Appendix

A Conditions for the thresholding rule

$$
\begin{aligned}
&(i) && |s_\lambda(z)| \leq c|y| \quad \forall \quad z, y \text{ satisfying } |z-y| \leq \lambda \text{ and some } c > 0 && \\
&(i)^* && |s_\lambda(z)| \leq |z| && \text{(shrinkage)} \\
&(ii) && s_\lambda(z) = 0 \text{ for } |z| \leq \lambda && \text{(thresholding)} \\
&(iii) && |s_\lambda(z) - z| \leq \lambda \quad \forall \quad z \in R && \text{(limits of shrinkage)}
\end{aligned}
$$

Cai and Liu (2011) conduct analyses for the class satisfying (i),(ii),(iii) but state that it is also possible to adapt this to (i)*. The hard-thresholding rule is ruled out by (i), but other thresholding functions, such as the soft-thresholding $s_\lambda(z) = \text{sgn}(z)(|z| - \lambda)_+$ and the adaptive LASSO rule $s_\lambda(z) = z(1 - |\lambda/z|^\eta)_+$ with $\eta \geq 1$, are included.

B Illustration of Cross-Validation Procedure

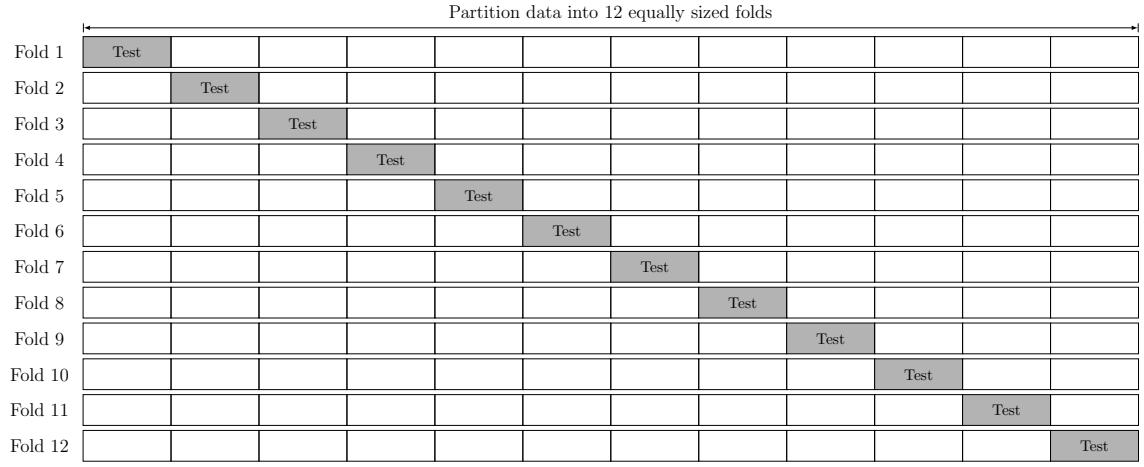

Figure B.8: Illustration of 12-fold CV. First we split the sample into 12 folds of equal length. Then we use 11 folds for training and validate the trained model's performance on the hold-out test fold.

C χ^2 Test Statistic for Equality of Covariance Matrices

We want to test the hypothesis if the generalized shock covariance matrix Ω is different from the identity matrix, i.e.,

$$H_0 : \Omega = I_N.$$

Let N be the number of time series in our panel and T be the number of time periods in the respective sample. Ledoit and Wolf (2002) develop a χ^2 test statistic to test H_0 specifically when N is large and T is small:

$$W = \frac{NT}{2}\left\{\frac{1}{N}\mathrm{tr}\big((\Omega - I_N)^2\big) - \frac{N}{T}\left[\frac{1}{N}\mathrm{tr}(\Omega)\right]^2 + \frac{N}{T}\right\}.$$

Under the null hypothesis H_0, W is distributed as $\chi^2\big[\frac{N(N+1)}{2}\big]$. Note that the degrees of freedom parameter for the χ^2-distribution corresponds to the number of distinct parameters in Ω. Finally, we reject H_0 at significance level α if W is greater than $\chi^2\big[\alpha, \frac{N(N+1)}{2}\big]$. A corresponding p-value follows as $p = 1 - \mathrm{CDF}_{\chi^2_{N(N+1)/2}}(W)$.

D Complementary Graphs

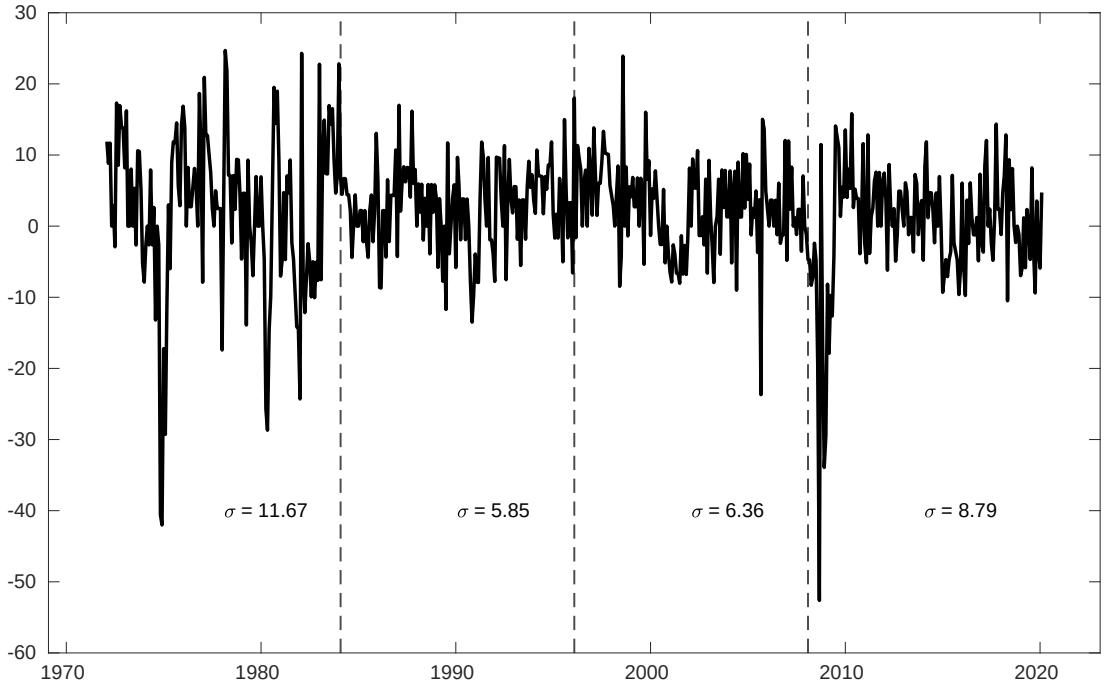

Figure D.9: Annualized growth rates of monthly aggregate industrial production (IP) in percentage points. σ denotes the sample volatility in the respective subsample.

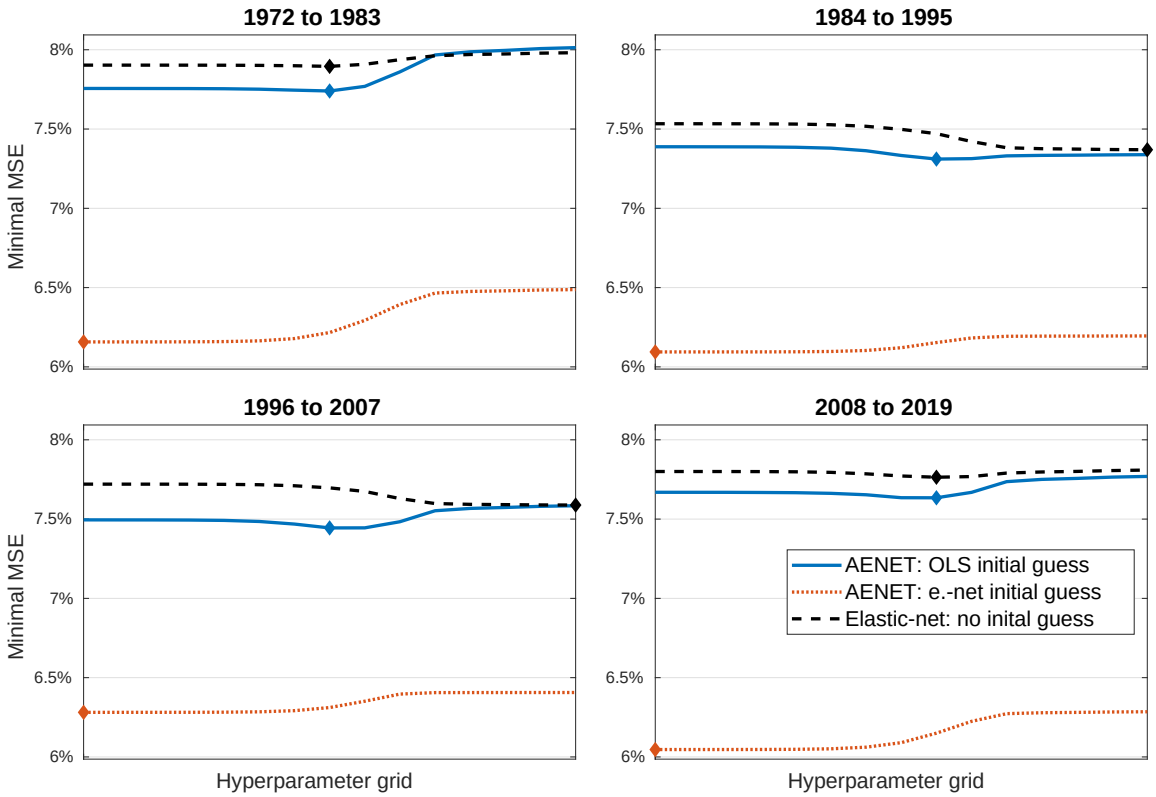

Figure D.10: 12-fold CV results for the tuning parameter α in the adaptive elastic-net. The scale shows the percentage improvement over the OLS case, e.g., 33% refers to a forecast error being only a third that of the OLS estimate. The blue solid curves show the minimal MSE with $\hat{A}_{OLS}$ as the initial estimate in the weights of the individual penalties for different αs, following the approach of Demirer et al. (2017). The red dotted curves show the same for $\hat{A}_{ENET}$ as an initial guess (selected by 12-fold CV, $\alpha = 0.5$, and $w_i = 1$), the black dashed lines correspond to the original (non-adaptive) elastic-net as in Zou and Hastie (2005).

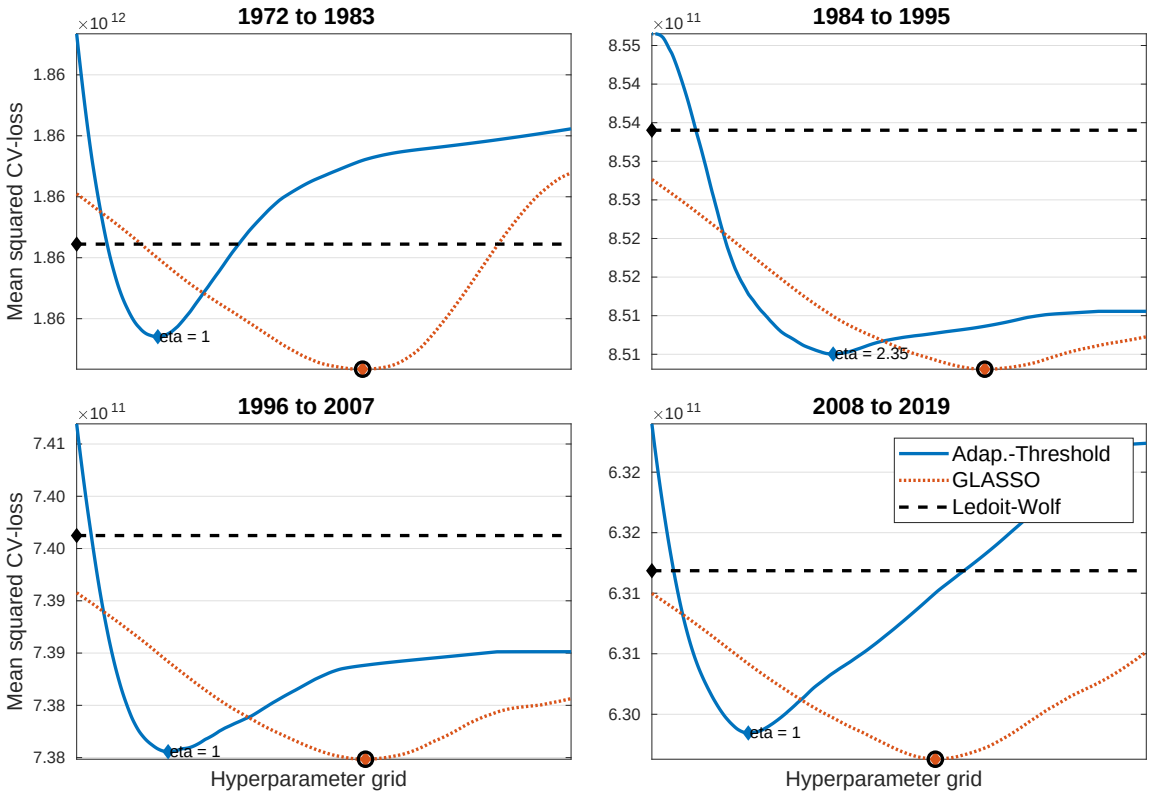

Figure D.11: The 12-fold CV results for the tuning parameter of different covariance regularization methods. Values of the tuning parameters are on the x-axis and the MSEs are shown on the y-axis. The two thresholding estimators tune δ^{AT} in (13), GLASSO tunes δ^{GL} in (16), and the manual Ledoit-Wolf tunes the otherwise automatic shrinkage parameter δ^{LW}. The x-axis is normalized for each regularization method such that the convexity of all of the MSE curves is visible.

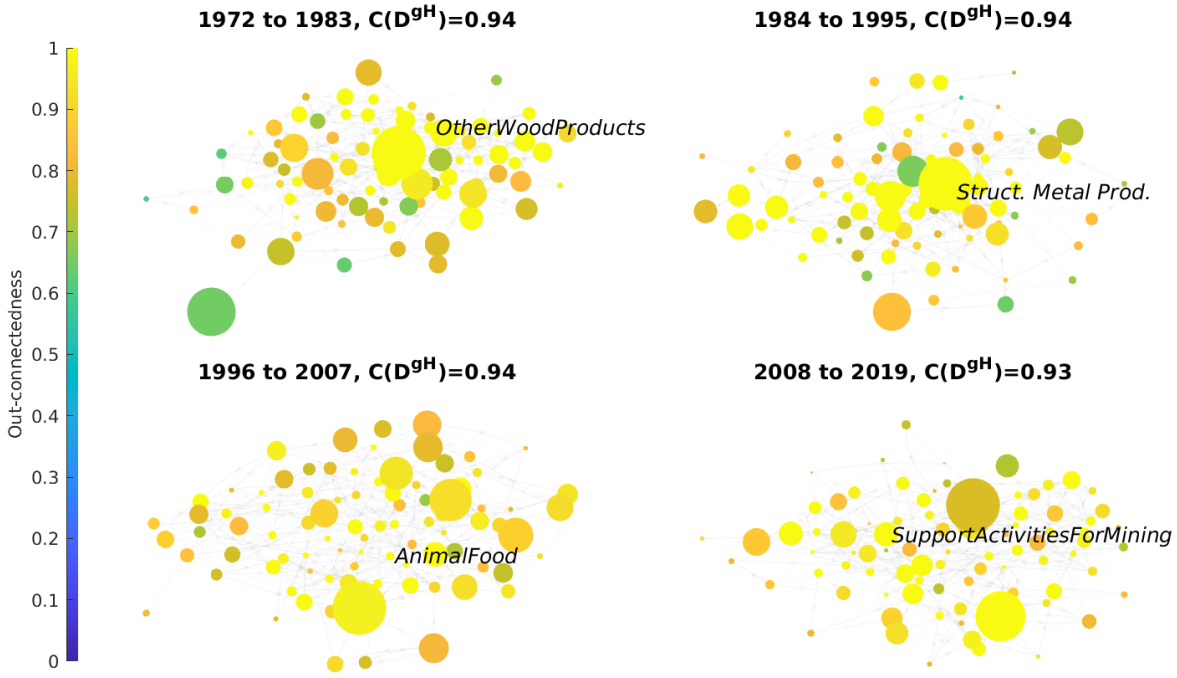

Figure D.12: Connectedness networks for the respective subsamples using the non-regularized estimates. The force-directed graph drawing algorithm arranges the nodes. That is, two nodes appear closer in the graph if they have stronger connections to each other. Although we initialize the subgraphs on the same scale, the algorithm cannot guarantee that the graphs are comparable in size. The size of the node relates to the respective average weight $\bar{w}_{i,t^*}$ of the sector in the IP index. The colors depict the out-connectedness. For the sake of the visualization, we cap the color scale at 1. The sectors with the highest out-connectedness are labeled.

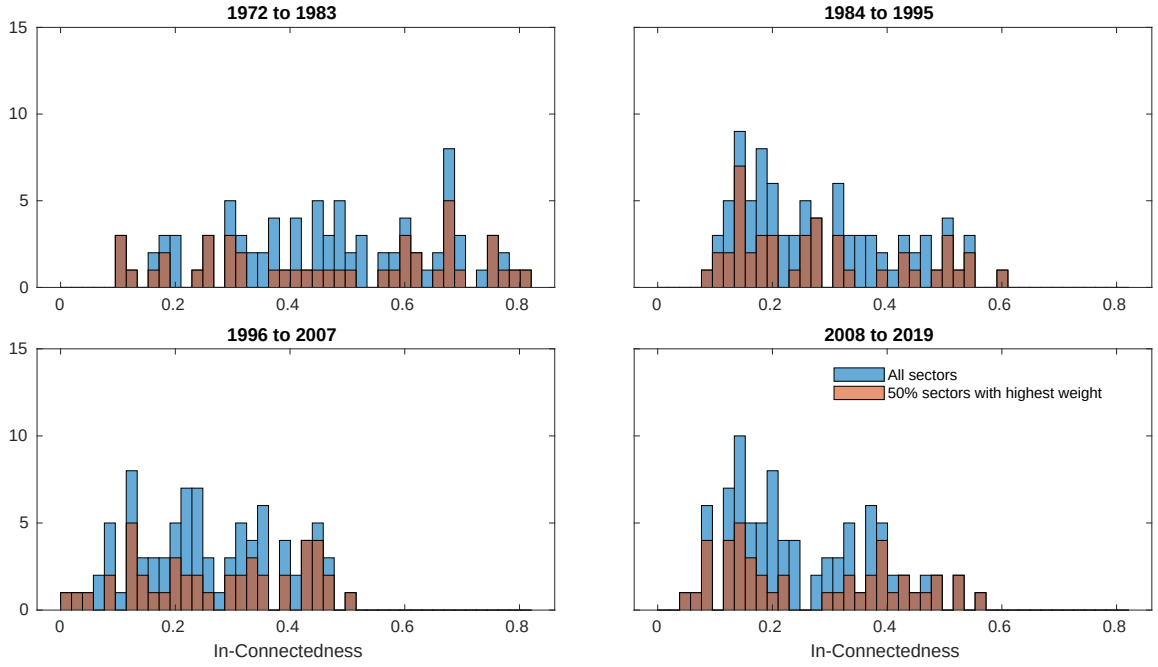

Figure D.13: Histogram of the in-connectedness of the 88 three-digit level sectors. The 50% of the sectors with the biggest weights are highlighted in red. The in-connectedness relates to incoming spillovers from other sectors.

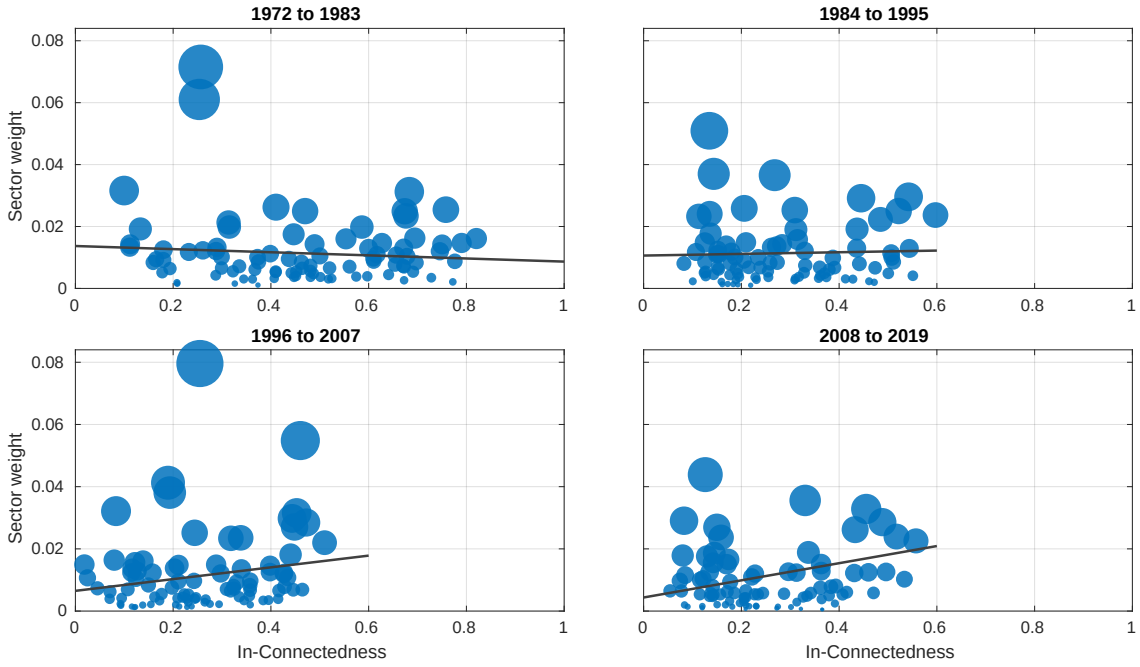

Figure D.14: Scatterplots of the in-connectedness of the 88 three-digit-level sectors against the sectoral weights in the aggregate IP index. The in-connectedness relates to incoming spillovers from other sectors. We plot the regression line per sample.

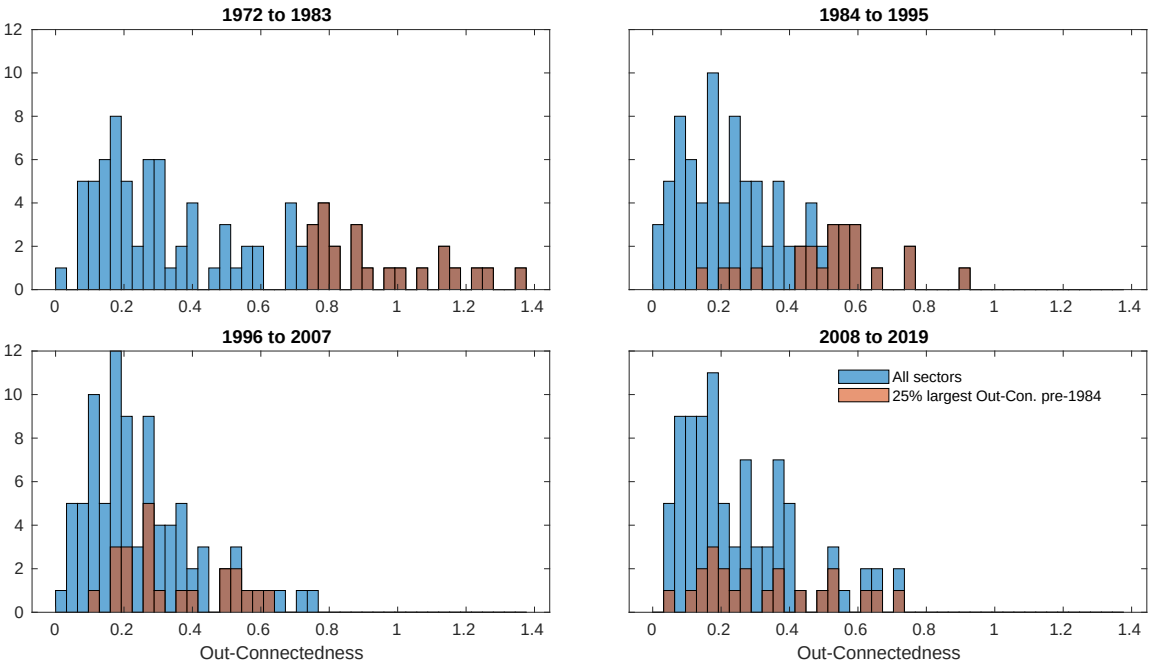

Figure D.15: Histogram of the out-connectedness of the 88 three-digit-level sectors. The 25% of the sectors with the biggest weights in the first subsample are highlighted in red. The out-connectedness relates to outgoing spillovers to other sectors.

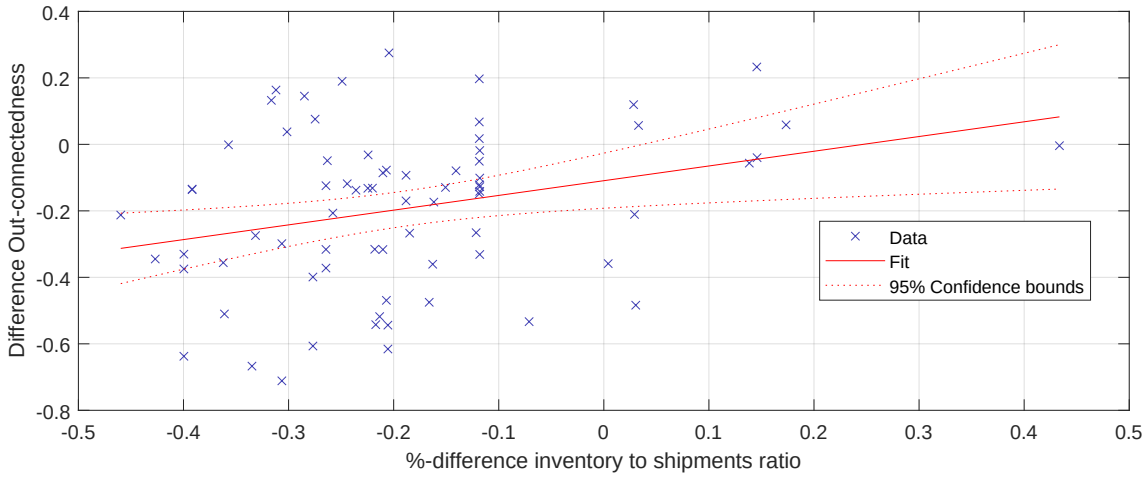

Figure D.16: Regression plot of sectoral changes in out-connectedness on percentage changes in inventories-to-shipments ratio. Changes in out-connectedness for sector i $\Delta OutCon_i$ are differences in the estimates from the pre-Great Moderation subsample (1972 to 1983) to the next subsample (1984 to 1995). Inventories-to-shipments ratios are obtained from the Manufacturers' Shipments, Inventories, and Orders (M3) survey and percentage changes $\Delta I2S_i$ follow from pre- and post-1984 means. 73 SIC categories are matched to NAICS categories via the official dispersement table of the US Census Bureau. Note that some sectors' series are unpublished and the matching is imperfect due to new sector definitions. The fitted regression equation reads $\Delta OutCon_i = -0.1093 + 0.4429\,\Delta I2S_i$. t-statistics (p-values) for the intercept and coefficient are $-2.6295(0.0105)$ and $2.6088(0.0111)$, respectively.

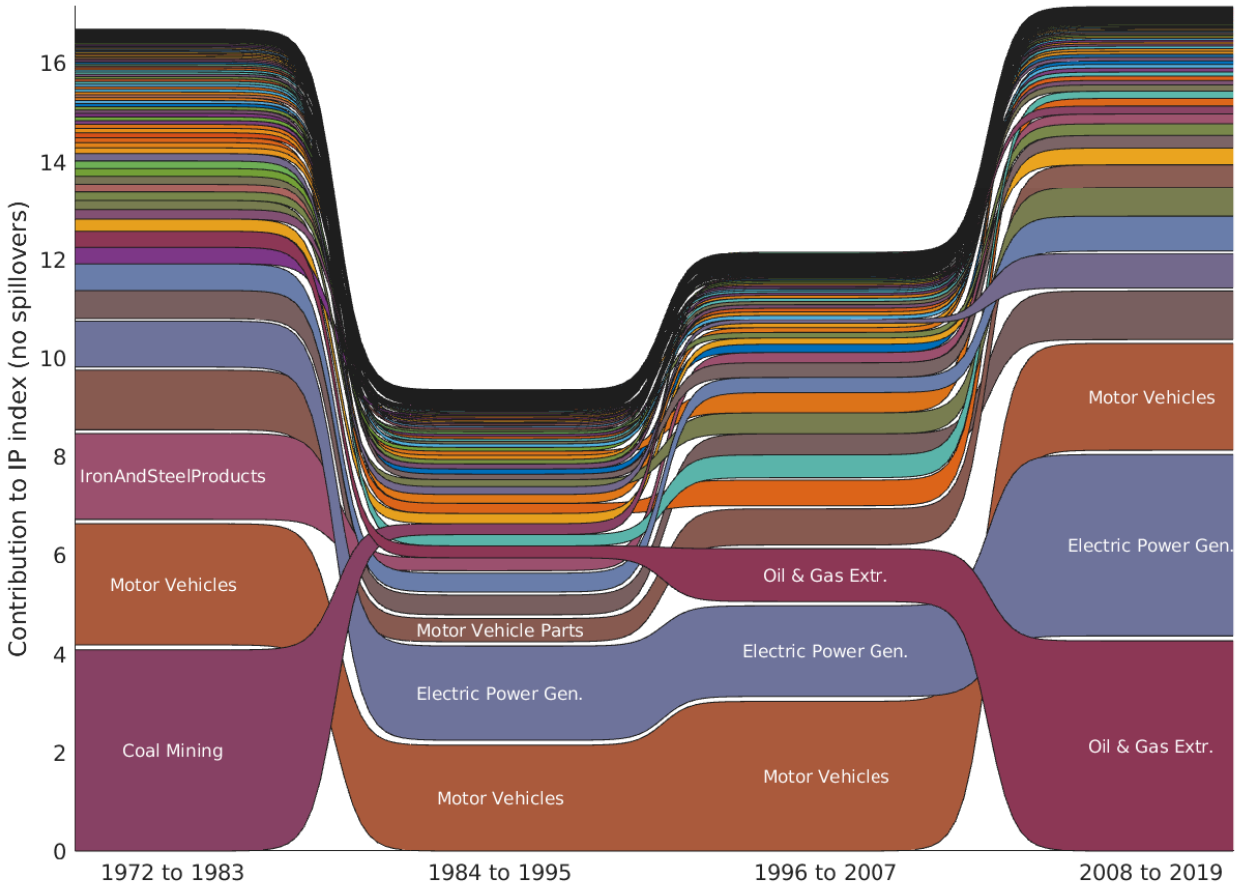

Figure D.17: Bump chart of sectoral contributions to the aggregate variance of the IP index without spillovers. Contributions are calculated as in (17). Data includes IP indices for 88 three-digit-level sectors. The top three contributors per subsample are labeled.

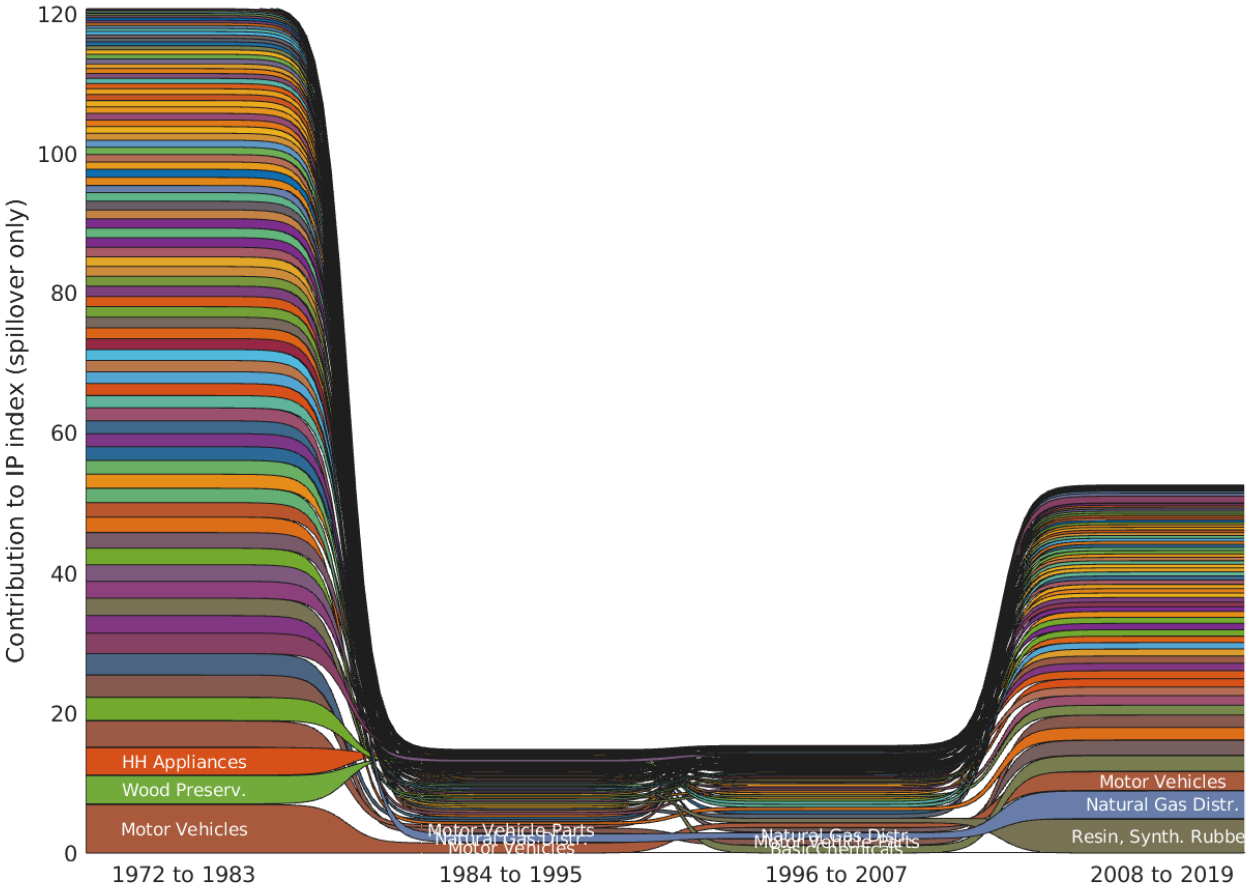

Figure D.18: Bump chart of sectoral contributions to the aggregate variance of the IP index with spillovers only. Contributions are calculated as in (17). Data includes IP indices for 88 three-digit-level sectors. The top three contributors per subsample are labeled.