

NBER WORKING PAPER SERIES

MARKUPS, GAPS, AND THE WELFARE COSTS OF BUSINESS FLUCTUATIONS

Jordi Galí
Mark Gertler
J. David López-Salido

Working Paper 8850
<http://www.nber.org/papers/w8850>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

March 2002

We thank Jeff Amato, Andy Levin, Susanto Basu, Jonathan Parker and Bob Hall for helpful comments, as well as participants in seminars at LSE, Bocconi, IIES, UPF, Bank of Spain, Bank of England, NBER Summer Institute, CERGE, EUI, UAB, and the ECB Workshop on DSGE Models. Galí is thankful to the Bank of Spain for financial support and Galí and Gertler acknowledge the support of the NSF and the C.V. Starr Center. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 2002 by Jordi Galí, Mark Gertler and J. David López-Salido. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Markups, Gaps, and the Welfare Costs of Business Fluctuations
Jordi Galí, Mark Gertler and J. David López-Salido
NBER Working Paper No. 8850
March 2002
JEL No. E3

ABSTRACT

In this paper we present a simple, theory-based measure of the variations in aggregate economic efficiency associated with business fluctuations. We decompose this indicator, which we refer to as “the gap”, into two constituent parts: a price markup and a wage markup, and show that the latter accounts for the bulk of the fluctuations in our gap measure. Finally, we derive a measure of the welfare costs of business cycles that is directly related to our gap variable, and which takes into account explicitly the existence of a varying aggregate inefficiency. When applied to postwar U.S. data, for plausible parametrizations, our measure suggests welfare losses of fluctuations that are of a higher order of magnitude than those derived by Lucas (1987). It also suggests that the major postwar recessions involved substantial efficiency costs.

Jordi Galí
CREI
Universitat Pompeu Fabra
Ramon Trias Fargas 25
08005 Barcelona Spain
and NBER
jordi.gali@econ.upf.es

Mark Gertler
Department of Economics
New York University
269 Mercer Street, 7th Floor
New York, NY 10003
and NBER
mark.gertler@nyu.edu

J. David López-Salido
Bank of Spain

1 Introduction

To the extent that there exist price and wage rigidities, or possibly other types of market frictions, the business cycle is likely to involve inefficient fluctuations in the allocation of resources. Specifically, the economy may oscillate between expansionary periods where the volume of economic activity is close to the social optimum and recessions that feature a significant drop in production relative to the first best. In this paper we explore this hypothesis by developing a simple measure of aggregate inefficiency and examining its cyclical properties. The measure we develop - which we call “the inefficiency gap” or “the gap”, for short - is based on the size of the wedge between the marginal product of labor and the marginal rate of substitution between consumption and leisure. Deviations of this gap from zero reflect an inefficient allocation of employment. By constructing a time series measure of the inefficiency gap, we are able to obtain some insight into both the nature and welfare costs of business cycles.

From a somewhat different perspective, as we discuss below, the inefficiency gap corresponds to the inverse of the markup of price over social marginal cost. Pro-cyclical movements in the inefficiency gap accordingly mirror counter-cyclical movements in this markup. Our approach, however, differs from much of the recent literature on business cycles and markups by allowing for the possibility that the movement in the overall markup depends on variations in a wage markup as well as in a price markup.¹ Put differently, in contrast to much of the existing literature, we allow for the possibility of labor market frictions that introduce a wedge between the wage and the household’s consumption/leisure tradeoff. By doing so we can obtain some sense of the relative importance of price versus wage rigidities for overall fluctuations in the inefficiency gap. In addition, focusing on the gap between the labor demand and supply curves leads directly to a welfare measure of the costs of business cycles, based on the lost surplus owing to aggregate fluctuations.

Our approach builds on a stimulating paper by Hall (1997) that analyzes the cyclical behavior of the neoclassical labor market equilibrium. Specifically Hall shows that the business cycle is associated with highly pro-cyclical movements in the difference between the observable component of the household’s marginal rate of substitution and the marginal product of labor. Hall interprets this difference –which we refer to as the Hall residual– as reflecting a preference shock². However, we present evidence that suggests this residual instead reflects counter-cyclical markup variation. We then proceed to show that, as an implication, business cycles may entail significant welfare

¹See Rotemberg and Woodford (1999) for a survey of the literature on business cycles and counter-cyclical markups. For business cycle models that feature a role for wage markups as well as price markups, see Blanchard and Kiyotaki (1987) and Erceg, Henderson and Levin (2000).

²To be fair, Hall does not take the preference shock interpretation literally, but rather as a starting point for analyzing the significance of the cyclical movement in the gap between the measured labor demand and supply curves. In his conclusion, he observes that cyclical movements in unemployment could underlie the measured preference shock.

losses.

In section 2 we develop a framework for measuring the inefficiency gap and its price and wage markup components in terms of observables, conditional on standard assumptions about preferences and technology. In section 3 we present empirical measures of this variable for the postwar U.S. economy. We show that the inefficiency gap exhibits large procyclical swings, thus confirming our basic hypothesis. In addition, most of its variation is associated with countercyclical movements in the wage markup.³ The price markup shows, at best, a weak contemporaneous correlation. Finally, we demonstrate the robustness of our gap measure and its decomposition to alternative assumptions about preferences and technology. In section 4 we consider the possibility that preference shocks underlie the variation in our gap measures. Specifically, we present VAR evidence that suggests that the Hall residual is endogenous and thus cannot simply reflect exogenous variation in preferences. The evidence is instead consistent with our maintained hypothesis that endogenous variation in markups is largely responsible for the movement in the inefficiency gap.

In Section 5 we examine the welfare consequences of business fluctuations. Our approach differs significantly from Lucas (1987), who examines the welfare costs of consumption variability associated with the cycle. We instead focus on the welfare costs associated with fluctuations in the efficiency of resource allocation, as implied by the time series variation in our gap measure. As we show, our framework implies that business contractions below the steady state produce greater efficiency costs than the efficiency gains arising from symmetric expansions above. This asymmetry, in turn, implies that fluctuations raise efficiency costs on average, even if these fluctuations are themselves symmetric. We show that under plausible parametrizations, business cycles of the magnitude observed in postwar U.S. data may indeed involve significant welfare costs. In addition to studying average performance over the postwar, we also examine particular episodes and find that the major recessions indeed involved considerable efficiency losses.

Concluding remarks are in section 6.

2 The Gap and its Components: Theory

Let the *inefficiency gap* (henceforth, *the gap*) be defined as follows:

$$gap_t = mrs_t - mpn_t \tag{1}$$

where mpn_t and mrs_t denote, respectively, the (log) marginal product of labor and (log) marginal rate of substitution between consumption and leisure.

³In this respect our results are consistent with recent evidence in Sbordone (1999, 2001), Galí and Gertler (1999), Galí, Gertler and Lopez-Salido (2001) and Christiano, Eichenbaum and Evans (1997, 2001) that in somewhat different contexts similarly points to an important role for wage rigidity.

As illustrated by Figure 1, our gap variable can be represented graphically as the vertical distance between the *perfectly competitive* labor supply and labor demand curves, evaluated at the current level of employment (or hours). In much of what follows we assume that $\{gap_t\}$ follows a stationary process with a (possibly nonzero) constant mean, denoted by gap (without any time subscript). The latter represents the steady state deviation between mrs_t and mpn_t . Notice that these assumptions are consistent with both mrs_t and mpn_t being non-stationary, as it is likely to be the case in practice as well as in the equilibrium representation of a large class of dynamic business cycle models.

We next relate the gap to the markups in the goods and labor markets. Under the assumption of wage-taking firms, and in the absence of labor adjustment costs, the nominal marginal cost is given by $w_t - mpn_t$, where w_t is (log) compensation per unit of labor input (including non-wage costs).⁴ Accordingly, we define the aggregate price markup as follows:

$$\mu_t^p = p_t - (w_t - mpn_t) \quad (2)$$

$$= mpn_t - (w_t - p_t) \quad (3)$$

The aggregate wage markup is given by:

$$\mu_t^w = (w_t - p_t) - mrs_t \quad (4)$$

i.e., it corresponds to the difference between the wage and the marginal disutility of work, both expressed in terms of consumption. Notice that the wage markup should be understood in a broad sense, including the wedge created by payroll taxes paid by the firm and labor income taxes paid by the worker.

There are a variety of frictions (perhaps most prominently, wage and price rigidities) which may induce fluctuations in the markups: It is in this respect that these frictions are associated with inefficient cyclical fluctuations, or more precisely, with variations in the aggregate level of (in)efficiency. In particular, given that the marginal rate of substitution is likely to be procyclical, rigidities in the real wage—resulting either from nominal or real rigidities—will give rise to countercyclical movements in the wage markup.⁵ Nominal price rigidity, in turn, may give rise to a countercyclical price markup in response to demand shocks since, holding productivity constant, the marginal product of labor is countercyclical.⁶

To formalize the link between markup behavior and the gap, we first express equation (1) as

⁴We show subsequently that our results are robust to allowing for labor adjustment costs.

⁵Models with countercyclical wage markups due to nominal rigidities include Blanchard and Kiyotaki (1987) and Erceg, Henderson and Levin (2000). Alexopoulos (2000) develops a model with a real rigidity due to efficiency wages that can generate a countercyclical wage markup.

⁶With productivity shocks, the markup could be procyclical (since the marginal product of labor moves procyclically in that instance).

$$gap_t = -\{[mpn_t - (w_t - p_t)] + [(w_t - p_t) - mrs_t]\} \quad (5)$$

Combining equations (3), (4), and (5) then yields a fundamental relation linking the gap to the wage and price markups:

$$gap_t = -(\mu_t^p + \mu_t^w) \quad (6)$$

In the steady state, further:

$$gap = -(\mu^p + \mu^w) < 0 \quad (7)$$

where variables without time subscripts denote steady state values.

It is natural to assume that $\mu_t^p \geq 0$ and $\mu_t^w \geq 0$ for all t , implying $gap_t \leq 0$ for all t . In this case the level of economic activity is inefficiently low (the gap is always negative), so that (small) increases in our gap measure will be associated with a smaller distortion (i.e., an allocation closer to the perfectly competitive one). Notice also that countercyclical movements in these markups imply that the gap is high in booms and low in recessions.

To the extent that we can measure the two markups (or, at least their variation), we can characterize the behavior of the gap, as well as its composition. Identifying the markups requires some assumptions about technology and preferences. We first consider a baseline case with reasonably conventional assumptions. We then show that our results are robust to a number of leading alternative restrictions.

Given equation (2), identification of price markup variations only requires an assumption on technology. Under the assumption of a technology with constant elasticity of output with respect to hours (say, α), we have (up to an additive constant):

$$mpn_t = y_t - n_t \quad (8)$$

where y_t is output per capita and n_t is hours per capita.⁷

Combining equations (2) and (8) yields:

$$\mu_t^p = (y_t - n_t) - (w_t - p_t) \quad (9)$$

$$\equiv -ulc_t \quad (10)$$

Hence the price markup can be measured (up to an additive constant) as *minus* the (log) *real* unit labor costs, denoted by ulc_t .

Let c_t be consumption per capita and $\bar{\xi}_t$ be a deterministic, low frequency preference shifter. Then, the (log) marginal rate of substitution can be written (up to an additive constant) as:

⁷As we discuss in section 5, this specification of production allows for variable capital utilization. Under certain conditions it is also compatible with variable labor utilization, particularly if labor effort moves roughly proportionately with hours, as the evidence suggests (see, e.g., Basu and Kimball, 1997).

$$mrs_t = \sigma c_t + \varphi n_t - \bar{\xi}_t \quad (11)$$

where the parameter σ is the coefficient of relative risk aversion and φ measures the curvature of the disutility of labor. Following Hall (1997), we allow for the possibility of low frequency shifts in preferences over consumption versus leisure, as represented by movements in $\bar{\xi}_t$. These preferences shifts may be interpreted broadly to include institutional or demographic changes that affect the labor market. We differ from Hall, though, by restricting these shifts to the low frequency. In section 4 we provide evidence to justify this assumption. It follows that the wage markup is given by:

$$\mu_t^w = (w_t - p_t) - (\sigma c_t + \varphi n_t) + \bar{\xi}_t \quad (12)$$

Given a measure of both the price and the wage markup, one can obtain a measure of the gap using equation (6). Alternatively, one can combine equations (9), (12) and (6) to obtain:

$$gap_t = (\sigma c_t + \varphi n_t - \bar{\xi}_t) - (y_t - n_t) \quad (13)$$

3 The Gap and Its Components: Evidence

We now use the theoretical relations in the previous section to construct measures of the gap and its two main components: the price and wage markups. Our evidence is based on quarterly postwar U.S. data over the sample period 1960:4 - 1999:4.

Identification of gap and wage markup variations requires that we make an assumption on the coefficient of relative risk aversion σ and on φ , a parameter which corresponds to the inverse of the (Frisch) wage elasticity of labor supply. A vast amount of evidence from micro-data suggests wage elasticities mostly concentrated in the range of 0.05 – 0.3, though the business cycle literature tends to use much higher values.⁸ We accordingly use a baseline value $\varphi = 5$, which corresponds to a labor supply elasticity of 0.2, which is slightly above the mean of the labor supply elasticity estimates from the micro data. However, we also experiment with other values, including values used in the business cycle literature (see, e.g., Cooley and Prescott, 1995.)

⁸MacCurdy (1981) estimates the Frisch elasticity of labor supply for men to be 0.15, a finding that has been largely confirmed by subsequent literature (e.g., Altonji, 1986, and more recently Pencavel, (forthcoming)). In his survey of the literature, Card (1994) concludes that this elasticity is “surely no higher than 0.5 and probably no higher than 0.2.” Though, for an alternative view, see Mulligan (1998). Finally, less is known about female labor supply elasticity, but Pencavel (1998) has estimated a Frisch elasticity of 0.21 for this group. Pencavel’s sample period covers the mid 1970s to mid 1990s. Apparently, as the gap between male and female labor force participation has narrowed, female labor supply elasticities have become similar to male labor supply elasticities.

There is a similar controversy over the choice of the coefficient of relative risk aversion, which corresponds to the inverse of the intertemporal elasticity of substitution. Direct estimates of the latter tend to fall in the range $0.1 - 0.3$.⁹ This evidence suggests a value of σ that varies from 10 to 3.3. The business cycle literature instead tends to assume log utility over consumption (i.e., $\sigma = 1$), based on the justification that these preferences are consistent with balanced growth. We also adopt this assumption for our baseline case, but then also experiment with parametrizations consistent with the direct evidence.

In addition, we need to make an assumption to identify the low frequency shifter $\bar{\xi}_t$. Let $\tilde{\mu}_t^w \equiv (w_t - p_t) - (\sigma c_t + \varphi n_t)$ be the observable component of the wage markup. It follows that

$$\tilde{\mu}_t^w = \mu_t^w - \bar{\xi}_t \quad (14)$$

From this perspective, the wage markup μ_t^w is the cyclical component of $\tilde{\mu}_t^w$ and $\bar{\xi}_t$ is (minus) the trend component. To identify μ_t^w from observations on $\tilde{\mu}_t^w$, we use a band-pass filter which discards fluctuations outside a frequency range between 2 and 60 quarters.¹⁰

Finally, before proceeding, we note that the relationships derived in the previous section hold only up to an additive constant. Accordingly, our framework only allows us to identify the *variations* over time in the markup and its components, but not their levels. Our baseline results thus employ measures of the price and wage markups and the gap constructed using, respectively, equations (10), (12), and (6), and expressed in terms of deviations from their respective sample means.

Figure 2 presents the times series measure of our gap variable under our baseline assumptions of $\sigma = 1$ and $\varphi = 5$. Notice that this variable comoves strongly with the business cycle, displaying large declines during NBER-dated recessions (represented by the shaded areas in the graph). In addition, a number of researchers have estimated a reduction in the volatility of the business cycle, beginning around 1984. Figure 2 indicates that associated with this reduction has been a decline in the variability of the inefficiency gap. It is also interesting to observe that the gap hovers near zero for most of the period post 1995. The resulting implication is that the rapid output growth over this period must have been due to real factors (e.g. technology improvements) as opposed to excess demand.

We next decompose the movements of the gap into its wage and price markup components. The wage markup measures were constructed using (12).¹¹ The price

⁹Using micro-data, Barsky *et. al* (1997) estimate an intertemporal elasticity of substitution of 0.18, implying a coefficient of relative risk aversion slightly above 5. Using macro-data, Hall (1988) concludes that the intertemporal elasticity of substitution ($1/\sigma$) is likely below 0.2.

¹⁰We choose a filter of this frequency to sort cycle from trend following Staiger, Stock and Watson (2001). Our results are virtually unchanged by reducing the length of the filter from 60 quarters to 32.

¹¹The results are robust to simple adjustments for compositional bias of the real wage, based on

markup corresponds to minus the log of real unit labor costs, as implied by (10). Figure 3 shows the behavior of the gap against the wage markup. To facilitate visual inspection, we plot the inverse of the wage markup (i.e., minus the log wage markup). By definition, the difference between the gap and the inverse wage markup is the inverse price markup. What is striking about the pictures is the strong co-movement between the gap and the (inverse) wage markup. Put differently, the inefficiency gap seems to be driven largely by countercyclical movements in the wage markup.¹²

To be clear, our conclusion that countercyclical wage markup variation drives the variation in the gap rests on the assumption that wages are allocational.¹³ While this assumption is standard in the literature on business cycles and markups (e.g., Rotemberg and Woodford, 1999), it is not without controversy. In the event that wages are not allocational, our gap variable instead reflects entirely countercyclical movements in the price markup.

Table 1 reports some basic statistics that support the visual evidence in Figure 3. In particular, the Table reports a set of second moments for the gap and its two components: the wage and price markup, and also for detrended (log) GDP, a common indicator of the business cycle. Note first that the percent standard deviation of the gap is large (relative to detrended output) and that departures of the gap from steady state are highly persistent. In addition, the wage markup is nearly as volatile as the overall gap and is strongly negatively correlated, not only with the gap but with detrended output as well. This confirms the visual evidence that movements in the gap are strongly associated with countercyclical movements in the wage markup. On the other hand, the price markup is less volatile than the wage markup and does not exhibit a strong contemporaneous correlation with the gap¹⁴.

In Figure 4 we demonstrate that the qualitative pattern of the gap is robust to reasonable alternative assumptions about labor supply elasticity and about the coefficient of relative risk aversion. While the micro-evidence suggests a small labor supply elasticity, the business cycle literature tends to assume a high elasticity, typically unity and above. We accordingly reconstruct the gap measure assuming $\varphi = 1$, which implies a Frisch labor supply elasticity of unity. The top panel of Figure 4 plots the behavior of the gap under this new parametrization against the baseline

Barsky, Solon and Parker (1994).

¹²As a somewhat cleaner way to illustrate the strong countercyclical relation between the gap and the wage markup, we show later that this pattern also holds conditional on a shock to monetary policy.

¹³Some indirect evidence that wages are allocational is found in Sbordone (1999) and Gali and Gertler (1999) who show that firms appear to adjust prices in response to measures of marginal cost based on wage data. In turn, as we showed in an earlier version of this paper, they do not respond to marginal cost measures that employ the household's marginal rate of substitution in place of the wage, as would be appropriate if wages were not allocational.

¹⁴However, the relatively weak co-movement of the price markup with detrended output is useful for understanding the dynamics of inflation and the recent evidence on the New Keynesian Phillips curve. See Sbordone (1999) and Gali and Gertler (1999).

case of $\varphi = 5$. Overall, the qualitative pattern is similar across the two cases. A higher supply elasticity, however, does imply quantitatively smaller fluctuations in the gap. Intuitively, a more elastic labor supply curve implies that any change in employment from its natural level will yield a smaller change than otherwise in the distance between the labor demand and labor supply curves.

We next explore adjusting the coefficient of relative risk aversion. We consider a value of 5 for this parameter, implying an intertemporal elasticity of substitution of 0.2, consistent with the evidence mentioned earlier. The bottom panel accordingly plots the gap variable for the case of $\sigma = 5$ versus the baseline case of $\sigma = 1$. (In each instance we keep φ at its baseline value of 5.) Again, the qualitative pattern is similar across the cases. The amplitude of the gap variable, though, increases with risk aversion. Intuitively, a rise in risk aversion makes labor supply more inelastic, which raises the sensitivity of the gap to employment fluctuations.

Though we do not report the results here, it also remains true that the movements in the gap for both the high labor supply elasticity case and the high risk aversion case are associated largely with countercyclical movements in the wage markup. This should not be surprising since the wage markup is computed simply as minus the difference between the gap and the price markup, where the measure of the latter is invariant to the labor supply elasticity.

To summarize: the results thus far suggest that the business cycle is associated with large co-incident movements in the efficiency gap. Thus, under our framework, the evidence suggests that countercyclical markup behavior is an important feature of the business cycle. A decomposition of the gap, further, suggests that the countercyclical movement in the wage markup is by far the most important source of overall variations in the gap. This in turn suggests that some form of wage rigidity, either real or nominal, may be central to business fluctuations.

Finally, we now demonstrate that our gap measure is robust to alternative assumptions about production (that yield alternative measures of the marginal product of labor.) Our baseline case assumes constant elasticity of output with respect to hours. We consider three alternative assumptions suggested by Rotemberg and Woodford (1999): (i) Cobb-Douglas modified to allow for overhead labor; (ii) CES; and (iii) Cobb-Douglas with labor adjustment costs. In each case we follow the parametrization recommended in Rotemberg and Woodford (1999). As Figure 5 indicates, our gap measures are quite robust to these alternative assumptions. Though we do not report the results here, it remains the case that the movements in the gap are strongly associated with a countercyclical wage markup. For completeness, Figure 6 shows that our results are robust to allowing for a measure of the marginal rate of substitution based on time dependent preferences in leisure, following Eichenbaum, Hansen and Singleton (1988).

4 Preference Shocks, Hall's Residual and the Gap

As we have discussed, the notion that the business cycle is associated with cyclical movements in the gap between the labor demand curve and the observable component of the labor supply curve originated with Hall (1997). In his baseline identification scheme, however, Hall associated this gap entirely with preference shocks.¹⁵ In this section we show that the high frequency movements in the gap cannot be simply due to preference shocks.

Let us follow Hall (1997) by assuming that the marginal rate of substitution is now augmented with a preference shock ξ_t that contains a cyclical component, $\tilde{\xi}_t$, as well as a trend component, $\bar{\xi}_t$:

$$mrs_t = c_t + \varphi n_t - \xi_t \quad (15)$$

with

$$\xi_t = \bar{\xi}_t + \tilde{\xi}_t$$

where we maintain our baseline assumption that the coefficient of relative risk aversion, σ , is unity. Hall then defines the residual x_t as the difference between the “observable” component of the marginal rate of substitution, $c_t + \varphi n_t$, and the marginal product of labor, $y_t - n_t$:

$$x_t \equiv (c_t + \varphi n_t) - (y_t - n_t) \quad (16)$$

The issue then is how exactly to interpret the movement in Hall's residual. Using the augmented specification of the marginal rate of substitution allowing for preference shocks (15), together with (8) and the definition of the inefficiency gap (1), it is possible to express Hall's residual as follows:

$$x_t \equiv (mrs_t - mpn_t) + \xi_t \quad (17)$$

Hall assumes that the frictionless competitive equilibrium obtains in both goods and labor markets, implying $mrs_t - mpn_t = 0$. In this instance the residual x_t corresponds exactly to the preference shock.¹⁶ Hall analyzes the behavior of the time series for x_t constructed according to (16) and concludes it is highly procyclical.

However, because we allow for departures from perfect competition and the possibility of variable price and wage markups, our analysis suggests that Hall's residual is also capturing phenomena unrelated to preference shifts. In particular, combining (17) with our definition of the inefficiency gap (1) we have

$$x_t = -(\mu_t^p + \mu_t^w) + \xi_t$$

¹⁵As we noted earlier, Hall does not take the preference shock hypothesis literally. However, since it is a possible alternative to our countercyclical markup interpretation, it is worth investigating this possibility.

¹⁶See also Baxter and King (1991).

Hall's assumption of perfect competition in both goods and labor markets implies $\mu_t^p = \mu_t^w = 0$. This allows him to interpret variable x_t as a preference shock, since under this assumption $x_t = \xi_t$. Notice that under these circumstances the efficiency gap is zero, as there are no imperfections in either goods or labor markets. On the other hand, if preferences are not subject to shocks ($\xi_t = 0$, all t), x_t will purely reflect movements in markups, i.e., $x_t = -(\mu_t^p + \mu_t^w)$. In the latter instance, Hall's residual corresponds exactly to our inefficiency gap, i.e., $x_t = gap_t$, for all t .

Note that if the Hall residual indeed reflects exogenous preference shocks, it should be invariant to any other type of disturbance. In other words, the null hypothesis of preference shocks implies that the Hall residual should be exogenous. We next present two tests that reject the null of exogeneity, thus rejecting the preference shock hypothesis.

First, we test the hypothesis of no-Granger causality from a number of variables to our gap measure. The variables used are: detrended GDP, the nominal interest rate, and the yield spread. Both the nominal interest rate and the yield spread may be thought of as a rough measure of the stance of monetary policy, while detrended GDP is just a simple cyclical indicator. Table 2 displays the *p-values* for several Granger-causality tests. These statistics correspond to bivariate tests using alternative lag lengths. They indicate that the null of no-Granger causality is rejected for all specifications, at conventional significance levels. This finding is robust to reasonable alternative parametrizations of σ and φ . Overall, the evidence of Granger causality is inconsistent with the hypothesis that Hall residual mainly reflects variations in preferences.

As a second test, we estimate the dynamic response of our gap variable to an identified exogenous monetary policy shock. The identification scheme is similar to the one proposed by Christiano et al. (1999), and others. It is based on a VAR that includes measures of output, the price level, commodity prices, and the Federal Funds rate, to which we add our gap measure (or, equivalently, Hall's residual) and the price markup.¹⁷ From the gap and the price markup response we can back out the behavior of the wage markup, using equation (6). We identify the monetary policy shock as the orthogonalized innovation to the Federal Funds rate, under the assumption that this shock does not have a contemporaneous effect on the other variables in the system. Figure 7 shows the estimated responses to a monetary contractionary policy shock. The responses of the nominal rate, output and prices are similar to those found in Christiano et al. (1999), Bernanke and Mihov (1998) and other papers in the literature. Most interesting for our purposes, the inefficiency gap declines significantly in response to the unanticipated monetary tightening. Its overall pattern of response closely mimics the response of output. This endogenous reaction, of course, is inconsistent with the preference shock hypothesis, but fully consistent with our hypothesis that countercyclical markups may underlie the cyclical variation in

¹⁷Here we use the unfiltered wage and price markups to avoid some econometric issues associated with using the filtered data.

the Hall residual. In this respect, note that the tight money shock induces a rise in the wage markup that closely mirrors the decline in the gap, both in the shape and the magnitude of the response. This countercyclical movement in the wage markup is consistent with evidence on unconditional variation presented in Table 1. The price markup also rise, though with a significant lag. Apparently, the sluggish response of wages, which gives rise to strong countercyclical movement in the wage markup, delays the rise in the price markup.¹⁸ In any event, the decline in the inefficiency gap is clearly associated with a countercyclical rise in markups.

To be clear, because preference shocks are not observable, it is not possible to directly determine the overall importance of these disturbances. While our evidence rejects the hypothesis that exogenous preference variation drives all the movement in our gap measure, it cannot rule out the possibility that some of this movement is due to preference shocks. Yet, to the extent that preference shocks are mainly a low frequency phenomenon, as seems plausible under the interpretation that they largely reflect institutional and demographic factors, then they are likely to be captured by the trend component associated with our band-pass filter. In this instance our filtered gap series, which isolates the high frequency movement in this variables, is likely to be largely uncontaminated by exogenous preference variations.

5 Welfare and the Gap

In this section we derive a simple way to measure the welfare costs of business fluctuations based on the inefficiency gap and then apply this methodology to postwar U.S. data. In addition to obtaining a measure of the average cost of business cycles, we also apply our methodology to examining the efficiency losses during particular episodes, including the major postwar recessions.

As we noted in the introduction, our approach differs from Lucas (1987) who considered the costs to risk averse households of the consumption variability associated with the cycle. For roughly the same reason that the baseline neoclassical model has difficulty accounting for the equity premium (i.e., the relatively low variability in aggregate consumption), the Lucas approach suggests very low costs of business fluctuations. For reasonable degrees of risk aversion, Lucas finds that households would be willing to sacrifice less than 0.1 percent of their consumption per period to eliminate fluctuations, clearly a small number. Many papers have extended the Lucas approach, either to allow for incomplete markets or to allow the business cycle to have more persistent effects on consumption variability. These papers also tend

¹⁸As Gali and Gertler (1999) and Sbordone (1999) observe, the sluggish behavior of the price markup helps explain the inertial behavior of inflation, manifested in this case by the delayed and weak response of inflation to the monetary shock. Staggered pricing models relate inflation to an expected discounted stream of real marginal costs, which corresponds to the inverse of the price markup. The sluggish response to the price markup translates into sluggish behavior of real marginal cost.

to find small welfare costs, though with a few exceptions (e.g., Barlevy (2000) and Beaudry and Pages (2001)).

Our approach instead measures the costs stemming from fluctuations in the degree of inefficiency of the aggregate resource allocation, as reflected by the movements in our gap variable.¹⁹ As in Ball and Romer (1987), the cycle generates efficiency costs on average within our framework because the welfare effects of employment fluctuations about the steady state are asymmetric. As Figure 1 illustrates, given that the steady state level of employment is inefficient (due to positive steady state price and wage markups), the efficiency costs of an employment contraction below the steady state will exceed the benefits of a symmetric increase. In particular, note that the vertical distance between the labor demand and supply curves rises as employment falls below the steady state and falls when employment moves above. The quantitative effect of this nonlinearity on the welfare cost of fluctuations ultimately depends on the slopes of the labor demand and supply curves. We show that under reasonable parametrizations it can be quite significant.

We now proceed to derive our welfare measure. First, we assume an equilibrium relationship between output and hours of the form

$$Y_t = F(X_t) N_t^a \quad (19)$$

where Y_t is output, N_t is hours, and X_t captures other factors like technology or capital. One may interpret equation (19) as a reduced form that embeds the possibility of variable capital utilization, as discussed in detail in King and Rebelo (1999). In this respect, note that variable capital utilization will tend to raise the effective output elasticity of employment, implying that a will be a number that exceeds the labor share.

Next, let $U(C_t, N_t)$ denote the period utility function of the representative consumer, where C_t is consumption and N_t is hours. We define the variable Δ_t as the change in utility that results from a one percent increase in hours in period t , i.e., the *marginal* welfare gain from expanding the level of activity in period t . It is possible to express Δ_t as a simple function of the inefficiency gap. In particular, and recalling that $gap_t \equiv \log \frac{MRS_t}{MPN_t}$, we have:

$$\begin{aligned} \Delta_t &\equiv \frac{dU(C_t, N_t)}{d \log N_t} \\ &= (U_{c,t} MPN_t + U_{n,t}) N_t \\ &= U_{c,t} MPN_t N_t (1 - \exp\{gap_t\}) \\ &= a U_{c,t} Y_t (1 - \exp\{gap_t\}) \end{aligned} \quad (20)$$

where $U_{c,t}$ is the utility gain from an additional unit of output (under the assumption that, at the margin, the household is indifferent between consuming and saving), and

¹⁹For a very early attempt to measure the welfare cost of inefficiently high unemployment, see Gordon (1973).

where $MPN_t \equiv a \frac{Y_t}{N_t}$, as implied by (19). It is convenient to express this gain in percentage units of output by dividing Δ_t by $U_{c,t} Y_t$ to obtain

$$\begin{aligned}\delta_t &\equiv \frac{\Delta_t}{U_{c,t} Y_t} \\ &= a (1 - \exp\{gap_t\}) \equiv \delta(gap_t)\end{aligned}\tag{21}$$

As we discussed in section 2, the fact that markups cannot be negative implies that $gap_t \leq 0$.²⁰ It follows from equation (21) that $\delta_t \geq 0$, i.e., the efficiency gain from raising hours is always non-negative in the relevant range. At the first-best (competitive) allocation, $gap_t = 0$, and hence $\delta_t = 0$. At any level of activity below the first best, including the steady state (recall that we assume a steady state with positive markups), $gap_t < 0$ and hence $\delta_t > 0$. Further, in this instance, δ_t varies inversely with gap_t , i.e., as resource allocation converges to the first best, the efficiency gains from further raising employment fall. Intuitively, as employment converges to the competitive equilibrium, the gap between the labor demand and supply curves shrinks, and hence the marginal welfare gains diminish. As a consequence, symmetric fluctuations in gap_t about the steady state lead to asymmetric movements in δ_t , with the marginal efficiency gain from raising employment being greater when gap_t is below its steady state relative to being above an equi-distant amount. We next proceed to show that, due to this asymmetry, efficiency losses are an increasing function of the variability of the gap.

To obtain a measure of the impact on welfare of a deviation of the gap from its steady state value in any given period t , we add up the marginal gains given by equation (21) over the relevant range. Let $\mu \equiv \mu^p + \mu^w$ equal the sum of the steady state price and wage markups, implying that the steady state gap of size $-\mu$ (see equation (7)). Then the welfare effect of a deviation from the gap from steady state in period t is given by the following integral:²¹

$$\begin{aligned}\omega_t &= \int_{-\mu}^{gap_t} \delta(z) dz \\ &= a [\widehat{gap}_t + \exp\{-\mu\} (1 - \exp\{\widehat{gap}_t\})] \equiv \omega(\widehat{gap}_t)\end{aligned}\tag{22}$$

where $\widehat{gap}_t \equiv gap_t - (-\mu) = gap_t + \mu$ denotes the deviation of our gap variable from its value in the steady state.

Equation (22) gives the net efficiency gain (or loss) in percentage units of output associated with deviations of the inefficiency gap from steady state. Note that ω_t has the same sign as \widehat{gap}_t , implying welfare gains if the gap is above steady state and,

²⁰In addition to its plausibility, that assumption would also seem necessary to justify that both firms and workers are always willing to accommodate changes in demand.

²¹Following convention, whenever $gap_t < -\mu$ we have $\int_{-\mu}^{gap_t} \delta(z) dz = -\int_{gap_t}^{-\mu} \delta(z) dz$.

conversely, welfare losses if it is below. In addition, in the range below the first best allocation, ω_t is increasing and concave in \widehat{gap}_t . To see, note that:

$$\omega'(\widehat{gap}_t) = \delta(gap_t) > 0$$

and

$$\omega''(\widehat{gap}_t) = -a \exp\{-\mu + \widehat{gap}_t\} < 0$$

By construction, the derivative of ω_t with respect to the gap is the marginal efficiency gain function δ introduced above, which is positive as long as the economy is below the first best (see equations (20) and (21)). The second derivative, in turn, corresponds to the first derivative of δ with respect to the gap. As we discussed, the latter is negative since the net surplus from subsequent employment increases shrinks as the economy converges to the first best. It is also interesting to observe that a rise in the average steady state markup reduces the absolute value of the second derivative, and hence the degree of curvature of ω . Intuitively, the gains and losses from fluctuations in the gap about the steady state are more symmetric, the higher is the steady state markup. Figure 8 displays the function $\omega(\widehat{gap}_t)$ for three different steady markup levels, and illustrates how the degree of concavity is decreasing in the average markup μ .

The concavity of $\omega(\widehat{gap}_t)$, of course, implies that symmetric fluctuations in the gap have asymmetric welfare effects, which in turn implies that variability in the gap reduces welfare. To gain a sense of the quantitative importance of that effect we construct a quarterly time series of ω_t based on equation (22). The gap variable we use as input is our baseline measure, constructed using $\sigma = 1$ and $\varphi = 5$. In addition, we set the elasticity of output with respect to hours, a , at unity, in line with the evidence (e.g., King and Rebelo, 1999). Figure 9 plots the resulting time series of ω_t over the sample 1960:IV-99:IV. The realization for each time t is interpretable as the efficiency gain or loss in percentage units of output associated with the deviation of the inefficiency gap from its steady state in that period. As the figure shows, sharp efficiency losses arise in recessions that do not appear to be offset by commensurate gains during booms. Note also that the efficiency losses are particularly large during the major recessions, ranging between five and nine percent of output per period around the time of the respective troughs.

We are now in a position to derive a measure of the average welfare effects of business fluctuations. To facilitate comparison with the literature, we express the cost of business cycles as a per period percent loss in trend consumption. Accordingly, to obtain a quantitative assessment of the *average* welfare effect of fluctuations over an arbitrary sample period $t = 0, 1, \dots, T$, we compute the following “adjusted” sample mean of $\{\omega_t\}$:

$$\bar{\Omega} = \frac{\Phi}{T} \sum_{t=0}^T \omega(\widehat{gap}_t) \quad (23)$$

where Φ denotes the average (steady state) ratio of GDP to (nondurables and services) consumption. Equation (23) gives the average efficiency cost per period of business fluctuations over the sample as its equivalent permanent percent change in consumption.

To gain some intuition about the factors that determine the magnitude of the welfare effects of fluctuations in the inefficiency gap, it is useful to consider the “population” analogue of equation (23). It follows from equation (22) that we can express the unconditional expected welfare cost, Ω , after adjusting by Φ , as

$$\Omega = a\Phi [E\{\widehat{gap}_t\} + \exp\{-\mu\} (1 - E\{\exp(\widehat{gap}_t)\})] \quad (24)$$

A second order approximation about the steady state yields²²

$$\begin{aligned} \Omega &\simeq a\Phi \exp\{-\mu\} \left(1 - \exp\left\{\frac{1}{2} \sigma_{gap}^2\right\}\right) \\ &\simeq -\frac{a\Phi}{2} \exp\{-\mu\} \sigma_{gap}^2 < 0 \end{aligned}$$

where σ_{gap}^2 is the variance of the inefficiency gap. Notice that, as a result of the concavity of ω , the expected welfare effects of fluctuations in the gap variable are negative, i.e. these fluctuations imply *losses* in expected welfare. That loss, further, is of “second order,” as it is linearly related to the variance of the inefficiency gap. It is, however, potentially large, depending on the values of parameters Φ , a , and μ and, in particular, σ_{gap}^2 . As section 3 suggests, σ_{gap}^2 is potentially large if labor supply is relatively inelastic or risk aversion is relatively high.

To be clear, our approach provides a lower bound on the measure of the total welfare costs of fluctuations. The reason is simple: it does not include the welfare costs from *efficient* fluctuations in consumption and employment. To see that suppose that the data were generated by a real business cycle model with frictionless, perfectly competitive markets. We should then expect to see no variation in our gap measure, as the resource allocation would always be efficient. Our metric would then indicate no welfare costs of fluctuations, while some losses would still be implied by the variability of consumption and leisure (under standard convexity assumptions on preferences).

5.1 Evidence

Table 3 presents measures of the welfare costs of economic fluctuations in the postwar U.S. for alternative values of parameters φ , σ , and μ . Recall that φ corresponds to the inverse of the Frisch elasticity of labor supply, whereas σ is the inverse of the intertemporal elasticity of substitution. We consider three values of each of these parameters:

²²The approximation is exact under the assumption that $\{\widehat{gap}_t\}$ follows a *normal* distribution with zero mean and variance σ_{gap}^2 .

1, 5, and 10, implying that each of the corresponding elasticities ranges from 1.0 to 0.1. For the parameter μ , the sum of the steady state wage and price markups, we consider values of 0.15, 0.25, and 0.40, which we think of as falling within a plausible range. (In addition, consistent with our assumption that the economy does not produce above the first best, we only report results for parameter configurations where δ_t , the marginal efficiency gain from raising employment, is always non-negative).

For a parametrization that corresponds to our baseline case of section 3 augmented with the intermediate value of the markup ($\sigma = 1$; $\varphi = 5$; $\mu = 0.25$), we estimate the welfare cost of postwar U.S. business fluctuations to be roughly 0.8 percent of consumption, a number considerably higher than Lucas' estimate of less than 0.1 percent. If we reduce the intertemporal elasticity of substitution from unity to a more empirically reasonable value of 0.2 (i.e., if we raise σ from 1.0 to 5.0) the estimated cost rises to 1.35 of consumption, respectively.

Overall, the estimated welfare costs are highly sensitive to the labor supply elasticity and the intertemporal elasticities of substitution. For high values of the two elasticities (corresponding in our case to $\sigma = 1$ and $\varphi = 1$), the welfare costs are small, on the order of Lucas' estimates. In this case, roughly speaking, the labor supply curve is relatively flat, implying small cyclical fluctuations in the inefficiency gap (see Figure 4.) At the other extreme, the case of low elasticities (corresponding in our case to $\sigma = 10$ and $\varphi = 10$), the cost becomes huge; roughly 4.6 percent. In this instance, the labor supply curve is very steep, implying very large fluctuations in the inefficiency gap.

Because there is considerable uncertainty over the appropriate parametrization of choices of σ and φ , we cannot say with any degree of precision what are the true efficiency costs of business fluctuations, other than to point out the critical role of these parameters. On the other hand, as we discussed in section 3, there is a considerable body of evidence suggesting that our intermediate value of 5.0 for each parameter provides a reasonable benchmark. As Table 3 indicates, under this benchmark, the welfare costs are far from negligible.

Our measure of the average cost of business cycles obscures the fact that individual recessionary episodes can be quite costly. What moderates the impact of these episodes on the overall welfare measure is the fact they have been relatively infrequent, particularly over the last several decades. One reason for this may be that stabilization policy has been reasonably effective. Another possibility is that the economy has been subject to smaller shocks. In either event, it is of interest to examine efficiency losses during downturns. Doing so provides a sense of the gains from avoiding future recessions (either by good policy or by good luck.)

We accordingly consider the four episodes in our sample where the economy experienced a boom followed by recession.²³ In each instance we measure the boom as the period where the gap variable climbs above zero up to the point where it returns to zero. The recession is the period that follows, where the gap turns negative up to

²³We combine the 1980 and 1981-82 recessions into a single episode.

the time it returns to steady state. For each episode, Table 4 reports the efficiency gains from the boom and the costs from the recession, followed by the net loss (the sum of the first two columns.) We measure the gains and losses as a percent of one year's consumption. Also, we restrict attention to our baseline case, with $\sigma = 1$ and $\varphi = 5$.

In each episode, the costs of the recession outweigh the benefits of the boom, as we might expect, given the asymmetric nature of the cyclical efficiency gains and losses. In addition, the net cost of the cycle (last column) as well as the gross cost of the recession (2nd column) is large in each case, particularly so during the major recessions. For example, the gross efficiency loss of the 1974-75 recession was equivalent to 22.2 percent of one year's consumption, while the net cost after deducting the gains from the preceding boom was 14.9 percent. For the 1980-82 recession, the gross and net costs were 16.8 percent and 8.1 percent, respectively.

Finally, we observe that our calculation ignores the costs of inflation associated with the economy moving above steady state output. For this reason, our metric may overstate the gains from booms. To the extent, the costs of inflation offset the efficiency gains from the boom, our measure of the gross efficiency loss of the recession may provide a more accurate indicator of the costs of these episodes.

6 Concluding Comments

At the risk of considerable oversimplification, it is possible to classify modern business cycle models into two types. The first class attempts to explain quantity fluctuations by appealing to high degrees of intertemporal substitution in an environment of frictionless markets. The second instead appeals to countercyclical markups owing to particular market frictions. Within this latter approach, large quantity fluctuations are possible even when intertemporal elasticities of both labor supply and consumption are low. Perhaps a central message of this paper is that the issue of the welfare costs of business cycles cannot be cleanly separated from the issue of which of these business cycle paradigms provides a better description of actual economic fluctuations. We find that with high degrees of intertemporal substitution, the costs of business fluctuations are relatively small, which perhaps should not be surprising since labor supply curves are relatively flat in this setting. On the other hand, with low substitution elasticities (implying that strongly countercyclical markups are needed to explain the data), we find significant welfare costs associated with cyclical fluctuations in the inefficiency of resource allocation. To be sure, the appropriate parametrization of these intertemporal elasticities remains an open question. For the time being, though, we note that as discussed in section 3, there is a considerable body of evidence consistent with the low intertemporal elasticities that we stressed in our analysis.

Finally, we emphasize that our estimates of the efficiency costs of business fluctuations are likely to be conservative because they do not take into account the welfare

costs of inflation variability that may be associated with cyclical fluctuations. Recent work by Woodford (1999) and others suggests that these efficiency costs may be highly significant. Accounting for this factor in our overall welfare measure is something we plan for future research.

Table 1. Basic Statistics: 1960-1999*Baseline Calibration ($\sigma = 1, \varphi = 5$)*

<i>Variable</i>	s.d.(%)	ρ	Correlation			
			GDP	Gap	Price Mkup	Wage Mkup
GDP	2.2	0.91	1			
Gap	10.2	0.91	0.81	1		
Price Markup	4.3	0.96	0.02	0.22	1	
Wage Markup	11.2	0.92	-0.81	-0.92	-0.41	1

Table 2. Granger Causality Tests (1960-1999)

<i>Baseline Calibration</i> ($\sigma = 1$, $\varphi = 5$)			
<i>Bivariate VAR</i>			
<i>Variable</i>	4-lags	5-lags	6-lags
Detrended GDP	0.000	0.000	0.000
Nominal Interest Rate	0.000	0.000	0.000
Yield Spread	0.003	0.004	0.008

Note: The values reported are p-values for the null hypothesis of no Granger causality from each variable listed to Hall's (F-test).

Table 3. The Welfare Costs of Fluctuations
(percent average consumption)

	$\varphi = 1$	$\varphi = 5$	$\varphi = 10$
$\sigma = 1$			
μ			
0.15	0.064	0.819	-
0.25	0.063	0.774	2.753
0.40	0.061	0.715	2.470
$\sigma = 5$			
μ			
0.15	0.373	-	-
0.25	0.365	1.348	-
0.40	0.355	1.243	3.285
$\sigma = 10$			
μ			
0.15	1.192	-	-
0.25	1.135	-	-
0.40	1.058	2.245	4.652

Note: Calibration $a = 1$. Sample Period: 1960:1/1999:4. The average output consumption ratio is 1.71. The welfare cost have been calculated only if $\delta_t > 0$, otherwise we use '-'.

Table 4. Costs(-) and Benefits(+) of Boom/Recession Episodes
(percent of one year's consumption)

<i>Time Period</i>			<i>Boom</i>	<i>Recession</i>	<i>Net</i>
Start	Turning Point	End			
68:2	70:2	72:3	6.5	-9.4	-2.9
72:4	74:4	77:3	7.3	-22.2	-14.9
77:4	80:2	83:4	8.7	-16.8	-8.1
87:4	90:4	94:1	9.3	-14.8	-5.5

Note: Baseline Calibration $a = 1$, $\sigma = 1$ and $\varphi = 5$. Sample Period: 1960:1-1999:4.

References

- [1] Alexopolous, Michelle, (2001), "Unemployment and the Business Cycle," mimeo, University of Toronto.
- [2] Altonji, Joseph G., (1986), "Intertemporal Substitution in Labor Supply: Evidence from Micro Data," *Journal of Political Economy*, 94, June, S176-S215.
- [3] Ball, Lawrence and David Romer, (1989), "Are Prices Too Sticky?" *Quarterly Journal of Economics*, 1987, 104, 507-524.
- [4] Barlevy, Gadi, 2000, "Evaluating Costs of Business Cycles of Business Cycles in Models of Endogenous Growth", mimeo.
- [5] Barsky, Robert, Jonathon Parker, and Gary Solon,(1994), "Measuring the Cyclicity of Real Wages: How Important is Composition Bias?," *Quarterly Journal of Economics* 109, 1-25.
- [6] Barsky, Robert, F.Thomas Juster, Miles S. Kimball and Matthew D. Shapiro, (1997), "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," *Quarterly Journal of Economics*, 112, 537-579.
- [7] Basu, S. and M. Kimball (1997), "Cyclical Productivity with Unobserved Input Variation", NBER WP 5915.
- [8] Baxter, M. and Robert King (1991), "Productive Externalities and Business Cycles", Discussion Paper no 53. Federal Reserve Bank of Minneapolis, Institute for Empirical Macroeconomics, November 1991.
- [9] Beaudry, Paul and Carmen Pages (2001): "The Cost of Business Cycle and the Stabilization Value of Unemployment Insurance," *European Economic Review* 45, no. 8, 1545-1572.
- [10] Bernanke, Ben and Ilyan Mihov (1998), "Measuring Monetary Policy," *Quarterly Journal of Economics*, CXIII, 869-902.
- [11] Blanchard Olivier, J. and Nobuhiro Kiyotaki, (1987), "Monopolistic Competition and the Effects of Aggregate Demand," *American Economic Review*, LXXVII, 647-66.
- [12] Card, David, (1994), "Intertemporal Labor Supply: An Assessment," in *Frontiers of Econometrics*, Christopher Sims, editor.
- [13] Christiano, Lawrence, Martin Eichenbaum and Charle Evans, (1997), "Sticky Prices and Limited Participation Models: A Comparison," *European Economic Review*, 41, 1201-1249.

- [14] Christiano, Lawrence, Martin Eichenbaum and Charle Evans, (1999), "Monetary Policy Shocks: What Have we Learned and to What End?," *Handbook of Macroeconomics*, Vol.1A,.Amsterdam North-Holland.
- [15] Christiano, Lawrence, Martin Eichenbaum and Charle Evans, (2001), "Nominal Rigidities and the Effects of a Shock to Monetary Policy." mimeo Northwestern University.
- [16] Cooley, Thomas and Edward Prescott, (1995) "Economic Growth and Business Cycles," in *Frontiers of Business Cycle Research*, Thomas Cooley (editor), Princeton University Press.
- [17] Eichenbaum, Martin, Lars Hansen and Kenneth Singleton, (1988): "A Time Series Analysis of Representative Agent Models of Consumption and Leisure Choice under Uncertainty," *Quarterly Journal of Economics* , 103 (1), February 1988, 51-78.
- [18] Erceg C., Dale Henderson, and Andrew Levin, (2000), "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics*, 46(2), 281-313.
- [19] Galí, J. and Mark Gertler (1999), "Inflation Dynamics: A Structural Econometric Analysis", *Journal of Monetary Economics*, 44, 195-222.
- [20] Galí, J., Mark Gertler and J.David López-Salido (2001), "European Inflation Dynamics", *European Economic Review*, 45 (7), June 2001, 1121-1150.
- [21] Gordon, Robert J. (1973), "The Welfare Cost of Higher Unemployment," *Brookings Papers on Economic Activity* no. 1, 133-195. .
- [22] Hall, Robert, E., (1988), "Intertemporal Substitution in consumption," *Journal of Political Economy*, 96 (2), 339-357.
- [23] Hall, Robert, E. (1997): "Macroeconomic Fluctuations and the Allocation of Time", *Journal of Labor Economics*, vol.15,no.1,pt.2. 1997, S223-S250.
- [24] King, Robert and Sergio Rebelo, (1999):"Resuscitating Real Business Cycles", *The Handbook of Macroeconomics*, Vol 1B, Amsterdam North-Holland.
- [25] Lucas, Robert E., (1987), *Models of Business Cycles*
- [26] MacCurdy, Thomas, (1981), "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy*, 89, December, 1059-85.
- [27] Mulligan, Casey B.(1998), "Substitution Over Time: Another Look at Life-Cycle Labor Supply," *NBER MacroAnnual*, vol.13, Ben Bernanke and Julio Rotemberg, editors.

- [28] Pencavel, John, (1998), "The Market Work Behavior and Wages of Women, 1975-94," *Journal of Human Resources*, 33, 771-804.
- [29] Pencavel, John, "A Cohort Analysis of the Association Between Work Hours and Wages Among Men, *Journal of Human Resources*, forthcoming.
- [30] Rotemberg, Julio, and Michael Woodford (1999), "The Cyclical Behavior of Prices and Costs," *The Handbook of Macroeconomics, Vol 1B*, Amsterdam North-Holland.
- [31] Sbordone, Argia, M. (1999), "Prices and Unit Labor Costs: A New Test of Price Stickiness". mimeo, Rutgers University.
- [32] Sbordone, Argia, M., (2000), "An Optimizing Model of Wage and Price Dynamics," mimeo Rutgers University.
- [33] Staiger, Douglas, James Stock and Mark Watson (2001), "Prices, Wages and the U.S. NAIRU in the 1990s," mimeo.
- [34] Woodford, M. (1999): "Inflation Stabilization and Welfare" NBER Working Paper # 5684.

Figure 1. The Gap: A Diagrammatic Exposition




Figure 2. The Gap
Baseline Calibration ($\sigma=1, \phi=5$)




Figure 3. The Gap and the Wage Markup
Baseline Calibration





Figure 4. The Gap under Alternative Calibrations

a. High Risk Aversion ($\sigma=5$)



b. High Labor Supply Elasticity ($\phi=1$)




Figure 5 . The Gap under Alternative Marginal Cost Measures
Baseline Calibration




Figure 6. The Gap under Non-Separable Preferences
Parameter Values: $\sigma=1$, $\varphi=5$, $\beta=0.99$, $b=0.8$




Figure 7. Dynamic Effects of Monetary Policy Shocks
Baseline Calibration

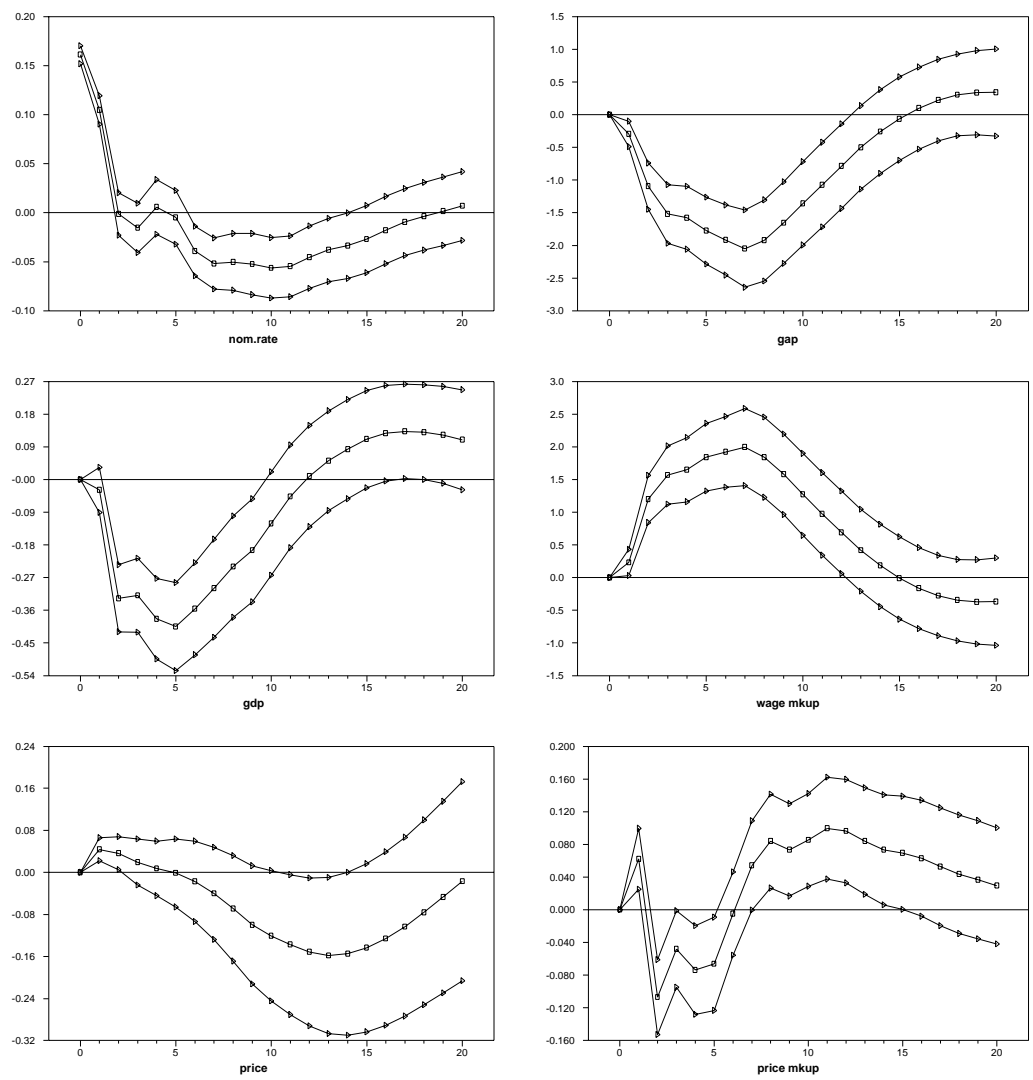


Figure 8. Welfare and the Gap




Figure 9. The Welfare Effects of Postwar U.S. Fluctuations
Baseline Calibration ($\sigma=1, \varphi=5, \mu=0.25$)

