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

Christian vom Lehn
Brigham Young University and IZA
cvomlehn@byu.edu

Thomas Winberry
Wharton and NBER
twinb@wharton.upenn.edu

April 6, 2021

Abstract

We argue that the network of investment production and purchases across sectors is an important propagation mechanism for understanding business cycles. Empirically, we show that the majority of investment goods are produced by a few "investment hubs" which are more cyclical than other sectors. We embed this investment network into a multisector business cycle model and show that sector-specific shocks to the investment hubs and their key suppliers have large effects on aggregate employment and drive down labor productivity. Quantitatively, we find that sector-specific shocks to hubs and their suppliers account for an increasing share of aggregate fluctuations over time, generating the declining cyclicality of labor productivity and other changes in business cycle patterns since the 1980s.

*We thank Rodrigo Adao, Gadi Barlevy, Bob Barsky, David Baqaee, David Berger, Steve Davis, Emmanuel Farhi, Simon Gilchrist, Mike Golosov, Francois Gourio, Veronica Guerrieri, Jonathan Heathcote, Erik Hurst, Anil Kashyap, Jennifer La'O, Ernest Liu, Lilia Maliar, Brent Neiman, Pablo Ottonello, Richard Rogerson, Matt Rognlie, Elisa Rubbo, Rob Shimer, Henry Siu, Joe Vavra, Christian Wolf, and seminar participants at various institutions for useful suggestions. Winberry gratefully acknowledges financial support from Chicago Booth.

1 Introduction

The defining feature of business cycles is the comovement of production across different sectors of the economy. However, recent work has shown that the degree of sectoral comovement has fallen since the early 1980s, suggesting that sector-specific shocks have become more volatile relative to aggregate shocks.[1] Our basic questions are how these sector-specific shocks are propagated to macroeconomic aggregates and whether their rising importance helps understand the changing nature of business cycles since the early 1980s. Of course, a large literature studies the role of the input-output network of intermediate goods in propagating sector-specific shocks. Given the importance of investment in business cycle fluctuations, we instead focus on the role of the *investment network* — the distribution of investment production and purchases across sectors — in propagating these shocks.

We argue that the investment network is an important propagation mechanism for understanding business cycle fluctuations in three main steps. First, we measure the investment network in the data and show that investment production is dominated by a small number of *investment hubs* which are substantially more cyclical than other sectors. Second, we embed our measured investment network into a standard multisector real business cycle model and show that sector-specific shocks to investment hubs and their key intermediates suppliers have large effects on aggregate employment, driving down labor productivity. Third, we measure the realized time series of sector-level shocks in the data, feed them into our model, and show that shocks to the hubs and their key suppliers account for a large and increasing share of aggregate fluctuations over time. This fact allows the model to generate the declining cyclicality of labor productivity and other changes in business cycle patterns since the early 1980s — despite the fact that the model has flexible prices and frictionless labor adjustment.

The first step in our analysis is to measure the investment network, which we define as the amount of investment goods that are produced in sector i and subsequently sold to sector j for each pair of sectors (i, j). While the BEA has released this information in its capital flows tables, those tables are only available for a small subset of years, do not include the majority

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

of intellectual property, and are not consistently coded across time. We therefore perform our own measurement of the investment network building on disaggregated asset-level data for each sector. Our network covers a 37-sector disaggregation of the entire private nonfarm economy, incorporates all of intellectual property, and is available each year between 1947-2018. For most of our analysis in this paper, we will average the network over time and refer to the averaged network as "the" investment network. We have constructed alternative investment networks which incorporate agriculture and government sectors; separate equipment, structures, and intellectual property products; and make other adjustments that may be of interest to other researchers. We have also constructed the network of capital rental services across sectors.

Our measured investment network is extremely sparse; four investment hubs — construction, machinery manufacturing, motor vehicles manufacturing, and professional/technical services — produce nearly 70% of total investment even though they only account for 15% of value added, employment, or intermediates production. Production and employment in these hubs are more volatile, more correlated with aggregates, and more strongly lead the aggregate cycle than in non-hub sectors, consistent with the hubs' central role in our model.

The second step of our analysis is to incorporate this measured investment network into a version of the multisector real business cycle framework from Horvath (2000). Each sector produces gross output using capital, labor, and a bundle of intermediate goods consisting of other sectors' output; this bundle is a Cobb-Douglas aggregate which characterizes the intermediates input-output network. Each sector also accumulates new capital using another Cobb-Douglas aggregator of investment goods, which characterizes the investment network. While other studies have also used this basic model structure, we discipline it with our new measurement of the investment network, explicitly study the network's role in propagating sector-specific shocks, and show that it quantitatively accounts for the declining cyclicality of labor productivity and other changes in business cycle patterns over time.

Our main new mechanism from this model is that shocks to investment hubs and their key suppliers generate large changes in aggregate employment while shocks to other sectors do not. This result reflects the fact that a sector-specific shock only affects employment if it increases the production of investment goods in the economy; shocks that only affect the production of consumption goods generate offsetting income and substitution effects, leaving employment unchanged. We show that the importance of each sector in producing investment goods can be summarized using the *Leontief-adjusted investment network*, which accounts for both directly producing investment goods and indirectly supplying intermediate goods to investment producers. Shocks to hubs and their key suppliers in this network thus act as aggregate investment supply shocks.[2] In contrast, shocks to other sectors act as idiosyncratic investment demand shocks, which do not generate large changes in employment.

Our third step is to quantify the importance of this mechanism in explaining the postwar U.S. time series by feeding the realizations of sector-level productivity shocks into a calibrated version of the model. Since the early 1980s, the covariance of productivity shocks across sectors has fallen by much more than the variance of shocks within sectors, which we interpret as reflecting a decline in the volatility of aggregate shocks relative to the volatility of sector-specific shocks. This change is consistent with the decline in aggregate volatility following the Great Moderation (see, e.g., Foerster, Sarte and Watson (2011)). In order to isolate the role of this change in the shock process in driving our results, we hold all other parameters of the model, including the investment network, fixed over time for our baseline analysis.

We find that the rising importance of sector-specific shocks, when propagated through the investment network, quantitatively generates the declining cyclicality of labor productivity and other business cycle changes since the 1980s. The pre-1980s sample features procyclical labor productivity because it is dominated by aggregate TFP shocks. However, since sector-specific shocks become more important after the 1980s, shocks to investment hubs and their suppliers account for a larger share of employment fluctuations over that period. These shocks drive down labor productivity because they increase aggregate employment by more than GDP, thereby generating the declining cyclicality of labor productivity observed in the data. Our model also generates the decline in the volatility of GDP and the increased volatility of employment relative to GDP.

[2] Our 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 their aggregate effect. Our model generates positive comovement through the intermediate inputs linkages in the Leontief-adjusted investment network, which we discuss in Appendix F.

These quantitative results are robust to a number of model extensions. First, they are robust to allowing for trend changes in the investment network and other structural parameters, indicating that the rising importance of sector-specific shocks is the key force driving these changes in business cycle patterns. Second, our results also hold in a second-order approximation with CES production functions and preferences, which allows for richer non-linearities (see, e.g. Baqaee and Farhi (2019)). Third, our results are robust to various forms of adjustment frictions in labor and capital markets.

Finally, we document two new empirical results which support the role of the investment network in accounting for the changes in business cycle patterns since the 1980s. First, we show that the volatility of investment relative to the volatility of GDP has substantially increased since the 1980s, consistent with the idea that sector-specific shocks to investment suppliers play a more important role over time. Second, we show that the changes in business cycle patterns have not occurred within individual sectors but are due to changes in the comovement of activity across sectors. For example, the entire decline in the cyclicality of aggregate labor productivity is due to changes in the covariance of value added and employment across sectors; sector-level labor productivity is still highly procyclical within sector. Our model matches these changing covariance patterns due to the declining importance of aggregate shocks and the sparseness of the investment network. In contrast, existing explanations for the declining cyclicality of labor productivity largely abstract from sectoral heterogeneity and therefore do not speak to this empirical result.

Related Literature Our paper builds on three lines of existing research. The first uses the multisector real business cycle model to study how connections between sectors propagate sector-specific shocks to macroeconomic aggregates. Our model's basic structure builds on Horvath (2000), as do many others in the literature (see, for example, Foerster, Sarte and Watson (2011) and Atalay (2017)).[3] We make three main contributions to this literature.

[3]Horvath (2000) varies the number of sectors $N \in \{6, 21, 36, 77\}$ and finds that the aggregate volatility generated by purely i.i.d. sectoral shocks declines more slowly than $1/N$, which Horvath (2000) interprets as the intermediate network "postponing" the law of large numbers. While we have found that the investment network plays an important role in that exercise (results available upon request), it is fundamentally different from our own, which fixes the number of sectors $N = 37$ at the finest level of disaggregation for which we can parameterize the model and studies the effects of the empirical shock process. As we describe in Section 4, the key mechanism driving our results is that, among sector-specific shocks, it is primarily those to the

First, we focus on the investment network rather than on the input-output network of intermediate goods. While a number of other papers also include an investment network, they do not analyze its role in propagating sector-specific shocks.[4] Second, our new measurement of the investment network provides an annual time series of the network, includes all of intellectual property, and is consistently coded over the entire postwar sample.[5] Third, we show that shocks to investment hubs and their key suppliers decrease labor productivity and that their rising importance over time accounts for the declining cyclicality of labor productivity and other changes in business cycle patterns observed in the post-1980s data.

The second line of related research is the fast-growing networks literature which studies how richer input-output networks in intermediate goods propagate idiosyncratic shocks to macroeconomic aggregates (see, for example, Acemoglu et al. (2012), Acemoglu, Ozdaglar and Tahbaz-Salehi (2017), Baqaee and Farhi (2019), Baqaee and Farhi (2020), Bigio and La'o (2020), or the survey in Carvalho and Tahbaz-Salehi (2019)). In order to allow for rich network structures, these papers use static models which abstract from investment. A natural benchmark in these models is a strong version of Hulten's theorem: under Cobb-Douglas preferences/production and competitive/frictionless markets, the effect of a sector-specific shock on real GDP is globally equal to the sector's Domar weight, which is constant. The literature has shown how deviations from Cobb-Douglas production (e.g. Baqaee and Farhi (2019)) or from competitive/frictionless markets (e.g. Baqaee and Farhi (2020) or Bigio and La'o (2020)) can break this version of Hulten's theorem. We show that the presence of investment also breaks Hulten's theorem because the capital accumulation technology is not Cobb-Douglas. Furthermore, we characterize how the investment network interacts with the

investment hubs and their key suppliers which drive aggregate employment fluctuations.

[4] In a recent complementary paper, Foerster et al. (2020) use the same basic model structure and show that long-run changes in productivity growth rates at investment-producing sectors have a large impact on long-run GDP growth by increasing capital accumulation. This analysis differs from ours in two key respects. First, we focus on short-run business cycle fluctuations rather than long-run growth. Second, shocks to investment producers are important in our model not because they generate large changes in the capital stock but because they generate large changes in employment; Foerster et al. (2020) abstract from this mechanism by taking employment (our main outcome of interest) as exogenous. In addition, like other papers in this literature, they measure the investment network using the 1997 capital flows table (see Footnote 5).

[5] Foerster, Sarte and Watson (2011) and Atalay (2017) calibrate the investment network using the BEA capital flows data from 1997, which excludes the majority of intellectual property. They are also forced to make an adjustment to ensure their model is invertible but which artificially reduces the importance of the network. We do not require any ad-hoc adjustment to our model.

intermediates input-output network using our Leontief-adjusted investment network.

The final line of related literature studies how business cycle patterns have changed since the 1980s and whether the real business cycle framework can explain those patterns. A large subset of this literature focuses, in particular, on the declining cyclicality of labor productivity and has suggested roughly three sets of explanations: the first is that the aggregate shock process has changed over time (see, for example, Galí and Gambetti (2009) or Barnichon (2010)), the second is that firms and/or workers can now more easily adjust labor inputs in response to shocks (see, for example, Galí and Van Rens (2021), Koenders and Rogerson (2005), Berger et al. (2012), or Bachmann (2012)), and the third is that there has been no actual change in the cyclicality of labor productivity, but that (mis)measurement of those objects has changed (see, for exmaple, Fernald and Wang (2016), McGrattan and Prescott (2014), or McGrattan (2020)). This literature typically constructs models without sectoral heterogeneity and therefore cannot speak to our empirical finding that the entire decline in the cyclicality of labor productivity is due to changes in the covariance of activity across sectors.[6] More generally, we show that the investment network can reconcile a real business cycle framework with key features of business cycles since the 1980s.

Road Map Our paper is organized as follows. We measure the empirical investment network and document the cyclical behavior of investment hubs in Section 2. We describe our version of the multisector real business cycle model and calibrate it to match the measured investment network in Section 3. In Section 4, we show that shocks to investment hubs and their suppliers have large effects on aggregate employment, driving down labor productivity, while shocks to other sectors have small aggregate effects. In Section 5, we feed the realized time series of sector-level productivity into the model and show that the rising importance of sector-specific shocks generates the declining cyclicality of labor productivity since the 1980s. We provide empirical support for this mechanism in Section 6, which shows that those aggregate changes have not occurred within sector but are driven by changes in sectoral comovement (consistent with our model). Section 7 concludes.

[6] We are aware of one paper which studies the declining cyclicality of labor productivity in a model with sectoral heterogeneity: Garin, Pries and Sims (2018). However, Garin, Pries and Sims (2018)'s mechanism relies on sector-specific shocks generating negative employment comovement across sectors, inconsistent with the fact that employment comovement is positive and stable over time (which we show in Appendix H).

Table 1
The 37 Sectors Used in Our Analysis

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

Notes: list of sectors used in our empirical analysis. Sectors are classified according to the NAICS-based BEA codes. See Appendix A.1 for details of the data construction.

2 Descriptive Evidence on the Investment Network

We combine three sources of sector-level data for our empirical work. We construct the investment network using the BEA Fixed Assets and Input-Output databases for a sample of 37 private non-farm sectors from 1947-2018 (our construction of the investment network is described below). We use the BEA GDP-by-Industry database to obtain annual observations of 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 must 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 - 2018 period. Appendix A.1 contains details about the construction of our dataset.[7]

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

[7]We must use annual data because quarterly observations of value added, employment, and investment are not available at the sectoral level over the entire postwar period.

that it covers the entire postwar sample, which is necessary to analyze changes in business cycle patterns over time. In addition, the partition of sectors provides fairly detailed coverage of the private nonfarm economy. We cannot disaggregate the sectors much more finely in a consistently-defined way over time and retain coverage of the entire postwar time period.

2.1 Empirical Investment Network

We define the *investment network* in year t as the share of the total investment expenditure of a given sector j that is purchased from another sector i for each pair of sectors (i, j) in the economy. While the BEA capital flows tables provides information about these pairwise flows, those tables have three key shortcomings. First, the BEA tables are only available for a handful of years, most recently 1997. Second, the sectoral disaggregation used in the various BEA tables is not consistently defined over time. Third, and most importantly, the BEA tables do not include all of intellectual property; in fact, the 1997 table is the only one which records any intellectual property at all, but it only includes software (which was a third of all intellectual property investment in that year).

We construct our own measurement of the investment network which overcomes these issues. Our construction is based on disaggregated asset-level data which records the purchases of 33 types of capital assets for each sector in each year. We estimate a series of "bridge files" which allocate the production of each of these assets to a mix of producing sectors. Appendix A.2 describes our procedure for estimating the bridge files, which follows BEA practice as closely as possible.[8]

To our knowledge, our investment network is the only version of the capital flows tables that is available in every year 1947-2018, is consistently defined over that period, and is consistent with modern national accounting practices regarding intellectual property. We also provide a number of alternative tables which may be of interest to other researchers. First, we provide an investment network which also includes agriculture, federal government, and state/local government sectors. Second, we provide an investment network that adds

[8]The investment expenditures data includes imported capital, so in this sense our investment network accounts for the fact that the share of imported capital has increased over time (see House, Mocanu and Shapiro (2017)).

FIGURE 1: Heatmap of Empirical Investment Network

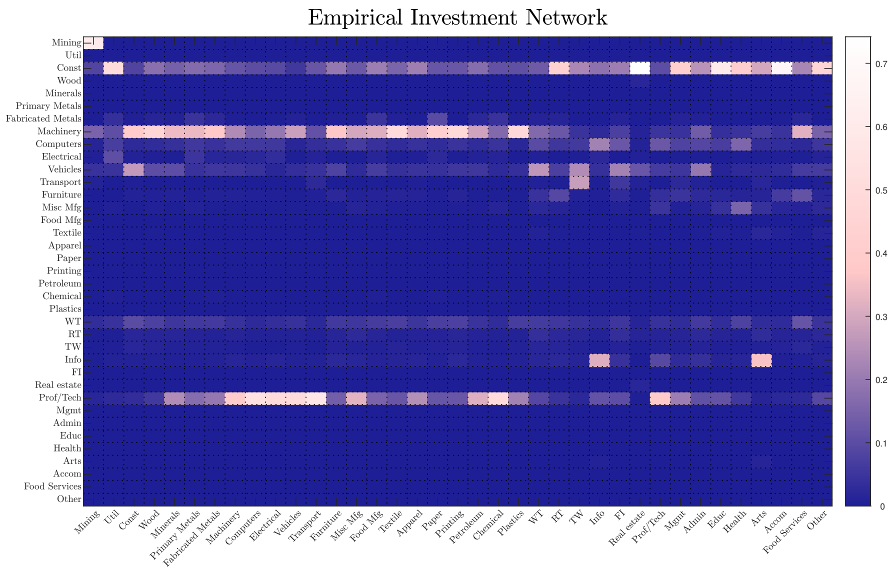

Notes: heatmap of empirical investment network. Entry (i, j) computes share of total investment expenditure in sector j that is purchased from sector i, averaged over the 1947 - 2018 sample.

an ad-hoc adjustment for estimates of maintenance investment following Foerster, Sarte and Watson (2011) and Atalay (2017).[9] Third, we provide analogous tables for capital rental services, which may be useful in calibrating static models with capital (but without investment) or in constructing a measure of national income along the lines of Barro (2021). Finally, we also provide the asset-level bridge files used to construct the network.

Investment Network is Highly Concentrated Figure 1 plots a heatmap of our investment network averaged over time. Four sectors supply the majority of investment goods to the rest of 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 the majority of intellectual property. We

[9]As described in Footnote 5, Foerster, Sarte and Watson (2011) and Atalay (2017) add an adjustment to the investment network implied by the 1997 BEA capital flows table to ensure their models are invertible (though Horvath (2000) does not). This adjustment is meant to account for maintenance investment that is done out of own-sector output. While there is evidence that maintenance investment are sizable (see McGrattan and Schmitz Jr (1999)), there are not estimates of how much of maintenance investment is done from own-sector output because maintenance is largely not recorded in the national accounts. Therefore, we prefer not to add an artificial adjustment for maintenance investment in our baseline analysis; however, Appendix G shows that our model results are robust to adding this correction.

TABLE 2
VOLATILITY OF ACTIVITY AT INVESTMENT HUBS

	Investment Hubs		**Non-Hubs**	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(\Delta y_{st})$	9.13%	9.18%	6.63%	5.51%
$\sigma(\Delta l_{st})$	6.14%	4.83%	3.81%	3.14%
$\sigma(y_{st}^{hp})$	5.64%	6.29%	3.91%	3.40%
$\sigma(l_{st}^{hp})$	4.08%	3.21%	2.29%	1.91%

Notes: standard deviation of sector-level value added or employment. y_{st} is logged real value added in sector s and l_{st} is logged employment in sector s. $\sigma(\Delta y_{st})$ and $\sigma(\Delta l_{st})$ refer to the standard deviation of the first-differences of these variables, while $\sigma(y_{st}^{hp})$ and $\sigma(l_{st}^{hp})$ refer to the standard deviation of the HP-filtered variables with smoothing parameter 6.25 for annual data. "Investment hubs" computes the unweighted average 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 - 2018 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 the HP-filtered statistics.

refer to these four sectors as *investment hubs*. Together, these hubs produce approximately 70% of the investment goods produced in the economy, even though they only account for approximately 15% of value added produced, intermediate goods produced, or workers employed. The fact that this small number of hubs produce the majority of investment implies that the investment network is highly concentrated; in fact, Appendix A.3 shows that the investment network is two to three times more concentrated than the intermediates network according the skewness of their eigenvalue centralities or weighted outdegrees.

Appendix A.3 analyzes how the investment network has changed over time. The primary change has been the rising importance of professional/technical services as an investment supplier due to the rising importance of intellectual property products. While these changes are important for long-run trends, we focus on the average investment network for the business cycle analysis in this paper.

Investment Hubs are Highly Cyclical Table 2 shows that employment and real value added produced at investment hubs are more volatile over the business cycle than those at non-hubs. We measure business cycle volatility using log-first differences or the HP filter.

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

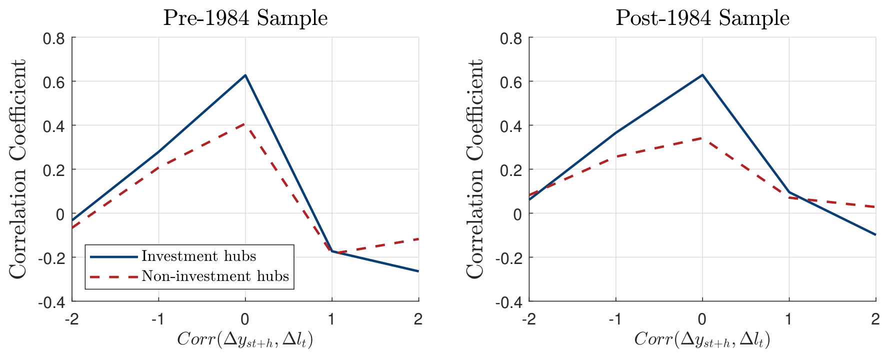

Notes: correlation of value added growth in sector s in year $t+h$, Δy_{st+h}, with aggregate employment growth in year t, Δl_t. Both y_{st+h} and l_t are logged and Δ denotes the first-difference operator. The x-axis varies the lead/lag $h \in \{-2, -1, 0, 1, 2\}$. "Investment hubs" compute the unweighted average 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 - 2018 subsample.

Under either transformation of the data, the investment hubs are approximately 1.5 - 2 times as volatile as non-hub sectors in both the pre- and post-1984 subsamples.[10] For the rest of the paper, we will use log-first differences to analyze business cycle fluctuations in order to avoid the issues with two-sided filters explained in e.g. Hamilton (2018). However, all our results are robust to using the HP filter, and we present those results from time to time to help compare our results to previous studies.

Figure 2 shows that these fluctuations at the investment hubs are more correlated with the aggregate business cycle. We compute the correlogram of sector-level real value added growth in year $t+h$ with aggregate employment growth in year t.[11] At most horizons, investment hubs' value added is more correlated with aggregate employment than is non-hubs' value added. The difference is larger in the post-1984 subsample, consistent with the idea that shocks to investment hubs have become more important for aggregate fluctuations

[10]We compute these statistics as the unweighted average across sectors in order to focus on the volatility of the average sector. Of course, aggregate value added and employment, which we analyze in Section 5, also depends on the share of activity in the various sectors.

[11]We use aggregate employment growth as our proxy for the aggregate cycle because our model predicts that shocks at investment hubs have a larger impact on aggregate employment than GDP. Nevertheless, Appendix B shows that similar, but slightly weaker, results hold when using GDP to proxy for the aggregate cycle.

over time. In addition, investment hubs more strongly lead the aggregate cycle than do non-hubs.[12]

3 Model and Calibration

We now develop and 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 and based on Horvath (2000).

Environment Time is discrete and infinite. There are a finite number of sectors indexed by $j = \{1, ..., N\}$, where $N = 37$ as in our data. Each sector produces gross output using the production function

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

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

$$\log A_{jt+1} = \rho_j \log A_{jt} + \varepsilon_{jt+1}, \tag{2}$$

where ρ_j is the persistence and ε_{jt} are innovations (which can be correlated across sectors).

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

$$M_{jt} = \Pi_{i=1}^{N} M_{ijt}^{\gamma_{ij}}, \quad \text{where } \sum_{i=1}^{N} \gamma_{ij} = 1, \tag{3}$$

[12] Appendix B shows that non-hub manufacturing sectors' behavior is more similar to the other non-hub sectors than it is to the investment hubs. This result allays the concern that our results are driven by the fact that two of our four investment hubs are manufacturing sectors, and that manufacturing may be more cyclical than other sectors for reasons outside of our model. Furthermore, Appendix F shows that the extent to which manufacturing sectors differ from other non-hub sectors is largely explained by their role as suppliers of intermediate goods to the investment hubs, consistent with our model.

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, each sector j observes the TFP shock A_{jt}, uses its pre-existing stock of capital K_{jt}, hires labor L_{jt} from a competitive labor market, and purchases intermediates M_{ijt} in competitive markets in order to produce gross output Q_{jt}.

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

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

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

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

where I_{ijt} is the amount of sector i's output used by sector j and λ_{ij} are parameters. Investment hub sectors i have high λ_{ij} for many purchasing sectors j. We denote the investment network matrix as $\Lambda = [\lambda_{ij}]$.

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

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

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.

Equilibrium We study the competitive equilibrium, which is efficient. Output market clearing for sector j ensures that gross output is used for final consumption, investment, or an intermediate in production:

$$Q_{jt} = C_{jt} + \sum_{i=1}^{N} I_{jit} + \sum_{i=1}^{N} M_{jit}. \tag{7}$$

Labor market clearing ensures that aggregate labor demand equals labor supply: $\sum_{j=1}^{N} L_{jt} = L_t$. We denote the price of sector j's output by p_{jt}, the rental rate on sector j's capital by r_{jt}, and the wage rate by W_t (which is common across sectors since labor is perfectly substitutable). We take the price index of the household's consumption bundle as the numeraire. See Appendix C for more details on the equilibrium conditions.[13]

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 nearly one hundred state variables). In addition, the linear solution features certainty equivalence, so that the covariance matrix of these innovations does not affect the decision rules. This property allows us to simply feed the empirical time series of realized shocks into the decision rules without needing to estimate how the entire covariance structure of shocks has changed over time. However, linearization implies that we do not capture potential nonlinearities, such as size- or state- dependent responses to shocks. We show that our results are robust to allowing for nonlinearities in Appendix G.

3.2 Remarks on Simplifying Assumptions

We have made a number of simplifying assumptions in our model specification. For example, Cobb-Douglas preferences impose that the elasticity of substitution across different sectors' consumption goods is one, while Cobb-Douglas production technologies impose that the elasticity between capital, labor, and intermediates are also one. We have also assumed that there are no adjustment frictions to capital or labor, either across sectors or over time (though we will add a simple capital reallocation friction in our quantitative analysis in Section 5).

As will become clear in Section 4, these stark assumptions allow us to clearly explain the contribution of the investment network in propagating sector-specific shocks. In fact, without investment, the Cobb-Douglas assumptions imply that employment is constant in

[13] While we assume that the household owns the capital stock, the equilibrium allocation would be identical in the version of the model in which firms own their own capital stock and make their own investment decisions. The reason is that, in this alternative model, the household still owns the firms, so the firms value output in all states of the world using the household's stochastic discount factor. All our propositions in Section 4 also go through in this alternative model provided that one defines the shadow rental rate $\widetilde{r}_{jt} = \alpha_j \theta_j \frac{p_{jt} Q_{jt}}{K_{jt}}$ to be the ex-post return to capital in sector j.

response to shocks and the effect of these shocks on real GDP is given by the sector's Domar weight. Hence, our model is a useful benchmark for understanding the role of investment — and the investment network — in driving employment fluctuations. Nevertheless, we show in Appendix G that our main results are robust to relaxing these simplifying assumptions.

We also assume that all structural parameters of the model are constant over time, so that the only force which generates changing business cycle patterns is changes in the process generating sector-level productivity. We show in Appendix G that our results are robust to allowing the other structural parameters to change over time as well. We interpret this finding as indicating that changes in these other parameters are secondary for understanding the aggregate business cycle fluctuations that we study (though they may be important for understanding long-run changes or other business cycle features).

3.3 Calibration

We calibrate the structural parameters of the model so that the model's steady state matches key empirical targets averaged over the postwar sample. A model period is one year. We identify the $N = 37$ sectors in our model with those in our empirical work, and therefore use the BEA input-output database to identify 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 average labor compensation (adjusted for taxes and self employment) as a share of total value added.

The parameters of the intermediates input-output network γ_{ij} are given by sector j's expenditure on intermediates from sector i as a share of its total intermediates expenditure, averaged over the years 1947-2018. Figure 3 plots the heatmap of our calibrated intermediates network. It has a strong diagonal element, capturing firms' purchases of intermediates from within their own sector, but is also richly populated off the diagonal, capturing intermediates purchased from other sectors.[14]

The parameters of the investment network, λ_{ij}, are equal to the share of sector j's total

[14] Our measured intermediates and investment purchases account for goods that are imported from sectors outside the U.S, but the model counterfactually assumes that all factor supply is domestically produced. While extending our model to an open economy framework would be an interesting exercise, it is outside the scope of this paper.

Figure 3: Heatmap of Intermediates Network

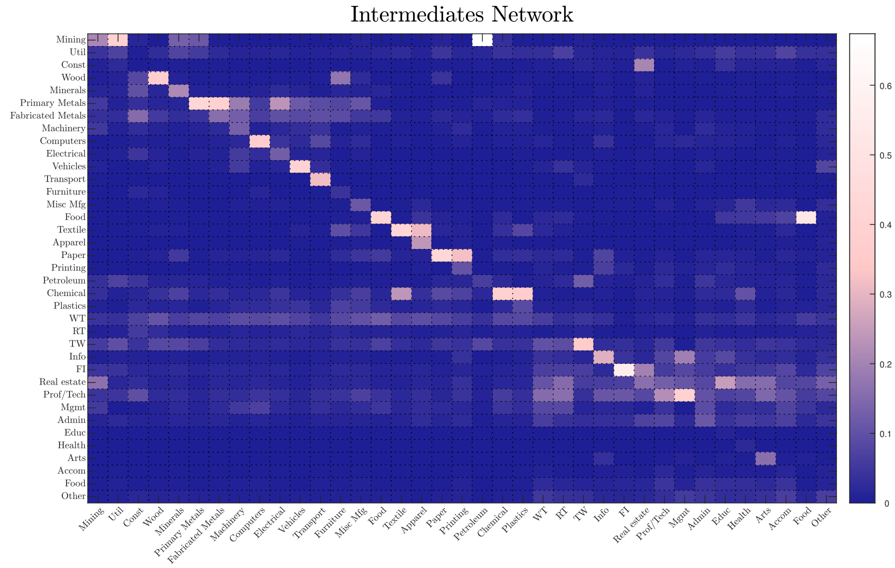

Notes: heatmap of intermediates input-output network γ_{ij} constructed as described in Appendix D. The (i, j) entry of each network corresponds to parameter γ_{ij}, i.e. the share of intermediate expenditures by sector j on goods produced by sector i averaged across the years 1947-2018.

investment expenditure that is purchased from sector i, averaged over time – already plotted in Figure 1. Capital depreciation rates δ_j are the implied depreciation rates for each sector, based on average annual depreciation of each capital good and the average amount of each type of good used in sector j.

The consumption shares ξ_j are given by the average consumption expenditure on sector j's output as a fraction of total consumption expenditure. We set the discount factor to $\beta = 0.96$. We normalize the disutility of labor parameter to $\chi = 1$. We take the Frisch elasticity $\eta \to \infty$ to capture indivisible labor at the individual level, as in Rogerson (1988), since we analyze fluctuations in employment rather than hours.

Appendix D shows that there is substantial heterogeneity in our calibrated parameters across sectors. The share of intermediates $1 - \theta_j$ in production ranges from 70-80% in many durable manufacturing sectors to only 20-30% in sectors like wholesale trade, retail trade, and real estate. Labor's share in production $1 - \alpha_j$ is lowest in real estate, whose primary value added is transferring structures capital to consumers, and highest in construction, professional/technical services, and the management of business and enterprises. Heterogeneity in the depreciation rate δ_j across sectors captures heterogeneity in the mix of capital goods used across sectors; structures-intensive sectors like real estate, utilities, and education have low deprecation rates while equipment-intensive sectors like computer and motor vehicle manufacturing have high depreciation rates. Finally, our calibrated consumption shares ξ_j show that the majority of the household's consumption basket comes from real estate, retail trade, health care, food manufacturing, and food services.

4 Role of Investment Network in Propagating Sector-Specific Shocks

Before turning to our quantitative analysis, we explain the theoretical mechanisms through which sector-specific shocks affect employment, GDP, and labor productivity.

4.1 Aggregation of Sector-Level Outcomes

Our first step is to define real GDP and aggregate employment in our multisector model. While it is straightforward to compute aggregate employment $L_t = \sum_{j=1}^{N} L_{jt}$, it is more difficult to compute real GDP because relative prices change over time. We follow national accounting practices and define real GDP using a Divisia index. The Divisia index begins with the definition of nominal GDP $P_t^Y Y_t = \sum_{j=1}^{N} p_{jt}^Y Y_{jt}$, where P_t^Y and p_{jt}^Y are price indices for aggregate and sector-level value added and Y_{jt} is sector-level real value added (defined in Appendix C). The Divisia index then computes real GDP growth as the log-change in nominal value added, holding prices fixed:[15]

$$d \log Y_t = \sum_{j=1}^{N} \left(\frac{p_{jt}^Y Y_{jt}}{P_t^Y Y_t} \right) d \log Y_{jt}. \tag{8}$$

In our model, sector-level value added is equal to payments to the primary inputs because there are no economic profits. Appendix C shows that these payments depend only on TFP

[15]The Divisia index is defined in continuous time while our model is in discrete time. For the purposes of simplifying exposition here, we do not take a stance on the exact discrete time approximation to the continuous time Divisia index used, but we use a Tornqvist index in our quantitative analysis.

and the primary inputs themselves: $d\log Y_{jt} = \frac{1}{\theta_j} d\log A_{jt} + \alpha_j d\log K_{jt} + (1-\alpha_j) d\log L_{jt}$. Plugging this expression into the Divisia index (8) implies the following proposition:

***Proposition* 1.** *Up to first order, the impact effect of a sector-specific shock A_{it} on real GDP Y_t is determined by*

$$d\log Y_t = \underbrace{\sum_{j=1}^{N} \left(\frac{p_j Q_j}{P^Y Y}\right)^* d\log A_{jt}}_{\equiv d\log TFP_t} + (1-\alpha^*) \underbrace{\sum_{j=1}^{N} \left(\frac{L_j}{L}\right)^* d\log L_{jt}}_{\equiv d\log L_t} . \tag{9}$$

where $\left(\frac{p_j Q_j}{P^Y Y}\right)^*$ is the ratio of sector j's sales to nominal GDP in steady state (its Domar weight*), $\left(\frac{L_j}{L}\right)^*$ is sector j's employment share in steady state, and $1-\alpha^* = \sum_{j=1}^{N}(1-\alpha_j)\left(\frac{p_j^Y Y_j}{P^Y Y}\right)^* = \left(\frac{WL}{P^Y Y}\right)^*$ is the aggregate labor income share in steady state.*

Proof. See Appendix E. ■

Proposition 1 shows that the effect of a sector-specific shock in some sector i, A_{it}, on the Divisia index can be decomposed into the shock's effect on aggregate TFP, $d\log TFP_t$, and its effect on aggregate employment, $d\log L_t$ (capital does not enter this expression since it is fixed upon impact). Aggregate TFP is the sum of sector-level TFP weighted by the sectors' steady state Domar weights $\left(\frac{p_j Q_j}{P^Y Y}\right)^*$ (Hulten, 1978).[16] The insight of Hulten's theorem is that the Domar weight is a sufficient statistic for capturing how a shock to a given sector propagates to the other sectors through the input-output network of intermediate goods.

Since Hulten's theorem for aggregate TFP is well understood, we will instead focus our analysis on understanding the endogenous response of aggregate employment. Under our preference specification, equilibrium employment in sector j is given by

$$L_{jt} = (1-\alpha_j)\theta_j \frac{p_{jt} Q_{jt}}{C_t}. \tag{10}$$

Employment is proportional to the *household's valuation of output* $\frac{p_{jt} Q_{jt}}{C_t}$, which converts

[16] In principle, the reallocation of activity across sectors may also affect aggregate TFP by changing the distribution of Domar weights across sectors. However, Proposition 1 shows that these reallocation effects are second order; up to first order, only the steady state Domar weights are relevant for computing aggregate GDP.

gross sales into utility units by multiplying by the marginal utility of consumption.[17]

4.2 What Determines Fluctuations in Employment?

In order to understand the effect of a shock on the household's valuation of output, and therefore on employment, we must define two objects summarizing the intermediates network. First, the *input-output matrix* summarizes the intermediates network across sectors:

$$\Gamma = \begin{bmatrix} \gamma_{11}(1-\theta_1) & \dots & \gamma_{1N}(1-\theta_N) \\ \vdots & & \vdots \\ \gamma_{N1}(1-\theta_1) & \dots & \gamma_{NN}(1-\theta_N) \end{bmatrix}.$$

Second, the *Leontief inverse* is

$$\mathcal{L} = (I - \Gamma)^{-1} = I + \Gamma + \Gamma^2 + ...$$

As described by Carvalho and Tahbaz-Salehi (2019), the (i, j)-th element of this matrix, ℓ_{ij}, captures all the direct and indirect paths through the input-output matrix Γ by which sector i supplies intermediate goods to sector j. The Leontief inverse is key in determining the allocation of employment across sectors:

***Proposition* 2.** *The allocation of employment across sectors satisfies*

$$L_{jt} \propto \sum_{k=1}^{N} \ell_{jk} \frac{p_{kt} C_{kt}}{C_t} + \sum_{k=1}^{N} \ell_{jk} \sum_{m=1}^{N} \lambda_{km} \frac{p^I_{mt} I_{mt}}{C_t}. \tag{11}$$

Furthermore, $\frac{p_{kt} C_{kt}}{C_t} = \xi_k$ *for all realizations of* $\{A_{it}\}_{i=1}^N$*. Therefore, the fluctuations in em-*

[17]The expression (10) uses our assumption that the Frisch elasticity of labor supply $\eta \to \infty$. With a finite Frisch elasticity, the expression becomes

$$L_{jt} = \alpha_j \theta_j \frac{p_{jt} Q_{jt}}{C_t} \frac{1}{L_t^{1/\eta}}.$$

All of our results hold using this more general preference specification, but the expressions become more complicated. Therefore, we prefer to use the $\eta \to \infty$ specification to keep this discussion as simple as possible.

ployment L_{jt} satisfy

$$d\log L_{jt} = \sum_{m=1}^{N} \widetilde{\omega}_{jm} d\log\left(\frac{p^I_{mt} I_{mt}}{C_t}\right), \tag{12}$$

where $\widetilde{\omega}_{jm} = \sum_{k=1}^{N} \ell_{jk}\lambda_{km}\left(\frac{p^I_m I_m}{p_j Q_j}\right)^*$.

Proof. See Appendix E. ∎

Equation (11) in Proposition 2 shows that the level of employment in a given sector j depends on how that sector supplies consumption goods to the household and investment goods to other firms, either directly or indirectly through the intermediates network. The contribution to consumption is characterized by the Leontief inverse, ℓ_{jk}, times the household's valuation of consumption produced by all sectors k, $\frac{p_{kt}C_{kt}}{C_t}$. The contribution to investment is characterized by the Leontief inverse, ℓ_{jk}, times the contribution of all sectors k in supplying investment goods to other sectors m through the investment network, λ_{km}, times the household's valuation of investment purchased by those sectors, $\frac{p^I_{mt}I_{mt}}{C_t}$.

Due to the household's Cobb-Douglas preferences, its valuation of consumption across sectors is constant over time and equal to the preference parameter ξ_k; therefore, shocks which only affect the household's valuation of consumption goods do not affect employment – regardless of the structure of the intermediates network. Instead, these shocks generate proportional increases in the marginal product of labor and in aggregate consumption. The resulting income and substitution effects on labor supply exactly offset because our preferences are consistent with balanced growth in the aggregate.[18]

In contrast, the household's valuation of investment goods $\frac{p^I_{mt}I_{mt}}{C_t}$ fluctuates over time because investment is a dynamic problem and the capital accumulation technology is not Cobb-Douglas.[19] Changes in the household's valuation of investment goods generate fluctu-

[18] It is fairly well-known in the one-sector RBC model that employment only responds to TFP shocks because the household would like to produce more investment goods (see the discussion in Benhabib, Rogerson and Wright (1991), for example). Basu et al. (2013) extend that logic to a two-sector model and show that shocks which only affect the production of consumption goods have no effect on employment, while shocks which affect investment production have a strong effect on employment. Our results further extend this logic to a multisector framework and show that the classification of consumption- and investment-producing sectors interacts with the intermediates network through the Leontief inverse.

[19] Appendix F shows that the linearity of the capital accumulation equation is the key departure from Cobb-Douglas which generates employment fluctuations. In particular, we show that, if the capital accumulation equation is also Cobb-Douglas $K_{jt+1} = K_{jt}^{1-\delta_j} I_{jt}^{\delta_j}$, then sector-level employment is constant over time (similar

FIGURE 4: The Leontief-Adjusted Investment Network Ω

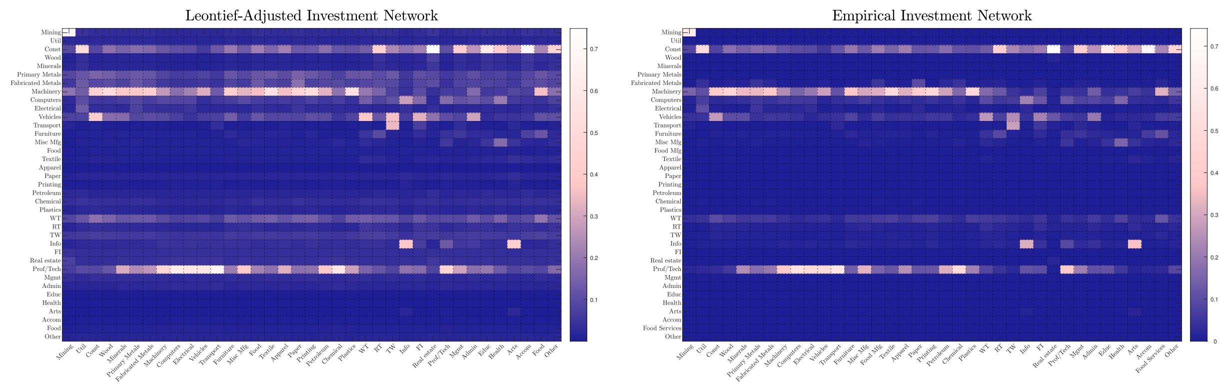

Notes: left panel plots plots the elements of the Leontief-adjusted investment network $\omega_{ij} = \sum_{k=1}^{N} \ell_{ik}\lambda_{kj}$, where ℓ_{ik} are elements of the Leontief inverse and λ_{kj} are elements of the investment network. Right panel plots the elements of our measured investment network λ_{ij} (reproduced from Figure 1 for convenience).

ations in employment because investment weakens the income effect on labor supply. The strength of this force is determined by $\omega_{jm} \equiv \sum_{k=1}^{N} \ell_{jk}\lambda_{km}$, which captures the role of sector j in supplying investment goods to sector m both directly through the investment network and indirectly through the intermediates network.

We call the matrix of these linkages the *Leontief-adjusted investment network* because it is the matrix product of the Leontief inverse with the investment network: $\Omega = \mathcal{L}\Lambda$. The left panel of Figure 4 shows that the Leontief-adjusted investment network is less concentrated than the raw investment network Λ (reproduced in the right panel of the figure). This occurs because the density of the intermediates network implies that many sectors supply intermediate goods to investment hubs; the durable manufacturing sectors near the top of the heatmap — primary metals, fabricated metals, and computers — as well as wholesale trade and transportation & warehousing are particularly important intermediate suppliers of the investment hubs.

to Rossi-Hansberg and Wright (2007)). In this case, the unitary elasticity of substitution between investment and undepreciated capital implies that investment is proportional to output, so shocks generate exactly offsetting income and substitution effects (as they do for the household's valuation of consumption). The linearity of the capital accumulation equation $K_{jt+1} = (1-\delta_j)K_{jt} + I_{jt}$ breaks this result because investment becomes perfectly substitutable with undepreciated capital in the production of new capital.

A related special case of the model is full depreciation $\delta_j = 1$. In this case, one can view capital as an intermediate good with one period time to build. This specification also implies constant employment because it falls within the Cobb-Douglas class.

The Leontief-adjusted investment network Ω shapes employment fluctuations in two important ways. First, given the fluctuations in the household's valuation of investment throughout the economy $\frac{p^I_{mt} I_{mt}}{C_t}$, the Leontief-adjusted investment network Ω characterizes how employment in sector j changes in response (see equation (12)). Since Ω is fairly dense, shocks A_{it} which affect the household's valuation of investment throughout the economy m will also generate employment fluctuations in many sectors j.

Second, the Leontief-adjusted investment network also determines which sector's shocks A_{it} generate large changes in the households valuation of investment $\frac{p^I_{mt} I_{mt}}{C_t}$. Unfortunately, just as with the one-sector RBC model, our model does not admit a closed-form solution to allow an analytical characterization of this mechanism. However, we can show numerically which sectors' shocks generate large changes and then use basic investment theory to explain those results. In addition, Appendix F confirms that these numerical results are primarily driven by the structure of the Leontief-adjusted investment network and not other features of the model.[20]

Specifically, Figure 5 shows that shocks to the key sectors in the Leontief-adjusted investment network — both the investment hubs and their key suppliers – have large effects on aggregate employment while shocks to other sectors do not. The figure computes a numerical elasticity of aggregate employment with respect to a sector-specific shock A_{it} in each sector.[21] The investment hubs have the four largest elasticities, and the next largest elasticities are from the key suppliers to hubs identified in the Leontief-adjusted investment network: durable manufacturing, wholesale trade, and transportation & warehousing. The remaining sectors have very small elasticities.

By equation (12), shocks to investment hubs and their suppliers have large effects on employment because those shocks generate large changes in the household's valuation of investment throughout the economy (we also confirm this fact numerically in Appendix F). In order to understand why that is the case, consider the Euler equation for investment in

[20]Specifically, Appendix F shows that the Leontief-adjusted investment network is the key propagation mechanism determining these numerical responses rather than other parameters. We sample 10,000 values of these other parameters from uniform distributions over a wide range of the parameter space and show that, in the vast majority of these simulations, the response of the household's valuation of investment to a sector-specific shock is highly correlated with that sector's role in supplying investment goods (as measured by its element in the Leontief-adjusted investment network ω_{ij}).

[21]We assume that the persistence of the shocks ρ_j are the calibrated values from Section 5.

Figure 5: Elasticity of Aggregate Employment to Sectoral Shocks

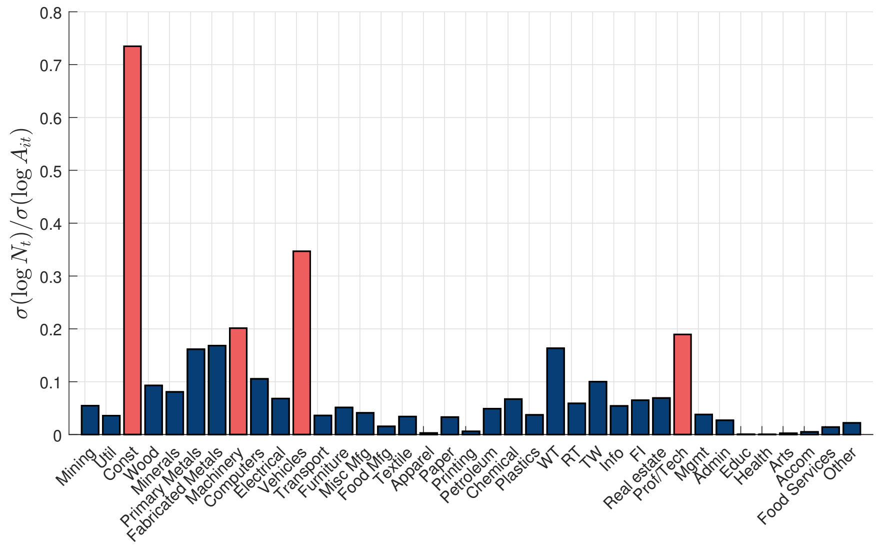

Notes: reduced-form elasticities of aggregate employment N_t to sector-specific shocks A_{it}. For each sector, we simulate the model with $\sigma(\varepsilon_{it}) = 1\%$ shocks to that sector only. The bars plot the volatility of aggregate employment $\sigma(\log N_t)$ divided by the volatility of sector-specific TFP $\sigma(\log A_{it})$. Investment hubs are highlighted in red.

some sector m:

$$\frac{p^I_{mt}}{C_t} = \beta \mathbb{E}_t \left[\alpha_m \theta_m \frac{p_{m,t+1}}{C_{t+1}} \frac{Q_{m,t+1}}{K_{m,t+1}} + (1-\delta_m) \frac{p^I_{m,t+1}}{C_{t+1}} \right]. \tag{13}$$

The marginal benefit of investment on the right-hand side of (13) is the present value of next period's marginal product of capital plus the value of undepreciated capital, times the household's marginal utility of consumption. The marginal cost of investment on the left hand side of (13) is equal to its price index $p^I_{mt} \equiv \Pi_{k=1}^N \left(\frac{p_{kt}}{\lambda_{km}}\right)^{\lambda_{km}}$, again times the marginal utility of consumption.

Proposition 3. *Up to first order, the effect of sector-specific shocks A_{it} on the investment price index for sector m, p^I_{mt}, holding primary input prices fixed, is:*

$$d \log p^I_{mt} = -\sum_{i=1}^N \omega_{im} d \log A_{it}, \tag{14}$$

where ω_{im} are the elements of the Leontief-adjusted investment network.

Proof. See Appendix E. ■

Proposition 3 shows that shocks to investment hubs and their key suppliers A_{it} act as aggregate investment supply shocks in the sense that they decrease the price index for investment goods p^I_{mt} for many sectors m. In fact, holding primary input prices fixed, the investment price index is the sum of all sectors' productivity weighted by the sectors' Leontief-adjusted investment network connections ω_{im}.[22] In response to a shock to one of the investment hubs or the key suppliers, many sectors m increase their optimal quantity of investment I_{mt}, and therefore the household's valuation of their investment goods and ultimately employment.

In contrast, shocks to other sectors act as idiosyncratic investment demand shocks in the sense that they primarily affect the marginal product of capital in their own sector. While these shocks may spill over to other sectors, Figure 5 shows that these spillovers are small in terms of their impact on employment. Therefore, going forward, we focus on the role of shocks to investment hubs and their key suppliers in driving employment fluctuations. We define the *key suppliers* as durable manufacturing, wholesale trade, and transportation & warehousing because these sectors have large weights in the Leontief-adjusted investment network.

Relationship to Networks Literature These results are related to the recent networks literature, which typically uses static models without investment to study how idiosyncratic shocks affect macroeconomic aggregates. Without investment, our model implies that employment is literally constant because shocks do not affect the household's valuation of consumption in Proposition 2. In this case, the Domar weight is also a sufficient statis-

[22] If we allow factor prices to adjust, equation (14) in Proposition 3 becomes

$$d\log p^I_{mt} = \sum_{i=1}^{N} \omega_{im} \left(-d\log A_{it} + \alpha_i\theta_i d\log r_{it} + (1-\alpha_i)\theta_i d\log W_t\right).$$

In this more general case, the contribution of sector i to the investment price index in sector m also depends on the marginal factor costs in sector i. Movements in these factor costs may dampen the pass-through of a productivity shock A_{it} to the price index p^I_{mt}, but those effects are still intermediated through the Leontief-adjusted investment network ω_{im}. Appendix F also shows that the correlation between the elements of the Leontief-adjusted investment network and the numerical elasticity of sector-specific investment prices to sector-specific shocks is very high across a wide range of parameter values, suggesting that these general equilibrium movements are of secondary importance for investment prices. Of course, our quantitative analysis accounts for these general equilibrium movements in factor prices.

tic for the effect of the shock on real GDP: $d\log Y_t = \sum_{j=1}^{N} \left(\frac{p_j Q_j}{P^Y Y}\right)^* d\log A_{jt}$. In addition, the Domar weights are constant over time, so Domar aggregation is no longer a first-order approximation but instead is globally true.

Our results in this section show that investment, and the investment network, breaks this strong version of Hulten's theorem in two ways. First, the Domar weight is not a sufficient statistic for the effect of a shock on real GDP because the shock also affects employment, and the response of employment is determined by the Leontief-adjusted investment network. Second, the Domar weights fluctuate over time due to changes in the household's valuation of investment.[23]

4.3 Implications for Changing Business Cycles Since the 1980s

We now briefly discuss how the key insight of this section — employment fluctuations are primarily driven by shocks to investment hubs and their suppliers — can qualitatively account for a number of changes in business cycle patterns since the early 1980s; we will quantify this mechanism in Section 5. In that section, we show that the key change in the early 1980s was that the correlation of shocks A_{it} across sectors fell dramatically. We interpret this change as reflecting the fact that the pre-1980s sample is dominated by aggregate shocks, which affect all sectors at once, while the post-1980s sample is dominated by idiosyncratic shocks, which affect specific sectors in isolation.

The left panel of Figure 6 plots the impulse responses of real GDP, aggregate employment, aggregate TFP, and aggregate labor productivity to a 1% aggregate shock (increasing TFP by 1% in each sector). The shock increases employment because it increases the productivity of investment hubs and their key suppliers, as discussed above. The shock simultaneously increases productivity at the other sectors, raising their production and therefore real GDP. The effect on these other sectors' productivity is reflected in a roughly 2% increase in aggregate TFP, equal to the sum of Domar weights across all sectors in the economy. Overall, real GDP increases by more than aggregate employment, so labor productivity rises upon

[23] Of course, the investment network also shapes the distribution of Domar weights in steady state. Appendix F shows that the average Domar weights of the investment hubs are comparable to the Domar weights of other sectors.

Figure 6: Impulse Responses to Aggregate vs. Hub + Supplier Shocks

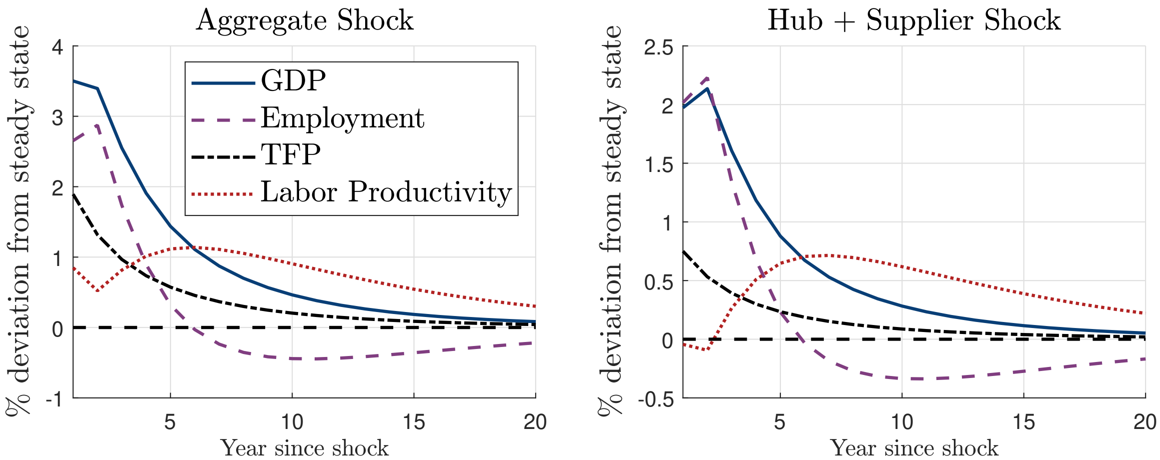

Notes: impulse responses of real GDP, aggregate employment, aggregate TFP, and aggregate labor productivity to combinations of sector-specific shocks A_{it}. Left panel: response to a 1% increase in A_{it} for all sectors i. Right panel: response to a 1% increase in A_{it} for the investment hubs and their key suppliers only.

impact of the shock — consistent with its procyclicality in the pre-1980s sample.

The right panel of Figure 6 plots the same impulse responses in response to a shock which affects only the investment hubs and their suppliers. As before, aggregate employment increases because these sectors are the primary suppliers of investment. But unlike before, this increase in employment is not accompanied by an increase in productivity of the other sectors; aggregate TFP only increases by 0.75% (since these sectors' Domar weights are nearly 40% of the aggregate). In total, aggregate employment increases by more than real GDP upon impact of the shock, driving down aggregate labor productivity. Section 5 shows that this mechanism generates acyclical labor productivity in the post-1980s sample because, as sector-specific shocks become more important in that period, it is primarily these shocks to investment hubs and their key suppliers which drive aggregate fluctuations.[24]

[24]Appendix F computes the cyclicality of labor productivity induced by each sectors' shocks in isolation and shows that shocks to nearly all of the investment hubs and their suppliers generate countercyclical labor productivity. The only exceptions are professional/technical services, wholesale trade, and transportation & warehousing. While shocks to these sectors have a large effect on employment, they are also important suppliers in the intermediates network; hence, they have large Domar weights, generating a larger effect on aggregate TFP and therefore real GDP.

4.4 Additional Results

Appendix F contains two sets of additional results. First, we relate our analysis to the literature which studies the effects of investment-specific technical shocks (e.g. Greenwood, Hercowitz and Krusell (2000) or Justiniano, Primiceri and Tambalotti (2010)). One can view these models as a two-sector version of our model without the intermediates input-output network. Our model provides two main contributions to this literature. First, our model provides a richer classification of sectors in which the correct concept of an "investment producer" is not only its direct production of investment goods but also its role in supplying intermediate goods to the investment hubs. Second, our model also solves the so-called "comovement" problem in this literature, which is that shocks to the investment-producing sector do not generate positive comovement in the consumption-producing sector. Our model generates comovement through the intermediates network, through which many non-investment hubs indirectly produce investment goods (captured by the Leontief-adjusted investment network).

The second set of additional results provides supporting evidence for key mechanisms described above. Similar to Section 2, we show that the key suppliers to investment hubs are more volatile and more correlated with the aggregate cycle than the other sectors, consistent with the role of key suppliers in our model described above.

5 Application: Changes in Business Cycles Since 1980s

We now apply the insights developed in Section 4 to study changes in business cycle patterns since the early 1980s.

5.1 Quantifying the Effects of Changes in Sector-Level Productivity Shocks

We assume that all parameters of the model are fixed, so the only force driving changes in business cycle patterns over time is changes in the process generating sector-level productivity shocks.[25] We measure sector-level productivity as the Solow residual of real gross output net of inputs:[26]

$$\log A_{jt} = \log Q_{jt} - \theta_{jt}\alpha_{jt} \log K_{jt} - \theta_{jt}(1-\alpha_{jt}) \log L_{jt} - (1-\theta_{jt}) \log M_{jt}. \tag{15}$$

Of course, changes in the measured Solow residual may reflect pure 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 this exercise as a natural first step in quantifying the role of the investment network in propagating sector-specific shocks.[27]

We need to detrend sector-level TFP because our model does not feature trend growth. However, a log-linear trend does not fit sector-level data well because sectors typically grow and shrink in nonlinear ways. We therefore 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 D 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.

The left panel of Table 3 characterizes how TFP shocks have changed over time by performing the following statistical decomposition:

$$\mathbb{V}ar(\Delta \log A_t) = \underbrace{\sum_{j=1}^{N} (\omega_{jt})^2 \mathbb{V}ar(\Delta \log A_{jt})}_{\text{variances}} + \underbrace{\sum_{j=1}^{N} \sum_{o \neq j} \omega_{jt}\omega_{ot} \mathbb{C}ov(\Delta \log A_{jt}, \Delta \log A_{ot})}_{\text{covariances}} \tag{16}$$

where $\log A_{jt}$ is log TFP, Δ denotes first differences, and ω_{jt} is the average Domar weight

[25] Section 5.4 discusses robustness of our results when allowing the structural parameters of the model, such as the investment network, to change over time as well. We interpret this finding to indicate that those structural changes are not first-order for the specific cyclical patterns we study (though of course they may be important for other outcomes).

[26] We 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. Our main model results are virtually identical if we instead assume the factor shares are constant when computing TFP.

[27] These 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.

Table 3

Decomposition of Shock Volatility

	Measured TFP		**Value Added**	
	Pre-84	*Post-84*	*Pre-84*	*Post-84*
$1000\mathbb{V}ar(x_t)$	0.41	0.10	1.01	0.39
Variances	0.08	0.06	0.12	0.08
Covariances	0.33	0.03	0.89	0.31

Notes: results of the decomposition (16) in the pre-1984 sample (1948-1983) and post-1984 sample (1984-2018). "Variances" refers to the variance component $1000\sum_{j=1}^{N}(\omega_{jt})^2\mathbb{V}ar(\Delta\log A_{jt})$, weighted by sector j's average Domar weight in the relevant subsample. "Covariances" refers to the covariance component $1000\sum_{j=1}^{N}\sum_{o\neq j}\omega_{jt}\omega_{ot}\mathbb{C}ov(\Delta\log A_{jt},\Delta\log A_{ot})$. "Measured TFP" refers to performing this analysis on log measured TFP growth $\Delta\log A_{jt}$. "Value added" refers to performing this analysis on log real value added growth; in this specification, we weight by value added shares rather than Domar weights. Totals may not appear to be exact sums due to rounding.

of sector j in the subsample (either pre- or post-1984). The volatility of aggregate TFP has fallen by more than 70% since 1984, consistent with the "Great Moderation" of aggregate volatility. Nearly the entire decline in aggregate volatility is accounted for by a decline in the covariance of TFP across sectors; the within-sector variances component has declined by much less.

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

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

where A_t is an aggregate shock common to all sectors and $\widehat{A}_{jt}$ 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 (16) maps directly into a decline in $\mathbb{V}ar(\Delta A_t)$. Appendix D 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 aggregate 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 3 shows that our results hold for value added as well.[28]

We use the following procedure to feed the realized series of TFP shocks into our model. First, we 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 linearized decision rules in our model because those decision rules do not depend on the covariance matrix of shocks. Second, given the values of ρ_j, we compute the innovations to our detrended productivity shocks in the data. We simulate the decision rules given this realized history of shocks, starting from the non-stochastic steady state in 1948.

Investment Production Frictions If we feed these measured shocks directly into our baseline model, the model produces a counterfactually high volatility in the distribution of investment expenditures across sectors. Table 4 measures this volatility as the average change in sector j's total investment expenditures as a fraction of aggregate investment expenditures, $\mathbb{E}[\Delta|\frac{p^I_{jt}I_{jt}}{\sum_{k=1}^N p^I_{kt}I_{kt}}|]$, or as the standard deviation of that change, $\sigma\left(\frac{p^I_{jt}I_{jt}}{\sum_{k=1}^N p^I_{kt}I_{kt}}\right)$. The left and middle panels of Table 4 shows that these changes are five times larger in the model than in the data. This result occurs because an investment-producing sector i sees its potential customer sectors j as perfect substitutes, given the linearity of the market clearing condition (33). Therefore, small changes in investment demand from these purchasing sectors generate counterfactually large changes in the composition of the producing sector's customers, and thus in the distribution of investment expenditures across sectors. In turn, this excess volatility generates an excessively high volatility of aggregate investment to GDP.

We introduce a simple friction to bring this excess volatility in line with the data. Following Huffman and Wynne (1999), we modify the market clearing condition to be

$$Q_{jt} = C_{jt} + \sum_{i=1}^{N} M_{ijt} + \left(\sum_{i=1}^{N} I_{ijt}^{-\rho}\right)^{-\frac{1}{\rho}}, \tag{17}$$

where $\rho \leq -1$ controls the degree of the investment production friction. The baseline model

[28]The "Great Moderation" literature has suggested two broad interpretations of this decline in aggregate volatility. The first is good luck; aggregate shocks have simply become less volatile over time (e.g., oil shocks became less severe and less frequent). The second is good policy; either public policy or private inventory management have allowed the economy to better absorb aggregate shocks. In this paper, we simply take the decline in aggregate volatility as given without taking a stand on why it has occurred.

Table 4
Volatility of Investment Expenditures Composition

	Data	Model w/o Frictions	Model w/ Frictions
$1000 \times \mathbb{E}[\Delta \lvert \frac{p^I_{jt} I_{jt}}{\sum_{k=1}^N p^I_{kt} I_{kt}} \rvert]$	2.0	10.1	2.1
$1000 \times \sigma \left(\frac{p^I_{jt} I_{jt}}{\sum_{k=1}^N p^I_{kt} I_{kt}} \right)$	2.7	15.1	2.7

Notes: measures of changes in the distribution of investment expenditures across sectors. "Data" refers to value of the statistic in the data. "Model w/o Frictions" refers to the model described in Section 3. "Model w/ Frictions" refers to the model augmented with Huffman and Wynne (1999) frictions, as described in the main text.

from Section 3 imposed $\rho = -1$, corresponding to an infinite elasticity of substitution between different purchasing sectors. When $\rho < -1$, investment becomes imperfectly substitutable across purchasing sectors, capturing the idea that the types of investment goods produced by sector i are specific to its customers j, at least in the short run.[29]

We set the parameter $\rho = -1.04$ in order to match the changes in investment composition in Table 4. Our calibrated ρ is similar to the value used in Huffman and Wynne (1999). While all our quantitative results going forward use this extended version of the model, Appendix G shows that our main results are even stronger without these additional frictions because investment expenditures and, by Proposition 2, employment are more responsive to investment supply shocks.

5.2 Changes in Aggregate Business Cycle Patterns in Calibrated Model

In this subsection, we show that the changing sectoral shock process quantitatively generates a number of changes in business cycle patterns since the early 1980s. In Section 5.3, we show that these results are driven by the structure of the investment network, consistent with the theoretical discussion in Section 4.3.

[29]This reallocation friction (17) does not affect the theoretical results derived in Section 4. The only difference is that the price of investment goods now has an endogenous component which reflects the imperfect substitutability of investment customers and dampens large changes in the composition of investment production. See Appendix G for more details.

Table 5
Changes in Business Cycle Patterns Since 1984

	Data		**Model**	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(\Delta y_t)$	3.18%	1.98%	3.95%	2.42%
$\rho(\Delta y_t - \Delta l_t, \Delta y_t)$	0.56	0.28	0.52	-0.01
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.83	1.01	0.90	1.03
$\sigma(\Delta i_t)/\sigma(\Delta y_t)$	2.25	3.10	3.78	4.11
$\sigma(y_t^{hp})$	2.03%	1.24%	2.52%	1.80%
$\rho(y_t^{hp} - l_t^{hp}, y_t^{hp})$	0.52	0.14	0.53	0.01
$\sigma(l_t^{hp})/\sigma(y_t^{hp})$	0.85	1.09	0.92	1.01
$\sigma(i_t^{hp})/\sigma(y_t^{hp})$	2.41	3.50	3.86	4.04

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2018). "Data" refers to our empirical dataset. "Model" refers to model simulation starting from steady state and feeding in realizations of measured TFP over the sample. y_t is log real GDP, l_t is log aggregate employment, and i_t is log real aggregate investment. Δ denotes the first difference operator, and the ^{hp} superscript denotes the HP-filtered series with smoothing parameter $\lambda = 6.25$. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample in computing the HP filtered statistics.

The left panel of Table 5 documents the key changes in aggregate business cycle patterns in the data. The top panel computes the statistics using first differences while the bottom panel uses the HP filter (for both the data and the model).[30] Using either procedure, the volatility of GDP is approximately 40% lower in the post-1984 sample than in the pre-1984 sample — again, consistent with the well-known Great Moderation of aggregate volatility. The cyclicality of labor productivity, measured as the correlation of GDP per worker with GDP, switched from being procyclical in the pre-1980s to being essentially acyclical in the post-1980s. In addition, the volatility of employment rose by approximately 1/3 relative to GDP over this time. Appendix H shows that this rising volatility of employment accounts for the entire decline in the cyclicality of labor productivity; intuitively, since employment

[30]The HP filter has the advantage of isolating business cycle frequencies, while first differences include both high-frequency noise as well as low-frequency changes in average growth rates. The main disadvantage of the HP filter is that its two-sided nature induces cyclical deviations that may not have been known to agents at the time. We partially address this concern by HP-filtering both the model and data series in order to ensure an apples-to-apples comparison. We also omit the first and last three years of data over the sample in order to avoid endpoint bias from the HP filter. We HP-filter the aggregate series directly, rather than aggregating the HP-filtered sector-level series. As in Section 2, we use a smoothing parameter of $\lambda = 6.25$.

and GDP are highly correlated in both subsamples, the time series behavior of their ratio depends on the more volatile component.[31]

Finally, the left panel of Table 5 shows that the volatility of investment relative to GDP has also risen substantially since 1984. This finding is consistent with the idea that the shocks to investment hubs and their suppliers account for a larger share of aggregate fluctuations since 1984. To our knowledge, we are the first to note the increased relative volatility of investment over this period.

The right panel of Table 5 shows that the model generates all of these changes in business cycle patterns. The model matches the decline in the volatility of real GDP because TFP shocks become less correlated over time, similar to the results in Foerster, Sarte and Watson (2011). More novel is the fact that the model's cyclicality of aggregate labor productivity also falls over this period; using first differences, the cyclicality of labor productivity in the model falls by 0.53 compared to 0.28 in the data, while using the HP filter it falls by 0.52 compared to 0.38 in the data. Consistent with this result, the standard deviation of employment relative to GDP rises similarly in the model as in the data (see Footnote 31). Finally, the model's relative volatility of investment also increases over time.[32]

Figure 7 shows that the model also matches the timing of the decline in the cyclicality 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 using either first differences or the HP filter; the correlation between the model and data's series of rolling windows is 0.76 using first differences and 0.91 using the HP filter. The cyclicality of labor productivity is fairly stable until the early 1980s, at which point it

[31] One can see the source of this result using the identity (derived in Appendix H):

$$\mathbb{C}orr(\Delta y_t, \Delta y_t - \Delta l_t) = \frac{1 - \frac{\sigma(\Delta l_t)}{\sigma(\Delta y_t)}\mathbb{C}orr(\Delta y_{t,t})}{\sqrt{1 + \frac{\sigma(\Delta l_t)^2}{\sigma(\Delta y_t)^2} - 2\frac{\sigma(\Delta l_t)}{\sigma(\Delta y_t)}\mathbb{C}orr(\Delta y_t, \Delta l_t)}}. \tag{18}$$

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

[32] The volatility of aggregate investment relative to GDP is somewhat higher in our model than in the data, especially in the pre-1984 period. While in principle we could allow for adjustment costs to the accumulation of capital within sector to match the overall level of volatility, we have found that these adjustment costs generate counterfactually low volatility in the composition of investment spending across sectors and counterfactually high comovement in investment fluctuations across sectors.

Figure 7: 14-Year Forward-Looking Rolling Windows of Labor Productivity Cyclicality

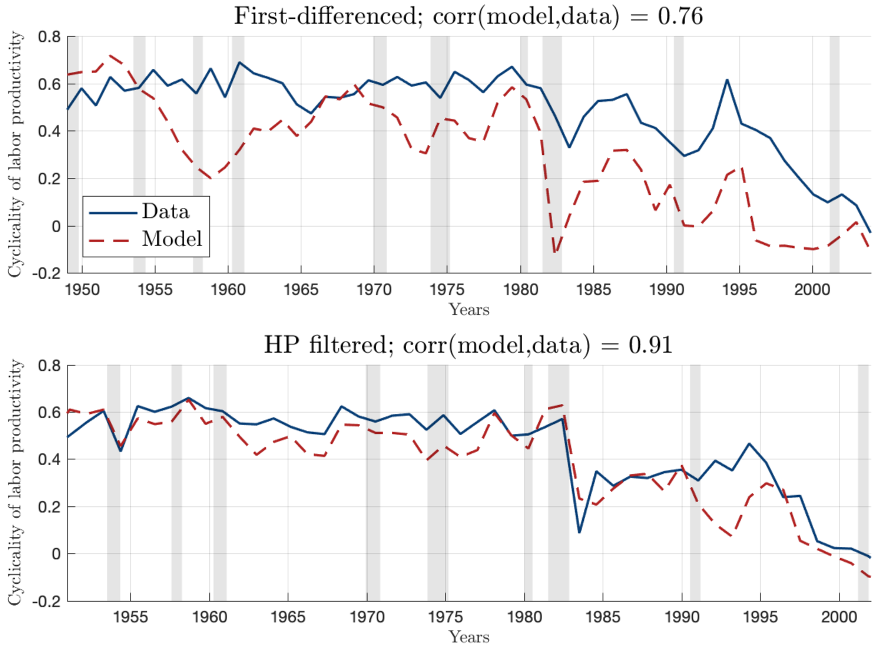

Notes: 14-year forward-looking rolling windows of the cyclicality of labor productivity (e.g. 1950 data point computes the cyclicality between 1950-1963). "Data" corresponds to aggregated version of our dataset. "Model" corresponds to aggregated version of model simulation under measured realizations of sector-level TFP shocks. Top panel computes the statistic using first differences: $\mathbb{C}orr(\Delta y_t - \Delta l_t, \Delta y_t)$ where y_t is log aggregate value added, l_t log aggregate employment, and Δ denotes the first-difference operator. The bottom panel computes the same statistic using the HP filter instead of first differences. We omit the first three and last three years of filtered data from the HP-filtered results to avoid endpoint bias.

drops sharply following the Volcker recession. The cyclicality further declines in the 2008 financial crisis and its aftermath; by the end of the sample, it has fallen by a similar amount in the model and in the data.

Appendix G compares the entire time series of aggregate GDP, consumption, investment, and employment in our model to the data. The average correlation between the model and data for these series is approximately 0.5. In addition, these key macro aggregates strongly comove over the business cycle, which is often difficult to generate in models driven by investment-specific shocks. However, consumption is too smooth and investment is too volatile in our model relative to the data, as they are in the one-sector RBC model. It is not surprising that the model does not perfectly fit each feature of the time series given that no aggregate series were targeted in the calibration and the model does not feature the host of

TABLE 6
ROLE OF INVESTMENT NETWORK IN DRIVING CHANGING BUSINESS CYCLES

Full Model	*Pre-84*	*Post-84*	***Identity Inv. Net***	*Pre-84*	*Post-84*
$\sigma(\Delta y_t)$	3.95%	2.42%	$\sigma(\Delta y_t)$	3.16%	1.72%
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.90	1.03	$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.88	0.90
$\mathbb{C}orr(\Delta y_t - \Delta l_t, \Delta y_t)$	0.52	-0.01	$\mathbb{C}orr(\Delta y_t - \Delta l_t, \Delta y_t)$	0.59	0.48
Uniform Variances	*Pre-84*	*Post-84*			
$\sigma(\Delta y_t)$	1.76%	1.29%			
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.88	1.03			
$\mathbb{C}orr(\Delta y_t - \Delta l_t, \Delta y_t)$	0.55	0.03			

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2018). y_t is log real GDP, l_t is log aggregate employment, i_t is log real aggregate investment, and Δ denotes the first difference operator. "Full model" corresponds to the model described in the main text. "Identity Inv. Net." assumes that sectors invest using only their own output, i.e. $\lambda_{ii} = 1$ for all i and $\lambda_{ij} = 0$ for all $j \neq i$. "Uniform Variances" standardizes the size of shocks to be 1% in both the pre- and post-1984 sample.

other frictions emphasized in the DSGE literature. The main takeaway from our analysis is simply that our model fits the post-1980s data about as well as it does the pre-1980s data, despite the changes in cyclical patterns over that period documented above.

5.3 Role of Investment Network and Sector-Specific Shocks in Driving Changing Business Cycles

We now show that the structure of the empirical investment network is the key propagation mechanism through which the the model generates these changes in business cycle patterns. Specifically, we set the investment network to the identity $\Lambda = I$ so that each sector only uses its own output as an investment good. In this case, there is no concentration of investment production at investment hubs and their key suppliers. And indeed, the top right panel of Table 6 shows that neither the relative volatility of aggregate employment nor the cyclicality of labor productivity significantly change since the 1980 in this case.

This comparative static is consistent with our theoretical discussion in Section 4.3, which showed that sector-specific shocks to investment hubs and their key suppliers have different aggregate effects than purely aggregate shocks. That discussion assumed that the size of aggregate vs. sector-specific shocks were the same in order to focus on the role of the investment network in propagating those shocks, but the sizes of empirical shocks that we feed in from the data may change over time, which in principle, may drive some of our quantitative results. The bottom left panel of Table 6 shows that this is not the case. We standardize the size of shocks in each sector to be 1% in both the pre- and post-1984 subsamples; therefore, the only change since 1984 is that the correlation of shocks fell by the level observed in the data. While the overall level of volatility is obviously different in this version of the model, the relative volatility of employment and cyclicality of labor productivity change by similar amounts as in the baseline exercise.

Another complementary way to show that our quantitative results are consistent with the theoretical analysis in Section 4.3 is to decompose our measured TFP shocks into aggregate and sector-specific components. Table 7 performs this decomposition by identifying the aggregate shock using the first principal component F_t of the sectoral TFP shocks ε_{jt} and identifying the sector-specific shocks as the residual.[33] We then assess the contribution of aggregate shocks or sector-specific shocks by feeding in only the relevant shocks and setting the remaining shocks to zero.

Table 7 shows that aggregate shocks account for the majority of employment fluctuations in the pre-1980s period and generate procyclical labor productivity, as in the left panel of Figure 6. However, sector-specific shocks account for the majority of employment fluctuations in the post-1980s period and aggregate labor productivity is countercyclical in response to these shocks, as in the right panel of Figure 6. Shocks to the investment hubs and their suppliers drive more than 95% fluctuations since shocks to the other sectors have a small effect on aggregate employment.

[33]This procedure yields the decomposition $\varepsilon_{jt} = \alpha F_t + e_{jt}$, where αF_t is the "aggregate shock" and e_{it} is the "sector-specific" shock. The results in Table 7 are not additively separable because the statistics reported are not linear and the sector-specific shocks e_{jt} are not orthogonal to F_t for all sectors j in both pre- and post-1984 time periods. While this approach is not the only one to decomposing aggregate vs. sector-specific shocks in our model, we use it because it is also commonly used in the data (see, for example, Foerster, Sarte and Watson (2011) and Garin, Pries and Sims (2018)).

Table 7
Decomposing the Effects of Aggregate vs. Sectoral Shocks

	All Shocks		**Agg. Shocks Only**		**Sectoral Shocks Only**	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(\Delta y_t)$	3.95%	2.42%	3.45%	1.66%	1.67%	1.40%
$\sigma(\Delta l_t)$	3.55%	2.48%	2.74%	1.41%	1.82%	1.61%
$\rho(\Delta y_t - \Delta l_t, \Delta y_t)$	0.52	-0.01	0.85	0.78	-0.18	-0.28
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.90	1.03	0.79	0.85	1.09	1.14
$\sigma(\Delta i_t)/\sigma(\Delta y_t)$	3.78	4.11	3.31	3.48	4.44	4.52

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2018). y_t is log real GDP, l_t is log aggregate employment, i_t is log real aggregate investment, and Δ denotes the first difference operator. "All Shocks" refers to the baseline model described in the main text. "Agg. Shocks Only" refers to feeding in only the aggregate shocks (as identified following Footnote 33). "Sectoral Shocks Only" refers to feeding in only the sector-specific shocks (again, as identified following Footnote 33).

5.4 Robustness

This subsection summarizes the exercises performed in Appendix G, which show that our main results are robust to relaxing a number of simplifying assumptions in our model.

Structural Change So far, we have held the parameters of the economic environment fixed in order to focus on changes in the shock process. Appendix G allows for those parameters to change over time, specifically: the share of primary inputs in production θ_{jt}, labor's share in production $1 - \alpha_{jt}$, the entries of the intermediates network γ_{ijt}, the entries of the investment network λ_{ijt}, capital depreciation rates δ_{jt}, and the consumption shares ξ_{jt}.

While these parameters have indeed changed in interesting ways, Appendix G shows that our main results are robust to allowing for these changes. We incorporate these parameter changes in two exercises. First, we solve a transition path where agents have perfect foresight over the trend path of parameters, while retaining uncertainty over realizations of productivity shocks using the approach developed in Maliar et al. (2020). Second, we compute the average values of the parameters in the pre-1984 and post-1984 subsamples and simulate the model's moments based on estimated covariance matrices of productivity shocks and the parameter configurations for each time period. Our main results about changes in business cycle patterns continue to hold in both of these exercises. Of course, a full analysis of the process of structural change, its driving forces, and how much is expected by economic agents at the time is outside the scope of our paper; the goal of these exercises is simply to show that allowing for these changes does not affect our main results. We conclude that structural change is not first-order for understanding the specific business cycle patterns we study in this paper and have therefore abstracted from it in the main text for the sake of parsimony.

Non-Cobb Douglas Production and Preferences Appendix G also extends the model to allow for a non-unitary elasticity of substitution between capital, labor, and intermediate goods, as well as a non-unitary elasticity across goods in preferences. We discipline these elasticities using the estimates from Atalay (2017) and Oberfield and Raval (2021) and show that our quantitative results are similar in this extended model. We also solve the extended model using a second order approximation in order to capture the rich nonlinearities described in Baqaee and Farhi (2019). This nonlinear model produces changes in business cycle patterns very similar our baseline analysis in the main text.

Other Robustness Checks Finally, Appendix H shows that our results are robust to a number of other extensions. First, we vary the strength of the Huffman and Wynne (1999) investment production frictions. Second, we allow for convex adjustment costs to capital. Third, we allow for maintenance investment in the investment network, as discussed in Footnote 9. Fourth, we allow for labor reallocation frictions across sectors.

6 Changes in Aggregate Cycles Driven by Changes in Sectoral Comovement

While we believe that our explanation for the changes in business cycle patterns is a natural one, there are many other possible explanations as well (such as those surveyed in the related literature section). Therefore, we now document a new empirical fact which motivates a sectoral explanation: the changes in business cycle patterns have not occurred *within* the average sector of the economy, but are instead due to changes in the comovement of activity *across* sectors. Hence, it would be counterfactual to explain the changes in aggregate business

Table 8
Divergence of Aggregate and Sectoral Cycles

Data	*Aggregate*		*Within-Sector*	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(\Delta y_t)$	3.18%	1.98%	5.42%	4.29%
$\rho(\Delta y_t - \Delta l_t, \Delta y_t)$	0.56	0.28	0.69	0.67
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.83	1.01	0.76	0.81

Model	*Aggregate*		*Within-Sector*	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(\Delta y_t)$	3.95%	2.42%	5.89%	4.93%
$\rho(\Delta y_t - \Delta l_t, \Delta y_t)$	0.52	-0.01	0.79	0.85
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.90	1.03	0.52	0.43

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2018). "Data" refers to our empirical dataset. "Model" refers to model simulation starting from steady state and feeding in realizations of measured TFP over the sample. y_t is log real value added, l_t is log employment and Δ denotes the first difference operator. "Aggregate" refers to outcomes for aggregate variables. "Within-Sector" computes the statistics for each sector and then averages them weighted by the average share of nominal value added within that sub-sample.

cycles with a change in the behavior of the average sector of the economy. We show that our model matches this new empirical finding due to the structure of the investment network and how it propagates the rising importance of sector-specific shocks over time. We focus this section on the cyclicality of labor productivity and relative volatility of employment.

Sector-Level Cycles Stable Over Time Table 8 shows that within-sector business cycles are stable over the postwar sample both in the data and in our calibrated model. These within-sector business cycle statistics first compute the statistics for each sector in the economy and then average those statistics across all sectors (Appendix H shows that these findings are robust to using various weighting schemes to compute the within-sector average and to using the HP filter). While the volatility of sector-level value added falls somewhat post-1984, its magnitude is about half as large as the decline in the volatility of GDP. More importantly, the cyclicality of sector-level labor productivity – the correlation of sector-level value added per worker with sector-level value added – and the relative volatility of sector-level employment are essentially constant across the two sub-samples. In our model, within-sector cycles are relatively stable because within-sector cycles are primarily driven by productivity within that sector, not the comovement of productivity across sectors.

Changes Driven by Sectoral Comovement 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:

$$\frac{\mathbb{V}ar(\Delta l_t)}{\mathbb{V}ar(\Delta y_t)} \approx \underbrace{\omega_t}_{\text{variance weight}} \underbrace{\frac{\sum_{j=1}^{N}(\omega_{jt}^{l})^2 \mathbb{V}ar(\Delta l_{jt})}{\sum_{j=1}^{N}(\omega_{jt}^{y})^2 \mathbb{V}ar(\Delta y_{jt})}}_{\text{variances}} + (1-\omega_t) \underbrace{\frac{\sum_{j=1}^{N}\sum_{o\neq j}\omega_{jt}^{l}\omega_{ot}^{l}\mathbb{C}ov(\Delta l_{jt}, \Delta l_{ot})}{\sum_{j=1}^{N}\sum_{o\neq j}\omega_{jt}^{y}\omega_{ot}^{y}\mathbb{C}ov(\Delta y_{jt}, \Delta y_{ot})}}_{\text{covariances}} \tag{19}$$

where y_{jt} is log real value added of sector j, l_{jt} is employment of sector j, and y_t and l_t are aggregate value added and employment. We focus on the rise in the relative volatility of employment over time because it accounts for the declining cyclicality of labor productivity (see Footnote 31) but is more amenable to a variance decomposition. This decomposition, derived in Appendix H, breaks down 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.

The left panel of Table 9 shows that 85% of the increase in the relative volatility of aggregate employment in the data is accounted for by an increase in the covariances term; the within-sector variances term is relatively stable, consistent with the results in Table 8. Appendix H shows that the changes in covariances reflect two patterns. First, the covariance of value added across sectors fell in the post-1984 sample, decreasing the volatility of GDP. Second, the covariance of employment across sectors remained comparatively stable, stabilizing its aggregate volatility and therefore raising its volatility relative to GDP.

The right panel of Table 9 replicates this decomposition on model-simulated data and shows that the covariance terms account for approximately 90% of the increase in the relative

Table 9
Decomposition of Relative Employment Volatility

	Data			**Model**		
	Pre-84	*Post-84*	*Contribution of entire term*	*Pre-84*	*Post-84*	*Contribution of entire term*
$\frac{\mathbb{V}ar(l_t)}{\mathbb{V}ar(y_t)}$	0.68	1.04	100%	0.81	1.05	100%
Variances	0.41	0.48	15%	0.75	0.57	10%
Covariances	0.72	1.19	85%	0.82	1.15	90%
Variance Weight	0.12	0.21		0.10	0.17	
($\omega_t = \sum_{j=1}^{N}(\omega_{jt}^{y})^2\mathbb{V}ar(y_{jt})/\mathbb{V}ar(y_t)$)						

Notes: results of the decomposition (19) in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). "Data" refers to our empirical dataset. "Model" refers to model simulation starting from steady state and feeding in realizations of measured TFP over the sample. "Variances" refers to the variance component $\frac{\sum_{j=1}^{N}(\omega_{jt}^{l})^2\mathbb{V}ar(\Delta l_{jt})}{\sum_{j=1}^{N}(\omega_{jt}^{y})^2\mathbb{V}ar(\Delta y_{jt})}$. "Covariances" refers to the covariance component $\frac{\sum_{j=1}^{N}\sum_{o\neq j}\omega_{jt}^{l}\omega_{ot}^{l}\mathbb{C}ov(\Delta l_{jt},\Delta l_{ot})}{\sum_{j=1}^{N}\sum_{o\neq j}\omega_{jt}^{y}\omega_{ot}^{y}\mathbb{C}ov(\Delta y_{jt},\Delta y_{ot})}$. "Variance Weight" refers to the weighting term $\omega_t = \sum_{j=1}^{N}(\omega_{jt}^{y})^2\mathbb{V}ar(\Delta y_{jt})/\mathbb{V}ar(\Delta y_t)$. "Contribution of entire term" computes the contribution of the first term of the decomposition (19) (in the variances row) or the contribution of the second term (in the covariances row).

volatility of employment. As in the data, this result reflects two patterns. First, the covariance of value added falls and drives down the volatility of GDP $\mathbb{V}ar(y_t)$ because the covariance of productivity shocks themselves falls. Second, the covariance of employment across sectors is relatively stable, stabilizing the variance of aggregate employment $\mathbb{V}ar(l_t)$.

The investment network is crucial to generating stable employment comovement; if we simulate the model with the identity investment network, the covariance of employment across sectors counterfactually falls by nearly 75%, generating almost no change in the relative variance of aggregate employment (from 0.77 to 0.81). The covariance of employment is relatively stable in our full model because employment fluctuations are primarily driven by shocks to the investment hubs and their key suppliers. Since these sectors only represent a small subset of all sectors in the economy, their shocks do not wash out in the aggregate as they become less correlated over time. Appendix H provides more details about these patterns and also shows that our model fits the covariance changes at the sector-pair level.

Appendix H 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 the HP filter rather than first differences to detrend the data. Third, it shows that the changes in covariances are reflected in changes in correlations rather than changes in variances. Fourth, it shows that the approximation inherent in the decomposition (19) is accurate. Fifth, it shows that the results of this decomposition also hold for a finer 450-sector disaggregation of manufacturing in the NBER-CES database.

7 Conclusion

In this paper, we have argued that the investment network plays an important role in propagating sector-specific shocks to macroeconomic aggregates. 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 concentrated network into a standard multisector business cycle model and showed that shocks to the investment hubs and their key suppliers have large effects on aggregate employment, driving down labor productivity. 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, generating a number of changes in business cycle patterns since the early 1980s.

In order to isolate the role of the investment network in our analysis, we 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 have also 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 into the model as well.

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," *Unpublished Manuscript, University of Notre Dame.*

Baqaee, D. R., and E. Farhi (2019): "The macroeconomic impact of microeconomic shocks: beyond Hulten's Theorem," *Econometrica*, 87(4), 1155–1203.

——— (2020): "Productivity and misallocation in general equilibrium," *The Quarterly Journal of Economics*, 135(1), 105–163.

Barnichon, R. (2010): "Productivity and unemployment over the business cycle," *Journal of Monetary Economics*, 57(8), 1013–1025.

Barro, R. J. (2021): "Double-Counting of Investment*," *The Economic Journal*, ueab005.

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. Fernald, J. Fisher, and M. Kimball (2013): "Sector-specific technical change," *Manuscript, Federal Reserve Bank of San Francisco.*

Basu, S., J. G. Fernald, and M. S. Kimball (2006): "Are technology improvements contractionary?," *American Economic Review*, 96(5), 1418–1448.

Benhabib, J., R. Rogerson, and R. Wright (1991): "Homework in macroeconomics: Household production and aggregate fluctuations," *Journal of Political economy*, 99(6), 1166–1187.

Berger, D., et al. (2012): "Countercyclical restructuring and jobless recoveries," *Manuscript, Yale*.

Bigio, S., and J. La'o (2020): "Distortions in production networks," *The Quarterly Journal of Economics*, 135(4), 2187–2253.

Carvalho, V. M., and A. Tahbaz-Salehi (2019): "Production networks: a primer," *Annual Review of Economics*, 11.

Fernald, J. G., and J. C. Wang (2016): "Why has the cyclicality of productivity changed? What does it mean?," *Annual Review of Economics*, 8, 465–496.

Foerster, A. T., A. Hornstein, P.-D. G. Sarte, and M. W. Watson (2020): "Aggregate Implications of Changing Sectoral Trends," Discussion paper.

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., and S. Klimek (2018): "The Effects of Industry Classification Changes on US Employment Composition," Discussion paper, US Census Bureau, Center for Economic Studies.

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 (2021): "The vanishing procyclicality of labour productivity," *The Economic Journal*, 131(633), 302–326.

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.

Hamilton, J. D. (2018): "Why you should never use the Hodrick-Prescott filter," *Review of Economics and Statistics*, 100(5), 831–843.

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. (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.

Huffman, G. W., and M. A. Wynne (1999): "The role of intratemporal adjustment costs in a multisector economy," *Journal of monetary Economics*, 43(2), 317–350.

Hulten, C. R. (1978): "Growth accounting with intermediate inputs," *The Review of Economic Studies*, 45(3), 511–518.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2010): "Investment shocks and business cycles," *Journal of Monetary Economics*, 57(2), 132–145.

Karabarbounis, L., and B. Neiman (2014): "The global decline of the labor share," *The Quarterly journal of economics*, 129(1), 61–103.

——— (2019): "Accounting for factorless income," *NBER Macroeconomics Annual*, 33(1), 167–228.

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.

King, R. G., and S. T. Rebelo (1999): "Resuscitating real business cycles," *Handbook of macroeconomics*, 1, 927–1007.

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.

Maliar, L., S. Maliar, J. B. Taylor, and I. Tsener (2020): "A tractable framework for analyzing a class of nonstationary Markov models," *Quantitative Economics*, 11(4), 1289–1323.

McDaniel, C. (2007): "Average tax rates on consumption, investment, labor and capital in the OECD 1950-2003," *Manuscript, Arizona State University*, 19602004.

McGrattan, E. R. (2020): "Intangible capital and measured productivity," *Review of Economic Dynamics*, 37, S147–S166.

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., and D. Raval (2021): "Micro Data and Macro Technology," *Econometrica*, 89(2), 703–732.

Rogerson, R. (1988): "Indivisible labor, lotteries and equilibrium," *Journal of monetary Economics*, 21(1), 3–16.

Rossi-Hansberg, E., and M. L. Wright (2007): "Establishment size dynamics in the aggregate economy," *American Economic Review*, 97(5), 1639–1666.

Sato, K. (1976): "The meaning and measurement of the real value added index," *The Review of economics and Statistics*, pp. 434–442.

vom Lehn, C. (2020): "Labor market polarization, the decline of routine work, and technological change: A quantitative analysis," *Journal of Monetary Economics*, 110, 62–80.

A Construction of Dataset and the Investment Network

This Appendix describes the details of our data set and our construction of the investment network.

A.1 Data Sources

Our analysis of business cycle fluctuations uses a dataset of gross output, intermediate inputs, value added, employment, and investment for non-government, non-farm sectors over the 1948 - 2018 sample. We define sectors using NAICS codes, resulting in the 37-sector partition in Table 1. Data on nominal and real measures of gross output, intermediate inputs, and value added are taken from the GDP by Industry database, while data on nominal and real investment expenditures are from the BEA Fixed Asset database.

The main challenge in compiling this dataset is constructing consistent measures of sector-level employment over the entire 1948 - 2018 sample. Starting in 1998, we observe sector-level employment in NIPA Table 6.4D, which reports the total number of full-time and part-time employees by sectors defined according to NAICS codes. Before 1998, the BEA Industry Accounts provide historical employment data converted to NAICS codes for 1948-1997. However, this data is only available for 17 out of the 37 sectors that we consider prior to 1977; the remaining sectors are in manufacturing, which the BEA collapses into broad durable and non-durable sectors over this time period. Fortunately, the BEA provides disaggregated manufacturing employment in SIC codes over this period in NIPA Tables 6.4B and 6.4C. We convert these data to the NAICS classification using the Fort and Klimek (2018) crosswalk. We ensure there is not a discontinuity at the 1977 breakpoint by cumulating the growth rates from the converted data in each sector to compute the levels of employment in the pre-1977 period rather than relying on the levels in the raw data.

Table A.1
Investment Flows Table Visualization

		Investment Purchasers				
		Mining	Utilities	Construction	$\cdots$	**Total Production**
Investment Producers	Mining					
	Utilities					
	Construction					
	$\vdots$					
	Total Expenditures					

A.2 The Investment Network

Our investment network records the share of new investment expenditures of sector j that were purchased from sector i for each pair of sectors (i, j) and for each year t in our sample. While the BEA capital flows tables provide some relevant information in some years, those tables are limited in three key ways for our analysis. First, they are only available for seven of the 72 years from 1947-2018: 1963, 1967, 1972, 1977, 1982, 1992 and 1997.[34] Second, they are not consistently defined over time because they use different vintages of SIC or NAICS codes. Finally, and most importantly for our analysis, the BEA tables do not include all of intellectual property; in fact, the 1997 table is the only one with any intellectual property, but it only has software (not R&D investment or artistic originals). To our knowledge, our investment network is the only version of the capital flows tables that is consistent with modern national accounting practices regarding intellectual property.[35]

We construct our investment networks in order to overcome these limitations, but otherwise try to follow the BEA methodology as closely as possible. To help explain our approach, Table A.1 visualizes the *investment flows table*, whose $(i, j)^{th}$ entry records the total investment expenditures by sector j purchased from sector i in a given year. Summing across columns for each row in this table generates total production of investment by each sector, while summing across rows for each column generates total investment expenditures for each sector. The investment network simply divides each column j of this table by total

[34]Only the 1982, 1992, and 1997 tables are currently published on the BEA website, but older tables can be obtained from archived issues of the Survey of Current Business.

[35]We have found that the presence of intellectual property is the key difference between the BEA 1997 capital flows table and our measured investment network in that year. Specifically, intellectual property makes up more than a quarter of total investment spending in the sectors for which our network is significantly different from the BEA's capital flows table.

expenditures in that sector in order to compute expenditure shares.

We construct the investment network in three steps: (i) separately construct the investment flows tables for residential investment, non-residential structures, non-residential equipment, and intellectual property, (ii) aggregate those four investment flows tables to total investment, and (iii) rescale them to compute the aggregate investment network. Steps (ii) and (iii) are straightforward matrix operations, so we focus this appendix on explaining how we perform step (i).

Unfortunately, there are no publicly available data on the pairwise investment flows between producers and purchasers necessary to fill in each element of Table A.1. Instead, we estimate these pairwise flows using the following data which the BEA does provide:

(i) Total investment expenditures by sector for each year from Table 3.7 of the Fixed Assets data (the "total expenditures" row in Table A.1).

(ii) Total investment production by sector for each year from the annual use tables from the Input-Output database (the "total production" column in Table A.1). Before 1997, these tables separately record the total production of structures (both residential and non-residential), equipment (both residential and non-residential), and intellectual property. After 1997, the tables record the total production of residential investment, non-residential structures, non-residential equipment, and intellectual property.

(iii) Aggregate residential structures and residential equipment expenditures for each year from NIPA Tables 5.4.5 and 5.5.5. Because we assume that the real estate sector is the sole purchaser of residential investment – following the BEA's methodology – there is no need for detailed residential investment expenditure data by sector.

(iv) Sector-level investment expenditure on 33 different types of assets for each year: residential structures, residential equipment, two types of non-residential structures (mining and all other), four different intellectual property assets, and 25 different equipment assets. We construct this data from the expenditures on residential structures and equipment described in point (iii) above and detailed data for Fixed Assets by Industry, which provides expenditures on the other types of assets (available at `https://apps.bea.gov/national/FA2004/Details/Index.htm`).

(v) Sector-level detail on the production of individual equipment assets for the years 1997-2018 (available at https://www.bea.gov/products/industry-economic-accounts/underlying-estimates), and 1987 and 1992 in SIC codes (available at https://www.bea.gov/industry/historical-benchmark-input-output-tables), which we convert to NAICS codes using the crosswalk in Fort and Klimek (2018).

Our approach primarily utilizes asset-level expenditure data to estimate the individual entries in Table A.1. We estimate those pairwise investment flows as

$$I_{ijt} = \sum_{a=1}^{A} \omega_{iat} I_{ajt}^{exp}, \tag{20}$$

where I_{ijt} is the $(i,j)^{th}$ element of the investment flows table in year t, I_{ajt}^{exp} is expenditures by sector j on capital asset a in year t, and ω_{iat} represents the fraction of capital asset a produced by sector i in year t.[36] The key assumption in equation (24) is that the the mix of sectors producing a given asset a is the same for all sectors j which purchase that asset, i.e. that ω_{iat} is independent of the purchasing sector j. The BEA also makes this assumption in constructing their capital flows tables.

The main challenge in our measurement procedure is to estimate the collection of ω_{iat} – which is called a *bridge file* – across assets a, sectors i, and years t. The remainder of

[36] To be consistent with input-output methodology, our investment network represents expenditures on new investment, not used or scrap transactions. However, the Fixed Assets investment expenditures data used to construct I_{ajt}^{exp} in (24) does include net purchases of used assets, which often enter recorded investment expenditure as a negative value. Thus, the reported expenditures may understate total expenditures on new assets. In terms of measured investment expenditures, the addition of net used transactions is only a concern for equipment assets; for structures and intellectual property, net used transactions are negligible. We adjust total investment spending from the Fixed Assets data to eliminate used and scrap transactions as follows:

- For all equipment assets aside from autos, we scale up investment expenditures uniformly across sectors in order to match the total production of new assets. In 1997-2018, the scaling factor ensures that total expenditures equals total production of that asset as reported in sector-level detail on the production of individual equipment assets. Before 1997, we use the median scaling factor from 1997-2018. Overall, this correction is non-negligible only for trucks and aircraft.
- For autos, we scale up expenditures on autos in the rental/leasing sector in order to be consistent with the observation in Meade, Rzeznik and Robinson-Smith (2003) that net sales of used autos are primarily from that sector to private households (the rental/leasing sector is part of real estate in our 37 sector partition). We again choose the scaling factor to ensure that total expenditures on autos equals total expenditures in 1997-2018 (when production data is available), and choose the median scale factor from that period for the pre-1997 data (when the production data is not available).

this subsection describes how we construct these annual bridge files ω_{iat} separately for non-residential structures, intellectual property, residential investment, and equipment.

Non-Residential Structures

We assume that all non-residential structures are produced by the construction sector except for mining structures, which we assume are produced by the mining sector. Therefore, for a = non-residential non-mining structures, we set $\omega_{iat} = 1$ if i = construction and zero otherwise. For a = non-residential mining structures, we set $\omega_{iat} = 1$ if i = mining and zero otherwise. This allocation rule is consistent with how the BEA constructs the capital flows tables.[37]

Intellectual Property

We have data on four types of intellectual property assets: prepackaged software, own and custom software, research and development, and artistic originals. We allocate the production of these assets to sector i based on the BEA practices described in McGrattan (2020).[38]

(i) We assume that own and custom software is produced by the professional/technical services sector, i.e. for a = own and custom software, $\omega_{iat} = 1$ if i = professional/technical services and zero otherwise.

(ii) We assume that R & D investment is also produced by the professional/technical services sector, i.e. for a = R & D investment, $\omega_{iat} = 1$ if i = professional/technical services and zero otherwise.

(iii) We assume that artistic originals are produced by the information sector (which includes radio & TV communication and motion picture publishing) and the arts &

[37] In 1997-2018, the construction and mining sectors produce 99.9% of non-residential structures investment net of brokers' commissions on structures (which are excluded from our investment network following the BEA's methodology for the 1997 capital flows table).

[38] McGrattan (2020) estimates a version of the investment network for the year 2007 that includes intellectual property products. However, McGrattan (2020) argues for an allocation of the production of intellectual property products different from current practice by the BEA. While we are sympathetic to those arguments, we chose to follow BEA practice in the spirit of following their approach as closely as possible throughout our analysis.

entertainment services sector. We assume that artistic originals is the only type of intellectual property produced by the arts & entertainment sector, and therefore estimate its production of artistic originals as its total production of intellectual property from the Input-Output tables. Hence, for a = artistic originals and i = arts & entertainment, we set $\omega_{iat} = \frac{I^{prod}_{ait}}{\sum_{j=1}^{N} I^{exp}_{ajt}}$ where I^{prod}_{ait} is the total production of intellectual property by i = arts & entertainment and $\sum_{j=1}^{N} I^{exp}_{ajt}$ are total economy-wide expenditures on artistic originals. We then set $\omega_{iat} = 1 - \omega_{i'at}$ for i = information (where i' = arts & entertainment) and $\omega_{iat} = 0$ for all other i.

(iv) Finally, we assume that all pre-packaged software is produced by the information sector. However, we must also take into account the fact that the wholesale trade, retail trade, and transportation & warehousing sectors play a role in delivering new pre-packaged software to customers (these delivery expenses are called *margin payments*).[39] We compute the margin payments on pre-packaged software as the total production of intellectual property for those sectors as recorded in the Input-Output Tables. Hence, for a = pre-packaged software and i = information, we set $\omega_{iat} = \frac{I^{prod}_{ait}}{\sum_{j=1}^{N} I^{exp}_{ajt}}$ where $i \in$ {wholesale trade, retail trade, transportation & warehousing } and $\sum_{j=1}^{N} I^{exp}_{ajt}$ is total economy-wide expenditure on pre-packaged software. We then set $\omega_{iat} = 1 - \sum_k \omega_{akt}$ for i = information and $k \in$ {wholesale trade, retail trade, transportation & warehousing}. Finally, we set $\omega_{iat} = 0$ for all other sectors i.

Residential Investment

Residential investment is the sum of residential structures and residential equipment (such as appliances or other consumer durables owned by landlords and included in residential leases). As described above, the BEA directly reports the total production of residential investment by sector in the Input-Output Tables between 1997-2018. However, that production data also includes margin payments on used residential structures transactions, which we need to

[39] We assume that pre-packaged software is the only intellectual property product with margin payments because that is the case in the benchmark 2007 and 2012 input-output tables, which has detailed observations on margin payments (available at https://www.bea.gov/products/industry-economic-accounts/underlying-estimates).

eliminate from our investment network. Our approach to estimating these margin payments for the 1997-2018 period depends on the sector:

(i) We assume that some sectors – real estate, finance/insurance, and legal services (part of professional/technical services) – only produce margin payments on residential structures and not on residential equipment.[40] For these sectors, we assume that 13.2% of their production of residential structures corresponds to margin payments on new transactions, based on the estimated fraction of real estate broker margins that were for new residential structure investment (as used in the 1997 BEA capital flows data and reported in Meade, Rzeznik and Robinson-Smith (2003)).

(ii) Other sectors produce margin payments for both residential structures and residential equipment (wholesale trade, retail trade, and transportation & warehousing). For these sectors, we estimate their total margin payments as the sum of their total production of residential equipment (corresponding to margin payments on residential equipment, observed in detailed equipment production data) and 13.2% of their production of residential structures (corresponding to margin payments on new residential structures, assumed to be the remainder of these sectors' total production of residential investment).

Unfortunately, the BEA does not separately report production of residential investment by sector prior to 1997. Our procedure to estimate residential investment in this period depends on the sector:

(i) We assume some sectors (wood products manufacturing, finance/insurance and professional/technical services) produce residential structures but do not produce non-residential structures or residential equipment in the post-1997 data.[41] We therefore estimate these sectors' total production of residential investment as their reported total production of structures pre-1997. We then eliminate margin payments on used transactions following the same procedure described in the previous paragraph.

[40] This assumption is validated by the fact that these sectors do not contribute to the production of total equipment (residential plus non-residential equipment) in the pre-1997 production data.

[41] This assumption is validated by the fact that these sectors report zero production of any equipment investment pre-1996 and no production of non-residential structures post-1997 in the Input-Ouput tables.

(ii) We estimate the other sectors' production of residential investment using the following procedure.

   - Residential structures: in the 1997-2018 period, we can infer sector-level production of residential structures directly in the Input-Output Tables (given that we separately observe production of residential equipment in the detailed data describing production of equipment over this period). In the pre-1997 period, when we do not observe production, we estimate it as the average share of residential structures produced by that sector in the 1997-2018 data times the aggregate spending on residential structures in a given year $t < 1997$.
   - Residential equipment: we follow a similar procedure as for residential structures; given observed sector-level production of residential equipment for later years, we use the average share produced by each sector times aggregate spending on residential equipment. Because we have detailed data on the sectoral production of residential equipment for the years 1987 and 1992 as well, for years $t < 1987$, we use 1987 data on the shares of sectoral production times the aggregate time series for residential equipment spending. For years between 1987 and 1992, we use a moving average of the residential equipment production shares from the 1987 and 1992 bridge files, and for years between 1992 and 1997, we use a moving average of the data in 1992 and 1997.

Since these procedures define the bridge files ω_{iat} recursively, we do not write out their formulas here.

Non-Residential Equipment

Constructing the bridge files for equipment assets is the most involved task because there are 25 detailed types of equipment assets reported in the Fixed Asset data. We describe our procedure separately for three time periods which have different data availability from the BEA:

(i) 1997 - 2018: The BEA already publishes detailed data on the production of individual equipment assets by each sector, i.e. we observe ω_{iat} from the data directly.

(ii) 1987 and 1992: The BEA publishes bridge files for each of these assets in 1987 and 1992, but the sectors correspond to SIC codes rather than NAICS codes. Therefore, we convert these bridge files from SIC to NAICS using the Fort and Klimek (2018) crosswalk.[42]

(iii) Remaining years: in the years for which we have no publicly available bridge files, we interpolate the existing bridge files and re-scale the interpolation so that it matches the total production of equipment investment by sector from the input-output data. To understand our procedure, first note that the total production of equipment capital by sector i is:

$$I_{it}^{prod} = \sum_a \omega_{iat} I_{at}^{exp}, \tag{21}$$

where I_{at}^{exp} is total expenditures on equipment asset a in year t (from the Fixed Assets data), I_{it} is the production of all equipment assets by sector i (from the input-output data), and ω_{iat} is our bridge file to be estimated.

We initialize our estimate of the bridge file, $\widehat{\omega}_{iat}$, as either the bridge data from the last available year or a moving average of the two nearest bridge files. This estimate $\widehat{\omega}_{iat}$ may not satisfy the relationship (21) given our observations of I_{it}^{prod} and I_{at}^{exp}. Let $\alpha_{it} = \frac{I_{it}^{prod}}{\sum_a \widehat{\omega}_{iat} I_{at}^{exp}}$ denote the ratio of true equipment production of sector i to its production implied by the bridge file estimate. We use α_{it} to arrive at our final estimate of the bridge file:

$$\omega_{iat} = \frac{\alpha_{it}\widehat{\omega}_{iat}}{\sum_{j=1}^{N} \alpha_{jt}\widehat{\omega}_{jat}}. \tag{22}$$

Equation (22) ensures that the total investment production in each sector i implied by the bridge file is equal to total investment production in the data. The key assumption

[42] If the converted bridge file implies that a sector produces an equipment asset that the sector is not observed to produce in the detailed equipment production data in the years 1997-2018, we modify the conversion of NAICS to SIC sectors such that this sector does not produce that good in the final converted bridge file. However, these older bridge files only contain limited detail on margin sectors, making careful conversion to NAICS sectors infeasible. In order to ensure that we do not have a discontinuous break in margin payments by each sector at 1997, we take the total reported margins for each asset in these older bridge files and multiply them by the share of margins produced by each margin sector for that asset from the detailed equipment bridge files in 1997-2001.

is that the production of assets a by sector i always occur in proportion to $\widehat{\omega}_{iat}$.

Additional Networks

Alternative sectoral disaggregation While we use a 37-sector disaggregation in the main text, we can also incorporate the agriculture, state/local government, and federal government sectors. The agriculture sector can be incorporated following the same steps as above without modification. The government sectors can be incorporated by using the Input-Output tables directly because investment by federal and state/local governments is a final use. We do not incorporate these sectors in the main text in order to focus on the private nonfarm economy.

We can also disaggregate the mining and real estate sectors more finely than in the main text. In particular, we can split the mining sector into oil & gas extraction, support activities for mining, and other mining, and we can split the real estate sector into real estate and rental/leasing services. We do not use these additional sectors in our baseline analysis because the way investment purchases and expenditures are allocated across these sectors is unusual (e.g., most of investment purchases by mining are produced by the support activities for mining sector, and the purely real estate sector is largely owner-occupied housing imputations). However, our results are robust to using this more detailed partition of sectors for the private non-farm economy.

Capital Rental Services We also construct a *capital rental services network*, defined as the fraction of capital rental service expenditures by sector j, $R_{jt}K_{jt}$, purchased from all other sectors i in the economy in year t. This rental services network may be useful for at least two reasons. First, it is consistent with the national accounting procedure described in Barro (2021), which constructs a measure of national income whose present value equals the present value of consumption over time. Second, the rental services network may be used to incorporate sectoral linkages in capital services in a static model.

As with the investment network, the rental services network combines rental expenditures of sector j on asset a, $R^a_t K^a_{jt}$, with bridge files to infer from which sectors those assets were purchased. We compute $R^a_t K^a_{jt}$ using data on the nominal capital holdings of each sector j

for each asset a in each year t, $P_t^a K_{jt}^a$ from the BEA Fixed Assets data.[43] We then combine that series with a time series for the real rental rate $\frac{R_t^a}{P_t^a}$, which we construct following the approach in Karabarbounis and Neiman (2019):

$$\frac{R_t^a}{P_t^a} = \frac{1+\tau_t^x}{1-\tau_t^k}\left[\left(\frac{(1+\tau_{t-1}^x)P_{t-1}^a}{(1+\tau_t^x)P_t^a}\right)(1+(1-\tau_t^k)r_t)-(1-\delta_t^a)-\frac{\tau_t^k\delta_t^a}{1+\tau_t^x}\right] \tag{23}$$

where τ_t^x is the tax rate on investment, τ_t^k is the tax rate on capital income, r_t is a measure of the real rate of return on capital, P_t^a is the price of a new unit of capital (investment) of asset type a, and δ_t^a is the depreciation rate of asset a. We follow the same broad steps as Karabarbounis and Neiman (2019) (on their more aggregated data) in order to measure these objects:

(i) Real interest rate r_t: measured as the nominal rate of return on 10 year Treasuries net of expected inflation (measured as a five year moving average of observed PCE inflation) plus a 3% risk premium (which avoids negative values of the rental rate).

(ii) Price P_t^a: directly observed in NIPA tables 5.4.4, 5.5.4, and 5.6.4.

(iii) Taxes τ_t^x and τ_t^k: directly from McDaniel (2007), which have been updated through 2017.

We then take a seven-year moving average of the real rental rate $\frac{R_t^a}{P_t^a}$ in order to smooth out high-frequency variations; Karabarbounis and Neiman (2019) use a five-year moving average, which leaves more noise in our disaggregated data.

We use the same bridge files ω_{iat} constructed above to allocate the production of new rental services for asset a to various sectors i. Our key assumption is that the composition of sectors which produced the existing capital asset a in the past is the same as the composition of sectors which produce new capital. This assumption may fail if the sectors producing that particular capital asset have substantially changed over time, but that is unlikely to be an important issue for two reasons. First, we use a fairly detailed partition of capital assets whose production patterns have not changed much over time. For example, for our estimated

[43] We assume that the rental rate of a given capital asset is specific to the asset a, not to the sector j renting it. This assumption is consistent with BEA data showing that the price of capital assets is almost identical for the same asset across sectors.

equipment bridge files, the average correlation of the distribution of sectors that produce each asset in 1947 and the distribution of sectors that produce that asset in 2018 is 0.91. Second, the assets for which production has changed the most also have the highest depreciation rates, implying that our bridge files for new investment correspond to a large fraction of the existing capital stock as well. That said, we will also provide the time series of each of our bridge files so that other researchers can relax this assumption and explicitly cumulate the pairwise purchases of investment over time using the perpetual inventory method.[44]

Given this modular approach, other researchers can construct rental services by asset in different ways — for example, reflecting a different formulation of the rental rate — and combine them with our bridge files to build their own rental services network. We provide networks with and without taxes (given that our model does not include taxes) as well as a network using rental rates net of depreciation.[45] Figure A.1 plots the heatmap of our gross rental services table without taxes and shows that it is very similar to our investment network considered in the main text; the network with taxes and the network using net rental rates is also similar.

A.3 Additional Analysis of Investment Network

This subsection presents two pieces of additional analysis of the investment network referenced in Section 2 of the main text.

Changes in Network Over Time Figure A.2 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.

[44]We do not make a correction for used goods when building the rental services tables because this correction is significantly more complex when considering the stock of all capital and not the period flows of investment.

[45]In the case of the net rental services matrix, net rental rates are measured without taxes, with $\frac{R_t^a}{P_t^a} = \left[\left(\frac{P_{t-1}^a}{P_t^a}\right)(1+r_t) - 1\right]$. We add an additional two percentage point risk premium and smooth changes in asset prices, $\frac{P_{t-1}^a}{P_t^a}$, using a five year moving average in order to avoid negative net rental rates.

Figure A.1: Heatmap of Empirical Rental Services Network

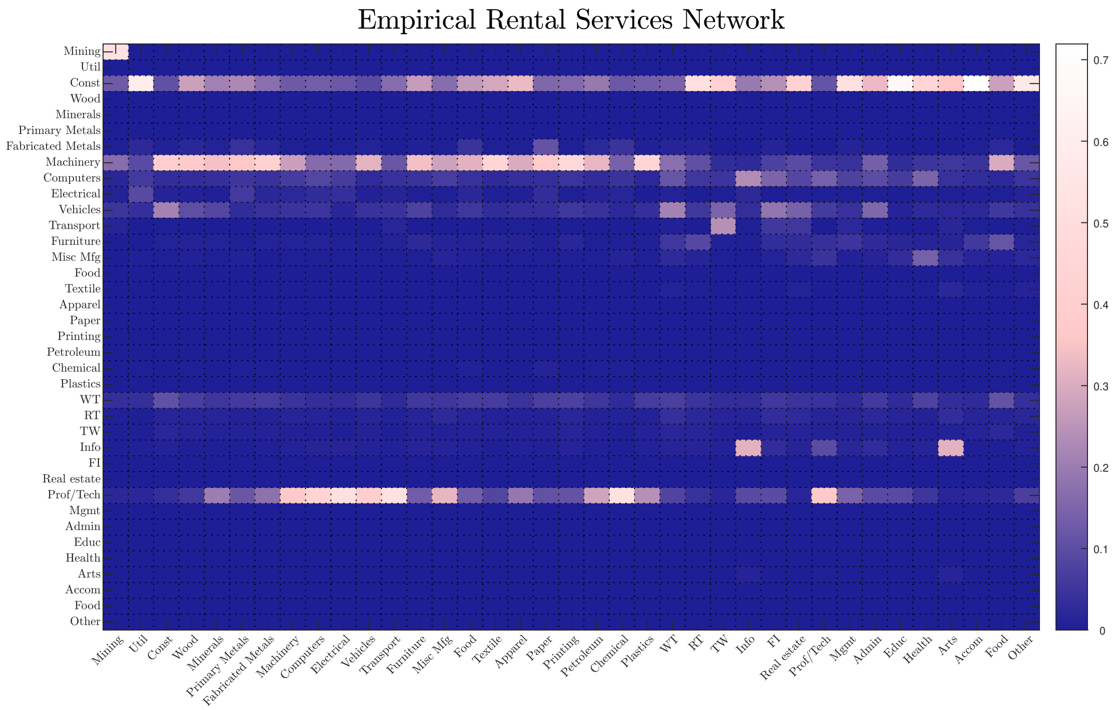

Notes: heatmap of empirical gross rental services network (without taxes). Entry (i, j) computes share of total rental expenditures by sector j that are produced by sector i, averaged over the 1947 - 2018 sample.

Figure A.3 compares the heatmaps of the equipment and intellectual property products networks in the pre- vs. post-1984 sample (the residential investment and non-residential structures networks do not substantially change over time). The left panel shows that computing manufacturing and professional/technical services becomes a larger supplier of non-residential equipment over time while the information sector becomes a larger supplier of intellectual property products, both of which reflect the rising importance of IT.

To study the time series of the key changes in the investment network, Figure A.4 summarizes the importance of each sector as a supplier of investment goods by the weighted outdegree of that sector. The left panel shows that our four investment hubs are systematically important suppliers of investment goods over the entire sample. The rise of intellectual property products over time implies that professional/technical services, which produces the majority of these products, has become relatively more important over time while the other three hubs have become relatively less important. The right panel shows that the non-hub

FIGURE A.2: Heatmaps of Investment Network, Pre/Post 1984

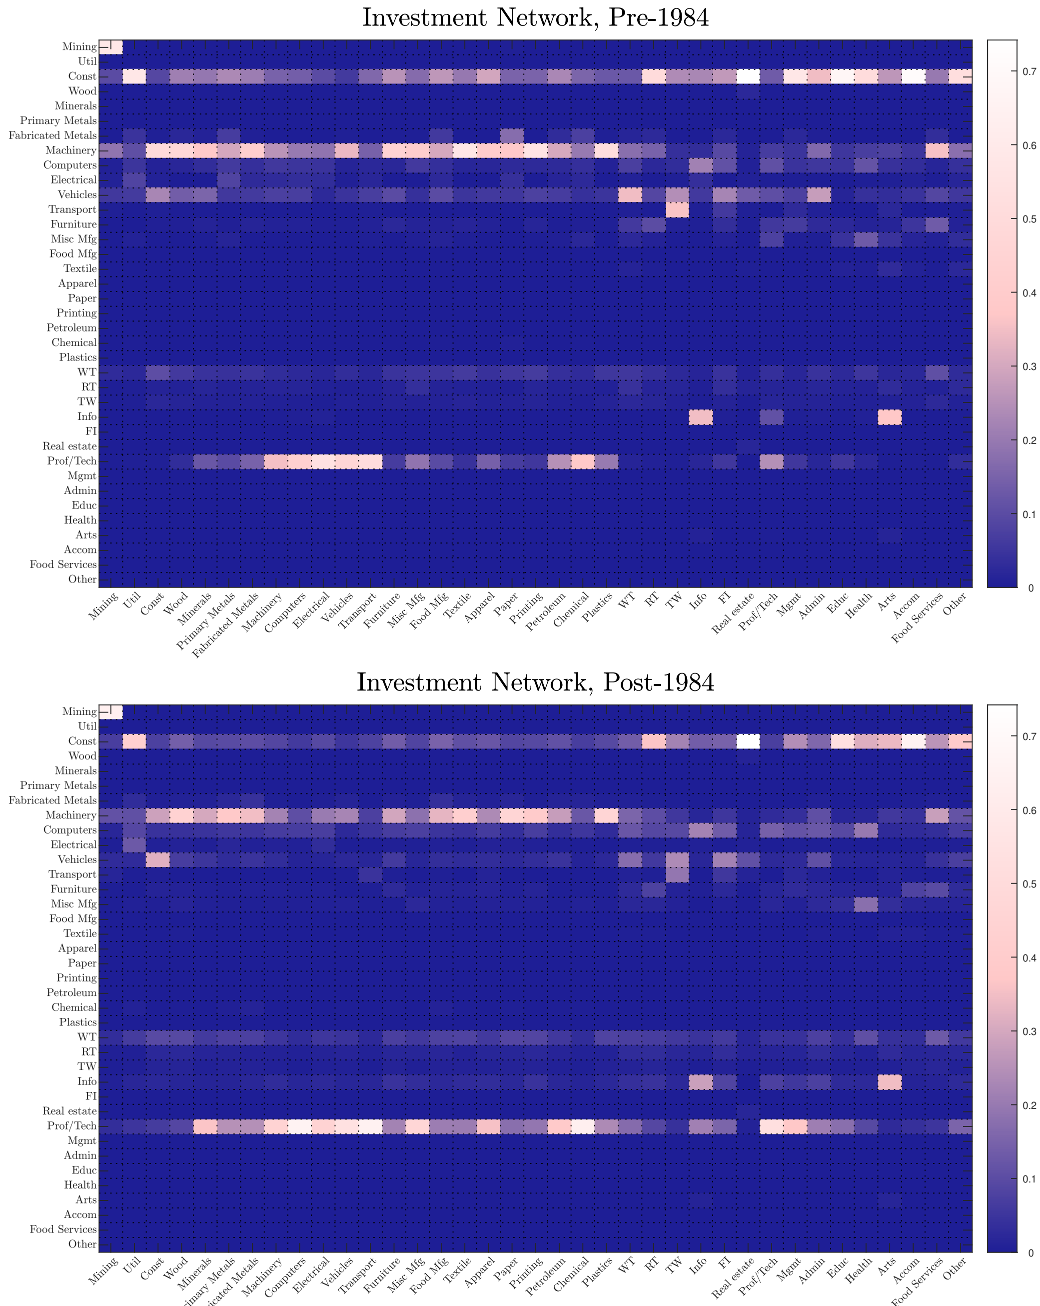

Notes: Heatmaps of the investment network λ_{ij} are constructed as described in the main text. The (i, j) entry of each network corresponds to parameter λ_{ij}, i.e., the amount of sector i's good used in sector j. The pre-84 network corresponds to the years 1947-1983 and the post-84 network corresponds to the years 1984-2018.

FIGURE A.3: Heatmaps of Equipment and Intellectual Property Investment Networks, Pre/Post 1984

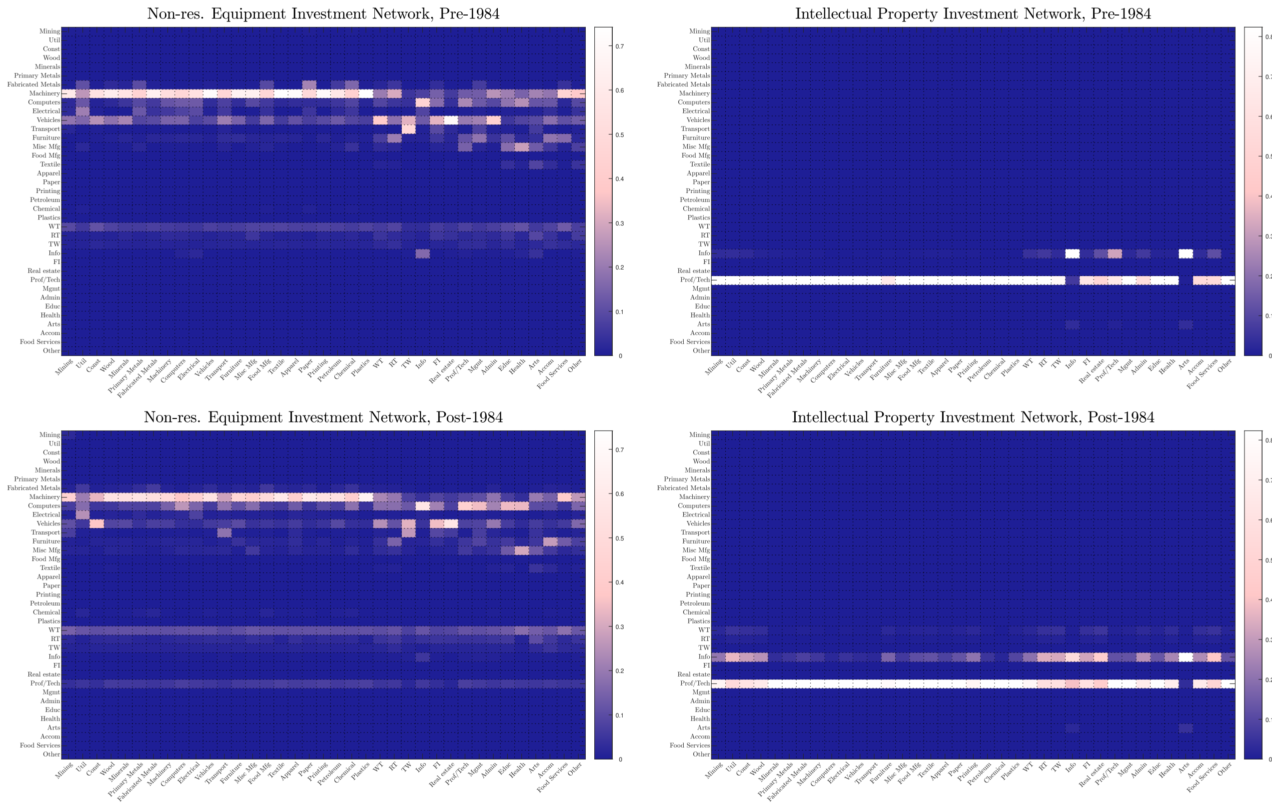

Notes: Heatmaps of the non-residential equipment and intellectual property investment networks are constructed from the bridge files as described in Appendix A.2. The pre-84 network corresponds to the years 1947-1983 and the post-84 network corresponds to the years 1984-2018.

sectors are relatively unimportant suppliers over the entire sample.[46]

A natural question is to what extent the time-series variation in the network is driven by changes in our estimated bridge files. To answer this question, recall that we construct the pairwise flows of investment expenditures between sectors i and j in year t as

$$I_{ijt} = \sum_{a=1}^{A} \omega_{iat} I_{ajt}^{exp}, \tag{24}$$

where I_{ajt}^{exp} is expenditures by sector j on capital asset a in year t (provided by the BEA) and ω_{iat} is the fraction of capital asset a produced by sector i in year t (our estimated

[46]The biggest exceptions are computers and electronic machinery manufacturing, information, and wholesale trade, all of whose outdegrees increase substantially over the sample due to the rise of IT.

Figure A.4: Weighted Outdegree of Hubs and Non-Hubs

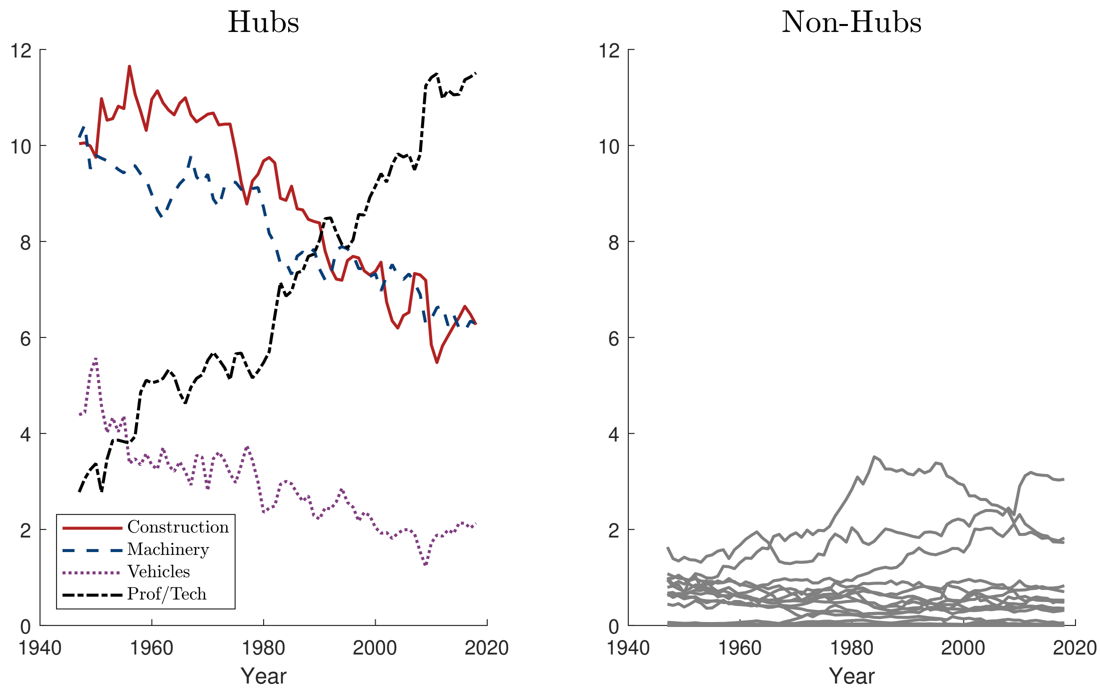

Notes: Time series of the sum over columns for the rows of the investment network corresponding to the four investment hubs: construction, machinery manufacturing, vehicles manufacturing, and professional/technical services.

"bridge files"). Recall that the investment network simply normalizes these pairwise flows of investment expenditures I_{ijt} by the total investment expenditure of sector j in year t.

Equation (24) shows that the time-variation in the pairwise investment flows I_{ijt} can come from one of two sources: variation in our estimated bridge files ω_{iat} or variation in investment expenditures I^{exp}_{ajt}. We assess the contribution of each of these components by computing two counterfactual investment flows series:

$$\widehat{I}^{\omega}_{ijt} = \omega_{iat} \overline{I}^{exp}_{aj} \tag{25}$$

$$\widehat{I}^{exp}_{ijt} = \overline{\omega}_{ia} I^{exp}_{ajt}, \tag{26}$$

where $\overline{I}^{exp}_{aj}$ is the expenditure of sector j on asset a averaged over time and $\overline{\omega}_{ia}$ is the bridge file for sector i producing asset a averaged over time. $\widehat{I}^{\omega}_{ijt}$ in equation (25) computes the

Table A.2
Sources of Variation in Investment Network

	Avg. variance (×1000)	*Bridge files only*	*Expenditures only*
Full Time Series	6.66	5.94 (89%)	0.26 (4%)
Business Cycle (% deviation)	9.37	7.84 (84%)	0.44 (5%)

Notes: First row is the average variance of the raw time series of measured investment flows (I_{ijt}) and counterfactual series where time variation is only present in bridge files ($\widehat{I}^{\omega}_{ijt}$) or investment expenditures ($\widehat{I}^{exp}_{ijt}$). Second row is the average variance of each of those series after being logged and HP filtered with parameter 6.25. Each variance is weighted by the average value of that element in the investment network. Investment network elements which take on a zero at any point in the time series are omitted; the first row results are insensitive to inclusion of these, however. Percentages do not sum to 100% since (1) the underlying relationships are not linear, so the decomposition is not exact, and (2) we do not consider covariance terms.

implied variation in I_{ijt} if investment expenditures were held fixed over time and therefore captures the contribution of time-variation in our estimated bridge files. Conversely, $\widehat{I}^{exp}_{ijt}$ in equation (26) computes the implied variation in I_{ijt} if the bridge file were held fixed over time, capturing the contribution of time-variation in investment expenditures. While these two objects do not sum up to the total investment flows I_{ijt}, their relative variation is a useful metric for quantifying the two sources of variation.

Table A.2 shows that the vast majority of time-variation in our estimated investment flows I_{ijt} is driven by variation in investment expenditures I^{exp}_{ajt}, not variation in the bridge file ω_{iat}. The top row computes the weighted average of variances of the counterfactual series $\widehat{I}^{\omega}_{ijt}$ and $\widehat{I}^{exp}_{ijt}$ compared to the variance of the total series I_{ijt}, while the bottom row first detrends the data using a log HP filter in order to focus on the business cycle fluctuations. In either case, the variance coming from our estimated bridge files, $\widehat{I}^{\omega}_{ijt}$, is only 5% of the variance in the overall capital flows series I_{ijt}.[47] Hence, the time-series variation in our estimated investment network parameters is not driven by variation in our estimated bridge files but rather variation in the investment expenditures reported by the BEA.[48]

[47]The average variance of the business cycle component is higher than the full time series because we compute the business cycle component as the log-deviation from trend, while the full time series computes variance of the level.

[48]This result also holds if we first compute the fraction of variance explained within each sector-pair for investment flows and then average that fraction across sector pairs.

Table A.3
Skewness of Investment and Intermediates Networks

	Eigenvalue Centrality	**Weighted Outdegree**
Investment network	3.32	2.70
Intermediates network	1.42	0.68

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.

Concentration of the Investment Network Table A.3 shows that the investment network is significantly more concentrated than the intermediates input-output network, measured using two metrics of network skewness.[49] Carvalho and Tahbaz-Salehi (2019) discuss both of these metrics; 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 sectors. A highly skewed set of centrality measures indicates that the network is dominated by a small number of highly important sectors. Across both measures of centrality, the investment network is on average roughly two to three times more skewed than the intermediates input-output network.

B Additional Results on Descriptive Evidence of Investment Hubs

This appendix present three pieces of additional analysis regarding the cyclical behavior of investment hubs referenced in Section 2 in the main text. First, Figure B.1 presents the correlogram between sector-level value added and aggregate GDP rather than aggregate employment as in Figure 2. Consistent with Figure 2, hubs are more correlated with aggregate GDP than are non-hubs, and this difference between hubs is larger in the post-1984 sample.

Second, we address the concern that the empirical behavior of investment hubs is driven by the fact that two of four hubs are manufacturing sectors, and that manufacturing may

[49] We describe our measurement of the intermediates network, which follows standard procedure, in Section 3.

Figure B.1: Correlogram of Sector-level Value Added with Aggregate GDP

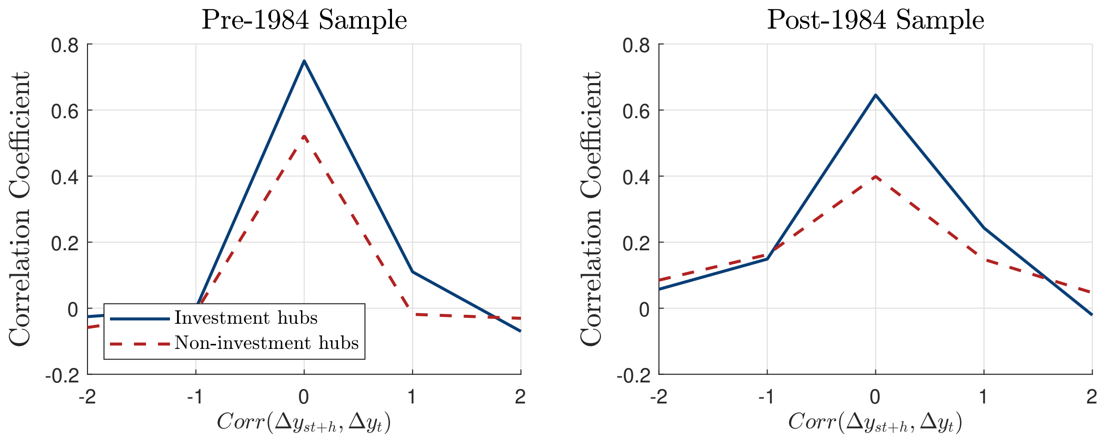

Notes: correlation of value added at sector s in year $t+h$, Δy_{st+h}, with aggregate employment in year t, Δy_t. Both y_{st+h} and y_t are logged and Δ denotes the first-difference operator. The x-axis varies the lead/lag $h \in \{-2, -1, 0, 1, 2\}$. "Investment hubs" compute the unweighted average 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 - 2018 subsample.

Table B.1
Volatility of Activity, Hubs vs. Manufacturing

	Investment Hubs		**Non-Hubs**		**Non-Hub Manuf.**	
	Pre-84	*Post-84*	*Pre-84*	*Post-84*	*Pre-84*	*Post-84*
$\sigma(\Delta y_{st})$	9.13%	9.18%	6.63%	5.51%	9.14%	6.97%
$\sigma(\Delta l_{st})$	6.14%	4.83%	3.81%	3.14%	5.12%	3.77%

Notes: standard deviation of business cycle component of sector-level value added or employment. y_{st} is logged real value added in sector s, l_{st} is logged employment in sector s, and Δ denotes the first difference operator. "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 - 2018 subsample.

FIGURE B.2: Correlogam of Sector-level Value Added with Aggregate Employment, Hubs vs. Manufacturing

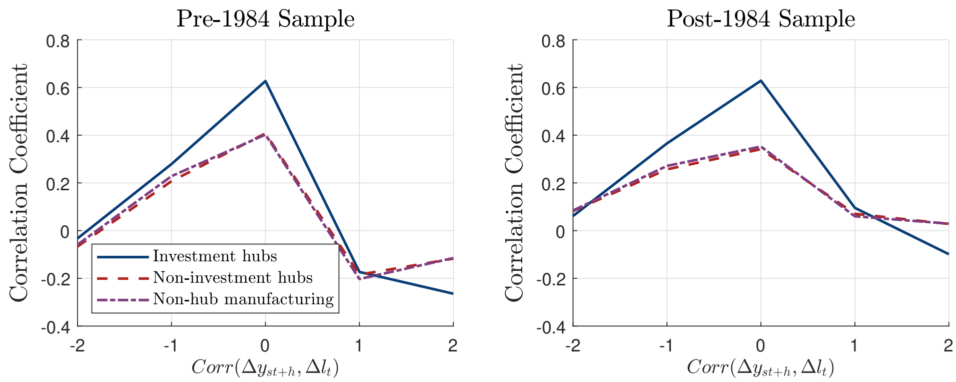

Notes: correlation of log real value added in sector s in year $t+h$, y_{st+h}, with log aggregate employment in year t, l_t. Δ denotes the first difference operator. The x-axis varies the lead/lag $h \in \{-2,-1,0,1,2\}$. "Investment hubs" compute the unweighted average the value of these statistics 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 - 2018 subsample.

be more cyclical than other sectors for reasons outside our model. We present two pieces of evidence against this concern. First, Table B.1 shows that non-hub manufacturing sectors are less volatile than investment hubs. Although non-hub manufacturing sectors are more volatile than non-hub non-manufacturing sectors, we show in Section 4 and Appendix F that this result is consistent with our model because durable manufacturing sectors are key intermediates suppliers to investment hubs. Second, Figure B.2 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 corresponding correlation of the investment hubs.

C Equilibrium Conditions

This appendix collects the equilibrium conditions of our model.

Households We simplify the household's problem in two ways. First, the intratemporal consumption allocation decision implies that $p_{jt}C_{jt} = \xi_j P_t^c C_t$, where $P_t^c = \Pi_{j=1}^N \left(\frac{p_{jt}}{\xi_j}\right)^{\xi_j}$ is the price index of the consumption bundle. We take the price of the consumption bundle

$P_t^c = 1$ as our numeraire. Second, the intratemporal investment allocation decision for sector j implies that $p_{it}I_{ijt} = \lambda_{ij}p_{jt}^I I_{jt}$, where $p_{jt}^I = \Pi_{i=1}^N \left(\frac{p_{it}}{\lambda_{ij}}\right)^{\lambda_{ij}}$ is the price index of the investment bundle for sector j.

With these simplifications, the household's problem is

$$\max_{C_t, K_{jt+1}, L_{jt}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \left(\log C_t - \chi L_t\right)\right] \text{ s.t. } C_t + \sum_{j=1}^N p_{jt}^I \left(K_{jt+1} - (1-\delta_j)K_{jt}\right) \leq W_t L_t + \sum_{j=1}^N r_{jt}K_{jt}.$$

The first order conditions for this problem are

$$\frac{p_{jt}^I}{C_t} = \beta \mathbb{E}_t \left[\frac{1}{C_{t+1}} \left(r_{jt+1} + p_{jt+1}^I(1-\delta_j)\right)\right] \tag{27}$$

$$\chi = \frac{W_t}{C_t}. \tag{28}$$

Firms Before solving the firm's profit maximization problem, we note that its cost-minimization problem with respect to intermediate input mix implies that $p_{it}M_{ijt} = \gamma_{ij}p_{jt}^M M_{jt}$, where $p_{jt}^M = \Pi_{i=1}^N \left(\frac{p_{it}}{\gamma_{ij}}\right)^{\gamma_{ij}}$ is the price index of the materials bundle for sector j. The profit maximization problem is then

$$\max_{L_{jt}, K_{jt}, M_{jt}} p_{jt}Q_{jt} - W_t L_{jt} - r_{jt}K_{jt} - p_{jt}^M M_{jt}.$$

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

The first order conditions for this problem are

$$W_t = \theta_j(1-\alpha_j)\frac{p_{jt}Q_{jt}}{L_{jt}} \tag{29}$$

$$r_{jt} = \theta_j \alpha_j \frac{p_{jt}Q_{jt}}{K_{jt}} \tag{30}$$

$$p_{jt}^M = (1-\theta_j)\frac{p_{jt}Q_{jt}}{M_{jt}}. \tag{31}$$

Note that constant returns to scale implies

$$W_t L_{jt} + r_{jt}K_{jt} + p_{jt}^M M_{jt} = p_{jt}Q_{jt}. \tag{32}$$

Therefore, the accounting definition of nominal value added is simply $p_{jt}Q_{jt} - p^M_{jt}M_{jt} = w_tL_t + r_{jt}K_{jt}$, which is by definition $p^Y_{jt}Y_{jt}$.

To obtain real value added, we use the Divisia index definition, which differentiates the accounting definition of nominal value added holding prices fixed:

$$\begin{aligned} p^Y_{jt}dY_{jt} &= p_{jt}dQ_{jt} - p^M_{jt}dM_{jt} \\ p^Y_{jt}Y_{jt}d\log Y_{jt} &= p_{jt}Q_{jt}d\log Q_{jt} - p^M_{jt}M_{jt}d\log M_{jt} \\ \theta_j d\log Y_{jt} &= d\log Q_{jt} - (1-\theta_j)d\log M_{jt} \\ d\log Y_{jt} = \frac{1}{\theta_j}&d\log A_{jt} + \alpha_j d\log K_{jt} + (1-\alpha_j)d\log L_{jt} \end{aligned}$$

Market Clearing Output market clearing for sector j ensures that gross output is used for consumption, investment, or an intermediate in production:

$$Q_{jt} = C_{jt} + \sum_{i=1}^{N} I_{jit} + \sum_{i=1}^{N} M_{jit}. \tag{33}$$

Using the firms' first order conditions for optimal investment and intermediates purchases, we can rewrite this condition to avoid the need to keep track of each intermediate purchase and consumption:

$$Q_{jt} = \frac{\xi_j C_t}{p_{jt}} + \sum_{i=1}^{N} \frac{\lambda_{ji}p^I_{it}I_{it}}{p_{jt}} + \sum_{i=1}^{N} \frac{(1-\theta_i)\gamma_{ji}p_{it}Q_{it}}{p_{jt}} \tag{34}$$

D Details of Model Calibration

This appendix presents additional details on our calibration of the model. As discussed in the main text, we choose all of the parameters other than the shock process so that the model's steady state corresponds to the average of the postwar U.S. economy. We then feed in the measured productivity shocks from the data.

Figure D.1: Calibrated Value Added Shares θ_j

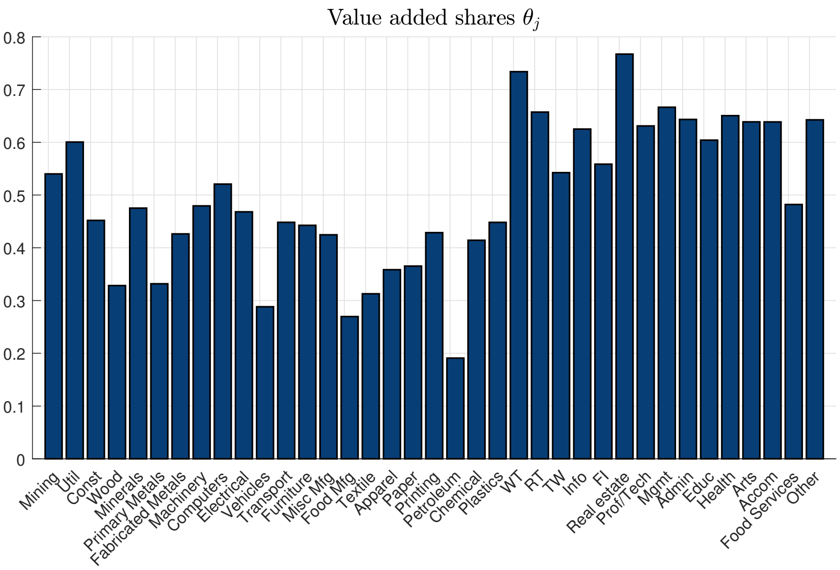

Notes: Values for the value-added shares θ_j are computed as the ratio of value added to gross output in each sector, averaged across the entire sample, 1947-2018.

D.1 Steady State Parameters

Figure D.1 plots our calibrated primary input shares θ_j for each sector j. 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 from the BEA input-output data. We obtain this ratio for each year in our 1947-2018 sample and then compute their average value over time.

Figure D.2 plots the calibrated labor shares $1 - \alpha_j$ for each sector, averaged over 1947 - 2018.[50] To correct for the fact that sector-level compensation in the BEA data does not include self-employed income, we multiply sectoral compensation by one plus the ratio of self-employed employment to total part-time and full-time employment in the sector.[51] We then

[50]For years prior to 1987, we convert SIC based data to NAICS using the crosswalk in Fort and Klimek (2018).

[51]This operation implicitly assumes that average compensation for self-employed workers is the same as non-self-employed workers. The BEA data on self-employment by sector covers a coarse set of sectors, so we apply the self-employment to employment ratio to each sector based on the finest available sector in the self-employed data. The one exception is for the management of companies and enterprises, for which we assume that there is no self-employment. If we allowed for self-employment in that sector, the implied labor share often exceeds one.

Figure D.2: Calibrated Labor Shares $1 - \alpha_j$

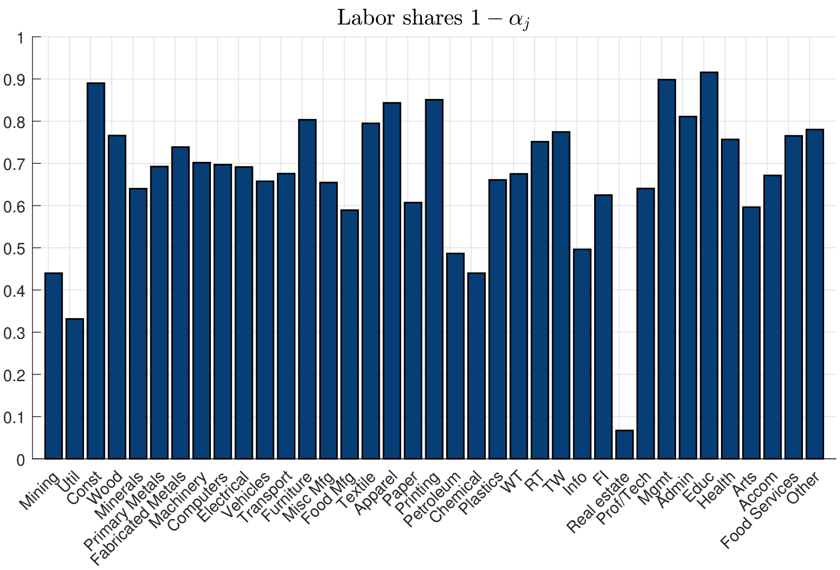

Notes: Values for the labor share $1 - \alpha_j$ are computed from sectoral data on compensation (adjusted for self-employment) divided by value added (with indirect taxes and subsidies removed), averaged across all years in the data, 1947-2018.

compute the labor share as the ratio of adjusted compensation to value added in that sector minus indirect taxes and subsidies. Our results are also robust to making no adjustments for self-employment.

Figure D.3 plots our calibrated depreciation rates, δ_j, which are equal to the average implied depreciation rate reported in the Fixed Assets database from 1947-2018. Figure D.4 plots our calibrated Cobb-Douglas preference parameters weighting consumption in different sectors' output, ξ_j. We measure ξ_j as the share of total consumption expenditures purchased from sector j.

D.2 Measured Sector-Level Productivity Series

We measure sector-level TFP using the Solow residual approach. In particular, we compute TFP for sector j in year t as

$$\log A_{jt} = \log Q_{jt} - \theta_{jt}\alpha_{jt} \log K_{jt} - \theta_{jt}(1 - \alpha_{jt}) \log L_{jt} - (1 - \theta_{jt}) \log M_{jt},$$

Figure D.3: Calibrated Depreciation Rates δ_j

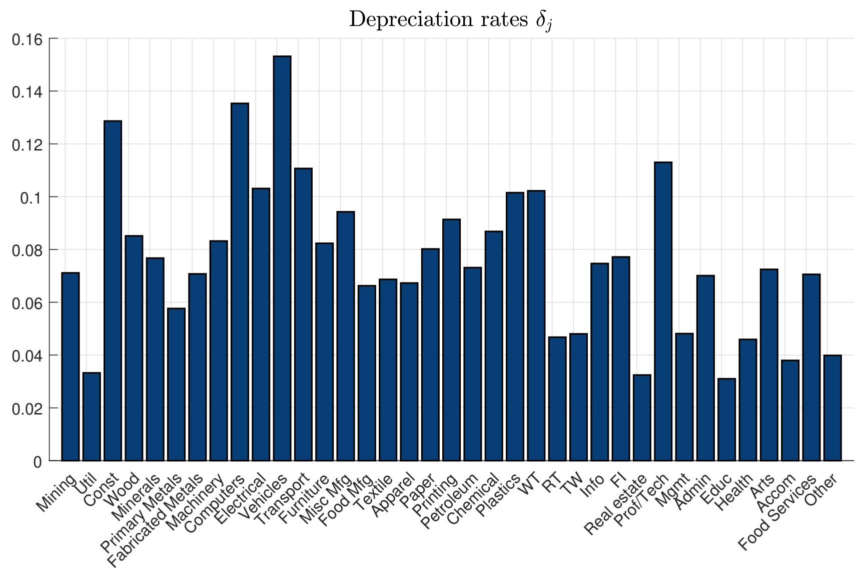

Notes: Values for sector-level depreciation rates δ_j are taken as each sector's average implied depreciation rate from BEA Fixed Assets data, averaged from 1947-2018.

Figure D.4: Calibrated Consumption Shares ξ_j

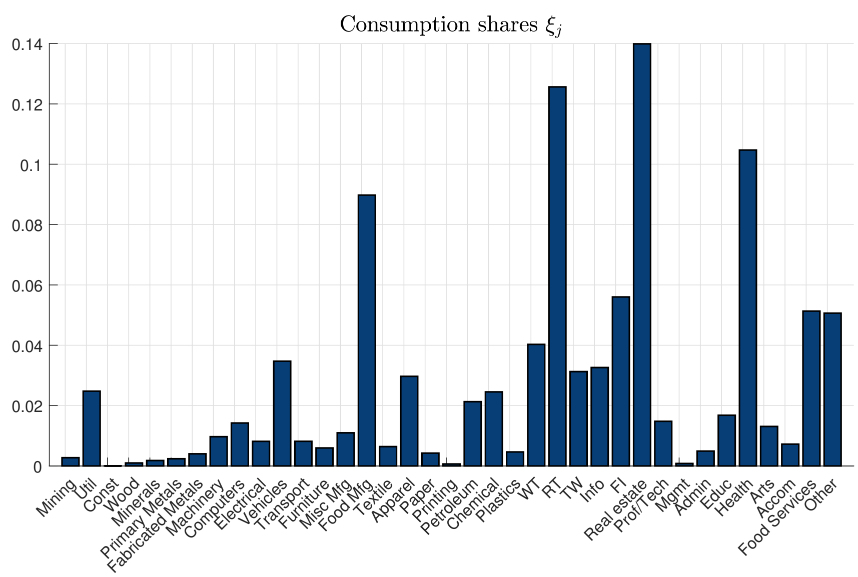

Notes: Values for consumption preference ξ_j are constructed as the fraction of total nominal consumption expenditures on each sector's goods or services, averaged over the entire sample 1947-2018.

Figure D.5: Detrending Sector-Level Data

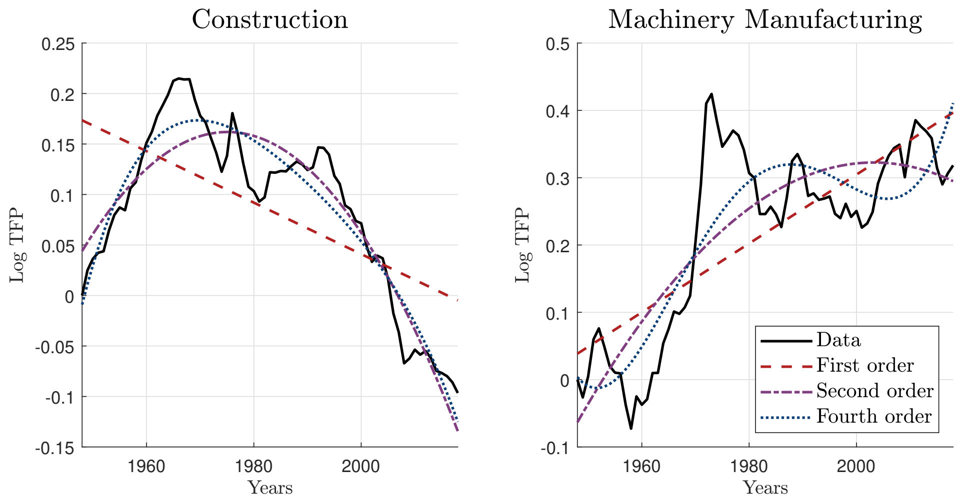

Notes: The figure reports log sector level TFP for the Construction and Machinery Manufacturing sectors, normalized to zero in the year 1948. We also report a fitted polynomial trend lines for polynomials of order 1, 2, and 4, estimated via OLS.

where the factor shares vary over time in order to capture changes in the production technology that are outside our model (our results are robust to fixing the factor shares over time). We construct the capital stock for each sector in each year via the perpetual inventory method, using the nominal year-end capital stock for each sector in 1948 as our starting point (from BEA Fixed Assets data).

We detrend our model using a log-polynomial trend because log-linear trends provide a poor fit to sector-level TFP. Figure D.5 plots the time series of sector-level TFP for two example sectors, construction and machinery manufacturing. Construction TFP evolves nonlinearly over time and a fourth order polynomial trend captures these nonlinearities.[52] 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, but we show in Appendix G that our main results are robust to using lower-order polynomials for detrending. Figure D.6 plots the persistence parameters ρ_j, which we estimate using maximimum likelihood on detrended

[52]We do not present the third order trends in this figure for parsimony, but they are generally more similar to fourth order trends than to the second order trends.

FIGURE D.6: Calibrated Persistence Parameters ρ_j

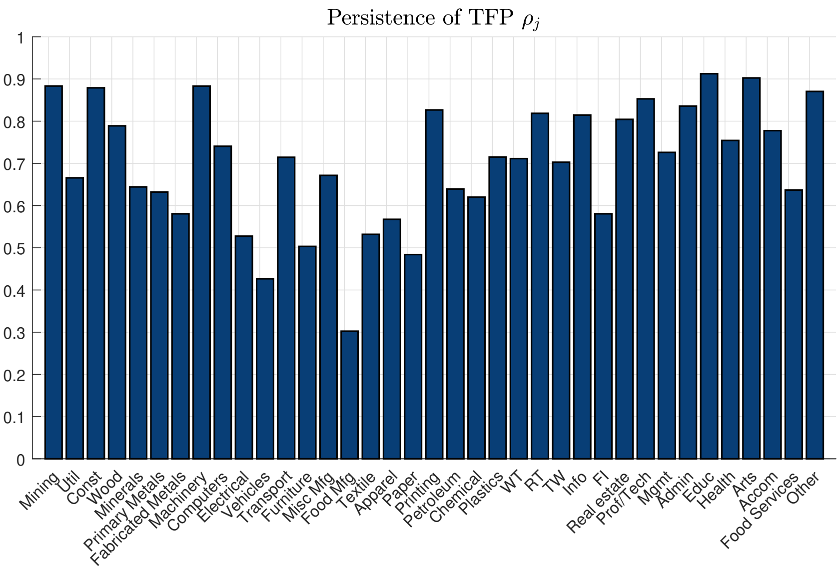

Notes: Persistence parameters ρ_j of sector-level TFP are estimated from detrended TFP using maximum likelihood.

log-TFP.

Interpreting changes in productivity over time using principal components In the main text, we interpreted the decline in the correlation of TFP across sectors as reflecting 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 decomposition similar to Garin, Pries and Sims (2018). 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 37 sectors and less than 37 years of data in each time period. We therefore collapse our data down to 30 sectors by aggregating all non-durable manufacturing sectors into one sector and then perform the principal components decomposition on log TFP growth for 30 sectors pre- and post-1984.[53]

[53]We could have alternatively collapsed a different set of sectors, but we prefer this approach because: aggregating within non-durable manufacturing does not affect the investment hubs or their key suppliers, many non-durable manufacturing sectors are small, and the aggregated sector of non-durable manufacturing is more intuitive than aggregates of alternative sets of service sectors.

Table D.1
Principal Components Analysis of Measured TFP

Sample period	$1000\mathbb{V}ar(\Delta \log A_t)$	Due to 1st component	Residual
1949-1983	0.41	0.31 (75%)	0.10 (25%)
1984-2017	0.09	0.03 (35%)	0.06 (65%)

Notes: the aggregate shock is equal to the vector product of the loadings associated with the first principal component with the vector of sector-level TFP. 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).

Table D.1 reports the results of this principal components exercise. The first principal component – which can be loosely interpreted as the aggregate shock – accounts for 75% of the variance of aggregate TFP in the pre-1984 sample, but only 35% 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 – declines by much less over time.

E Proofs

This appendix proves the three propositions in Section 4.

Proof of Proposition 1 Plug in the definition of sector-level real value added growth $d\log Y_{jt}$ (omitting capital, because it is fixed upon impact) to the Divisia index to get

$$d\log Y_t = \sum_{j=1}^{N} \left(\frac{p^Y_{jt} Y_{jt}}{P^Y_t Y_t}\right)\left(\frac{1}{\theta_j} d\log A_{jt} + (1-\alpha_j) d\log L_{jt}\right). \tag{35}$$

The intermediates first order condition (31) and the zero profit condition (32) imply that θ_j is equal to the ratio of value added to gross output: $\theta_j = \frac{p^Y_{jt}Y_{jt}}{p_{jt}Q_{jt}}$. Therefore, the weight on TFP in the sum (35) is $\frac{p^Y_{jt}Y_{jt}}{P^Y_t Y_t}\frac{p_{jt}Q_{jt}}{p^Y_{jt}Y_{jt}} = \frac{p_{jt}Q_{jt}}{P^Y_t Y_t}$ – the Domar weight.

The labor first order condition (29) can be rearranged to $(1-\alpha_j)\theta_j p_{jt} Q_{jt} = W_t L_{jt}$. But again, the zero profits condition implies that $\theta_j p_{jt} Q_{jt} = p^Y_{jt} Y_{jt}$, so this condition becomes $(1-\alpha_j) p^Y_{jt} Y_{jt} = W_t L_{jt}$. Divide this expression by nominal GDP to get $(1-\alpha_j)\frac{p^Y_{jt}Y_{jt}}{P^Y_t Y_t} = \frac{W_t L_{jt}}{P^Y_t Y_t}$.

Then sum over sectors j to get $1-\alpha_t \equiv \sum_{j=1}^{N}(1-\alpha_j)\frac{p_{jt}^{Y}Y_{jt}}{P_t^{Y}Y_t} = \frac{W_t L_t}{P_t^{Y}Y_t}$, the aggregate labor share. Then multiply this expression by $\frac{L_{jt}}{L_t}$ and combine with the previous expression to get $(1-\alpha_t)\frac{L_{jt}}{L_t} = \frac{W_t L_{jt}}{P_t^{Y}Y_t} = (1-\alpha_j)\frac{p_{jt}^{Y}Y_{jt}}{P_t^{Y}Y_t}\frac{L_t}{L_{jt}}$.

Plugging all this into the expression for real GDP growth (35) gives

$$d\log Y_t = \sum_{j=1}^{N}\left(\left(\frac{p_{jt}Q_{jt}}{P_t^{Y}Y_t}\right)d\log A_{jt} + (1-\alpha_t)\frac{L_{jt}}{L_t}d\log L_{jt}\right).$$

Under a first-order approximation, fluctuations in either the Domar weight $\left(\frac{p_{jt}Q_{jt}}{P_t^{Y}Y_t}\right)$ or the employment share $(1-\alpha_t)\frac{L_{jt}}{L_t}$ multiply TFP growth or employment growth, which are zero in steady state. This insight yields the result in Proposition 1.

Proof of Proposition 2 The market clearing condition for sector j in terms of overall expenditures is $p_{jt}Q_{jt} = p_{jt}C_{jt} + \sum_{i=1}^{N} p_{jt}I_{jit} + \sum_{i=1}^{N} p_{jt}M_{jit}$. Due to the Cobb-Douglas production functions, sector i's expenditures on intermediates from sector j is simply proportional to sector i's total sales: $p_{jt}M_{jit} = (1-\theta_i)\gamma_{ji}p_{it}Q_{it}$. Similarly, sector i's expenditures on investment goods from sector j is $p_{jt}I_{jit} = \lambda_{ji}p_{it}^{I}I_{it}$, where $p_{it}^{I} = \Pi_{k=1}^{N}\left(\frac{p_{kt}}{\lambda_{ki}}\right)^{\lambda_{ki}}$ is the price index for investment. Therefore, total expenditure on sector j is

$$p_{jt}Q_{jt} = p_{jt}C_{jt} + \sum_{i=1}^{N}\lambda_{ji}p_{it}^{I}I_{it} + \sum_{i=1}^{N}(1-\theta_i)\gamma_{ji}p_{it}Q_{it}. \tag{36}$$

For notational convenience, define

$$\widehat{Q}_t = \begin{bmatrix} p_{1t}Q_{1t} \\ \vdots \\ p_{Nt}Q_{nt} \end{bmatrix}, \ \widehat{C}_t = \begin{bmatrix} p_{1t}C_{1t} \\ \vdots \\ p_{Nt}C_{nt} \end{bmatrix}, \text{ and } \widehat{I}_t = \begin{bmatrix} p_{1t}I_{1t} \\ \vdots \\ p_{Nt}I_{nt} \end{bmatrix}.$$

Then the market clearing condition (36) can be written in matrix form as $\widehat{Q}_t = \widehat{C}_t + \Lambda'\widehat{I}_t + \Gamma'\widehat{Q}_t$, where Λ is the investment network matrix. Solve out this expression for $\widehat{Q}_t$ to get

$$\widehat{Q}_t = (I-\Gamma')^{-1}\left(\widehat{C}_t + \Lambda'\widehat{I}_t\right).$$

Writing this equation for element j, dividing by aggregate consumption C_t, and noting that $(I - \Gamma')^{-1}$ gives

$$\frac{p_{jt}Q_{jt}}{C_t} = \sum_{k=1}^{N} \ell_{jk} \frac{p_{kt}C_{kt}}{C_t} + \sum_{k=1}^{N} \ell_{jk} \sum_{m=1}^{N} \lambda_{km} \frac{p^I_{mt} I_{mt}}{C_t}.$$

Now plug this expression into the equilibrium labor supply relationship from equation (10), $L_{jt} = (1-\alpha_j)\theta_j \frac{p_{jt}Q_{jt}}{C_t}$, to get

$$L_{jt} = (1-\alpha_j)\theta_j \left[\sum_{k=1}^{N} \ell_{jk} \frac{p_{kt}C_{kt}}{C_t} + \sum_{k=1}^{N} \ell_{jk} \sum_{m=1}^{N} \lambda_{km} \frac{p^I_{mt} I_{mt}}{C_t} \right], \tag{37}$$

which is proportional to equation (11) in the main text. To arrive at equation (12), simply take log-deviations from steady state and note that $\frac{p_{kt}C_{kt}}{C_t} = \xi_k$ is constant over time.

Proof of Proposition 3 Using the first order conditions for the profit maximization problem, equations (29)-(31)), we can write the price of each sector j's final good as:

$$\begin{aligned} p_{jt} &= \frac{1}{A_{jt}} \left(\frac{r_{jt}}{\alpha_j \theta_j} \right)^{\alpha_j \theta_j} \left(\frac{W_t}{(1-\alpha_j)\theta_j} \right)^{(1-\alpha_j)\theta_j} \left(\frac{p^M_{jt}}{1-\theta_j} \right)^{1-\theta_j} \\ &= \frac{1}{A_{jt}} \left(\frac{r_{jt}}{\alpha_j \theta_j} \right)^{\alpha_j \theta_j} \left(\frac{W_t}{(1-\alpha_j)\theta_j} \right)^{(1-\alpha_j)\theta_j} \left(\frac{\prod_{i=1}^{N} \left(\frac{p_{it}}{\gamma_{ij}} \right)^{\gamma_{ij}}}{1-\theta_j} \right)^{1-\theta_j} \end{aligned}$$

using the fact that $p^M_{jt} = \prod_{i=1}^{N} \left(\frac{p_{it}}{\gamma_{ij}} \right)^{\gamma_{ij}}$.

Taking the log of both sides gives us:

$$\log p_{jt} = -\log A_{jt} + \alpha_j \theta_j \log r_{jt} + (1-\alpha_j)\theta_j \log W_t + \sum_{i=1}^{N} (1-\theta_j)\gamma_{ij} \log p_{it} + \Phi_j$$

where $\Phi_j = \log \left(\left(\frac{1}{\alpha_j \theta_j} \right)^{\alpha_j \theta_j} \left(\frac{1}{(1-\alpha_j)\theta_j} \right)^{(1-\alpha_j)\theta_j} \left(\frac{\prod_{i=1}^{N} \left(\frac{1}{\gamma_{ij}} \right)^{\gamma_{ij}}}{1-\theta_j} \right)^{1-\theta_j} \right)$.

To assess the direct effect of a TFP shock on output prices, we totally differentiate the above expression, holding fixed any response of the rental rates or wages, obtaining:

$$d\log p_{jt} = -d\log A_{jt} + \sum_{i=1}^{N}(1-\theta_j)\gamma_{ij}d\log p_{it}$$

Or in matrix notation,

$$\begin{aligned} d\log p_t &= -d\log A_t + \Gamma' d\log p_t \\ d\log p_t &= -\mathcal{L}' d\log A_t \end{aligned}$$

where $d\log p_t$ is an $N \times 1$ vector of sector-level prices and $d\log A_t$ is the vector of sector-level productivity.

To relate this to the investment price index, we use the fact that $p^I_{jt} = \prod_{i=1}^N \left(\frac{p_{it}}{\lambda_{ij}}\right)^{\lambda_{ij}}$ and thus:

$$\begin{aligned} d\log p^I_t &= \Lambda' d\log p_t \\ &= -(\mathcal{L}\Lambda)' d\log A_t \end{aligned}$$

In non-matrix notation, this implies the result that $d\log p^I_{mt} = -\sum_{i=1}^N \omega_{im} d\log A_{it}$, yielding the proposition in the text.

F Additional Results For Section 4

This appendix describes additional results mentioned in Section 4 of the main text.

F.1 Relationship 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)). This literature typically works with two-sector models in which one sector produces only consumption goods and the other only produces investment goods with no intermediate goods connections between them.

Our model provides a richer sectoral disaggregation to bring the model to the data because the correct notion of the "investment producers" includes the key suppliers of investment hubs in the Leontief-adjusted investment network. As we show in the main text, productivity shocks in these sectors generate substantial aggregate fluctuations. In contrast, a common approach to measuring shocks in the investment-specific shock literature is to use the aggregate price index of investment relative to consumption. It is difficult to generate large business cycle fluctuations with this price series because it is only weakly correlated with the aggregate cycle.

An equally important but more subtle issue is that the investment-specific shock literature struggles to generate positive comovement between the consumption- and investment-producing sectors. To help understand this issue, rewrite equation (12) without the intermediates network:

$$d \log L_{jt} = \sum_{m=1}^{N} \lambda_{jm} \left(\frac{p_m^{I*} I_m^*}{p_j^* Q_j^*} \right) \left(d \log p_{mt}^I I_{mt} - d \log C_t \right),$$

which is the same as (12) in the main text except that the Leontief-adjusted investment network is equal to the raw network: $\Omega = \Lambda$. Following the same logic in the main text, only employment in the investment hubs will meaningfully fluctuate over time because the other sectors have a small role in producing investment goods (i.e. λ_{jm} is small for non-hub sectors j).

Table F.1 illustrates the comovement problem in the version of our model without the intermediates network (i.e. setting the input-output network $\Gamma = I$). The table computes the correlation between employment fluctuations at our four investment hub sectors and a set of sectors we define as "consumption hubs:" food manufacturing, finance & insurance, real estate, retail trade, health services, food services, and other services (these seven sectors comprise roughly 60% of consumption expenditures in our model calibration). Without the intermediates network, the correlation between employment in these two sets of sectors is counterfactually low.[54] The intermediates network Γ in our full model solves this comovement

[54]With an an infinite Frisch elasticity ($\eta \to \infty$), the correlation would be exactly zero if none of the

Table F.1
Comovement of Investment and Consumption Sectors, Data and Model

	Data		**Model**		**Model, w/o Intermediates**	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*	*Pre-1984*	*Post-1984*
$\overline{\rho(\Delta l_{ct}, \Delta l_{it})}$	0.51	0.50	0.76	0.81	0.18	0.22

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2018). l_{ct} is log employment in consumption hub sectors, defined as: food manufacturing, finance & insurance, real estate, retail trade, health services, food services, and other services. l_{it} is log employment in investment hub sectors. $\overline{\rho(\Delta l_{ct}, \Delta l_{it})}$ denotes the average correlation between the employment fluctuations for consumption and investment hub sector. "Model" corresponds to simulations in our model where shocks are measured from the data and then fed into model described in the main text, albeit without the investment production frictions described in Section 5. "Model, w/o Intermediates" corresponds to simulations from the same exercise, except where there are no intermediate inputs in production (i.e. setting the share of value added in gross output equal to 1, $\theta_j = 1$.)

problem because it implies consumption-producing sectors also supply intermediate goods to investment-producing sectors. In contrast, the investment-specific shocks literature uses other nominal or real rigidities to overcome the negative comovement problem.[55]

The fact that the model without intermediates generates too little comovement implies that aggregate employment will be less volatile as well. Figure F.1 illustrates this issue by computing the elasticity of aggregate employment with respect to a each sector-specific shock A_{it} in the version of our model without intermediates ($\Gamma = I$). Only the shocks to the investment hubs, highlighted in red, have a meaningful impact on aggregate employment. Furthermore, their effect on aggregate employment is primarily limited to employment in the hubs themselves.

Figure F.1: Elasticity of Aggregate Employment to Sectoral Shocks Without Intermediates Network

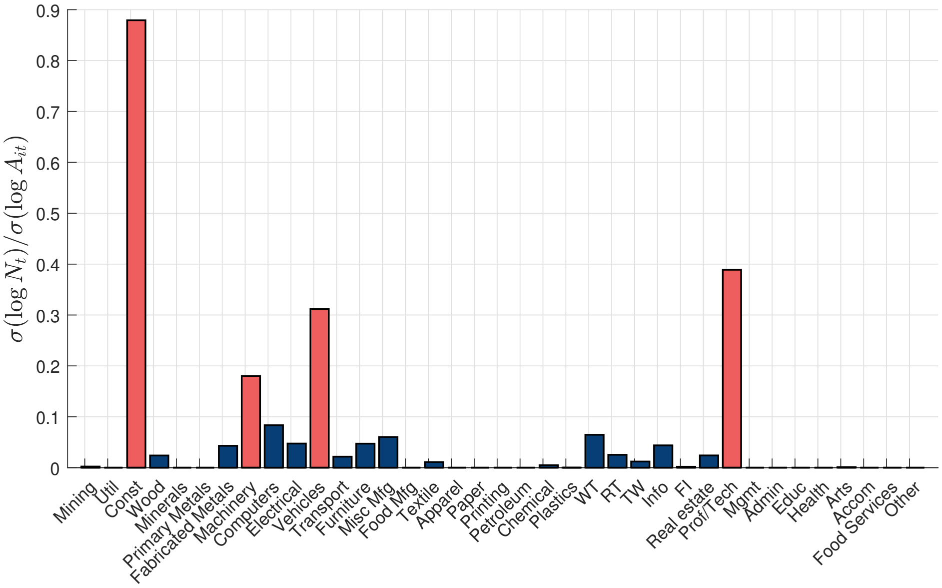

Notes: reduced-form elasticities of aggregate employment N_t to sector-specific shocks A_{it} in a version of the model without intermediate goods (i.e. $\theta_j = 0$ for all j). For each sector, we simulate the model with $\sigma(\log A_{it}) = 1\%$ shocks to that sector only. The bars plot the volatility of aggregate employment $\sigma(\log N_t)$. Investment hubs are highlighted in red.

Table F.2
Volatility of Activity, Hubs vs. Intermediate Suppliers

	Investment Hubs		**Suppliers**		**Others**	
Data	*Pre-84*	*Post-84*	*Pre-84*	*Post-84*	*Pre-84*	*Post-84*
$\sigma(\Delta y_{st})$	9.13%	9.18%	8.03%	6.72%	5.94%	4.90%
$\sigma(\Delta l_{st})$	6.14%	4.83%	6.04%	4.04%	2.70%	2.69%
Model	*Pre-84*	*Post-84*	*Pre-84*	*Post-84*	*Pre-84*	*Post-84*
$\sigma(\Delta y_{st})$	12.92%	9.63%	9.02%	7.03%	5.57%	4.93%
$\sigma(\Delta l_{st})$	9.37%	6.65%	5.93%	4.28%	1.68%	1.18%

Notes: standard deviation of business cycle component of sector-level value added or employment. y_{st} is logged real value added in sector s, l_{st} is logged employment in sector s, and Δ denotes the first difference operator. "Investment hubs" compute the unweighted average the value of these statistics over $s =$ construction, machinery manufacturing, motor vehicles manufacturing, and professional/technical services. "Suppliers" computes the weighted average over the non-hub sectors of durable manufacturing, wholsesale trade, and transportation & warehousing. "Others" computes the unweighted average over the sectors not classified as investment hubs or suppliers. "Pre-1984" performs this analysis in the 1948 - 1983 subsample and "post-1984" performs this analysis in the 1984 - 2018 subsample.

F.2 Supporting Evidence for Mechanism in the Data

We present supporting evidence for the role of the key suppliers to investment hubs discussed in 4. Table F.2 shows that the key suppliers to investment hubs are more volatile over the business cycle than other non-hub sectors, consistent with the role of the Leontief-adjusted investment network in propagating shocks. The model provides a good fit for the behavior of these sectors, especially for employment. The table also shows that the suppliers are less volatile than the hubs themselves, again consistent with the model. Figure F.2 shows that the key suppliers to investment hubs are more correlated with the aggregate business cycle than other non-hub sectors, consistent with their role in propagating sector-specific shocks to aggregates.

F.3 Cobb-Douglas Capital Accumulation

We now show that employment is constant in the version of the model in which we replace the standard linear capital accumulation rule, $K_{jt+1} = (1-\delta_j)K_{jt} + I_{jt}$, with a Cobb-Douglas one:

$$K_{jt+1} = K_{jt}^{1-\delta_j} I_{jt}^{\delta_j}. \tag{38}$$

While this alternative rule (38) is inconsistent with national accounting practice, and thus not suitable for a quantitative model, it is nonetheless useful in explaining how investment drives our results. In particular, we will show that the Cobb-Douglas form (38) implies that investment expenditure is proportional to total income, which in turn implies that sector-specific shocks generate exactly offsetting income and substitution effects which leave employment unchanged.[56]

consumption hub sectors produced investment because their employment would be constant (see Proposition 2). However, retail trade, finance & insurance and real estate produce small amounts of investment, so employment in those sectors does fluctuate over time. With a finite Frisch elasticity, an increase in an investment hub sector also increases the marginal disutility of supplying labor to non-hub sectors, which would decrease employment in those sectors and generate negative comovement. See Kim and Kim (2006) for further discussion of the role of the Frisch elasticity in determining sectoral comovement.

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

[56] We thank Ernest Liu and Matt Rognlie for pointing this property out to us.

FIGURE F.2: Correlogram of Sector-level Value Added with Aggregate Employment, Hubs vs. Intermediate Suppliers

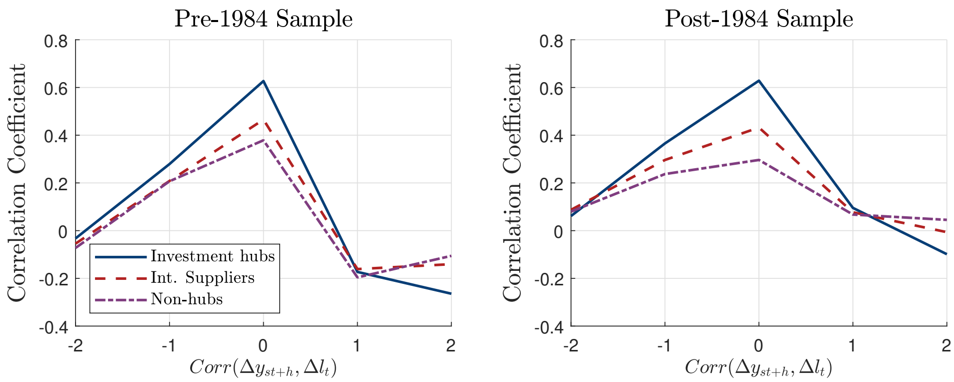

Notes: correlation of log real value added in sector s in year $t-h$, y_{st+h}, with log aggregate employment in year t, l_t. Δ denotes the first difference operator. The x-axis varies the lead/lag $h \in \{-2,-1,0,1,2\}$. "Investment hubs" compute the unweighted average the value of these statistics over $s =$ construction, machinery manufacturing, motor vehicles manufacturing, and professional/technical services. "Intermediate Suppliers" computes these statistics for the remaining durable manufacturing sectors, wholesale trade, and transportation & warehousing. "Non-hubs" computes 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 - 2018 subsample.

The alternative capital accumulation rule changes the Euler equation for capital (27) into

$$\frac{p^I_{jt} I_{jt}}{C_t} \frac{1}{\delta_j K_{jt+1}} = \beta \mathbb{E}_t \left[\frac{1}{C_{t+1}} \left(\alpha_j \theta_j \frac{p_{jt+1} Q_{jt+1}}{K_{jt+1}} + \frac{(1-\delta_j)}{\delta_j} \frac{p^I_{jt+1} I_{jt+1}}{K_{jt+1}} \right) \right].$$

which can be rearranged into

$$\frac{p^I_{jt} I_{jt}}{C_t} = \beta \mathbb{E}_t \left[\delta_j \alpha_j \theta_j \frac{p_{jt+1} Q_{jt+1}}{C_{t+1}} + (1-\delta_j) \frac{p^I_{jt+1} I_{jt+1}}{C_{t+1}} \right]. \tag{39}$$

We now guess and verify that the household's valuation of output and investment are constant over time. Denote those constants as $I^*_j = \frac{p^I_{jt} I_{jt}}{C_t}$ and $Q^*_j = \frac{p_{jt} Q_{jt}}{C_t}$. The Euler equation (39) relates these two objects through

$$I^*_j = \frac{\beta \delta_j \alpha_j \theta_j}{1-\beta(1-\delta_j)} Q^*_j, \tag{40}$$

Now define $B_j = \frac{\beta \delta_j \alpha_j \theta_j}{1-\beta(1-\delta_j)}$ and B to be the matrix with B_j on the diagonals and zero off-diagonal.

We plug (40) into the expression for the household's value of output (11) in order to solve for Q_j^* and I_j^*. We write the market clearing condition in matrix form using the notation

$$Q^* = \begin{bmatrix} Q_1^* \\ \vdots \\ Q_N^* \end{bmatrix} \text{ and } \widehat{\xi} = \begin{bmatrix} \xi_1 \\ \vdots \\ \xi_N \end{bmatrix}.$$

Using this notation, and plugging in our guesses, the household's value of output from (11) becomes

$$\widehat{Q}^* = \mathcal{L}\xi + \beta\mathcal{L}\Lambda' B\widehat{Q}^*,$$

where the second term on the right-hand side uses the fact that $I_j^* = B_j Q_j^*$ from (40). Solving this equation for Q^* yields

$$\begin{aligned} Q^* &= (I - \beta\mathcal{L}\Lambda' B)^{-1}\mathcal{L}\xi \\ &= \sum_{s=0}^{\infty} (\beta\mathcal{L}\Lambda' B)^s \mathcal{L}\xi \end{aligned} \tag{41}$$

Equation (41) shows that the household's valuation of output equals the discounted sum of its value of consumption, taking into account the ability to transfer resources over time using investment. The only condition we need to verify is that our guessed equilibrium is consistent with constant labor supply. Given our growth-consistent preferences, we indeed have that $L_j^* = \frac{\theta_j(1-\alpha_j)}{\chi} Q_j^*$ is constant over time.

Hence, the Cobb-Douglas capital accumulation equation (38) implies that investment — and the investment network — are irrelevant for aggregate dynamics beyond their impact on the steady state Domar weights. Intuitively, the Cobb-Douglas capital accumulation equation implies that investment expenditures are proportional to total income, which in turn is proportional to gross output. Therefore, sector-specific shocks generate equal-sized income and substitution effects, just as in the model without investment. Our full model with the linear capital accumulation rule breaks this irrelevance result by increasing the elasticity of the capital stock with respect to current investment.[57] In this case, changes in current investment have a larger effect on the capital stock, breaking the result that investment expenditures are proportional to output. This property allows the household's valuation of output, and therefore employment, to fluctuate over time.

Relationship to Full Depreciation It is well-known that the one-sector RBC model, with the standard linear capital accumulation rule, admits a closed-form solution with constant employment in the case of full depreciation. The discussion above makes clear that full depreciation is just a special case of the Cobb-Douglas capital accumulation rule (38) with $\delta_j = 1$; indeed, it is the only value of δ_j for which the linear and Cobb-Douglas capital accumulation rules are the same.

F.4 Other Analysis Mentioned in Main Text

This subsection collects a number of miscellaneous results referenced in Section 4.

Distribution of Domar weights Figure F.3 shows that the model fits the stationary distribution of Domar weights fairly well. In the model, a sector's Domar weight is related to its role in supplying consumption and investment goods. The Domar weights at investment hubs are not abnormally large because investment is a smaller fraction of overall spending than consumption.

Cyclicality of Labor Productivity Due to Sectoral Shocks Subtracting aggregate employment from our expression for real GDP in Proposition 1, the impact effect of a sector-specific shock A_{it} on aggregate labor productivity LP_t is

$$d\log LP_t = \sum_{j=1}^{N} \left(\frac{p_j Q_j}{P^Y Y}\right)^* d\log A_{jt} - \alpha^* \sum_{j=1}^{N} \left(\frac{L_j}{L}\right)^* d\log L_{jt}$$

All else equal, higher aggregate TFP increases labor productivity because it increases the productivity of all factors; on the other hand, higher aggregate employment decreases labor

[57] Of course, with the linear accumulation rule, that elasticity becomes infinite; more generally, we conjecture that increasing the elasticity beyond the Cobb-Douglas case will generate fluctuations in employment.

Figure F.3: Stationary Distribution of Domar Weights

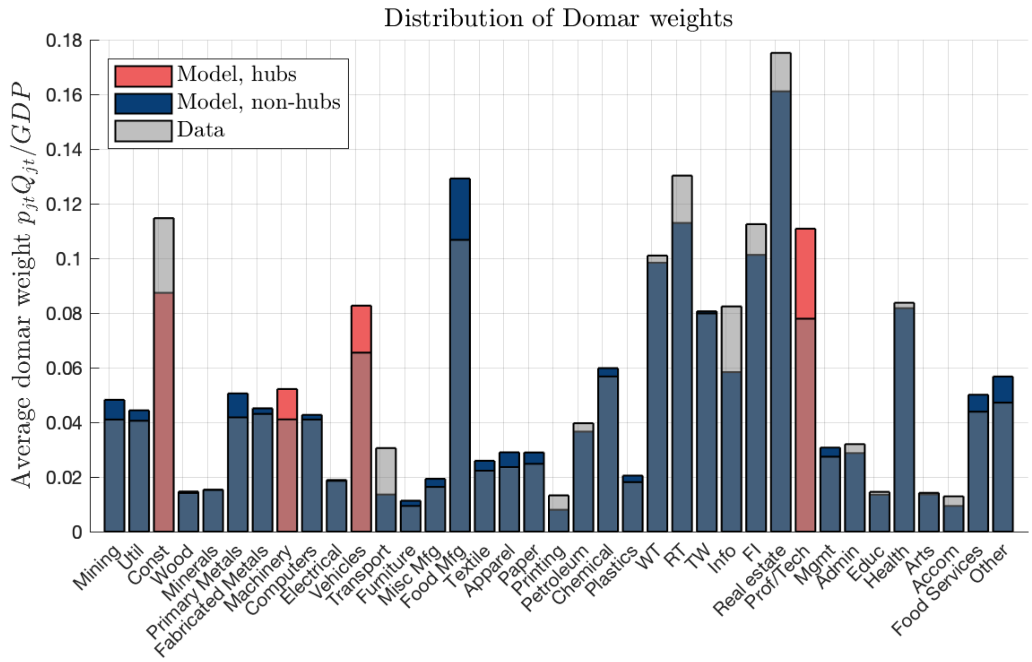

Notes: average values of the Domar weights $\mathbb{E}[\frac{p_{jt}Q_{jt}}{P_t^Y Y_t}]$ in the model (the steady state) and the data (averaged over the entire sample 1948-2018).

productivity because of decreasing returns to scale in labor (which implies that the aggregate capital share $\alpha^* > 0$). Hence, shocks which increase weighted employment $\alpha^* d \log L_t$ by more than the sector's steady state Domar weight, which determines the response of aggregate TFP, will decrease labor productivity.

Figure F.4 shows that shocks to nearly all of the investment hubs and their intermediate suppliers decrease labor productivity. The figure plots the cyclicality of aggregate labor productivity in response to 1% sector-specific shocks to each sector in isolation. Shocks to investment hubs and their suppliers generally increase aggregate employment substantially more than their sectors' Domar weights to decrease labor productivity. The exceptions are professional/technical services, wholesale trade, and transportation & warehousing in the right of the figure. While these sectors have sizeable effects on employment, they are also well-connected in the intermediates network and therefore also have large Domar weights.

Figure F.4: Cyclicality of Labor Productivity Due to Sectoral Shocks

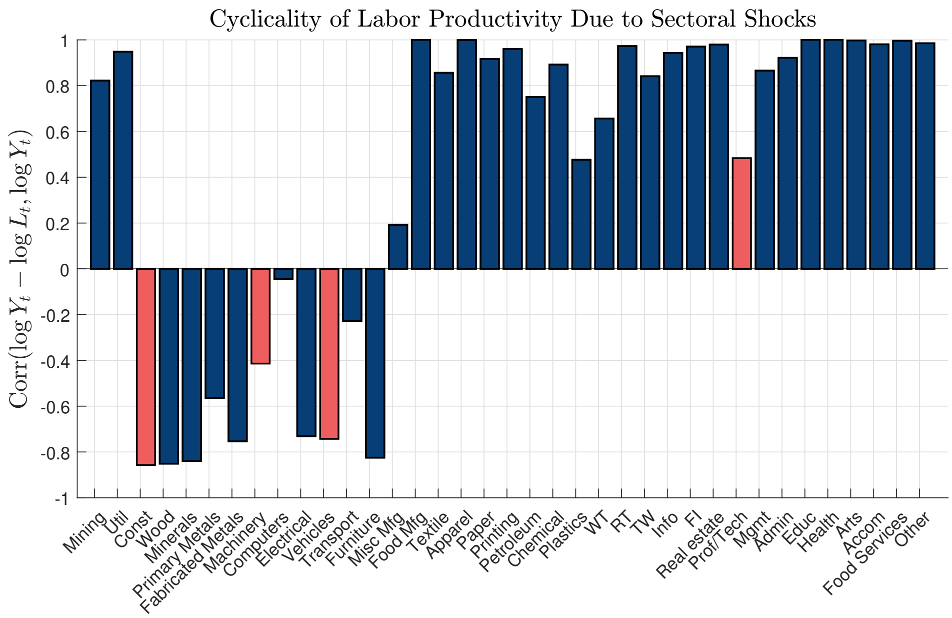

Notes: cyclicality of labor productivity, $\mathbb{C}orr(\log Y_t - \log L_t, \log Y_t)$ in response to sector-specific shocks A_{it}). For each sector, we simulate the model with $\sigma(\log A_{it}) = 1\%$ shocks to that sector only. Investment hubs are highlighted in red.

Numerical Exploration of Sectoral Investment Response to Shocks We now provide numerical comparative statics to understand the mapping from sectoral shocks to the household's valuation of aggregate investment, $\frac{p_t^I I_t}{C_t} = \frac{\sum_j p_{jt}^I I_{jt}}{C_t}$ (Proposition 2 shows how employment responds to changes in the household's valuation of investment). Figure F.5 plots the elasticity of the household's valuation of aggregate investment in response to sector-specific shocks to each sector. The blue bars show that this elasticity is very similar to the elasticity of aggregate employment in response to the shocks plotted in Figure 5 in the main text, consistent with Proposition 2.[58]

Figure F.5 also shows that the distribution of these elasticities across sectors is primarily determined by the Leontief-adjusted investment network. In particular, the grey bars in

[58]Given the result of Proposition 2, we can write changes in aggregate employment as $dL_t = \sum_j dL_{jt} = \sum_j \sum_m \omega_{jm} d\left(\frac{p_{mt}^I I_{mt}}{C_t}\right) = \sum_m d\left(\frac{p_{mt}^I I_{mt}}{C_t}\right) \sum_j \omega_{jm}$. The result that the numerical response of aggregate employment is proportional to the numerical response of the aggregate household's valuation of investment, $dL_t = \phi d\frac{p_t^I I_t}{C_t}$ will obtain if the sum over the rows of each column of ω_{jm} is the same across sectors. This is exactly true in the case where there are no intermediate goods, since in that case, the Leontief-adjusted network is equal to the investment network (whose rows sum to 1 by construction).

Figure F.5: Elasticity of Household's Valuation of Aggregate Investment to Sectoral Shocks

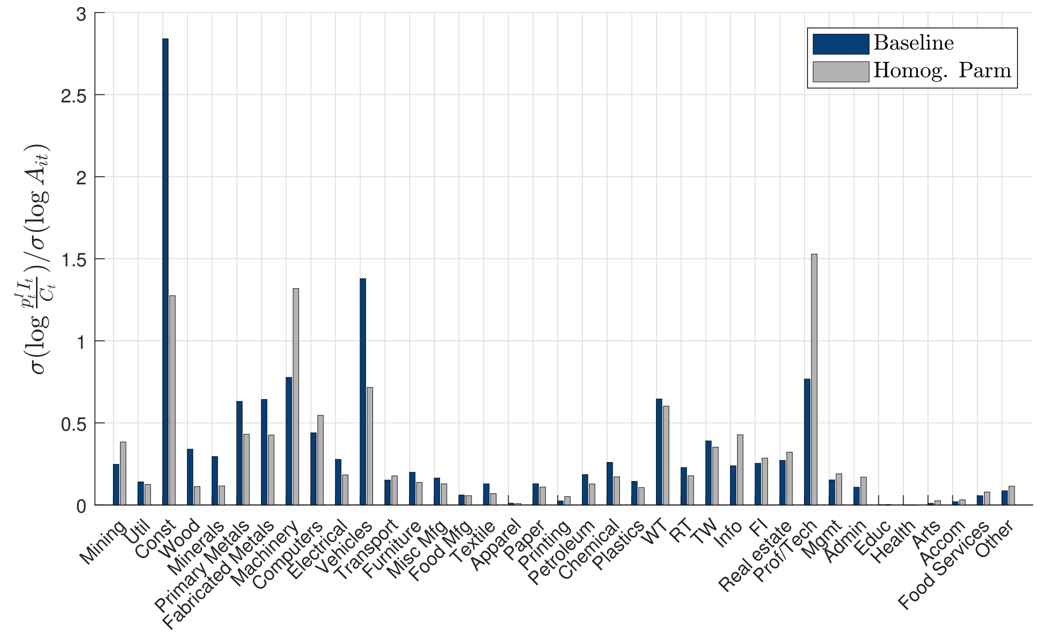

Notes: reduced-form elasticities of the household's valuation of aggregate investment, $\frac{p_t^I I_t}{C_t}$ to sector-specific shocks A_{it}. For each sector, we simulate the model with $\sigma(\epsilon_{it}) = 1\%$ shocks to that sector only. The bars plot the volatility of household's valuation of aggregate investment $\sigma(\log \frac{p_t^I I_t}{C_t})$ divided by the volatility of sector-specific TFP $\sigma(\log A_{it})$. The grey bars show this elasticity where all non-network parameters are set to the mean across sectors.

Figure F.5 plot the elasticities in which all these other parameters are set equal to the average value across sectors.[59] In this case, variation in the elasticities is solely determined by heterogeneity in the Leontief-adjusted investment network. The blue and grey bars are fairly similar, consistent with the idea that heterogeneity in the Leontief-adjusted investment network is the primary source of differences across sectors in our full model. The main exception is the effect of a shock to construction, which is also shaped by the low depreciation rate of residential structures and the abnormally high capital share in real estate (detailed results available upon request).

While that exercise shows that the Leontief-adjusted investment network is a key determinant of these elasticities starting from our calibrated parameters, we now show that it is also the key determinant for a large region of the parameter space. We explore the parameter space by sampling 10,000 alternative parameterizations of the model in which we

[59]The parameters are depreciation rates (δ_j), capital shares (α_j), value added shares of gross output (θ_j), the persistence of TFP shocks (ρ_j), and consumption shares (ξ_j).

hold the Leontief-adjusted investment network fixed but draw the other parameters from the following distributions (all independent of each other):[60]

- Capital shares $\alpha_j \sim$ uniform$[0.08, 0.93]$
- Consumption shares $\xi_j \sim$ uniform$[0, 1]$, rescaled to ensure $\sum_{j=1}^{N} \xi_j = 1$
- Depreciation rates $\delta_j \sim$ uniform$[0.03, 0.15]$
- Persistence parameters $\rho_j \sim$ uniform$[0.30, 0.91]$.

The support of these uniform distributions are determined by the lowest and highest values of each parameter in our baseline calibration of the model. Hence, these 10,000 draws represent an exhaustive exploration of the parameter space (although that almost all of these parameterizations are not representative of the data).

For each of these 10,000 parameterizations, we compute the following two statistics which summarize the role of the Leontief-adjusted investment network in determining the elasticity of the investment to consumption ratio to sector-specific shocks:

- $\mathbb{C}orr(\frac{d \log P_t^I I_t / C_t}{d \log A_{it}}, \sum_{j=1}^{N} \omega_{ij})$: the correlation of the elasticity of the aggregate investment-to-consumption ratio with respect to a given sector's shock and that sector's weighted outdegree in the Leontief-adjusted investment network (which captures its importance in supplying investment goods). If the Leontief-adjusted investment network is the only factor determining the elasticities, then this correlation would be close to 1.[61]
- $\mathbb{C}orr(\frac{d \log p_{jt}^I}{d \log A_{it}}, \omega_{ij})$: the correlation of the passthrough of one sector's shock to another sector's price index of investment with the pair's entry in the Leontief-adjusted investment network. Proposition 3 shows this correlation should be 1 if primary input prices are held fixed, but those prices may fluctuate over time in general equilibrium, driving the correlation below 1.

[60]We do not re-sample the parameters of the intermediates network because those parameters also determine the Leontief-adjusted investment network. We also do not need to sample the covariance matrix of shocks because, at first order, the covariance matrix is irrelevant for the elasticity of the investment to consumption ratio with respect to sectoral TFPs due to certainty equivalence.

[61]In particular, the correlation is 0.99 in the version of the model in which all parameters are homogeneous across sectors except for the Leontief-adjusted investment network.

FIGURE F.6: Role of Investment Network in Driving Numerical Results

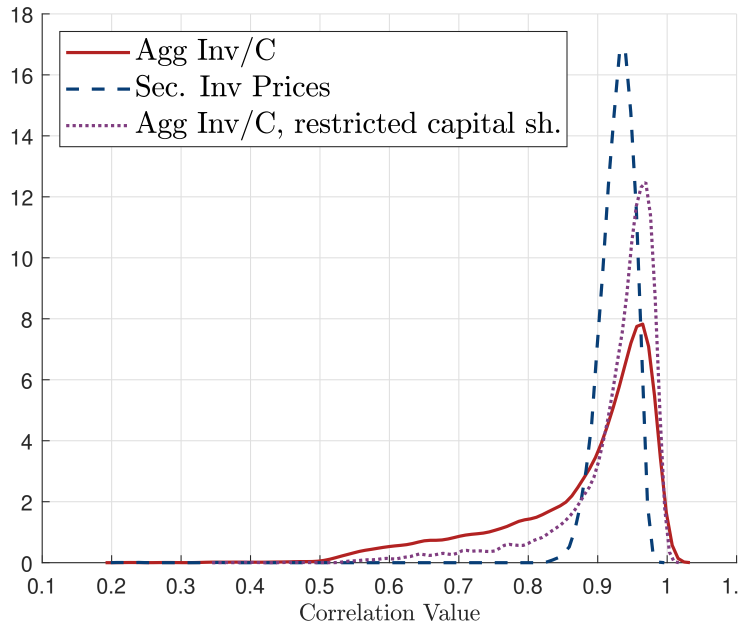

Notes: kernel density estimates from computing correlation statistics from 10,000 calibrations of our model where in each calibration, α_j, ξ_j, δ_j, and ρ_j are drawn from uniform distributions as described in the text above. The red solid line represents the kernel density of the correlation between each sector's weighted outdegree in the Leontief-adjusted Investment Network and the numerical elasticity of the aggregate investment to consumption ratio in response to a shock to that sector, $\mathbb{C}orr(\frac{d\log P^I_t I_t/C_t}{d\log A_{it}}, \sum_{j=1}^N \omega_{ij})$. The blue dashed line presents the kernel density estimate of the correlation between each the elasticity of each sector's investment price with respect to each sector's shock and that pair's entry in the Leontief-adjusted investment network. $\mathbb{C}orr(\frac{d\log p^I_{jt}}{d\log A_{it}}, \omega_{ij})$. The purple dotted line represents the kernel density of $\mathbb{C}orr(\frac{d\log P^I_t I_t/C_t}{d\log A_{it}}, \sum_{j=1}^N \omega_{ij})$ using tighter bounds on the support of possible values of α_j to eliminate wildly counterfactual outliers. Kernel density estimates are constructed using 100 equally spaced points and normal kernel.

Figure F.6 plots the kernel density of these statistics across the 10,000 draws of parameter values. The dashed blue line shows that the price pass-through correlation $\mathbb{C}orr(\frac{d\log p^I_{jt}}{d\log A_{it}}, \omega_{ij})$ is above 0.8 in 99.97% of the parameterizations and above 0.9 in 87.7% of the parameterizations. This result indicates that the general equilibrium effects which may break Proposition 3 are relatively unimportant, validating our interpretation of shocks to the investment hubs and their key suppliers as aggregate investment supply shocks is valid.

The solid red line in Figure F.6 shows that $\mathbb{C}orr(\frac{d\log P^I_t I_t/C_t}{d\log A_{it}}, \sum_{j=1}^N \omega_{ij})$ is also high across the majority of simulations, indicating that the Leontief-adjusted investment network is the primary determinant of the elasticity of the aggregate investment-to-consumption ratio in response to shocks. However, the relationship is less tight than for investment prices; 78.6%

of parameterizations have a correlation above 0.8 and 57.6% have a correlation above 0.9. The dashed purple line shows that those lower correlations are primarily determined by extreme and unrealistic values of the consumption shares α_j; if we constrained those draws to be within the more reasonable range of $\alpha_j \sim \text{uniform}[0.10, 0.67]$, then the distribution of correlation statistics become more tightly bunched around 0.9: 92.3% of parameterization have a correlation greater than 0.8 and 77.6% of parameterizations have a correlation greater than 0.9.

G Additional Results on Changing Business Cycles

We now provide several additional results referenced in Section 5 of the main text.

G.1 Investment Production Frictions

In this subsection, we provide details about the investment production frictions from Section 5 impact the equilibrium conditions of our model and then show that our results are robust to varying the strength of these frictions.

Equilibrium conditions The investment production frictions change the output market clearing condition to be:

$$Q_{jt} = C_{jt} + \sum_{i=1}^{N} M_{ijt} + \left(\sum_{i=1}^{N} I_{ijt}^{-\rho}\right)^{-\frac{1}{\rho}} \tag{42}$$

Define the total production of investment goods by sector j as $Z_{jt} = \left(\sum_{i=1}^{N} I_{jit}^{-\rho}\right)^{-\frac{1}{\rho}}$. Then the intratemporal investment allocation decision becomes:

$$p_{it}\left(\frac{Z_{jt}}{I_{ijt}}\right)^{1+\rho} = \lambda_{ij} p_{jt}^{I} I_{jt} \tag{43}$$

The corresponding cost-minimization problem implies that the price index of a new unit of investment for sector j is now:

$$p^I_{jt} = \prod_{i=1}^{N} \left(\frac{p_{it}}{\lambda_{ij}}\right)^{\lambda_{ij}} \prod_{i=1}^{N} \left(\frac{Z_{it}}{I_{jt}}\right)^{(1+\rho)\lambda_{ij}} \tag{44}$$

Therefore, the price of purchasing an investment good is now specific to the producer-purchaser pair; an increase in investment demand from a given sector will put upward pressure on its price index for investment goods, dampening fluctuations in investment.

Importantly, this extension of the model does not change the results in any of the propositions presented in Section 4. Proposition 1 only relies on the definition of value added, which is unaffected by this friction. Proposition 2 does rely on the resource constraint, which has now been modified, but that modification does not change those results; to see this fact, note that we can solve for Z_{jt} using equation (44) above:

$$\begin{aligned} Z_{jt} &= \left(\sum_{i=1}^{N} I_{jit}^{-\rho}\right)^{-\frac{1}{\rho}} \\ &= \left(\sum_{i=1}^{N} \frac{\lambda_{ji} p^I_{it} I_{it}}{p_{jt} Z_{jt}^{1+\rho}}\right)^{-\frac{1}{\rho}} \\ &= \sum_{i=1}^{N} \frac{\lambda_{ji} p^I_{it} I_{it}}{p_{jt}} \end{aligned}$$

This result, together with equation (17), implies that we can still write the resource constraint as in equation (36) in the proof for Proposition 2. Essentially, because investment expenditures by each sector remain Cobb-Douglas over each intermediate investment good and markets are competitive, the expenditures on each intermediate investment good remain proportional to total expenditure.

Finally, the result in Proposition 3 is also unchanged as long as the conditions for isolating the direct effect of TFP shocks on investment prices are extended to include holding fixed investment production and expenditures.

Table G.1
Robustness with Respect to Investment Production Frictions

	Baseline		**No Frictions**		**Large Frictions**	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(\Delta y_t)$	3.95%	2.42%	3.97%	2.64%	3.86%	2.38%
$\rho(\Delta y_t - \Delta l_t, \Delta y_t)$	0.52	-0.01	0.38	-0.31	0.57	0.10
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.90	1.03	0.93	1.12	0.88	1.00
$\sigma(\Delta i_t)/\sigma(\Delta y_t)$	3.78	4.11	5.63	9.22	3.74	4.16

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2018). y_t is log real GDP, l_t is log aggregate employment, i_t is log real aggregate investment, and Δ denotes the first difference operator. "Baseline" refers to the baseline model described in the main text, which uses $\rho = -1.04$. "No Frictions" refers to the model without reallocation frictions, i.e. $\rho = -1$. "Large Frictions" refers to the model with $\rho = -1.5$.

Robustness to varying ρ In our baseline results, we set the parameter $\rho = -1.04$ to match movements in the distribution of investment expenditures across sectors. Table G.1 that without these frictions (setting $\rho = -1$), investment is more volatile than in the baseline model, especially in the post-1984 sample. This excess volatility in turn implies higher volatility of employment by (12), so the cyclicality of labor productivity falls by nearly 0.7 and becomes countercyclical in the post-1984 sample. On the other hand, Table G.1 shows that increasing ρ to -1.5 does not materially impact our primary findings. These results indicate that while breaking the perfect substitutability matters for our results, the precise degree of imperfect substitutability does not.

G.2 Time Series Fit of the Model

We now compare the model's implied time series of real GDP, aggregate employment, aggregate investment, and aggregate consumption to the data. Recall that no features of these series were targeted in our calibration; instead, we simply feed in the realized series of sector-level TFP shocks and let the model endogenously produce these macroeconomic outcomes.

Figure G.1 plots the first-differenced series and Figure G.2 plots the HP-filtered ones. In both cases, the average correlation between the model's and data's time series is approximately 0.5.[62] Importantly, aggregate consumption comoves with the business cycle, which

[62] The weakest correlation between model and data is in employment, although this largely seems to be due

Figure G.1: Aggregate Time Series in Model and Data: First Differences

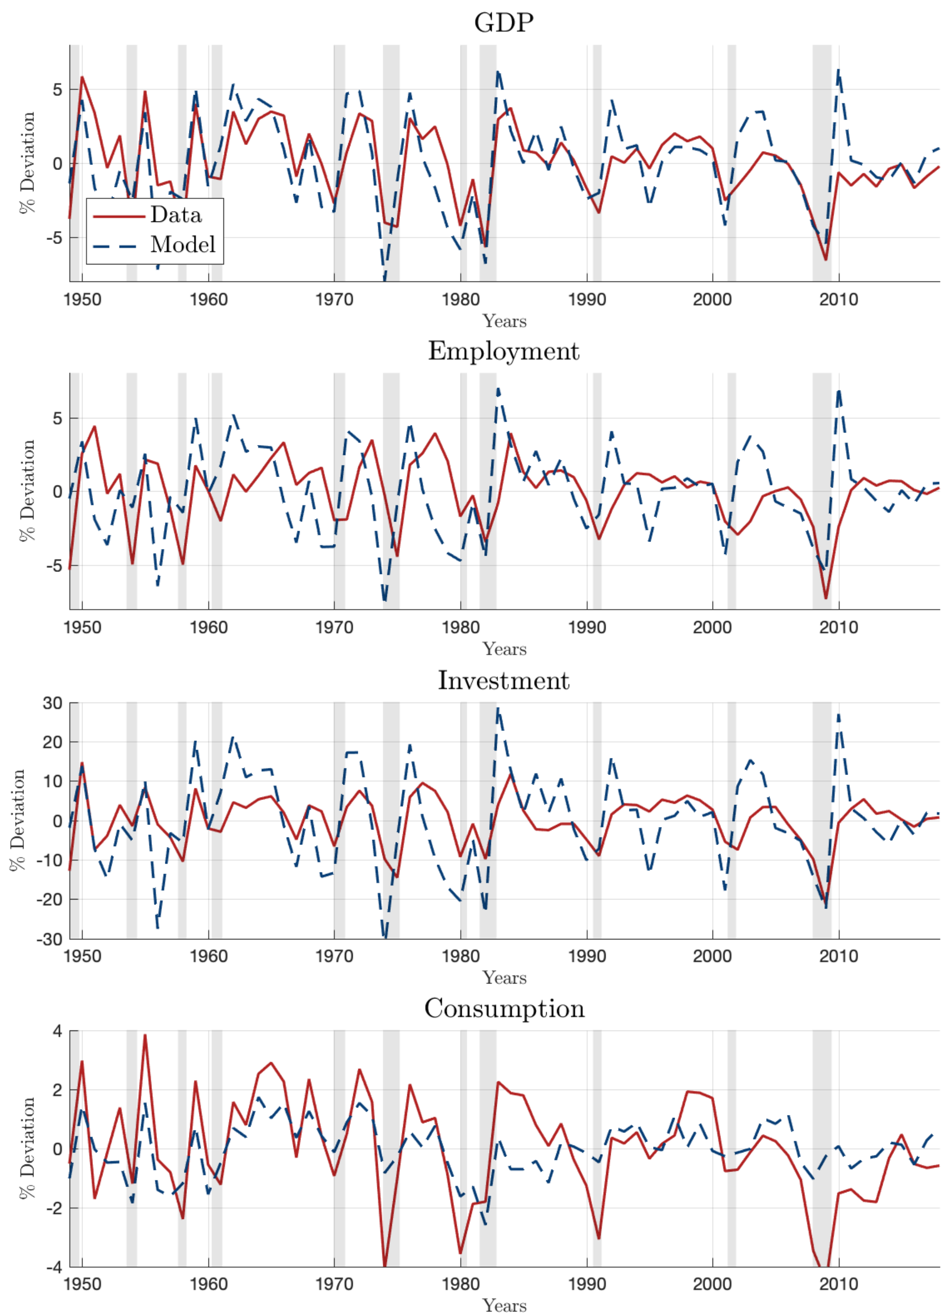

Notes: time series of aggregate GDP, employment, investment, and consumption in the model and the data. Each series has been logged and first-differenced.

Figure G.2: Aggregate Time Series in Model and Data: HP Filter

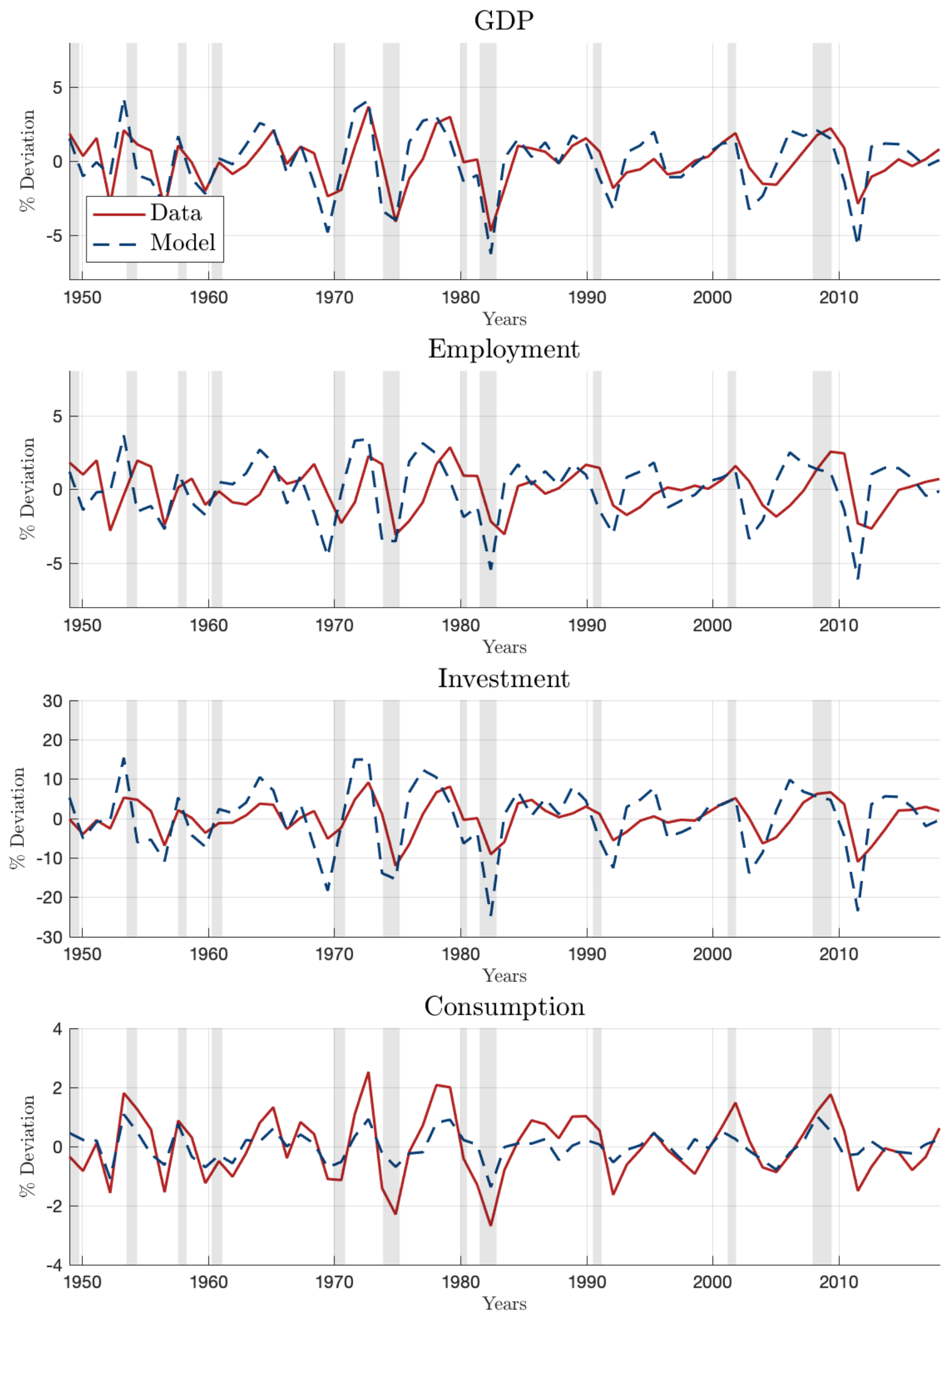

Notes: time series of aggregate GDP, employment, investment, and consumption in the model and the data. Each series has been logged and HP filtered with smoothing parameter $\lambda = 6.25$. To avoid endpoint bias from the HP filter, we omit the first and last three years of data of the entire sample.

is a challenge in models driven by shocks to investment supply; as we discuss in Appendix F, our model generates comovement through the intermediates network. The model also matches the severity of the 1982 and 2008 recessions, which are typically challenging to explain with supply shocks.

All that said, there are important differences between the model and data, indicating scope for additional shocks, nominal rigidities, or nominal rigidities to improve model fit. First, given the complete markets structure, consumption is much smoother than in the data (as in the one-sector RBC model – see King and Rebelo (1999)). Second, investment is more volatile in our model than in the data, again as in the one-sector RBC model. Third, the first-differenced model series predict robust recoveries following the post-1980s recessions which manifest as higher-than-average growth rates following these recessions. These high growth rates did not materialize in the data partly because the average growth rate fell over this period, which is outside our model. The HP filter eliminates these trend changes in the growth rate, bringing the model closer to the data.

G.3 Structural Change

Our baseline analysis focuses on changes in the sectoral shock process by holding the parameters of the economy fixed over time. However, there have been substantial trend changes in many of these parameters over time, such as structural transformation from manufacturing to services (impacting the distribution of consumption expenditures and production networks), the rise of intellectual property products (generating a higher depreciation rate), and the decline in labor share. While a full analysis of the impact of these structural changes on business cycle fluctuations is beyond the scope of this paper, we present two complementary exercises to show that the main outcomes of interest in our analysis are robust to accounting for structural change in the parameters of our model.

Transition Path Our first exercise allows for the structural parameters to change smoothly over time along a perfect foresight transition path. In particular, we assume that, starting

to a timing difference in the model and the data; if the model time series is shifted one time period forward, the correlation between model and data is much higher. This timing reflects the fact that employment slightly lags GDP in the data (this lag is smaller in quarterly data but magnified in annual data).

in 1948, agents become aware of the trend path of all structural parameters of the economy over the 1948-2018 period.[63] We further assume that these trends continue through 2043 and then gradually converge to their new steady state by 2068.[64]

We solve for the equilibrium over this path using a variant of the solution algorithm developed in Maliar et al. (2020). This algorithm assumes that, while agents have perfect foresight over the changes in the parameters of the economy, there is still uncertainty over the realization of TFP shocks each period. We first solve for policy functions for log capital at T, when parameter changes have ceased and the economy is stationary. We then iterate backward, solving for the policy functions in $T-1$ taking the policy function in period T as given. We iterate over this procedure until we have policy functions for the entire sample (vom Lehn (2020) implements this algorithm in a similar way). We assume that the initial condition of the economy is the steady state corresponding to the parameter values observed in the year 1948.

We use a Smolyak grid of points to solve for the decision rules. We limit ourselves to a first-order Smolyak grid and approximate the policy function for the log of capital as linear in the state variables. For tractability, and given that our policy functions for capital are log-linear, we assume that certainty equivalence holds and evaluate expectations with a first-order quadrature. We solve for the capital accumulation policy functions in each period and then feed in the time series of measured TFP shocks used in our baseline analysis.[65]

Table G.2 shows that our main results continue to hold along this transition path: the

[63]Specifically, the set of parameters which we allow to change over time are: the investment network (λ_{ij}), the intermediates network (γ_{ij}), depreciation rates (δ_j), capital shares (α_j), the share of primary inputs in production (θ_j), and the consumption shares (ξ_j). We identify the trends in parameter values using a fourth-order polynomial, consistent with our approach to detrending TFP in Section 5. Since the consumption share, investment network, and intermediate network parameters must sum to 1, we do not compute trends for those parameters directly. Instead, we first compute trends in the levels of consumption expenditures, intermediates expenditures, and investment expenditures, and then compute expenditure shares based on those trends.

[64]We project forward these trends conservatively, on the basis of linear trends for the moving averages of parameters for the last 5, 10, 15, or 20 years of data, selecting which yields the smallest trend growth in absolute value. We do this to minimize the likelihood of extreme trends following the last year of observed data.

[65]We set the parameter governing the investment production frictions to $\rho = -1.3$ because the changes in parameter values increase the volatility of investment. Our approximated decision rules imply negative investment in 1% of observations, which is inconsistent with our investment production frictions. In these cases, we set investment to 10% of the depreciated capital stock in that period; our results are robust to varying this boundary value.

Table G.2
Allowing for Structural Change via Transition Path

	Baseline Model		Structural Change	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(\Delta y_t)$	3.95%	2.42%	4.57%	2.10%
$\rho(\Delta y_t - \Delta l_t, \Delta y_t)$	0.52	-0.01	0.50	0.05
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.90	1.03	0.89	1.02
$\sigma(\Delta i_t)/\sigma(\Delta y_t)$	3.78	4.11	5.43	4.66

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2018). y_t is log real GDP, l_t is log aggregate employment, i_t is log real aggregate investment, and Δ denotes the first difference operator. "Baseline" corresponds to the model described in the main text. "Structural change" corresponds to the results from the transition path exercise.

cyclicality of labor productivity falls by 0.45 (compared to 0.53 in the baseline) and the relative volatility of employment rises by 0.13 (the same as in the baseline). The main difference from our baseline result is that the relative volatility of investment is now higher, reflecting the fact that our forward-looking agents change their investment decisions in response to changes in the path of structural parameters (as well as the simple fact that nonlinearities in the solution method also increase volatility).

Simulation Exercises While the previous exercise allowed for smooth changes in structural parameters over time, it relied on strong assumptions regarding how firms adjust to these parameter changes, including an unrealistic degree of foresight on the part of agents. Our second exercise sidesteps these issues by simply simulating the model separately for parameterizations corresponding to the pre- and post-1984 period. In particular, instead of feeding in the realized time series of sectoral TFP shocks as in the main text, we estimate the covariance matrix of these shocks separately for the pre vs. post 1984 subsamples and compute population moments from those two estimates. The advantage of this approach is that it avoids assumptions on the particular path of parameters over time. However, the disadvantages are that it instead assumes that parameters change once and for all, that agents understand that abrupt change, and that there are no meaningful transitional dynamics between the two sets of parameters.

An additional disadvantage of this exercise is that we cannot estimate a full-rank co-

TABLE G.3
ALLOWING FOR STRUCTURAL CHANGE VIA SIMULATION

	Baseline (30 Sectors)		**Simulation**		**Structural Change**	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(\Delta y_t)$	3.21%	1.94%	3.63%	2.13%	4.07%	2.04%
$\rho(\Delta y_t - \Delta l_t, \Delta y_t)$	0.75	0.32	0.76	0.42	0.80	0.43
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.79	0.95	0.83	0.92	0.83	0.92
$\sigma(\Delta i_t)/\sigma(\Delta y_t)$	3.32	3.70	3.54	3.79	3.76	3.71

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2018). y_t is log real GDP, l_t is log aggregate employment, i_t is log real aggregate investment, and Δ denotes the first difference operator. "Baseline (30 Sectors)" corresponds to the exercise where shocks are measured from the data and then fed into model described in the main text, albeit with 30 sectors (non-durable manufacturing is collapsed into a single sector). "Simulation" corresponds to the simulation exercises based on estimated covariance matrices for 30 sectors. "Structural change" corresponds to the simulation exercises, where model parameters are estimated separately for the pre-1984 and post-1984 period.

variance matrix with 37 sectors and less than 37 years of data both pre- and post-1984 subsamples. Therefore, following the same procedure described for the principal components analysis in Appendix D, we collapse our data to 30 sectors by aggregating all non-durable manufacturing sectors into a single sector. We show in the left panel of Table G.3 that the changes in business cycles observed in the model with 37 sectors are still observed with this coarser disaggregation of the economy.[66]

We estimate the covariance matrix of innovations to TFP separately for each subsample using the sample covariance matrix and then separately simulate the model for 10,000 periods under each of the estimated covariance matrices, discarding the first 100 periods in each case. The middle panel of Table G.3 shows that, if we hold all the structural parameters fixed over time, this simulation approach generates similar changes in aggregate business cycle patterns to feeding in the realized series. The right panel of Table G.3 shows that our key outcomes of interest do not change very much relative to the simulation benchmark when we allow

[66]The relative volatility of employment is lower than in the baseline model because we aggregate some relatively important suppliers of investment hubs (e.g. petroleum manufacturing and chemicals manufacturing) with other sectors that produce primarily consumption goods (e.g. food/beverage manufacturing, apparel manufacturing); in this case, a shock to one of the relatively important suppliers also increases consumption and therefore generates an income effect on labor supply which is absent in the baseline model. The fact that employment is less volatile implies that labor productivity is more procyclical than in the baseline model as well (see Footnote 31 for the precise relationship between the relative volatility of employment and the cyclicality of labor productivity).

for structural change.[67] The main exceptions are that the model no longer generates an increase in the volatility of investment over time and implies a somewhat larger decline in the volatility of GDP.

G.4 Non-Cobb Douglas Production and Preferences

While our baseline analysis imposed Cobb-Douglas production and utility functions for analytical tractability, we now show numerically that our results are robust to allowing for constant elasticity of substitution (CES) functional forms. Specifically, we generalize the production function to become

$$Q_{jt} = \left[\theta_j^{\frac{1}{\sigma_y}} Y_{jt}^{\frac{\sigma_y - 1}{\sigma_y}} + (1-\theta_j)^{\frac{1}{\sigma_y}} M_{jt}^{\frac{\sigma_y - 1}{\sigma_y}}\right]^{\frac{\sigma_y}{\sigma_y - 1}} \tag{45}$$

where

$$Y_{jt} = A_{jt}\left[\alpha_j^{\frac{1}{\sigma_k}} K_{jt}^{\frac{\sigma_k - 1}{\sigma_k}} + (1-\alpha_j)^{\frac{1}{\sigma_k}} L_{jt}^{\frac{\sigma_k - 1}{\sigma_k}}\right]^{\frac{\sigma_k}{\sigma_k - 1}} \tag{46}$$

and

$$M_{jt} = \left(\sum_{i=1}^{N} \gamma_{ij}^{\frac{1}{\sigma_m}} M_{jt}^{\frac{\sigma_m - 1}{\sigma_m}}\right)^{\frac{\sigma_m}{\sigma_m - 1}}. \tag{47}$$

We assume that productivity shocks affect the primary inputs because, as shown in Sato (1976), there would otherwise not exist a unique function for real value added. Therefore, in these exercises, we feed in productivity measured as value added net of primary inputs (rather than measured as gross output net of all inputs as in the main text). We also generalize the consumption aggregate which enters utility to be:

$$C_t = \left(\sum_{j=1}^{N} \xi_j^{\frac{1}{\sigma_c}} C_{jt}^{\frac{\sigma_c - 1}{\sigma_c}}\right)^{\frac{\sigma_c}{\sigma_c - 1}}. \tag{48}$$

We choose values for the elasticities of substitution from Oberfield and Raval (2021) and Atalay (2017). We set the elasticity of substitution between consumption goods to $\sigma_c = 0.75$,

[67] We allow the same set of parameters as in the transition path exercise to change over time, plus the persistence of TFP shocks ρ_j. We compute the average value of these parameters separately for the pre vs. post 1984 subsamples and compute simulated moments given the covariance matrix of shocks estimated as above.

which is on the low end of the range of values considered in Oberfield and Raval (2021) (0.75-1.15).[68] We set the elasticity between intermediate inputs to Atalay (2017)'s preferred value $\sigma_m = 0.1$. We set the elasticity between primary inputs and intermediates to the midpoint of the range of estimates in Oberfield and Raval (2021) (0.6-1), i.e., $\sigma_y = 0.8$. Finally, we set the elasticity between capital and labor to Oberfield and Raval (2021)'s midrange of $\sigma_k = 0.6$.[69] Given these parameter values, we then re-calibrate the share parameters in the production function in order to match the expenditure shares in the model's steady state to the data.

Table G.4 reports a number of results using these alternative functional forms. First, for the sake of comparability, the top left panel shows the results from our baseline Cobb-Douglas model are very similar when we measure productivity as value added net of primary inputs (which we must do in the CES case given (46)).[70] Second, the top middle panel shows that allowing for the CES production and utility functions barely affect the changes in business cycle statistics over time; for example, the cyclicality of labor productivity declines by 0.64 with CES functional forms compared to 0.61 with Cobb-Douglas. However, the overall level of employment and GDP volatility is higher with the CES functional forms, consistent with the idea that complementarity amplifies overall volatility.

The next four panels of Table G.4 decompose the role of each elasticity of substitution in isolation, and show that the higher volatility of the CES model is driven by the complementarity between capital and labor. This finding indicates that, in the CES model, investment fluctuations have a large impact on labor demand, which mirrors our main result in the Cobb-Douglas model that they have a large impact on labor supply.

Finally, the bottom panels of Table G.4 investigate the role of nonlinearities by computing a second-order approximation of the model.[71] Baqaee and Farhi (2019) show how a second-

[68] We choose the low end of this range because Oberfield and Raval (2021) looks at finely disaggregated manufacturing industries, which have greater similarity, and thus potentially a higher degree of substitutability, than the 37 sectors we consider covering the entire private non-farm economy.

[69] We have also tried using Karabarbounis and Neiman (2014)'s estimate $\sigma_k = 1.25$ and found that this higher elasticity does not substantially impact our results (available upon request).

[70] Of course, the two notions of productivity are theoretically isomorphic under Cobb-Douglas production: $\widetilde{A}_{jt} = A_{jt}^{\frac{1}{\theta_j}}$ where $\widetilde{A}_{jt}$ is TFP measured as value added net of primary inputs and A_{jt} is measured as gross output net of all inputs. However, this relationship may not hold in the data if production is not Cobb-Douglas or there is measurement error.

[71] We need to specify the covariance matrix of TFP shocks in order to solve for the decision rules because certainty equivalence does not hold in a second-order approximation. We use the sample covariance matrix for our measured innovations to TFP for the entire period 1948-2018.

Table G.4
Allowing for Non-Cobb Douglas Functional Forms

	CD		**All CES**		**σ_c only**	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(\Delta y_t)$	3.89%	2.79%	4.31%	2.94%	3.87%	2.79%
$\rho(\Delta y_t - \Delta l_t, \Delta y_t)$	0.52	-0.15	0.30	-0.38	0.54	-0.14
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.90	1.05	0.96	1.08	0.89	1.04
$\sigma(\Delta i_t)/\sigma(\Delta y_t)$	3.78	4.09	3.73	4.17	3.79	4.11
	σ_k only		**σ_y only**		**σ_m only**	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(\Delta y_t)$	4.43%	2.98%	3.85%	2.76%	3.86%	2.79%
$\rho(\Delta y_t - \Delta l_t, \Delta y_t)$	0.24	-0.35	0.50	-0.20	0.55	-0.12
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.97	1.08	0.90	1.06	0.89	1.04
$\sigma(\Delta i_t)/\sigma(\Delta y_t)$	3.71	4.01	3.81	4.13	3.77	4.14
	CD (2nd order)		**All CES (2nd order)**		**Ident Inv., CES, 2nd**	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(\Delta y_t)$	3.89%	2.79%	4.38%	3.07%	3.71%	2.23%
$\rho(\Delta y_t - \Delta l_t, \Delta y_t)$	0.53	-0.11	0.29	-0.43	0.57	0.45
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.89	1.04	0.96	1.10	0.92	0.92
$\sigma(\Delta i_t)/\sigma(\Delta y_t)$	3.98	4.45	3.98	4.52	2.78	2.76

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2018). y_t is log real GDP, l_t is log aggregate employment, i_t is log real aggregate investment, and Δ denotes the first difference operator. "CD" corresponds to the baseline model, but instead measuring productivity shocks as value added net of primary inputs rather than gross output net of all inputs. "All CES" corresponds to the model with all functional forms (as described in the text) allowed to be CES. "σ_c only" corresponds to only having a CES nest in consumption aggregation. "σ_k only" corresponds to only having a CES nest in capital and labor. "σ_v only" corresponds to only having a CES nest between value added and intermediates. "σ_m only" corresponds to only having a CES nest in intermediate bundling. "CD (2nd order)" corresponds to solving the model using a 2nd order approximation when using value added based measures of TFP. "All CES (2nd order)" corresponds to solving the model with a second order approximation with all functional forms are CES, as described in the text. "Ident Inv., CES, 2nd" corresponds to solving the model with a second order approximation with all functional forms are CES and where the investment network is set to the identity matrix.

order approximation allows the model to capture rich substitution patterns which exist with CES production functions. However, we find that these nonlinearities do not have a large effect on the changes in aggregate fluctuations on which we focus in this paper. In fact, with an identity investment network, there is almost no change in aggregate fluctuations, as was the case in the first-order Cobb-Douglas specification of the main text.

G.5 Other Robustness Checks

Adding Other Frictions We now show that our results are robust to allowing for frictions to reallocating labor across sectors and to accumulating capital within sectors. The labor reallocation frictions we consider modify the disutility of labor to become $\left(\sum_j L_{jt}^{\frac{\tau+1}{\tau}}\right)^{\frac{\tau}{\tau+1}}$ (as in Horvath (2000)), which implies that workers are imperfect substitutes across sectors. We set the value of $\tau = 4.5$ to match the volatility of employment relative to GDP in the pre-1984 period. The capital adjustment costs modify the capital accumulation equation in each sector to take the following form:

$$K_{jt+1} = (1-\delta_j)K_{jt} + I_{jt} - \frac{\phi}{2}\left(\frac{I_{jt}}{K_{jt}} - \delta_j\right)^2 K_{jt} \tag{49}$$

We calibrate the size of adjustment costs ϕ to match the volatility of investment within sectors using a decomposition for aggregate investment variance like the one for employment in Equation (19), in a model without investment production frictions (i.e. $\rho = -1$). This generates a value of $\phi = 0.5$. We also consider results where we use this value for the adjustment costs and include investment production frictions with $\rho = -1.04$.

Table G.5 shows that including these frictions does not significantly impact our main findings. While both of sets of frictions decrease the relative volatility of employment – and therefore increase the overall cyclicality of labor productivity – the cyclicality still falls over time by as much in the data.

Maintenance As discussed in footnotes 5 and 9, some previous studies using the 1997 BEA capital flows table were forced to make a correction to the investment network in order

Table G.5
Robustness with Respect to Other Frictions

	Baseline		Labor Reallocation		Convex AC only	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(\Delta y_t)$	3.95%	2.42%	3.63%	2.21%	3.62%	2.20%
$\rho(\Delta y_t - \Delta l_t, \Delta y_t)$	0.52	-0.01	0.71	0.33	0.70	0.36
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.90	1.03	0.83	0.95	0.83	0.94
$\sigma(\Delta i_t)/\sigma(\Delta y_t)$	3.78	4.11	3.49	3.81	3.56	3.88
	All Capital Frictions					
	Pre-1984	*Post-1984*				
$\sigma(\Delta y_t)$	3.59%	2.19%				
$\rho(\Delta y_t - \Delta l_t, \Delta y_t)$	0.72	0.40				
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.82	0.93				
$\sigma(\Delta i_t)/\sigma(\Delta y_t)$	3.54	3.90				

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2018). y_t is log real GDP, l_t is log aggregate employment, i_t is log real aggregate investment, and Δ denotes the first difference operator. "Baseline" refers to the baseline model described in the main text. "Labor Reallocation" refers to adding labor reallocation frictions from Horvath (2000). "Convex AC only" refers to adding only quadratic capital adjustment costs without investment production frictions (i.e. setting $\rho = -1$). "All Capital Frictions" corresponds to including both investment production frictions and convex adjustment costs.

Table G.6
Robustness with Respect to Maintenance Investment

	Baseline		12.5% Maintenance	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(\Delta y_t)$	3.95%	2.42%	3.81%	2.30%
$\rho(\Delta y_t - \Delta l_t, \Delta y_t)$	0.52	-0.01	0.57	0.10
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.90	1.03	0.88	1.00
$\sigma(\Delta i_t)/\sigma(\Delta y_t)$	3.78	4.11	3.77	4.06

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2018). y_t is log real GDP, l_t is log aggregate employment, i_t is log real aggregate investment, and Δ denotes the first difference operator. "Baseline" refers to the baseline model described in the main text, which uses $\rho = -1.04$. "12.5% maintenance" adjusts the investment network to allow for an additional 12.5% of investment expenditures to be purchased from within each sector.

to ensure the model is invertible.[72] 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)). However, a key challenge in adjusting for maintenance is that the mix of sectors which produce this maintenance investment is not observable in the data. One extreme assumption is that maintenance is produced by the same mix of sectors as the new investment recorded in our investment network; in this case, the investment network would not change. The opposite extreme assumption is that all maintenance investment is produced using own-sector output. We follow Foerster, Sarte and Watson (2011) and assume that 50% of maintenance investment is produced proportionally to the investment network process and 50% is produced using own-sector resources. Given that McGrattan and Schmitz Jr (1999) identify maintenance expenditures to be, on average, 30% as big as new investment in national accounts (and thus roughly 20-25% of a combination of all new and maintenance investment), we account for maintenance investment by adding a correction to the diagonal amounting to 12.5% of total investment. Table G.6 shows that with this adjustment to the investment network our results continue to hold. The fact that each sector now uses its own output for investment weakens the strength of the investment hubs, but quantitatively, the model still generates a decrease in the correlation of labor productivity and aggregate GDP similar to our baseline results.

Detrending As discussed in the main text, we detrend measured TFP using a log-polynomial trend before feeding it into our model. Table G.7 shows that our main results are robust to using a second-order or fifth-order polynomial trend, rather than a fourth-order one as in the main text.[73]

H Changes in Aggregate Cycles Driven by Changes in Sectoral Comovement

This Appendix contains additional results referenced in Section 6 in the main text.

[72] In numerical simulations we have done, it appears a key reason this correction may be necessary is because TFP shocks are assumed to follow a random walk.

[73] Our results are very similar when using a third-order trend as well; we omit those results for parsimony.

Table G.7

Robustness with Respect to Other Levels of Detrending

	Baseline (4th order)		**2nd order trend**		**5th order trend**	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(\Delta y_t)$	3.95%	2.42%	3.75%	2.30%	3.88%	2.66%
$\rho(\Delta y_t - \Delta l_t, \Delta y_t)$	0.52	-0.01	0.66	0.23	0.47	0.07
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.90	1.03	0.85	0.98	0.92	1.01
$\sigma(\Delta i_t)/\sigma(\Delta y_t)$	3.78	4.11	3.65	3.95	3.87	4.07

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2018). y_t is log real GDP, l_t is log aggregate employment, i_t is log real aggregate investment, and Δ denotes the first difference operator. Different columns present results for different degrees of the polynomial trend that we take out of measured TFP before feeding it into the model. "Baseline (4th order)" refers to the baseline model described in the main text, detrends using a fourth-order polynomial. "2nd order trend" refers to using a quadratic trend and "5th order trend" refers to using a 5th order polynomial.

H.1 Proof of Footnote 31

We first show that the decline in the cyclicality 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 Footnote 31 in the main text). Of course, the definition of the correlation between labor productivity and output is $\mathbb{C}orr(\Delta y_t, \Delta y_t - \Delta l_t) = \frac{\mathbb{C}ov(\Delta y_t, \Delta y_t - \Delta l_t)}{\sigma(\Delta y_t)\sigma(\Delta y_t - \Delta l_t)}$ where y_t denotes logged and GDP and l_t is logged aggregate employment (the proof also holds for logged and HP filtered data). Using the linear properties of covariance and rearranging, we can write this as:

$$\frac{\mathbb{C}ov(\Delta y_t, \Delta y_t - \Delta l_t)}{\sigma(\Delta y_t)\sigma(\Delta y_t - \Delta l_t)} = \frac{\mathbb{C}ov(\Delta y_t, \Delta y_t)}{\sigma(\Delta y_t)\sigma(\Delta y_t - \Delta l_t)} - \frac{\mathbb{C}ov(\Delta y_t, \Delta l_t)}{\sigma(\Delta y_t)\sigma(\Delta y_t - \Delta l_t)} =$$
$$\frac{\sigma(y_t)}{\sigma(\Delta y_t - \Delta l_t)} - \frac{\sigma(\Delta l_t)}{\sigma(\Delta y_t - \Delta l_t)}\mathbb{C}orr(\Delta y_t, \Delta l_t) = \frac{\sigma(y_t)}{\sigma(\Delta y_t - \Delta l_t)}\left(1 - \frac{\sigma(\Delta l_t)}{\sigma(\Delta y_t)}\mathbb{C}orr(\Delta y_t, \Delta l_t)\right)$$

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

$$\begin{aligned}
\sigma(\Delta y_t - \Delta l_t) &= \sqrt{\sigma(\Delta y_t)^2 + \sigma(\Delta l_t)^2 - 2\mathbb{C}ov(\Delta y_t, \Delta l_t)} \\
&= \sigma(\Delta y_t)\sqrt{1 + \left(\frac{\sigma(\Delta l_t)}{\sigma(\Delta y_t)}\right)^2 - 2\left(\frac{\sigma(\Delta l_t)}{\sigma(\Delta y_t)}\right)\mathbb{C}orr(\Delta y_t, \Delta l_t)}
\end{aligned}$$

Table H.1
Components of Aggregate Labor Productivity Cyclicality

	Pre-1984	*Post-1984*
$\mathbb{C}orr(\Delta y_t - \Delta l_t, \Delta y_t)$	0.56	0.28
$\mathbb{C}orr(\Delta y_t, \Delta l_t)$	0.80	0.83
$\mathbb{C}orr(\Delta y_t, \Delta l_t)$ only	0.56	0.56
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.83	1.01
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$ only	0.56	0.30

Notes: decomposition of the cyclicality of labor productivity in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2018). y_t is log aggregate value added, l_t is log aggregate employment, and Δ is the first-difference operator. "$\mathbb{C}orr(\Delta y_t, \Delta l_t)$ only" computes the cyclicality of labor productivity from (18) using the actual value of $\mathbb{C}orr(\Delta y_t, \Delta l_t)$ in each subsample but holding fixed $\sigma(\Delta l_t)/\sigma(\Delta y_t)$ at its value in the pre-1984 subsample. "$\sigma(\Delta l_t)/\sigma(\Delta y_t)$ only" computes labor productivity from (18) using the actual value of $\sigma(\Delta l_t)/\sigma(\Delta y_t)$ in each subsample but holding fixed $\mathbb{C}orr(\Delta y_t, \Delta l_t)$ at its value in the pre-1984 subsample.

Combining this expression with the previous one yields:

$$\frac{\sigma(y_t)}{\sigma(\Delta y_t - \Delta l_t)}\left(1 - \frac{\sigma(\Delta l_t)}{\sigma(\Delta y_t)}\mathbb{C}orr(\Delta y_t, \Delta l_t)\right) = \frac{1 - \frac{\sigma(\Delta l_t)}{\sigma(\Delta y_t)}\mathbb{C}orr(\Delta y_t, \Delta l_t)}{\sqrt{1 + \frac{\sigma(\Delta l_t)^2}{\sigma(\Delta y_t)^2} - 2\frac{\sigma(\Delta l_t)}{\sigma(\Delta y_t)}\mathbb{C}orr(\Delta y_t, \Delta l_t)}}$$

which is expression (18) 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 ($\mathbb{C}orr(\Delta y_t, \Delta l_t)$) and the relative standard deviation of employment and GDP ($\frac{\sigma(\Delta l_t)}{\sigma(\Delta y_t)}$).

Table H.1 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 the cyclicality 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.

H.2 Robustness of business cycle moments

We now show that the aggregated and within sector business cycle moments from Table 8 are robust to various choices in the statistical methodology. Table H.2 show that those

Table H.2

Changes in Business Cycles, HP Filter

Data	*Aggregated*		*Within-Sector*	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(y_t)$	2.03%	1.24%	3.27%	2.62%
$\rho(y_t - l_t, y_t)$	0.52	0.14	0.63	0.66
$\sigma(l_t)/\sigma(y_t)$	0.85	1.09	0.83	0.77
Model	*Aggregated*		*Within-Sector*	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(y_t)$	2.52%	1.80%	3.66%	3.22%
$\rho(y_t - l_t, y_t)$	0.53	0.01	0.78	0.82
$\sigma(l_t)/\sigma(y_t)$	0.92	1.01	0.55	0.47

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2018). 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.

Table H.3

Within-Sector Business Cycle Statistics with Different Weights

	Time-Varying (Baseline)		**Fixed Weights**		**Unweighted**	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*	*Pre-1984*	*Post-1984*
$\sigma(\Delta y_t)$	5.42%	4.29%	4.98%	4.62%	6.90%	5.90%
$\rho(\Delta y_t - \Delta l_t, \Delta y_t)$	0.69	0.67	0.68	0.69	0.76	0.76
$\sigma(\Delta l_t)/\sigma(\Delta y_t)$	0.76	0.81	0.78	0.78	0.66	0.63

Notes: business cycle statistics in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2018). 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. 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. In "Unweighted," we construct moments as the simple mean across all sectors.

results, in both the model and the data, continue to hold using the HP filter rather than first differences to detrend the data. Table H.3 shows that the average value of the within sector statistics is similar when using fixed weights or no weights, compared to using time-varying weights (as in the main text).

H.3 Derivation of Decomposition (19)

To derive the decomposition presented in equation (19), we start by decomposing the variance of aggregate employment into within-sector variances and between-sector covariances. We take a first-order Taylor approximation of aggregate employment growth, which yields

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

where ω_{jt}^{l} is the average share of sectoral employment in the aggregate for the time period studied, l_t is log aggregate employment, and l_{jt} is log sector-level employment. 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:

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

We perform a similar decomposition for aggregate GDP, and then we consider the ratio of these two decompositions.[74] This ratio is given by:

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

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

Table H.4
Accuracy of the Decomposition

	Pre-84	*Post-84*
Actual, variance	0.68	1.02
Approximation, variance	0.68	1.04
Actual, standard deviation	0.83	1.01
Approximation, standard deviation	0.83	1.02

Notes: variance and standard deviation of real GDP to aggregate employment. "Actual" refers to the actual values of those statistics in the aggregate data. "Approximation" refers to the right-hand side of the decomposition (19).

This expression can be rewritten as:

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

And then, defining the "variance weight" as $\omega_t = \sum_{j=1}^{N}(\omega_{jt}^{y})^2 \mathbb{V}ar(\Delta y_{jt})/\mathbb{V}ar(\Delta y_t)$, we obtain the final relationship (19) in the main text.

H.4 Additional Quantitative Results

Accuracy of the Decomposition Table H.4 shows that the approximate decomposition (19) is accurate in our data. In particular, the relative variance and the standard deviation of employment implied by the decomposition are close to their actual values in the data.

Changes in Comovement Patterns In the main text, we asserted that the comovement of value added across sectors fell in the post-1980s data but the comovement of employment did not. We now support this assertion by computing the change in the average correlation of value added and employment growth across pairs of sectors:

$$\rho_{\tau}^{x} \equiv \frac{\sum_{i=1}^{N}\sum_{j=i+1}^{N}\omega_i^x\omega_j^x \mathbb{C}orr(\Delta x_{jt}, \Delta x_{jt} | t \in \tau)}{\sum_{i=1}^{N}\sum_{j=i+1}^{N}\omega_i^x\omega_j^x} \tag{50}$$

Table H.5
Average Pairwise Correlations, Model vs. Data

	Data		**Model**	
	Employment	*Value added*	*Employment*	*Value added*
Pre-1984	0.50	0.29	0.98	0.32
Post-1984	0.49	0.17	0.95	0.14
Difference	-0.01	-0.12	-0.03	-0.18

Notes: average pairwise correlations ρ^x_τ in (50). "Pre-1984" computes ρ^x_τ in the 1948-1983 subsample and "post-1984" computes ρ^x_τ in the 1984-2017 subsample. "Data" refers to the data and "Model" to the model.

where x_{jt} is either employment or value added and ω_j are value added or employment shares.

Table H.5 shows that the correlation of value added falls nearly in half, generating most of the decline in the covariances in the decomposition (19); in contrast, the correlation of employment is essentially stable, generating the stability of the between sector covariances as well.[75] To our knowledge, our model is the only explanation for the declining cyclicality of aggregate labor productivity that is consistent with these facts in the data.

Sector Pair Covariance Changes In the main text, we asserted that the changes in covariance patterns are broad-based and not driven by outliers. We illustrate these patterns in Figure H.1, which provides a scatter plot of the change in employment and value added covariances for each sector pair. The covariance of value added declines for most pairs of sectors in the data. Further, while there is substantial heterogeneity in changes in the covariance of employment, these changes are generally of a smaller magnitude than the changes in value added covariance. The figure also shows that these patterns are not driven by outliers but are occurring across many sector pairs.

Model Fit to Covariance Changes across Sector Pairs We now show that the model matches changes in covariance patterns across individual sector pairs. We summarize the sector-pair level change with the "diff-in-diff" $\Delta\mathbb{C}ov(l_{jt}, l_{ot}) - \Delta\mathbb{C}ov(y_{jt}, y_{ot})$. On average,

[75] The 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 $\eta \to \infty$, as described in Footnote 54. However, allowing for a finite Frisch still implies that the correlation of employment across sectors is stable over time (details available upon request).

Figure H.1: Scatterplot of Changes in Sector-Pair Covariances

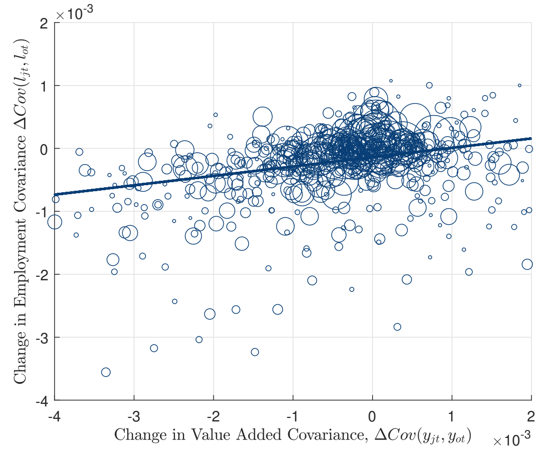

Notes: This figure plots changes in the covariance for each pair of sectors (j, o) in our dataset. The horizontal axis computes the change in the covariance of value added $\mathbb{C}ov(y_{jt}, y_{ot})$ in the post-1984 sample (1984-2018) relative to the pre-1984 sample (1948-1983). Each point is weighted 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. Employment and value added are in log first differences.

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. We plot this diff-in-diff in the data and in the model in Figure H.2. Although neither of these objects were targeted in the calibration, the model explains 50% of the cross-sectional variation in the data.[76]

Decomposition in Finer Disaggregation of Manufacturing Table H.6 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 462 manufacturing sectors from 1958-2011.[77] 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.

[76]The weighted regression line for the data and the model is slightly less steep than the 45-degree line (a regression coefficient of 0.85), indicating that the magnitude of the differences in differences is slightly larger in the model than in the data. However, even the R^2 of the 45-degree line remains high at $R^2 = 0.39$.

[77]There are seven sectors which we omit because they report zero employment at some point in the sample frame.

Figure H.2: Model Fit of Sector-Pair Level $\Delta\mathbb{C}ov(l_{jt}, l_{ot}) - \Delta\mathbb{C}ov(y_{jt}, y_{ot})$ ($R^2 = 50\%$)

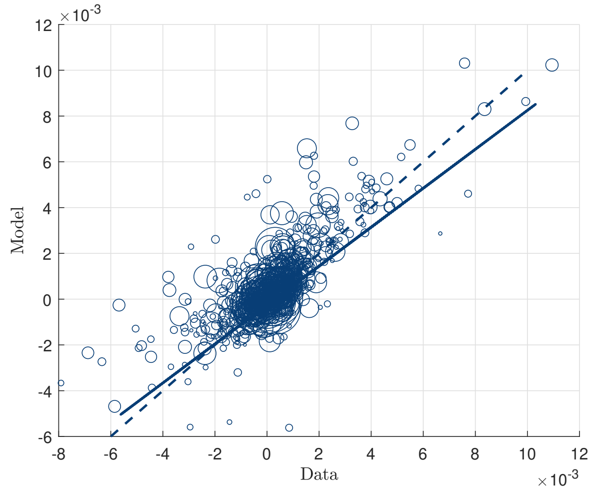

Notes: model fit to sector-pair (j, o) value of $\Delta\mathbb{C}ov(l_{jt}, l_{ot}) - \Delta\mathbb{C}ov(y_{jt}, y_{ot})$, where $\Delta\mathbb{C}ov(l_{jt}, l_{ot})$ is the covariance of log first differenced employment in the post-1984 sample relative to the pre-1984 sample, and $\Delta\mathbb{C}ov(y_{jt}, y_{ot})$ is the covariance of log first differenced 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 0.53. The dashed line is the 45-degree line. In the plot, circle size is proportional to the product of the pair's share of value added over the entire sample.

Table H.6

Decomposition of Relative Employment Volatility, NBER-CES

	Pre-84	*Post-84*	*Contribution of entire term*
$\frac{\mathbb{V}ar(l_t)}{\mathbb{V}ar(y_t)}$	0.37	0.57	100%
Variances	0.33	0.21	1.4%
Covariances	0.37	0.60	98.6%
Variance Weight	0.03	0.06	
($\omega_t = \sum_{j=1}^{N} (\omega^y_{jt})^2 \mathbb{V}ar(y_{jt}) / \mathbb{V}ar(y_t)$)			

Notes: results of the decomposition (19) using NBER-CES data for 462 manufacturing sectors. "Variances" refers to the variance component $\frac{\sum_{j=1}^{N} (\omega^l_{jt})^2 \mathbb{V}ar(l_{jt})}{\sum_{j=1}^{N} (\omega^y_{jt})^2 \mathbb{V}ar(y_{jt})}$. "Covariances" refers to the covariance component $\frac{\sum_{j=1}^{N} \sum_{o \neq j} \omega^l_{jt} \omega^l_{ot} \mathbb{C}ov(l_{jt}, l_{ot})}{\sum_{j=1}^{N} \sum_{o \neq j} \omega^y_{jt} \omega^y_{ot} \mathbb{C}ov(y_{jt}, y_{ot})}$. "Variance weight" refers to the weighting term $\omega_t = \sum_{j=1}^{N} (\omega^y_{jt})^2 \mathbb{V}ar(y_{jt}) / \mathbb{V}ar(y_t)$. "Contribution of entire term" computes the contribution of the first term of the decomposition (19) (in the variances row) and the contribution of the second term (in the covariances row). Real value added is constructed using the gross output price deflator.

Table H.7
Decomposition of Relative Employment Volatility, Equal Weights

	Baseline			Equal Weights		
	Pre-84	*Post-84*	*Contribution of entire term*	*Pre-84*	*Post-84*	*Contribution of entire term*
$\frac{\mathbb{V}ar(l_t)}{\mathbb{V}ar(y_t)}$	0.68	1.04	100%	0.72	0.94	100%
Variances	0.41	0.48	15%	0.44	0.41	11%
Covariances	0.72	1.19	85%	0.76	1.06	89%
Variance Weight	0.12	0.21		0.12	0.19	
($\omega_t = \sum_{j=1}^{N} (\omega_{jt}^y)^2 \mathbb{V}ar(\Delta y_{jt}) / \mathbb{V}ar(\Delta y_t)$)						

Notes: results of the decomposition (19) in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2018). "Baseline" refers to the decomposition from the main text. "Equal weights" sets all the weights $\omega_{jt}^y = \omega_{jt}^l = 1$.

Table H.8
Decomposition of Relative Employment Volatility, HP Filter

	First Differences			HP Filter		
	Pre-84	*Post-84*	*Contribution of entire term*	*Pre-84*	*Post-84*	*Contribution of entire term*
$\frac{\mathbb{V}ar(l_t)}{\mathbb{V}ar(y_t)}$	0.68	1.04	100%	0.72	1.09	100%
Variances	0.41	0.48	15%	0.48	0.49	13%
Covariances	0.72	1.19	85%	0.75	1.25	87%
Variance Weight	0.12	0.21		0.11	0.20	
($\omega_t = \sum_{j=1}^{N} (\omega_{jt}^y)^2 \mathbb{V}ar(\Delta y_{jt}) / \mathbb{V}ar(\Delta y_t)$)						

Notes: results of the decomposition (19) in the pre-1984 sample (1948 - 1983) and post-1984 sample (1984-2017). "First differences" refers to first differencing the data as in the main text. "HP filter" refers to using HP-filtered data. To avoid endpoint bias with the HP filter, we eliminate the first and last three years of the sample.

Equal Weights in the Decomposition Since our decomposition (19) is weighted by sector size, the changes over time may be driven by changes in the distribution of weights rather than changes in comovement patterns. However, H.7 shows that this is not the case; the results are nearly identical if we use constant, equal weights over time.

HP Filter Table H.8 shows that the decomposition results are robust to using the HP filter rather than first-differences.