Working Paper Series

What Inventory Behavior Tells Us About How Business Cycles Have Changed

WP 14-06R

Thomas Lubik
Federal Reserve Bank of Richmond

Pierre-Daniel Sarte
Federal Reserve Bank of Richmond

Felipe Schwartzman
Federal Reserve Bank of Richmond

THE FEDERAL RESERVE BANK OF RICHMOND
RICHMOND ▪ BALTIMORE ▪ CHARLOTTE

What Inventory Behavior Tells Us About How Business Cycles Have Changed

Thomas Lubik Pierre-Daniel Sarte

Felipe Schwartzman*

Federal Reserve Bank of Richmond

August 6, 2014

Working Paper 14-06R

Abstract

Beginning in the mid-1980s, the nature of U.S. business cycles changed in important ways, as made evident by distinctive shifts in the comovement and relative volatilities of key economic aggregates. These include labor productivity, hours, output, and inventories. Unlike the widely documented change in absolute volatility over that period, known as the Great Moderation, these shifts in comovement and relative volatilities persist into the Great Recession. To understand these changes, we exploit the fact that inventory data are informative about sources of business cycles. Specifically, they provide additional information relative to aggregate investment regarding firms' intertemporal decisions. In this paper, we show that the "investment wedge" estimated with inventories, unlike previous measures, correlates well with established independent measures of credit market frictions. Furthermore, contrary to previous findings, our generalized investment wedge informed by inventory behavior plays a key role in explaining the shifts in U.S. business cycles observed after the mid-1980s.

Keywords: Business Cycles, Inventories, Investment Wedge, Financial Frictions

*We thank Fabio Schiantarelli, Mark Watson, Esteban Rossi-Hansberg, Julia Thomas, Aubhik Khan, and Jonathan Willis for helpful and detailed comments. We also thank seminar participants at the University of Rochester, the University of Virginia, PUC-Rio, FGV-EESP, the 2012 CEF meetings, the 2013 Summer Econometric Society Meetings, and the 2013 Sao Paulo Macro Workshop for their comments. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System. All errors are our own.

JEL Code: E32, E44

1 Introduction

The post-war U.S. economy is characterized by a noticeable fall in the volatility of key macroeconomic aggregates after the early 1980s. This has come to be known as the Great Moderation. However, this period is also characterized by notable changes in the comovement properties of some of these aggregates. In particular, labor productivity ceases to be procyclical and, strikingly, starts to comove negatively with hours worked over the post-1984 period.[1] These shifts in comovement cannot be easily reconciled with conventional real business cycle models driven by temporary Hicks-neutral productivity shocks. Over the same period, as we document in the paper, the inventory-sales ratio has become more procyclical, and the volatility of both inventories and hours relative to output has increased. Unlike the fall in absolute volatility, these shifts in comovement patterns and relative volatilities continue into the Great Recession. This paper argues that a change in the cyclical relevance of financial frictions have played a key role in explaining not just the Great Recession but, more generally, post-1984 business cycles.

Our work builds on a literature exemplified by Bils and Kahn (2000), that uses inventory data to inform our understanding of business cycles.[2] Inventories are informative about factors driving business cycles because they help differentiate between changes in investment induced by variations in i) the physical return to fixed investment, which is typically governed by technological opportunities (Fisher, 2006; Greenwood, Hercowitz, and Krusell, 2000); and ii) the wedge between this return and the discount rate implied by household consumption behavior, which is often associated with credit market frictions. This distinction points to variations in the "investment wedge" as having played an increasingly prominent role in explaining inventory behavior after 1984. Moreover, investment wedge fluctuations also help explain the reduced comovement between output and labor productivity, commonly referred to as the "labor productivity puzzle". This finding arises because fluctuations in the investment wedge affect labor supply through intertemporal substitution.

The investment wedge that we estimate using inventory data resembles indicators of credit market frictions emphasized in the literature. Unlike conventional measures of invest-

[1]See Gordon (2010), and Garin, Pries, and Sims (2013)

[2]The study by McConnell and Perez-Quiros (2000) on the decline in US GDP volatility in the early 1980s generated a large literature that investigated changes in inventory management as the primary driver (see also Kahn, McConnell, and Perez-Quiros (2002)). Using the U.S. automobile industry as a case study, Ramey and Vine (1989); **?** show that this can explain the changing behavior of US time series both in terms of volatility and comovement. More recent work, such as Chang, Hornstein, and Sarte (2009) extended this line of work to account for changes in pricing behavior and mark-ups, building on Bils and Kahn (2000).

ment wedges that abstract from information on inventories, it comoves robustly with credit frictions indicators, for instance, the spread between the returns of high and low credit risk bonds, or dividend payouts to business owners. Similarly to our investment wedge, conventional credit friction indicators also comove more strongly with output and fixed investment after 1984 relative to the pre-Great Moderation period. Finally, the negative comovement of our investment wedge with output and investment is concentrated in mainly two episodes: the savings and loans crisis of the mid-to-late 1980s, and the period following the two large financial regulation overhauls of 1994 and 1999.

Following the methodology introduced by Chari, Kehoe, and McGrattan (2007), we investigate deviations of the data from the implications of a frictionless model in order to separate the role of shocks affecting the production structure of the economy (efficiency wedges) from those distorting labor supply and investment decisions (the labor wedge and the investment wedge). To allow flexibility in the way that efficiency wedges affect business cycles, the model extends the canonical frictionless environment to include multiple stages of production, productivity changes that affect sectors differently, and both permanent and temporary shifts in production possibilities, which we discipline with the use of inventory data and observed input-output linkages. This rich production structure incorporates previous developments that have emphasized different dimensions of productivity shocks. In particular, it nests the time-to-build technology of Kydland and Prescott (1982), as well as input-output linkages.[3] In addition, the framework shares key features of conventional inventory models characterized by a linear storage technology and a long-run target inventory-sales ratio, such as the stock out avoidance model in Kahn (1992), which Iacoviello, Schiantarelli, and Schuh (2011) and Wen (2011) extended to a general equilibrium setting. This rich structure is important since the measurement of wedges is sensitive to details of the production technology (Christiano and Davis, 2006).[4] Specifically, the more flexible structure allows for effects of the production technology on allocations that might otherwise be assigned to behavioral wedges. We discuss in detail in Section 4 how the production structure modifies the calculation of standard wedges.

Our analysis of the shifts in properties of U.S. business cycles starts with an assessment of the extent to which they can remain consistent with the early Real Business Cycles

[3]For example, Long and Plosser (1982), Dupor (1999), and Foerster, Watson, and Sarte (2011).

[4]Buera and Moll (2012) develop a model where aggregate savings derive largely from entrepreneurs with heterogeneous productivity and their own credit constraint. In that environment, shocks that intensify credit frictions may lead to movements in aggregate TFP or the labor wedge without necessarily affecting the investment wedge.

(RBC) view that emphasizes exogenous shocks to efficiency wedges. Since we consider a rich production structure, productivity shocks can, in principle, affect labor supply and investment decisions in more diverse ways relative to previous literature. This flexibility opens up the possibility that the changes in comovement and relative volatilities after 1984 reflect changes in the nature of productivity shocks, as opposed to an increased role of distortions in labor and financial markets. However, our findings do not support this view. As indicated by the early RBC literature, efficiency shocks are able to account for the comovement of output, hours, labor productivity, and inventories, and much of their variation prior to 1984. After that date, this is no longer the case even when productivity shocks operate along the various margins we allow in this paper.

While the labor productivity puzzle, the negative correlation between labor productivity and hours, and the disappearance of the countercyclicality of the inventory-sales ratio, remain a challenge for the notion of productivity-driven business cycles, the investment wedge estimated using inventory data plays a critical role in explaining these post-1984 facts. Our findings, taken together with the similarities in behavior between the investment wedge and measures of credit conditions, suggest that models emphasizing credit frictions, as well as other factors influencing the modified investment wedge, are important for understanding shifts in business cycles that occurred over the last three decades.

A large literature explores the role of financial factors for business cycles, such as Christiano, Motto, and Rostagno (2010), Gilchrist, Sim, and Zakrajsek (2010), Jermann and Quadrini (2012), Shourideh and Zetlin-Jones (2012), and Khan and Thomas (2013). Our paper contributes to this literature and connects it to two other lines of research. The first examines changes in the composition of shocks affecting the economy after the Great Moderation, for instance, Justiniano, Primiceri, and Tambalotti (2010). We share with this latter paper the finding that business cycles after the Great Moderation were greatly influenced by shocks that affect investment, but which are not necessarily related to the relative price of investment goods. While these authors consider a New Keynesian setting, we come to this conclusion with a model that adheres to the spirit of the early RBC literature in that it explores the implications of technological constraints for business cycle patterns.[5] The second is concerned with using wedges as helpful indicators of market frictions, such as Hall (1996), Chari, Kehoe, and McGrattan (2007) and Hsieh and Klenow (2009). In particular,

[5] Aside from the literature on financial factors and business cycles, our work simultaneously speaks to a growing literature motivated by the labor productivity puzzle. For a recent example that looks at alternative candidate explanations for that puzzle, see Garin, Pries, and Sims (2013).

we show that an aggregate investment wedge can be correlated with indicators of financial frictions and, moreover, that the frictions implied by this wedge have played an increasingly prominent role.

There is relatively little work that underscores changes in inventory behavior before and after 1984. Notable exceptions are Benati and Lubik (2014), which describes changes in inventory behavior over the last century in a time-varying parameter VAR framework, and Iacoviello, Schiantarelli, and Schuh (2011) who study those changes in the context of an estimated general equilibrium model with an emphasis on whether these changes can account for the observed fall in output volatility. We focus instead on whether changes in factors affecting inventory investment also help explain changes in the comovement between output and labor.

This paper is organized as follows. Section 2 highlights some key ways in which the nature of post-war U.S. business cycles changed after 1984. Section 3 develops a frictionless model of business cycles with a generalized production technology and a role for inventory formation. Section 4 contrasts the composition of wedges under this generalized approach with those that emerge in the standard one-sector growth model. Section 5 discusses the data and estimation. Section 6 presents, deconstructs, and provides external validation for, our empirical findings. Section 7 concludes.

2 Business Cycles Old and New

Table 1 summarizes changes in the behavior of key U.S. economic aggregates over the post-war era. We focus on output, consumption, investment, hours, and inventories.[6] Inventories are measured as total inventories, including raw materials, work in process, and finished goods. Except for hours, all data is obtained from the Bureau of Economic Analysis and cover the period from 1953Q3 to 2011Q4. Hours data are obtained from the Bureau of Labor Statistics and are constructed as described in Francis and Ramey (2005). We focus on business cycle frequencies and pass the raw data through a standard HP-filter.[7]

Table 1 shows that prior to 1984, labor productivity, defined as output per hour, is strongly procyclical. Its correlation with output is 0.65. Moreover, labor productivity co-

[6]In this paper, output is defined as the sum of final private expenditures of domestic residents, and includes personal consumption expenditures and private investment in both fixed assets and inventories. We exclude government consumption and the trade balance.

[7]The salient changes in the nature of business cycles we consider are robust to different detrending methods, including expressing the data in growth rates.

moves positively with labor input. In addition, consumption is only a third as volatile as output while, conversely, the standard deviation of investment is more than twice that of output. These basic properties of U.S. economic aggregates are all easily reproduced in a prototypical one-sector growth model driven solely by Hicks-neutral productivity shocks. Indeed, the ascendancy of the RBC literature starting in the mid-1980s is intimately linked with the behavior of key economic aggregates up to that time, as exemplified in King, Plosser, and Rebelo (1988).

Table 1 also shows that the inventory-sales ratio prior to 1984 is countercyclical. Its correlation with output is negative at -0.57, thus lining up with findings in Bils and Kahn (2000). While the early RBC literature mostly abstracts from inventory behavior, we show in Section 6.1 that the countercyclicality of the inventory-sales ratio can also be supported by productivity shocks in an extension of the RBC framework that allows for inventory formation. Intuitively, a positive productivity shock raises the marginal return to fixed investment relative to inventory investment, so that inventories rise relatively less rapidly than other aggregates.

The behavior of economic aggregates changes considerably after 1984. The standard deviation of output declines by almost half over the period 1984-2008, as do the standard deviations of consumption and investment, while not quite matching the fall in output volatility. On the whole, the period 1984-2008 has been described as a Great Moderation relative to the period that preceded it. This reduction in volatility dissipates after 2008, with the onset of the Great Recession.

Table 1 further indicates that the Great Moderation period is characterized by significant changes in the comovement patterns compared with the pre-1984 period. The procyclicality of labor productivity all but vanishes. Its correlation with output after 1984 falls to 0.06. At the same time, labor productivity switches from comoving positively with hours before 1984 to comoving negatively with hours in the post-1984 era. This fact holds irrespective of the definition of labor input and also holds for employment. Both these changes directly challenge the notion of a business cycle driven primarily by Hicks-neutral productivity shocks in a neoclassical environment. Moreover, Table 1 shows that the volatility of labor input relative to that of output increases by more than 50 percent after 1984, which suggests changes in frictions governing labor markets and their associated wedge. Finally, the countercyclicality of the inventory-sales ratio also vanishes in the post-1984 period. At the same time, the volatility of inventories relative to output increases by more than 50 percent.[8]

The shifts in relative volatilities and comovement patterns documented in Table 1 hold throughout the Great Recession, suggesting that this last recession was typical of post-1984 business cycles, albeit on a much larger scale. To address these changes, the next section sets up an environment with multiple sectors and multiple stages of production that features inventories while, at the same time, nesting the rich production structure of detailed RBC models first introduced by Kydland and Prescott (1982) or Long and Plosser (1982). This detailed production structure allows for flexibility in the dimensions along which changes in efficiency wedges might account for business cycles.

3 A Frictionless Framework with Multi-Sector and Multi-Stage Production

The business cycle accounting approach proposed by Chari, Kehoe, and McGrattan (2007) can be thought of as answering the following question: For a given description of preferences and technology, in what ways do allocations depart from optimality? These "wedges" represent manifestations of frictions that might include nominal rigidities, government regulations that distort labor markets or hinder production possibilities, or commitment problems that constrain the allocation of credit and financial intermediation. Given the post-1984 shifts in comovement and relative volatilities of economic aggregates depicted in Table 1, it is natural to suppose that the relative importance of those frictions might have changed over time.

This approach is more informative with respect to frictions the more realistic the assumptions regarding preferences and technology. Therefore, to add an important layer of realism to the baseline model, we adopt a multi-sector and multi-stage production structure.[9] One distinct contribution of our paper is our description of wedges in an economy with multiple sectors in that sectoral heterogeneity has important effects on the behavior of aggregate wedges. For our purposes, distinguishing between durable and non-durable production turns out to be crucial, as it allows for the possibility of fluctuations in the return to

[8] This increase in relative volatility is distinct from the well-documented secular decline in the inventory-sales ratio in the post-war period.

[9] Other dimensions, such as firm-level heterogeneity, can also significantly affect how credit friction translate into aggregate wedges, as underscored by Buera and Moll (2012). Under certain special assumptions, the positive relationship between credit shocks and the investment wedge may even vanish. In Section 6.3, we verify that our measured investment wedge in fact correlates well with the different measures of credit frictions highlighted in Section 5.

fixed investment that are not well captured by changes in aggregate total factor productivity (TFP).

We interpret the stage-of-production technology broadly. Firms cannot increase production overnight. Goods delivered on a given date are the result of a lengthy process starting from planning and design decisions, then involving the coordination of suppliers, the physical production process itself, and the distribution to final retailers and consumers. Throughout the later stages of the process, materials, works in progress, and finished goods sit in trucks or in storage as they are transferred between different suppliers, wait for complementary parts to arrive, or for the final consumer to enter the store. A conservative measure of the overall duration of this process is given by the ratio between inventories and the production of finished goods. In the U.S., this measure stands at about two months worth of final sales held in inventories. This number is conservative in that it accounts only for the actual production and distribution stages, but not the early design and planning stages. It also varies substantially between sectors.

3.1 Economic Environment

Consider an economy with N distinct sectors of production. At date t, each sector transforms inputs produced in previous dates, $Z_{j,t-s|t}$, into sales, $Y_{j,t}$, of a good j. The subscript $t-s|t$ refers to an input used at date $t-s$ in order to generate sales at date t. Thus, s is the time elapsed between the purchase of the input and the sale of the good. The technology necessary for distribution and sales requires combining inputs produced at various previous dates, $Z_{j,t-s|t}$:

$$Y_{j,t} = \left(B_j \sum_{s=0}^{S} \omega_j(s)^{\frac{1}{\varrho}} Z_{j,t-s|t}^{\frac{\varrho-1}{\varrho}} \right)^{\frac{\varrho}{\varrho-1}}, \;\; \varrho > 0, \tag{1}$$

where $s = 0, ..., S$ and B_j is a normalizing constant.

As an example, consider the case $S = 2$, where j is an aircraft, to be used by a final customer at $t+2$. This requires the use of inputs at date t , $t+1$, and $t+2$, aggregated in, respectively, $Z_{j,t|t+2}$, $Z_{j,t+1|t+2}$, and $Z_{j,t+2|t+2}$. The first stage, taking place at t, necessitates acquiring varieties of raw materials from different suppliers such as aluminum and other metals and alloys, and moving them to warehouses for storage until different varieties are available to be combined. The combination of inputs used at date t defines $Z_{j,t|t+2}$. In the following stage, at date $t+1$, $Z_{j,t|t+2}$ is cut and formed into components, such as the fuselage, wings, or vertical tail, that will make up the final aircraft using labor and specialized machinery. This labor and machinery constitute $Z_{j,t+1|t+2}$. In the final stage, at $t+2$, the different components of the aircraft are assembled into a finished product that is then extensively tested to gain certification before being flown to final customers. The inputs involved in the testing, distribution, and delivery make up $Z_{j,t+2|t+2}$.[10]

The parameter ϱ in the aggregator function (1) governs the substitutability of the different stages of production. At one extreme, as $\varrho \to 0$, there is no substitution. Raw materials have to be available one full period before they can be cut and formed into components which, in turn, take another full period to be assembled, tested, and delivered to the final consumers. This assumption reproduces the time-to-build technology of Kydland and Prescott (1982). At the other extreme, with $\varrho \to \infty$, it is possible to compensate fully for smaller amounts of materials acquired at t with a more elaborate assembly process at $t+1$. A storage technology is a special case of $\varrho \to \infty$ where the same input mix is used at each stage of production. That is, in order to have an aircraft available at $t+2$, production can either take place at t where the aircraft is then stored until $t+2$ (with a proportional depreciation cost governed by $\omega_j(s)$), or take place at $t+2$ with the aircraft sold right away.

In the general case, where $0 < \varrho < \infty$, the technology allows for limited substitution between stages of production. Because the extent of substitution is limited, it is seldom optimal to reduce production to a single stage despite the opportunity cost of time. Consequently, there always exists a positive steady-state target for the inventory-sales ratio. In that sense, our technology also captures a set up, common in the inventory literature, that combines a linear storage technology with a cost of deviating from some target inventory-sales ratio.

The production of inputs available in a given sector, $Z_{j,t|t+s}$, takes place using a Cobb-Douglas technology that combines capital, $K_{j,t|t+s}$ (warehouses, moving trucks, or machinery), labor, $L_{j,t|t+s}$ (including distribution and sales), and materials produced in other sectors, $M_{ij,t+s}$:

$$Z_{j,t|t+s} = K_{j,t|t+s}^{\alpha_j} \left(\prod_{i=1}^{N} M_{ij,t|t+s}^{\gamma_{ij}} \right) \left(A_{j,t} L_{j,t|t+s} \right)^{1-\alpha_j - \Sigma_{i=1}^{N} \gamma_{ij}}. \tag{2}$$

As emphasized in Long and Plosser (1982), the fact that each sector uses materials from other sectors represents a source of dynamic linkages in the model. The ij subscript in $M_{ij,t|t+s}$ denotes materials from sector i used in the production of sector j goods.

In equation (2), $A_{j,t}$ captures changes in the efficiency with which labor contributes to

[10] As we discuss in detail below, the cost of inputs included in $Z_{t|t+2}$ and $Z_{t+1|t+2}$ are added together into inventories. Thus, we do not strictly differentiate between raw materials, work-in-process, and finished goods. Similarly to Ramey (1989), we treat all inventories as necessary to a broader production process. For a more detailed treatment of different types of inventories see Humphreys, Maccini, and Schuh (2001).

production. These changes are allowed to operate along three dimensions,

$$A_{j,t} = u_t A_t a_{j,t}, \tag{3}$$

where u_t and $a_{j,t}$ have unit unconditional means. A_t is a common component in all sectors and has permanent effects on Hicks-neutral total factor productivity (TFP). As such, the growth rate, $g_t = A_t/A_{t-1}$, captures long-run labor-augmenting technological progress that varies over time around a constant mean $g > 1$. Second, a common disturbance, u_t, accounts for temporary changes in the level of efficiency with which labor is used in production. Finally, some of the changes in production efficiency in a given sector reflect considerations idiosyncratic to that sector, denoted by a_{jt}.

In each sector, j, the capital stock evolves according the law of motion

$$K_{j,t+1} = X_{j,t} + (1-\delta)K_{j,t}, \tag{4}$$

where $X_{j,t}$ denotes investment in sector j and δ is the depreciation rate. Investment in each sector j is produced using the amount $I_{ij,t}$ of sector i output by way of a constant returns to scale technology,

$$X_{j,t} = \Xi_j \prod_{i=1}^{N} I_{ij,t}^{\theta_{ij}}, \ \sum_{i=1}^{N} \theta_{ij} = 1, \tag{5}$$

where Ξ_j is a constant. As in models with investment adjustment costs, equation (5) allows for a non-constant rate of transformation between consumption and investment. Furthermore, just as in Kydland and Prescott (1982), multiple stages of production imply a constraint on the speed with which sales of investment goods can adjust to disturbances.

A representative household derives utility from the consumption of goods produced in every sector according to

$$E_t \sum_{t=0}^{\infty} \left(\beta^t \prod_{v=0}^{t-1} \zeta_v \right) \left[\kappa \ln C_t + (1-\kappa) \ln(1 - \Upsilon_t L_t) \right], \tag{6}$$

where C_t denotes the consumption aggregate

$$C_t = \prod_{i=1}^{N} \Lambda C_{j,t}^{\eta_j}, \ \sum_{j=1}^{N} \eta_j = 1, \tag{7}$$

while L_t represents aggregate labor input and thus must satisfy

$$L_t = \sum_{j=1}^{N} L_{j,t}. \tag{8}$$

The η_j's are the shares of consumption of good j in aggregate consumption. The variable ζ_v in equation (6) represents a random disturbance that changes the discount rate used to evaluate investment projects. The term Υ_t shifts the supply of labor and enters the definition of a modified labor wedge that, under special assumptions, reduces to that measured in previous work. Moreover, each sector is subject to a resource constraint,

$$C_{j,t} + \sum_{i=1}^{N} I_{ji,t} + \sum_{i=1}^{N} M_{ji,t} = Y_{j,t}. \tag{9}$$

Our environment simplifies to the prototype growth model of King, Plosser, and Rebelo (1988) when production does not use dated inputs, $\omega(s) = 0$ for $s > 0$, interlinkages in materials or investment are unimportant, $\gamma_{ij} = \theta_{ij} = 0$, and the economy contains one sector only, $j = 1$. In that case, aggregate output is produced via an aggregate production function described by a Cobb-Douglas technology that uses only contemporaneous inputs.

3.2 National Accounting and Relationships to Data

We now outline how the model maps into data by way of corporate and national accounting identities. A key feature of the environment described in the previous section pertains to the implications of our multi-stage production structure for the behavior of inventories. While we analyze our model using the solution to a planner's problem, a decentralized version of the economic environment would feature relative prices defined by the ratio of Lagrange multipliers. Therefore, when measuring variables at constant prices, we rely on these multipliers evaluated at the steady state. We choose the normalizing constants B_j and Ξ_j in the production technology, and Λ in preferences, such that relative prices are equal to one in the steady state.

Given that the total value of goods produced in the economy, Z_t, is the sum of output across all sectors and all stages of production with future completion dates t to $t + S$, we have

$$Z_t = \sum_{j=1}^{N} \sum_{s=0}^{S} Z_{j,t|t+s}. \tag{10}$$

Aggregate value added, or gross domestic product (GDP), V_t, is the total value of production less the value of intermediate inputs, $V_t = Z_t - \sum_{i=1}^N \sum_{j=1}^N M_{ij,t}$. Sales in sector j are given by $Y_{j,t}$, so that aggregate sales are $\mathcal{S}_t = \sum_{j=1}^N Y_{j,t}$,while sales of final goods, or final sales, $\mathcal{FS}_t$, amount to total sales less the value of materials $\mathcal{FS}_t = \sum_{j=1}^N Y_{j,t} - \sum_{i=1}^N \sum_{j=1}^N M_{ij,t}$. Let N_t denote the stock of inventories at date t, and let ΔN_t be inventory investment. The NIPA definition of inventory investment, namely gross domestic product less final sales, is then given by

$$\Delta N_t = V_t - \mathcal{FS}_t. \tag{11}$$

Given the technology in (1), sales in period t rely on inputs produced in previous periods. In particular, at date t, the corporate accounting definition of the cost of goods sold, denoted $\mathcal{C}_t$, reflects the cost of inputs used between dates $t - S$ and t,

$$\mathcal{C}_t = \sum_{j=1}^{N} \sum_{s=0}^{S} Z_{j,t-s|s}. \tag{12}$$

Since our technology satisfies constant returns to scale, the cost of goods sold in our model is simply the value of sales, $\mathcal{C}_t = \mathcal{S}_t$.[11] Adding and subtracting materials to the right-hand-side of equation (11), inventory investment can then be expressed as $\Delta N_t = Z_t - \mathcal{C}_t$. Substituting in the definitions of Z_t and $\mathcal{C}_t$ shows that inventory investment is closely tied to the multi-stage nature of production,

$$\Delta N_t = \sum_{j=1}^{N} \sum_{s=0}^{S} Z_{j,t|t+s} - \sum_{j=1}^{N} \sum_{s=0}^{S} Z_{j,t-s|s}. \tag{13}$$

By iterating equation (13) forward, we can recover the stock of inventories. The equations described in this section relate key variables in the model to NIPA measures. With this mapping in hand, we now turn to the identification of wedges implied by our model, and how those compare with previous work.

[11] In theory, these might differ slightly because conventional corporate accounting does not take into account the opportunity cost of committing resources to production for a good to be sold at a later date reflected in the interest rate. In practice, this distinction is quantitatively small.

4 Wedge Accounting with Multiple Production Stages and Sectors

We now define the wedges that we use to account for the salient features of U.S. business cycles described in Table 1. We consider a labor wedge, an investment wedge, and Hicks-neutral efficiency wedges that vary across sectors. Moreover, we detail how our approach nests previous work on the measurement of wedges. In our model specification, we assume for simplicity that factor shares are identical across stages of production. This is largely motivated by the fact that we do not have data on factor intensity for each stage. However, it also simplifies the wedge derivation since it allows composite goods associated with each production stage to be aggregated into a general output, $Z_{j,t}$:

$$Z_{j,t} = \sum_{s=0}^{S} Z_{j,t|t+s}. \tag{14}$$

Furthermore, $Z_{j,t}$ can be expressed as a function of aggregate inputs, $K_{j,t} = \sum_{s=0}^{S} K_{j,t|t+s}$, $M_{ij,t} = \sum_{s=0}^{S} M_{ij,t|t+s}$, and $L_{j,t} = \sum_{s=0}^{S} L_{j,t|t+s}$, by way of the Cobb-Douglas production function (2).

The proof for this aggregation result, as well as all derivations underlying the results in this section and the sections below, can be found in online-only supplementary notes, Lubik, Sarte, and Schwartzman (2014). The notes also show how, in a decentralized setup, the preference shocks Υ_t and ζ_t can be interpreted in terms of taxes on labor and capital respectively.

4.1 The Labor Wedge

The labor wedge is defined as the deviation of the marginal product of labor from the marginal rate of substitution between consumption and leisure. This object is of interest since it measures the role of other forces affecting labor markets, given a correct specification of preferences and technology and a lack of exogenous shocks or shifts in the primitives. The existence of this wedge is typically motivated by distortions, such as labor taxes. As an example, given the preferences adopted in this paper, the labor wedge in the prototypical one-sector growth model can be expressed as:

$$\widetilde{\tau}_t^L = \underbrace{\left(\widehat{Z}_t - \widehat{L}_t\right)}_{\text{Labor Productivity}} - \left(\widehat{C}_t + \frac{L}{1-L}\widehat{L}_t\right), \tag{15}$$

where the caret represents percentage deviations from steady state, and variables without a "t" subscript are evaluated at the steady state.[12]

Let $\phi_{j,t}$ denote the ratio between the marginal cost of the inputs involved in the production of good j and its marginal value to a household consuming it. This ratio varies over time since the marginal rate of transformation of output into current sales depends on past production decisions. We can then derive the appropriate labor wedge in our framework with multiple sectors and multiple stages of production as:

$$\tau_t^L = \sum_{j=1}^{N} \underbrace{\eta_j\left(\widehat{Z}_{j,t} - \widehat{L}_{j,t}\right)}_{\text{Sectoral Labor Productivity}} - \left(\widehat{C}_t + \frac{L}{1-L}\widehat{L}_t\right) + \sum_{j=1}^{N}\eta_j\widehat{\phi}_{j,t}. \tag{16}$$

There are two key differences between equations (15) and (16). First, the relevant notion of a labor wedge involves the marginal contribution of labor to the generation of current and future sales, rather than just the marginal product of labor. This fact is captured by the term $\widehat{\phi}_{j,t}$ and relies on the assumption underlying equation (1), namely that, in the absence of distribution and sales, current output cannot be directly used for consumption or investment. Second, in a multi-sector economy, the marginal product of labor in a given sector affects aggregate labor supply decisions more the larger the share of that sector is in aggregate consumption. Hence, calculating the labor wedge as a function of aggregate consumption requires averaging labor productivity in different sectors by using their consumption shares, η_j, as weights. In the one-sector case with no production lags ($N = 1$ and $S = 0$), $\widehat{\phi}_{j,t} = 0$ and the labor wedge in (16) reduces to the conventional wedge (15).

We note that in our framework, the labor wedge is approximated by $\tau_t^L = \widehat{\Upsilon}_t/(1-L)$. Time-variation in the preference parameter Υ_t gives our model the necessary flexibility to generate a time-varying labor wedge. Alternatively, a labor tax in the decentralized version of the model would play the same role. We describe in Section 6 how the labor wedge defined in our model has behaved in the U.S. economy and how this behavior compares with the more standard labor wedge in the prototypical growth model.

[12]The formal derivation of this equation assumes a steady state value for $\Upsilon_t = 1$.

4.2 The Investment Wedge

The investment wedge measures the deviation between households' intertemporal rate of substitution and the physical return to investment. This wedge is typically associated with distortions to credit markets that arise from informational or limited commitment problems, or more simply taxes on investment. Given a one-sector model with logarithmic preferences and no production lags, the investment wedge is

$$\widetilde{\tau}_t^X = \underbrace{E_t\left[\left(1-\widetilde{\beta}\right)\left(\widehat{V}_{t+1}-\widehat{K}_{t+1}\right)\right]}_{\text{Marginal Return to Investment}} - E_t\left(\Delta\widehat{C}_{t+1}\right), \tag{17}$$

where $\Delta\widehat{C}_{t+1} = \widehat{C}_{t+1} - \widehat{C}_t$ and $\widetilde{\beta} = \beta(1-\delta)$.

The investment wedge takes a more complicated form in the multi-sector, multi-stage production setup,

$$\begin{aligned} \tau_t^X \quad = \quad & \underbrace{E_t\left(1-\widetilde{\beta}\right)\sum_{j=1}^{N}\eta_j\left(\widehat{Z}_{j,t+1}-\widehat{K}_{j,t+1}+\widehat{\phi}_{j,t+1}\right)-\sum_{j=1}^{N}\sum_{i=1}^{N}\left(\eta_j\theta_{ij}-\eta_j\eta_i\right)\left(\widehat{\lambda}_{i,t}-\widetilde{\beta}\widehat{\lambda}_{i,t+1}\right)}_{\text{Marginal Return to Investment}} \\ & -E_t\left(\Delta\widehat{C}_{t+1}\right), \end{aligned} \tag{18}$$

where $\widehat{\lambda}_{j,t}$ is the Lagrange multiplier associated with the resource constraints described in equation (9).

There are three main differences between the conventional wedge (17) and the one in equation (18). The first two differences are analogous to those for the labor wedge. They are related to the distinction between output and sales and the calculation of the marginal product of capital in a multi-sector model. These differences are captured by the terms $\widehat{\phi}_{j,t}$ and the consumption weights η_j. In addition, different sectors contribute differently to investment in other sectors given the technology described by equation (5). The marginal return on investment therefore incorporates changes in the cost of producing investment goods relative to consumption in the current and following period. This is reflected by the θ_{ij}'s in the last term of the marginal return to investment. It is also the case in our model that $\tau_t^X = -\widehat{\zeta}_t$, so that shifts in the households' discount factor can be thought of as a stand-in for a tax on investment. Observe that the investment wedge in the generalized model cannot be readily identified from aggregate consumption, investment, and output data because of the variable, $\widehat{\phi}_{j,t+1}$, in the marginal return to investment.

Equations (17) and (18) give the investment wedge in terms of the Euler equation governing fixed investment. However, to the degree that the wedge captures distortions in credit markets, it will also appear in the Euler equation for inventory investment. It is in this sense that inventory data are informative with respect to the investment wedge. Given a process for τ_t^X, optimal inventory investment in sector j requires that

$$\widehat{\phi}_{j,t} + \frac{1}{\varrho}\left(\frac{\Delta N_{j,t}}{Y_{j,t}}\right) = E_t\left(\sum_{s=1}^{S}\psi_j(s)\left[\Delta\widehat{\lambda}_{j,t+s} + \frac{1}{\varrho}\Delta\widehat{Y}_{j,t+s} - \left(\sum_{u=0}^{s-1}\tau_{t+u}^X\right)\right]\right), \tag{19}$$

where $\psi_j(s)$ is the steady state ratio of inputs dedicated to production s periods hence to current production, $\psi_j(s) = Z_{t|t+s}/Z_t = Z_s/Z$. The ratio $\Delta N_{j,t}/Y_{j,t}$ is that of inventory investment to sales in sector j.

The left-hand side shows the marginal cost of increasing inventories. For given sales, inventory investment requires raising output, with associated marginal cost given by $\widehat{\phi}_{j,t}$. The marginal inventory investment necessary to generate future sales increases as total investment rises, a fact captured by the term $\frac{1}{\varrho}\left(\frac{\Delta N_{j,t}}{Y_{j,t}}\right)$. The right-hand side of (19) summarizes the marginal benefits of inventory investment. These benefits are weighted by the share of currently accumulated inventories that is dedicated to production in each period, $\psi_j(s)$. Inventory investment helps increase sales in future periods relative to the current period, an effect that on the margin is valued relative to current sales at $\Delta\widehat{\lambda}_{j,t+s}$. Moreover, the marginal contribution of current inventory investment to future sales is larger when future sales are expected to be large, an effect reflected in the term $\frac{1}{\varrho}\Delta\widehat{Y}_{j,t+s}$. Finally, accumulated expected future investment wedges, $\sum_{u=0}^{s-1}\tau_{t+u}^X$, lower the benefits of increased inventory investment.

4.3 Efficiency Wedges

Efficiency wedges capture all factors that influence the efficiency with which inputs, that is, capital, labor and materials, are transformed into output. As such, they reflect changes in production possibilities associated with technological progress, as well as changes in taxes and regulations which distort the composition of intermediate inputs or the allocation of resources across firms. In a multi-sector model, efficiency wedges are defined separately for each sector,

$$\tau_{j,t}^A = \widehat{Z}_{j,t} - \alpha_j\widehat{K}_{j,t} - \xi_j\widehat{L}_{j,t} - \sum_i \gamma_{ij}\widehat{M}_{ij,t}. \tag{20}$$

In our environment, given the specification of production shocks (3), these wedges are a function of productivity components

$$\tau_{j,t}^{A} = \xi_j \left(\widehat{u}_t + \widehat{A}_t + \widehat{a}_{j,t} \right), \tag{21}$$

where $\widehat{u}_t$ is a transitory aggregate component, $\widehat{A}_t$ is a permanent component, and $\widehat{a}_{j,t}$ is a sector-specific component. The RBC literature has traditionally placed great emphasis on efficiency wedges as drivers of business cycles, albeit typically in the form of a single, transitory aggregate Hicks-neutral component. In specifying a more flexible form of efficiency wedges, combined with a production structure that generalizes technological trade offs, we position the model to explain shifts in comovement and relative volatilities over the post-war period with productivity shocks alone.

5 Calibration and Estimation

Our empirical approach combines elements of calibration and likelihood-based estimation. We choose to calibrate a large number of the preference and technology parameters since information on their values is readily available and there is substantial consensus in the literature. Parameters that are not easily calibrated are estimated using a Bayesian approach. Our empirical analysis remains firmly wedded to the calibration literature and thereby assures comparability of the conclusions. At the same time, we are able to make inferences about the structural shocks in a consistent manner.

Our calibration follows common practice in the business cycle literature. We set the per-capita steady-state growth rate equal to 2 percent, the depreciation rate of fixed capital to 10 percent, and the discount factor β to 0.96, all on an annual basis. The parameter κ governing the share of leisure in utility is chosen so that in steady state leisure is approximately two-thirds of the available time endowment. We rely on the input-output tables and the capital flow tables to parameterize the contribution of different sectors to inputs and capital formation. In line with the discussion on National Accounting in Section 3.2, we follow Schwartzman (2012) and use data on the inventory-sales ratio to calibrate the weights associated with the different stages of production, $\omega_j(s)$. We assume that these weights decline exponentially with s and choose a rate of decline $\frac{\omega_j(s+1)}{\omega_j(s)}$ that matches the average inventory-sales ratio in each sector over the entire sample period. The rate of decline is slower in the durable goods sector, since this sector can make greater use of inventories.

With the exception of cross-sectoral weights, all calibrated parameters are stated in Table 2.[13]

We then estimate the remaining parameters using a Bayesian approach, specifically the parameters of the unobservable exogenous shock processes and the elasticity of substitution between stages of production, ϱ. The latter cannot be easily calibrated as there is virtually no prior information. The structural estimate of ϱ will therefore provide independent insight into the nature of the intersectoral linkages and the production processes in the U.S. economy. Furthermore, our estimation allows us to characterize the time paths for the shocks and set them in relation to the wedges as discussed above. We assume that all shocks follow AR(1) processes and we allow for correlation in their innovations.

The benchmark specification consists of two sectors and four stages of production, as in Iacoviello, Schiantarelli, and Schuh (2011).[14] One sector corresponds to durables, the other to nondurables. As described in sections 3.1 and 4.3, we allow for five shocks: three shocks to efficiency wedges, one shock capturing changes in the labor wedge, and a shock to capture changes in the investment wedge. The unobservable shocks are jointly identified off the corresponding equilibrium conditions. The exogenous forces that drive measured TFP therefore consist of: (i) a transitory disturbance that augments labor by the same proportion in both sectors; (ii) a permanent disturbance that augments labor proportionally in both sectors; and (iii) a sectoral shock that only affects labor productivity in the durable sector. This specification for TFP allows for considerably more flexibility than in previous work that explores the role of efficiency wedges. Including both transitory and permanent features to the process governing TFP addresses the observation by Gorodnichenko and Ng (2010) that constraining the model to only one component can materially affect estimation results.

We make use of four aggregate series in the estimation: output, consumption, hours, and inventories. Output is the sum of consumption, fixed investment, and changes in inventories. Our definition of output, therefore, excludes government consumption and net exports. Separate data on fixed investment is redundant since it is spanned by the other time series. All variables are in terms of the working age non-institutionalized population, converted to natural logs, and enter in first differences. With the exception of hours data, all data are extracted from the Haver Analytics database. Series for output, consumption, and inven-

[13] A table with all the calibrated and estimated parameters is included in the technical appendix

[14] While there are no conceptual difficulties to including more sectors or stages of production, estimation of the model becomes computationally expensive as these dimensions increase. Specifically, the number of equations to be solved increases linearly with the number of sectors, but with the square of the number of production stages.

tories are constructed using NIPA data produced by the Bureau of Economic Analysis. All variables are measured in real terms, and we aggregate price deflators using Divisia indices whenever appropriate. Inventories are taken directly from the NIPA data. Since we do not model durable consumption explicitly, consumption only includes nondurable goods and services. We also add to consumption the services produced by durable goods, calculated using the formula in the technical appendix of Chari, Kehoe, and McGrattan (2007). We treat durable consumption as part of fixed investment and, thus, part of output. Finally, hours data were calculated by the Bureau of Labor Statistics following the procedure in Francis and Ramey (2005). Our data series are quarterly and cover the period from 1953:Q3 to 2011:Q4.

We estimate the structural model using the Bayesian approach detailed in An and Schorfheide (2007). This requires choosing a prior distribution for the parameters to be estimated, which is then combined with the likelihood function for the linear rational expectations model laid out above. The joint density is then evaluated using the Kalman-filter and the posterior distribution is evaluated using a random-walk Metropolis-Hastings algorithm for 500,000 draws. We assume that the prior distribution of the substitution elasticity ϱ is Gamma with a mean of 4 and a high variance. This assumes a reasonable amount of substitution between output produced in different periods, as would be the case in typical models of inventory investment. As for the shock processes, we adopt a Beta prior for the persistence parameters (denoted $\rho_u, \rho_g, \rho_{a_D}, \rho_\Upsilon$ and ρ_ζ with subscripts representing the corresponding wedge) with a mean of 0.9, except for the growth rate of the permanent component of TFP, and standard deviation 0.5. For the process governing the growth rate of the permanent component of TFP, we choose a lower prior mean of 0.5, and a larger standard deviation of 0.2.

We report the estimation results in Table 2. The posterior mean of ϱ is 18.9, which indicates that the data are informative with respect to this parameter. This implies a high degree of substitutability between stages of production. The estimated production function thus approximates a storage technology, while keeping a long-run target level for the inventory-sales ratio. In other words, the technology reproduces features similar to "stock-out avoidance" models of inventory behavior that combine a linear storage technology with a cost of deviating from a long-run target. The estimated shock processes are highly persistent with auto-correlations approaching unity, the exception being the shock process that captures the permanent component of TFP with a posterior mean for the autocorrelation coefficient of 0.66.

6 Empirical Findings

We now present the main contribution of the paper, that is, the role that different wedges play in explaining post-war U.S. business cycles. In the spirit of the early RBC literature, we first attempt to gauge how well efficiency wedges alone can account for the changing nature of business cycles. Our key finding is that efficiency wedges can simultaneously explain various comovement properties and relative volatilities of economic aggregates prior to 1984, but they cannot account for shifts in these properties after 1984 even with the flexible technological framework adopted here. Although they still explain a large percentage of output variations after 1984, they are unable to account for the bulk of the variation in hours after that date and, therefore, changes in the behavior of labor productivity. We compare our benchmark findings with those that emerge in the more restricted canonical one-sector growth model. We also contrast our findings with those from a variant of our multi-stage multi-sector model that abstracts from inventories. Our analysis ultimately leads us to focus on role of the investment wedge, rather than the labor wedge emphasized in more conventional frictionless models.

6.1 Impulse Response Functions

We report impulse response functions of key variables to the structural shocks in Figures 1 and 2. The responses are computed at the posterior means and normalized such that the average response of GDP over 32 quarters is positive and equal to 1. Figure 1 depict the responses of output, hours, and labor productivity to disturbances affecting the different wedges. The panels on the left show the responses of these variables to efficiency wedge shocks. A positive transitory aggregate shock and positive sectoral shock have similar implications, although the idiosyncratic sectoral shock mutes the positive effect on labor productivity. In particular, they imply a distinctly positive comovement of output, hours, and labor productivity. A permanent increase in aggregate TFP stands out by generating a slight hump-shaped response in output and labor productivity. It also generates a decline in hours on impact through a wealth effect. This implies negative comovement between labor productivity and hours as well as output and hours over short horizons, which reproduces the findings by Sims (2011).

The right-hand side panels show the response of the variables of interest to shocks that reduce the labor and investment wedge. Both shocks have similar implications, but differ from productivity shocks in two ways. First, for a given response of output, there is a larger response in labor input since these shocks do not directly affect productivity. Second, they induce labor productivity to comove negatively with both hours and output. Declines in both the labor and investment wedge affect output by altering the amount of labor supplied for a given marginal product of labor. Conditional on a decrease in the labor wedge, the intratemporal return to work rises; in the case of a decrease in the investment wedge, current leisure becomes less valuable relative to the future. Therefore, compared with a productivity shock, unanticipated decreases in the labor or investment wedge induce a larger movement in hours given the change in output. This suggests that in order to explain the increase in the relative volatility of hours after 1984 together with he reduction in the cyclicality of labor productivity, shocks to either wedge must have become more prevalent relative to efficiency wedges.

We now turn to the responses of inventories and the inventory-sales ratio relative to output in Figure 2. Two observations stand out. First, shocks that increase aggregate TFP or reduce the labor wedge generate a countercyclical inventory-sales ratio. These shocks promote a boom in fixed investment and thereby increase sales and the opportunity cost of holding inventories, so that fewer inventories are held. The opposite is true of shocks to the investment wedge; a shock that lowers the investment wedge and increases output also increases inventories and the inventory-sales ratio. Second, the investment wedge has a direct first-order effect on the Euler equation governing inventory investment (19). It therefore induces a larger response of inventories than do other shocks, including shocks to the labor wedge. The distinctively different response of inventories to variations in the investment wedge relative to other wedges is what allows inventory data to identify changes in the investment wedge separately from changes in the labor wedge.

The responses of inventories and the inventory-sales ratio to shocks reducing the labor wedge are very similar to those induced by a transitory increase in aggregate productivity. The intuition is that similar to a temporary rise in aggregate productivity, a decrease in the labor wedge affects inventory investment decisions by changing the relative cost of producing in different periods. However, these shocks are fairly persistent in our estimation. It follows that the change in relative cost is small and does not generate a strong motive to accumulate inventories. In contrast, the investment wedge shock has a fundamentally different effect, that alters the opportunity cost of holding inventories and, therefore, has more direct and substantive implications for inventory accumulation. It follows that while both the labor and investment wedge can potentially account for the joint behavior of output, hours, and labor productivity after 1984, the investment wedge is uniquely positioned to account for the change in the relative volatility and cyclicality of both inventories and the inventory-sales ratio after that date.

6.2 Wedges and Business Cycles

Table 3 shows the cumulative effects of our estimated wedges on moments of the data of interest both prior to and after 1984. Reading the table from left to right describes how comovement properties and relative volatilities of the data in Table 1 change as each additional wedge is taken into account. The first three columns of Table 3 contain the successive effects of the efficiency wedges.

Considering the transitory aggregate component alone, we find that temporary productivity shocks induce labor productivity movements that are strongly correlated with both output, at 0.99, and hours worked, at around 0.95. These disturbances also produce an inventory-sales ratio that comoves negatively with output as in the pre-1984 data. As indicated by the impulse response function depicted in Figure 2, this finding is consistent with the negative overall effect of temporary TFP shocks on the inventory-sales ratio. In addition, temporary changes in TFP produce an inventory series whose standard deviation relative to output closely matches the data at around 0.79. In summary, the benchmark model driven by our estimated aggregate transitory efficiency wedge is able to simultaneously explain all key comovement properties and relative volatilities of the data in Table 1 before 1984. For the same reason, this wedge loses its explanatory power when the comovement properties and relative volatilities of the data shift dramatically after that date.

The second column of Table 3 shows the effects of shocks that affect the permanent component of the aggregate efficiency wedge on moments of the data. This aspect of efficiency wedges is emphasized by Sims (2011) as an important component of business cycles. Permanent changes to productivity generate a fall in labor supply initially through a wealth effect, which potentially reduces the positive comovement between labor productivity, hours, and output. We find, however, that the reduction in comovement is not enough to account for the changing nature of the business cycle. The third column of Table 3 shows how these moments change with the addition of the sectoral efficiency wedge, which we associate with shifts in production possibilities of the durable sector. The sectoral efficiency wedge lowers the positive comovement between labor productivity and hours worked further.[15] Ultimately, however, the labor productivity puzzle remains as labor productivity continues to be strongly procyclical both before and after 1984.

More generally, Table 3 demonstrates that efficiency wedges, while capable of accounting for the nature of economic fluctuations prior to 1984, are not enough to explain salient features of U.S. business cycles over the last three decades. Aside from comovement patterns, we note that efficiency wedges explain a remarkable 98 percent of output variations and 67 percent of the variation in hours before 1984, but account for 80 percent of fluctuations in output, and only 37 percent in hours worked, after that date.

This observation has led researchers to focus on the labor wedge, which is in part motivated by the growing emphasis on nominal rigidities in dynamic stochastic general equilibrium models. Thus, the fourth column of Table 3 shows how the moments of the data change with the addition of our estimated labor wedge. Comovement of labor productivity with output does decrease considerably but, in contrast to the data, ultimately remains distinctly procyclical in both the pre-1984 and post-1984 samples. Similarly, labor productivity comoves much less with hours, especially after 1984, but not to the point of resolving the labor productivity puzzle after 1984. As we saw in the impulse responses in Figure 1, shocks driving the labor wedge generate opposite responses of labor productivity and hours and thus reduce the procyclicality of the former. However, in moving the comovement properties of labor productivity closer to the data, the labor wedge simultaneously moves the moments associated with inventory behavior in a counterfactual direction. After 1984, the inventory-sales ratio becomes even more countercyclical, while the standard deviation of inventories falls instead of increasing markedly as in the data. Figure 2 illustrates that an increase in the labor wedge raises output but lowers the inventory-sales ratio, which is opposite of what is needed to explain the shift in inventory behavior after 1984.

The last column of Table 3 considers the effects of the investment wedge. It plays a significant role in lowering the procyclicality of labor productivity after 1984. The correlation of labor productivity and output goes from 0.64 in the post-84 period to 0.13 when the investment wedge is taken into account. In the post-84 period, the investment wedge also moves the correlation between labor productivity and hours from positive into negative territory. Figure 1 shows that, as with the labor wedge, an increase in the investment wedge

[15] This insight relates to Garin, Pries, and Sims (2013), who emphasize the importance of reallocative shocks in explaining changes in the nature of economic fluctuations over the last three decades.

moves labor productivity and hours in opposite directions. It is also notable that the investment wedge changes the cyclicality of the inventory-sales ratio from being countercyclical to acyclical after 1984. From the impulse response functions we see that a shock that raises the investment wedge also induces an increase in both output and the inventory-sales ratio, thus contributing to a reduction in the countercyclicality of the inventory-sales ratio. The key finding is that the investment wedge plays a distinctly important role in explaining shifts in the nature of business cycles post-1984.

Figure 3 shows the paths of output, hours, and inventories following the 1980, 1991, and 2007 recessions when driven by efficiency wedges, the labor wedge, and the investment wedge separately. Shocks to efficiency wedges alone can explain virtually all of GDP movements following the 1980 recession. However, they are not enough to explain output variations in the 1991 recession and the latter phases of the Great Recession. Moreover, whereas shocks to efficiency wedges explain a fair share of movements in hours and inventories during the 1980 recession, they can only explain a much smaller portion of those movements in the 1991 and 2007 recessions. The key difference between the 1991 and 2007 recessions lies in the relative roles of the labor and investment wedge. While during the 1991 recession, the investment wedge plays an important role in delaying the recovery, especially with respect to inventories, the labor wedge is more prominent overall, as it explains virtually all of the fluctuation in hours. In contrast, during the 2007 recession, the investment wedge captures not only the largest part of output deviations from trend starting from mid-2009, but also a substantial fraction of hours deviations, as well as virtually all of the fall in inventories.

In order to highlight the importance of inventory data in identifying the investment wedge, Figure 4 shows the same historical decompositions but as obtained from estimating the canonical one-sector growth model without stages of production (i.e. with $N = 1$ and $S = 0$) while ignoring the inventory data. Findings from the 1980 recession share key features with those of Chari, Kehoe, and McGrattan (2007) despite a different calibration and a longer sample period that includes the Great Recession. In addition, we allow for permanent and transitory disturbances to TFP. In the prototypical one-sector growth framework, most of the decline in output is explained by shocks to TFP in all recessions, as is a significant portion of the decline in hours. Interestingly, the canonical decomposition of aggregates during the 1980 recession is similar to that of our benchmark model, with shocks to TFP explaining virtually all of the fall in output and shocks to the labor wedge playing a key role in explaining the decline in hours. However, the decomposition of macroeconomic aggregates from the prototypical framework is notably different from that of our benchmark for the 1991 and 2007 recessions. In the latter recessions, our benchmark model assigns a lesser role to efficiency and labor wedges and a distinctively greater role to the investment wedge. This difference between models is particularly evident during the Great Recession. In this last recession, the investment wedge switches from being a key driver of aggregate movements, in our benchmark case, to being essentially irrelevant in the one-sector growth model.

6.3 The Investment Wedge and Alternative Measures of Credit Market Frictions

Business cycle models with credit market frictions, such as Carlstrom and Fuerst (1997) and Bernanke, Gertler, and Gilchrist (1999), can in principle give rise to an investment wedge. This suggests that its behavior can be used to assess the relevance of credit market frictions. Subsequent research, however, has shown that the mapping from credit frictions to *an aggregate* investment wedge is not necessarily a feature of models with credit market imperfections (Buera and Moll, 2012). We now compare our generalized investment wedge with other available measures of credit market frictions. To the degree that the investment wedge captures in part the same information, this wedge then provides us with an additional indicator of financial frictions. The advantage of this alternative indicator is twofold. First, the investment wedge is less reliant on the details of financial market institutions relative to more conventional measures of credit market frictions. Second, and more importantly, we can use the investment wedge to infer the effects of those frictions on business cycles.

Table 4 shows the correlation between two measures of financing conditions and macroeconomic aggregates. Panel a) shows the correlations with Moody's index of bond yields for bonds rated "Baa" and Treasury Bonds, a measure that is used by Hall (2010) as an indicator of credit frictions, lagged four quarters.[16] A substantial literature documents that corporate bond spreads are highly correlated with future output, and recent empirical and theoretical work links those spreads to frictions affecting financial intermediaries rather than default probabilities alone.[17] Lagged values of the corporate bond spread reflect the notion, emphasized by Chari, Christiano, and Kehoe (2008), and Gilchrist and Zakrajsek (2011),

[16] Hall (2011) focuses on the spread between Baa rated bonds and 20-year Treasury bonds, whereas our data concerns the spread relative to 10-year Treasuries, since those are the longest maturity available throughout the whole sample period.

[17] For empirical evidence, see Gilchrist and Zakrajsek (2012) and Adrian, Moench, and Shin (2010a,b). For a theoretical treatment, see He and Krishnamurthy (2012, 2013)

that over the short run, firms tend to rely on unused revolving credit lines. Increases in bond spreads have their highest impact on the marginal cost of funds only after some time. Panel b) focuses on dividend payouts to business owners, as calculated by Jermann and Quadrini (2006, 2012). It depicts the correlations with the negative of those payouts, since higher payouts are associated with smaller credit market frictions.

As shown in Table 4, the correlation of indicators of financial frictions with different macroeconomic aggregates becomes more negative over time, reaching their most negative values in the Great Recession period. Furthermore, of the different aggregates, inventories are the most strongly correlated with the two measures of credit frictions. This validates our empirical strategy of using inventory data to infer the investment wedge, since it suggests that if shocks to credit conditions affect macroeconomic aggregates, the effect is likely to manifest itself prominently in inventory holdings.

Figure 5, panel a), depicts the investment wedge over time compared with the lagged Baa-Treasury bond spread. The investment wedge not only correlates closely with this spread but, remarkably, the two series move over the same range. The comovement between the two series is especially evident after 1984. The fact that this comovement increases over time is consistent with changes in financial markets that ultimately make bond spreads more informative measures of financial frictions over time. In particular, the period after 1984 has seen significant deepening of bond markets along with a much broader range of firm participation in those markets.

As a measure of robustness, Table 5, panel a), shows the business-cycle frequency correlation of the investment wedge with various alternative indicators of credit market conditions. With the exception of Debt Repurchases in the pre 1984 period, these business cycle correlations are always positive. Table 5, panel b), controls for the possibility that the different measures of credit market conditions share the same cyclical properties. It does so by depicting the partial correlations of the investment wedge with the various indicators of credit market conditions after extracting the common component that correlates with GDP. The resulting partial correlations are not only uniformly positive, but they are also generally more pronounced for the 1953-1984 period than the total correlations.

Figure 5, panel b), shows a 24-quarter rolling correlation of the investment wedge with key macroeconomic aggregates. The correlation of our investment wedge with both GDP and fixed investment has become more negative over time, especially after the mid-1990s. As we saw earlier, prior to 1984, output is mainly driven by productivity shocks which mutes its negative relation to the investment wedge. In addition, and similar to the more conventional indicators of credit frictions in Table 4, the investment wedge is more negatively correlated with inventories than with either GDP or fixed investment. Specifically, the correlation between the investment wedge and output goes from 0.29 pre-1984 to -0.38 between 1984 and 2008, and to -0.69 thereafter. In comparison, its correlation with inventories goes from an already large -0.71 pre-1984 to -0.84 between 1984 and 2008, and to -0.97 thereafter.

There are two instances where the correlation between the investment wedge and GDP becomes especially negative. The first occurs in the mid-to-late 1980s, coinciding with the Savings and Loans crisis that took place between 1986 and 1989. This fall-off is temporary and is played out as the crisis is resolved in the early 1990s. The second notable decline, more permanent, takes place in the mid to late 1990s and coincides with two large-scale regulatory reforms that affect both the spatial and functional segmentation of financial markets (the Riegle-Neal Act of 1994 and the Gramm-Leach-Bliley Act of 1999, respectively). Strikingly, the negative correlation between credit friction indicators and macroeconomic aggregates remains stable in the years approaching the 2008 recession. In that sense, the Great Recession episode distinguishes itself from the period immediately preceding it by its size rather than by changes in the relationship between credit market frictions and economic activity.

In summary, the aggregate investment wedge that emerges from our model informed by inventory behavior correlates well with a wide array of measures of credit market frictions. Moreover, the timing of changes in its correlation with GDP matches important events in financial markets and shifts in regulation. We take both of these observations as evidence that variations in the generalized investment wedge capture some measure of distortions imposed by changes in financial market conditions. None of these conclusions hold for the simpler investment wedge computed using the one-sector growth model, as suggested by the differences in behavior between these two wedges when comparing Figures 3 and 4.

7 Concluding Remarks

Until the early 1980s, conventional wisdom held that, in the words of Jr. (1977), "Business Cycles are all alike". This informed the development of real business cycle theories, which were highly successful in explaining business cycles along key dimensions. After 1984, however, business cycles changed. The volatility of output became smaller and the relative volatilities of labor input and inventories increased substantially, with labor productivity moving from pro-cyclical to counter-cyclical and the inventory-sales ratio from counter-cyclical to acyclical. We investigate which changes in the underlying mechanisms behind business fluctuations could have led to these changes. We find that the changing business cycle facts can be well explained by an increased amplifying role of financial factors in distorting intertemporal decisions.

References

Adrian, Tobias, Emmanuel Moench, and Hyun S. Shin. 2010a. "Financial Intermediation, Asset Prices, and Macroeconomic Dynamics." *Federal Reserve Bank of New York Staff Report* (422).

———. 2010b. "Macro Risk Premium and Intermediary Balance Sheet Quantities." *IMF Economic Review* 1 (58):179–207.

Benati, Luca and Thomas Lubik. 2014. "Sales, Inventories, and Real Interest Rates: A Century of Stylized Facts." *Forthcoming in the Journal of Applied Econometrics* .

Bernanke, Ben, Mark Gertler, and Simon Gilchrist. 1999. "The Financial Accelerator in a Quantitative Business Cycle Framework." In *Handbook of Macroeconomics*, vol. 1, edited by John Taylor and Michael Woodford, chap. 21. Elsevier.

Bils, Mark and James A. Kahn. 2000. "What Inventory Behavior Tells Us about Business Cycles." *The American Economic Review* 90 (3):458–481.

Buera, Francisco J. and Benjamin Moll. 2012. "Aggregate Implications of the Credit Crunch." *National Bureau of Economic Research* (17775).

Carlstrom, Charles and Timothy S. Fuerst. 1997. "Agency Costs, Net Worth and Business Fluctuations: A Computable General Equilibrium Analysis." *American Economic Review* 87 (5):893–910.

Chang, Yongsung, Andreas Hornstein, and Pierre-Daniel Sarte. 2009. "On the Employment Effects of Productivity Shocks: The Role of Inventories, Demand Elasticity and Sticky Prices." *Journal of Monetary Economics* 56 (3):328–343.

Chari, V. V., Lawrence Christiano, and Patrick Kehoe. 2008. "Facts and Myths about the Financial Crisis of 2008." *Federal Reserve Bank of Minneapolis Working Paper Series* (666).

Chari, V. V., Patrick Kehoe, and Ellen McGrattan. 2007. "Business Cycle Accounting." *Econometrica* 3 (75):781–836.

Christiano, Lawrence J. and Joshua M. Davis. 2006. "Two Flaws in Business Cycle Accounting." *National Bureau of Economic Research* (12647).

Christiano, Lawrence J., Roberto Motto, and Massimo Rostagno. 2010. “Financial Factors in Business Cycles.” *Mimeo* .

Dupor, Bill. 1999. “Aggregation and irrelevance in multi-sector models.” *Journal of Monetary Economics* 43 (2):391–409.

Fisher, Jonas D.M. 2006. “The Dynamic Effects of Neutral and Investment-Specific Technology Shocks.” *Journal of Political Economy* 114 (June):413–451.

Foerster, Andrew, Mark W. Watson, and Pierre-Daniel G. Sarte. 2011. “Sectoral vs. Aggregate Shocks: A Structural Factor Analysis of Industrial Production.” *Journal of Political Economy* 119 (1):1–38.

Francis, Neville and Valeria Ramey. 2005. “Is the Technology-Driven Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited.” *Journal of Monetary Economics* 52 (8):1379–1399.

Garin, Julio, Michael Pries, and Eric Sims. 2013. “Reallocation and the Changing Nature of Economic Fluctuations.” .

Gilchrist, Simon, Jae W. Sim, and Egon Zakrajsek. 2010. “Uncertainty, Financial Frictions, and Investment Dynamics.” .

Gilchrist, Simon and Egon Zakrajsek. 2011. “Credit Spread and Business Cycle Fluctuations.” In *Approaches to the Evolving World Economy. The Proceedings of the 16th World Congress of the International Economic Association Vol. III: The Global Macro Economy and Finance.*, edited by Franklin Allen, Masahiko Aoki, Nobuhiro Kyotaki, Richard Gordon, and Joseph Stiglizt.

———. 2012. “Credit Spread and Business Cycle Fluctuations.” *American Economic Review* 102 (4).

Gordon, Robert J. 2010. “Okun’s Law and Productivity Innovations.” *American Economic Review: Papers and Proceedings* 100:11–15.

Gorodnichenko, Yuriy and Serena Ng. 2010. “Estimation of DSGE Models when Data are Persistent.” *Journal of Monetary Economics* 57:325–340.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell. 2000. “The Role of Investment-Specific Technological Change in the Business Cycle.” *Europeah Economic Review* 44 (1):91–115.

Hall, George. 1996. "Overtime, Effort, and the Propagation of Business Cycle Shocks." *Journal of Monetary Economics* 38 (1):139–160.

Hall, Robert. 2010. "Why does the Economy Fall to Pieces after a Financial Crisis?" *Journal of Economic Perspectives* 24 (4):3–20.

———. 2011. "The High Sensitivity of Economic Activity to Financial Frictions." *Economic Journal* 121 (552):351–378.

He, Zhiguo and Arvind Krishnamurthy. 2012. "A Model of Capital and Crises." *Review of Economci studies* 79 (2):735–777.

———. 2013. "Intermediary Asset Pricing." *American Economic Review* 103 (2):732–770.

Hsieh, Chang-Tai and Peter Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *Quarterly Journal of Economics* 124 (4):1403–1448.

Humphreys, Brad, Louis Maccini, and Scott Schuh. 2001. "Input and output inventories." *Journal of Monetary Economics* 47 (2):347–375.

Iacoviello, Matteo, Fabio Schiantarelli, and Scott Schuh. 2011. "Input and output inventories in general equilibrium." *International Economic Review* 52 (4):1179–1213.

Jermann, Urban and Vicenzo Quadrini. 2006. "Financial Innovations and Macroeconomic Volatility." *National Bureau of Economic Research* (12308).

———. 2012. "Macroeconomic Effects of Financial Shocks." *American Economic Review* 102 (1):238–271.

Jr., Robert E. Lucas. 1977. ""Understanding Business Cycles"." *Carnegie-Rochester Conference Series on Public Policy* 5:7–29.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti. 2010. "Investment Shocks and Business Cycles." *Journal of Monetary Economics* 57 (2):132–145.

Kahn, James A. 1992. "Why is Production More Volatile than Sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory Holding." *Quarterly Journal of Economics* 107 (2):481–510.

Kahn, James A., Margaret Mary McConnell, and Gabriel Perez-Quiros. 2002. "On the Causes of Increased Stability in the U.S. Economy." *Economic Policy Review* 8 (1):1–20.

Khan, Aubhik and Julia Thomas. 2013. "Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity." *Unpublished Manuscript* .

King, Robert G., Charles I. Plosser, and Sergio T. Rebelo. 1988. "Production, growth and business cycles: I. The basic neoclassical model." *Journal of Monetary Economics* 21 (2-3):211–248.

Kydland, Finn E. and Edward C. Prescott. 1982. "Time to Build and Aggregate Fluctuations." *Econometrica* 50 (6):1345–1370.

Long, John B. and Charles I. Plosser. 1982. "Real Business Cycles." *The Journal of Political Economy* 91 (1):39–69.

Lubik, Thomas, Pierre-Daniel Sarte, and Felipe Schwartzman. 2014. "Supplementary Material: What Inventory Behavior Tells Us About How Business Cycles Have Changed." *mimeo* .

McConnell, Margaret M. and Gabriel Perez-Quiros. 2000. "Output Fluctuations in the United States: What has Changed since the Early 1980's?" *The American Economic Review* 90 (5):1464–1476.

Ramey, Valerie. 1989. "Inventories as Factors of Production and Economic Fluctuations." *The American Economic Review* 79 (3):338–354.

Ramey, Valerie and Daniel J. Vine. 1989. "Inventories as Factors of Production and Economic Fluctuations." *The American Economic Review* 79 (3):338–354.

Schwartzman, Felipe F. 2012. "Time to Produce and Emerging Market Crises." *Federal Reserve Bank of Richmond Working Paper Series* 10-15R.

Shourideh, Ali and Ariel Zetlin-Jones. 2012. "External Financing and the Role of Financial Frictions over the Business Cycle: Measurement and Theory." *mimeo* .

Sims, Eric R. 2011. "Permanent and Transitory Technology Shocks and the Behavior of Hours: A Challenge for DSGE Models." *mimeo* .

Wen, Yi. 2011. "Input and output inventory dynamics." *American Economic Journal - Macroeconomics* 3:181–212.

Table 1: **Changes in Business Cycle Properties in the Post-War Era**

	1953-1983	1984-2008	2008-2012
	a. Cross Correlations		
Output per Hour and Output	0.65	0.06	0.06
Output per Hour and Hours	0.13	-0.47	-0.33
Inventory/Sales and Output	-0.57	-0.03	0.18
	b. Standard Deviations		
Output	2.59	1.43	2.57
Consumption Relative to Output	0.37	0.44	0.37
Investment Relative to Output	2.04	2.35	2.52
Hours Relative to Output	0.77	1.12	1.04
Inventories to Output	0.75	1.13	1.22
Inventory-Sales Ratio to Output	0.87	0.71	0.67

Note: Correlations and Standard Deviations calculated from HP-Filtered data. See text for details of variable definitions.

Table 2: **Model Parameters**

	Estimated Value	Posterior Std.	Prior Mean	Prior Std.	Prior Distribution
g	1.01	-	-	-	Calibrated
δ	0.03	-	-	-	Calibrated
κ	0.40	-	-	-	Calibrated
β	0.99	-	-	-	Calibrated
ϱ	18.91	0.4124	4	4	Gamma
$\omega_D(s+1)/\omega_D(s)$	0.61	-	-	-	Calibrated
$\omega_{ND}(s+1)/\omega_{ND}(s)$	0.15	-	-	-	Calibrated
ρ_u	0.9978	0.0028	0.9	0.05	Beta
ρ_g	0.6555	0.0535	0.5	0.20	Beta
ρ_{a_D}	0.9980	0.0027	0.9	0.05	Beta
ρ_Υ	0.9890	0.0045	0.9	0.05	Beta
ρ_ζ	0.9984	0.0026	0.9	0.05	Beta

Note: “Estimated Values” refer to posterior mode estimates.

Table 3: **Cumulative Effects Wedges on Data Moments**

	TFP Transitory	TFP Permanent	TFP Sectoral	Labor Wedge	Investment Wedge	Data
			a. corr(GDP, GDP/Hours)			
Pre 84	0.99	0.93	0.89	0.57	0.67	0.65
Post 84	0.99	0.98	0.87	0.66	0.13	0.13
			b. corr(Hours, GDP/Hours)			
Pre 84	0.96	0.74	0.51	0.27	0.16	0.14
Post 84	0.96	0.91	0.18	0.17	-0.38	-0.38
			c. corr(GDP, Inventories/Final Sales)			
Pre 84	-0.23	0.26	-0.26	-0.27	-0.61	-0.58
Post 84	-0.30	-0.13	-0.45	-0.50	0.08	0.01
			d. Std(GDP)			
Pre 84	2.27	2.49	2.57	3.33	2.62	2.61
Post 84	1.61	2.05	1.34	1.59	1.67	1.67
			e. Std(Hours)/Std(GDP)			
Pre 84	0.37	0.54	0.53	0.85	0.75	0.77
Post 84	0.37	0.43	0.50	0.76	1.07	1.07
			f. Std(Consumption)/Std(GDP)			
Pre 84	0.42	0.33	0.98	0.77	0.36	0.37
Post 84	0.42	0.36	1.34	1.21	0.41	0.41
			g. Std(Inventories)/Std(GDP)			
Pre 84	0.79	1.11	0.98	0.92	0.73	0.74
Post 84	0.77	0.87	0.86	0.78	1.13	1.15

Note Calculated based on HP-filtered time-series obtained from historical decompositions.

Table 4: **Cyclical Correlations of Credit Conditions and Macroeconomic Aggregates in the Post-War Era**

	1953-1980	1984-2007	2008-2012
	a. Lagged Bond Spread (Baa - 10 year Treas.)		
Output	-0.079	-0.380	-0.627
Fixed Investment	-0.091	-0.451	-0.704
Inventories	-0.607	-0.649	-0.841
	b. (-) Payouts to Business Owners (Total)		
Output	-0.119	-0.546	-0.781
Fixed Investment	-0.119	-0.528	-0.797
Inventories	-0.204	-0.553	-0.825

Note: HP-Filtered Data

Table 5: **Business Cycle Correlation of Investment Wedge with Aggregates and Indicators of Credit Conditions**

	1953-1983	1984-2007	2008-2012
		a. Total Correlations	
Lagged Baa-Treasury Spread	0.301	0.593	0.879
Lagged Baa-Aaa Spread	0.175	0.289	0.833
Lagged Aaa-Treasury Spread	0.387	0.484	0.843
Lagged GZ Spread	-	0.349	0.860
Lagged GZ Excess Bond Premium	-	0.378	0.821
(-) Payouts (total)	0.084	0.421	0.836
(-) Payouts (corporate)	0.147	0.397	0.664
Debt Repurchases	-0.074	0.268	0.836
		b. Partial Correlations	
Lagged Baa-Treasury Spread	0.523	0.513	0.491
Lagged Baa-Aaa Spread	0.451	0.212	0.454
Lagged Aaa-Treasury Spread	0.476	0.526	0.501
Lagged GZ Spread	-	0.390	0.405
Lagged GZ Excess Bond Premium	-	0.320	0.699
Payouts (total)	0.212	0.474	0.785
Payouts (corporate)	0.358	0.470	0.500
Debt Repurchases	0.185	0.345	0.726

Note: HP-Filtered Data. Bond spreads are lagged four quarters.

Figure 1: **Impulse Response Functions for Output, Hours and Labor**

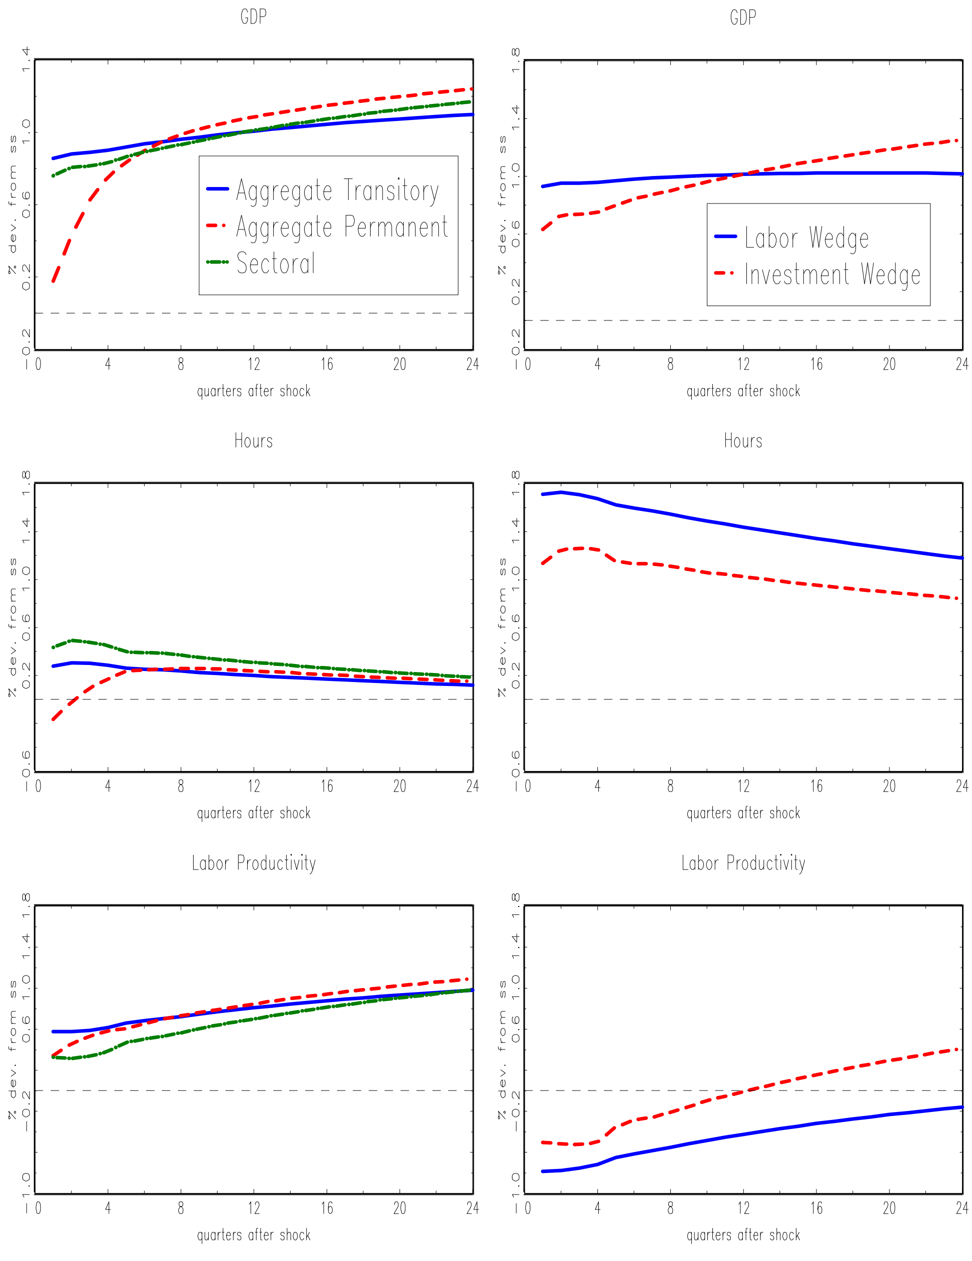

Note: Based on posterior mode estimate of model parameters.

Figure 2: **Impulse Response Functions for Output, Inventories and Inventory/Sales**

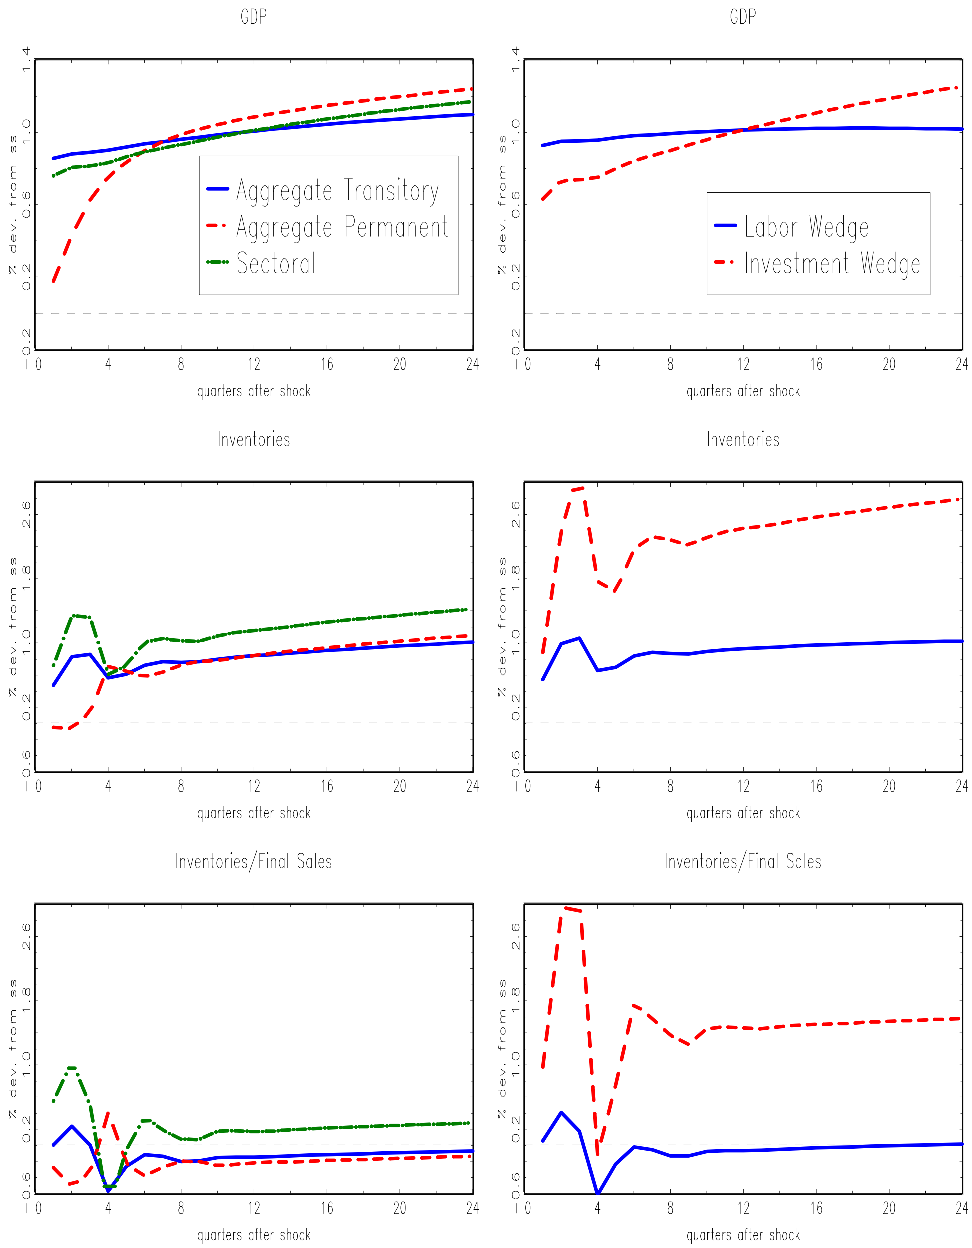

Note: Based on posterior mode estimate of model parameters.

Figure 3: **Historical Decomposition of Three Recessions - Benchmark**

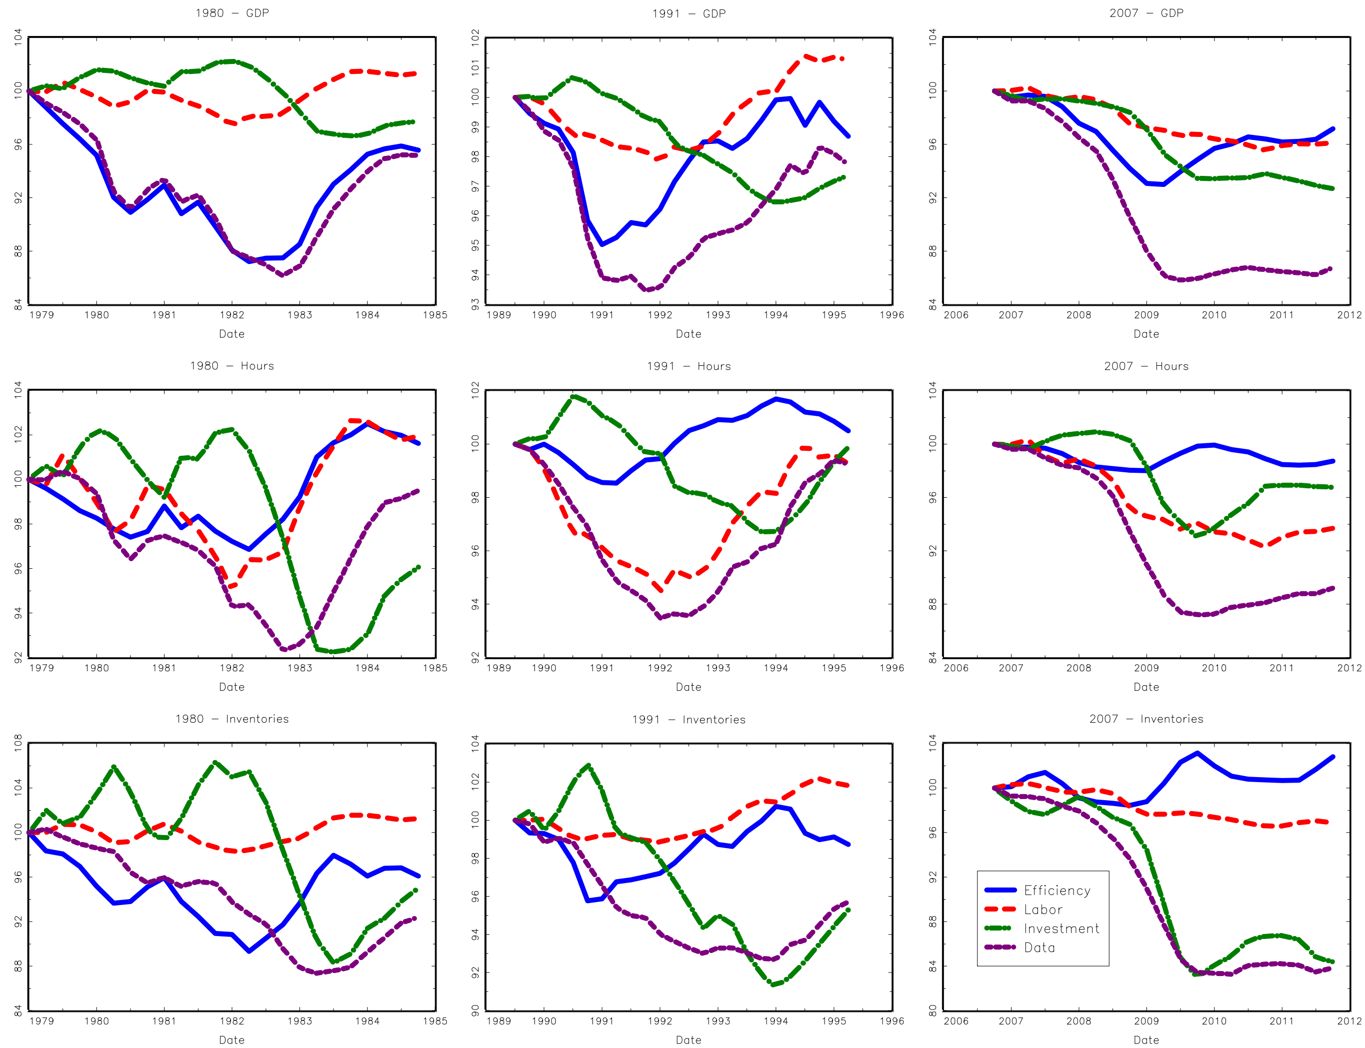

Note: Based on posterior mode estimate of model parameters.

Figure 4: **Historical Decomposition of Three Recessions - One Sector/Zero Stages of Production**

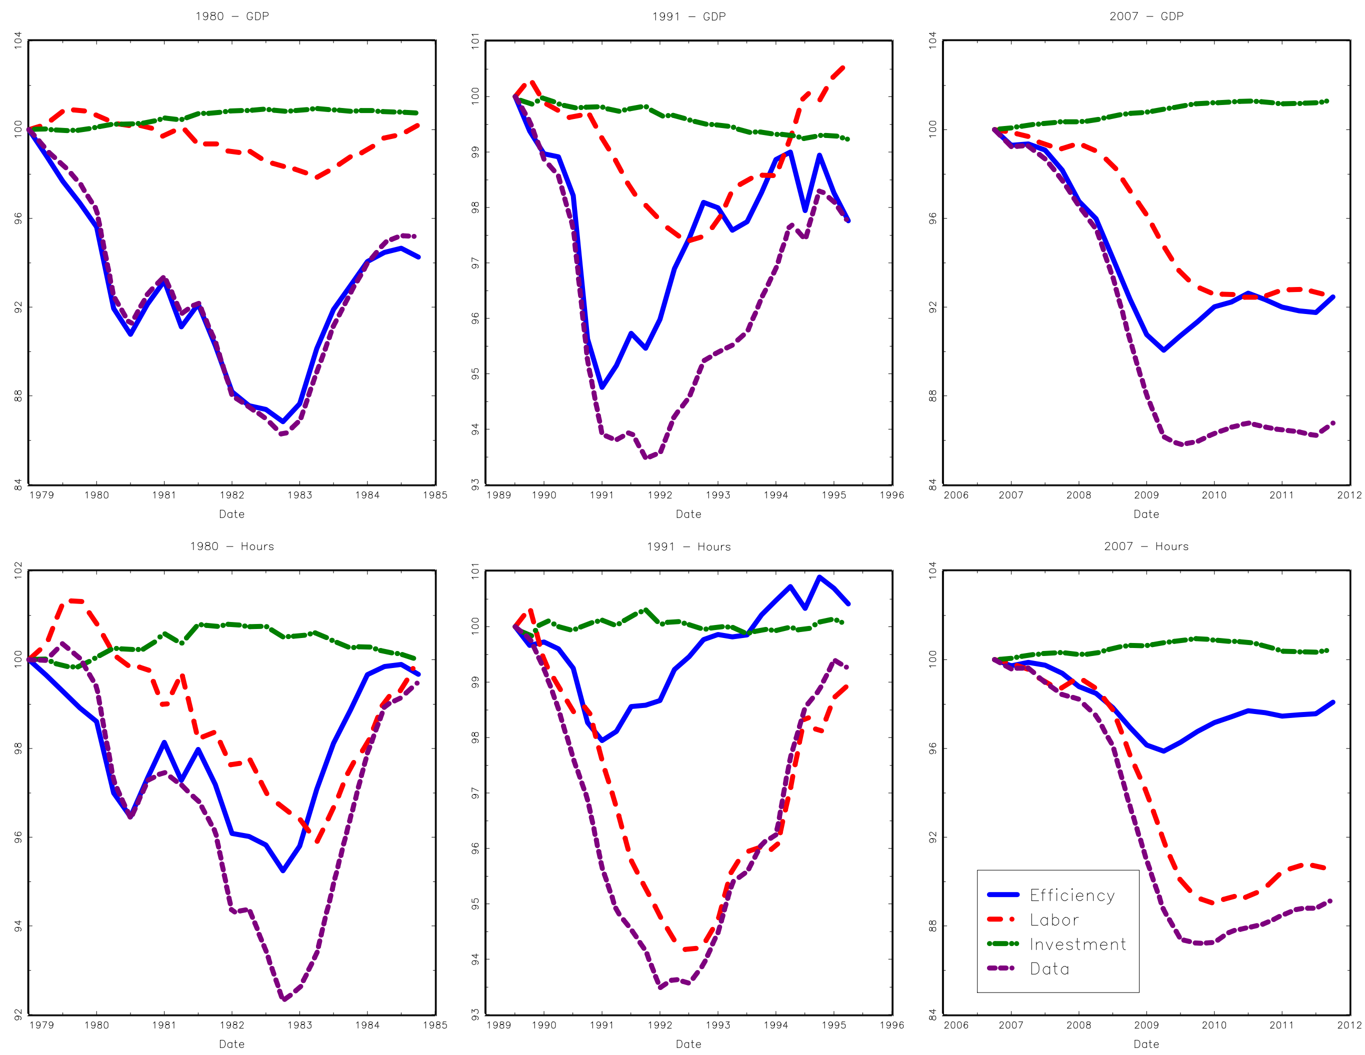

Note: Based on posterior mode estimate of model parameters for the model with $S = 0$ and $N = 1$ and using data for aggregate consumption, output and hours worked.

Figure 5: **External Validation and History of the Investment Wedge**

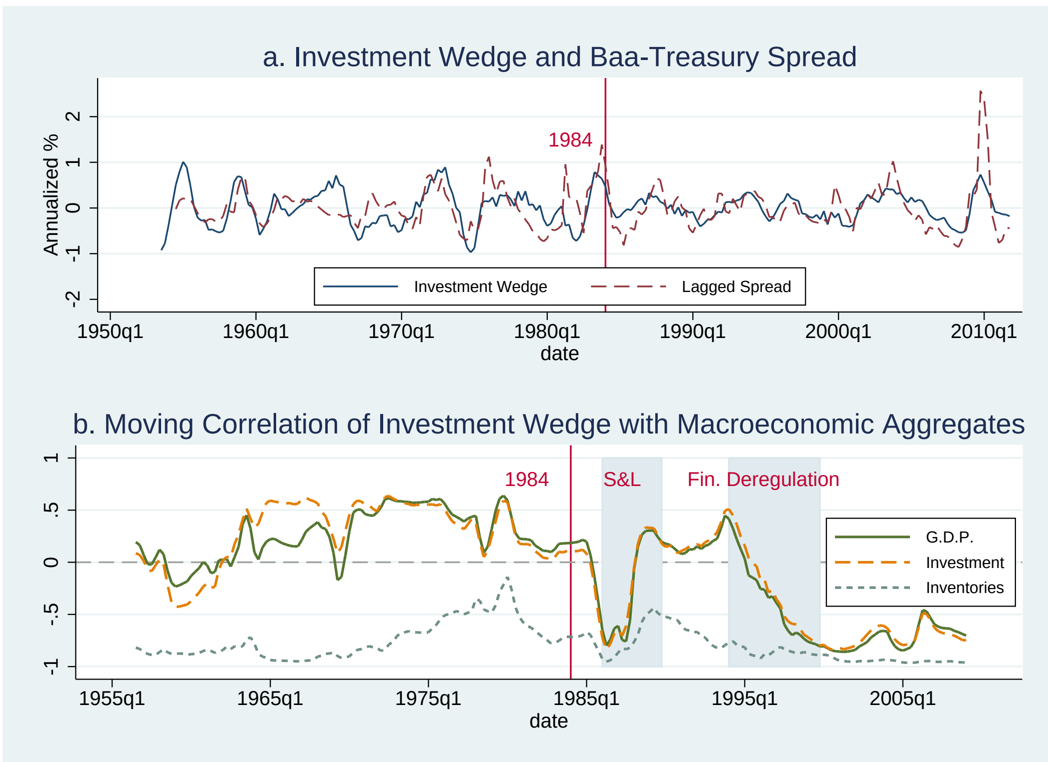

Note: Based on posterior mode estimate. Panel a) compares the estimated investment edge with the spread between and index of Moody's Baa rated corporate bonds and 10-year Treasuries. The spread is lagged four quarters. Panel b) shows centered moving correlations using windows which are 24 quarters wide.