The Relative Importance of Aggregate and Sectoral Shocks and the Changing Nature of Economic Fluctuations*

Julio Garín†
University of Georgia

Michael Pries‡
University of Notre Dame

Eric Sims§
University of Notre Dame & NBER

This Version: March 4, 2017

Abstract

The reduction in aggregate volatility in the U.S. that occurred in the mid 1980s was accompanied by a shift in the relative importance of aggregate and sectoral shocks. Using a principal components decomposition of sectoral IP data, we document that the contribution of aggregate shocks to the variance of aggregate output declined from about 70 percent in the period 1967-1983 to about 30 percent in the period 1984-2014. We develop an "islands" model with two sectors and costly labor reallocation to investigate how this change in the relative importance of shocks alters standard business cycle moments. Calibrating the aggregate and sectoral productivity shocks to match the overall volatility of output and the relative importance of the shocks found in the pre- and post-1984 data, we find that the post-1984 version of the model with relatively more important sectoral shocks results in a sizeable decline in the cyclicality of average labor productivity and is consistent with changes in several other business cycle moments observed in the post-1984 data.

JEL Classification: E24, E32.
Keywords: Reallocation; Productivity; Business Cycles.

*Corresponding author: Eric Sims. A closely related paper previously circulated under the title "Reallocation and the Changing Nature of Economic Fluctuations." We are thankful for helpful comments and suggestions from Richard Rogerson, Nelson Mark, Robert Lester, Robert Flood, Daniel Cooper, Regis Barnichon, Steve Davis, Michael Reiter, several anonymous referees, and seminar participants at the Univ. of Notre Dame, the Univ. of Houston, DePaul Univ., the Minneapolis Fed, Texas A&M, Columbia University, the Philadelphia Fed, the Atlanta Fed, the 2011 SED Annual Meeting, the 2011 Cologne Workshop on Macroeconomics, the 2011 NBER Summer Institute, the Univ. of British Columbia, and the Fall 2012 Midwest Macroeconomic Meetings.

†E-mail address: jgarin@uga.edu.
‡E-mail address: mpries@nd.edu.
§E-mail address: esims1@nd.edu.

1 Introduction

What are the shocks that drive economic fluctuations? The answer to this question has important implications for policymakers. For example, the appropriate policy response to a downturn caused by a disruption of credit would clearly differ from the optimal policy response to an increase in unemployment following a reallocative shock, such as a negative shock to a particular region or industry.

In this paper, we explore the idea that the relative importance of the different of shocks that drive fluctuations in the U.S. has changed over time. In particular, we focus on two broad categories of shocks: "aggregate shocks" that affect all sectors and regions of the economy in the same way, and "reallocative shocks" that affect some sectors or regions, but not others. Using principal components analysis of U.S. industrial production (IP) data, we find that in the last three decades (the period from 1984 to 2014), the volatility of aggregate shocks has declined considerably (relative to the earlier period from 1967 to 1983), while the volatility of reallocative shocks has remained relatively unchanged. These findings are broadly consistent with related work by Foerster, Sarte, and Watson (2011).

We then use an islands model of labor reallocation to examine whether this changing shock structure can account for several significant changes that have occurred with respect to the nature of economic fluctuations in the United States. Contrary to the famous assertion by Lucas (1977) that "business cycles are all alike," the behavior of the business cycle in the United States since the early 1980s looks quite different than the earlier post-WWII period to which Lucas was referring. One well-known change has been the marked decline in the volatility of output and other economic aggregates. This broad-based decline in volatility has been dubbed the "Great Moderation" and has been the subject of a substantial body of research. Average labor productivity, which was once robustly procyclical, is now slightly countercyclical. We also document changes in the skewness of output growth and in the time series of the labor and efficiency wedges as defined by Chari, Kehoe, and McGrattan (2007). Another well-researched change is the so-called "Jobless Recovery" phenomenon—job growth following the three most recent US recessions has been considerably more anemic, relative to output growth, when compared to earlier recessions. We investigate the extent to which a declining volatility of aggregate shocks relative to sector-specific shocks can account for these changes.

In our principal components analysis, for twelve sectors of the IP data (which comprise the vast majority of the aggregate) we extract the first principal component from the sectoral output growth rates to capture movements that are common across sectors. We interpret movements in this first principal component as fluctuations that stem from aggregate, or common, shocks. The movements in the residual components that are orthogonal to the first principal component are associated with sectoral shocks. A decomposition of the variance of the growth rate of aggregate IP shows that in the 1967-1983 sub-sample, the contribution of aggregate shocks (i.e. the first principal component) accounted for 70.9 percent of the variance, whereas in the 1984-2014 sub-sample the contribution of aggregate shocks accounted for only 31.5 percent of the variance. Moreover, the decline in the contribution of the aggregate shocks by itself fully accounts for the overall decline in IP volatility between the two periods. That is, the volatility of aggregate shocks declined, while the volatility of sectoral shocks remained roughly the same.

To explore the implications of this change in the relative importance of aggregate and sectoral shocks, we construct a model with two different sectors, or "islands," of production. Output is produced on each island using capital and labor, with islands subjected to both island-specific and aggregate productivity shocks. Households consume a composite good made by combining the outputs of each island, and supply labor indivisibly, but with employment lotteries as in Rogerson (1988). Because we focus on labor reallocation, we assume that capital is costless to transfer between islands, but labor is not. While workers can costlessly transition between work and non-work within an island, we assume that when workers move between islands they are unable to work for a (stochastic) period of time, as they acquire new skills or relocate geographically.

The responses of the economy to aggregate shocks and to island-specific shocks differ starkly along several dimensions. In response to a negative *aggregate* productivity shock (one that affects both islands), the model behaves very much like a one-sector real business cycle model. With a low marginal product of labor, there is a reduction in employment (since labor is indivisible) on both islands, and hence a reduction in aggregate employment. Labor productivity—both island-specific and aggregate—falls. As soon as aggregate conditions improve, workers on each island quickly transition back to work, and the observed aggregate employment recovery is relatively fast. In contrast to an aggregate productivity shock, an island-specific productivity shock triggers a change in the *relative* productivities of the two islands and precipitates a movement of workers between the islands. Aggregate employment and output decline following a reallocative shock of this sort, while aggregate labor productivity rises, since employment falls immediately on the adversely affected island and workers then gradually transition to the island with improved productivity. Thus, labor productivity is countercyclical. Furthermore, the more time-consuming is the process of reallocating workers across islands, the more prolonged is the response of aggregate employment following a reallocative shock.

The two shocks also have significantly different implications for the relative importance of the "efficiency wedge" (essentially, the Solow residual measured from the model's simulated aggregate data) and the "labor wedge" (essentially, the wedge in the static first-order condition for labor supply from a one-sector neoclassical growth model measured from data simulated from the model) that Chari, Kehoe, and McGrattan (2007) propose as a tool for identifying the sources of fluctuations. In particular, aggregate shocks in the model translate into movements in the efficiency wedge, with no movements in the labor wedge, whereas reallocative shocks generate fluctuations in the labor wedge, since less than the optimal amount of labor is employed during the process of reallocation that follows a shock to the relative productivities of the islands (reallocative shocks also generate some movement in the efficiency wedge).

The degree of symmetry in output fluctuations also differs across the two shocks. Aggregate shocks generate relatively symmetric fluctuations, since output spikes up in response to a positive shock and jumps down in response to a negative shock. In contrast, reallocative shocks always precipitate a reduction in employment, and thus in output, with no corresponding upward spikes. Thus, output growth rates are negatively skewed when fluctuations result from reallocative shocks, but have no significant skewness when fluctuations are the result of aggregate shocks.

While the two shocks by themselves have starkly different implications for the nature of fluctuations, in reality observed fluctuations are the result of a combination of both reallocative shocks and aggregate shocks. Clearly then, the relative importance of the two shocks will determine whether fluctuations have features that more closely resemble an economy with only aggregate shocks, or an economy with only reallocative shocks. Given the reduced importance of aggregate shocks that we identify with our analysis of the IP data, our model suggests that fluctuations in the post-1984 period should have features that more closely resemble the economy with reallocative shocks. We show this to be true in the U.S. data. First, labor productivity went from strongly procyclical in the pre-1984 period (with a correlation between output and labor productivity of 0.64) to slightly countercyclical in the post-1984 period (with a correlation of -0.08). Second, the volatility of the measured efficiency wedge declines (the standard deviation falls from 0.017 to 0.009), while the volatility of the labor wedge is basically unchanged (the standard deviation goes from 0.014 to 0.015). Third, output growth rates have become more negatively skewed, with skewness declining from -0.025 to -1.114. Finally, following recessions in the post-1984 period, total hours worked has recovered slowly when compared to output—which is related to the "jobless recovery" phenomenon.

To assess how well our model can quantitatively account for these observed changes in business cycle moments, we run simulations of two versions of the model—one intended to represent the 1967-1983 sub-sample, and another intended to capture the 1984-2014 sub-sample. The only model parameters that we allow to differ between the two calibrations are the standard deviations of the aggregate and reallocative shock processes. Specifically, for both the 1967-1983 and the 1984-2014 versions of the model, we calibrate the volatility of the aggregate and island-specific shocks in the model to match, for each period, two moments: (1) the variance of the growth rate of the aggregate U.S. IP data, and (2) the fraction of the volatility of output growth rates that is accounted for by the aggregate shock, for which we use the estimates from our principal components analysis. We then examine the extent to which the reduced importance of aggregate shocks, by itself, can account for the changing nature of U.S. business cycles. While it may seem odd to consider a version of the model with smaller aggregate shocks given the recent large recession, it is important to note that a smaller standard deviation does not preclude a particularly large negative shock realization. Indeed, our principal components analysis identifies an unusually large negative aggregate shock hitting the economy in September of 2008, the dating of which coincides with the conventional dating of the height of the financial crisis.[1] Our analysis also finds negative reallocation shocks played a dampening role in the economy in the early part of 2009.

When calibrating the model with a smaller aggregate productivity shock, we find that the volatilities of output and the efficiency wedge decline by about one-half, just as they do in the data. That the volatility of output is lower in the model with smaller aggregate productivity shocks is essentially by construction given our calibration approach. A stronger test of the quantitative relevance of the model therefore focuses on other moments. Our model does well on these dimensions. Labor productivity goes from strongly procyclical in the 1967-1983 calibration to mildly countercyclical in the 1984-2014 calibration. The efficiency wedge remains strongly procyclical but is less positively correlated with output than in the calibration with a more important aggregate shock. The labor wedge becomes more countercyclical in the 1984-2014 calibration and its volatility is roughly unchanged. Output growth becomes significantly more left-skewed when calibrating the model with a smaller aggregate productivity shock. All of these features are broadly in-line with changes observed in the data. Our model also generates slower employment recoveries (relative to output) after periods identified as recessions in model simulations. In particular, the half-life of employment relative to output increases in the calibration of the model with a smaller aggregate productivity shock in a way qualitatively in-line with what is observed in data.

Our paper fits into several different literatures. Lilien (1982), Abraham and Katz (1986), and Davis (1987) are early papers that study the role of labor reallocation in the business cycle. Hornstein (2012), Wiczer (2013), Sahin, Song, Topa, and Violante (2012), Fujita and

[1]Further to this point, Gadea, Gomez-Loscos, and Perez-Quiros (2014) have argued that the Great Recession does not signal an end to the period of low output volatility that began in the early 1980s.

Moscarini (2013), and Mehrotra and Sergeyev (2012) are more recent works that explore the role of "mismatch," structural factors, and sectoral shifts in accounting for employment dynamics. Hall (2007), Gali and Gambetti (2009), Barnichon (2010), Hagedorn and Manovskii (2011), and Gali and van Rens (2014) point out that the cyclicality of labor productivity has changed in an important and stark way since the early 1980s. Bachmann (2009) studies the so-called jobless recovery phenomenon. Berger (2015) seeks to provide a joint explanation for the declining cyclicality of productivity and the jobless recovery phenomenon.

The remainder of this paper is organized as follows. Section 2 applies principal components techniques to the U.S. sectoral Industrial Production data to assess the relative importance of aggregate and sectoral shocks for two different periods: 1967-1983 and 1984-2014. Section 3 lays out the islands model and describes the model's equilibrium and its characteristics. Section 4 carries out quantitative exercises to show that the model, when calibrated to account for the changing relative importance of aggregate and sectoral shocks, can account for several observed changes in business cycle moments. The final section concludes by briefly discussing the broader implications of our findings.

2 The Relative Importance of Aggregate and Sector-Specific Shocks

The most well-known change in the business cycle since the early 1980s is the so-called Great Moderation, first documented in Kim and Nelson (1999) and McConnell and Perez-Quiros (2000). Stock and Watson (2003) identify 1984 as the break point for the decline in the volatility of output; we take this breakpoint as given for the remainder of the paper. In this section, we employ principal components analysis on sectoral IP data to investigate the extent to which the decline in aggregate volatility is driven by a decline in common variation across sectors or a decline in sector-specific sources of variation. Our analysis forms the basis for some of our calibration targets for the quantitative exercises in the model laid out in the next section.

We focus on twelve sectors that make up the bulk of the aggregate IP index.[2] These sectors include the second-level sub-categories under "Final products and non-industrial supplies," the second-level sub-categories under "Materials," and the total manufacturing

[2]It is possible to go to a finer level of disaggregation, but it is problematic in that not all of the subsectors that make up the twelve sectors on which we focus have data going back as far as 1967. Since our ultimate objective is to compare business cycle moments both before and after the conventional dating of the Great Moderation, it is important for our data to extend back as far as possible. The sample could be extended further back in time by focusing on less disaggregated data, but this would leave only a few sectors.

sector.[3] We have these data at a monthly frequency going back to 1967 and our sample ends in July of 2014. To check that these twelve sectors are representative of the aggregate, we form a synthetic aggregate IP series by aggregating the growth rates of sectoral IP using employment shares as weights. The correlation of the growth rate of the synthetic series with the growth rate of actual aggregate IP is 0.99.

To assess the relative contribution of aggregate and sectoral shocks, we extract principal components from the growth rates of sectoral IP and examine what fraction of the variance of aggregate IP growth these principal components can explain. To facilitate comparison with our model in the next section, which only features one aggregate shock, we focus on only the first principal component of the sectoral IP growth rates.

To be more specific, the aggregate growth rate of IP, $\Delta \ln IP_t$, can be expressed as approximately equal to the share-weighted sum of sectoral growth rates:

$$\Delta \ln IP_t = \sum_{i=1}^{H} \omega_{i,t} \Delta \ln IP_{i,t}. \tag{1}$$

Where $\Delta \ln IP_{i,t}$ denotes the growth rate of sector i's IP. $i = 1, \ldots, H$ indexes sectors and ω_i are employment weights by sector. Let X_t denote a vector of sectoral IP growth rates:

$$X_t = \left[\Delta \ln IP_{1,t} \;\; \ldots \;\; \Delta \ln IP_{H,t}\right], \tag{2}$$

and let V denote the variance-covariance matrix of X_t. This variance-covariance matrix can be decomposed as:

$$V = \Gamma \Lambda \Gamma'. \tag{3}$$

Here Λ is a diagonal matrix of the eigenvalues of V, where the eigenvalues are sorted in descending order according to modulus. Γ is matrix of eigenvectors with columns sorted in accord with the ordering of eigenvalues in Λ. The first principal component of the sectoral IP growth rates is then defined as

$$F_t = X_t \Gamma_1, \tag{4}$$

where Γ_1 is the first eigenvector.

In other words, the first principal component is the linear combination of sectoral IP growth rates that maximally explains the variance-covariance matrix of sectoral IP growth rates. We then decompose the variance of the growth rate of the aggregated IP into the

[3] The series, and their ID numbers, are B51100 (Durable consumer goods), B51200 (Nondurable consumer goods), B52110 (Transit equipment), B52120 (Information processing and related equipment), B52130 (Industrial and other equipment), B52300 (Defense and space equipment), B54100 (Construction supplies), B54200 (Business supplies), B53100 (Durable goods materials), B53200 (Nondurable goods materials), B53300 (Energy materials), and B00004 (Manufacturing (SIC)).

variance attributable to that principal component and the variance attributable to the residual movements of the sectoral IP growth rates (the two components sum to the full variance, since by construction the residuals are orthogonal to the principal component). Table 1 shows the results. Using the full sample period (1967-2014), we find that the first principal component of the twelve sectoral IP growth rates accounts for 42.4 percent of the variance of aggregate IP growth. We also carry out the decomposition for the pre-1984 and post-1984 subsamples. To do this, we compute the principal components separately for each sub-sample. In the pre-1984 sample, the first principal component accounts for 70.9 percent of the total variance, whereas in the post-1984 sample it accounts for only 31.5 percent.

Table 1: Contribution of Aggregate Shock to Variance of Aggregate IP Growth

Sample Period	Var(IP Growth)	Due to 1st Component	Residual Variance
1967-2014	0.0055	0.0023 (42.4)	0.0032 (57.6)
1967-1983	0.0086	0.0061 (70.9)	0.0025 (29.1)
1984-2014	0.0038	0.0012 (31.5)	0.0026 (68.5)

Notes: This table shows the total variance (times 100) of aggregate IP growth in three different samples, as well as the variance of aggregate IP growth accounted for by the first principal component (times 100) and the residual variance (also times 100). Numbers in parentheses denote shares of the aggregate variance.

We interpret the portion attributable to the first principal component as a measure of the contribution of aggregate shocks to the variance of total output, and the remainder as the contribution of sectoral shocks.[4] While the contribution of aggregate shocks has clearly declined from the pre-1984 sample to the post-1984 period, it is also worth noting that the decline in this aggregate component fully accounts for the overall decline in the variance of IP growth rates. That is, the variance of the sectoral component did not decline. If we convert the variances to standard deviations, the pre-1984 standard deviation of the aggregate component was 0.0078 and the post-1984 standard deviation was 0.0035, whereas the pre-1984 standard deviation of the residual component was 0.0050 and the post-1984 standard deviation was 0.0051. In other words, the volatility of the aggregate component declined by about half, and that decline fully accounts for the decline in the overall variance of aggregate IP growth rates.

A declining volatility of aggregate IP growth in conjunction with a decline in the relative importance of the first principal component of the twelve sectoral IP growth rates suggests that much of the Great Moderation can be attributed to a decline in the volatility of aggregate shocks. This is consistent with the analysis in Foerster, Sarte, and Watson (2011). Their

[4]While this interpretation seems natural, it is subject to the caveat that the first principal component of the sectoral IP growth rates may partly pick up purely sectoral shocks which are propagated across sectors via sectoral linkages and complementarities.

paper combines a structural factor model with a multi-sector neoclassical growth model allowing for input-output leakages. They too argue that the Great Moderation is associated with a decline in the importance of aggregate shocks while the volatility of sectoral shocks was largely unchanged. Concretely, they find that sectoral shocks account for about 20 percent of the variance of aggregate IP growth in the pre-1984 sample and about 50 percent post-1984. These numbers are in-line with ours. Because their analysis explicitly accounts for the possibility of sectoral linkages, this suggests that the caveat mentioned in footnote 4 is not driving our results.[5]

3 Model

This section introduces an island model of labor reallocation, in the spirit of Lucas and Prescott (1974), that features both island-specific productivity shocks and aggregate productivity shocks. Frictions impede the reallocation of workers that is desired following island-specific shocks that change the relative productivity of islands.

We begin our discussion of the model by describing the production side of the economy. There are two islands where production takes place. These islands could represent different sectors, regions, industries, or occupations. Though it would be conceptually straightforward to extend the analysis to several islands, for computational tractability we focus on just two. On each island i there is a representative firm that produces an intermediate good using the technology $X_{i,t} = A_t z_{i,t} K_{i,t}^{\alpha} L_{i,t}^{1-\alpha}$, where A_t is an aggregate shock that is common to both islands, $z_{i,t}$ is the productivity shock specific to island i, and $L_{i,t}$ and $K_{i,t}$ are the labor and capital utilized on island i. A_t and $z_{i,t}$ both follow Markov processes. The intermediate goods from the two islands are transformed into a final good by a competitive firm utilizing the following CES technology:

$$Y_t = \left(X_{1,t}^{\frac{\sigma-1}{\sigma}} + X_{2,t}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}. \tag{5}$$

The economy is populated by a continuum of infinitely lived households of measure 1. These households are located in one of three states: on one of the two islands or in an intermediate state between the islands. Household labor can only be supplied on the island in which the household currently resides; moving to the other island requires first passing through the intermediate state for a period of time. The lost productive time in the intermediate state could represent retraining or geographic relocation.

[5]There are a couple of other differences worth mentioning. First, they focus on a more disaggregated level of IP data, including 117 sectors. Second, the beginning of their sample is 1972, whereas ours goes back to 1967 (which is why we focus on less disaggregated data due to data limitations prior to 1972). Third, their sample period ends in 2007, whereas ours includes the Great Recession and its aftermath.

The households face a consumption/saving decision and a labor supply decision, and seek to maximize the present discounted value of flow utility,

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\ln c_t - \nu(l_t) \right], \tag{6}$$

where c_t is the household's consumption, l_t is the household's labor supply, and $\beta < 1$ is the household's discount factor.[6] The disutility of labor, $\nu(l_t)$, is an increasing and convex function in l_t. Finally, a household's labor supply is assumed to be indivisible: $l_t \in \{0, 1\}$.

We assume that there are complete asset markets that allow households to insure perfectly 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.). In addition, we assume that there are employment lotteries as in Rogerson (1988). These assumptions allow us to identify the competitive equilibrium as the outcome of a social planning problem in which the planner has preferences

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\ln C_t - \phi L_t \right], \tag{7}$$

where $\phi = \nu(1) - \nu(0)$. When discussing household preferences we used lowercase c_t and l_t, but here, for the social planner, we can express preferences in terms of aggregates, which are denoted using uppercase variables. This is appropriate because the lotteries and perfect insurance, along with the fact that there is a unit measure of households, mean that household and aggregate consumption are the same. Likewise, because there is a unit measure of households, L_t is both aggregate employment as well as the fraction of households with $l_t = 1$.

Complete asset markets and employment lotteries render household heterogeneity irrelevant: we need not know the identities of the households that are allocated to the two islands, nor the identities of the households that are unemployed. However, while the identities of the households in these different situations do not matter, the overall distribution of households across states does matter. That is, the key decisions for the social planner relate to the distribution of workers over different employment states: (i) what fraction of the workers to allocate to the two islands and (ii), what fraction of the workers on each island to assign to be employed. Of course, in making these decisions the social planner faces the same frictions that individual households face—reallocating workers from one island to the other is time-consuming. 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

[6] Given the assumed separability between consumption and leisure, we assume log utility so that the model is consistent with balanced growth.

first pass through the intermediate state, and a spell of "reallocation unemployment," while engaged in activities that make them employable on the other island. Workers stochastically escape this reallocation spell at the exogenously given rate λ, and we assume that when they escape they can reallocate to either island.[7]

In this environment, workers at a point in time are in one of three situations: (i) located on an island and employed, (ii) located on an island and unemployed, but with the possibility of a frictionless transition back to employment on that island, and (iii) not located on an island, but rather in the state of "reallocation unemployment." To understand the second situation, note that the planner may choose not to employ some workers on an island, while also not reallocating those workers to the other island. For example, holding island-specific productivities constant, if aggregate productivity is temporarily low, then it may be optimal for the planner to reduce employment on each island (due to the indivisibility of labor and the disutility associated with work) without initiating any reallocation.

To maintain tractability, and to focus attention on the role of labor market frictions, we assume that capital can be transferred from one island to another with no frictions. Thus, the social planner enters a period with an aggregate stock of capital, chosen in the previous period, which it can allocate to the two islands after observing the current period's island-specific productivities. Capital depreciates at rate δ.

The timing of events within a period is as follows. Period t begins with $N_{1,t-1}$ and $N_{2,t-1}$ workers located on each island and an aggregate capital stock K_t. After aggregate productivity A_t and the island-specific productivities $z_{1,t}$ and $z_{2,t}$ are revealed, the planner makes several simultaneous decisions. First, the planner decides how many workers to allocate to the two islands in the current period, $N_{1,t}$ and $N_{2,t}$. This decision is constrained by the fact that at most $\lambda(1 - N_{1,t-1} - N_{2,t-1})$ in total can be added to the two islands. Second, the planner decides how many of the $N_{1,t}$ and $N_{2,t}$ workers on each island will be employed, i.e. $L_{1,t}$ and $L_{2,t}$, and how much of the aggregate capital stock, K_t, to utilize on each island, i.e. $K_{1,t}$ and $K_{2,t}$. Finally, given these choices of inputs, the total output of the final good is determined and the planner must decide how to allocate it between consumption C_t and the next period's aggregate capital stock K_{t+1}.

The social planner's problem then is to choose history-contingent sequences for the vector of choice variables $\{L_{1,t}, L_{2,t}, N_{1,t}, N_{2,t}, K_{1,t}, K_{2,t}, K_{t+1}, C_t\}$ in order to maximize:

[7]This assumption simplifies the model by eliminating the need to keep track of which island each worker in the reallocation process originally came from. Workers will move in response to island-specific shocks, and will move to the island where productivity is increasing, 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.

$$\max \ \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\ln C_t - \phi(L_{1,t} + L_{2,t})\right]$$

subject to:

$$C_t + K_{t+1} \leq \left(X_{1,t}^{\frac{\sigma-1}{\sigma}} + X_{2,t}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} + (1-\delta)(K_{1,t} + K_{2,t}) \tag{8}$$

$$X_{i,t} = A_t z_{i,t} K_{i,t}^{\alpha} L_{i,t}^{1-\alpha}, \quad i = \{1,2\} \tag{9}$$

$$K_t = K_{1,t} + K_{2,t} \tag{10}$$

$$L_{1,t} \leq N_{1,t} \tag{11}$$

$$L_{2,t} \leq N_{2,t} \tag{12}$$

$$N_{1,t} \leq N_{1,t-1} + \lambda(1 - N_{1,t-1} - N_{2,t-1}) \tag{13}$$

$$N_{2,t} \leq N_{2,t-1} + \lambda(1 - N_{1,t-1} - N_{2,t-1}) \tag{14}$$

$$N_{1,t} + N_{2,t} \leq N_{1,t-1} + N_{2,t-1} + \lambda(1 - N_{1,t-1} - N_{2,t-1}) \tag{15}$$

$$N_{1,0},\ N_{2,0}, \text{and } K_0 \text{ given.} \tag{16}$$

The first constraint is the aggregate resource constraint. The next constraint imposes the production technologies for the two intermediate goods. The third constraint states that the sum of the capital on the two islands in period t must be equal to the period t aggregate capital stock (which was chosen in period $t-1$). The fourth and fifth constraints require that employment on each island not exceed the number of workers allocated to that island. The sixth and seventh constraints state that the total number (measure) of workers on an island must be less than the sum of the number of workers already there in the previous period and the workers who successfully exited the reallocation process in the previous period. The eighth constraint imposes that the workers available to be reallocated can only be reallocated to one island or the other. Finally, the initial allocation of workers is, like the initial capital stock K_0, exogenously given and must satisfy $N_{1,0} + N_{2,0} \leq 1$.

It is straightforward to express this social planning problem as a 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, the aggregate capital stock K, and the values of the aggregate and idiosyncratic shocks. To simplify notation, let $\xi_t = \{A_t, z_{1,t}, z_{2,t}\}$ denote the vector of exogenous shocks. We can express the problem recursively as a Bellman equation:

$$V(N_1, N_2, K, \xi) = \max_{L_1, L_2, N_1', N_2', K_1, K_2, K', C} \ln C - \phi(L_1 + L_2) + \beta \mathbb{E} V(N_1', N_2', K', \xi')$$
$$\text{s.t. (8)–(16).}$$

We solve this problem numerically using standard techniques. Specifically, we create a grid of values for N_1, N_2, and K and then iterate on the Bellman equation above until it converges. Evaluating the value function at points between the gridpoints for N_1, N_2, and K requires interpolation; we use a simplicial 3-D interpolation (see Judd 1998, p. 242).

Before turning to a quantitative analysis of the model, it is worthwhile to discuss its qualitative features. Figure 1 graphically depicts the basic structure of the model economy. The two types of shocks will generate different responses of aggregate employment, productivity, and output. Consider a negative aggregate shock, holding the relative productivities of the two islands constant. The indivisible labor and employment lotteries assumptions mean that employment on both islands will decline while aggregate productivity remains low. In terms of Figure 1, some workers move to the rectangles on the sides—they temporarily move out of employment, but remain on the same island. However, when aggregate productivity recovers, these unemployed workers quickly return to employment on the same island. As a result, employment, output, and productivity all decline and then quickly recover together. In essence, in the absence of reallocative shocks the model becomes a two-sector version of the Hansen (1985) general equilibrium model of indivisible labor with employment lotteries. As such, changes in A_t will map into a standard Solow residual or "efficiency wedge" in the terminology of Chari, Kehoe, and McGrattan (2007).

Figure 1: Graphical depiction of the "island" model

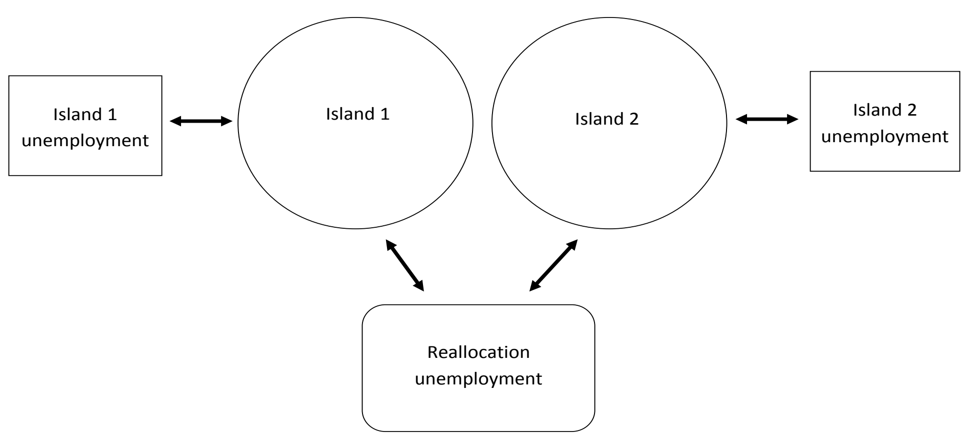

Now consider a purely reallocative shock that increases the productivity of one island relative to the other, while leaving aggregate productivity unaffected. This type of shock will generate what we call "reallocative unemployment." That is, the desire to reallocate workers from the island with reduced productivity to the island with increased productivity means that workers must pass through a time-consuming process of reallocation, represented in the figure by the cell on the bottom. To the extent that workers must spend a significant period of time (in expectation) in the reallocation spell, employment might recover much more slowly following a reallocative shock. The response of average labor productivity to a reallocation shock will also look different than the response to an aggregate shock. Labor can be quickly reduced on the adversely affected island, thus driving back up the marginal product of labor on that island. Because labor only slowly moves to the island with increased productivity, the marginal product there also remains high. As a result, average productivity will temporarily increase in response to the reallocative shock. Therefore, conditional on a reallocative shock, aggregate productivity will be less strongly correlated with aggregate output and employment than in the case of an aggregate shock.

While the aggregate productivity shock will generate what looks like an efficiency wedge, in our model a reallocative shock will generate a labor wedge, which is an important feature of US business cycle data. Chari, Kehoe, and McGrattan (2007) define the labor wedge in the data as the residual from the static first-order condition that equates the marginal rate of substitution between labor and consumption with the marginal product of labor that would obtain in a standard RBC model. In our model, the first-order conditions of the planner's problem for the choice of labor on each island are:

$$u'(C)\frac{\partial Y}{\partial L_1} = \phi + \mu_1 \tag{17}$$

$$u'(C)\frac{\partial Y}{\partial L_2} = \phi + \mu_2. \tag{18}$$

Here μ_1 and μ_2 are the Lagrange multipliers on the constraints that employment on islands 1 and 2 not exceed the population on that island. If there were no reallocation friction, these constraints would never bind—hence $\mu_1 = \mu_2 = 0$ would hold at all times and the marginal products of labor on each island would be equalized. Because of the assumption of Cobb-Douglas production, the marginal and average products are proportional to one another. Equal average products on each island in turn imply equality to the aggregate average product of labor. This would mean that there would exist an aggregate representation of the conventional static first-order condition for labor supply in which the marginal rate of substitution between labor and consumption, i.e. $\phi/u'(C)$, would always be equal to the aggregate marginal product of labor—there would be no labor wedge. If there is a desire to reallocate, caused by a change in the z_i, however, these constraints may bind, which would lead to the marginal products of labor on each island being different from one another and leading to an observed residual from the static first-order condition of a neoclassical growth model. That is, there will be a measured labor wedge whenever workers are not optimally allocated across the two islands, given the relative productivity of the two islands. Thus, the labor wedge should be more volatile when reallocation is more difficult (λ is smaller) or when the desired amount of reallocation is greater (σ is larger). The quantitative exercises below confirm this intuition.

4 Empirical Facts and Quantitative Results

In this section we document several empirical facts about how key business cycle moments have changed in the post-1984 period relative to the earlier pre-war sample. We then assess the extent to which our model can quantitatively replicate these facts.

4.1 Empirical Facts

The most well-known change in the business cycle is the decline in output volatility, or the so-called "Great Moderation." While the Great Moderation has been the subject of a great deal of research, there are several other changes in business cycle moments which are equally, if not more, stark than the decline in output volatility. For example, labor productivity switched from strongly procyclical to mildly countercyclical around the same time as the decline in the volatility in output.[8] Here we document this and several other facts about changes in the nature of economic fluctuations.

We focus on the following moments: the standard deviations of output, the efficiency wedge, and the labor wedge; and the correlations between output and average labor productivity, between the efficiency wedge and aggregate output, and between the labor wedge and aggregate output. In the data, output is defined as real GDP from the NIPA accounts and total hours worked is hours per capita in the non-farm business sector. In both the model and data we measure the efficiency and labor wedges as they are defined in Chari, Kehoe, and McGrattan (2007). That is, we measure the efficiency wedge as an aggregate Solow residual, assuming a Cobb-Douglas aggregate production function.[9] The labor wedge is defined as the

[8] Hall (2007), Gali and Gambetti (2009), Barnichon (2010), Hagedorn and Manovskii (2011), and Gali and van Rens (2014) are among the first authors to have documented the sharp switch in the cyclicality of labor productivity.

[9] In particular, we posit the existence of an aggregate production function of the form: $Y_t = \theta_{e,t} K_t^{\alpha} L_t^{1-\alpha}$,

residual from the standard static first-order condition for labor supply that would emerge as part of the planner's solution to a neoclassical growth model with indivisible labor.[10] Labor productivity is measured as output per hour in the non-farm business sector.

The two columns in Table 3 (shown below alongside the model's simulation results) under the heading "data" present moments for the 1967-1983 "early" period as well as the 1984-2014 "later" period.[11] In terms of volatilities, we see that the volatility of output is about half as large in the later period as in the early period (i.e. the Great Moderation). Likewise, the efficiency wedge is also about half as volatile in the later sample period. Interestingly there is essentially no difference in the volatility of the labor wedge for the two periods. In terms of correlations, in the early period output is strongly positively correlated with labor productivity (correlation of 0.64). The efficiency wedge is also strongly procyclical (correlation with output of 0.86), whereas the labor wedge is countercyclical (correlation with output of -0.53). When moving from the early period to the later period, labor productivity switches from strongly procyclical to mildly countercyclical (correlation with output of -0.08 in the later sample). There is a modest reduction in the cyclicality of the efficiency wedge, but it remains procyclical and the decline in correlation between it and output is not nearly as large as the decline in the correlation between labor productivity and output. The labor wedge becomes more strongly countercyclical in the later sample (correlation with output of -0.74). There is also a sharp change in the skewness of output growth—it goes from negligible skewness to significantly left-skewed in the post-1984 sample.

4.2 Calibration

Before studying the quantitative properties of our model, in this subsection we discuss our calibration of the model's parameters. These parameters include: β, the discount factor; ϕ, the scaling parameter on the disutility of labor; σ, the elasticity of substitution among

and measure the efficiency wedge, $\theta_{e,t}$, as $\ln \theta_{e,t} = \ln Y_t - \alpha \ln K_t - (1-\alpha)\ln L_t$. In the data, we use the total factor productivity series produced in Fernald (2014) (the measure which is not adjusted for capacity utilization). To compute the efficiency wedge on model generated data, we assume that $\alpha = 1/3$ and measure aggregate output, capital, and labor in the model.

[10] The labor wedge is defined assuming log utility over consumption and linear utility from labor, with the same aggregate production function discussed above. In particular, we assume that the static FOC in the planner's problem is given by $\phi = (1-\theta_{l,t})\frac{1}{C_t}(1-\alpha)\frac{Y_t}{L_t}$. $\theta_{l,t}$ measures the labor wedge and is isomorphic to a distortionary tax on labor income. In the data we measure the labor wedge by taking logs of this static FOC; ignoring constants, one gets $\theta_{l,t} \approx -\ln\left(\frac{C_t}{Y_t}\right) - \ln L_t$. We measure the consumption-output ratio in the data by taking the ratio of nominal non-durable and services consumption to total nominal output, and measure L_t as total hours worked in the non-farm business sector. We construct the model's labor wedge in the same way using the corresponding data concepts in the model.

[11] Quarterly observations on these series are available going back to 1947. We instead fix the start date of the sample to 1967, as this corresponds with the first observations in the sectoral IP data that we use to calibrate the shocks.

intermediate inputs from the two islands; δ, the depreciation rate on capital; λ, the parameter that determines the pace of reallocation between the two islands; and the parameters governing the stochastic processes of the aggregate and island-specific productivity shocks.

Table 2: Parameter Values

Parameter	Value	Description
β	0.99	Discount factor
ϕ	$L^* = 0.975$	Disutility of labor
δ	0.025	Depreciation rate
α	0.33	Capital elasticity of intermediate good production
σ	3	Elasticity of substitution
ρ_A	0.92	AR(1) aggregate productivity
ρ_Z	0.95	AR(1) island-specific productivity
1967-1983 sample		
s_A	0.0072	Standard deviation of innovation, aggregate shock
s_Z	0.0174	Standard deviation of innovation, island-specific shock
1984-2014 sample		
s_A	0.0026	Standard deviation of innovation, aggregate shock
s_Z	0.0161	Standard deviation of innovation, island-specific shock

Notes: This table presents the values assigned to model parameters in our baseline calibration strategy. The motivations for these values are described in the text.

Table 2 lists the parameter values assumed in our baseline calibration. The unit of time in the model is taken to be a quarter. Accordingly, we set $\beta = 0.99$. The depreciation rate is set to $\delta = 0.025$. The parameter ϕ controls the disutility of work, and thus determines average non-employment (note that the workers who are reallocating between islands can be considered "unemployed," as they wish to work but are not working, whereas the workers who remain on the island but do not work would more appropriately be labeled as non-participants). Given the difficulty of finding an appropriate empirical counterpart for the average non-employment of the model's stochastic equilibrium (because it is a mix of workers who are experiencing reallocation and other workers who are choosing not to work, and because it does not include unemployment that stems from the conventional search frictions associated with transitions between firms/jobs in the same sector), we set ϕ so that employment in the non-stochastic version[12] of the model, L^*, is equal to 0.975. To understand this choice, note that it is important that ϕ is low enough so that L^* is not too low; otherwise the equilibrium would feature "reserves" of non-employed workers on each island who can quickly transition into and out of work on an island in response to island-specific shocks, thereby making reallocation unnecessary. At the same time, ϕ must be high enough so that there will be some non-employment ($L^* < 1$); otherwise, workers on an island would never be idle and

[12] In the non-stochastic version of the model, there is no reallocation and labor is allocated evenly across the two islands and thus one can solve for ϕ, given L^*, using the static first-order condition for labor supply.

employment would not respond to aggregate shocks. Thus, the key consideration in choosing L^* is that it must be such that employment in the model will respond to both aggregate and reallocative shocks. The chosen value, 0.975, achieves that, and the results are not very sensitive to small changes in this value (given that other parameters, such as the standard deviations of the shocks, are also re-calibrated when the target L^* is altered), though clearly a sufficiently large change in L^* would make the model unresponsive to one of the shocks.

The parameter σ governs the degree of substitutability between the outputs of each island, and thus governs the extent to which it is desirable to reallocate resources across islands in response to changes in relative productivities. A value of $\sigma = \infty$ indicates perfect substitutes; in this case it would always be optimal, in the absence of reallocation frictions, to shift all inputs to the more productive island. A value of $\sigma = 0$, on the other hand, indicates perfect complements; in this case, it would be optimal to reallocate resources toward the less productive island, so as to equalize the production on each island. Finally, a value of $\sigma = 1$ corresponds to Cobb-Douglas; in this case, it would be optimal to have equal inputs on both islands regardless of island-specific productivities. In order to get reallocation of workers toward the more productive island, we need $\sigma > 1$. We set $\sigma = 3$ based on evidence in Broda and Weinstein (2006), who provide estimates of this parameter from SITC data for the US. This value is similar to parameterizations used in other quantitative models featuring Lucas and Prescott (1974) islands.[13] In the Appendix, we provide sensitivity analysis that shows that our quantitative results are fairly robust to smaller or larger values of σ.

The rate at which workers stochastically escape reallocation and become employable, λ, will be lower the greater are the frictions that make it difficult for workers to switch islands. Numerous factors can impede workers in this way, such as the need to acquire new skills or the need to re-locate to a different region. Empirical work—such as Ruhm (1991) and Jacobson, LaLonde, and Sullivan (1993)—that looks at the impact on earnings of a job displacement, which is probably in many ways similar to the experience of a worker who is forced to switch islands in our model, indicates that the impact is felt for many years. In particular, Jacobson et al. (1993) find that earnings two years after a displacement are roughly 50 percent below pre-displacement earnings. We choose λ to match this number in the model for average earnings losses among workers who switch from employment on one island to reallocation unemployment. In our model this implies a value of $\lambda = 0.1$. While this may seem low in light of the average duration of unemployment spells in the data, it is important to keep in mind that unemployment resulting from the need to switch sectors is rather different than the "average" spell of unemployment, which includes frictional and seasonal unemployment,

[13] See Alvarez and Shimer (2011) for a further discussion of some of this literature.

and which is typically very short in duration.[14] Below we also examine the robustness of our results to different values of this parameter.

We next turn to a discussion of our parameterization of the stochastic processes for aggregate productivity, A_t, and island-specific productivity, z_t. To be consistent with the principal components analysis above, which we use here to calibrate these processes, we assume that z_t and A_t are independent of one another (by construction, the first principal component, which captures the aggregate shock, is orthogonal to sector-specific shocks). The Markov process for aggregate productivity A_t is parameterized to approximate, via the Tauchen (1986) procedure, a stationary AR(1) process with mean normalized to unity:

$$A_t = (1 - \rho_A) + \rho_A A_{t-1} + \varepsilon_{A,t}. \tag{19}$$

The parameter ρ_A is restricted to lie between 0 and 1, and the innovation has variance equal to s_A^2.

The island-specific productivities also follow Markov processes. Because we have just two islands, and what matters is the *relative* productivity of the two islands, we do not need to consider two completely independent processes, but rather can assume that they are perfectly negatively correlated (this comes without any loss of generality but has the benefit of reducing the size of the state space). Specifically, if $z_{i,j}$ denotes the productivity of island i in state j, then when $z_{1,j}$ differs from the mean of z by an amount x, the value of $z_{2,j}$ differs from the mean of z by $-x$. Given this structure, we specify the Markov process for one island (and thus, the other as well) as an approximation to a stationary AR(1) process with mean normalized to unity:

$$z_t = (1 - \rho_z) + \rho_z z_{t-1} + \varepsilon_{z,t}. \tag{20}$$

The parameter ρ_z is restricted to lie between 0 and 1, and the innovation has variance equal to s_z^2.

We set $\rho_A = 0.92$ and $\rho_z = 0.95$. The former is a fairly common value in quantitative real business cycle models. For the latter, it is important that island-specific shocks are sufficiently persistent so as to trigger reallocation of labor across islands. Below we report robustness exercises that examine how sensitive the results are to different values of ρ_z. We parameterize the shock magnitudes, s_A and s_z, as follows. We target two moments discussed in Section 2—the volatility of aggregate IP growth rates (higher values of both s_A and s_z contribute to higher volatility) and the fraction of the variance of aggregate IP growth rates that is

[14] Clark and Summers (1979) show that a large fraction of unemployment spells end very quickly, which significantly reduces mean and median unemployment durations. They also argue that longer term unemployment, as well as separations followed by exit from the labor force, account for a significant fraction of unemployment.

accounted for by the aggregate shock (this pins down the *relative* values of s_A and s_z). We use this approach to calibrate the values separately for the pre- and post-1984 periods. That is, for the calibration of the pre-1984 period, we target the observed volatility of IP growth in the early sample, i.e. 0.0093, and the measured contribution of the aggregate shock (70.9 percent). For the calibration of the later period, we target a volatility of IP growth of 0.0062 and a contribution of aggregate shocks of 31.5 percent. This approach yields $s_A = 0.0071$ and $s_z = 0.0169$ for the early sample, and $s_A = 0.0027$ and $s_z = 0.0159$ in the later sample. In this calibration the magnitudes of both aggregate and island-specific shocks decline when moving from the 1967-1983 sample to the 1984-2014 sample, though the decline is much larger for the aggregate shock. Put differently, in our calibration strategy the *relative* importance of island-specific shocks increases in the later sample.

4.3 Results

Our objective in this section is to quantitatively examine the implications of a declining relative importance of aggregate shocks in our island model with costly labor reallocation. We consider two calibrations of shock volatilities—one meant to match the volatility and relative importance of aggregate shocks in the 1967-1983 sample, and the other meant to match the volatility and relative importance of aggregate shocks in the 1984-2014 sample. Parameters unrelated to the shock processes are held fixed at their values described in the subsection above.

For both the pre-1984 and post-1984 parameterizations, we solve the model and then simulate 15,000 different samples, with the number of observations in each sample equal to the number of quarterly observations in the sample period in question (i.e. either 1967-1983 or 1984-2014). We calculate each moment of interest for each sample and report the average moment across the 15,000 samples. To reduce sensitivity to starting values, we use a 400-period "burn-in" when creating each simulated sample. For both the model data and U.S. data, we log and HP-filter the sample before calculating moments.

Moments from our quantitative simulations are presented in Table 3 under the heading "Model." The first thing to note is that the volatilities of output, the efficiency wedge, and the labor wedge in the model are lower, in both periods, than in the data. To understand why this is the case, note that the volatilities of the two shock processes were calibrated to match the volatility of the *growth rates* of the Industrial Production data, rather than the volatility of the HP filtered log *level* of GDP (which is what is reported for the data in Table 3), since the principal components analysis is carried out using growth rates. The fact that we do not match the *levels* of the volatilities is not so important, as our focus is more specifically on accounting for the *change* in the volatilities when moving from the early to the later sample. In this regard, the model performs well, as the volatility of output declines by almost half, as in the data. Similarly, the volatility of the efficiency wedge declines by about one-half in both the data and in the quantitative model. The model also accounts well for the small increase in the volatility of the labor wedge in the data.

Table 3: Data and Model Moments: Pre- and Post-1984

Moment	Data		Model	
	1967-1983	1984-2014	1967-1983	1984-2014
Std(Output)	0.021	0.011	0.013	0.007
Std(Efficiency Wedge)	0.017	0.009	0.010	0.004
Std(Labor Wedge)	0.014	0.015	0.010	0.011
Corr(Output, Productivity)	0.638	-0.079	0.552	-0.130
Corr(Efficiency Wedge, Output)	0.865	0.584	0.807	0.456
Corr(Labor Wedge, Output)	-0.535	-0.744	-0.445	-0.857
Skewness(Output Growth)	-0.025	-1.114	-0.645	-2.296

Notes: This table shows moments from both the data and the model. The model moments are produced using two different calibrations of the parameters governing the shock processes, one meant to capture the pre-1984 data and the other the post-1984 data. Moments are generated by simulating 15,000 different samples with the same number of observations as we have in the two data samples of 1967-1983 and 1984-2014. Each simulation uses a 400 period "burn-in." The moments presented in the table are the averages of the moments across the 15,000 different samples. The data moments are produced using standard definitions of variables as described in the text. All series are logged and HP-filtered with smoothing parameter 1600.

Turning to the correlations, the model performs well at matching the pre-1984 cyclicality of labor productivity, the efficiency wedge, and the labor wedge. While there are some quantitative discrepancies relative to the data, qualitatively the model does a fairly good job of matching the pre-Moderation moments of interest. The model also performs very well in capturing the changes, from the pre-1984 period to the post-1984 period, in the cyclicality of these variables. As in the data, labor productivity in the model becomes mildly countercyclical (correlation with output of -0.13), while the efficiency wedge becomes less correlated with output but remains procyclical (correlation with output of 0.46) and the labor wedge becomes even more countercyclical (correlation with output of -0.86). Overall, the model calibrated to feature a smaller relative contribution of aggregate shocks does a very good job at matching the observed changes in these business cycle moments. It is worth emphasizing that the model's ability to match these correlations (both the levels and the changes) is not a result of the way the model was calibrated, as none of these moments were targeted in the calibration procedure. In particular, the changing cyclicalities of productivity, the efficiency wedge, and the labor wedge are attributable entirely to the changing relative importance of the two shocks.

Reallocative shocks and aggregate shocks also have different implications for the skewness of output growth rates. Fluctuations driven solely by aggregate shocks generate relatively symmetric output fluctuations, as the increase in output that accompanies a positive aggregate shocks is similar in magnitude to the decline in output that results from a similarly sized negative shock. In contrast, fluctuations driven entirely by reallocative shocks generate a strongly left-skewed distribution of output growth rates, since reallocative shocks generally result in a drop in output (as employment declines while workers are being reallocated). Thus, a decrease in the relative importance of aggregate shocks should make the distribution of output growth rates more left skewed. The bottom row of Table 3 confirms this intuition. As noted above, in the data, output growth rates were basically symmetric in the pre-1984 period (skewness = -0.025) but became negatively skewed in the post-1984 period (skewness = -1.114).[15] The model's simulated data are more negatively skewed in both periods, relative to the data, but importantly the model exhibits an increase in the negative skewness in the calibration meant to match the post-1984 period.

Although it is not a primary focus of this paper, it is also worth examining the model's implications for the volatility of consumption and investment. In the baseline calibration of the model, consumption volatility goes from 0.0096 in the first part of the sample to 0.009. This is relatively close to what we observe in the data, where consumption volatility goes from 0.01 to 0.007. The volatility of investment in the model goes from 0.025 in the pre-1984 calibration, to 0.005 for the post-1984 period, whereas in the data investment volatility declines from 0.059 to 0.038. While the model's *decline* in investment volatility is on par with what is observed in the data, the overall *level* of volatility (in both periods) is too low. In particular, although investment in the model is more volatile than output in the pre-1984 period, as it should be, it is counterfactually less volatile than output in the post-1984 period.

It seems likely that the model's reduced investment volatility, relative to the data, can be traced to the assumption that capital is perfectly mobile across islands. If instead capital could not be seemlessly transferred across islands, then a reallocative shock would precipitate an increase in investment on the island with increased productivity that is currently absent in the model (instead capital currently just moves from one island to the other). This response of investment to reallocative shocks would raise the overall level of investment volatility (in both the pre-1984 and post-1984 calibrations), presumably bringing it more in line with the

[15] If one excludes the Great Recession by limiting the sample to the period 1984-2007, the distribution remains negatively skewed, with skewness= -0.25.

data.[16]

To get a better sense of the forces at work in the model, we next show the response of selected variables, in both the pre- and post-1984 calibrations, to recession events. In particular, we consider the following exercise. We simulate 200,400 periods of data from the model starting from its steady state (under either the pre- or post-1984 calibration). The first 400 periods are dropped as a "burn-in," leaving 200,000 observations. We then HP filter (with smoothing parameter 1600) the simulated log output series. We define periods of recession as periods in which HP filtered output is in its bottom 20th percentile.[17] For each simulated period in which the economy enters a recession so defined, we record the paths of the variables over the subsequent 20 periods. We then average over each of these recession events, which gives us the average dynamic paths of variables conditional on entering a recession. We express these paths relative to the unconditional mean of the variable in question (and thus the paths for the variables should, if given enough time, return to zero). While the idea of this exercise is similar to an impulse response, it differs in that we are not conditioning on the realization of a particular shock. Rather, we are conditioning on the economy entering a recession as we have defined it, which in the model can be driven by either an aggregate shock, a reallocative shock, or a combination of the two. Conditional on entering a recession, we are tracing out the expected time paths of variables.

Figure 2 plots the responses of output, labor hours, and productivity in the model using the pre-1984 parameterization. Output, hours, and average labor productivity all decline and then revert back toward their means after several periods. These responses are very similar to the responses to a negative productivity shock in a standard one-sector real business cycle model, where output, hours, and productivity all decline. Like that model, our model does not generate as large of a decline in hours as output. These responses are also qualitatively consistent with the behavior of these variables around recessions prior to the mid-1980s, where output, hours, and productivity all tend to decline together.[18]

[16]While it would be appealing to solve the model with immobile capital (or partially immobile, with an adjustment cost associated with movements of capital across islands), to do so would carry a significant computational burden. The state space, which already is quite large, would increase by one dimension (since both the capital stocks of both islands would need to be tracked, whereas currently only the aggregate capital stock must be tracked), as would the set of choice variables (currently the allocation of capital across the islands is a static problem, the first order conditions of which can be substituted into the problem, thus reducing the effective number of choice variables).

[17]The definition of a recession event is admittedly somewhat arbitrary. We have experimented with different definitions of a recession event (e.g. conditioning on large declines in output, using the same cutoff value in terms of the cyclical component of output to define a recession in the two separate simulations, etc.), and the resulting responses are qualitatively similar in terms of the co-movements of output, hours, and productivity.

[18]This comparison to actual recessions is only meant to be qualitative. Like the basic neoclassical model, our model lacks a strong propagation mechanism and fails to generate inertial, hump-shaped behavior of

Figure 2: Responses following a Recession: Pre-1984 Calibration

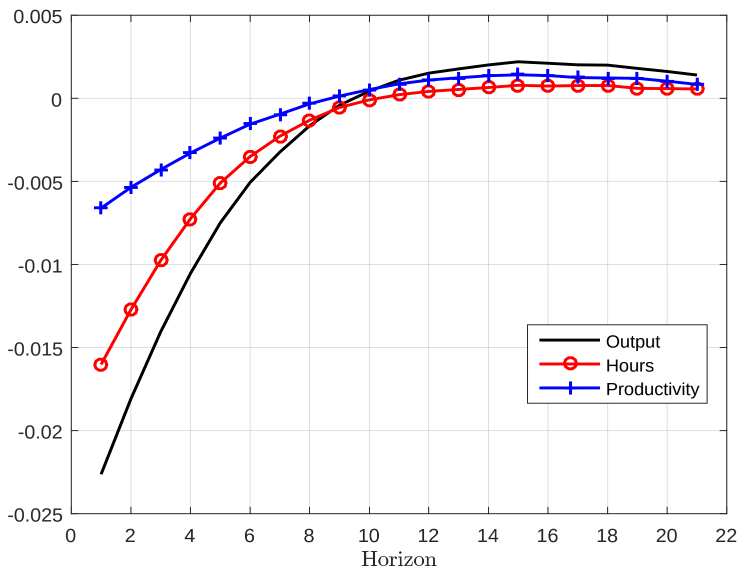

Notes: This figure plots the responses of output, hours, and labor productivity following periods identified as recessions in the model using the pre-1984 calibration. The definition and construction of these responses is described in the text. All variables are logged and expressed relative to the their unconditional means, so the units of the vertical axis are percentage deviations from the unconditional means. The units of the horizontal axis are quarters.

Next, we consider the responses following a "recession" in the model using the post-1984 parameterization. These responses are shown in Figure 3. The definition of a recession is identical, though it is based on the bottom 20th percentile of HP filtered output in the post-1984 simulation, which is not as negative as in the pre-1984 simulation, given the lower overall volatility of output. These responses differ in important ways relative to the responses from the pre-1984 calibration. First, productivity rises even though output falls. Second, this occurs because the decline in hours is larger than the decline in output, whereas in the pre-1984 calibration the reverse is true. Third, the hours response seems somewhat more protracted (relative to the output response) in the post-1984 calibration compared to the pre-1984 version of the model. These features qualitatively correspond with features of post-1984 recessions in the data, where average productivity does not decline much (or in fact rises) and hours worked and other labor market indicators recover slowly relative to output.

these variables.

Figure 3: Responses following a Recession: Post-1984 Calibration

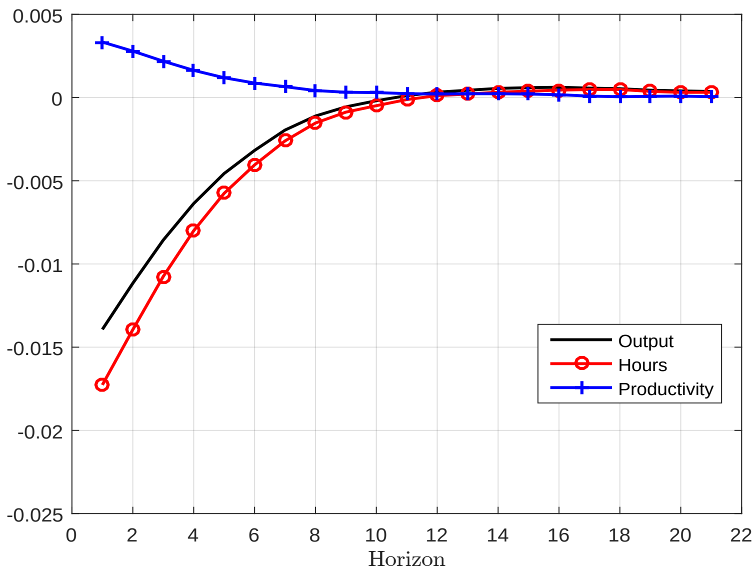

Notes: This figure plots the responses of output, hours, and labor productivity following periods identified as recessions in the model using the post-1984 calibration. The definition and construction of these responses is described in the text. All variables are logged and expressed relative to the their unconditional means, so the units of the vertical axis are percentage deviations from the unconditional means. The units of the horizontal axis are quarters.

We next pursue in a more formal way the idea, evident in the figure, that the response of hours in a recession has become more persistent, relative to the response of output, in the post-1984 period. We begin by looking for evidence of this feature in the data. We proceed as follows. In our sample period, there are three recessions (as identified by the NBER) in the pre-1984 sample, and three recessions in the post-1984 sample, where we treat the "double-dip" recessions of 1980 and 1981-1982 as a singular event.[19] For each of these episodes, we identify the quarter within a ten period window in which the cyclical component of HP filtered output is lowest, which we take to measure the trough. It is worth noting that the identified trough does not necessarily line up with the NBER dates.[20] We then compute the

[19]The NBER dates for the three pre-1984 recessions are 1969Q4-1970Q4, 1973Q4-1975Q1, and 1980Q1-1982Q4 (again, noting that we treat the two separate recessions dated 1980Q1-1980Q3 and 1981Q3-1982Q4 as a singular event). The dates for the post-1984 recessions are 1990Q3-1991Q1, 2001Q1-2001Q4, and 2007Q4-2009Q2.

[20]The identified trough dates in this exercise are 1970Q4, 1975Q2, 1982Q4, 1992Q2, 2003Q3, and 2010Q2. For the pre-1984 recessions, the identified trough is either the last period of the NBER defined recession, or the period after. For the post-1984 recessions, the identified trough is several quarters after the end of the recession as defined by the NBER in each case. This finding accords with the analysis in Gali, Smets, and Wouters (2012) that recoveries since the mid-1980s have been slower compared to the period before the

number of quarters, rounded to the nearest integer, that it takes output and hours to recover halfway back to trend. We refer to this number of periods as the half-life of output or hours conditional on being at a trough.

Table 4: Estimated Half-Lives of Hours and Output in the Data and the Model

Sample Period	Half-Life of L_t	Half-Life of Y_t	Relative Half-Life
Data			
1967-1983	2.67	2.33	1.15
1984-2014	4.67	2.67	1.75
Model			
1967-1983	1.99	2.60	0.76
1984-2014	2.38	2.49	0.96

Notes: This table shows the approximate half-lives of output and hours conditional on being at the trough of a recession, both for the data and the model. The construction of the half-life statistics is described in the text.

The upper panel of Table 4 shows the average half-lives of output and hours for the three recessions identified in each subsample. For the pre-1984 period, the average half-life of hours is 2.67 quarters, while it is 2.33 quarters for output. The relative half-life of hours, equal to the ratio of the two half-lives, is 1.15. For the post-1984 period, the average half-life of hours is substantially larger, at 4.67 quarters. The half-life of output is also larger compared to pre-1984 recessions, but only slightly. The relative half-life of hours therefore increases to 1.75 in the post-1984 period.

The lower panel of Table 4 shows half-life statistics for both output and hours in the two different calibrations of the model. We use the same definition of a recession as we used to construct Figures 2 and 3. Given this definition of a recession event, the construction of the half-lives is identical to what we do with the data. The numbers presented in the table are the average half-lives across all recession events in the simulation. For the pre-1983 period, the model generates a half-life of hours that is a bit smaller than in the data, and a half-life of output that is a bit too big. Overall, the simulated half-lives are nevertheless qualitatively in-line with what is observed in the data, particularly when factoring in that the numbers from the data are averages based on only three recessions. When shifting to the post-1984 calibration with smaller aggregate shocks, one observes that there is an increase in the half-life of hours, while the half-life of output is roughly unchanged. The 0.2 increase in the relative half-life of hours in the model is qualitatively in-line with what is observed in the data, though quantitatively our model does not generate as large an increase in the relative half-life as is seen in the data.

mid-1980s.

Our finding that employment recoveries are slower relative to output recoveries, both in the data and in our model, is related to the so-called "Jobless Recovery" phenomenon, which has been a subject of debate in the literature. There does not appear to be a well-accepted definition of what is meant by a jobless recovery. If one were to take a strict definition in which labor market variables continue to decline after output has begun to recover, our model does not produce a jobless recovery—as can be seen in Figure 3, upon entering a recession event hours begins to recover as soon as output does. A weaker definition of a jobless recovery would be one in which hours recovers more slowly relative to output. This seems to be a feature of both the data and our model.[21]

In summary, our relatively simple islands model calibrated to match the relative importance of aggregate shocks in the pre- and post-1984 samples is capable of capturing several salient changes in business cycle moments. In the data, the post-1984 "Great Moderation" is associated not only with a drop in aggregate output volatility, but also with a declining procyclicality of labor productivity, an increasing importance of the labor wedge relative to the efficiency wedge, and slower employment recoveries in the wake of recessions. Our model captures all of these features well. The model also generates an increase in the left-skewness of output growth.

4.4 Robustness

This subsection examines the robustness of our model's quantitative predictions to parameter values. We are especially interested in understanding how sensitive the quantitative results are to different values of the parameters that were more difficult to calibrate, yet which seem potentially important for results. In particular, we are interested in the parameter that controls the ease or speed with which workers can reallocate—i.e. λ—and the parameters that control the persistence of the two types of shocks—i.e. ρ_A and ρ_Z. One approach to this sensitivity analysis would be similar to a comparative statics exercise: consider a small change in the parameter of interest, holding other parameters fixed at their baseline values, and assess the impact that the change has on the moments of interest. This approach has the benefit of offering clean insights into the workings of the model. On the other hand, given the potentially non-linear nature of the model, these insights are only valid "locally." An

[21] Gali et al. (2012) argue that there is no jobless recovery phenomenon in the post-1984 data, but rather that recoveries are generally slower. They base this analysis on cumulative growth rates of output and hours after troughs, noting that output growth and hours growth are much weaker in post-1984 recessions compared to earlier recessions, but that the difference between output growth and hours growth is relatively unchanged. It is not obvious whether looking at differences in growth rates is the appropriate comparison. For example, prior to the 1990-1991 recession, hours growth is on average about one-half as large as output growth in the one to two years after a trough. Starting with the 1990-1991 recession, hours growth is between zero and twenty percent as large as output growth during the recovery phase.

alternative approach would be to consider a larger change in the parameter of interest, while at the same time re-parameterizing the model so that it continues to match the moments that guide the original calibration, and then assess how the non-targeted moments are affected. In Appendix A we report results of the local, or comparative statics, approach.[22] Here, we report results from exercises that consider larger changes in the parameter. In these exercises, after changing the parameter of interest, we re-calibrate the standard deviations of the shock processes to again match (as described in the calibration discussion above) the volatility of IP production and the relative contribution of the two types of shocks, for each of the two periods.

We first focus on λ. For this exercise, we consider a value of $\lambda = 0.2$, which is double the value we assume in the baseline. Figure 4 presents responses following recession events in two different shock calibrations: one, on the left, meant to match the pre-1984 data, and another, on the right, meant to match the post-1984 data.

Figure 4: Responses Following a Recession: Sensitivity to λ

(a) Pre-1984 (b) Post-1984

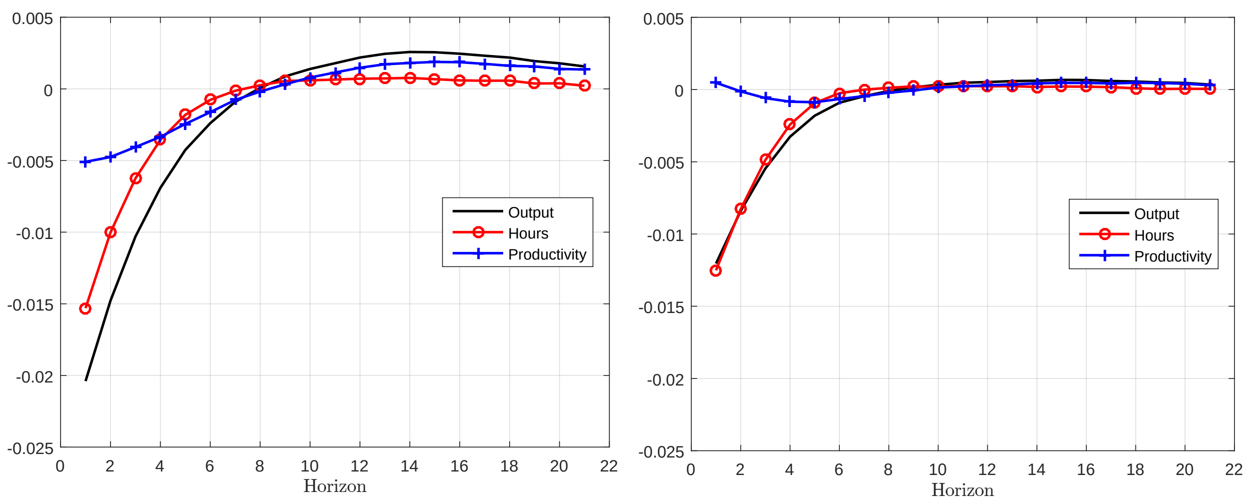

Notes: These figures plots the responses of output, hours, and productivity for pre-1984 and post-1984 calibrations of the model with a much higher value of λ (i.e. 0.2) than in the baseline. The volatilities of the shock processes were re-calibrated using the same approach used in the baseline calibration.

There is not much evident difference in the left panel compared to Figure 2. There are, however, some differences when comparing the right panel to Figure 3. Although productivity does still increase while output falls, the initial increase in productivity is much smaller and soon turns slightly negative. Correspondingly, for a higher value of λ the model generates a

[22] In addition to perturbations of λ, ρ_A and ρ_Z, the comparative statics exercises in the Appendix also consider change in the value of σ. We also considered larger changes in the value of σ but we do not report those results in the main text as they do not differ significantly from the local exercises presented in Appendix A.

smaller decline in the cyclicality of productivity. This can be seen in Table 5, which shows a comparison of pre- and post-1984 moments using a larger value of λ. In particular, with this larger value of λ, the correlation of productivity with output falls from 0.59 in the pre-1984 sample to 0.28 in the post-1984 sample. While a decline of 0.3 points is considerable, it is about half as large as the decline that was observed for $\lambda = 0.1$. With a larger value of λ, while the model produces an increase in the relative half-life of hours after the trough of a recession, it is smaller than when $\lambda = 0.1$. In particular, the relative half-life of hours goes from 0.66 in the pre-1984 calibration to 0.74 in the post-1984 calibration, a significantly smaller increase than what is shown in Table 4.

Table 5: Sensitivity to a larger value of λ

	$\lambda = 0.1$		$\lambda = 0.2$	
Moment	1967-1983	1984-2014	1967-1983	1984-2014
Std(Output)	0.013	0.007	0.012	0.007
Std(Efficiency Wedge)	0.010	0.004	0.010	0.005
Std(Labor Wedge)	0.010	0.011	0.094	0.009
Corr(Output, Productivity)	0.552	-0.130	0.584	0.280
Corr(Efficiency Wedge, Output)	0.807	0.456	0.826	0.659
Corr(Labor Wedge, Output)	-0.445	-0.857	-0.448	-0.759
Skewness(Output Growth)	-0.645	-2.296	-0.561	-1.251

Notes: This table reproduces the model's key moments with $\lambda = 0.1$ and compares them to a larger value of λ, under the heading $\lambda = 0.2$. For the higher value of λ, the standard deviations of the shock processes are re-calibrated.

We next consider, separately, larger values of the autoregressive parameters governing the aggregate and reallocative shock processes. (Need a sentence here about why this is an interesting exercise—i.e. hard to directly measure the persistence of these shocks, you might try to choose them to match persistence of the model's measured productivity, but it turns out you can't match that due to lack of internal propagation.) Moments for the pre- and post-1984 calibrations are summarized in Table 6.

Relative to our baseline (reported in Table 3), when the aggregate productivity shock is more persistent ($\rho_A = 0.98$, compared with a baseline value of 0.92), the volatilities of output and the efficiency wedge are slightly smaller in the pre-1984 calibration, but the changes in these moments moving to the post-1984 calibration are the same as in our baseline. The volatility of the labor wedge and its change across samples is roughly the same as when

$\rho_A = 0.92$. Compared to when ρ_A is smaller, the correlations of output with average labor productivity and the efficiency wedge are both smaller in the pre-1984 calibration, though both still decline significantly when calibrating the shock standard deviations to match post-1984 moments. The levels and changes in the correlation of the labor wedge with output are virtually identical to our baseline calibration. Relative to our baseline calibration, output growth is more negatively skewed in the pre-1984 calibration, but still becomes significantly more left-skewed when moving to the post-1984 calibration.

Table 6: Sensitivity to Higher Persistence of Shocks

	$\rho_A = 0.98$		$\rho_z = 0.98$	
Moment	1967-1983	1984-2014	1967-1983	1984-2014
Std(Output)	0.012	0.007	0.013	0.007
Std(Efficiency Wedge)	0.007	0.003	0.010	0.004
Std(Labor Wedge)	0.010	0.010	0.006	0.007
Corr(Output, Productivity)	-0.007	-0.398	0.679	0.280
Corr(Efficiency Wedge, Output)	0.449	0.305	0.872	0.701
Corr(Labor Wedge, Output)	-0.532	-0.855	0.020	-0.597
Skewness(Output Growth)	-1.848	-2.637	-0.886	-3.576

Notes: This table reports the model's key moments with $\rho_A = 0.98$ and $\rho_z = 0.98$. The standard deviations for the stochastic processes are disciplined as in the baseline exercise.

We next turn attention to the case when reallocative shocks are more persistent, with $\rho_z = 0.98$ instead of 0.95. Results are summarized in the right columns of Table 6. In terms of volatilities, the principal effect of this change in calibration is to lower the volatility of the labor wedge relative to our baseline analysis, though its change across pre- and post-1984 calibrations is unaffected. Focusing next on correlations, a higher value of ρ_z results in smaller changes in the cyclicalities of labor productivity and the efficiency wedge relative to our baseline analysis, though qualitatively these correlations still change in the same direction as they do in the data. With a more persistent island-specific productivity shock, in the pre-1984 sample the labor wedge is acyclical (as opposed to strongly countercyclical, both in the data as well as in the baseline calibration of our model), though it changes to strongly countercyclical in the post-1984 calibration. Finally, there is a larger increase in the left skewness of output growth when going from the pre- to post-1984 calibrations when $\rho_z = 0.98$, though qualitatively the change in skewness is the same as in the data and our baseline calibration of the model.

5 Conclusion

This paper provides evidence that the so-called "Great Moderation" period is associated with a decline in the importance of aggregate shocks relative to sector-specific shocks. It develops an island model with both aggregate and island-specific shocks. Moving labor between islands is costly in terms of time spent unable to work. Calibrating the shock processes to match empirical facts about the relative importance of aggregate and sector-specific shocks, we show that the model calibrated to the post-1984 data is capable of generating several features consistent with the data. Labor productivity switches from procyclical to mildly countercyclical, the labor wedge becomes more negatively correlated with output and significantly more volatile relative to the efficiency wedge, employment recoveries after recessions are slower compared to the recovery in output, and output growth becomes significantly left-skewed.

Our analysis has potentially important implications for economic policy. If aggregate shocks are a less important driver of aggregate volatility, then this raises the question of how aggressive countercyclical demand management policies ought to be. Stimulating demand through aggressive monetary easing or fiscal expansion may only serve to postpone the necessary reallocation of resources. As written, our model is quite stylized and lacks the kinds of frictions (e.g. nominal rigidities) which would allow one to sensibly discuss policy. The role of policy in a world where aggregate shocks are relatively less important is an area that merits further investigation.

References

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

Alvarez, F. and Shimer, R. (2011), "Search and Rest Unemployment", *Econometrica*, **79**, 75–122.

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

Barnichon, R. (2010), "Productivity and Unemployment over the Business Cycle", *Journal of Monetary Economics*, **57**, 1013–1025.

Berger, D. (2015), "Countercyclical Restructuring and Jobless Recoveries", Technical report, Northwestern University.

Broda, C. and Weinstein, D. E. (2006), "Globalization and the Gains from Variety", *The Quarterly Journal of Economics*, **121**, 541–585.

Chari, V. V., Kehoe, P. J., and McGrattan, E. R. (2007), "Business Cycle Accounting", *Econometrica*, **75**, 781–836.

Clark, K. B. and Summers, L. H. (1979), "Labor Market Dynamics and Unemployment: A Reconsideration", *Brookings Papers on Economic Activity*, pages 13–72.

Davis, S. J. (1987), "Fluctuations in the pace of labor reallocation", *Carnegie-Rochester Conference Series on Public Policy*, **27**, 335–402.

Fernald, J. (2014), "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity", Federal Reserve Bank of San Francisco Working Paper Series 2012-19.

Foerster, A. T., Sarte, P.-D. G., and Watson, M. W. (2011), "Sectoral versus Aggregate Shocks: A Structural Factor Analysis of Industrial Production", *Journal of Political Economy*, **119**, 1 – 38.

Fujita, S. and Moscarini, G. (2013), "Recall and Unemployment", Working paper series, Yale University.

Gadea, M., Gomez-Loscos, A., and Perez-Quiros, G. (2014), "The Two Greatest. Great Recession vs. Great Moderation", Working papers 1423, Banco de Espana.

Gali, J. and Gambetti, L. (2009), "On the Sources of the Great Moderation", *American Economic Journal: Macroeconomics*, **1**, 26–57.

Gali, J., Smets, F., and Wouters, R. (2012), "Slow Recoveries: A Structural Interpretation", *Journal of Money, Credit, and Banking*, **44**, 9–30.

Gali, J. and van Rens, T. (2014), "The Vanishing Procyclicality of Labor Productivity", CEPR Working Papers 9853.

Hagedorn, M. and Manovskii, I. (2011), "Productivity and the Labor Market: Comovement over the Business Cycle", *International Economic Review*, **52**, 603–619.

Hall, R. E. (2007), "How Much Do We Understand about the Modern Recession?", *Brookings Papers on Economic Activity*, **38**, 13–30.

Hansen, G. D. (1985), "Indivisible labor and the business cycle", *Journal of Monetary Economics*, **16**, 309–327.

Hornstein, A. (2012), "Accounting for unemployment: the long and short of it", Working Paper 12-07, Federal Reserve Bank of Richmond.

Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G. (1993), "Earnings Losses of Displaced Workers", *American Economic Review*, **83**, 685–709.

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.

Lucas, R. J. and Prescott, E. C. (1974), "Equilibrium search and unemployment", *Journal of Economic Theory*, **7**, 188–209.

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.

Mehrotra, N. R. and Sergeyev, D. (2012), *Sectoral Shocks, the Beveridge Curve and Monetary Policy*, Ph.D. thesis, Columbia University.

Rogerson, R. (1988), "Indivisible labor, lotteries and equilibrium", *Journal of Monetary Economics*, **21**, 3–16.

Ruhm, C. J. (1991), "Are Workers Permanently Scarred by Job Displacements?", *American Economic Review*, **81**, 319–24.

Sahin, A., Song, J., Topa, G., and Violante, G. L. (2012), "Mismatch Unemployment", Working Paper 18265, National Bureau of Economic Research.

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

Tauchen, G. (1986), "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions", *Economics Letters*, **20**, 177–181.

Wiczer, D. (2013), *Long-Term Unemployment: Attached and Mismatched?*, Ph.D. thesis, University of Minnesota.

A Additional Robustness Exercises

The robustness exercises that we present are essentially local comparative statics: we perturb one parameter at a time by a small amount, holding all other parameters (including the shock magnitudes) fixed at their baseline values, and examine how the model-generated moments change.

We first consider small changes in the parameter σ. The key moments of interest, for $\sigma = 2.8$ and $\sigma = 3.2$ (small perturbations, relative to the baseline value of $\sigma = 3$), are summarized in the top panel of Table 7. A higher value of σ tends to increase the volatility of the labor wedge and to reduce the correlation of output with average labor productivity. To understand this, note that a higher value of σ means that island-specific outputs are more substitutable so that, conditional on island-specific shocks, there is relatively more reallocation of labor. As discussed in Section 3, this tends to make the labor wedge more volatile. Moreover, because the reallocation process lowers output while raising productivity, greater amounts of reallocation (due to a higher σ) lowers the overall correlation between aggregate productivity and output. While the value of σ has these effects on the *levels* of these moments, there is little difference (compared to the baseline results in Table 3) when focusing on how these moments *change* when going from the pre- to post-1984 calibrations. In particular, for values of σ between 2.8 and 3.2, the labor wedge continues to be approximately equally volatile in the pre- and post-1984 calibrations, and the correlation between output and productivity still declines by about 0.7.

We next consider, in the second panel of Table 7, small perturbations in the parameter λ. This parameter governs how time-consuming, and thus costly, it is to reallocate labor across islands. A higher value of λ results in more reallocation and thus one would expect the volatility of the labor wedge and the cyclicality of labor productivity to be lower for higher values of λ. This is in fact what we see in the table, though the differences relative to the baseline case are not large. As with σ, small differences in the value of λ have little effect on how the moments of interest *change* across the two calibrations of the shock magnitudes.

The third panel of Table 7 shows the results from small changes in ρ_A, the parameter governing the persistence of the aggregate shock. Higher values of ρ_A tend to reduce the *levels* of the correlation between productivity and output and of the correlation between the measured efficiency wedge and output. The intuition for this result is that hours (and thus output) react less to aggregate productivity shocks the more persistent are those shocks, via a standard wealth effect argument. The *changes* in moments, when moving from the pre- to the post-1984 calibration, is roughly invariant to the perturbations of ρ_A.

The bottom panel of Table 7 displays the model's moments for small changes in ρ_z, the

Table 7: Local Comparative Statics

Moment	1967-1983	1984-2014	1967-1983	1984-2014
	σ = 2.8		σ = 3.2	
Std(Output)	0.013	0.007	0.014	0.008
Std(Efficiency Wedge)	0.010	0.004	0.010	0.004
Std(Labor Wedge)	0.009	0.010	0.011	0.012
Corr(Output, Productivity)	0.626	-0.064	0.522	-0.207
Corr(Efficiency Wedge, Output)	0.839	0.490	0.796	0.427
Corr(Labor Wedge, Output)	-0.384	-0.846	-0.460	-0.874
Skewness(Output Growth)	-0.502	-2.002	-0.717	-2.453
	λ = 0.08		λ = 0.12	
Std(Output)	0.013	0.007	0.013	0.007
Std(Efficiency Wedge)	0.010	0.004	0.010	0.004
Std(Labor Wedge)	0.010	0.010	0.010	0.011
Corr(Output, Productivity)	0.574	-0.096	0.549	-0.127
Corr(Efficiency Wedge, Output)	0.816	0.458	0.809	0.467
Corr(Labor Wedge, Output)	-0.405	-0.844	-0.445	-0.862
Skewness(Output Growth)	-0.545	-2.086	-0.646	-2.168
	ρ_A = 0.90		ρ_A = 0.94	
Std(Output)	0.014	0.007	0.012	0.007
Std(Efficiency Wedge)	0.011	0.004	0.009	0.003
Std(Labor Wedge)	0.010	0.011	0.010	0.011
Corr(Output, Productivity)	0.615	-0.051	0.437	-0.261
Corr(Efficiency Wedge, Output)	0.836	0.499	0.754	0.391
Corr(Labor Wedge, Output)	-0.373	-0.840	-0.524	-0.885
Skewness(Output Growth)	-0.491	-2.112	-0.891	-2.509
	ρ_z = 0.93		ρ_z = 0.97	
Std(Output)	0.013	0.007	0.012	0.006
Std(Efficiency Wedge)	0.010	0.004	0.010	0.004
Std(Labor Wedge)	0.011	0.012	0.007	0.007
Corr(Output, Productivity)	0.486	-0.228	0.739	0.260
Corr(Efficiency Wedge, Output)	0.774	0.362	0.893	0.658
Corr(Labor Wedge, Output)	-0.462	-0.872	-0.257	-0.760
Skewness(Output Growth)	-0.512	-1.433	-0.404	-2.240

Notes: This table reports the model's key moments, for both the pre-1984 and post-1984 calibration of shocks, when four model parameters are increased and decreased relative to their baseline values. The baseline values of those parameters are $\sigma = 3$, $\lambda = 0.1$, $\rho_A = 0.92$, and $\rho_z = 0.95$.

persistence parameter for the island-specific shock. When island-specific shocks are more persistent, there is greater incentive to reallocate workers to the more productive island, as it is likely to remain more productive for longer. However, because shifts in relative productivity are longer-lived, labor is allocated suboptimally across islands for a smaller fraction of time, and as such one would expect a less volatile labor wedge. The results confirm this intuition, as the standard deviation of the labor wedge is smaller when ρ_z is greater. In a similar way, because workers are on average better allocated across islands, changes in productivity and output are driven mostly by aggregate shocks and are therefore more positively correlated. In terms of the impact of ρ_z on the changes in the moments when moving from the pre-1984 shock calibration to the post-1984 calibration, the declines in the correlations of output with labor productivity, or with the measured efficiency wedge, are slightly smaller when z is more persistent.