Journal of Monetary Economics 152 (2025) 103766

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

A plucking model of business cycles[☆]

Stéphane Dupraz[a], Emi Nakamura[b], Jón Steinsson[b,*]

[a] Banque de France, France
[b] University of California, Berkeley, United States of America

ARTICLE INFO

Dataset link: https://eml.berkeley.edu/~jsteinsson/papers/pluckingReplication.zip

JEL classification:
E24
E30
E52

Keywords:
Downward Nominal Rigidity
Stabilization Policy
Labor Search

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 a decrease does not predict the amplitude of the following increase. In addition, business cycles last seven years on average and unemployment rises much faster during recessions than it falls during expansions. We augment a standard labor search model with downward nominal wage rigidity and show how it can fit the plucking property.

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 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, United States (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, United States 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.

* Corresponding author.
E-mail address: jsteinsson@berkeley.edu (J. Steinsson).

[1] The 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)

https://doi.org/10.1016/j.jmoneco.2025.103766
Received 7 August 2022; Received in revised form 4 March 2025; Accepted 5 March 2025
Available online 14 March 2025
0304-3932/

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.

To match the facts about plucking that we document in the data, we introduce downward nominal wage rigidity into a tractable version of the Diamond–Mortenson–Pissarides model with endogenous separations analyzed by Fujita and Ramey (2012). We show that the model can quantitatively match the plucking dynamics of US unemployment. 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.

The plucking dynamics of our model imply 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 3.1%.

An alternative interpretation of the "plucking" property of unemployment is that it derives from exogenous shocks to the determinants of the unemployment rate that themselves have this property. We have chosen to model the exogenous shocks in our model as symmetric 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).

While our baseline model fits the plucking property, it fails to fit some other salient features of unemployment dynamics. Empirically, business cycles last around 7 years from peak to peak, and the unemployment rate rises much more rapidly during downturns than it falls during expansions. In principle, search models such as the one we introduce provide an intuitive mechanism for slow recoveries: firms can shed workers rapidly, but it takes time due to search and matching frictions to expand employment. In practice, however, the large number of workers who flow between employment and unemployment in every month implies that the DMP model does not have appreciable internal propagation (see, e.g. Cole and Rogerson, 1999). As a result, unemployment falls quickly once negative shocks dissipate. Our baseline model inherits this feature of standard search models, generating short business cycles.

We explore the potential for additional non-standard features to explain this speed asymmetry, including insecure short-term jobs and a hump-shaped driving process for productivity shocks (an AR(2) process). Intuitively, the presence of insecure short-term jobs implies that 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). While these simple deviations from our baseline model cannot match all of the empirical facts we present, we argue they provide suggestive evidence for what kind of models will be able to unify these empirical patterns. We leave a full theoretical solution to this problem to future work.

Our work is related to several strands of existing literature. Petrosky-Nadeau et al. (2018) show how a DMP model features asymmetries that can generate business-cycles disasters – large drops in production – despite symmetric shocks. This force generates some plucking. But when we consider business cycles of the size we have experience in our sample period, it generates much less plucking than we document, motivating the role of DNWR. Kim and Nelson (1999) and Sinclair (2010) are two of the very few modern attempts to assess the specific asymmetry emphasized by Friedman. Caballero and Hammour (1998), Bordo and Haubrich (2012), and Fatás and Mihov (2013) explore related ideas. Ferraro (2018) provides an alternative explanation for the speed asymmetry in the unemployment rate.

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 how this model can match the plucking property, and demonstrates that stabilizing fluctuations can reduce the average level of unemployment. Section 5 considers whether extensions to the baseline model can help match the longer duration of unemployment cycles in the data, and/or help better match the asymmetry in the duration of contractions and expansions. Section 6 concludes.

2. The plucking property and the dynamics of unemployment cycles

We start by demonstrating Friedman's plucking property for post-WWII US unemployment data; namely, 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. We also demonstrate a speed asymmetry in unemployment: the unemployment rate rises more quickly than it falls, implying that the duration of recoveries is typically much longer than the duration of recessions, as emphasized by Neftçi (1984) among others. Finally, we review the empirical evidence on wage rigidity, and show that the incidence of wage rigidity is countercyclical.

To implement our empirical analysis, we define business cycle peaks and troughs such that they line up exactly with peaks and troughs of the unemployment rate. This yields business cycle dates that are very similar to but not identical to those identified by the NBER Business Cycle Dating Committee. Fig. 1 plots the unemployment rate over our sample period – which runs from January

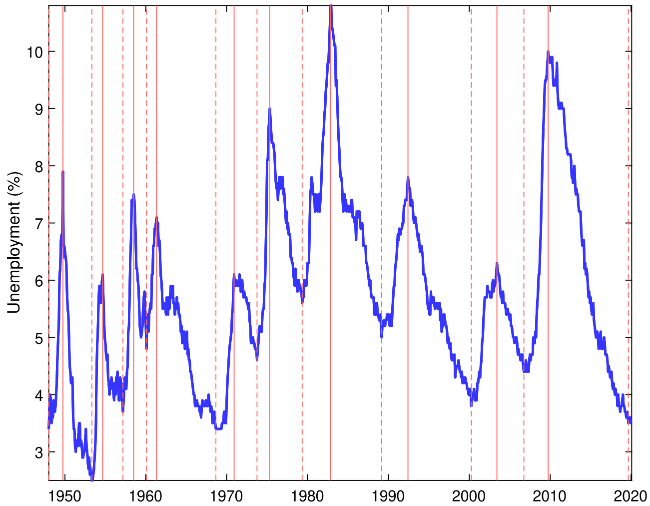

Fig. 1. Peaks and Troughs in the Unemployment Rate.
Note: The unemployment rate is plotted in blue. Business cycle peaks are denoted by dashed red vertical lines, while business cycle troughs are denoted by solid red vertical lines. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

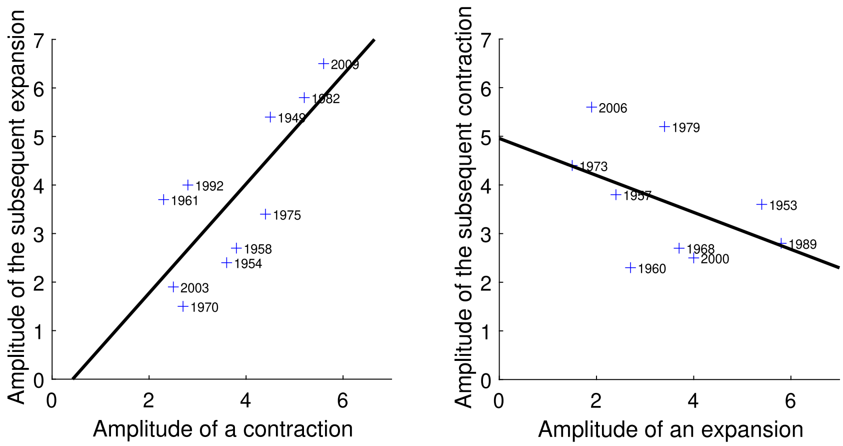

Fig. 2. The Plucking Property of the Unemployment Rate.
Note: The points in the left panel are labeled with the year the contraction in question ended and expansion in question began. The points in the right panel are labeled with the year the expansion in question ended and contraction in question began. OLS regression lines are plotted in each panel.

1948 to February 2020 – with vertical lines indicating the times that we identify as business cycle peaks and troughs. For details on our algorithm, see Appendix A.[2]

2.1. The plucking property

Fig. 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 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.

[2] We identify ten peaks and ten troughs. To these we add a peak at the beginning of our sample. One might worry 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 end our sample at the onset of the Covid-19 recession, another clear peak in the data.

Table 1
Plucking property of unemployment and speed asymmetries.

	β	R^2
Subsequent expansion on contraction	1.12	0.59
	(0.33)	
Subsequent contraction on expansion	−0.38	0.22
	(0.27)	
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: 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. The speed of expansions and contractions in the third and fourth rows is measured in percentage points of unemployment per year.

Table 1 reports the estimated coefficient from an OLS regression of the size of the subsequent expansion on the size of a contraction. 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 1.1 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 Fig. 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 statistically significant relationship in this case. The size of an expansion does not forecast the size of the next contraction. One cannot reject that, in Friedman's language, each contractionary pluck that the economy experiences is independent of what happened before. Table 1 reports the estimated coefficient from a linear regression of the size of the subsequent contraction on the size of an expansion. The relationship is actually slightly negative, but is far from statistically significant. Moreover, the R^2 of the regression is only 0.22.[3]

2.2. The speed asymmetry

Unemployment rises more quickly during contractions than it falls during expansions, a point made quantitatively in early work by Neftçi (1984).[4] Table 1 reports the average speed of expansions and contractions to illustrate this asymmetry. We measure the change in unemployment (in percentage points) over the spell and the length of time the spell lasts for. The speed for an 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.

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 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 Fig. 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.

---

3 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.

4 Sichel (1993) refers to this asymmetry as the "steepness" asymmetry, while McKay and Reis (2008) refer to 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".

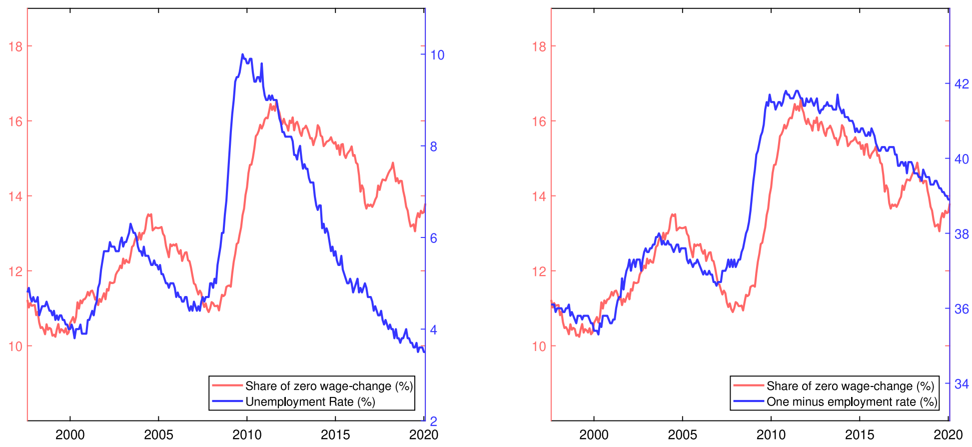

Fig. 3. San Francisco Fed Wage Rigidity Meter.
Note: The figure plots the share of wage freezes of all job-stayers (paid by the hour or not) with respect to the wage one year prior, with no correction for measurement errors. This series is constructed by the Federal Reserve Bank of San Francisco using data used from the Current Population Survey.

2.3. Wage rigidity

A large literature has presented microeconomic evidence of downward nominal wage rigidity in US data.[5] Olivei and Tenreyro (2010) present evidence that such wage rigidity matters for real outcomes. A large theoretical literature has studied the implications of downward nominal wage rigidity for business cycles.[6]

An intuitive fact about the data is that wage rigidity has been countercyclical in recent US business cycles—i.e., wage freezes tend to occur during recessions. Fig. 3 plots the fraction of "Wage Freezes" from the San Francisco Fed's Wage Rigidity Meter for the period 1997–2019 along with the unemployment rate (Left Panel) and the non-employment rate for the working age population (Right Panel).[7] The correlation is striking. The fraction of wage freezes rises rapidly in each of the three recessions that occur in this sample period. Before 1997 the data are less complete, and the correlation is weaker because of the confounding effects of changes in trend inflation.

3. A plucking model with downward nominal wage rigidity

Next, we present a model designed to fit the plucking property discussed in Section 2. This model augments the workhorse DMP model with endogenous separation developed by Fujita and Ramey (2012) with downward nominal wage rigidity. We begin by describing the model with Nash Bargaining, then go on to present a version of the same model with downward nominal wage rigidity.

3.1. Search and matching with endogenous separation

The model consists of an infinite mass of atomistic firms, and a mass of workers with inelastic labor supply normalized to one. At the beginning of a period, workers are either already matched with a firm, or looking for a job.[8] If a firm and a job-seeker match in period t, we assume the worker starts working right away.[9] To match with a new worker, a firm must post a vacancy, at a cost c

---

[5] See, in particular, McLaughlin (1994), Kahn (1997), Card and Hyslop (1997), Bewley (1999), Altonji and Devereux (2000), Kurmann and McEntarfer (2017), Hazell and Taska (2018), Grigsby et al. (2021). Correcting for measurement error might reveal an even higher prevalence of wage rigidity. Early work using data from the Panel Study of Income Dynamics (PSID) and the Current Population Survey (CPS) includes McLaughlin (1994), 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 et al. (2014) report similar findings based on their adjustments for measurement error in the Survey of Income and Program Participation.

[6] Prominent contributions include, e.g., Akerlof et al. (1996), Kim and Ruge-Murcia (2009), Benigno and Ricci (2011), Abbritti and Fahr (2013), Schmitt-Grohe and Uribe (2016), Chodorow-Reich and Wieland (2020). 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 et al. (2020).

[7] The Wage Freeze measure reports the fraction of job-stayers whose wages are unchanged versus one year prior (Daly et al., 2011, 2012; Daly and Hobijn, 2014). See https://www.frbsf.org/economic-research/indicators-data/nominal-wage-rigidity/.

[8] Only unemployed workers look for jobs. Fujita and Ramey (2012) also develop a version of the model with on-the-job search. We consider their model with endogenous separation but no on-the-job search.

[9] This timing of the labor market follows, e.g., Blanchard and Gali (2010), but differs from Fujita and Ramey (2012). It is not important for our results, but avoids the unpalatable assumption that the number of job-seekers at the beginning of period t (which includes at least all the workers who separated from their jobs at the end of $t-1$) is the same as the number of unemployed workers u_t in period t, even when search frictions shrink to zero and the job finding rate is 1.

per period in which the vacancy remains open. An open vacancy fills with probability q_t, which is taken as exogenous by the firm. In the aggregate, the probability q_t is determined by a matching function $q(\theta_t)$, where $\theta_t = V_t/U_{0,t}$ denotes labor market tightness. Labor market tightness is the ratio of the number of vacancies posted V_t to the number of job-seekers $U_{0,t}$ at the beginning of the period. The matching function also determines the probability for a job-seeker of finding a job. This is equal to the ratio of matches $q(\theta_t)V_t$ to job-seekers $U_{0,t}$, $f(\theta_t) \equiv q(\theta_t)V_t/U_{0,t} = \theta_t q(\theta_t)$.

When a worker and a firm are matched, they produce output $A_t x_t$, where A_t and x_t are aggregate and match-specific productivity factors. We assume that both follow AR(1) exogenous processes in logs,

$$\log(A_t) = \rho^a \log(A_{t-1}) + \varepsilon_t^a, \tag{1}$$

$$\log(x_t) = \rho^x \log(x_{t-1}) + \varepsilon_t^x, \tag{2}$$

where ε_t^a and ε_t^x are Gaussian shocks with standard deviations σ_ε^a and σ_ε^x. All new matches start at the same match productivity level x^{hire}, which we take to be average match productivity $x = 1$.[10]

An ongoing match continues into the next period unless it is exogenously terminated – which occurs with probability δ – or endogenously terminated. Let $J_t(x_t)$ be the value to a firm of an ongoing match with match-specific productivity x_t.[11] The firm can terminate the match if it yields a negative value. This implies that $J_t(x_t)$ is given by:

$$J_t(x_t) = \max\{J_t^c(x_t), 0\} \tag{3}$$

where J_t^c is the value of the match to the firm if it is continued, which solves the recursion

$$J_t^c(x_t) = x_t A_t - w_t(x_t) + \beta(1-\delta)E_t\left(J_{t+1}(x_{t+1})\right), \tag{4}$$

where $w_t(x_t)$ is the real wage paid in the match at period t.

Let $W_t(x_t)$ be the value to a worker of a match with match-specific productivity x_t, and U_t be the value to a worker of being unemployed. Since the worker can terminate the match if it yields it less than being unemployed, the value of being employed in a match at x_t is given by

$$W_t(x_t) = \max\{W_t^c(x_t), U_t\} \tag{5}$$

where W_t^c is the value if the match is continued, which solves the recursion

$$W_t^c(x_t) = w_t(x_t) - \zeta + \beta E_t\left((1-\delta)W_{t+1}(x_{t+1}) + \delta(1-f_{t+1})U_{t+1} + \delta f_{t+1} W_{t+1}(x^{hire})\right), \tag{6}$$

where ζ is the disutility cost of working relative to being unemployed. The value of being unemployed solves the recursion

$$U_t = b + \beta E_t\left((1-f_{t+1})U_{t+1} + f_{t+1}W_{t+1}(x^{hire})\right), \tag{7}$$

where b is unemployment benefits. Subtracting (7) from (6), the value of being employed relative to being unemployed, $V_t = W_t - U_t$, solves the recursion

$$V_t^c(x_t) = w_t(x_t) - z + \beta(1-\delta)E_t\left(V_{t+1}(x_{t+1}) - f_{t+1}V_{t+1}(x^{hire})\right), \tag{8}$$

where

$$V_t(x) = \max\{V_t^c(x), 0\} \tag{9}$$

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

Assuming free entry, the cost of posting a new vacancy must in equilibrium be equal to the benefit to the firm of posting a new vacancy. This equates the value of a new hire to the firm to the expected cost of a new hire:

$$J_t(x^{hire}) = \frac{c}{q_t}. \tag{10}$$

The model is closed with an assumption on wage-setting. Fujita and Ramey (2012) assume that wages are set according to Nash-bargaining, as is standard in search models. Define the total value of a continuing match to be $S_t^c(x_t) = J_t^c(x_t) + V_t^c(x_t)$. Under Nash-bargaining, $V_t^{c,Nash} = \gamma S_t^{c,Nash}$, where $\gamma \in [0,1]$ is the bargaining power of workers. Combining Eqs. (4) and (8) and the Nash-bargaining assumption gives

$$J_t^{c,Nash}(x_t) = (1-\gamma)(A_t x_t - z) + \beta(1-\delta)E_t\left(J_{t+1}^{Nash}(x_{t+1}) - f_{t+1}^{Nash}\gamma J_{t+1}^{Nash}(x^{hire})\right). \tag{11}$$

Eqs. (10) and (11) allow us to solve for q_t and $J_t^{c,Nash}(x_t)$, as detailed in Appendix B.

[10] Because Fujita and Ramey (2012) also consider a model with on-the-job search, they assume that new matches start at the highest productivity level to make sure all job offers are accepted. Since we do not consider on-the-job search, job offers are accepted even if the productivity in new matches is not the highest productivity level. Our process for match-specific shocks is also a more standard AR(1) process than the memoryless Poisson process with infrequent large shocks assumed in Fujita and Ramey (2012).

[11] We make explicit the dependence of J_t and other value functions only in the idiosyncratic state x_t. The dependence of J_t on the aggregate state is kept implicit in the t time-index of J_t.

Combining (4) and (11) allows us to recover the Nash wage as

$$w_t^{Nash}(x_t) = \left(\gamma A_t x_t + (1-\gamma)z\right) + \beta(1-\delta)E_t\left(\gamma f_{t+1}^{Nash} J_{t+1}^{Nash}(x^{hire})\right). \tag{12}$$

3.2. Downward nominal wage rigidity

We extend the Fujita–Ramey model to allow for downward nominal wage rigidity (DNWR).[12] The presence of match heterogeneity implies that wages differ for new and ongoing matches. We assume DNWR for both groups. For an ongoing match at time t, the nominal wage is either the flexible nominal wage – which we assume to be the wage that would obtain if the firm set wages by Nash bargaining in this and future periods – or, if this requires the nominal wage to fall, the nominal wage remains unchanged from the previous period, i.e.:

$$w_t(x_t, w_{t-1}) = \max\left\{ w_t^{Nash}(x_t), \frac{w_{t-1}}{\Pi_t} \right\}, \tag{13}$$

where Π_t is the inflation rate.

For new matches, we assume that the hiring wage at time t is either the flexible nominal wage – which we again assume to be the wage that would obtain under Nash bargaining – or, if this requires the nominal hiring wage to fall, the nominal hiring wage in the previous period, i.e.:

$$w_t^{new}(x^{hire}, w_{t-1}^{new}) = \max\left\{ w_t^{Nash}(x^{hire}), \frac{w_{t-1}^{new}}{\Pi_t} \right\}. \tag{14}$$

To generate fluctuations in firms' hiring decisions, wage rigidity in new matches is particularly important, as emphasized by Pissarides (2009). Assumption (14) is therefore essential to introducing DNWR in a search and matching model with match heterogeneity. Hazell and Taska (2018) document such DNWR in new hires' wages.[13]

In both Eqs. (13) and (14), inflation relaxes the constraint on downward *real* wage adjustments: it greases the wheels of the labor market. 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 $\overline{\Pi}$. Eqs. (10) and (11) still hold under DNWR, except that the firm's value function now depends on lagged wages. We first solve the model under Nash-bargaining to recover the Nash wage (12), then use Eqs. (10) and (11) to solve for J_t^c and q_t, as detailed in Appendix B.

3.3. Worker flow accounting

Let s_t be the destruction rate, defined as the fraction of matches that get destroyed at the beginning of period t. Because matches can be endogenously terminated, the destruction rate s_t depends on the cross-sectional distribution of employment across the state of matches. Under Nash-bargaining, the state of a match reduces to match productivity x_t. Fujita and Ramey (2012) show how to keep track of the distribution of employment across matches to calculate the destruction rate in this case. Under DNWR, the state of a match includes both match productivity x_t and the lagged wage w_{t-1}. Appendix B shows how to keep track of the joint distribution of employment across the joint state (x_t, w_{t-1}) and calculate the destruction rate under DNWR.

Workers' job-finding rate f_t is a direct function of the vacancy-filling rate q_t through the matching function and can therefore be easily recovered from q_t. Since we assume that job-seekers who match with a firm at the beginning of period t start working at t, workers who separate from a firm between $t-1$ and t and join the pool of job-seekers may find a new job at t without spending time unemployed, instead transitioning directly from job to job. The exit rate from employment to unemployment is therefore distinct from the destruction rate s_t and equal to

$$\bar{s}_t = (1 - f_t)s_t. \tag{15}$$

The law of motion of the unemployment rate u_t is

$$u_t = (1 - f_t)u_{t-1} + \bar{s}_t(1 - u_{t-1}). \tag{16}$$

[12] We assume DNWR by directly specifying the wage-rule. Other work directly specifying a wage rule in a search and matching framework includes Blanchard and Gali (2010), Shimer (2010), and Michaillat (2012). This approach has the advantage that we can more easily investigate what features the wage process needs to generate realistic unemployment dynamics.

[13] Our model does not feature preemptive wage moderation of the type that is present in wage setting models (e.g., Kim and Ruge-Murcia, 2009; Elsby, 2009; Benigno and Ricci, 2011). 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.

Table 2
Calibration.

β	$0.96^{1/12}$	
η	0.5	
c	0.30	
z	0.95	
δ	1.9%	
ρ^a	0.98	
ρ^x	0.98	
	Nash	DNWR
$\overline{\Pi}$	–	$1.02^{1/12}$
γ	0.57	0.43
σ_ε^a	st. $\sigma^a = 1.6\%$	st. $\sigma^a = 1.5\%$
σ_ε^x	st. $\sigma^x = 2.1\%$	st. $\sigma^x = 1.5\%$

Note: The abbreviation "st." stands for "such that".

3.4. Calibration

Table 2 provides a summary of our calibration. 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 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. We calibrate the flow value of unemployment following Hagedorn and Manovskii (2008) to $z = 0.95$, so that the model generates significant fluctuations in the job-finding rate under Nash bargaining, which we will consider as a benchmark.[14] This calibration also generates the realistic prediction that firms' surplus is increasing in productivity under DNWR, as explained in Appendix H. We calibrate the exogenous destruction rate δ to match the average monthly share of quits in the non-farm sector in JOLTS between January 2000 and February 2020 of 1.9%. The JOLTS data likely overstates the number of exogenous separations since many quits are job-to-job flows. We therefore provide robustness results under a lower calibration of $\delta = 1\%$ in Appendix F.

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.1 for further discussion of this point.) We normalize μ to 1. We calibrate c so that the cost of hiring a worker $\frac{c}{q} = J$ is 10% of monthly wages in a steady state with $u = 5.7\%$ and $\bar{s} = 2\%$, in line with what Silva and Toledo (2009) report based on the Employer Opportunity Pilot Project survey in the US. This yields $c = 0.30$.

We set the auto-regressive root of the aggregate productivity process ρ^a to 0.98 following Shimer (2010). We set the auto-regressive root of the match-specific productivity process ρ^x to 0.98, following Foster et al. (2008)'s estimates of the persistence of plants' TFP (0.8 on an annual basis). When we consider the model under downward nominal wage rigidity, we set inflation to 2% per year. Inflation is immaterial in the version of the model without DNWR.

This leaves γ, σ_ε^a and σ_ε^x. We pick them to match the average level of the unemployment rate (5.7% in the data), the standard deviation of the unemployment rate (1.6% in the data), and the average of the rate of exit from employment $\bar{s}$ (2% as reported by Fujita and Ramey (2006, 2012)). 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. These choices yield $\gamma = 0.57$, $\sigma_a = 1.6\%$ and $\sigma_x = 2.1\%$ under Nash bargaining, and $\gamma = 0.43$, $\sigma_a = 1.5\%$ and $\sigma_x = 1.5\%$ under DNWR.

4. Quantitative analysis of the baseline model

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 discusses the algorithm we use in detail. We simulate 5000 samples of 866 periods (the length of our sample of real-world data) and calculate the statistics reported in the empirical Table 1 in each of these simulated samples. We then 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.

Table 3 reports the results. The top panel presents results on the plucking property. We first consider the model with flexible wages (Nash bargaining). This version of the model generates non-trivial plucking, but substantially less than in the data. The regression coefficient for the size of expansions on the previous contraction is 0.35 versus −0.04 for contractions on the previous expansion. These coefficients are 1.12 and −0.38 in the data. Next, we present results for the model with DNWR. This case produces substantially more plucking that the model with flexible wages. The regression coefficient for the size of expansions on the previous contraction is 0.64 versus −0.05 for contractions on the previous expansion. Appendix F shows that the plucking property is equally present – if anything slightly more so – when the exogenous separation rate is calibrated to a lower value of $\delta = 1\%$.

[14] Fujita and Ramey (2012) show that their model under Nash bargaining can generate fluctuations in the unemployment rate away from a high calibration of z by making unemployment fluctuate through fluctuations in separations. As they show, however, away from a high calibration of z, the model cannot generate fluctuations in the job-finding rate.

Table 3
Simulation results: Plucking property, speed, and duration.

	Data	AR(1)		AR(2)		AR(2) + Job Ladder	
		Nash	DNWR	Nash	DNWR	Nash	DNWR
Subsequent expansion on contraction, β	1.12	0.35	0.64	0.40	0.64	0.39	0.47
		(0.43)	(0.27)	(0.26)	(0.15)	(0.27)	(0.19)
Subsequent contraction on expansion, β	−0.38	−0.04	−0.05	0.02	−0.02	0.03	0.00
		(0.44)	(0.31)	(0.28)	(0.15)	(0.29)	(0.19)
Subsequent expansion on contraction, R^2	0.59	0.16	0.42	0.17	0.42	0.16	0.22
		(0.24)	(0.26)	(0.18)	(0.17)	(0.18)	(0.15)
Subsequent contraction on expansion, R^2	0.22	0.07	0.04	0.04	0.01	0.04	0.02
		(0.17)	(0.11)	(0.10)	(0.03)	(0.10)	(0.05)
Speed of expansions (pp/year)	0.87	1.56	2.53	0.73	2.07	0.61	1.19
		(0.64)	(0.73)	(0.18)	(0.29)	(0.11)	(0.26)
Speed of contractions (pp/year)	1.89	1.56	4.56	0.70	3.66	0.60	1.56
		(0.78)	(1.51)	(0.20)	(0.83)	(0.11)	(0.28)
Duration of expansions (months)	59.1	35.8	27.4	84.6	37.6	91.9	58.6
		(12.1)	(6.7)	(16.6)	(3.7)	(16.7)	(7.2)
Duration of contractions (months)	26.9	36.1	18.3	86.5	23.5	92.5	40.8
		(12.1)	(5.5)	(16.9)	(2.7)	(16.4)	(6.3)

Note: The table compares data from the model under AR(1) shocks, from the model under AR(2) shocks, and from the model under AR(2) shocks and a job ladder, in each case both with Nash bargaining and with downward nominal wage rigidity (DNWR). 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 next two rows report the 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. Expansions and contractions of more than 6.5 percentage points are excluded from the samples. Appendix E provides results including all expansions and contractions. For the version of the model without a job ladder, 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). For the model under AR(2) shocks and a job ladder, results are shown for 1000 samples of 5*866 periods each, to avoid samples with no or few cycles of less than 6.5 percentage points. The standard error reported in parentheses is the standard deviation of the estimates across the 5000 (or 1000) samples.

There is a large amount of noise in the relationship between the amplitude of expansions and contractions in the model. This is particularly the case in the model with flexible wages. In this model, the R^2 of the size of expansions on the previous contraction is only 0.16, compared to 0.59 in the data. The model with DNWR does substantially better on this metric with an R^2 of 0.42.

4.1. Mechanisms for plucking

Two features of the model contribute to the plucking property. First, non-linearity in the worker-flow Eqs. (15)–(16) generates plucking. Fig. 4 plots the steady-state relationship between the job finding rate and unemployment implied by the worker-flow relationship (15)–(16) taken in steady-state.[15] The relationship is convex because it gets harder and harder to lower the unemployment rate the lower it gets. Earlier work has emphasized this non-linearity (Petrosky-Nadeau and Zhang, 2017; Petrosky-Nadeau et al., 2018).[16]

The reason why this non-linearity generates only modest amounts of plucking is that over the empirically relevant range of unemployment rates – between 2.5% and 10.8% for the sample period 1948–2019 – the degree of non-linearity is modest. This is illustrated in Fig. 4. In Table 3, we make sure not the overstate the strength of this non-linearity by presenting results for a "top-truncated" sample that only includes expansions and contractions that are less than 6.5 percentage points in size—i.e., the size of the 2009–2019 expansion which is the largest one in our sample.[17]

The second feature of our model that contributes to plucking is DNWR. When wages are downward rigid, negative shocks typically result in higher unemployment while positive shocks yield wage increases. As a result, unemployment sometimes rises far above its steady state, but rarely falls below. Table 3 shows that this feature can generate a substantial amount of plucking.

But even the model with DNWR produces less plucking than the data. The top panel of Fig. 5 helps us understand why. This panel plots a simulated path of unemployment in our model with DNWR. A prominent feature of this simulated series is that expansions are often interrupted by a new contraction before unemployment has had enough time to fully recover. This can introduce a great deal of noise in our measure of plucking (as shown in Figure C.1 of Appendix C). Some expansions will be cut short. Others will appear excessively long because their length is governed not only by the size of the previous contraction but also by earlier contractions the recovery from which were cut short.

15 The analytical expression is $u = s/(s + (f/(1 - f)))$ under our assumption that workers who separate from a firm get a chance to find a new job right away and spend no time unemployed. Assuming instead that workers must necessarily spend one period unemployed before finding a new job, the relationship would be $u = s/(s + f)$. In this case too, the steady-state worker-flow relationship is convex but the non-linearities it induces are quantitatively small.

16 Hairault et al. (2010), Jung and Kuester (2011), and Lepetit (2020) also emphasize this source of non-linearity.

17 Counterfactually large business cycle fluctuations are a feature of the search and matching model under constant returns to labor and would disappear under decreasing returns to labor. See Appendix H on how decreasing returns affect the model.

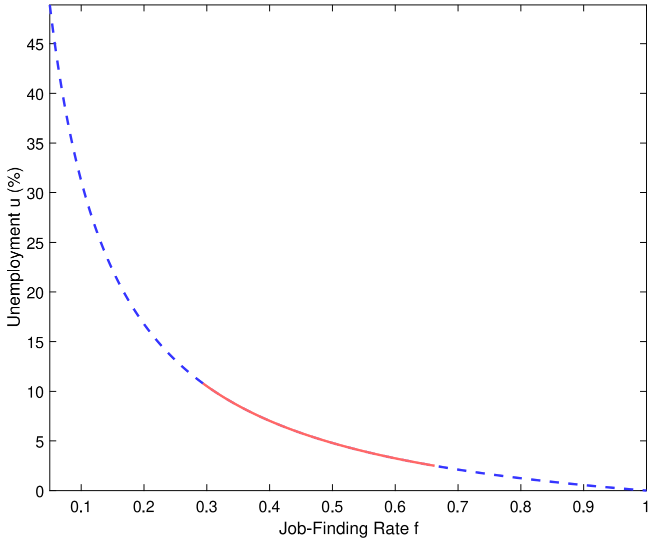

Fig. 4. Steady-State Worker-Flow Relationship.
Note: The figure plots the steady-state relationship between the job-finding rate and the unemployment rate implied by the worker-flow Eqs. (15)–(16). In plain line the relationship is plotted over the range of unemployment rates observed in the US between 1948 and 2019: from 2.5% to 10.8%.

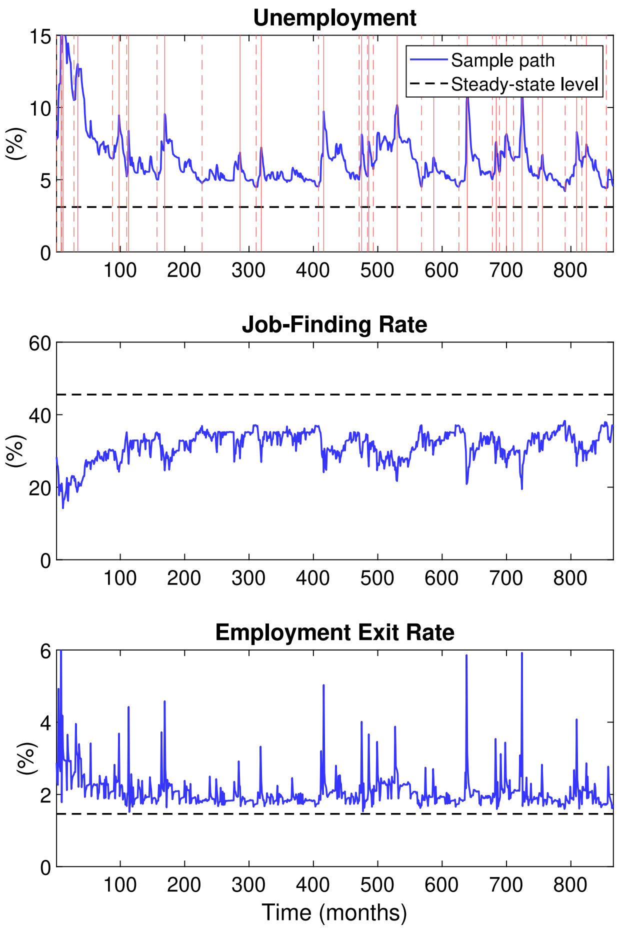

Fig. 5. Simulated Paths for the Unemployment Rate, Job-Finding Rate and Employment Exit Rate.
Note: The figure plots sample path of 72 years (the same length as our empirical sample) for the unemployment rate u, the job-finding rate f and employment exit rate $\bar{s}$ in the Fujita–Ramey model with downward nominal wage rigidity of Section 3. The dashed lines indicate the steady-state level of each variable.

A second weakness of the model with DNWR is that the unemployment rate is much less persistent than in the data. The low persistence in the model is documented in the lower half of Table 3. The model generates unemployment cycles that are far too short. The average duration of expansions is 27.4 months versus 59.1 in the data, and the average duration of contraction is 18.3 months, versus 26.9 in the data.

The model with flexible wages does considerably better on these metrics. With our Hagedorn and Manovskii (2008) calibration, movements in real wages are muted and the model's endogenous variables inherit the behavior of the shock process. This yields quite a bit of persistence. With DNWR, however, it is nominal wages that are rigid in downturns. Real wages fall due to increases in the price level. This implies that gaps between A_t and w_t get eroded faster than in the flexible wage model with the Hagedorn and Manovskii (2008) calibration.

4.2. Stabilization policy and output gaps in the plucking model

In a thought-provoking exercise, Lucas (1987, 2003) showed that the welfare benefits of eliminating all economic fluctuations are trivial in a simple benchmark model. Crucially, Lucas assumed in this thought experiment that the average level of economic activity was unaffected by the elimination of business cycles. Our plucking model violates this assumption.

In a plucking model, recessions are asymmetrically periods when the economy drops below potential. Eliminating these business cycle fluctuations raises the average level of economic activity. The top panel of Fig. 5 shows that in our model, eliminating all fluctuations reduces the average unemployment rate from 5.7% to 3.1%. Conversely, increasing the standard deviation of aggregate shocks by 50% (from 1.5% to 2.25%) increases the average unemployment rate to 12.3%. Figure D.2 in Appendix D plots the average level of the unemployment rate in our plucking model as a function of the volatility of aggregate shocks.

The plucking model also implies that standard measures of the output gap are biased. Aiyar and Voigts (2019) show that common estimation methods of the output gap implicitly assume a zero-mean output gap. In the plucking model, however, output does not fluctuate symmetrically around a natural rate. Standard methods systematically underestimate the amount of slack because the output gap is on average negative.

4.3. Entry and exit over the business cycle

The bottom two panels of Fig. 5 give simulated paths for the job-finding rate f_t and the rate of exit from employment $\bar{s}_t$ in our model. Both the hiring and separation margins contribute to the sharp rise in unemployment during recessions, consistent with evidence documented by Elsby et al. (2009) for US data.[18] Negative shocks decrease the job-finding rate and increase the rate of exit from employment, while positive shocks mostly lead to increases in wages.

We quantify the relative importance of hiring versus separations by calculating their contributions to the volatility of unemployment. For the contribution of the job-finding rate f, we simulate the unemployment rate from Eq. (16) with $\bar{s}$ fixed at its mean, then calculate the standard deviation of the resulting counterfactual unemployment rate, and divide it by the standard deviation of actual unemployment. Analogously, for the employment exit rate, we do the same analysis with the job-finding rate fixed at its mean. According to this metric, fluctuations in the job-finding rate explain 43% of fluctuations in unemployment, while fluctuations in the employment exit rate explain 54%.

Table 4 reports additional statistics: the volatility, the autocorrelation, and the correlation with productivity of the unemployment rate, of the job-finding rate, and of the employment exit rate. We calculate these statistics for the HP-filtered log of simulated series from the model. Relative to the model with Nash bargaining, the model with DNWR lowers the correlation of all three variables with productivity and brings it closer to the data. The autocorrelation statistics are, however, substantially smaller than in the data—a manifestation of the lack of persistence also captured with the speed and duration statistics in Table 3. Also, to match the volatility of the unemployment rate, the model requires shocks that imply a volatility of the separation rate that is about twice as high as in the data.[19]

5. Extensions

Our baseline model does not match the long duration of unemployment cycles in the data, and the asymmetry in the duration of contractions and expansions. Next, we consider whether extensions to the baseline model can help match these features of the data. We analyze the effect of (1) subjecting the model to AR(2) shocks, and (2) assuming that new workers often turn out to be poor matches for the firm. The second of these features yields a job ladder with a "slippery" first rung, increasing the time it takes a firm to rebuild its workforce after layoffs.

[18] The separation margin is not necessary to replicate the plucking property, however. In a previous version of the paper, we considered a version of the model with constant exogenous separation rate and decreasing returns to scale, and showed that it, too, generated the plucking property.

[19] Here, our calibration strategy differs from Fujita and Ramey (2012), who calibrate the idiosyncratic shock process to match the autocorrelation and volatility of the separation rate and, as a result, underperform on the volatility of the unemployment rate.

Table 4
Second order moments.

X_t	u_t	f_t	$\bar{s}_t$
Data			
$\sigma(X)$	0.096	0.077	0.058
$corr(A_t, X_t)$	−0.460	0.369	−0.535
$corr(X_t, X_{t-1})$	0.926	0.803	0.631
Nash			
$\sigma(X)$	0.115	0.063	0.080
$corr(A_t, X_t)$	−0.964	0.965	−0.875
$corr(X_t, X_{t-1})$	0.815	0.771	0.508
DNWR			
$\sigma(X)$	0.130	0.070	0.115
$corr(A_t, X_t)$	−0.81	0.792	−0.539
$corr(X_t, X_{t-1})$	0.666	0.567	0.194

Note: The table reports the second order moments of the Fujita–Ramey model of Section 3 under Nash bargaining and Downward Nominal Wage Rigidity. The moments are reported after data have been converted to quarterly, logged, and HP-filtered. The data panel uses the data from Fujita and Ramey (2012) and is therefore identical to the results they report. For the model panels, the table reports medians of each statistics over our 5000 samples of 866 months each.

5.1. Ar(2) shocks

We first assume that instead of following an AR(1) process, aggregate shocks follow an AR(2):

$$\log(A_t) = (I - \rho_1^a L)^{-1}(I - \rho_2^a L)^{-1}\varepsilon_t^a, \tag{17}$$

still with Gaussian innovations, $\varepsilon_t^a \sim \mathcal{N}(0, \sigma_\varepsilon^a)$. We are motivated by the DSGE literature, which often includes equations that 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. An AR(2) process can generate high persistence at business cycle frequencies without extreme levels of persistence at very low frequencies, and also helps fit the fact that the dynamic responses of economic activity to many shocks is hump-shaped (e.g. Romer and Romer, 2004; Christiano et al., 2005).[20]

We estimate the two roots ρ_1^a and ρ_2^a of aggregate productivity process from the BLS quarterly series on labor productivity. We first apply a three-period moving-average filter to smooth out high-frequency variations including measurement errors. We then detrend the series by removing a quadratic trend. The quadratic trend allows us to capture the productivity slowdown from the 1970s onward. We then estimate an AR(2) on the cyclical component of labor productivity. After converting the roots to a monthly frequency, this yields $\rho_1^a = 0.985$ and $\rho_2^a = 0.88$.[21]

When simulating the model, we recalibrate the bargaining power of workers γ and the volatility of aggregate shocks σ_a so that we continue to match the average and standard deviation of unemployment in the data. This gives $\gamma = 0.55$ and $\sigma_a = 0.016$ under Nash bargaining, and $\gamma = 0.37$ and $\sigma_a = 0.010$ under DNWR. One downside of the AR(2) model is that, with only aggregate shocks, the DNWR binds rarely, because the aggregate productivity process is so smooth. In an earlier version of this paper, we included sectoral shocks, which generated binding DNWR constraints and a large amount of plucking even with 2% inflation and 2% aggregate productivity growth. For simplicity, here, we recalibrate a model with only aggregate shocks to have a 0.5% inflation target to generate a non-trivial role for the DNWR constraint.[22]

The middle columns of Table 3 present the results for the AR(2) version of the model. The model under DNWR features about as much plucking as the baseline model under an AR(1).[23] This model does generate somewhat longer contractions and expansions (37.6 and 23.5 months)—closer to what we see in the data, though still less persistent. The model under Nash bargaining generates less plucking but much longer expansions and contractions. The reason is the same as with AR(1) shocks: under Nash bargaining the gap between wages and productivity inherits the dynamics of the shock process, while under DNWR the gap erodes faster. Neither model generates asymmetries in the speed of contractions versus expansions like those we see in the data.

5.2. Job ladder

We next assume that on top of AR(2) shocks, a newly hired worker can turn out to be a poor match to the firm so that the firm can choose to lay them off and look for another worker. The fact that newly hired workers face much higher separation rates than

[20] Fujita 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.

[21] The autoregressive coefficients of the AR(2) estimation are $\phi_1^a = 1.64$ and $\phi_2^a = -0.65$. They are related to the roots ρ_1^a and ρ_2^a through the equation $I - \phi_1^a L - \phi_2^a L^2 = (I - \rho_1^a L)(I - \rho_2^a L)$. This gives roots 0.96 and 0.68 on a quarterly frequency. The monthly roots are the quarterly roots raised to the power 1/3.

[22] Our analysis of the earlier model with sectoral shocks is available in the NBER working paper version of this paper.

[23] As in our baseline analysis, we report results for a "top-truncated" sample that only includes expansions and contractions that are less than 6.5 percentage points in size to avoid overstating the role of non-linearities arising from counterfactually large business cycles.

workers with longer tenure has been emphasized by Hall (1995), Pries (2004), Krolikowski (2017), Jung and Kuhn (2019), Jarosch (2023), and Hall and Kudlyak (2021), among others. The need to assess workers' suitability on the job instead of through job interviews increases the time needed to find a good match. It can potentially increase the duration of recoveries.

We model this process by adding a job ladder with a slippery first rung to the model of Section 3. We assume that during the first M months of a new hire's tenure at a firm, the newly hired worker faces a probability d of seeing its productivity collapse to zero. The firm then chooses to lay them off. During this trial period, the worker's idiosyncratic productivity if it does not collapse remains constant at its initial value. Once their trial period is over, the worker's idiosyncratic productivity starts to fluctuate according to the idiosyncratic productivity process (2). Appendix G details how we formalize this assumption.

For illustrative purposes, we take the trial period to last $M = 12$ months, and the monthly probability of turning into a poor match to be $d = 10\%$. When simulating the model, we recalibrate the bargaining power of workers γ and the volatility of aggregate shocks σ_a to still match the average and standard deviation of unemployment in the data. This gives $\gamma = 0.29$ and $\sigma_a = 0.013$ under Nash bargaining, and $\gamma = 0.25$ and $\sigma_a = 0.012$ under DNWR.

The rightmost columns of Table 3 present the results for this version of the model. While the model under DNWR features longer expansions than without the job ladder (58.6 months, close to the empirical value of 59.1 months), it also features longer contractions, so that the asymmetry between expansions and contractions is much less than in the data. In addition, the model under DNWR now features somewhat less plucking than in our baseline specification.[24]

6. 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. We show that a workhorse labor search model augmented with downward nominal wage rigidity can fit these facts. In this model, eliminating business cycles lowers the average unemployment rate. Since output is more often below than above the natural rate, standard methods systematically underestimate the amount of slack in the economy.

While our benchmark model with match heterogeneity and downward nominal wage rigidity succeeds in fitting the plucking property, it fails to match the overall duration of unemployment cycles, and the fact that expansions are on average twice as long as contractions. We consider additional non-standard features that may be required to match the asymmetry and duration of contractions and expansions—namely, AR(2) shocks and a job ladder with a "slippery" first rung. However, none of these simple extensions to the basic model can fit the constellation of facts we emphasize. Additional features such as decreasing returns to scale, and sectoral heterogeneity may be important in fitting these facts. A full integration of these features is a promising topic for future research.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jmoneco.2025.103766.

Data availability

Replication package is available as supplementary material at doi https://eml.berkeley.edu/~jsteinsson/papers/pluckingReplication.zip.

References

Abbritti, M., Fahr, S., 2013. Downward wage rigidity and business cycle asymmetries. J. Monet. Econ. 60, 871–886.
Aiyar, S., Voigts, S., 2019. The negative mean output gap. IMF Working Paper No. 19/183.
Akerlof, G., Dickens, W., Perry, G., 1996. The macroeconomics of low inflation. Brookings Pap. Econ. Act. 1996 (1), 1–76.
Altonji, J., Devereux, P., 2000. The extent and consequences of downward nominal wage rigidity. In: Polachek, S. (Ed.), Worker Well-Being. Elsevier, New York, pp. 383–431.
Barattieri, A., Basu, S., Gottschalk, P., 2014. Some evidence on the importance of sticky wages. Am. Econ. Journal: Macroecon. 6 (1), 70–101.
Benigno, P., Ricci, L.A., 2011. The inflation-output trade-off with downward wage rigidities. Am. Econ. Rev. 101 (4), 1436–1466.
Bewley, T., 1999. Why Wages Don't Fall During a Recession. Harvard University Press, Cambridge, MA.
Blanchard, O., Gali, J., 2010. Labor markets and monetary policy: A new Keynesian model with unemployment. Am. Econ. Journal: Macroecon. 2 (2), 1–30.
Bordo, M., Haubrich, J., 2012. Deep recessions, fast recoveries, and financial crises: Evidence from the American record. NBER Working Paper No. 18194.
Caballero, R., Hammour, M., 1998. The macroeconomics of specificity. J. Political Econ. 106 (4), 724–767.
Card, D., Hyslop, D., 1997. Does inflation grease the wheels of the labor market. In: Romer, C.D., Romer, D.H. (Eds.), Reducing Inflation: Motivation and Strategy. University of Chicago Press, Chicago.
Chodorow-Reich, G., Wieland, J., 2020. Secular labor reallocation and business cycles. J. Political Econ. 128 (6), 2245–2287.
Christiano, L.J., Eichenbaum, M., Evans, C.L., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. J. Political Econ. 115, 1–45.

[24] The reduction in plucking occurs mostly in the results for truncated cycles that we report in Table 3. The reduction is negligible when including all cycles, as reported in Appendix E. The model generates very large fluctuations in unemployment that fit the plucking property, but these fluctuations are too large to be included in the untruncated results. The large fluctuations in unemployment are a feature of the search and matching model under constant returns to labor and would disappear under decreasing returns to labor, as detailed in Appendix H. The NBER Working Paper version of this paper considered a version of the model with decreasing returns to scale.

Cole, H., Rogerson, R., 1999. Can the Mortensen-Pissarides matching model match the business-cycle facts. Int. Econ. Rev. 40 (4), 933–959.
Daly, M., Hobijn, B., 2014. Downward nominal wage rigidities bend the phillips curve. J. Money, Credit. Bank. 46 (2), 51–93.
Daly, M.C., Hobijn, B., Lucking, B., 2012. Why has wage growth stayed strong?. Federal Reserve Bank of San Francisco Economic Letter 2012-11.
Daly, M.C., Hobijn, B., Wiles, T.S., 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., Summers, L.H., 1988. How does macroeconomic policy affect output? Brookings Pap. Econ. Act. 1988 (2), 433–480.
Elsby, M., 2009. Evaluating the economic significance of downward nominal wage rigidity. J. Monet. Econ. 56, 154–169.
Elsby, M., Michaels, R., Solon, G., 2009. The ins and outs of cyclical unemployment. Am. Econ. Journal: Macroecon. 1 (1), 84–110.
Fatás, A., Mihov, I., 2013. Recoveries. CEPR Discussion Paper No. DP9551.
Ferraro, D., 2018. The asymmetric cyclical behavior of the U.S. labor market. Rev. of Econ. Dynamics 30, 145–162.
Foster, L., Haltiwanger, J., Syverson, C., 2008. Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? Am. Econ. Rev. 98 (1), 394–425.
Friedman, M., 1964. Monetary studies of the National Bureau. In: The National Bureau Enters Its 45th Year. National Bureau of Economic Research, New York, NY, pp. 7–25, 44th Annual Report, available at http://www.nber.org/nberhistory/annualreports.html.
Friedman, M., 1993. The "Plucking model" of business fluctuations revisited. Econ. Inq. 31, 171–177.
Fujita, S., Ramey, G., 2006. The Cyclicality of Job Loss and Hiring. Working Papers 06–17, Federal Reserve Bank of Philadelphia.
Fujita, S., Ramey, G., 2007. Job matching and propagation. J. Econom. Dynam. Control 31, 3671–3698.
Fujita, S., Ramey, G., 2012. Exogenous versus endogenous separation. Am. Econ. Journal: Macroecon. 4 (4), 68–93.
Gertler, M., Huckfeldt, C., Trigari, A., 2020. Unemployment fluctuations, match quality, and the wage cyclicality of new hires. Rev. Econ. Stud..
Gertler, M., Trigari, A., 2009. Unemployment fluctuations with staggered Nash wage bargaining. J. Political Econ. 117 (1), 38–86.
Gottschalk, P., 2005. Downward nominal-wage flexibility: Real or measurement error? Rev. Econ. Stat. 87 (3), 556–568.
Grigsby, J., Hurst, E., Yildirmaz, A., 2021. Aggregate nominal wage adjustment: New evidence from administrative payroll data. Am. Econ. Rev. 111, 428–471.
Hagedorn, M., Manovskii, I., 2008. The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited. Am. Econ. Rev. 98 (4), 1692–1706.
Hairault, O., Langot, F., Osotimehin, S., 2010. The (un)importance of unemployment fluctuations for the welfare cost of business cycles. Rev. Econ. Dyn. 13, 759–779.
Hall, R.E., 1995. Lost jobs. Brookings Pap. Econ. Act. 1995 (1), 221–273.
Hall, R., 2005. Employment fluctuations with equilibrium wage stickiness. Am. Econ. Rev. 95 (1), 50–65.
Hall, R.E., Kudlyak, M., 2021. Why has the US economy recovered so consistently from every recession in the past 70 years? NBER Macroecon. Annu. 36, 1–55.
Hazell, J., Taska, B., 2018. Downward rigidity in the wage of new hires. Working Paper, MIT.
Jackson, M., Tebaldi, P., 2017. A forest fire theory of the duration of a boom and the size of a subsequent bust. Working Paper, Columbia University.
Jarosch, G., 2023. Searching for job security and the consequences of job loss. Econometrica 91, 903–942.
Jung, P., Kuester, K., 2011. The (un)importance of unemployment fluctuations for the welfare cost of business cycles. J. Econom. Dynam. Control 35, 1744–1768.
Jung, P., Kuhn, M., 2019. Earings losses and labor mobility over the life cycle. J. Eur. Econ. Assoc. 17 (3), 678–724.
Kahn, S., 1997. Evidence of nominal wage stickiness from microdata. Am. Econ. Rev. 87 (5), 993–1008.
Kim, C.-J., Nelson, C.R., 1999. Friedman's plucking model of business fluctuations: Tests and estimates of permanent and transitory components. J. Money, Credit. Bank. 31 (3), 317–334.
Kim, J., Ruge-Murcia, F., 2009. How much inflation is necessary to grease the wheels? J. Monet. Econ. 56, 365–377.
Krolikowski, P., 2017. Job ladders and earnings of displaced workers. Am. Econ. Journal: Macroecon. 9 (2), 1–31.
Kurmann, A., McEntarfer, E., 2017. Downward wage rigidity in the United States: New evidence from administrative data. Working Paper.
Lepetit, A., 2020. Asymmetric unemployment fluctuations and monetary policy trade-offs. Review of Economic Dynamics 36, 29–45.
Lucas, R., 1987. Models of Business Cycles. Oxford:Blackwell.
Lucas, R., 2003. Macroeconomic priorities. Am. Econ. Assoc. 93 (1), 1–14.
McKay, A., Reis, R., 2008. The brevity and violence of contractions and expansions. J. Monet. Econ. 55, 738–751.
McLaughlin, K., 1994. Rigid wages? J. Monet. Econ. 34, 383–414.
Michaillat, P., 2012. Do matching frictions explain unemployment? Not in bad timess. Am. Econ. Rev. 102 (4), 1721–1750.
Mitchell, W.C., 1927. Business Cycles: The Problem and Its Setting. National Bureau of Economic Research, New York.
Neftçi, S., 1984. Are economic time series asymmetric over the business cycle? J. Political Econ. 92 (2), 307–328.
Olivei, G., Tenreyro, S., 2010. Wage-setting patterns and monetary policy: International evidence. J. Monet. Econ. 57, 785–802.
Petrongolo, B., Pissarides, C., 2001. Looking into the black box: A survey of the matching function. J. Econ. Lit. 39, 390–431.
Petrosky-Nadeau, N., Zhang, L., 2017. Solving the Diamond-Mortensen-Pissarides model accurately. Quant. Econ. 8 (2), 611–650.
Petrosky-Nadeau, N., Zhang, L., Kuehn, L., 2018. Endogenous disasters. Am. Econ. Rev. 108 (8), 2212–2245.
Pissarides, C., 2009. The unemployment volatility puzzle: Is wage stickiness the answer? Econometrica 77 (5), 1339–1369.
Pries, M., 2004. Persistence of employment fluctuations: A model of recurring job loss. Rev. Econ. Stud. 71 (1), 193–215.
Romer, C.D., Romer, D.H., 2004. A new measure of monetary shocks: Derivation and implications. Am. Econ. Rev. 94 (4), 1055–1084.
Schmitt-Grohe, S., Uribe, M., 2016. Downward nominal wage rigidity, currency pegs, and involuntary unemployment. J. Political Econ. 124, 1466–1514.
Shimer, R., 2005. The cyclical behavior of equilibrium unemployment and vacancies. Am. Econ. Rev. 95 (1), 25–49.
Shimer, R., 2010. Labor Market and Business Cycles. Princeton University Press, Princeton, NJ.
Sichel, D., 1993. Business cycle asymmetry: A deeper look. Econ. Inq. 31, 224–236.
Silva, I.J., Toledo, M., 2009. Labor turnover costs and the cyclical behavior of vacancies and unemployment. Macroecon. Dyn. 13, 76–96.
Sinclair, T., 2010. Asymmetry in the business cycle: Friedman's plucking model with correlated innovations. Stud. Nonlinear Dyn. Econ. 14 (1), 1–31.
Tasci, M., Zevanove, N., 2019. Do longer expansions lead to more severe recessions?. Federal Reserve Bank of Cleveland Economic Commentary.

Further references in online appendix

Bry, G., Boschan, C., 1971. Cyclical Analysis of Time Series: Selected Procedures and Computer Programs. NBER, New York.
Rouwenhorst, K.G., 1995. Asset pricing implications of equilibrium business cycle models. In: Cooley, T.F. (Ed.), Frontiers of Business Cycle Research. Princeton University Press, pp. 294–330.
Tauchen, G., 1986. Finite state Markov-Chain approximations to univariate and vector autoregressions. Econ. Lett. 20, 177–181.