UvA-DARE (Digital Academic Repository)

Imperfect information, lagged labor adjustment and the Great Moderation

Willems, T.; van Wijnbergen, S.

Imperfect Information, Lagged Labor Adjustment And The Great Moderation

Tim Willems
Sweder van Wijnbergen*

University of Amsterdam.

* Tinbergen Institute

Imperfect information, lagged labor adjustment and the Great Moderation[*]

Tim Willems[†] and Sweder van Wijnbergen
University of Amsterdam and Tinbergen Institute

January 8, 2010

Abstract

This paper first documents the increase in the time lag with which labor input reacts to the economy's driving structural shocks ("the labor adjustment lag") that is visible in US data since the mid-1980s. We show that lagged labor adjustment is optimal in a setting where there is uncertainty about the persistence of shocks and where labor input is costly to adjust. We then present evidence that both the nature of shocks hitting the economy as well as labor adjustment costs have changed since the 1980s in a direction that could explain the increase in the lag. Finally, we argue that the increase in the labor adjustment lag has the potential to explain some macroeconomic puzzles that characterize post-1984 US data, such as the reduced procyclicality of labor productivity and the reduction in output volatility (known as the Great Moderation).

Key words: imperfect information, labor adjustment, jobless growth, option value of waiting, Great Moderation

JEL-classifications: E24, E32, J23, J24

1 Introduction

The last two economic recoveries (the ones that started in 1991 and 2001) were markedly different from their predecessors: contrary to all previous recoveries, they

were not accompanied by a simultaneous increase in employment - very much to the confusion and surprise of most economists. Although we do have a name for this phenomenon (it is generally referred to as a "jobless recovery" or "jobless growth"), there is no agreement yet on its roots.

What is less well-known however, is that besides jobless recoveries, the 1990s and 2000s also brought recessions that were initially relatively job preserving, but had a more prolonged negative effect on labor input. Putting these two observations together, the empirics seem to suggest that the lag with which labor input responds to business cycle fluctuations (henceforth referred to as "the labor adjustment lag") went up in the 1980s. The documentation, causes and consequences of this development lie at the heart of this paper.

To explain the lagged response of labor input to business cycle fluctuations, we develop a theory based on the option value of waiting. As is well-known from the investment literature, there exists an option value to waiting in a stochastic imperfect information setting where the initial costs of investment are at least partially sunk. Although most applications of this idea have so far focused on investment decisions in physical capital (*cf.* the text book treatment of Dixit and Pindyck (1994)), this story similarly applies to labor markets. After all, in the presence of labor adjustment costs, hiring and firing decisions also become costly to undo. We show that it is optimal for employers to respond sluggishly to driving shocks in this environment as the positive option value to waiting induces them to gather more information on the nature of the shock (persistent or transitory?), before making any investment decisions with some element of irreversibility in them. It is shown that, in such a setting, the labor input cycle will lag the business cycle.

Uncertainty about the persistence of economic shocks plays an important role in reality. There, it is always trivial to identify the nature of shocks in hindsight, but when they occur, persistent and transitory shocks have proved to be hard to distinguish from each other. This is for example nicely illustrated by the following quote from Jean-Philippe Cotis, former chief economist of the OECD, on the recovery of the 2001 recession. In December 2002 he wrote in the *OECD Observer*:

> "The world recovery appears more hesitant and less widespread than expected. Activity bounced back early in 2002 but then lost momentum, in a context of weakening consumer and business confidence. This pattern of fits and starts is not unusual in the initial stages of a recovery but it has been associated with a further deterioration of equity and financial markets."[1]

---

[1] http://www.oecdobserver.org/news/fullstory.php/aid/811/A_hesitant_recovery.html

The truth is however that the NBER Business Cycle Dating Committee announced in July 2003 that the recession had ended in November 2001 already. The adverse shocks that hit the economy during 2002 thus turned out to be of a transitory nature, which was apparently not clear to the OECD on impact.

Similarly, it has also proved to be difficult to recognize persistent negative shocks when they actually occur. Consider the following quote by Federal Reserve Chairman Ben Bernanke from March 2007:

> "Although the turmoil in the subprime mortgage market has created severe financial problems for many individuals and families, the implications of these developments for the housing market as a whole are less clear (...) At this juncture, however, the impact on the broader economy and financial markets of the problems in the subprime market seems likely to be contained."[2]

That the above prediction has not quite materialized goes without saying.

To address this particular type of uncertainty, which is neglected in most existing analyses, we build a simple two-period model with imperfect information in Section 4 of this paper. After working through the appropriate cases in Section 5, Section 6 shows that only the *combination* of informational imperfections and labor adjustment costs leads to a lagged labor input response. In Section 7, we will try to give an answer to the quintessential question why the labor adjustment lag has increased in the 1980s. Subsequently, Section 8 discusses the link between the lengthening of the lag and some macroeconomic puzzles, such as the Great Moderation and the vanishing procyclicality of labor productivity. Finally, Section 9 concludes.

But first, Section 2 illustrates the existence of and increase in the labor adjustment lag since the mid 80s, after which Section 3 discusses the related literature.

2 The labor adjustment lag

2.1 Raw data

As noted in the introduction, basically all recoveries before the 1990/1 recession were accompanied by a simultaneous increase in the number of people employed. This is shown in Figure 1, which illustrates the evolution of employment over the cycle. It displays the employment level several months before and after the start of the recovery ("the trough", indicated by the vertical line), relative to the level

---

[2]http://www.federalreserve.gov/newsevents/testimony/bernanke20070328a.htm

of employment at the trough, for the average pre-1990 recession, as well as for the recessions of 1990/1 and 2001.[3] The difference between the three series is clear: while all recoveries before 1990 were accompanied by an immediate, strong increase in employment, employment showed a lagged response to the last two recoveries, as a result of which they were initially jobless.[4,5]
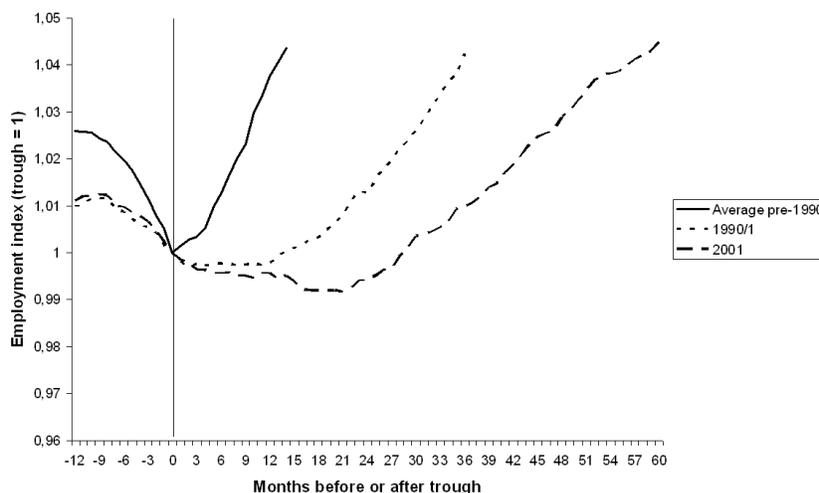
Figure 1: Employment over the cycle

Figure 1 also shows that the employment response to recessions differs between the two sample periods: recessions in the early sample were accompanied by a stronger immediate reduction in employment than the two post-1990 recessions. Moreover, since the 1990s, the trough in employment no longer coincides with the trough in production (which roughly coincides with the NBER-announced end of the recession; see Figure 3 below): after the 1991 trough, employment for example continued to fall for about a year. This is even clearer for the 2001 recession, where employment reached its trough 1.5 years after output did. The incidence of jobless recoveries thus seems to have been accompanied by the introduction of recessions that were *initially*

---

[3]This section uses the NBER postwar peak-to-trough recession dates being: November 1948-October 1949; July 1953-May 1954; August 1957-April 1958; April 1960-February 1961; December 1969-November 1970; November 1973-March 1975; January 1980-July 1980; July 1981-November 1982; July 1990-March 1991; and March 2001-November 2001.

[4]The monthly employment data are taken from the St. Louis Fed database and contain total nonfarm payrolls. The data are only seasonally adjusted. As dispersion around the average pre-1990 line is small, we only display the average.

[5]Note that the trend in employment could play a distorting role here, which will be corrected for in Section 2.2 (as well as in the Appendix).

relatively job preserving, but where employment *continued* to fall for a longer period of time.

It is however important to realize that the lack of an employment response on the extensive margin does not need to have important implications for output: after all, changes in labor input on the extensive margin are to a large extent substitutable with labor input changes on the intensive margin of hours worked. Therefore, we repeat the above analysis using data on total hours worked in the economy.[6] This results in Figure 2, which looks very similar to Figure 1: the trough in total hours worked lags the output trough by at least one year since 1990.[7]
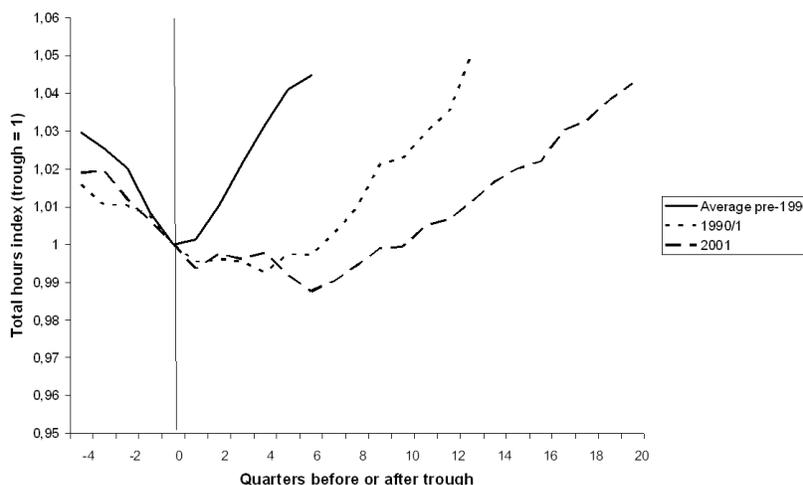
Figure 2: Total hours worked over the cycle

But what have these changes implied for the behavior of GDP over the cycle? Figure 3 gives the answer. As this figure shows, the change in labor input dynamics has had a clear impact on the path followed by total output around recessions: both post-1990 recoveries are much weaker than their predecessors. Given the changed behavior of labor input, this is not that surprising: after all, since the 1980s an autonomous increase in GDP (for example due to a positive shock to total factor productivity) is no longer accompanied by a simultaneous increase in labor input as

---

[6] These BLS-based data are taken from Francis and Ramey (2009). In contrast to standard BLS-data, they contain the number of hours worked in the *total* economy (rather than the number of hours worked in the private business sector). Unfortunately, they are only available at a quarterly frequency. Again, dispersion around the pre-1990 line is small.

[7] Again, the trend in total hours worked could play an undesirable role here which will be addressed in the next section (as well as in the Appendix).

a result of which output (which is a function of labor input, among other things) shows a much weaker response.[8]
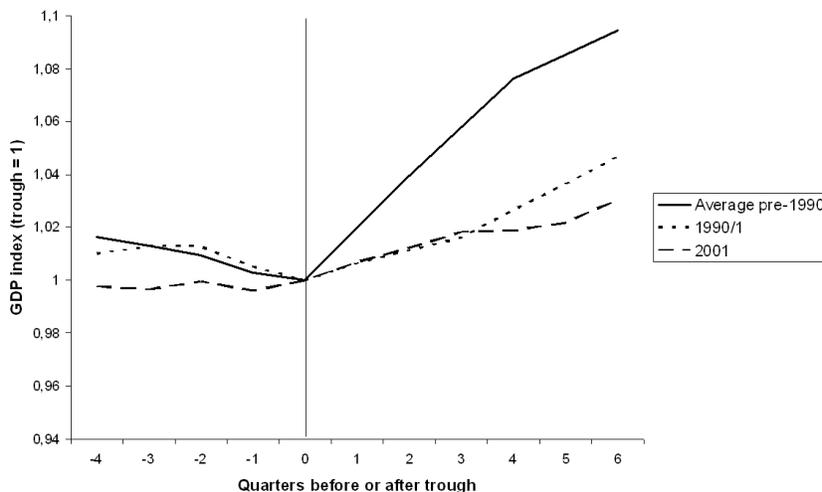
Figure 3: GDP over the cycle

In sum, the raw data seem to suggest that labor input responds with a greater lag to shocks that drive GDP-fluctuations since the 1980s. In the next section, we will investigate this conjecture by looking at the impulse response functions of labor input to the shocks that drive fluctuations in output.

2.2   Labor adjustment in response to structural shocks

Next to looking at the raw data, we can also try to identify the structural shocks that drive GDP fluctuations and analyze the response of labor input to these.

We identify the structural shocks by imposing long run restrictions as in Blanchard and Quah (1989). Following Christiano, Eichenbaum and Vigfusson (2003) we define a vector $Y \equiv (\Delta lp, h)'$, where Δlp is the quarterly change in the natural log of labor productivity (output per hour worked in the non-farm business sector, taken from the BLS) and h is the natural log of the number of hours worked per capita over the age of 16 (again taken from the dataset used in Francis and Ramey

---

[8]Essentially, this is also the explanation we put forward in Section 8.1 for the reduction in output volatility that is visible in US data since the 1980s (often referred to as the "Great Moderation").

(2009)).[9] Moreover, we take the vector $\varepsilon = \left(\varepsilon_t^T, \varepsilon_t^{NT} \right)'$ to describe the two structural disturbances: technology (T) and non-technology (NT) shocks.

The vector Y can then be expressed as a distributed lag of both types of structural shocks. That is:

$$\begin{bmatrix} \Delta lp_t \\ h_t \end{bmatrix} = \begin{bmatrix} A_{11}(L) & A_{12}(L) \\ A_{21}(L) & A_{22}(L) \end{bmatrix} \begin{bmatrix} \varepsilon_t^T \\ \varepsilon_t^{NT} \end{bmatrix}$$
$$\Leftrightarrow Y_t = A(L)\varepsilon_t$$

Here, the assumption that the structural shocks are orthogonal to each other, together with a standard normalization, implies that $\mathbb{E}\varepsilon_t\varepsilon_t' = I$.

Now it is possible to identify both shocks by assuming (as in Galí (1999)) that technology shocks are the only ones to affect the level of labor productivity in the long run. This restriction requires the matrix $A(1)$ to be lower triangular (*i.e.* $A_{12}(1) = 0$; see Blanchard and Quah (1989) for exact details on the identification process).

Having identified the shocks in this way, the second row of the A-matrix gives us the impulse response functions (IRFs) of labor input to both technology and non-technology shocks. Figures 4 and 5 display these.[10]

As can be seen from Figure 4, the persistence in the labor input response to a standardized pre-1984 technology shock went up over time: although the response in the first six quarters has not changed much over the years, the post-1989 IRF shows that since then, a bigger share of total labor adjustment is being made at longer lags.

---

[9]Note that this specification differs from the one used in the classical study by Galí (1999). As he could not reject the null hypothesis of a unit root in total hours worked in the economy (using an augmented Dickey Fuller (ADF) test), he stationarized the series by applying the first difference filter to the natural log of total hours worked. It is however well known that ADF tests have a low power in rejecting the null hypothesis. Therefore Christiano *et al.* (2003) employ a version of Hansen's covariate augmented Dickey Fuller test, which has a much higher power. Using this test, they find strong evidence against the null hypothesis of a unit root in the natural log of hours worked per capita.

The consequences of choosing either specification are far from innocent, for other reasons than of interest for the issue under investigation in the current paper: using Galí's (1999) specification, one finds that hours worked *fall* after a technology shock, while the level-specification used in Christiano *et al.* (2003) results in a *positive* response of hours worked to a technology shock. However, as will be argued in the conclusion, whether hours worked rise or fall after a technology shock is not essential for our theory.

[10]Both figures are based upon a $VAR(4)$, where the lag length was selected via Akaike's information criterion. The underlying data run from 1948Q1 to 2007Q4. The break points are chosen such that the recovery from the 1982 recession still is included in the first sample, while the recessions to which the recoveries were jobless are included in the second sample. Both graphs are robust to picking any other break point in the 1982-1989 interval.
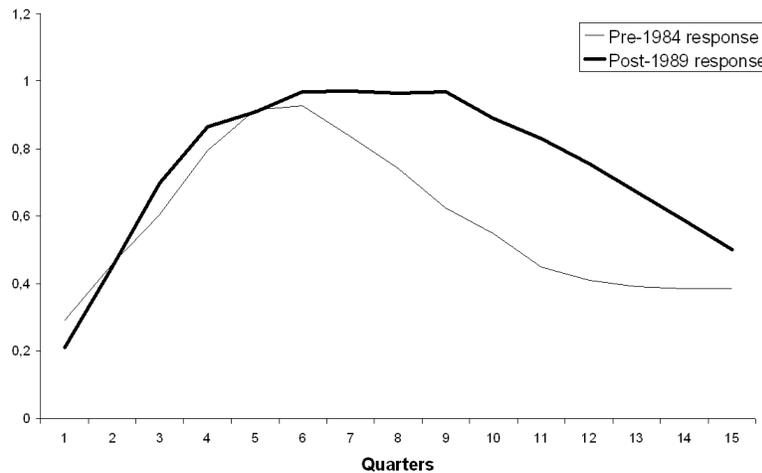
Figure 4: Impulse response functions of total hours to a technology shock

Figure 5 shows the IRFs to a standardized pre-1984 non-technology shock. Here, the response on impact has fallen from 0.73 to 0.53, while the post-1984 response is also much more muted and lagged: whereas the pre-1984 IRF peaks after three quarters already, the post-1989 IRF needs seven quarters to reach its maximum.
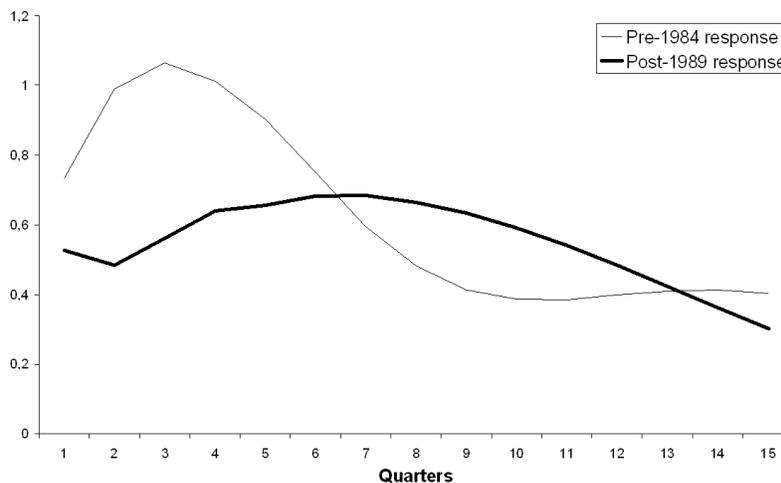
Figure 5: Impulse response functions of total hours to a non-technology shock

2.3   Summing up

In this section we have presented evidence that the lag with which labor input is adjusted in response to structural economic shocks went up somewhere in between

the 1982 recovery and the one starting in 1991. The lag is present at the extensive margin and robust to including the intensive margin. We therefore conclude that the increase in this lag is not the result of an increased reliance of employers to adjust labor input on the hours margin. This confirms the finding of several other studies, namely that the intensive margin is less important than the extensive one for changing labor input in the US.[11] Consequently, we will abstract from the hours-margin and focus at the extensive margin instead in the remainder of the paper.

3   Related literature

The wide attention that has been given to the labor adjustment lag (and jobless growth in particular) in newspaper articles is inversely related to the amount of attention it has received in academia. There, only a few papers have dealt with this issue - all of them particularly focusing at jobless recoveries. First of all, Groshen and Potter (2003) have pointed at the role played by sectoral reallocation. They explain the occurrence of jobless recoveries by hypothesizing that the changing economic environment over the last decades has forced fired workers to undergo time consuming career changes. The problem with their explanation is however that sectoral reallocation has actually *declined* over the last two recessions compared to all previous ones (Aaronson, Rissman and Sullivan, 2004), as a result of which this theory predicts the exact opposite of what has been observed in practice.

Second, Andolfatto and MacDonald (2004) attribute jobless growth to the fact that a new technology impacts different sectors of the economy unevenly and is slow to diffuse. However, as there is no evidence that the impact of innovations suddenly changed somewhere between the last "job driven" recovery in 1982 and the first jobless-one in 1991, their model cannot explain why there was no jobless growth before the 1991-recession.

Bachmann (2009) has constructed a DSGE-model in which employers can generate a jobless recovery by adjusting their labor input along the intensive (rather than on the extensive) margin. However, Figure 2 in this paper shows that this does not explain the recent two jobless recoveries as they do not show an increase in the total number of hours worked either.

Finally, Koenders and Rogerson (2005) relate the occurrence of jobless growth to organizational issues within firms. They start from the observation that the two most recent recessions are special in that they both followed upon unusually

---

[11]See for example Gertler, Sala and Trigari (2008) who estimated a Frisch-elasticity of zero for the intensive margin.

long expansions. Assuming that organizational inefficiencies accumulate during these expansions, Koenders and Rogerson (2005) argue that the two most recent recessions were used by firms to eliminate the accumulated inefficiencies by reducing their levels of employment. But because reorganizations had been postponed for so long, the length of the last two recessions was not sufficient to remove all inefficiencies - thereby leading to jobless growth. The problem with this explanation is that it also predicts that the last two recessions should have been accompanied by relatively large reductions in levels of employment (see Koenders and Rogerson (2005), Figure 1). However, our Figure 1 shows that the exact opposite is the case.

To overcome the difficulties with existing theories, we construct a model based upon imperfect information to explain the stylized facts. The key of our story is that there exists an option value to waiting for employers when employers are uncertain about the persistence of shocks. Combining this with labor adjustment costs, results in a labor adjustment lag that materializes in the form of jobless recoveries and recessions that are initially relatively job preserving, as the data suggest.

We view our explanation as complementary to the one offered by Van Rens (2004). As Koenders and Rogerson (2005), he also stresses the importance of organizational issues within firms, but assumes that reorganization requires the *input* of labor rather than the *shedding* of labor as in Koenders and Rogerson (2005). In Van Rens (2004), production requires the input of both productive activities (say working on the assembly line) as well as organizational/supportive ones (cleaning up the factory hall). By allowing firms to relocate workers from organizational tasks to productive activities, Van Rens adds an extra margin along which firms can temporarily increase production at the beginning of a recovery without hiring extra workers. In his model, the postponement of labor adjustment is optimal as firms want to minimize the present value of labor adjustment costs. However, after a while, the depletion of the stock of organizational capital starts to harm production (the factory hall has become a mess) as a result of which labor is shifted back to organizational activities and additional hiring is necessary after all. Different from our explanation, his does not include the option value of waiting and even holds in a completely deterministic setting.

A contribution conceptually related to the current one is that of Bentolila and Bertola (1990). In a stochastic continuous-time setting, they note the existence of another option value for firms in the presence of labor adjustment costs: in their model, hiring implies that a firm gives up the call option to delay the hiring decision, while it acquires the put option to fire the new employee in the future. In that setting, they show that labor adjustment costs lower the volatility of the level of employment - a result that is contained by our model (see Section 6). Contrary to the current paper, Bentolila and Bertola (1990) however do not touch upon the role

played by imperfect information as a result of which they do not establish the link with lagged labor adjustment.

4 Model

As in Van Wijnbergen (1985, who applies the idea of the option value of waiting to trade reform issues), we analyze the problem in a simple two period setting. The structure of the model is as follows: the model is populated by a large pool of potential employees (sufficiently large to prevent the stock of employees from becoming a binding constraint) along with a continuum of small firms (each too small to affect prices in the economy). The wage rate is determined exogenously.[12]

At the beginning of period 1, the economy is hit by a positive productivity shock.[13] In particular, the level of productivity unexpectedly rises from $A_0 = A$ to $A_1 = A + \xi$.[14] *After* the occurrence of the shock, each firm gets to set its level of employment N_1 with which it enters the production stage. It produces according to the following production function:

$$Y_t = A_t N_t^\alpha$$

Here $\alpha \in (0,1)$ captures the notion of decreasing returns to labor. For convenience, we abstract from capital accumulation in this short-run analysis. Output is sold at a competitive goods market at a fixed price, normalized to 1.

When the productivity shock ξ materializes in period 1, it is however not immediately clear whether this shock is permanent or transitory. This will only become clear at the beginning of period 2 when A_2 materializes. In particular, with exogenous probability θ (which is known to the rational agents in the model) the shock is fully persistent which implies that $A_2 = A + \xi$. However, there is a probability $(1 - \theta)$ that the shock was purely transitory in which case $A_2 = A$ again. In this way, we model the fact that agents are uncertain about the duration of booms and recessions

---

[12]This assumption could easily be relaxed without affecting our results. The only condition that is important is that the uncertainty about the nature of shocks is not fully shifted to the workers via the wage rate, as employers then do not care about it anymore. Nash bargaining, widely used for the determination of wages in labor market theories (see *e.g.* Pissarides (2000, chapter 1)), for example satisfies this condition.

[13]Note that we follow the RBC-literature here in assuming that business cycles are driven by productivity shocks. The main results however also hold in a demand shock driven model as a result of which the answer to the question whether business cycles are driven by supply or demand shocks is only of minor importance here.

[14]Every statement in the paper applies *mutatis mutandis* to negative productivity shocks.

in reality; only by observing next period's state, they obtain more information about the persistence of a shock.

Graphically, the shock structure of the model can be depicted as follows:
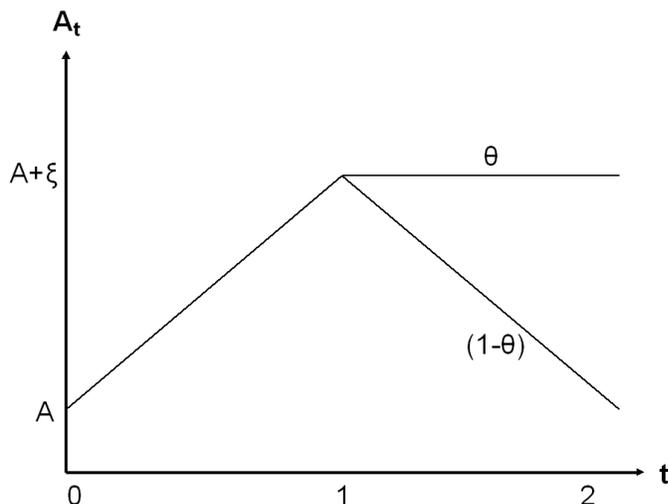
Figure 6: Shock structure of the model

The employer discounts the future at factor $\beta \in [0,1)$. The firm's objective function then reads:

$$\max_{N_1,N_2} \ A_1 N_1^\alpha - wN_1 + \beta\mathbb{E}_1 \left\{ A_2 N_2^\alpha - wN_2 \right\},$$

where $\mathbb{E}_1$ represents the expectations operator that conditions on all information available at $t = 1$.

From here on, we will add the informational imperfection and costly labor adjustment. To identify the separate contributions of each of these elements, we develop a very simple and flexible framework which allows us to add the two frictions one-by-one. This leads to four cases that are to be considered (see Table 1). First, we will analyze the benchmark case in which there are no labor adjustment costs and where the employer has full information. In this setting, the firm can costlessly reset its employment level at the beginning of period 2, after the nature of the shock has been revealed. Subsequently, we will introduce a role for the informational imperfection by considering the case in which the employer has to commit. In this setting, he must set both his first- and second-period level of employment in period 1 already. Hence, at the time of his decision, he does not know yet whether the shock is persistent or

12

not. Next, we will add labor adjustment costs. Motivated by the available empirical evidence (see *e.g.* Hamermesh and Pfann (1996) and Bloom (2009)), we model labor adjustment costs in a non-convex way. In particular, we consider the limiting case in which the employer loses the opportunity to adjust his labor input at the beginning of the second period. Effectively, this represents infinite, non-convex labor adjustment costs. Finally, we will also consider the remaining case in which there are infinite, non-convex labor adjustment costs, but where employers are perfectly able to distinguish persistent shocks from transitory ones.

	Perfect information	Imperfect information
No adjustment costs	Case I	Case II
Infinite adjustment costs	Case IV	Case III

Table 1: The cases

5 The cases

5.1 Case I: flexibility, no labor adjustment costs

Suppose that there are no labor adjustment costs and that the employer has full flexibility in setting his desired level of employment each period. The employer then sets his period 1 employment level directly after the occurrence of the shock. When doing this, he already knows that he will get the opportunity to reset his employment level costlessly at the beginning of period two, after the nature of the shock has been revealed. Consequently, the informational imperfection is not binding and the employer solves two independent static optimization problems. In this setting, his period 2 problem reads:

$$\max_{N_2} \ A_2 N_2^\alpha - w N_2$$

The first-order condition implies that, conditional upon observing A_2:

$$N_2^f(A_2) = \left[\frac{\alpha A_2}{w} \right]^{\frac{1}{1-\alpha}}, \tag{1}$$

where the superscript f indicates that this is the solution under flexibility.

Given the employer's optimal response in period 2 (equation (1)), his period 1

problem reads:

$$\max_{N_1} \ (A+\xi) N_1^\alpha - wN_1 + \beta \left[\theta \left(A+\xi\right) \left[\frac{\alpha \left(A+\xi\right)}{w}\right]^{\frac{\alpha}{1-\alpha}} + \left(1-\theta\right) A \left[\frac{\alpha A}{w}\right]^{\frac{\alpha}{1-\alpha}} \right]$$

$$- \beta \left[\theta w \left[\frac{\alpha \left(A+\xi\right)}{w}\right]^{\frac{1}{1-\alpha}} + \left(1-\theta\right) w \left[\frac{\alpha A}{w}\right]^{\frac{1}{1-\alpha}} \right],$$

which leads to the following first-order condition:

$$N_1^f = \left[\frac{\alpha(A+\xi)}{w}\right]^{\frac{1}{1-\alpha}} \tag{2}$$

Note that his period 1 hiring decision is independent of θ, the probability that the shock survives into the second period. This is due to the fact that the employer can costlessly reset his employment level after the revelation of the nature of the shock. Consequently, he does not have to take the persistence issue into account in period 1 already.

5.2 Case II: commitment, no labor adjustment costs

We now introduce a role for the informational imperfection by assuming that the employer must set his period 2 level of employment in period 1 already. Consequently, when the employer picks his level of employment for periods 1 and 2, he does not know yet whether the productivity shock ξ is going to be persistent or not. So in this setting, the informational imperfection binds as N_2 is a state variable at the beginning of period 2 and cannot be changed in response to the revelation of the nature of the shock at the beginning of this period.

The employer's problem then reads:

$$\max_{N_1, N_2} \ [A+\xi] N_1^\alpha - wN_1 + \beta \mathbb{E}_1 \left\{ A_2 N_2^\alpha - wN_2 \right\}$$

As $\mathbb{E}_1 \{A_2\} = A + \theta\xi$ under our rational expectations assumption, we can rewrite this problem as:

$$\max_{N_1, N_2} \ [A+\xi] N_1^\alpha - wN_1 + \beta \left[(A+\theta\xi) N_2^\alpha - wN_2\right]$$

The associated first-order conditions imply that the employer chooses to set the following levels of employment in period 1:

$$N_1^c = \left[\frac{\alpha (A + \xi)}{w} \right]^{\frac{1}{1-\alpha}} \tag{3}$$

$$N_2^c = \left[\frac{\alpha (A + \theta\xi)}{w} \right]^{\frac{1}{1-\alpha}}, \tag{4}$$

where the superscript c refers to the fact that this is the solution under commitment. Note that labor adjustment is still costless in this case, as the employer is free to set N_2 different from N_1.

5.3 Case III: infinite labor adjustment costs under imperfect information

Now we add labor adjustment costs to the model. In particular, we assume that there only exist two-period employment contracts and that additional hiring at the beginning of period 2 is not possible. Effectively this represents the limiting case of infinite, non-convex labor adjustment costs over the time period to which the informational imperfection applies (*i.e.* over periods 1 and 2). The employer is now no longer able to change his labor input for the second period as a result of which $N_2 = N_1$. His problem then reads:

$$\max_{N_1} \ (A + \xi) N_1^\alpha - wN_1 + \beta\mathbb{E}_1 \left\{ A_2 N_1^\alpha - wN_1 \right\}$$

As $\mathbb{E}_1 \{A_2\} = A + \theta\xi$, we can rewrite this problem as:

$$\max_{N_1} \ (A + \xi) N_1^\alpha - (1 + \beta)wN_1 + \beta (A + \theta\xi) N_1^\alpha$$

The associated first-order condition implies that the employer sets his level of employment for both periods equal to:

$$N_1^{\infty,im} = N_2^{\infty,im} = \left[\frac{\alpha \left[(1 + \beta)A + (1 + \beta\theta)\xi \right]}{(1 + \beta)w} \right]^{\frac{1}{1-\alpha}}, \tag{5}$$

where the superscript ∞, im refers to the fact that this is the solution under infinite labor adjustment costs and imperfect information. Note that first-period hiring is increasing in the probability that the positive productivity shock is persistent, θ: after all, if it is highly likely that next period's productivity level will be high again,

the employer decides to guard himself against this by hiring more employees in period 1 already as he cannot hire additional ones at the beginning of period 2.

5.4 Case IV: infinite labor adjustment costs under perfect information

Finally we also consider the remaining possible case, namely the case in which firms do have perfect information but face infinite labor adjustment costs. Here, we have to distinguish between two subcases: one in which employers *know* that the shock is fully persistent, and one in which employers *know* that the shock is purely transitory.[15]

Subcase 1: *Persistent shock*

In this case the problem reads:

$$\max_{N_1} \ (1+\beta)\,(A+\xi)\,N_1^\alpha - (1+\beta)wN_1$$

The first-order condition then implies:

$$N_1^{\infty,p} = N_2^{\infty,p} = \left[\frac{\alpha\,(A+\xi)}{w}\right]^{\frac{1}{1-\alpha}}, \tag{6}$$

where the superscript ∞, p indicates that this is the level of employment set under infinite labor adjustment costs in response to a fully persistent shock.

Subcase 2: *Transitory shock*

In this case the problem reads:

$$\max_{N_1} \ [(1+\beta)A + \xi]\,N_1^\alpha - (1+\beta)wN_1$$

The first-order condition implies that the associated level of employment equals:

$$N_1^{\infty,t} = N_2^{\infty,t} = \left[\frac{\alpha\,[(1+\beta)A + \xi]}{(1+\beta)\,w}\right]^{\frac{1}{1-\alpha}}, \tag{7}$$

where the superscript ∞, t refers to the fact that this is the solution under infinite labor adjustment costs to a purely transitory shock.

Together, equations (6) and (7) show that the employer's behavior under infinite labor adjustment costs in the presence of imperfect information (represented by equation (5)) is basically a weighted average between the corresponding solutions under perfect information. The particular weight placed on the two cases is the probability that the shock is persistent or not, θ.

---

[15]Note that these cases are effectively already covered under Case III by setting $\theta = 0$ or $\theta = 1$.

6 Results

We are now ready to state the main results of this paper.

Proposition 1 (Option value of waiting) *In the presence of informational imperfections (i.e. as long as $\theta \in (0,1)$), there exists a positive option value to having the opportunity to await the nature of the productivity shock.*

Proof. The option value of waiting can be obtained by subtracting the expected value under commitment from the expected value under flexibility. The expected value for the employer under flexibility reads:

$$
\begin{aligned}
V(N^f) &= (1+\beta\theta)(A+\xi)\left[\frac{\alpha(A+\xi)}{w}\right]^{\frac{\alpha}{1-\alpha}} - (1+\beta\theta)w\left[\frac{\alpha(A+\xi)}{w}\right]^{\frac{1}{1-\alpha}} \\
&\quad + \beta(1-\theta)\left[A\left[\frac{\alpha A}{w}\right]^{\frac{\alpha}{1-\alpha}} - w\left[\frac{\alpha A}{w}\right]^{\frac{1}{1-\alpha}}\right]
\end{aligned}
\tag{8}
$$

Similarly, the expected value for the employer under commitment reads:

$$
\begin{aligned}
V(N^c) &= (A+\xi)\left[\frac{\alpha(A+\xi)}{w}\right]^{\frac{\alpha}{1-\alpha}} - w\left[\frac{\alpha(A+\xi)}{w}\right]^{\frac{1}{1-\alpha}} \\
&\quad + \beta\left[(A+\theta\xi)\left[\frac{\alpha(A+\theta\xi)}{w}\right]^{\frac{\alpha}{1-\alpha}} - w\left[\frac{\alpha(A+\theta\xi)}{w}\right]^{\frac{1}{1-\alpha}}\right]
\end{aligned}
\tag{9}
$$

Subtracting (9) from (8) results in the following expression, representing the option value of waiting:

$$
\begin{aligned}
\Omega &\equiv V(N^f) - V(N^c) \\
&= (1+\beta\theta)(A+\xi)\left[\frac{\alpha(A+\xi)}{w}\right]^{\frac{\alpha}{1-\alpha}} - (1+\beta\theta)w\left[\frac{\alpha(A+\xi)}{w}\right]^{\frac{1}{1-\alpha}} \\
&\quad + \beta(1-\theta)\left[A\left(\frac{\alpha A}{w}\right)^{\frac{\alpha}{1-\alpha}} - w\left(\frac{\alpha A}{w}\right)^{\frac{1}{1-\alpha}}\right] \\
&\quad - (A+\xi)\left[\frac{\alpha(A+\xi)}{w}\right]^{\frac{\alpha}{1-\alpha}} + w\left[\frac{\alpha(A+\xi)}{w}\right]^{\frac{1}{1-\alpha}} \\
&\quad - \beta\left[(A+\theta\xi)\left[\frac{\alpha(A+\theta\xi)}{w}\right]^{\frac{\alpha}{1-\alpha}} - w\left[\frac{\alpha(A+\theta\xi)}{w}\right]^{\frac{1}{1-\alpha}}\right]
\end{aligned}
$$

17

The equation for the option value has its roots at $\theta = 0$ and $\theta = 1$ (which is intuitive as there is no uncertainty at the endpoints). As the second derivative of Ω with respect to θ is negative due to the concavity of the production function ($\alpha < 1$), the option value is strictly positive over the domain $\theta \in (0, 1)$ by Rolle's theorem. ∎

As with all options, the value of this one reaches its maximum for the value of θ at which the variance of the underlying process is maximized. Since we assumed a Bernoulli distribution, which has variance $\theta (1 - \theta)$, the option value reaches its maximum at $\theta = \frac{1}{2}$: at that value, the uncertainty about the persistence of the shock is maximized, and so is therefore the value of being able to wait.

Figure 7 shows the option value for different values of θ and α. Note that the option value also increases in α since a higher α implies a higher marginal product of labor and therefore a higher cost of just guessing wrong.
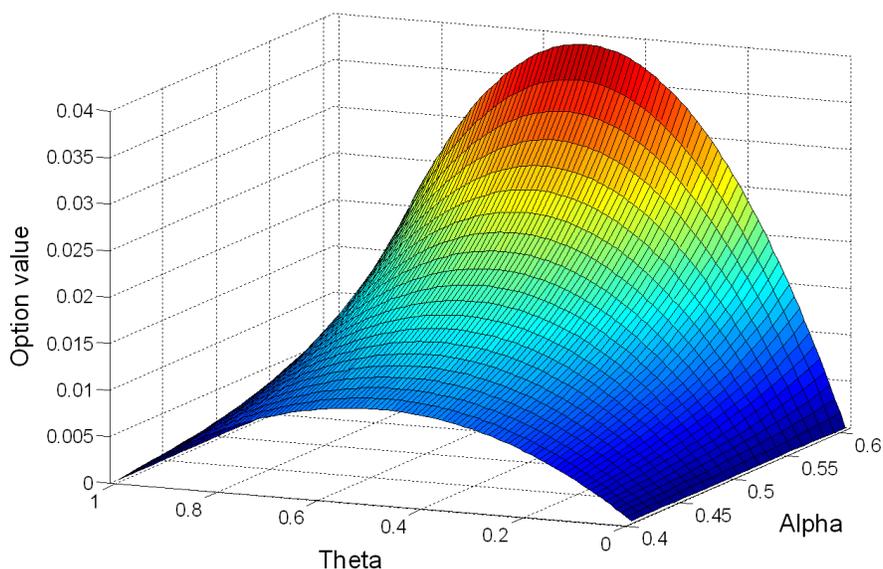
Figure 7: Option value of waiting for the employer in the (α, θ)-space

Proposition 2 (Labor adjustment lag) *The combination of informational imperfections with non-convex labor adjustment costs leads to a lagged response of labor input to productivity shocks.*

Proof. Compare the firm's first-period hiring decision under infinite labor adjustment costs and imperfect information (equation (5)), with the firm's first-period hiring decision under perfect information without labor adjustment costs (equation

18

(2)). As $\theta < 1 \Rightarrow N_1^{\infty,im} < N_1^f$ in case of a positive productivity shock. Equivalently, in the converse case of a negative productivity shock $\theta < 1$ ensures that $N_1^{\infty,im} > N_1^f$. However, as soon as it becomes clear that the shock is fully persistent, we are in the full information case (Case IV) and when employers get the opportunity to reset their labor input *with* knowledge of the nature of the shock we end up at the fully flexible solution (compare equations (1) and (6)). Repeating this result (*i.e.* pasting it behind itself iteratively to mimic an n-period setting) shows that the labor input cycle lags the cycle for the technology shock. ∎

Figure 8 gives a graphical illustration of the proof. In the figure, a dashed line indicates the time period during which the nature of the shock is not known yet. Hence, the first positive productivity shock hits at (a), but the fact that the shock is permanent only becomes clear at (b). Consequently, employers set their level of employment at point (a) without knowing whether the shock is persistent or not, while they do have this information at point (b). As can be seen from the figure, the model equations imply that labor input does not adjust immediately to the optimal level as long as the nature of the shock is unknown (see points (a), (c) and (e) - consistent with equation (5)). However, as soon as the nature of the shock is revealed, labor input is set at its optimal level (points (b) and (d) - consistent with equation (6)).
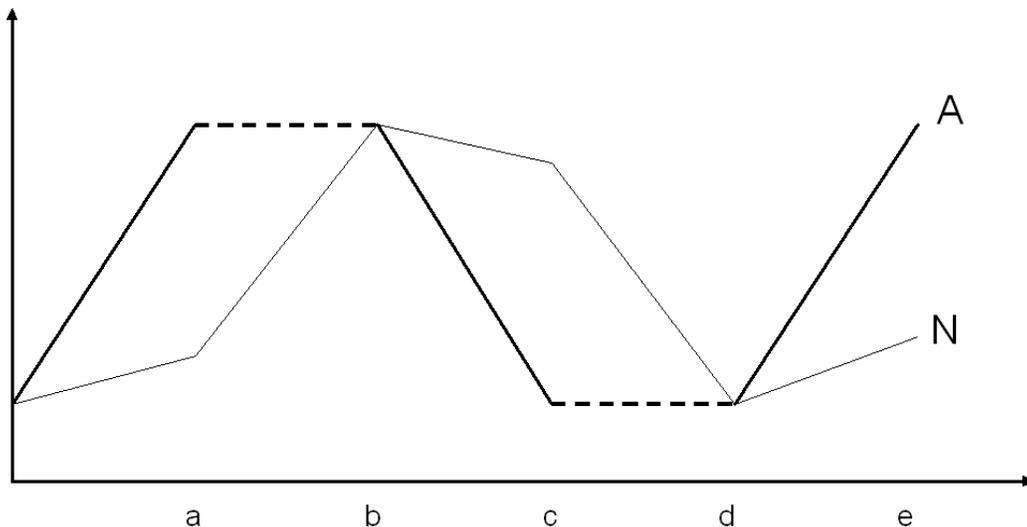

Figure 8: Lagged labor input response under imperfect information and labor adjustment costs

Note by comparing equations (2) and (3) that the period 1 solution under commitment is equal to the solution under full flexibility. Hence, *the informational im-*

perfection alone is not enough to generate a labor adjustment lag: for that it needs to be combined with labor adjustment costs. The informational imperfection alone only affects the period 2 decision, to which it applies. The period 1 decision remains unaffected as it can costlessly be undone in the absence of labor adjustment costs. Consequently employers can freely respond to the first-period productivity shock, without worrying about its persistence.

Equivalently, equations (6) and (7) show that *non-convex labor adjustment costs alone aren't sufficient to generate lagged labor adjustment either*. In that case, firms either increase employment all the way (if they realize that a shock is persistent) or, if they know that the shock is only temporary, they set their level of employment somewhere in between the level of employment chosen absent any shocks and the employment level they would have picked after a persistent shock of equal size. This shows that non-convex labor adjustment costs in isolation only lead to a muted labor input response following shocks that do not have a unit root. This result is consistent with the theoretical result obtained in Bentolila and Bertola (1990) and the empirical evidence presented in Bertola (1990).[16]

It is thus *the combination* of informational imperfections and non-convex labor adjustment costs that makes labor input lag the cycle. Unfortunately, only the perfect information environment is studied in most labor market models, as a result of which both the option value of waiting and the labor adjustment lag are basically neglected in current analyses.

7 Why has the labor adjustment lag increased?

Now that we have pointed at the existence of the labor adjustment lag and investigated its potential causes, the key question is: why has this lag increased since the 1980s?

As shown in the previous section, a labor adjustment lag results from two ingredients: labor adjustment costs and imperfect information on the persistence of shocks. To answer the question why the lag went up in the 1980s, it is therefore natural to look for changes in these two ingredients that occurred around that time. In this

---

[16]Do note that convex labor adjustment costs would generate lagged (or maybe it is better to speak of "gradual") labor adjustment - even in a completely deterministic setting. However, as noted before, empirical studies hardly ever find evidence for convexities in labor adjustment costs. Moreover, convex labor adjustment costs would also imply that hiring peaks on impact of the shock, which is again not in line with the empirical evidence (where it is reported that hiring shows a hump-shaped response; *cf.* Ravn and Simonelli (2008) and Fujita (forthcoming)).

section, we consider both of them in turn.[17]

7.1 Labor adjustment costs

As shown in Figure 9, both the job finding and separation process have changed considerably since the 1980s. From this figure, three changes seem noteworthy for our story. First, observe that the separation probability shows a clear downward trend since the early 80s. Second, the separation rate has essentially been acyclical during the last two recessions (Hall, 2005; Shimer, 2005): where all pre-1990 recessions were accompanied by *both* a drop in hirings *and* an upward spike in the separation rate, the latter development is no longer observed in the two post-1990 recessions. Finally, and in line with our empirical analysis presented in Section 2, the job finding probability has started to lag the business cycle since the 1980s: after the end of the recession of 1990/1 (in March 1991), the job finding probability continued to fall for five quarters before reaching its trough. For the 2001 recession this lag even equaled seven quarters.
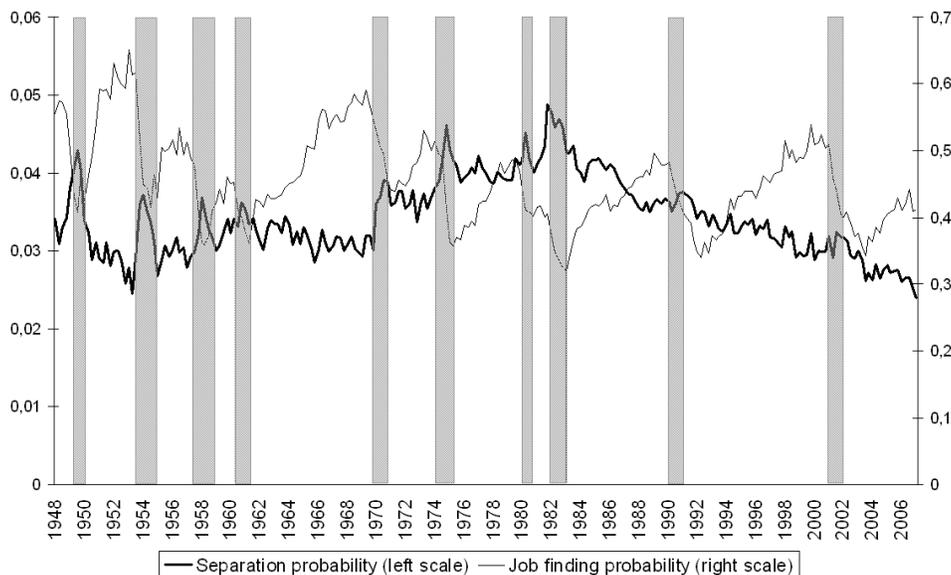

Figure 9: Separation and job finding probabilities (data are from Shimer (2005)). Shaded areas correspond with the NBER recession dates.

---

[17]Alternatively, one could also view this section as a litmus test for our theory, as it should be able to explain why the labor adjustment lag has increased in the 1980s.

Together, these observations hint at an increased reluctance of firms to quickly adjust their labor force since the 1980s, which could be brought about by changes in the labor adjustment costs. The remainder of this section discusses the changes that could underlie this development.

Direct labor adjustment costs   The most obvious candidates in the labor adjustment costs-category are the direct hiring and firing costs or restrictions that firms face when adjusting their labor force. The OECD (1999, Table 2.1) mentions one change in the US over the relevant time period. It became effective in 1989. In that year, the Worker Adjustment and Reassignment Notification Act was installed. This act requires firms with more than 100 workers affected by plant closures or mass layoffs to give a 60 days' notice to all employees involved. Although it is implausible that this act alone is responsible for the increase in the lag with which labor adjustments are made (especially since there have probably also been developments working in the opposite way),[18] it is something that works in the observed direction and hence could have contributed to it.

The data also point at another potentially important development, related to the composition of US unemployment. As shown in Figure 10, all pre-1990 increases in unemployment where at least partially caused by an increase in temporary layoffs. However, since the 1980s, almost all cyclical variation in the temporary layoff rate has vanished and unemployment rises exclusively due to permanent layoffs in recessions.

---

[18]The development of the internet is for example something that is generally believed to have eased the process of hiring new workers for employers. In this light it should be noted however that the available empirical evidence suggests otherwise: using data from www.monster.com, Brencic and Norris (2008) find that, controlling for observables, online search tools do *not* shorten the duration of search for employers.
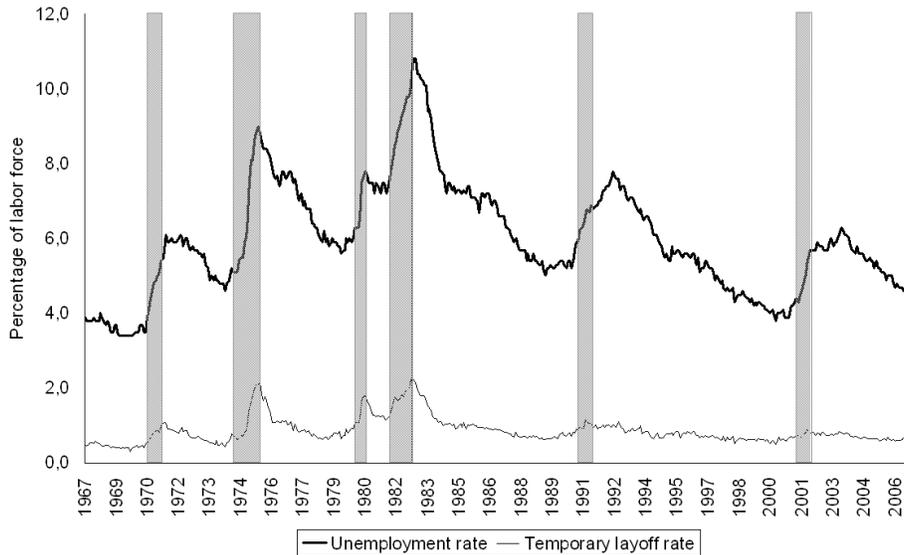
Figure 10: contribution of temporary layoffs to the unemployment rate (data are from the BLS). Shaded areas correspond with the NBER recession dates.

This development could have contributed to the increase in the labor adjustment lag. After all, when employers are able to layoff workers only temporarily, worker-firm matches stay intact and the whole option value of waiting argument applies to a much lesser extent: the firm still knows the worker and can reactivate him by a simple phone call, rather than going through the costly and uncertain process of hiring a new worker. In the early sample, employers thus basically passed the option value of waiting risk on to the US tax payer, as they would just push an employee into the welfare system whenever the economy started to lose momentum, without paying any significant firing taxes.

US policy makers realized this in the early 80s (see the introduction of Burdett and Wright (1989)) and increased the reliance upon "experience rating" in response (Topel, 1990). Under a system of experience rating, the height of the firing tax a firm has to pay to layoff workers at time t depends positively upon the number of workers a firm has laid off in the recent past. Consequently, a system of experience rating strongly reduces firms' incentives to layoff workers temporarily (see Feldstein (1976) for the theoretical argument and Card and Levine (1994) for empirical evidence on this claim). As a result, it became less attractive for US firms to pass the option risk on to the tax payer and the worker-firm matches that were destroyed in the last two recessions, were destroyed permanently. These matches could only be restored by overcoming the option value of waiting-barrier (the region of inaction) during the

23

subsequent upturn - thereby potentially contributing to lagged labor adjustment.

Implicit labor adjustment costs   Besides direct labor adjustment costs, a more subtle notion of labor adjustment costs could have played a role as well. Over the years, the knowledge intensive service sector has become of greater total relevance for the US economy (structural change). Consequently, the - what we will refer to as - *implicit* labor adjustment costs (*i.e.* a firm's intrinsic reluctance to hire or fire workers) may have gone up: after all, in a knowledge intensive industry, firms have invested in the accumulation of firm specific human capital of their employees, which would be lost if the employee is fired. This provides another incentive for firms to hold on to their employees during downturns (especially since the use of temporary layoffs has been discouraged in the 1980s by the increase in experience rating), and could have contributed to the developments displayed in Figure 9.

7.2   Clarity of information

Another reason why the labor adjustment lag could have gone up (a reason that could also explain the aforementioned apparent increased reluctance of US firms to quickly adjust their labor force), is by an increase in the uncertainty about the persistence of shocks. In the language of our two-period model: as the variance of next period's expected productivity level is maximized for $\theta = \frac{1}{2}$ (when all shocks occurring are equally likely to be persistent or transitory through the eyes of the agents), so is the option value of awaiting the nature of the shock (*cf.* Figure 7), as a result of which it becomes more attractive for employers to postpone decisions with some element of irreversibility in it.

There is indeed evidence that this type of uncertainty, which has to be distinguished from volatility, has increased since the 1980s:[19] Campbell (2007) notes, by looking at the Survey of Professional Forecasters (SPF), that macroeconomic predictability has declined significantly since 1984. Similar findings are reported by Schuh (2001, who also uses the SPF), D'Agostino, Giannone and Surico (2006, who look at both the SPF and the Federal Reserve's Greenbook) and Tulip (2005, who uses forecasts from the Greenbook). The latter notes that "the predictable component of output growth has virtually disappeared. Although output was highly variable in the 1970s and early 1980s, most of this variation was predicted. In contrast, variations since the late 1980s have been surprises" (Tulip, 2005: p. 12).

---

[19] A variable can be highly volatile, but perfectly predictable. Consider for example the seasonal fluctuations in output that account for about 85 percent of total output variability, but which are almost perfectly predictable (Beaulieu and Miron, 1992).

Campbell (2007: p. 199) on his turn reports that "the information content of current conditions has (...) declined" since 1984, which is consistent with an increase in the confusion about the persistence of shocks.[20]

Next to this, there is also evidence that firm-specific uncertainty has increased over the last three decades: Campbell *et al.* (2001) report that firm-level volatility of annualized daily stock returns of US firms increased in the 80s and 90s compared to the 60s and 70s, whereas market-level volatility remained roughly constant over this time period. A similar conclusion is reached by Comin and Philippon (2006), who - besides analyzing stock returns - also look at sales volatility and credit spreads of US firms.

But has this increased uncertainty really had an impact on agents in the economy? Well, in any case it seems to have played a complicating role for the economists on the NBER Business Cycle Dating Committee: where the pre-1990 peaks and troughs were announced with an average lag of 7.5 months, it took about twice as long before the committee was certain that a recession had started or ended in the post-1990 period: since then, the announcement lag equals 14.8 months on average.[21] In this light it seems reasonable to assume that "ordinary" agents, who monitor the economic situation less carefully and probably use the information provided by the NBER, are more uncertain about the current state and direction of the US economy as well.[22]

To sum all evidence up, it seems that shocks are surrounded by noisier signals on their persistence since the 1980s. In terms of our two-period model this implies that θ has moved towards $\frac{1}{2}$. Consequently, the option value of waiting has increased and employers simply have to wait longer (in order to obtain more signals) before they can make well-informed decisions.

Why the shocks became noisier is still an open question, however. They could have changed just because of "bad luck", while it is - somewhat paradoxically - also possible that this is a result of the increased communication of both monetary and real statistics over the years: Amador and Weill (2008) for example show that more public information on aggregate fundamentals can actually *increase* confusion and uncertainty about them - thereby leading to less accurate forecasts. The reason is that individuals put less weight on private information after public announcements

---

[20]A similar development seems to have been going on with respect to inflation. Atkeson and Ohanian (2001) and Orphanides and Van Norden (2005) for example both find that the predictability of inflation has gone down since 1984.

[21]See the announcement dates on http://www.nber.org/cycles.html.

[22]Note that all this is consistent with the finding of Kim and Nelson (1999). They report that the gap in the mean growth rate during booms and recessions has narrowed since the mid-1980s, thereby making the different phases of the business cycle (booms and recessions) harder to distinguish from each other.

(as the acquisition of information is costly), which, Amador and Weill show, can reduce the informational efficiency of the price system to such a large extent that uncertainty about fundamentals goes up.[23]

8   Is the labor adjustment lag the answer to some macroeconomic puzzles?

As will be argued below, theory predicts that an increase in the lag with which labor input is adjusted in response to structural shocks should have substantial implications for certain macroeconomic variables, such as labor productivity and the volatility of output. In this section we will describe the theoretical predictions with respect to these variables and show that these predictions are consistent with the puzzling behavior that these variables have displayed in reality since the 1980s.

8.1   The Great Moderation

As is well-known, the volatility of several macroeconomic variables (in particular that of output) has reduced significantly since the 1980s. This observation is generally referred to as "the Great Moderation". The fact that the timing of the Great Moderation (often estimated to lie somewhere around 1984; *cf.* McConnell and Perez-Quiros (2000)) roughly coincides with the introduction of the labor adjustment lag (which must lie somewhere in between the 1982 recession and the 1990/1 one), suggests that the two developments may be related. The explanation could be as follows.

Before the 1990/1 recession, a persistent positive productivity shock would increase output (assumed to be produced according to $y_t = \exp(z_t)k_t^\alpha h_t^{1-\alpha}$) for two distinct reasons: first, directly because productivity z_t went up, and second because labor input h_t would soon follow. After all, before the 1990s, the cycles for productivity and labor input stood roughly in phase as a result of which they amplified

<hr/>

[23]Similar arguments are developed in Morris and Shin (2005) and Wong (2008). The existence of this trade-off is well acknowledged by central bankers. Tommaso Padoa-Schioppa (a former member of the ECB's executive board) for example stated in an interview with the *Wall Street Journal* published on July 15th 2004 that he sees "a danger that the market becomes lazy" and that the market can be spoiled to the point where "it relies not on its own analysis, but on the analysis of the central bank". His former colleague Lorenzo Bini Smaghi added to this in a 2006 speech that "a trade-off needs to be made between central banks' desire to closely guide market behavior and the need to let markets make their own assessment and leave them the responsibility for their own expectations" and that there is a risk that central banks become like "a dog chasing its tail" (see http://www.ecb.int/press/key/date/2006/html/sp060508.en.html).

each other - thereby leading to a rather high volatility in output.

However, since the mid-1980s labor input responds with a greater lag to structural shocks (see the IRFs presented in Section 2.2 and the evidence in Stiroh (2009), who reports that the contemporaneous correlation between productivity shocks and labor input has declined since 1984). Consequently, when a productivity shock hits, output shows a much more muted response. Initially output just increases directly (through the increase in z_t) and only after a while (after it has become clear that the shock indeed is persistent) output goes up because of additional hiring. So due to the increase in the labor adjustment lag, the cycles for technology and labor input stand no longer in phase as a result of which the cycles of these production inputs amplify each other less than they did before.

But is a lagged labor input response *quantitatively* able to explain the observed reduction in output volatility, or are the effects only of minor importance? To address this question, we conducted the following experiment. Consider a basic RBC-model (see King and Rebelo (2000)), where the representative household faces the standard optimization problem:

$$
\max_{\{c_t,h_t,k_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{c_t^{1-\nu} - 1}{1-\nu} + \frac{\vartheta}{1-\eta} \left[(1-h_t)^{1-\eta} - 1 \right] \right]
$$

$$
s.t. \ y_t = \exp(z_t) k_t^{\alpha} h_t^{1-\alpha}
$$

$$
k_{t+1} = (1-\delta)k_t + i_t
$$

$$
y_t = c_t + i_t
$$

$$
z_{t+1} = \rho z_t + \varepsilon_{t+1}, \ \varepsilon_t \sim \mathcal{N}(0,\sigma^2)
$$

The associated first-order conditions read:

$$
\partial c_t \ : \ c_t^{-\nu} = \lambda_t \tag{10}
$$

$$
\partial h_t \ : \ \vartheta \left(1 - h_t\right)^{-\eta} = (1-\alpha) \lambda_t \exp(z_t) k_t^{\alpha} h_t^{-\alpha} \tag{11}
$$

$$
\partial k_{t+1} \ : \ \lambda_t = \beta \mathbb{E}_t \left\{ \lambda_{t+1} \left[\alpha \exp(z_{t+1}) k_{t+1}^{\alpha-1} h_{t+1}^{1-\alpha} + (1-\delta) \right] \right\} \tag{12}
$$

Here, λ_t measures the time t marginal utility of wealth.

The model (in which each period corresponds to one quarter in reality) is calibrated using standard values in the literature (see Table 2). The preference parameter for leisure ϑ is set equal to 3.48, to match the fact that US agents spend about 20 percent of their available time on market production. Finally, the standard deviation of the productivity shock (σ) is set such that the standard deviation of the simulated

27

series for output matches its US pre-1984 data equivalent of 0.02 (as reported by Galí and Gambetti (2009)).

Symbol	Interpretation	Value
α	Output elasticity of capital	0.33
β	Quarterly discount rate	0.99
δ	Quarterly depreciation rate	0.025
η	Labor supply elasticity	1
ϑ	Preference for leisure	3.48
ν	Coefficient of relative risk aversion	1
ρ	AR-coefficient on productivity process	0.95
σ	Standard deviation of productivity shock	0.015

Table 2: Calibration

However, as already noted by Van Rens (2004, p.27), the standard RBC-model fails in not generating a labor adjustment lag. This can for example be seen by the fact that there is almost perfect comovement between employment and the business cycle: in the standard RBC-model described above $corr(y, h) = 0.9802$, whereas post-1984 US data for example imply a value of 0.8148. To address this deficiency of the standard model, we propose the following fix to bring the procyclicality of labor input down to its post-1984 value: we introduce a labor adjustment lag in a very mechanical way, namely by altering the production function so that it only includes *lagged* labor input. Total output then follows $y_t = \exp(z_t)k_t^\alpha h_{t-i}^{1-\alpha}$, where i is the lag-length. As a result, first-order conditions (11) and (12) change into:

$$\partial h_t \quad : \quad \vartheta\left(1 - h_t\right)^{-\eta} = (1 - \alpha)\,\beta^i \lambda_{t+i} \exp(z_{t+i}) k_{t+i}^\alpha h_t^{-\alpha} \tag{11a}$$

$$\partial k_{t+1} \quad : \quad \lambda_t = \beta \mathbb{E}_t \left\{ \lambda_{t+1} \left[\alpha \exp(z_{t+1}) k_{t+1}^{\alpha-1} h_{t+1-i}^{1-\alpha} + (1 - \delta) \right] \right\} \tag{12a}$$

The implications for the standard deviation of output are shown in Figure 11.
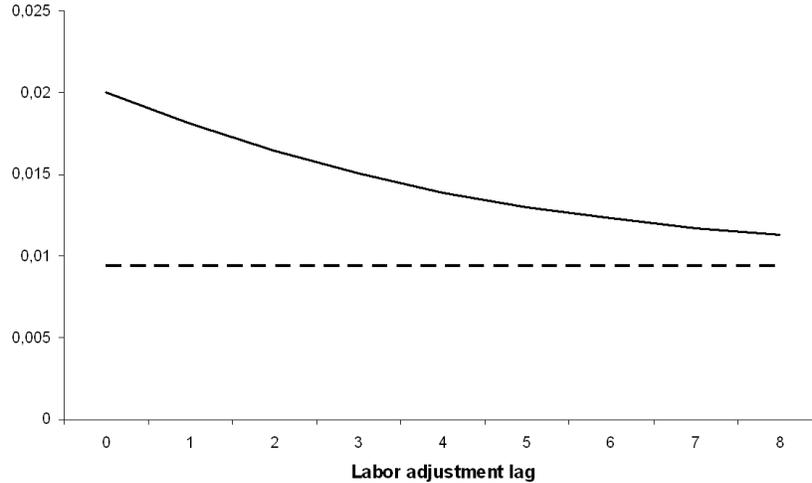
28

Figure 11: Standard deviation of output as function of the labor adjustment lag

As the figure shows, a labor adjustment lag indeed reduces output volatility. Interestingly, when the model's correlation between output and labor input approximately matches its post-1984 data equivalent of 0.8148 at $i = 8$, the model's standard deviation of output roughly matches its post-1984 data equivalent as well (equal to 0.0094 (Galí and Gambetti, 2009) and indicated by the dashed line in Figure 11). The reduced contemporaneous correlation between output and labor input (brought about by an increase in the labor adjustment lag) is thus quantitatively able to explain the Great Moderation in a standard RBC-setting.

However, a lag of eight quarters seems implausibly long. This therefore asks for a more realistic and structural model that generates a lagged labor input response. As shown in Andolfatto (1996, Table 1) introducing labor market matching frictions à la Pissarides (2000) does not solve the problem for a standard RBC-model: in that case, the contemporaneous correlation between output and labor input still equals 0.96. One possibility to fix this problem would be to build a DSGE-model featuring labor adjustment costs along with permanent and transitory shocks.[24] When agents are only able to distinguish between the two types of shocks by solving a signal extraction problem via collecting more information over time, we suspect (based upon our analytical results in Section 6 of this paper) that a labor adjustment lag will result. We leave this for future work.

Relating our finding that the labor adjustment lag has the potential to explain the Great Moderation to our discussion in Section 7 suggests two novel hypotheses

---

[24]Essentially, this is an infinite horizon version of our model in Section 4.

on the sources of the Great Moderation. First, it may be a result of higher (implicit) labor adjustment costs since the 1980s due to for example the increased importance of experience rating and high-skilled labor over time.[25] Part of this hypothesis can be tested cross-sectionally as it implies that countries with higher labor adjustment costs should display lower output volatility. And as reported by Merkl and Schmitz (2009), this is indeed the case.

Second, and contrasting with the "good luck hypothesis" advocated by among others Stock and Watson (2005),[26] our results point toward the possibility of a "bad luck hypothesis": potentially because of bad luck, shocks have been surrounded by more noise on their persistence since the mid-80s, as a result of which a given shock has induced smaller contemporaneous fluctuations in factor inputs - and hence output - since.[27] Note that - as the good luck hypothesis - ours also conditions on the size of the shocks and thus does not rule out heavy business cycle fluctuations since the 1980s (a claim that is pretty unwarranted in the light of the recent "Great Recession"); it only implies that a *given* persistent shock leads to greater macroeconomic fluctuations if the fact that the shock is persistent, is realized upon impact.

Hereby, we take an intermediate stand in the debate on the question whether the Great Moderation is caused by smaller shocks hitting the economy since the 1980s or by weaker propagation of a given shock.[28] According to our analysis, the *nature* of the shocks hitting the economy has changed (from clear to noisy), thereby weakening the propagation of these shocks through the economy. Note that this bad luck hypothesis is able to give an explanation for the paradoxical situation that the reduction in macroeconomic volatility seems to have been accompanied by a decrease in predictability: in our hypothesis, the decrease in predictability has actually caused the reduction in macroeconomic volatility! This also suggests that the Great Moderation has been accompanied by a reduction in macroeconomic efficiency, as agents have not been able to respond efficiently to shocks on impact due to a lack of clear information.

---

[25]Note that especially this latter explanation views the observed drop in volatility more as being part of a longer-term downward trend rather than as a structural break around 1984 (*cf.* Blanchard and Simon (2001) for a similar view).

[26]Proponents of the "good luck hypothesis" argue that the Great Moderation is simply a result of good luck in that the economy was not hit by any major shocks over the moderation period.

[27]A similar hypothesis is developed in independent work by Bullard and Singh (2009), who explain the Great Moderation via a learning mechanism. They argue that "regimes (*recessions and expansions*) which are closer together pose a more difficult inference problem for agents. The agents then take actions which are not as extreme as they would be under complete information. The result is a moderating force in the economy" (p.16, part in italics added).

[28]Stock and Watson (2002) argue the former, while Giannone, Lenza and Reichlin (2008) favor the latter explanation.

8.2 The cyclicality and volatility of labor productivity

Besides the volatility of output, the volatility and cyclicality of US labor productivity have also shown puzzling behavior since the mid-1980s. In particular, Galí and Gambetti (2008) and Galí and Van Rens (2009) report two stylized facts relevant to our story:

1. the procyclicality of labor productivity has dropped over the years. In particular, the correlation between output and labor productivity went down from its pre-1984 value of 0.87 to a significantly lower post-1984 value of 0.72;

2. the standard deviation of labor productivity relative to that of output is higher in the post-1984 sample than it was in the pre-1984 sample (went up from 0.45 to 0.55).

As labor productivity is nothing more than an accounting identity (by definition it equals total output divided by labor input), simple arithmetic already implies that an increase in the lag between output and labor input should have implications for its behavior. This can be shown most easily by conducting a simple sine-experiment, similar to the one in Wen (2004). First, assume that the cyclical component of the employment series follows:

$$n_t = \sin(\omega t)$$

Similarly, the cyclical component of output is assumed to be generated by:

$$y_t = \mu \sin(\omega t + \phi)$$

Here, the parameter μ allows us to capture the empirical fact that output is slightly more volatile than employment, while setting $\phi \neq 0$ results in a labor input lag.

Now consider Figure 12, which is generated by setting $\omega = 0.2$ and $\phi = 0$ such that output and labor input are in phase with each other, which approximates the situation before the 1990/1 recession. We set μ equal to 1.2 such that the standard deviation of employment relative to that of output roughly matches its empirical counterpart of 0.85.
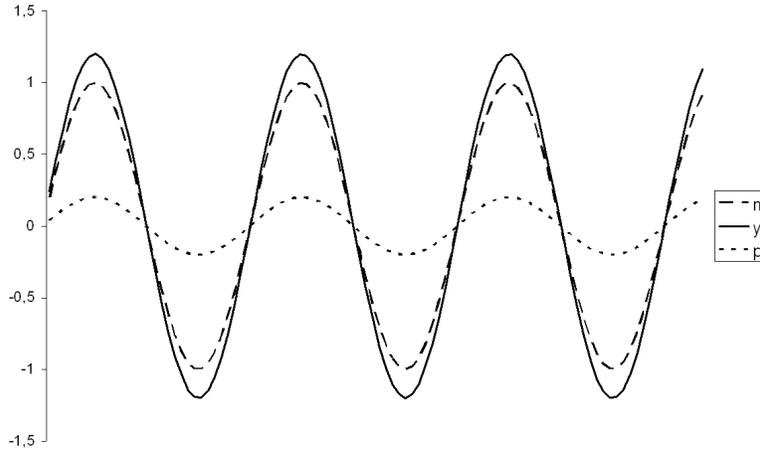
Figure 12: Cyclicality without labor adjustment lag

Next to the series for output and employment, Figure 12 also shows the implied behavior of labor productivity $p \equiv y - n$. As can be seen from the figure, labor productivity is perfectly procyclical, *i.e. $corr(y, p) = 1$.*

Next, we approximate the situation after 1984 by introducing a labor input lag by setting $\phi > 0$, say equal to 0.5. As can be seen from Figure 13, this simple experiment indeed predicts that the introduction of the lag has important consequences for the behavior of labor productivity: in particular, the cyclicality of labor productivity goes down (in this case $corr(y, p) = 0.58$), while the volatility of labor productivity relative to that of output goes up - exactly in line with what the data show (*cf.* stylized facts 1 and 2).[29]

---

[29]Note that we can match the observed decline in the correlation between output and labor input exactly by setting $\phi = 0.115$ to approximate the situation before 1984 and picking $\phi = 0.22$ to capture the situation after.
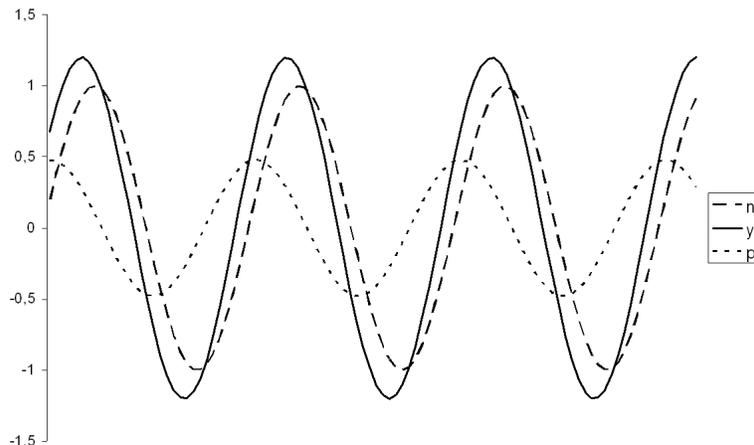
Figure 13: Cyclicality with labor adjustment lag

9   Conclusion

In this paper, we have presented evidence that the lag with which labor input reacts to structural economic shocks went up in the 1980s, thereby bringing jobless recoveries and recessions that were relatively job preserving to the US economy. We have shown that this lagged response is optimal in a setting where labor input is costly to adjust and where employers are uncertain about the persistence of shocks that drive the business cycle.

To explain the *increase* in this lag, we have therefore looked for changes in the adjustment process of labor input and macroeconomic predictability since the mid-1980s. And in both cases there is evidence for these changes to be there. For example, since the early 1980s the use of temporary layoffs has been discouraged by US government policy - thereby leading to worker-firm break-ups that are of a more permanent nature and hence more difficult to restore. In addition, macroeconomic predictability has shown a remarkable decline since the mid-80s, suggesting that the clarity of information about the persistence of shocks has gone down. Both developments increase employers' incentives to postpone hiring and firing decisions that are costly to reverse and could thereby explain the increase in the labor adjustment lag.[30]

---

[30]Both factors can also explain the simultaneous increase in the importance of temporary employment in adjusting labor input (documented in Peck and Theodore (2007)): because the aforementioned developments make firms less eager to commit to any decisions that are costly to reverse, they are now willing to make use of temp agencies that charge fees of 25 percent or more of the wage.

We finally show that the labor adjustment lag has the potential to explain two macroeconomic puzzles that characterize US data since the mid-1980s. First, it is able to explain the reduction in the procyclicality of labor productivity that has kicked in over this period. Regarding the simultaneous reduction in output volatility, we show that this could be explained by the fact that agents have become slower in adjusting their labor input in response to driving shocks. Consequently, the exogenous shock-part and the endogenous factor input-part of the production function are no longer in phase with each other, as a result of which they amplify each other less than they did before, thereby leading to smaller fluctuations in output. Note that the analysis employed in this paper similarly applies to capital investments, as they also contain some element of irreversibility. Therefore, it would be interesting to see whether capital adjustments also respond with a longer lag to structural shocks since the mid-1980s. If so, then this would add to the moderating force already exercised by the lagged response of labor input. Here, one runs however into the problem that capital is hard to measure in practice.

As emphasized before, we have adopted the view that labor input rises in response to a positive technology shock in this paper. Whether this view is correct is however still subject to a lively debate (see Galí and Rabanal (2005) for an overview). This is however not essential for our story. What is essential, is that agents in the economy are uncertain about the persistence of *shocks that drive the business cycle.* Whether they occur at the supply or demand side is not important.

Finally, this paper also points at several directions for future research. For example, why exactly has macroeconomic predictability come down? Is this just a result of bad luck or is there a link with the change in the conduct of monetary policy since the 1980s? And will predictability remain low in the future or will it increase again?

Hopefully, we will be able to answer these questions in future research.

10 Appendix: analysis of filtered data

10.1 Time domain

Section 2.1 used raw data (that was only seasonally adjusted) and may therefore be distorted by the presence of a trend. To get around this issue, and to ensure that moments are well-defined, this section focuses on the cyclical component of the data by looking at the cross-correlation between the HP-filtered series of GDP and

---

Essentially, these agencies have thus stepped into the space created by the greater value employers attach to remaining flexible since the 1980s. Also see the discussion in Dixit and Pindyck (1994, p.15-16).

employment.[31] Using all available postwar data (from 1947Q1 to 2008Q4) results in the following cross-correlogram.[32]
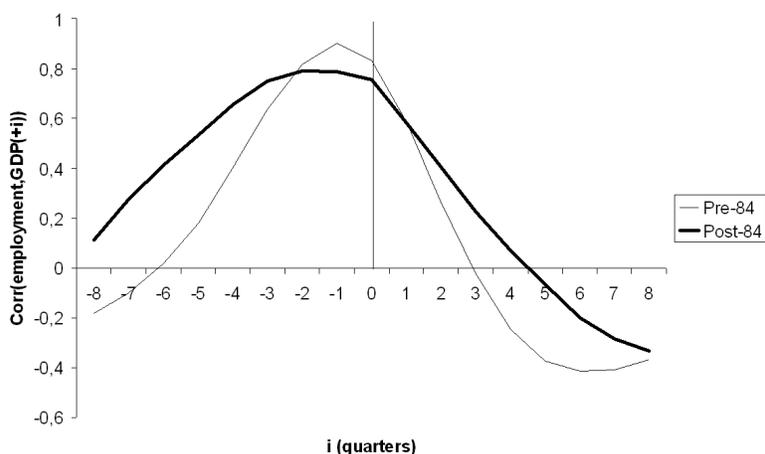

Figure A1: Correlation between employment and lagged/led GDP

This figure indeed suggests that employment responds with a greater lag to GDP fluctuations since 1984Q1: while employment before 1984Q1 correlates mostly with GDP one quarter ago, the peak correlation occurs at two quarters in the post-1984Q1 sample. More markedly, the persistence in the employment-lagged GDP correlation is also much stronger since 1984.

Repeating the analysis using data on total hours worked produces a picture that looks very similar:

---

[31]The quarterly data are taken from the St. Louis Fed database. For GDP we use seasonally adjusted output in chained 2000-Dollars, while we use the same employment measure as before - only now the quarterly averages of the monthly data. The correlations are based upon the cyclical component of the data only, obtained by logging and subsequently HP-filtering (with $\lambda = 1600$) the raw data.

[32]As this appendix uses all available data, we picked a break point somewhere in between the "normal" 1982 recovery and the first jobless one in 1991. In particular, we chose 1984Q1 as this is the estimated starting date of the so-called 'Great Moderation' (*cf.* McConnell and Perez-Quiros (2000)). The possible relation between the two concepts is discussed in Section 8.1.
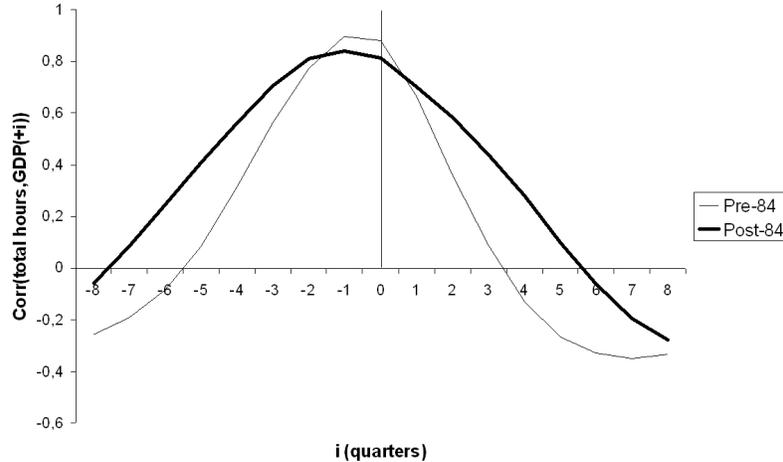
Figure A2: Correlation between total hours worked and lagged/led GDP

10.2 Frequency domain

The lead-lag relationship between output and employment can also be analyzed in the frequency domain. This may be the more natural thing to do as we are focusing on the cyclical properties of the data and enables us to capture the lead-lag relationship at each frequency in a single number, called the phase statistic. In addition, the phase statistic also allows us to distinguish explicitly between for- and backward phase shifts. After all, the evidence presented in the previous section is somewhat open to interpretation as the cross-correlograms displayed in Figures A1 and A2 also show an increased correlation between employment and *led* GDP.

To do this, we first need to calculate the cross-spectrum between GDP and labor input.[33] The cross-spectrum between two time-series x_t and y_t at frequency ω equals the Fourier-transform on the cross-covariances $\gamma_{xy,j}$ of the two series (where j indicates the lag) divided by 2π (note that i is the imaginary number $\sqrt{-1}$):

$$S_{xy}(\omega) = \frac{1}{2\pi} \sum_{j=-\infty}^{\infty} \gamma_{xy,j} e^{-i\omega j}$$

The cross-spectrum can be decomposed into a real part (the co-spectrum, $c_{xy}(\omega)$) and an imaginary part (the quadrature spectrum, $q_{xy}(\omega)$):

$$S_{xy}(\omega) = c_{xy}(\omega) + i q_{xy}(\omega)$$

---

[33]See Koopmans (1974), upon which this section heavily draws, for a comprehensive textbook treatment of (cross-)spectral analysis.

Twice the co-spectral density equals the covariance between the components that are in phase at frequency ω, while twice the quadrature spectral density gives the covariance between the components that are in quadrature (*i.e.* a quarter of the cycle out-of-phase in either direction) at that frequency.

From this decomposition we can calculate the phase statistic as:

$$\varphi_{xy}(\omega) = \arctan\left(\frac{q_{xy}(\omega)}{c_{xy}(\omega)}\right)$$

This statistic gives us the lead of y over x at frequency ω. It is expressed as a fraction of the cycle. Applying this to the HP-filtered series for output and employment yields the following graph.[34,35]
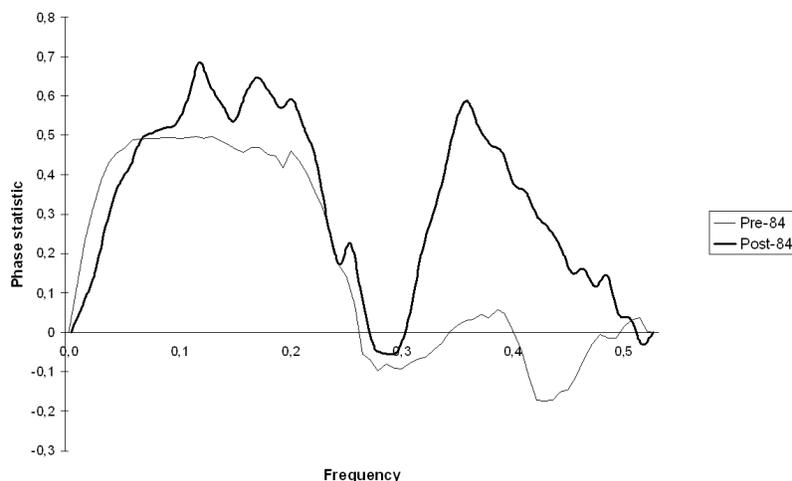

Figure A3: Phase spectrum of output and employment

To interpret the graph, note that a positive value of the phase statistic means that output leads employment at that frequency; a negative value implies the opposite. As can be seen from the figure, output leads employment over most frequencies. Furthermore observe that the phase lead of output over employment went up over

---

[34]As the HP-filter is a symmetric two-sided filter it does not induce a phase shift (provided that one drops some observations at the beginning and at the end of the sample to mitigate the start up/end of sample problems any two-sided filter has). HP-filtering then thus does not distort the data along the dimension we are interested in.

[35]The phase statistics were calculated by using the Tukey-Hamming smoothing algorithm. The window span was set equal to a standard value of approximately 15% of the number of data points. This implies a span length of 21 quarters for the pre-1984 sample and 15 quarters for the post-1984 sample.

all frequencies since 1984 (except for the very low ones). Following the evidence from the time domain, this suggests that the lag with which employment responds to output fluctuations went up in the mid-80s. In particular, the average phase statistic over all frequencies roughly doubled: it went up from an average value of 0.17 for the pre-1984 period, to a value of 0.32 for the period after 1984.

Figure A4 shows the corresponding graph for output and total hours worked.
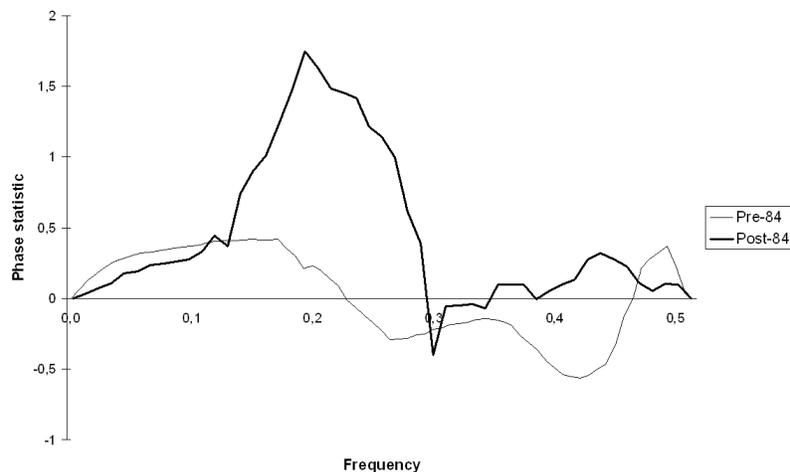

Figure A4: Phase spectrum of output and total hours worked

In this case, the average phase statistic goes up from 0.02 for the pre-1984 sample (indicating that output and total hours worked moved almost contemporaneously over this period), to a value of 0.44 since 1984.

11   Literature

Aaronson, D., E.R. Rissman and D.G. Sullivan (2004), "Can sectoral reallocation explain the jobless recovery?", *Federal Reserve Bank of Chicago Economic Perspectives*, QII, pp. 36-49.

Amador, M. and P.O. Weill (2008), "Learning from prices: Public communication and welfare", NBER Working Paper No. 14255.

Andolfatto, D. (1996), "Business cycles and labor-market search", *American Economic Review*, 86 (1), pp. 112-132.

Andolfatto, D. and G. MacDonald (2004), "Jobless recoveries", *mimeo*, Simon Fraser University.

Atkeson, A. and L.E. Ohanian (2001), "Are Phillips curves useful for forecasting inflation?", *Federal Reserve Bank of Minneapolis Quarterly Economic Review*, 25, pp. 2-11.

Bachmann, R. (2009), "Understanding jobless recoveries", *mimeo*, Yale University.

Beaulieu, J.J. and J.A. Miron (1992), "A cross country comparison of seasonal cycles and business cycles", *The Economic Journal*, 102 (413), pp. 772-788.

Bentolila, S. and G. Bertola (1990), "Firing costs and labor demand: How bad is Eurosclerosis?", *Review of Economic Studies*, 57 (3), pp. 381-402.

Bertola, G. (1990), "Job security, employment and wages", *European Economic Review*, 34 (4), pp. 851-879.

Blanchard, O.J. and D. Quah (1989), "The dynamic effects of aggregate demand and supply disturbances", *American Economic Review*, 79 (4), pp. 655-673.

Blanchard, O.J. and J. Simon (2001), "The long and large decline in US output volatility", *Brookings Papers on Economic Activity*, 1, pp. 135-164.

Bloom, N. (2009), "The impact of uncertainty shocks", *Econometrica*, 77 (3), pp. 623-685.

Brencic, V. and J.B. Norris (2008), "Online job boards as an employer recruitment tool", Paper prepared for the 2008 World Congress on National Accounts and Economic Performance Measures for Nations, Arlington VA.

Bullard, J. and A. Singh (2009), "Learning and the Great Moderation", Federal Reserve Bank of St. Louis Working Paper.

Burdett, K. and R. Wright (1989), "Unemployment insurance and short-time compensation: The effects on layoffs, hours per worker, and wages", *Journal of Political Economy*, 97 (6), pp. 1479-1496.

Campbell, S.D. (2007), "Macroeconomic volatility, predictability, and uncertainty in the Great Moderation: Evidence from the survey of professional forecasters", *Journal of Business & Economic Statistics*, 25 (2), pp. 191-200.

Campbell, J.Y., M. Lettau, B.G. Malkiel and Y. Xu (2001), "Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk", *Journal of Finance*, 56 (1), pp. 1-43.

Card, D. and P.B. Levine (1994), "Unemployment insurance taxes and the cyclical and seasonal properties of unemployment", *Journal of Public Economics*, 53 (1), pp. 1-29.

Christiano, L.J., M. Eichenbaum and R. Vigfusson (2003), "What happens after a technology shock?", NBER Working Paper No. 9819.

Comin, D. and T. Philippon (2005), "The rise in firm-level volatility: Causes and consequences", *NBER Macroeconomics Annual*, 20, pp. 167-227.

D'Agostino, A., D. Giannone and P. Surico (2006), "(Un)predictability and macro-economic stability", ECB Working Paper No. 605.

Dixit, A.K. and R.S. Pindyck (1994), *Investment under uncertainty*, Princeton, NJ: Princeton University Press.

Feldstein, M. (1976), "Temporary layoffs in the theory of unemployment", *Journal of Political Economy*, 84 (5), pp. 937-957.

Francis, N. and V. Ramey (2009), "Measures of per capita hours and their implications for the technology-hours debate", *Journal of Money, Credit and Banking*, 41 (6), pp. 1071-1097.

Fujita, S. (forthcoming), "Dynamics of worker flows and vacancies: Evidence from the sign-restriction approach", *Journal of Applied Econometrics*.

Galí, J. (1999), "Technology, employment, and the business cycle: do technology shocks explain aggregate fluctuations?", *American Economic Review*, 89 (1), pp. 249-271.

Galí, J. and L. Gambetti (2009), "On the sources of the Great Moderation", *American Economic Journal: Macroeconomics*, 1 (1), pp. 26-57.

Galí, J. and P. Rabanal (2005), "Technology shocks and aggregate fluctuations: How well does the RBC model fit postwar US data?, in: M. Gertler and K. Rogoff (*eds.*), *NBER Macroeconomics Annual 2004*, Cambridge, Massachusetts: MIT Press.

Galí, J. and T. van Rens (2009), "The vanishing procyclicality of labor productivity and the Great Moderation", *mimeo*, Universitat Pompeu Fabra.

Gertler, M., L. Sala and A. Trigari (2008), "An estimated monetary DSGE model with unemployment and staggered nominal wage bargaining", *Journal of Money Credit and Banking*, 40 (8), pp. 1713-1764.

Giannone, D.M., M. Lenza and L. Reichlin (2008), "Explaining the Great Moderation: It is not the shocks", *Journal of the European Economic Association*, 6 (2-3), pp. 621-633.

Groshen, E.L. and S. Potter (2003), "Has structural change contributed to a jobless recovery?", *Current Issues in Economics and Finance*, 9 (8), pp. 1-7.

Hall, R.E. (2005), "Employment efficiency and sticky wages: evidence from flows in the labor market", *Review of Economics and Statistics*, 87 (3), pp. 397-407.

Hamermesh, D.S. and G.A. Pfann (1996), "Adjustment costs in factor demand", *Journal of Economic Literature*, 34 (3), pp. 1264-1292.

Kim, C-J and C. Nelson (1999), "Has the U.S. economy become more stable? A Bayesian approach based on a Markov-switching model of the business cycle", *The Review of Economics and Statistics*, 81 (4), pp. 608-616.

King, R.G. and S.T. Rebelo (2000), "Resuscitating Real Business Cycles", in: M. Woodford and J.B. Taylor (*eds.*), *Handbook of Macroeconomics*, Vol. 1A, Amster-

dam; New York and Oxford: Elsevier Science, North-Holland.

Koenders, K. and R. Rogerson (2005), "Organizational dynamics over the business cycle: a view on jobless recoveries", *Federal Reserve Bank of St. Louis Review*, 87 (4), pp. 555-579.

Koopmans, L.H. (1974), *The spectral analysis of time series*, New York, NY: Academic Press.

McConnell, M.M. and G. Perez-Quiros (2000), "Output fluctuations in the United States: What has changed since the early 1980s?", *American Economic Review*, 90 (5), pp. 1464-1476.

Morris, S. and H. Shin (2005), "Central bank transparency and the signal value of prices", *Brookings Papers on Economic Activity*, 2, pp. 1-43.

OECD (1999), *OECD Employment Outlook 1999*, Paris: Organisation for Economic Cooperation and Development.

Orphanides, A. and S. van Norden (2005), "The reliability of inflation forecasts based on output gap estimates in real time", *Journal of Money, Credit and Banking*, 37 (3), pp. 583-601.

Peck, J. and N. Theodore (2007), "Flexible recession: the temporary staffing industry and mediated work in the United States", *Cambridge Journal of Economics*, 31 (2), pp. 171-192.

Pissarides, C.A. (2000), *Equilibrium unemployment theory*, Second Edition, Cambridge, Mass.: MIT Press.

Ravn, M.O. and S. Simonelli (2008), "Labor market dynamics and the business cycle: Structural evidence for the United States", *Scandinavian Journal of Economics*, 109 (4), pp. 743 - 777.

Schuh, S. (2001), "An evaluation of recent macroeconomic forecast errors", *New England Economic Review*, January/February 2001, pp. 35-56.

Shimer, R. (2005), "The cyclicality of hires, separations, and job-to-job transitions", *Federal Reserve Bank of St. Louis Review*, 87 (4), pp. 493-507.

Stiroh, K.J. (2009), "Volatility accounting: A production perspective on increased economic stability", *Journal of the European Economic Association*, 7 (4), pp. 671-696.

Stock, J.H. and M.W. Watson (2002), "Has the business cycle changed and why?", *NBER Macroeconomics Annual*, 17, pp. 159-218.

Stock, J.H. and M.W. Watson (2005), "Understanding changes in international business cycle dynamics", *Journal of the European Economic Association*, 3 (5), pp. 968-1006.

Topel, R.H. (1990), "Financing unemployment insurance: History, incentives, and reform", in: W.L. Hansen and J.F. Byers (*eds.*), *Unemployment Insurance*, University

of Wisconsin Press, Madison, pp. 108–135.

Tulip, P. (2005), "Has output become more predictable? Changes in Greenbook forecast accuracy", Finance and Economics Discussion Series No. 31.

Van Rens, T. (2004), "Organizational capital and employment fluctuations", *mimeo*, Universitat Pompeu Fabra.

Van Wijnbergen, S. (1985), "Trade reform, aggregate investment and capital flight: On credibility and the value of information", *Economics Letters*, 19 (4), pp. 369-372.

Wen, Y. (2004), "What does it take to explain procyclical productivity?", *Contributions to Macroeconomics*, 4 (1), pp. 1-38.

Wong, J. (2008), "Information acquisition, dissemination, and transparency of monetary policy", *Canadian Journal of Economics*, 41 (1), pp. 46-79.