A Construction of Dataset and Investment Network

A.1 Dataset

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

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

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

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

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

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

A.2 Construction of Investment Network

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

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

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

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

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

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

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

Structures

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

Intellectual Property

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

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

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

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

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

Equipment

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

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

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

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

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

do not have much detail in terms of the margin sectors, and failure to adjust these results in unreasonable differences in the margin allocations compared to the data from 1997 onward. As a result, we take the total reported margins for each commodity and allocate it according to the distribution of margins reported in the years 1997-2001.[60]

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

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

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

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

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

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

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

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

1993-1996).

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

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

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

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

Following this approach, we construct bridge data for equipment investment for all years prior to 1997.[63] We then combine all the bridge files we have constructed to construct the aggregate investment matrix.

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

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

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

Table 16
Skewness of Investment and Intermediates Networks

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

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

A.3 Additional Analysis of Investment Network

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

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

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

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

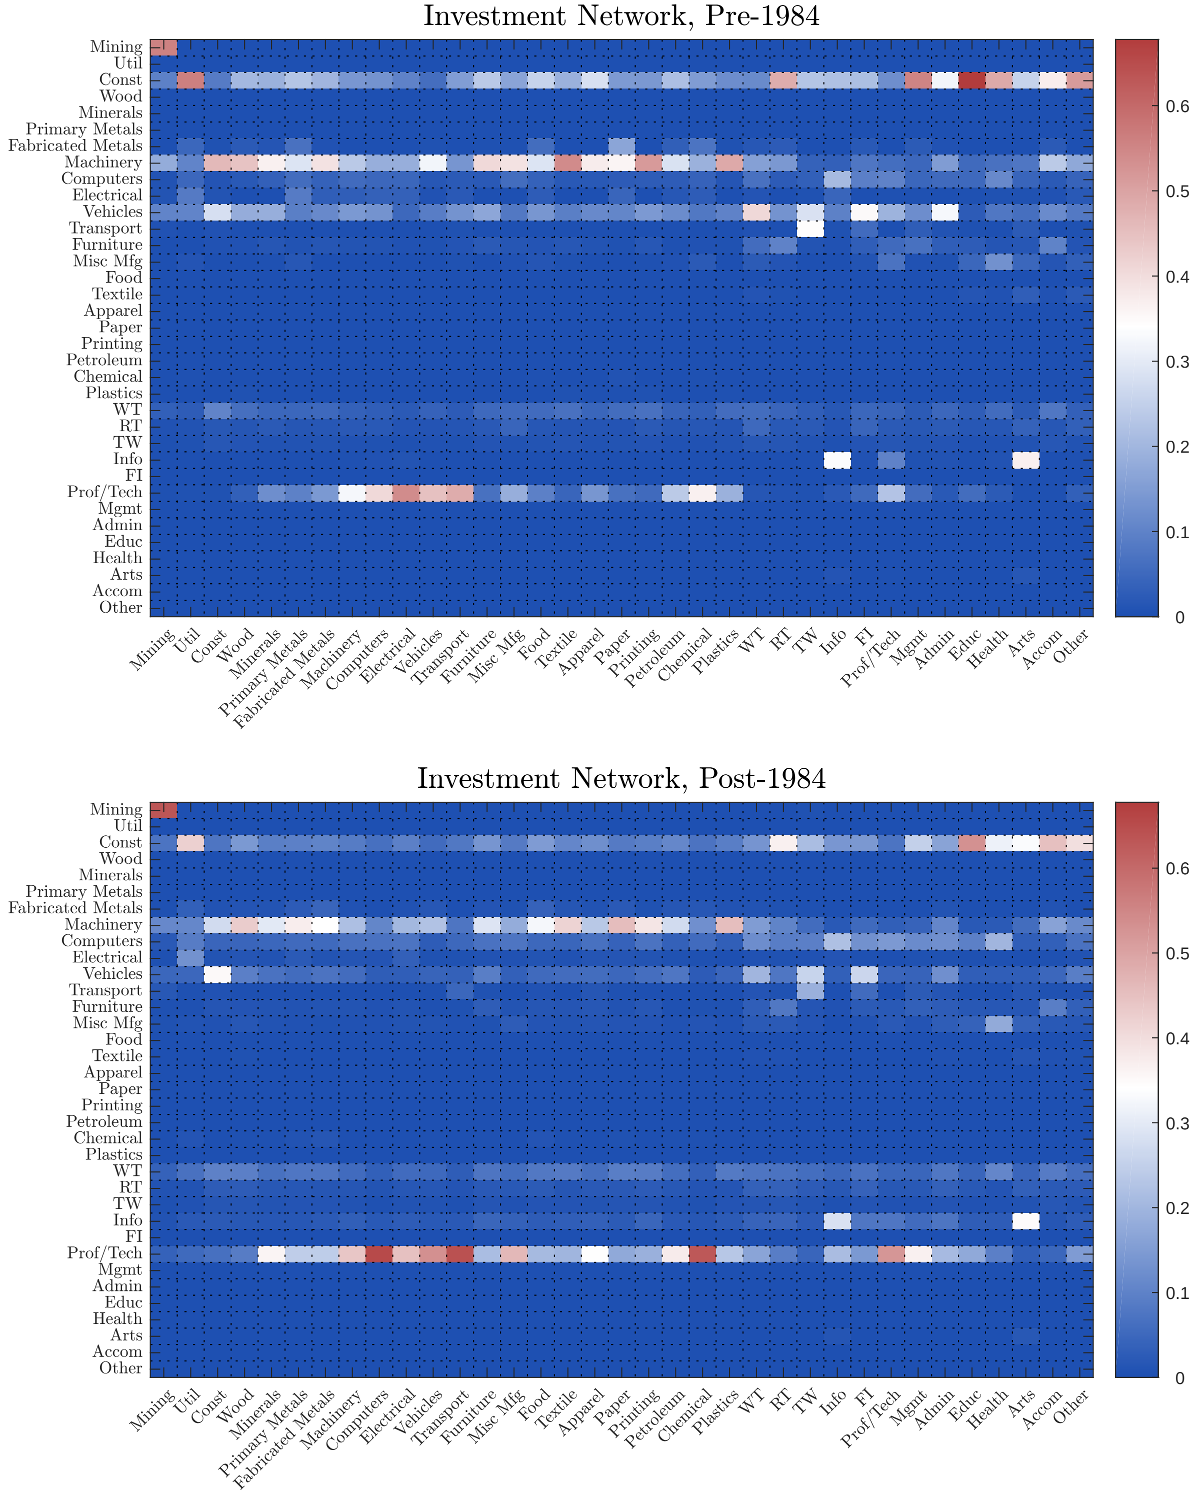

Notes: Heatmaps of the investment input-output network λ_{ij} are constructed as described in the main text. The (i, j) entry of each network corresponds to parameter γ_{ij} and λ_{ij}, i.e., the amount of sector i's good used in sector j. The pre-84 network corresponds to the years 1947-1983 and the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

Figure 15: Correlogam of Sector-level Value Added with Aggregate GDP

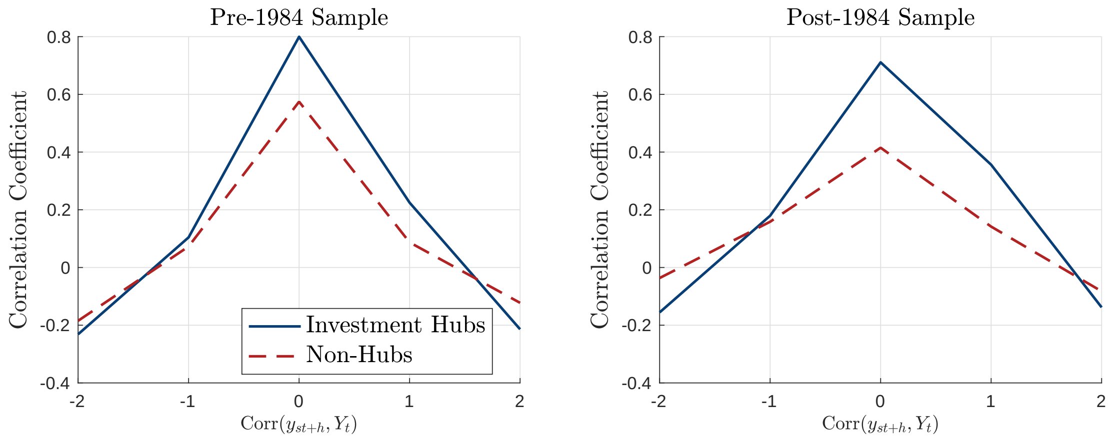

Notes: correlation of value added at sector s in year $t-h$, y_{st+h}, with aggregate GDP in year t, Y_t. Both y_{st+h} and Y_t are logged and HP-filtered with smoothing parameter 6.25. x-axis varies the lag $h \in \{-2,-1,0,1,2\}$. "Investment hubs" compute the unweighted average the value of these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

sectoral shocks than the input-output network.

B Additional Results on Descriptive Evidence of Investment Hubs

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

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

Table 17
Volatility of Activity, Hubs vs. Manufacturing

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

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

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

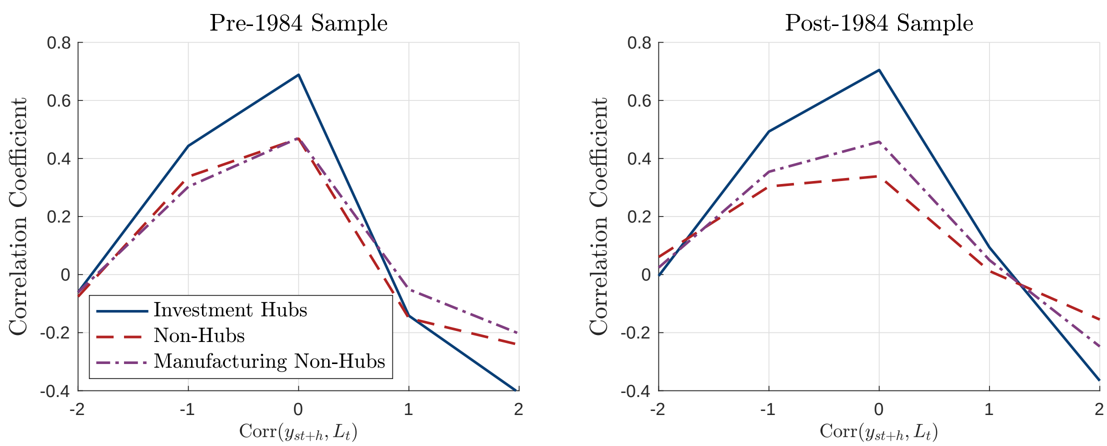

Notes: correlation of value added at sector s in year $t-h$, y_{st+h}, with aggregate employment in year t, L_t. Both y_{st+h} and L_t are logged and HP-filtered with smoothing parameter 6.25. x-axis varies the lag $h \in \{-2, -1, 0, 1, 2\}$. "Investment hubs" compute the unweighted average the value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

TABLE 18
SECTORAL COMOVEMENT WITH INVESTMENT HUBS

	Hubs		**Non-Hubs**	
	Pre-1984	*Post-1984*	*Pre-1984*	*Post-1984*
Value added	0.46	0.32	0.31	0.18
Employment	0.62	0.62	0.53	0.52

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

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

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

C Details of Model Calibration

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

C.1 Parameters Other Than Shocks

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

Figure 17: Calibrated Value Added Shares θ_j

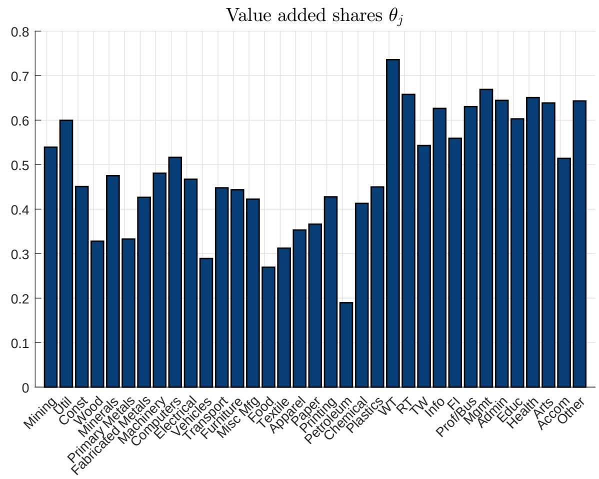

Notes: Values for the value-added shares θ_j are computed as the ratio of value added to gross output in each sector, averaged across the entire sample, 1947-2017.

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

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

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

FIGURE 18: Calibrated Labor Shares $1 - \alpha_j$

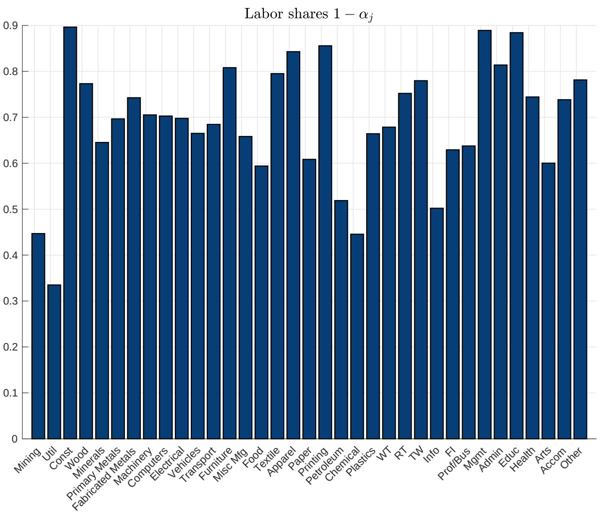

Notes: Values for the labor share $1 - \alpha_j$ are computed from sectoral data on compensation (adjusted for self-employment) divided by value added (with indirect taxes and subsidies removed), averaged across all years in the data, 1947-2017.

compensation for self-employed workers is the same as non-self-employed workers).[65] However, our results are robust to making no adjustments for self-employment. Labor share by industry in each year is then computed as the ratio of adjusted compensation to value added in that industry minus indirect taxes and subsidies.

Figure 18 plots the calibrated labor shares $1 - \alpha_j$ for each sector, averaged over 1947 - 2017. Of course, it is well-known that the labor share has also changed over time; Appendix G shows that our key model results are robust to allowing theses parameters to change over time.

Depreciation rates Depreciation rates by industry are taken directly from the Fixed Assets database, which computes implied depreciation rates by industry for each year from 1947-2017. We set the depreciation rate in each industry, δ_j, equal to the average implied

[65]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 industry based on the finest available industry 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 industry, the implied labor share is frequently in excess of one.

Figure 19: Calibrated Depreciation Rates δ_j

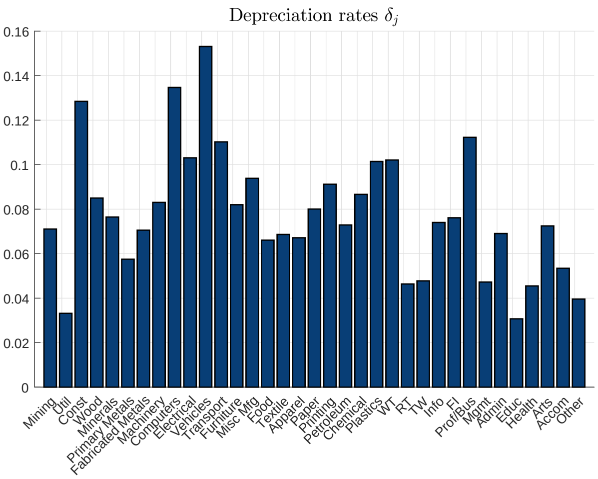

Notes: Values for sector-level depreciation rates δ_j are taken as each sector's average implied depreciation rate from BEA Fixed Assets data, averaged from 1947-2017.

depreciation rate from 1947-2017; these are plotted in Figure 19. These values have slightly risen over time due to the rise of intellectual property products, which have higher depreciation rates. Appendix G shows that our key model results are robust to allowing theses parameters to change over time.

Consumption shares We pin down the the Cobb-Douglas preference parameters weighting consumption in different sectors' output, ξ_j using the BEA Input-Output data on total private consumption by sector. We also account for consumption of new residential structures; without this correction, the consumption share of the construction sector would be essentially zero, and significantly raise its volatility relative to the data.[66] In order to make this correction, we use data from 1997-2017 which separately splits out the final use of each industry's output into residential and non-residential structures. We compute the average fraction of structures production in each industry which is allocated to residential structures. We then add to private consumption by each industry this industry-specific fraction of structures investment that is in residences to capture new housing expenditures in consumption.

[66]Our results are robust – in fact even stronger – if we do not make this correction for consumption of residential structures.

Figure 20: Calibrated Consumption Shares ξ_j

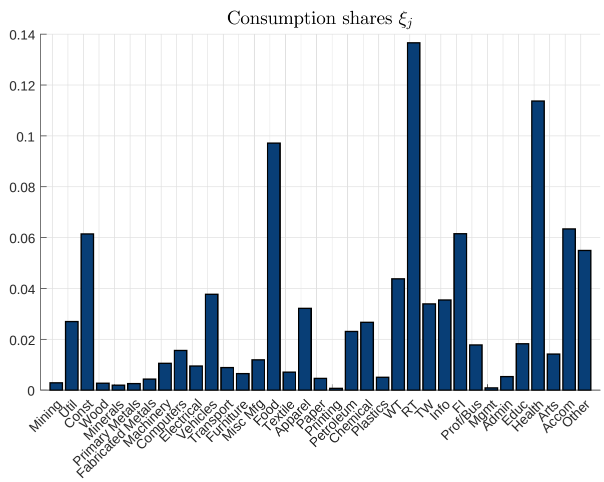

Notes: Values for consumption preference ξ_j are constructed as the fraction of total nominal consumption expenditures on each sector's goods or services, averaged over the entire sample 1947-2017.

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

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

C.2 Measured sector-level productivity

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

TFP directly with value added:

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

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

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

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

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

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

Figure 21: Detrending Sector-Level Data

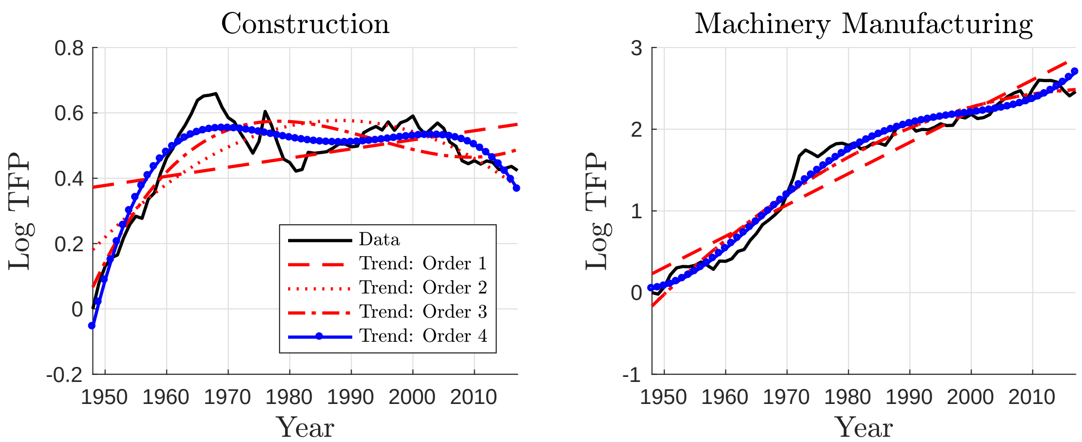

Notes: The figure reports log sector level TFP for the Construction and Machinery Manufacturing industries, normalized to zero in the year 1948. We also report a fitted polynomial trend lines for polynomials of order 1-4, estimated via OLS.

Figure 22: Calibrated Persistence Parameters ρ_j

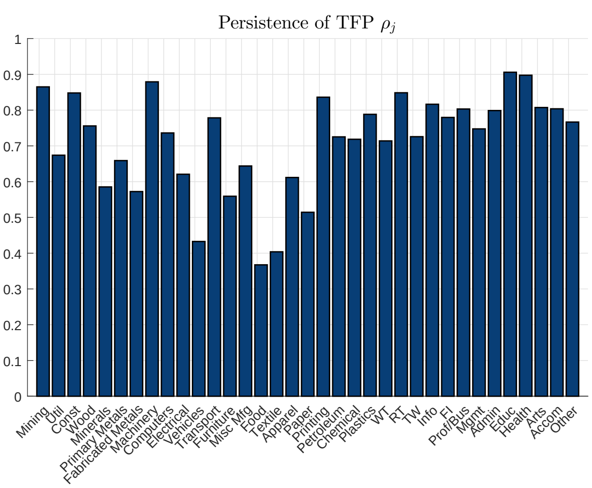

Notes: Persistence parameters ρ_j of sector-level TFP are estimated from log polynomial (of order 4) detrended TFP data using maximum likelihood.

Table 19
Principal Components Analysis of Measured TFP

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

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

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

However, performing that principal components exercise requires us to estimate a full rank covariance matrix for TFP pre- and post-1984, which we cannot do with 35 sectors and less than 35 years of data post-1984.[67] As a result, we collapse our data down to 28 sectors by condensing all non-durable manufacturing sectors into one sector, and then do principal components on log TFP growth for 28 sectors pre- and post-1984.[68]

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

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

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

over time.

D Proof of Theorem that Employment is Constant without Capital

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

$$Y_{jt} = A_{jt} L_{jt}^{\theta_j} X_{jt}^{1-\theta_j}$$
$$Y_{jt} = C_{jt} + \sum_{i=1} M_{jit}$$

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

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

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

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

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

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

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

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

$$
\mu_t\vec{Y}_t = (I-\Phi)^{-1}\xi
$$

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

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

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

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

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

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

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

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

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

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

The key assumptions needed for this result are that consumption preferences and production technologies be Cobb-Douglas, implying constant expenditure shares, and that aggregate preferences over consumption are given by log consumption, which implies that income and substitution effects will exactly offset.

■

E Investment Hubs Forecast Employment Better than GDP

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

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

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

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

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

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

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

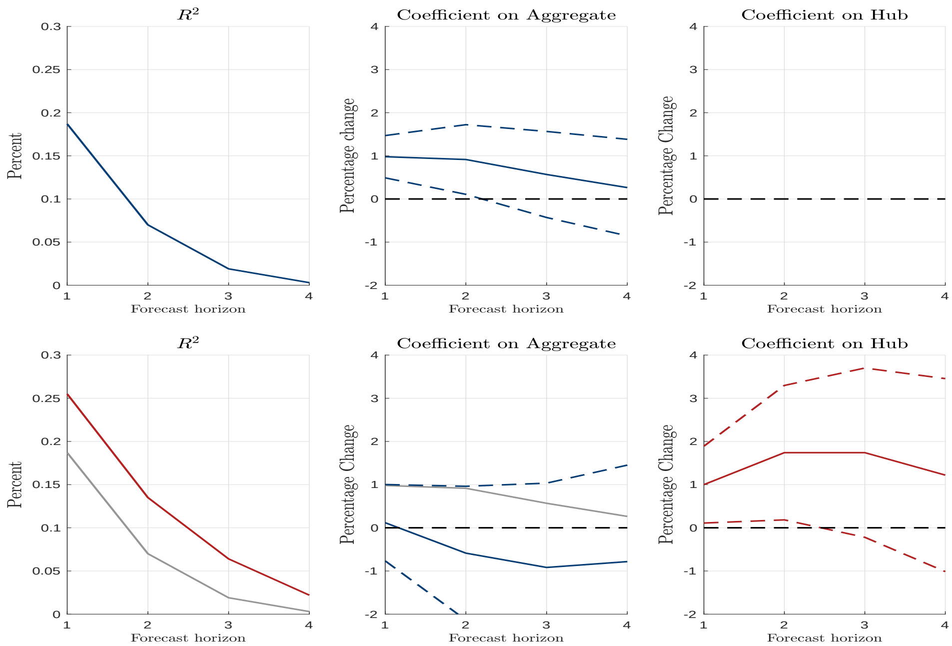

Notes: top panel plots the results from estimating the forecasting regression (19) using aggregate GDP only. Bottom panel plots the results from running the forecasting regression 19 using both aggregate GDP and investment hubs' value added (the grey lines represent point estimates from the results in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

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

F Changes in Aggregate Cycles Driven by Changes in Sectoral Comovement

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

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

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

Declining cyclicality of labor productivity driven by rising volatility of employment 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 15 in the main text). Of course, the definition of the correlation between labor productivity and output is $\mathbb{C}orr(y_t, y_t - l_t) = \frac{\mathbb{C}ov(y_t, y_t - l_t)}{\sigma(y_t)\sigma(y_t - l_t)}$. Using the linear properties of covariance and rearranging, we can write this as:

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

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

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

Combining this expression with the previous one yields:

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

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

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

Table 20

Components of Aggregate Labor Productivity Cyclicality

	Pre-1984	*Post-1984*
$\mathbb{C}orr(y_t - l_t, y_t)$	0.65	0.26
$\mathbb{C}orr(y_t, l_t)$	0.81	0.83
$\mathbb{C}orr(y_t, l_t)$ only	0.65	0.66
$\sigma(l_t)/\sigma(y_t)$	0.76	1.02
$\sigma(l_t)/\sigma(y_t)$ only	0.65	0.26

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

Table 21

Changes in Business Cycles, First Differences

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

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

the 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.

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

Table 22

Within-Sector Business Cycle Statistics with Different Weights

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

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

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

F.0.2 Sectoral Decomposition of Rising Relative Volatility of Employment

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

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

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

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

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

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

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

This expression can be rewritten as:

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

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

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

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

Table 23
Accuracy of the Decomposition

	Pre-84	*Post-84*
Actual, variance	0.58	1.04
Approximation, variance	0.57	0.94
Actual, standard deviation	0.76	1.02
Approximation, standard deviation	0.75	0.97

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

Figure 24: Scatterplot of Changes in Sector-Pair Covariances

(a) Sector-Pair Covariances (b) Within-Sector Variances

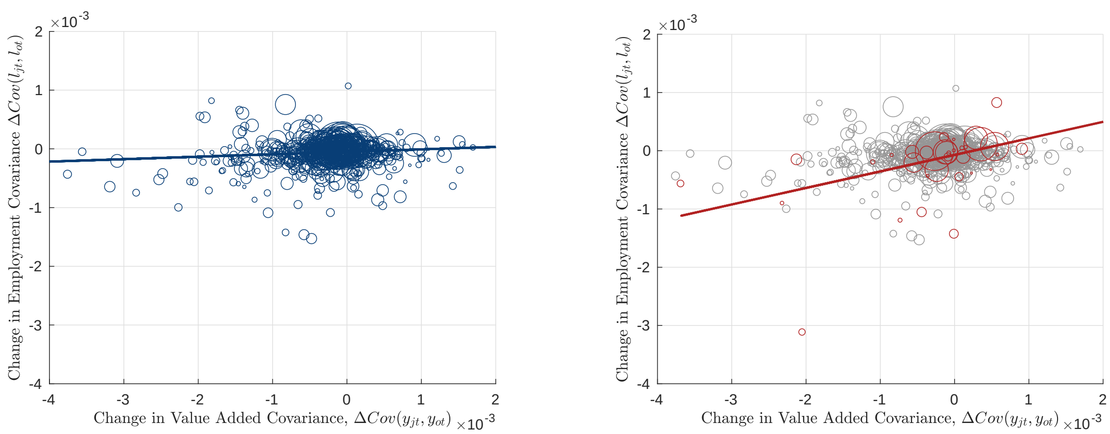

Notes: panel (a) plots changes in the covariance for each pair of sectors (j, o) in our dataset. The horizontal axis computes the change in the covariance of value added $\mathbb{C}ov(y_{jt}, y_{ot})$ in the post-1984 sample (1984-2017) relative to the pre-1984 sample (1948-1983). The vertical axis computes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

TABLE 24
MEASURING COMOVEMENT WITH CORRELATIONS

	Employment	**Value added**
1951 - 1983	0.55	0.36
1984 - 2014	0.51	0.17
Difference	-0.04	-0.18

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

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

The right panel of Figure 24 illustrates the stability of the within-sector variance of employment relative to value added. It shows a positive relationship between the change in the variance of employment and the variance of value added; in fact, the coefficient of the regression line is approximately 0.3, similar to the overall level of their ratio in Table 9.[70]

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

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

Table 25
Decomposition of Relative Employment Volatility, First Differences

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

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

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

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

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

Table 26

Decomposition of Relative Employment Volatility, Fixed Weights

	Pre-84	*Post-84*	*Contribution of entire term*
$\frac{\mathbb{V}ar(l_t)}{\mathbb{V}ar(y_t)}$	0.60	0.81	100%
Variances	0.44	0.32	8%
Covariances	0.62	0.93	92%
Variance Weight	0.11	0.20	
($\omega_t = \sum_{j=1}^{N}(\omega_{jt}^{y})^2\mathbb{V}ar(y_{jt})/\mathbb{V}ar(y_t)$)			

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

Table 27

Decomposition of Relative Employment Volatility, NBER-CES

	Pre-84	*Post-84*	*Contribution of entire term*
$\frac{\mathbb{V}ar(l_t)}{\mathbb{V}ar(y_t)}$	0.40	0.57	100%
Variances	0.34	0.20	1.4%
Covariances	0.37	0.60	92.6%
Variance Weight	0.03	0.06	
($\omega_t = \sum_{j=1}^{N}(\omega_{jt}^{y})^2\mathbb{V}ar(y_{jt})/\mathbb{V}ar(y_t)$)			

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

Table 28
Decomposition of Relative Employment Volatility, Goods vs. Services

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

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

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

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

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

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

TABLE 29
DECOMPOSITION OF INVESTMENT VOLATILITY

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

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

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

G Additional Quantitative Model Results

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

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

Table 30
Population Moments, Baseline Model vs. Changing Parameters

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

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

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

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

Table 31

Quantitative Results, Alternative Orders of Polynomial Detrending of TFP

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

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

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

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

Table 32
Quantitative Results, Baseline Model vs. Maintenance Investment

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

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

Table 33
Quantitative Results, Baseline Model vs. Adjustment Costs

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

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

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

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

FIGURE 25: Aggregate Dynamics Driven by Shocks to Individual Sectors

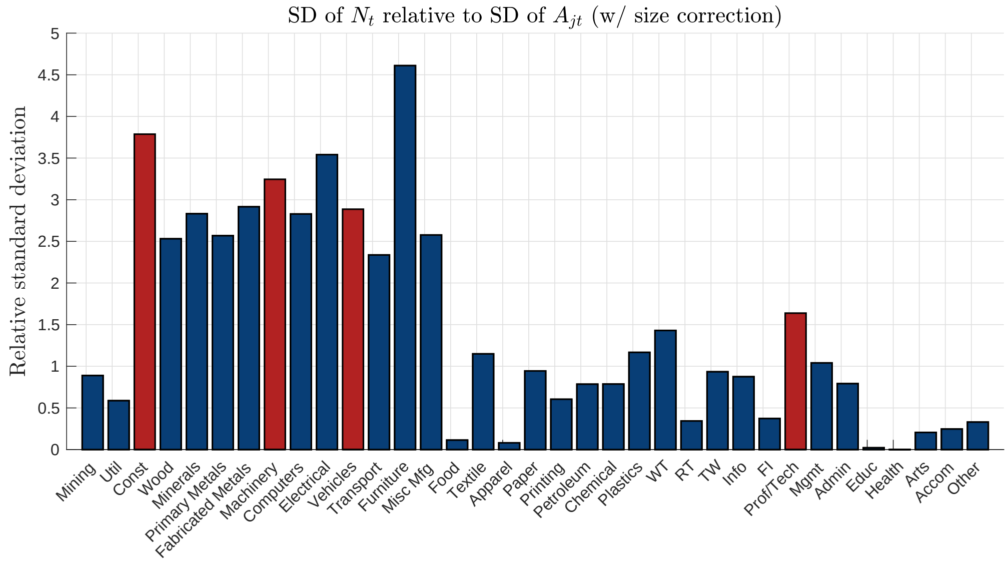

Notes: results from simulating model with empirical shocks to only one sector at a time (the remaining sectors' shocks are set to zero). Plots standard deviation of aggregate employment n_t relative to the standard deviation of the sectors' productivity A_{jt}. To correct for sector size, the figure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

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

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

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

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

FIGURE 26: Distributional Effect of Investment Stimulus

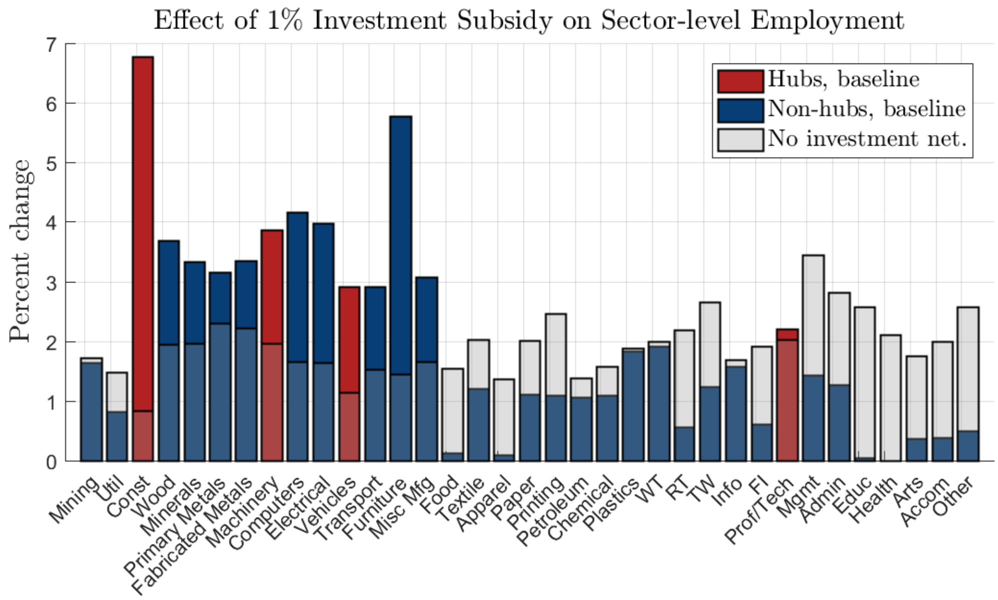

Notes: effect of a one-time $sub_t = 0.01$ shock to the stimulus policy shock described in the main text. Each bar plots the percentage change in employment at that particular sector. Red bars are the investment hubs' response in our baseline model, blue bars are the non hubs' response in our baseline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

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

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