmitt-Grohe and Uribe, 2016).

We show that the dynamics of the US unemployment rate strongly favor the plucking model of business cycles. An implication of the plucking model—highlighted by Friedman (1964)—is that the dynamics of unemployment should display the following asymmetry: economic contractions are followed by expansions of a similar amplitude—as if the economy is recovering back to its maximum level—while the amplitude of contractions are not related to the previous expansion—each pluck seems to be a new event. We refer to this asymmetry as the plucking property. We present strong evidence that the US unemployment rate displays the plucking property: The increase in unemployment during a contraction forecasts the amplitude of the subsequent expansion one-for-one, while the fall in unemployment during an expansion has no explanatory power for the size of the next contraction.

Our empirical evidence suggests that macroeconomic policy evaluation should be conducted using models that can capture the plucking property. To this end, we consider whether the workhorse Diamond-Mortensen-Pissarides (DMP) search model of the labor market can account for the plucking property we document in the data. We find that the standard version of this

¹The term “plucking” originates in Friedman’s image of a string (output) attached to the underside of a board (potential output): “Consider an elastic string stretched taut between two points on the underside of a rigid horizontal board and glued lightly to the board. Let the string be plucked at a number of points chosen more or less at random with a force that varies at random, and then held down at the lowest point reached.” (Friedman, 1964)

model cannot. The aggregate labor-demand schedule of this model features some non-linearity that can generate a modest amount of the plucking property, but much less than in the data.

To match the facts about plucking that we document in the data, we introduce downward nominal wage rigidity into the DMP model. Earlier research—notably [Benigno and Ricci \(2011\)](#) and [Schmitt-Grohe and Uribe \(2016\)](#)—has highlighted the potential importance of downward nominal wage rigidity for business cycle analysis. Relative to this work, we show how embedding downward nominal wage rigidity within a search model makes it consistent with optimizing behavior, and thus robust to [Barro’s \(1977\)](#) critique that wage rigidity should neither interfere with the efficient formation of employment matches nor lead to inefficient job separations.

We show that our DMP search model augmented with downward nominal wage rigidity can quantitatively match the plucking property we document in the data. Intuitively, our model reproduces the plucking property because good shocks mostly lead to increases in wages, while bad shocks mostly lead to increases in unemployment. The dynamics of the unemployment rate thus become asymmetric; the unemployment rate rises far above its steady state level in response to adverse shocks, but falls much less in response to favorable shocks.

Another important asymmetry in the unemployment rate is that it rises much faster during contractions than it falls during expansions. A DMP-style search and matching model may seem like a good candidate for matching this asymmetry: workers can be fired rapidly but search frictions imply that it takes time to expand employment. However, the standard DMP model does not deliver this result. Job finding is sufficiently rapid in this model that the model has minimal internal propagation. As a consequence, unemployment dynamics are much too rapid and there is no speed asymmetry.

Our plucking model is able to match the asymmetry in the speed of expansions and contractions. To match this we must slow down unemployment dynamics and create bursts of separations during downturns. To slow down unemployment dynamics, we augment the basic DMP model to allow for insecure short-term jobs: most new matches turn out to be poor matches and separate quickly; however, some survive and become stable matches. As a consequence, workers who become unemployed cycle through several jobs before finding stable employment, a pattern emphasized by [Hall \(1995\)](#), [Pries \(2004\)](#), [Krolkowski \(2017\)](#), [Jung and Kuhn \(2019\)](#), [Jarosch \(2021\)](#), and [Hall and Kudlyak \(2021\)](#), among others. The bursts of separations during recession occur when downward nominal wage rigidity starts to bind sufficiently that firms decide to endogenously lay-off some of their long-term workers. Intuitively, wage rigidity during recessions leads

unemployment to shoot up as long-term workers are laid off; and unemployment remains high for much longer than in the standard DMP model since unemployed workers must cycle through several jobs before they find stable employment.

The plucking nature of our model has important normative implications. It implies that fluctuations in unemployment are fluctuations above a resting point of low unemployment, not symmetric fluctuations around a natural rate. As a consequence, a reduction in the volatility of aggregate shocks not only reduces the volatility of the unemployment rate, but also reduces its average level, as in the models of [Benigno and Ricci \(2011\)](#) and [Schmitt-Grohe and Uribe \(2016\)](#). Eliminating all aggregate shocks in our calibrated model reduces the average unemployment rate from 5.7% to 3.8%. The welfare benefits of stabilization policy are an order of magnitude larger in our model than in standard models in which stabilization policy cannot affect the average level of output and unemployment.

Our plucking model implies that a modest amount of inflation can “grease the wheels of the labor market” by allowing real wages to fall in response to adverse shocks even though nominal wages are downward rigid. Increasing the average inflation rate from 2% (our baseline calibration) to 4% yields a drop in average unemployment from 5.7% to 3.9%. The benefits of inflation diminish at higher levels of inflation but are quite large at low levels.

The notion that good shocks mostly lead to increases in wages, while bad shocks mostly lead to increases in unemployment has a long history within macroeconomics going back at least to [Tobin \(1972\)](#). The main theoretical challenge for this line of thinking has been how to justify the notion that wages do not fall in recessions despite obvious incentives of unemployed workers to bid wages down. To make downward nominal wage rigidity robust to this critique, we build on recent insights from the labor search literature. Intuitively, because of search frictions, unemployed workers cannot freely meet with firms and offer to replace employed workers at a lower wage. Instead, unemployed workers and potential employers must engage in a costly matching process. But after the worker and employer have matched, the worker has some monopoly power and therefore no longer has an incentive to bid the wage down. Without further assumptions, wages are not uniquely pinned down ([Hall, 2005](#)). The literature has often assumed Nash bargaining to pin down wages. The asymmetric nature of unemployment dynamics suggests that a bargaining procedure that gives rise to asymmetric wage adjustment is more promising.

Our work is related to several strands of existing literature. [Kim and Nelson \(1999\)](#) and [Sinclair \(2010\)](#) are two of the very few modern attempts to assess the specific asymmetry emphasized by

Friedman (see also Caballero and Hammour, 1998; Bordo and Haubrich, 2012; Fatás and Mihov, 2015). Ferraro (2017) provides an alternative explanation for skewness and the speed asymmetry in the unemployment rate. Our assumption of downward nominal wage rigidity is motivated by microeconomic evidence on the prevalence of wage freezes in US data (see, e.g., McLaughlin, 1994; Kahn, 1997; Card and Hyslop, 1997; Altonji and Devereux, 2000; Kurmann and McEntarfer, 2017; Hazell and Taska, 2018; Grigsby, Hurst, and Yildirmaz, 2018). A large literature has considered models with downward nominal wage rigidity (see, e.g., Akerlof, Dickens, and Perry, 1996; Kim and Ruge-Murcia, 2009; Benigno and Ricci, 2011; Abbritti and Fahr, 2013; Schmitt-Grohe and Uribe, 2016; Chodorow-Reich and Wieland, 2018). The importance of wage rigidity for generating realistic fluctuations in unemployment has been stressed by Shimer (2005), Hall (2005), Gertler and Trigari (2009), and Gertler, Huckfeldt, and Trigari (2016). Bewley (1999) presents survey evidence on why firms don't lower wages in recessions. Olivei and Tenreyro (2010) present evidence that wage rigidity is allocative. Recent work has explored several ways in which Lucas's (1987, 2003) calculations may underestimate the costs of business cycle fluctuations (see, e.g., Imrohoroglu, 1989; Obstfeld, 1994; Tallarini, 2000; Ramey and Ramey, 1991; Barlevy, 2004; Gali, Gertler, and Lopez-Salido, 2007; Barro, 2009; Krusell et al., 2009). Finally, the gains from targeting a positive level of inflation have been reassessed in recent years following the Great Recession (see, e.g., Coibion, Gorodnichenko, Wieland, 2012; Andrade et al. 2016; Blanco 2018).

The paper proceeds as follows. Section 2 presents our empirical results on the asymmetric dynamics of the unemployment rate. Section 3 lays out our plucking model of business cycles. Section 4 shows that a basic DMP version of the model does not generate any of the business cycle asymmetries we document. Section 5 shows how our full model can generate all of these asymmetries. Section 6 shows that stabilizing fluctuations and raising inflation each reduce the average level of unemployment. Section 7 concludes.

2 Three Business-Cycle Asymmetries

We begin our analysis by using data on post-WWII US unemployment to document three business cycle asymmetries. The first of these is Friedman's plucking property: the amplitude of a contraction forecasts the amplitude of the subsequent expansion, while the amplitude of an expansion does not forecast the amplitude of the subsequent contraction. The second asymmetry is that the distribution of the unemployment rate is right-skewed, as emphasized by Sichel (1993).

The third asymmetry is that the unemployment rate rises more quickly than it falls, as emphasized by Neftçi (1984) among others. In addition to these asymmetries, we highlight that business cycle expansions are often very long.

2.1 Defining Expansions and Contractions

To implement our empirical analysis, we need dates of business cycle peaks and troughs. One option is to make use of the dates identified by the NBER Business Cycle Dating Committee. A downside of using these dates is that the NBER dates do not line up exactly with peaks and troughs of the unemployment rate. The reason for this is that the NBER Business Cycle Dating Committee uses a wide variety of cyclical indicators to date turning points. Since our empirical analysis is based on the amplitude and speed of cyclical movements in unemployment, we need dates that line up with turning points for the unemployment rate. For this reason, we develop a simple algorithm that defines business cycle peak and trough dates for the unemployment rate. The basic idea is to find local minima and maxima of the unemployment rate. However, we ignore small “blips” or “wiggles” in the unemployment rate and focus instead on delineating substantial swings in the unemployment rate in a similar manner as the peaks and troughs identified by the NBER Business Cycle Dating Committee. For details on our algorithm, see Appendix A.

Figure 1 plots the unemployment rate over our sample period—which runs from January 1948 until February 2020—with vertical lines indicating the times that we identify as business cycle peaks and troughs. We identify ten peaks and ten troughs. To these we add a peak at the beginning of our sample. A concern regarding this date is that the contraction at the beginning of our sample may have started earlier. We are however reassured on this point by the fact that the NBER identified November 1948 as a peak. We are also confident about our last peak since we end our sample at the onset of the Covid-19 recession. Table A.1 presents the peak and trough dates we identify and compares them with the peak and trough dates identified by the NBER.

2.2 The Plucking Property

Figure 2 presents scatter plots illustrating the plucking property for the unemployment rate. The left panel plots the amplitude of a contraction on the x-axis and the amplitude of the subsequent expansion on the y-axis. The amplitude of contractions is defined as the percentage point increase in the unemployment rate from the business cycle peak to the next trough. The amplitude of

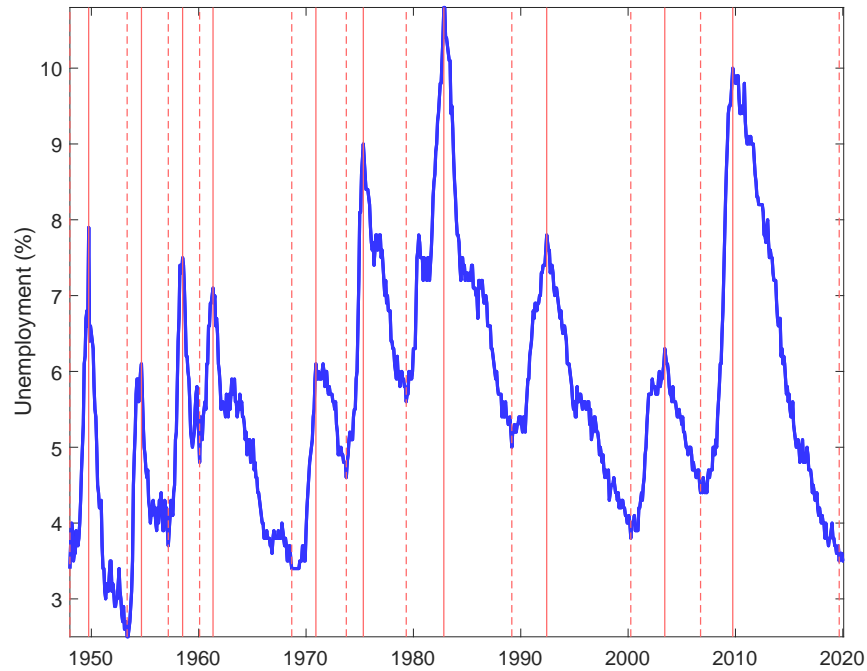


Figure 1: Peaks and Troughs in the Unemployment Rate

Note: The unemployment rate is plotted in blue. Business cycle peaks are denoted by dashed red vertical lines, while business cycle troughs are denoted by solid red vertical lines.

expansions is defined analogously. There is clearly a strong positive relationship between the amplitude of a contraction and the amplitude of the subsequent expansion in our sample period. In other words, the size of a contraction strongly forecasts the size of the subsequent expansion. We have included an OLS regression line in the panel. Table 1 reports the regression coefficient from this regression. The relationship is roughly one-for-one. For every percentage point increase in the amplitude of a contraction, the amplitude of the subsequent expansion increases by 1.12 percentage points on average. Despite the small number of data points, the relationship is highly statistically significant. Furthermore, the explanatory power of the amplitude of the previous contraction is large. The R^2 of this simple univariate regression is 0.59.

The right panel of Figure 2 plots the amplitude of an expansion on the x-axis and the amplitude of the subsequent contraction on the y-axis. In sharp contrast to the left panel, there is no relationship in this case. The size of an expansion does not forecast the size of the next contraction. In Friedman's language, each contractionary pluck that the economy experiences is independent of what happened before. The linear regression line in the panel is actually slightly downward

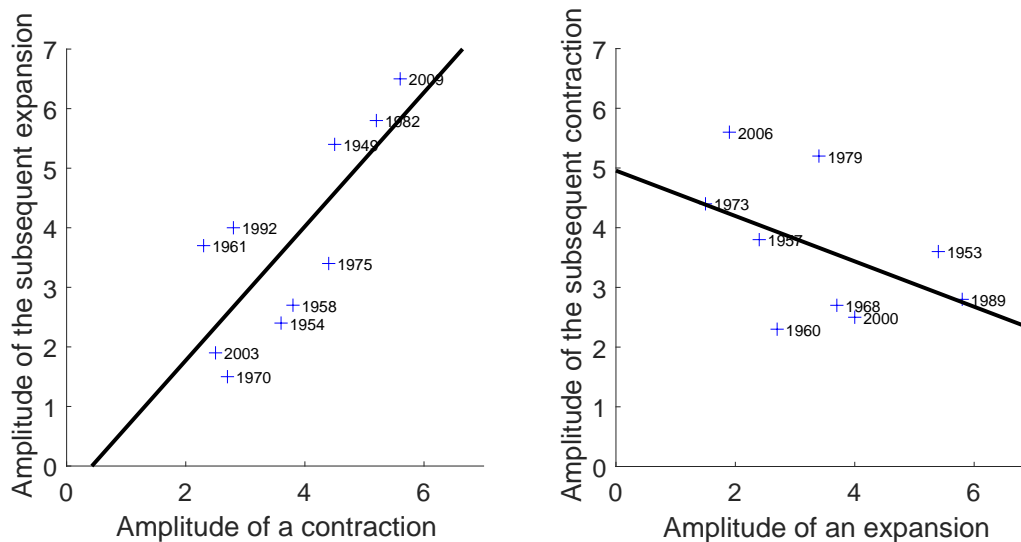


Figure 2: The Plucking Property of the Unemployment Rate

Note: The points in the left panel are labeled with the year the contraction in question ended and expansion in question began. The points in the right panel are labeled with the year the expansion in question ended and contraction in question began. OLS regression lines are plotted in each panel.

sloping. But the association is far from statistically significant and the R^2 of the regression is only 0.22.

Jackson and Tebaldi (2017) suggest that the *duration* (not size) of an expansion is predictive of the size of the following contraction. They motivate this idea by analogy to forest fires: the longer the expansion, the more “underbrush” builds up—e.g., low quality matches and entrants—that becomes fuel in the subsequent contraction. We find no evidence of the forest fire theory at the aggregate level: the duration of an expansion is no more predictive of the size of the following contraction than the size of the expansion is. The relationship is actually negative (but not significantly so), driven by the fact that the three longest post-WWII expansions (1961-1968, 1982-1989, 1992-2000) were followed by relatively mild recessions. Tasci and Zevanove (2019) confirm these results and also present state level results for the plucking model and forest fire theory. Their state level results are similar to our results at the aggregate level: There is strong evidence for the plucking property but no evidence for the forest fire theory.

Table 1: Plucking Property of Unemployment

	β	R^2
Subsequent expansion on contraction	1.12 (0.33)	0.59
Subsequent contraction on expansion	-0.38 (0.27)	0.22

Note: The first row reports the coefficient in an OLS regression of the size of the subsequent expansion (percentage point fall in unemployment rate) on the size of a contraction (percentage point increase in unemployment rate). The second row reports the coefficient in an analogous regression of the size of the subsequent contraction on the size of an expansion.

2.3 Two Additional Unemployment Rate Asymmetries

Table 2 presents results on two additional asymmetries in the unemployment rate. The first of these is the “deepness” asymmetry, which Sichel (1993) defines as the characteristic that “troughs are deeper than peaks are tall.” Sichel suggests assessing this idea by testing for skewness in the distribution of the level of a business cycle variable such as unemployment. The skewness of the unemployment rate over our sample period is 0.64. We can reject a null of no skewness with a p-value of 0.066.²

A second additional asymmetry is that the unemployment rises more quickly during contractions than it falls during expansions. Neftçi (1984) is an early paper to assess this asymmetry statistically.³ A particularly simple way to illustrate this asymmetry is to calculate the average speed of expansions and contractions in percentage points of unemployment per year. Table 2 reports estimates of the average speed of expansions and contractions.⁴ We find that the unemployment rate rises roughly twice as quickly during contractions (1.9 percentage points per year) as it falls during expansions (0.9 percentage points per year). This difference is highly statistically

²The combination of high autocorrelation and non-linearity makes it difficult to calculate a confidence interval for the skewness of the unemployment rate. For this reason, we focus on testing against a null of zero skewness. We construct the reported p-value by estimating an AR(2) model for the unemployment rate and then bootstrapping this model to construct a distribution for estimates of skewness under the null of zero skewness. Before estimating the AR(2), we apply a three-period moving-average filter to smooth the very-high-frequency variations in the unemployment rate. This makes sure the results are similar to those obtained on (non-smoothed) quarterly data. In section 5, we report a simulated distribution for the skewness of the unemployment rate from a much more sophisticated model.

³Sichel (1993) refers to this asymmetry as the “steepness” asymmetry, while McKay and Reis (2008) refer it as the greater “violence” of contractions, in reference to Mitchell (1927). Given that contractions and expansions are of about the same average size (3.7 percentage points), the fact that contractions are “steeper” or “more violent” than expansions is equivalent to the fact that they are briefer. Awareness of this asymmetry dates back at least to the 1920s. Mitchell (1927) notes that “business contractions appear to be briefer and more violent than business expansions.”

⁴For each expansion and contraction, we calculate the change in unemployment over the spell and the length of time the spell lasts for. The speed for that expansion or contraction is the ratio of those two numbers. We then take a simple average across all expansions and separately a simple average across all contractions.

Table 2: Further Unemployment Asymmetries

Skewness of unemployment	0.64
P-value for no skewness	0.066
Speed of expansions (pp/year)	0.87
Speed of contractions (pp/year)	1.89
P-value for equal speed	0.002
Duration of expansions (months)	59.1
Duration of contractions (months)	26.9

Note: We report the skewness of the distribution of the level of the unemployment rate. The speed of expansions and contractions is measured in percentage points of unemployment per year.

significant. We run a regression of the absolute value of the speed of expansions and contractions on a dummy variable for a spell being a contraction and find that the p-value for the dummy is 0.002.

Looking back at Figure 1, we can clearly see that when the unemployment rate starts falling, it usually falls relatively steadily for a long time. As a consequence, expansions are quite long. The average length of expansions in our sample is roughly 59 months, or almost five years. Contractions are also quite persistent, but less so. The average length of contractions in our sample is roughly 27 months, a bit more than two years. Perhaps most strikingly, in a few cases—the 1960s, 1980s, 1990s, and 2010s—the unemployment rate has fallen steadily for six to ten years without reversal. We will argue later in the paper that these long and steady expansions and contractions place interesting restrictions on the types of models or shock processes that can drive business cycles.

3 An Equilibrium Model of Downward Nominal Wage Rigidity and Insecure Short-Term Jobs

What type of model of the labor market can generate the asymmetries we document in section 2? We show below that the canonical Diamond-Mortensen-Pissarides (DMP) model with symmetric real wage rigidity to address the [Shimer \(2005\)](#) puzzle cannot. For this reason, we develop an augmented version of the DMP model that incorporates two additional features: downward nominal wage rigidity and insecure short-term jobs through which unemployed workers must

transition (in some cases quite a few times) before securing a long-term, stable job. We show that incorporating downward nominal wage rigidity yields the plucking property, but both features are needed to also match the asymmetric speed of expansions and contractions. Our full model also features sectoral heterogeneity, which allows downward nominal wage rigidity to bind more or less stringently (i.e., for a larger or smaller fraction of sectors), and to bind for some firms even when aggregate productivity growth is quite positive on average (Benigno and Ricci, 2011). In this section, we present our model in its most general form.

3.1 Heterogeneous Labor Inputs

There is a continuum of labor types denoted by $i \in [0, 1]$. To each labor type i corresponds a sector i in which identical firms (or equivalently, a representative firm) have access to a decreasing-returns-to-scale technology that uses labor type i as its single input. The production function in sector i is

$$Y_t^i = A_t^i F(N_t^i), \tag{1}$$

where Y_t^i is output, N_t^i is employment, and A_t^i is an exogenous productivity shifter. The productivity shifter should be viewed as a stand-in for any shifter of labor demand in sector i . Michailat (2012) emphasizes the importance of decreasing-returns-to-scale for the business cycle behavior of DMP search models. For simplicity, we restrict sectoral heterogeneity to labor markets: consumers perceive goods produced in different sectors as identical and therefore value them equally. All goods are sold in a competitive product market at a common price P_t .

A given worker provides a particular type of labor, and can therefore only seek to work at a firm in one sector. This implies that there is a distinct labor market for every type of labor and workers cannot flow across labor markets. We think of these labor types as occupations in a particular location, e.g., lawyers in Houston. Switching occupations is difficult due to occupation-specific human capital. Mobility constraints limit the willingness of workers to switch locations.⁵

3.2 Labor Demand with Insecure Short-Term Jobs

A firm in sector i starts period t with the stock of workers it inherits from the previous period, denoted by M_t^i . These are workers securely attached to the firm: at the end of each period only

⁵The assumption of such differentiated labor inputs is standard in the New-Keynesian literature. There, differentiated labor inputs are an important source of strategic complementarity in price setting. See, e.g., Woodford (2003, ch. 3).

a small fraction $\delta \in (0, 1)$ of them separate from the firm for exogenous reasons. Before starting production, the firm typically hires workers, in number H_t^i . These workers start working at time t , so that the level of employment at the firm at t is

$$N_t^i = M_t^i + H_t^i. \quad (2)$$

Newly-hired workers separate from the firm at the end of period t at a higher rate $d \geq \delta$. If they do not separate at the end of period t , however, they join the pool of securely attached workers and face the low exogenous separation rate δ from period $t + 1$ onward. The stock of workers attached to the firm at the beginning of the next period is then

$$M_{t+1}^i = (1 - \delta)M_t^i + (1 - d)H_t^i \quad \text{if } H_t^i \geq 0. \quad (3)$$

The difference between the separation rates of securely attached workers and newly-hired workers captures the fact that many newly formed matches turn out to be poor matches for various reasons and are therefore terminated quickly. An unemployed worker will typically transition between several jobs with intervening unemployment spells before one of these jobs turns out to be a good match and thus turns into a long-term, secure position. The fact that newly hired workers face much higher separation rates than workers with longer tenure has been emphasized by [Hall \(1995\)](#), [Pries \(2004\)](#), [Krolikowski \(2017\)](#), [Jung and Kuhn \(2019\)](#), [Jarosch \(2021\)](#), and [Hall and Kudlyak \(2021\)](#), among others. The fact that there are effectively only two types of jobs in our model—very short-term jobs that last one period and very secure jobs—is a simplifying assumption that we make for tractability. This assumption alleviates the need to keep track of the stock of short-term workers working at the firm.

We allow for the firm not to hire in period t and instead endogenously lay off securely-attached workers, above and beyond the fraction δ that exogenously left at the end of period $t - 1$. In this case, $H_t^i < 0$ and all the workers N_t^i who work at the firm in period t are securely attached, so that the firm starts the next period with a number of workers

$$M_{t+1}^i = (1 - \delta)N_t^i \quad \text{if } H_t^i \leq 0. \quad (4)$$

We assume that the firm pays its securely attached and newly-hired workers the same wage, for instance, because paying them differently would adversely affect morale at the firm ([Bewley, 1999](#)). We denote the nominal wage in sector i by W_t^i and the real wage by $w_t^i = W_t^i/P_t$. To hire workers, the firm must post vacancies. Posting a vacancy costs cA_t^i units of goods, where c is a

constant. A vacancy translates into a hire if it matches with a job-seeker. A match happens with probability q_t^i , which firm i takes as given. Firms are large enough that they can abstract from the randomness associated with hiring new workers: hiring one worker requires the firm i to post $1/q_t^i$ vacancy and has the certain real cost $A_t^i c/q_t^i$.

Firm i 's real profits at time t are

$$A_t^i F(N_t^i) - w_t^i N_t^i - \frac{A_t^i c}{q_t^i} H_t^i I |_{H_t^i \geq 0}. \quad (5)$$

The first term in this expression is revenue; the second term is the cost of labor; and the third term is hiring costs. Like all agents in the model, the firm is risk-neutral and discounts the future with a factor $\beta \in (0, 1)$. The firm is forward-looking and chooses how many workers to hire to maximize intertemporal real profits

$$\Pi_t^i(M_t^i) = \max_{H_t^i} \{A_t^i F(M_t^i + H_t^i) - w_t^i (M_t^i + H_t^i) - \frac{A_t^i c}{q_t^i} H_t^i 1_{H_t^i \geq 0} + \beta E_t(\Pi_{t+1}^i(M_{t+1}^i))\}, \quad (6)$$

subject to the law-of-motion of its workforce $M_{t+1}^i(N_t^i)$ given by equations (3)-(4).

Let $J_t^{i*}(M_t^i) = \frac{\partial \Pi_t^i(M_t^i)}{\partial M_t^i}$ denote the equilibrium marginal value to the firm of a long-term worker inherited from the previous period. For a level of employment N_t^i —not necessarily equal to M_t^i —the marginal value to the firm of a worker already attached to the firm is

$$J_t^i(N_t^i) \equiv A_t^i F'(N_t^i) - w_t^i + \beta(1 - \delta) E_t(J_{t+1}^{i*}(M_{t+1}^i(N_t^i))) \quad (7)$$

and the marginal value to the firm of a newly-hired worker is

$$J_t^{i,new}(N_t^i) \equiv A_t^i F'(N_t^i) - w_t^i + \beta(1 - d) E_t(J_{t+1}^{i*}(M_{t+1}^i(N_t^i))). \quad (8)$$

Each is equal to the marginal product of labor net of the real wage, plus the continuation value discounted with the appropriate separation rate. Notice also that $J_t^i(N_t^i) \geq J_t^{i,new}(N_t^i)$.

Consider a firm that starts off period t with M_t^i workers inherited from the previous period. This firm will hire additional workers $N_t^i > M_t^i$ up to a point where

$$J_t^{i,new}(N_t^i) = \frac{A_t^i c}{q_t^i}, \quad (9)$$

if at the inherited level of employment M_t^i the marginal value of a newly-hired worker is greater than the hiring cost $J_t^{i,new}(M_t^i) > \frac{A_t^i c}{q_t^i}$. The firm will freeze employment

$$N_t^i = M_t^i. \quad (10)$$

if at the inherited level of employment M_t^i the marginal value of an already-attached worker is positive but the marginal value of a newly-hired worker is less than the hiring cost, i.e., $J_t^i(M_t^i) \geq 0$ and $J_t^{i,new}(M_t^i) \leq \frac{A_t^i c}{q_t^i}$. The firm will endogenously fire some of its existing workers $N_t^i < M_t^i$ up to

$$J_t^i(N_t^i) = 0. \quad (11)$$

if at the inherited level of employment M_t^i the marginal value of an already-attached worker is negative $J_t^i(M_t^i) < 0$. When already-hired workers and newly-hired workers face the same separation rate $\delta = d$, then $J_t^{i,new} = J_t^i$ and labor demand reduces to the labor demand of the standard DMP model. Appendix B.1 provides a derivation.

3.3 Workers

There is a fixed supply of workers in each sector, equal across sectors. We abstract from intensive-margin labor-supply decisions by workers. Workers supply (or at least try to supply) an exogenous quantity of labor—which we normalize to 1. Workers in sector i can be in one of three states at t . A number $N_{0,t}^i \equiv \min(N_t^i, M_t^i)$ are employed in a secure job. A number $N_t^i - N_{0,t}^i$ are newly employed. The remaining $U_t^i \equiv 1 - N_t^i$ are unemployed. We denote $U_{t,0}^i \equiv 1 - N_{t,0}^i$ the number of workers who are either unemployed or employed in a short-term job, and refer to it as the “broad” unemployment rate. $U_{t,0}^i$ is also the number of job-seekers in our model (see below). The N_t^i employed workers—either in short-term jobs or secure jobs—earn the real wage w_t^i and forego leisure relative to being unemployed, which they value as $A_t^i \zeta$ units of consumption. Unemployed workers earn unemployment benefits $A_t^i b$.⁶

Unemployed workers transition to being newly employed with the job-finding probability f_t^i , and remain unemployed with probability $1 - f_t^i$. We assume that employed workers who lose their jobs between periods $t - 1$ and t get a chance to find a new job at the beginning of period t and therefore to work in period t , spending no period unemployed. Securely employed workers separate from their jobs with a probability δ_t^i . (The probability δ_t^i is equal to δ when firms hire workers, but it can be higher when firms endogenously lay off workers). They therefore transition to unemployment with a probability $\delta_t^i(1 - f_t^i)$, transition to being newly hired with probability $\delta_t^i f_t^i$, and remain securely employed with a probability $(1 - \delta_t^i)$. Newly employed

⁶Making unemployment benefits and the utility of leisure proportional to productivity allows us to consider non-stationary sectoral productivity shocks without increasing the size of the state-space beyond what is computationally feasible.

workers separate from their jobs with a probability d_t^i . (The probability d_t^i is equal to d when firms hire workers, but it can be higher when firms endogenously lay off workers). They therefore transition to being securely employed with a probability $1 - d_t^i$, transition to unemployment with a probability $d_t^i(1 - f_t^i)$, and remain newly employed—in a new job—with probability $d_t^i f_t^i$.

Workers, like firms, are risk-neutral. Their valuations of being unemployed U_t^i , securely employed W_t^i , and newly employed $W_t^{i,new}$ solve the recursive equations

$$U_t^i = A_t^i b + \beta E_t \left((1 - f_{t+1}^i) U_{t+1}^i + f_{t+1}^i W_{t+1}^{i,new} \right), \quad (12)$$

$$W_t^i = w_t^i - A_t^i \zeta + \beta E_t \left(\delta_{t+1}^i (1 - f_{t+1}^i) U_{t+1}^i + \delta_{t+1}^i f_{t+1}^i W_{t+1}^{i,new} + (1 - \delta_{t+1}^i) W_{t+1}^i \right), \quad (13)$$

$$W_t^{i,new} = w_t^i - A_t^i \zeta + \beta E_t \left(d_{t+1}^i (1 - f_{t+1}^i) U_{t+1}^i + d_{t+1}^i f_{t+1}^i W_{t+1}^{i,new} + (1 - d_{t+1}^i) W_{t+1}^i \right). \quad (14)$$

When already-hired workers and newly-hired workers face the same separation rate $\delta = d$, then $W_t^{i,new} = W_t^i$ and the model reduces to the standard DMP model.

3.4 Matching Function

The probability of filling a vacancy q_t^i is determined in equilibrium by a matching function $q(\theta_t^i)$, where $\theta_t^i = \max(0, H_t^i/q_t^i)/U_{0,t}^i$ denotes labor market tightness in labor-market i . Labor market tightness is the ratio of the number of vacancy posted $\max(0, H_t^i/q_t^i)$ to the number of job-seekers at the beginning of the period. Our assumption that workers who separate from their jobs at $t - 1$ get a chance to find a new job at the beginning of t implies that the pool of job-seekers at the beginning of period t is all workers not securely employed at t , $U_{0,t}^i$. The matching function also determines the probability for an unemployed worker of type i of finding a job: it is equal to the ratio of hires to job-seekers $f(\theta_t^i) \equiv \max(0, H_t^i)/U_{0,t}^i = \theta_t^i q(\theta_t^i)$.

When firms hire in sector i , they do not endogenously lay off workers so $N_{0,t}^i = M_t^i$ and the number of job-seekers is equal to $U_{0,t}^i = 1 - M_t^i$. The job-finding rate in sector i is therefore the following function of employment N_t^i and the inherited stock of employment M_t^i

$$f_t^i = \max \left(0, \frac{N_t^i - M_t^i}{1 - M_t^i} \right). \quad (15)$$

3.5 Wage Band

Search frictions imply that unemployed workers cannot instantly meet with firms and offer to replace employed workers at a lower wage. Rather, an unemployed worker can only meet a firm

by engaging in a costly search process. The effort involved in matching implies that workers have some monopoly power once they have matched with a firm and, therefore, no longer have an incentive to bid the wage down. This logic implies that there is no reason why wages would be driven down to their market-clearing level. In other words, nothing forces the equilibrium to be at the crossing of the labor-demand curve—equation (9), (10), or (11) depending on whether firms in the sector are hiring, freezing employment, or firing workers—and labor-supply curve $N_t^i = 1$.

A reasonable restriction on an equilibrium wage however is that it is privately efficient. We call a wage privately efficient if there is no other wage that would make both the firm and workers better off. A privately efficient wage is equivalently a wage immune to the Barro (1977) critique that optimizing individuals should always be able to find a mutually advantageous alternative. As shown by Hall (2005), private efficiency does not uniquely pin down the wage in a search model but only yields an upper and lower bound on possible equilibrium wages.⁷

We derive the wage band of our model under our assumption that all workers in a firm must be paid the same wage. We furthermore assume that the firm has a clearly understood, deterministic policy for the order in which workers will be laid off (e.g., last-in-first-out). Conflicting interests among workers in a firm occur when a high wage would lead to some workers being laid off but not others: laid off workers would rather remain employed at a lower wage, while those not slated to be laid off are better off with an unchanged wage. We assume that workers in a firm support an unchanged wage if a fraction \mathcal{M} of them are better off with an unchanged wage than with a lower wage. We assume that \mathcal{M} is large, i.e., wage rigidity is only allowed if a large fraction of workers support it (see below for details).⁸

Under this assumption, the firm may on occasion lay off some workers rather than lower the wages of all its workers without conflicting with private efficiency. Specifically, it implies that a privately efficient wage (in the sense of being supported by the firm and a fraction \mathcal{M} of its workers) must be lower than the upper-bound

$$\bar{w}_t^i(\mathcal{M}) = A_t^i F'(\mathcal{M}M_t^i) + \beta E_t((1 - \delta)J_{t+1}^{i*}(M_{t+1}^i(\mathcal{M}M_t^i))). \quad (16)$$

In simulations of our full model with downward nominal wage rigidity, endogenous separations

⁷We consider the case where the total surpluses of matches are positive, $S_t = V_t^i + J_t^i \geq 0$, $S_t^{i,new} = V_t^{i,new} + J_t^{i,new} \geq 0$.

⁸Standard Pareto efficiency corresponds to unanimity voting, i.e., $\mathcal{M} = 0$. However, we consider $\mathcal{M} > 0$, i.e., a case where it is more difficult for workers to oppose a wage cut. Whenever the production function satisfies the Inada condition $\lim_{N \rightarrow 0} F'(N) = \infty$, the upper-bound on the wage band under unanimity voting is infinite: because the marginal productivity of the worker least at risk of being laid off tends to infinity when the workforce shrink to zero, he always retains his job and always vetoes a wage cut.

are consistent with private efficiency for a very substantial super-majority almost always (i.e., \mathcal{M} substantially larger than 0.5).⁹ Appendix C discusses the role of two assumptions in this result: that workers must be paid the same wage, and that there are decreasing returns to scale in production.

A privately efficient wage must also be such that $V_t^{i,new} \equiv W_t^{i,new} - U_t^i \geq 0$. Otherwise workers would opt for their outside option of unemployment and both the firm and all its workers would be better off with a higher wage that allows all matches to be maintained.¹⁰ This implies that a privately efficient wage must be greater than the lower-bound

$$\underline{w}_t^i = A_t^i z - \beta E_t \left((1 - d_{t+1}^i) \left(V_{t+1}^i - f_{t+1}^i V_{t+1}^{i,new} \right) \right), \quad (17)$$

where $z = b + \zeta$ is the flow value of unemployment in terms of both unemployment benefits and leisure.

3.6 Wage-Setting

Further assumptions are needed to uniquely pin down the equilibrium wage. Nash bargaining has been a popular assumption in the literature, as have various other bargaining protocols (e.g., [Hall and Milgrom, 2008](#)). However, little direct evidence exists about how firms and workers bargain. Bargaining protocols are therefore usually judged by a combination of their theoretical appeal and the wage process that they generate (e.g., the Hall-Milgrom protocol yields wages that are less sensitive to unemployment). An alternative approach is to directly specify a wage process ([Blanchard and Gali, 2010](#); [Shimer, 2010](#); [Michaillat, 2012](#)). An appeal of this approach is that one can more easily investigate what features the wage process needs to have to be consistent with the real-world behavior of unemployment. (Once this is understood, it can be an input into research on bargaining protocols.)

We take this second approach and consider two wage-setting rules: a rule with symmetric real wage rigidity and a rule with downward nominal wage rigidity. In both cases, wage rigidity is defined relative to a flexible-wage benchmark that we take to be the Nash bargained wage. Specifically, it is the wage bargained by the firm and a securely attached worker, assuming Nash

⁹For example, with $\mathcal{M} = 75\%$, an episode occurs where downward nominal wage rigidity violates private efficiency once every 11,133 months (928 years). In performing these calculations we consider the fraction of laid-off workers over the entire episode of layoffs, which may last more than one month. Calculating the fraction of laid-off workers over a single period makes endogenous separations even more consistent with private efficiency.

¹⁰Since a secure job is more valuable to a worker than a short-term job $W_t^i \geq W_t^{i,new}$, the participation of short-term workers guarantees the participation of all workers, $V_t^i \equiv W_t^i - U_t^i \geq 0$.

bargaining prevails today and in all subsequent periods. Appendix B.2 shows that this wage satisfies

$$w_{Nash,t}^i = A_t^i \left(\gamma F'(N_{Nash,t}^i) + (1 - \gamma)z \right) + \beta E_t \left((1 - \delta)(1 - \gamma) f_{Nash,t+1}^i V_{Nash,t+1}^{i,new} \right), \quad (18)$$

where $\gamma \in [0, 1]$ parameterizes the bargaining power of workers.

The DMP model with wages determined by Nash-bargaining—equation (18)—fails to generate significant fluctuations in unemployment for plausible calibrations of the volatility of labor demand shocks (Shimer, 2005). Symmetric real wage rigidity allows the DMP model to generate a volatile unemployment rate—i.e. avoid the Shimer puzzle (Hall, 2005). Following Shimer (2010), the symmetric real wage rigidity rule that we consider is a weighted average of the past real wage and the present flexible-wage target given by equation (18):

$$\log(w_t^i) = \rho \log(e^g w_{t-1}^i) + (1 - \rho) \log(w_{Nash,t}^i), \quad (19)$$

where ρ is a weight between 0 and 1, and g is the rate of trend growth in productivity.¹¹ We assume that the wage w_t^i , while it adjusts slowly to shocks, still adjusts instantly to trend growth in productivity e^g . We do so to guarantee that the real wage is equal to the flexible-wage target in a non-stochastic steady-state with constant growth.

We show in section 4, that symmetric real wage rigidity does not generate the plucking property we have documented. To remedy this failure of the basic model, we consider a wage rule that features downward nominal wage rigidity rather than symmetric real wage rigidity. We assume that the nominal wage is set to the flexible wage—given by equation (18)—except if this requires the nominal wage to fall: $W_t^i = \max\{P_t w_{Nash,t}^i, W_{t-1}^i\}$. Expressed in terms of real wages, and denoting the inflation rate by $\Pi_t = P_t/P_{t-1}$, the wage-setting equation becomes

$$w_t^i = \max \left\{ w_{Nash,t}^i, \frac{w_{t-1}^i}{\Pi_t} \right\}. \quad (20)$$

With this wage rule, inflation relaxes the constraint on downward *real* wage adjustments: it greases the wheels of the labor market.

¹¹Since we assume a symmetric process for the logarithm of A_t^i , we take the average to be geometric—arithmetic for the logarithm of wages—in order not to introduce an ad hoc source of asymmetry in the model.

3.7 Equilibrium

To close the model, we assume that the goods market clears. This implies that production is equal to households' demand for consumption, plus firms' demand for hiring services:

$$\int_0^1 Y_t^i di = C_t + \int_0^1 cA_t^i \frac{H_t^i}{q(\theta_t^i)} I_{H_t^i \geq 0} di. \quad (21)$$

Under risk-neutrality, this market-clearing condition determines consumption residually given other equilibrium outcomes. We can, therefore, abstract from it (as well as from output) in defining the equilibrium.

An equilibrium is given by a processes for employment N_t^i , labor market tightness θ_t^i , and the real wage w_t^i for each sector $i \in [0, 1]$ such that in all sectors i , firm i is on its labor demand schedule—equation (9), (10), or (11) depending on the situation—the job-finding rate satisfies equation (15), and wages are set according to the wage setting rule for that model (equation (19) or equation (20)).

An equilibrium is conditional on exogenous processes for productivity A_t^i , initial conditions for employment M_0^i , and a monetary policy. We specify monetary policy as directly setting a path for the inflation rate Π_t , which we take to be constant at some target value $\bar{\Pi}$. We assume that sectoral productivity $\log(A_t^i)$ in sector i is the sum of a time trend g , an aggregate component $\log(A_t)$, and an idiosyncratic component $\log(Z_t^i)$:

$$\log(A_t^i) = g \times t + \log(A_t) + \log(Z_t^i). \quad (22)$$

We assume the idiosyncratic component follows an AR(1) in growth rates

$$\Delta \log(Z_t^i) = \rho_{\Delta z} \Delta \log(Z_{t-1}^i) + \varepsilon_t^{\Delta z, i} \quad (23)$$

with Gaussian innovations: $\varepsilon_t^{\Delta z, i} \sim \mathcal{N}(0, \sigma_\varepsilon^{\Delta z})$. We assume the aggregate component follows either an AR(1) or an AR(2) in levels

$$\log(A_t) = (I - \rho_1^a L)^{-1} \varepsilon_t^a, \quad (24)$$

$$\log(A_t) = (I - \rho_1^a L)^{-1} (I - \rho_2^a L)^{-1} \varepsilon_t^a, \quad (25)$$

where again the innovations are Gaussian: $\varepsilon_t^a \sim \mathcal{N}(0, \sigma_\varepsilon^a)$. We furthermore assume that all these processes are independent: $\log(A_t) \perp \log(Z_t^i), \log(Z_t^i) \perp \log(Z_t^j)$ for $i \neq j$.

3.8 Worker Flow Accounting

From economy-wide worker flows, we can define the economy-wide rate of job destruction s_t , the economy-wide rate of inflow from employment into unemployment \bar{s}_t , and the economy-wide rate of outflow from unemployment to employment f_t (equal to the economy-wide job-finding rate). In sector i at time t , there are $N_{0,t}^i = \min(N_t^i, M_t^i)$ workers who work in the same job as at time $t - 1$. The number of workers who find a new job at the beginning of time t in sector i is therefore $N_t^i - N_{0,t}^i$. In the overall economy it is $N_t - N_{0,t}$, where the aggregates N_t and $N_{0,t}$ are defined as sums across sectors, $N_t = \int_i N_t^i di$ and $N_{0,t} = \int_i N_{0,t}^i di$. The economy-wide job-finding rate is therefore

$$f_t = \frac{N_t - N_{0,t}}{1 - N_{0,t}}. \quad (26)$$

The number of workers who separate from their jobs between $t - 1$ and t in sector i is $N_{t-1}^i - N_{0,t}^i$. In the total economy it is $N_{t-1} - N_{0,t}$. The economy-wide job destruction rate is therefore:

$$s_t = \frac{N_{t-1} - N_{0,t}}{N_{t-1}}. \quad (27)$$

Our assumption that a worker who separates from his job at the end of period $t - 1$ has a chance f_t of finding a new job at the beginning of period t implies that the job-destruction rate s_t is not equal to the rate of inflow from employment into unemployment \bar{s}_t , or employment exit rate. Among the workers who separates from their jobs at the end of $t - 1$, the fraction f_t that starts a new job at the beginning of t does not transition from employment to unemployment but from job to job. Only the fraction $1 - f_t$ transitions to unemployment. The economy-wide rate of inflow from employment into unemployment is therefore:

$$\bar{s}_t = (1 - f_t)s_t. \quad (28)$$

From combining equations (26), (27), and (28), the law of motion of economy-wide unemployment $U_t \equiv 1 - N_t$ is:

$$U_t = (1 - f_t)U_{t-1} + \bar{s}_t(1 - U_{t-1}),. \quad (29)$$

3.9 Calibration

Table 3 provides a summary of our calibration. We present results for four versions of our model. First, we consider a version of the model with symmetric real wage rigidity, jobs that are secure

from the start ($\delta = d$), no sectoral shocks, and an AR(1) process for productivity. We refer to this version of the model as the “simple” model. In contrast, our “full model” (last column of Table 3) has downward nominal wage rigidity, jobs that are initially insecure ($\delta < d$), multiple sectors, and an AR(2) process for productivity. We also consider the simple model with downward nominal rate rigidity (second column of Table 3) and our full model with symmetric real wage rigidity (third column of Table 3).

We calibrate the model to a monthly frequency. We set the discount factor β to correspond to an annual interest rate of 4%. We assume a constant-elasticity production function $F(N) = N^\alpha$ and set $\alpha = 2/3$. We assume a Cobb-Douglas matching function $q(\theta) = \mu\theta^{-\eta}$ and set the elasticity of the matching function to $\eta = 0.5$, in the middle of the range reported in [Petrongolo and Pissarides \(2001\)](#)’s survey.

The parameters μ and c jointly determine hiring costs. One of the two is redundant as only the composite parameter $c\mu^{\frac{-1}{1-\eta}}$ is relevant for the equilibrium. (See [Appendix B.3](#) for further discussion of this point.) We normalize μ to 1. In the simple model, we set c so that the cost of hiring a worker $\frac{c}{q} = J$ is 10% of the monthly steady-state wage \bar{w} in a steady state with $u = 5.7\%$, in line with what [Silva and Toledo \(2009\)](#) report based on the Employer Opportunity Pilot Project survey in the US. This yields $c = 0.15$. In the full model, we calibrate c so that the total steady-state cost of hiring a worker in a long-term job J is still 10% of the monthly steady-state wage. This is the case for $c = 6.2 \times 10^{-3}$.

We calibrate the monthly separation rates δ and d so that, based on steady-state relationships, the average separation rate s is consistent with the average of the unemployment rate (5.7%) and the average of the job-finding rate f as estimated by [Shimer \(2012\)](#) based on CPS data (45%). Since in steady-state $1/s = (1/u - 1)(1/f - 1)$, this implies we target $s = 5.0\%$. In our simple model, this means we set $\delta = d = 5.0\%$. In our full model, we set $\delta = 0.2\%$, which is the order of magnitude needed to generate a slow rebuilding of long-term firm-worker relationships during expansions. In steady-state the average separation rate satisfies $s = \frac{\delta}{\delta + 1 - d}$, so that matching $s = 5.0\%$ sets $d = 96.2\%$.

We calibrate the flow value of unemployment z to 70% of the wage (specifically, 70% of the labor share α) following [Hall and Milgrom \(2008\)](#): 25% through unemployment benefits b and 45% through less foregone leisure ζ . We calibrate the bargaining power of workers γ to get a steady-state wage \bar{w} consistent with a steady-state level of unemployment equal to the average level of unemployment in the data (5.7%). The resulting values for γ for the four versions of

Table 3: Calibration

	Single type of job, AR(1) No Sectoral Shocks		Two types of jobs, AR(2) Sectoral Shocks	
	SRWR	DNWR	SRWR	DNWR
β				0.96 ^{1/12}
α				2/3
η				0.5
z				0.47
g	0.023/12	0	0.023/12	0.023/12
$\bar{\Pi}$	–	1.02 ^{1/12}	–	1.02 ^{1/12}
ρ	0.9	–	0.9	–
γ	0.87	0.80	0.99	0.97
δ	5.0%		0.2%	
d	5.0%		96.2%	
c	0.15		6.2×10^{-3}	
$\rho_{\Delta z}$	–		0.98	
$\sigma_{\Delta \varepsilon}^z$	–		7×10^{-4}	
ρ_a^1	0.98		0.96	
ρ_a^2	–		0.84	
σ_ε^a	st. $\sigma^a = 2.0\%$	st. $\sigma^a = 3.1\%$	st. $\sigma^a = 2.6\%$	st. $\sigma^a = 3.9\%$

Note: The abbreviation “st.” stands for “such that.”

the model we present results for are listed in Table 3. For the cases we consider with symmetric real wage rigidity, we calibrate the persistence parameter in the wage process ρ to 0.9, following Shimer (2010). For the cases we consider with downward nominal wage rigidity, we set inflation to 2% per year. (Inflation is immaterial in the symmetric real wage rigidity case.)

We set the growth rate of productivity g to 2.3% annually, the average growth of US labor productivity from 1948 to 2018.¹² We present results for two different processes for aggregate productivity: AR(1) and AR(2). In the AR(1) case, we set the auto-regressive root of the aggregate productivity process ρ_1^a to 0.98 following Shimer (2010). In the AR(2) case, we calibrate the two roots of the aggregate productivity process to match the frequency of unemployment cycles in the data. Because there is little internal propagation in our model beyond the forces that slow down employment recoveries (but not employment contractions), we calibrate the roots to the ones obtained from estimating an AR(2) directly on the US unemployment rate series. We first apply a three-period moving-average filter to the level of the unemployment rate to smooth out

¹²We deviate from this in one case. When we add downward nominal wage rigidity to the simple model, we set $g = 0$. We discuss this further below.

high-frequency variations. This yields $\rho_1^a = 0.96$ and $\rho_2^a = 0.84$.¹³ We get the same results when estimating the roots on the quarterly unemployment series without any smoothing and converting them to a monthly frequency.

We calibrate the standard deviation of innovations to the aggregate productivity process σ_ε^a to get the standard deviation of unemployment equal to its value in the data (1.6 percentage points). The resulting values for the standard deviation of the productivity process σ^a for the four versions of the model we present results for range from 2.0% to 3.9% (see Table 3). This calibration implies that the wage rules that we use go most of the way towards resolving the Shimer puzzle: Our calibration for the standard deviation of productivity is quite similar to the 2% value reported by Shimer (2005) for the cyclical component of productivity in the data. We choose to match the standard deviation of unemployment exactly (as opposed to calibrating to the standard deviation of productivity in the data) so that we can apply our definition of expansions and contractions to our simulated samples in the same way as we do to the real world data.

Our full model also features sectoral productivity shocks. We calibrate the persistence of the idiosyncratic productivity process based on KLEMS annual data on US sectoral productivity from 1947 to 2010 (Jorgenson, Ho, and Samuels, 2012). The KLEMS dataset provides labor productivity series (value added per hour) for 31 sectors. We take $\log(Z_t^i)$ to be the log difference between the sectoral labor productivity series and the BLS series for aggregate labor productivity. Here again, we first apply a three-period moving-average filter to the level of these series to smooth out high-frequency variations in $\log(Z_t^i)$. We then first-difference the resulting series and estimate AR(1) models for $\Delta \log(Z_t^i)$ in each sector. The average estimated autoregressive root across sectors is $\rho_{\Delta z} = 0.62$ at an annual frequency. We therefore calibrate $\rho_{\Delta z} = 0.62^{\frac{1}{12}} = 0.96$ in our monthly calibration. We calibrate the volatility of idiosyncratic productivity growth $\sigma_\varepsilon^{\Delta z}$ to roughly match the average of the fraction of constrained firms in the data as measured by the San Francisco Fed's Wage Rigidity Meter: 13% (see Appendix D). The value we use is $\sigma_\varepsilon^{\Delta z} = 7 \times 10^{-4}$.

Given the asymmetries and non-linearities our model is intended to capture, we rely on global methods to numerically solve for the equilibrium. Appendix B.5 discusses the algorithm we use in detail.

¹³The autoregressive coefficients are $\phi_1^a = 1.801$ and $\phi_2^a = -0.807$. These are related to the roots ρ_1^a and ρ_2^a through the equation $I - \phi_1^a L - \phi_2^a L^2 = (I - \rho_1^a L)(I - \rho_2^a L)$.

4 Does a Simple DMP Search Model Generate Plucking?

The first theoretical question we seek to assess is whether a simple DMP search model can generate the plucking property as well as the other business cycle asymmetries that we document in section 2. The second column of Table 4 presents results on the business cycle statistics that we document in section 2 for our simple model (symmetric real wage rigidity, jobs that are secure from the start ($\delta = d$), no sectoral shocks, and an AR(1) process for productivity). For each model we present results for in Table 4, we simulate 5000 samples of 866 periods (equal to the length of our sample of real-world data). We then calculate the business cycle statistics that we document in section 2 in each of these simulated samples and report the median estimate across samples for each statistic as a point estimate and the standard deviation of the estimates across samples in parentheses below each point estimate. The first column of Table 4 reports the estimated value of these business cycles statistics in the U.S. data for comparison.

The simple model generates a positive but very modest amount of the plucking property. The regression coefficient for the size of subsequent expansions on the size of contractions is only slightly larger (0.37) than the regression coefficient for the size of subsequent contractions on the size of expansions (0.18). This difference is far smaller than in the data (1.12 vs. -0.38). Similarly, the former regression has slightly more explanatory power than the latter regression: the R^2 is 0.13 versus 0.04. Again, this difference is much smaller than in the data (0.59 vs. 0.22). Figure E.1 provides a visual illustration of these results that is analogous to Figure 2. The simple model also generates little skewness, virtually no speed asymmetry, and durations of expansions and contractions that are both much too short and symmetric. In short, the simple model fails to match any of the features of unemployment rate dynamics we emphasize in section 2.

The fact that the simple model is not able to match the plucking property and the other asymmetries in the data indicates that this model does not feature significant non-linearity. In Appendix F, we show that the primary source of non-linearity in the simple model is the worker-flow relationship (28)-(29). Figure 3 plots the steady-state relationship between the job finding rate and unemployment implied by the worker-flow relationship (28)-(29).¹⁴ When the job finding rate is low—and therefore unemployment high—an increase in the job-finding rate applies to many job-

¹⁴The analytical form of this relationship is $u = s/(s + (f/(1 - f)))$. Consistent with the timing of our model, it takes s as constant and therefore makes $\bar{s} = (1 - f)s$ vary with f . Under an alternative assumption on the timing of the labor market in which workers separate from their jobs at a rate \bar{s} and must then necessarily spend one period unemployed, the relationship would be $u = \bar{s}/(\bar{s} + f)$. In this case too, the steady-state worker-flow relationship is convex but the non-linearities it induces are quantitatively small.

Table 4: Plucking Property, Skewness, Speed, and Duration

	Data	Single, AR(1)		Dual, AR(2)	
		SRWR	DNWR	SRWR	DNWR
Subsequent expansion on contraction, β	1.12	0.37 (0.22)	0.84 (0.22)	0.43 (0.25)	0.76 (0.22)
Subsequent contraction on expansion, β	-0.38	0.18 (0.22)	-0.06 (0.17)	0.24 (0.29)	-0.08 (0.32)
Subsequent expansion on contraction, R^2	0.59	0.13 (0.14)	0.68 (0.27)	0.21 (0.19)	0.74 (0.25)
Subsequent contraction on expansion, R^2	0.22	0.04 (0.09)	0.02 (0.04)	0.07 (0.13)	0.04 (0.10)
Skewness	0.64	0.26 (0.22)	2.13 (0.49)	0.28 (0.26)	1.55 (0.49)
Speed of expansions (pp/year)	0.87	3.59 (0.36)	3.92 (0.52)	2.22 (0.26)	1.46 (0.27)
Speed of contractions (pp/year)	1.89	3.41 (0.33)	4.03 (0.55)	2.22 (0.25)	2.74 (0.49)
Duration of expansions (months)	59.1	15.0 (1.8)	13.4 (1.7)	25.0 (3.8)	41.5 (7.4)
Duration of contractions (months)	26.9	15.6 (1.9)	16.9 (2.6)	24.5 (3.6)	21.5 (4.3)

Note: The table compares real world data with data from four versions of our model. The first column—labeled “Data”—reports empirical results based on data from the U.S. economy from section 2. The second column—labeled “Single, AR(1)” and “SRWR”—reports results for our simple model. “Single, AR(1)” means a single type of job ($\delta = d$), no sectoral heterogeneity, and aggregate productivity following an AR(1) process. “SRWR” means symmetric real wage rigidity. The last column—labeled “Dual, AR(2)” and “DNWR”—reports results for our full model. “Dual, AR(2)” means two types of jobs ($\delta < d$), sectoral heterogeneity, and aggregate productivity following an AR(2) process. “DNWR” means downward nominal wage rigidity model. The third column reports results for the simple model but with DNWR and the fourth column reports results for the full model but with SRWR. The first (third) row reports the coefficient (R^2) in an OLS regression of the size of an expansion (percentage point fall in unemployment rate) on the size of the previous contraction (percentage point increase in unemployment rate). The second (fourth) row report the coefficient (R^2) in an analogous regression of the size of a contraction on the size of the previous expansion. The fifth row reports the skewness of the distribution of the unemployment rate. The next two rows report the spell-weighted average speed of expansions and contractions, measured in percentage points of unemployment per year. The final two rows report the average duration of expansions and contractions, measured in months. For the models, the reported point estimate is the median value of the statistic over 5000 samples of 866 periods each (the length of our sample of real-world data). The standard error reported in parentheses is the standard deviation of the estimates across the 5000 samples. The Dual AR(2) models are simulated with 1000 sectors.

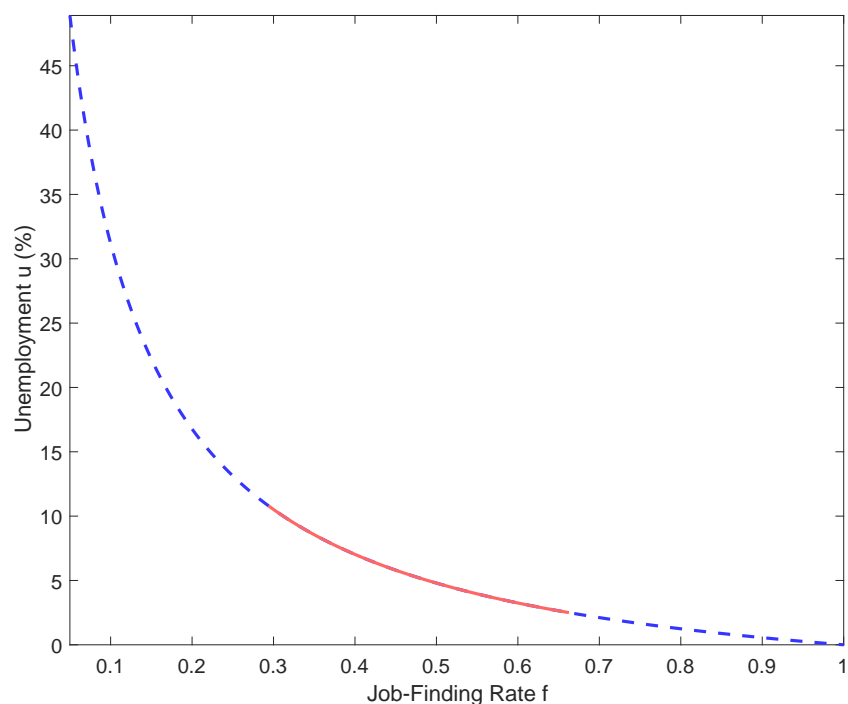


Figure 3: Steady-State Worker-Flow Relationship

Note: The figure plots the steady-state relationship between the job-finding rate and the unemployment rate implied by the worker-flow equation. In plain line the relationship is plotted over the range of unemployment rates observed in the US between 1948 and 2019: from 2.5% to 10.8%.

seekers and therefore decreases unemployment substantially. When the job-finding rate is high—and therefore unemployment low—an increase in the job-finding rate applies to fewer job-seekers and therefore does not decrease unemployment as much. Intuitively, it gets harder and harder to lower the unemployment rate the lower it gets. [Hairault, Langot, and Osotimehin \(2010\)](#), [Jung and Kuester \(2011\)](#), and [Lepetit \(2018\)](#) emphasize this source of non-linearity.

The substantial degree of convexity of the worker-flow relationship (28)-(29) plotted in Figure 3 suggests that the simple model might be able to generate the plucking property and other business cycle asymmetries. The reason this is not the case is that fluctuations in unemployment in the US over our sample period have not been large enough. Over the sample period 1948-2019, unemployment fluctuated between 2.5% and 10.8%. This is the red solid segment in Figure 3. Over this relatively small segment, the worker-flow relationship is close to linear. For realistic fluctuations in the unemployment rate, the simple model therefore generates a minimal degree of business cycle asymmetry. Only for counter-factually large fluctuations in the job-finding rate and

unemployment does the convexity start to become quantitatively important.¹⁵

5 Matching Unemployment Asymmetries

The right-most column in Table 4 presents results on the business cycle statistics we document in section 2 for our full model with downward nominal wage rigidity, jobs that are initially insecure ($\delta < d$), multiple sectors, and an AR(2) process for productivity. For all the business cycle asymmetries we document, the full model generates vastly more asymmetry than the simple model and comparable asymmetry to the data. Consider first the degree of plucking. The regression coefficient for the size of subsequent expansions on the size of contractions is 0.76, while the regression coefficient for the size of subsequent contractions on the size of expansions is -0.08. This large difference is statistically insignificantly different from its data counterpart. The former regression yields a large R^2 (0.74), while the R^2 of the latter regression is very small (0.04). Again, this difference is statistically insignificantly different from its data counterpart.¹⁶ Each of the four statistics that comprise these differences is also statistically insignificantly different from its data counterpart. The full model yields a substantial amount of skewness of unemployment, somewhat more than in the data.

The full model also does a good job of matching the speed and duration of expansions and contractions in sharp contrast to the simple model. In the full model the unemployment rate falls by 1.46 percentage points per year and rises by 2.74 percentage points per year. The speed difference between expansions and contractions is 1.28 percentage points per year, which is close to the speed difference in the data (1.02). The average duration of expansions is 41.5 months in the full model, while the average duration of contractions is 21.5 months. The full model therefore roughly matches the fact that expansions are twice as long as contractions.

¹⁵The limited degree of non-linearity that arises from the worker-flow relationship (28)-(29) implies that the simple model's aggregate labor demand schedule is close to log-linear. We analyze this in more detail in Appendix F.1. It is common in the DMP literature to abstract from mechanisms other than search costs that make the labor demand schedule downward sloping, such as decreasing returns to labor. This tends to magnify the degree of non-linearity in these models. For example, Petrosky-Nadeau, Zhang, and Kuehn (2018) show how a DMP model with constant returns to labor features asymmetries that can generate business-cycles disasters—large drops in production—despite symmetric shocks. We show in Appendix F.1 that a version of our simple model with constant returns to scale can generate more plucking, but that this is driven by recessions more extreme than we have seen in our post-WWII sample.

¹⁶Figure E.1 provides a visual illustration of these results that is analogous to Figure 2.

5.1 Downward Nominal Wage Rigidity and the Plucking Property

For matching the plucking property, the crucial feature of our full model is downward nominal wage rigidity (DNWR). The simplest way to see this is to compare the second and third columns of Table 4. The third column in Table 4 adds DNWR to the simple model without adding any of the other features in the full model. We see that this model generates a comparable amount of the plucking property as the full model.¹⁷

Figure 4 illustrates the role of DNWR by plotting a simulated path for unemployment in our full model (top panel) and comparing this to a simulated path for unemployment (with the same shocks) from a version of our model that differs from the full model only in that it has symmetric real wage rigidity (bottom panel). The difference is striking. While the model with symmetric real wage rigidity (SRWR) yields roughly symmetric cycles, the full model with DNWR yields highly asymmetric cycles.

The reason for this is intuitive. With DNWR, a string of positive shocks eventually yield a sharp increase in wages, while negative shocks result in higher unemployment since wages are rigid downward. As a result, unemployment sometimes rises far above its steady state, but rarely falls below its steady state.¹⁸ With SRWR, in contrast, wage variation is dampened symmetrically relative to changes in productivity. Variation in unemployment is therefore symmetric about that model's steady state.

We parameterized both models in Figure 4 such that the average unemployment rate is equal to its data counterpart (5.7%). This requires a much lower steady-state level of unemployment in the full model (3.8%) than in the model with symmetric real wage rigidity (5.7%). As we discuss in greater detail in section 6, this implies that reducing aggregate fluctuations has vastly larger welfare consequences in our full model than in models that don't capture the plucking property.

Inflation and trend growth in productivity are two secular forces that naturally "grease the

¹⁷A caveat to this is that our calibration of the simple model with DNWR assumes 0% average aggregate productivity growth. Without sectoral shocks, the downward nominal wage rigidity constraint does not bind a large enough fraction of the time in our model to generate the plucking property we document in the data when average productivity growth is 2.3%. A second unappealing feature of the model with downward nominal wage rigidity but without sectoral shocks is that the downward nominal wage rigidity constraint binds for either all or no firm at any given point in time. Figure D.1 shows that in the US a fraction of workers face wage freezes at any given point in time and this fraction is highly counter-cyclical. Figure D.2 shows that our full model can qualitatively match this feature of the data. The share of workers facing wage freezes in our model is highly correlated with the unemployment rate as in the data.

¹⁸In the full model, unemployment can fall slightly below its steady state level. The reason for this is that the sectoral shocks in this model imply that in steady state (i.e., absent aggregate shocks) there is a fraction of firms with wages above the flexible wage. These firms received positive sectoral shocks at an earlier date that raised wages but were then reversed. A sequence of positive aggregate shocks can therefore lower unemployment below the steady state in this model by reducing the fraction of firms that have wages above the flexible wage to a level below its steady state level.

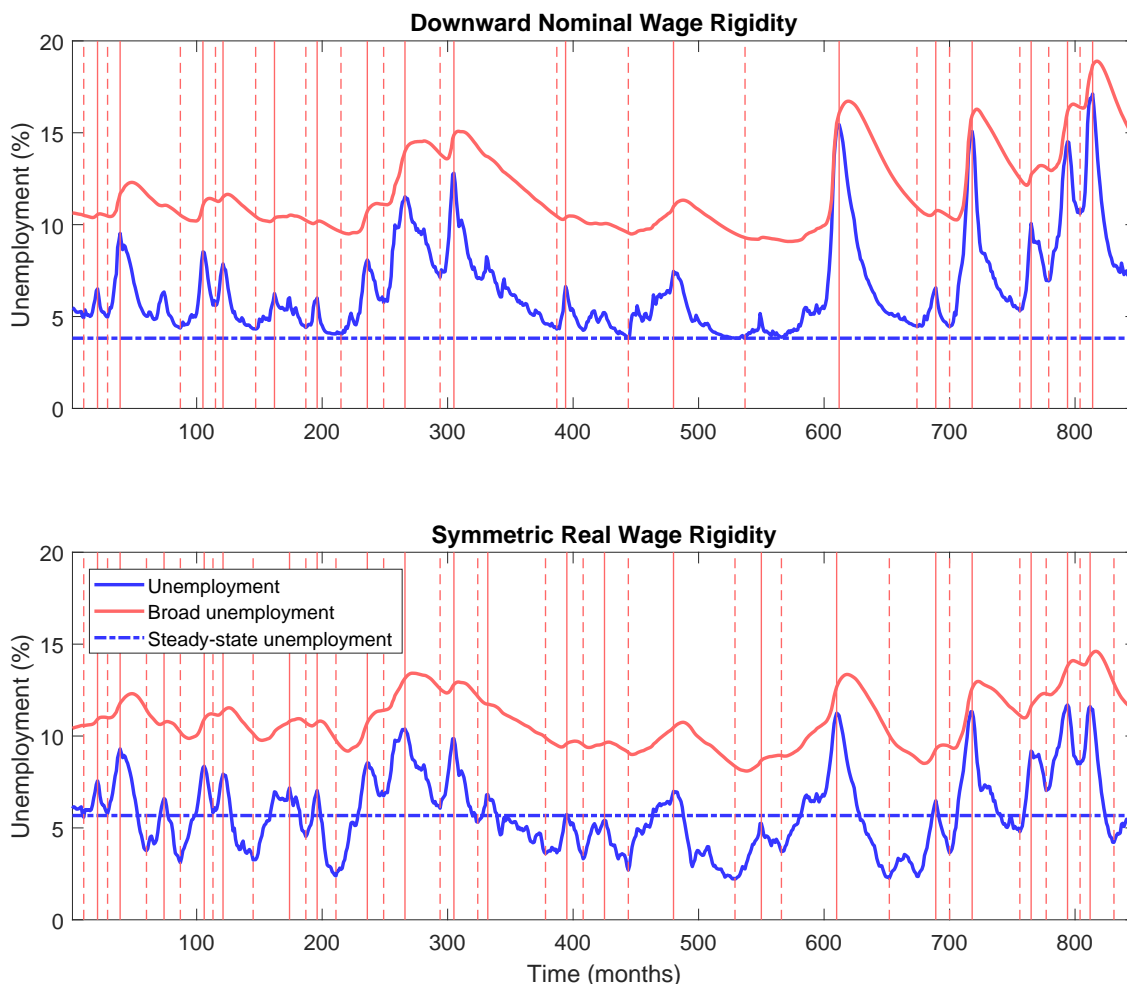


Figure 4: Simulated Paths for the Unemployment Rate

Note: The figure plots sample paths of 72 years (the same length as our empirical sample) for the unemployment rate and broad unemployment rate for two versions of our model with secure and insecure jobs, multiple sectors and an AR(2) aggregate productivity process. The upper panel plots unemployment in the model with downward nominal wage rigidity. The lower panel plots unemployment in this model with symmetric real wage rigidity. The sequence of shocks are the same in both cases. The solid vertical lines identify business-cycle troughs, while the dashed vertical lines identify peaks.

wheels of the labor market” by eroding gaps between wages and the flexible wage when wages are too high. We calibrate our full model to have 2% annual inflation and 2.3% annual trend aggregate productivity growth. These forces impart a substantial positive trend to the flexible nominal wage. Negative shocks to labor demand must be big enough to more than offset this trend for downward nominal wage rigidity to bind. As [Benigno and Ricci \(2011\)](#) emphasize, sectoral shocks are an important source of volatility in labor demand. In our full model, sectoral shocks are an important reason why the downward nominal wage rigidity constraint binds despite the

substantial positive trend of the flexible nominal wage.

As we note above, our model yields more skewness of the unemployment rate than in the data. Unemployment falls somewhat too quickly at the beginning of an expansion in our model, slowing down its decrease only later in the expansion (see figure 4). In contrast, the falls in unemployment in the data are more linear, as emphasized by [Hall and Kudlyak \(2021\)](#). As a result, unemployment in our model spends little time at very high values, which results in a distribution with a larger skewness than the data.

5.2 The Speed and Duration of Expansions and Contractions

An appealing aspect of search models—at least in principle—is that they provide an intuitive mechanism to match the asymmetry in the speed of expansions and contractions: firms can shed workers rapidly, but it takes time due to search and matching frictions to expand employment. In practice, however, the standard DMP model does not deliver this result. We can see this by considering the implications of our simple model for the speed and duration of business cycles in Table 4. The simple model yields expansions and contractions that are much too rapid and symmetrically so. The reason for the lack of asymmetry is that job flows are very large implying that the model does not have appreciable internal propagation (see, e.g., [Cole and Rogerson, 1999](#)). The unemployment rate, therefore, inherits the dynamics of the wage-to-productivity ratio w_t/A_t .

To deliver on the promise of search models regarding the asymmetry in the speed of expansions and contractions, we must do two things. First, we must slow down unemployment dynamics in the model; and second, we must create bursts of separations that speed up contractions relative to expansions.

Our full model has two features that slow down unemployment dynamics. First, we assume that labor demand shocks follow an AR(2) process. An AR(2) process can generate very high persistence at business cycle frequencies without extreme levels of persistence at very low frequencies.¹⁹ The second feature of our full model that slows down unemployment dynamics is

¹⁹A robust lesson from the recent empirical literature in macroeconomics is that the dynamic responses of economic activity to many shocks is hump-shaped (e.g. [Romer and Romer, 2004](#); [Christiano, Eichenbaum, and Evans, 2005](#)). Medium scale DGSE models fit the data better when they are modified to include equations which yield AR(2) dynamics—e.g., investment adjustment costs, habits in consumption, and lagged terms in the price and wage Phillips curves combined with AR(1) shocks ([Smets and Wouters, 2007](#)). In contrast, an AR(1) model would require an extremely high persistence parameter to fit the persistence facts we emphasize. Under such a calibration, the model would generate too much variation in unemployment at very low frequencies.

short-term insecure jobs. The standard DMP model abstracts from the notion that the failure rate of new matches in the labor market is higher than the failure rate of mature matches, i.e., unemployed workers typically cycle through several jobs before they find a secure, long-term position. Our full model incorporates this feature of reality. It slows down the process of workers returning to (stable) employment after recessions.²⁰

While these two features impart added persistence, they do not generate asymmetry on their own. We can see this from the fact that the full model modified to have SRWR does not generate speed and duration asymmetry (fourth column of Table 4). To get asymmetry, we need bursts of separations during downturns. The features of our model that generate such bursts are DNWR and endogenous separations. With SRWR, endogenous separation of securely attached workers can happen but is very limited.²¹ With DNWR, a string of negative labor demand shocks eventually results in significant endogenous separations of securely attached workers—i.e., a burst of separations. These bursts of separations are short-lived and modest in size. But they contribute importantly to speeding up unemployment contractions.²²

Taken together these features yield a substantial degree of asymmetry in the speed of contractions and expansions (as well as the duration of expansions and contractions). This speed asymmetry occurs for intuitive reasons. During a recession, DNWR results in a burst of separations of workers in otherwise secure jobs. These workers then cycle through several short-term jobs before they eventually find new stable employment. This process of cycling through short-term jobs contributes to preventing the unemployment rate from quickly returning to its steady state level. The slow speed of recoveries is also due to slow underlying dynamics in labor demand and gradual dissipation of constraints on wages.

The speed asymmetry is even more pronounced for “broad” unemployment (i.e., all workers not in secure jobs). Figure 4 illustrates this by plotting the dynamics of “broad” unemployment. Clearly, broad unemployment falls much more gradually than it rises. It also falls more linearly than the standard unemployment rate in our model, and as a result has a lower skewness (1.07).

²⁰Fujita and Ramey (2007) explore an alternative mechanism for increasing the propagation of shocks in the labor market. They assume that the cost of opening a vacancy is non-zero and increasing in the number of new vacancies opened. This makes vacancy creation sluggish.

²¹With SRWR, the monthly rate of separation from secure employment rises above 0.5% once every 4555 months (380 years). With DNWR, this rate rises above 0.5% once every 98 months (8 years).

²²Appendix G provides further details on the behavior of the separation margin in our model. While endogenous separations contribute significantly to the sharp rise in unemployment during recessions, fluctuations in unemployment are primarily determined by fluctuations in the job-finding rate, consistent with what Elsby, Michaels, and Solon (2009) and Shimer (2012) document on US data.

6 Costs of Business Cycles and Benefits of Stabilization Policy

We now turn to the normative implications of our model. We present two sets of results. First, smaller shocks to labor demand yield a lower average level for the unemployment rate. Second, raising the average inflation rate lowers the average level of the unemployment rate.

6.1 First-Order Effect of Stabilization Policy

In a thought-provoking exercise, Lucas (1987, 2003) calculated the welfare benefits of eliminating economic fluctuations. Assuming log-utility and trend-stationary fluctuations with Gaussian innovations, Lucas (2003) calculated that the representative agent would be willing to forgo no more than 0.05% of his or her consumption to be rid of fluctuations. Lucas concluded from this that the benefits of stabilization policy are trivial.

An important maintained hypothesis in Lucas' analysis is that eliminating fluctuations does not affect the average level of economic activity.²³ In our plucking model, however, fluctuations are drops below potential rather than cycles around a natural rate. Eliminating fluctuations therefore raises the average level of economic activity.

The left panel of figure 5 illustrates this by plotting the average level of the unemployment rate in our plucking model as a function of the volatility of aggregate shocks. Eliminating all fluctuations in our model reduces the average unemployment rate from 5.7% to 3.8%. Conversely, increasing the standard deviation of aggregate shocks by 50% (from 3.9% to 5.8%) increases the average unemployment rate to 8.4%. Since the efficient level of unemployment is far below the average level of unemployment for any level of aggregate shocks in our plucking model, stabilization policy that is able to reduce aggregate fluctuations yields appreciable welfare gains.²⁴

The consumption-equivalent welfare gains from eliminating aggregate fluctuations are, however, not equal to the average reduction in the unemployment rate. One factor making the welfare gains smaller is decreased leisure time associated with more work. Another factor is that more resources are consumed matching workers and firms when the labor market is tighter. In appendix H, we calculate the consumption equivalent welfare gain from eliminating aggregate fluctuations

²³Aiyar and Voigts (2019) show that common estimation methods of the output gap implicitly take this view by assuming a zero-mean output gap, and therefore systematically underestimate the amount of slack—and resulting benefits of stabilization policies—under the plucking view.

²⁴The efficient level of unemployment is somewhat sensitive to the exact timing assumptions in our model. But it is far below the average level of unemployment no matter which timing assumption we adopt. We discuss this in more detail in appendix H.

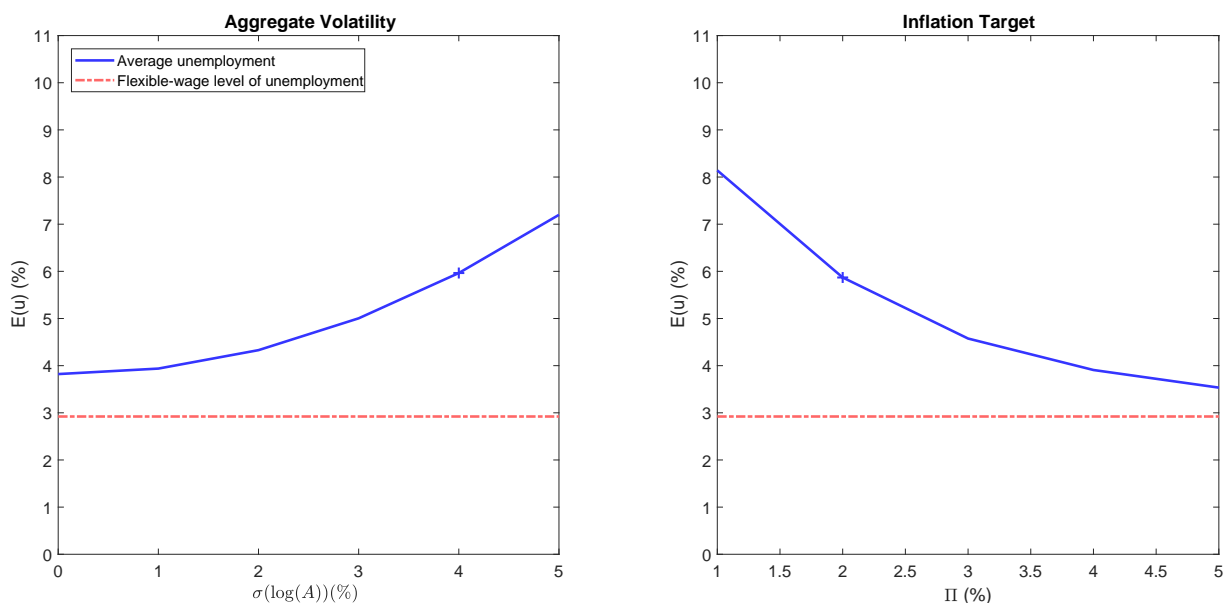


Figure 5: Average Unemployment, the Volatility of Aggregate Shocks and Inflation

Note: The figure gives the average rate of unemployment as a function on the standard deviation of aggregate shocks (left panel) and as a function of the inflation target (right panel) in our full model with downward nominal wage rigidity, two types of jobs, sectoral heterogeneity and an AR(2) process for productivity.

in our model. We do this for two different assumptions about the value of leisure time during unemployment. Assuming that all time while unemployed is leisure implies that the welfare gain of eliminating aggregate fluctuations in our model is 0.65%. In contrast, if we assume that unemployment does not yield additional leisure time—i.e., that people enjoy their time unemployed as much/little as their time while working—the welfare gain of eliminating aggregate fluctuations is 1.20% in our model. In both cases, therefore, our plucking model implies that the welfare gains from eliminating economic fluctuations are an order of magnitude larger than in standard macroeconomic models.

6.2 Greasing the Wheels of the Labor Market

Lucas’s thought experiment of eliminating all fluctuations is intended to give an upper-bound of the benefits of stabilization policies, abstracting from constraints that may exist on what outcomes policy can achieve. We next consider a particular policy: the choice of the inflation target. In our model with downward nominal wage rigidity, inflation greases the wheels of the labor market by easing the downward adjustment of real wages.

The right panel of figure 5 gives the average unemployment rate as a function of the inflation

target in our model. Increasing the inflation target from 2% to 4% decreases average unemployment from 5.7% to 3.9%, and to 3.5% for an inflation target of 5%. Strikingly, extra inflation can reduce unemployment more than the elimination of all aggregate fluctuations. The reason for this is that inflation eases not only adjustment to aggregate shocks but also adjustment to idiosyncratic shocks. As the inflation target increases further, marginal benefits fall. Average unemployment asymptotes to its values absent any frictions on wage adjustments and absent any (idiosyncratic and aggregate) shocks: 2.9%.²⁵

These estimates of the effect of the inflation target on unemployment rely on the assumption that wage-setting remains unchanged in the face of the new monetary policy. For a high enough inflation target, it is however likely that some workers would shift to thinking in real terms. Any reluctance to bear wage cuts would then also manifest itself through downward *real* wage rigidity, and inflation would no longer have an effect on unemployment. The middle part of our sample period includes the Great Inflation, where the US experienced inflation in the high single digits and low double digits. The fact that the US unemployment rate displayed the plucking property during this period of relatively high inflation suggests that barriers to downward adjustment are not fully eliminated by high inflation but rather that they also manifest themselves as downward real wage rigidity in these circumstances perhaps due to widespread implicit or explicit cost-of-living adjustments in wage setting.²⁶

Although the distinction between real and nominal downward rigidity is critical when assessing the benefits of a higher inflation target, this distinction is not crucial for the ability of our model to replicate the plucking property. Indeed, our wage-setting equation (20) reduces to the case of downward real wage rigidity for an inflation target of zero (up to a reinterpretation of the parameters). Nor do the implications of our model for the welfare costs of business cycles depend on whether the constraint on wage cuts bears on nominal or real wages. If *real* wages cannot fall as easily as they can rise, it is still the case that more volatile shocks increase average unemployment.

²⁵Our results contrast sharply with those of [Coibion, Gorodnichenko, and Wieland \(2012\)](#). They find that optimal inflation in a model with downward nominal wage rigidity is 0.3% and is smaller than in a model with only price rigidity. They derive this result in a model in which the difference between the marginal value of output and the marginal disutility of labor is small implying that the benefits of extra output are small. In our model, this difference is large and the value of extra output is therefore large. In their model the costs of inflation due to price dispersion rise rapidly when inflation increases. We do not model costs of inflation due to price dispersion. ([Nakamura et al. \(2018\)](#) fail to find evidence supporting rapidly increasing price dispersion when inflation rises.) We focus on the plucking property as the main source of the potential benefits to downward wage rigidity, whereas their model abstracts from several of the features we emphasize that generate the plucking property. On the other hand, we abstract from the zero lower bound on nominal interest rates, which is the main source of benefits of inflation in their model.

²⁶[Benigno, Ricci, and Surico \(2015\)](#) show that a model with real wage rigidity can match the dynamics of the unemployment rate reasonably well during the Great Inflation period.

The real effect of downward nominal wage rigidity can be attenuated if firms and workers preemptively moderate wage increases in booms, in order to reduce the probability of a painful adjustment during a downturn. Such forward-looking wage moderation is present in wage setting models (e.g., [Kim and Ruge-Murcia, 2009](#); [Elsby, 2009](#); [Benigno and Ricci, 2011](#)). However, [Benigno and Ricci \(2011\)](#) find large benefits of inflation in greasing the wheels of the labor market despite their model featuring such preemptive wage moderation. Our model does not feature preemptive wage moderation. Yet this does not mean firms in our model are myopic. They rationally maximize intertemporal profits. What they preemptively moderate in anticipation of a fall in productivity is hires, not wages. Either wages or hires can respond to concerns about the future. In our model it is hires that are moderated, because wages are not set by firms.

7 Conclusion

We build a plucking model of the business cycle that captures the asymmetry in the predictive power of contractions and expansions emphasized by Milton Friedman. In our model, the asymmetry arises from downward nominal wage rigidity. In contrast to earlier models of downward nominal wage rigidity, our model is consistent with optimizing behavior and therefore robust to the [Barro \(1977\)](#) critique. In contrast to standard DMP-style search and matching models, our model can also match the large asymmetry between the speed of contractions and the speed of expansions.

We show that in our model eliminating business cycles has large welfare benefits since it lowers the average unemployment rate. Our simulations imply that eliminating all aggregate fluctuations could lower the average unemployment rate by about 1.9 percentage points. Downward nominal wage rigidity provides one rationale for a positive inflation rate. Our results imply that moving from a 2% inflation target to a 4% inflation target would lower the average unemployment rate by 1.8 percentage points by easing the adjustment to both idiosyncratic and aggregate shocks. Lowering the inflation target to 1% would raise the average unemployment rate by 2.3 percentage points.

A Defining Expansions and Contractions

A.1 An Algorithm for Defining Expansions and Contractions

Let u_t denote the unemployment rate at time t . The algorithm begins by taking the first month of our sample as a candidate for a business cycle peak, cp . If, in all the following months until unemployment becomes X percentage points higher than u_{cp} , unemployment is higher than u_{cp} , we confirm that cp is a business cycle peak. If, instead, the unemployment rate falls below u_{cp} before it is confirmed as a peak, the month in which this happens becomes the new candidate peak. Once we have identified a peak, we switch to looking for a trough (in the analogous manner) and so on until we reach the end of the sample. Formally, starting with $t = 1$ the algorithm is:

1. Set $cp = t$ and set $t = t + 1$ (i.e., move to the next time period).
2. If $u_t < u_{cp}$ go back to step 1
3. If $u_{cp} \leq u_t < u_{cp} + X$ set $t = t + 1$ and go back to step 2
4. If $u_t \geq u_{cp} + X$ add cp to the set of peaks
5. Set $ct = t$ and set $t = t + 1$
6. If $u_t > u_{ct}$ go back to step 5
7. If $u_{ct} \geq u_t > u_{ct} - X$ set $t = t + 1$ and go back to step 6
8. If $u_t \leq u_{ct} - X$ add ct to the set of troughs, and go back to step 1

We set $X = 1.5$ percentage points. With this value, our algorithm generates the same set of expansions and contractions as the NBER Business Cycle Dating Committee with one exception: Our algorithm considers the 1979-1982 double-dip recession as a single contraction as opposed to two contractions interrupted by a brief and small expansion (unemployment decreased by 0.6 percentage points in 1980-1981). Values for X between 0.8 and 1.5 percentage points identify exactly the same cycles. Values of X larger than 1.5 drop the 1970-1973 expansion.

An advantage of our algorithm is that it does not impose a duration upon expansions and contractions but only a size X , in contrast to other algorithms based on turning points like the [Bry and Boschan \(1971\)](#) routine. Our algorithm can therefore also be used to define expansions and contractions in our model simulations, even for models that do not match the duration of expansions and contractions in the real-world data.

Table A.1: Business Cycle Peaks and Troughs

	Unemployment		NBER	
	Peak	Trough	Peak	Trough
1	[1/1948]	10/1949	11/1948	10/1949
2	5/1953	9/1954	7/1953	5/1954
3	3/1957	7/1958	8/1957	4/1958
4	2/1960	5/1961	4/1960	2/1961
5	9/1968	12/1970	12/1969	11/1970
6	10/1973	5/1975	11/1973	3/1975
7a	5/1979		1/1980	7/1980
7b		11/1982	7/1981	11/1982
8	3/1989	6/1992	7/1990	3/1991
9	4/2000	6/2003	3/2001	11/2001
10	10/2006	10/2009	12/2007	6/2009
11	9/2019		2/2020	

Note: Business cycle peaks and troughs defined solely based on the unemployment rate and, for comparison, business cycle peaks and troughs as defined by the Business Cycle Dating Committee of the National Bureau of Economic Research.

A.2 Peak and Trough Dates from 1948 to 2020

Table A.1 presents the peak and trough dates we identify. For comparison purposes, we also present the peak and trough dates identified by the NBER. We identify the same set of expansions and contractions as the NBER Business Cycle Dating Committee with one exception: we consider the 1979-1982 double-dip recession as a single contraction as opposed to two contractions interrupted by a brief and small expansion (unemployment decreased by 0.6 percentage points in 1980-1981). The exact timing of the NBER peaks and troughs do not line up exactly with ours for the reasons discussed above. However, in most cases, our dates are quite similar to theirs. The NBER peaks tend to lag our peaks by a few months and the NBER troughs tend to precede our troughs by a few months. This implies that our estimate of the average duration of contractions is about one year longer than what results from the NBER's dating procedure. We identify September 2019 as a peak as opposed to February 2020 because the unemployment rate first hit 3.5% in September 2019. When several months are tied for the lowest unemployment rate at the end of an expansion, our algorithm picks the first of these months as the peak (and similarly for troughs). Table A.2 lists the duration of all expansions and contractions over our sample period.

Table A.2: The Duration of Expansions and Contractions

	Dates		Length in Months	
	Peak	Trough	Expansion	Contraction
1	[1/1948]	10/1949		21
2	5/1953	9/1954	43	16
3	3/1957	7/1958	30	16
4	2/1960	5/1961	19	15
5	9/1968	12/1970	88	27
6	10/1973	5/1975	34	19
7	5/1979	11/1982	48	42
8	3/1989	6/1992	76	39
9	4/2000	6/2003	94	38
10	10/2006	10/2009	40	36
11	9/2019		119	
Mean			59.1	26.9

B Model Solution

B.1 Labor Demand

The employment decision of the firm is

$$\Pi_t^i(M_t^i) = \max_{H_t^i} \{A_t^i F(M_t^i + H_t^i) - w_t^i(M_t^i + H_t^i) - \frac{A_t^i c}{q_t^i} H_t^i 1_{H_t^i \geq 0} + \beta E_t(\Pi_{t+1}^i(M_{t+1}^i))\}, \quad (\text{B.1})$$

subject to the law-of-motion of its workforce $M_{t+1}^i(N_t^i)$ given by equations (3)-(4).

The envelope condition is

$$J_t^{i*}(M_t^i) = A_t^i F'(N_t^{i*}) - w_t^i + \beta(1 - \delta) E_t(J_{t+1}^{i*}(M_{t+1}^i(N_t^{i*}))). \quad (\text{B.2})$$

where we define $J_t^{i*}(M_t^i) = \frac{\partial \Pi_t^i}{\partial M_t^i}(M_t^i)$, the equilibrium marginal value to the firm of a securely attached worker inherited from the previous period.

The first-order conditions give:

- If $H_t^{i*} > 0$ (i.e., if the firm hires):

$$\frac{A_t^i c}{q_t^i} = A_t^i F'(N_t^{i*}) - w_t^i + \beta(1 - d) E_t(J_{t+1}^{i*}(M_{t+1}^i(N_t^{i*}))). \quad (\text{B.3})$$

This happens when $A_t^i F'(M_t^i) - w_t^i + \beta(1 - d) E_t(J_{t+1}^{i*}(M_{t+1}^i(M_t^i))) > \frac{A_t^i c}{q_t^i}$.

- If $H_t^{i*} < 0$ (i.e., if the firm fires):

$$0 = A_t^i F'(N_t^{i*}) - w_t^i + \beta(1 - \delta) (J_{t+1}^{i*}(M_{t+1}^i(N_t^{i*}))). \quad (\text{B.4})$$

This happens when $A_t^i F'(M_t^i) - w_t^i + \beta(1 - \delta)E_t (J_{t+1}^{i*}(M_{t+1}^i(M_t^i))) < 0$.

- When neither condition is satisfied, i.e. when

$$A_t^i F'(M_t^i) - w_t^i + \beta(1 - d)E_t (J_{t+1}^{i*}(M_{t+1}^i(M_t^i))) \leq \frac{A_t^i c}{q_t^i}, \quad (\text{B.5})$$

$$A_t^i F'(M_t^i) - w_t^i + \beta(1 - \delta)E_t (J_{t+1}^{i*}(M_{t+1}^i(M_t^i))) \geq 0, \quad (\text{B.6})$$

then the firm freezes hires $H_t^{i*} = 0$.

To simplify and interpret this labor demand schedule, equation (7) defines the marginal value to the firm of a worker already attached to the firm and equation (8) defines the marginal value to the firm of a newly-hired worker. The labor-demand schedule can then be rewritten as in equations (9)-(11).

In equilibrium, because all firms in sector i are identical, if firms are freezing employment then $A_t^i c/q_t^i = 0$. Therefore, the aggregate labor demand schedule in sector i is:

- If $J^{i,new}(M_t^i) > 0$, then firms hire up to:

$$J^{i,new}(N_t^{i*}) = \frac{A_t^i c}{q(N_t^{i*})}. \quad (\text{B.7})$$

- If $J^i(M_t^i) < 0$ then firms fire up to:

$$J^i(N_t^{i*}) = 0. \quad (\text{B.8})$$

- If $J^i(M_t^i) \geq 0 \geq J^{i,new}(M_t^i)$ firms freeze employment:

$$N_t^{i*} = M_t^i. \quad (\text{B.9})$$

B.2 The Nash-Bargained Wage

We solve for the Nash-bargaining equilibrium in order to obtain the Nash wage. We take the Nash wage to be the one that prevails when wages are Nash-bargained between the firm and securely attached workers, both today and in all subsequent periods. For workers, we get from equations (12)-(14) that the values of having a secure job relative to being unemployed $V_t^i \equiv W_t^i - U_t^i$ and the value of having a short-term job relative to being unemployed $V_t^{i,new} \equiv W_t^{i,new} - U_t^i$ solve:

$$V_t^{i,new} = w_t^i - A_t^i z + \beta E_t \left((1 - d_{t+1}^i) \left(V_{t+1}^i - f_{t+1}^i V_{t+1}^{i,new} \right) \right), \quad (\text{B.10})$$

$$V_t^i = w_t^i - A_t^i z + \beta E_t \left((1 - \delta_{t+1}^i) \left(V_{t+1}^i - f_{t+1}^i V_{t+1}^{i,new} \right) \right), \quad (\text{B.11})$$

where $z = b + \zeta$ is the flow value of unemployment in terms of both unemployment benefits and more leisure.

For firms, we get from equations (7)-(8) that:

$$J_t^{i,new} = A_t^i F'(N_t^i) - w_t^i + \beta E_t((1-d)J_{t+1}^i), \quad (\text{B.12})$$

$$J_t^i = A_t^i F'(N_t^i) - w_t^i + \beta E_t((1-\delta)J_{t+1}^i). \quad (\text{B.13})$$

Summing up to get the total surpluses of a secure job $S_t^i \equiv J_t^i + V_t^i$ and of a temporary job $S_t^{i,new} \equiv J_t^{i,new} + V_t^{i,new}$:²⁷

$$S_t^{i,new} = A_t^i (F'(N_t^i) - z) + \beta E_t((1-d)(S_{t+1}^i - f_{t+1}^i V_{t+1}^{i,new})), \quad (\text{B.14})$$

$$S_t^i = A_t^i (F'(N_t^i) - z) + \beta E_t((1-\delta)(S_{t+1}^i - f_{t+1}^i V_{t+1}^{i,new})). \quad (\text{B.15})$$

Nash-bargaining between the firm and long-term workers implies that $J_t^i = (1-\gamma)S_t^i$ and $J_{t+1}^i = (1-\gamma)S_{t+1}^i$. Combined with equation (B.15) this implies:

$$J_{Nash,t}^i = (1-\gamma)A_t^i (F'(N_{Nash,t}^i) - z) + \beta E_t \left((1-\delta)J_{Nash,t+1}^i - (1-\gamma)(1-\delta)f_{Nash,t+1}^i V_{Nash,t+1}^{i,new} \right). \quad (\text{B.16})$$

Combining equations (B.13) and (B.16) to eliminate $J_{Nash,t}^i$ gives the expression of the wage given in equation (18).

Injection of the expression for the Nash wage—equation (18)—into equation (B.12) gives:

$$J_{Nash,t}^{i,new} = (1-\gamma)A_t^i (F'(N_{Nash,t}^i) - z) + \beta E_t \left((1-d)J_{Nash,t+1}^i - (1-\gamma)(1-\delta)f_{Nash,t+1}^i V_{Nash,t+1}^{i,new} \right). \quad (\text{B.17})$$

Injecting the expression for the Nash wage (18) into equation (B.10) gives:

$$V_{Nash,t}^{i,new} = \gamma A_t^i (F'(N_{Nash,t}^i) - z) + \beta E_t \left((1-d) \frac{\gamma}{1-\gamma} J_{Nash,t+1}^i - \left(\gamma(1-\delta) + \delta - d \right) f_{Nash,t+1}^i V_{Nash,t+1}^{i,new} \right). \quad (\text{B.18})$$

An equilibrium under Nash bargaining is then processes for N_t^i , J_t^i , $J_t^{i,new}$ and $V_t^{i,new}$ that solve (B.16), (B.17) and (B.18), and the labor demand schedule (9)-(11), and where f_t^i is the function of N_t^i given by equation (15). This can be calculated recursively as follows. Given J_{t+1}^i , f_{t+1}^i and $V_{t+1}^{i,new}$ the expressions (B.16) and (B.17) solve for employment N_t^i when intersected with the labor demand schedule (9)-(11), and so for f_t^i . Equations (B.16) and (B.18) then allow us to calculate J_t^i and $V_t^{i,new}$. The Nash wage can then be recovered, e.g., by equation (B.13).

²⁷We use the fact that whenever $\delta_{t+1}^i \neq \delta$ and $d_{t+1}^i \neq d$ there are endogenous layoffs so $S_{t+1}^i = J_{t+1}^i = 0$ and $f_{t+1}^i = 0$, so we can replace δ_{t+1}^i with δ and d_{t+1}^i with d .

B.3 Normalization of μ

Recall that the matching function is Cobb-Douglas. The vacancy-filling rate is therefore $q_t^i = \mu\theta^{-\eta}$. Furthermore, the job finding rate is $f_t^i = \theta_t^i q(\theta_t^i)$. Combining these equations allows us to express the vacancy-filling rate as a function of the job-finding rate:

$$q_t^i = \mu^{\frac{1}{1-\eta}} (f_t^i)^{\frac{-\eta}{1-\eta}}. \quad (\text{B.19})$$

We can now see that there is a one-to-one mapping between the cost of hiring a worker $C_t^i \equiv c/q_t^i$ and the job-finding rate f_t^i :

$$C_t^i \equiv \frac{c}{q_t^i} = \left(c\mu^{\frac{-1}{1-\eta}} \right) (f_t^i)^{\frac{\eta}{1-\eta}}. \quad (\text{B.20})$$

This mapping can be used to write the equilibrium conditions of the model in terms of either the cost of hiring a worker or the job-finding rate, without reference to the other (and without reference to labor market tightness). When the model is written in this way (e.g., in terms of the cost of hiring a worker), the parameters c and μ only enter the model through the composite term $c\mu^{\frac{-1}{1-\eta}}$. This implies that we can normalize either c or μ without loss of generality. We choose to normalize $\mu = 1$. Intuitively, only the cost of hiring a worker matters to a firm. It is immaterial to the firm whether this cost consists of posting few vacancies that fill with a high probability but are expensive to post, or of posting many vacancies that fill with a low probability but are inexpensive to post.

B.4 Steady-State

In a non-stochastic steady-state equilibrium with $A_t = 1$ and $\Delta \log(Z_t^i) = 0$, equations (7) and (8) reduce, for $\tilde{J} = J^i/A^i$, $\tilde{J}^{new} = J^{i,new}/A^i$, and $\tilde{w} = w^i/A^i$, to

$$\tilde{J}(N, \tilde{w}) = \frac{1}{1 - \beta e^g(1 - \delta)} (F'(N) - \tilde{w}), \quad (\text{B.21})$$

$$\tilde{J}^{new}(N, \tilde{w}) = (1 - \beta e^g(d - \delta)) \tilde{J}(N, \tilde{w}). \quad (\text{B.22})$$

In steady-state, firms hire workers such that the law of motion of the stock of attached workers is given by equation (3). Combined with the definition of hires—equation (2)—this gives the steady-state relationship between M and N

$$M(N) = \frac{1 - d}{1 - d + \delta} N. \quad (\text{B.23})$$

Combined with equation (15), this gives $f(N)$ as a function of N alone, and through equation (B.19) it gives the hiring costs $c/q(N)$ as a function of N alone.

In steady-state, firms hire workers such that labor demand is given by equation (9). Combined with equation (B.22), this gives

$$\frac{c}{q(N)} = (1 - \beta e^g (d - \delta)) \tilde{J}(N, \tilde{w}). \quad (\text{B.24})$$

To calibrate the model, we assume that hiring costs $\frac{c}{q}$ are 10% of \tilde{J} , and obtain the steady-state wage as the only one consistent with equation (B.24) and a level of employment $N = 1 - U = 1 - 5.7\%$.

In steady-state, the wage is equal to the Nash-bargained wage. (Under downward nominal wage rigidity, we consider cases where $\log(\bar{\Pi}) + g \geq 0$ to make sure this can be the case.) We back out the value of a workers' bargaining power γ as the only value that make the steady-state Nash wage equal to the steady-state wage we have obtained.

B.5 Solution Method

The hiring decision of a firm in sector i is a function of four or five state variables depending on the process for aggregate productivity. These are: aggregate productivity A (and lagged aggregate productivity A_{-1} if aggregate productivity follows an AR(2)), idiosyncratic productivity growth ΔZ^i , the lagged wage w_{-1}^i , and the stock of workers inherited from the previous period, M^i . We have introduced sectoral heterogeneity in such a way that a firm does not need to forecast any endogenous aggregate variable in order to decide how many workers to hire. Therefore, we do not need to keep track of the endogenous aggregate state of the economy in order to solve for the hiring decision of a firm.

A solution to the model can be described as a pair of policy functions for N^i and $\tilde{J}^i = J^i/A^i$ over this (four or) five dimensional state space. We make the following change of variables. First, we define the AR(1) process:

$$\eta_t \equiv (1 - \rho_2^a L)^{-1} \varepsilon_t^a, \quad (\text{B.25})$$

so that:

$$\log(A_{t+1}) = \rho_1^a \log(A_t) + \eta_{t+1}. \quad (\text{B.26})$$

Second, we define $\zeta_{-1}^i \equiv w_{-1}^i/A^i$ the ratio of the lagged wage to current total productivity. Given these definitions, the five-dimensional state can be parameterized as $(\log(A), \eta, \Delta \log(Z^i), \zeta_{-1}^i, M^i)$.

We form a discrete grid of the state-space with 11 points along the exogenous dimensions $(\log(A), \eta, \Delta \log(Z^i))$ and 21 along the endogenous dimensions (ζ_{-1}^i, M^i) . We approximate the AR(1) processes for the exogenous variables η_t and $\Delta \log(Z_t^i)$ using the [Rouwenhorst \(1995\)](#) discretization method. The Rouwenhorst method is more accurate than the [Tauchen \(1986\)](#) method for persistence processes. [Petrosky-Nadeau and Zhang \(2017\)](#) emphasize this point in the context of the DMP model. Our approach to adapting the Rouwenhorst method to an AR(2) process is close to the one used by [Galindev and Lkhagvasuren \(2010\)](#), who consider the more general case of a VAR(1).

We solve for the policy functions at each point on the grid by policy function iteration. Specifically, we start from an initial guess for the policy functions for N^i and \tilde{J}^i . At each point of the grid, we then use this guess to calculate the expectation term in equation (7)-(8). In calculating the expectation term, we need to evaluate the policy function for \tilde{J}^i at points that are not on the grid. We do so through linear interpolation. Having calculated values for the expectation term, we compute the values of N^i and \tilde{J}^i that solve equations (9)-(11), (26), and (B.19), and store the resulting values in new policy functions. Once this has been done for all points on the grid, we have a new set of guesses for the policy functions. We repeat this process until the policy functions converge. In simulating the heterogeneous model, we assume 1000 sectors.

In solving for the model under wage-rigidity, we need to first solve for the Nash wage for each point of the grid. We do so by first solving the model under the assumption of Nash-bargaining using the same iterative method. The policy function under Nash bargaining does not depend on ζ_{-1}^i . Under symmetric real wage rigidity, wages in a sector may be so low relative to productivity in that sector that firms are willing to hire more workers than there are in the sector. We deal with such cases by assuming that firms hire all workers but no more.

C Common Wages, Constant Returns to Scale, and the Wage Band

We derive the upper-bound of the wage band in our model under the assumptions that workers must be paid the same wage, and that there are decreasing returns to scale in production. Here, we stress the role of these two assumptions by deriving the upper-bound of the wage band under the

alternative assumption that workers can be paid different wages or that there are constant returns to scale in production.

If the firm can pay its workers different wages, the only two parties involved in deciding the wage of an individual worker are the firm and this particular worker. The firm can approach the worker and threaten to lay him—and only him—off if he—and only he—does not accept a lower wage. This is a credible threat as long as the value of the worker to the firm is negative $J_t^i(M_t^i) < 0$. Since the worker would prefer keeping his job at a lower wage to the alternative of unemployment, such a lay-off would lead to private inefficiency. It implies that when workers can be paid different wages a privately efficient wage must be lower than the upper-bound

$$\bar{w}_t^i = A_t^i F'(M_t^i) + \beta E_t((1 - \delta) J_{t+1}^{i*}(M_{t+1}^i(M_t^i))). \quad (\text{C.1})$$

For a wage lower than \bar{w}_t^i the firm never endogenously lays off a worker whose match generates a positive total surplus.

The upper-bound (C.1) is lower than the upper-bound (16) in our baseline model. When the firm must pay its workers the same wage, approaching one individual worker to offer him a lower wage is impossible: all wages must be cut together. This dramatically changes the threat point of the firm. As long as there are decreasing returns to scale, the firm cannot credibly threaten to lay off all its workers if they do not collectively take a wage cut. Workers understand that if the firm lays off only a few workers, the marginal product of labor will increase, making the value to the firm of the remaining workers positive again and making the firm unwilling to lay off its entire workforce. This leads to the higher upper-bound (16).

Decreasing returns to scale in production are therefore essential to the expression of the upper-bound (16). Under constant returns to scale, if the wage is so high that $J_t^i < 0$, the firm lays off its entire workforce. As a consequence, *all* the firm's workers would be better off with a lower wage that allows all matches to be maintained. Any wage that leads to some lay-offs is privately inefficient, even if workers must be paid the same wage. Formally, under constant returns to scale both upper-bound (16) and upper-bound (C.1) reduce to

$$\bar{w}_t^i = A_t^i + \beta E_t((1 - \delta) J_{t+1}^{i*}). \quad (\text{C.2})$$

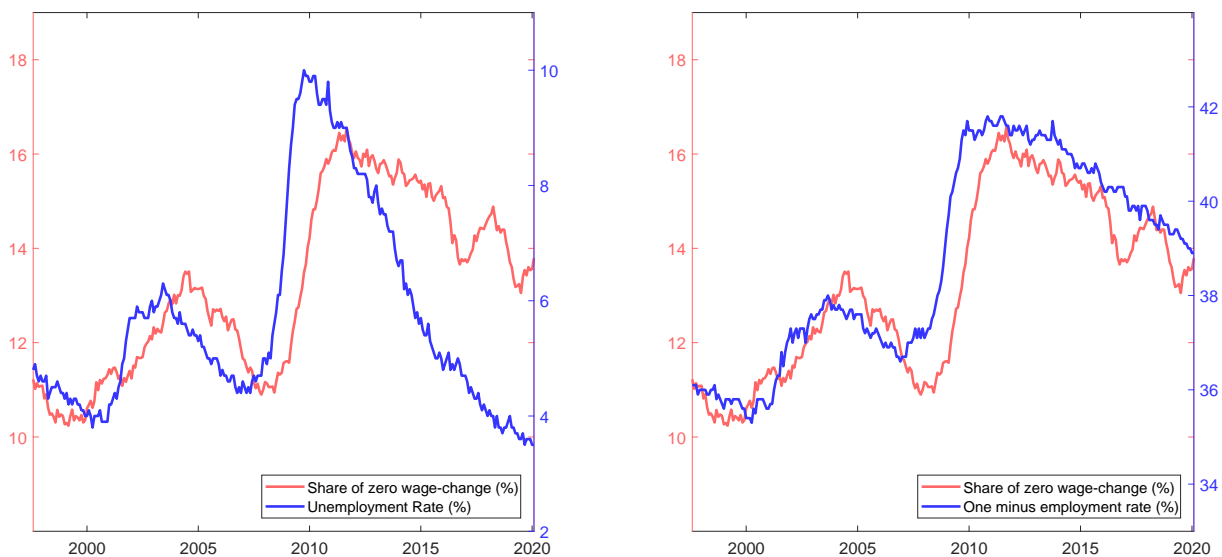


Figure D.1: San Francisco Fed Wage Rigidity Meter

Note: The figure plots the share of wage freezes of all job-stayers (paid by the hour or not) with respect to the wage one year prior, with no correction for measurement errors. This series is constructed by the Federal Reserve Bank of San Francisco using data used from the Current Population Survey.

D Wage Rigidity

Downward nominal wage rigidity appears to be highly pro-cyclical in the US. The Federal Reserve Bank of San Francisco releases a Wage Rigidity Meter, which measures the fraction of job-stayers that experience a one-year wage freeze (Daly, Hobijn, and Wiles, 2011; Daly, Hobijn, and Lucking, 2012; Daly and Hobijn, 2014).²⁸ The left panel of Figure D.1 plots this series along with the unemployment rate for the period 1997-2019.²⁹ The right panel of Figure D.1 plots the same series along with the non-employment rate for the working age population (one minus the employment rate). The correlation between wage freezes and these series is striking. The fraction of wage freezes rises rapidly in each of the three recessions that occur in this sample period.

There is a large literature extending back several decades that seeks to estimate the extent of downward nominal wage rigidity. An important empirical challenge in this literature is measurement error in reported wages. The Wage Rigidity Meter reports the fraction of wage freezes with no correction for measurement error. Early work using data from the Panel Study of Income Dynamics (PSID) and the Current Population Survey (CPS) includes McLaughlin (1994),

²⁸<https://www.frbsf.org/economic-research/indicators-data/nominal-wage-rigidity/>

²⁹Before 1997, the Wage Rigidity Meter does not report a continuous data series (e.g., no data are reported for 1994-1997). In addition to this, inflation is higher before this, leading to a weaker relationship.

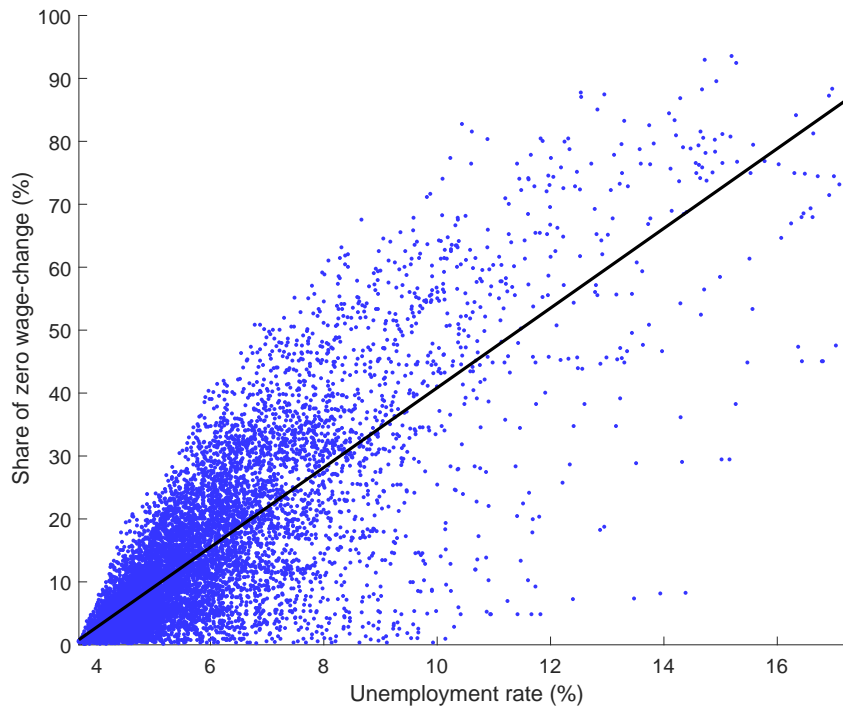


Figure D.2: Relationship between the Unemployment Rate and the Fraction of Wage Freezes

Note: The figure plots a scatter plot of the relationship between the unemployment rate and the fraction of wage freezes in our model with secure and insecure jobs, multiple sectors, an AR(2) in aggregate productivity and downward nominal wage rigidity. The regression line is added in black. A wage freeze is defined as a wage that has remained constant over the past twelve months. The simulated sample is 10,000 months.

Kahn (1997), and Card and Hyslop (1997). Altonji and Devereux (2000) report larger amounts of downward nominal wage rigidity and virtually no wage cuts in the PSID after correcting for measurement error. Gottschalk (2005) and Barattieri, Basu, and Gottschalk (2014) report similar figures based on analysis of the Survey of Income and Program Participation and correcting for measurement error in a different way. More recently, Hazell and Taska (2018) and Grigsby, Hurst, and Yildirmaz (2018) find substantial downward nominal wage rigidity in firm wage posting data and administrative payroll data, respectively.

Figure D.2 presents a scatter plot of the unemployment rate and the fraction of sectors with wage freezes in our full model. It shows that the share of workers facing wage freezes in our full model is highly correlated with the unemployment rate as in the data. The contemporaneous correlation is 0.80. For comparison, the contemporaneous correlation between the San Francisco Wage Rigidity Meter and unemployment is 0.52, and reaches 0.89 when the Wage Rigidity Meter is lagged by 18 months.

E Scatter Plots Displaying Plucking Property in Models

The results presented in Table 4 in the main text indicate that our simple model (with symmetric real wage rigidity, a single type of job, a single sector, and an AR(1) process for productivity) generates a minimal amount of Friedman’s plucking property, while our full model (with downward nominal wage rigidity, two types of jobs, multi sectors, and an AR(2) process for productivity) can roughly match the data in this regard. Figure E.1 provides a visual illustration of these findings. We display a scatter plot for each model of the amplitude of subsequent expansions on the amplitude of contractions (left panels) and of the amplitude of subsequent contractions on the amplitude of expansions (right panels). This figure is the counterpart for our simple and full models of Figure 2 in the main text.

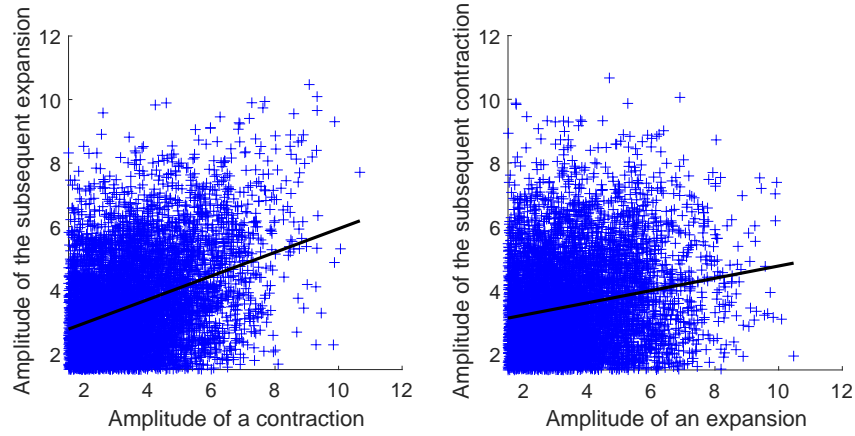
F Asymmetry in the Simple Model

Our simple model does not feature substantial asymmetry. To understand this better, it is useful to go through the different components of the model that might give rise to asymmetry. One possibility is that asymmetry arises outside of the labor market. Any such non-labor-market asymmetry would enter our model through non-linearity in the exogenous productivity process that drives cycles in our model—we think of this productivity process as a reduced form representation of all sources of shifts to labor demand. However, we have chosen to model this exogenous process as being linear, and therefore shut down this source of asymmetry.³⁰ A second potential source of asymmetry is the wage-setting rule. However, in the simple model, we have assumed a symmetric wage-setting rule.

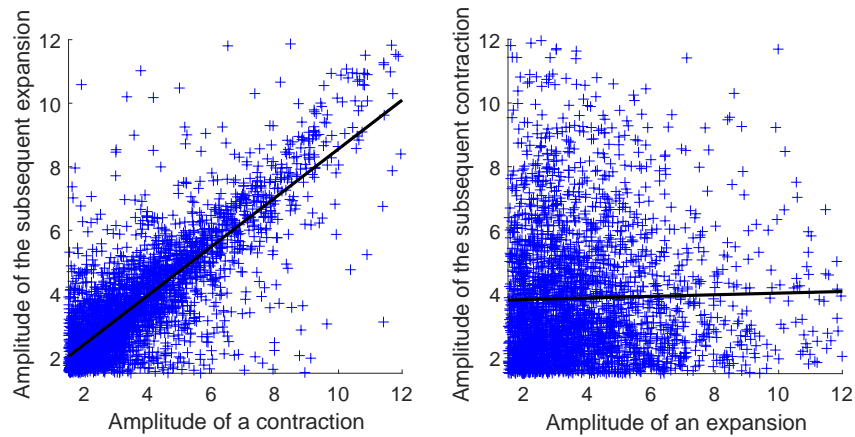
This leaves three components of the model as potential sources of asymmetry: the labor-demand schedule of individual firms—equations (9)-(11)—the relationship between the vacancy-filling probability q_t and the job-finding probability f_t implied by the matching function, and the worker-flow relationship—equations (28) and (29). It is useful to combine these three relationships into an aggregate labor demand schedule.

To this end, notice that in our simple model with a single type of job $\delta = d$, and with a separation rate δ calibrated to match the large inflows of workers into unemployment each month, firms

³⁰We do this for two reasons. First, we are interested in exploring the ability of the DMP model to generate asymmetry within the labor market. Second, prior empirical work has found asymmetries to be more pronounced in the unemployment rate than in other macroeconomic data suggesting that the source of asymmetry is the labor market (e.g., McKay and Reis, 2008).



(a) Simple Model



(b) Full Model

Figure E.1: Plucking Scatter Plots

Note: The figure displays the scatter plots associated to the plucking regressions for our simple model (panel a) and our full model (panel b). The plots feature all the expansion/contraction pairs obtained by pooling together 250 samples of 866 months: 6803 and 6553 in the upper panels and 3249 and 2999 in the lower panels. OLS regression lines are plotted in each panel.

virtually never lay off workers above and beyond the workers who leave through exogenous separation. As a result, a very good approximation is to assume that $H_t \geq 0$ in all periods, so that the law of motion of employment is equation (3) in all periods.

Given this, equation (26) gives the job-finding rate as a function of past and present employment

$$f_t = \frac{N_t - (1 - \delta)N_{t-1}}{1 - (1 - \delta)N_{t-1}}. \quad (\text{F.1})$$

Combining this equation with equation (B.20) yields

$$C(N_t, N_{t-1}) \equiv \frac{c}{q_t} = c\mu^{\frac{-1}{1-\eta}} \left(\frac{N_t - (1 - \delta)N_{t-1}}{1 - (1 - \delta)N_{t-1}} \right)^{\frac{\eta}{1-\eta}}. \quad (\text{F.2})$$

Next consider the labor demand schedule. Since $\delta = d$ and firms never endogenously lay off workers. Labor demand therefore reduces to equations (8) and (9), which can be combined into

$$A_t C(N_t, N_{t-1}) = A_t F'(N_t) - w_t + \beta(1 - \delta)E_t (A_{t+1} C(N_{t+1}, N_t)). \quad (\text{F.3})$$

Dividing by A_t yields

$$\frac{w_t}{A_t} = F'(N_t) - C(N_t, N_{t-1}) + \beta(1 - \delta)E_t \left(\frac{A_{t+1}}{A_t} C(N_{t+1}, N_t) \right), \quad (\text{F.4})$$

The difference between the individual labor demand schedule—equations (9)-(11)—and the *aggregate* labor demand schedule (F.4) captures the fact that in the search model, search costs are *external* labor adjustment costs. Therefore, a given firm sees the cost of hiring a worker as constant, but in equilibrium the marginal cost of hiring a worker is increasing in employment because of search externalities.

Because our model features hiring costs, the aggregate labor demand schedule—equation (F.4)—is forward-looking: what matters to the hiring decision of a firm is not only today's real wage and today's marginal productivity of a worker, but also the expected future value of these variables. One way to see this formally is to iterate equation (F.4) forward:

$$C(N_t, N_{t-1}) = E_t \left(\sum_{k=0}^{\infty} \beta^k (1 - \delta)^k \frac{A_{t+k}}{A_t} \left(F'(N_{t+k}) - \frac{w_{t+k}}{A_{t+k}} \right) \right). \quad (\text{F.5})$$

However, it is well known that the DMP model has minimal internal propagation. It is therefore useful to abstract initially from the intertemporal dimension of hiring by considering the steady-state aggregate labor demand schedule induced by the model. This corresponds to the demand for labor should the real wage remain constant and productivity grow deterministically. It is given by:

$$\frac{w}{A} = F'(N) - K(N)(1 - \beta e^g (1 - \delta)) \quad (\text{F.6})$$

where

$$K(N) \equiv C(N, N) = c\mu^{\frac{-1}{1-\eta}} \left(\frac{\delta N}{1 - (1 - \delta)N} \right)^{\frac{\eta}{1-\eta}}. \quad (\text{F.7})$$

Equation (F.6) makes apparent the two reasons why labor demand is downward sloping in a search model. First, there are decreasing returns to scale in production (first right-hand-side term in equation (F.6)). This is what makes labor demand downward sloping in a model without search frictions. Second, the marginal search cost is increasing in employment (second right-hand-side term in equation (F.6)). This is the motive that is specific to a search model.

With a constant-elasticity production function, the first term on the right-hand side of equation (F.6) is log-linear. In contrast, the function $K(N)$ that arises from search costs makes the log of employment a convex function of the log of the real wage. The function $K(N)$ combines two features of the model: the matching function and the worker-flow relationship (28)-(29). Our assumption of a Cobb-Douglas matching function with $\eta = 0.5$ implies that the relationship between the cost of hiring c/q_t and the job-finding rate f_t is linear (see equation (B.20)). The convexity of the relationship between log employment and log real wages therefore arises primarily from the worker-flow relationship.

F.1 Non-Linearity and Returns to Scale

The limited degree of non-linearity that arises from the worker-flow relationship (28)-(29) implies that the simple model's aggregate labor demand schedule is close to log-linear. The left panel of Figure F.1 plots the steady state labor demand schedule for the simple model as well as the labor demand schedule with no search frictions ($K(N) = 0$).³¹ The labor demand schedule in the simple model is not completely log-linear, but the amount of non-linearity is quantitatively minimal. The left panel of Figure F.1 also includes a scatter plot of the relationship between the log of the wage-to-productivity ratio and unemployment away from steady-state. The cloud of points is slightly more convex than the steady-state relationship, but even this degree of non-linearity is quantitatively minimal.

The extent of non-linearity in the aggregate labor-demand schedule of our simple model depends significantly on the assumed shape of the production function. To illustrate this, we consider a version of our baseline model with constant returns to scale as opposed to decreasing returns to scale. We recalibrate the bargaining power of workers γ and the volatility of aggregate

³¹We plot these relationships with unemployment on the y-axis. But the results are virtually indistinguishable with $-\log(N)$ on the y-axis, because the linear approximation $u \simeq -\log(N)$ holds well over the range of values we consider.

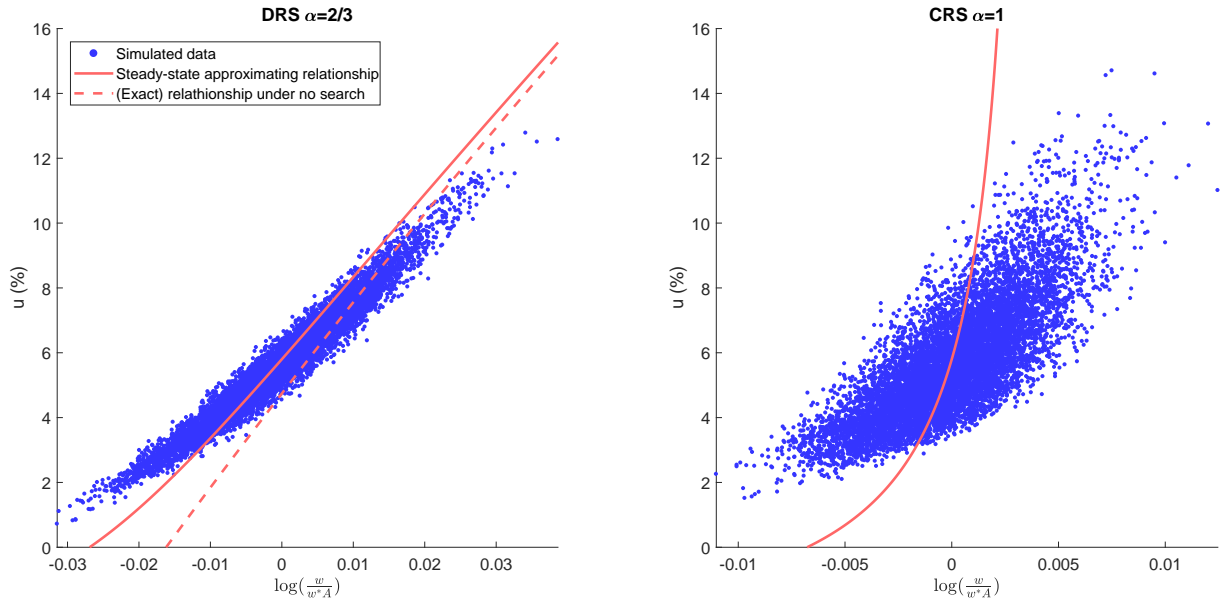


Figure F.1: Aggregate Labor Demand Schedule under Constant vs. Decreasing Returns to Scale

Note: The figure plots the relationships between the log of the wage-to-productivity ratio (normalized by the steady-state wage w^*) and unemployment in the simple DMP model with one type of job, no sectoral heterogeneity, symmetric real wage rigidity and AR(1) productivity shocks. The left panel considers the case of decreasing returns to scale (DRS). The right panel considers the case of constant returns to scale (CRS).

shocks σ^a so that the average level of unemployment in this version of the model still equals the average level of unemployment in the data (5.7%) and the standard deviation of unemployment in the model still equals its value in the data (1.6 percentage points). Because employment responds more strongly to changes in the wage-to-flexible-wage ratio with a constant returns to scale production function, the resulting value of σ^a is substantially smaller (0.75%) than in the baseline model with decreasing returns (2.0%). To solve the model under constant returns to scale, we assume that in extreme nodes where firms would like to endogenously lay off workers (which almost never occur on an equilibrium path in our calibration since the monthly exogenous separation rate is $s = 5\%$), they simply freeze hires.

Table F.3 presents results on the business cycle statistics we focus on for the simple model with constant returns to scale and compares these to our simple model with decreasing returns to scale. The constant-returns model generates substantially more plucking than the decreasing-returns model: the regression coefficient for the size of a subsequent expansion on the size of a contraction is now 0.54, substantially larger than the regression coefficient for the size of a subsequent contraction on the size of an expansion (0.06). The difference in the explanatory power of the two regressions is also larger. The R^2 of the former regression is 0.31, while that of the later regression

is only 0.02. Furthermore, the model with constant returns to scale generates more skewness in the unemployment rate (0.72).

To provide further insight about the greater asymmetry of the search model with constant returns to scale, the right panel of Figure F.1 presents results on the shape of the aggregate labor demand schedule in this model. The relationship between unemployment and the logarithm of the wage-to-flexible wage ratio is clearly substantially more non-linear in the constant returns to scale case. With constant returns to scale the shape of this relationship is governed entirely by search costs which yield convexity.

Figure F.1 also illustrates that with constant returns to scale, the unemployment rate rises to very high values with a higher frequency than in the baseline model. This occurs despite the fact that we calibrate the constant-returns version of the model to match the standard deviation of unemployment in the data. Petrosky-Nadeau and Zhang (2017); Petrosky-Nadeau, Zhang, and Kuehn (2018) stress this property of the constant-returns model.

The greater extent of plucking in the constant-returns model is mostly driven by these extreme events, which have no equivalent in post-WWII US data. Over our sample period, the largest expansion or contraction—the 2009-2019 expansion—was 6.5 percentage points in size. Table F.3 presents results on the degree of plucking for a “top-truncated” sample that only includes expansions and contractions that are less than 6.5 percentage points in size. The degree of plucking in the model with constant returns to scale is considerably reduced in the truncated sample. Therefore, although the baseline DMP model with constant returns to scale can generate an appreciable amount of plucking, it can do so only through counter-factually large expansions and contractions.

G Endogenous Separation

In our full model, unemployment fluctuates due to variation in the job-finding rate f_t and the employment exit rate \bar{s}_t . The employment exit rate, which averages 2.8% in our simulations, undergoes significant fluctuations over the business cycle. The standard deviation of the employment exit rate is 0.4 percentage points, close to the 0.5 percentage point standard deviation of the empirical series calculated by Shimer (2012).³² As a result inflows into unemployment partly explain fluctuations in unemployment in our model, although fluctuations in unemployment remain primarily determined by fluctuations in the job-finding rate, consistent with what Elsby, Michaels,

³²The 2.8% average is also close to the average of the empirical series calculated by Shimer: 3.4%.

Table F.3: Decreasing vs. Constant Returns to Scale

	Data		Simple Model		
			Decreasing Returns		Constant Returns
				Trunc.	Trunc.
Subsequent expansion on contraction, β	1.09	0.37	0.25	0.54	0.40
		(0.22)	(0.21)	(0.23)	(0.23)
Subsequent contraction on expansion, β	-0.38	0.18	0.13	0.06	0.05
		(0.22)	(0.21)	(0.21)	(0.23)
Subsequent expansion on contraction, R^2	0.58	0.14	0.07	0.31	0.16
		(0.14)	(0.11)	(0.20)	(0.16)
Subsequent contraction on expansion, R^2	0.22	0.04	0.03	0.02	0.03
		(0.09)	(0.08)	(0.06)	(0.08)
Skewness	0.63	0.26	0.20	0.72	0.57
		(0.22)	(0.20)	(0.37)	(0.28)
Speed of expansions (pp / year)	0.88	3.59	3.59	3.27	3.25
		(0.36)	(0.36)	(0.39)	(0.39)
Speed of contractions (pp / year)	1.89	3.42	3.42	3.08	3.04
		(0.33)	(0.34)	(0.35)	(0.36)
Duration of expansions (months)	57.9	15.0	14.2	16.5	15.8
		(1.8)	(1.8)	(2.3)	(2.2)
Duration of contractions (months)	26.9	15.6	14.7	17.3	16.5
		(1.9)	(1.8)	(2.2)	(2.2)

Note: The table compares data from our simple model (i.e. with symmetric real wage rigidity, a single type of job ($\delta = d$), no sectoral heterogeneity, and aggregate productivity following an AR(1) process) under decreasing returns to scale (DRS) and constant returns to scale (CRS) along four dimensions. The first (third) row reports the coefficient (R^2) in an OLS regression of the size of an expansion (percentage point fall in unemployment rate) on the size of the previous contraction (percentage point increase in unemployment rate). The second (fourth) row report the coefficient (R^2) in an analogous regression of the size of a contraction on the size of the previous expansion. The fifth row reports the skewness of the distribution of the unemployment rate. The next two rows report the average speed of expansion and contractions, measured in percentage points of unemployment per year. The final two rows report the average duration of expansions and contractions, measured in months. For the models, the reported point estimate is the median value of the statistic over 5000 samples of 866 periods each (the length of our sample of real-world data). The standard error reported in parentheses is the standard deviation of the estimates across the 5000 samples. For both constant and decreasing returns to scale, we also report the results for the “top-truncated” samples that only include expansions and contractions of less than 6.5 percentage points.

and Solon (2009) and Shimer (2012) document on US data.

Endogenous fluctuations in \bar{s}_t occur for two reasons in our model. First, our model allows for the endogenous layoff of some securely attached workers, when the wage in a sector becomes too high relative to productivity. To quantify such separations, we can define the aggregate rate of exit from secure employment

$$s_{0,t} = \delta + \frac{M_t - N_{0,t}}{N_{0,t-1}}, \quad (\text{G.1})$$

where the second term accounts for endogenous separation.³³ Figure G.1 gives a simulated path for the rate of exit from secure employment $s_{0,t}$ in our full model. Most of the time, the rate of exit from secure employment is close to the exogenous rate of separation of securely attached workers $\delta = 0.2\%$ per month. But it spikes up at the onset of recessions, when a large negative shock pushes firms in many sectors to resort to layoffs. In large recessions $s_{0,t}$ can then reach 1.5% per month. These bursts of layoffs play an important role in the fast increases of unemployment in contractions, and therefore in the speed asymmetry between expansions and contractions.

A second reason why the employment exit rate \bar{s}_t fluctuates is that the bursts of layoffs discussed above result in a shift in the composition of workers between long-term secure jobs and short-term insecure jobs. Immediately following a layoff burst, a larger fraction of workers are in short-term insecure jobs and therefore have high exit rates from employment. This results in an elevated exit rate from employment for some time at the beginning of expansions while workers are churning through short-term jobs in search of a good match. As time passes, more and more such workers find long-term stable employment and the employment exit rate falls.

The main shortcoming of our model when it comes to inflows and outflows is that the job finding rate falls more sharply at the beginning of severe recessions in our model than in the data. Since very few people are hired at these times, the fraction of workers in short-term insecure jobs is briefly very low which results in a brief dip in the employment exit rate. This shortcoming stems from the fact that in our model there is no heterogeneity within sector, which implies that sectors with endogenous layoffs have zero hiring. In reality, within-sector heterogeneity implies that some firms are hiring while others are laying off workers even at the onset of a recession.

³³Given our assumption that workers who lose their job at the end of a period directly join the pool of job-seekers, workers who exit from a secure job may in principle find a short-term job straight away and not exit employment. However, in our model when firms in a sector lay off securely attached workers they do not hire. This implies that all laid-off workers become unemployed.

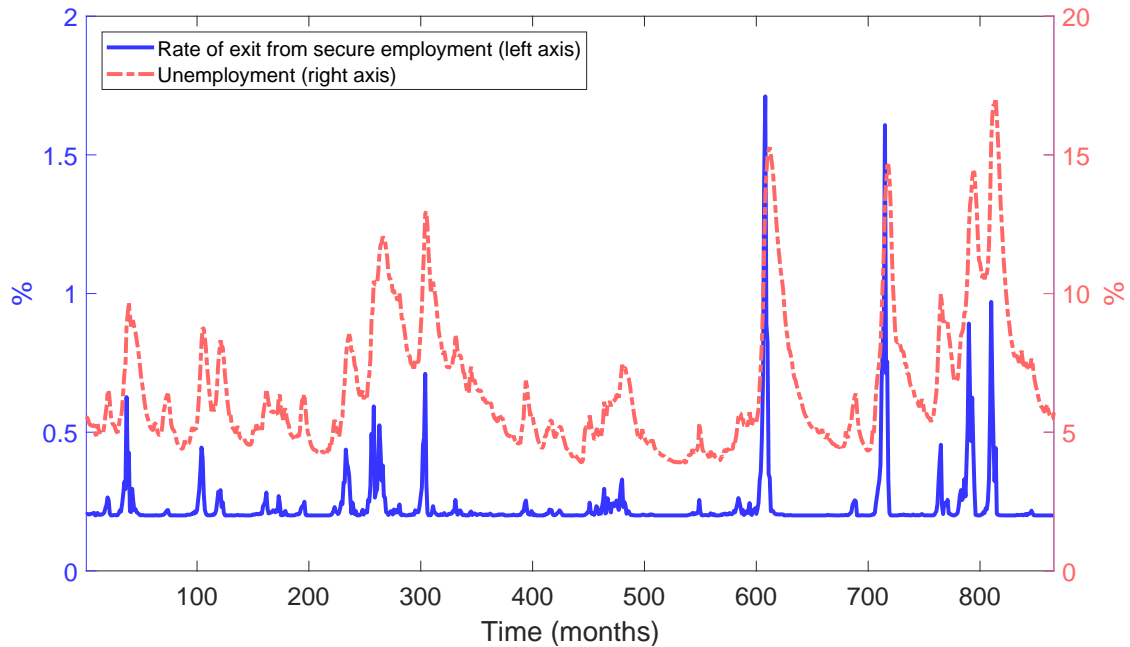


Figure G.1: Simulated Paths for the Employment Exit Rates

Note: The figure plots sample paths of 72 years (the same length as our empirical sample) for the rate of exit from secure employment $s_{0,t}$ and the unemployment rate for our full model with downward nominal wage rigidity, secure and insecure jobs, multiple sectors and an AR(2) aggregate productivity process.

H Efficient Level of Unemployment and Welfare

In our DMP-style search model, unemployment is potentially a productive state since it facilitates the creation of matches in the labor market. The efficient level of unemployment is therefore potentially larger than zero. Below, we calculate the efficient level of unemployment for the steady state of our model. Sectors in our model differ only by the shocks that hit them. To characterize the efficient allocation in steady-state, we can, therefore, consider the one-sector version of the model. We say that an allocation is efficient if it maximizes the equally-weighted sum of workers' individual utilities.

An individual worker $j \in [0, 1]$ has intertemporal linear utility

$$E_0 \sum_{t=0}^{\infty} \beta^t \left(c_t^j - \xi_t^j A_t \right), \quad (\text{H.1})$$

where c_t^j is his consumption and $\xi_t^j A_t$ is the cost of the leisure he foregoes, expressed in units of consumption. The cost of foregone leisure ξ_t^j is ξ_e when the worker is employed and ξ_u when the worker is unemployed and searching for a job. Consumption also takes two values: $c_{t,e}$ when the worker is employed and $c_{t,u}$ when the worker is unemployed.

Aggregate consumption is equal to aggregate production net of hiring costs

$$N_t c_{t,e} + (1 - N_t) c_{t,u} = C_t = A_t F(N_t) - A_t c V_t,$$

where $V_t = H_t/q_t$ are vacancies. The aggregate cost of foregone leisure is

$$\int_i \xi_t^j di = \xi_u + \zeta N_t,$$

where we define $\zeta \equiv \xi_e - \xi_u$ the leisure cost of working relative to being unemployed, as used in equation (12).

Using these expressions, we can write the social-welfare objective as

$$\max_{V_t, N_t, M_t} E_0 \sum_{t=0}^{\infty} \beta^t A_t \left(F(N_t) - c V_t - \zeta N_t \right), \quad (\text{H.2})$$

$$\text{s.t. } N_t = M_t + q \left(\frac{V_t}{1 - M_t} \right) V_t, \quad (\text{H.3})$$

$$\text{s.t. } M_t = (1 - \delta) M_{t-1} + (1 - d)(N_{t-1} - M_{t-1}), \quad (\text{H.4})$$

$$\text{s.t. } N_t \leq 1, \quad (\text{H.5})$$

which takes into account that the law of motion of the workforce (H.4) is equation (3) because firms hire in steady-state, $H_t > 0$. Note that unemployment benefits do not appear in the social planner's program. Indeed, since unemployment benefits must be paid for by a decrease in the earnings of employed workers, they are purely redistributive and with linear utility fully cancel out in aggregate.

The social planner chooses employment N_t and M_t and vacancies V_t to maximize the objective (H.2). The first-order conditions of the social planners problem are

$$V_t : J_t^{new} = \frac{c}{q(\theta_t) + \theta_t q'(\theta_t)}, \quad (\text{H.6})$$

$$N_t : J_t^{new} = F'(N_t) - \zeta - \mu_t + \beta E_t \left(\frac{A_{t+1}}{A_t} (1 - d) J_{t+1} \right) = 0, \quad (\text{H.7})$$

$$M_t : J_t = J_t^{new} (1 + \theta_{t+1}^2 q'(\theta_{t+1})) + \beta E_t \left(\frac{A_{t+1}}{A_t} (d - \delta) J_{t+1} \right), \quad (\text{H.8})$$

where $A_t J_t^{new}$ denotes the Lagrange multiplier on the constraint (H.3), $A_t J_t$ denotes the Lagrange multiplier on the constraint (H.4), and $A_t \mu_t$ denotes the Lagrange multiplier on the constraint (H.5).

Combine the three conditions to eliminate J_t and J_t^{new} . The steady-state version of the resulting condition with a Cobb-Douglas matching function $q(\theta) = \theta^{-\eta}$ is

$$\frac{c}{q(N)(1 - \eta)} = \frac{F'(N) - \zeta}{1 - \frac{\beta e^g (1 - d)}{1 - \beta e^g (d - \delta)} (1 - \eta f(N))}, \quad (\text{H.9})$$

if the solution is such that $N \leq 1$, and $N = 1$ otherwise, where $f(N)$ is defined by (15) and (B.23), and $q(N) = f(N)^{\frac{-\eta}{1-\eta}}$. Equation (H.9) determines the efficient steady-state level of employment as the level that equates the benefits of higher employment in terms of higher production to its cost in terms of higher marginal hiring costs. In equalizing costs and benefits, the social planner takes into account how the numbers of vacancies and job-seekers affect the probability of matching.

The flow value of unemployment that is relevant for the efficient level of unemployment—the social value of unemployment—is ζ , the difference in disutility of foregone leisure between employment and unemployment (job-seeking). Contrary to the private value of unemployment relevant to the outside option of a worker bargaining his wage in a decentralized equilibrium $z = b + \zeta$, it does not include unemployment benefits b since they do not constitute extra resources in aggregate.

We consider two assumptions on the calibration of ζ . The first one follows our calibration of the flow value of unemployment $z = b + \zeta$ when solving for the equilibrium of the model, as 70% of the labor share α . This calibration follows Hall and Milgrom (2008), who further decompose z as 45% of the wage for the difference in the disutility of foregone leisure ζ and 25% of the wage for unemployment benefits b . We therefore calibrate ζ as 45% of the labor share α , i.e. 0.3. This calibration is based on the assumption that a job-seeker forgoes no leisure, so that unemployment is akin to inactivity. In the second case we assume that a job-seeker forgoes as much leisure as an employed worker does. In this case, $\zeta = \xi_e - \xi_u = 0$ since employment is no more unpleasant than job-seeking.³⁴

For both calibrations of ζ , the level of employment that solves equation (H.9) is greater than 1. In other words, the efficient unemployment rate is zero. However, the efficient level of unemployment is somewhat dependent on the exact timing assumptions of our model. Our timing assumes that workers who lose their jobs at the end of period t immediately join the pool of job-seekers and get a chance of spending no period unemployed. An alternative is to assume that workers who lose their jobs in period $t - 1$ can only start looking for a new job at $t + 1$ and necessarily spend period t unemployed. The pool of job-seekers at t is then $1 - N_{t-1}$ instead of $1 - M_t$. In this case, the expression of the job-finding rate is not (15) but

$$f_t = \frac{N_t - M_t}{1 - N_{t-1}}. \quad (\text{H.10})$$

³⁴This alternative calibration of $\zeta = 0$ makes virtually no change to the equilibrium Nash-wage, and therefore virtually no change to the equilibrium of our model.

In steady-state, it implies:

$$f = \frac{\frac{\delta}{1-d+\delta}}{\frac{1}{N} - 1} = \frac{s}{\frac{1}{N} - 1}. \quad (\text{H.11})$$

As a consequence, in steady state $u = s/(f + s)$ is necessarily greater than $s/(1 + s)$ since the job-finding probability f cannot be greater than 1. This implies that the unemployment rate has a lower bound that is larger than zero. This lower bound on unemployment is a direct consequence of the assumption that finding a new job takes at least one month: even if all job-seekers get employed ($f = 1$), there are still about s unemployed workers in the economy. A characterization of the efficient level of unemployment under this alternative timing can be obtained along the same lines as the derivation above. The efficient allocation hits the constraint $f = 1$, so that the efficient level of unemployment is the lower bound $s/(1 + s)$. It is $u = 4.8\%$ if we still calibrate $s = 5.0\%$. Under this alternative timing however, the value of s consistent with $f = 45\%$ and $u = 5.7\%$ is now $s = 2.7\%$. For this alternative calibration, the lower bound on unemployment is 2.6%.

A continuous-time version of the model avoids the discrete timing assumption discussed above and is in effect a midway assumption on the easiness of finding a job. The law-of-motion of employment in the continuous-time version of the model is

$$\dot{N}_t = -\delta N_t + (1 - d + \delta)H_t. \quad (\text{H.12})$$

The efficient level of employment in the continuous-time version of the model satisfies

$$\frac{c}{q(N)(1 - \eta)} = \frac{F'(N) - \zeta}{\frac{r-g+\delta}{1-d+\delta} + \eta f(N)}. \quad (\text{H.13})$$

where $f(N)$ is given by (H.11), $q(N) = f(N)^{\frac{-\eta}{1-\eta}}$, r is the interest rate, and f is now the arrival rate of a Poisson process, and can therefore be greater than one. The solution to (H.13) is $u = 0.7\%$ when $\zeta = 0.3$, and $u = 0.5\%$ when $\zeta = 0$. Table H.1 summarizes our results for the efficient level of unemployment.

We now calculate the welfare gains of eliminating aggregate fluctuations in our model of downward nominal wage rigidity. We perform two simulations of our model. First, we simulate the model subject to both idiosyncratic and aggregate shocks. Second, we simulate it subject to idiosyncratic shocks only. Following Lucas (2003), in the latter case we set the constant level of productivity to $\exp(\frac{1}{2}\sigma_a^2)$ to guarantee that the average level of productivity in levels $E(A_t)$ is the

Table H.1: Efficient Level of Unemployment

ζ	Efficient Unemployment		
	Timing 1	Continuous-Time	Timing 2
0.3	$u = 0\%$	$u = 0.7\%$	$u = 2.6\%$
0	$u = 0\%$	$u = 0.5\%$	$u = 2.6\%$

Note: Timing 1 is the timing assumed in the core of the paper. It assumes workers losing their jobs in period t can start looking for a new job at t and get a chance of spending no period unemployed. Timing 2 is the alternative assumption that workers losing their jobs in period t can only start looking for a new job at $t + 1$ and necessarily spend period t unemployed. This issue does not arise in the continuous-time version of the model, which is in effect halfway between the two timing assumptions in discrete time.

Table H.2: Welfare Gains of Eliminating Business Cycles

ζ	Consumption-Equivalent
	Welfare Gains
0.3	0.65%
0	1.20%

same between the two simulations. In both cases we calculate the consumption-equivalent utility level at t as:

$$\mathcal{W}_t = \frac{1 - \beta e^g}{\int_0^1 A_t^i di C^*} E_t \left(\sum_{k=0}^{\infty} \beta^k \int_0^1 A_{t+k}^i e^{gk} \left(F(N_{t+k}^i) - cV_{t+k}^i I_{V_{t+k}^i \geq 0} - \zeta N_{t+k}^i \right) di \right), \quad (\text{H.14})$$

where $C^* = 0.97$ is the steady-state level of consumption. We run both simulations for 100,000 months and 1000 sectors. We compute the average utility level \mathcal{W} as the time-average of \mathcal{W}_t over the simulated series \mathcal{W}_t . In computing \mathcal{W}_t , we truncate the infinite sum in equation (H.14) to a horizon of 1000 periods. The welfare gains of eliminating aggregate fluctuations are the difference between the level of welfare we calculate from the two simulations. These welfare gains results are reported in Table H.2. In the case where we assume that workers forgo no leisure during unemployment, the welfare gains from eliminating aggregate fluctuations are 0.65% of steady-state consumption. In the case, where workers forgo as much leisure when unemployed as when they are working—i.e., enjoy their time unemployed as much/little as when they are working—the welfare gains from eliminating aggregate fluctuations are 1.20% of steady-state consumption.

References

- ABBRIITI, M. AND S. FAHR (2013): "Downward Wage Rigidity and Business Cycle Asymmetries," *Journal of Monetary Economics*, 60, 871–886.
- AIYAR, S. AND S. VOIGTS (2019): "The Negative Mean Output Gap," IMF Working Paper No. 19/183.
- AKERLOF, G., W. DICKENS, AND G. PERRY (1996): "The Macroeconomics of Low Inflation," *Brookings Papers on Economic Activity*.
- ALTONJI, J. AND P. DEVEREUX (2000): "The Extent and Consequences of Downward Nominal Wage Rigidity," in *Worker Well-Being*, ed. by S. Polachek, New York: Elsevier, 383–431.
- ANDRADE, P., J. GALI, H. LE BIHAN, AND J. MATHERON (2016): "The Optimal Inflation Target and the Natural Rate of Interest," NBER Working Paper No. 24328.
- BARATTIERI, A., S. BASU, AND P. GOTTSCHALK (2014): "Some Evidence on the Importance of Sticky Wages," *American Economic Journal: Macroeconomics*, 6, 70–101.
- BARLEVY, G. (2004): "The Cost of Business Cycles Under Endogenous Growth," *American Economic Review*, 94, 964–989.
- BARRO, R. (1977): "Long-Term Contracting, Sticky Prices, and Monetary Policy," *Journal of Monetary Economics*, 3, 305–316.
- BARRO, R. J. (2009): "Rare Disasters, Asset Prices, and Welfare Costs," *American Economic Review*, 99, 243–264.
- BENIGNO, P. AND L. A. RICCI (2011): "The Inflation-Output Trade-Off with Downward Wage Rigidities," *American Economic Review*, 101, 1436–1466.
- BENIGNO, P., L. A. RICCI, AND P. SURICO (2015): "Unemployment and Productivity in the Long Run: The Role of Macroeconomic Volatility," *Review of Economics and Statistics*, 97, 698–709.
- BEWLEY, T. (1999): *Why Wages Don't Fall During a Recession*, Harvard University Press.
- BLANCHARD, O. AND J. GALI (2010): "Association Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment," *American Economic Journal: A New Keynesian Model with Unemployment*, 2, 1–30.
- BLANCO, A. (2018): "Optimal Inflation Target in an Economy with Menu Costs and a Zero Lower Bound," University of Michigan Working Paper.
- BORDO, M. AND J. HAUBRICH (2012): "Deep Recessions, Fast Recoveries, and Financial Crises: Evidence from the American Record," NBER Working Paper No. 18194.
- BRY, G. AND C. BOSCHAN (1971): *Cyclical Analysis of Time Series: Selected Procedures and Computer Programs*, NBER.
- CABALLERO, R. AND M. HAMMOUR (1998): "The Macroeconomics of Specificity," *Journal of Political Economy*, 106, 724–767.
- CARD, D. AND D. HYSLOP (1997): "Does Inflation Grease the Wheels of the Labor Market," *Reducing Inflation: Motivation and Strategy*.
- CHODOROW-REICH, G. AND J. WIELAND (2018): "Secular Labor Reallocation and Business Cycles," Working Paper.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): "Nominal Rigidities and the

- Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 115, 1–45.
- COIBION, O., Y. GORODNICHENKO, AND J. WIELAND (2012): "The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound?" *Review of Economic Studies*, 79, 1371–1406.
- COLE, H. AND R. ROGERSON (1999): "Can the Mortensen-Pissarides Matching Model Match the Business-Cycle Facts," *International Economic Review*, 40, 933–959.
- DALY, M. AND B. HOBIIJN (2014): "Downward Nominal Wage Rigidities Bend the Phillips Curve," *Journal of Money, Credit and Banking*, 46, 51–93.
- DALY, M. C., B. HOBIIJN, AND B. LUCKING (2012): "Why Has Wage Growth Stayed Strong?" Federal Reserve Bank of San Francisco Economic Letter 2012-11.
- DALY, M. C., B. HOBIIJN, AND T. S. WILES (2011): "Dissecting Aggregate Real Wage Fluctuations: Individual Wage Growth and the Composition Effect," Federal Reserve Bank of San Francisco Working Paper 2011-23.
- DE LONG, J. B. AND L. H. SUMMERS (1988): "How does macroeconomic policy affect output?" *Brookings Papers on Economic Activity*, 1988, 433–480.
- ELSBY, M. (2009): "Evaluating the economic significance of downward nominal wage rigidity," *Journal of Monetary Economics*, 56, 154–169.
- ELSBY, M., R. MICHAELS, AND G. SOLON (2009): "The Ins and Outs of Cyclical Unemployment," *American Economic Journal: Macroeconomics*, 1, 84–110.
- FATÁS, A. AND I. MIHOV (2015): "Recoveries," Working Paper, INSEAD.
- FERNANDEZ-VILLAYERDE, J., F. MANDELMAN, Y. YU, AND F. ZANETTI (2019): "Search Complementarities, Aggregate Fluctuations, and Fiscal Policy," NBER Working Paper No. 26210.
- FERRARO, D. (2017): "The Asymmetric Cyclical Behavior of the U.S. Labor Market," Working Paper, Arizona State University.
- FRIEDMAN, M. (1964): "Monetary Studies of the National Bureau," in *The National Bureau Enters Its 45th Year*, New York, NY: National Bureau of Economic Research, 7–25, 44th Annual Report, available at <http://www.nber.org/nberhistory/annualreports.html>.
- (1993): "The "Plucking Model" of Business Fluctuations Revisited," *Economic Inquiry*, 31, 171–177.
- FUJITA, S. AND G. RAMEY (2007): "Job Matching and Propagation," *Journal of Economic Dynamics & Control*, 31, 3671–3698.
- GALI, J., M. GERTLER, AND D. LOPEZ-SALIDO (2007): "Markups, Gaps, and the Welfare Costs of Business Fluctuations," *Review of Economics and Statistics*, 89, 44–59.
- GALINDEV, R. AND D. LKHAGVASUREN (2010): "Discretization of highly persistent correlated AR(1) shocks," *Journal of Economic Dynamics and Control*, 34, 1260–1276.
- GERTLER, M., C. HUCKFELDT, AND A. TRIGARI (2016): "Unemployment Fluctuations, Match Quality, and the Wage Cyclicity of New Hires," NBER Working Paper No. 22341.
- GERTLER, M. AND A. TRIGARI (2009): "Unemployment Fluctuations with Staggered Nash Wage Bargaining," *Journal of Political Economy*, 117, 38–86.
- GOTTSCHALK, P. (2005): "Downward Nominal-Wage Flexibility: Real or Measurement Error?" *Review of Economics and Statistics*, 87, 556–568.

- GRIGSBY, J., E. HURST, AND A. YILDIRMAZ (2018): "Aggregate Nominal Wage Adjustment: New Evidence from Administrative Payroll Data," Working Paper, University of Chicago.
- HAIRAULT, O., F. LANGOT, AND S. OSOTIMEHIN (2010): "The (un)importance of unemployment fluctuations for the welfare cost of business cycles," *Review of Economic Dynamics*, 13, 759–779.
- HALL, R. (2005): "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review*, 95, 50–65.
- HALL, R. E. (1995): "Lost Jobs," *Brookings Papers on Economic Activity*, 1995, 221–273.
- HALL, R. E. AND M. KUDLYAK (2021): "Why Has the US Economy Recovered So Consistently from Every Recession in the Past 70 Years?" *NBER Macroeconomics Annual*, 36, forthcoming.
- HALL, R. E. AND P. R. MILGROM (2008): "The Limited Influence of Unemployment on the Wage Bargain," *American Economic Review*, 98, 1653–1674.
- HARTLEY, J. (2021): "Friedman's plucking model: New international evidence from Maddison Project data," *Economics Letters*, 199.
- HAZELL, J. AND B. TASKA (2018): "Posted Wage Rigidity," Working Paper, MIT.
- IMROHOROGLU, A. (1989): "Cost of Business Cycles with Indivisibilities and Liquidity Constraints," *Journal of Political Economy*, 97, 1364–1383.
- JACKSON, M. AND P. TEBALDI (2017): "A Forest Fire Theory of the Duration of a Boom and the Size of a Subsequent Bust," Working Paper.
- JAROSCH, G. (2021): "Searching for Job Security and the Consequences of Job Loss," Working Paper, Princeton University.
- JORGENSON, D., M. HO, AND J. SAMUELS (2012): "A Prototype Industry-Level Production Account for the United States, 1947-2010," Second World KLEMS conference, Harvard University.
- JUNG, P. AND K. KUESTER (2011): "The (un)importance of unemployment fluctuations for the welfare cost of business cycles," *Journal of Economic Dynamics & Control*, 35, 1744–1768.
- JUNG, P. AND M. KUHN (2019): "Earings Losses and Labor Mobility over the Life Cycle," *Journal of the European Economic Association*, 17, 678–724.
- KAHN, S. (1997): "Evidence of Nominal Wage Stickiness from Microdata," *American Economic Review*, 87, 993–1008.
- KIM, C.-J. AND C. R. NELSON (1999): "Friedman's Plucking Model of Business Fluctuations: Tests and Estimates of Permanent and Transitory Components," *Journal of Money, Credit and Banking*, 31, 317–334.
- KIM, J. AND F. RUGE-MURCIA (2009): "How much inflation is necessary to grease the wheels?" *Journal of Monetary Economics*, 56, 365–377.
- KROLIKOWSKI, P. (2017): "Job Ladders and Earnings of Displaced Workers," *American Economic Journal: Macroeconomics*, 9, 1–31.
- KRUSELL, P., T. MUKOYAMA, A. ŞAHIN, AND A. A. SMITH, JR. (2009): "Revisiting the Welfare Effects of Eliminating Business Cycles," *Review of Economic Dynamics*, 12, 393–404.
- KURMANN, A. AND E. MCENTARFER (2017): "Downward Wage Rigidity in the United States: New Evidence from Administrative Data," Working Paper.
- LEPETIT, A. (2018): "Asymmetric Unemployment Fluctuations and Monetary Policy Trade-offs,"

Working Paper.

- LUCAS, R. (1987): *Models of Business Cycles*, Oxford:Blackwell.
- (2003): “Macroeconomic Priorities,” *American Economic Association*, 93, 1–14.
- MCKAY, A. AND R. REIS (2008): “The Brevity and Violence of Contractions and Expansions,” *Journal of Monetary Economics*, 55, 738–751.
- MCLAUGHLIN, K. (1994): “Rigid Wages?” *Journal of Monetary Economics*, 34, 383–414.
- MICHAILLAT, P. (2012): “Do Matching Frictions Explain Unemployment? Not in Bad Times,” *American Economic Review*, 102, 1721–1750.
- MITCHELL, W. C. (1927): *Business Cycles: The Problem and Its Setting*, National Bureau of Economic Research.
- NAKAMURA, E., J. STEINSSON, P. SUN, AND D. VILLAR (2018): “The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation,” *Quarterly Journal of Economics*, 133, 1933–1980.
- NEFTÇI, S. (1984): “Are Economic Time Series Asymmetric over the Business Cycle?” *Journal of Political Economy*, 92, 307–328.
- OBSTFELD, M. (1994): “Evaluating risky consumption paths: The role of intertemporal substitutability,” *European Economic Review*, 38, 1471–1486.
- OLIVEI, G. AND S. TENREYRO (2010): “Wage-setting patterns and monetary policy: International evidence,” *Journal of Monetary Economics*, 57, 785–802.
- PETRONGOLO, B. AND C. PISSARIDES (2001): “Looking into the Black Box: A survey of the Matching Function,” *Journal of Economic Literature*, XXXIX, 390–431.
- PETROSKY-NADEAU, N. AND L. ZHANG (2017): “Solving the Diamond-Mortensen-Pissarides model accurately,” *Quantitative Economics*, 8, 611–650.
- PETROSKY-NADEAU, N., L. ZHANG, AND L. KUEHN (2018): “Endogenous Disasters,” *American Economic Review*, 108, 2212–2245.
- PRIES, M. (2004): “Persistence of Employment Fluctuations =: A Model of Recurring Job Loss,” *The Review of Economic Studies*, 71, 193–215.
- RAMEY, G. AND V. RAMEY (1991): “Technology Commitment and the Cost of Economic Fluctuations,” NBER Working Paper No. 3755.
- ROMER, C. D. AND D. H. ROMER (2004): “A New Measure of Monetary Shocks: Derivation and Implications,” *American Economic Review*, 94, 1055–1084.
- ROUWENHORST, K. G. (1995): “Asset Pricing Implications of Equilibrium Business Cycle Models,” in *Frontiers of Business Cycle Research*, ed. by T. F. Cooley, Princeton University Press, chap. 10, 294–330.
- SCHMITT-GROHE, S. AND M. URIBE (2016): “Downward Nominal Wage Rigidity, Currency Pegs, and Involuntary Unemployment,” *Journal of Political Economy*, 124, 1466–1514.
- SHIMER, R. (2005): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *American Economic Review*, 95, 25–49.
- (2010): *Labor Market and Business Cycles*, Princeton University Press.
- (2012): “Reassessing the ins and outs of unemployment,” *Review of Economic Dynamics*, 15, 127–148.

- SICHEL, D. (1993): "Business Cycle Asymmetry: A Deeper Look," *Economic Inquiry*, 31, 224–236.
- SILVA, I. J. AND M. TOLEDO (2009): "Labor Turnover Costs and the Cyclical Behavior of Vacancies and Unemployment," *Macroeconomic Dynamics*, 13, 76–96.
- SINCLAIR, T. (2010): "Asymmetry in the Business Cycle: Friedman's Plucking Model with Correlated Innovations," *Studies in Nonlinear Dynamics and Econometrics*, 14, 1–31.
- SMETS, F. AND R. WOUTERS (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97, 585–606.
- TALLARINI, T. (2000): "Risk-sensitive real business cycles," *Journal of Monetary Economics*, 45, 507–532.
- TASCI, M. AND N. ZEVANOVE (2019): "Do Longer Expansions Lead to More Severe Recessions?" Federal Reserve Bank of Cleveland Economic Commentary.
- TAUCHEN, G. (1986): "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," *Economic Letters*, 20, 177–181.
- TOBIN, J. (1972): "Inflation and Unemployment," *American Economic Review*, 62, 1–18.
- WOODFORD, M. (2003): *Interest and Prices, Foundations of a Theory of Monetary Policy*, Princeton University Press.