



Working Papers

www.cesifo.org/wp

Innovation Network

Daron Acemoglu
Ufuk Akcigit
William Kerr

CESIFO WORKING PAPER NO. 6173

CATEGORY 6: FISCAL POLICY, MACROECONOMICS AND GROWTH

NOVEMBER 2016

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

ISSN 2364-1428

Innovation Network

Abstract

Technological progress builds upon itself, with the expansion of invention in one domain propelling future work in linked fields. Our analysis uses 1.8 million U.S. patents and their citation properties to map the innovation network and its strength. Past innovation network structures are calculated using citation patterns across technology classes during 1975-1994. The interaction of this pre-existing network structure with patent growth in upstream technology fields has strong predictive power on future innovation after 1995. This pattern is consistent with the idea that when there is more past upstream innovation for a particular technology class to build on, then that technology class innovates more.

JEL-Codes: O310, D850, O320, O330, O340.

Keywords: innovation, networks, patents, growth.

Daron Acemoglu
Department of Economics
Massachusetts Institute of Technology
USA - 02142-1347 Cambridge MA
daron@mit.edu

Ufuk Akcigit
Department of Economics
University of Chicago
USA - Chicago, IL, 60637
uakcigit@uchicago.edu

William Kerr
Harvard Business School
Soldiers Field Road
USA – Boston, MA 02163
wkerr@hbs.edu

October 2016

We are grateful to Ben Jones and Paula Stephan for their very helpful directions and to seminar participants for their comments. We also thank Alexis Brownell for excellent research assistance. Finally, we gratefully acknowledge financial support from the National Science Foundation, the Alfred P. Sloan Foundation, the Ewing Marion Kauffman Foundation, and Harvard Business School.

Technological and scientific progress propels economic growth and long-term well-being. Prominent theories depict this process as a cumulative one in which new innovations build on past achievements, using Newton’s descriptive phrase of “standing on the shoulders of giants” (e.g., refs 1 and 2). Several studies provide evidence supporting this view, and more generally, knowledge development is embedded in a landscape of individual scientists, research institutes, private sector actors, and government agencies that shape the fundamental rate and direction of new discoveries.¹ Despite this burgeoning literature, our understanding of how progress in one technological area is linked to prior advances in upstream technological fields is limited. Open but important questions include the long-term stability of how knowledge is shared across technological fields, the pace and timing of knowledge transfer, and how closely connected upstream fields need to be to have material impact on a focal technology. This paper provides some quantitative evidence on these and related questions.

We show that a stable *innovation network* acts as a conduit of this cumulative process of technological and scientific progress. We analyze 1.8 million U.S. patents and their citation properties to map the innovation network and its strength. Past innovation network structures are calculated using citation patterns across technology classes during 1975–1994. The interaction of this pre-existing network structure with patent growth in “upstream” technology fields has strong predictive power on future “downstream” innovation after 1995. Remarkably, 55% of the aggregate variation in patenting levels across technologies for 1995–2004 can be explained by variation in upstream patenting; this explanatory power is 14% when using panel variation within each field (R-squared value from regressions tabulated below). Detailed sectors which have seen more rapid patenting growth in their upstream technology fields in the last ten years are much more likely to patent today.

This pattern is consistent with the idea that when there is more past innovation for a particular technology class to build on, then that technology class innovates more. As an example using patent subcategories defined below, “Chemicals: Coating” and “Nuclear & X-rays” display similar patenting rates in 1975–1984. Before 1995, citation patterns indicate that “Nuclear & X-rays” drew about 25% of its upstream innovation inputs from “Electrical Measuring & Testing”, while “Chemicals: Coating” had a similar dependence on “Chemicals: Misc.”. The former upstream field grew substantially less during 1985–1994 than the latter in terms of new patenting. In the ten year period after 1995, “Chemicals: Coating” exhibits double the growth of “Nuclear & X-rays”. The network heterogeneity further indicates that knowledge development

¹For example, references (3)–(13).

is neither global, in the sense that fields collectively share an aggregate pool of knowledge, nor local, in the sense that each field builds only upon itself.

It is useful to motivate our approach with the standard endogenous growth and technological progress models in economics, which posit a production function of new ideas of the form

$$\Delta N(t) = f(N(t), R(t)),$$

where $N(t)$ is the stock of ideas, $\Delta N(t)$ is the flow of new ideas produced, and $R(t)$ is the resources that are utilized to produce these new ideas (e.g., scientists). Though some studies estimate the impact of the stock of ideas, $N(t)$, on the flow of new ideas (e.g., whether there are increasing returns or “fishing out” externalities), most of the literature takes the input into the production function of new ideas in every field to be either their own idea stock or some aggregate stock of knowledge spanning across all fields. We take a step towards opening this black box and measuring the heterogeneous dependence of new idea creation on the existing stock of ideas through studying innovation networks.

We suppose that new innovations in technology $j \in \{1, 2, \dots, J\}$ depend on past innovations in all other fields through an innovation network. Suppressing the resource variable $R(t)$ for simplicity and assuming a linear form, we can write

$$\Delta N_{J \times 1}(t) = \alpha \cdot M_{J \times J} \cdot N_{J \times 1}(t),$$

where $\Delta N_{J \times 1}(t)$ and $N_{J \times 1}(t)$ are, respectively, the $J \times 1$ vector of innovation rates and the stock of knowledge in the J technology classes at time t , and $M_{J \times J}$ is a $J \times J$ matrix representing the innovation network—how much one class builds on the knowledge stocks of other classes. Given the scalar α and our focus on relative growth for technologies, we can normalize the row sums of $M_{J \times J}$ to one. The case in which new innovations depend symmetrically upon an economy-wide technology stock is represented by all entries in $M_{J \times J}$ being equal to $1/J$; the case in which fields only build upon their own knowledge stock is given by the identity matrix.

We analyze utility patents granted between 1975–2009 by the United States Patent and Trademark Office (USPTO). Each patent record provides information about the invention (e.g., technology classifications, citations of patents on which the current invention builds) and the inventors submitting the application. We analyze 1.8 million patents applied for in 1975–2004 with at least one inventor living in a U.S. metropolitan area. The 2004 end date allows for a five-year window for patent reviews. In our data, 98% of patent reviews are completed within this window.

Figure 1 describes the 1975–1984 innovation network in matrix form.² The year restriction refers to the dates of cited patents, and forward citing patents are required to be within ten years of the cited patent. The ten-year window for forward citations keeps a consistent number of observations per diffusion age. USPTO technologies are often grouped into a three-level hierarchy: 6 categories, 36 subcategories, and 484 classes. This matrix lists subcategories and their parent categories; our empirical analysis considers subcategory- and class-level variation.

Each row provides the composition of citations made by the citing technology field, summing to 100% across the row. Own-citations (citations that fields make to themselves) account for a majority of citations, and for visual purposes are given a dark shading in Figure 1. In our empirical work, we face a dilemma: a complete growth accounting includes how cumulative technological progress in one field affects the field’s own future development. In fact, own-technology spillovers are usually the most important channel of cumulative knowledge development and also connect to the concept of absorptive capacity, where research in one’s own field prepares one to absorb external knowledge from other fields (e.g., refs 18 and 19). However, it is very difficult to convincingly establish the importance of the innovation network when looking within individual fields, as technological progress for a field over time can be endogenously related to its past and future progress, as well as outside factors, and also display serial correlation for other reasons (e.g., rising government funding levels, dynamic industry conditions). A contribution of our network-based analysis that uses upstream technology progress outside of an individual field, as moderated by a pre-existing network structure, to predict future innovation is to demonstrate the importance of this knowledge development process in an empirical setting that minimizes these difficult identification challenges.

We thus present our findings below in two ways. One route is to consider the external network only, which excludes own-citations and within-field spillovers to better isolate network properties. We write our upcoming equations for this case. To afford the complete growth perspective, we also report results for the complete network that includes own-field spillovers. Formally, an entry in matrix $M_{J \times J}$ from a citing technology j (row) to a cited technology j' (column) is

$$m_{j \rightarrow j'} = \frac{\text{Citations}_{j \rightarrow j'}}{\sum_{k \neq j} \text{Citations}_{j \rightarrow k}}.$$

In this representation, the notation $j \rightarrow j'$ designates a patent citation from technology j to

²Hall et al. (14) further describe the patent data. Studies of cross-sector spillovers date to at least Scherer (15) and Verspagen (16). Schmitzer and Watzinger (17) provide a recent example.

j' , which in turn means knowledge flowing from technology j' to j . For the complete network calculation, the denominator summation includes $k = j$.

Figure 1 highlights the heterogeneity in technology flows. The block diagonals indicate that subcategories within each parent category tend to be inter-related, but these flows vary substantially in strength and show important asymmetries. For example, patents in “Computers: Peripherals” tend to pull more from “Computers: Communications” than the reverse, as “Computers: Communications” builds more on electrical and electronic subcategories. There are also prominent examples of connections across technology categories, such as the link between “Organic Compounds” and “Drugs.” Figure 2 depicts this information in a network format, which groups in two-dimensional space the stronger relationships in nearer proximity.

The innovation network is quite stable. Calculating $M_{J \times J}$ for the ten-year periods of 1975–1984, 1985–1994, and 1995–2004, the correlations and rank correlations of cell values over ten-year horizons are both above 0.9; across a 20-year horizon, both are above 0.8. Appendix Figures 1–3 show comparable network structures when using more-stringent thresholds for including network edges/connections, when examining raw citations without the normalization such that every technology’s outbound citations are equally weighted, and when using longer data horizons. Appendix Figures 4–6 show three frequently calculated diagnostics for network nodes: in-degree importance, closeness, and betweenness. A common theme, which is also evident in Figure 2, is that many high-profile technology areas (e.g., Drugs) are at the periphery of the innovation network. Technologies like “Electrical Devices” and “Materials Processing/Handling” occupy more-central positions.

We take advantage of the considerable heterogeneity in the speed at which knowledge diffuses—how many years after invention patents in technology j' typically receive citations from technology j . We construct our innovation network matrix to model separately each year of the diffusion process:

$$\text{CiteFlow}_{j \rightarrow j', a} = \frac{\text{Citations}_{j \rightarrow j', a}}{\text{Patents}_{j'}}$$

where CiteFlow quantifies the rate at which patents in technology j cite patents in j' ($\text{Patents}_{j'}$) for each of the first ten years after the latter’s invention.³ This augmented structure extends the simple theoretical model described with the M matrix to allow for more complex knowledge

³Time lags consistently broaden the downstream technology impact. One year after invention, 81% of downstream citations are from the same category (62% are from the same patent class, 10% are from another patent class within the same subcategory, and 9% from another subcategory within the same category). After ten years, 75% of citations occur within the same patent category (respectively, 51%, 12%, 12%).

diffusion processes that depend upon invention age.⁴

To predict forward patents, we combine the pre-existing network with technology development that occurs within a ten-year window before the focal year t . Define $\hat{P}_{j,t}$ to be the expected patenting in technology j for a year t after 1994. Our estimate of $\hat{P}_{j,t}$ combines patents made in the prior ten years with an added diffusion lag of $a = [1, 10]$ years,

$$\hat{P}_{j,t} = \sum_{k \neq j} \sum_{a=1}^{10} \text{CiteFlow}_{j \rightarrow k, a} P_{k, t-a},$$

where $P_{k, t-a}$ is the patenting in technology k at a diffusion lag a from the year t . As an example, for a patent from technology j' applied for in 1990, we model its impact for technology j in 1997 by looking at the average impact that occurred with a seven-year diffusion lag during the pre-period. The double summation in the calculation of $\hat{P}_{j,t}$ repeats this process for each potential upstream technology class and diffusion lag. In addition to the network being estimated from pre-period interactions, our calculation requires that upstream patents pre-date downstream predictions by at least one year (i.e., $a \geq 1$). For the complete network calculation, the first summation term again includes $k = j$.

The first row of Panel A in Figure 3 reports the strong levels relationship between the predicted values ($\hat{P}_{j,t}$) and actual values ($P_{j,t}$) using subcategory variation in a log format. This estimate includes 360 observations through the analysis of 36 subcategories in each year during 1995–2004; each subcategory is weighted by its initial level of patenting. A 10% increase in expected patenting is associated with an 8% increase in actual patenting when considering the external network. We report standard errors that are robust against serial correlation within a subcategory. This specification explains about 55% of the aggregate variation in 1995–2004 patenting levels. The empirical strength of the complete network estimation is even stronger, with a 10% increase in expected patenting associated with a 9% increase in actual patenting.

While powerful, there are several potential concerns with the simple approach. First, persistence in the relative sizes of technological fields may lead to overstatements of network importance. Likewise, aggregate fluctuations in the annual patenting rates of all fields could result in over-emphasis on the importance of upstream fields. To address, we consider a panel regression

⁴While Figures 1 and 2 are normalized to sum to 100% for a citing technology using the network matrix M , we leave this measure relative to baseline patenting to allow direct use with the forward patenting rates by technology. Patents differ substantially in the number of citations that they make, and we weight citations such that each citing patent receives the same importance. Our results are robust to different approaches for dealing with patents that make no citations and instances where patents list multiple technologies.

that includes field and time controls,

$$\ln(P_{j,t}) = \beta \ln(\hat{P}_{j,t}) + \phi_j + \eta_t + \varepsilon_{j,t},$$

where $P_{j,t}$ and $\hat{P}_{j,t}$ are actual and expected patent rates for technology j in year t ($\varepsilon_{j,t}$ is an error term). The estimation includes fixed effects for subcategories (ϕ_j) that remove their long-term sizes; likewise, fixed effects for years (η_t) remove aggregate changes in USPTO grant rates common to all technologies, so that the identification of the β parameter comes only from variations within fields. Intuitively, β captures whether the actual patenting in technology j is abnormally high relative to its long-term rate when it is predicted to be so based upon past upstream innovation rates. A β estimate of one would indicate a one-to-one relationship between predicted and actual patenting after conditioning on these controls.

We estimate in the second row of Panel A a statistically significant and economically substantial value of β : 0.85 (standard error = 0.17). Though less than one, the estimated coefficient shows a very strong relationship between predicted and actual patenting. Appendix Figure 7 provides visual representations of these subcategory-level estimations. This figure shows that our results are not driven by outliers or weighting strategy.⁵

Panel B of Figure 3 shows very similar patterns when using variation among more-detailed patent classes. We consider in this estimation 353 patent classes that maintain at least five patents per annum. The levels variation is very similar to that found using subcategories in Panel A. The panel estimates are smaller, suggesting a 3%-4% increase in patenting for every 10% increase in expected patenting, but remain quite important economically and statistically. Appendix Figures 8-9 provide visual representations of these class-level estimations.

Panel C of Figure 3 shows a second approach to quantifying the innovation network strength. We regress cumulative actual patenting during 1995–2004 for each class on its expected value based upon the innovation network and a control for historical patenting levels,

$$\ln(P_j^{95-04}) = \beta \ln(\hat{P}_j^{95-04}) + \gamma \ln(P_j^{85-94}) + \varepsilon_j.$$

This approach allows greater variation in how the lag structure of the innovation network impacts current technological change; we now estimate a 10% increase in upstream innovation corresponds to a 3.5% increase in forward patenting. Appendix Figure 10 provides a visual depiction.

⁵For the panel estimations, we plot in these appendix figures the residualized values of actual patenting against predicted patenting. Residualized values are calculated as the unexplained portions of a regression of $\ln(P_{j,t})$ on the fixed effects ϕ_j and η_t (similar process for predicted patenting series). Conveniently, the slope of the trend line in this figure is equal to β .

This cumulative approach is a good platform for robustness checks and extensions. Our first check is to compare our expected patenting growth due to upstream stimulus with a parallel metric developed using downstream stimulus. Our account emphasizes the upstream contributions flowing through the innovation network, but it is natural to worry whether our estimates are instead picking up broad local shocks in technology or a demand-side pull. As the innovation network is asymmetric, we can test this directly, and we confirm in Figure 3 that the upstream flows are playing the central role. Appendix Table 1 documents many additional robustness checks: controlling for parent technology trends, adjusting sample weights, using growth formulations, considering second-generation diffusion,⁶ and so on. The results are robust to dropping any single subcategory, although they depend upon at least some computer and communication fields being retained. We also find these results when using the International Patent Classification system.

Finally, when introducing the $M_{J \times J}$ matrix, we noted two polar cases common to the literature: all entries being equal to $1/J$ (fields building upon a common knowledge stock) or the identity matrix (fields building only on own knowledge). The bottom of Figure 3 and Appendix Table 2 quantify that the truth lies in between—technologies building upon a few key classes that provide them innovation stimulants. We find a robust connection of innovation to the ten most important upstream patent classes, which diminishes afterwards. This relationship is also shown using the subcategory-category structure, although this approach is cruder given the knowledge flows across technology boundaries.⁷ This network heterogeneity indicates that knowledge development is neither global, in the sense that fields collectively share an aggregate pool of knowledge, nor local, in the sense that each field builds only upon itself.

To conclude, our research finds upstream technological developments play an important and measurable role in the future pace and direction of patenting. A better accounting for the innovation network and its asymmetric flows will help us model the cumulative process of scientific discovery in a sharper manner. A better understanding of these features can be an aid to policy makers. For example, the finding that upstream research is highly salient for growth implies that if R&D slackens in one period the effects will be felt years later. This paper has approached

⁶While some network analyses consider high-order relationships (e.g., Leontief inverse in production theory), first-order relationships are sufficient when directly observing intermediating outcomes. As an example, consider $j \rightarrow j' \rightarrow k$, with technology k being upstream from j' . As we directly model patenting in technology j' to downstream outcomes in j , we have already included any potential upstream stimulus from k . Appendix Table 1 shows similar results using second-order diffusion when excluding the first-order relationship.

⁷The top 20 upstream classes account for 80% of citations and are distinct from subcategories. Among the top ten, 27% of citations come from the same subcategory, and another 27% come from other subcategories within the same category. Among the next ten, these figures are 16% and 30%, respectively.

these issues in a setting that considers all patents and inventions, the development of which might be thought of as normal or regular science and innovation. An interesting path for future research is to consider whether large leaps behave in a similar format to that depicted here. We also believe this can be pushed to consider regional and firm-level variation, which can further help us understand the causal impact of patenting on economic and business outcomes.

References

- [1] Romer P (1990) Endogenous technological change. *Journal of Political Economy* 98(5):1002–1037.
- [2] Aghion P, Howitt P (1992) A model of growth through creative destruction. *Econometrica* 60(2):323–351.
- [3] Fleming, L (2001) Recombinant uncertainty in technological search. *Management Science* 47(1):117–132.
- [4] Wuchty S, Jones B, Uzzi B (2007) The increasing dominance of teams in production of knowledge. *Science* 316(5827):1036–1039.
- [5] Azoulay P, Ding W, Stuart T (2009) The impact of academic patenting on the rate, quality and direction of (public) research output. *Journal of Industrial Economics* 57(4):617–676.
- [6] Furman J, Stern S (2011) Climbing atop the shoulders of giants: The impact of institutions on cumulative research. *American Economic Review* 101(5):1933–1963.
- [7] Lerner J, Stern S (2012) *The Rate and Direction of Inventive Activity Revisited* (National Bureau of Economics, Cambridge, MA).
- [8] Stephan P (2012) Foreign born scientists: mobility patterns for 16 countries. *Nature Biotechnology* 30:150-153.
- [9] Bloom N, Schankerman M, Van Reenen J (2013) Identifying technology spillovers and product market rivalry. *Econometrica* 81(4):1347–1393.
- [10] Williams H (2013) Intellectual property rights and innovation: evidence from the human genome. *Journal of Political Economy* 121(1): 1–27.

- [11] Freeman R, Huang W (2014) Strength in diversity. *Nature* 513(7518):305.
- [12] Konig M, Liu Z, Zenou Y (2014) R&D networks: Theory, empirics and policy implications. University of Zurich Working Paper No. 142.
- [13] Glasco A, Schankerman M (2015) Patents and cumulative innovation: causal evidence from the courts. *The Quarterly Journal of Economics* 130(1):317–369.
- [14] Hall B, Jaffe A, Trajtenberg M (2001) The NBER patent citations data file: lessons, insights, and methodological tools. National Bureau of Economic Research Working Paper No. 8498.
- [15] Scherer FM (1984) Using linked patent data and R&D data to measure technology flows. *R & D, Patents and Productivity*, ed Griliches Z (The University of Chicago Press, Chicago, IL), pp 417-464.
- [16] Verspagen B (1997) Measuring intersectoral technology spillovers: estimates from the European and US patent office databases. *Economic Systems Research* 9(1):47-65.
- [17] Schnitzer M, Watzinger M (2015) Measuring the spillovers of venture capital. Working Paper.
- [18] Cohen W, Levinthal D (1990) Absorptive capacity: a new perspective on learning and innovation. *Administrative Science Quarterly* 35(1):128–152
- [19] Belderbos R, Mohen P (2012) Intersectoral and international R&D spillovers. SIMPATIC Working Papers 798, Bruegel.

Figure 1: Citation matrix 1975-1984. Each row describes the field composition of citations made by the technology subcategory indicated on the left hand side. Entries across cited technology fields for each citing technology subcategory sum to 100%. The diagonals—citations of one’s own field, the majority of citations—are excluded from the calculation but given dark shading for reference. Appendix Figure 1 shows the 1975-2004 network and additional sub-periods.

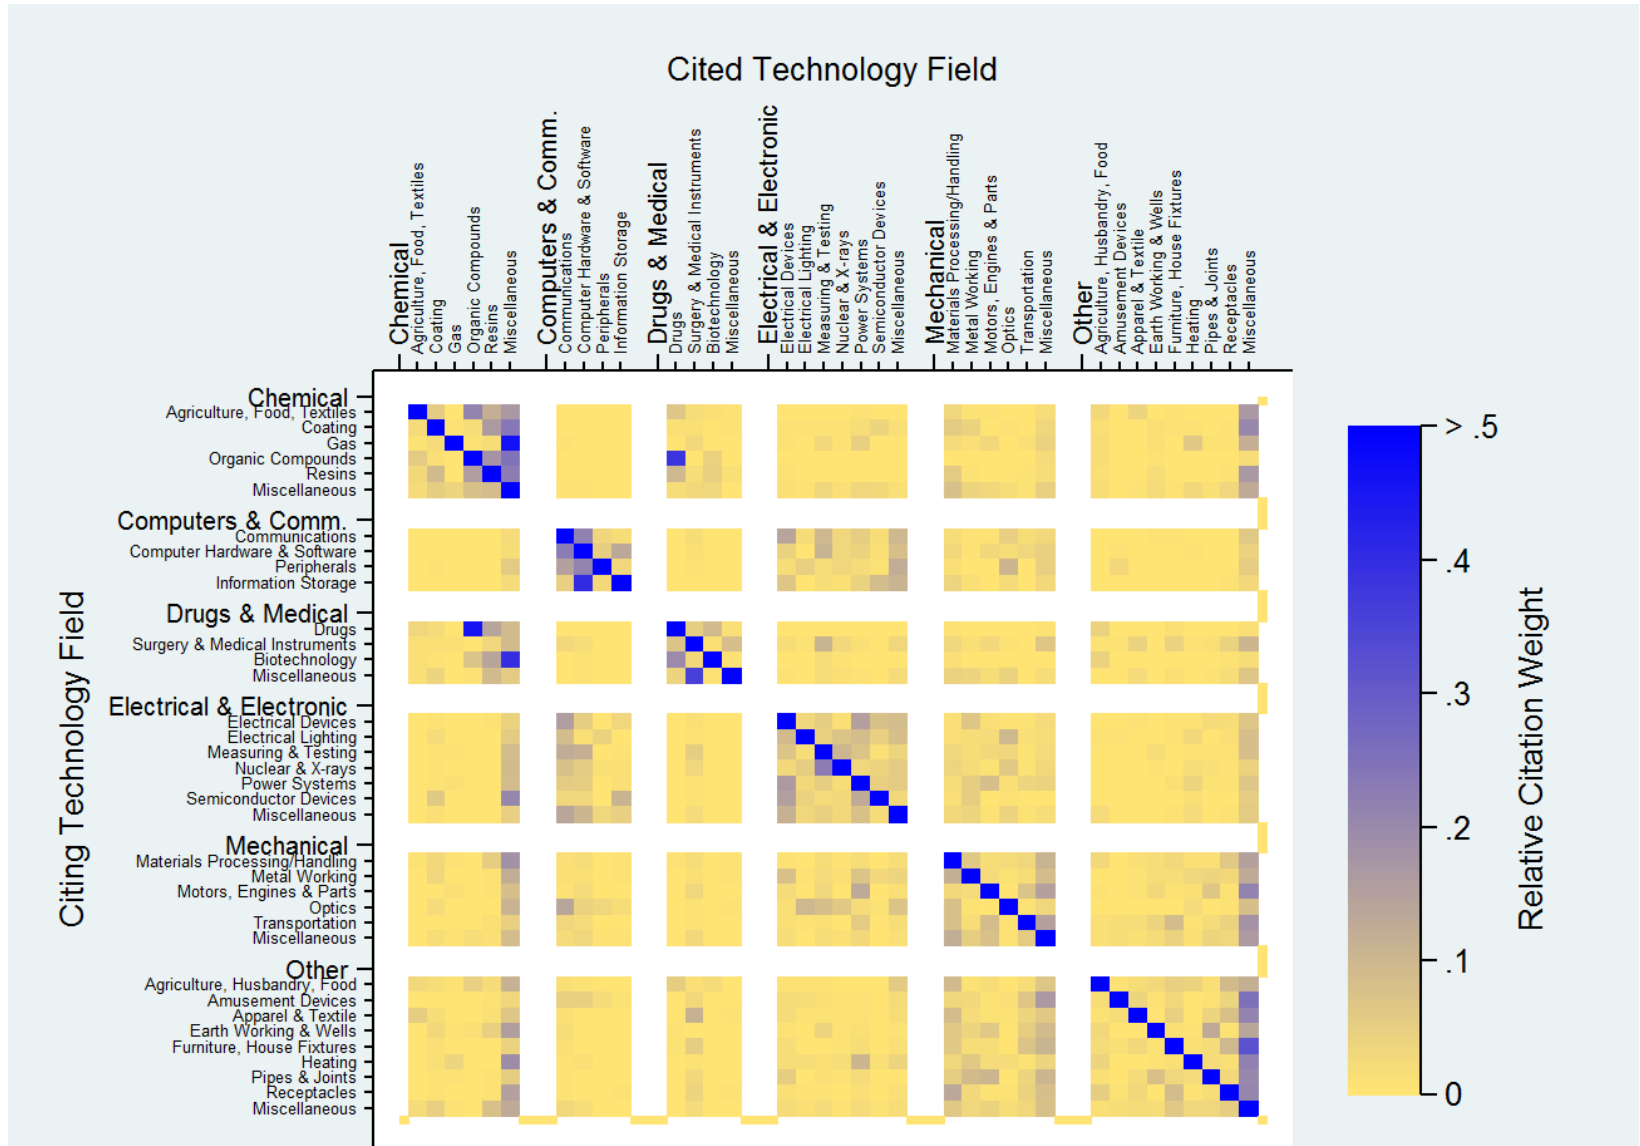
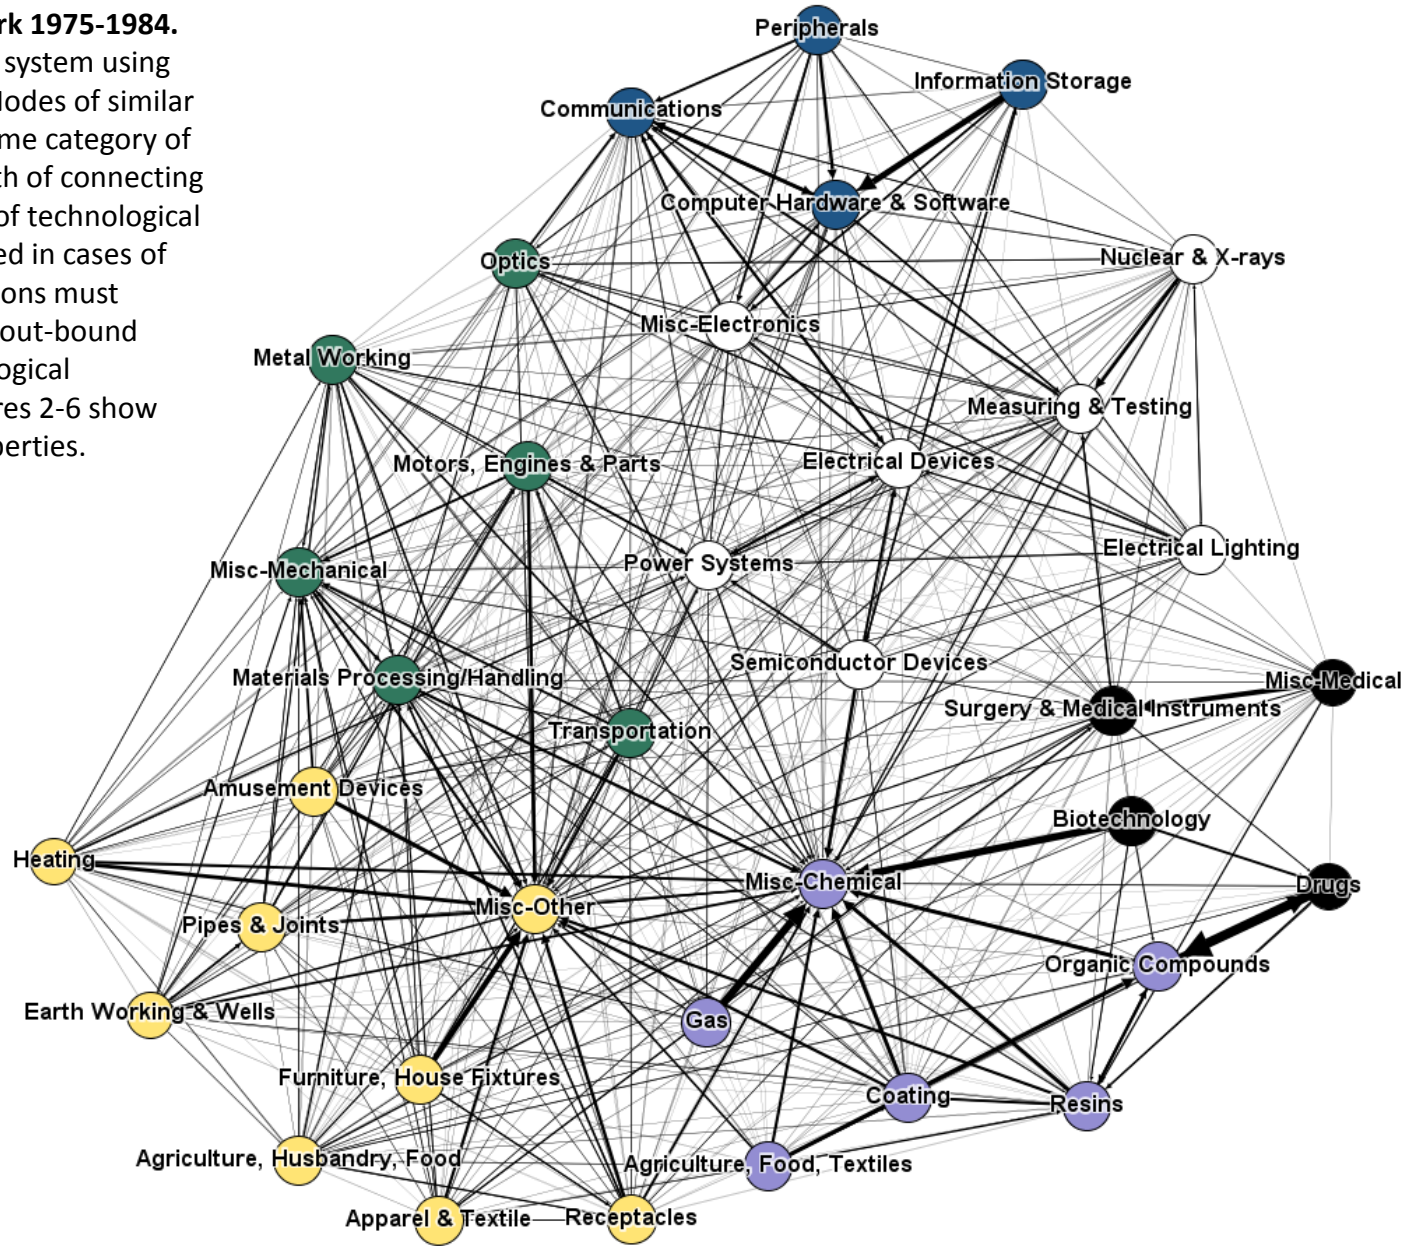


Figure 2: Innovation network 1975-1984.

Network mapping of patent system using technology subcategories. Nodes of similar color are pulled from the same category of the USPTO system. The width of connecting lines indicates the strength of technological flows, with arrows being used in cases of strong asymmetry. Connections must account for at least 0.5% of out-bound citations made by a technological subcategory. Appendix Figures 2-6 show variations and network properties.



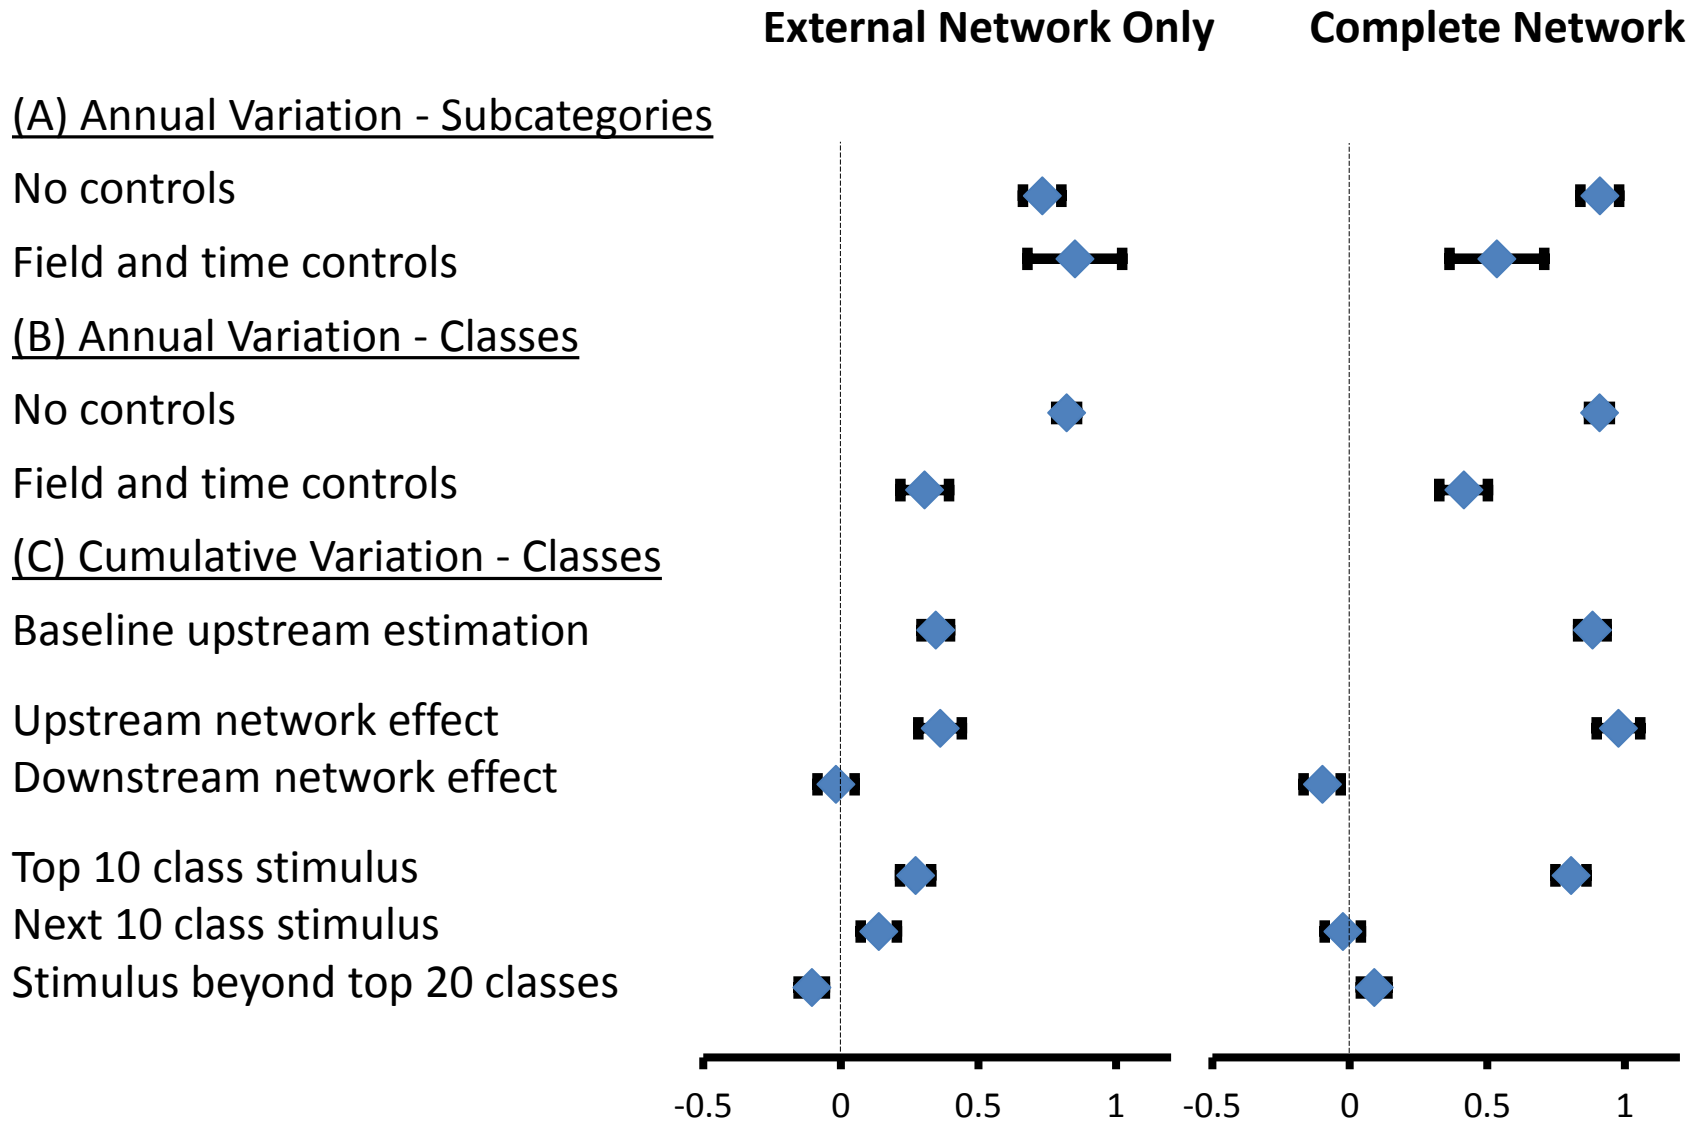
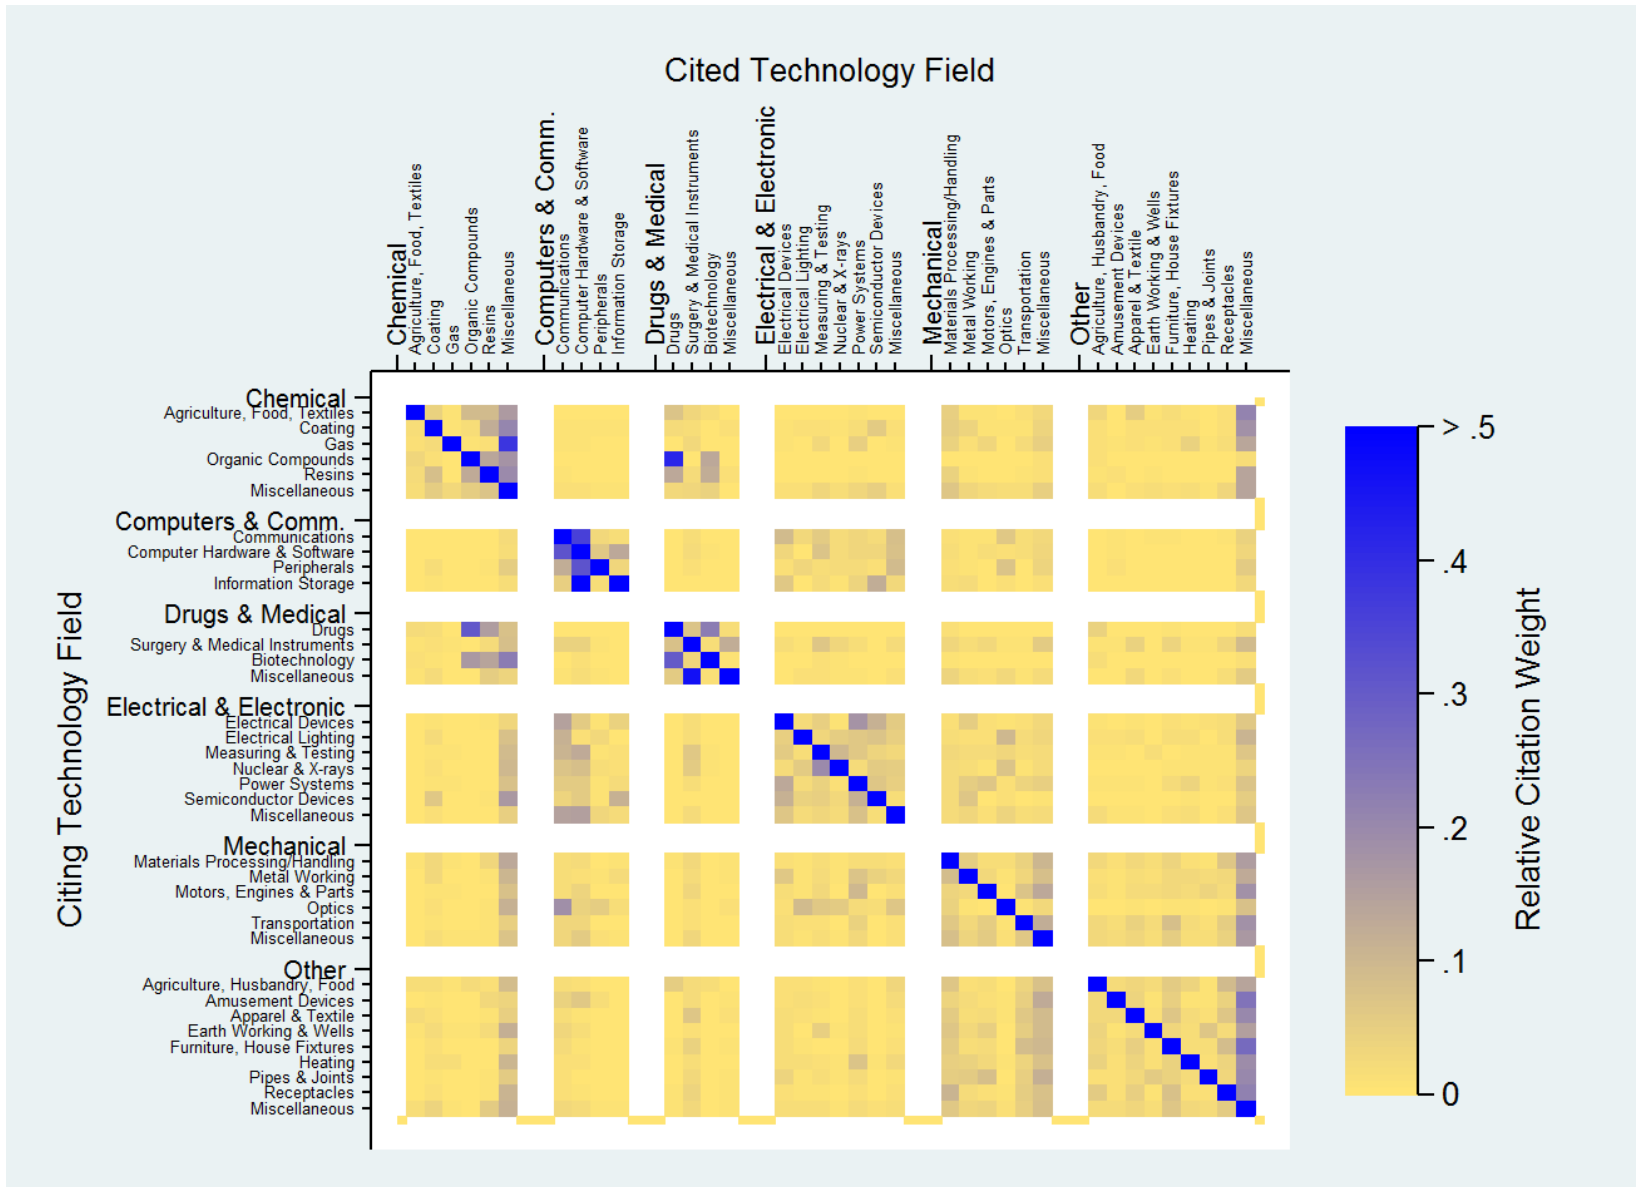


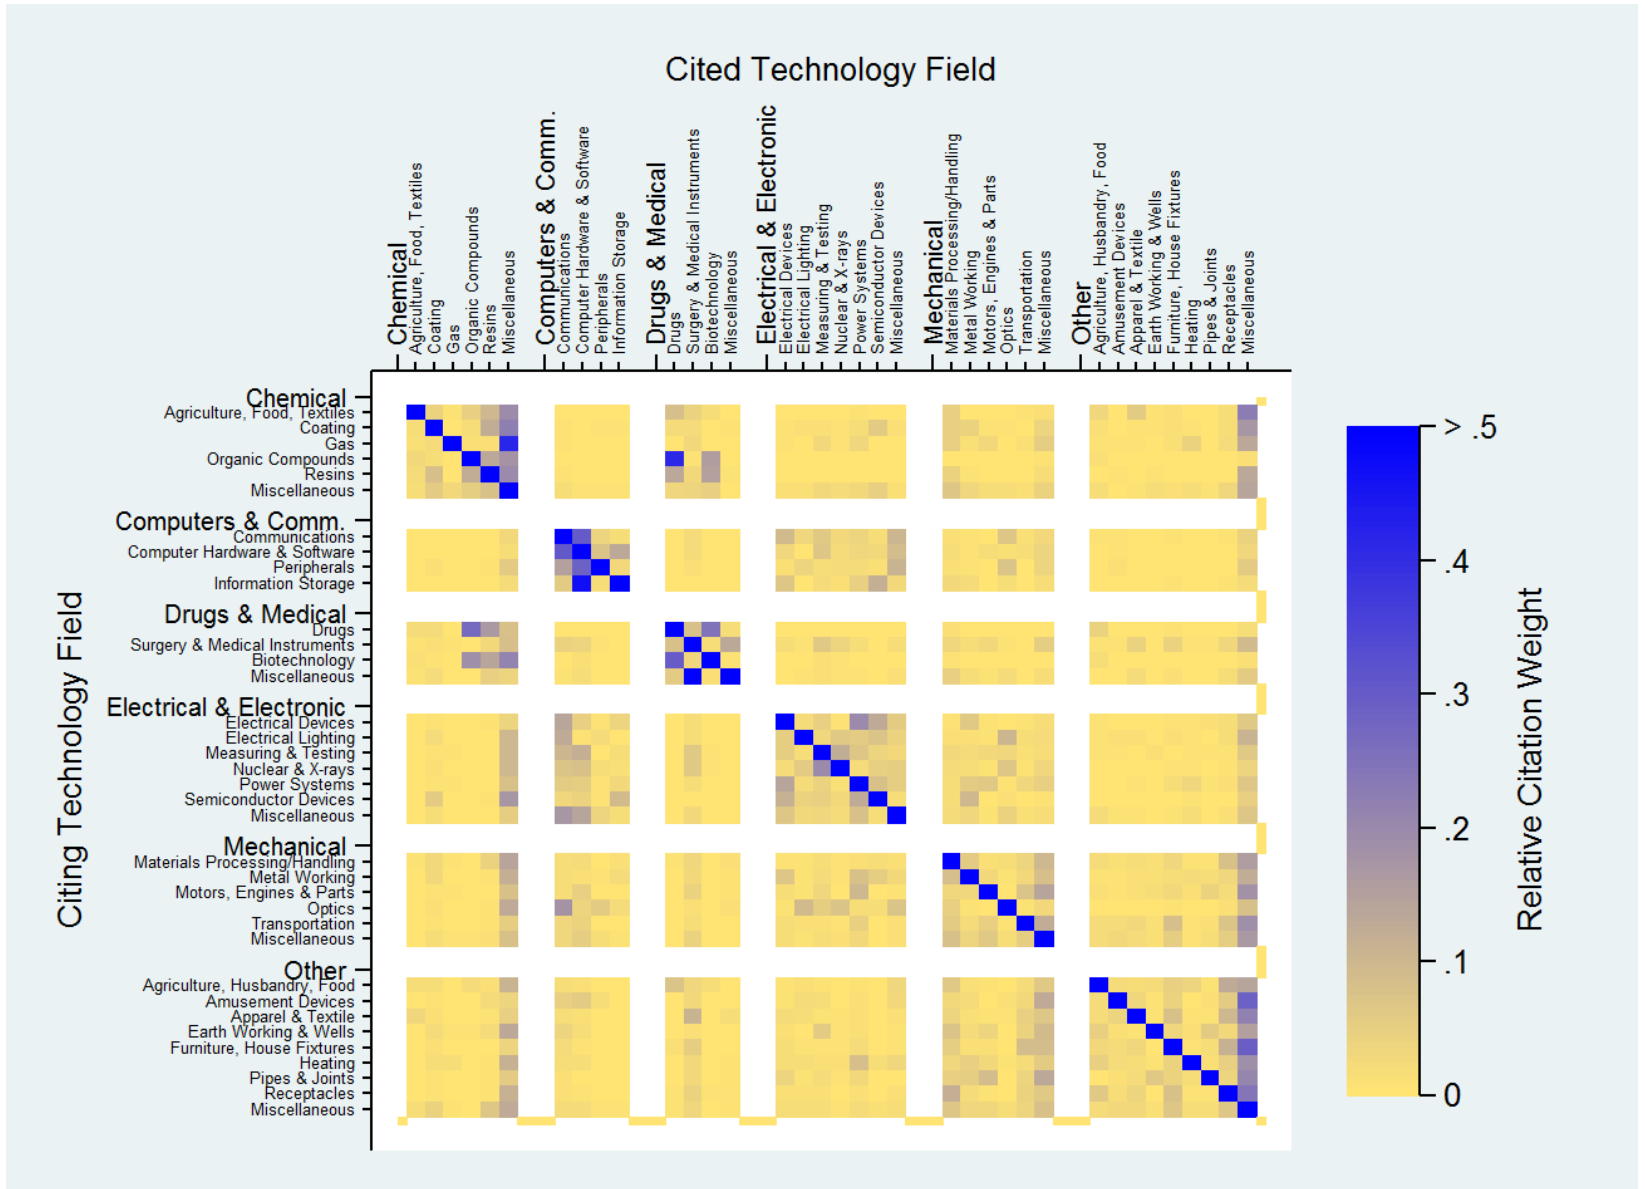
Figure 3: (A) reports regressions of actual patenting during 1995-2004 on predicted patenting calculated using the 1975-1994 innovation network and the growth in upstream technology subcategories predating the focal year. “Field and time controls” analysis reports a panel data analysis where we first remove averages from each subcategory and each year from actual and predicted values. In External Network Only analyses, we consider predicted patenting due to upstream patenting outside of the focal patent subcategory. (B) repeats this analysis for detailed patent classes maintaining over five patents per annum. (C) reports regressions using the patent class sample where we calculate cumulative actual and predictive patenting during 1995-2004 for a patent class. After reporting baseline effects in the cumulative format, we contrast the focal upstream effect with a reverse downstream effect. We next disaggregate the stimulus to demonstrate localized spillovers.

Appendix Materials for “The Innovation Network”
Acemoglu, Akcigit and Kerr

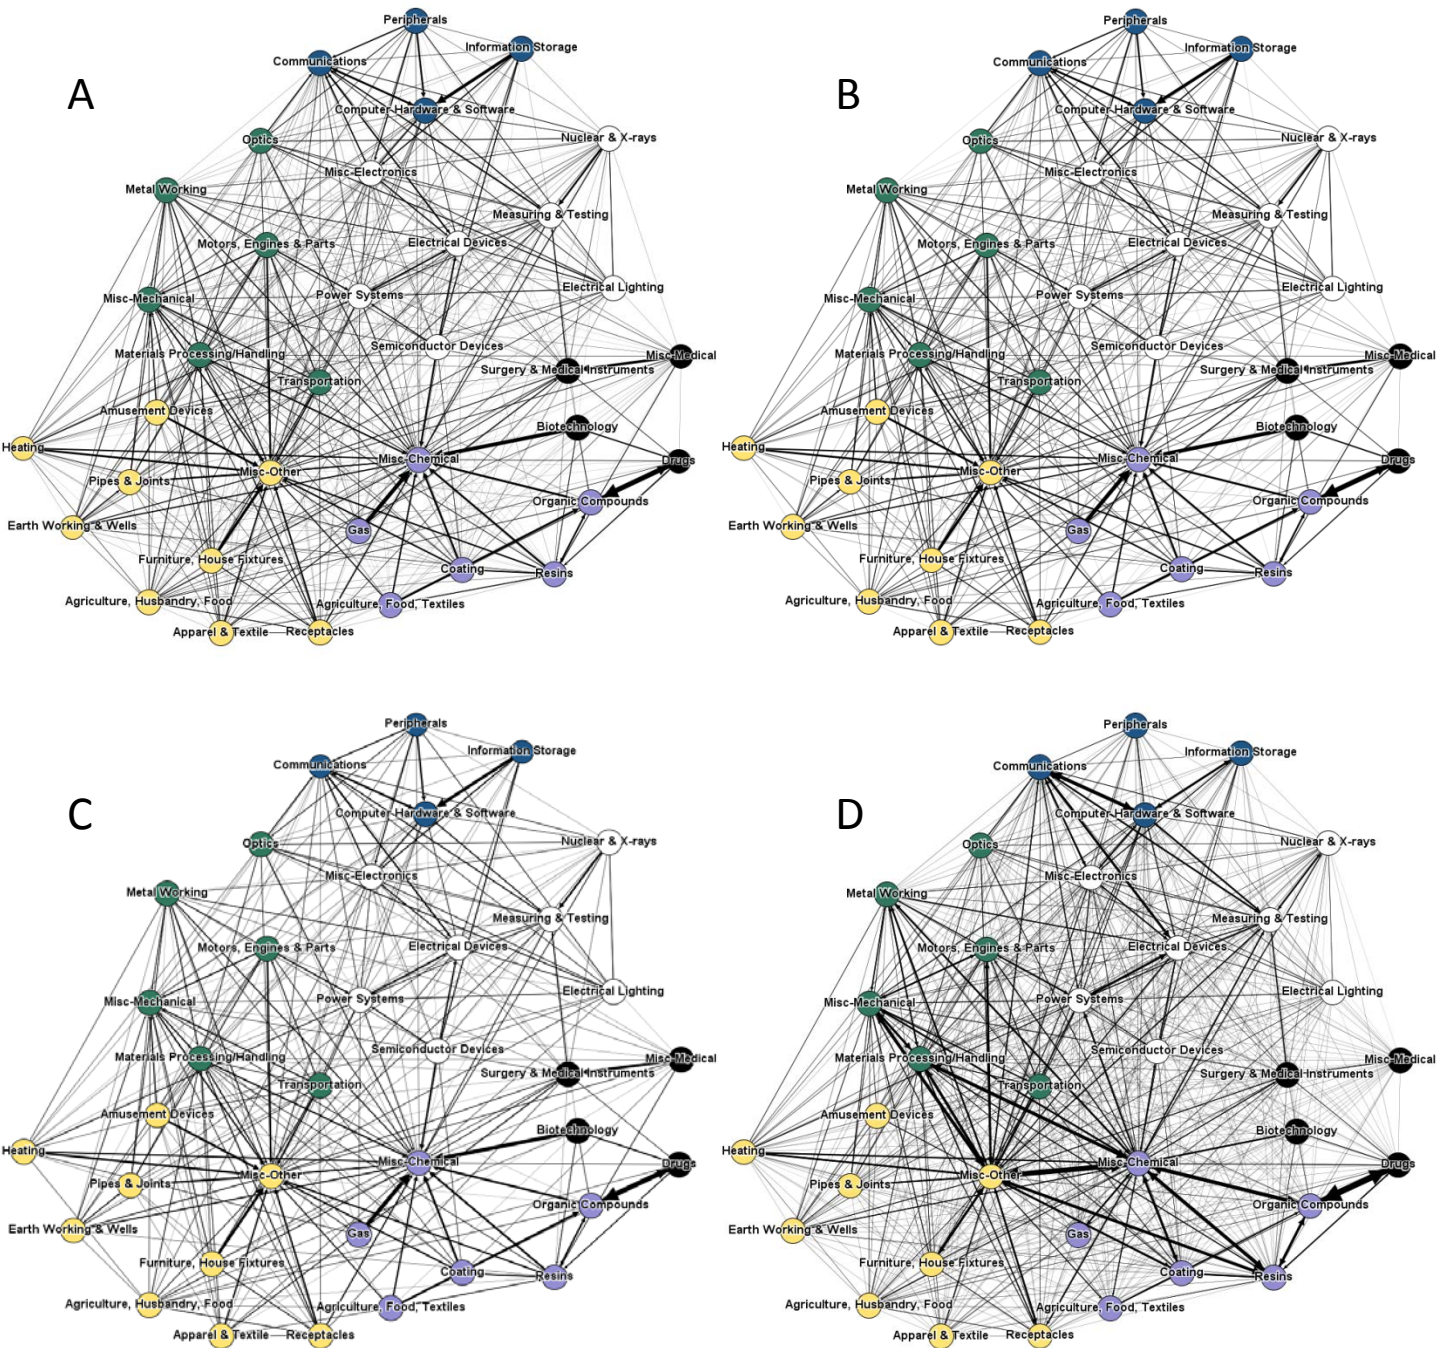
Appendix Figure 1a: Citation matrix 1975-2004. Extends Figure 1 to use all citation information over 1975-2004.



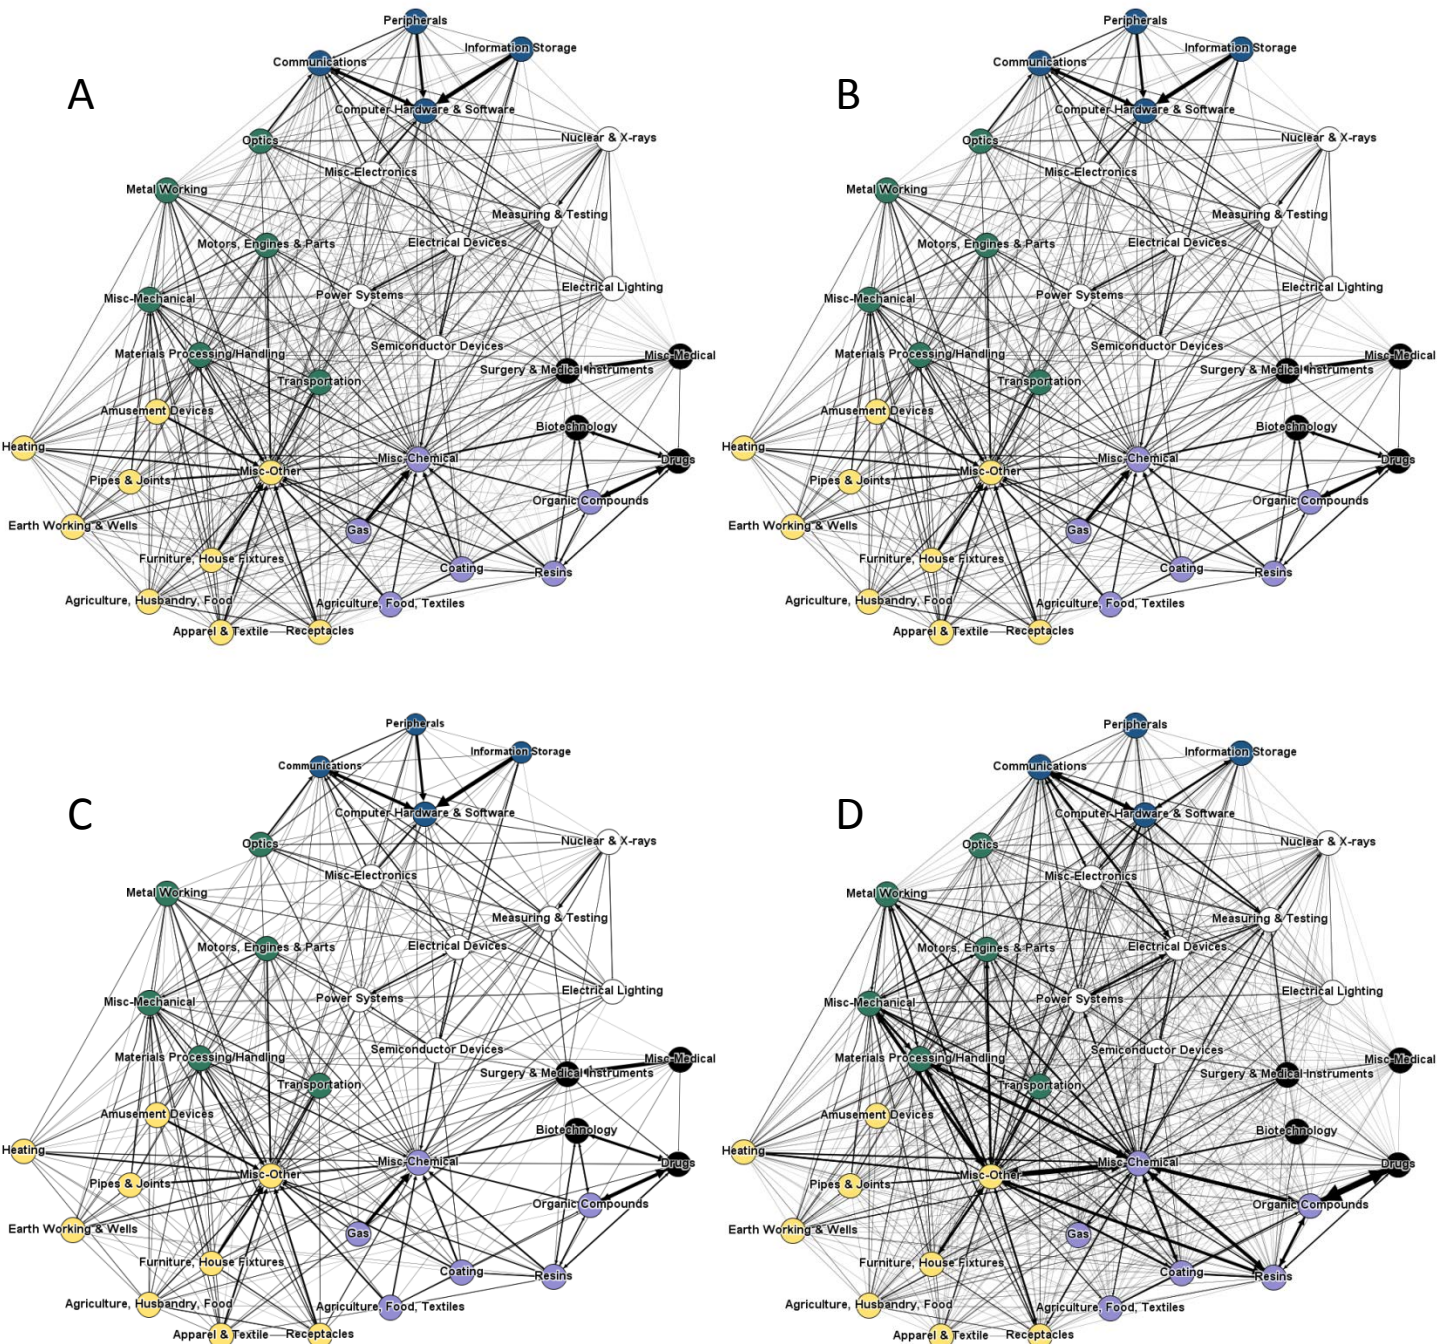
Appendix Figure 1b: Citation matrix 1985-1994.



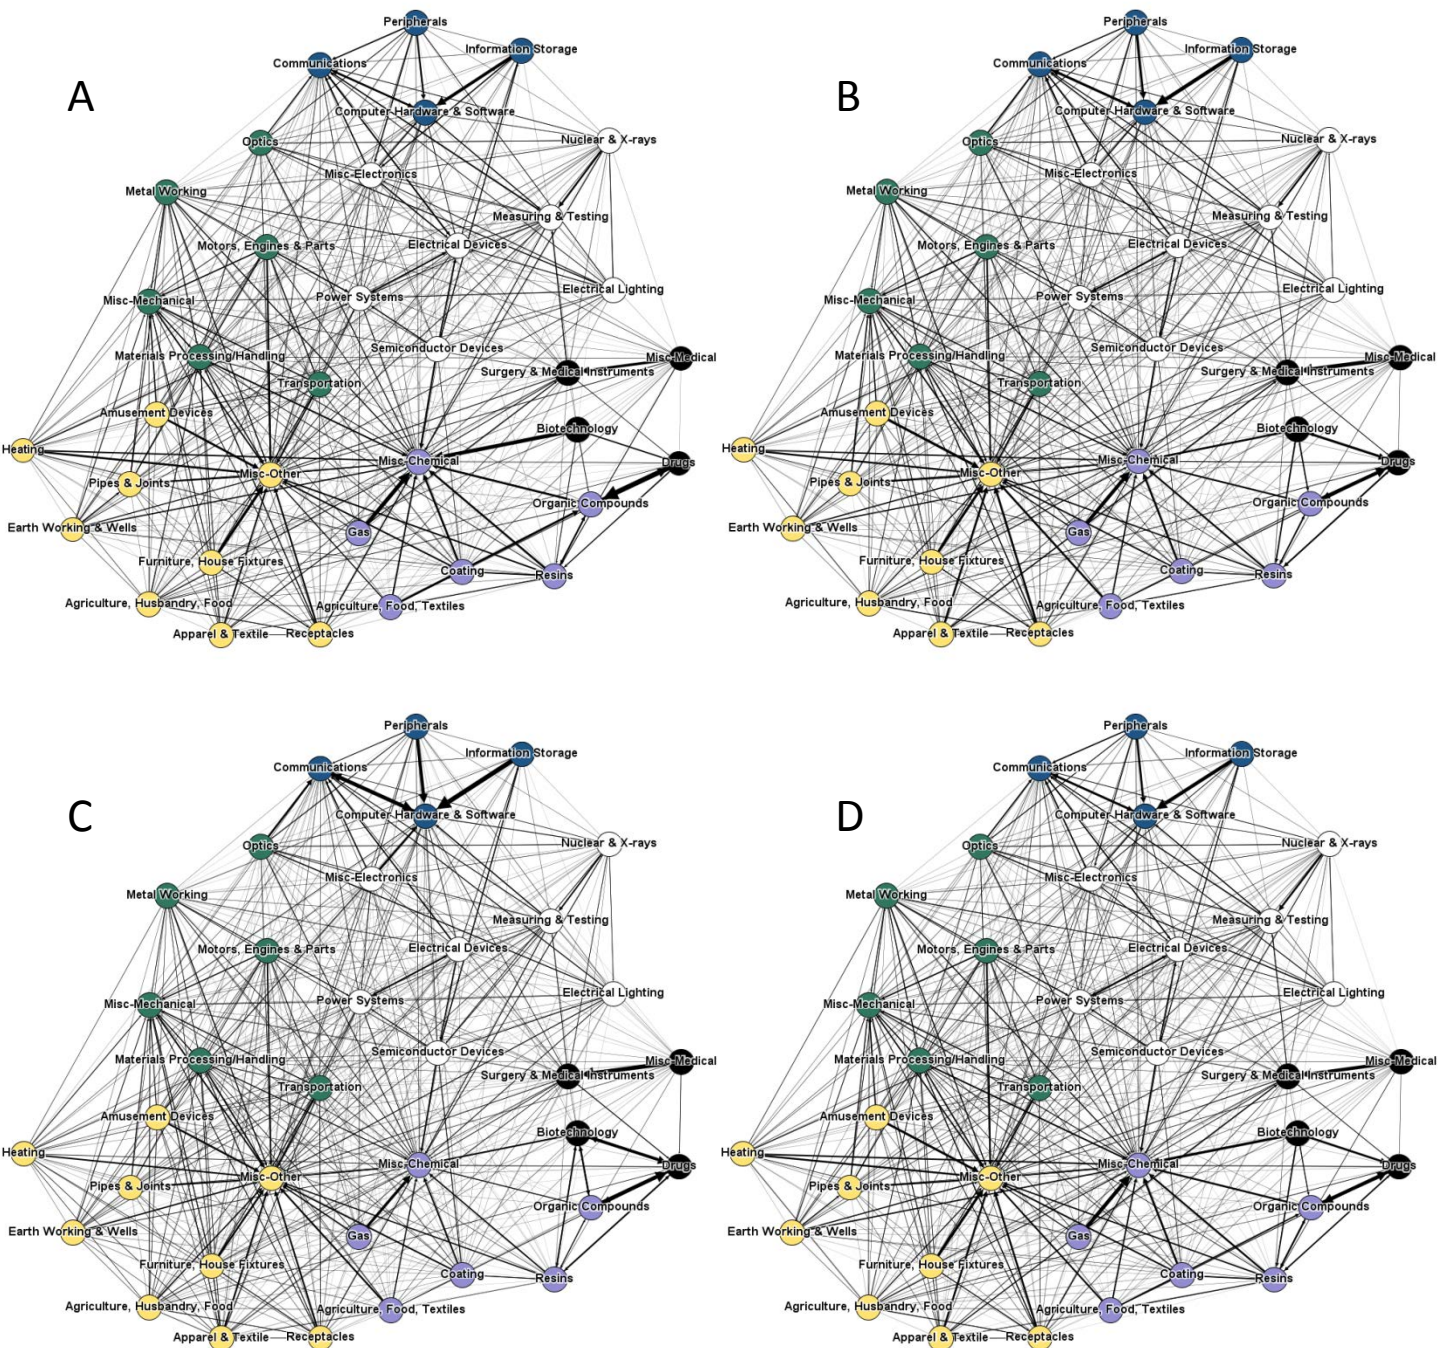
Appendix Figure 2: Innovation network 1975-1984, different edge cut-offs. Extends Figure 2 to illustrate different network cut-offs for the 1975-1984 innovation network. Nodes are colored by their primary technology category. **(A)** uses a 0.5% cut-off for edge inclusion, as shown in the main text. **(B)** uses a 1.0% cut-off for edge inclusion. **(C)** uses a 2.0% cut-off for edge inclusion. **(D)** shows a network without normalization such that every technology's outbound citations are equally weighted.



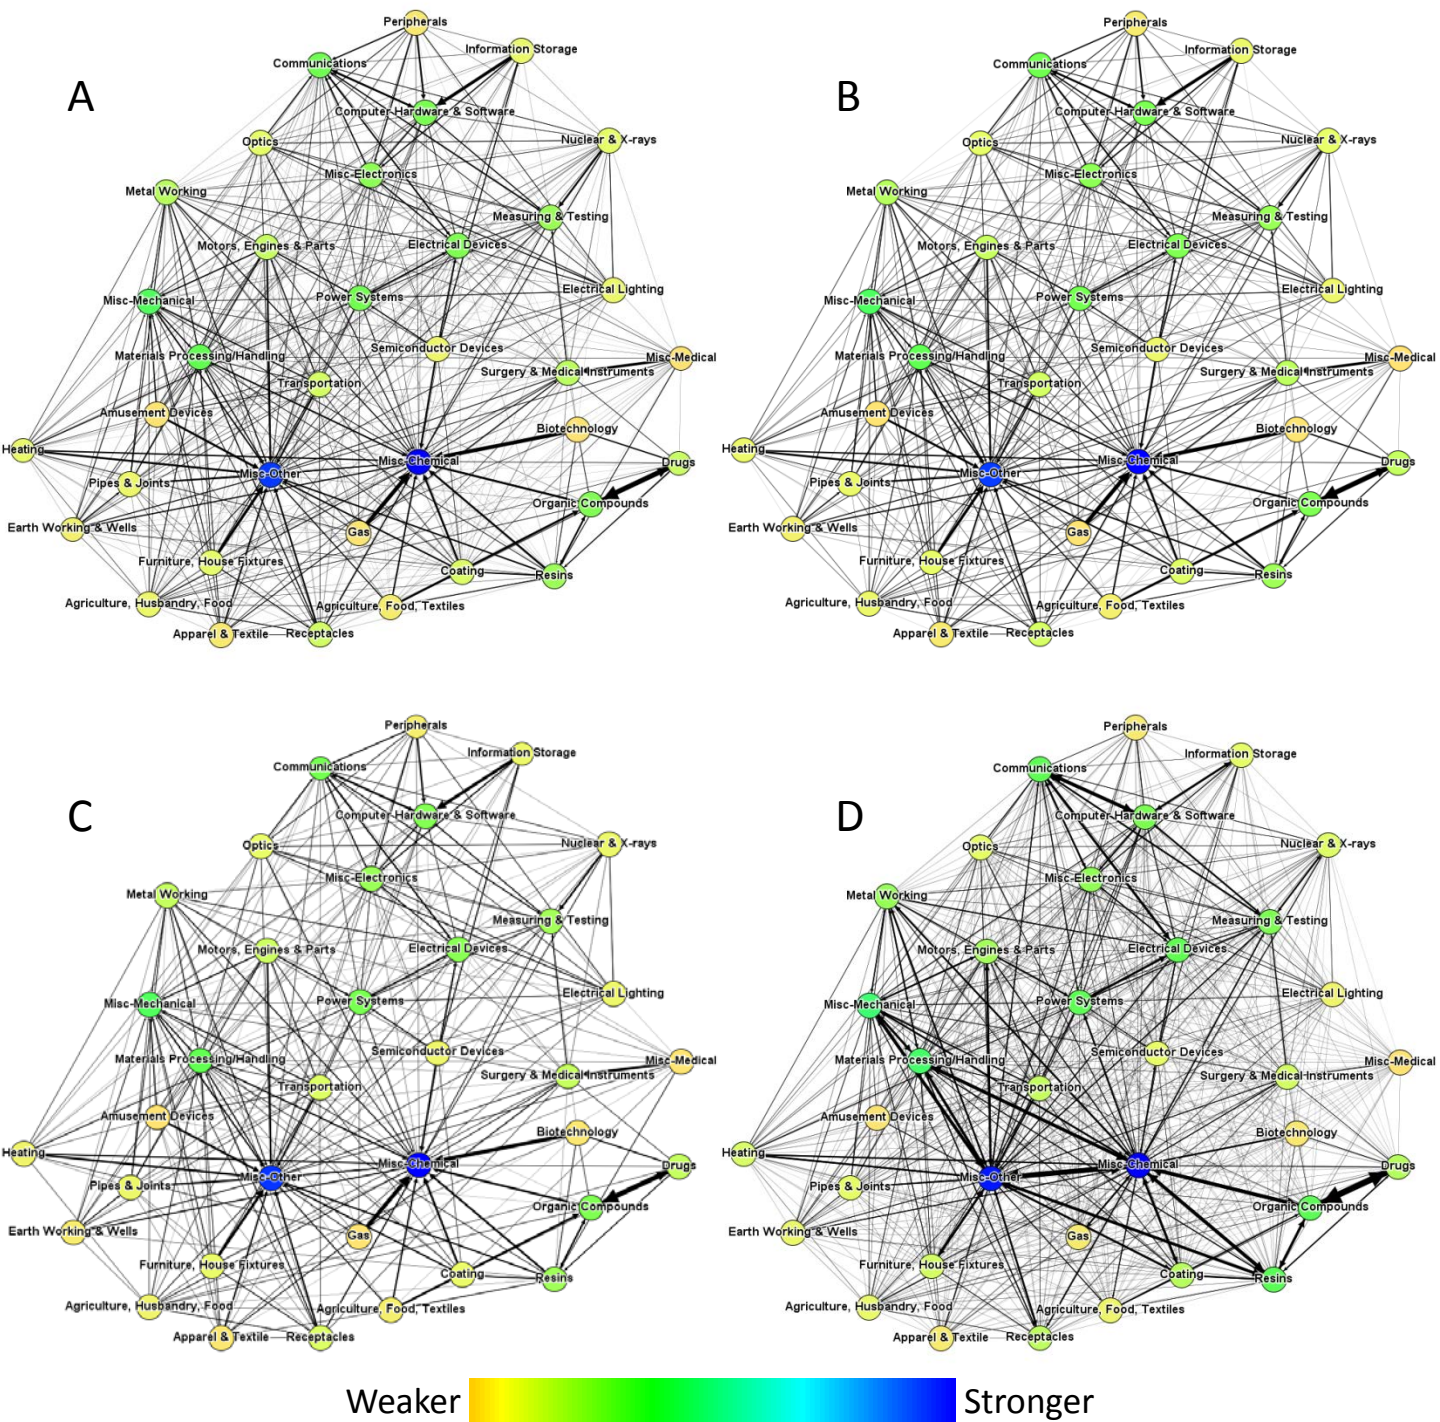
Appendix Figure 3a: Innovation network 1975-2004, full data period. Extends Figure 2 and Appendix Figure 2 to illustrate different network cut-offs for the 1975-2004 innovation network. Nodes are colored by their primary technology category. **(A)** uses a 0.5% cut-off for edge inclusion. **(B)** uses a 1.0% cut-off for edge inclusion. **(C)** uses a 2.0% cut-off for edge inclusion. **(D)** shows a network without normalization such that every technology's outbound citations are equally weighted.



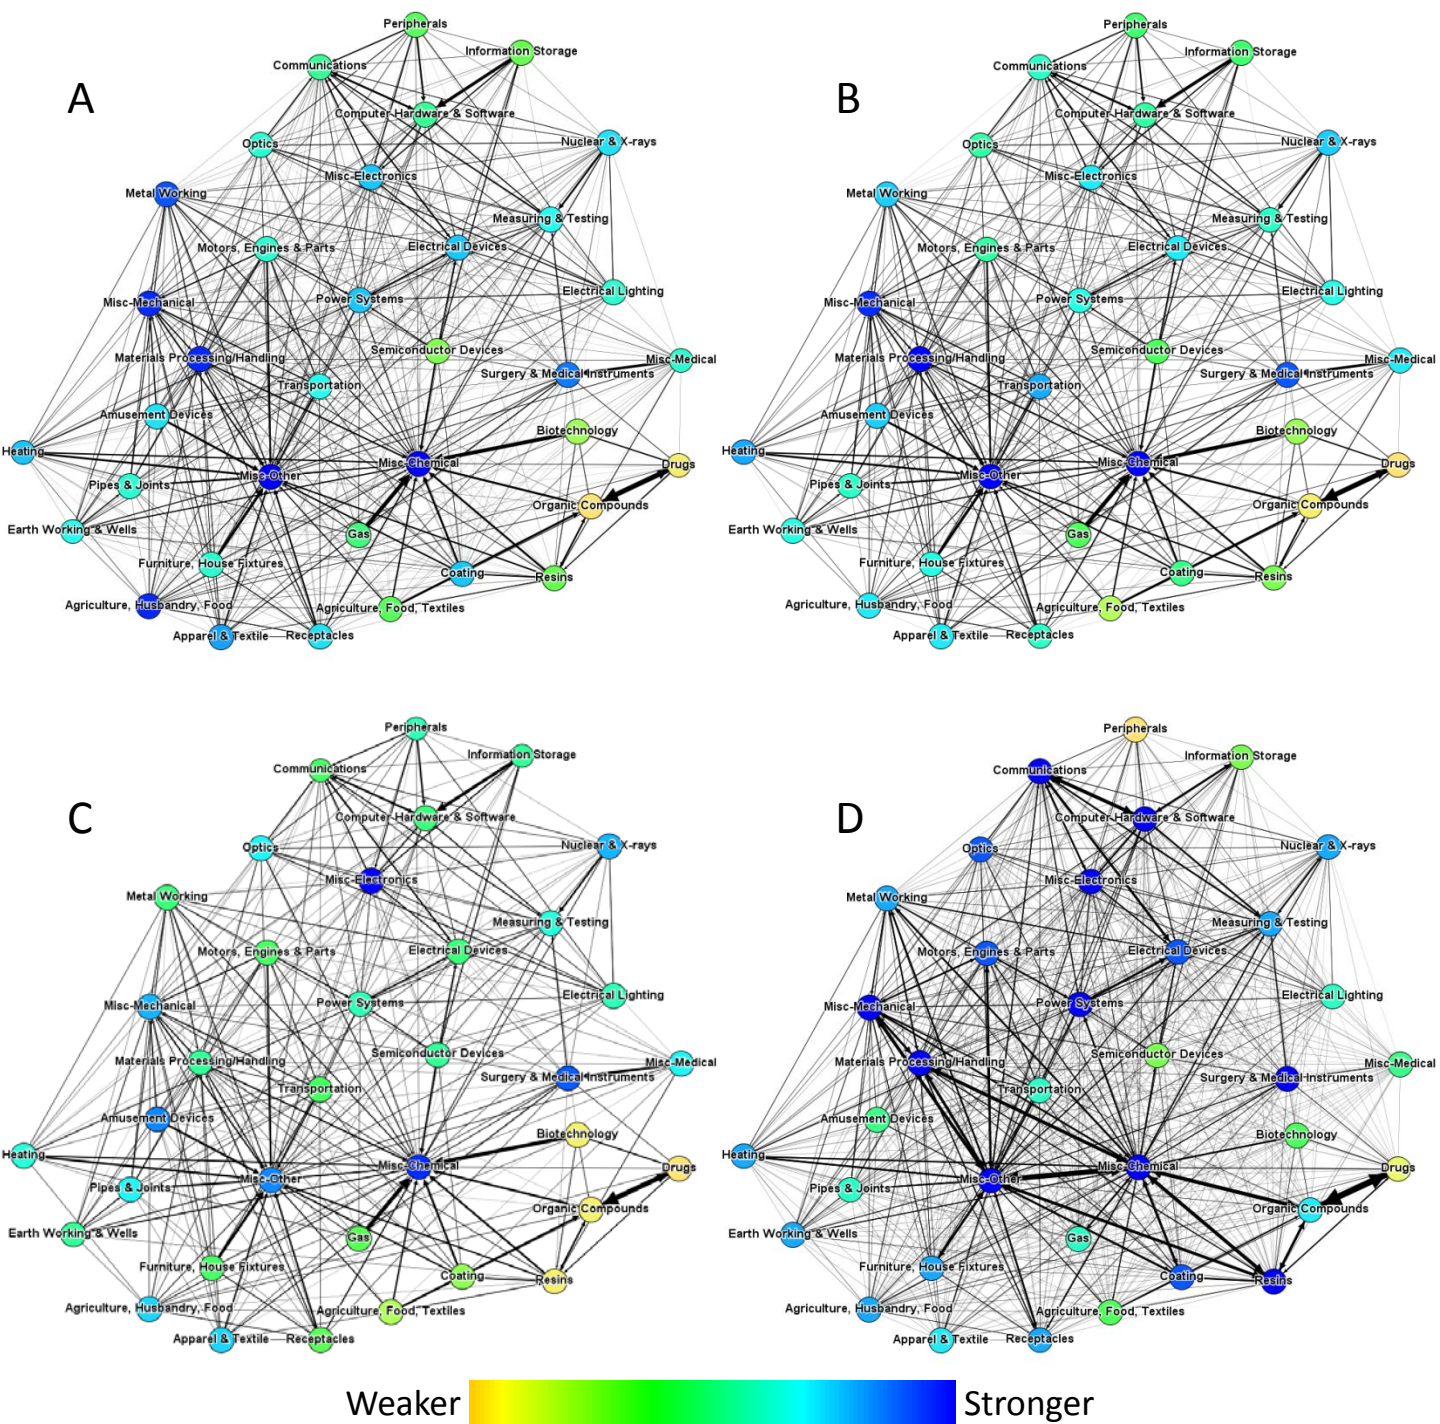
Appendix Figure 3b: Innovation network 1975-2004, by time period. Extends Figure 2 and Appendix Figure 3a to illustrate different time periods and innovation network stability. **(A)** uses 1975-1984, as shown in the main text. **(B)** uses 1985-1994. **(C)** uses 1995-2004. **(D)** uses 1975-1994.



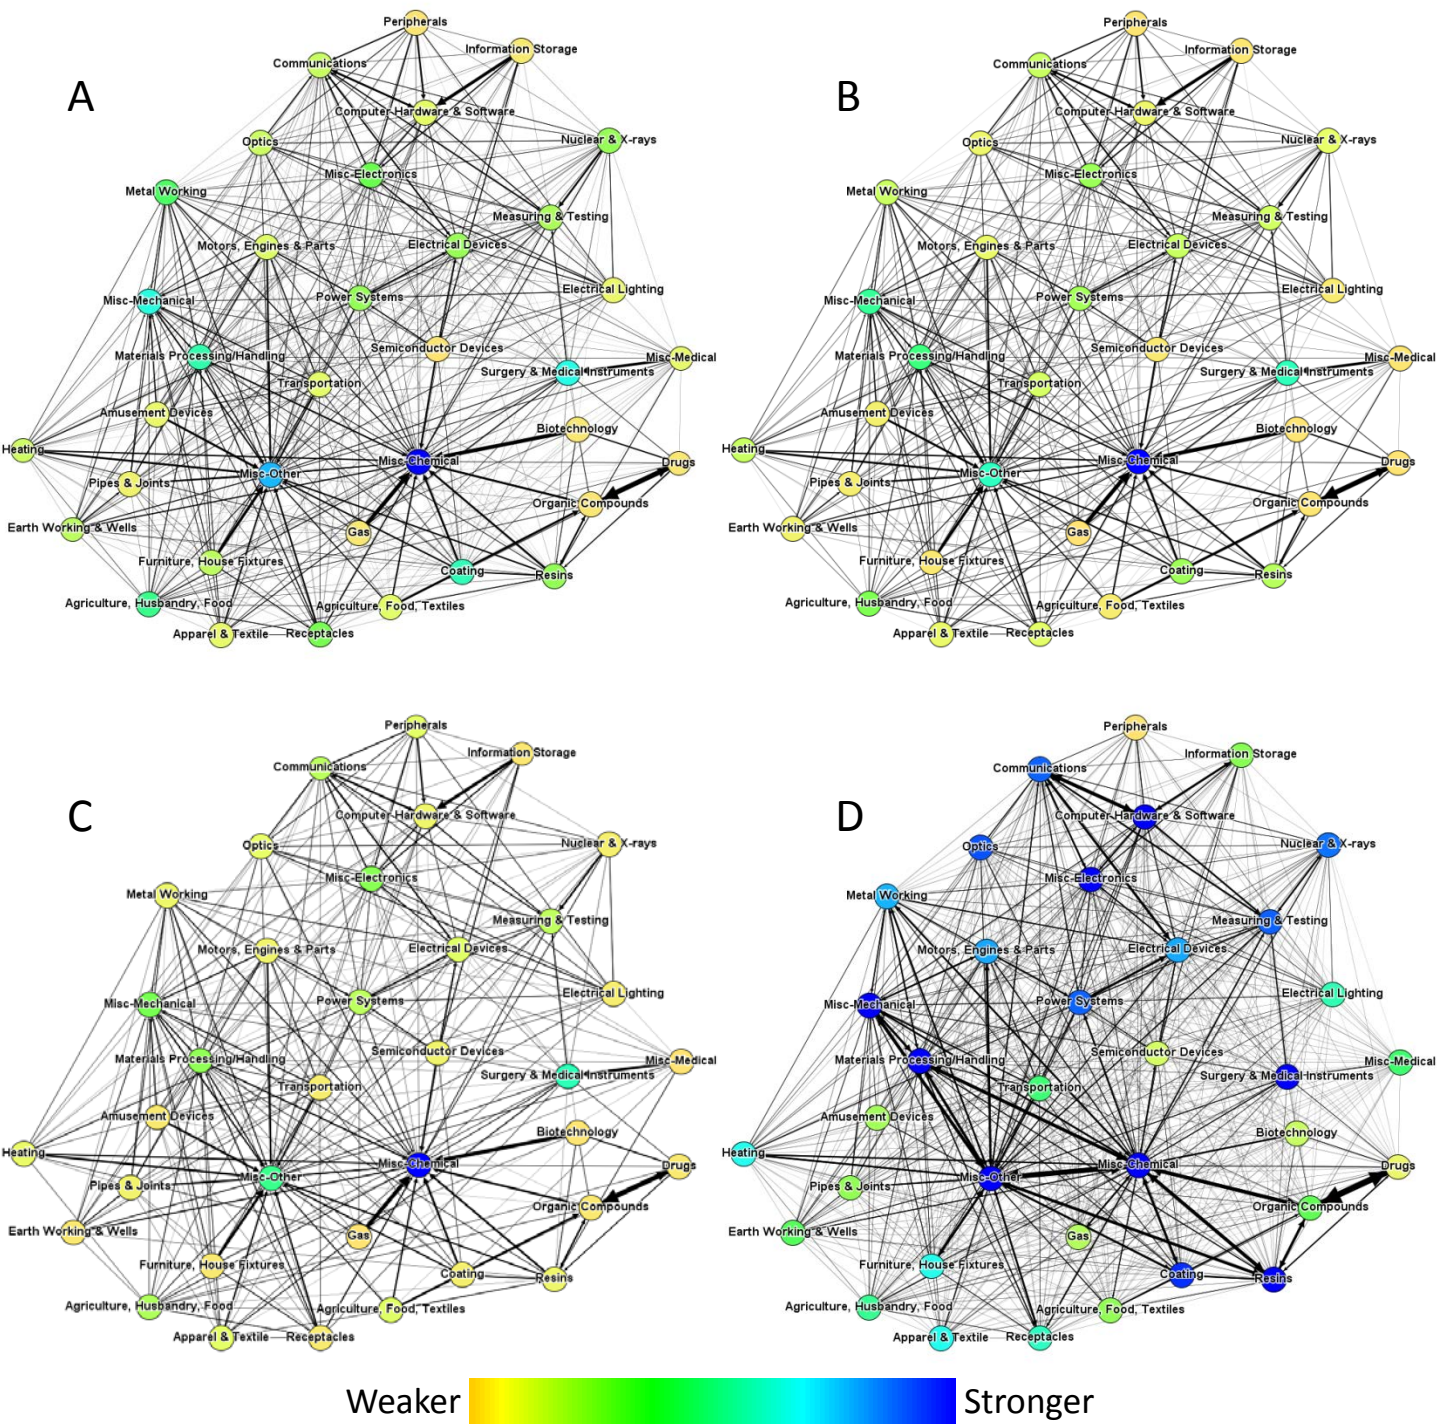
Appendix Figure 4: Innovation network 1975-1984, in-degree importance. The in-degree of a node in the innovation network is the weighted share of patents in other subcategories that are citing patents in the focal node. (The out-degree of a node is not defined in our primary case given the outbound citation normalization for each technology.) Nodes are colored by their importance using the scale at the bottom. **(A)** uses a 0.5% cut-off for edge inclusion. **(B)** uses a 1.0% cut-off for edge inclusion. **(C)** uses a 2.0% cut-off for edge inclusion. **(D)** shows a network without normalization such that every technology's outbound citations are equally weighted.

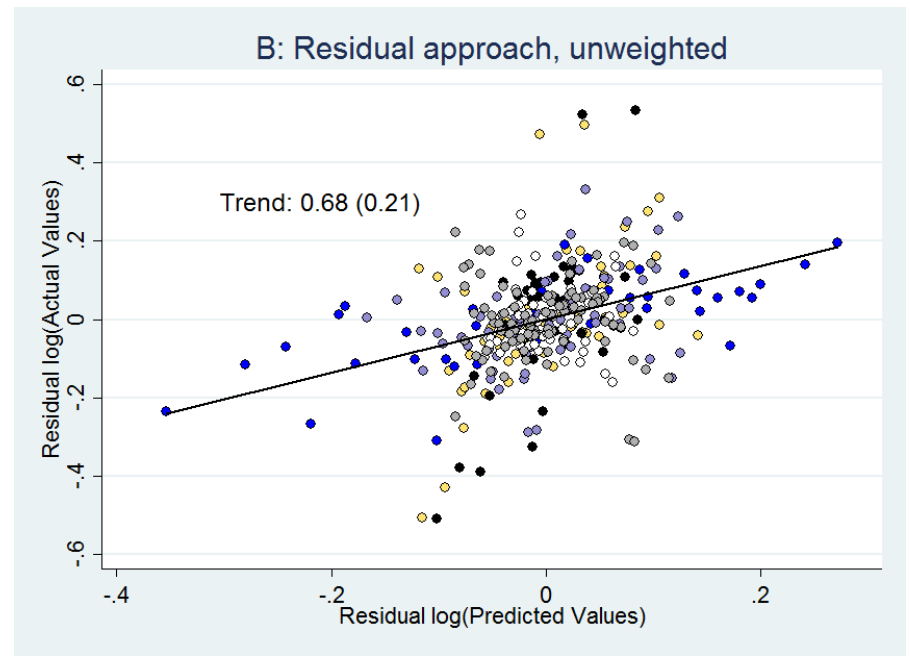
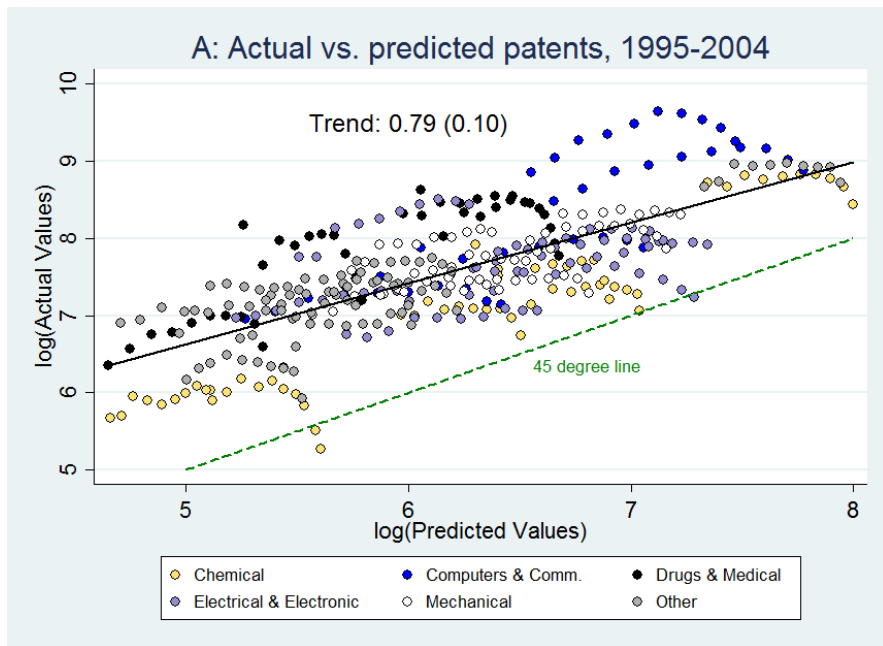


Appendix Figure 5: Innovation network 1975-1984, closeness. The closeness of a node is the average shortest-path distance between that node and all the other nodes in the network. Specifically, $Closeness_n = [1/(g-1)] \sum_m path(n,m)$, where g is the number of nodes in the network and m is any node other than n . A lower closeness score indicates a more important node in the sense of being closer on average to the rest of the nodes in the network. Nodes are colored by their importance using the scale at the bottom. **(A)** uses a 0.5% cut-off for edge inclusion. The closeness scores in the innovation network range from 1.06 (Misc-Chemical and Misc-Other) to 1.75 (Organic Compounds). **(B)** uses a 1.0% cut-off for edge inclusion. **(C)** uses a 2.0% cut-off for edge inclusion. **(D)** shows a network without normalization such that every technology's outbound citations are equally weighted.

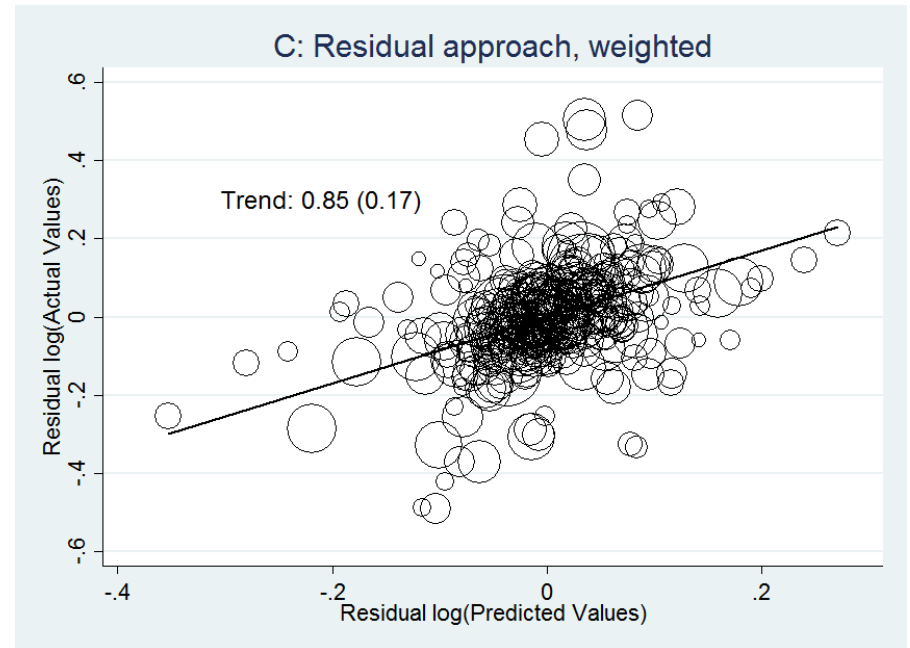


Appendix Figure 6: Innovation network 1975-1984, betweenness. Betweenness measures how often a node links two other nodes. Identifying the shortest paths between every pair of nodes in the graph, the betweenness score of a given node n is the number of those shortest paths that include n (not including paths where n is one of the endpoints). If there are multiple different shortest paths between two nodes, the contribution to the betweenness score of a node on any of the paths is inversely proportional to the number of paths. Nodes are colored by their importance using the scale at the bottom. **(A)** uses a 0.5% cut-off for edge inclusion. The betweenness scores in the innovation network range from 0.21 (Information Storage) to 38.23 (Misc-Chemical). **(B)** uses a 1.0% cut-off for edge inclusion. **(C)** uses a 2.0% cut-off for edge inclusion. **(D)** shows a network without normalization such that every technology's outbound citations are equally weighted.

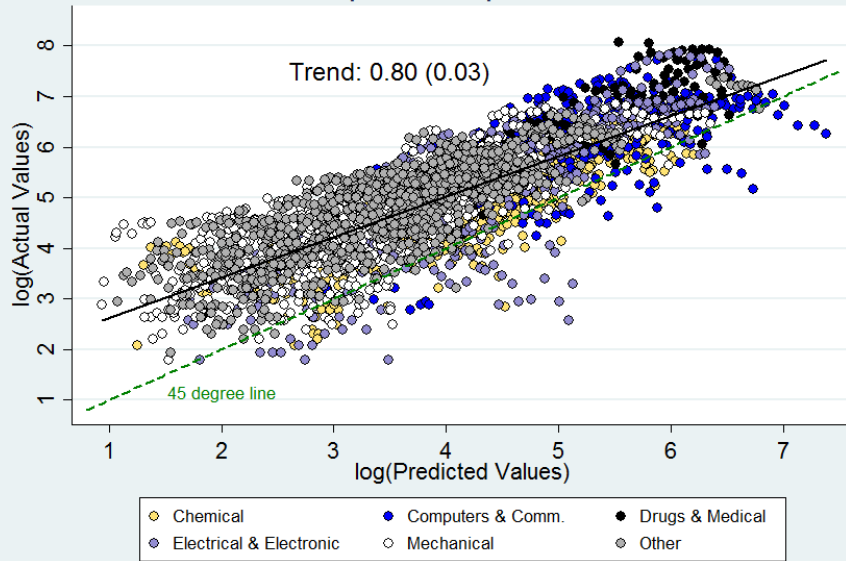




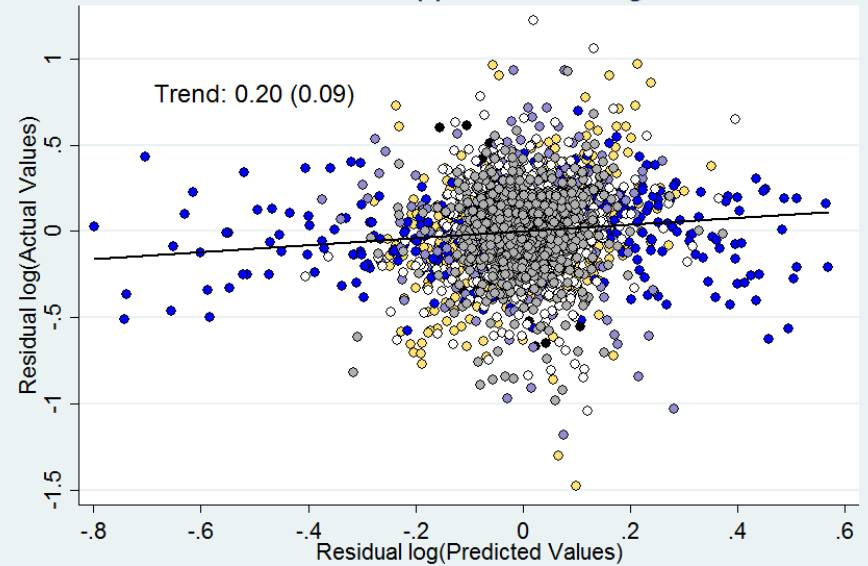
Appendix Figure 7: Network strength at subcategory level. (A) shows actual patenting during 1995-2004 against predicted patenting calculated using the 1975-1994 innovation network and the growth in upstream technology subcategories predating the focal year. Observations represent a subcategory-year pair, color-coded by parent technological category, and the 45-degree line is shown for reference. (B) shows a panel data analysis where we first remove averages from each subcategory and each year from actual and predicted values and plot the residuals. The color scheme is the same as (A). (C) shows a version of (B) that weights observations by their initial size in terms of patenting during 1985-1994. Bubble size indicates weight. This is equivalent to the panel regression findings reported in the text. Reported trend lines are calculated with standard errors clustered by subcategory.



A: Actual vs. predicted patents, 1995-2004

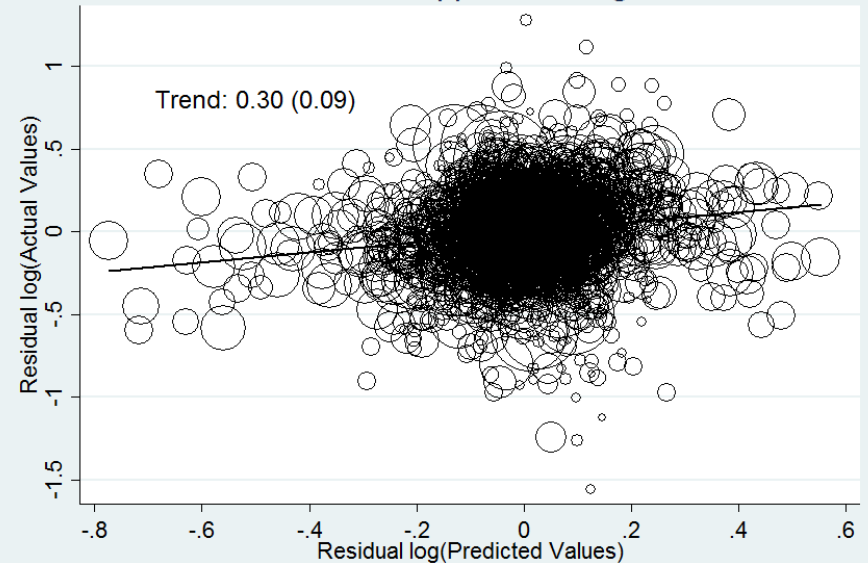


B: Residual approach, unweighted

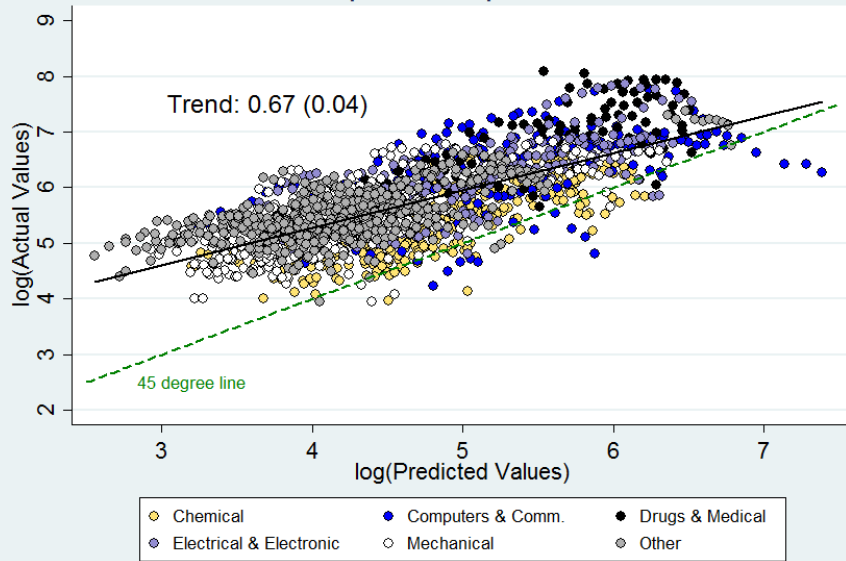


Appendix Figure 8: Network strength at patent class level. Figure repeats Appendix Figure 7 at the patent class level for patent classes maintaining over five patents per annum during 1975-2004.

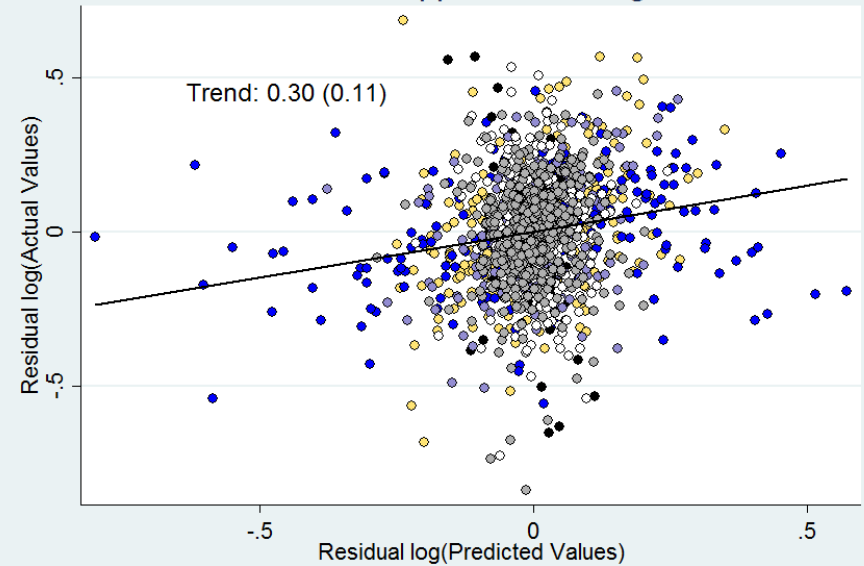
C: Residual approach, weighted



A: Actual vs. predicted patents, 1995-2004

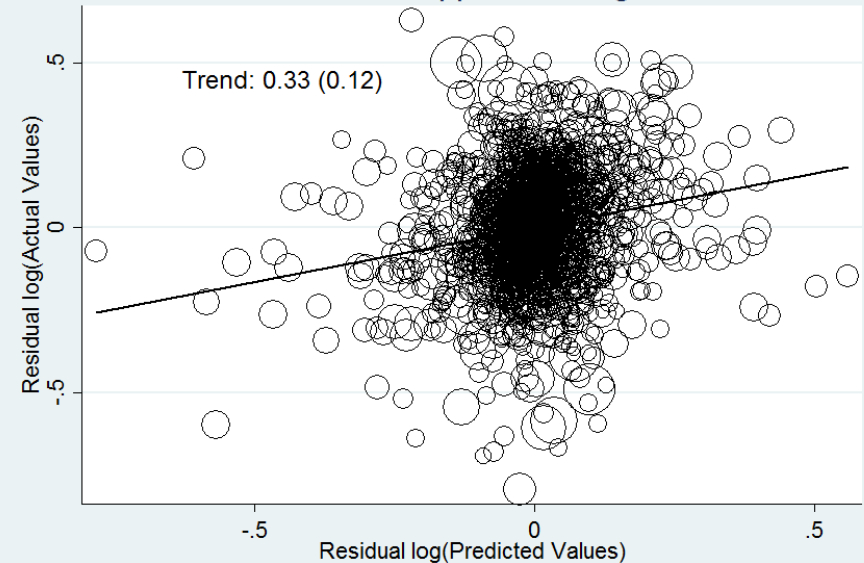


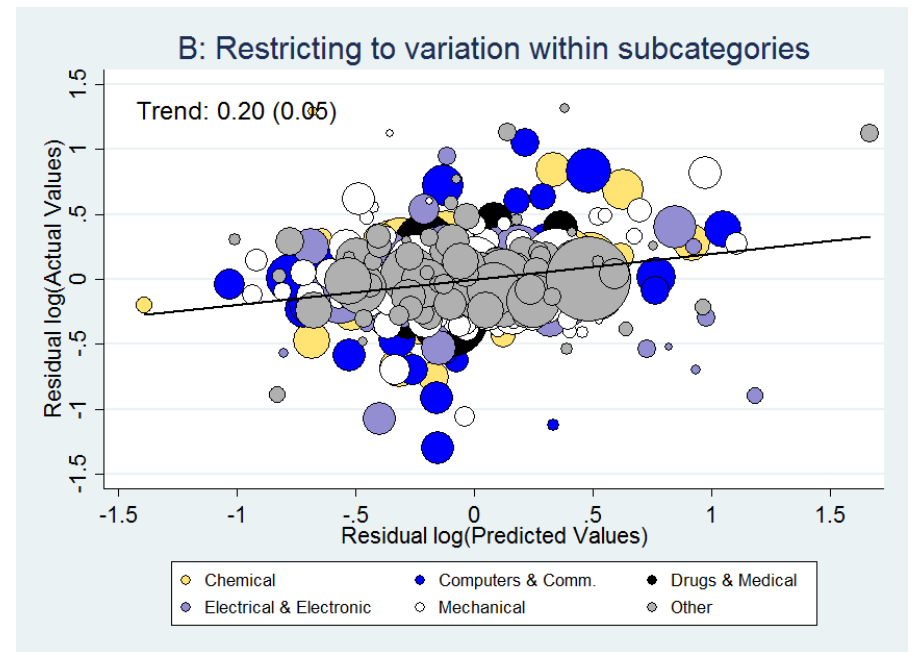
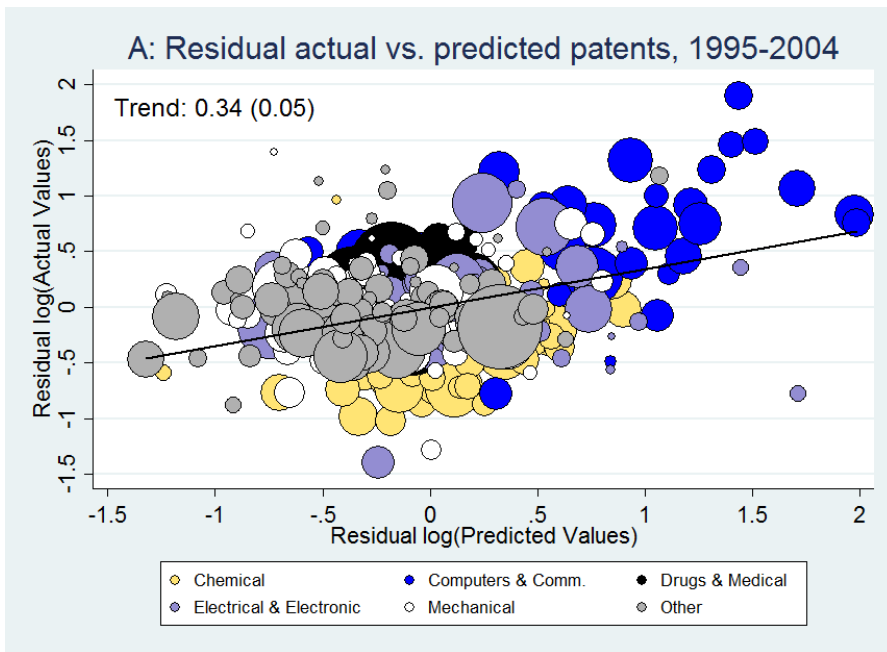
B: Residual approach, unweighted



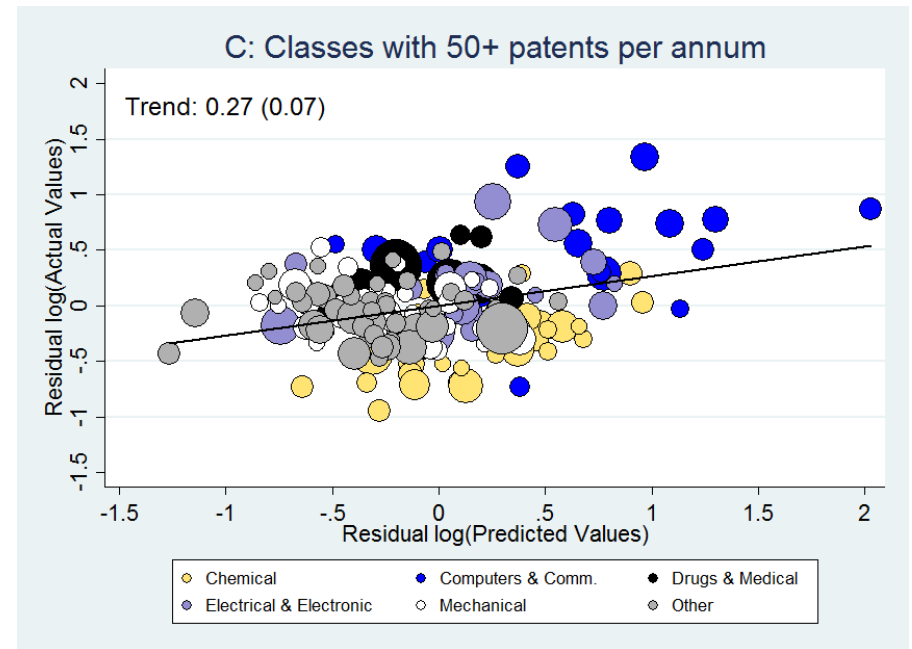
Appendix Figure 9: Network strength at patent class level among large patent classes. Figure repeats Appendix Figure 8 for patent classes maintaining over 50 patents per annum.

C: Residual approach, weighted





Appendix Figure 10: Network strength in cumulative format. (A) shows the panel result in a format where we calculated cumulative actual and predictive patenting during 1995-2004 for a patent class. Figure presents residual values from regressions that control for initial patenting levels by classes during 1985-1994. **(B)** shows a version where residuals are prepared after also including subcategory fixed effects in the regressions, effectively isolating class-level variation within each subcategory. **(C)** restricts panel (A) to patent classes that have 50+ patents granted per year during 1975-2004.



Appendix Table 1: Regression analysis of innovation network

	Base regression	Clustering standard errors by patent subcategory	Including patent category fixed effects	Including patent subcategory fixed effects	Using second application of innovation matrix	Using second-generation citation matrix	Using log patent class weights	Dropping patent class weights	Using constrained growth formulation
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent variable is log cumulative patent counts in patent class 1995-2004									
A. Analysis with 353 patent classes maintaining 5 or more patents per annum									
Log cumulative patents 1985-1994	0.678 (0.061)	0.678 (0.102)	0.779 (0.047)	0.768 (0.048)	0.800 (0.062)	0.785 (0.063)	0.722 (0.057)	0.744 (0.058)	[1.0]
Log expected patenting from network stimulus	0.345 (0.052)	0.345 (0.094)	0.190 (0.045)	0.197 (0.048)	0.217 (0.054)	0.229 (0.051)	0.271 (0.053)	0.242 (0.055)	0.104 (0.028)
B. Analysis with 175 patent classes maintaining 50 or more patents per annum									
Log cumulative patents 1985-1994	0.785 (0.083)	0.785 (0.097)	0.829 (0.060)	0.739 (0.070)	0.923 (0.083)	0.915 (0.082)	0.816 (0.072)	0.827 (0.071)	[1.0]
Log expected patenting from network stimulus	0.265 (0.066)	0.265 (0.083)	0.145 (0.052)	0.184 (0.060)	0.130 (0.067)	0.133 (0.060)	0.232 (0.060)	0.223 (0.060)	0.121 (0.038)

Notes: Table quantifies the relationship between log patenting 1995-2004 by patent class and the stimulus provided by the innovation network. Estimations control for log initial patenting in the patent class during 1985-1994. Estimations report robust standard errors and weight classes by the patent count in the class during 1985-1994. Column headers report specification variants. In Column 5, second-generation stimulus is estimated through a second application of the core innovation network matrix to the predicted patenting. In Column 6, second-generation stimulus is estimated through direct measure of a second-generation citation matrix restricted to ten years.

Appendix Table 2: Disaggregated analysis of innovation network

	Full sample with >5 per annum			Restricted sample with >50 per annum		
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable is log cumulative patent counts in patent class 1995-2004						
Log cumulative patents 1985-1994	0.678 (0.061)	0.732 (0.055)	0.840 (0.066)	0.785 (0.083)	0.806 (0.069)	0.905 (0.079)
Log expected patenting from network stimulus	0.345 (0.052)			0.265 (0.066)		
Log expected patenting from network stimulus due to top 10 upstream classes		0.298 (0.056)			0.294 (0.066)	
Log expected patenting from network stimulus due to next 10 upstream classes		0.067 (0.074)			-0.041 (0.091)	
Log expected patenting from network stimulus outside of top 20 upstream classes		-0.061 (0.054)			-0.032 (0.068)	
Log expected patenting from network stimulus within subcategory			0.077 (0.029)			0.043 (0.037)
Log expected patenting from network stimulus within rest of category			0.105 (0.037)			0.097 (0.047)
Log expected patenting from network stimulus outside of category			-0.004 (0.048)			0.012 (0.054)

Notes: See Appendix Table 1. In Columns 2 and 5, we separate the upstream stimulus provided by the ten most-important upstream classes, the next ten upstream classes, and those beyond. These upstream classes are defined by citation shares made by the focal class during the pre-1995 period. In two cases, the patent class cites fewer than ten upstream categories and is excluded. Columns 3 and 6 alternatively rely on the USPTO classification system of subcategories and categories, which is naturally cruder since it is less-tailored to an individual technology's citation patterns. The disaggregated results do not add up to the total network effect due to the log transformations.