

Amsterdam
School



for Social science Research

ASSR Working paper 05/02

February 2005

**SPECIALIZATION AND BROKERAGE:
A THEORY OF KNOWLEDGE GROWTH**

Gianluca Carnabuci and Jeroen Bruggeman

Gianluca Carnabuci is PhD Student at the ASSR

Jeroen Bruggeman is assistant professor at the department of Sociology at the University of Amsterdam

ASSR WORKING PAPER SERIES
[HTTP://WWW.ASSR.NL/](http://www.assr.nl/)
EMAIL: ASSRWORKINGPAPERS@FMG.UVA.NL

ABSTRACT:

Some scholars, mainly economists, argue that knowledge growth is enhanced by specialization, whereas others, mainly sociologists, argue in favor of knowledge brokerage. By casting the process of knowledge growth in a network representation, we are able to point out formally that specialization and brokerage are contradicting strategies. We then resolve this contradiction in a dynamic theory of knowledge growth, which we subsequently test on data describing all technological knowledge patented in the United States between 1975 and 1999. Our theory explains why domains of technological knowledge grow at different rates.

We are grateful to Bernhard Kittel and Willem Saris.

Send correspondence to Gianluca Carnabuci, Amsterdam School for Social science Research (ASSR), University of Amsterdam. Kloveniersburgwal 48, 1012 CX, Amsterdam, The Netherlands.

E-mail: g.carnabuci@uva.nl; j.p.bruggeman@uva.nl.

1. INTRODUCTION

Over any time span in human history reaching the present, the role of knowledge in economic growth has increased. This trend appears to have accelerated in the last decades with the emergence of the knowledge economy, and nowadays many scholars would agree that knowledge overshadows capital and labour as an asset of economic production (Romer 1993). Accordingly, modern organizations strive to generate novel ideas “in every stage of the production process, from the R&D lab to the factory floor to the interface with customers,” (Powell and Snellman 2004, p.201). These processes of knowledge growth have been subject to systematic investigation (Jaffe and Trajtenberg 2002), but the underlying mechanisms have not. In this paper we want to shed light on some of these mechanisms, with the purpose to answer the question: Why do certain domains of knowledge advance faster than others?

To enhance knowledge growth, the literature suggests two different strategies. At the core of economic thinking lies the view that specialization favours the advancement of knowledge. This conjecture was already central in the work of Adam Smith (1776), Amasa Walker (1867), Allyn Young (1928), and Alfred Marshall (1936)¹. More recently, and mostly in the sociological literature, the notion that brokerage of diverse knowledge sources propels the advancement of knowledge appears to have become widespread (Burt 2004; Stuart and Podolny 1996; Sutton and Hargadon 1996)². Surprisingly, the concepts of knowledge specialization and knowledge brokerage have not been related in the literature. However, they appear to exclude each other, because specializing in a particular domain impairs brokering many, and vice-versa.

Moreover, neither for brokerage nor for specialization the literature tells us under which conditions their beneficial effects can be expected, hence we can't explain variations in knowledge growth without further ado. For this reason, we begin our argument by fleshing out the mechanisms of knowledge growth that underlie specialization and brokerage. To clarify and then to resolve the specialization-

¹ Xiaokai Yang and Siang Ng (1998) report that the advantages of specialization for invention were spelled out even before the work of Adam Smith, for example in the French *Encyclopedia* published in 1701, in Maxwell (1721), and Tucker (1755, 1774).

² Interestingly, also this idea dates back to at least Adam Smith's work: “When the mind is employed about a variety of objects it is some how expanded and enlarged” (cited in Burt 2004, p. 350).

brokerage paradox, we formalize these two concepts network-analytically, and propose a theory of knowledge growth that reconciles their roles in a dynamic explanation. We believe that our theory holds at multiple levels of aggregation, because at all levels the dynamics of hybridisation are ultimately driven by the same mechanism – human ingenuity – and by the same resource – existing knowledge. In this paper, however, we test our theory at the level of technological domains, because at this level the best feasible data are available and in largest number: patent data. These data are relevant because they describe the growth of codified technological knowledge for a large and highly influential part of our knowledge society. Having treated the data, we will briefly present the topology of the knowledge network, and subsequently test our hypotheses about knowledge domains. We finish this paper by discussing possible policy implications, as well as the scope of our findings.

2. MECHANISMS OF KNOWLEDGE GROWTH

Building on research in cognitive psychology (e.g. Weisberg 1993), and subsequent applications in economics (e.g. Weitzman 1998), we adopt the well-known conception of knowledge growth as a hybridising process: “[N]ew ideas are essentially successful reconfigurations of existing ideas that have not previously been combined with each other” (Weitzman 1996, p.209). This view, going back to mathematician Henri Poincaré (1921) at least, was nicely cast in a mathematical model by Weitzman (1998), making possible to sharp-focus both conceptual development and empirical research. Weitzman’s model is unconstrained, though, allowing for a combinatorial explosion of ideas across the board, thus leaving moot why knowledge tends to grow in *trajectories* (Dosi 1982), and why certain trajectories are more fruitful than others.

Our approach differs from Weitzman’s in that we don’t assume that all hybridisations are possible or useful; rather, we expect the actual hybridisation patterns underlying knowledge growth to differ across people, firms, nations, as well as across the units of analysis of the present study: knowledge domains. As we will show, our focus on the structural configuration of hybridisation patterns makes possible to model the dynamics of knowledge growth network-analytically. Moreover, the concept of hybridisation provides a common denominator for specialization and brokerage.

Specialization is a mode of knowledge production in which ideas are hybridised from a set of closely related domains of knowledge; brokerage, in contrast, means that new knowledge is generated by hybridising ideas from unrelated domains.

To understand the benefits of specialization and brokerage, we need to understand how their characteristic modes of knowledge hybridisation enhance knowledge growth. A sensible strategy to tackle this issue is to identify the cognitive mechanisms of creativity associated with specialization on the one hand, and brokerage on the other.

3. CREATIVITY, SPECIALIZATION AND BROKERAGE

A popular view sees creative thinking as a mode of reasoning in which enlightened thoughts spring to mind without being cued, but this turns out to be more of a myth than a representation of reality. Mastery of the subject matter is necessary both to appreciate and to pursue the potential of a creative thought (Simonton 2000; Walberg 1988), and it conduces to absorptive capacity (Cohen and Levinthal 1990). Mastery, albeit necessary, can be achieved only at the price of high fixed learning costs (Hayes 1989; Simonton 1991). It typically takes years of preparation before even the most talented individuals can become proficient in any given domain (Weisberg 1993)³. It follows that it is efficient to spread fixed learning costs over a larger knowledge output, thus to hybridise as many ideas as possible from an unchanging set of strictly related knowledge domains. The reverse is also true. Hybridisation of ideas from a broad set of unrelated knowledge domains is an inefficient mode of knowledge production because it involves higher learning costs associated with a given level of knowledge output. For these reasons, we conclude that:

Proposition 1 Specialization yields efficiency in knowledge production

While mastery of the subject matter makes it possible to appreciate and pursue the potential for developing new ideas, the novelty of hybridisations resides in the useful relations they establish between yet unrelated ideas (Weisberg 1993). Cross-fertilizing

³ Mastery does not exclude serendipity but as Pasteur said, “chance favors the prepared mind.”

transfers of ideas typically occur by shifting mental models through analogies, metaphors, or other cognitive mechanisms (Holyoak & Thagard 1995). Regardless of the mechanisms involved, it is the exposition to ideas from different domains and applications that prompts unexplored mental representations and makes novel hybridisations visible (Anderson & Thompson, 1989). A well-known example is Gutenberg's printing press, which resulted from creatively pooling his knowledge of paper, metallurgy, press, ink, movable types, and the alphabet, among others (Diamond 1997). The literature abounds of similar accounts of other inventions, both old and new (Mokyr 2002). More systematic empirical evidence is reported by Dunbar (1996), who showed that scientists in laboratories with a greater diversity of scientific backgrounds are better able to solve problems by conceiving new ideas. Benefits of brokering a diversity of knowledge sources, and developing a proficiency at it, can be gained beyond laboratories. Hargadon (2002) and Burt (2004) found that individuals and organizations brokering between unrelated regions of their social network come up with more innovative ideas. Because the diversity of ideas someone is exposed to is determined by the diversity of the knowledge sources one learns from and strives to improve upon, we expect that hybridising ideas from unrelated knowledge domains develops the ability to identify novel hybridisations.

Proposition 2 Brokerage yields new hybridisation potential

Before we elaborate on the implications of Propositions 1 and 2, we raise the level of analytical precision by developing a network representation of knowledge growth.

4. A NETWORK REPRESENTATION OF KNOWLEDGE GROWTH

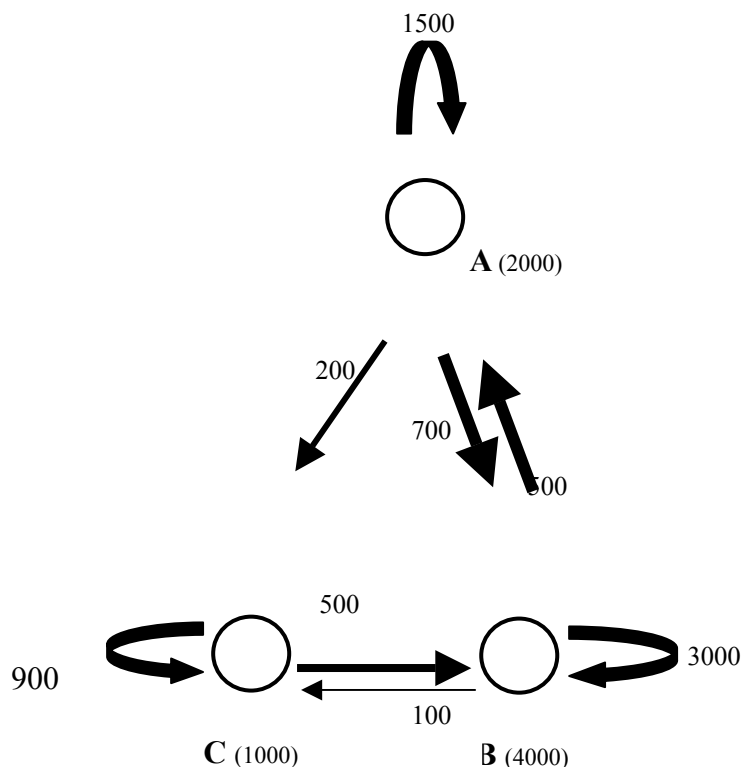
As said, our units of analysis are domains of knowledge⁴. Accordingly, we view knowledge growth as resulting from the evolving network of hybridisations that take place within and between domains of knowledge. Formally, a network N_t at time interval t is a four-tuple, $N_t = \langle J_t, L_t, V_b, A_t \rangle$, which consists of a finite set of nodes,

⁴ These domains, such as Power Plants and Motor Vehicles, are institutionalised artefacts of "loose communities" (Merton 1968) of individuals and organizations sharing a normative engineering culture (Constant 1987, p.227; White 1992, pp.202-203). By analysing knowledge domains, we presume that the knowledge sources that communities tap from in their creative endeavour are revealed by the hybridisation patterns observable at the level of the pertaining domains. In this respect, our study is similar to network analyses of industries or other aggregates.

$J_t = \{i, \dots, k, q, \dots, j\}$; a set of ties between the nodes, $L_t = \{l_{ik,t}, \dots, l_{qj,t}\}$; a function $V_t(.)$ mapping ties on pertaining values, usually called “weights”; and, a function $A_t(.)$ mapping nodes on node values. Nodes represent knowledge domains; directed ties represent hybridisations of ideas from one domain used in ideas belonging to the same domain or in other domains⁵; tie values indicate the number of hybridisations between two domains or within a single domain; node values represent knowledge output accumulated within domains.

In the example in Figure 1, A, B, and C are knowledge domains comprising, respectively, 2000, 4000, and 1000 accumulated ideas, within a given time period. Ideas in domain A have been developed by hybridising 1500 times ideas from A itself, 500 times ideas from domain B, and are hybridised 200 times to develop ideas in domain C.

Figure 1 Example Network



⁵ An idea that spills over in the knowledge production process is one that becomes hybridised as an input of a newer idea. While the often-used terms *spillover* and *diffusion* refer to the transfer of knowledge, *hybridisation* refers to the inventive mechanism that underlies such transfer.

This network representation makes possible a rigorous and straightforward approach to model processes of knowledge growth empirically. It imposes no *a priori* assumptions concerning the number or proportion of ideas that can be hybridised. Thus, it includes the limiting case where all possible hybridisations are carried out, as in Weitzman's combinatorial model (1998). Furthermore, it imposes no *a priori* assumptions concerning the distribution of hybridisations across the stock of knowledge. At one extreme, ideas could be hybridised randomly across domains, corresponding to a random network. At the opposite extreme, ideas could be hybridised exclusively within domains, which would generate a disconnected network of isolated trajectories. In actuality, growth processes are likely to be somewhere in between these two extremes - which we want to investigate. Our proposed network representation also offers an analytical framework for theorizing about how specific network properties relate to knowledge advancement processes. In the following section, we use it to relate knowledge specialization and brokerage formally.

5. THE NETWORK STRUCTURE OF SPECIALIZATION AND BROKERAGE

In line with the concept of hybridisation, we regard knowledge domains as specialized when they advance by hybridising closely related ideas, whereas we view as brokering those whose progress springs from the combination of unrelated ideas. The starting point for our formalization of these concepts is the representation of a domain's *niche* as a *sub-network* (of the entire knowledge network) comprising the focal domain, its contact domains, the valued and directed ties linking the focal domain to its contacts, and the valued and directed ties linking the domain's contacts among each other⁶. Structural holes (Burt 1992) in a focal node's niche indicate a lack of connections between domains in the sub-network. In knowledge networks, structural holes in a domains' niche indicate the degree to which that domain hybridises ideas from domains that do not hybridise ideas from one another or from the focal domain. Then, the brokerage of a domain is captured by the structural holes in its niche, and specialization, in turn, is captured by the lack of structural holes.

⁶ This conception of niche as a relationally defined position in a network is similar to the one employed, at the level of the patent, by Podolny and Stuart (1995).

To measure the level of brokerage⁷, B_{it} , of a focal domain, i , within a given time interval⁸, t , we adapt Burt's model (1992) as in equation (1):

$$B_{it} = 1 - p_{iit} + \sum_j (p_{ijt} + \sum_q p_{iqt} p_{qjt})^2 \quad (1)$$

where:

$$p_{iit} = \frac{h_{iit}}{\sum_j h_{ijt}}$$

$$p_{ijt} = \frac{h_{ijt}}{\sum_q h_{iqt}} * \frac{h_{jit}}{\sum_x h_{jxt}}$$

$$p_{iqt} = \frac{h_{iqt} + h_{qit}}{\sum_j (h_{ijt} + h_{jit})}$$

$$p_{qjt} = \frac{h_{qjt} + h_{jqt}}{\sum_z (h_{qzt} + h_{zqt})}$$

In this model, j and q are variables for the nodes in i 's niche, and $i \neq j \neq q$; for nodes in j 's niche x is used, $x \neq j$; and z stands for nodes in q 's niche, $z \neq q \neq j$; h is the number of ideas belonging to the right-hand subscript domain, used in hybridisations belonging to the left-hand subscript domain. Then, p_{iit} is the proportion of self-hybridisations within domain i , relative to all hybridisations carried out by i ; p_{ijt} is the proportion of ideas belonging to domain j that are hybridised in advancing i , weighed by the proportion of ideas that j hybridises from i ; and, the term within the inner summation indicates the redundancy of i 's contact domains, i.e. how closely related they are to both i and i 's other contacts. B_{it} is zero when i advances by hybridising ideas only from one fully specialized knowledge source, either itself or another, so when there are no structural holes in its niche; it approaches one when i hybridises ideas from many mutually-unrelated knowledge sources⁹.

⁷ To avoid unnecessary complications, in the following we will often use either the term *specialization* or the term *brokerage*, leaving their duality implicit. It remains understood that *specialization* means *lack of brokerage*, and *brokerage* means *lack of specialization*.

⁸ Note that t indicates a time interval, not a point in time.

⁹ Our model of network brokerage differs from Burt's in two ways. First, we account for the role of *loops*, i.e. domains' self-hybridisations, with the term p_{iit} . Second, we use a multiplicative term to

The concept of structural holes makes it possible for any given network to compare the degree of specialization versus brokerage of each of the nodes. In order to understand and model actual knowledge development, it is useful to look into the dynamic version of this cross-sectional view. The specialization of a domain, i , between two subsequent time intervals, t_{-1} and t , increases with hybridisations carried out that reduce the structural holes in i 's niche relative to t_{-1} . In turn, hybridisations reduce i 's structural holes when they are *path-dependent*, i.e. reinforce already established ties in i 's niche. As can be seen in equation (1), this can happen in three ways. First, the proportion of self-hybridisations increases ($p_{iit} - p_{iit-1} > 0$). Second, i 's hybridisations to other domains become more concentrated ($\sum_j (p_{ijt})^2 - \sum_j (p_{ijt-1})^2 > 0$). Third, the domains that i hybridises from become more redundant, i.e. they hybridise larger proportions from one another or from i ($\sum_q p_{iqt} p_{qjt} - \sum_q p_{iqt-1} p_{qjt-1} > 0$).

5. TOWARDS A DYNAMIC THEORY OF SPECIALIZATION AND BROKERAGE

The evolution of knowledge was first a subject of philosophy. The evolutionary theory prevailing in contemporary philosophy of science (Kuhn 1962; Lakatos 1970) states that knowledge typically advances by relatively long trajectories of cumulative refinements, called normal science. In this mode of knowledge production, mopping-up and puzzle-solving operations push forward the frontier of knowledge incrementally (Kuhn 1962). Although for some time a great deal of progress can be gained by normal science, it invariably exhausts creativity. When stagnation looms, manifesting itself through an increasing number of anomalies and insoluble problems, only radical departure from the core principles – a paradigm shift – can bring new momentum to the progress of knowledge. Kuhn's theory captures a dynamics of knowledge growth that appears to take place well beyond the realm of science. A central tenet of evolutionary economics is that the growth of technological knowledge follows an analogous pattern of advancement (Dosi 1982; Nelson and Winter 1982).

calculate p_{ijt} , to make the role of j 's hybridisation of ideas from i conditional on the proportion of ideas that i hybridises from j . Burt calculates p_{ijt} by an additive term; as a consequence, if i does not hybridise ideas from j while j does hybridise ideas from i , changes in j may affect the computed value of i 's brokerage while i 's actual brokerage is unaffected.

This view is backed by empirical evidence from many industries that long periods of incremental technical improvement are interrupted by short periods of radical and disruptive innovation – a pattern coined “punctuated equilibrium”(Abernathy and Utterback 1978; Tushman and Anderson 1986; Utterback and Suarez 1993).

In our network representation, the trajectory of development of a knowledge domain (in technology, science, or other realms) at any point in time is determined by the domain’s niche of knowledge sources. Normal knowledge advancement consists of path-dependent hybridisations that progressively exploit already established ties in a domain’s niche, thereby increasing its specialization. This mode of knowledge advancement is efficient (Proposition 1) – which explains why it is the normal (i.e. predominant) one. However, by progressively depleting the structural holes in a domain’s niche, it exhausts its creative sources (Proposition 2). Paradigm shifts, on the contrary, result from hybridisations that initiate relations with knowledge sources previously outside a domain’s niche, thereby expanding its structural holes. By connecting ideas from yet unrelated domains of knowledge, path-breaking hybridisations prompt unexplored perspectives, which may yield potentially fruitful veins for new hybridisations of ideas. Advancement is in no way guaranteed, though: path-breaking hybridisations are both costly and risky from a cognitive and an institutional perspective, and many holes are likely to be inconsequential dead-ends. In some cases, however, a new process of normal knowledge advancement catches on that reaps the value inherent in the new knowledge structure opened up. Then a sequence of cumulative hybridisations materializes the inventive potential by exploiting the newly created tie(s), signalling a paradigm shift.

By focusing on the creation and depletion of structural holes, Propositions 1 and 2 not only make possible to explain macro-dynamics of knowledge growth in terms of precise network conditions, they also complement the notion that at the micro level, knowledge production is maximized by balancing processes of knowledge *exploitation* and *exploration* (March 1991). Exploitation, which consists of harvesting as much value as possible from one’s established knowledge base, appears to be the predominant strategy of knowledge production in both humans and firms. In organizational sociology and psychology the acquisition of new knowledge is regarded as conditional upon one’s existing competencies (March and Simon 1958).

Similarly, according to Cohen and Levinthal's (1990) concept of *absorptive capacity*, firms' ability to assimilate and integrate new technological knowledge depends on their past R&D. In the long run, however, exploitation cannot sustain knowledge growth; hence, some degree of exploration beyond one's established competencies is necessary (Kogut and Zander 1992; Henderson and Cockburn 1994; Rosenkopf and Nerkar 2001; Teece et al 1997). Albeit useful, the notions of exploitation and exploration are markedly abstract, which hinders their use for scientific theory and research. The network-structural perspective fleshes out these two notions: knowledge exploitation consists of hybridisations that deplete the existing structural holes in one's path-dependent niche of knowledge sources, while knowledge exploration consists of hybridisations that break outside one's niche, thus enlarging its structural holes.

While connecting with extant literature, our network approach provides a novel explanatory framework for analysing the process of knowledge advancement. In the following section, we state four testable hypotheses about the effects of knowledge specialization and brokerage, which resolve the contradiction in the literature and lay the foundations for a full-fledged theory of knowledge growth.

6. HYPOTHESES

We have argued that path-dependent hybridisations are an efficient mode of knowledge production. These are hybridisations that deplete the structural holes in a domain's niche, thereby increasing its specialization. Hence, we expect that domains that increase their specialization between two points in time produce more knowledge in that time period:

Hypothesis 1 Changes in specialization are positively associated with domains' growth rates

However, the structural holes of any niche are finite. This leads to two further testable implications. The first is that higher specialization is associated with slower growth, because more hybridisation opportunities have already been exploited in specialized

domains, and fewer opportunities are left for subsequent hybridisations. Conversely, brokering domains should have greater potential for upcoming hybridisations.

Hypothesis 2 Levels of specialization are negatively associated with domains' growth rates

The second implication is that path-dependent hybridisations should yield marginally decreasing returns. Unlike hypothesis 2, the argument here is not that specialized knowledge offers fewer possibilities to hybridise, but that possible hybridisations have a more limited innovative potential. The higher the level of specialization, the more hybridisation possibilities consist of refinements of current ideas in a mopping-up fashion. On the other hand, hybridisations in brokering domains spring, on average, from more distant cross-fertilizations, resulting in greater knowledge advancement.

Hypothesis 3 The effect of changes in specialization on domains' growth rates decreases with levels of specialization

Another way to look at specialization versus brokerage is to also take into account the associated risk. According to hypothesis 3, specialized domains advance slower than brokering domains *on average*. However, we expect a relatively larger portion of brokering ideas to be inconsequential dead-ends because of cognitive and institutional obstacles associated with the development of path-breaking inventions. Contributions to specialized domains, in contrast, follow a more established and more legitimate path of knowledge advancement, hence run lower risks of being rejected or ignored. For this reason, we hypothesize,

Hypothesis 4 The variance in knowledge growth is larger among brokering than among specialized domains

By our network approach and our dynamic theory of knowledge, highlighted by the four hypotheses, we have resolved the specialization-brokerage paradox on the fly.

7. DATA

We test the hypotheses at the level of technological domains. We use data of all patents (over two million), granted by the largest international patent office worldwide, USPTO, between 1975 and 1999. These data, collected by the National Bureau for Economic Research (NBER), are thoroughly described by Hall and colleagues (2001). To use the data appropriately, we must be able (1) to clearly define the nodes, i.e., the knowledge domains; (2) to measure the hybridisations within and between domains, and (3) to quantify domains' knowledge growth.

First, as to the definition of our network nodes, the USPTO officers categorize each granted patent into one of 418 internationally standardized technological domains. Although these categories do not represent objective knowledge boundaries, the officers of USPTO undoubtedly have the best possible judgement on this issue¹⁰. On the basis of their classification, we make a homomorphic mapping of the individual patents and citations to the 418 technological domains, and obtain a network with a fixed number of technological domains over the period of observation, related by patterns of hybridisations¹¹.

Second, patent data contain so-called prior art, i.e. a list of all previous patents that the patented invention draws from. The accuracy of the prior art in citing all relevant previous patents is to a large extent guaranteed by the fact that being accurate is legally binding and economically advantageous¹². Moreover, one of the most important tasks of patent officers at USPTO is to cross-check and integrate the prior art compiled by the applicants.

¹⁰ USPTO officers categorize patents also into more than 120.000 subclasses. Next to reasons of tractability, we prefer to use the broader 3-digit partitioning because we believe that this higher level of data aggregation offers a more reliable representation.

¹¹ Notably, there are tie-loops of fields to themselves; these loops are not possible at the level of individual patents because a patent can't cite itself.

¹² The discovery of inaccuracies in a patent's prior art lengthens the process that leads to its grant, and therefore hampers the legal protection of the invention. This can have two negative consequences for the applicant. First, the benefits springing from the patent such as royalties or licenses are delayed. Second, it might happen that another patent is granted in the meanwhile that makes the initial invention obsolete and no longer patentable.

Since the digitalisation of patents, patent citations have become widely used data for scientific research on technological knowledge. The citations contained in the prior art have been used as indicator of hybridisations in a variety of studies, which in the literature are typically referred to as knowledge spillovers (Jaffe et al 1993; Jaffe et al 2000). At the level of individual patents, there appears to be too much noise to interpret each citation as knowledge hybridisation, but at the aggregate level of domains, this noise vanishes in the large numbers (Jaffe et al 2000). Hence we regard the ties at this level as indicative for the patterns of hybridisation within and between domains.

Third, patent data reveal the quantity of technological knowledge generated. For an invention to be patented, it must consist of knowledge that is new, non trivial, and usable for some productive application. Therefore a patent is an idea that contributes to the advancement of codified technological knowledge. Accordingly, patent counts are a valuable proxy for measuring this advancement, in particular at the level of domains, if the quality of each patent is taken into account. An established way to measure quality is to weigh patents by the citations they receive within a given time period after their grant (Griliches 1990), which also we do in our study.

8. EMPIRICAL MODELING

Based on the framework proposed in the previous sections, we model the NBER data as a valued and directed network of patent citations between technological domains. To characterize the topology of the network for the 25-year period compressed, we follow Duncan Watts and Steven Strogatz (1998), who in their famous paper on small-worlds observe networks' clustering coefficient and characteristic path-length¹³. Our knowledge network has a clustering value of 0.702 (for a random graph with same number of nodes and ties it is 0.19, the expected value is 0.512), and a

¹³ For these particular calculations, we simplified our network ties to the unweighted and undirected ones used in Watts and Strogatz's measures. The clustering coefficient for one node is defined as the ratio of the actual number of ties between its neighbours and the maximum number; for the network, the clustering coefficient is the mean of the individual coefficients (Watts and Strogatz 1998, p.441). Path length between two nodes is the minimum number of steps along ties to progress from one node to the other.

characteristic path length of 1.490 (the expected value¹⁴ for a random graph is 1.488). Therefore, as is the case with small worlds, our knowledge network features a higher degree of local clustering than would be expected in a random graph, which does not result in a higher global distance across nodes. However, unlike small worlds, our network also features high connectivity: on average, a domain is linked by at least one patent citation to more than half of all other domains.

To capture the evolution of the network, we use the method of comparative statics, and group the patents into periods of five years. Hence, one cross sectional image of the network is based on all citations between patents granted within the same five-year window. The choice of five years may seem arbitrary, but it is nonetheless appropriate for two reasons. First, other studies employing patent citations data adopt a five-year window (Podolny et al 1996), which makes our study comparable. Second, and most importantly, a crucial criterion in defining the length of time periods is whether domains' citation patterns to older patents are directed to different domains than their citation patterns to recent patents. By means of QAP regression analysis, we can assess such differences (Simpson 2001). We regress the inter-domain adjacency matrix in which both cited and citing patents are granted in the interval 1995-1999, on the adjacency matrix constructed using citing patents granted in the same five years (1995-1999) but cited patents granted during the entire observation period, 1975-1999. This allows us to assess the discrepancy in network representation resulting from using a 5-year window rather than a 25-year one. Table 1 shows the results, which indicate that the two network configurations are virtually identical. Thus, it can be concluded that a five-year window represents the network structure of patent citations between technological domains with great precision¹⁵.

¹⁴ Newman, Strogatz, and Watts (2001) showed that, for random networks, the characteristic path length is approximated by $\ln(n) / \ln(k)$. This approximation, however, is unbiased only in the limiting case of $n \gg k$ (in which case, moreover, the standard deviation is quite large). In the network described here, the condition $n \gg k$ is not satisfied. To calculate L_r on a more sound basis, we generated 50 random networks with n and k set equal to those observed in our hybridisation network, and measured L_r as the average of their characteristic path lengths. Note that the resulting distribution of L_r had a standard deviation smaller than 0.01

¹⁵ Strictly speaking, this conclusion is valid only for the patents granted between 1995 and 1999. However, we carried out similar QAP analyses for the preceding time intervals (comparing the network representation based on 1990-1994 patents against the one based on patents granted between 1975 and 1994; the network based on 1985-1989 patents against the one based on patents granted between 1975 and 1989; etc.) and we obtained equivalent results.

Table 1 Model fit and regression coefficients of QAP regression, with 2000 permutations

R-square	Adj. R-square	Probability	Observations
1.000	1.000	0.000	174306
Independent coefficients	Un-std. Coefficient	Std. Coefficient	Significance
Intercept	-0.000792		
5-year-window network	1.000049	1.000000	0.000

We measure domains' knowledge output as the number of citations received within the five-year period. This poses a problem, though, because citation lags are very long: less than 30 percent of all citations made is to patents between 0 and 5 years older than the citing patent, and it takes 50 years to cover 90 percent of all citations received by a patent. To assess the magnitude of the problem, we focused our attention on the patents granted in each domain between 1975 and 1979. We first counted the citations received by domains during the same five-year interval, then we counted the citations they received from all patents granted up to 1999, finally we correlated the two vectors of scores. We calculated both *Pearson's* correlation and *Spearman's rho* correlation of rank orders, which result in coefficients of 0.973 ($p < 0.0001$) and 0.976 ($p < 0.0001$), respectively. On the basis of these values, we concluded that the measurement of knowledge output based on five-year windows is warranted¹⁶.

Given the choice of five-year intervals, a maximum of twenty-one networks, each overlapping for 80% with its neighbour, could be used as observations for the time-series. After an examination of the data, however, we decided to reduce the observation instances to five non-overlapping networks. This decision is driven by two considerations. First, modelling the networks from non-overlapping data subsets guarantees that relationships between networks are not artificial. Second, an estimated structural equation model with multiple lagged explanatory variables indicated that

¹⁶ We carried out similar analyses focusing on the other time intervals (i.e., for the patents granted in each domain between 1979 and 1984, we correlated the citations received by domains within the same 5-year interval, with the citations they received between 1979 and 1999; for the patents granted between 1985 and 1989, we correlated the citations they received within the same 5-year interval, with the citations they received between 1985 and 1999; etc.), which led to the same conclusion.

effects are strongest with lags in the neighbourhood of five years, which suggests that our choice is an appropriate simplification to capture the hypothesised effects. As a consequence of this decision, the structure of patent citations between technological domains in the years from 1975 to 1999 is modelled as a time series of five networks.

One of the advantages of testing hypotheses by means of panel data is that the researcher is not constrained to infer causal relationships from *level* variables, but can use *change* variables too. To model the hypothesized effects, we make a combined use of level and change variables. To model hypothesis 1, we relate domains' knowledge growth to domains' change in specialization, between t_{-1} and t . To model hypothesis 2, we relate domains' knowledge growth between t_{-1} and t to domains' level of specialization at t_{-1} . And, to model hypothesis 3, we relate domains' knowledge growth between t_{-1} and t to the interaction of domains' specialization at t_{-1} and change of specialization between t_{-1} and t .

8.1 The Inferential Model

Our data are non-independent in two ways. First, being interconnected in a network, domains are by definition non-independent units. Second, our observations are repeated; hence, they are nested within units. Non-independence biases the standard error of a statistical relationship because it induces correlated residuals across units. However, if such a correlation is appropriately modelled, units' non-independence is no problem. For this reason, we will address non-independence in the statistical model.

8.2 Statistical Model

When observations are taken from units embedded in a network structure, *network autocorrelation* needs to be accounted for (Doreian 1981). In the context of this study, we reckon that there exist three possible sources of network autocorrelation. First, time-varying factors can affect observations of all units in a similar way. For example, macro-economic fluctuations may influence the growth dynamics of all technological domains. We model away this effect by introducing period-dummy variables. By doing so, we also sort out the possibility that the estimates pertaining to our hypothesized relationships are altered by artifactual changes in the data. For example, both the number of patents and the number of citations have increased rapidly during

the observation period. However, it is hard to establish on the basis of the data what part of this increase reflects changes in the dynamics of knowledge growth, and what part is due to changing practices among USPTO examiners, or to changes in the propensity to patent or to cite. By controlling for period effects, we get rid of all variation that could be caused by such artifacts. Second, observations might correlate within subsets of units. For example, Hall *et al.* (2001) have shown similarities in patent statistics within the six macro-technological areas defined by the NBER 1-digit classification. We control for this effect by means of area-dummy variables. Third, network autocorrelation may have a dyadic structure, i.e. a node value may affect other node values to varying degrees. This kind of autocorrelation is the most awkward from a modelling point of view, and is often ignored in the literature. A methodologically sound solution to this problem consists of using available information on the network structure of the data to model network autocorrelation explicitly into the statistical equation (Doreian 1980, 1981). For this purpose, a weighting matrix of dyadic interdependences, W , needs to be specified. Typically, W reflects the network connections, whereby each (valued) dyadic tie determines the extent to which the response variable of a contact node affects the response variable of the focal node. The response variable of each node is then weighted on the basis of this dyadic inter-dependence score, and summed across all contacts for each of the nodes. The outcome is a node-level score of network autocorrelation, which is then employed as a regressor in the statistical equation. As a result of this procedure, the degree to which network connections induce network autocorrelation can be estimated and controlled for. Applied to our domain-by-domain knowledge network, this procedure estimates the degree to which the percentage growth of focal domain i is affected by the percentage growth of contact domains j . The weight that j 's growth has on i 's growth is determined by the number of citations between i and j . We computed this variable and included it in two alternative model specifications, one in which the growth of contacts j is assumed to have a contemporaneous effect on i 's growth, and the other in which the effects of j 's growth are supposed to be strongest after one period¹⁷.

¹⁷ Because the estimates turn out not to differ at all between these two specifications, we will only report the model with contemporaneous network autocorrelation.

Because observations are repeated on the same units, in our case domains, panel data typically have a nested structure: observations within units are more likely to be similar than observations between units. Ignoring this nested structure leads to biased estimates and incorrect standard errors. A general form of statistical models that account for unit heterogeneity is:

$$Y_{it} = \alpha_i + \mathbf{X}_{it}'\boldsymbol{\beta} + \varepsilon_{it} \quad (2)$$

where $i = 1, \dots, N$; $t = 1, \dots, T$

The main difference between (2) and a standard regression equation is represented by term α_i , denoting unit-specific factors. We consider a number of alternative approaches to modelling α_i . First, we remove all between-domain variance from the estimation by taking unit-demeaned variable transformations. This approach, known as *fixed-effects* or *within* estimation, is equivalent to modelling α_i as a set of $i = \{1, \dots, N\}$ domain-dummy variables. The second model we employ eliminates the problem of unit heterogeneity in a way opposite to the within estimator, i.e. by regressing unit means. Because it removes all within-unit variance, the resulting model is often termed *between-effects*. The third approach we use is to model α_i as random draws from a wider population, which results in a random-effects model and leads to

$$\alpha_i = \alpha + \eta_i \quad (3)$$

and, in turn, to

$$Y_{it} = \alpha + \mathbf{X}_{it}'\boldsymbol{\beta} + \eta_i + \varepsilon_{it} \quad (4)$$

The domain-specific factors are now denoted η_i (instead of α_i) to emphasize that they are modelled as a random component of the same type as the error ε_{it} – i.e. as a distribution characterized by its mean and variance – rather than constant parameters as in the fixed-effects specification. The random-effects model estimates the coefficients from weighted unit-demeaned variable transformations, where the weight,

$0 < \theta < 1$, is determined by:

$$\theta = 1 - \left(\frac{\sigma_{\varepsilon}^2}{\sigma_{\varepsilon}^2 + T\sigma_{\eta}^2} \right) \quad (5)$$

Because θ is calculated by partitioning the variance into a within and a between component, the random effects model can be regarded as a weighted combination of the fixed-effects and the between-effects specifications.

By positing that knowledge productivity varies more across brokering domains than across specialized ones, Hypothesis 4 suggests that our data are heteroskedastic, hence that ordinary computation of standard errors could lead to biased results. Moreover, in our first three models, estimation might also be altered by unaccounted serial correlation. As a conservative method to check if heteroskedasticity and autocorrelation influence our estimates of interest, in our fourth model we assume that these problems exist, and test the hypothesized relationships on this basis. To do this, we model panel-level heteroskedasticity and a first-order autoregressive (AR1) error covariance structure¹⁸, and employ a Feasible Generalized Least Squares (FGLS) estimator.

Equation 1 measures domain's specialization on the basis of the network-form of domain's niche, and is normalized to be independent of domain's size. This makes possible to define the concept of specialization in a way that can be applied meaningfully to all domains in our population. However, the actual effect of specialization depends on the quantity of structural holes in a domain's niche, which is co-determined by form and size. A small domain connected to two disconnected nodes has fewer hybridisation opportunities than a large domain with the same niche configuration. Because the number of citations a domain receives is indicative of its size, a straightforward way to account for this discrepancy is to model domains' knowledge growth in percentage change, and changes in specialization in first-differences. This amounts to estimating the effects of changes in the configuration of

¹⁸ We follow Beck and Katz (1995), who make a strong argument in favour of modelling the coefficient of the first-order temporal correlation as common to all panels.

Table 2 The dynamic effects of specialization on knowledge growth under the fixed-effects, between-effects, random-effects, and FGLS specifications.

	Dep. variable: % knowledge growth			
	Fixed-effects	Between-effects	Random-effects	FGLS
Intercept	0.71*** (0.13)	0.52*** (0.06)	0.31*** (0.06)	0.28*** (0.02)
Network autocorrelation	0.00 (0.00)	0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)
Domain size	-9.3e ⁻⁰⁷ (5.2e ⁻⁰⁶)	-1.7e ⁻⁰⁶ (2.9e ⁻⁰⁶)	1.3e ⁻⁰⁶ (8.8e ⁻⁰⁷)	-9.3e ⁻⁰⁷ (5.2e ⁻⁰⁶)
1985 thru 1989	0.35*** (0.04)		0.37*** (0.04)	0.33*** (0.01)
1990 thru 1994	0.11** (0.05)		0.13*** (0.04)	0.11*** (0.01)
1995 thru 1999	0.29*** (0.05)		0.32*** (0.04)	0.24*** (0.02)
Computers & Communications		0.82*** (0.08)	0.84*** (0.07)	0.66*** (0.05)
Drugs & Medical		0.91*** (0.11)	0.94*** (0.11)	0.69 (0.06)
Electronic		0.07 (0.07)	0.10 (0.07)	0.08*** (0.02)
Mechanical		0.05 (0.06)	0.07 (0.06)	0.03 (0.02)
Others		0.12** (0.06)	0.14*** (0.06)	0.10*** (0.02)
Specialization	-1.01*** (0.36)	-0.28** (0.13)	-0.37*** (0.12)	-0.20*** (0.05)
Δ Specialization	1.49*** (0.39)	3.11*** (0.83)	1.92*** (0.34)	2.16*** (0.30)
Specialization * Δ Specialization	-1.65** (0.69)	-2.99** (1.10)	-1.69*** (0.56)	-2.23*** (0.70)
Number of units	413	413	413	413
Periods	4	4	4	4
Num. Observations	1652	1652	1652	1652
F	17.90	23.83		
Prob. > F	0.0000	0.0000		
Wald chi-square			374.46	1116.45
Prob. > chi-square			0.0000	0.0000
R-Squared within	0.1042	0.0237	0.1015	
R-Squared between	0.0198	0.3721	0.3676	
R-Squared overall	0.0608	0.1615	0.2117	
Corr. (η_i, X_b)	-0.1212			

domains' niche between $t-1$ and t , holding domains' size at $t-1$ constant¹⁹. We also model domains' size as a regressor, thereby accounting for dynamics of error correction and mean reversion (Finkel 1995).

9. ANALYSIS

Table 2 displays the results of our analyses. Before we examine them from a theoretical point of view, let us discuss their reliability. The first noteworthy element is that directions and significance levels of the estimated coefficients are nearly identical in all four models. This is a reassuring starting point, because the estimates are based upon different assumptions and modelling specifications. To be sure, both fixed-effects and between-effects models are conservative approaches that generate consistent estimates with panel data. Random-effects models, in contrast, are not consistent if the unobserved unit-specific factors, η_i , are significantly correlated with the regressors, X_b . Nonetheless the random-effects specification is more efficient, and it allows estimating jointly time-invariant and unit-invariant variables. As shown in Table 2, the estimated correlation coefficient between η_i and X_b , is relatively low, amounting to a mere -.112. Whether this correlation induces inconsistent estimates in the random-effects specification can be formally established by means of a Hausman specification test (1978). This tests the null hypothesis that there is no systematic difference between the coefficients estimated by a random-effects specification and its fixed-effects counterpart. As the null hypothesis cannot be rejected (Chi-square = 6.63; Prob. > Chi-square = 0.4687), it can be safely concluded that also the random-effects specification yields consistent estimates.

Let us now turn to the interpretation of the results. The first rows of Table 2 report estimates for the intercept and the control variables. Network autocorrelation seems not to be an issue, since in all model specifications the pertaining coefficient is entirely insignificant. Domains' size affects negatively their growth rate, although the effect is not significant. Reflecting the generalized increase in patenting and citation

¹⁹ With this specification, we estimate the semi-elasticity of domains' knowledge output to the form of domains' niche. An alternative way to account for size is to weigh changes in domains' specialization by the size of each domain, and estimate their effects on first-difference in domains' knowledge output. The estimates obtained do not differ from the ones obtained under this alternative specification.

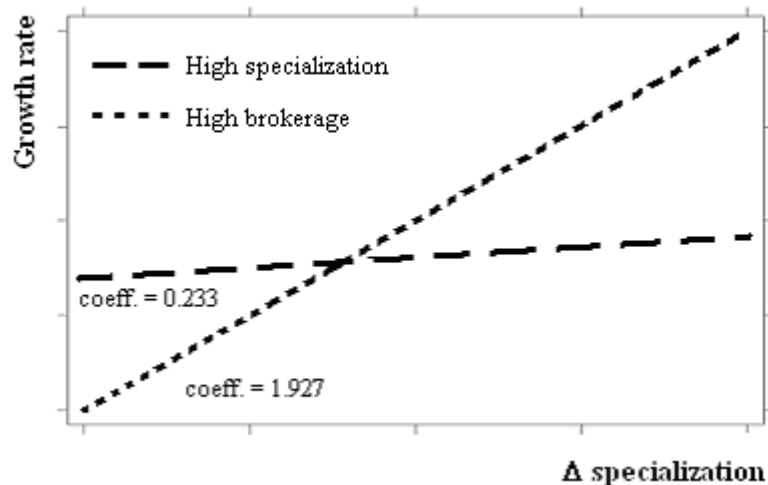
rates, domains appear to have grown faster over the years relative to their size. This trend however is not smooth over the observation period, as domain's percentage growth rates are lowest in the beginning of the nineties. As could be expected, domains in the areas of Computer and Communications have grown fastest during the observation period, followed by those belonging to Drugs and Medicals (reference category: Chemicals). Somewhat more surprisingly, third in rank is the miscellaneous category Others, while in the area of Electrical and Electronics, domains have not grown significantly faster than in the Chemical sector. This appears to contradict the distinction made by Hall *cum suis* (2001) between traditional and emergent areas, whereby the area Others would belong to the former while Electrical and Electronics to the latter. The difference in results is explained by difference in focus; while we investigate percentage growth of individual domains within each macro area, Hall *cum suis* study the growth of the macro areas themselves. Jointly, these results seem to indicate that, compared to Electrical and Electronics, the growth observed in the area Others is to a larger extent due to the growth of small domains. As smaller domains are likely to be less mature, the analysis presented here provides an interesting nuance to our understanding of emergent fields.

The last three rows of the upper part of Table 2 test the first three hypotheses. In all model specifications, all three hypotheses are supported by the data. As expected, specialized domains grow on the average slower than brokering ones (hypothesis 2), although increasing specialization is positively associated with growth (hypothesis 1); moreover, the positive effect of increasing specialization decreases as domains specialize (hypothesis 3). The support of the hypotheses under all estimated models is not only indicative of their robustness, but also of their generality. As captured by the fixed-effects specification, the “dialectic” relationship of knowledge specialization and brokerage explains how the growth rate of each individual domain changes over time. In addition to changes in each domain's growth rates, however, certain domains grow on average faster than others. The between-effects model shows that the hypothesized dynamics of knowledge specialization and brokerage contribute significantly to explaining also these differences. The random-effects model shows that both kinds of effects are there net of each other and of the complete set of control variables. Furthermore, while the fixed-effects model makes possible to generalize the estimated relationships to the population of observed effects, generalizations within

the random-effects framework refer to the universe of all possible effects. Finally, the Feasible Generalized Least Squares model demonstrates that the hypotheses hold even when the data are assumed to be heteroskedastic and serially correlated.

A more in-depth analysis of the interactive effect of levels and changes in domains' specialization (hypothesis 3) reveals theoretically important nuances that emerged from the estimation²⁰. Figure 2 shows the effects of changes in specialization on domains' growth when levels of specialization are highest or lowest. In both cases, changes in specialization are positively associated with domains' growth rates. However, while the benefits of increases in specialization are weak and non-significant when domains are highly specialized (t-value = 0.67), they are strong and highly significant when domains broker between disconnected sources (t-value = 5.61). This result shows precisely the predictions of our theory: path-dependent hybridisations are productive, but their productivity disappears in the absence of structural holes.

Figure 2 Effects of changes in specialization on knowledge growth for highest 10% specialized and highest 10% brokering domains.



²⁰ Estimates are based on the random-effects specification.

9.1 Risk and progress

To test hypothesis 4 (Table 3), we compare the variance in growth rates of two groups, the 10% domains having highest average brokerage and the 10 % domains having highest average specialization during the period of observation. To account for possible fixed effects in the growth patterns of domains, we express our dependent variable in deviations from unit means; hence, the variance we observe is within domains over time. As we predicted, the variance in growth rates is strikingly larger for brokering domains than it is for specialized ones, and the difference is highly significant according to Levene's test. The high variance of brokering shows the associated risk, whereas specialization is relatively safe. A glance at the confidence intervals for the means gives a sense of how differently the two strategies of knowledge production operate: within a 95% confidence interval, the growth rates observed for brokering domains deviate nearly twice as much from their mean as specialized domains do.

Table 3 Levene test of homogeneity of variance. Highest 10% brokerage versus highest 10 % specialization

Dependent variable: Unit-demeaned percentage growth

Groups	N	T	Obs	Mean	Std. dev.	Std. error	95% confidence interval	
							Min	Max
High brokerage	41	4	164	0.000	0.698	0.054	-0.10	0.10
High specialization	41	4	164	0.000	0.425	0.033	-0.06	0.06
Levene statistic	13.05							
Significance	0.000							

10. DISCUSSION

At the level of knowledge domains, we showed that specialization and brokerage are dynamically intertwined, which on the one hand resolves the paradox pointed out, and on the other hand makes clear why some domains advance faster than others. The dynamic perspective makes possible to speculate on some policy and productivity implications. We showed that domains featuring high levels of specialization have lower growth rates. A straightforward conclusion is that among domains of similar size it is wiser to invest in domains richer in structural holes. For the same reason, policies should be designed that stimulate specialized domains to draw ideas from sources outside their path-dependent niche. Obviously, investments in path-breaking connections carry a greater risk, and there is little guarantee as to whether or when they will pay off. When the goal of the policy-maker is to favour the advancement of technological knowledge at large, our perspective suggests that productivity gains would result from maintaining a balance between dense areas of knowledge hybridisation and structural holes between these dense areas. It might be the case that this balance is achieved by invisible hand, because small-world networks tend to be governed by a self-organizing process of proportionate growth, and such processes can hardly be influenced externally (Barabási and Albert 1999). Nonetheless, policies could aim at modifying the structure locally, in favour of immediate exploitation of structural holes, or to invest in the generation of new structural holes for future benefits.

Concluding, the concept of structural hole is just one, albeit possibly the most powerful, of network concepts that may help to explain the process of knowledge advancement. We have little doubt that other interesting aspects of knowledge growth can be analysed at all levels of aggregation within a network-analytic framework.

REFERENCES

- Abernathy, W. and J. M. Utterback (1978). "Patterns of innovation in technology." Technology Review **80**(7): 40-47.
- Amabile, T. M. (1988). "A model of creativity and innovation in organizations." Research in Organizational Behavior **10**: 123-167.
- Anderson, J. R., and R. Thompson. 1989. "Use of analogy in a production system architecture." Pp. 267-297 in *Similarity and analogical reasoning*, edited by S. Vosniadou and A. Ortony. Cambridge, UK: Cambridge University Press.
- Barabási, A. L. and R. Albert (1999). "Emergence of scaling in random networks." Science(286): 509-512
- Beck, N., and J. N. Katz. 1995. "What to do (and not to do) with time -series-cross-section data." *American Political Science Review* 89:634-647
- Burt, R. S. 1992. *Structural holes: the social structure of competition*. Cambridge, Mass.: Harvard University Press.
- Burt, R. S. 2004. "Structural holes and good ideas." *American Journal of Sociology*:349-399.
- Cohen, W. M. and D. A. Levinthal (1990). "Absorptive-Capacity - a New Perspective on Learning and Innovation." Administrative Science Quarterly **35**(1): 128-152.
- Diamond, J. 1997. *Guns, germs, and steel: the fates of human societies*. New York and London: W.W. Norton
- Doreian, P. 1980. "Linear models with spatially distributed data: spatial disturbances or spatial effects?" *Sociological Methods & Research* 9:29-60.
- . 1981. "Estimating linear models with spatially distributed data." *Sociological Methodology*:359-388.
- Dosi, G. (1982). "Technological paradigms and technological trajectories." Research Policy(11): 147-162.
- Dunbar, K. (1996). How scientists think: online creativity and conceptual change in science. Conceptual structures and processes: emergence, discovery and change. T. B. Ward, S. M. Smith and S. Vaid. Washington DC, APA Press.
- Finkel, S. E. 1995. *Causal analysis with panel data*. Thousand Oaks: SAGE Publications
- Hall, B., A. Jaffe, and M. Trajtenberg. 2001. "The NBER patent citations data file: lessons, insights and methodological tools." NBER

Hargadon, A. 2002. "Brokering knowledge: linking learning and innovation." *Research in organizational behaviour* 24:41-85.

Hausman, J. (1978). "Specification tests in econometrics." *Econometrica* 46: 1251 - 1271.

Hayes, J. (1989). The complete problem solver. Hillsdale, NJ, Erlbaum.

Henderson, R. and I. Cockburn (1994). "Measuring Competence - Exploring firm effects in pharmaceutical research." *Strategic Management Journal* 15: 63-84.

Hsiao, C. (2003). Analysis of panel data. Cambridge; New York, Cambridge University Press.

Jaffe, A., and M. Trajtenberg. 2002. *Patents, citations, and innovations: a window on the knowledge economy*. Cambridge, Mass.: MIT Press.

Jaffe, A., M. Trajtenberg, and M. S. Fogarty. 2000. "Knowledge spillovers and patent citations: evidence from a survey of inventors." *American Economic Review* 90:215-218.

Jaffe, A., M. Trajtenberg, and R. Henderson. 1993. "Geographic localization of knowledge spillovers as evidenced by patent citations." *Quarterly Journal of Economics* 108:577-598.

Kogut, B. and U. Zander (1992). "Knowledge of the firm, combinative capabilities, and the replication of technology." *Organization Science* 3: 383-397.

Kuhn, T. (1962). The structure of scientific revolutions. [Chicago], University of Chicago Press.

Lakatos, I. (1970). Falsification and the methodology of scientific research programmes, Cambridge University Press.

March, J. (1991). "Exploration and exploitation in organizational learning." *Organization Science* 2(1): 71-87.

March, J. and H. Simon (1958). Organizations. New York, Wiley.

Marshall, A. (1936). Principles of economics, an introductory volume. London, Macmillan and co. limited.

Maxwell, H. 1721. *Reasons offered for erecting a bank in Ireland*. Dublin

Merton, R. K. 1968. *Social theory and social structure*. N.Y.: Free Press.

Mokyr, J. 2002. *The gifts of Athena: historical origins of the knowledge economy*. Princeton, [N.J.]: Princeton University Press

- Nelson, R. and S. Winter (1982). An evolutionary theory of economic change. Cambridge, Mass., Belknap Press of Harvard University Press.
- Poincaré, H. 1921. *The foundations of science: science and hypothesis, the value of science, and science and method*. New York: Science Press.
- Podolny, J. M., and T. E. Stuart. 1995. "A role-based ecology of technological change." *American Journal of Sociology* 100:1224-1260
- Podolny, J. M., T. E. Stuart, and M. T. Hannan. 1996. "Networks, knowledge, and niches: competition in the worldwide semiconductor industry, 1984-1991." *American Journal of Sociology* 102:659-689.
- Powell, W. W., and K. Snellman. 2004. "The knowledge economy" *Annual Review of Sociology* 30:199-220
- Romer, P. 1993. "Idea gaps and object gaps in economic development." *Journal of Monetary Economics* 32:543-573.
- Rosenkopf, L. and A. Nerkar (2001). "Beyond local search: boundary-spanning, exploration, and impact in the optical disk industry." Strategic Management Journal 22(4): 287-306.
- Simonton, D. K. (1991). "Emergence and realization of genius - the lives and eorks of 120 classical composers." Journal of Personality and Social Psychology 61(5): 829-840.
- Simonton, D. K. (2000). "Creativity - Cognitive, personal, developmental, and social aspects." American Psychologist 55(1): 151-158.
- Simpson, W. (2001). The Quadratic Assignment Procedure (QAP). North American Stata Users' Group Meetings 20011.2, Stata Users Group.
- Smith, A. (1776). An inquiry into the nature and causes of the wealth of nations. London, W. Strahan & T. Cadell.
- Stuart, T. E. and J. M. Podolny (1996). "Local search and the evolution of technological capabilities." Strategic Management Journal 17: 21-38.
- Sutton, R. I. and A. Hargadon (1996). "Brainstorming groups in context: effectiveness in a product design firm." Administrative Science Quarterly 41(4): 685-718.
- Teece, D. J., G. Pisano, and A. Shuen. 1997. "Dynamic capabilities and strategic management." *Strategic Management Journal* 18:509-533.
- Tucker, J. 1755. *The elements of commerce and theory of taxes*. London: privately printed.
- . 1774. *Four tracts on political and commercial subjects*. Gloucester: Taikes, R.

Tushman, M. L. and P. Anderson (1986). "Technological discontinuities and organizational environments." Administrative Science Quarterly(31): 439-465.

Utterback, J. M. and F. F. Suárez (1993). "Innovation, competition, and industry structure." Research Policy **22**(1): 1-21.

Walberg, H. J. (1988). Creativity and talent as learning. The Nature of Creativity: Contemporary Psychological Perspectives. R. J. S. (ed.). Cambridge, Cambridge University Press: 340-361.

Walker, A. (1867). The science of wealth: a manual of political economy; embracing the laws of trade, currency, and finance. Boston, Little Brown and Company.

Watts, D., and S. H. Strogatz. 1998. "Collective dynamics of 'small world' networks." *Nature*:440-442

Weisberg, R. W. (1993). Creativity: beyond the myth of genius. New York, W. H. Freeman and Company.

Weitzman, M. L. 1996. "Hybridizing growth theory." *American Economic Review* 86:207-212.

-. 1998. "Recombinant growth." *Quarterly Journal of Economics*:331-360.

White, H. C. 1992. *Identity and control: a structural theory of social action*. Princeton: Princeton University Press

Young, A. (1928). "Increasing returns and economic progress." The Economic Journal **152**: 527-542.