

The Investment Network, Sectoral Comovement, and the Changing U.S. Business Cycle*

Christian vom Lehn
Brigham Young University
cvomlehn@byu.edu

Thomas Winberry
Chicago Booth and NBER
Thomas.Winberry@chicagobooth.edu

December 5, 2019

Abstract

We argue that the input-output network of investment goods across sectors is an important propagation mechanism for understanding business cycles. First, we show that the empirical network is dominated by a few “investment hubs” that produce the majority of investment goods, are highly volatile, and are strongly correlated with the cycle. Second, we embed this network into a multisector model and show that shocks to investment hubs have large aggregate effects while shocks to non-hubs do not. Finally, we measure realized sector-level productivity shocks in the data, feed them into our model, and find that hub shocks account for a large and increasing share of aggregate fluctuations. This fact allows the model to match the decline in the cyclicity of labor productivity and other business cycle changes since the 1980s. Our model also implies that investment stimulus policies increase employment throughout the economy but have unequal effects across sectors.

*We thank Rodrigo Adao, Gadi Barlevy, Bob Barsky, David Berger, Steve Davis, Emmanuel Farhi, Simon Gilchrist, Mike Golosov, Francois Gourio, Veronica Guerrieri, Jonathan Heathcote, Erik Hurst, Anil Kashyap, Jennifer La’O, Brent Neiman, Lilia Maliar (discussant), Brent Neiman, Pablo Ottonello, Rob Shimer, Henry Siu, Joe Vavra, and Christian Wolf for useful suggestions. We also thank seminar participants at the Southern Economic Association, North American Summer Meetings of the Econometric Society, International Conference for Computing in Economics and Finance, Federal Reserve Bank of Minneapolis, Federal Reserve Bank of Chicago, Norges Bank, Bundesbank, Penn-Wharton Budget Model, Chicago Booth, and University of Michigan. Winberry gratefully acknowledges financial support from Chicago Booth.

1 Introduction

The defining feature of business cycles is the comovement of production across the different sectors of the economy. A recent body of work has shown that the degree of sectoral comovement has fallen since the early 1980s and suggests that sector-specific shocks have become more volatile relative to aggregate shocks.¹ Our basic questions are how sector-specific shocks are propagated to macroeconomic aggregates and whether the resulting fluctuations resemble empirical business cycles. Of course, a large literature has studied how sector-specific shocks are propagated through the input-output network of intermediate goods across sectors. However, the input-output network of investment goods – which sectors produce investment goods and to which sectors those goods are sold – remains understudied, despite the fact that investment is the most volatile component of GDP over the business cycle.

We argue that the investment network is an important propagation mechanism for understanding business cycle fluctuations. We make this argument in three steps. First, we measure the empirical investment network and show that it is dominated by a small number of “investment hubs” that produce the majority of investment goods, are highly volatile, and are highly correlated with the aggregate cycle. Second, we embed this empirical network into a standard multisector real business cycle model and show that shocks to investment hubs have large aggregate effects while shocks to non-hubs do not. Finally, we measure realized sector-level productivity shocks in the data, feed them into our model, and show that shocks to investment hubs account for a large and increasing share of fluctuations over time. This fact allows the model to match a number of changes in business cycle patterns since the early 1980s, such as the declining cyclicalities of labor productivity.

The first step of our analysis is to measure the input-output network of investment goods. The investment network computes the amount of investment goods that are produced in sector i and subsequently sold to sector j for each pair of sectors (i, j) in the economy. While the BEA has released this information in the 1997 capital flows table, that table does not include all types of capital goods and is not readily available for other years. We therefore construct our own investment network using disaggregated data on sector-level purchases

¹See, for example, [Foerster, Sarte and Watson \(2011\)](#) and [Garin, Pries and Sims \(2018\)](#).

and production of various types of capital goods at the 35-sector level.

Our measured investment network is extremely sparse; four investment hubs – construction, machinery manufacturing, motor vehicles manufacturing, and professional/technical services – produce nearly 2/3 of total investment even though they only account of 10-15% of value added, employment, or intermediates production. These hubs are also more volatile, more correlated with aggregates, and more strongly lead the aggregate cycle than non-hub sectors, consistent with the central role of investment in cyclical fluctuations.

The second step of our analysis is to embed the empirical investment network into the multisector real business cycle model in order to understand the role of the network in propagating shocks. Each sector in our model produces gross output using capital, labor, and a bundle of intermediate goods consisting of other sectors' output; this bundle is computed by a Cobb-Douglas aggregator which characterizes the intermediates input-output network. Each sector also accumulates new capital using another Cobb-Douglas aggregator of investment goods, which characterizes our investment network. While other studies have used this basic model structure, our innovations are to discipline the investment network with our new measurement and explicitly study its role in propagating sector-specific shocks.

Our main new result from this model is that sector-specific shocks to investment hubs have large effects on aggregate employment while shocks to non-investment hubs do not. A positive shock to an investment hub directly increases production and employment at the hub; because the shock also raises the supply of investment goods for the rest of the economy, other sectors increase employment in order to produce more intermediate inputs for the hub and facilitate their own capital accumulation. In contrast, a shock at a non-hub has a small effect on investment supply and therefore generates smaller spillovers to the rest of the economy.²

Changes in investment supply are key to propagating sector-specific shocks because they

²Our investment hub shocks are reminiscent of the investment-specific technology shocks studied in, for example, [Greenwood, Hercowitz and Krusell \(2000\)](#) or [Justiniano, Primiceri and Tambalotti \(2010\)](#). A common problem in that literature is that investment-specific shocks generate negative comovement between investment- and consumption-producing sectors, decreasing the aggregate effect of these shocks. Our model generates positive comovement through the intermediates network, which implies that investment-hub shocks have large aggregate effects. In addition, this literature often measures investment-specific shocks using the relative price of investment goods, which is weakly correlated with the aggregate cycle; we measure investment hub shocks using sector-level productivity, which is strongly correlated with the cycle.

weaken general equilibrium effects which dampen the response of employment to the shock. In fact, in the version of our model without investment, employment is constant in response to shocks due to these general equilibrium effects, regardless of the structure of the intermediates network. In that model, a positive sector-specific shock leads to equal-sized but opposing changes in the marginal product of labor and in the marginal utility of consumption (due to our Cobb-Douglas preferences and technology as well as the fact that consumption is the only final use of output). In equilibrium, the marginal utility of consumption is equal to the price of the sector's output, so these two forces leave the marginal revenue product of labor – and therefore equilibrium employment – unchanged. In our model with investment, a sector-specific shock increases both consumption and investment, which dampens the decline in the marginal utility of consumption – especially for shocks to investment hubs. Hence, the investment network propagates sector-specific shocks in a fundamentally different way than the intermediates network.

We show that three key testable implications of our main result hold in the data. First, the volatility of sector-level employment is higher at investment hubs than at non-investment hubs, consistent with the idea that investment hub shocks generate larger changes in employment than non-hub shocks. Second, value added growth at investment hubs predicts changes in aggregate employment much more strongly than value added growth at non-hubs, consistent with the idea that hub shocks have larger effects on aggregate employment. Third, sectors which supply intermediate inputs to the investment hubs comove more strongly with the hubs than sectors which do not supply the hubs with intermediates. This finding is consistent with the role of the intermediates network in propagating investment hub shocks to other sectors.

The third step of our analysis is to use this model to understand the aggregate effects of changes in the process for sector-level productivity shocks since 1984. We measure the realized time series of sector-level productivity in our data using the Solow residual approach. We find that the covariance of productivity across sectors has fallen substantially but that the variance within sectors has not, which we interpret as a rise in the volatility of sector-specific shocks relative to the volatility of aggregate shocks (consistent with [Foerster, Sarte and Watson \(2011\)](#) and [Garin, Pries and Sims \(2018\)](#)'s findings using value added). We feed

the realized time series of these shocks into our model and simulate the model's decision rules over the entire postwar sample. In order to isolate the role of the change in the shock process, we hold all other parameters of the model (including the investment network) fixed over time; changes in these parameters are second-order for our results in the sense that our results are robust to allowing these parameters to change as well.

We find that this rising importance of sector-specific shocks, when propagated through the investment network, quantitatively generates a number of changes in aggregate business cycle patterns since 1984, including the declining cyclical of labor productivity. This result can be understood in two steps. First, the rising importance of sector-specific shocks implies that shocks to investment hubs account for the majority of aggregate fluctuations post-1984 (because sector-specific shocks to non-hubs have small aggregate effects). Second, aggregate labor productivity is countercyclical in response to shocks to investment hubs because these shocks have strong spillovers onto the production of others sectors, as described above. These other sectors increase their employment by more than value added because their productivity is unchanged and labor is subject to decreasing returns to scale; in total, aggregate employment increases by more than GDP, decreasing labor productivity. In contrast, the pre-1984 sample was dominated by aggregate TFP shocks, which generates procyclical labor productivity almost by construction.

We document two new empirical results that support the idea that the changes in business cycle patterns reflect the rising importance of shocks to investment hubs. First, the changes in business cycle patterns have not occurred within individual sectors of the economy, but are due to changes in the comovement of activity across sectors. For example, sector-level labor productivity is still highly procyclical within sector; instead, the entire decline in the cyclical of aggregate labor productivity is due to changes in the covariance of value added and employment across sectors. This finding is consistent with the idea that the comovement of shocks across sectors is the key force driving changes since 1984. Existing explanations for the declining cyclical of labor productivity largely abstract from sectoral heterogeneity and therefore do not speak to this result. Our second finding is that the volatility of investment relative to the volatility of GDP has substantially increased since 1984, consistent with the idea sector-specific shocks to investment hubs play an increasingly important role over time.

We show that our model quantitatively replicates both of these empirical findings.

We conclude that the investment network is an important mechanism in propagating sector-specific shocks; we also briefly use the model to illustrate how the investment network propagates investment stimulus policies, such as investment tax credits or the bonus depreciation allowance. These policies increase the demand for investment goods, which directly increases production and employment at the hubs; in turn, the hubs demand intermediates from other sectors, which increases their employment as well. However, the change in other sectors' employment is relatively small, so most of the increase in aggregate employment is concentrated among investment hubs and their intermediate goods suppliers. Hence, despite the fact that the stimulus subsidizes investment purchases equally for all sectors, the sparseness of the investment network implies that its effect on investment production and employment is unequally distributed across sectors – resembling industrial policy.

Related Literature Our paper is most closely related to three lines of research. The first is the large and growing literature which studies how the rich structure of the intermediates input-output network amplifies the effect of idiosyncratic shocks. The first wave of these papers studies the role of sector-level linkages while the second further studies firm-to-firm linkages.³ Another set of papers endogenizes the network as an optimal firm-level choice of input suppliers.⁴ In order to allow for a rich intermediates network, these papers typically use static models which abstract from investment. Our contribution to this literature is to use a dynamic model to study the role of investment and its network structure across sectors.⁵ We show that the investment network propagates shocks fundamentally different from the intermediates network; without investment, sector-specific shocks would have no effect on employment due to general equilibrium effects, no matter how concentrated the

³See, for example, [Acemoglu et al. \(2012\)](#), [Acemoglu, Ozdaglar and Tahbaz-Salehi \(2017\)](#), [Baqae and Farhi \(2017\)](#), [Baqae and Farhi \(2019\)](#), [Bigio and La’o \(2019\)](#), or the survey in [Carvalho and Tahbaz-Salehi \(2019\)](#).

⁴See, for example, [Oberfield \(2012\)](#), [Lim \(2018\)](#), or [Taschereau-Dumouchel \(2019\)](#).

⁵A natural benchmark in these static models is Hulten’s theorem, which states that the Domar weight of a sector is a sufficient statistic to compute the effect of a shock to that sector on aggregate GDP. In our model, a sector-specific shock has a different effect on aggregate GDP, aggregate employment, and aggregate investment, suggesting that one weight cannot be a sufficient statistic for all three effects. While it may be of interest to try to show whether a Hulten’s theorem-type result holds in our model, and how the investment network enters that result, we do not pursue it here.

intermediates network is.

The second strand of related literature uses the multisector real business cycle model to embed input-output networks in a dynamic setting.⁶ Foerster, Sarte and Watson (2011) and Atalay (2017) quantify versions of this model and argue that sector-specific shocks can generate aggregate fluctuations. Foerster, Sarte and Watson (2011) additionally argue that the volatility of sector-specific shocks has risen relative to the volatility of aggregate shocks since the 1980s and that this fact accounts for the decline in GDP volatility. These papers primarily focus on the role of the intermediates network while we focus on the role of the investment network.⁷ We make three main contributions to this literature. First, we calibrate the model using our empirical investment network, which is more sparse than other calibrations in the literature (see footnote 7). Second, we show that shocks to investment hubs have large aggregate effects while shocks to non-hubs do not. Third, we show that the rising importance of sector-specific shocks, when filtered through our investment network, generates a number of changes in business cycle patterns beyond the decline in GDP volatility.

The final strand of related literature studies how business cycle patterns have changed since the 1980s.⁸ A large subset of that literature focuses on the declining cyclicity of labor productivity in particular and has suggested roughly three sets of explanations. The first set of explanations is that the aggregate shock process has changed over time.⁹ The second set

⁶Seminal contributions to this literature include Long Jr and Plosser (1983), Horvath (1998, 2000), and Dupor (1999).

⁷Both Foerster, Sarte and Watson (2011) and Atalay (2017) include an investment network in their models, but because of the nature of their questions, neither paper focuses on the role of this network in driving their results and, therefore, neither paper studies the mechanism we emphasize in this paper. In addition, Foerster, Sarte and Watson (2011) and Atalay (2017) calibrate the investment network using the BEA capital flows data from 1997; under this network, they are forced to make an adjustment to ensure their model is invertible but which reduces the importance of investment hubs. A contribution of our paper is to carefully measure the network over the entire sample in order to precisely quantify the strength of investment hubs. Our network does not require any adjustment to ensure that the model is invertible; hence, it may be of interest to other modelers in this literature as well.

In a complementary recent paper, Foerster et al. (2019) use a related multisector model to study trends in sector-level productivity rather than deviations around trend. Quantitatively, they find that capital accumulation in investment-producing sectors plays an important role in aggregating sector-specific trends into the aggregate growth rate, complementary with the role of investment hubs in propagating shocks which we study in this paper. Like Foerster, Sarte and Watson (2011), Foerster et al. (2019) calibrate the investment using the 1997 capital flows table.

⁸The first strand of this literature studies the declining volatility of aggregate GDP since the early 1980s; see, for example, the early review in Stock and Watson (2003) and the papers that follow.

⁹For example, Galí and Gambetti (2009) and Barnichon (2010) argue that the volatility of demand shocks has changed since the 1980s, while Garin, Pries and Sims (2018) argue that reallocation shocks have become

is that firms and/or workers can now more easily adjust labor inputs in response to shocks.¹⁰ The third is that there has been no actual change in the cyclical behavior of labor productivity, but that mismeasurement of those objects has changed.¹¹ This literature typically constructs models without sectoral heterogeneity and therefore cannot speak to our empirical fact that the entire decline in the cyclical behavior of labor productivity is due to changes in the covariance of activity across sectors.¹²

Road Map Our paper is organized as follows. In Section 2, we measure the empirical investment network and document the cyclical behavior of investment hubs. We then describe our version of the multisector real business cycle model and calibrate it to match our measured investment network in Section 3. We use a simplified two-sector version of the model to show that shocks to investment hubs have large aggregate effects while shocks to non-hubs do not in Section 4. We also show that key testable implications of that mechanism hold in the data. In Section 5, we show two sets of empirical results that motivate our main quantitative application of the model: the correlation of sector-level productivity has fallen since 1984 and the changes in aggregate business cycle patterns since 1984 have not occurred within sectors but are driven by changes in the comovement of activity across sectors. In Section 6, we feed in the realized series of sector-level productivity shocks into our model and show that the model endogenously replicates the changes in business cycle patterns. We briefly analyze the effect of investment stimulus policy in Section 7. Section 8 concludes.

more important.

¹⁰Barnichon (2010) and Galí and Van Rens (2014), among others, argue that labor market frictions have declined since the 1980s. Koenders and Rogerson (2005), Berger (2012), and Bachmann (2012) build models in which frictions individual firms' hiring and firing decisions interact with the declining aggregate volatility to generate the declining cyclical behavior of labor productivity and/or the emergence of jobless recoveries.

¹¹One strand of this literature argues that utilization is not correctly measured in productivity data; see, for example, Fernald and Wang (2015). Another strand argues that intangible capital is not correctly measured in the output data; see, for example, McGrattan and Prescott (2014) or McGrattan (2017).

¹²We are aware of one paper which studies the declining cyclical behavior of labor productivity in a model with sectoral heterogeneity: Garin, Pries and Sims (2018). We view their paper as complementary to our paper; we both study the rise of sector-specific shocks, but focus on different mechanisms which propagate those shocks to the aggregate. In Garin, Pries and Sims (2018)'s two sector model, a negative sector-specific shock induces costly worker reallocation to the other sector, so employment falls by more than value added and labor productivity increases. This mechanism implies that employment in the two sectors comoves negatively, especially post-1984; however, in the data, employment comovement is positive and stable. Hence, to our knowledge, our model is the only explanation for the declining cyclical behavior of labor productivity that is driven by the changes in comovement patterns that we document in the data.

TABLE 1
THE 35 SECTORS USED IN THE ANALYSIS

Mining	Utilities
Construction	Wood products
Non-metallic minerals	Primary metals
Fabricated metals	Machinery
Computer and electronic manufacturing	Electrical equipment manufacturing
Motor vehicles manufacturing	Other transportation equipment
Furniture and related manufacturing	Misc. Manufacturing
Food and beverage manufacturing	Textile manufacturing
Apparel manufacturing	Paper manufacturing
Printing products manufacturing	Petroleum and coal manufacturing
Chemical manufacturing	Plastics manufacturing
Wholesale trade	Retail trade
Transportation and warehousing	Information
Finance and insurance	Professional and technical services
Management of companies and enterprises	Administrative and waste management services
Educational services	Health care and social assistance
Arts, entertainment, and recreation services	Accommodation and food services
Other services	

Notes: list of sectors used in our empirical analysis. See Appendix A for details of the data construction.

2 Descriptive Evidence on the Investment Network

2.1 Data Sources

We combine three sources of sector-level data for our empirical work. First, we construct the investment network using the BEA Fixed Assets and Input-Output databases for a sample of 35 private non-farm sectors from 1947-2017 (our construction of the investment network is described below). Second, we use the BEA GDP-by-Industry database to obtain value added and employment for the same set of sectors; however, since this data only records employment at our level of disaggregation starting in 1977, we extend the data back to 1948 using historical supplements to the data. Our combined dataset contains annual observations of value added, investment, and employment for the 1948 - 2017 period. Appendix A contains details of the construction of our dataset.

Table 1 lists the sectors available in our dataset.¹³ The main advantage of this data is

¹³We exclude real estate because its value added is dominated by imputations of owner-occupied housing.

that it is a long time series, which is necessary to analyze changes in business cycle patterns over time. In addition, the partition of sectors provides fairly broad coverage of the economy. We cannot disaggregate the sectors more finely in a consistently-defined way over time.

For the remainder of the paper, references to “value added” and “investment” generally refer to real value added and real investment, unless otherwise noted. We measure real value added and real investment using the index numbers provided by the BEA.

2.2 Empirical Investment Network


The investment network records the share of total investment expenditure of a given sector j that is purchased from another sector i for each pair of sectors (i, j) in the data. While the BEA capital flows tables report these pairwise shares, they are only available for a handful of years, are not coded for a consistent set of sectors, are not computed for a consistent definition of investment goods over time, and do not include the majority of intellectual property. We therefore construct our own measure of the investment network which is consistently defined over time, includes intellectual property, and is available for the entire 1947-2017 sample.¹⁴ Our procedure allocates the production of roughly thirty types of disaggregated capital goods to a particular mix of sectors, and allocates the purchases from a particular sector using that industry mix. Appendix A contains the details of our procedure.¹⁵ We compute the network for each year in the sample and then average it over time. Appendix A discusses how the network has changed over time and Appendix G shows how those changes over time impact our model results.¹⁶

¹⁴We include all of intellectual property capital in our analysis following the most recent BEA convention. However, our results – both in the data and the quantitative model – are stronger if we focus only on equipment and structures.

¹⁵Our measured investment network includes imports from outside of the U.S. as well. Hence, our measurement incorporates the fact that the share of imported capital has increased over time (see House, Mocanu and Shapiro (2017)). See footnote 24 for a discussion of how this measurement relates to our closed-economy model.

¹⁶As discussed in footnote 7, Foerster, Sarte and Watson (2011) and Atalay (2017) add a correction to the investment network implied by the 1997 BEA capital flows table to ensure their models are invertible. This correction is meant to account for maintenance investment that is done out of own-sector output. While there is evidence that maintenance investment is sizable (see McGrattan and Schmitz Jr (1999)), there are not comprehensive estimates of how much comes from own-sector output. Therefore, we do not add an artificial correction for maintenance investment in our baseline analysis. Appendix G does add such a correction and shows that our key model results also hold in this version of the model.

FIGURE 1: Heatmap of Empirical Investment Network



Notes: heatmap of empirical investment network. Entry (i, j) computes share of total investment expenditure in sector j that is purchased from sector i , averaged over the 1948 - 2017 sample.

Figure 1 plots a heat map of our estimated investment network. Four sectors supply investment goods to most other sectors in the economy: construction, which supplies the majority of structures; machinery manufacturing and motor vehicle manufacturing, which supply the majority of equipment; and professional/technical services, which supplies a majority of intellectual property. We refer to these four sectors as *investment hubs*.¹⁷

Table 2 shows that the investment hubs account for approximately 2/3 of all investment produced in the economy both before and after the 1984 breakpoint in business cycle patterns.¹⁸ In contrast, hubs only account for approximately 15% of value added production, intermediates production, or employment, and even less of investment purchased. The fact

¹⁷Professional/technical services is different from the other hubs in two ways. First, it primarily produces intellectual property, while the other hubs produce equipment and structures. Second, its share of investment supply has grown substantially over the sample relative to the other hubs. All of our empirical results regarding investment hubs hold using only construction, machinery manufacturing, and motor vehicles.

¹⁸We split the sample in 1984 because it is a commonly used break point in the literature on changes in business cycle patterns. We do not have the statistical power to estimate a structural break because our data is at the annual frequency.

TABLE 2
SHARES OF ACTIVITY AT INVESTMENT HUBS

	Investment Hubs		Other Sectors	
	<i>Pre-1984</i>	<i>Post-1984</i>	<i>Pre-1984</i>	<i>Post-1984</i>
Investment produced	0.67	0.61	0.33	0.39
Value added produced	0.16	0.17	0.84	0.83
Intermediates produced	0.12	0.15	0.87	0.85
Employment	0.14	0.14	0.86	0.86
Investment purchased	0.10	0.12	0.90	0.88

Notes: share of nominal investment produced, nominal value added produced, nominal intermediates produced, employment, and nominal investment purchased by investment hub sectors and other sectors in the economy. Investment hubs are construction, machinery manufacturing, motor vehicles manufacturing, and professional/technical services. “Pre-1984” refers to average values in 1948 - 1983 subsample and “post-1984” refers to averages in 1984-2017 subsample.


TABLE 3
VOLATILITY OF ACTIVITY AT INVESTMENT HUBS

	Investment Hubs		Non-Hubs	
	<i>Pre-1984</i>	<i>Post-1984</i>	<i>Pre-1984</i>	<i>Post-1984</i>
$\sigma(y_{st})$	6.62%	8.27%	4.78%	4.02%
$\sigma(l_{st})$	4.33%	3.39%	2.73%	2.22%

Notes: standard deviation of business cycle component of sector-level value added or employment. y_{st} is logged real value added at sector s , HP-filtered with smoothing parameter 6.25. l_{st} is logged real value added at sector s , HP-filtered with smoothing parameter $\lambda = 6.25$. “Investment hubs” compute the unweighted average the value of these statistics over $s =$ construction, machinery manufacturing, motor vehicles manufacturing, and professional/technical services. “Non-hubs” compute the unweighted average over the remaining sectors. “Pre-1984” performs this analysis in the 1948 - 1983 subsample and “post-1984” performs this analysis in the 1984 - 2017 subsample. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

that this small number of hubs produce the majority of investment indicates that the investment network is highly skewed; in fact, Appendix A shows that the investment network is three times more skewed than the intermediates network according the skewness of their eigenvalue centralities or weighted outdegrees.

FIGURE 2: Correlogram of Sector-level Value Added with Aggregate Employment



Notes: correlation of value added at sector s in year $t - h$, y_{st+h} , with aggregate employment in year t , L_t . Both y_{st+h} and L_t are logged and HP-filtered with smoothing parameter 6.25. x-axis varies the lag $h \in \{-2, -1, 0, 1, 2\}$. “Investment hubs” compute the unweighted average the value of these statistics over $s =$ construction, machinery manufacturing, motor vehicles manufacturing, and professional/technical services. “Non-hubs” compute the unweighted average over the remaining sectors. “Pre-1984” performs this analysis in the 1948 - 1983 subsample and “post-1984” performs this analysis in the 1984 - 2017 subsample. Aggregate employment is defined as the sum of employment in each of our sectors and is logged and HP-filtered. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

2.3 Investment Hubs are Highly Cyclical

We now show that these four investment hubs are more cyclical than other sectors, consistent with their central role in driving fluctuations in our model. Table 3 shows that value added and employment at the investment hubs are more volatile over the business cycle than value added and employment at non-hub sectors; in both the pre- and post-1984 subsamples, the investment hubs are approximately 1.5 - 2 times as volatile as non-hub sectors. This difference in value added volatility is larger in the post-1984 subsample.

Figure 2 shows that investment hubs are also more correlated with the aggregate business cycle. We compute the correlogram of HP-filtered sector-level value added in year $t + h$ with HP-filtered aggregate employment in year t (Appendix B shows that similar results hold using aggregate GDP rather than aggregate employment). Investment hubs’ value added is more correlated with aggregate employment at most horizons and the difference is stronger in the post-1984 subsample, consistent with the idea that hubs shocks have become more

important for aggregate fluctuations over time. In addition, investment hubs more strongly lead aggregate employment than non-hubs.

Appendix B contains two additional results about the cyclical behavior of investment hubs. First, we show that non-hub manufacturing sectors' behavior is more similar to the non-hub sectors than they are to the investment hubs. This result allays the concern that our results may be fundamentally driven by higher cyclicalities of manufacturing sectors, and the fact that manufacturing just happens to be over-represented in our set of investment hubs. Second, we show that, on average, each sectors' value added and employment comoves more strongly with investment hubs than non-hubs.

3 Model and Calibration

We now calibrate a version of the multisector real business cycle model in order to match our empirical investment network.

3.1 Model Description, Equilibrium, and Solution

The specification of the model is standard (see, for example, [Foerster, Sarte and Watson \(2011\)](#)).

Environment Time is discrete and infinite. There are a finite number of sectors indexed by $j = \{1, \dots, N\}$, where $N = 35$ as in our data. Each sector produces gross output using the production function

$$Q_{jt} = A_{jt} \left(K_{jt}^{\alpha_j} L_{jt}^{1-\alpha_j} \right)^{\theta_j} X_{jt}^{1-\theta_j} \quad (1)$$

where Q_{jt} is output, A_{jt} is total factor productivity, K_{jt} is capital, L_{jt} is labor, X_{jt} is a bundle of intermediate goods, and α_j and θ_j are parameters. TFP A_{jt} follows the AR(1) process

$$\log A_{jt+1} = \rho_j \log A_{jt} + \varepsilon_{jt+1}, \quad (2)$$

where ρ_j is the persistence parameter and ε_{jt} are innovations (which can be correlated across sectors). We solve the model by linearization, so the covariance matrix of these innovations

does not affect the decision rules. In our quantitative analysis in Section 6, we simply feed in the empirical time series of realized shocks into the decision rules.

The bundle of intermediate inputs X_{jt} consists of other sectors' output, aggregated through the economy's intermediates input-output network

$$X_{jt} = \prod_{i=1}^N M_{ijt}^{\gamma_{ij}}, \quad \text{where } \sum_{i=1}^N \gamma_{ij} = 1, \quad (3)$$

where M_{ijt} is the amount of sector i 's output used by sector j and γ_{ij} are parameters. Constant returns to scale in intermediate bundling implies that, within sector j , the parameters γ_{ij} sum to one. Each period, firms in sector j observe the TFP shock A_{jt} , use their pre-existing stock of capital K_{jt} , hire labor L_{jt} from the competitive labor market, and purchase intermediates X_{jt} in competitive markets in order to produce gross output Q_{jt} .¹⁹

After production, each sector accumulates capital for the next period using a bundle of inputs that are aggregated through the economy's investment network. The capital accumulation technology is

$$K_{jt+1} = (1 - \delta_j)K_{jt} + I_{jt} \quad (4)$$

where δ_j is the depreciation rate of capital in sector j and I_{jt} is the bundle of investment goods.²⁰ The bundle is given by

$$I_{jt} = \prod_{i=1}^N I_{ijt}^{\lambda_{ij}}, \quad \text{where } \sum_{i=1}^N \lambda_{ij} = 1, \quad (5)$$

where I_{ijt} is the amount of sector i 's output used by sector j and λ_{ij} are parameters. Invest-

¹⁹The specification of our production function imposes two simplifications. First, the elasticity of substitution between capital and labor is one. We view this as a useful benchmark; most empirical estimates of the elasticity are less than one, but some estimates are greater than one. An elasticity less than one would strengthen our results because it makes capital and labor more complementary in production, and therefore changes in investment would spill over more strongly to changes in employment. Second, the elasticity of substitution between intermediates and the primary inputs is also one. [Atalay \(2017\)](#) argues that this elasticity is likely below one; however, we nevertheless view our choice as a useful benchmark for two reasons. First, [Atalay \(2017\)](#)'s estimates are not readily available at our level of sectoral aggregation. Second, our results would be stronger with a lower elasticity because an investment hub shock would generate a larger increase in intermediates demand and, therefore, even stronger spillovers to other sectors in the economy. See Section 4 for a detailed discussion of these channels.

²⁰We do not include capital adjustment costs in our baseline for the sake of expositional clarity; however, Appendix G shows that our main results also hold with adjustment costs.

ment hub sectors i have high λ_{ij} for many purchasing sectors j .

There is a representative household who owns all the firms in the economy and supplies labor to those firms. The household's preferences are represented by the utility function

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\log C_t - \chi \frac{L_t^{1+1/\eta}}{1+1/\eta} \right), \quad \text{where } C_t = \prod_{j=1}^N C_{jt}^{\xi_j} \text{ and } \sum_{j=1}^N \xi_j = 1 \quad (6)$$

where β is the discount factor, χ controls the disutility of labor supply, η is the Frisch elasticity of labor supply, and ξ_j are parameters governing the importance of each sector's consumption good in aggregate consumption.²¹

Equilibrium We study a competitive equilibrium. There are two sets of market clearing conditions in that equilibrium. Output market clearing for sector j ensures that gross output is used for final consumption, investment, or an intermediate in production: $Q_{jt} = C_{jt} + \sum_{i=1}^N I_{jit} + \sum_{i=1}^N M_{jit}$. Labor market clearing ensures that aggregate labor demand equals labor supply: $\sum_{j=1}^N L_{jt} = L_t$. The equilibrium is efficient, so we characterize it using the social planner's problem.

Solution Method We solve the model by log-linearization. A key advantage of linearization is that it is efficient enough to handle a model of this size (with several hundred endogenous variables). As discussed above, the solution features certainty equivalence, which allows us feed in the realized sector-level shocks without estimating how the entire covariance structure of shocks has changed over time. However, linearization implies that we do not capture potential nonlinearities that arise due to the investment network, such as size-

²¹Our preferences impose two additional simplifying assumptions. First, there are no frictions to reallocating workers across sectors; in our quantitative work, we find that we do not need these frictions to match the comovement of employment across sectors at the business cycle frequency. We interpret the comovement across sectors in our model as primarily reflecting movement in and out of non-employment within sectors. Second, the elasticity of substitution between consumption of different sectors' output is one; we view this choice as a useful benchmark and is broadly consistent with estimates in the literature. For example, [Herrendorf, Rogerson and Valentinyi \(2013\)](#) estimate the elasticity between agriculture, manufacturing, and services to be around 0.9, while [Oberfield and Raval \(2014\)](#) estimate the elasticity between finely disaggregated manufacturing sectors to be between 0.8 and 1.1. Note that our elasticity of substitution is lower than the elasticity between detailed varieties of goods, which is the typical level of aggregation in DSGE models, because we have a relatively small number of sectors.

or state- dependent responses to shocks.²²

3.2 Calibration

We calibrate the parameters of the model so that the model’s steady state matches key empirical targets averaged over the postwar sample. For now, we leave the process for sector-level shocks unspecified; we will feed in the realizations of the shocks from the data in Section 6.

A model period is one year. We identify the sectors in our model with those in our empirical work, and therefore use the BEA input-output database to infer the parameters of the production function. The share of primary inputs in production θ_j is given by the ratio of sector j ’s value added to its gross output, averaged over time. The labor share $1 - \alpha_j$ is given by the average labor compensation (adjusted for taxes and self employment) as a share of total costs.²³ See Appendix C for the calibrated values of these parameters. Finally, the parameters of the intermediates network γ_{ij} are given by sector j ’s average expenditure on intermediates from sector i as a share of its total intermediates expenditure.

We infer the parameters of the investment technology using the BEA fixed asset tables and our measured investment network. The depreciation rates δ_j are average annual depreciation of capital goods, aggregated to the sector level weighted by the average amount of each type of good used in sector j . We construct the investment network using the procedure described in Section 2 (and in detail in Appendix A).


Figure 3 compares the heatmaps of our calibrated intermediates network and investment network (the investment network is reproduced from Figure 1). The intermediates network has a strong diagonal element, capturing firms’ purchases of intermediates from within their own sector, but is also populated off the diagonal, capturing intermediates purchased from other sectors. In comparison, the investment network is much more sparse, as discussed in Section 2.²⁴

²²See Baqaee and Farhi (2017) for an analysis of the types of nonlinearities that arise in a rich static model with an intermediates input-output network.

²³Labor compensation is only available in the BEA GDP by Industry database starting in 1987. In order to compute labor shares before 1987, we extend the series back to 1987 using historical tables and harmonize the two series using the Fort and Klimek (2016) crosswalk. See Appendix C for details.

²⁴Our measured intermediates and investment input-output networks account for goods that are imported

FIGURE 3: Heatmaps of Input-Output Networks



Notes: heatmaps of intermediates input-output network γ_{ij} and the investment input-output network λ_{ij} constructed as described in the main text and Appendix C. The (i, j) entry of each network corresponds to parameter γ_{ij} and λ_{ij} , i.e., the amount of sector i 's good used in sector j .

Finally, we infer the parameters of the utility function using the BEA’s final use table. The consumption shares ξ_j are given by the average consumption expenditure on sector j output as a fraction of total consumption expenditure. We set the discount factor to $\beta = 0.96$. We normalize the disutility of labor parameter, $\chi = 1$. We take the Frisch elasticity $\eta \rightarrow \infty$ to capture indivisible labor at the individual level, as in [Rogerson \(1988\)](#).²⁵

4 Role of Investment Network in Propagating Sector-Specific Shocks and Testable Implications for the Data

We now show the main implication of our concentrated investment network: sector-specific shocks to the investment hubs have large aggregate effects while shocks to non-hubs do not. We also show that key testable implications of this result hold in the data and relate the result to the investment-specific technology shock literature (e.g. [Greenwood, Hercowitz and Krusell \(2000\)](#) and [Justiniano, Primiceri and Tambalotti \(2010\)](#)).

Special Case of the Model To clarify the exposition, we use a simple $N = 2$ sector version of the model; we perform a quantitative analysis of the full model in Section 6. The network structure of this special case mirrors the structure of the empirical networks; sector 1 is the investment hub of the economy which produces all its investment goods ($\lambda_{11} = \lambda_{12} = 1$), but both sector 1 and sector 2 use each others’ goods in equal proportion ($1 - \theta_j = 0.4$, $\gamma_{11} = \gamma_{22} = 0$). We set the remaining parameters, which are unimportant to this discussion, to values similar to the full calibrated model ($\beta = 0.96$, $\xi_1 = 0.1$, $\delta = 0.10$, $\eta \rightarrow \infty$).

Each of the two sectors j has productivity

$$\log A_{jt} = \log A_t + \log \widehat{A}_{jt} \tag{7}$$

from sectors outside the U.S. Therefore, our model’s decision rules for factor demand correctly account for trade with the result of the world; however, the model counterfactually assumes that all factor supply is produced domestically. While extending our model to an open economy framework would be an interesting exercise, it is outside the scope of this paper. Our measured productivity shocks, discussed in Section 6, are derived from purely domestic sources; hence, foreign demand shocks do not directly enter the exogenous shocks that we feed into the model.

²⁵We use indivisible labor preferences due to our focus on employment fluctuations in the empirical analysis.

where A_t is an aggregate shock common across the two sectors and \widehat{A}_{jt} is an idiosyncratic shock independently distributed across the two sectors. We assume that these shocks follow AR(1) processes in logs, so that $\log A_t = \rho \log A_{t-1} + \varepsilon_t$ and $\log \widehat{A}_{jt} = \rho \widehat{A}_{jt-1} + \varepsilon_{jt}$. We set the persistence of the process to $\rho = 0.7$, similar to the average persistence of TFP across sectors in the data. For illustrative purposes, we set the volatility of the shocks to $\sigma(\varepsilon_t) = 0.01$ and $\sigma(\varepsilon_{jt}) = 0.01$.

4.1 Role of Investment Network in Propagating Shocks

We primarily analyze the effects of sector-specific shocks ε_{jt} on employment because their effects on value added are fairly mechanical. After optimizing its choice of intermediates, sector j 's value added is given by:

$$\log Y_{jt} = \frac{1}{\theta_j} \log A_{jt} + \alpha_j \log K_{jt} + (1 - \alpha_j) \log L_{jt} \text{ for } j \in \{1, 2\}. \quad (8)$$


The shock exogenously increases value added through its effect on the production function; by construction, this component only affects value added in sector j . In addition, the shock endogenously increases value added by increasing the use of the primary inputs K_{jt} and L_{jt} ; because capital K_{jt} is predetermined upon impact, these effects primarily operate through employment L_{jt} .

Main Result Figure 4 shows our main result: employment in either sector, and therefore in the aggregate, is substantially more responsive to the sector 1-specific shock than the sector 2-specific shock. While this result is quantitative in nature, it holds for a large set of parameter values. To explain the economic mechanisms underlying the result, we will use the following condition, which characterizes the economy's allocation of labor:

$$\xi_1 \frac{MPL_{1t}}{C_{1t}} = \chi (L_{1t} + L_{2t})^{\frac{1}{\eta}} = \xi_2 \frac{MPL_{2t}}{C_{2t}} \quad (9)$$

where MPL_{jt} is the marginal product of labor in sector j . This condition equates the marginal disutility of supplying labor to the market with the marginal product of labor

FIGURE 4: Impulse Responses of Employment to Sectoral Shocks, Two-Sector Model



Notes: response of sector-level employment to a sector-specific TFP shocks $\varepsilon_{jt} = 0.01$ in the two-sector model. We solve the model by linearization. See main text for description of the model.

in each sector times the marginal utility of consumption of that sector's output. In equilibrium, these marginal utilities are decentralized through the relative prices of the two goods $p_{jt} = \xi_j / C_{jt}$ and the wage $w_t = \chi (L_{1t} + L_{2t})^{\frac{1}{\eta}}$. In our quantitative work, we take the Frisch elasticity $\eta \rightarrow \infty$, so the marginal disutility of labor supply is a constant χ .

Role of Investment In order to understand the role of the investment network in driving this result, it is useful to first consider the model without investment. In this case, we show a stark irrelevance result; sector-specific shocks have literally zero effect on employment in either sector:

Theorem 1. *Suppose $\alpha_j = 0$ for all j , i.e. there is no capital in the economy and*


$$Q_{jt} = A_{jt} L_{jt}^{\theta_j} X_{jt}^{1-\theta_j}$$

$$Q_{jt} = C_{jt} + \sum_{i=1}^N M_{jit}$$

where $X_{jt} = \prod_{i=1}^N M_{ijt}^{\gamma_{ij}}$. Then employment in each sector L_{jt} is constant in response to aggregate or sector-specific shocks to A_{jt} .

Proof. See Appendix D. ■

FIGURE 5: Decomposing Effect of Sector 1-Specific Shock



Notes: response of sector-level variables to a sector 1-specific shock $\varepsilon_{1t} = 0.01$ in the two-sector model. We solve the model by linearization. Components of employment response are given the individual terms in the decomposition (10). Employment response is equal to the sum of the other responses.

In this benchmark model without investment, a shock generates no change in employment because general equilibrium forces completely offset the effects of the shock. This result is the consequence of two sets of assumptions. First, Cobb-Douglas production and preferences, together with the fact that consumption is the only final use of production, implies that consumption C_{jt} is proportional to gross output Q_{jt} in each sector. Second, under our growth-consistent preferences, this property implies that a shock leads to equally sized income and substitution effects on labor supply which leaves employment in (9) unchanged.

Our model with investment breaks this irrelevance result by weakening the dampening effects of general equilibrium, especially for investment hub shocks. Because consumption is not the only final use of production, consumption C_{jt} is no longer strictly proportional to gross output Q_{jt} . A sector-specific shock increases both investment and consumption, so consumption C_{jt} increases less than proportionally to gross output Q_{jt} . This fact then weakens the income effect relative to the substitution effect on labor supply, generating an increase in employment. Because this force is stronger for the investment hub sector 1, shocks to that sector have a larger effect on employment than shocks to sector 2.

Channels of Propagation To more deeply understand how a sector 1-specific shock is propagated in our model, Figure 5 rearranges (9) into:

$$d \log L_{jt} = \frac{1}{1 - \theta_j(1 - \alpha_j)} [-d \log C_{jt} + d \log A_{jt} + \alpha_j \theta_j d \log K_{jt} + (1 - \theta_j) d \log X_{jt}], \quad (10)$$

where “ $d \log$ ” denotes log-deviations from steady state. In the model without investment, employment is constant in response to the shock because $d \log C_{jt} = d \log A_{jt} + (1 - \theta_j) d \log X_{jt}$. Due to investment, we now have $d \log C_{jt} < d \log A_{jt} + (1 - \theta_j) d \log X_{jt}$, so employment increases.

The intermediates network generates spillovers from the sector 1-specific shock onto sector 2’s employment. These spillovers can be understood through the two first-order conditions for intermediates:


$$\frac{\xi_j}{C_{jt}} = MPX_{-jt} \frac{\xi_{-j}}{C_{-jt}} \quad (11)$$

where there is one equation for each $j \in \{1, 2\}$, MPX_{jt} is the marginal product of intermediates, and $-j$ denotes the sector other than j . These conditions equate the marginal cost of foregone consumption with the marginal benefit of using the output as an intermediate in the other sector. These two conditions can be used to understand an “upstream” and a “downstream” effect of the shock on intermediates usage. The upstream effect occurs because the marginal product of intermediates in sector 1, MPX_{1t} , increases, which then increases its use of intermediate inputs from sector 2; because employment is complementary to intermediates, this force further increases employment.²⁶ The downstream effect occurs because the marginal utility of consumption in sector 1 (and therefore the price of its output) falls, inducing sector 2 to use more intermediates and increasing its employment as well.

Figure 6 shows that the sector 2-specific shock generates larger changes in the marginal utility of consumption and therefore stronger general equilibrium dampening effects on employment. These changes also dampen the response of intermediates usage through the equations (11). In our quantitative model, individual non-hub sectors have an even smaller spillover onto aggregate investment, and therefore their shocks have an even smaller effect

²⁶Appendix D also shows that without investment, this upstream effect would be zero because the change in the marginal utility of consumption would offset the increase in the marginal product of intermediates.

FIGURE 6: Decomposing Effect of Sector 2-Specific Shock



Notes: response of sector-level variables to a sector 1-specific shock $\varepsilon_{2t} = 0.01$ in the two-sector model. We solve the model by linearization. Components of employment response are given the individual terms in the decomposition (10). Employment response is equal to the sum of the other responses.

TABLE 4
SIMULATION OF TWO-SECTOR MODEL

	Aggregate shocks ε_t	Sector 1 shocks ε_{1t}	Sector 2 shocks ε_{2t}
$\sigma(l_t)/\sigma(y_t)$	0.77	1.17	0.42
$\text{Corr}(y_t - l_t, y_t)$	0.88	-0.52	0.99

Notes: simulated business cycle statistics in the two-sector model. All variables have been logged and HP-filtered. y_t refers to aggregate value added and l_t refers to aggregate employment. “Aggregate shocks” refers to a simulation in which there are only aggregate shocks ($\sigma(\varepsilon_t) = 0.01$). “Sector 1 shocks” refers to a simulation in which there are only sector 1-specific shocks ($\sigma(\varepsilon_{1t}) = 0.01$). “Sector 2 shocks” refers to a simulation in which there are only sector 2-specific shocks ($\sigma(\varepsilon_{2t}) = 0.01$).

on employment.

Implications for Aggregate Labor Productivity Table 4 simulates the model in order to relate this mechanism to the dynamics of aggregate labor productivity. An aggregate productivity shock generates procyclical labor productivity essentially by construction; because the shock enters the production function of both sectors, it increases value added by more than employment in both sectors, increasing aggregate labor productivity. In contrast, a sector 1-specific shock only affects the production function in sector 1, but increases employ-

TABLE 5
SECTOR-LEVEL EMPLOYMENT VOLATILITY AND LABOR PRODUCTIVITY CYCLICALITY

	Investment Hubs	Non-Hubs
$\frac{\sigma(l_{st})}{\sigma(y_{st})}$	0.81	0.59
$\text{Corr}(y_{st} - l_{st}, y_{st})$	0.53	0.79

Notes: business cycle statistics at investment hubs and non-hubs. y_{st} is logged real value added at sector s , HP-filtered with smoothing parameter 6.25. l_{st} is logged real value added at sector s , HP-filtered with smoothing parameter $\lambda = 6.25$. “Investment hubs” compute the unweighted average the value of these statistics over $s =$ construction, machinery manufacturing, motor vehicles manufacturing, and professional/technical services. “Non-hubs” compute the unweighted average over the remaining sectors. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

ment in both sectors; therefore, aggregate employment responds by more than value added, decreasing labor productivity. The sector 2-specific shock has a small effect on aggregate employment relative to value added, generating procyclical labor productivity.


4.2 Testable Implications for the Data

We now examine three testable implications of our main result for the data.

Employment More Volatile at Hubs The first implication is that the response of employment to shocks is larger at investment hubs than at non-hubs (because shocks to investment hubs generate weaker general equilibrium dampening effects). Table 5 shows that the volatility of employment relative to the volatility of value added is approximately 1/3 higher at investment hubs than at non-investment hubs. Since employment is more volatile at investment hubs, labor productivity – value added per worker – should also be less correlated with value added at investment hubs. Table 5 shows that this is indeed the case; that correlation is more than 1/3 lower at investment hubs than non-hubs.

Hubs Forecasting Aggregate Employment Better than Non-Hubs The second implication we test is that value added growth at investment hubs should forecast future aggregate employment better than value added growth at non-hubs to the extent that value added growth reflects productivity shocks. We estimate a simple Jordà (2005)-style forecast-

FIGURE 7: Forecasting Power of Hubs vs. Non-Hubs for Aggregate Employment



Notes: results from estimating the forecasting regression (12) on our dataset. Left panel plots the coefficient γ_h on the growth rate of investment hubs’ value added growth as a function of the forecasting horizon h , together with a 95% confidence interval. Middle panel plots the coefficient β_h on the growth rate non-investment hubs’ value added together with a 95% confidence interval. Right panel plots the marginal increase in R^2 from including either term. “Due to hub” is the difference in R^2 between the joint regression and the regression with non-hubs only. “Due to non-hub” is the difference in R^2 between the joint regression and the regression with hubs only.

ing regression


$$\log N_{t+h} - \log N_t = \alpha_h + \gamma_h(\log y_{hub,t} - \log y_{hub,t-1}) + \beta_h(\log y_{non,t} - \log y_{non,t-1}) + \varepsilon_{t+h} \quad (12)$$

where $h = 1, 2, 3, 4$ is the forecasting horizon, N_t is aggregate employment, $y_{hub,t}$ is real aggregated value added across the non-hubs, and $y_{non,t}$ is real aggregated value added across the hubs. In order to make the coefficients interpretable, we standardize the growth rates of the two right-hand side variables. We estimate these forecasting regressions separately for the different forecasting horizons h .

Figure 7 shows that investment hubs’ value added has much stronger predictive power for aggregate employment than non-hubs’ value added. A one standard deviation increase in investment hubs’ value added growth predicts a persistent one percent increase in aggregate employment over the next few years. In contrast, non-hubs’ value added is a statistically insignificant predictor of aggregate employment, and including non-hubs in the regression has no impact on its R^2 .

Appendix E shows that investment hubs’ value added also predicts future aggregate

FIGURE 8: Comovement of Intermediate Suppliers with Investment Hubs




Notes: correlation of sector-level activity (logged and HP-filtered) with investment hubs' value added. The x-axis of each graphs computes, for each sector s , the share of gross output that is supplied to an investment hub as an intermediate, weighted by the share of investment produced by that hub. Left panel plots the correlation of sector s employment with the hubs' value added. Right panel plots the correlation of sector s intermediates production with the hubs' value added. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

employment better than aggregate GDP; in fact, aggregate GDP is statistically insignificant once investment hubs are included in the regression. Hence, the data suggest that all the predictive power of GDP for employment growth is driven by the investment hubs.

Intermediates Suppliers Comove Strongly with Hubs The third implication we test is that the intermediates network propagates investment hub shocks to other sectors through upstream effects, i.e. by increasing demand for intermediates at the hubs. We measure “supplier importance” as the average share of a given sector’s gross output that is sold as intermediate inputs to investment hubs. We then compute the correlation of sector-level employment or intermediates production with the investment hubs’ value added and study how this correlation depends on the measure of supplier importance.

Figure 8 shows that sectors which supply intermediates to investment hubs indeed comove more strongly with the hubs. The left panel shows a clear positive relationship between the supplier importance to the hubs and the correlation of sector-level employment with hubs’ value added; the R^2 of the regression line is approximately 23%. The right panel shows an even stronger positive relationship of supplier importance and the correlation of intermediates

FIGURE 9: Impulse Responses of Employment to Sector 1-Specific Shock, Our Two-Sector Model vs. “Investment-Specific Shock” Model



Notes: response of sector-level employment to a sector-specific TFP shocks $\varepsilon_{jt} = 0.01$ in the two-sector model. We solve the model by linearization. See main text for description of the model. Left panel plots response in our two-sector model. Right panel plots response in “Investment-specific shock” model, which refers to special case of our model in which sector 1’s output is not used for consumption ($\xi_1 = 0$) and there is no intermediate network ($\theta_j = 0$).

production with hubs’ value added; the R^2 is approximately 30%.

4.3 Comparison to Investment-Specific Shock Literature

The role of investment hub shocks in driving fluctuations in our model is reminiscent of the large literature on investment-specific technology shocks (see, for example, [Greenwood, Hercowitz and Krusell \(2000\)](#) or [Justiniano, Primiceri and Tambalotti \(2010\)](#)). In fact, that class of models can be viewed as a special case of our two-sector model in which there is no intermediate network ($\theta_j = 0$).²⁷ This literature studies whether idiosyncratic shocks to the investment-producing sector can generate large aggregate effects.

A key issue in this literature is that an investment-specific shock does not generate positive comovement between the consumption- and investment-producing sectors. The right panel of Figure 9 illustrates this problem in the special case of our model. Without the intermediates network, a sector 1-specific shock has literally zero effect on employment in sector

²⁷One could also set the consumption share of sector 1 to $\xi_1 = 0$, but this change has a small quantitative effect on the results.

TABLE 6
SIMULATION OF TWO-SECTOR MODEL VS. “INVESTMENT-SPECIFIC SHOCK” MODEL

	Sector 1-Specific Shocks Only	
	Baseline Model	IST Shocks Model
$\sigma(l_t)$	0.96%	0.56%
$\text{Corr}(y_t - l_t, y_t)$	-0.52	0.63

Notes: simulated business cycle statistics in the two-sector model with sector-specific shocks to sector 1 only. All variables have been logged and HP-filtered. y_t refers to aggregate value added and l_t refers to aggregate employment. “Baseline model” refers to two-sector model described in main text. “IST Shocks Model” refers to special case of baseline model in which sector 1’s output is not used for consumption ($\xi_1 = 0$) and there is no intermediate network ($\theta_j = 0$).

2.²⁸ Therefore, value added of the two sectors – corresponding to measured consumption and investment – will not comove either (similar to the [Barro and King \(1984\)](#) comovement problem). Due to this lack of comovement, the investment-specific shock also has a smaller effect on aggregate employment as well; in fact, [Table 6](#) shows that aggregate employment is about half as responsive to the shock as in our full model.

Our model solves the comovement problem through the intermediates network; the hub-specific shock also increases the supply of and demand for intermediates from the non-hub sector, raising that sector’s employment and value added.²⁹ [Table 6](#) shows that the hub shock generates twice as much fluctuations in aggregate employment in our model than in the model without intermediates; in fact, the fluctuations in employment are larger than those in GDP, generating countercyclical movement in labor productivity. In contrast, the investment-specific shocks literature uses other nominal or real rigidities overcome the negative comovement problem.³⁰

²⁸The fact that comovement is zero rather than negative reflects our use of an infinite Frisch elasticity $\eta \rightarrow \infty$; with finite Frisch η , [equation \(9\)](#) shows that the increase in sector 1-employment also increases the marginal disutility of supplying labor to sector 2, which would then decrease employment in sector 2 and generate negative comovement. See [Kim and Kim \(2006\)](#) for further discussion of the role of the Frisch elasticity η in determining sectoral comovement.

²⁹See [Hornstein and Praschnik \(1997\)](#) and [Ascari, Phaneuf and Sims \(2019\)](#) for related models which solve the “[Barro and King \(1984\)](#) curse” using roundabout production.

³⁰Another debate in this literature concerns how to measure the investment-specific technology shock. One approach is to use the price of investment goods relative to consumption goods; however, this price series is only weakly correlated with the aggregate cycle, so it is difficult to generate large fluctuations with it. In our model, investment-specific shocks can be directly measured as the productivity at investment hub sectors. In [Section 6](#), we show that these shocks generate substantial business cycle fluctuations.

5 Application: Changes in Business Cycles Since 1984

We now apply the insights developed in Section 4 to study the role of the investment network in determining the changes in business cycle patterns since 1984. The key driving force that sector-level shocks have become less correlated across sectors over time, which we document in Section 5.1. We show in Section 6 that this change, when propagated through the investment network, endogenously generates a number of well-known changes in business cycles patterns. We document those business cycle patterns in Section 5.2 and show two new results which support our explanation of them. First, the changes in business cycle patterns have not occurred within sector; instead, they are driven by changes in the comovement of activity across sectors. Second, the volatility of aggregate investment has risen relative to the volatility of GDP.

5.1 Changes in Sector-Level Productivity Shocks

We measure sector-level productivity as the Solow residual of value added net of the primary inputs:³¹

$$\log \tilde{A}_{jt} = \log Y_{jt} - \alpha_{jt} \log K_{jt} - (1 - \alpha_{jt}) \log L_{jt} \quad (13)$$

where Y_{jt} is value added.³² In our model, this value added-based productivity measure is isomorphic to the gross output-based productivity A_{jt} through the equation $\tilde{A}_{jt} = A_{jt}^{\theta_j}$. We measure productivity using the value added approach in order to make our results more comparable to existing literature. Abusing notation, we will refer to the value added-based productivity as “TFP” A_{jt} for the rest of the paper.

Of course, changes in the measured Solow residual may reflect changes in technology shocks or changes in other non-technology forces, such as allocational efficiency or the utilization of resources (see, for example, [Basu, Fernald and Kimball \(2006\)](#)). We view our simple exercise as a natural first step in quantifying the role of the investment network in

³¹We allow the factor shares α_{jt} to change year-by-year to ensure that changes in our measured productivity are not driven by changes in the production technology. This choice creates a slight inconsistency with our model, in which the factor shares are constant over time. Appendix G shows that our main model results are robust to allowing these parameters to change over time.

³²This equation assumes optimality of firm’s intermediates input choice in order to express value added as a function of the primary inputs capital and labor.

TABLE 7
DECOMPOSITION OF SHOCK VOLATILITY

	Measured TFP		Value Added	
	<i>Pre-84</i>	<i>Post-84</i>	<i>Pre-84</i>	<i>Post-84</i>
1000 $\text{Var}(x_t)$	0.21	0.07	0.52	0.19
Variances	0.04	0.03	0.06	0.04
Covariances	0.18	0.03	0.46	0.14

Notes: results of the decomposition (14) in the pre-1984 sample (1948-1983) and post-1984 sample (1984-2017). “Variances” refers to the variance component $\sum_{j=1}^N (\omega_{jt}^y)^2 \text{Var}(\log A_{jt})$, weighted by sector j ’s average value added share in the relevant subsample. “Covariances” refers to the covariance component $\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^y \omega_{ot}^y \text{Cov}(\log A_{jt}, \log A_{ot})$. “Measured TFP” refers to doing the analysis on HP-filtered log measured TFP $\log A_{jt}$. “Value added” refers to doing the analysis on HP-filtered (smoothing parameter $\lambda = 6.25$), log-value added y_{jt} . To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

propagating sector-specific shocks. The insights we develop here are relevant in the propagation of other non-technology shocks as well.³³

We need to detrend sector-level TFP because our model does not feature trend growth. However, a log-linear trend does not fit well to sector-level data because sectors typically grow and shrink in nonlinear ways. We take out a log-polynomial trend in order to capture these nonlinearities. We choose degree 4 in order to strike a balance between flexibility in the trend and not overfitting the data; Appendix C shows how various degrees fit the data and justifies our use of a fourth-order trend. Furthermore, Appendix G shows that our main results hold for other degrees of this polynomial trend. Foerster et al. (2019) study how these types of nonlinear trends aggregate to determine the aggregate growth rate in the economy and how that growth rate has changed over time.

The left panel of Table 7 characterizes how the covariance of TFP shocks have changed over time by performing the following statistical decomposition:

$$\text{Var}(\log A_t) = \underbrace{\sum_{j=1}^N (\omega_{jt}^y)^2 \text{Var}(\log A_{jt})}_{\text{Variances}} + \underbrace{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^y \omega_{ot}^y \text{Cov}(\log A_{jt}, \log A_{ot})}_{\text{Covariances}} \quad (14)$$

³³These is also a practical reason that we do not correct for utilization; consistent measures of hours-per-worker in each sector, which are required to perform the Basu, Fernald and Kimball (2006) correction, are not available in our data.

where $\log A_{jt}$ is HP-filtered log TFP. We compute this decomposition separately for the pre vs. post 1984 subsamples. The volatility of aggregate TFP has fallen by 2/3 since 1984, consistent with the “Great Moderation” of aggregate volatility. The entire decline in aggregate volatility is accounted for by a decline in the covariance of TFP across sectors; the within-sector component has remained comparatively stable.

We interpret these changes as reflecting a decline in the variance of aggregate shocks with a relatively stable variance of sector-specific shocks. A helpful special case of our shock process to develop that intuition is

$$\log A_{jt} = \log A_t + \log \hat{A}_{jt},$$

where A_t is an aggregate shock common to all sectors and \hat{A}_{jt} is independent across sectors. In this special case, the only source of covariance is the aggregate shock A_t , so the decline in covariances in the decomposition (14) maps directly into a decline in $\text{Var}(A_t)$. Appendix C performs a more general principal components analysis and yields a similar conclusion; the volatility of the first principal component – the “aggregate shock” – declines substantially since 1984 and accounts for the entire decline in volatility. Foerster, Sarte and Watson (2011) and Garin, Pries and Sims (2018) make a similar argument based on the comovement patterns of sector-level value added rather than measured productivity; the right panel of Table 7 shows that our results hold for value added as well.

5.2 Changes in Aggregate Business Cycle Patterns

The left panel of Table 8 documents the key changes in aggregate business cycle patterns that we will analyze using our model in Section 6. The volatility of GDP is 40% lower in the post-1984 sample than in the pre-1984 sample (the “Great Moderation”). The cyclicity of labor productivity, measured as the correlation of HP-filtered GDP per worker with HP-filtered GDP, fell by more than 50% in the post-1984 sample.³⁴ In addition, the volatility

³⁴The post-1984 cyclicity of labor productivity in our data is higher than in other datasets. There are at least three possible reasons for this difference. First, we measure labor productivity using output per worker, while many studies use output per hour. Second, our data measures value added using income-side accounting, while other data often uses expenditure-side accounting. Third, our data includes nonprofit organizations, while other data typically includes on nonfarm business companies.

TABLE 8
CHANGES IN BUSINESS CYCLE PATTERNS

	Aggregated		Within-Sector	
	<i>Pre-1984</i>	<i>Post-1984</i>	<i>Pre-1984</i>	<i>Post-1984</i>
$\sigma(y_t)$	2.27%	1.36%	3.58%	3.00%
$\text{Corr}(y_t - l_t, y_t)$	0.65	0.26	0.73	0.71
$\sigma(l_t)/\sigma(y_t)$	0.75	1.02	0.65	0.65
$\sigma(i_t)/\sigma(y_t)$	1.94	2.91	2.76	2.84

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). y_t is log value added, l_t is log employment, and i_t is log investment. “Aggregated” aggregates value added and investment across sectors using a Tornqvist index weighted by nominal value added shares, aggregates employment as the simple sum, HP-filters each series with smoothing parameter $\lambda = 6.25$, and computes the statistics. “Within-sector” HP-filters each sector-level series with smoothing parameter $\lambda = 6.25$, computes the statistics, and then averages them weighted by the average share of nominal value added within that sub-sample. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

of both employment and investment rose by approximately 1/3 relative to the volatility of GDP. To our knowledge, we are the first to note the increased relative volatility of investment. Appendix F shows that the rising volatility of employment accounts for the entire decline in the cyclicity of labor productivity; intuitively, since employment and GDP are highly correlated in both subsamples, the time-series behavior of their ratio depends on the more volatile component.³⁵ The fact both employment and investment became more volatile is consistent with the idea that shocks to investment hubs account for a larger share of fluctuations in both variables since 1984.

The right panel of Table 8 shows that these changes in aggregate cycles do not occur within the average sector of the economy. The cyclicity of sector-level labor productivity – the correlation of sector-level value added per worker with sector-level value added – and the relative volatilities of sector-level employment and investment are essentially constant across the two sub-samples. While the volatility of sector-level value added does fall post-

³⁵One can see the source of this result using the identity (derived in Appendix F):

$$\text{Corr}(y_t, y_t - l_t) = \frac{1 - \frac{\sigma(l_t)}{\sigma(y_t)} \text{Corr}(y_t, l_t)}{\sqrt{1 + \frac{\sigma(l_t)^2}{\sigma(y_t)^2} - 2 \frac{\sigma(l_t)}{\sigma(y_t)} \text{Corr}(y_t, l_t)}}. \quad (15)$$

Since output and employment are highly correlated both before and after 1984, the decline in the cyclicity of labor productivity is driven by the increase in the relative volatility of employment.

1984, its magnitude is about half as large as the decline in the volatility of GDP. Appendix F shows that these findings are robust to using various weighting schemes to compute the within-sector average and to using first-differences to detrend the data.

Since the changes in the aggregate cycle do not occur within sector, they must be driven by changes in the covariances of activity across sectors. We formalize this argument using the following decomposition (derived in Appendix F):

$$\frac{\mathbb{V}ar(l_t)}{\mathbb{V}ar(y_t)} \approx \underbrace{\omega_t}_{\text{variance weight}} \underbrace{\frac{\sum_{j=1}^N (\omega_{jt}^l)^2 \mathbb{V}ar(l_{jt})}{\sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt})}}_{\text{variances}} + (1 - \omega_t) \underbrace{\frac{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^l \omega_{ot}^l \text{Cov}(l_{jt}, l_{ot})}{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^y \omega_{ot}^y \text{Cov}(y_{jt}, y_{ot})}}_{\text{covariances}} \quad (16)$$

where y_{jt} is HP-filtered log value added of sector j , l_{jt} is HP-filtered log employment of sector j , and y_t and l_t are aggregate value added and employment. We present this decomposition for the relative volatility of employment because its rise accounts for the declining cyclical volatility of labor productivity. The decomposition (16) breaks the variance of employment relative to the variance of GDP into two components. The first “variances” component is the average variance of employment relative to the average variance of value added within sectors. The second “covariances” component is the average covariance of employment across all pairs of sectors relative to the average covariance of value added across pairs. The “variance weight” $\omega_t = \sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt}) / \mathbb{V}ar(y_t)$ ensures that the averages of these ratios add up to the ratio of aggregate variances.

Table 9 shows that 90% of the increase in the relative volatility of aggregate employment is accounted for by an increase in the covariances term; in contrast, the within-sector average variances are stable, consistent with the results in Table 8. Appendix F shows that the changes in covariances reflect two patterns in the data. First, the covariance of value added across sectors fell in the post-1984 sample, decreasing the volatility of aggregate GDP. In contrast, the covariance of employment across sectors remained comparatively stable, stabilizing its aggregate volatility and therefore raising its volatility relative to output. Appendix F shows that similar results hold when decomposing the volatility of investment; approximately 80% of the increase in that relative volatility is accounted for by the covariance terms.

TABLE 9
DECOMPOSITION OF RELATIVE EMPLOYMENT VOLATILITY

	<i>Pre-84</i>	<i>Post-84</i>	<i>Contribution of entire term</i>
$\frac{\mathbb{V}ar(l_t)}{\mathbb{V}ar(y_t)}$	0.57	0.94	100%
Variances	0.40	0.39	13%
Covariances	0.59	1.10	87%
Variance Weight	0.11	0.23	
$(\omega_t = \sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt}) / \mathbb{V}ar(y_t))$			

Notes: results of the decomposition (16) in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). “Variances” refers to the variance component $\frac{\sum_{j=1}^N (\omega_{jt}^l)^2 \mathbb{V}ar(l_{jt})}{\sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt})}$. “Covariances” refers to the covariance component $\frac{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^l \omega_{ot}^l \text{Cov}(l_{jt}, l_{ot})}{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^y \omega_{ot}^y \text{Cov}(y_{jt}, y_{ot})}$. “Variance weight” refers to the weighting term $\omega_t = \sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt}) / \mathbb{V}ar(y_t)$. “Contribution of entire term” column computes the contribution of the first term of the decomposition (16) (in the variance row) and the contribution of the second term (in the covariance row). To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

We conclude from this evidence that a good explanation for the declining cyclical volatility of aggregate labor productivity should be driven by changes in sectoral comovement, not changes occurring within sectors. Our model naturally generates these facts because the volatility of sector-specific shocks has remained relatively constant, but the volatility of aggregate shocks has fallen and generated the changes in comovement patterns across sectors. In contrast, nearly all existing explanations for the declining cyclical volatility of labor productivity in the literature abstract from sectoral heterogeneity, so they do not speak to these empirical results.

Appendix F contains five additional pieces of analysis of this decomposition in order to ensure that the results are robust features of the data. First, it shows that the changes in covariance patterns we discuss are broad-based and not driven by outliers. Second, it shows that the results also hold using first-differences rather than the HP-filter to detrend the data. Third, it shows that the changes in covariances are reflected in changes in correlations, rather than variances. Fourth, it shows that the approximation inherent in the decomposition (16) is accurate. Fifth, it shows that the results of this decomposition also hold for the finer 450-sector partition of manufacturing in the NBER-CES database.

6 Quantitative Analysis of Changes Since 1984

We now argue that the rising importance of sector-specific shocks caused the changes in business cycle patterns using by feeding in our measured sector-level productivity series into the model. We assume all the other parameters are fixed over time in order to isolate the role of the shocks in driving the changes in business cycle patterns. However, Appendix G shows that our results are similar if we allow those parameters to change over time, indicating that changes in the shock process is the key change over this period.

We use the following procedure to feed the realized TFP shocks into our model. We first estimate the persistence ρ_j using maximum likelihood over the entire sample. These parameters, along with the others parameters calibrated in Section 3, are sufficient to compute the decision rules in our model because the decision rules do not depend on the covariance matrix of shocks (due to certainty equivalence of the linearized solution). We simulate the decision rules given the realized history of shocks, starting from the non-stochastic steady state.³⁶

6.1 Model Matches Changes in Aggregate Cycles

Table 10 shows that the model matches key business cycle patterns both before and after the 1984 breakpoint. In our model, the cyclicalities of aggregate labor productivity falls by 0.45 points in the post-1984 sample, compared to 0.39 points in the data.³⁷ Consistent with that fact, the standard deviation of aggregate employment relative to GDP increases by approximately 25% in our model, compared to a 34% increase in the data (recall that the rise in this relative volatility accounts for the decline in labor productivity cyclicalities). The relative volatility of investment also increases, consistent with the rising importance of

³⁶An alternative approach is to estimate the covariance matrix of innovations separately for the pre vs. post 1984 subsamples of our data and then compute the implied population moments of the model. However, we cannot estimate full-rank covariance matrices since the number of sectors is larger than the number of time-series observations in the two subsamples. In Appendix G, we collapse the partition of sectors to $N = 28$, which allows us to estimate full rank covariances matrices and compute population moments corresponding to the two subperiods. Those results are similar to the results in the main text.

³⁷The pre-1984 level of labor productivity cyclicalities is substantially higher in our model than the data. Other mechanisms could decrease the overall level, such as other aggregate shocks or even measurement error in the data, without affecting our conclusions here.

TABLE 10
CHANGES IN BUSINESS CYCLE PATTERNS, MODEL VS. DATA

	Data		Model	
	<i>Pre-1984</i>	<i>Post-1984</i>	<i>Pre-1984</i>	<i>Post-1984</i>
$\sigma(y_t)$	2.27%	1.36%	2.60%	2.24%
$\rho(y_t - l_t, y_t)$	0.65	0.26	0.90	0.45
$\sigma(l_t)/\sigma(y_t)$	0.76	1.02	0.74	0.92
$\sigma(i_t)/\sigma(y_t)$	1.94	2.91	3.92	4.67

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). “Data” refers to our empirical dataset (and replicates the results of Table 8). “Model” refers to model simulation starting from steady state and feeding in realizations of measured TFP over the sample. y_t is log GDP, l_t is log aggregate employment, and i_t is log aggregate investment. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.


investment hub shocks.³⁸

Figure 10 shows that the model also matches the timing of the decline in the cyclical-ity of labor productivity. We compute the dynamics of this statistic using 14-year forward-looking rolling windows in both the data and in our model. The two series track each other quite closely. The cyclical-ity of labor productivity is fairly stable until the early 1980s, when it drops sharply after the Volcker recession. The cyclical-ity then recovers somewhat in the 1990s but then drops again during the 2008 financial crisis and its aftermath. By the end of the sample, the cyclical-ity of labor productivity has fallen by a similar amount in the model and in the data.

Role of Investment Network Table 11 compares a simulation of the model with shocks only to the investment hubs to a simulation with shocks only to the non-investment hubs. The hub shocks account for a larger share of aggregate fluctuations in the post-1984 data and generate countercyclical fluctuations in aggregate labor productivity. Hence, a reallocation of volatility to investment hub-specific shocks will decrease the cyclical-ity of aggregate labor

³⁸The model only generates a 0.36 percentage point decline in the volatility of GDP compared to a 0.91pp decline in the data. We therefore conclude that changes in the shock process alone only capture approximately 40% of the “Great Moderation” of aggregate volatility. Appendix G shows that allowing the other parameters of the model to change accounts for the remaining decline in aggregate volatility without significantly affecting our other results.

FIGURE 10: Dynamics of Labor Productivity Cyclicity Over Time



Notes: rolling windows of $\text{Corr}(y_t - l_t, y_t)$ where y_t is HP-filtered log aggregate value added and l_t is HP filtered log aggregate employment (each filtered with smoothing parameter $\lambda = 6.25$). Rolling windows have length of 14 years and are forward-looking (e.g. 1950 data point computes the correlation between 1950-1963). “Data” corresponds to aggregated version of our dataset. “Model” corresponds to aggregated version of model simulation under measured realizations of sector-level TFP shocks. Both series have been normalized to zero in 1950. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

TABLE 11
ROLE OF INVESTMENT HUB SHOCKS IN CHANGING BUSINESS CYCLES

	Non-Hub Shocks Only		Hub Shocks Only	
	<i>Pre-84</i>	<i>Post-84</i>	<i>Pre-84</i>	<i>Post-84</i>
$\sigma(y_t)$	1.83%	1.29%	0.87%	1.35%
$\sigma(l_t)$	1.23%	1.04%	0.82%	1.37%
$\text{Corr}(y_t - l_t, y_t)$	0.93	0.68	0.41	-0.11

Notes: business cycle statistics in our model in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). “Hub Shocks Only” refers to setting all non-investment hub shock realizations to zero. “Non-Hub Shocks Only” refers to setting all the investment hub shock realizations to zero. y_t is log GDP and l_t is log aggregate employment.

productivity, as shown above.^{39,40}

Figure 11 further decomposes the effects of sector-specific shocks for individual sectors, obtained by separately feeding in measured TFP realizations for each individual sector. The top panel plots the standard deviation of aggregate employment relative to the standard deviation of the shock, which can be loosely interpreted as a reduced-form elasticity of aggregate employment to a sector-specific shock. The four investment hubs have high elasticities due to their large effects on aggregate employment described Section 4. Also consistent with that discussion, shocks to sectors which supply intermediates to investment hubs – primarily the manufacturing sectors in the left of figure – generate larger effects on employment than other non-hub sectors.⁴¹ The bottom panel of the figure plots the cyclicalities of aggregate labor productivity induced by shocks to each of these sectors. A shock to most investment hubs generate countercyclical fluctuations in labor productivity because they increase aggregate employment by more than GDP. The exception is professional/technical services, which generates procyclical labor productivity; however, those fluctuations are substantially less procyclical than other service sectors.⁴²

We therefore conclude that changes in the process for sectoral TFP shocks, when filtered through the investment network, are the key driver of these changes in business cycle patterns over time. Much of the literature has argued that the declining cyclicalities of labor productivity implies that TFP shocks are unimportant in explaining business cycle fluctuations. Our results suggest caution in that interpretation. In our model, the declining cyclicalities reflect changes in how sector-level TFP shocks are propagated through the economy’s investment network, not that they are unimportant altogether.


³⁹The cyclicalities of labor productivity in response to non-hub shocks also falls in the post-1984 sample. This finding reflects the fact that some non-hubs are still investment producers and that shocks to these sectors and shocks to suppliers of investment hubs also become more important over time; see Figure 11 below for further discussion.

⁴⁰The hub shocks generate procyclical labor productivity in the pre-1984 sample because shocks to professional/technical services account for a larger share of those fluctuations in that period, and labor productivity is procyclical in response to shocks to professional/technical services. See Figure 11 below for further discussion.

⁴¹Shocks to the wholesale trade sector also have a large effect on employment; Appendix G shows that this result is because wholesale trade is a large share of value added.

⁴²Professional/technical services generates procyclical labor productivity because it also supplies intermediate goods to many sectors in the economy. Therefore, a professional/technical services-specific shock generates a larger increase in value added than in employment, increasing labor productivity.

FIGURE 11: Aggregate Dynamics Driven by Shocks to Individual Sectors



Notes: results from simulating model with empirical shocks to only one sector at a time (the remaining sectors' shocks are set to zero). Top panel: standard deviation of aggregate employment n_t relative to the standard deviation of the sectors' productivity A_{jt} . Bottom panel: the correlation of aggregate labor productivity with aggregate GDP, $\text{Corr}(y_t - l_t, y_t)$. All variables have been logged and HP-filtered with smoothing parameter 6.25. Investment hubs are highlighted in red.

TABLE 12
DIVERGENCE OF AGGREGATE AND SECTORAL CYCLES

Data	<i>Aggregated</i>		<i>Within-Sector</i>	
	<i>Pre-1984</i>	<i>Post-1984</i>	<i>Pre-1984</i>	<i>Post-1984</i>
$\sigma(y_t)$	2.27%	1.36%	3.58%	3.00%
$\rho(y_t - l_t, y_t)$	0.65	0.26	0.73	0.71
$\sigma(l_t)/\sigma(y_t)$	0.76	1.02	0.65	0.65
Model				
$\sigma(y_t)$	2.60%	2.24%	4.03%	4.18%
$\rho(y_t - l_t, y_t)$	0.90	0.45	0.82	0.80
$\sigma(l_t)/\sigma(y_t)$	0.74	0.92	0.48	0.51

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). “Data” refers to our empirical dataset (and replicates the results of Table 8). “Model” refers to model simulation starting from steady state and feeding in realizations of measured TFP over the sample. y_t is log value added and l_t is log employment. “Aggregated” aggregates value added across sectors using a Tornqvist index weighted by nominal value added shares, aggregates employment as the simple sum, HP-filters both series with smoothing parameter $\lambda = 6.25$, and computes the statistics. “Within-sector” HP-filters each sector-level series with smoothing parameter $\lambda = 6.25$, computes the statistics, and then averages them weighted by the average share of nominal value added within that sub-sample. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

6.2 Changes in Cycles Driven by Changes in Comovement

We now show that the changes in business cycle patterns are driven by changes in the sectoral comovement and not changes within sectors, as we documented in Section 5. We focus on the cyclicity of labor productivity and relative volatility of employment. Table 12 shows that the changes in business cycle patterns do not occur within sectors of our model, consistent with the data; both sector-level labor productivity and the relative volatility of sector-level employment are stable over both subsamples. In our model, the sector-level patterns are stable because sector-specific shocks are the dominant source of fluctuations within sector and the volatility of these shocks has remained stable over time. However, the propagation of these shocks to other sectors has become relatively more important post-1984 because aggregate shocks become less volatile. The resulting changes in sectoral comovement then account for the changes in aggregate cycles.

Table 13 formalizes this argument by replicating the decomposition (16) on our model-

TABLE 13

DECOMPOSITION OF RELATIVE EMPLOYMENT VOLATILITY, MODEL VS. DATA

	Data			Model		
	Pre-84	Post-84	Contribution of entire term	Pre-84	Post-84	Contribution of entire term
$\frac{\text{Var}(l_t)}{\text{Var}(y_t)}$	0.57	0.94	100%	0.55	0.84	100%
Variations	0.40	0.39	13%	0.47	0.47	11%
Covariances	0.59	1.10	87%	0.56	0.92	89%
Variance Weight ($\omega_t = \sum_{j=1}^N (\omega_{jt}^y)^2 \text{Var}(y_{jt}) / \text{Var}(y_t)$)	0.11	0.23		0.11	0.18	
	Model, No Investment Net.					
	Pre-84	Post-84	Contribution of entire term			
$\frac{\text{Var}(l_t)}{\text{Var}(y_t)}$	0.55	0.59	100%			
Variations	0.45	0.39	250%			
Covariances	0.57	0.69	-150%			
Variance Weight	0.15	0.33				

Notes: results of the decomposition (16) in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). “Data” refers to our empirical dataset (and replicates the results of Table 8). “Model” refers to model simulation starting from steady state and feeding in realizations of measured TFP over the sample. “Model, No Investment Net.” refers to version of the model in which we have eliminated the investment network by assuming sectors accumulate capital out of their own output only. “Variations” refers to the variance component $\frac{\sum_{j=1}^N (\omega_{jt}^l)^2 \text{Var}(l_{jt})}{\sum_{j=1}^N (\omega_{jt}^y)^2 \text{Var}(y_{jt})}$. “Covariances” refers to the covariance component $\frac{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^l \omega_{ot}^l \text{Cov}(l_{jt}, l_{ot})}{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^y \omega_{ot}^y \text{Cov}(y_{jt}, y_{ot})}$. “Variance weight” refers to the weighting term $\omega_t = \sum_{j=1}^N (\omega_{jt}^y)^2 \text{Var}(y_{jt}) / \text{Var}(y_t)$. “Contribution of entire term” column computes the contribution of the first term of the decomposition (16) (in the within-sector row) and the contribution of the second term (in the between-sector row). To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

simulated data; for convenience, the decomposition is

$$\frac{\text{Var}(l_t)}{\text{Var}(y_t)} \approx \underbrace{\omega_t}_{\text{variance weight}} \underbrace{\frac{\sum_{j=1}^N (\omega_{jt}^l)^2 \text{Var}(l_{jt})}{\sum_{j=1}^N (\omega_{jt}^y)^2 \text{Var}(y_{jt})}}_{\text{variances}} + (1 - \omega_t) \underbrace{\frac{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^l \omega_{ot}^l \text{Cov}(l_{jt}, l_{ot})}{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^y \omega_{ot}^y \text{Cov}(y_{jt}, y_{ot})}}_{\text{covariances}}$$

In the data, the covariance terms account for 87% of the increase in the relative volatility of employment (and, therefore, the decrease in the cyclicity of labor productivity); in our model, the covariance terms account for 89% of the increase. The changes in these covariance terms reflect two patterns in both the model and the data. First, the covariance of value added

TABLE 14
AVERAGE PAIRWISE CORRELATIONS, MODEL VS. DATA

	Data		Model	
	<i>Employment</i>	<i>Value added</i>	<i>Employment</i>	<i>Value added</i>
Pre-1984	0.55	0.36	0.88	0.35
Post-1984	0.51	0.17	0.84	0.19
<i>Difference</i>	-0.04	-0.19	-0.04	-0.17
Model, no investment net.				
	<i>Employment</i>	<i>Value added</i>		
Pre-1984	0.39	0.28		
Post-1984	0.20	0.10		
<i>Difference</i>	-0.19	-0.18		

Notes: average pairwise correlations (17). “Pre-1984” computes ρ_τ^x in the 1948-1983 subsample and “post-1984” computes ρ_τ^x in the 1984-2017 subsample. “Data” refers to the data, “Model” to the model, and “Model, no investment net.” to a version of the model in which the investment network has been eliminated (by assuming that all investment is done using own-sector output). To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

across sectors falls after 1984 because the covariance of productivity falls; this fact decreases the volatility of aggregate GDP. Second, the covariance of employment across sectors is stable over this period, which stabilizes the volatility of aggregate employment. Together, these two facts drive up the relative volatility of employment and therefore drive down the cyclical volatility of labor productivity.


Table 14 further studies these comovement patterns by computing the change in the average correlation of value added and employment across pairs of sectors:

$$\rho_\tau^x \equiv \frac{\sum_{i=1}^N \sum_{j=i+1}^N \omega_i^x \omega_j^x \text{Corr}(x_{jt}, x_{it} | t \in \tau)}{\sum_{i=1}^N \sum_{j=i+1}^N \omega_i^x \omega_j^x} \quad (17)$$

where x_{jt} is either employment or value added and ω_j are value added or employment shares. The correlation of value added falls nearly in half, generating most of the decline in the covariances in the decomposition (16); in contrast, the correlation of employment is essentially stable, generating the stability of the between sector covariances as well.⁴³ To our

⁴³The fact that the correlation of employment across sectors is higher in our model than the data is driven by our choice of an infinite Frisch elasticity $\eta \rightarrow \infty$; this assumption implies that the marginal disutility of labor supply is constant, so an increase in one sector’s employment does not affect the incentives to supply labor to other sectors. With a finite Frisch elasticity $\eta < \infty$, an increase in one sector’s employment increases

FIGURE 12: Model Fit of Sector-Pair Level $\Delta\text{Cov}(l_{jt}, l_{ot}) - \Delta\text{Cov}(y_{jt}, y_{ot})$ ($R^2 = 27\%$)



Notes: model fit to sector-pair (j, o) value of $\Delta\text{Cov}(l_{jt}, l_{ot}) - \Delta\text{Cov}(y_{jt}, y_{ot})$, where $\Delta\text{Cov}(l_{jt}, l_{ot})$ is the covariance of log HP-filtered employment in the post-1984 sample relative to the pre-1984 sample, and $\Delta\text{Cov}(y_{jt}, y_{ot})$ is the covariance of log HP-filtered value added in the post-1984 sample relative to the pre-1984 sample. Horizontal axis is the value of that statistic in the data while the vertical axis is the value in the model. The solid line is the regression line across all sectors, which has an R^2 of 0.27. In the plot, circle size is proportional to the product of the pair's share of value added over the entire sample. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

knowledge, our model is the only explanation for the declining cyclicity of aggregate labor productivity that is consistent with these facts in the data.

Tables 13 and 14 also show that the investment network is crucial to generating stable employment comovement and, therefore, the declining cyclicity of labor productivity over this period. We eliminate the investment network by assuming that all sectors invest out of their own output, i.e. the network is the identity matrix. In this case, the correlation of employment across sectors counterfactually falls by as much as the correlation of value added. Therefore, aggregate employment volatility falls by as much as GDP volatility, and there is no decline in the cyclicity of labor productivity (it slightly falls from 0.93 to 0.82).

Finally, Figure 12 shows that our model provides a good fit of not only the average change

the disutility of supply labor to other sectors, decreasing the level of employment comovement. However, allowing for a finite Frisch still implies that the correlation of employment across sectors is constant over time (in results not reported). We focus on the infinite Frisch as our baseline in order to focus on changes in employment; Appendix G shows that our main results are robust to allowing for $\eta < \infty$.

in covariances over this period, but the changes at the sector-pair level as well. We summarize the sector-pair level change with the “diff-in-diff” $\Delta\text{Cov}(l_{jt}, l_{ot}) - \Delta\text{Cov}(y_{jt}, y_{ot})$. On average, this object is positive because employment covariances change by less than the value added covariances, and larger values correspond to a larger divergence between employment and value added covariances over time. Although neither of these objects were targeted in the calibration, the model explains 27% of the variation in the data.

Appendix G contains a number of additional exercises to show that these model results are robust. First, we allow the other parameters of the model – those governing the production structure, networks, and consumption shares – to change over time.⁴⁴ Second, we estimate the covariance matrix of shocks and compute the model’s population moments over the two sub-samples. Third, we use alternative degrees of polynomial trends when measuring TFP. Fourth, we allow for adjustment costs of capital and maintenance investment in the computation of the investment network. Our results continue to hold in all of these cases.

7 Implications of Network for Stimulus Policy

The analysis so far has focused on how the investment network propagates sector-specific productivity shocks; we now briefly study how it propagates investment stimulus policies, such as investment tax credits or bonus depreciation allowance. We model investment stimulus as an exogenous shock to the cost of capital:

$$(1 - sub_t) \times \nu_{jt}, \tag{18}$$

where ν_{jt} is the marginal cost of producing investment goods and sub_t is the policy shock. Winberry (2018) shows that a number of actual policies map into this reduced-form shock.⁴⁵

We assume that the policy shock is financed from outside the economy in order to focus on

⁴⁴An alternative approach would be to estimate a process for the changes in parameters over time, and use the extended path algorithm from Maliar et al. (2015) to compute the changes in business cycle fluctuations. While that approach has merit, we do not pursue it here for the sake of simplicity.

⁴⁵The key intuition behind this result is that, without financial frictions, the present value of tax savings per unit of investment is a sufficient statistic to capture the effects of these policies on investment.

TABLE 15
EFFECTS OF 1% INVESTMENT PURCHASE SUBSIDY

	Baseline	No intermediates
Δi_t	8.82%	8.52%
Δn_t	1.77%	1.68%
Δn_t^{hubs}	3.93%	4.65%
$\Delta n_t^{\text{non-hubs}}$	1.28%	0.67%

Notes: effect of a one-time $sub_t = 0.01$ shock to the stimulus policy shock described in the main text. “Baseline” refers to full model and “No intermediates” refers to model without intermediate goods (i.e. $\theta_j = 1$ for all sectors j). Δi_t is the percentage change in aggregate investment, Δn_t is the percentage change in aggregate employment, Δn_t^{hubs} is the percentage change in employment at the investment hubs, and $\Delta n_t^{\text{non-hubs}}$ is the percentage change in employment at the non-hubs.

how it affects investment incentives.⁴⁶


Table 15 shows that the investment stimulus increases employment in many sectors of the economy. A 1% subsidy shock increases aggregate investment by nearly 9%. Most of this increased investment is produced by investment hubs, whose employment increases by about 4%. Employment at non-hubs also increases by about 1.3% in order to supply intermediates to the investment hubs (similar to the effects of hub-specific productivity shocks described above). The right column of Table 15 shows that, without these spillovers from the intermediates network, employment at the non-hubs increases by about half as much.⁴⁷ Hence, the intermediates network propagates the effects of the stimulus throughout the economy.

Figure 13 shows that the effects of the stimulus shock on employment are unevenly distributed across sectors of the economy. Over 40% of the increase in aggregate employment is concentrated in the four investment hubs because they produce the majority of investment. There is also a sizable increase in employment in the sectors which supply intermediates to investment hubs (primarily the manufacturing sectors in the left of the figure). However, the sectors which do not supply intermediates to the hubs see virtually no change in their employment. Figure 13 also shows that, in a counterfactual version of the model without

⁴⁶Of course, since the equilibrium of our model is efficient, all stimulus policies are strictly welfare-reducing. We think our model nevertheless provides useful insights about the positive effects of these policies. These effects will be important forces in a normative exercise using richer models in which the policies may be welfare-improving.

⁴⁷The fact that non-hubs’ employment increases even without the intermediates network reflects the fact that they also produce some investment goods.

FIGURE 13: Distributional Effects of Investment Stimulus



Notes: effect of a one-time $sub_t = 0.01$ shock to the stimulus policy shock described in the main text. Each bar plots the change in employment at that particular sector, divided by the change in aggregate employment. The sum of all the bars equals 100% of the change in aggregate employment. Red bars are the investment hubs’ response in our baseline model, blue bars are the non hubs’ response in our baseline model, and transparent grey bars are the responses in a version of the model in which we eliminate the investment network by assuming all investment is done out of own-sector output.

the investment network, the effect of the policy is more uniformly distributed across sectors.⁴⁸ Hence, the sparseness of the investment network implies that investment stimulus policies have very unequal effects on employment across various sectors in the economy – resembling industrial policy.⁴⁹ These distributional implications occur despite the fact that policy subsidizes the purchases of investment equally across sectors, and may be seen as a negative consequence to policymakers who wish to avoid the appearance of conducting industrial policy.⁵⁰

⁴⁸Without the investment network, the service sectors (in the right of the plot) account for a larger share of the aggregate response than the non-service sectors. This result simply reflects the fact that service sectors are larger and therefore mechanically account for a larger share of employment fluctuations; Appendix G shows that the percentage change in employment within sectors, which is not mechanically related to size, is fairly uniformly distributed across sectors in the model without the investment network.

⁴⁹These uneven effects across sectors are minimized by our choice of an infinite Frisch elasticity $\eta \rightarrow \infty$; as discussed in footnote 43, a finite Frisch elasticity $\eta < \infty$ implies that an increase in one sector’s employment increases the disutility of supplying labor to other sectors, increasing the dispersion of employment growth across sectors.

⁵⁰House, Mocanu and Shapiro (2017) build a model in which investment stimulus has a different effect on the purchases and production of investment due to imported investment goods. While our model does not

8 Conclusion

In this paper, we have argued that the investment network plays an important role in propagating sectoral shocks into aggregate fluctuations. Our argument had three main components. First, we showed that the empirical investment network is dominated by four investment hubs that produce the majority of investment goods, are highly volatile at business cycle frequencies, and are strongly correlated with the aggregate cycle. Second, we embedded this network into a standard multisector business cycle model and showed that shocks to the investment hubs have strong propagation onto other sectors and, therefore, onto aggregates. Third, we measured sector-level productivity shocks in the data, fed them into the model, and found that shocks to investment hubs accounted for a large and increasing share of aggregate fluctuations. Since shocks to these hubs generate more volatile movement in employment than value added, this shift accounts for the decline in the cyclicalities of aggregate labor productivity and other changes in business cycle patterns since the early 1980s. We also showed that investment stimulus policies, applied equally across sectors, primarily increase production at investment hubs and their intermediate goods suppliers but not in other sectors of the economy.

In order to isolate the role of the investment network, we have embedded it into a purposely simple multisector real business cycle model. A natural next step would be to add the rich set of nominal and real rigidities which the DSGE literature has argued are relevant for business cycle analysis. We kept our quantitative exercise simple by focusing on sector-level productivity shocks measured as a simple Solow residual. While we do not think that the role of the investment network as a propagation mechanism is specific to productivity shocks – other non-technology shocks may have similar effects – another next step would be to understand what drives the variation in our measured shocks and incorporate other shocks. Our analysis of investment stimulus policy was purely positive because our model does not feature the externalities which motivate the use of these policies in a first place. A natural next step for policy analysis would be to incorporate these frictions and perform a normative analysis. The insights developed in our model would be useful for that exercise.

incorporate imports, it allows us to study the distributional effects of the policy across sectors within the domestic economy.

References

- ACEMOGLU, D., V. M. CARVALHO, A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2012): “The network origins of aggregate fluctuations,” *Econometrica*, 80(5), 1977–2016.
- ACEMOGLU, D., A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2017): “Microeconomic origins of macroeconomic tail risks,” *American Economic Review*, 107(1), 54–108.
- ASCARI, G., L. PHANEUF, AND E. SIMS (2019): “Can New Keynesian Models Survive the Barro-King Curse?,” .
- ATALAY, E. (2017): “How important are sectoral shocks?,” *American Economic Journal: Macroeconomics*, 9(4), 254–80.
- BACHMANN, R. (2012): “Understanding the Jobless Recoveries After 1991 and 2001,” .
- BAQAEE, D. R., AND E. FARHI (2017): “Productivity and Misallocation in General Equilibrium,” Discussion paper, National Bureau of Economic Research.
- (2019): “The macroeconomic impact of microeconomic shocks: beyond Hulten’s Theorem,” *Econometrica*, 87(4), 1155–1203.
- BARNICHON, R. (2010): “Productivity and unemployment over the business cycle,” *Journal of Monetary Economics*, 57(8), 1013–1025.
- BARRO, R. J., AND R. G. KING (1984): “Time-separable preferences and intertemporal-substitution models of business cycles,” *The Quarterly Journal of Economics*, 99(4), 817–839.
- BASU, S., J. G. FERNALD, AND M. S. KIMBALL (2006): “Are technology improvements contractionary?,” *American Economic Review*, 96(5), 1418–1448.
- BERGER, D. (2012): “Countercyclical restructuring and jobless recoveries,” .
- BIGIO, S., AND J. LA’O (2019): “Distortions in production networks,” *Work. Pap., Columbia Univ., New York*.

- CARVALHO, V. M., AND A. TAHBAZ-SALEHI (2019): “Production networks: a primer,” *Annual Review of Economics*, 11.
- DUPOR, B. (1999): “Aggregation and irrelevance in multi-sector models,” *Journal of Monetary Economics*, 43(2), 391–409.
- FERNALD, J. G., AND J. C. WANG (2015): “Why has the cyclicity of productivity changed?: what does it mean?,” .
- FOERSTER, A., A. HORNSTEIN, P.-D. SARTE, AND M. W. WATSON (2019): “Aggregate Implications of Changing Sectoral Trends,” Discussion paper, National Bureau of Economic Research.
- FOERSTER, A. T., P.-D. G. SARTE, AND M. W. WATSON (2011): “Sectoral versus Aggregate Shocks: A Structural Factor Analysis of Industrial Production,” *Journal of Political Economy*, 119(1), 1–38.
- FORT, T. C., AND S. D. KLIMEK (2016): “The effects of industry classification changes on US employment composition,” *Tuck School at Dartmouth*.
- GALÍ, J., AND L. GAMBETTI (2009): “On the sources of the great moderation,” *American Economic Journal: Macroeconomics*, 1(1), 26–57.
- GALÍ, J., AND T. VAN RENS (2014): “The vanishing procyclicality of labor productivity,” .
- GARIN, J., M. J. PRIES, AND E. R. SIMS (2018): “The relative importance of aggregate and sectoral shocks and the changing nature of economic fluctuations,” *American Economic Journal: Macroeconomics*, 10(1), 119–48.
- GREENWOOD, J., Z. HERCOWITZ, AND P. KRUSELL (2000): “The role of investment-specific technological change in the business cycle,” *European Economic Review*, 44(1), 91–115.
- HERRENDORF, B., R. ROGERSON, AND A. VALENTINYI (2013): “Two perspectives on preferences and structural transformation,” *American Economic Review*, 103(7), 2752–89.
- HORNSTEIN, A., AND J. PRASCHNIK (1997): “Intermediate inputs and sectoral comovement in the business cycle,” *Journal of Monetary Economics*, 40(3), 573–595.

- HORVATH, M. (1998): “Cyclicalities and sectoral linkages: Aggregate fluctuations from independent sectoral shocks,” *Review of Economic Dynamics*, 1(4), 781–808.
- (2000): “Sectoral shocks and aggregate fluctuations,” *Journal of Monetary Economics*, 45(1), 69–106.
- HOUSE, C. L., A.-M. MOCANU, AND M. D. SHAPIRO (2017): “Stimulus effects of investment tax incentives: Production versus purchases,” Discussion paper, National Bureau of Economic Research.
- JORDÀ, Ò. (2005): “Estimation and inference of impulse responses by local projections,” *American economic review*, 95(1), 161–182.
- JUSTINIANO, A., G. E. PRIMICERI, AND A. TAMBALOTTI (2010): “Investment shocks and business cycles,” *Journal of Monetary Economics*, 57(2), 132–145.
- KIM, K., AND Y. S. KIM (2006): “How important is the intermediate input channel in explaining sectoral employment comovement over the business cycle?,” *Review of Economic Dynamics*, 9(4), 659–682.
- KOENDERS, K., AND R. ROGERSON (2005): “Organizational dynamics over the business cycle: a view on jobless recoveries,” *Review-Federal Reserve Bank of Saint Louis*, 87(4), 555.
- LIM, K. (2018): “Endogenous production networks and the business cycle,” *Work. Pap., Univ. Toronto, Can.*
- LONG JR, J. B., AND C. I. PLOSSER (1983): “Real business cycles,” *The Journal of Political Economy*, pp. 39–69.
- MALIAR, L., S. MALIAR, J. TAYLOR, AND I. TSENER (2015): “A tractable framework for analyzing a class of nonstationary Markov models,” Discussion paper, National Bureau of Economic Research.
- MCGRATTAN, E. R. (2017): “Intangible capital and measured productivity,” Discussion paper, National Bureau of Economic Research.

- MCGRATTAN, E. R., AND E. C. PRESCOTT (2014): “A reassessment of real business cycle theory,” *American Economic Review*, 104(5), 177–82.
- MCGRATTAN, E. R., AND J. A. SCHMITZ JR (1999): “Maintenance and repair: Too big to ignore,” *Federal Reserve Bank of Minneapolis. Quarterly Review-Federal Reserve Bank of Minneapolis*, 23(4), 2.
- MEADE, D. S., S. RZEZNIK, AND D. ROBINSON-SMITH (2003): “Business Investment by Industry in the US Economy for 1997,” *Survey of current business*, 83(11), 18–70.
- OBERFIELD, E. (2012): “Business networks, production chains, and productivity: A theory of input-output architecture,” .
- OBERFIELD, E., AND D. RAVAL (2014): “Micro data and macro technology,” Discussion paper, National Bureau of Economic Research.
- ROGERSON, R. (1988): “Indivisible labor, lotteries and equilibrium,” *Journal of monetary Economics*, 21(1), 3–16.
- STOCK, J. H., AND M. W. WATSON (2003): “Has the business cycle changed and why?,” in *NBER Macroeconomics Annual 2002, Volume 17*, pp. 159–230. MIT press.
- TASCHEREAU-DUMOUCHEL, M. (2019): “Cascades and fluctuations in an economy with an endogenous production network,” *Available at SSRN 3115854*.
- WINBERRY, T. (2018): “Lumpy Investment, Business Cycles, and Stimulus Policy,” *Working Paper*.

A Construction of Dataset and Investment Network

A.1 Dataset

We use data from the BEA’s GDP by Industry Database to construct a dataset covering 35 non-government, non-farm sectors covering the years 1948-2017. Our data also omits the real estate sector, as it is not possible to consistently separate imputations for owner-occupied housing from value added in this sector over time. Industries are defined according to NAICS codes; a full list of all 35 industries is available in Table 1.

Data on nominal and real measures of value added by industry are directly available from the BEA.⁵¹ For employment data from 1998 to the present, we use data directly measured in NIPA Table 6.4D, which reports the total number of full-time and part-time employees by industry, where industries are defined according to NAICS codes. Unfortunately, data in NIPA Tables 6.4B and 6.4C report employment by industry according to SIC definitions of industries and are not directly comparable.

We construct NAICS-denominated employment data using historical employment data from the BEA from 1948-1997. These data only provide employment for 16 of the 35 sectors we consider for the full time series. The reason that these data have fewer industries is that, prior to 1977, manufacturing is collapsed into durable and non-durable manufacturing sectors (compared to the 19 subsectors we use in our analysis). To obtain a full time series of employment in these 19 manufacturing subsectors, we take SIC coded employment data from the GDP by Industry Database covering 1947-1997 and convert the data to NAICS industries using the concordance provided in [Fort and Klimek \(2016\)](#).⁵² We then combine this data with the existing data for these manufacturing subsectors from 1977-1997, scaling our converted data pre-1977 to be consistent with post-1977 data.⁵³

⁵¹Because the base year used for the quantity index measures of real value added differs across data for pre- and post-1997, we rescale these data to be consistent over time.

⁵²This concordance is based off of LBD data with detailed SIC and NAICS codes. We convert SIC-coded employment to NAICS coded proportionally on the basis of employment in each SIC industry.

⁵³We rescale by backcasting the 1977-1997 data using the growth rate of employment in each sector as obtained from the SIC converted data.

A.2 Construction of Investment Network

The investment network that we construct records the share of investment expenditures of sector j that were purchased from sector i for each pair (i, j) of the 35 sectors in our dataset and for each year in our 1947 - 2017 sample. While the BEA has published this information in its publicly available capital flows tables, these tables are only available for a handful of years, they do not include the majority of intellectual property, their classification procedures are not consistently defined across years, and they are recorded at inconsistent levels of sectoral disaggregation over time.

We therefore construct our own investment network to overcome these issues with the capital flows tables. Our investment network is available for each year in the 1947-2017 sample, includes all of intellectual property, are based on consistent classification procedures, and are consistently recorded at our 35-sector level of disaggregation. The main challenge is to estimate the share of total production in each sector i that is bought by sector j . The publicly available BEA does not report these pairwise transactions; instead, we only observe the total expenditures on each 30 type of disaggregated investment goods (from the fixed asset table) and total production of broadly defined structures, equipment, and intellectual properties (from the input-output data) for each sector in our data.

We address this challenge in three main steps. First, we construct a *bridge file* that computes, for each of 30 types of capital goods, the share of total production of that good accounted for by a given sector i . Second, we compute the total amount of each good purchased by sector j , and assume that the sector purchases that good proportionally from the sectors which produce it.⁵⁴ Finally, we aggregate across the different types of capital goods to arrive at our investment network.⁵⁵ Our construction of these bridge files closely follows the methodology the BEA has used to produce similar the capital flows tables and is consistent with the approach used in [McGrattan \(2017\)](#) to create a capital flows table for the year 2007.

The remainder of this appendix describes how we construct the bridge files which com-

⁵⁴The BEA also makes this assumption in the construction of its capital flows tables.

⁵⁵Following BEA practice, our investment network covers private new investment but excludes used and scrapped goods. We also exclude residential investment in structures because production of residential investment is not reported separately in the input-output data prior to the year 1997. However, residential structures is almost entirely sold to the real estate sector, which we omit from our analysis (due to its value being dominated by imputations for owner-occupied housing).

putes the share of production of each type of capital good that is accounted for by each sector in the economy. We describe this procedure separately for equipment, structures, and intellectual property capital goods. Given these bridge files, it is straightforward to aggregate up to the investment network displayed in Figure 1 in the main text.

Structures

We allocate production of all structures investment commodities to the construction sector except for the production of mining structures (which are allocated to the mining sector). This allocation is consistent with how the BEA constructs its own capital flows tables, as described in [McGrattan \(2017\)](#). We have verified that the production of structures implied by this bridge file is consistent with the production of structures recorded in the input-output data.⁵⁶

Intellectual Property

We allocate the production of intellectual property products based on the BEA practices discussed in [McGrattan \(2017\)](#). There are four detailed types of intellectual property commodities: prepackaged software, own and custom software, research and development and artistic originals. We allocate the production of pre-packaged software to the Information sector (particularly NAICS sector 5112 Software Publishers), the production of own and custom software and R&D investment to the Professional and Technical Services sector, and split the production of artistic originals between the Information sector (which includes sub-sectors like radio and TV communication and motion picture publishing) and the Arts and Entertainment Services Sector.⁵⁷

We also adjust our allocation to account for the fact that Wholesale Trade, Retail Trade,

⁵⁶The only exception concerns the real estate sector, which is omitted from our analysis anyway. In particular, between the years of 1997-2017 roughly 2.5% of reported production of structures investment comes from sectors other than construction (almost exclusively real estate). These contributions account for broker's commission payments for sales of non-residential structures. These contributions are also omitted from the BEA capital flows tables (see [Meade, Rzeznik and Robinson-Smith \(2003\)](#)).

⁵⁷We determine the split of artistic originals using the yearly Input-Output data on total intellectual property products produced by the Arts and Entertainment sector and the Fixed Assets data on total purchases of artistic originals to construct the fraction of all artistic originals produced by Arts and Entertainment Services. The residual fraction of artistic originals (the majority) is allocated to Information.

and Transportation/Warehousing Services play a role in the delivery of new intellectual property to customers (the expenses incurred in this delivery are often called “margins” on purchases). The 2007 and 2012 data provide detail about the allocation of margins, and show that pre-packaged software publishing is the only intellectual property product with margins. We therefore allocate a fraction of pre-packaged software purchases to Wholesale Trade, Retail Trade, and Transportation/Warehousing Services in order to account for the margins. Our implied production of intellectual property products by each sector aligns with the data in the Input-Output database.

Equipment

Constructing a bridge file for equipment investment commodities is the most difficult task because there are 25 detailed equipment commodities reported in the Fixed Assets data and the share of production of these commodities are allocated to a different mix of sectors in each year. Because of varying data availability from the BEA, we describe how we construct the bridge file for equipment investment separately for 1997-2017, for 1987 and 1992, and for the remaining years (1947-1981, 1988-1991, 1993-1996).

For 1997-2017, the BEA already provides a detailed bridge file for equipment commodities as part of its Input-Output data; therefore, no further work on the bridge file is needed.⁵⁸

The BEA also provides a bridge file for 1987 and 1992, but it is coded using SIC rather than NAICS sector classifications; we convert the SIC to NAICS definitions using the crosswalks defined in [Fort and Klimek \(2016\)](#).⁵⁹ One disadvantage of this bridge data is that we

⁵⁸For these years, and for all years of data, however, we do need to make one adjusted to the Fixed Assets data on investment purchases. The Fixed Assets data does not separately split out the purchases of used or scrap goods, and thus it does not neatly correspond to the Input-Output accounting framework, which separately splits out used and scrap goods. These differences are generally small, however, for some equipment commodities, especially vehicles (autos, trucks, boats, etc.), the discrepancies are substantial, since the volume of capital transactions in used vehicle markets is high. To address this inconsistency, we symmetrically scale up in each sector the total purchases of equipment capital to match the total production amounts reported in the bridge data file for each year from 1997-2017. For equipment commodities, this adjustment is negligible. For years prior to 1997, we used the median scaling factor for each sector to make this adjustment.

⁵⁹We allocate SIC codes to NAICS codes on the basis of total shipments by sector in the [Fort and Klimek \(2016\)](#) crosswalk. We make one small adjustment to this crosswalk to be consistent with Input-Output data. Converting from SIC to NAICS, for a few commodities, the bridge implies they contribute to equipment production, whereas the Input Output data reports a zero for total equipment investment production. Any time the converted bridge file implies that a sector is contributing to the production of a capital good when

do not have much detail in terms of the margin sectors, and failure to adjust these results in unreasonable differences in the margin allocations compared to the data from 1997 onward. As a result, we take the total reported margins for each commodity and allocate it according to the distribution of margins reported in the years 1997-2001.⁶⁰

To obtain bridge data for all years prior to 1987 and for 1988-1991 and 1993-1996 (in which there are no publicly available bridge files), we interpolate existing data and discipline that interpolation to ensure it matches the total production of equipment investment by sector from the Input-Output data. We can write the total production of equipment capital by sector i as the following:

$$I_{it} = \sum_c \omega_{cit} \hat{I}_{ct}$$

where \hat{I}_{ct} is the level of investment purchased of equipment commodity c , ω_{cit} is the fraction of that commodity produced by sector i (the bridge data) and I_{it} is the production of all equipment capital by sector i . We have data on \hat{I}_{ct} and I_{it} for all years from Fixed Assets and Input-Output data, but we do not know the value of the bridge data ω_{cit} for most years prior to 1997.

To obtain interpolated values for the bridge data ω_{cit} for the years without a bridge file, we start with an approximation of production of equipment investment by sector, given by:

$$\tilde{I}_{it} = \sum_c \bar{\omega}_{ci} \hat{I}_{ct}$$

where $\bar{\omega}_{ci}$ is an approximate guess for the bridge relationship. In practice, $\bar{\omega}_{ci}$ is either the bridge data from the last available year (for years prior to 1987) or a moving average of the two nearest bridge files (1987 and 1992 for years 1988-1991 and 1992 and 1997 for years

the Input-Output data reports that sector generated zero production of equipment capital, we edit the SIC to NAICS crosswalk to not allocate value to that sector.

⁶⁰Fortunately, although there are substantial changes in sector codes, there are minimal changes in the classification of capital commodities over these years, making it straightforward to map these margin allocations over time. The only difference between investment commodity classifications between the 1997 and earlier data is that there is no further detail on instruments (nonmedical vs. medical in 1997), trucks (light vs. other in 1997), or tractors (construction vs. agriculture in 1997). We are constrained to assume that the production allocation of these equipment commodities is identical pre-1997. In the cases of tractors, since these are allocated to either construction or agricultural machinery in the level of detail available for the bridge files for 1997 onward, we allocate these expenditures proportionally based on the amount of expenditures on each type of tractor by each sector in each year.

1993-1996).

However, the true level of investment production is given by $I_{it} = k_{it}\tilde{I}_{it} = k_{it} \sum_c \bar{\omega}_{ci} \hat{I}_{ct}$, where k_{it} is a scaling factor. The idea here is that some change in the production of equipment investment by sector may be captured by changes in the distribution of equipment purchases across commodities according to some fixed bridge file (changes in the distribution of \hat{I}_{ct} over c), and that anything not captured by those changes must show up in this residual scaling constant, k_{it} . Given that we can directly observe actual production of all equipment capital sector and construct the approximate production of equipment, we can solve for k_{it} in every year.⁶¹

Given a value of k_{it} , we can construct the approximate bridge relationship for each equipment commodity, sector and year, $\tilde{\omega}_{cit}$, as:

$$\tilde{\omega}_{cit} = \frac{k_{it}\omega_{cit}}{\sum_j k_{jt}\omega_{cjt}}$$

Thus, although this relies heavily on the most recent observations of the bridge data, it is constrained to be consistent with overall changes in the distribution of equipment production by sector. The key assumption here is that any changes in this distribution of equipment capital by sector apart from those captured by changes in the distribution of investment purchases over commodities (changes in the distribution of \hat{I}_{ct} over c), are generated by symmetrically proportional changes in the bridge file for each sector.⁶²

Following this approach, we construct bridge data for equipment investment for all years prior to 1997.⁶³ We then combine all the bridge files we have constructed to construct the aggregate investment matrix.

⁶¹In practice, we do this by matching to the growth rate of actual investment. We do this since the margin totals may not exactly equate to the fixed asset data for the years 1987 and 1992, and thus we do not want to strictly impose this in levels.

⁶²An additional assumption needed here is that there are no sectors which previously contributed to production of any equipment commodity that no longer contribute in the first year that bridge data is available.

⁶³We also use this approach to update the existing bridge files for 1987 and 1992 to be consistent with Input-Output data.

TABLE 16
SKEWNESS OF INVESTMENT AND INTERMEDIATES NETWORKS

	Eigenvalue Centrality	Weighted Outdegree
Investment net.	3.09	2.40
Intermediates net.	1.29	0.74

Notes: Eigenvalue centrality is defined as the eigenvector associated with the largest eigenvalue of the matrix. The weighted outdegree is defined as the sum over columns of the network matrix. Skewness of each of these centrality measures is computed as the sample skewness.


A.3 Additional Analysis of Investment Network

We now present additional analysis of the investment network referenced in Section 2 of the main text.

Changes in the network over time Figure 14 compares the heatmaps of the investment network in the pre and post 1984 samples. Our four investment hubs are the primary suppliers of investment goods in each subsample. The main difference across subsamples is that professional/technical services accounts for a larger share of investment production in the post-1984 period, consistent with the well-known rise of intellectual property. Appendix G shows that our main model results are robust to allowing the network to differ in the pre and post 1984 sample.


Comparing skewness of investment and intermediates networks Table 16 shows that the investment network is significantly more skewed than the intermediates input-output network according two typical measures of network skewness. [Carvalho and Tahbaz-Salehi \(2019\)](#) discuss both of these measures of skewness; intuitively, they compute a measure of centrality for each sector, which determines how important of a supplier it is to other sectors, and then compute the skewness of these centrality measures across industries. A highly skewed set of centrality measures indicates that the network is dominated by a small number of highly important sectors, or hubs. For both measures of centrality, the investment network is roughly three times more skewed than the intermediates input-output network, suggesting that it is much more sparse and thus potentially more powerful in propagating

FIGURE 14: Heatmaps of Investment Network, Pre/Post 1984



Notes: Heatmaps of the investment input-output network λ_{ij} are constructed as described in the main text. The (i, j) entry of each network corresponds to parameter γ_{ij} and λ_{ij} , i.e., the amount of sector i 's good used in sector j . The pre-84 network corresponds to the years 1947-1983 and the post-84 network corresponds to the years 1984-2017.

FIGURE 15: Correlogram of Sector-level Value Added with Aggregate GDP



Notes: correlation of value added at sector s in year $t - h$, y_{st+h} , with aggregate GDP in year t , Y_t . Both y_{st+h} and Y_t are logged and HP-filtered with smoothing parameter 6.25. x-axis varies the lag $h \in \{-2, -1, 0, 1, 2\}$. “Investment hubs” compute the unweighted average the value of these statistics over $s =$ construction, machinery manufacturing, motor vehicles manufacturing, and professional/technical services. “Non-hubs” compute the unweighted average over the remaining sectors. “Pre-1984” performs this analysis in the 1948 - 1983 subsample and “post-1984” performs this analysis in the 1984 - 2017 subsample. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

sectoral shocks than the input-output network.

B Additional Results on Descriptive Evidence of Investment Hubs

This appendix present additional pieces of evidence regarding the behavior of investment hubs referenced in Section 2 in the main text.


Correlogram with GDP Figure 15 presents the correlogram between sector-level value added and aggregate GDP rather than aggregate employment as in Figure 2. As in that figure, hubs are more correlated with aggregate GDP fluctuations than non-hubs, and this difference between hubs is larger in the post-1984 sample.

TABLE 17
VOLATILITY OF ACTIVITY, HUBS VS. MANUFACTURING

	Investment Hubs		Non-Hubs		Non-Hub Manuf.	
	<i>Pre-84</i>	<i>Post-84</i>	<i>Pre-84</i>	<i>Post-84</i>	<i>Pre-84</i>	<i>Post-84</i>
$\sigma(y_{st})$	6.62%	8.27%	4.78%	4.02%	6.03%	4.84%
$\sigma(l_{st})$	4.33%	3.39%	2.73%	2.22%	2.97%	2.53%

Notes: standard deviation of business cycle component of sector-level value added or employment. y_{st} is logged real value added at sector s , HP-filtered with smoothing parameter 6.25. l_{st} is logged real value added at sector s , HP-filtered with smoothing parameter $\lambda = 6.25$. “Investment hubs” compute the unweighted average the value of these statistics over $s =$ construction, machinery manufacturing, motor vehicles manufacturing, and professional/technical services. “Non-hubs” compute the unweighted average over the remaining sectors. “Non-hub manufacturing” computes the average over manufacturing sectors other than machinery and motor vehicles. “Pre-1984” performs this analysis in the 1948 - 1983 subsample and “post-1984” performs this analysis in the 1984 - 2017 subsample. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

FIGURE 16: Correlogram of Sector-level Value Added with Aggregate Employment, Hubs vs. Manufacturing



Notes: correlation of value added at sector s in year $t - h$, y_{st+h} , with aggregate employment in year t , L_t . Both y_{st+h} and L_t are logged and HP-filtered with smoothing parameter 6.25. x-axis varies the lag $h \in \{-2, -1, 0, 1, 2\}$. “Investment hubs” compute the unweighted average the value of these statistics over $s =$ construction, machinery manufacturing, motor vehicles manufacturing, and professional/technical services. “Non-hubs” compute the unweighted average over the remaining sectors. “Non-hub manufacturing” computes the average over manufacturing sectors other than machinery and motor vehicles. “Pre-1984” performs this analysis in the 1948 - 1983 subsample and “post-1984” performs this analysis in the 1984 - 2017 subsample. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

TABLE 18
SECTORAL COMOVEMENT WITH INVESTMENT HUBS

	Hubs		Non-Hubs	
	<i>Pre-1984</i>	<i>Post-1984</i>	<i>Pre-1984</i>	<i>Post-1984</i>
Value added	0.46	0.32	0.31	0.18
Employment	0.62	0.62	0.53	0.52

Notes: computes $\rho_{\tau}^i(x) \equiv \frac{\sum_{j \neq i} \omega_j^x \text{Corr}(x_{it}, x_{jt} | t \in \tau)}{\sum_{j=i+1}^N \omega_j^x}$ where x_{jt} is logged + HP-filtered variable of interest, $\tau \in \{\text{pre 1984, post 1984}\}$ is time period, and $\omega_{i\tau}^x$ are sectoral shares. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

Results not driven by manufacturing sectors One concern with our results may be that the behavior of hubs is driven by the fact that two of four hubs are manufacturing sectors, and that manufacturing is more cyclical than other sectors. We present two pieces of evidences against this concern. First, Table 17 shows that the non-hub manufacturing sectors are significantly less volatile than the investment hubs. Second, Figure 16 shows that the correlation of non-hub manufacturing sectors with aggregate employment is close to that of the other non-hubs, and lower than the correlation of the investment hubs.

Other sectors comove more strongly with hubs than non-hubs Table 18 shows that the average correlation between fluctuations in hubs' value added or employment with all other sectors is higher than the correlation of non-hubs' value added or employment with those sectors. This finding is consistent with evidence from correlograms, suggesting that fluctuations in hub sectors are closely aligned with fluctuations in other sectors and the aggregate.


C Details of Model Calibration

We now describe the details of our model's calibration.

C.1 Parameters Other Than Shocks

We begin with all the parameters other than those pertaining to the TFP shocks.

FIGURE 17: Calibrated Value Added Shares θ_j




Notes: Values for the value-added shares θ_j are computed as the ratio of value added to gross output in each sector, averaged across the entire sample, 1947-2017.

Value added shares We calibrate the share of intermediate inputs in production, $1 - \theta_j$, using the BEA input-output database. Given the Cobb-Douglas structure of our production function, the shares θ_j are pinned down by the ratio of value added to gross output at the firm level. We obtain this ratio for each year in our 1947-2017 sample and then compute their average value over time (empirically, the shares are fairly stable over time anyway). Figure 17 plots our calibrated share θ_j .

Labor shares We pin down the labor share $1 - \alpha_j$ using the ratio of sector-level labor costs to sector-level income (i.e. nominal value added). The BEA provides data on compensation by sector in NAICS codes back to the year 1987; for years prior to 1987, we convert SIC based data to NAICS using the crosswalk in Fort and Klimek (2016).⁶⁴ We also correct for the fact that sector-level compensation in the BEA data does not include self-employed income; we use BEA data on the number of self-employed workers by industry from 1987-2017 to and we multiply industry compensation by one plus the ratio of self-employed employment to total part-time and full-time employment in the industry (implicitly assuming that average

⁶⁴To allocate compensation by industry in the Fort and Klimek (2016) crosswalk, we use the crosswalk data for payroll by industry.

FIGURE 20: Calibrated Consumption Shares ξ_j



Notes: Values for consumption preference ξ_j are constructed as the fraction of total nominal consumption expenditures on each sector’s goods or services, averaged over the entire sample 1947-2017.

The only sector this meaningfully impacts is the construction sector. With this addition to private consumption, ξ_j is set to the fraction of all consumption expenditures accounted for by each industry averaged over the years 1947-2017. These values are plotted in Figure 20.

Investment and Intermediates Networks We discuss the calibration of the investment network in the main text and in Appendix A. The parameters of the intermediates network γ_{ij} are pinned down by the share of total intermediates expenditures in sector j that is purchased from some other sector i from the BEAs’ use table. We compute these shares for each year of our data and then average it across the years 1947-2017.

C.2 Measured sector-level productivity

We now describe how we measure sector-level productivity. We perform this measurement using a Solow residual approach using value added net of primary inputs. For notational convenience, we consider a renormalization of our original production function to associate

TFP directly with value added:

$$Q_{jt} = \left(A_{jt} K_{jt}^{\alpha_j} L_{jt}^{1-\alpha_j} \right)^{\theta_j} X_{jt}^{1-\theta_j}$$

Using this approach, we then estimate TFP from the data as the Solow residual based on the following equation:


$$\log A_{jt} = \log Y_{jt} - \alpha_{jt} \log K_{jt} - (1 - \alpha_{jt}) \log L_{jt}$$

Annual industry value added and employment are measured as described in Appendix A. We construct the capital stock for each industry in each year using the perpetual inventory method, using the nominal year-end capital stock for each industry in 1947 as our starting point (from BEA Fixed Assets data). We then use the annual implied depreciation rates and real quantities of investment for each industry to iterate forward the capital accumulation process and generate a time series of capital for each industry. We allow the parameter α_{jt} to vary period by period (taking the two period average constructing TFP in growth rates) in order to isolate changes in productivity rather than the production function; however, our results are robust to using a fixed parameter value to compute TFP.

We detrend our model using a log-polynomial trend because log-linear trends provide a poor fit to sector-level TFP. Figure 21 plots the time-series of sector-level TFP for two examples, construction and machinery manufacturing. Construction TFP evolves nonlinearly over time, and a third or fourth order polynomial trend is required to capture these nonlinearities. In contrast, machinery manufacturing evolves more linearly, but a polynomial trend continues to fit better than a linear one. We choose a fourth order trend for the main text in order to balance these nonlinearities against overfitting the data. However, we show in Appendix G that our main results are robust to using lower-order polynomials for detrending.


Given detrended values for log TFP, we estimate the autocorrelation parameter for each sector's TFP, ρ_j , by maximum likelihood. We then extract the innovations to log TFP, which are then fed into the model. The values for these ρ_j parameters are plotted in Figure 22.

FIGURE 21: Detrending Sector-Level Data



Notes: The figure reports log sector level TFP for the Construction and Machinery Manufacturing industries, normalized to zero in the year 1948. We also report a fitted polynomial trend lines for polynomials of order 1-4, estimated via OLS.

FIGURE 22: Calibrated Persistence Parameters ρ_j



Notes: Persistence parameters ρ_j of sector-level TFP are estimated from log polynomial (of order 4) detrended TFP data using maximum likelihood.

TABLE 19
PRINCIPAL COMPONENTS ANALYSIS OF MEASURED TFP

Sample period	$1000\text{Var}(\Delta \log A_t)$	Due to 1st component	Residual
1949-1983	0.40	0.32 (81%)	0.08 (19%)
1984-2017	0.27	0.15 (56%)	0.12 (44%)

Notes: We estimate the loadings associated with the first principal component, and then multiply these by each sector’s TFP series to construct our aggregate shock (1st component). We then regress aggregate TFP on this constructed aggregate shock and report the explained sum of squares and R^2 (the variance attributable to the 1st component) and the sum of squared errors (the variance attributable to the residual, interpreted as sectoral shocks). We do this separately for the pre- and post-1984 periods.

Interpreting changes in sector-level productivity over time using principal components analysis The key feature of the data which drives the changes in business cycle patterns since 1984 is that sector-level productivity has become less correlated across sectors. In the main text, we interpret this change as a decline in the variance in the volatility of aggregate shocks which affect all sectors in the economy. We now provide further support for this interpretation using a principal components exercise similar to [Garin, Pries and Sims \(2018\)](#).

However, performing that principal components exercise requires us to estimate a full rank covariance matrix for TFP pre- and post-1984, which we cannot do with 35 sectors and less than 35 years of data post-1984.⁶⁷ As a result, we collapse our data down to 28 sectors by condensing all non-durable manufacturing sectors into one sector, and then do principal components on log TFP growth for 28 sectors pre- and post-1984.⁶⁸

The results of this principal components exercise are reported in Table 19. The first principal component – which can be loosely interpreted as the aggregate shock – accounts for 81% of the variance of aggregate TFP in the pre-1984 sample, but only 56% of the variance in the post-1984 sample. Furthermore, the variance of the residual component – which can be loosely interpreted as the sector-specific shocks – has remained fairly stable

⁶⁷There are exactly 35 years of data pre-1984 when studying growth rates, since employment data only begins in 1948.

⁶⁸We could have alternatively collapsed other sectors to shrink the number of sectors low enough to generate covariance matrices. We prefer this approach as aggregating within non-durable manufacturing does not affect hubs, many non-durable manufacturing sectors are small, and the aggregate sector of non-durable manufacturing is more intuitive than alternative aggregated sectors in services.

over time.

D Proof of Theorem that Employment is Constant without Capital

Theorem 1 Suppose that there is no capital in the economy, implying that production in each sector and the market clearing condition for output are given by:

$$Y_{jt} = A_{jt} L_{jt}^{\theta_j} X_{jt}^{1-\theta_j}$$

$$Y_{jt} = C_{jt} + \sum_{i=1} M_{jit}$$

where $X_{jt} = \prod_{i=1}^N M_{ijt}^{\gamma_{ij}}$. With period by period household preferences of $U(C_t, L_t) = \log(C_t) - \chi \frac{L_t^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}}$ and aggregate consumption given by the Cobb-Douglas aggregate of $C_t = \prod_{j=1}^N C_{jt}^{\xi_j}$, then employment in each sector, and thus total employment, will always be constant.

Proof. We prove this in two steps. First, we show that the level of nominal output in each sector is constant over time given these assumptions. Then we show that a constant level of nominal output implies a constant level of employment in each sector.

From the planner's problem, the first order conditions for individual intermediate goods and for sectoral consumption are given by:

$$\mu_{it} = \mu_{jt} \theta_j \gamma_{ij} \frac{Y_{jt}}{M_{ijt}}$$

$$\frac{\xi_j}{C_{jt}} = \mu_{jt}$$

Using these first order conditions for intermediates and consumption, we can rewrite the resource constraint as:

$$\begin{aligned}
Y_{jt} &= C_{jt} + \sum_{i=1}^N M_{jit} \\
Y_{jt} &= \frac{\xi_j}{\mu_{jt}} + \sum_{i=1}^N \left(\frac{(1 - \theta_j) \gamma_{ji} \mu_{it} Y_{it}}{\mu_{jt}} \right) \\
\mu_{jt} Y_{jt} &= \xi_j + \sum_{i=1}^N ((1 - \theta_j) \gamma_{ji} \mu_{it} Y_{it})
\end{aligned}$$

The product of the multiplier μ_{jt} and real gross output Y_{jt} gives us a measure of nominal gross output for each sector. Given the above expression, we can stack market clearing conditions for all sectors and write this in matrix form, getting:

$$\begin{aligned}
\mu_t \vec{Y}_t &= \underbrace{\begin{bmatrix} \xi_1 \\ \xi_2 \\ \vdots \\ \vdots \\ \xi_3 \end{bmatrix}}_{\xi} + \underbrace{\begin{bmatrix} (1 - \theta_1) \gamma_{11} & (1 - \theta_1) \gamma_{12} & \cdots & (1 - \theta_1) \gamma_{1N} \\ (1 - \theta_2) \gamma_{21} & (1 - \theta_2) \gamma_{22} & & \vdots \\ \vdots & & \ddots & \vdots \\ (1 - \theta_N) \gamma_{N1} & \cdots & \cdots & (1 - \theta_N) \gamma_{NN} \end{bmatrix}}_{\Phi} \mu_t \vec{Y}_t \\
\mu_t \vec{Y}_t &= (I - \Phi)^{-1} \xi
\end{aligned}$$

Thus, $\mu_{jt} Y_{jt}$ in each sector is a constant, depending on the consumption parameters ξ and the Leontief inverse representing the input-output network for intermediates, $(I - \Phi)^{-1}$. Given that μ_{jt} is equal to ξ_j / C_{jt} , this implies that the consumption to output ratio in each sector will be perfectly constant as well.

We now relate this to employment. Turning to the first order conditions for employment,

we have:

$$\mu_{jt}\theta_j\frac{Y_{jt}}{L_{jt}} = \chi\left(\sum_{i=1}^N L_{it}\right)^{\frac{1}{\eta}}$$

$$\theta_j\mu_{jt}Y_{jt} = \chi L_{jt}\left(\sum_{i=1}^N L_{it}\right)^{\frac{1}{\eta}}$$

Given that nominal output, $\mu_{jt}Y_{jt}$ is a constant, this implies that $L_{jt}\left(\sum_{i=1}^N L_{it}\right)^{\frac{1}{\eta}}$ must also be constant.

Now, consider the ratio of the first order conditions for labor in two sectors k and j :

$$\frac{\theta_j\mu_{jt}Y_{jt}}{\theta_k\mu_{kt}Y_{kt}} = \frac{\chi L_{jt}\left(\sum_{i=1}^N L_{it}\right)^{\frac{1}{\eta}}}{\chi L_{kt}\left(\sum_{i=1}^N L_{it}\right)^{\frac{1}{\eta}}}$$

$$\frac{\theta_j\mu_{jt}Y_{jt}}{\theta_k\mu_{kt}Y_{kt}} = \frac{L_{jt}}{L_{kt}}$$

Given that nominal output is constant for all sectors, this implies that ratio of employment in any two sectors will also be constant.

Finally, using this relationship, we can rewrite the original first order condition for labor purely in terms of employment in sector j as follows:

$$\theta_j\mu_{jt}Y_{jt} = \chi L_{jt}\left(\sum_{i=1}^N L_{it}\right)^{\frac{1}{\eta}}$$

$$\theta_j\mu_{jt}Y_{jt} = \chi L_{jt}\left(\sum_{i=1}^N \frac{\theta_i\mu_{it}Y_{it}}{\theta_j\mu_{jt}Y_{jt}} L_{jt}\right)^{\frac{1}{\eta}}$$

$$L_{jt}^{1+\frac{1}{\eta}} = \frac{\theta_j\mu_{jt}Y_{jt}}{\chi \sum_{i=1}^N \frac{\theta_i\mu_{it}Y_{it}}{\theta_j\mu_{jt}Y_{jt}}}$$

Given that the right hand side is a constant, employment in each sector must be constant as well. Further, since aggregate employment is given by $L_t = \sum L_{it}$, this further implies that aggregate employment is also constant.

The key assumptions needed for this result are that consumption preferences and produc-

tion technologies be Cobb-Douglas, implying constant expenditure shares, and that aggregate preferences over consumption are given by log consumption, which implies that income and substitution effects will exactly offset. ■

E Investment Hubs Forecast Employment Better than GDP

We compare the forecasting power of investment hubs and aggregate GDP using the Jordà (2005)-style forecasting regression

$$\log N_{t+h} - \log N_t = \alpha_h + \gamma_h(\log y_{hub,t} - \log y_{hub,t-1}) + \beta_h(\log Y_t - \log Y_{t-1}) + \varepsilon_{t+h} \quad (19)$$


where Y_t is aggregate GDP and $y_{hub,t}$ is value added at investment hubs, aggregated as in Section 4 above. In order to make the coefficients interpretable, we standardize the growth rates of the two right-hand side variables. We estimate these forecasting regressions separately for the different forecasting horizons h .

The top panel of Figure 23 shows the results from forecasting using only aggregate GDP in (19). A one standard deviation increase in aggregate GDP growth predicts a one percent increase in aggregate employment which reverts to zero after four years. Aggregate GDP explains about 20% of one-year-ahead aggregate employment growth.

The bottom panel of Figure 23 shows that aggregate GDP becomes insignificant once investment hubs are also included in (19). The coefficient on aggregate GDP falls to zero and becomes statistically insignificant, implying that its univariate forecasting power came from its correlation with investment hubs. In contrast, a one standard deviation increase in investment hubs' value added growth predicts a persistent 1-2 percent increase in aggregate employment over the next four years. In addition, investment hubs explain more than 25% of the variation in one-year-ahead aggregate employment growth.

Of course, economic forecasters use more sophisticated tools than our simple bivariate regression, but we nevertheless believe that our results contain two useful insights. First, they

FIGURE 23: Forecasting Power of Hubs vs. Aggregate GDP for Aggregate Employment



Notes: top panel plots the results from estimating the forecasting regression (19) using aggregate GDP only. Bottom panel plots the results from running the forecasting regression 19 using both aggregate GDP and investment hubs' value added (the grey lines represent point estimates from the results in the top panel). The horizontal axis in each plot is the forecasting horizon h . The left panels plot the R^2 of the regression. The middle panels plot the coefficient on aggregate employment growth together with a 95% confidence interval. The bottom right panel plots the coefficient on investment hubs' value added growth together with a 95% confidence interval.

show how moving beyond aggregate data can improve forecasts. Second, they show how the structure of the investment network can help interpret the predictive power of various sectors. These sectors are often combined together into reduced-form “factors” in Factor-Augmented VAR models.

F Changes in Aggregate Cycles Driven by Changes in Sectoral Comovement

We now show a number of results referenced in Section 5 of the main text.

F.0.1 Changes in aggregate vs. within-sector cycles

We start by showing additional results concerning the finding that the changes in the aggregate cycle are not reflected in changes in sector-level cycles within sector.

Declining cyclicity of labor productivity driven by rising volatility of employment

We first show that the decline in the cyclicity of aggregate labor productivity is entirely accounted for, in a statistical sense, by the increase in the volatility of employment relative to the volatility of output (as shown in equation 15 in the main text). Of course, the definition of the correlation between labor productivity and output is $\text{Corr}(y_t, y_t - l_t) = \frac{\text{Cov}(y_t, y_t - l_t)}{\sigma(y_t)\sigma(y_t - l_t)}$. Using the linear properties of covariance and rearranging, we can write this as:

$$\begin{aligned} \frac{\text{Cov}(y_t, y_t - l_t)}{\sigma(y_t)\sigma(y_t - l_t)} &= \frac{\text{Cov}(y_t, y_t)}{\sigma(y_t)\sigma(y_t - l_t)} - \frac{\text{Cov}(y_t, l_t)}{\sigma(y_t)\sigma(y_t - l_t)} = \\ \frac{\sigma(y_t)}{\sigma(y_t - l_t)} - \frac{\sigma(l_t)}{\sigma(y_t - l_t)}\text{Corr}(y_t, l_t) &= \frac{\sigma(y_t)}{\sigma(y_t - l_t)} \left(1 - \frac{\sigma(l_t)}{\sigma(y_t)}\text{Corr}(y_t, l_t) \right) \end{aligned}$$

We can write $\sigma(y_t - l_t)$ as:

$$\sigma(y_t - l_t) = \sqrt{\sigma(y_t)^2 + \sigma(l_t)^2 - 2\text{Cov}(y_t, l_t)} = \sigma(y_t) \sqrt{1 + \left(\frac{\sigma(l_t)}{\sigma(y_t)}\right)^2 - 2\left(\frac{\sigma(l_t)}{\sigma(y_t)}\right)\text{Corr}(y_t, l_t)}$$

Combining this expression with the previous one yields:

$$\frac{\sigma(y_t)}{\sigma(y_t - l_t)} \left(1 - \frac{\sigma(l_t)}{\sigma(y_t)}\text{Corr}(y_t, l_t) \right) = \frac{1 - \frac{\sigma(l_t)}{\sigma(y_t)}\text{Corr}(y_t, l_t)}{\sqrt{1 + \frac{\sigma(l_t)^2}{\sigma(y_t)^2} - 2\frac{\sigma(l_t)}{\sigma(y_t)}\text{Corr}(y_t, l_t)}}$$

which is expression (15) in the main text. This expression makes clear that the correlation of labor productivity with GDP depends only on two statistics: the correlation between output and employment ($\text{Corr}(y_t, l_t)$) and the relative standard deviation of employment and GDP ($\frac{\sigma(l_t)}{\sigma(y_t)}$).

Table 20 shows that the correlation of employment and GDP is stable over time; therefore, the rising volatility of employment relative to GDP accounts for the entire decline in

TABLE 20
COMPONENTS OF AGGREGATE LABOR PRODUCTIVITY CYCLICALITY

	<i>Pre-1984</i>	<i>Post-1984</i>
$\text{Corr}(y_t - l_t, y_t)$	0.65	0.26
$\text{Corr}(y_t, l_t)$	0.81	0.83
$\text{Corr}(y_t, l_t)$ only	0.65	0.66
$\sigma(l_t)/\sigma(y_t)$	0.76	1.02
$\sigma(l_t)/\sigma(y_t)$ only	0.65	0.26

Notes: decomposition of the cyclicalty of labor productivity in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). y_t is log aggregate value added and l_t is log aggregate employment, both HP-filtered with smoothing parameter $\lambda = 6.25$. “ $\text{Corr}(y_t, l_t)$ only” computes the cyclicalty of labor productivity from (15) using the actual value of $\text{Corr}(y_t, l_t)$ in each subsample but holding fixed $\sigma(l_t)/\sigma(y_t)$ at its value in the pre-1984 subsample. “ $\sigma(l_t)/\sigma(y_t)$ only” computes labor productivity from (15) using the actual value of $\sigma(l_t)/\sigma(y_t)$ in each subsample but holding fixed $\text{Corr}(y_t, l_t)$ at its value in the pre-1984 subsample. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

TABLE 21
CHANGES IN BUSINESS CYCLES, FIRST DIFFERENCES

	Aggregated		Within-Sector	
	<i>Pre-1984</i>	<i>Post-1984</i>	<i>Pre-1984</i>	<i>Post-1984</i>
$\sigma(\Delta y_t)$	3.39%	2.30%	5.71%	5.01%
$\rho(\Delta y_t - \Delta l_t, \Delta y_t)$	0.68	0.40	0.77	0.74
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.74	0.93	0.62	0.63
$\sigma(\Delta i_t)/\sigma(\Delta y_t)$	1.82	2.51	2.62	2.50

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). y_t is log value added, l_t is log employment, and i_t is log investment. “Aggregated” aggregates value added and investment across sectors using a Tornqvist index weighted by nominal value added shares, aggregates employment as the simple sum, first-differences each variable, and computes the statistics. “Within-sector” first-differences each variable, computes the statistics, and then averages them weighted by the average share of nominal value added within that sub-sample. For consistency with the HP-filtered moments, we again restrict our sample by removing the first three and last three years of data from computed moments.

the cyclicalty of labor productivity. Intuitively, since GDP and employment are so highly correlated, the time-series behavior of their ratio just depends on which component is more volatile.

Robustness of business cycle moments We now show that the business cycle moments from Table 8 are robust to various choices in methodology. Table 21 show that those results

TABLE 22
WITHIN-SECTOR BUSINESS CYCLE STATISTICS WITH DIFFERENT WEIGHTS

	Time-Varying (Baseline)		Fixed Weights	
	<i>Pre-1984</i>	<i>Post-1984</i>	<i>Pre-1984</i>	<i>Post-1984</i>
$\sigma(y_t)$	3.58%	3.00%	3.32%	3.23%
$\rho(y_t - l_t, y_t)$	0.73	0.71	0.72	0.73
$\sigma(l_t)/\sigma(y_t)$	0.65	0.64	0.65	0.65
$\sigma(i_t)/\sigma(y_t)$	2.76	2.84	2.85	2.75

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). y_t is log value added, l_t is log employment, and i_t is log investment. “Baseline” first-differences each variable, computes the statistics, and then averages them weighted by the average share of nominal value added within that sub-sample. For consistency with the HP-filtered moments, we again restrict our sample by removing the first three and last three years of data from computed moments. In “Fixed Weights,” we use each sector’s value added share averaged for the entire sample window to weight sectoral moments both pre- and post-1984.

hold using first-differences rather than the HP-filter to detrend the data. Table 22 shows that the average value of the sector-level statistics is similar using time-varying weights rather than fixed weights (as in the main text).

F.0.2 Sectoral Decomposition of Rising Relative Volatility of Employment

We now show additional results concerning the decomposition of the rising volatility of aggregate employment relative to GDP in (16) in the main text.

Derivation of sectoral decomposition To derive the decomposition presented in Equation (16), we start by decomposing the aggregate variance of employment into within-sector variances and between-sector covariances. We first take a first-order Taylor approximation of aggregate employment, which yields

$$l_t = \sum_{j=1}^N \omega_{jt}^l l_{jt}$$

where ω_{jt}^l is the average share of sectoral employment in the aggregate for the time period studied. The approximation reflects the facts that the log of the sum is not equal to the sum of the logs and that the shares ω_{jt}^l are not constant over time. Given this linear expression

for aggregate employment, standard rules of variance and covariance imply the following decomposition of aggregate employment variance:

$$\mathbb{V}ar(l_t) \approx \sum_{j=1}^N (\omega_{jt}^l)^2 \mathbb{V}ar(l_{jt}) + \sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^l \omega_{ot}^l \mathbb{C}ov(l_{jt}, l_{ot})$$

We perform a similar decomposition for aggregate GDP, and then we consider the ratio of these two decompositions.⁶⁹ This ratio is given by:

$$\begin{aligned} \frac{\mathbb{V}ar(l_t)}{\mathbb{V}ar(y_t)} &= \frac{\sum_{j=1}^N (\omega_{jt}^l)^2 \mathbb{V}ar(l_{jt})}{\sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt}) + \sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^y \omega_{ot}^y \mathbb{C}ov(y_{jt}, y_{ot})} \\ &\quad + \frac{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^l \omega_{ot}^l \mathbb{C}ov(l_{jt}, l_{ot})}{\sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt}) + \sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^y \omega_{ot}^y \mathbb{C}ov(y_{jt}, y_{ot})} \end{aligned}$$

This expression can be rewritten as:

$$\begin{aligned} \frac{\mathbb{V}ar(l_t)}{\mathbb{V}ar(y_t)} &= \frac{\sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt})}{\mathbb{V}ar(y_t)} \frac{\sum_{j=1}^N (\omega_{jt}^l)^2 \mathbb{V}ar(l_{jt})}{\sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt})} \\ &\quad + \frac{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^y \omega_{ot}^y \mathbb{C}ov(y_{jt}, y_{ot})}{\mathbb{V}ar(y_t)} \frac{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^x \omega_{ot}^l \mathbb{C}ov(l_{jt}, l_{ot})}{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^y \omega_{ot}^y \mathbb{C}ov(y_{jt}, y_{ot})} \end{aligned}$$

And then, defining the “variance weight” as $\omega_t = \sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt}) / \mathbb{V}ar(y_t)$, we obtain the final relationship (16) in the main text.

Accuracy of sectoral decomposition As discussed in the derivation above, our decomposition approximates the log of sums with a sum of the logs and assumes that the shares of particular sectors in the aggregate are constant year-to-year. Table 23 shows that these approximations are accurate. We compare the relative variance and standard deviation of employment implied by the decomposition to the actual values in the data, and show that the two are close to each other.


⁶⁹Since aggregate GDP is obtained via a Tornqvist index, log changes in GDP are already given as a weighted sum of log changes in industry GDP. Thus, the approximation only reflects the fact that the weights are not constant over time.

TABLE 23
ACCURACY OF THE DECOMPOSITION

	<i>Pre-84</i>	<i>Post-84</i>
Actual, variance	0.58	1.04
Approximation, variance	0.57	0.94
Actual, standard deviation	0.76	1.02
Approximation, standard deviation	0.75	0.97

Notes: Aggregate employment and value added are constructed from industry data as the sum of all employment and a Tornqvist index of sectoral value added. The decomposition-based aggregate variance is based on the decomposition derived in the Appendix text.

FIGURE 24: Scatterplot of Changes in Sector-Pair Covariances



Notes: panel (a) plots changes in the covariance for each pair of sectors (j, o) in our dataset. The horizontal axis computes the change in the covariance of value added $Cov(y_{jt}, y_{ot})$ in the post-1984 sample (1984-2017) relative to the pre-1984 sample (1948-1983). The vertical axis computes the change in the covariance of employment $Cov(l_{jt}, l_{ot})$ over the same periods. Each point is weighted by the product of the two sector-pair's average nominal value added share over the whole sample. The blue solid line is the OLS regression line. Panel (b) adds the change in within-sector variances to the plot (i.e. sets $j = o$). The red solid line is the OLS line through the within-sector variances only.

TABLE 24
MEASURING COMOVEMENT WITH CORRELATIONS

	Employment	Value added
1951 - 1983	0.55	0.36
1984 - 2014	0.51	0.17
<i>Difference</i>	-0.04	-0.18

Notes: computes $\rho_{\tau}^i(x) \equiv \frac{\sum_{j \neq i} \omega_j^x \text{Corr}(x_{it}, x_{jt} | t \in \tau)}{\sum_{j=i+1}^N \omega_j^x}$ where x_{jt} is logged + HP-filtered variable of interest, $\tau \in \{\text{pre 1984, post 1984}\}$ is time period, and $\omega_{i\tau}^x$ are sectoral shares. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

Scatterplot of changes in sector-pair covariances In the main text, we argued that the increase in the covariance term in the decomposition (16) reflects the fact that the covariance of value added across sectors fell post-1984 while the covariance of employment did not. Figure 24 illustrates those two patterns. First, the fact that covariance of value added fell for most pairs of sectors in our data implies that most sectors are to the left of zero on the x-axis; this force contributes to a decline in the volatility of aggregate value added. Second, for 82% of pairs in our data, the change in the covariance of employment is smaller than the change in the covariance of value added. The stable employment covariances stabilize its aggregate volatility and, therefore, account for its increase relative to the volatility of aggregate GDP.

The right panel of Figure 24 illustrates the stability of the within-sector variance of employment relative to value added. It shows a positive relationship between the change in the variance of employment and the variance of value added; in fact, the coefficient of the regression line is approximately 0.3, similar to the overall level of their ratio in Table 9.⁷⁰

Measuring comovement with correlations Table 24 reports the average comovement between sectors' employment and value added measured using the average correlation between sectors over time (weighting the average by nominal value added shares). Consistent with Figure 24, the average correlation of sectors' employment has remained stable pre- and post-1984 but the average correlation in value added has dropped substantially.

⁷⁰The difference in the slope of the regression line and the levels in the decomposition is due to differences in how sectors are weighted.

TABLE 25

DECOMPOSITION OF RELATIVE EMPLOYMENT VOLATILITY, FIRST DIFFERENCES

	<i>Pre-84</i>	<i>Post-84</i>	<i>Contribution of entire term</i>
$\frac{\mathbb{V}ar(l_t)}{\mathbb{V}ar(y_t)}$	0.55	0.87	100%
Variances	0.35	0.39	15%
Covariances	0.58	1.01	85%
Variance Weight	0.12	0.23	
$(\omega_t = \sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt}) / \mathbb{V}ar(y_t))$			

Notes: results of the decomposition (16) in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). “Variances” refers to the variance component $\frac{\sum_{j=1}^N (\omega_{jt}^l)^2 \mathbb{V}ar(l_{jt})}{\sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt})}$. “Covariances” refers to the covariance component $\frac{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^l \omega_{ot}^l \text{Cov}(l_{jt}, l_{ot})}{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^y \omega_{ot}^y \text{Cov}(y_{jt}, y_{ot})}$. “Variance weight” refers to the weighting term $\omega_t = \sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt}) / \mathbb{V}ar(y_t)$. “Contribution of entire term” column computes the contribution of the first term of the decomposition (16) (in the variance row) and the contribution of the second term (in the covariance row). In this case, the variables are first-differenced rather than HP-filtered. For consistency with the HP-filtered moments, we again restrict our sample by removing the first three and last three years of data from computed moments.

Other robustness of sectoral decomposition Table 25 shows that the results from this decomposition hold when we use first-differences rather than the HP filter to detrend the data. Table 26 shows that the result hold when we hold the sector-level weights fixed over time in order to isolate the changes in the sector-level behavior.

Decomposition results at other levels of sectoral disaggregation Table 27 shows that our results hold using a finer disaggregation of sectors within the manufacturing sector only. These data are from the NBER-CES database, which covers 469 manufacturing sectors from 1959-2011. We still observe at this finely disaggregated level that the rise in the relative variance of employment to GDP is almost exclusively due to changes in the covariance of activity across sectors.

In contrast, Table 28 shows that these results no longer hold at a broadly aggregated goods vs. services split of sectors. At this two sector level, the contribution of the variances and covariances to the aggregate change in volatility is roughly the same. This result is consistent with the fact that many changes are occurring within these sectors but across our

TABLE 26
DECOMPOSITION OF RELATIVE EMPLOYMENT VOLATILITY, FIXED WEIGHTS

	<i>Pre-84</i>	<i>Post-84</i>	<i>Contribution of entire term</i>
$\frac{\text{Var}(l_t)}{\text{Var}(y_t)}$	0.60	0.81	100%
VariANCES	0.44	0.32	8%
CovARIANCES	0.62	0.93	92%
Variance Weight	0.11	0.20	
$(\omega_t = \sum_{j=1}^N (\omega_{jt}^y)^2 \text{Var}(y_{jt}) / \text{Var}(y_t))$			

Notes: results of the decomposition (16) in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). “VariANCES” refers to the variance component $\frac{\sum_{j=1}^N (\omega_{jt}^l)^2 \text{Var}(l_{jt})}{\sum_{j=1}^N (\omega_{jt}^y)^2 \text{Var}(y_{jt})}$. “CovARIANCES” refers to the covariance component $\frac{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^l \omega_{ot}^l \text{Cov}(l_{jt}, l_{ot})}{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^y \omega_{ot}^y \text{Cov}(y_{jt}, y_{ot})}$. “Variance weight” refers to the weighting term $\omega_t = \sum_{j=1}^N (\omega_{jt}^y)^2 \text{Var}(y_{jt}) / \text{Var}(y_t)$. “Contribution of entire term” column computes the contribution of the first term of the decomposition (16) (in the variance row) and the contribution of the second term (in the covariance row). To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures. In this case, the weights used in computing the variance, covariance and within weight terms above correspond to the average shares of nominal value added or employment for the entire period 1948-2017.

TABLE 27
DECOMPOSITION OF RELATIVE EMPLOYMENT VOLATILITY, NBER-CES

	<i>Pre-84</i>	<i>Post-84</i>	<i>Contribution of entire term</i>
$\frac{\text{Var}(l_t)}{\text{Var}(y_t)}$	0.40	0.57	100%
VariANCES	0.34	0.20	1.4%
CovARIANCES	0.37	0.60	92.6%
Variance Weight	0.03	0.06	
$(\omega_t = \sum_{j=1}^N (\omega_{jt}^y)^2 \text{Var}(y_{jt}) / \text{Var}(y_t))$			

Notes :results of the decomposition (16) using NBER-CES data for 469 manufacturing sectors. “VariANCES” refers to the variance component $\frac{\sum_{j=1}^N (\omega_{jt}^l)^2 \text{Var}(l_{jt})}{\sum_{j=1}^N (\omega_{jt}^y)^2 \text{Var}(y_{jt})}$. “CovARIANCES” refers to the covariance component $\frac{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^l \omega_{ot}^l \text{Cov}(l_{jt}, l_{ot})}{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^y \omega_{ot}^y \text{Cov}(y_{jt}, y_{ot})}$. “Variance weight” refers to the weighting term $\omega_t = \sum_{j=1}^N (\omega_{jt}^y)^2 \text{Var}(y_{jt}) / \text{Var}(y_t)$. “Contribution of entire term” column computes the contribution of the first term of the decomposition (16) (in the variance row) and the contribution of the second term (in the covariance row). To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures. Pre-1984 corresponds to the period 1958-1983; post-1984 corresponds to the period 1984-2011. Real value added is constructed using the gross output price deflator.

TABLE 28

DECOMPOSITION OF RELATIVE EMPLOYMENT VOLATILITY, GOODS VS. SERVICES

	<i>Pre-84</i>	<i>Post-84</i>	<i>Contribution of entire term</i>
$\frac{\mathbb{V}ar(l_t)}{\mathbb{V}ar(y_t)}$	0.58	1.05	100%
Variances	0.56	0.96	51%
Covariances	0.61	1.17	49%
Variance Weight	0.57	0.58	
$(\omega_t = \sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt}) / \mathbb{V}ar(y_t))$			

Notes: results of the decomposition (16) in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017) for two-sector level of aggregation: goods sectors include mining, construction, all manufacturing sectors and utilities; services sectors are all remaining sectors. Employment is aggregated to the two sector level by a simple sum; value added is aggregated using a Tornqvist index. “Variances” refers to the variance component $\frac{\sum_{j=1}^N (\omega_{jt}^l)^2 \mathbb{V}ar(l_{jt})}{\sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt})}$. “Covariances” refers to the covariance component $\frac{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^l \omega_{ot}^l \text{Cov}(l_{jt}, l_{ot})}{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^y \omega_{ot}^y \text{Cov}(y_{jt}, y_{ot})}$. “Variance weight” refers to the weighting term $\omega_t = \sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt}) / \mathbb{V}ar(y_t)$. “Contribution of entire term” column computes the contribution of the first term of the decomposition (16) (in the variance row) and the contribution of the second term (in the covariance row). To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

more finely disaggregated set of sectors. Hence, this result suggests that our results are not driven by a broad process of structural transformation, and that the precise structure of the input-output networks are key to understanding our results.

Sectoral Decomposition of Investment Volatility Finally, Table 29 performs a similar decomposition of the increase in the relative volatility of investment:

$$\frac{\mathbb{V}ar(i_t)}{\mathbb{V}ar(y_t)} \approx \underbrace{\omega_t}_{\text{variance weight}} \underbrace{\frac{\sum_{j=1}^N (\omega_{jt}^i)^2 \mathbb{V}ar(i_{jt})}{\sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt})}}_{\text{variances}} + (1 - \omega_t) \underbrace{\frac{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^i \omega_{ot}^i \text{Cov}(i_{jt}, i_{ot})}{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^y \omega_{ot}^y \text{Cov}(y_{jt}, y_{ot})}}_{\text{covariances}} \quad (20)$$

where i_{jt} is real investment in sector j , i_t is aggregated real investment, and ω_{jt}^i is the share of total investment accounted for by sector j . As with employment, the average variance of investment within-sector is fairly stable relative to the average variance of value added; instead, more than 80% of the increase in the variance of aggregate investment is driven

TABLE 29
DECOMPOSITION OF INVESTMENT VOLATILITY

	<i>Pre-84</i>	<i>Post-84</i>	<i>Contribution of entire term</i>
$\frac{\mathbb{V}ar(i_t)}{\mathbb{V}ar(y_t)}$	3.77	8.49	100%
Variances	4.89	6.14	19%
Covariances	3.64	9.18	81%
Variance Weight	0.11	0.23	
$(\omega_t = \sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt}) / \mathbb{V}ar(y_t))$			

Notes: results of the decomposition (20) in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017) where i_{jt} is real investment and i_t is aggregated real investment (aggregated using a Tornqvist index). “Variances” refers to the variance component $\frac{\sum_{j=1}^N (\omega_{jt}^i)^2 \mathbb{V}ar(i_{jt})}{\sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt})}$. “Covariances” refers to the covariance component $\frac{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^i \omega_{ot}^i \text{Cov}(i_{jt}, i_{ot})}{\sum_{j=1}^N \sum_{o \neq j} \omega_{jt}^y \omega_{ot}^y \text{Cov}(y_{jt}, y_{ot})}$. “Variance weight” refers to the weighting term $\omega_t = \sum_{j=1}^N (\omega_{jt}^y)^2 \mathbb{V}ar(y_{jt}) / \mathbb{V}ar(y_t)$. “Contribution of entire term” column computes the contribution of the first term of the decomposition (16) (in the variance row) and the contribution of the second term (in the covariance row). To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing these figures.

by an increase in the comovement term. Furthermore, as with employment, the increase in this covariance term is driven by the fact that the covariance of investment across sectors is stable over time (not reported).

G Additional Quantitative Model Results

In this section, we present several additional results from our full quantitative model to ensure that its conclusions are robust.

Computing population moments The results presented in the main text feed in the realized time-series of sectoral TFP shocks; here, we show that the results also hold if we estimate the covariance matrix of those shocks separately for the pre vs. post 1984 subsamples and compute population moments from those two estimates. The main challenge is that we cannot estimate a full-rank covariance matrix with 35 sectors and less than 35 years of data pre- and post-1984. Therefore, following the same procedure described for the principal

TABLE 30
POPULATION MOMENTS, BASELINE MODEL VS. CHANGING PARAMETERS

	Population Moments		Changing Structure	
	<i>Pre-1984</i>	<i>Post-1984</i>	<i>Pre-1984</i>	<i>Post-1984</i>
$\sigma(y_t)$	2.68%	2.12%	3.13%	1.85%
$\rho(y_t - l_t, y_t)$	0.85	0.47	0.85	0.54
$\sigma(l_t)/\sigma(y_t)$	0.77	0.91	0.79	0.88
Variance contribution to change		15%		38%
Covariance contribution to change		85%		62%

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). “Population moments” refers to estimated the covariance matrix of TFP shock innovations separately on each subsample (as described in the Appendix text) and analytically computing HP-filtered moments. “Changing structure” refers to computing those population moments, but also allowing the other parameters of the model to change: the investment and intermediate networks, consumption shares, labor shares, intermediates shares, depreciation rates, and persistence of TFP. “Variance contribution to change” refers to the first term in the decomposition (16). “Covariance contribution to change” refers to the second term in (16).

components analysis in Appendix C, we collapse our data to 28 sectors by aggregating all non-durable manufacturing sectors into a single sector. We then estimate the covariance matrix of innovations to TFP separately for each subsample and analytically compute the model’s HP-filtered population moments separately for the pre-1984 and post-1984 periods. Table 30 show that the results are similar to those in the main text.

Allowing other parameters to change over time Table 30 show that our results are robust to allowing the non-shock parameters of the model - those governing the investment and intermediate networks, consumption shares, labor shares, intermediates shares, depreciation rates, and persistence of TFP – to change over time. In particular, we compute the average value of these parameters separately for the pre vs. post 1984 subsamples and compute the implied population moments given the covariance matrix of shocks estimated as above. Of course, a full analysis of these parameter changes would consider their full time path as part of the structural transformation process, but that analysis is outside the scope of this paper. These results simply show that our main results are robust to the simplest way in which to estimate how parameter values have changed over time.

TABLE 31
 QUANTITATIVE RESULTS, ALTERNATIVE ORDERS OF POLYNOMIAL DETRENDING OF TFP

	3rd order		2nd order	
	<i>Pre-1984</i>	<i>Post-1984</i>	<i>Pre-1984</i>	<i>Post-1984</i>
$\sigma(y_t)$	2.55%	2.13%	2.45%	2.01%
$\rho(y_t - l_t, y_t)$	0.91	0.52	0.93	0.62
$\sigma(l_t)/\sigma(y_t)$	0.72	0.90	0.69	0.87
Variance contribution to change		12%		11%
Covariance contribution to change		88%		89%

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). “3rd order” refers to detrending the raw Solow residuals using a third-order log-polynomial before feeding in detrended productivity to the model. “2nd order” refers to detrending the raw Solow residuals using a second-order log-polynomial before feeding in detrended productivity to the model. “Variance contribution to change” refers to the first term in the decomposition (16). “Covariance contribution to change” refers to the second term in (16).

Robustness to alternative order of detrending for TFP Table 31 show that our main results are robust to using a third or fourth order polynomial to detrend sector-level TFP, rather than a fourth order polynomial as in the main text.

Maintenance Investment As discussed in footnote 7, previous studies using the 1997 BEA capital flows table were forced to make a correction to the investment network in order to ensure the model is invertible. A motivation for this correction is to account for “maintenance investment” that may be a large part of investment activity but which is not accounted for in the BEA data (see McGrattan and Schmitz Jr (1999)). A key challenge in adjusting for maintenance is that it is not clear from which sector maintenance investment is purchased. One extreme is that maintenance is purchased from the same mix of sectors as the new investment recorded in our investment network; in this case, the investment network would not change. Another extreme is that all maintenance investment is purchased out of own-sector output; Foerster, Sarte and Watson (2011) make this assumption, and add a correction amounting to 25% of investment. Table 32 takes a case between these two extremes – adding a 10% of maintenance investment from own-sector output – and shows that our results continue to hold (similar results obtain with a larger 25% correction). The

TABLE 32
 QUANTITATIVE RESULTS, BASELINE MODEL VS. MAINTENANCE INVESTMENT

	Baseline Results		Maintenance Investment	
	<i>Pre-1984</i>	<i>Post-1984</i>	<i>Pre-1984</i>	<i>Post-1984</i>
$\sigma(y_t)$	2.60%	2.24%	2.58%	2.16%
$\rho(y_t - l_t, y_t)$	0.90	0.45	0.93	0.52
$\sigma(l_t)/\sigma(y_t)$	0.74	0.92	0.73	0.9
Variance contribution to change		11%		9%
Covariance contribution to change		89%		11%

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). “Baseline results” refers to exercise in main text. “Maintenance investment” refers to that model, plus an adjustment to the investment network to specify that 10% of investment is purchased from within own-sector output. “Variance contribution to change” refers to the first term in the decomposition (16). “Covariance contribution to change” refers to the second term in (16).

TABLE 33
 QUANTITATIVE RESULTS, BASELINE MODEL VS. ADJUSTMENT COSTS


	Baseline Results		Adjustment Costs	
	<i>Pre-1984</i>	<i>Post-1984</i>	<i>Pre-1984</i>	<i>Post-1984</i>
$\sigma(y_t)$	2.60%	2.24%	2.57%	2.18%
$\rho(y_t - l_t, y_t)$	0.90	0.45	0.91	0.56
$\sigma(l_t)/\sigma(y_t)$	0.74	0.92	0.73	0.89
Variance contribution to change		11%		10%
Covariance contribution to change		89%		90%

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). “Baseline results” refers to exercise in main text. “Adjustment costs” refer to uniform quadratic adjustment cost with parameter $\phi = 0.4$ chose to match the correlation of investment across sectors. “Variance contribution to change” refers to the first term in the decomposition (16). “Covariance contribution to change” refers to the second term in (16).

fact that each sector now uses its own output for investment weakens the strength of the investment hubs somewhat, but quantitative the model still generates a sizable decrease in the correlation of labor productivity and aggregate GDP and a sizable increase in relative employment volatility.

Capital adjustment costs Our baseline model does not include capital adjustment costs. While our model matches the overall level of investment volatility, the correlation of invest-

FIGURE 25: Aggregate Dynamics Driven by Shocks to Individual Sectors



Notes: results from simulating model with empirical shocks to only one sector at a time (the remaining sectors’ shocks are set to zero). Plots standard deviation of aggregate employment n_t relative to the standard deviation of the sectors’ productivity A_{jt} . To correct for sector size, the figure then divides by steady state share of nominal value added. All variables have been logged and HP-filtered with smoothing parameter 6.25. Investment hubs are highlighted in red.


ment across sectors is counterfactually low, reflecting the fact that investment is extremely responsive to sector-specific shocks. We solve this problem by incorporating capital adjustment costs at the sector-level:

$$K_{jt+1} = (1 - \delta_j)K_{jt} + I_{jt} - \frac{\phi}{2} \left(\frac{I_{jt}}{K_{jt}} - \delta_j \right)^2 K_{jt} \quad (21)$$

We set the adjustment cost parameter $\phi = 0.4$ to match the roughly match the average correlation of HP-filtered real investment across sectors. Table 33 shows that our results are robust to allowing for capital adjustment costs.

Size-adjusted reduced-form elasticity of aggregate employment with respect to sector-specific shocks Figure 11 in the main text plots the implied reduced-form elasticity of aggregate employment with respect to shocks to individual sectors. Figure 25 computes the same object but divides by the sector’s average share of value added in order to account for the fact that some sectors are larger than others and, therefore, will mechanically have a

FIGURE 26: Distributional Effect of Investment Stimulus



Notes: effect of a one-time $sub_t = 0.01$ shock to the stimulus policy shock described in the main text. Each bar plots the percentage change in employment at that particular sector. Red bars are the investment hubs' response in our baseline model, blue bars are the non hubs' response in our baseline model, and transparent grey bars are the responses in a version of the model in which we eliminate the investment network by assuming all investment is done out of own-sector output.

larger effect on aggregate employment. The investment hubs continue to have a substantial effect on aggregate employment. There are two main differences from Figure 11 in the main text. First, the suppliers of investment hubs (the manufacturing sectors in the right of the x-axis) now have a larger effect on aggregate employment. Second, the service sectors (to the right of the x-axis) have a smaller effect on aggregate employment, reflecting the fact that these sectors account for a larger share of value added.

Distributional effects of investment stimulus policy Figure 26 plots the percentage change in sector-level employment in response to a 1% investment subsidy shock. Similarly to Figure 13 in the main text, Figure 26 shows that employment at investment hubs and their suppliers are the most responsive to investment stimulus. Without the investment network, the effect of the investment stimulus is fairly uniformly distributed across sectors. The fact that the effect is less uniformly distributed Figure 13 simply reflects the fact that the service sectors (to the right of the x-axis) simply account for a larger share of economic activity.