FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Output and Unemployment Dynamics

Mary C. Daly Federal Reserve Bank of San Francisco

John G. Fernald Federal Reserve Bank of San Francisco

Òscar Jordà Federal Reserve Bank of San Francisco

Fernanda Nechio Federal Reserve Bank of San Francisco

November 2014

Working Paper 2013-32 http://www.frbsf.org/publications/economics/papers/2013/wp2013-32.pdf

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Board of Governors of the Federal Reserve System.

Output and Unemployment Dynamics^{*}

Mary C. Daly

John G. Fernald Öscar Jordà

Fernanda Nechio

Federal Reserve Bank of San Francisco

November, 2014

Abstract

The evolution of the secular and business-cycle comovement between different components of the production function and unemployment, Okun's law, provides important stylized facts for macro modelers. We show that total hours worked adjust two-to-one to changes in the unemployment rate. The cyclicality of productivity has changed over time and as a function of the type of shock hitting the economy. Even the responses of different margins to shocks vary over time. We document these and other features of the data using the growth-accounting decomposition in Fernald (2014).

JEL classification codes: E23, E24, E32, J20

Keywords: growth accounting, output and employment fluctuations, cyclical productivity, Okun's Law

^{*}We thank Valerie Ramey and participants at the SED 2013, 2013 NBER Summer Institute, and the 2014 European Economic Association meetings for helpful comments on this paper. We also thank Bart Hobijn, Ron Smith, and seminar participants at the Bank of England for helpful comments on an earlier and much different draft. Last but not least, we thank Daniel Molitor and Kuni Natsuki for excellent research assistance. This paper was previously circulated under the title "Okun's Macroscope and the Changing Cyclicality of Underlying Margins of Adjustment." The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of San Francisco or the Federal Reserve System. E-mails: mary.daly@sf.frb.org; john.fernald@sf.frb.org; oscar.jorda@sf.frb.org; fernanda.nechio@sf.frb.org. Corresponding author: John Fernald, email: john.fernald@sf.frb.org.

1 Introduction

The relationship between output and unemployment is among the most important in studied relationships in macroeconomics. In 1962, Arthur Okun proposed that this key relationship could be easily summarized in a simple regression relating the change in the unemployment rate to the change in real output growth. Specifically, Okun estimated, using data from 1947-1960, that a 2% increase in output corresponds to a 1% decline in the rate of cyclical unemployment. According to his estimates, this aggregate relationship represented associated, underlying changes in labor force participation (0.5% increase), hours worked per employee (0.5% increase), and labor productivity (1% increase).¹

Since 1962, various researchers have reevaluated Okun's original output/unemployment estimate and found that it has remained remarkably stable over time. This stability combined with the simplicity of a single summary number has made Okun's insight, often termed as Okun's Law, a guidepost if not the foundation for many central bank, government, and private forecasting models, and also a standard reference in textbooks.²

Even more remarkably perhaps, research investigating the stability of this relationship in the aftermath of the Great Recession has uncovered little evidence of a new Okun relationship (Ball, Leigh and Loungani 2013, and Daly et al. 2014). Importantly, the stability of the reduced form relationship is frequently assumed to imply that the underlying relationships have also remained the same. This is a point we come back to in our paper.

Given the enormous changes that have occurred in the economy over the past 50 years, the relatively stability of the Okun observation is surprising and an empirical puzzle. As Okun laid out in his original paper, the relationship he estimated, between output and unemployment, depends on a number of underlying relationships that are themselves affected by institutions, cyclical fluctuations and economic shocks. Over the last 50 years, the profile of recessions and recoveries has changed, the labor market has become much more flexible and the mix of shocks hitting the economy has changed. Moreover, the response to shocks by the central bank and the fiscal authority has changed.

In this paper, we examine and refine the stylized fact that Okun's relationship has remained stable over time. To do this we first link Okun's reduced form empirical observations to a growth accounting/production-theory framework consistent with modern macro modeling. Using this framework we show how the average relationship between output and unemployment has remained stable over time, despite considerable changes in the economy. We investigate the properties of the rela-

¹Okun documented this relationship in studies focused on measuring potential output (Okun 1962, 1965). His original work used data from 1947-1960. He regressed unemployment change on real GNP growth and reported that a 1 percent increase in GNP was associated with a 0.30 percentage point decrease in unemployment. As Plosser and Schwert (1979) point out, the coefficient on the reverse regression – i.e., the expected change in output from a 1 percentage point increase in unemployment – is not the inverse of 0.3, but depends on the R^2 . Using Okun's original estimates, where the R^2 is 0.62, implies that a 1 percentage point increase in unemployment predicts a 2.1 percent decline in output.

²Okun's Law is used as a forecasting guidepost by many central banks and the U.S. Congressional Budget Office.

tionships underlying the Okun's Law over time, the business cycle, and as a function of the type of shock using the novel quarterly growth-accounting dataset for the U.S. business sector from Fernald (2014).

We show that the relationships underlying output and unemployment fluctuate considerably over time but for the most part these changes have been offsetting such that output and unemployment itself is rather stable. In particular, changes to the response of number of employees, capital utilization and technology to shocks have offset each other such that the Okun coefficient has remained nearly constant in our sample.

Our findings have implications for macro modeling and forecasting. Beneath the stability of output and unemployment dynamics hide important fluctuations in the underlying margins of adjustment in the economy. More specifically, the empirical estimates that we provide reveal several features of the data that have important implications for a wide variety of recent macroeconomic models. First, the data suggest a larger elasticity of hours worked to the unemployment rate. The response we measure is roughly two-to-one, rather than one-to-one, as typically assumed in dynamic stochastic general equilibrium (DSGE) models with unemployment (e.g., Galí, Smets, Wouters 2011 or Christiano, Eichenbaum, and Trabandt, 2013). This finding suggests that in addition to the direct channel (unemployment rises because fewer people are employed and hours and unemployment move oneto-one), other margins may also adjust, such as labor force participation, multiple-job holders, and hours per worker.

A second empirical result highlights the changing cyclicality of labor productivity, as noted by other authors (e.g., Stiroh 2009; Galí and Gambetti 2009; Barnichon 2010; and Galí and van Rens 2010). In the decades prior to the Great Recession, productivity shifted from procyclical to counter-cyclical with respect to unemployment. Since the mid-2000s, the strong countercyclicality has largely disappeared. In our data, the pre-Great recession shift in the cyclicality of labor productivity mainly reflects a sharply reduced role for variations in factor utilization in the 1990s.

Third, we find that productivity tends to be more procyclical in recessions (with the exception of the Great Recession in which financial factors may have played an outsize role) than in normal times. The source of this change in cyclicality depends on the type of shock hitting the economy. Specifically, we investigate an exogenous monetary contraction as measured by Christiano, Eichenbaum and Evans (1999); an exogenous oil price shock as in Hamilton (2003); and shocks to total factor productivity in the investment and consumer goods sectors based on Fernald (2014). For most of these shocks, changes in factor utilization appear to be the main explanation for the change in cyclicality.

Fourth, we investigate the stability of how margins of adjustment respond to each of these shocks over time. For example, monetary policy shocks had large effects on output early in the sample but their effect has faded considerably over time. Oil and TFP shocks have had more negative and longer lasting effects on total hours over time. The timing of this shift coincides with the changing cyclicality in productivity.

Fifth, despite the sensitivity of the margins of adjustment to the type of shock in the short-run, we find that the loadings on each margin remain relatively constant in the long-run. A decomposition of the response of the margins of adjustment by type of shock over the long-term does not reveal substantial differences with respect to the benchmark static results.

The relative stability of this relationship, documented in Ball, Leigh and Loungani (2013) and that we also confirm, hides considerable time variation in its components. Fluctuations in one margin tend to be compensated by fluctuations in another margin working across purposes so that the net effect masks the underlying currents. Our results suggest that while this is indeed a good approximation in the long-run, the dynamic behavior of the different margins of adjustment in the short-run is quite sensitive to the shocks hitting the economy and the stage of the business cycle in which they occur. The policy response has to therefore be appropriately modulated to account for these sensitivities.

2 Margins of adjustment, growth accounting and Okun's Law

Business cycle fluctuations cause firms and workers to adapt through different margins, such as adjusting the workweek, staffing levels, resource intensity, and so on. At the aggregate level, fluctuations in output and un/employment are routinely used to assess the macroeconomic implications of these fluctuations and as a platform to inform policy decisions.

Okun's law, as the relationship between output growth and changes in un/employment is called since it was introduced by Arthur Okun in (1962), is a key ingredient of the policy debate because it neatly summarizes the contributions of the individual margins of adjustment into one coefficient. Therefore, we use Okun's law as the foundation from which to examine the broad contours of the margins of adjustment in U.S. data.³

Economic theory determines our specification of Okun's law by linking together the modern macro literature and a large literature on measurement and production theory. In particular, the growthaccounting perspective of this section implies a relationship between output growth and input growth – and changes in the unemployment rate are closely related to input growth. Consider the relationship between output growth and the change in the unemployment rate given by:

$$\Delta y = \mu + \beta \Delta U + \varepsilon, \tag{1}$$

where lower-case letters denote log-levels, and Δ denotes the time difference operator. Hence Δy is the growth rate (log change) in real output and ΔU is the change in the rate of unemployment. This

³Prachowny (1993) also links Okun's Law to production theory but not through a growth-accounting decomposition.

is one of the classical specifications of Okun's law.

We find it useful to decompose Δy using several accounting identities and a production-theory view for some of the components. First, note that output growth can be written as the sum of growth in total hours, Δl , and labor productivity, $(\Delta y - \Delta l)$:

$$\Delta y \equiv \Delta l + (\Delta y - \Delta l) \,. \tag{2}$$

Furthermore, total hours can be decomposed into the sum of growth in hours per worker, Δh and total workers, Δn , so that expression (2) becomes:

$$\Delta y \equiv \Delta h + \Delta n + (\Delta y - \Delta l). \tag{3}$$

Next, we decompose productivity based on the work of Basu, Fernald and Kimball (2006), henceforth BFK. In particular, we specify a production function that controls for non-technology factors in a manner that is more detailed than usually specified in typical macroeconomic models:

$$Y = F\left(W \times K, L \times Q \times E, A\right),$$

where Y is output, W is the workweek of capital (the number of hours the capital is actually in operation), and K is the stock of capital. The effective labor input depends on hours L; the average "quality" of each hour (including age, experience, and other observables), Q; and the effort E per quality-adjusted hour. Note that variation in utilization shows up in W and labor hoarding shows up in E. A is technology.

In the more general case in which the production function takes the translog form it provides a flexible second-order approximation to any function. In that case, the factor shares/output elasticities α and $(1 - \alpha)$ are time-varying and are properly taken as the average of shares between one period and the next. Basu and Fernald (2001) show that allowing for time-variation in factor shares has little impact on the analysis that we pursue, however.⁴ For this reason we prefer to proceed with the more straightforward Cobb-Douglas case, where the shares are constant over time.⁵

Under this assumption, this production function takes the form (in growth rates):

$$\Delta y = \alpha \left(\Delta k + \Delta w \right) + (1 - \alpha) \left(\Delta l + \Delta e + \Delta q \right) + \Delta a, \tag{4}$$

⁴Basu and Fernald (2001) discuss the more general case in which an aggregate constant-returns production function may not exist and how, in practice, the effects are likely to show up as procyclical movements in the cyclicality of the aggregate Solow residual (measured TFP, the empirical counterpart of A).

⁵Recent research by Elsby, Hobijn and Şahin (2013) discusses the recent decline of the labor share in the last twenty years. This observation has attracted attention because of the implications, among other things, for inequality, the equilibrium long-term interest rate, and asset price valuations. However, none of these factors are central to the analysis that we pursue.

where, again, we use lower case to indicate the logs of the variables and Δ to denote first differences.

Expression (4) can be rearranged as follows:

$$\Delta y - \Delta l = \alpha \left(\Delta k - \Delta l \right) + (1 - \alpha) \Delta q + (\alpha \Delta w + (1 - \alpha) \Delta e) + \Delta a$$

$$\equiv \alpha \left(\Delta k - \Delta l \right) + (1 - \alpha) \Delta q + \Delta v + \Delta a.$$
(5)

Therefore, labor productivity, $(\Delta y - \Delta l)$, can change because of capital-deepening, given by $\alpha (\Delta k + \Delta l)$, labor quality, given by $(1 - \alpha) \Delta q$, cyclical variations in utilization, given by $\Delta v \equiv \alpha \Delta w + (1 - \alpha) \Delta e$, or technology, Δa .

We define the standard measure of total factor productivity (TFP) growth, Δz , as output growth that is not explained by observed growth in inputs. This is given by the following expression:

$$\Delta z = \Delta y - \alpha \Delta k - (1 - \alpha) \left(\Delta l + \Delta q \right).$$

Using expression (5), we can write this as $\Delta z = \Delta v + \Delta a$. We will also refer to the empirical counterpart to Δa as "utilization-adjusted TFP" (as a reminder that it is technology only under the conditions that there is a constant returns aggregate production function).⁶

Summarizing, we now combine expressions (3) and (5) as stand-ins for the term Δy in expression (1) in order to obtain a decomposition of the Okun coefficient β . This is done to capture the contribution of each of the margins of adjustment. Specifically,

$$\Delta y = \underbrace{\Delta h + \Delta n}_{\Delta l} + \underbrace{\alpha \left(\Delta k - \Delta l\right) + (1 - \alpha) \,\Delta q + \Delta \upsilon + \Delta a}_{(\Delta y - \Delta l)}.\tag{6}$$

The Okun coefficient β in expression (1) can now be decomposed using expression (6) to gauge the relative contribution of each margin of adjustment. That is, the ordinary least squares (OLS) estimate $\hat{\beta}$ in expression (1) turns out to be the sum of the projection coefficients of each of the terms in expression (6) on to ΔU . For example, consider the broad decomposition of output growth, Δy , into total hours, Δl and productivity, $(\Delta y - \Delta l)$, as in expression (2). In the regression of Δy on to ΔU , $\hat{\beta}$ is the ratio of $cov(\Delta y, \Delta U)$ to $var(\Delta U)$. Then notice that $cov(\Delta y, \Delta U)$ can be expressed as the sum of $cov(\Delta l, \Delta U)$ and $cov(\Delta y - \Delta l, \Delta U)$. Using similar arguments, we can think of the contributions of each of the components in expression (6).

We reflect on the economic implications of the decomposition to provide a guide to the results reported in the next section. Consider the components of total hours worked, Δl , first. Other things equal, a one percentage point increase in the unemployment rate tends to reduce employment and

⁶BFK find that the cyclical properties of the Solow residual can be quite different when one controls for nontechnological factors, such as utilization.

hours per worker by a little over one percent. When unemployment rises, we would expect hours per worker to fall. In regard to employment, we note that multiple job holders in such situations lose second or third jobs, and labor-force participation might fall (reflecting a shift towards home production or, for other reasons, a reduction in labor force attachment such as we have seen since the Great Recession).⁷ Therefore we expect the employment projection to exceed unity in absolute value.

Next, consider the productivity term in expression (6). Capital deepening and labor quality tend to respond countercyclically. For example, in recessions unemployment rises and hours worked fall. Since capital is relatively smooth, capital deepening tends to rise. This capital-deepening effect, which reflects the diminishing returns to labor alone, pushes labor productivity up. The labor quality term has a similar effect. Since firms disproportionately retain the more skilled workers in recessions, labor quality tends to rise.

On the other hand, declining utilization in recessions pushes measured productivity down. When unemployment is high, firms hoard labor and reduce the workweek of capital (e.g., going from two shifts a day to one).⁸ This tends to affect the Okun coefficient β negatively.

Finally, the effects of utilization-adjusted TFP growth, Δa , are theoretically ambiguous but are often estimated to be positive. In traditional real-business-cycle-type models, positive technology shocks not only raise labor productivity but would typically be expected to reduce unemployment. Procyclical productivity affects the Okun coefficient negatively. In models with nominal or real rigidities, however, labor productivity and unemployment may be positively correlated conditional on a technology shock. Galí (1999), Francis and Ramey (2005), BFK and others argue that this may be the empirically relevant case.

A different way to gain further intuition is to consider what would happen in a simplified economy with constant-returns, where unemployment changes did not generate systematic changes in hours per worker, and with no changes in the employment-worker gap, labor-force participation, or population. Moreover, suppose that unemployment and technology are not systematically related, as would be the case with a demand shock. It is easy to show that the Okun coefficient is the negative of the labor's share in income, $-(1 - \alpha)$.

Specifically, $\Delta l = -\Delta U$ and the contribution of total hours to the Okun coefficient is normalized

⁷Variations in the number of workers naturally affect the unemployment rate through the following relationship: $N = \frac{N}{Emp} \frac{Emp}{LabForce} \frac{LabForce}{Pop} Pop = \frac{N}{Emp} (1-U) \frac{LabForce}{Pop} Pop$, where N is the number of workers, Emp is the number of people employed, LabForce is the labor force, and Pop is the overall working-age population so that the ratio is the labor force participation rate. The first term on the right-hand side, $\frac{N}{Emp}$, reflects the fact that the number of workers is potentially different from the number of people employed. The second term is employment as a share of the labor force, which is by definition equal to (1-U). This decomposition is similar to Gordon (2011).

⁸Basu and Fernald (2001) discuss the importance of procyclical fluctuations in utilization margins in productivity measurement.

to be one. If at the same time, capital, labor quality, and utilization do not change systematically with unemployment, then, from equation (5), $\Delta y - \Delta l = -\alpha (\Delta l) = \alpha \Delta U$ and hence the contribution of productivity to the Okun coefficient is α . Therefore, the Okun coefficient becomes the negative of the labor's share in income or $\beta = -(1 - \alpha)$, as described above. This result says that a one percentage point increase in unemployment reduces labor hours by one percent, while leaving all other inputs and technology unaffected. Therefore output falls by the labor share.

In practice, as Okun (1962) showed, the magnitude of Okun's coefficient is substantially larger than the labor share. The decomposition shows that this larger coefficient necessarily reflects the systematic cyclicality of other margins of adjustment, as reflected in (4) and (5). The components of these equations point us where to look for potential answers, something we investigate in the next section.

3 Margins of adjustment: stylized facts

In this section we turn our focus to the decomposition based on the margins of adjustment introduced above. First we report the broad contributions of each margin to the overall understanding of the comovements between output and unemployment. Next we investigate the stability of these relations over time, both over the medium run and at business cycle frequencies. The following section investigates the differences in the margins of adjustment caused by the type of shock hitting the economy and relates back to the business cycle results reported in this section.

3.1 Data

Okun (1962) took a "leap from the unemployment rate to potential output rather than [taking] a series of steps involving several underlying factors," because he was limited by the data. We overcome these data limitations by using relatively new, detailed quarterly growth-accounting data for the U.S. business sector from Fernald (2014). Our dataset runs from 1947Q2 through 2014Q1. Although the unemployment rate available corresponds to that in the overall economy, it does not pose a problem because most of the economy's cyclicality arises from the business sector. The Appendix describes the data in greater detail and provides a comparison of the relationship between the business sector and the overall economy as an additional check.

Recently, Nalewaik (2010) raises the question of whether gross domestic product (GDP) or gross domestic income (GDI) provides a more accurate reading on economic activity, especially around turning points. As a robustness check, we will use both expenditure-side and income-side measures of output. Specifically, the "standard" measure of GDP and business-sector output from the Bureau of Economic Analysis (BEA) is from the expenditure side. Nalewaik (2010) argues that GDI may better capture the business cycle variations in output growth and that it correlates more strongly with other business cycle variables, before and after data revisions. Nevertheless, Greenaway-McGrevy (2011) and Aruoba, Diebold, Nalewaik, Schofheide and Song (2012) suggest that both GDP and GDI provide independent information, and recommend taking a weighted average of the two.

The standard measure of business output in the NIPAs and in the BLS productivity releases comes from the expenditure side, in particular, GDP less non-business output. Fernald (2014) constructs a corresponding income-side measure as GDI less non-business output. Our benchmark measures of business-sector output, labor productivity, and TFP use an equally weighted average between the GDP and GDI based measures of business output. However, in several places, we discuss where the distinction matters. In the Appendix we also provide further insight on the different implications based on GDP measured from the expediture and the income sides.

An important aspect of the Fernald (2014) dataset is its empirical measure of factor utilization, which is a quarterly version of the BFK measure. BFK wrote down a dynamic cost-minimizing model of the firm where labor and capital are quasi-fixed. If the firm wants more input in the short run, it can adjust an observable intensity margin of hours per worker; or unobserved margins of labor effort and the workweek of capital. The first-order conditions imply that the firm uses all margins simultaneously. Hence, the observable hours per worker can proxy for the unobservable utilization margins and BFK estimate the parameters relating them. BFK and Fernald (2014) implement this measure using detrended industry hours per worker, with different parameters across industries. Hence, variations in utilization are not perfectly correlated with aggregate hours per worker.

3.2 A static decomposition

In this section we project the different margins of adjustment on to the unemployment rate. The sum of the coefficients will be approximately equivalent to the coefficient in the typical Okun's law regression of output growth on the change in the unemployment rate in expression (1). Specifically, for any variable X_{jt} with j = 1, ..., J denoting each of the J margins of adjustment that we consider in expression (6), let $x_{jt} = \log X_{jt}$. Then we take the year on year difference,⁹ denoted Δ_4 , to calculate a smooth yearly rate of change, which in the logs is approximately a percentage change. The regressions then take the form:

$$\Delta_4 x_{jt} = \mu_j + \beta_j \Delta_4 U_t + \varepsilon_{jt}. \tag{7}$$

Table 1 reports the estimate of β in the usual Okun regression, such as that in (1), as well as the estimates of the β_j , for each margin of adjustment considered. We note that we use the averaging of output measures recommended by Aruoba et al. (2012) in the table. For robustness, we also

⁹Quarter-to-quarter changes yield qualitatively similar results.

calculated the results using each measure of output separately. Therefore, row (1) of the table shows the estimate of the Okun coefficient for β , which is -2.25. The estimate based on the real expenditure measure delivers an estimate of β equal to -2.20 whereas the real income measure based estimate is -2.31. All three measures have a standard error of about 0.09 indicating that the estimates do not significantly vary by using one measure or another. Therefore, we proceed by using the average between the two and provide robustness checks in a separate table in the Appendix (Table 4). To anticipate results on the secular behavior of Okun's law reported below, it turns out that these differences become much more accentuated later in the sample, something we will discuss in more detail.

Rows labeled as (2) and (3) decompose the Okun coefficient into the part attributable to total hours and the part attributable to labor productivity, as in equation (2) above. The total hours coefficient in row (2) has a value of -2.09 and shows that most of the Okun coefficient is associated with the decline in hours. Instead, the productivity coefficient in row (3), which has a value of -0.16 and is statistically significant (at the 90% confidence level), shows that the response of labor productivity is roughly an order of magnitude smaller than the response of total hours. Over the full sample, productivity is modestly procyclical since the coefficient is negative.¹⁰ This finding is consistent with the stylized facts from the macro literature (see, e.g., Basu and Fernald 2001 for a discussion and references).

The remaining rows of Table 1 further decompose hours and labor productivity into their constituent elements. Consider total hours first, rows (2a), total employees, and (2b), hours per employee, show that about 80 percent (1.68 percentage points) of that decline in total hours reflects a decline in the number of workers, and about 20 percent reflects a decline in hours per worker. This means that most of the adjustments to total hours take place at the extensive, rather than the intensive, margin. Nevertheless, both margins matter quantitatively.

Rows (3a)-(3c) of Table 1 report the estimates on the growth-accounting based decomposition of labor-productivity growth using equation (5). As expected, capital deepening, $\alpha (\Delta k - l)$ and with a coefficient estimate of 0.62, and labor quality, $(1 - \alpha) \Delta q$ and with a coefficient estimate of 0.06, are countercyclical and contribute positively to labor productivity (both are statistically significant). In contrast, row (3c) shows that measured TFP growth, Δz and with a coefficient estimate of -0.84, is strongly procyclical.

The final two rows of Table 1 further decompose TFP growth. Row (3c.1) shows that utilization, Δv and with a coefficient of -1.03, drives the procyclicality of TFP, whereas technology, Δa , whose coefficient estimate of 0.19 is reported in row (3c.2) is mildly countercyclical. In other words, the

¹⁰Note that we define cyclicality with respect to labor as reflected in the unemployment rate. Hence, procylical productivity growth means that in a boom, labor productivity rises when unemployment falls – so they covary negatively.

procyclicality of TFP mainly reflects the (endogenous) procyclicality of factor utilization, i.e., labor effort and capital's workweek. The utilization margin is crucial for understanding why TFP is strongly procyclical and labor productivity weakly so. Indeed, after controlling for utilization, row (3c.2) shows that utilization-adjusted TFP, Δa , is actually countercyclical with respect to unemployment. These findings are in line with BFK and Galí (1999). They find that technology improvements are contractionary on impact with respect to inputs.

Out of all these findings, perhaps the most important is the large response of total hours worked to a change in unemployment – as row (2) shows, the response is roughly two-to-one. This is much more than in typical DSGE models with unemployment (e.g., Galí, Smets, and Wouters (2011), or Christiano, Eichenbaum, and Trabandt (2013), which generally assume that the number of people employed moves close to one-to-one with unemployment.

As a robustness check, in the Appendix Table 5, we reestimate the total hours response using two broader datasets with wider coverage than just the business sector based data of Table 1. One dataset is based primarily on the establishment survey augmented with data on active military employment as well as agricultural employment, self-employment. The second dataset is based and household employment from the household survey. The estimates based on these two more detailed datasets confirm that the roughly two-to-one response of total hours growth to changes in the unemployment rate hold. More details are provided in the Appendix.

In very broad strokes, the main two takeaways from the analysis are: (1) the total hours response to fluctuations in the unemployment rate is twice as large as is typically assumed in the DSGE literature (for example); (2) much of the procyclicality of TFP comes not from technology (which is countercyclical) but from the utilization margin. However, it is unclear how stable these findings are. The controversy about the stability of Okun's law in the literature (discussed earlier), and the documented changes in the cyclicality of productivity (see, e.g., Basu and Fernald 2001), highlight the need to evaluate the stability of our estimated relationships over the sample. We do this in the next section.

3.3 Secular trends

As a first pass, the simplest diagnostic about the stability of the margins of adjustment can be constructed using 40-quarter-rolling-window regression estimates that mirror those reported in Table 1. We organize these results into Figures 1, 2, and 3. In Figure 1 we report estimations of the regressions intercepts along with the Okun coefficient and the two broad components: total hours and productivity. In Figure 2 we further breakdown the total hours component into hours per worker and total workers. In Figure 3 we breakdown productivity into its components as in expression (6).

Figure 1 panel (b) shows that the Okun relationship has been fairly stable for much of the post-

WW2 era. However, starting around the mid-1990s the relationship became more volatile and the magnitude of the Okun coefficient itself declined somewhat from about an average value of nearly -2.5 to about -2 in more recent times. The mid-1990s coincide with the point in time where the coefficient on total hours markedly declined, and the productivity component switched from being procyclical to being countercyclical. Note also that Figure 1 panel (a) shows that the conditional mean (estimated intercepts) has been declining since the 1980s, while labor productivity has shown a more recent decline.

Figure 2 provides a more detailed look at the total hours component. The conditional averages and the slope coefficients show that the bulk of the decline in total hours is coming from the employment margin. The hours per worker component has been remarkably stable and its contribution is relatively smaller than the employment margin. Possible explanations for this decline in the employment factor include: changes in employer-employee relationships and unionization, globalization of supply chains, and capital substituting technology. Many of these explanations have been explored elsewhere, for example, in Elsby et al. (2013).

Next, we turn to productivity in Figure 3. Earlier we highlighted the important difference between the utilization margin and technology itself. Figure 3 panel (b) shows that the response of capital deepening and labor quality are relatively stable and countercyclical. More importantly, the key driver of the changing cyclicality of labor productivity growth is TFP. Since the mid-1990s the contribution of the utilization margin waned somewhat, and it was the utilization-adjusted TFP (technology) component which became more countercyclical. These changes experienced in the 1990s coincidentally match the decline in the employment component discussed earlier. During the mid-1990s, the correlation between changes in utilization and unemployment was small. The absence of a relationship between utilization and unemployment during this period meant the other variables, which are all countercyclical, pushed labor productivity itself to be countercyclical during this period.

Other components of productivity experienced smaller fluctuations over time. The small increase in the labor-quality component in the mid-1990s likely reflects that the strength of the labor market, when unemployment was falling, pulled lower-skilled workers into the labor force.

In sum, the decompositions of Okun's coefficient over time point to greater responsiveness of total hours – mostly employment – to changes in unemployment, and a largely offsetting change in the cyclicality of labor productivity. The labor productivity changes, in turn, reflect especially the response of utilization and technology itself.

3.4 Business cycle fluctuations

Figures 1-3 suggest that some of the variation in the margins of adjustment could be related to the business cycle. Economically, this also makes sense. In a downturn, firms may curtail hours before

staffing until the contraction in demand becomes more apparent, for example. In this section we explore these issues using two approaches.

First, we track the behavior of output growth and the changes in the unemployment rate (Okun's law) over the recession and recovery phases in reference to their full sample average relationship. We do the same for the relationship between total hours and productivity against changes in the unemployment rate using expression (2). However, we split the sample in 1985 for the productivity panel. Figure 1 showed that productivity goes from procyclical to countercyclical toward the latter part of the sample. The choice of a 1985 break point is loosely justified on this figure and made to roughly coincide with the period often known as the Great Moderation (see, e.g., McConnell and Pérez-Quirós 2000). Figure 4 is organized into four panels: panel (a) displays the Okun relationship; panel (b) displays the total hours versus unemployment rate relationship; and panels (c) and (d) display the productivity versus unemployment rate relationship for the pre- and post-1985 subsamples, respectively.

The simple scatter plot in panel (a) of Figure 4 shows that initially, output drops more quickly than the unemployment rate increases. As the recession wanes and the recovery ensues, the opposite appears to be the case. We observe this pattern in every recession since the beginning of the sample, including the Great Recession. This finding is consistent with Daly, Fernald, Jordà and Nechio (2014).

Total hours does not appear to be the main contributor of this loop-pattern that we just discussed. Panel (b) in Figure 4 indicates that total hours tend to move along the estimated long-run relationship rather than around it. The cyclicality depicted in panel (a) is most closely associated with what happens with productivity.

Panels (c) and (d) in Figure 4 highlight two interesting features. First, the change in the slope, from negative pre-1985 to positive post-1985 corroborates the changes in the cyclicality of productivity displayed earlier in Figure 1. Second, despite the change in the slope across samples, the loop-pattern seen in panel (a) of this figure holds relatively constant. Productivity appears to be the main source of fluctuations around the business cycle.

Next, we do a more detailed and formal analysis. We examine how each of the margins of adjustment varies depending on the stage of the business cycle using National Bureau of Economic Research (NBER) recession dates. Moreover, in order to isolate possible distortions coming from the Great Recession, we allow for the coefficients in that period to vary from those in other recessions. We allow the data to choose whether it wants to assign the same values or not. Therefore, the analysis is based on the margins projections discussed in Section 2 augmented with a set of dummies, that is, we use the following reduced-form regressions:

$$\Delta_4 x_{jt} = \sum_{k=1}^{K} \mu_{kj} I_{kt} + \sum_{k=1}^{K} \beta_{kj} I_{kt} \Delta_4 U_t + \varepsilon_{jt},$$

where $I_{kt} \in \{0, 1\}$ for $k \in \{expansion, recession, Great Recession recovery from Great Recession\}$. That is, we allow the intercept and slope coefficients to vary as a function of whether the economy is in expansion, recession, Great Recession or recovering from the Great Recession. The index j refers to each margin considered, just as in expression (7).

The results of estimating these regressions are reported in Table 2, which are organized as follows. In addition to reporting coefficient estimates and the standard error in parentheses, we also report a p-value in squared brackets. This p-value corresponds to a test of the null hypothesis that the corresponding coefficients differs significantly from the coefficient in expansions, which plays the role of a natural benchmark of comparison.

Several results deserve comment. Consider the Okun coefficient in a regression of output growth on changes in the unemployment rate. It is striking how stable this coefficient is: it fluctuates from a value of -1.98 in expansions, -2.08 in recessions, and -1.83 in the Great Recession. The null that the coefficients in the recessions are equivalent to the coefficient in expansions cannot be rejected at conventional confidence levels. There is a bit more separation in the recovery of the Great Recession with a coefficient of -1.20. This finding is consistent with the lackluster performance of output growth in the last few years, relative to the steady decline in the unemployment rate since the Great Recession ended. That is not to say that the intercept remains constant as well. As one would expect, the average output growth conditional on no changes in the unemployment rate is lower in recessions than in expansions. Perhaps more interesting is the fact that the intercepts for the recession and the recovery from the Great Recession.

We investigate the sources of the stability in the Okun relation by examining the behavior of total hours worked and labor productivity. Within the total hours component, we find considerable stability in its two components, number of workers and hours per employee. The coefficient on hours per employee hovers between -1.65 to -1.90 across all categories, whereas variation in the coefficient on hours per employee fluctuates between -0.19 and -0.40. In neither of these categories are we able to find any significant differences with respect the expansion benchmark. When we turn to the intercept values, there is a noticeable shift between the pre- and post-Great Recession eras with respect to the number of workers component. The average growth rate has declined from approximately 1.5 percent per year to nearly zero since the recovery from the Great Recession began.

Turning now to the component associated with labor productivity, there are some interesting differences between the variation in the value of the intercepts and the slopes. On average, productivity grows more slowly in recessions that in expansions, as would be expected. However, when it comes the covariation between productivity and the unemployment rate, the most significant change is located in the recovery from the Great Recession, with productivity becoming particularly countercyclical.

The sources of this variation are not to be found in the capital deepening component. Whether in terms of average behavior or in terms of its covariance with changes in the unemployment rate, there is striking stability in the coefficients. In contrast, there are some interesting differences associated with the behavior of TFP, which we break down into its components, utilization and utilization-adjusted TFP. Utilization rates exhibit considerable stability with respect to fluctuations in the unemployment rate. In terms of average behavior when the unemployment rate is stable, utilization rates decline in recessions at a rate of about 1.24 percent annually (in comparison to a 0.26 percent growth in expansions).

The more interesting changes come in the utilization-adjusted TFP component. The utilizationadjusted TFP covaries countercyclically with respect to the unemployment rate, dramatically so during the Great Recession (where the coefficient is statistically different than in expansions and about three times as large). Interestingly, when considering average growth rates of utilizationadjusted TFP when the unemployment rate remains constant, expansions and recession experience a similar rate of utilization-adjusted TFP growth at about between 1.2 and 1.5 percent per year. During the Great Recession that number crashed to -1.76 percent per year and has only recovered to about half its normal value in expansions in the recovery from the Great Recession.

The next step in the analysis recognizes that recessions can be generated by different shocks, and economies may adjust differently depending on the shock. Below we offer a more granular depending on the type of shock to explore this issue further. We do this in two ways, by analyzing short-term dynamic adjustment multipliers and then by considering the secular long-term averages.

4 Dynamic adjustment multipliers by shock

This section investigates which margins of adjustment are more important as a function of the shock firms and households are reacting to. In particular, we consider the following four shocks: (1) a monetary policy shock, (2) an oil price shock, (3) a shock to TFP.

Rather than specifying macroeconomic models from which these shocks can be drawn, we borrow these shocks directly from the literature. Specifically, the monetary shock comes from Christiano, Eichenbaum and Evans (1999).¹¹ The oil price shock follows Hamilton (1996). The TFP shock

¹¹To generate an update series of monetary shocks, we use the sum for the preceding year of quarterly VAR monetary innovations, following Christiano, Eichenbaum, and Evans (1999), Burnside (1996), and others. Following Burnside (1996), we measure monetary policy as innovations to the 3-month Treasury bill rate from a VAR with GDP, the GDP deflator, an index of commodity prices, the 3-month T-bill rate, and M1.

come from Fernald (2014). A detailed explanation of how these shocks are obtained and modified is provided in the Appendix. In addition and as a robustness check, the Appendix also reports results using alternative measures of shocks: (1) the monetary policy shock based on Romer and Romer (2004); (2) the fiscal shock from Ramey (2011) based on government defense expenditures; (3) as well as a shock based on an alternative measure of government defense expenditures; (4) a shock to TFP for consumer goods producers; (5) and a shock to TFP for investment goods producers.

To start, we use the whole sample and calculate how each shock affects each margin of adjustment over time using a simple regression. Specifically, we regress each margin of adjustment on the shock and up to its twelve lags (for example, Ramey and Shapiro 1998 use a similar approach to compute impulse response coefficients). We estimate the dynamic multipliers associated with the coefficients of the shock terms in this regression. Figures 5-7 report the accumulated value of the dynamic multipliers so as to smooth the trajectories displayed and to be able to read off the overall effect at each period. The Figures also report 95% confidence bands.

Figure 5 reports multipliers in response to the monetary shock. In response to contractionary monetary policy, output declines over the first year and then begins to stabilize. Both total hours and productivity margins decline. Most of the decline in total hours is due to adjustments on the number of employees. Hours per worker barely moves. The decline in productivity is mainly driven by a reduction in utilization rates displayed in Figure 5.

Oil price increases, shown in Figure 6, have a similar effect on output, that is, they contract economic activity. Like the monetary shock, the adjustment comes from both total hours and productivity. The reduction in total hours, however, is driven by a decline in hours-per-worker, instead of number of employees. Again, the productivity decline is not driven by total factor productivity itself but rather a decline in capital utilization rates.

Figure 7 shows how the behavior of the margins of adjustment to a productivity shock. These multipliers offer a different perspective than Figures 5 and 6, with some interesting differences and similarities between them. A shock to productivity has no impact immediately, but over time becomes expansionary. The increase in output is mainly driven by the gains in labor productivity. In the short run, total hours worked declines due to both a reduction in the number of employees and hours-per-worker.

4.1 Impulse response functions over time

Not only these margins can adjust differently depending on which shock hits the economy but it is also possible that the reaction of each of those margins changed over time. To explore this possibility, we estimate impulse response functions while constraining the data to rolling windows of 40 quarters at a time. More specifically, for each 40-quarter rolling window, we regress each margin of adjustment on the shock and up to its eight lags.¹²

We report the results in Figures 8-10. Figure 8 highlights the sizable negative medium-term effects of monetary shocks during the 1970s and 1980s. More recently, those effects have become more modest and are close to zero. Oil and TFP shocks effects on all variables, reported in Figures 9 and 10, show much larger variation in both size and signs over time.

A comparison between the set of Figures 5-7 and the set of Figures 8-10 show substantial variation in the response of margins to shocks over time. Interestingly, some of the commonly-known patterns explored in the literature seem to be mostly driven by responses during certain periods of time. For example, a comparison between Figures 5 and 8 suggest that the shape of the response of output to monetary shocks in Figure 5 seem to driven by the earlier sample.

More importantly in trying to explain the somewhat puzzling stability of the Okun coefficient, Figures 9 and 10 seem to provide some clue. Recall that our OLS time series analysis of Figures 1-3 showed that while the Okun coefficient remained stable, its labor productivity and total hours components changed substantially, particularly during the 1990s and early 2000s, with the former slope coefficient increasing and the latter slope coefficient decreasing (see Figure 1 panel (b)). As the panel (b) of Figures 2 and 3 showed, movements in total hours came from a decrease in the slope coefficient of number of employees, while changed in labor productivity were linked to variations in utilization and technology shocks themselves.

Figures 9 and 10 show that during this same period (between 1990s and early 2000s) the effects of oil and TFP shocks on the number of employees (and hence in total hours) become more negative and more long lasting. In contrast, the effects of these two shocks on labor productivity became more positive and more long lasting.

4.2 When the dust settles

The nature of the shock clearly has different implications for the margins by which the economy adjusts in the short-run. However, in previous sections we also have highlighted the relative stability of the static analysis. It seems natural to combine elements from both and examine whether the lessons from the static analysis would vary depending on the origin of the shocks considered. We investigate this feature in the next section.

The starting point of the analysis in this section is expression (6) and the sequence of regressions based on this expression and summarized in equation (7). One way to isolate fluctuations in the change of the unemployment rate that are due to the each shock considered above is to follow an instrumental variable (IV) approach in expression (7), where the instrument is the shock and the instrumented variable is the change in the unemployment rate. The coefficient estimates that result

 $^{^{12}}$ We use eight lags instead of twelve due to the smaller sample size.

from this procedure are as an average of the effects delineated in the previous section. This is the manner in which we can connect the static results reported earlier and the dynamic multipliers estimated in the previous section.

Specifically, we reestimate the OLS coefficients reported in Table 1 using each shock at a time (and twelve of its lags) as an instrumental variable. Coefficients are estimated using an two-step instrumental variable approach in expression, where the instrument is the shock and the instrumented variable is the change in the unemployment rate. For each variable x, we estimate $\Delta_4 x_{jt} = \mu_j + \beta_j \Delta_4 \hat{U}_t + \varepsilon_{jt}$, where $\Delta_4 x_{jt}$ is the four-quarter growth rate of x and $\Delta_4 \hat{U}_t$ is the fitted value obtained from the regression of four-quarter percentage-point change in unemployment rate on twelve lags of each shock (instrument).

The results are reported in Table 3. The broad message that all these estimates convey is simple: The static margins of adjustment do not vary substantially across different types of shock over the long-term.

The estimated values reported in Table 3 are close to the estimates reported in Table 1 based on OLS. With few exceptions, variation in the loadings on the different margins do not vary substantially with the type of shock. Consider first the total hours versus productivity split as in expression (2). For most shocks, total hours respond more than productivity except for the TFP shock.

The IV results help us think about the overall implications of the variability of the margins over time reported in Figure 1-3; the variability over the business cycle reported in Table 2 and Figure 4; and the dynamic adjustment multipliers in Figures 5-7. The average estimates of the margins of adjustment in Table 3 are very similar to those reported in Table 1 suggesting that despite some short-term variability, the economy tends to settle into specific patterns as it adjusts to different perturbations. One way to visualize this result is with panel (a) in Figure 4. Even though the Okun relationship displayed is tightly estimated, the recurrent cyclical loops characterizing each recession illustrate how this periodic source of variability dissipates into long-standing stable relations.

5 Conclusion

Macroeconomists and policymakers struggle to understand the channels by which modern economies adjust to shocks of different nature. Whether adjustments are made through staffing levels or hours, capital deepening or utilization rates, or any other available margin has important implications for identifying suitable policy responses and characterizing the macroeconomic environment. This paper uses the novel growth-accounting framework of Fernald (2014) to provide a more detailed look into the margins of adjustment than has been hitherto the case.

Our contributions can be grouped into two broad categories. On one side we establish a number

of stylized facts that enrich some of the assumptions typically made in the DSGE literature. On the other side, our findings enhance the rules-of-thumb upon which policymakers have come to rely on. These rules-of-thumb work well "on average," but our research shows that in the short-run the responses of each margin of adjustment can vary considerably from norm, albeit in predictable patterns.

We highlight several lessons for macroeconomists. First, the responsiveness of hours to changes in unemployment is much larger than typically allowed in DSGE models. We find that hours worked falls about two percent when unemployment rises by one percentage point. The structure of typical DSGE models with unemployment imposes a relationship that is rarely much larger than one-to-one.

Second, output declines faster than the unemployment rate increases at the start of the recession. This decline reverses course as the recession progresses and the recovery begins. The brunt of this cyclical pattern is explained by a similar pattern in productivity rather than in total hours. In recessions, productivity goes from being countercyclical to procyclical. The contribution of each productivity margin depends heavily on the type of shock hitting the economy.

Third, in response to all the shocks we analyze, the adjustment rests primarily with the utilization rate. These results tie into a large literature that emphasizes the importance of unobserved variations in factor intensity as an explanation of movements in productivity (see Basu and Fernald 2001 and references therein). In addition, this result ties into many DSGE models that find that a utilization margin helps to propagate shocks.

The main lesson for policymakers is that the standard Okun law result, while a reliable guidepost in general, conceals a much more nuanced reality. Even this rule-of-thumb is subject to sizeable fluctuations depending on the stage of the business cycle. Within the broad relationship between output and unemployment, there is considerable variation in the margins by which the economy adjusts to different shocks. And all these lessons do not even deal with the difficulties introduced by data available in real-time relative to later revisions (see, e.g. Daly et al. 2014).

References

- [1] Aaronson, Daniel and Daniel Sullivan (2001). "Growth in Worker Quality." Economic Perspectives, Federal Reserve Bank of Chicago.
- [2] Aruoba, Boragan, Francis X. Diebold, Jeremy Nalewaik, Frank Schorfheide and Dongho Song (2012). "Improving GDP Measurement: A Forecast Combination Perspective." In Chen, X., N. Swanson eds., Recent Advances and Future Directions in Causality, Prediction, and Specification Analysis: Essays in Honor of Halbert L. White Jr. Springer, pp. 1-25.

- [3] Barnichon, Regis (2010). "Productivity and Unemployment over the Business Cycle." Journal of Monetary Economics, Nov 2010.
- [4] _____ (2012). "The Shimer Puzzle and the Endogeneity of Productivity." Manuscript, June.
- [5] Ball, Laurence M, Daniel Leigh, and Prakash Loungani (2013). "Okun's Law: Fit at Fifty?" NBER Working Paper No. 18668, January.
- [6] Basu, Susanto, and J. Fernald (2001). "Why Is Productivity Procyclical? Why Do We Care?," New Directions in Productivity Analysis, edited by Edwin Dean, Michael Harper and Charles Hulten, Studies in Income and Wealth (63), University of Chicago Press.
- [7] Basu, Susanto, J. Fernald, and M. Kimball (2006). "Are Technology Improvements Contractionary?," American Economic Review, 96(5): 1418-1448.
- [8] Burnside, Craig (1996). "Production Function Regressions, Returns to Scale, and Externalities," Journal of Monetary Economics, 37(2): 177–201.
- [9] Chang, Pao-Li and Shinichi Sakata (2007). Estimation of Impulse Response Functions Using Long-Run Autoregression. The Econometrics Journal, 10(2): 453-469.
- [10] Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (1999). "Monetary policy shocks: What have we learned and to what end?," Handbook of Macroeconomics, in: J. B. Taylor & M. Woodford (ed.), First Edition, 1(2): 65-148.
- [11] Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt (2013). "Unemployment and Business Cycles." Manuscript, Northwestern University.
- [12] Daly, Mary C., John Fernald, Òscar Jordà, and Fernanda Nechio (2014). "Interpreting Deviations from Okun's Law," FRBSF Economic Letter 2014-12.
- [13] Dean, Edwin R., and Michael J. Harper (2001). "The BLS Productivity Measurement Program," New Developments in Productivity Analysis, 2001, pp 55-84, National Bureau of Economic Research, Inc.
- [14] Fernald, John (2012). "Productivity and Potential Output before, during, and after the Great Recession." Federal Reserve Bank of San Francisco Working Paper 2012-18, September.
- [15] Fernald, John (2014). "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." Federal Reserve Bank of San Francisco Working Paper 2012-19, April.

- [16] Francis, Neville and Valerie A. Ramey (2005). "Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited," Journal of Monetary Economics, Elsevier, vol. 52(8), pages 1379-1399, November.
- [17] Gali, Jordi (1999). "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?." American Economic Review, 89 (March): 249-271
- [18] Gali, Jordi and Luca Gambetti (2009). "On the sources of the great moderation," American Economic Journal: Macroeconomics 1 (1), 26–57
- [19] Gali, Jordi, Frank Smets, and Rafael Wouters (2011). "Unemployment in an Estimated New Keynesian Model," NBER Macroeconomics Annual.
- [20] Gali, Jordi and Thies van Rens (2010). "The Vanishing Procyclicality of Labor Productivity," IZA Discussion Paper No. 5099.
- [21] Gordon, Robert J. (1998). "Foundations of the goldilocks economy: Supply shocks and the time-varying NAIRU," Brookings Papers on Economic Activity (2): 297-346.
- [22] ______ (2011). "The Evolution of Okun's Law and of Cyclical Productivity Fluctuations in the United States and in the EU-15," for presentation at EES/IAB workshop, Labor Market Institutions and the Macroeconomic, Nuremberg, June 17-18.
- [23] Greenaway-McGrevy, Ryan (2011). "Is GDP or GDI \mathbf{a} better measure of output? А statistical approach." Manuscript, Bureau of Economic Analysis, http://www.bea.gov/papers/pdf/Is_GDP_or_GDI_a_better_measure_of_output.pdf
- [24] Hagendorn, Marcus and Iourii Manovskii (2013). "Productivity and the Labor Market: Comovement over the Business Cycle." International Economic Review, Forthcoming.
- [25] Hall, Robert and J.B. Taylor (1993). Macroeconomics, 4th Edition, New York: W.W. Norton.
- [26] Hamilton, James D. 2003. What is an Oil Shock? Journal of Econometrics 113: 363-398.
- [27] Hamilton, James D. 2011. Nonlinearities and the Macroeconomic Effects of Oil Prices. Macroeconomic Dynamics vol. 15, Supplement 3, pp. 364-378.
- [28] Hamilton, James D. and Ana María Herrera. 2004. Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy. Journal of Money, Credit, and Banking 36: 265-286.
- [29] Jordà, Oscar (2005). Estimation and Inference of Impulse Responses by Local Projections. American Economic Review, 95(1): 161-182.

- [30] Lewis, R. A. and Gregory C. Reinsel (1985). Prediction of Multivariate Time Series by Autoregressive Model Fitting. of Multivariate Analysis, 16(33): 393-411.
- [31] Nalewaik, Jeremy J. (2010). "The Income- and Expenditure-Side Measures of Output Growth," Brookings Papers on Economic Activity, vol. 1, pp. 71-106.
- [32] Okun, Arthur M. (1962). "Potential GNP: Its Measurement and Significance," Proceedings of the Business and Economics Statistics Section of the American Statistical Association, 98–104.
- [33] ______ (1965). "The Gap between Actual and Potential Output," in Arthur M. Okun, ed., The Battle against Unemployment. New York: W.W. Norton & Co.
- [34] Plosser, Charles I. and William G. Schwert (1979). "Potential GNP: Its measurement and significance : A dissenting opinion," Carnegie-Rochester Conference Series on Public Policy, Elsevier, vol. 10(1), pages 179-186, January.
- [35] Prachowny, Martin F. J. (1993). "Okun's Law: Theoretical Foundations and Revised Estimates." The Review of Economics and Statistics, Vol. 75, No. 2 (May, 1993), pp. 331-336.
- [36] Ramey, Valerie R. (2011). "Identifying Government Spending Shocks: It's All in the Timing." Quarterly Journal of Economics, February.
- [37] Ramey, Valerie R. (2012). "The Impact of Hours Measures on the Trend and Cycle Behavior of Labor Productivity." Manuscript, downloaded August 22, 2013 from http://www.econ.ucsd.edu/~vramey/research/Ramey-Hours-Productivity.pdf.
- [38] Ramey, Valerie R. and Matthew D. Shapiro (1998). Costly capital reallocation and the effects of government spending. Carnegie–Rochester Conference Series on Public Policy, 48: 145–194.
- [39] Romer, Christina D., and David H. Romer (2004). "A New Measure of Monetary Shocks: Derivation and Implications." American Economic Review, 94(4): 1055-1084.
- [40] Stiroh, Kevin (2009). "Volatility Accounting: A Production Perspective on Increased Economic Stability." Journal of the European Economic Association 7 (4), 671-696.

Dependent variables:	Slope estimates:
(1) Output (Δy)	-2.25^{***} (0.09)
(2) Hours (Δ l)	-2.09^{***} (0.06)
(2a) Employees (Δn)	-1.68^{***} (0.05)
(2b) Hours per employee (Δh)	-0.41^{***} (0.03)
(3) Labor productivity (Δy - Δl)	-0.16^{*} (0.09)
(3a) Capital deepening $(\alpha(\Delta k-\Delta l))$	0.62^{***} (0.02)
(3b) Labor quality ((1- α) Δ q)	0.06^{***} (0.01)
(3c) TFP (Δz)	-0.84^{***} (0.09)
(3c.1) Utilization (Δv)	-1.09^{***} (0.08)
(3c.2) Utilization-adjusted TFP (Δa)	0.25^{***} (0.08)

Table 1: Margins of adjustment

For each variable x, the entries shown are the slope coefficients from estimating $\Delta_4 x_{jt} = \mu_j + \beta_j \Delta_4 U_t + \varepsilon_{jt}$, where $\Delta_4 x_{jt}$ is the four-quarter growth rate of x and $\Delta_4 U_t$ is the four-quarter percentage-point change in unemployment. The column of entries measures output as the average of real business expenditure and income. The sample runs from 1949Q1 to 2014Q1.

		In	tercept				Slope	
Dependent variables:	Expansion	Recession	Great Recession	Recovery	Expansion	Recession	Great Recession	Recovery
(1) Output (Δy)	$\begin{array}{c} 3.94^{***} \\ (0.11) \\ [0.00] \end{array}$	$\begin{array}{c} 2.24^{***} \\ (0.34) \\ [0.00] \end{array}$	$\begin{array}{c} 0.74 \\ (0.93) \\ [0.13] \end{array}$	2.28^{***} (0.34) [0.00]	-1.98*** (0.11) [0.00]	-2.08^{***} (0.21) [0.66]	-1.83^{***} (0.42) [0.59]	-1.20^{***} (0.26) [0.01]
(2) Hours (Δl)	$\begin{array}{c} 1.16^{***} \\ (0.08) \\ [0.00] \end{array}$	$\begin{array}{c} 0.75^{***} \\ (0.24) \\ [0.11] \end{array}$	$\begin{array}{c} 0.63\\ (0.67)\\ [0.86] \end{array}$	$\begin{array}{c} 0.26 \\ (0.24) \\ [0.00] \end{array}$	-2.06*** (0.08) [0.00]	-1.85^{***} (0.15) [0.21]	-2.23^{***} (0.30) [0.25]	-2.23^{***} (0.18) [0.38]
(2a) Employees (Δn)	$\begin{array}{c} 1.30^{***} \\ (0.07) \\ [0.00] \end{array}$	$\begin{array}{c} 1.58^{***} \\ (0.23) \\ [0.24] \end{array}$	$\begin{array}{c} 0.81 \\ (0.63) \\ [0.25] \end{array}$	$\begin{array}{c} 0.10 \\ (0.23) \\ [0.00] \end{array}$	-1.69^{***} (0.07) [0.00]	-1.65^{***} (0.14) [0.80]	-1.84^{***} (0.29) [0.56]	-1.90^{***} (0.17) [0.27]
(2b) Hours per employee (Δh)	-0.15^{***} (0.04) [0.00]	-0.83^{***} (0.14) [0.00]	-0.18 (0.38) [0.11]	$\begin{array}{c} 0.16 \\ (0.14) \\ [0.04] \end{array}$	-0.36^{***} (0.04) [0.00]	-0.19^{**} (0.08) [0.08]	-0.40^{**} (0.17) [0.29]	-0.33^{***} (0.10) [0.77]
 Labor productivity (Δy-Δl) 	$\begin{array}{c} 2.78^{***} \\ (0.12) \\ [0.00] \end{array}$	$\begin{array}{c} 1.48^{***} \\ (0.37) \\ [0.00] \end{array}$	$\begin{array}{c} 0.11 \\ (1.01) \\ [0.20] \end{array}$	2.02^{***} (0.37) [0.05]	$\begin{array}{c} 0.08 \\ (0.12) \\ [0.52] \end{array}$	-0.24 (0.22) $[0.22]$	0.41 (0.46) [0.21]	$\begin{array}{c} 1.04^{***} \\ (0.28) \\ [0.00] \end{array}$
(3a) Capital deepening $(\alpha(\Delta \mathbf{k} \cdot \Delta \mathbf{l}))$	$\begin{array}{c} 0.84^{***} \\ (0.03) \\ [0.00] \end{array}$	$\begin{array}{c} 1.13^{***} \\ (0.10) \\ [0.00] \end{array}$	$\begin{array}{c} 0.91^{***} \\ (0.27) \\ [0.43] \end{array}$	$\begin{array}{c} 0.38^{***} \\ (0.10) \\ [0.00] \end{array}$	0.59^{***} (0.03) [0.00]	0.50^{***} (0.06) [0.18]	0.66^{***} (0.12) [0.25]	0.69^{***} (0.07) [0.21]
(3b) Labor quality ((1- α) Δ q)	$\begin{array}{c} 0.21^{***} \\ (0.02) \\ [0.00] \end{array}$	$\begin{array}{c} 0.37^{***} \\ (0.05) \\ [0.00] \end{array}$	$\begin{array}{c} 0.46^{***} \\ (0.14) \\ [0.53] \end{array}$	0.32^{***} (0.05) [0.04]	$\begin{array}{c} 0.02 \\ (0.02) \\ [0.24] \end{array}$	$\begin{array}{c} 0.01 \\ (0.03) \\ [0.84] \end{array}$	$\begin{array}{c} 0.06\\ (0.06)\end{array}$	$\begin{array}{c} 0.14^{***} \\ (0.04) \\ [0.00] \end{array}$
(3c) TFP (Δz)	$\begin{array}{c} 1.74^{***} \\ (0.12) \\ [0.00] \end{array}$	-0.02 (0.37) [0.00]	-1.29 (1.01) [0.24]	$\begin{array}{c} 1.34^{***} \\ (0.37) \\ [0.30] \end{array}$	-0.52^{***} (0.12) [0.00]	-0.74^{***} (0.22) [0.39]	-0.29 (0.46) [0.37]	$\begin{array}{c} 0.23 \\ (0.28) \\ [0.01] \end{array}$
(3c.1) Utilization (Δv)	0.26^{**} (0.11) [0.02]	-1.24^{***} (0.34) [0.00]	$\begin{array}{c} 0.47 \\ (0.94) \\ [0.09] \end{array}$	0.58^{*} (0.34) [0.38]	-0.90^{***} (0.11) [0.00]	-0.82^{***} (0.21) [0.76]	-1.43*** (0.43) [0.21]	-0.50^{*} (0.26) $[0.16]$
(3c.2) Utilization-adjusted TFP (Δa)	$\begin{array}{c} 1.47^{***} \\ (0.11) \\ [0.00] \end{array}$	$\begin{array}{c} 1.22^{***} \\ (0.33) \\ [0.46] \end{array}$	-1.76^{*} (0.91) [0.00]	0.76^{**} (0.33) [0.04]	0.38^{***} (0.11) [0.00]	$\begin{array}{c} 0.08 \\ (0.20) \\ [0.20] \end{array}$	$\begin{array}{c} 1.14^{***} \\ (0.41) \\ [0.02] \end{array}$	$\begin{array}{c} 0.72^{***} \\ (0.25) \\ [0.20] \end{array}$
For each business-sector variable x , the entries $\{Expansion, Recession, GreatRecession, Recovery\}$. report a p-value in squared brackets. This p-value corrismeasured as the average of real expenditure and rea	shown are the The index j recesponds to a teces al income. Reces	coefficients fifers to each m st of the null h sion dates are	com estimating $\Delta_4 x$, argin considered. Th ypothesis that the cor obtained from NBER.	$it = \sum_{k=1}^{K} \mu_{l}$ e table reports responding coeff The sample run	$y_{ij}I_{kt} + \sum_{k=1}^{K} \beta_{ij}$ coefficient estimation icients differs signation from 19490.1 t	$k_j I_{kt} \Delta_4 U_t +$ ates and the s- nificantly from to 2014Q1.	ε_{jt} , where $I_{kt} \in \{$ candard error in parent, the coefficient in expansion	$\begin{array}{c c} 0,1 \end{array} \begin{array}{c} \text{for } k \\ \hline \\ \text{theses, we also} \\ \text{nsions. Output} \end{array}$

Table 2: Margins of adjustment – Normal times versus recessions

	Slope estim	nates by type o	of shock (instrument):
Dependent variables:	Monetary	Oil	TFP
(1) Output (Δy)	-2.50^{***} (0.22)	-1.84^{***} (0.23)	(0.32)
(2) Hours (Δl)	-2.27^{***}	-1.92^{***}	-2.64^{***}
	(0.15)	(0.15)	(0.22)
(2a) Employees (Δn)	-1.99^{***}	-1.31^{***}	-1.94^{***}
	(0.15)	(0.15)	(0.19)
(2b) Hours per employee (Δh)	-0.29^{***}	-0.62^{***}	-0.70^{***}
	(0.09)	(0.09)	(0.12)
(3) Labor productivity $(\Delta y - \Delta l)$	-0.22 (0.22)	0.08 (0.23)	1.02^{***} (0.38)
(3a) Capital deepening $(\alpha(\Delta k-\Delta l))$	$\begin{array}{c} 0.78^{***} \\ (0.07) \end{array}$	0.71^{***} (0.06)	$\begin{array}{c} 0.83^{***} \\ (0.09) \end{array}$
(3b) Labor quality ((1- α) Δ q)	$\begin{array}{c} 0.11^{***} \\ (0.03) \end{array}$	$0.05 \\ (0.03)$	0.03 (0.04)
(3c) TFP (Δz)	-1.11^{***}	-0.65^{***}	0.17
	(0.23)	(0.24)	(0.37)
(3c.1) Utilization (Δv)	-0.94^{***}	-1.00^{***}	-2.38^{***}
	(0.21)	(0.21)	(0.37)
(3c.2) Utilization-adjusted TFP (Δa)	-0.17	0.35^{*}	2.55^{***}
	(0.21)	(0.2)	(0.53)

\mathbf{T}	Table 3: Margir	s of adju	stment - instrume	ental variables	bv	type of	f shock
--------------	-----------------	-----------	-------------------	-----------------	----	---------	---------

Coefficients are estimated using an two-step instrumental variable approach in expression, where the instrument is the shock and the instrumented variable is the change in the unemployment rate. For each variable x, the entries shown are the slope coefficients from estimating $\Delta_4 x_{jt} = \mu_j + \beta_j \Delta_4 \hat{U}_t + \varepsilon_{jt}$, where $\Delta_4 x_{jt}$ is the four-quarter growth rate of x and $\Delta_4 \hat{U}_t$ is the fitted value obtained from the regression of four-quarter percentage-point change in unemployment rate on twelve lags of each shock (instrument). Output is measured as the average of real expenditure and real income. The monetary shock is obtained from an update of Christiano et al. (1999). The oil shock is obtained from Hamilton (1996) and the sample runs from 1949Q1 to 2014Q1. The TFP shock is obtained from Fernald (2014) and the sample runs from 1949Q1 to 2014Q1.



Figure 1: Margins of adjustment – output, total hours and labor productivity

The figure reports the relationships between unemployment and output, total hours, and productivity over time. Each series corresponds to the 40-quarter-rolling-window regression estimates of the intercept, in panel (a), and the slope coefficient, in panel (b), in $\Delta_4 x_{jt} = \mu_j + \beta_j \Delta_4 U_t + \varepsilon_{jt}$, where $\Delta_4 x_{jt}$ is the four-quarter growth rate of x (output, total hours, labor productivity) and $\Delta_4 U_t$ is the four-quarter percentage-point change in unemployment. The sample runs from 1949Q1 to 2014Q1.



Figure 2: Margins of adjustment – total hours components

The figure reports the relationships between unemployment and number of employees, and hours-per-worker over time. Each series corresponds to the 40-quarter-rolling-window regression estimates of the intercept, in panel (a), and the slope coefficient, in panel (b), in $\Delta_4 x_{jt} = \mu_j + \beta_j \Delta_4 U_t + \varepsilon_{jt}$, where $\Delta_4 x_{jt}$ is the four-quarter growth rate of x (number of employees, hours-per-worker) and $\Delta_4 U_t$ is the four-quarter percentage-point change in unemployment. The sample runs from 1949Q1 to 2014Q1.



Figure 3: Margins of adjustment – labor productivity components

The figure reports the relationships between unemployment and capital deepening, labor quality, capital utilization and utilizationadjusted TFP over time. Each series corresponds to the 40-quarter-rolling-window regression estimates of the intercept, in panel (a), and the slope coefficient, in panel (b), in $\Delta_4 x_{jt} = \mu_j + \beta_j \Delta_4 U_t + \varepsilon_{jt}$, where $\Delta_4 x_{jt}$ is the four-quarter growth rate of x (capital deepening, labor quality, capital utilization, utilization-adjusted TFP) and $\Delta_4 U_t$ is the four-quarter percentage-point change in unemployment. The sample runs from 1949Q1 to 2014Q1.





on three NBER recession starting dates: 1973Q4, 2001Q1 and 2007Q4. Panel (a) displays the Okun relationship; panel (b) displays the total hours versus unemployment rate relationship; and post-1985 subsamples, respectively. The scatter in panels (a) The figure tracks the behavior of output growth, total hours and labor productivity and the changes in the unemployment rate over the recessions and recovery phases focusing

and (b) shows data between 1949Q1 and 2014Q1, in panel (c) correspond to data between 1949Q1 and 1984Q4, and panel (d) correspond to data between 1985Q1 and 2014Q1.



Figure 5: Impulse response functions – monetary shocks

The figure reports the accumulated value of the dynamic multipliers associated with the coefficients of the monetary shock from the regression of each margin of adjustment on the shock, and up to its twelve lags. Monetary shock is obtained from Christiano, Eichenbaum and Evans (1999). Sample runs from 1949Q1 to 2014Q1. Dashed lines correspond to 95% confidence bands.





The figure reports the accumulated value of the dynamic multipliers associated with the coefficients of the oil shock from the regression of each margin of adjustment on the shock, and up to its twelve lags. Oil shock is obtained from Hamilton (1996). Sample runs from 1974Q1 to 2014Q1. Dashed lines correspond to 95% confidence bands.





The figure reports the accumulated value of the dynamic multipliers associated with the coefficients of the TFP shock from the regression of each margin of adjustment on the shock, and up to its twelve lags. TFP are obtained from Fernald (2014). Sample runs from 1949Q1 to 2014Q1. Dashed lines correspond to 95% confidence bands.





The figure reports the accumulated value of the dynamic multipliers in the first, fourth and eitghth quarters following the monetary shock. These effects are obtained from the regression of each margin of adjustment on the shock, and up to its eighth lags using rolling windows of 40 quarters. Monetary shock is obtained from Christiano, Eichenbaum and Evans (1999). Sample runs from 1949Q1 to 2014Q1. Shades correspond to 95% confidence bands.



Figure 9: Impulse response functions – monetary shocks

The figure reports the accumulated value of the dynamic multipliers in the first, fourth and eitghth quarters following the oil shock. These effects are obtained from the regression of each margin of adjustment on the shock, and up to its eighth lags using rolling windows of 40 quarters. Oil shock is obtained from Hamilton (1996). Sample runs from 1974Q1 to 2014Q1. Shades correspond to 95% confidence bands.





The figure reports the accumulated value of the dynamic multipliers in the first, fourth and eitghth quarters following the TFP shock. These effects are obtained from the regression of each margin of adjustment on the shock, and up to its eighth lags using rolling windows of 40 quarters. TFP are obtained from Fernald (2014). Sample runs from 1949Q1 to 2014Q1. Shades correspond to 95% confidence bands.

A Appendix

A.1 Expenditure versus income sides

Recently, Nalewaik (2010) raises the question of whether gross domestic product (GDP) or gross domestic income (GDI) provides a more accurate reading on economic activity, especially around turning points. As a robustness check, we use both expenditure-side and income-side measures of output. Specifically, the "standard" measure of GDP and business-sector output from the Bureau of Economic Analysis (BEA) is from the expenditure side. Nalewaik (2010) argues that GDI may better capture the business cycle variations in output growth and that it correlates more strongly with other business cycle variables, before and after data revisions. Nevertheless, Greenaway-McGrevy (2011) and Aruoba et al (2012) suggest that both GDP and GDI provide independent information, and recommend taking a weighted average of the two.

In particular, the first column in Table 4 shows results using unpublished BLS data on the total economy;¹³ these data are primarily from the establishment survey, but are augmented with data on active military employment as well as agricultural employment, self-employment, and household employment from the household survey. Some recent literature argues for using the household survey as the primary source, instead of the establishment survey; see Ramey (2012) and Hagedorn and Manovski (2013). The second column in Table 4 shows results using these data. The row numbers used in the table correspond to those in Table $1.^{14}$

Table 4 shows that the high response of total hours to an unemployment rate change is not explained by our focus on the more cyclically sensitive business sector data, nor is it limited to the establishment survey. In both columns, hours fall just under two percent when the unemployment rate rises by one percentage point. In the household survey, more of the response comes from hours per worker, and less from a change in the number of people working, but the total response is similar.¹⁵ This suggests that our results accurately reflect the fact that a wide range of margins are important. For example, the final line of Table 4 shows that labor-force participation is quantitatively important in explaining the response of the number of people working. Other margins, such as changes in number of multiple-job holders, also contribute. Together, the results imply that models that ignore these margins potentially miss quantitatively important aspects of the economy's adjustment to shocks.

We now consider how the output-unemployment relationship has changed over time. Figure 11 plots 40-quarter rolling estimates of the Okun coefficient β_j when output is measured using either real income or real expenditures. Figure 11 shows the striking divergence in the estimated Okun coefficient between the income- and expenditure-based measures of output beginning in the early 1990s. Note the two series also diverged briefly in the second half of the 1970s. The crosses on each series represent periods when the differences in the two coefficients are statistically significant. Looking first at the expenditure series, the magnitude of the Okun coefficient has declined over time

¹³We thank John Glaser (from the Bureau of Labor Statistics) for providing us with the data.

¹⁴CPS hours work data start in 1976, restricting the sample to 1976Q3 to 2012Q4.

¹⁵Multiple job-holding may explain the different mix between individuals and hours between the two surveys. Suppose the unemployment rate rises and some individuals lose a second job. In the establishment survey, that shows up as one employee fewer. In the household survey, the person still has a job – but would report fewer hours worked.

and has been smaller in the last two decades than it was in the previous three decades. This pattern is not present in the income data where the coefficient has been more stable.

Nalewaik (2010) argues that over the past few decades, the income-based measures of output are more correlated with other measures of activity (say, ISM surveys); and are more closely related to what forecasters are saying. More pointedly, he argues that the product-based measure of output has become increasingly unreliable. If this is true, the magnitude of the Okun coefficient on the expenditure side should fall, since the covariance with unemployment should fall. Nalewaik's claim is thus consistent with our results.

Despite this divergence, Table 4 shows that our results are qualitatively unchanged when using one measure or the other.

A.2 Total economy hours worked

Table 5 reestimated the total hours response using two broader datasets with wider coverage than just the business sector based data of Table 1. One dataset is based primarily on the establishment survey augmented with data on active military employment as well as agricultural employment, selfemployment. The second dataset is based and household employment from the household survey. The estimates based on these two more detailed datasets confirm that the roughly two-to-one response of total hours growth to changes in the unemployment rate hold.

A.3 Alternative shocks and instruments

Figures 12-15 and Table 6 replicate the results reported in the paper on Figures 5-7 and Table 3 but considers, instead alternative measures shocks: (1) the monetary policy shock based on Romer and Romer (2004),¹⁶ (2) the fiscal shock from Ramey (2011) based on government defense expenditures, (3) as well as a shock based on an alternative measure of government defense expenditures. It also considers (4) a shock to TFP for consumer goods producers, and (5) a shock to TFP for investment goods producers.

While Figures 12 and 13 bring similar lessons than Figures 5 and 6, the productivity shocks on the consumer and investment sectors differ substantially from the aggregate TFP shock reported in Figure 7.

A shock in the consumption sector is associated with a positive response of output whereas the investment productivity shock is contractionary. An explanation for this difference comes from the behavior of total hours. In contrast, productivity improves in both sectors, not surprisingly given the nature of the shock. In the case of the investment TFP shock, both employees and hours per worker decline along with utilization even as adjusted TFP is improving. Meanwhile, the response in employees, hours per worker and utilization rates is more muted when the shock comes from the consumption sector.

¹⁶We use the sum for the preceding year of quarterly VAR monetary innovations, following Christiano, Eichenbaum, and Evans (1999), Burnside (1996), and others. Following Burnside (1996), we measure monetary policy as innovations to the 3-month Treasury bill rate from a VAR with GDP, the GDP deflator, an index of commodity prices, the 3-month T-bill rate, and M1.

One way to reconcile these differences is to consider a sticky-price, two-sector model such as Basu, Fernald and Liu (2014). In that model, although a productivity shock will lower the price of investment goods eventually, in the short-run price-stickiness causes investment to be relatively expensive. As a result, the initial effect of the shock is contractionary. When the productivity shock is to the consumption sector instead, short-run price stickiness makes investment goods relatively cheaper, which in turn expands investment and output. In addition, there is a wealth effect coming from expected future gains in the consumption sector that also pushes economic activity up today.

A.4 Fernald (2014) Quarterly Growth-Accounting Data

These data are available at http://www.frbsf.org/economics/economists/jfernald/quarterly_tfp.xls. They include quarterly growth-accounting measures for the business-sector, including output, hours worked, labor quality (or composition), capital input, and total factor productivity from 1947:Q2 on. In addition, they include a measure of factor utilization that follows BFK. They are typically updated one to two months after the end of the quarter (for example, data through 2011:Q4 were posted on February 2, 2012, following the release of BLS Productivity and Cost data for the fourth quarter). Once aggregated to an annual frequency, they are fairly close to the annual BLS multifactor productivity estimates, although there are some differences in coverage and implementation.¹⁷ The data are described in greater detail in Fernald (2014).

Key data sources for estimating (unadjusted) quarterly TFP for the U.S. business sector are:

- (i) Business output: We use income and expenditure side measures of real output. The expenditure side, which corresponds to GDP is reported in NIPA tables 1.3.5 and 1.3.6 (gross value added by sector). Nominal business income (the business counterpart of GDI) is GDI less nominal non-business output from table 1.3.5. Real GDI and business income uses the expenditure-side deflators.
- (ii) Hours: From the quarterly BLS productivity and cost release.
- (iii) Capital input: Weighted growth in 13 types of disaggregated quarterly capital. Weights are estimated factor payments (which, in turn, use estimated user costs). The quarterly national income and product accounts (produced by the Bureau of Economic Analysis, BEA) provide quarterly investment data for 6 types of non-residential equipment and software; and for 5 types of non-residential structures. I use these data to create perpetual-inventory series on (end of previous quarter, i.e., beginning of current quarter) capital stocks by different type of asset. In addition, I use quarterly NIPA data on inventory stocks and interpolate/extrapolate the annual BLS estimates of land input. Note that the data also allow me to calculate sub-aggregates, such as equipment and software capital, or structures capital.

¹⁷To name six minor differences: (i) BLS covers private business, Fernald covers total business. (ii) BLS uses expenditure-side measures of output, whereas Fernald combines income and expenditure-side measures of output. (iii) BLS assumes hyperbolic (rather than geometric) depreciation for capital. (iv) BLS uses the more disaggregated investment data available at an annual frequency. (v) Fernald does not include rental residential capital. (vi) There are slightly different methodologies for estimating labor quality. Some of these differences reflect what can be done quarterly versus annually. For a review of the methodology and history of the BLS measures, see Dean and Harper (2001).

- (iv) Factor shares: Interpolated and, where necessary, extrapolated from the annual data on factor shares, α and (1α) , from the BLS multifactor productivity database.
- (v) Labor composition: Interpolated and extrapolated from annual measures in the BLS multifactor productivity data.

To estimate a quarterly series on utilization, the key data source is the following:

(vi) Hours-per-worker $\left(\frac{H^i}{N^i}\right)$ by industry *i* from the monthly employment report of the BLS. These are used to estimate a series on industry utilization $\Delta \ln \Upsilon_i = \beta_i \Delta \ln \left(\frac{H^i}{N^i}\right)$, where β^i is a coefficient estimated by BFK. Fernald (2014) then calculates an aggregate utilization adjustment as $\Delta \ln \Upsilon = \Sigma_i w_i \Delta \ln \Upsilon_i$, where is the industry weight from BFK (taken as the average value over the full sample).

The resulting utilization-adjusted series differs conceptually from the BFK purified technology series along several dimensions. BFK use detailed industry data to construct estimates of industry technology change that control for variable factor utilization and deviations from constant returns and perfect competition. They then aggregate these residuals to estimate aggregate technology change. Thus, they do not assume the existence of a constant-returns aggregate production function. The industry data needed to undertake the BFK estimates are available only annually, not quarterly. As a result, the quarterly series estimated here does not control for deviations from constant returns and perfect competition.¹⁸

For this paper, we modify the TFP and utilization-adjusted TFP measures in two ways relative to the figures in the downloadable spreadsheet. First, we create separate income-and output-side labor- and total-factor productivity measures, rather than simply using the geometric average in Fernald (2014). Second, the Fernald dataset uses two measures of labor "quality" to adjust for the composition of the workforce by age, education, and other observable demographics. The first measure is interpolated from the annual estimates available from the BLS and is available for the entire sample. The second is a true quarterly measure from the Current Population Survey, which implements the quarterly composition adjustment from Aaronson and Sullivan (2001). Although theoretically preferable, this second measure is available only since 1979. Especially when we look at time variation in coefficients, it is important to have a consistent measure. Hence, we adjust TFP and utilization adjusted TFP to use the consistent, interpolated BLS measure.

A.5 Relating Business and Non-Business Sectors to the Total Economy

Unemployment is for the total economy, whereas our growth-accounting data are for the business sector. The business sector accounts for about 3/4 of GDP (from the national accounts) and employment.¹⁹ The non-business sector is mainly government services, nonprofits, and household workers.

¹⁸The output data also differ, both in vintage and data source, from the annual data used by BFK.

¹⁹We thank John Glaser (from the Bureau of Labor Statistics) for providing us with the data.

The Tornquist approximation to chained GDP implies the following:

$$\Delta y_{total} = w_{bus} \Delta y_{bus} + (1 - w_{bus}) \Delta y_{nbus}$$

The logic of the growth-accounting decompositions from the text implies:

$$\beta_{total} = w_{Bus}\beta_{Bus} + (1 - w_{bus})\beta_{nbus},$$

where β_j is from $\Delta y_{jt} = \mu_j + \beta_j \Delta U_t + \varepsilon_{jt}$.

Figure 16 shows these estimates. Total economy is the average of real GDP and real GDI (where real GDI is nominal GDI deflated with the GDP deflator). Non-business output is from NIPA Table 1.3.3 (accessed August 13, 2014). Clearly, the cyclicality of output for the overall economy comes almost entirely from the business sector. Indeed, the non-business sector displays little cyclicality with respect to unemployment, apart from a brief period in the early 1970s.

Table 4: Margins of adjustment –	expenditure versus inc	ome sides
	Real Expenditure	Real Income
(1) Output (Δy)	-2.20***	-2.31***
	(0.09)	(0.09)
(2) Hours (Δ l)	-2.09***	-2.09***
	(0.06)	(0.06)
(2a) Employees (Δn)	-1.68***	-1.68***
	(0.05)	(0.05)
(2b) Hours per employee (Δh)	-0.41***	-0.41***
	(0.03)	(0.03)
(3) Labor productivity $(\Delta y - \Delta l)$	-0.11	-0.22**
	(0.09)	(0.09)
(3a) Capital deepening $(\alpha(\Delta k - \Delta l))$	0.62***	0.62***
	(0.02)	(0.02)
(3b) Labor quality $((1-\alpha)\Delta q)$	0.06***	0.06***
	(0.01)	(0.01)
(3c) TFP (Δz)	-0.84***	-0.84***
	(0.09)	(0.09)
$(3c.1)$ Utilization (Δv)	-1.09***	-1.09***
	(0.08)	(0.08)
(3c 2) Utilization-adjusted TFP (Δa) 0.25***	0 25***
	(0.08)	(0.08)

For each variable x, the entries shown are the slope coefficients from estimating $\Delta_4 x_{jt} = \mu_j + \beta_j \Delta_4 U_t + \varepsilon_{jt}$, where $\Delta_4 x_{jt}$ is the four-quarter growth rate of x and $\Delta_4 U_t$ is the four-quarter percentage-point change in unemployment. The column of entries measures real business output from the expenditure and income sides. The sample runs from 1949Q1 to 2014Q1.

	Slope estin	mates:
Dependent variables:	Establishment survey	Household survey
(2) Hours (Δ l)	-2.09***	-1.71***
	(0.06)	(0.08)
(2a) Employees (Δn)	-1.68***	-1.16***
	(0.05)	(0.05)
(2b) Hours per employee (Δh)	-0.41***	-0.56***
	(0.03)	(0.05)
Additional:		
Civilian employment		-1.12***
1 0		(0.05)
Labor force participation rate		-0.07***
1		(0.02)

Table 5: Comparing Establishment and Household Survey Hours

For each variable x, the entries shown are the slope coefficients from estimating $\Delta_4 x_{jt} = \mu_j + \beta_j \Delta_4 U_t + \varepsilon_{jt}$, where $\Delta_4 x_{jt}$ is the four-quarter growth rate of x and $\Delta_4 U_t$ is the four-quarter percentage-point change in unemployment. The table uses alternative data sources on employment and hours, which correspond to the total economy. The "Establishment Survey" column is from unpublished BLS total-economy data, which relies predominately on the establishment survey for employment and (for production workers) hours per worker. The "Household Survey" column shows number of persons at work and average hours of all persons at work from the household survey (data obtained from Haver Analytics). The memo items on civilian employment and labor-force participation are from the summary tables in the monthly employment report. The sample runs from 1976Q3 to 2014Q1.

		Slope e	stimates by ty ₁	pe of shock ((instrument):	
Dependent variables:	Monetary	Oil	Government spending	Defense spending	TFP Consumption	TFP Investment
(1) Output (Δy)	-2.55^{***} (0.26)	-2.18^{***} (0.2)	-1.84^{***} (0.44)	-2.13^{***} (0.52)	-2.13^{***} (0.16)	-2.54^{***} (0.25)
(2) Hours (ΔI)	-2.44^{***} (0.15)	-1.96^{**} (0.13)	-1.80^{***} (0.29)	-3.15^{**} (0.45)	-2.38^{***} (0.11)	-2.44^{***} (0.17)
(2a) Employees (Δn)	-2.05^{***} (0.15)	-1.42^{**} (0.13)	-1.39^{***} (0.28)	-2.01^{***} (0.33)	-1.69^{***} (0.1)	-1.96^{***} (0.16)
(2b) Hours per employee (Δh)	-0.39^{***} (0.1)	-0.54^{***} (0.08)	-0.40^{**} (0.16)	-1.14^{***} (0.32)	(70.0)	-0.48^{***} (0.09)
(3) Labor productivity $(\Delta y - \Delta l)$	-0.11 (0.26)	-0.22 (0.2)	-0.04 (0.44)	$1.02 \\ (0.64)$	0.25 (0.17)	-0.1 (0.25)
(3a) Capital deepening $(\alpha(\Delta k - \Delta l))$	0.74^{***} (0.05)	0.64^{**} (0.05)	0.44^{***} (0.13)	1.37^{***} (0.28)	0.94^{***} (0.06)	0.65^{***} (0.07)
(3b) Labor quality $((1-\alpha)\Delta q)$	0.15^{***} (0.05)	0.02 (0.03)	(0.06)	0.03 (0.08)	0.04^{*} (0.02)	0.14^{***} (0.04)
(3c) TFP (Δz)	-0.99^{***} (0.29)	-0.87^{***} (0.21)	-0.57 (0.45)	-0.37 (0.56)	-0.73^{***} (0.17)	-0.87^{***} (0.26)
(3c.1) Utilization (Δv)	-1.27^{***} (0.25)	-0.99^{***} (0.18)	-1.42^{***} (0.41)	-1.76^{***} (0.53)	-1.81^{***} (0.17)	-1.32^{***} (0.23)
(3c.2) Utilization-adjusted TFP (Δa)	$\begin{array}{c} 0.29\\ (0.24)\\ \end{array}$	$\begin{array}{c} 0.12\\ (0.18)\end{array}$	0.85^{**} (0.41)	1.39^{**} (0.63)	$\begin{array}{c} 1.08^{***} \\ (0.17) \\ \end{array}$	0.45^{**} (0.22)
Coefficients are estimated using an two-step instrumental	variable appros	ach in express	on, where the ins	trument is the	shock and the inst	rumented variable

is the change in the unemployment rate. For each variable x, the entries shown are the slope coefficients from estimating $\Delta_4 x_{jt} = \mu_j + \beta_j \Delta_4 U_t + \varepsilon_{jt}$,

where $\Delta_4 x_{jt}$ is the four-quarter growth rate of x and $\Delta_4 \hat{U}_t$ is the fitted value obtained from the regression of four-quarter percentage-point change in unemployment rate on twelve lags of each shock (instrument). Output is measured as the average of real expenditure and real income. The first column reports results using a monetary shock obtained fromRomer and Romer (2004). The second column reports oil shocks obtained from Hamilton (1996). The third column reports the results using a fiscal shock obtained from Ramey (2011). The fourth column reports results using a fiscal shock obtained from changes in defense spending. The last two columns reports results obtained from a consumption and an investment TFP shocks from Fernald (2014).

The sample runs from 1949Q1 to 2014Q1.

Table 6: Margins of adjustment – instrumental variables by type of shock



Figure 11: Okun coefficients: real income and real expenditure sides

The figure reports the relationships between unemployment and output over time. Output is measured from either the income or the expenditure sides. Each series corresponds to the 40-quarter-rolling-window regression estimates of the intercept, in panel (a), and the slope coefficient, in panel (b), in $\Delta_4 x_{jt} = \mu_j + \beta_j \Delta_4 U_t + \varepsilon_{jt}$, where $\Delta_4 x_{jt}$ is the four-quarter growth rate of x (output measured from income and expenditure sides) and $\Delta_4 U_t$ is the four-quarter percentage-point change in unemployment. The crosses on each series represent periods when the differences in the two coefficients are statistically significant. The sample runs from 1949Q1 to 2014Q1.





The figure reports the accumulated value of the dynamic multipliers associated with the coefficients of the monetary shock from the regression of each margin of adjustment on four lags of itself, the shock, and up to twelve lags of the shock. Monetary shock is obtained from Romer and Romer (2004) and the sample runs from 1969Q1 to 1996Q4. Cample runs from the correspond to 95% confidence bands.





The figure reports the accumulated value of the dynamic multipliers associated with the coefficients of the oil shock from the regression of each margin of adjustment on four lags of itself, the shock, and up to twelve lags of the shock. Oil shock is obtained from Hamilton (1996). Sample runs from 1974Q1 to 2014Q1. Dashed lines correspond to 95% confidence bands.



Figure 14: Impulse response functions – consumption-goods TFP shock

The figure reports the accumulated value of the dynamic multipliers associated with the coefficients of the consumption-goods TFP shock from the regression of each margin of adjustment on four lags of itself, the shock, and up to twelve lags of the shock. Consumption-goods TFP are obtained from Fernald (2014). Sample runs from 1949Q1 to 2014Q1. Dashed lines correspond to 95% confidence bands.



Figure 15: Impulse response functions – investment-goods TFP shock

The figure reports the accumulated value of the dynamic multipliers associated with the coefficients of the investment-goods TFP shock from the regression of each margin of adjustment on four lags of itself, the shock, and up to twelve lags of the shock. Investment-goods TFP are obtained from Fernald (2014). Sample runs from 1949Q1 to 2014Q1. Dashed lines correspond to 95% confidence bands.





Figure 16: Okun coefficient: business and non-business sectors

The figure reports the relationships between unemployment and output over time. Each series corresponds to the 40-quarter-rollingwindow regression estimates of the intercept, in panel (a), and the slope coefficient, in panel (b), in $\Delta_4 x_{jt} = \mu_j + \beta_j \Delta_4 U_t + \varepsilon_{jt}$, where $\Delta_4 x_{jt}$ is the four-quarter growth rate of x (business, non-business and total output) and $\Delta_4 U_t$ is the four-quarter percentagepoint change in unemployment. Non-business output is from NIPA Table 1.3.3. The sample runs from 1949Q1 to 2014Q1.