# The brevity and violence of contractions and expansions*

Alisdair McKay[a], Ricardo Reis[a,b,†]

[a]Department of Economics, Princeton University, Princeton, NJ 08544, USA

[b]Woodrow Wilson School, Princeton University, Princeton, NJ 08544, USA

April 2008

**Abstract**

Early studies of business cycles argued that contractions in economic activity were briefer (shorter) and more violent (rapid) than expansions. This paper systematically investigates this claim and in the process discovers a robust new business cycle fact: contractions in employment are briefer and more violent than expansions but we cannot reject the null of equal brevity and violence for expansions and contractions in output. The difference arises because employment typically lags output around peaks but they coincide in their troughs. We discuss the performance of existing business cycle models in accounting for this fact, and conclude that none can fully account for it. We then show that a business cycle model with asymmetric adjustment costs on employment and a choice of when to scrap old technologies can account for the business cycle fact both qualitatively and quantitatively.

*JEL classification:* E32, E23, E24, J60.

*Keywords:* Business cycles, Economic expansions and contractions, Asymmetric cycles, Unemployment.

*We are grateful to Mark Bils, Bill Dupor, Robert Hall, Jonathan Parker, David Romer, Glenn Rudebusch, Mark Watson, Jon Willis and several seminar participants for useful comments.

†Corresponding author. E-mail address: rreis@princeton.edu

# 1 Introduction

In a series of studies, Wesley Mitchell (1913, 1927, 1946 with Burns) collected a set of facts on the U.S. business cycle. Most of them have been thoroughly scrutinized since then and have survived the test of time. Today, it is well-established that: fluctuations occur in aggregate activity and not in particular sectors; cycles are recurrent but not periodic; cycles have at least two different stages, expansions and contractions; once the economy enters one of the stages, it stays there for some time, so detecting turning points is important for forecasting; and there are regular and predictable co-movements between variables over the cycle that can be expressed as relative variances and lead-and-lag correlations.[1]

There is another fact emphasized by Mitchell that has not received as much attention. In his words: "Business contractions appear to be briefer and more violent than business expansions." (Mitchell, 1927: 333). The aim of this paper is to investigate this claim by looking at the growth cycles in employment and output across different series. Our approach has three stages: first, we examine empirically whether this belongs among the set of business cycle stylized facts, second, we ask whether existing models fit the fact, and third, we propose a simple model that can account for it.

In our empirical investigation, we consider different measures of business activity in the product and labor markets, different procedures to de-trend the data and to detect peaks and troughs, different measures of violence and brevity, and different ways to visually and statistically compare contractions and expansions. After going over hundreds of different combinations of these methods, we reach one robust conclusion: confirming the initial claim, contractions in employment are briefer and more violent than expansions. However, contrary to the initial claim, there is little evidence to reject the null hypothesis that expansions and contractions in output are equally brief and violent, and strong evidence that the asymmetry in output is weaker than that in employment. The difference between output and employment comes from a difference in the timing of turning points: peaks in employment typically lag peaks in output, whereas the troughs in both series are roughly coincident. Because we find that these patterns are very robust, we propose them as a new business cycle fact.

Next, we ask whether existing business cycle models can account for this fact. We

[1] Zarnowitz (1992) and Stock and Watson (1999) summarize the established business cycle facts.

conclude that there is no available theory that can simultaneously account for all of its parts. While some theories can explain why output and unemployment can move in opposite directions during parts of the business cycle, or why contractions in employment are briefer than expansions, or why contractions in employment are more violent than expansions, there is no existing single theory that can account for all three parts of the fact.

We then present a simple business-cycle model that can account for the fact by relying on two mechanisms. The first of these is endogenous technology adoption, or the optimal timing of creative destruction. If firms can sustain ageing technologies for a while by hoarding labor and cutting on hours worked, then expansions in employment can persist even when output has started declining. This can explain why employment lags output at peaks, but not at troughs, and the asymmetry in brevity. To explain the asymmetry in the violence of employment changes, we use a fact that has been extensively documented in the literature on employment adjustments: job separations can occur abruptly while job creations take time because employers can fire quickly but need time to train new workers. Beyond qualitatively explaining the brevity and violence of contractions and expansions, we show that this simple business-cycle model can quantitatively account reasonably well for the U.S. data.

Section 2 of the paper presents the new business cycle fact in a baseline case, and section 3 discusses the robustness of the empirical findings. Section 4 uses our findings to evaluate existing theories of asymmetric business cycles and sections 5 and 6 present a new model to match the facts. Section 7 concludes and discusses the implications of two topics of current interest: European unemployment hysteresis and U.S. jobless recoveries.

*The related empirical literature on asymmetric business cycles*

There is a large empirical literature on asymmetric business cycles that we cannot do full justice to here. Relative to this paper, the literature fits broadly into three branches. The first branch, starting with Neftci (1984) and DeLong and Summers (1986) looks at skewness in either the level or in the changes in economic activity.[2] Skewness in levels would imply that the economy spends more time above or below trend. Our emphasis is instead on the behavior of the economy when it is expanding or contracting. Skewness in

[2] See also Falk (1986), Sichel (1989, 1993), Rothman (1991), Verbrugge (1997), Belaire-Franch and Peiro (2003), and Bai and Ng (2005).

changes evaluates whether economic activity is more likely to increase or to fall. Yet, while economic activity is generally rising during an expansion, there are some periods where it actually falls. In a typical U.S. expansion, output actually falls about one-fifth of the times, and in a typical contraction, output actually rises during one-fourth of the periods. An asymmetry between the dates when economic activity falls and rises does not imply nor is implied by an asymmetry between the business cycle phases of expansion and contraction. While skewness is interesting in its own right, it does not address the brevity and violence of expansions and contractions.

A second branch in the literature, starting with the seminal contribution of Hamilton (1989), estimates regime-switching models and examines whether there are differences between the two regimes.[3] The typical finding in these studies is that the dynamics of recessions are significantly different from the dynamics of booms. Our paper distinguishes itself from this literature because we are not looking at whether contractions are generally different from expansions. Rather, we focus on a more specific difference: whether they are briefer and more violent. This focus allows us to be more precise and to have more powerful tests of this particular type of asymmetry. It implies, of course, that even if we fail to find differences in brevity and violence, there may still be other forms of asymmetry.

A third branch of the empirical literature has focused on specific types of asymmetries. McQueen and Thorley (1993) found that peaks tend to be round, while troughs are sharp. Diebold and Rudebusch (1990) found no evidence that expansions and contractions are duration dependent and Diebold and Rudebusch (1992) and Watson (1994) compared the duration of business cycles in the post-war and pre-war data.[4] Relative to these articles, this paper focusses on a different type of asymmetry and compares post-war expansions and contractions.

Relative to previous empirical work, this paper therefore contributes: (i) the investigation of new types of asymmetries (brevity and violence), (ii) contrasting two states of the business cycle (contractions and expansions), and (iii) systematically comparing the

[3]See Acemoglu and Scott (1994), Ramsey and Rothman (1995), and Hamilton (2005) who use close variants of the Hamilton model. Beaudry and Koop (1993), Hussey (1992), Hess and Iwata (1997), Montgomery et al (1998), Rothman (1998), and Koop and Potter (1999) use other non-linear models to look for business cycle asymmetries. Clements and Krolzig (2003) bridge the two first branches of the literature, using a regime-switching model to look for skewness.

[4]See also McCulloch (1975), Sichel (1991), Durland and McCurdy (1994), and Lam (2004).

behavior of output vis-a-vis employment.[5]

## 2 The new business cycle fact

We start our empirical investigation by looking at one specific case. A common measure of business activity is the state of the labor market, and its most used measure is the unemployment rate. We take the log of 1 minus the quarterly unemployment rate from 1948:1 to 2005:1 and de-trend it using Rotemberg's (1999) modified HP-filter. Then, we identify peaks and troughs using the standard algorithm of Bry and Boschan (1971).[6] The bottom panel of figure 1 shows the periods of expansion and contraction in employment.

To measure brevity, we compute the average number of quarters during expansions and contractions. The average expansion in employment lasted 18 quarters, whereas the average contraction lasted 8 quarters. The test that these are equal has a t-statistic of 2.63, and a p-value of 0.00 against the one-sided alternative that contractions are briefer. In fact, looking at the whole distribution, there isn't a single expansion in the entire sample that lasted shorter than the median duration of a contraction. The data overwhelmingly points to shorter contractions than expansions in employment.

We measure violence by the average change in the series. The average growth rate of employment during an expansion, averaged across expansions, was 0.23% whereas the average growth rate during a contraction, averaged across contractions, was more than double: -0.48%. The t-statistic is 3.22, which overwhelmingly rejects equality in the absolute value of growth rates in favor of the alternative that contractions are more violent. Looking at the whole distribution, every single expansion in employment in the post-war had an average growth rate lower than the median absolute growth rate during a contraction. We thus agree with the initial claim:

**Result 1**: *Contractions in employment are briefer and more violent than expansions.*

To investigate further, we look at another series: the log of GDP. This captures the state of product markets, the other common measure of business activity. The top panel of figure

[5] We survey previous theoretical work in section 4.

[6] Both the Rotemberg modified HP filter and the Bry and Boschan algorithm are described in detail in the companion appendix, McKay and Reis (2008). Briefly, the Rotemberg filter chooses the smoothing parameter in the HP filter to ensure that differences in the trend growth rate are uncorrelated with the cycle, while the Bry-Boschan algorithm applies a series of smoothing filters to the data and looks for local maxima and minima.

1 shows the expansions and contractions in this series, again after using the modified-HP filter and the Bry and Boschan algorithm.

The results for GDP are strikingly different from those for employment. The average expansion in GDP lasted 11.5 quarters, whereas the average contraction lasted 9.5 quarters. A simple t-test of equal duration versus the alternative of longer expansions has a p-value of 0.22. Looking at the whole distribution of durations, we find that only 60% of expansions were longer than the median duration of a contraction. Moving to violence, the average growth of GDP during an expansion is 0.69%, while average growth during a contraction is -0.69%. A t-test that these numbers add up to zero has a p-value of 0.50 so, at typical significance levels, we accept the null hypothesis of equal violence. Moreover, again only 60% of expansions were less violent than the median contraction, and the distribution of violence in expansions is very similar to the distribution in contractions. If one came with a weak prior that expansions and contractions may be equally brief and violent, these data would strongly confirm this view.

Alternatively, perhaps one came with a strong prior that contractions are briefer and more violent than expansions. While the data would certainly shake this view, they may not reverse it. However, we can also test whether output is as asymmetric as employment. The t-test that expansions in output are longer than contractions by as much as in employment (10 quarters) is decisively rejected with a test statistic of 2.97 and a p-value of 0.00. The test that the difference in the growth rates of output is as pronounced as that in employment has a p-value of 0.01. Therefore, focussing on output leads to a very different conclusion from looking at employment.[7]

**Result 2**: *It is difficult to reject the view that expansions and contractions in output are equally brief and violent, but easy to reject the view that output is as asymmetric as employment.*

To understand why output and employment are so different, starting from each employment trough (peak), we recorded the date of the nearer trough (peak) in output. On

[7] Part of the difference between output and employment is due to the two brief cycles in the late 1960s and mid 1990s that appear in GDP but not in the employment rate. Excluding those two cycles, the case for briefer and more violent contractions in output slightly strengthens but remains very weak. In this case, the average expansion now lasts 15.9 quarters and the average contraction 10.2 quarters, with a p-value of 0.06 in a test of equality between the two. Growth during an average expansion is now 0.51% and during an average contraction -0.59%, with a p-value of 0.25. Result 2 is not just due to more output cycles.

average, troughs in employment lag output by only 0.13 quarters or about 12 days. The difference in duration between the two series therefore comes almost exclusively from peaks in employment lagging those in output by about 2.25 quarters or 203 days. This can be seen in figure 2, where we plotted the average dynamics of output and employment in a window of 7 quarters centered around the peak and trough for output. Employment is therefore a lagging indicator of output cycles only when coming down but not when going up.

**Result 3**: *Employment lags output at peaks but coincides with it at troughs.*

There is an alternative way to examine the data that merges the three results into a single moment. It consists of looking at the skewness of 4-quarter changes in the series. Negative skewness would reflect more rare and less intense instances of 1-year long periods of contraction and expansion.[8] This moment speaks less directly to the fact that we are investigating, since 1-year changes are imperfect measures of contractions and expansions. However, it has the virtues that it does not require de-trending the series or detecting turning points, so it is robust to the methods used in these. In the U.S. data, the skewness of 4-quarter changes in GDP is -0.06 with a p-value for the one-sided test of zero skewness of 0.41. For the employment rate, the skewness coefficient is -0.70 and the p-value is 0.02. Therefore, looking at skewness confirms results 1 to 3.

To conclude, figure 3 summarizes the peak-to-peak dynamics of output and employment suggested by the three results. Starting from a trough, employment in a recovery rises at a slower pace than output. Output eventually reaches its peak and starts falling, while employment keeps rising at a tame pace. Only almost 7 months after the peak in output does employment finally reach its peak, after which it falls sharply catching up with output at the next trough.

## 3 Is the fact robust?

To establish the business cycle fact, our empirical strategy consisted of five steps. We have investigated the robustness of each of them, and summarize here the results.[9]

*(i) Choosing a measure of business activity*

[8] Still, one-year changes are a better measure of changes during a business cycle stage than one-quarter changes as used by Neftci (1984) and discussed in section 1. Our results are robust to using 6-quarter changes instead.

[9] The details are outlined in a companion appendix, McKay and Reis (2008).

We looked at alternative series for output (GDP, industrial production, non-farm business output, real sales, real personal income, consumption, investment), and for employment (total payroll employment, total household employment, employment for 16-24 year olds, employment for workers over 24) and found similar dates for turning points, and thus similar conclusions on brevity and violence. A different time period for output (pig-iron production 1877-1929) and different frequency (monthly) also did not change the results. Total hours workers behaved in a similar way to output, while hours per worker had very different turning points than those for either employment or output.

*(ii) De-trending the data*

A series (like output) that trends up, will automatically have longer expansions and shorter contractions, since it rarely declines. The question of brevity and violence refers to "growth cycles" as opposed to "classical cycles" (Zarnowitz, 1992), and to investigate it requires de-trending the data.[10] We considered alternative de-trending filters aside from the modified HP-filter: linear trend with breaks, polynomial trend, and a band-pass filter that extracts cycles of durations 6-32 quarters or 2-80 quarters. All led to similar results.

*(iii) Dating turning points*

Aside from the Bry and Boschan (1971) algorithm, we considered three other methods to detect peaks and troughs: a "window method" that smooths the series with a 5-quarter centered moving average, at each date forms a symmetric 11-quarter window, and then sees whether the date is a maximum (for peaks) or minimum (for troughs) in the window; a "reversal method" that defines a peak as a date preceded by 3 consecutive quarters of increases and followed by 2 consecutive quarters of decreases (and the reverse for troughs); and a "Markov regime-switching" method that, following Chauvet and Hamilton (2005), identifies contractions and expansions as unobserved states that follow a Markov chain. We also experimented changing the number of quarters in the implementation of each and tested the algorithms on simulated data (with a close to 100% success rate). All of the methods found roughly similar turning point dates, all different from the NBER's chronology. The reason for the difference has to do with detrending, since the Bry and Boschan (1971) algorithm can reproduce almost exactly the NBER dates if the output

[10] One alternative to de-trending is to look at the skewness of 4-quarter changes as in the previous section. Another alternative is to fix the dates of peaks and troughs (and so brevity), and then use trending data to investigate violence. Both confirm our results.

series is not de-trended. Comparing the turning points for employment and output, typically peaks in employment lag peaks in output by between 1 and 3 quarters, whereas troughs in employment are typically within one quarter of troughs in output. The contractions in both output and employment end around the same time.

*(iv) Measuring brevity and violence*

Brevity seems uncontroversial, but there are alternative measures of violence. We considered two: the square root of the average squared change in the series, and the least-squares coefficient on a linear trend from a regression of the series on the trend and an intercept. If during a contraction (or expansion) a series falls exactly linearly, then all three measures are identical, but otherwise, the squared change adds a measure of the volatility of the series, while the least-squares coefficient is more robust to the exact location of the turning points. All methods confirmed our conclusions.

*(v) Systematically comparing expansions and contractions*

We took six approaches to infer whether contractions are different from expansions. First, we inspected the cumulative distribution functions (cdf's) across $i$, looking to see if the cdf for the duration of contractions tends to lie to the left of the cdf for expansions (and the reverse for violence). We systematically found that the cdf for the duration of employment during expansions first-order stochastically dominates that during contractions, while the distributions for output lay on top of each other. For violence, the results were not as overwhelming, but there was still a very clear contrast between employment and output.

Second, we used a t-statistic to test the null hypotheses of equal average duration (violence) against the one-sided alternative of shorter (more violent) contractions. We used both the asymptotic distribution as well as a bootstrap to obtain 5% critical values. We almost always reject the null for employment, while almost never for output.

Third, we tested the null hypothesis that output is as asymmetric as employment, again using a t-test. The null was typically rejected at 5% or 10% significance levels.

Fourth, we computed the skewness of 4-quarter changes in the series, and tested whether it is zero, using the test in Bai and Ng (2005). We could never reject zero skewness for output, but did so in the majority of the cases for employment.

Fifth, we tested the null hypothesis that the distributions of duration and violence are the same for expansions and contractions using a Wilcoxon rank-sum test with the exact p-values for each sample size (Diebold and Rudebusch, 1992). The results were consistent with the t-tests.

Sixth, we took into account the fact that the series for duration and violence are the product of algorithms by estimating symmetric models and generating artificial times-series, on which the algorithms and tests are then applied. Surprisingly, p-values typically became lower, so the rejections of symmetry for employment are stronger than before, while for output we could still typically not reject symmetry at the 5% level.

*The bottom line:* After trying hundreds of different combinations of the available methods and looking into the details of how each works, we found that the results in section 2 are very robust. The pattern that emerges from the data is clear: contractions in employment are briefer and more violent than expansions. Contractions and expansions in output are either equally brief and violent or slightly different, but definitely less asymmetric than employment. Employment and output differ because employment typically lags output at peaks but they roughly coincide in their troughs.

## 4 Can existing theories account for the fact?

There are a few existing models that generate asymmetric business cycles. Since the precise asymmetry that we found in the U.S. data is new, these models were of course not designed to fit it. In this section, we ask whether they are able to do it.

Credit constraints that bind during booms but not recessions are a source of asymmetry in Kocherlakota (2000). Large negative shocks can lead to large cuts in production since agents cannot borrow, while positive shocks are attenuated using savings. Credit constraints can explain the different reaction to positive and negative shocks, but they do not account for the difference between expansions and contractions. Moreover, credit constraints should affect both output and employment equally. The same problem arises with theories that emphasize capacity constraints. Gilchrist and Williams (2000) and Hansen and Prescott (2005) argue that during booms firms hit capacity constraints so expanding production requires expending resources to set up more plants. In recessions instead, some plants are not used and can be re-activated or de-activated at no cost. This model generates asymmetries in both output and employment.

Jovanovic (2006) focuses instead on mismatches between skills and technologies. In his model, firms must adopt technologies without knowing whether they are a good fit for their production process. As bad fits lower output by more than good fits raises it, output is negatively skewed. While skewness is an important feature of the data, it is conceptually distinct from brevity and violence of expansions and contractions.

Caballero and Hammour (1996) analyze an economy in which firms at each date face the option of paying a cost to scrap their old technology and adopt a new one. With technological progress, they show that this creative destruction should be bunched around recessions, when the marginal profitability of production is lower. If new technologies are embodied in jobs, then there is a sharp increase in unemployment around recessions. This model generates violent and short-lived contractions in employment. However, output follows the same dynamics as employment.

Increasing returns to scale can be another source of asymmetry. Acemoglu and Scott (1997) argue that investment in maintenance today not only raises productivity today but also lowers the cost of adopting new technologies tomorrow. Past shocks therefore affect the profitability of current investments and thus the economy's response to shocks. While their model is flexible enough to account for different types of asymmetries between prolonged expansions and prolonged contractions, it emphasizes investment as the source of asymmetries and output as its reflection. Our findings emphasize that employment is the key.

Chalkley and Lee (1998) and van Nieuwerburgh and Veldkamp (2006) argue that when output is high, investors face less uncertainty about productivity. Around peaks, they therefore respond to bad shocks quickly, leading to violent contractions, at least initially. Around troughs, there is less precision of information so the response to positive shocks is slow. These theories can account for differences in violence. However, they lead to a difference between contractions and expansions in output and investment, but not necessarily in employment. Moreover, they do not generate asymmetries in brevity.

From the perspective of the labor market, Burgess (1992) argues that the cost of adjusting employment for a firm depends on the tightness of the labor market. In booms, the labor market is tight, it is costly to fill a vacancy, so employment moves slowly. In slumps, the market is slack, it is easy to fill vacancies, so employment moves quickly. Our finding however was that expansions were different from contractions, rather than booms different than slumps. In the model of Burgess, the initial stage of contractions would be more violent than its later stages (and the reverse for expansions), but on average, expansions and contractions would be equally violent.

The model of job creation and job destruction of Mortensen and Pissarides (1994) can lead to different violence during expansions and contractions. In their model, job destruction occurs immediately once the value to the firm and the worker of being matched is negative. Job creation on the other hand takes place only with some probability. Thus, employment can fall quickly and violently, but it must expand slowly. However, output equals employment so it is asymmetric as well and expansions are as brief as contractions.[11]

Overall, we conclude that existing theories cannot account for all of: (1) briefer and more violent contractions than expansions in employment (2) approximately symmetric expansions and contractions in output, and (3) employment lagging output at peaks but not at troughs. The new business cycle fact in this paper puts forward a new challenge for business cycle models.

# 5 A model to fit the facts

A simple business cycle model that merges the neoclassical propagation mechanism with features from the literatures on labor adjustment and technological adoption can account for the facts.

## 5.1 The standard setup

There is a representative agent that maximizes:

$$\mathbb{E}_0 \left[ \int_0^\infty e^{-\rho t} \left( \ln C_t - b H_t \right) dt \right], \tag{1}$$

where $C_t$ is aggregate consumption and $H_t$ are total hours. Behind these aggregates are a continuum of households with unit mass. The aggregate budget constraint is:

$$C_t + dK_t/dt = Y_t - \delta K_t, \tag{2}$$

[11] Millard et al. (1997) investigate the performance of some of the models in this section at fitting the persistence of unemployment in response to shocks during recessions and booms.

so consumption plus savings in capital $K_t$ equal income net of depreciation at rate $\delta$. Output comes from a Cobb-Douglas production function:

$$Y_t = \left(e^{\gamma_t} L_t\right)^{\alpha} K_t^{1-\alpha}, \tag{3}$$

where $L_t$ is labor input and $\gamma_t$ is productivity, which is hit by random shocks around the average trend growth $\gamma$.

Labor input is given by a slightly more complicated expression that adds three familiar terms:

$$L_t = \int_1^{\bar{A}} A_t(A_j q_{j,t} + l_{j,t} - z_t I(q_{j,t}))d\Phi(A_j) - MC_t - AC(F_t, H_t). \tag{4}$$

The first of the three terms on the right-hand side links hours worked to labor input. Following Bils (1987), Cho and Cooley (1994) and many others, we explicitly consider the benefits of specialization and diminishing returns to effort. A central feature of the labor market is the tension between having, at any point in time, more employed working less hours, or fewer employed working longer. At the heart of this trade-off is the fact that initially people may be very productive at specific tasks to which they are suited, but as they work longer hours they are increasingly occupied with tasks with lower marginal returns. We model this in the simplest possible way by assuming a step function for productivity as a function of hours. The first hours at work, $q_{j,t} \in [0, 1]$, can be employed at a specific task $j$ in production up until a limit (say the first 8 hours), while extra hours, or overtime $l_{j,t} \geq 0$, are used in other activities with the other workers. Total hours by worker at task $j$ are therefore $h_{j,t} = q_{j,t} + l_{j,t}$. Specialized tasks are more productive than overtime by a factor $A_j \geq 1$, and there is a distribution of tasks $\Phi(A_j)$ in the interval $[1, \bar{A}]$. Specialization is limited because it costs $z$ units of non-specialized labor to manage each task operated, where $I(q_{j,t})$ is an indicator function equal to one if $q_{j,t} > 0$ and zero otherwise.[12] Finally, this total effective labor input is multiplied by $A_t$, a measure of overall technological progress common to all tasks.

The second term on the right-hand side, $MC_t$, refers to the overhead costs of maintaining a technology in operation. The third term $AC(F_t, H_t)$ measures adjustment costs in employment, as in Hansen and Sargent (1988), Cho and Bils (1994), and many oth-

[12] We assume that $z$ is not too large, so it is always optimal to have some workers employed.

ers. The total number of workers ($N_t$) can change without cost through separations at rate $\chi$, but voluntary hiring ($H_t \geq 0$) and firing ($F_t \geq 0$) requires diverting labor away from production and thus leads to lower labor input. The dynamics of employment are $dN_t/dt = H_t - F_t - \chi N_t$.

### 5.2 New ingredients

To this standard setup, we add two ingredients that have been extensively discussed in other literatures, but are somewhat novel to the business cycle literature.

The first ingredient concerns the adoption of new technologies. A large literature has consistently found that the adoption of new technologies (Rogers, 1995), the productivity gains from a new technology (Evenson and Westphal, 1995), or the spread of new products (Mahajan et al., 2000) all follow an S-shaped curve, initially slow, then fast and finally slow again. Innovation typically comes in 3 stages. First, there is a period of initial adaptation, with agents learning how to exploit the full potential of the new technology. Take-up is slow and productivity gains are small. Then, we enter a second stage with a quick diffusion of the technology, impressive productivity gains, and mass production. Finally, these benefits peter off, as productivity stagnates and the technology becomes increasingly obsolete. There are several theories that try to explain this S shape, ranging from learning by doing to the spread of information in networks. We take the life-cycle as given, focusing on its consequences for when technologies are adopted in the business cycle.[13]

Specifically, we assume that firms choose when to adopt a new technology, which determines the productivity of tasks ($\Phi(.)$ and $d\ln A_t/dt$) and the cost of maintenance ($MC_t$). Initially, labor productivity grows slower by the rate $g_1$, and because firms have not yet fully developed the tasks that exploit the new technology, the productivity of the tasks is lower so $\Phi(A_j) = \Phi^N(A_j)$. Starting the instant after adoption, this first stage can end with the arrival of a Poisson counter with parameter $\mu_1$. In the second stage, productivity rises fast, above $\gamma_t$ by $g_2$ and the firm starts operating the technology at its full potential

[13] While we focus on the business cycles, Comin and Gertler (2006) show that the interaction of technology adoption and the product cycle with R&D and embodied and disembodied productivity may also be important to understand medium-term fluctuations. We use our simple setup instead of the richer setup in Comin and Gertler (2006) to keep the model and its ingredients more transparent. However, we suspect that using that model would lead to similar predictions for the facts that we focus on. Empirically, Bernard et al (2006) use micro-data on firms to document widespread product changes at business-cycle frequencies: two-thirds of firms change their product mix in a 5-year period.

employing the tasks in $\Phi^E(A_j)$, such that $\int_x^{\bar{A}} A_j d\Phi^E(A_j) > \int_x^{\bar{A}} A_j d\Phi^N(A_j)$ for any $x$. With production at full speed, $\kappa A_t$ units of labor must be used for maintenance, so the cost is constant in efficiency units. The last stage arrives also through a Poisson counter with rate $\mu_2$. The product becomes progressively obsolete, with labor productivity growing at only $\gamma_t - g_3$ while the labor to maintain production remains constant at $\kappa A_\tau$ where $\tau$ is the date at which stage 3 started. Business cycles take place as the economy adopts new products and goes through their stages.[14]

The second ingredient is asymmetric costs of adjusting employment. A large empirical literature has documented significant asymmetries in the costs of adjusting the number of workers (Hamermesh and Pfann, 1996). We assume that hiring new workers involves training them, which is subject to strong decreasing returns to scale, while the marginal cost of firing a worker is close to independent of how many workers the firm fires. To make this distinction clear, we assume that the marginal cost of firing is constant at $\beta$ labor hours, while hiring new workers involves decreasing returns to scale at rate $\upsilon$. Therefore, $AC(F_t, H_t) = \beta F_t + H^{1/\nu}$.

These two premises therefore have solid foundations in empirical work and have been separately studied before. The innovation here is to combine particular versions of creative destruction and training costs in a tractable framework that can account for the brevity and violence of contractions and expansions. The particular functional forms and other details are solely for analytical convenience.

# 6 The properties and predictions of the model

We first solve two simpler versions of our model analytically to develop intuition on its properties. Then, we take the full model to the data, numerically solving and simulating it in order to compare its quantitative predictions with the facts.

## 6.1 The model with only technology adoption

First, we assume away adjustment costs. Moreover, to obtain analytical results, we assume that capital is fixed and there are no shocks to trend-productivity. McKay and Reis (2008)

[14] The arrival of production stages, which we model as exogenous, is likely affected by shocks to fiscal policy, monetary policy or productivity. The model is silent as to what causes business cycles, but instead focusses on explaining the interesting dynamics of employment and output that we found in the data.

shows that:

**Proposition 1.** *If* $AC_t = 0$, $K_t = 1$, *and* $\gamma_t = \gamma t$ *then:*

*a) Output is* $\ln Y_t = \alpha \ln(\alpha/b) + \alpha\,(\gamma + (-1)^s g_s)\,t$*;*

*b) Firms operate all tasks with productivity* $A_j > 1 + z$*, so employment is* $1 - \Phi^E(1+z)$ *in the production stages and* $1 - \Phi^N(1+z)$ *in the research stage;*

*c) When the productivity of an obsolete technology falls below a threshold fraction* $A^*$ *of what it was at the peak, the economy adopts a new technology.*

In this economy, a trough arrives when there is a switch from the research to the production stages of product development. Employment jumps up, as there are more productive uses for workers, and de-trended output starts rising as productivity grows and overtime expands. Eventually, the technology becomes obsolete and the cycle reaches its peak. While de-trended output starts falling, employment remains high since there is a benefit of maintaining production at its relatively high productivity. Only when productivity becomes too low, does the economy switch to researching a new technology. Output continues to fall and employment jumps down as the specialized workers that used the old technology became obsolete. A new expansion will start once this technology enters the production stage.

This stylized version of the model leads to the peak-to-peak dynamics in figure 4. Note that this simple version of the model is already able to match the facts that employment lags output at peaks but not at troughs, and that contractions in employment are briefer than expansions. Cutting on overtime allows firms to keep workers even though output is falling, and sticking to an obsolete technology allows them to maintain a product beyond its peak. Combining the ability to hoard labor by varying overtime with the ability to choose when to creatively destroy technologies already goes a long way towards fitting the business cycle facts. The missing feature is the violence of contractions in employment, which leads us to include adjustment costs.

## 6.2 The model with both technology adoption and adjustment costs

Adding adjustment costs, while maintaining the assumptions on capital and technological progress that allow for analytical solutions, McKay and Reis (2008) show:

**Proposition 2.** *If* $K_t = 1$*, and* $\gamma_t = \gamma$ *then:*

*a) Output is the same as in Proposition 1* *b) Employment in state $s$ can be in three regions: if $N_t^s > \bar{N}^s$, there is firing and a jump to $\bar{N}^s$, if $N_t^s \in \left[\hat{N}^s, \bar{N}^s\right]$ employment falls, and if $N_t^s < \hat{N}^s$ employment rises.*

c) *There is a productivity threshold for technology adoption $A^{**}$.*

Starting from a trough, employment is around $\hat{N}^1$. The economy enters the production stage and workers must be hired, but there are increasing marginal costs of doing so. Hiring therefore proceeds gradually and employment rises slowly even as output expands with productivity with the help of overtime. The peak arrives when the technology becomes obsolete. At this point, employment remains high, even as output drops until finally the economy switches to a new technology. Then, if specialized tasks are sufficiently unproductive during the research stage relative to the production stages, it is optimal to cut employment. Since the marginal cost of firing the extra worker is constant, this comes through a burst of firing.[15] Employment does not fall all the way though because exogenous separations (which are costless) can be left to deplete the stock of remaining workers. Employment continues to fall, now at a declining rate, until a new trough arrives. The peak-to-peak dynamics of this economy are in figure 4.

### 6.3 Solving and calibrating the full model

Using $\hat{L}_t$ to denote labor in efficiency units $L_t/A_t$, the solution of the full model is:

**Proposition 3.** *The dynamics of the economy in equilibrium are:*

*a) Output, consumption, and the capital stock solve a standard stochastic growth problem:*

$$\max_{C_t,\hat{L}_t} \mathbb{E}_0 \left[\int_0^\infty e^{-\rho t}\left(\ln C_t - b\hat{L}_t\right) dt\right] \tag{5}$$

$$s.t. \quad : \quad dK_t/dt = \left(e^{\gamma_t} A_t \hat{L}_t\right)^\alpha K_t^{1-\alpha} - C_t - \delta K_t, \tag{6}$$

*b) Employment can be in three regions as in Proposition 2. Its dynamics are depicted in the phase diagram in figure 5. The boundary conditions are the initial level of employment and a transversality (or smooth pasting) condition.*

*c) There is a productivity threshold for technology adoption determined by a value matching condition.*

McKay and Reis (2008) prove this proposition and describe our method to numerically

[15] The condition for there to be a burst of firing is that $\Gamma(\Phi^{-1}(.))$ is sufficiently large.

solve it. The solution method involves nesting a log-linear approximation of the stochastic growth model with a shooting algorithm to find the saddle path of the phase diagram.

We choose parameters by calibration using as time unit one quarter. Our general approach was to set parameters to fit the relevant first and second-order moments in employment, hours, and output in the data. Table 1 has the parameter values. The test of the model is then whether it can generate the asymmetries, or higher-order moments, that we have found in the data.

We set the first set of parameters, on preferences and the production function, in the conventional way in the literature. The second set of parameters, on trend productivity, were calibrated to match the trend in GDP computed by the modified HP-filter. We assumed that trend productivity follows a mean-reverting process (a continuous-time AR(1)): $d\gamma_t = -\xi(\gamma_t - \gamma)dt + \sigma dW_t$ where $dW_t$ is a standard Wiener process. We picked the long-run productivity growth rate ($\gamma$), the speed of mean reversion ($\xi$), and the standard deviation of shocks ($\sigma$) to fit the properties of trend output.

The parameters determining task-specific productivity determine the volatility of the economy as it goes through the business cycle. We set $g_1 = g_2 = g_3 = g$ and picked the common $g$ to match the variance of de-trended output. The arrival rates of the stages in the diffusion of technology, $\mu_1$ and $\mu_2$, were picked to match the number of business cycles in the data and to ensure that the trend in output comes from $\gamma_t$ only. We assumed the distribution of tasks is uniform with top technology either $\bar{A}^N$ or $\bar{A}^E$. The absolute level of these does not affect the dynamics of employment, so we normalize $\bar{A}^E$ to equal 3. What is important for the model dynamics is the difference between the technologies, since it determines the difference in the number of workers that are worth keeping around. We therefore set $\bar{A}^N$ to match the standard deviation of output per hour.

The fourth set of parameters determine the use of labor along its intensive and extensive margins. We set the administrative costs per task to match the average unemployment rate. The overhead maintenance costs were set to match the variance of hours.

We picked the fifth set of three parameters on the adjustment cost function and employment dynamics to match the volatility of employment. Specifically, we set the rate of exogenous separations and the marginal cost of firing to match the variance of the level of employment and its changes. Finally, following the convention, we assumed quadratic hiring costs.

### 6.4 The model's performance

We now ask whether the model can quantitatively fit the facts. Using a random draw for the shocks, we simulate the model to generate time-series for raw output and employment for 229 quarters, the duration of our sample. We then treat these simulated time-series as data and apply our algorithms (de-trending, picking turning points, measuring brevity and violence). We repeat this for 1000 draws and report in table 2 the average moments.

First looking at output, the model generates business cycles that are equally brief in contractions and expansions, and somewhat more violent expansions than contractions. In the data, expansions tended to be slightly longer and more violent than contractions, although we could not statistically reject that they are the same.

Turning next to employment, our model is able to generate asymmetric fluctuations matching our three empirical results. Noticeably, the model is able to generate an asymmetry in brevity of 6 quarters in employment despite 0 asymmetry in output (versus 10 and 3 in the data). And, it generates contractions that are 6 times sharper than expansions in employment despite approximately same violence in output cycles.

To quantitatively see the key role played by technology adoption and asymmetric adjustment costs, the last column of table 2 displays the predictions for brevity and violence if we shut off these mechanisms. In this frictionless economy, firms must adopt immediately a new technology once it gets discovered and can costlessly hire and fire. In this case, the model predicts symmetry in the brevity and violence of both output and employment across expansions and contractions.

Finally, to further test our model, we looked in more detail at the behavior of hours around peaks. Our story is that, after output peaks and starts falling, firms keep their workers employed but cut on their hours at work for approximately 2.25 quarters. Only then, do they start laying off workers. The model is ambiguous about whether hours fall or rise while employment is falling, since this depends on the relative rates at which employment and output fall but, in between the peaks of output and employment, it unambiguously predicts that we should see hours per worker falling. Figure 6 plots the path of hours per worker around the peak, averaged across all cycles in the data. There is some evidence in favor of our mechanism: hours per worker fall rapidly in between the peaks of output and employment.

To conclude, the model does not generate quite enough asymmetry in brevity and a little too much asymmetry in violence relative to the data. However, it gets close to quantitatively matching the facts, in spite of our very simple functional forms and stochastic processes. Moreover, a crucial part of the model's dynamics predicts that hours per worker should fall between the peaks of output and employment, and this seems to be the case in the data. This leads us to expect that more sophisticated models with asymmetric employment adjustment costs and technology adoption will be successful at matching both the well-known second-order properties of the data as well as the new higher-order moments that this paper uncovered.

## 7 Conclusion

This paper investigated whether business contractions are briefer and more violent than business expansions. We started by looking at U.S. post-war data to find that contractions in employment are briefer and more violent than expansions, but for output we could not reject symmetry. The difference between the two series arises because typically employment peaks 2 or 3 quarters after output, but the two coincide in their troughs. We performed hundreds of sensitivity checks on these results and found them very robust.

We discussed existing models of asymmetric business cycles and found that none of them could fully account for the facts. We then proposed a neoclassical model in which business cycles are driven by the endogenous adoption of new technologies and firms face costs of adjusting employment but can vary overtime hours. A calibration found that the model could roughly fit the data.

These results open a few new questions on business cycles. For instance, one might wonder how our findings inform the current debate on jobless recoveries in the United States. We have found that on average, in the post-war, troughs in employment and output have coincided so jobless recoveries are not the norm. However, we have also found that starting from a trough, employment expands at a slow pace in the beginning of a recovery. This may lead to an impression of joblessness at the start of a recovery. The decline in volatility in the last 20 years may have made recoveries even tamer, which may have made recoveries start seeming jobless.[16]

[16] According to their recent reports, the NBER business cyle dating commitee has particularly struggled to

Another question is whether our results extend to other countries. We have looked at the G-7 countries for which there is good quarterly data from 1960. Our business cycle fact applies to Canada, while in Japan, it is difficult to identify any cycles in the unemployment rate. In Europe, we cannot reject symmetry for output, but unemployment is quite different. After rising abruptly during a few years in the 1980s, unemployment in many European countries has been slowly declining since then. Brief and sharp increases in unemployment followed by protracted reductions is exactly what we also found for the United States. The difference between U.S. and European unemployment dynamics is on the much slower pace of decrease in unemployment in Europe, consistent with European labor markets being more rigid in their adjustments.

A further question is whether there are alternatives to our model. Given a new fact and the absence of theories to explain it, we proposed one theory and showed it could account for the fact. There may be plausible alternatives to the ingredients in our model. For instance, aside from being able to vary workers or hours, firms may be able to vary capital utilization (Greenwood et al, 1998), organizational capital (van Rens, 2004), or organizational restructuring (Koenders and Rogerson, 2005). Aside from choosing when to adopt technologies, firms may be able to choose when to switch between modes of governance (Philippon, 2006), or may learn at different speeds about positive versus negative productivity changes. Finally, aside from adjustments costs in employment, there is also the time it takes to match workers to firms in the labor market (Mortensen and Pissarides, 1994) or to learn about the quality of employer-employee matches (Pries, 2004).[17] While these articles do not fit our facts, one might ask if they (or others) can be modified to do it and, if so, how do they compare with our model. Another question is how to extend our model to include heterogeneous firms and workers and idiosyncratic productivity shocks in order to generate a richer description of the labor market and new predictions for the cross-section of firms. Addressing these extra questions goes beyond what a single article can achieve, but we hope that future work will explore them.

reconcile business cycles in employment and in output. Our findings suggest that it will indeed be difficult to look for common turning points employment and output, since the two have different dynamics. Bachmann (2007) uses a model with some similar features to ours to investigate jobless recoveries.

[17] In a previous draft, we considered job search and matching has an alternative to adjustment costs. Michelacci and Lopez-Salido (2007) explore further their interaction with the adoption of technologies.

**References**


Acemoglu, D., Scott A., 1994. Asymmetries in the cyclical behaviour of UK labour markets. Economic Journal 104, 1303-1323.

Acemoglu, D., Scott A., 1997. Asymmetric business cycles: theory and time-series evidence. Journal of Monetary Economics 40, 501-533.

Bachmann, R., 2007. Understanding jobless recoveries: a tale of two margins. Unpublished manuscript.

Bai, J., Ng, S., 2005. Tests of skewness, kurtosis, and normality in time series data. Journal of Business and Economics Statistics 23, 49-60.

Beaudry, P., Koop, G., 1993. Do recessions permanently affect output? Journal of Monetary Economics 31, 149-163.

Belaire-Franch, J., Peiro, A., 2003. Conditional and unconditional asymmetry in U.S. macroeconomic time series. Studies in Nonlinear Dynamics & Econometrics 7, article 4.

Bernard, A.B., Redding, S.J., Schott, P.K., 2006. Multi-product firms and product switching. NBER working paper No. 12293.

Bils, M., 1987. The cyclical behavior of marginal cost and price. American Economic Review 77, 838-855.

Burgess, S.M., 1992. Asymmetric employment cycles in Britain: evidence and an explanation. Economic Journal 102, 279-290.

Bry, G., Boschan, C., 1971. Cyclical Analysis of Time Series: Selected Procedures and Computer Programs. National Bureau of Economic Research, New York.

Burns, A.F., Mitchell, W.C., 1946. Measuring Business Cycles. National Bureau of Economic Research, New York.

Caballero, R.J., Hammour, M.L., 1996. On the timing and efficiency of creative destruction. Quarterly Journal of Economics 111, 805-852.

Chalkley, M., Lee, I.H., 1998. Learning and asymmetric business cycles. Review of Economic Dynamics 1, 623-645.

Chauvet, M., Hamilton, J.D., 2005. Dating business cycle turning points. NBER Working Paper No. 11422.

Cho, J.O., Bils, M., 1994. Cyclical factor utilization. Journal of Monetary Economics 33, 319-354.

Cho, J.O., Cooley, T.F., 1994. Employment and hours over the business cycle. Journal of Economic Dynamics and Control 18, 411-432.

Clements, M.P., Krolzig, H.M., 2003 Business cycle asymmetries: characterization and testing based on Markov-switching autoregressions. Journal of Business and Economic Statistics 21, 196-211.

Comin, D., Gertler, M., 2006. Medium-term business cycles. American Economic Review 96, 523-551.

DeLong, J.B., Summers, L.H., 1986. Are business cycles symmetrical? In: Gordon, R.J. (Eds.) The American Business Cycle. Continuity and Change. University of Chicago Press, Chicago, pp. 166-179.

Diebold, F.X., Rudebusch, G.D., 1990. A nonparametric investigation of duration dependence in the American business cycle. Journal of Political Economy 98, 596-616.

Diebold, F.X., Rudebusch, G.D., 1992. Have postwar economic fluctuations been stabilized? American Economic Review 82, 993-1005.

Durland, J.M., McCurdy, T.H., 1994. Duration-dependent transitions in a Markov model of U.S. GNP growth. Journal of Business and Economic Statistics 12, 279-288.

Evenson, R.E., Westphal, L.E., 1994. Technological change and technology strategy. In Srinivasan, T.N., Behrman, J. (Eds.), Handbook of Development Economics 3. North Holland: New York.

Falk, B., 1986. Further evidence on the asymmetric behavior of economic time series over the business cycle. Journal of Political Economy 94, 1096-1109.

Gilchrist, S., Williams, J., 2000. Putty-clay and investment: a business cycle analysis. Journal of Political Economy 108, 928-960.

Greenwood, J., Hercowitz, Z., Huffman, G.W., 1988. Investment, capacity utilization and the real business cycle. American Economic Review 78, 402-417.

Hamermesh, D.S., Pfann, G.A., 1996. Adjustment costs in factor demand. Journal of Economic Literature 34, 1264-1292.

Hamilton, J.D., 1989. A new approach to the economic analysis of nonstationary time series and the business cycle. Econometrica 57, 357-384.

Hamilton, J.D., 2005. What's real about the business cycle? Federal Reserve Bank of Saint Louis Review 87, 435-452.

Hansen, G.D., Prescott, E.C., 2005. Capacity constraints, asymmetries, and the business cycle. Review of Economic Dynamics 8, 850-865.

Hansen, G.D., Sargent, T.J., 1988. Straight time and overtime in equilibrium. Journal of Monetary Economics 21, 281-308.

Hess, G.D., Iwata, S., 1997. Asymmetric persistence in GDP? A deeper look at depth. Journal of Monetary Economics 40, 535-554.

Hussey, R., 1992. Nonparametric evidence on asymmetry in business cycles using aggregate employment time series. Journal of Econometrics 51, 217–231.

Jovanovic, B., 2006. Asymmetric cycles. Review of Economic Studies 73, 145-162.

Kocherlakota, N., 2000. Creating business cycles through credit constraints. Federal Reserve Bank of Minneapolis Quarterly Review, 2-10.

Koenders, K., Rogerson, R., 2005. Organizational dynamics over the business cycle: a view on jobless recoveries. Federal Reserve Bank of Saint Louis Review 87, 555-587.

Koop, G., Potter, S.M., 1999. Dynamic asymmetries in U.S. unemployment. Journal of Business and Economic Statistics 17, 298-312.

Lam, P.-S., 2004. A Markov-switching model of GNP growth with duration dependence. International Economic Review 45, 175-204.

Mahajan, V., Muller, E., Wind, Y., 2000. New-Product Diffusion Models. Springer.

McCulloch, J.H., 1975. The monte carlo cycle in business activity. Economic Inquiry 13, 303-321.

McKay, A., Reis, R., 2008. Measuring brevity and violence: a comparison of approaches. Unpublished manuscript.

McQueen, G., Thorley, S., 1993. Asymmetric business cycle turning points. Journal of Monetary Economics 31, 341-362.

Michelacci, Claudio and David Lopez-Salido (2007) "Technology shocks and job flows," *Review of Economic Studies*, forthcoming.

Millard, S., Scott, A., Sensier, M., 1997. The labour market over the business cycle: can theory fit the facts? Oxford Review of Economic Policy 13, 70-92.

Mitchell, W.C., 1913. Business Cycles. University of California Press, Berkeley.

Mitchell, W.C., 1927. Business Cycles: The Problem and Its Setting. National Bureau of Economic Research, New York.

Montgomery, M. L., Zarnowitz, V., Tsay, R., Tiao, C., 1998. Forecasting the U.S. unemployment rate. Journal of the American Statistical Association 93, 478-493.

Mortensen, D.T., Pissarides, C.A., 1994. Job creation and job destruction in the theory of unemployment. Review of Economic Studies 61, 397-415.

Neftci, S.N., 1984. Are economic time series asymmetric over the business cycle? Journal of Political Economy 92, 307-328.

Philippon, T., 2006. Corporate governance over the business cycle. Journal of Economic Dynamics and Control 30, 2117-2141.

Pries, M.J., 2004. Persistence of employment fluctuations: a model of recurring job loss. Review of Economic Studies 71, 193-215.

Ramsey, J.B., Rothman, P., 1995. Time irreversibility and business cycle asymmetry. Journal of Money, Credit and Banking 28, 1-21.

Rogers, E.M., 1995. Diffusion of Innovations. The Free Press, New York.

Rotemberg, J.J., 1999. A heuristic method for extracting smooth trends from economic time series. NBER Working Paper No. 7439.

Rothman, P., 1991. Further evidence on the asymmetric behavior of unemployment rates over the business cycle. Journal of Macroeconomics 13, 291-298.

Rothman, P., 1998. Forecasting asymmetric unemployment rates. Review of Economics and Statistics 80, 164-168.

Sichel, D.E., 1989. Are business cycles asymmetric? A correction. Journal of Political Economy 97, 1255-1260.

Sichel, D.E., 1991. Business cycle duration dependence: a parametric approach. Review of Economics and Statistics 73, 254-260.

Sichel, D.E., 1993. Business cycle asymmetry: a deeper look. Economic Inquiry 31, 224-236.

Stock, J.H., Watson, M.W., 1999. Business cycles. In Woodford, M., Taylor, J. (Eds.), Handbook of Macroeconomics 1, Elsevier, Amsterdam.

Van Nieuwerburgh, S., Veldkamp, L., 2006. Learning asymmetries in real business cycles. Journal of Monetary Economics 53, 753-772.

Van Rens, T., 2004. Organizational capital and employment fluctuations. Unpublished manuscript.

Verbrugge, R., 1997. Investigating cyclical asymmetries. Studies in Nonlinear Dynamics & Econometrics 2, 15-22.

Watson, M.W., 1994. Business cycle duration and postwar stabilization of the U.S. economy. American Economic Review 84, 24-46.

Zarnowitz, V., 1992. Business Cycles: Theory, History, Indicators, and Forecasts, University of Chicago Press, Chicago.

Table 1. Model parameters

| Parameter | | Value | Moment to match | Value |
|---|---|---|---|---|
| Preferences and production | | | | |
| Discount rate | $\rho$ | 0.010 | Average real interest rate | 0.010 |
| Weight on labor | b | 0.675 | Fraction of the day at work | 0.400 |
| Labor share | $\alpha$ | 0.640 | Labor share of income | 0.640 |
| Depreciation rate of capital | $\delta$ | 0.020 | Investment-capital ratio | 0.020 |
| Trend productivity | | | | |
| Long-run trend productivity growth | $\gamma$ | 0.014 | Average output growth rate | 0.014 |
| Mean-reversion of productivity shocks | $\xi$ | 0.993 | Serial correlation of output trend | 0.993 |
| Standard deviation of productivity shocks | $\sigma$ | 0.001 | Variance of output trend | $4\text{x}10^{-6}$ |
| Task-specific productivity | | | | |
| De-trended absolute productivity growth | g | 0.004 | Variance of output | $7\text{x}10^{-4}$ |
| Arrival rate of stage 2 of production | $\mu_1$ | 0.125 | Number of business cycles | 10 |
| Arrival rate of stage 3 of production | $\mu_2$ | 0.087 | No trend in detrended output | |
| Average productivity across tasks, s= 2,3 | $(A_E+1)/2$ | 2.000 | Normalization | |
| Average productivity across tasks, s= 1 | $(A_N+1)/2$ | 1.525 | Standard deviation output/hours | 0.02 |
| Production and employment | | | | |
| Administrative costs per task | z | 0.045 | Average unemployment rate | 0.056 |
| Overhead maintenance costs | $\kappa$ | 0.281 | Variance of hours | 0.001 |
| Employment adjustments | | | | |
| Rate of exogenous separations | $\chi$ | $10^{-5}$ | Variance of change in emp. | $2\text{x}10^{-5}$ |
| Returns to scale in hiring | $\upsilon$ | 0.5 | Quadratic adjustment costs | 2 |
| Units of labor per fire | $\beta$ | 0.05 | Variance of employment | $2\text{x}10^{-4}$ |

Notes: See the text for explanations.

Table 2. Quantitative performance of the model

| Moment | Data | Full model | Frictionless Model |
|---|---|---|---|
| Output | | | |
| Brevity | | | |
| Average duration of expansions | 7.33 | 11.56 | 13.45 |
| Average duration of contractions | 10 | 11.35 | 13.26 |
| Difference | 2.67 | 0.21 | 0.19 |
| Violence | | | |
| Average % growth during expansions | 1.65 | 1.11 | 0.76 |
| Average % growth during contractions | 1.61 | 0.78 | 0.77 |
| Ratio | 1.02 | 1.43 | 0.98 |
| | | | |
| Employment | | | |
| Brevity | | | |
| Average duration of expansions | 18 | 13.84 | 11.90 |
| Average duration of contractions | 8 | 7.82 | 11.89 |
| Difference | 10 | 6.02 | 0.01 |
| Violence | | | |
| Average % growth during expansions | 0.23 | 0.19 | 1.93 |
| Average % growth during contractions | 0.48 | 0.69 | 1.93 |
| Ratio | 0.47 | 0.27 | 1.00 |

Figure 1: Contractions and expansions in the baseline case for output, employment and the NBER

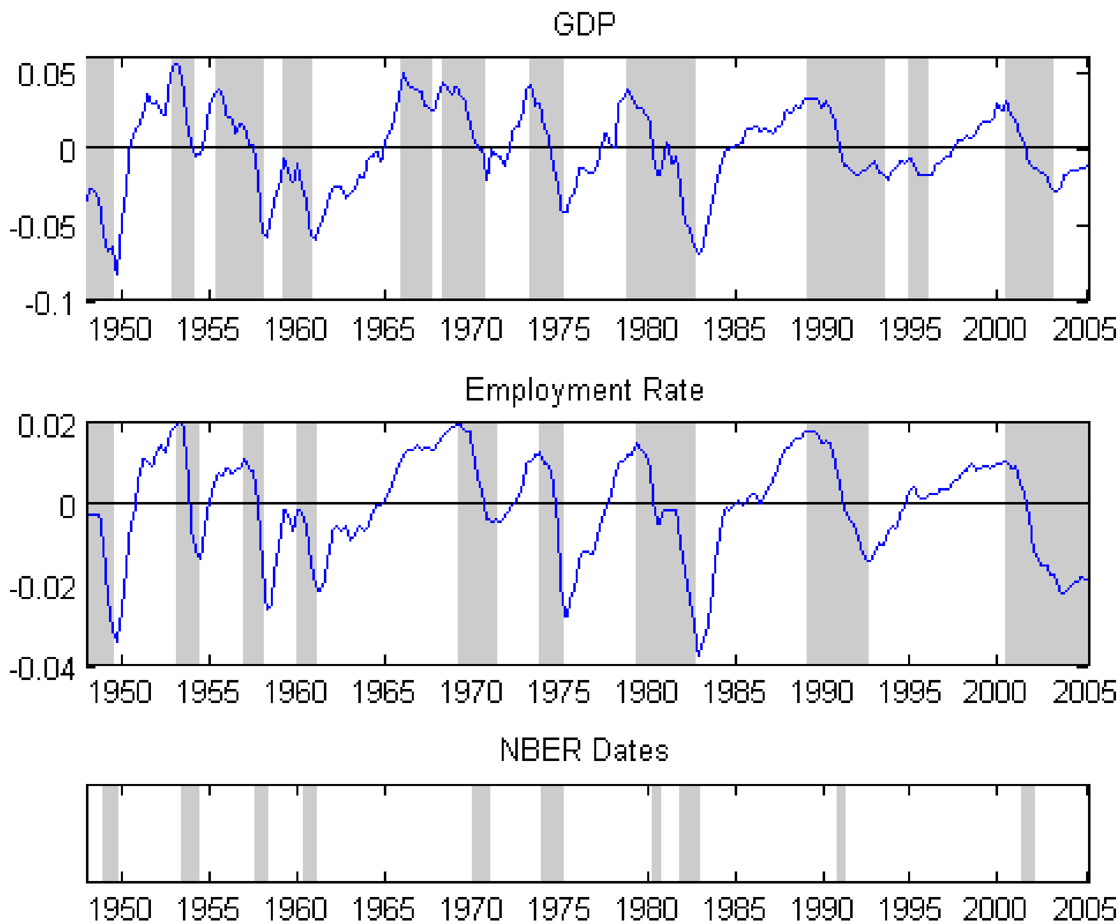

Figure 2: Average business cycle dynamics for output and employment near peaks and troughs in the baseline case

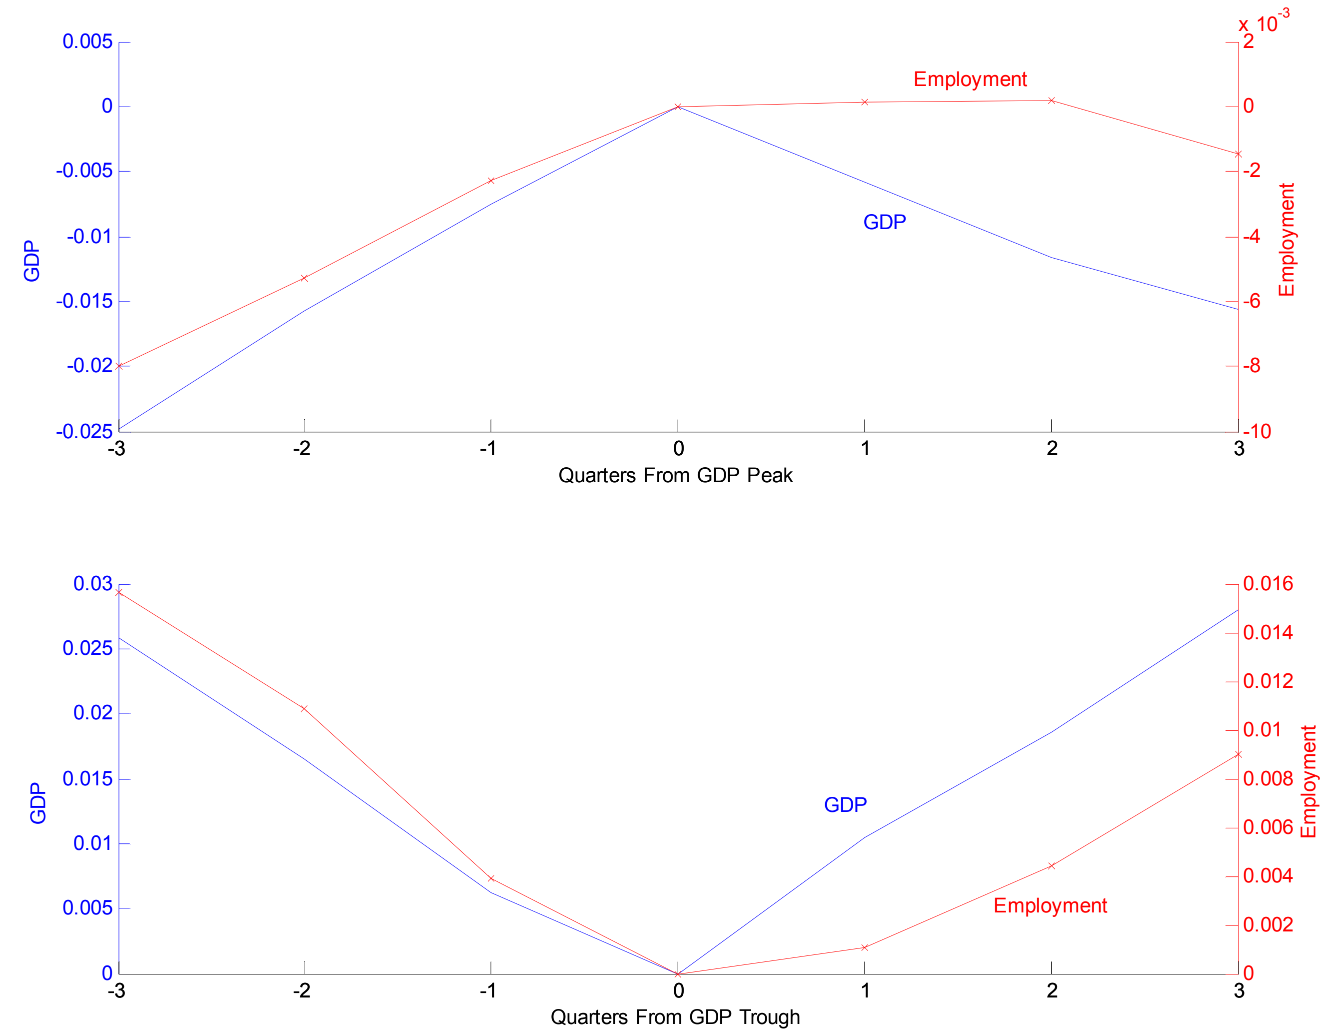

Figure 3: Representative peak-to-peak business cycle dynamics

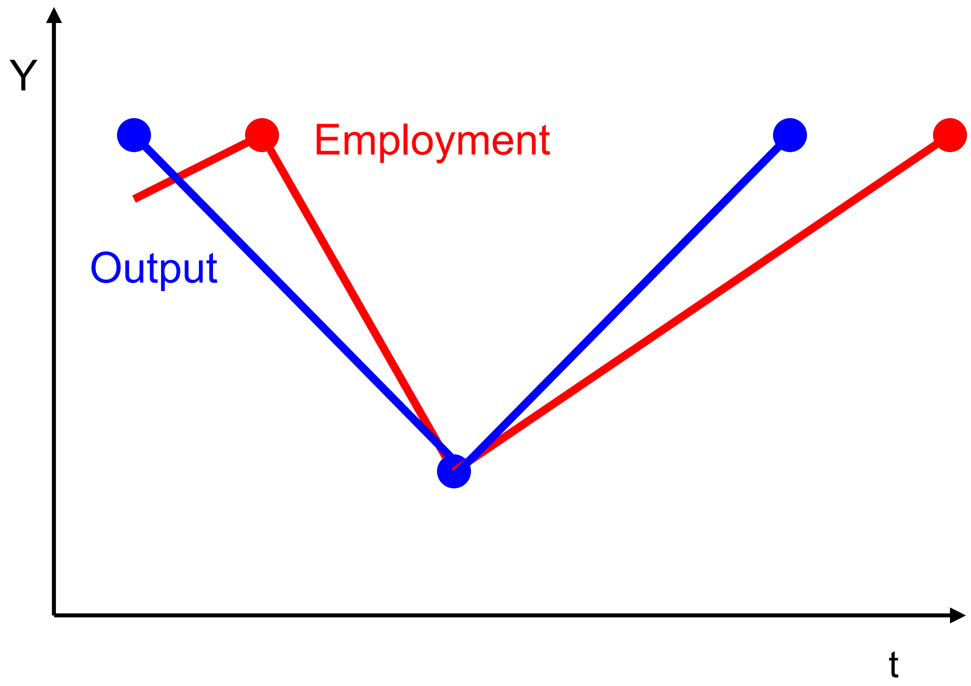

Figure 4. Representative peak-to-trough dynamics

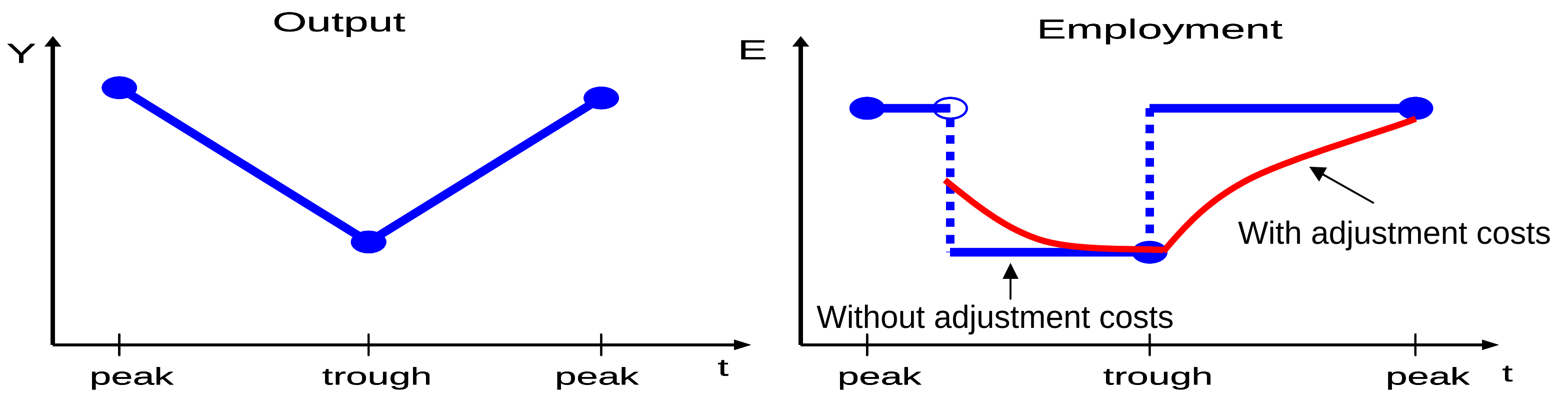

Figure 5: Phase diagram for employment and hiring

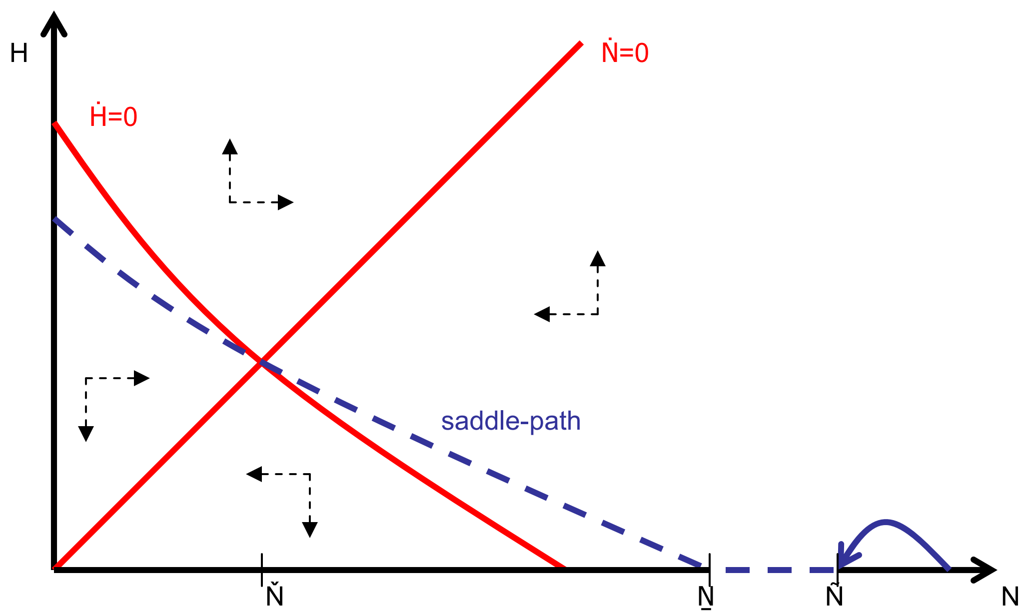

Figure 6: Average business cycle dynamics of hours per worker near the peaks