Reallocation and the Changing Nature of Economic Fluctuations*

(***Preliminary and incomplete—do not cite***)

Julio Garin
University of Notre Dame

Michael Pries
University of Notre Dame

Eric Sims
University of Notre Dame & NBER

September 2010

Abstract

We document three changes in the nature of economic fluctuations that coincide with the onset of the Great Moderation and the Jobless Recoveries phenomenon. Labor productivity turned from strongly procyclical to countercyclical; sectoral reallocation of labor increased and became less temporally concentrated in the initial stages of a recession; and much of the reduction in volatility of various aggregate time series was concentrated in the higher frequency component of those series, while the switch in cyclicality is concentrated in the medium frequency component. We construct a model of labor reallocation that can account for these facts.

Keywords: Reallocation; Great Moderation; Business Cycles
JEL Classifications: E24, E32

1 Introduction

Robert Lucas (1977) famously wrote that "One is led by the facts to conclude that, with respect to the qualitative co-movements among series, *business cycles are all alike*." This

*Corresponding Author: Michael Pries, mpries@nd.edu. We thank Nelson Mark, Robert Flood, Daniel Cooper, and seminar participants at Notre Dame for helpful comments and suggestions. Any remaining errors are our own.

general observation has served as a guidepost for modern research on economic fluctuations. Models are typically evaluated in terms of how well they can reproduce a number of empirical moments, under the implicit assumption that these moments themselves are time invariant. However, the facts now lead one to conclude that this assumption is no longer true: key features of business cycles have changed. This paper documents the changing nature of economic fluctuations in the United States, develops a model that can account for the changes, and analyzes the policy implications of the changes.

Some of the changes are more widely known than others. It has been well-documented that the volatility of output, inflation, and other economic aggregates declined substantially in the mid-1980s. This broad-based decline in volatility has been dubbed the "Great Moderation" and has been the focus of a substantial body of research (for example: Kim and Nelson (1999), McConnell and Perez-Quiros (2000), Stock and Watson (2003), Davis and Kahn (2008)). A second well-known change is the "Jobless Recovery" phonemenon: job growth following the last three recessions has been considerably more anemic when compared to earlier post-war recessions (Groshen and Potter (2003), Bachmann (2009), Gordon and Baily (1993), and Aaronson et al. (2004)).

This paper is motivated by three other significant changes in the nature of economic fluctuations that occurred alongside the Great Moderation and the Jobless Recoveries, but which have received little or no attention:

1. Labor productivity switched from strongly procyclical prior to 1984 to countercyclical in the period since 1984. In addition, the co-movements between productivity and labor market aggregates (e.g. hours, employment, and unemployment) switched signs at the same time.

2. The extent of reallocation of labor across sectors of the economy has increased in post-1984 recessions, while at the same time becoming less concentrated during the initial stages of a recession. That is, recessions in the post-1984 period are characterized by less sharp, but more protracted, episodes of reallocation when compared with recessions prior to 1984.

3. Much of the decline in volatility of various economic aggregates that has occurred since 1984 is attributable to reductions in high-frequency movements rather than medium-

to low-frequency movements. In addition, the changing sign of the correlation between labor productivity and other variables appears attributable to movements at medium- to low-frequencies.

As we show below, these changes are stark, and they coincide with the conventional dating of the onset of the Great Moderation. We contend that these changes are related by more than mere temporal coincidence and are in fact united as part of a broader change in the functioning of the economy. As such, they call for a single, unifying explanation.

Understanding and explaining these changes in the nature of economic fluctuations is an important task for macroeconomic research. Leading state-of-the-art DSGE models now in common use for policy analysis are ill-equipped to account for many of these changes. Accounting for the shift from procyclical to countercyclical productivity is particularly important. For example, the highly-touted Smets and Wouters (2007) estimated New Keynesian model yields highly procyclical labor productivity, even when the model is estimated only using post-1984 data. That is, the model is such that when they estimate it, the estimated productivity shocks that can account for the co-movements of the key aggregates are highly procyclical, even though in the data labor productivity has now turned highly countercyclical.

In many ways, procyclical productivity has been a mainstay of modern macroeconomic models for at least two decades. Indeed, standard one-sector DSGE models in which fluctuations are driven by productivity shocks will, essentially by construction, feature procyclical productivity. More generally, even in macro models with other kinds of shocks, but with standard neoclassical firms, countercyclical productivity is difficult to achieve. With a Cobb-Douglas production function, average labor productivity is proportional to the marginal product of labor and so if firms set the marginal product equal to the wage, then labor productivity can rise during recessions only if wages rise too. But standard models of fluctuations do not feature countercyclical wages, nor is this a feature of the data (either pre- or post-1984; see, e.g., Solon et al. (1994)).

While there has been considerable work aimed at understanding the Great Moderation and the Jobless Recovery phenomenon, there has been little research into the three other changes that we highlighted above. Moreover, attempts to explain the Great Moderation or Jobless Recoveries have implications that are often at odds with the other ways in which

economic fluctuations have changed. For example, one of the more popular explanations for the Great Moderation is the so-called "good luck hypothesis" (Stock and Watson, 2003), which holds that there has been a reduction in the magnitude of shocks hitting the economy. However appealing this explanation may or may not be, by itself it is ill-equipped to account for switching signs of the correlations among aggregate variables. Popular explanations for the Jobless Recovery phenomenon often center on labor hoarding stories (e.g. Bachmann (2009)). The difficulty with this explanation is that hoarding tends to make productivity decline in recessions, whereas productivity has clearly increased during each of the last three recessions. One possible explanation for the declining cyclicality of productivity is that demand shocks might have become more important relative to supply shocks in the post-1984 period. This story would not, however, be consistent with our frequency analysis nor with the facts concerning reallocation across sectors. Gali and van Rens (2010) propose an explanation of the vanishing procyclicality of productivity based on reduced hiring frictions. The difficulty with this explanation is that it is inconsistent with the Jobless Recovery phenomenon. It also leads to only very small reductions in output volatility in a quantitative model.

We seek a unifying explanation for these new facts. Our hypothesis is that the changing nature of reallocation lies at the heart of the various other changes. The changing importance and changing pace of reallocation can help to explain the shift in the cyclicality of productivity: as labor moves from less productive to more productive sectors during an economic recovery, all else equal, aggregate productivity will rise through a compositional effect. Since reallocating and re-training workers takes time, the greater, and more protracted, reallocation can also help to account for the Jobless Recovery phenomenon. Moreover, the changing nature of reallocation is consistent with (1) a reduction in high-frequency volatility relative to low-frequency volatility, (2) lower-frequency movements accounting for the changing cyclicality of productivity, and (3) the changing behavior of layoffs.

Motivated by these observations, we construct a model of labor reallocation that is capable of jointly explaining the facts highlighted above. The model features two islands, or sectors, and island-specific productivity shocks precipitate reallocation of workers from one island to the other. There are frictions associated with reallocating workers that make it a time-consuming activity and result in a pool of "reallocation unemployment." In addition

to island-specific shocks, there are aggregate shocks that are common to both islands. Labor is assumed to be indivisible so that when an aggregate shock hits, some workers on an island become unemployed—not due to frictions or because they are reallocating, but because the disutility of labor effort makes it optimal to reduce labor in response to the aggregate shock. Furthermore, we assume that the island-specific productivity shocks are temporally correlated with negative aggregate productivity shocks, so that when a negative aggregate productivity shock hits, there is also in increase in reallocative activity.

We use this basic model to explore the hypothesis that the above cited changes in the nature of fluctuations resulted from a shift in the relative importance of aggregate and reallocative shocks. In a world in which aggregate shocks are dominant, the model should behave very much like a one-sector neoclassical growth model. While there will be some reallocation across sectors during economic downturns, the aggregate shocks dominate and so productivity will be procyclical and employment declines and recoveries will be sharp and quick. We interpret this environment in which aggregate shocks dominate as a description of the pre-1984 world. The post-1984 version of the model is one in which aggregate shocks are smaller relative to island-specific shocks. This could be due either to good luck or, in a reduced form sense, to more aggressive countercyclical fiscal and monetary policy. With more dampened aggregate shocks, reallocative shocks play a relatively more important role during downturns. Clearly the dampened aggregate shocks will deliver a "Great Moderation." It is also apparent that productivity could become countercyclical (fact 1 above) because during downturns, when employment and output are low, labor productivity rises as labor is reduced in the relatively less productive island and gradually increased in the relatively more productive island. Moreover, because a larger fraction of the unemployment during a downturn is "reallocation unemployment", which is associated with frictions that make the reallocation time-consuming, employment recoveries should be more drawn out ("Jobless Recovery"), as should the level of reallocative activity (fact 2 above). We use the model to explore these predictions and find that that indeed it can deliver the changes outlined above.

Given a model that can account for the facts, we (intend to) examine some policy implications of our findings. For example, what are the welfare effects of monetary or fiscal policy that successfully dampens the impact of aggregate shocks? Does it merely draw out the reallocation of resources that would normally occur in a shorter timeframe if policy did

not counteract the shock? What role do unemployment insurance and worker retraining programs play in affecting the nature of fluctuations? Are unemployment benefits part of the reason why labor reallocation seems to take longer, or should benefits be increased/extended to help "soften the blow" to those affected by necessary reallocation?

The remainder of this paper is organized as follows. Section 2 provides evidence on the three key facts mentioned above. Section 3 lays out the model and section 4 carries out quantitative exercises to show that it can account for the facts. The final section discusses the broader implications of our research and speculates about the potential policy implications of our findings.

2 Evidence on Changes in the Nature of Economic Fluctuations

In this section we provide evidence on the three facts discussed above. Because the facts about the Great Moderation and the Jobless Recovery phenomenon are more broadly known, we do not re-state that evidence here. Unless otherwise noted, all of the data used in the analysis here cover the period 1947 to 2010, are publicly available on US government agency websites, are quarterly in frequency, are seasonally adjusted, and are expressed in natural logs where appropriate.[1]

2.1 The Changing Cyclicality of Productivity

The left panel of figure 1 shows HP detrended average labor productivity and output in the non-farm private business sector. Shaded areas indicate NBER recession dates. It is visually apparent that the two series are positively correlated initially before turning slightly negatively correlated. This visual impression is confirmed in the right panel of figure 1, which shows the rolling 40-quarter forward correlations between output and labor productivity.[2]

[1] For the most part, we isolate the cyclical components of series using the HP filter. Similar results obtain when looking at growth rates.

[2] A 40-quarter window is fairly standard in the literature on the Great Moderation. A similar picture emerges if we use a wider or narrower window for the rolling correlation. A narrower window leads to a somewhat more gradual decline because with a 40-quarter window the 1980-81 recession, which exhibits positive correlation of the two series, rolls out of the window at the same time that the 1991 recession, which exhibits negative correlation, rolls into the window. Although the decline is slightly more gradual with the

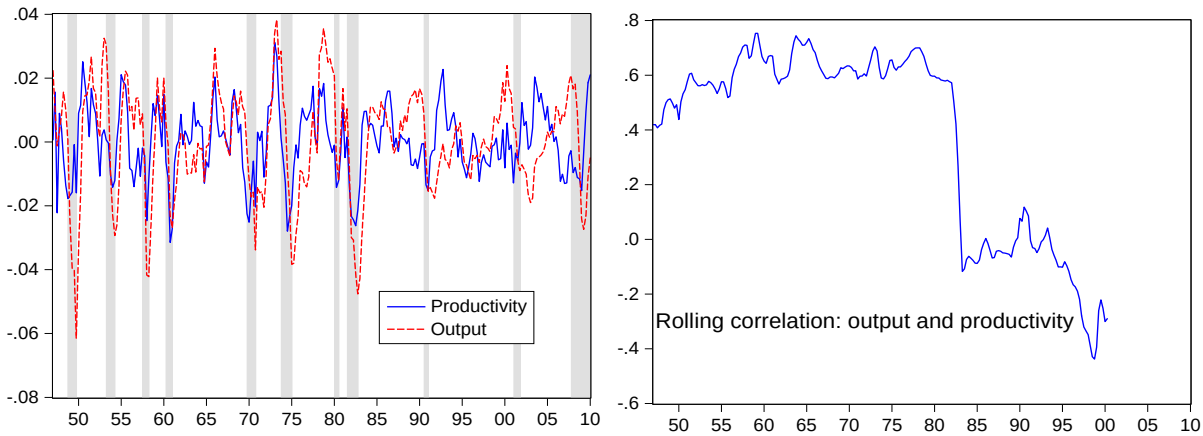

Figure 1: The left panel plots HP filtered real GDP and productivity together, with shaded lines denoting recessions. The right panel plots the 40-quarter forward rolling correlation between HP filtered GDP and productivity.

Productivity shifts rather suddenly from strongly procyclical to moderately countercyclical. The timing of this abrupt change coincides closely with the conventional timing of the onset of the Great Moderation (i.e. 1984).

The change in the correlation between labor market indicators and productivity is equally, if not more, striking.[3] The left panel of figure 2 shows HP detrended labor productivity and total hours in the non-farm business sector. It is visually apparent that the co-movement of these series becomes quite different in the later part of the sample. This is confirmed in the right panel, which shows the analogous rolling 40-quarter forward correlations between the series. The correlation goes from roughly 0.2 pre-1984 to -0.6 post-1984.

Table 1 summarizes these changing correlations between productivity and aggregate variables. We split the sample at the end of 1983 to coincide with the conventional date for the onset of the Great Moderation.

narrower window, the magnitude of the decline is the same.

[3]Hagedorn and Manovskii (forthcoming) have argued that constructing productivity using the employment level from the CPS, as opposed to total hours from the CES, yields different conclusions about the cyclicality of productivity. In their empirical analysis, however, they fail to exclude government and self-employed workers from the employment measure so as to properly measure the denominator in constructing productivity (they use private sector business output in the numerator). Even though they claim that this adjustment "does not matter," it clearly does. Properly accounting for workers does yield a correlation between labor input and output that is weakly positive in the later period (as opposed to modestly negative in the CES data) and a weakly negative correlation between labor input and productivity (as opposed to strongly negative in the CES). It is nevertheless true that there is a large break in the correlations in the CPS data. At a more fundamental level, there is little justification for measuring labor input or productivity from the household survey. The CPS is designed to measure labor force participation and unemployment; the CES is designed to gauge actual labor input.

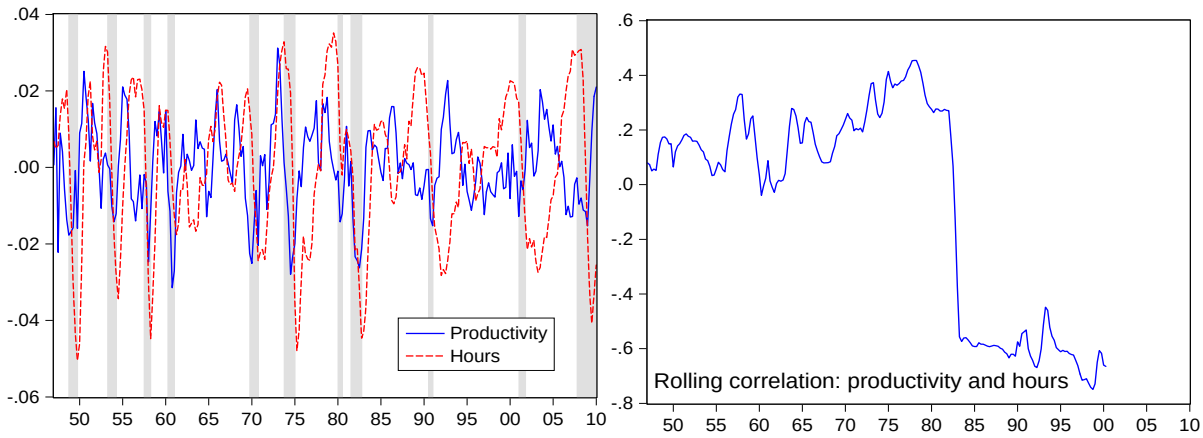

Figure 2: The left panel plots HP filtered total hours and productivity together, with shaded lines denoting recessions. The right panel plots the 40-quarter forward rolling correlation between HP filtered hours and productivity.

	1947-1983	1984-2010
Corr(Output, Productivity)	0.579	-0.127
Corr(Hours, Productivity)	0.211	-0.631
Corr(Employment, Productivity)	0.096	-0.601
Corr(Output, Hours)	0.898	0.862

Table 1: Correlations of productivity with other key aggregates.

In addition to the contemporaneous correlation between productivity and labor market variables, the dynamic relationships have changed as well. In the earlier period, the onset of a recession was accompanied by a sharp drop in productivity, followed soon after by a sharp decline in hours. In the initial periods following these earlier recessions, productivity and hours typically experienced robust recoveries (with the recovery in hours lagging that of productivity by a quarter or two). In the later period, however, the onset of a recession has been associated with small or non-existent declines in productivity, accompanied by relatively large declines in hours. During the recession and into the recovery, productivity grows robustly, yet hours continue to fall before rebounding slowly.

This marked change can be clearly seen in figure 3, which plots the cross-correlogram between productivity and hours in the two sub-periods. In the pre-1984 period, productivity positively leads hours—high productivity is followed by high hours shortly thereafter. In the post-1984 period, productivity positively leads hours with a delay of nearly two years—peaking productivity is associated with declines in hours that persist for at least a year.

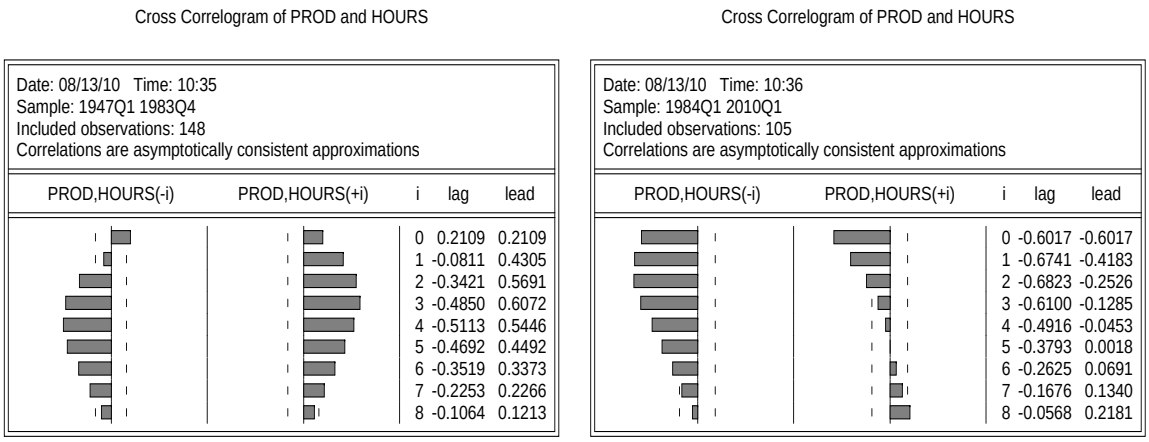

Figure 3: In each panel, the left column shows the correlation between current productivity and lagged hours at different lags. The right column shows the correlation between current productivity and led hours at different leads. The left panel is for the period 1947-1983. The right panel is for 1984-2010.

2.2 The Changing Nature of Sectoral Reallocation

The idea that recessions may be associated with elevated levels of reallocation of workers and resources across different sectors of the economy goes back to at least Lilien (1982). He proposes the dispersion of employment growth rates across sectors as an indicator of the extent of reallocative activity. Specifically, let $g_{i,t} = \ln N_{i,t} - \ln N_{i,t-1}$ be the quarterly growth rate of the level of employment in sector i at time t. Let g_t be defined similarly for aggregate private employment. Further, let $s_{i,t} = \frac{N_{i,t}}{N_t}$ be sector i's share of aggregate employment at time t and μ_i be the time series average growth rate of sectoral employment relative to aggregate employment growth, $g_{i,t} - g_t$ (e.g., this will be negative for a declining sector like manufacturing). Lilien's measure of reallocation is then the share-weighted standard deviation of net sectoral growth rates:

$$\sigma_t^L = \left[\sum_{i=1}^{M} s_{i,t} (g_{i,t} - g_t - \mu_i)^2 \right]^{\frac{1}{2}} \tag{1}$$

where M denotes the number of sectors. We take the ten one-digit SIC industries as a benchmark, so $M = 10$. Figure 4 shows this measure using quarterly data.

Lilien argued that if recessions were caused by an increase in reallocative activity then we should expect spikes in this measure of dispersion near recessions. The evidence in figure 4 is consistent with this story, as the dispersion measure is clearly countercyclical. However,

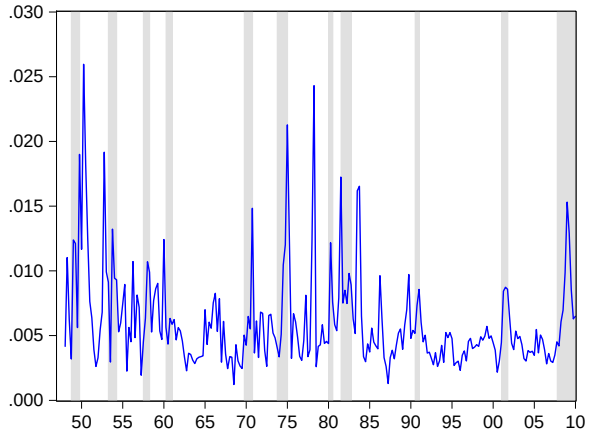

Figure 4: This figure shows the Lilien (1982) measure of reallocation, constructed as described in the text.

subsequent researchers, such as Abraham and Katz (1986), questioned this logic, pointing out that even if recessions were purely the result of aggregate disturbances, the dispersion measure would still spike if some sectors were naturally more cyclically sensitive.

We are not as interested in the question of whether aggregate or reallocative shocks are the fundamental impulse that drive economic fluctuations. The model outlined below incorporates both aggregate shocks and reallocative sectoral shocks. Even if aggregate shocks are fundamentally what drives economic fluctuations, there is still reason to expect that reallocative activity might increase during downturns when the opportunity cost of undertaking reallocative activity is lower.[4] Moreover, this reallocative activity can have important implications for the nature of fluctuations.

Thus, what is of interest to us is the question of whether the cyclical properties of reallocation have changed. If we interpret the Lilien dispersion measure as an indicator of reallocative activity, figure 4 suggests that post-1984 recessions have been accompanied by *less* reallocation (the spikes during these later recessions are smaller). We contend, however, that the Lilien measure is a poor indicator of the level of reallocative activity because it does not reflect the *permanence* of changes in sectoral employment. For example, if a sector of the economy contracted sharply in one year and recovered back to its initial employment in the following year, the Lilien measure would be high in each year, even though no lasting

[4]See Ramey (1991), Hall (1991), Caballero and Hammour (1996), Aghion and Saint-Paul (1998) for more on the idea that recessions are similar to the yellow caution flag in auto racing, in that they are the optimal time for a "pit stop."

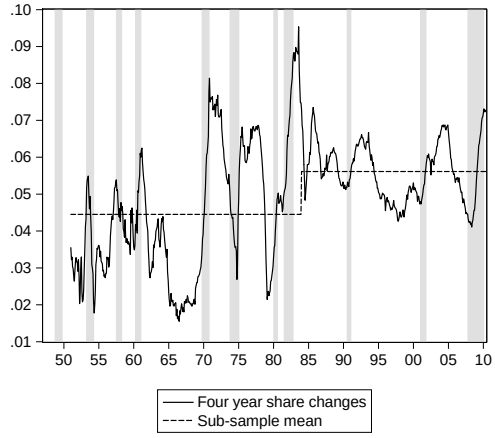

Figure 5: This figure shows the rolling four-year change in employment shares, as described in the text. The dashed line shows the mean for each sub-sample (pre- and post-1984).

reallocation actually took place.

We therefore propose two different approaches to assessing the extent to which changes in sectoral employment actually represent enduring reallocation of resources across different sectors.[5] First, we measure, at each point in time, the fraction of employment that has been reallocated to another sector relative to the distribution of employment four years earlier. More formally, using monthly sectoral employment shares we compute on a rolling basis $\sum_{i=1}^{10} |s_{i,t} - s_{i,t-48}|$. The time series for this statistic is plotted in figure 5.

The figure reveals that sectoral reallocation prior to 1984 looks very different than reallocation in the later period. First, the average level of this statistic is substantially higher in the post-1984 period than before (as indicated by the dashed line in the figure). That is, in the later period the distribution of workers across sectors changed more, on average, over four-year intervals than it did in the pre-1984 period. Secondly, while the statistic peaks in and around recessions in the entire sample, the dynamic movements of the statistic are quite different. In the earlier period, recessions were associated with a sharp increase in the rolling four-year change in sectoral shares, followed shortly thereafter by an equally sharp decline. This suggests that, during recessions, the distribution of employment would change dramatically relative to the distribution prior to the onset of the recessions, but would quickly revert

[5] Groshen and Potter (2003) suggest as a measure of reallocative activity the correlation of sectoral employment growth during recessions with sectoral growth in the first year following the recessions. While this measure does offer some indication of whether the reallocation is enduring, it suffers from the fact that it only looks at employment changes in the one year following the end of a recession. In practice, reallocation is likely to take much longer to play out.

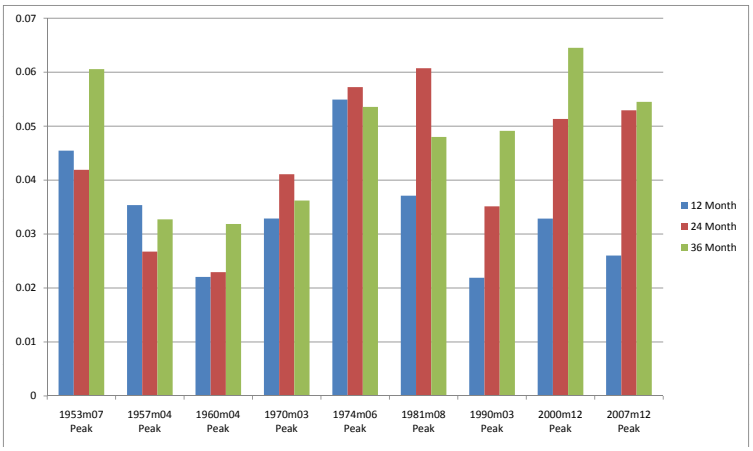

Figure 6: Total change in employment shares at 12 months, 24 months, and 36 months from employment peaks preceding different recessions. The 2007 recession is currently truncated at 30 months, so the 36-month bar indicates the 30-month change in shares. We omit the 1980 recession peak from the picture because there are insufficient "recovery" observations due to the 1981 recession following shortly thereafter. We therefore take August, 1981, as the employment peak.

to a distribution that was similar to the pre-recession distribution. Put differently, recessions prior to 1984 were associated with sharp, but not very persistent, changes in sectoral employment shares. In contrast, recessions in the later period result in reallocation that appears to be less sharp but more persistent. The hump-shaped increase in the four-year changes suggest that changes in the sectoral distribution of employment do not quickly revert, but rather are more enduring.

We now turn to our second approach to assessing whether changes in sectoral employment shares that occur during recessions represent enduring or transitory reallocation. For each post-war recession, we measure the change in each sector's employment share relative to the month in which private non-farm business sector employment peaks. We then take absolute values of these changes and sum over sectors. Letting τ denote the number of periods since the employment peak, we calculate:

$$s_{t,\tau} = \sum_{i=1}^{10} |s_{i,\tau} - s_{i,t}| \tag{2}$$

Figure 6 plots this statistic for $\tau = \{12, 24, 36\}$.

The evidence in figure 6 reinforces the evidence in figure 5. Prior to 1984, employment shares shifted rapidly across sectors, with a large change in the first 12 months following the

aggregate employment peak. However, in general there was little additional change in months 13-24 or 25-36, and in fact in several instances some of the share changes that occurred in the initial 12 months had subsequently reverted by month 36. One notable exception is the 1953 recession. Because the beginning of that recession coincided with the end of the Korean War, it is possible that the prolonged labor reallocation that followed that recession was associated with the transition from a wartime economy to a peacetime economy.

Post-1984 recessions, in contrast, are characterized by modest changes in employment shares during the first 12 months of the recession followed by continued increases by months 24 and 36. Moreover, the cumulative (through 36 months) level of share changes in these three recession is typically larger than in the earlier recessions. This is particularly interesting in light of the fact that the 1990 and 2001 recessions were milder (more modest declines in real GDP) compared to the earlier episodes. In summary, instead of a short burst of reallocative activity, post-1984 recessions exhibit a long, surging wave of reallocation.

We conclude from the evidence presented here that the nature of sectoral reallocation changed in a significant way around 1984. In recessions prior to 1984, sharp disruptions to the distribution of employment across sectors were not very long-lasting. In contrast, in recessions since 1984, reallocative activity was gradual and more protracted.

2.3 Frequency Analysis

Figures 1 and 2 show that the pre-1984 period, as compared to the post-1984 period, is marked by more pronounced spikes in employment, output, and productivity. That is, much of the reduction in the volatility of aggregate variables in the post-1984 period has been concentrated at higher frequencies. To demonstrate this more formally, and to examine whether other changes can be associated with fluctuations of either higher or lower frequencies, we use the bandpass filter of Baxter and King (1999) to decompose hours, output, and productivity into components with different periodicities. We take the conventional definition of business cycle frequencies as being between 6 and 32 quarters. We define "high" business cycle frequency movements as those with periodicities between 6 and 19 quarters, and "medium" business cycle frequency movements as those with periodicities between 19 and 32 quarters. Note that the nature of this filter is such that adding these two components together yields the cyclical component with periodicities between 6 and 32 quarters. "Low"

frequency movements would be those with periodicity above 32 quarters.

Output	Std. of High Freq. Component	Std. of Medium Freq. Component
1947-1983	0.0129	0.0091
1984-2010	0.0049	0.0056
% Change	-62%	-38%
Hours		
1947-1983	0.0133	0.0095
1984-2010	0.0060	0.0092
% Change	-55%	-3%

Table 2: High and Medium Frequency Volatilities

Table 2 shows the standard deviations of the high- and medium-frequency components of both output and hours, both pre- and post-1984. There has been a very large reduction in the volatility of the high-frequency component of output (62%), and a more modest reduction (38%) at medium frequencies. There is a similar reduction in the volatility of the high-frequency component of hours (55%). Interestingly, there is essentially no change in the volatility of the medium-frequency component of hours. Clearly the reduced volatility of high-frequency movements accounts for a significant majority of the reduction in macroeconomic volatility post-1984.

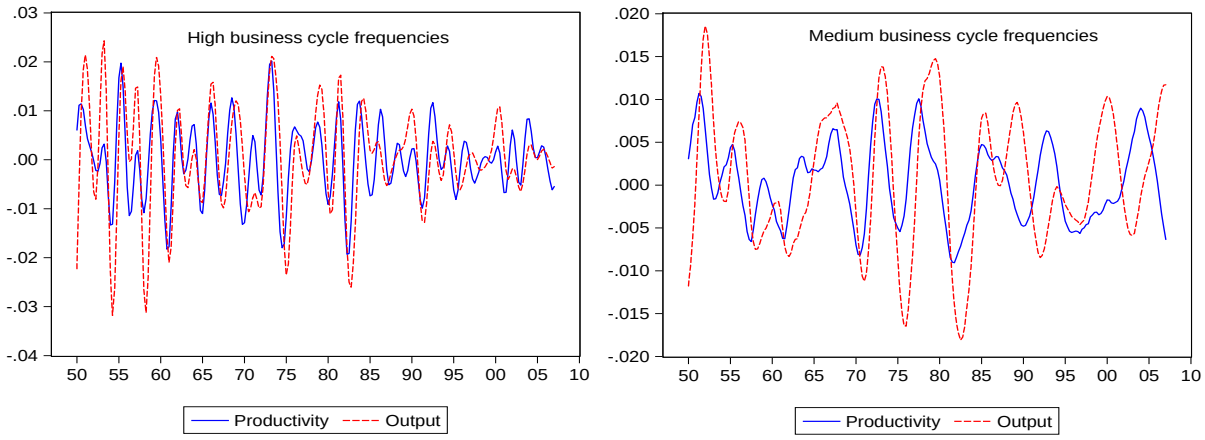

Figure 7: The left panel shows the bandpass filtered series with periodicities between 6 and 19 quarters. The right panel shows periodicities between 19 and 32 quarters.

To assess the role that the different frequencies play in accounting for the changing cyclicality of productivity, figure 7 plots the high- and medium-frequency components of output and productivity together across time. While there is a clear reduction in the volatility of the high-frequency component of both output and productivity in the left panel, there is no discernible change in the co-movement between the series. In contrast, there is a clear break in the correlation between the medium-frequency components of output and productivity in the right panel.

In table 3 we provide quantitative evidence of this changing behavior. Prior to 1984, the correlations between productivity and output were roughly the same at low and high frequencies. Both decline after 1984, but only the correlation of the medium-frequency components of these series declines substantially, switching from 0.58 in the pre-1984 period to -0.24 thereafter. Similar results hold for the relationship between labor market variables and productivity—much of the change in the cyclicality of productivity is driven by the medium-frequency movements.

High Frequency	1947-1983	1984-2010
corr(output, productivity)	0.63	0.37
Medium Frequency		
corr(output, productivity)	0.58	-0.24

Table 3: Medium- and High-frequency Correlations

2.4 Additional Evidence

Data from the BLS Household Survey sheds additional evidence on the nature of reallocation. The BLS asks unemployed respondents to report their reason for unemployment. The responses allow one to identify "temporary" unemployment, which typically counts furloughed workers, and so-called "permanent" unemployment, which measures workers who have lost jobs and do not expect to be re-hired by the same firm. These data are available from the BLS beginning in 1967 and are shown in figure 8 below. The striking feature of figure 8 is the near disappearance of the cyclicality of temporary layoffs in the post-1984 sample period (Faberman (2008) makes a similar point). In the early part of the sample,

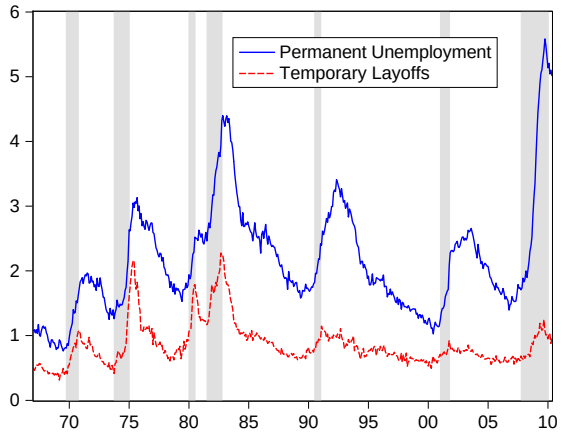

Figure 8: "Temporary" and "Permanent" unemployment. Shaded areas indicate recessions.

both temporary and permanent layoffs spike during recessions. After 1984, there is virtually no spike in temporary layoffs around recessions, while there remain large and persistent increases in the level of permanent layoffs. The increases in permanent layoffs around the time of the 1990-1991 and 2001 recessions are large relative to past recessions given the relative mildness of the output declines.

The near disappearance of temporary layoffs and the increasing importance of permanent layoffs in the post-1984 recessions are consistent with the above evidence regarding the increased, and more persistent, sectoral reallocation. It could also reflect reallocation of workers across firms within sector or the reallocation of workers across geographic regions; the aggregate data make it difficult to discriminate between these possibilities. In future work, we plan to examine the CPS micro data in detail to help evaluate these competing hypotheses.

The newest unemployment data from the most recent recession provide additional clues concerning reallocation. The fraction of total unemployed workers who have been unemployed for half a year or longer is currently twice as high as at any point in the post-war period. Because of the increasing duration of unemployment spells, unemployment has been concentrated among a relatively small class of workers. For example, over the course of 2008 only 13.2 percent of the total labor force experienced unemployment. In 1980, which was by comparison a much milder recession in terms of output decline, 18.8 percent of the labor force experienced unemployment. These facts point to much of the current downturn as being "structural," with workers having difficulty reallocating to growing parts of the economy,

as opposed to cyclical.

3 Modeling the Changing Nature of Fluctuations

This section introduces an island model of labor reallocation that features both island-specific productivity shocks and aggregate productivity shocks. Though highly stylized, the model is nevertheless able to account for the changes we documented above.

There are two islands, which represent different sectors of the economy or different geographic regions. The output produced on island i is given by $y_{i,t} = A_t z_{i,t} L_{i,t}^{\alpha}$, where A_t is an aggregate shock that is common to both islands, $z_{i,t}$ is the island-specific productivity shock on island i, and $L_{i,t}$ is employment on the island. A_t and $z_{i,t}$ both follow Markov processes, to be discussed in more detail below.

There is a continuum of infinitely lived households of measure 2. Each household consists of one worker who has one unit of time that it either fully supplies in the labor market or fully enjoys as leisure (labor is indivisible). Within-period preferences are given by:

$$u(c_t, l_t) = \frac{c_t^{1-\sigma} - 1}{1 - \sigma} - \phi l_t$$

where c_t is consumption and l_t is the labor supplied by the household. Households discount the future at rate ρ (and $\beta = 1/(1+\rho)$ denotes the discount factor).

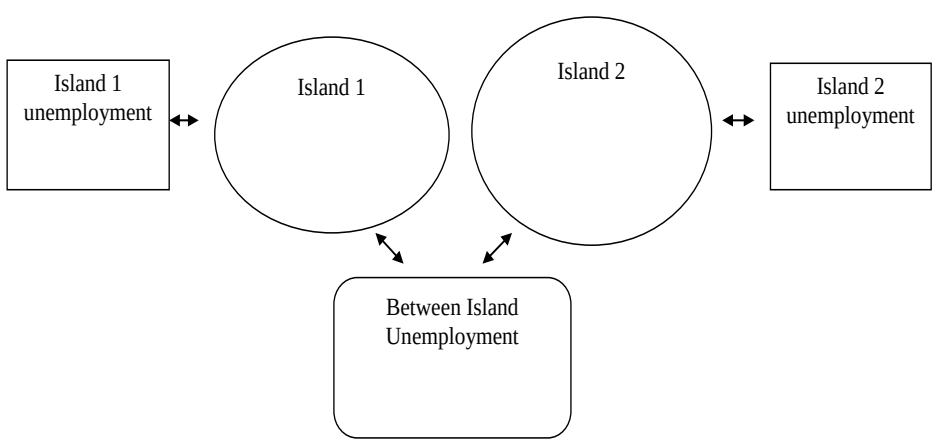

Figure 9: Graphical depiction of the "island" model

The basic structure of the model economy is depicted graphically in figure 9. There are frictions that impede the reallocation of workers from one island to the other. An example of such a friction is the need for re-training in sector-specific skills. We model these types of frictions and the time-consuming nature of reallocation by assuming that when workers move from one island to the other, they must pass through a period of "reallocation unemployment" while engaged in activities that make them employable in the other sector. Workers stochastically escape this reallocation process at the exogenously given rate λ, and we assume that when they escape they can choose which island to move to (this simplifies the model by eliminating the need to keep track of which island each worker in the reallocation process originally came from).[6]

To keep the model tractable, and to focus more clearly on labor market dynamics, we assume that there are no means of transferring resources across time. Moreover, we assume that households have perfect insurance against the idiosyncratic risks that they face (due to shocks to the productivity of the island on which they work, loss of income while reallocating, etc.). This allows us to identify the optimal allocation as the outcome of a social planner problem.

The planner faces the same frictions that households face—reallocating workers from one island to the other is time-consuming. Note that the planner may choose not to employ some workers on an island, while also not reallocating those workers to the other island. That is, holding island-specific productivities constant, if aggregate productivity is temporarily low, then it is optimal to reduce employment on each island (due to the disutility associated with work) without initiating any reallocation.

In this environment, workers at a point in time are in one of three states: (i) working on an island, (ii) on an island but unemployed, with frictionless transitions between working and not working, and (iii) in the state of "reallocation unemployment," where frictions make the transition from one island to another a time-consuming process.

If we let $N_{i,t}$ be the number (measure) of workers allocated to island i, and $L_{i,t}$ be the

[6]Workers will move in response to island-specific shocks, and will move from the less productive to the more productive island, but it is possible that the island-specific shocks will reverse again before the worker escapes the reallocation process, in which case the worker would like to return to the original island.

number of workers who are employed on island i, then the social planner's problem is:

$$\max_{L_{1,t},L_{2,t},N_{1,t},N_{2,t}} \sum_{t=0}^{\infty} \beta^t u(c_t, l_t)$$

$$\text{s.t.: } c_t = A_t(z_{1,t}L_{1,t}^{\alpha} + z_{2,t}L_{2,t}^{\alpha})$$

$$L_{1,t} \leq N_{1,t}$$

$$L_{2,t} \leq N_{2,t}$$

$$N_{1,t} + N_{2,t} \leq 2$$

The first constraint is the resource constraint, the second and third constraints require that employment on an island not exceed the number of workers allocated to that island, and the fourth constraint simply states that the total number (measure) of workers on the two islands must be less than the total number of workers. Note that there will be three different types of unemployed workers: workers located on island 1, but not employed ($U_{1,t} = N_{1,t} - L_{1,t}$), workers located on island 2, but not employed ($U_{2,t} = N_{2,t} - L_{2,t}$), and workers in "reallocation unemployment" ($U_{r,t} = 2 - N_{1,t} - N_{2,t}$).

This planner's problem can be expressed as a straightforward dynamic programming problem. The state variables are the number of workers allocated to the two islands at the beginning of the period, N_1 and N_2, as well as the values of the aggregate and idiosyncratic shocks. To simplify notation, let $x_t = \{A_t, z_{1,t}, z_{2,t}\}$ denote the vector of shocks. The Bellman equation can then be expressed as:

$$V(N_1, N_2, x) = \max_{L_1,L_2,N_1',N_2'} \frac{(A(z_1L_1^{\alpha} + z_{2,t}L_2^{\alpha}))^{1-\sigma} - 1}{1-\sigma} - \phi(L_1 + L_2) + \beta EV(N_1', N_2', x')$$

$$\text{s.t.: } L_1 \leq N_1$$

$$L_2 \leq N_2$$

$$N_1' + N_2' \leq 2$$

We solve this problem numerically using standard techniques. Specifically, we create a grid of values for $N_{1,t}$ and $N_{2,t}$ and then iterate on the Bellman equation above until it converges. Evaluating the value function at points between the gridpoints for $N_{1,t}$ and $N_{2,t}$ requires interpolation. We use a simplicial 2-D linear interpolation (see Judd (1998), p. 242).

4 Quantitative Analysis

In this section we undertake some preliminary quantitative analysis of the model, and show that it is capable of replicating many features of actual US data documented above.

It is first necessary to specify stochastic processes for the exogenous state variables. We assume that the aggregate productivity state follows an AR(1) process with a time-varying mean:

$$A_t = (1-\rho)\mu_t + \rho A_{t-1} + \varepsilon_t \qquad \varepsilon_t \sim N(0, \sigma_\varepsilon^2)$$

The mean is assumed to follow a two-state Markov process, taking on values $\mu^H > \mu^L$. Loosely speaking, one can think of the time-varying mean as governing the "regime" of the aggregate economy, with μ^H interpreted as "normal" times and μ^L interpreted as a "recession". Shocks to ε introduce higher frequency, period-to-period movements in A_t.

We assume that island-specific productivities follow a three-state Markov process, with values $\{z^H, z^M, z^L\}$, and that the two islands' productivities being perfectly negatively correlated. That is, when island 1 has idiosyncratic productivity $z_{1,t} = z^H$, island 2 has idiosyncratic productivity $z_{1,t} = z^L$, and vice-versa.

For now, to provide expositional simplicity, we set $\sigma_\varepsilon^2 = 0$, so that the aggregate state is completely summarized by μ_t. Since μ_t takes on only two values, we can summarize the aggregate state with two values: A^H and A^L. With two aggregate exogenous states and three island-specific ones, there are a total of six configurations, or states, for the vector of productivities, x:

$$x = \begin{bmatrix} A^H, z^H, z^L \\ A^H, z^M, z^M \\ A^H, z^L, z^H \\ A^L, z^H, z^L \\ A^L, z^M, z^M \\ A^L, z^L, z^H \end{bmatrix}$$

There will be a 6×6 transition matrix, P, that will describe the probabilities of moving between states. In particular, $P_{i,j}$ denotes the probability of transitioning from state i to state j. We impose the following structure on the transition matrix: the idiosyncratic pro-

ductivities can only change when the economy transitions to the "low" aggregate state. That is, the onset of a recession is a mix of aggregate and reallocative shocks. This assumption is meant to capture the fact that some sectors typically fare better than others during downturns in the actual data. Under this assumption, the transition matrix takes the following form:

$$P = \begin{bmatrix} p_{1,1} & 0 & 0 & p_{1,4} & p_{1,5} & p_{1,6} \\ 0 & p_{2,2} & 0 & p_{2,4} & p_{2,5} & p_{2,6} \\ 0 & 0 & p_{3,3} & p_{3,4} & p_{3,5} & p_{3,6} \\ p_{4,1} & 0 & 0 & p_{4,4} & 0 & 0 \\ 0 & p_{5,2} & 0 & 0 & p_{5,5} & 0 \\ 0 & 0 & p_{6,2} & 0 & 0 & p_{6,6} \end{bmatrix}$$

The upper right 3×3 matrix denotes the probabilities of leaving the high aggregate state. It is only when entering a recession that the island-specific productivities can change.

In order to investigate the quantitative properties of the model it is first necessary to assign values to the parameters. We caution that this parameterization should not yet be thought of as a rigorous calibration exercise, but rather as an illustrative example. We take the unit of time to be one year. As such, we set the household's subjective discount factor to $\beta = 0.95$. We assume that island-specific productivities can take on one of three values, with a mean of 1: $z^H = 1.05$, $z^M = 1$, and $z^L = 0.95$. We set the curvature parameter in the Cobb-Douglas production function to $\alpha = 0.67$. The parameter σ is set to 1, so that the flow utility function from consumption is the natural log. The disutility from labor, ϕ, is set to $\frac{\alpha}{2}$, or 0.33. This ensures that in a non-stochastic version of the model, with A_t, $z_{1,t}$ and $z_{2,t}$ all set to their mean values, employment on each island would be 1 and hence there would be "full" employment in aggregate. We set $\lambda = 0.5$, meaning that half of all workers in "reallocation" unemployment escape to an island each period. The Markov transition matrix is specified as follows:

$$P = \begin{bmatrix} 0.88 & 0 & 0 & 0 & 0.06 & 0.06 \\ 0 & 0.88 & 0 & 0.06 & 0 & 0.06 \\ 0 & 0 & 0.88 & 0.06 & 0.06 & 0 \\ 0.10 & 0 & 0 & 0.90 & 0 & 0 \\ 0 & 0.10 & 0 & 0 & 0.90 & 0 \\ 0 & 0 & 0.10 & 0 & 0 & 0.90 \end{bmatrix}$$

This parameterization implies that the average duration of "good times" is a little over 8 years, while the average duration of recessions is a little over one year. Both of these durations are loosely consistent with post-war data. The zeros on the "upper-right" diagonal mean that, whenever the economy enters a recession, there must be a change in island-specific productivities.

As we describe below, we consider two separate parameterizations of the aggregate state, A_t, one of which is meant to capture the pre-1984 period and the other the post-1984 period.

4.1 Some Results

We first consider how a change in the relative magnitudes of the shocks to idiosyncratic productivity and to aggregate productivity affect the key macro aggregates of interest, to determine whether the change in the relative importance of the two shocks can explain the changes in the nature of economic fluctuations that were highlighted above. Specifically, we consider two separate parameterizations of the aggregate state, one of which is meant to capture the pre-1984 period and the other the post-1984 period. For the early period, we impose that the aggregate productivity series takes a value of 1.02 in good times and 0.98 in bad times. In the post-1984 period, aggregate shocks are smaller. In particular, $a^H = 1.005$ and $a^L = 0.995$.

The table below presents some statistics from a simulation of the model described above. The bottom half of the table presents statistics from actual US data, both pre- and post-1984. So as to make the data consistent with the annual calibration in the model, we HP filter the logs of the annual series with smoothing parameter of 100. As noted previously, in the early part of the sample output and hours are about as volatile as one another and productivity is strongly correlated with both employment and output. In the later part of

	Large Agg. Shock	Small Agg. Shock	Change
s.d.(output)	0.022	0.016	-47%
s.d.(output)/s.d.(employment)	1.642	0.855	-28%
corr(output, productivity)	0.716	-0.289	-0.99
corr(employment, productivity)	0.179	-0.568	-0.74
	pre-1984 Data	post-1984 Data	Change
s.d.(output)	0.025	0.017	-33%
s.d.(output)/s.d.(employment)	1.087	0.680	-39%
corr(employment, productivity)	0.636	-0.010	-0.65
corr(employment, productivity)	0.273	-0.504	-0.77

Table 4: Comparison of the simulated model with the data.

the sample, the volatility of output falls relative to that of employment, and productivity switches from positively to negatively correlated with both output and hours. The absence of a decline in the volatility of hours in the post-1984 period is an important feature of the "Great Moderation" phenomenon, which has not been associated with declining volatility in labor markets.

The top half of the table shows analogous statistics from a simulation of the model. For the simulation we generated 30,000 observations from the parameterized model as described above. We then took logs and HP filtered (smoothing parameter 100) the relevant series.

Qualitatively, the model does a very good job of matching most of these statistics. With large aggregate shocks, the model yields output that is slightly more volatile than employment and correlations between productivity and output that are positive, with the correlation between output and productivity larger than the correlation between productivity and employment.

With smaller aggregate shocks, the model's moments move in the same direction (and in similar magnitudes) as observed in the actual data. The changing magnitude of the aggregate shock induces a very similar drop in output volatility to that which is observed in the data. Importantly, the volatility of output falls substantially relative to that of employment, in a manner similar to that observed in the actual data. Further, the correlations of productivity with output and hours both switch to negative, with the productivity/employment correlation more negative than the productivity/output correlation. As the "change" column

indicates, the quantitative magnitude of the changes in all of these moments in the model are very similar to the US data.

The parameterized model has other features which closely align with the changes observed in aggregate US data. First, when the magnitude of aggregate shocks becomes smaller, the extent to which employment lags output increases. That is, the correlation between employment and output lagged several periods increases. This features aligns well with the "jobless recovery" phenomenon. Second, smaller aggregate shocks are accompanied by a slight reduction in the volatility of the Lilien (1982) style measure of reallocative activity and a slight increase in the extent to which that measure is negatively correlated with aggregate output. Both of these features are apparent in the data (see figure 4 above). Since in our model with smaller aggregate shocks we know that reallocative activity is relatively more important, this reinforces our contention, above, that the Lilien (1982) style measure is not a good metric for the importance of reallocation over the business cycle.

We conduct a number of robustness checks using different values for some of the parameters. Figure 10 below shows how some of the simulated moments of the model change when (i) the parameter λ changes and (ii) the amount of dispersion in the aggregate shock varies. The labeling is such that "ξ" means that the "high" aggregate state is $A^H = 1 + \xi$ and the "low" aggregate state is $A^L = 1 - \xi$.

We see that the productivity-output correlation, the productivity-employment correlation, and output volatility relative to employment volatility are all clearly decreasing in the magnitude of aggregate shocks. Moreover, there is almost no difference in the mean and volatility of aggregate unemployment in the samples for different levels of aggregate dispersion holding λ fixed. This also roughly consistent with the aggregate data, since in the data the mean and standard deviation of unemployment do not differ significantly across the two time periods

Table 10 also reveals that the parameter λ, which loosely governs the "cost" of reallocating workers across sectors, has only small impacts on these statistics. Different values of λ mainly contribute to different levels of average unemployment in the simulations. The value of λ will mainly influence the pace and nature of reallocative activity.

A realistic addition to the model would be to endogenize the job-finding rate, λ. It particular, it seems plausible that λ would be a decreasing function of total reallocation

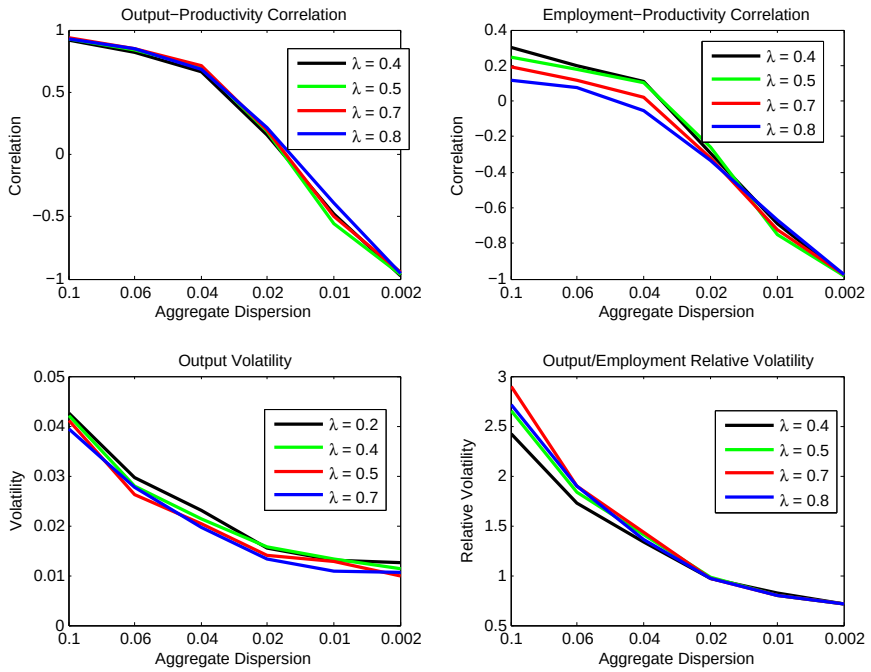

Figure 10: This figure shows various statistics of interest for different parameterizations of the model. The horizontal axis refers to the magnitude of the dispersion in the aggregate shock.

unemployment—the greater the number of workers in the process of reallocation, the harder it is for them to find a job. This feature would slow down the reallocation process since it would become optimal to spread out the inflow of workers into reallocation. This would allow the model to better capture the long surging wave of reallocation shown above in figure 6. When aggregate shocks are large and most of the employment decline during recessions is "island-specific", employment will recovery quickly following a downturn. When aggregate shocks are smaller and more of the employment decline is due to reallocative motives, an endogenous λ would work to slow down the reallocative process in much the same manner we observe in the data. An endogenous λ would also likely lead to an increase in the relative cyclicality of reallocation unemployment relative to island-specific unemployment, which is a feature of the data (see figure 8) but not readily apparent in the current version of the model.

The model as currently specified is too stylized to be able to capture some of the other facts documented in the empirical part of the paper. In the current version of the model, the aggregate state follows a step function (i.e. it is either high or low). This makes applications of the band-pass filter difficult to interpret, and therefore is not well-suited to the frequency

analysis presented above. Allowing the variance of the "higher" frequency aggregate shock, σ_ε^2 to be greater than zero should make this exercise more interesting.

We have not yet taken into consideration alternative parameterizations of the transition matrix or the magnitudes (and number) of the island-specific states. We suspect that including a greater number of island-specific states will enrich the dynamics of the model and make some of the frequency analysis more meaningful. We hope also to bring disaggregated data (say, from the 10 one digit SIC sectors) to bear to help discipline the magnitudes of these states and the probabilities of the Markov transition matrix.

5 Broader Implications and Policy Considerations

The business cycle has changed in dramatic ways in the last twenty-five years. This paper has (i) document the various important dimensions along which the business cycle has changed and (ii) developed a model of labor reallocation for understanding those changes.

The facts documented here should be of particular interest to macroeconomists. Macro models are typically calibrated or estimated to fit US data for the entire post-war period. When modeling exercises do attempt to account for the Great Moderation, the objective is usually only to match the reduced volatility of output and/or inflation. Our analysis suggests that models should also be judged on how well they can replicate the changing co-movements and frequency properties. Leading state-of-the-art macro models of which we are aware do not account for these new facts. For example, in the estimated Smets and Wouters (2007) New Keynesian model of the US economy labor productivity is procyclical and strongly correlated with hours worked.

Our analysis has important implications for economic policy. If recessions are increasingly about reallocation, then this raises the question of how aggressive countercyclical demand management policies should be. Stimulating demand through aggressive monetary easing or fiscal expansion may only serve to postpone the necessary reallocation of resources; it could also have longer term adverse consequences concerning productivity growth and human capital accumulation.

Part of the current political debate concerns unemployment benefits. Should we have such benefits? For how long should recipients receive benefits? Our theoretical model should be

able to shed some light on these questions. Since reallocation of workers across sectors or geographic boundaries naturally takes time, it may make sense to have more generous and long-lasting unemployment benefits to help cushion the blow of those disaffected workers. On the other hand, unemployment benefits reduce the incentive for workers to search for new employment and to acquire new skills. To the extent that reallocation is an increasingly important part of the business cycle, our model should be able to draw some conclusions about the role of unemployment insurance and to therefore help inform policy-making.

References

Aaronson, D., Rissman, E., and Sullivan, D. G. (2004), "Can sectoral reallocation explain the jobless recovery?", *Economic Perspectives, Federal Reserve Bank of Chicago*, **28**, 36–49.

Abraham, K. G. and Katz, L. F. (1986), "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?", *Journal of Political Economy*, **94**, 507–22.

Aghion, P. and Saint-Paul, G. (1998), "On the Virtue of Bad Times: An Analysis of the Interaction between Economic Fluctuations and Productivity Growth", *Macroeconomic Dynamics*, **2**, 322–344.

Bachmann, R. (2009), *Understanding Jobless Recoveries*, Ph.D. thesis, Yale University.

Baxter, M. and King, R. G. (1999), "Measuring Business Cycles: Approximate Band-Pass Filters For Economic Time Series", *The Review of Economics and Statistics*, **81**, 575–593.

Caballero, R. J. and Hammour, M. L. (1996), "On the Timing and Efficiency of Creative Destruction", *The Quarterly Journal of Economics*, **111**, 805–52.

Davis, S. J. and Kahn, J. A. (2008), "Interpreting the Great Moderation: Changes in the Volatility of Economic Activity at the Macro and Micro Levels", NBER Working Papers 14048, National Bureau of Economic Research.

Faberman, R. J. (2008), "Job flows, jobless recoveries, and the Great Moderation", Working Papers 08-11, Federal Reserve Bank of Philadelphia.

Gali, J. and van Rens, T. (2010), "The Vanishing Procyclicality of Labor Productivity", IZA Discussion Papers 5099, Institute for the Study of Labor (IZA).

Gordon, R. J. and Baily, M. N. (1993), "The Jobless Recovery: Does it Signal a New Era of Productivity-Led Growth?", *Brookings Papers on Economic Activity*, **1993**, 271–316.

Groshen, E. L. and Potter, S. (2003), "Has structural change contributed to a jobless recovery?", *Current Issues in Economics and Finance, Federal Reserve Bank of New York*, **9**.

Hagedorn, M. and Manovskii, I. (forthcoming), "Productivity and the Labor Market: Co-Movement over the Business Cycle", *International Economic Review*.

Hall, R. E. (1991), "Labor Demand, Labor Supply, and Employment Volatility", *NBER Macroeconomics Annual*, **6**, 17–62.

Judd, K. L. (1998), *Numerical Methods in Economics*, number 0262100711 in MIT Press Books (The MIT Press).

Kim, C.-J. and Nelson, C. R. (1999), "Has The U.S. Economy Become More Stable? A Bayesian Approach Based On A Markov-Switching Model Of The Business Cycle", *The Review of Economics and Statistics*, **81**, 608–616.

Lilien, D. M. (1982), "Sectoral Shifts and Cyclical Unemployment", *Journal of Political Economy*, **90**, 777–793.

Lucas, R. E. (1977), "Understanding business cycles", *Carnegie-Rochester Conference Series on Public Policy*, **5**, 7–29.

McConnell, M. M. and Perez-Quiros, G. (2000), "Output Fluctuations in the United States: What Has Changed since the Early 1980's?", *American Economic Review*, **90**, 1464–1476.

Ramey, V. A. (1991), "Nonconvex Costs and the Behavior of Inventories", *Journal of Political Economy*, **99**, 306–34.

Smets, F. and Wouters, R. (2007), "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach", *The American Economic Review*, **97**, 586–606.

Solon, G., Barsky, R., and Parker, J. A. (1994), "Measuring the Cyclicality of Real Wages: How Important Is Composition Bias?", *The Quarterly Journal of Economics*, **109**, 1–25.

Stock, J. H. and Watson, M. W. (2003), "Has the business cycle changed?", *Proceedings*, pages 9–56.