Nanotechnology and the Emergence of a General Purpose Technology

Stuart J.H. Graham[a,c] and Maurizio Iacopetta[b*]

[a] Georgia Institute of Technology, College of Management
800 West Peachtree Street, Atlanta, GA 30308 USA
stuart.graham@mgt.gatech.edu Tel: 404.385.5107 Fax: 404.894.6030

[b] Georgia Institute of Technology, School of Economics
781 Marietta Street, Atlanta, GA 30318 USA
maurizio.iacopetta@econ.gatech.edu Tel: 404.894.4913 Fax: 404.894.1890

[c] University of California, Berkeley Center for Law and Technology

31 December 2008

Abstract: This article examines how closely nanotechnologies share some identifiable characteristics with previously-identified General Purpose Technologies (GPT). Using US patent data from 1975-2006, we test for nanotechnology's "pervasiveness" and for the likelihood that it is spawning follow-on innovation. Introducing several methodological innovations, we employ concentration indices such as the Gini Index and Lorenz curve, and construct a 'knowledge dissemination curve' (KDC) for different technologies, thereby offering evidence that nanotechnology shares GPT characteristics with other candidates like information technologies (IT). We also for the first time use three different definitions of a "nanotechnology patent" and calculate "generality" indices, finding that nanotechnology patents are significantly more general than are IT patents. In a further contribution, we suggest that materials—principally steel—may demonstrate the characteristics of a GPT, and provide a historical parallel between the advancement of steel technology and nanotechnology (138 words).

Keywords: Nanotechnology, General Purpose Technologies, Patent Analysis

EL Classifications: 0330, 0300, 0340

*Acknowledgements: The authors contributed equally to this work. We would like to thank Richard Freeman, Philip Shapira, Jan Youtie, Lynn Zucker, and several anonymous reviewers, as well as participants at the NBER-NSF Nanotechnology and Nanoindicators Conference held in May 2008. This research was supported by the National Science Foundation under grants NSF NSEC CNS SES #0531194 and NSF NSEC CNS SES #0531146. Graham also acknowledges support from the Ewing Marion Kauffman Foundation. All opinions, mistakes, and omissions are the sole responsibility of the authors.

1. INTRODUCTION

While modern economies are driven by knowledge expansion and innovation, there is disagreement on how to characterize the process of technological change. On one hand, a significant part of the endogenous growth literature describes the process of innovation as a sequence of incremental changes that either improve the quality of inputs or expand the menu of technologies (see Barro and Sala-i-Martin, 2005 for a review). But in the mid 1990s a group of scholars, alert to the contribution of economic historians, began a systematic effort to give an alternative formal representation to technical progress (see Helpman, 1998). Their approach conceptualized the process as non-linear, in contrast to the linear one hypothesized in the earlier endogenous growth literature.

The non-linear model cycle developed by these scholars starts with a major breakthrough technology which opens up new opportunities to develop incremental innovation and which, in turn, facilitates the use of the radical innovation ex post. This conceptualization of innovation implies a substantially different view of the long-run dynamics of an economy, in which phases of development are organized along the introduction and the diffusion of radical "game changing" innovations, such as the steam or the combustion engine, electricity and the dynamo, and information technologies and the computer. Bresnahan and Trajtenberg, (1992, 1995) introduced the term "General Purpose Technology" (GPT) to describe the innovation at the heart of technological change, suggesting that such innovations would be a driver of economic dynamism in modern economies. In contrast to the endogenous growth literature, the newer GPT literature builds on the insights of Rosenberg (1976) and other economic historians who have suggested that, if we want to adequately understand the process of technology development, we ought to pay attention not only the rate of technical change but also to its direction.

In this paper, we pose the question of how closely the burgeoning "nanotechnologies" share some identifiable characteristics with past breakthrough innovations. Understanding the main features of an emerging GPT is relevant for making predictions about its effects on the dynamics of the economy, and for informing government policies on how best to allocate public resources on the development and rapid diffusion of new technology. Bresnahan and Trajtenberg (1995) have theoretically demonstrated that if innovation is driven by GPTs, then in a decentralized economy firms under-invest in the development and adoption of such new technologies. Predicting the key technologies in which to invest is a growing concern of countries that are trying to "catch up" with nations performing closer to the world's constantly-moving technological frontier. As an example of such planning, Taiwan and Malaysia have made huge investments to attract foreign firms to perform R&D in many technologies, intending that science and technology will spill over into other sectors of their economies (Liu and Chen, 2003; Bunnell, 2003).

We will conduct our analysis of the emergence of a General Purpose Technology by examining patent data, including citations made to other patents. The data we use originate primarily in the US Patent and Trade Office (USPTO), although we also use some data from the European Patent Office (EPO). In completing our analysis, we rely upon three aspects of a GPT emphasized in the literature: technological dynamism (the characteristic that the technology should improve over time), pervasiveness (that it should be adopted in many sectors), and its propensity to spawn other innovations (that it should accelerate the invention of follow-on processes, products, or materials). Our investigation will focus mainly on these last two aspects because these are the primary source of positive externalities.[1]

As a methodological contribution, we propose a novel method to evaluate the feature of innovation spawning through the use of patent citation data. In the existing literature exploiting patent data, a high value in a patent's generality index has been read to imply *both* greater pervasiveness and a higher likelihood to spawn follow-on innovations in other sectors. Our paper for the first time we are aware separates these two aspects of a GPT and empirically tests for innovation spawning by analyzing the distribution of "knowledge spillovers" across patent classes. We also add to a growing literature that tests the GPT characteristics of different technologies using patent data. For instance, Hall and Trajtenberg (2004) selected twenty-eight of the most highly cited patents—regardless of the technological

[1] As pointed out by Lucas (2002), externalities are likely to be a key aspect of any explanation of modern development. Recently Klenow and Rodriguez-Claire (2005) have tested several classes of models which feature externalities in the accumulation of knowledge possessed by firms, by workers or by researchers, and other models in which externalities are absent. Through these tests, they were able to confirm Lucas' intuition that models without externalities cannot match basic macroeconomic stylized facts.

fields to which they belong—and evaluated their generality indices. Similarly, Moser and Nicholas (2007) used the generality index to specifically test how closely advances in electricity (arguably the most cited historical example of a breakthrough technology) during the early 20th century exemplify a GPT. Our paper examines nanotechnology and compares this new technology with GPT candidates identified in these previous efforts.

Our article is a follow-on to an earlier work (Youtie et al. 2008) which remains, to the best of our knowledge, the only study that has tried to test with patent data whether nanotechnology is a GPT.[2] In that earlier paper, we concluded that the generality index associated with nanotechnology is comparable with that calculated for IT, thus finding some preliminary evidence of nanotechnology as at GPT. Our evidence was limited, however, both in terms of its scope, coverage, and timing. With this paper we complement that study, bringing new and updated data, and thus add to that evidence while also investigating more specifically whether we can find evidence that nanotechnology demonstrates the GPT characteristic of spawning innovation in other fields.

Our paper also offers a contribution in terms of nanotechnology data. In previous work that examined nanotechnology patenting in the United States, it has been common to rely exclusively on one definition—whether it be the classification-based definition of the USPTO or a simple keyword search of patent titles or abstracts, such as a Boolean search of any term including the stem "nano-." In our study, we rely upon three different definitions of US-granted "nanotechnology patents:" The technology classification applied by the USPTO (3-digit), a comprehensive keyword definition created at the Georgia Institute of Technology's School of Public Policy, and, for the first time to our knowledge, the European Patent Office's experimental patent classification for nanotechnology, Y01N.

Our methodology is inspired by authors who have used industry data instead of patents, and in so doing were able to examine pervasiveness and innovation spawning separately. We build on Jovanovic and Rousseau (2005), in which they test for the pervasiveness of electricity and of information technology

[2] Shea (2005) and Palmberg and Nikulainen (2006) also hypothesize that nanotechnology may be a GPT because it is likely a disruptive and radical technology but follows a different methodology to verify the conjecture.

(IT) by using diffusion curves across industry sectors (they used the share of horse power electrified, and the IT shares of capital stock, respectively). They inferred that follow-on innovations were spawned by the GPT, using as evidence the frequency of initial public offerings (IPOs) by firms that embraced these new technologies. We similarly evaluate pervasiveness by building dissemination curves for the three selected technologies, but instead of using stock market data—which is not as ubiquitously available for nanotechnology firms—we rely upon patent citation patterns. We exploit the historical aspect of the patent data and proxy for knowledge flows running from the nanotechnology field to other innovation fields. For evaluation, we compare the results for nanotechnology patents with those we obtain by examining a technology often considered to be a GPT (i.e., information technology (IT)) and a technology that had specific application primarily to one industry (i.e., the combustion engine (CE) in the automobile industry). The main aspect that we want to uncover is whether these external knowledge flows are limited to a handful of patent classes, or whether the effect of these flows can be seen more generally across many classes.

In a further contribution, we also provide a historical parallel between the advancement of steel technologies and nanotechnology. We argue that one characterization of GPT provided in the seminal article by Bresnahan and Trajtenberg (1995) is both restrictive and difficult to demonstrate: That of having a "generic function." This principle is intended as a general one around which new complementary technologies are developed. Examples of technologies with a generic function are continuous rotary motion for the steam engine and transistorized binary logic for integrated circuits (Lipsey et al. 1998).

Our reading of the literature on the technological features of nanotechnology, however, leads us to believe that this technology's main potential is to bring about a revolution in the creation and use of new materials. This view is consistent with the historical record. Studies of GPTs have generally focused on the technologies of steam, electricity, and IT (all of which seem to have a generic function). Steam and electricity are clear examples of new methods of delivering energy for work, and the latter was so pervasive that some have called the first quarter of the 20th century as the "Age of Electricity" (see Moser and Nicholas (2007) for a contrarian view on the revolutionary aspects of electricity). Similarly, the last quarter of the 20th century has been labeled as the "Age of Information Technology." [3] Historians have referred to the last third of the 19th century as the "Age of Steel" (see for instance Landes, 1969, pp. 249-259). Like nanotechnology, steel was a new material, which, we contend, fulfills the three GPT definitional criteria.

Our study is subject to the usual limitations of studies that attempt to infer the features of innovation through patents. First, not all innovative activity is reflected in the patent system. Secondly, a given patent class (assigned by patent examiners) does not necessarily have a clear correspondence with a technological field. While there is little we can do to correct for the first point, we do address the second issue by using alternative classification systems for robustness (we use both USPTO and WIPO patent classes), and by using alternative methods to define information technology, computer / software, and nanotechnology patents.

The balance of this paper is organized as follows. Section two discusses links between our paper and the GPT literature, making the case that the latter has neglected the historical role of 'new materials' in propelling dynamism in economies. Section three moves to an analysis of patenting, demonstrating that patenting activity is highly concentrated in a few technology classes. Section four discusses the concept of "pervasiveness" and examines its presence by using knowledge-flow dissemination curves. In section five we provide a comment on the question of technology "convergence." Section six builds and explains our patent citation "generality" indices, finding strong and persistent evidence of "pervasiveness" in nanotechnology patents. We present our conclusions in section seven.

2. GENERAL PURPOSE TECHNOLOGIES

2.1 Developing the Concept

The starting point of the GPT literature is the seminal NBER working paper of Bresnahan and Trajtenberg (1992), who criticize the smooth view of the innovation process underlying the theoretical

---

[3] See the Wikipedia entry at http://en.wikipedia.org/wiki/Information_Age for a list of anthologies, reviews, and journals, the titles of which include the expression "Information Age."

models by Romer (1986, 1990) and Aghion and Howitt (1992).[4] According to Bresnahan and Trajtenberg (1992), much of modern economic growth unfolds in a particular way. First there is a major innovation, which is relatively rough and subject to gradual improvement. This basic technology spurs new secondary innovations in a like-tree structure. As the number of downstream technology applications increase, there are greater incentives to improve the basic technology, making it more and more efficient. At the same time that the basic technology is being perfected, a wider breadth sectors find it beneficial to adopt it.

Contrary to standard economic growth theory, the GPT literature pays attention to the dissemination process. The basic idea of the model proposed by Helpman and Trajtenberg (1995) is that GPTs do not come "ready to use off the shelf"—they must be complemented by the development of a new family of equipment (and processes) which requires the diverting of resources from production into development. A new GPT will be adopted only after the number of new secondary technologies hits a critical mass. During this "sowing" phase, measured output declines as the economy is preparing itself to replace existing equipment associated with a pre-existing GPT with new equipment complementary to the new GPT, a phenomenon that may generate a productivity slowdown.[5] This mechanism has not been corroborated empirically, at least at low frequency, because in the post-WWII period US R&D has expanded at a relatively constant pace, both as a percentage of GDP and as a fraction of the labor force, whereas the economy has shown periods of expansion and depression. Alternatively, some have conjectured that the adoption of a GPT causes a productivity slowdown because there is a hidden adoption cost: Firms must be reorganized and workers need to acquire new skills specific to the new GPT (David, 1990).

Several technologies have been considered as candidates for being GPTs (see Lipsey et al. 2006 for a review). Interestingly, most authors have been attracted by the revolutionary role of new forms or sources of energy (steam, combustion, electricity, motors, and engines), new forms of transportation

[4] To be fair, Bresnahan and Trajtenberg (1992) is not the first formal work to characterize growth as a mix of major innovations, each of which was followed by a family of incremental innovations. The pioneering work by Jovanovic and Rob (1990) generates waves in production by assuming that a groundbreaking technological change is followed by secondary innovation.
[5] For a test of this slowdown, see Basu et al. 2008.

(ships, railroads), or some combination of them (steam-powered rail engines). There is a notable exception, however, in that 'new materials' have not been courted in the same manner. This omission is surprising given that, for example, economic historians have recounted how the rise of German, Swiss, Danish, Italian, and Polish industrial might was centered on advances in its chemical industry.[6] Even today, chemicals account for a large fraction of the most-cited patents, and many of these patents score highly in the 'generality index' (see Hall and Trajtenberg, 2004, Table 8, and Moser and Nicholas, 2007, Table 1).

2.2 A Comparison between Nanotechnology and Steel Technology

Materials, we argue, may serve as a GPT. In his classic work, Landes (1969) divides the industrialization process of Western Europe into 'technological eras,' each of which is driven by a technological prime-mover. James Watt's steam engine was the prime-mover of the first phase of the Industrial Revolution, and its diffusion throughout the economy revolutionized the organization of existing sectors, such as metallurgy, textiles, and transportation.[7] But Landes attributes the 'continental' European emulation for the creation and use of new materials: That is, the rise of the new chemical industry. Indeed, he puts at the center stage of modern German industrial development that nation's advances in metallurgy, along with the adoption of new sources of power (steam, combustion engine), and the distribution of energy (electricity). Metallurgy is properly a branch of applied chemistry, but given that modern economies are built on steel, historians have tended to consider it as separate from other chemical manufacture. Landes names the last third of the 19th century the "Age of Steel."

[6] A recent collection of works in Homburg et al (1998) argue that the period from 1850 to 1914 was extremely important in the development of the chemical industry. Chemistry combined technology and science to become one of the most important industries in the Second Industrial Revolution. A similar argument is developed in Landes (1969). Despite this, growth economists seem to have been less interested in chemistry than some other sciences.

[7] Of course, hundreds of other innovations occurred in England during this period, and only some are directly associated with the steam engine, as witnessed by the Crystal Palace Exposition in 1851–which represented probably the apex of British manufacturing hegemony–and other expositions organized later in other European cities.

It seems to us that steel, a material, easily satisfies the three main criteria of a GPT: pervasiveness, the spawning of downstream innovations, and technological dynamism.[8] In terms of pervasiveness, it goes without much argument that modern industry is built on a framework of steel. Furthermore, we can see it used in many household appliances and widely present in transportation infrastructure. A parsimonious way of looking at the dissemination of this material in its technological infancy, and the concomitant construction of new machines, would be to examine total steel production by the top producer countries of Britain, France, Germany and Belgium. In 1861, before the Bessemer process of mass steel production was adopted, aggregate steel production in these four nations was approximately 125,000 tons per year. In 1913 after the Bessemer process had taken hold, total production amounted to 32,000,000 tons, a gain of 83 fold, or a growth of approximately 10 per cent per year (Landes, 1969, p. 259). This evidence supports the historical record—steel was being widely applied to uses throughout these economies.

There is little doubt that steel has also spawned complementary innovations. Steel's strength in proportion to weight and volume makes possible the creation of lighter, smaller and yet more precise and rigid–hence faster–machines and engines. This strength also makes it an excellent construction material, especially in shipbuilding, where the weight of the vessel and the resulting space left for cargo allow transportation efficiencies. Hence steel allowed the creation of new more efficient machines and engines, induced architects and engineers to create lighter designs for industrial plants, buildings and houses, and made possible the creation of a large array of equipment used by both industry and household alike.

If steel is to fit the definition of a GPT it should also exhibit "technological dynamism," a characteristic for which we can test using two methods. We can either measure improvements in the quality of the product, or we may calculate a reduction in the product's production cost (this latter hypothetically reflecting corresponding process innovations). To be parsimonious, we report data on the

[8] Previous characterizations of GPTs often included a 'generic function' such as the steam-engine's 'rotary motion.' But this generic function imposes an unnecessarily restrictive definition on a major breakthrough technology. Accordingly, we find it unnecessary to consider it further.

prices of steel from Landes (1969).[9] While in 1815 the price of steel was £ 700 per ton in England, by the middle of the 19th century it had fallen to £ 55 in Sweden. This reduction was due primarily to the introduction of the Huntsman production technique. The "puddling" process followed quickly, helping to reduce the price even further to about £ 22 per ton in 1850. By the time the Bessemer and Siemens-Martin processes were introduced in the late 1850s, the industry was able to squeeze the price to £ 7 per ton. Hence, in 1860 the price of steel was approximately two orders of magnitude less costly to produce than it had been in 1815, amounting to a rate of price decline of about 10 percent per year. In the following 35 years, the price fell an additional 90 per cent, corresponding to a decline of 2.5 per cent per annum.

How does this decline in the price of steel compare with other GPT candidates such as IT, electricity, and motors? According to our calculation based on Jovanovic and Rousseau (2005, Fig. 11), over the span of a century the price of motors and vehicles declined at a similar rate (2.3 per cent per year). The rate of price decline in IT equipment since the earlier 1960s is ten times larger, but exhibits a quite exceptional phenomenon from a historical perspective.

The parallel between steel and nano-materials must be drawn at a technology level, since hard data on prices, investments and the like are not yet readily available for nanotechnology. There are qualitative indications that the three main GPT attributes (pervasiveness, technological dynamism, and fostering innovation in downstream sectors) may be present in nanotechnology. As with steel, there is not a generic function that can be associated with nanotechnology. Nevertheless, we see in nanotechnology a class of materials that has the potential to radically change the manufacturing process, in a manner possibly as far reaching as did steel in the second phase of the industrial revolution.

Steel's compactness and strength were the two defining characteristics of its utility for follow-on and complementary equipment. Similarly, nanotechnology is defined by its scale, ranging from one to

[9] Landes (1969, pp-253-55) provides what appear to be reliable data on the price of cast steel.

100 nanometers (nm).[10] A switch to nano scales is finding application in virtually every aspect of production.[11]

Signs of nanotechnology's pervasiveness already abound. Nano-materials are commonly used to replace existing larger-scale ones and to solve new technical problems. An example of this "replacement" effect can be found at the core of information technologies. Harriott (2000) reports that concerns in many industries about the possibility of Moore's Law reaching its physical limit have begun to be addressed by nanotechnology's potential to sustain circuit density increases through small scale lithography alternatives such as nano-imprint lithography or eventually self-assembly (see also Arnold, 1995).[12] Another example can be found in medical applications: Matter at the nano-scale exhibits novel properties which cannot be projected from either larger or smaller scales (Kostoff et al., 2006; Tannenbaum 2005). For instance, it has been discovered that the release of nano-scale agents can be triggered by differences in the acidity or alkalinity of the surrounding medium, a mechanism unique to materials of this scale.

In our analysis of the patent data below, we present a quantitative assessment of nanotechnology as a force that induces further innovation. But our quantitative evidence is supported by other qualitative data, including suggestions concerning the body of innovation spurred by nanotechnology. Lux Research (2006), for instance, proposes a specific value chain comprised of an initial set of nano-materials (such as carbon nano-tubes) which may be used as inputs into intermediate products. These intermediates can assimilate such nano-materials into coatings to enhance properties of finishes, and ultimately into final products which can integrate these coatings into a diverse set of product offerings. These final products may include automobiles, airplanes, electronics displays, nano-treated clothes, refrigerator surfaces with microorganism growth inhibitors, self-cleaning windows that oxidize organic matter, and the like. Lux

[10] One nanometer equals one billionth of a meter.
[11] Although we emphasize the similarity between steel and nanotechnology, there are important differences as well. Phil Shapira comments that steel is produced by a specific sector but is used by many sectors: Manufacturing, construction and engineering. In contrast, nanotechnology is both produced and used in many sectors. We discuss in this paper how a technology may be considered a GPT if it satisfies a set of parameters calculated on the application side rather than on the production side.
[12] Moore's law is not a physical law per se, but instead was a prediction made by Intel's Gordon Moore in the 1970s that CPU transistor counts, and thus computing power, would double every two years.

suggests that this value chain is supported by a set of complementary tools including scanning probe microscopes, nanofabrication tools, and computer modeling systems.[13]

Scope for improvement is likely associated with a combination of size reduction, lower production costs, and greater complexity of nano-materials. Nano-applications in semiconductor manufacturing technology aided in reducing processing from 90 nm to 65 nm 2005, and again to 45 nm in 2007 (Kanellos, 2005, 2006). An interesting case of reduction is documented by Lux Research (2006) for AFM instruments, whose prices—adjusted for the number of features—have declined due to the application of nanotechnology. Perhaps the most elaborate prediction on the breath of nanotechnology's technological improvement is to be found in the works of Roco (2004, 2005), where it is predicted that the field will evolve a level of complexity bringing benefits equal to those of information and communication technologies (ICTs) or biotechnology.

Such predictions aside, we contend that, to adequately examine the role played by 'materials' like nanotechnology as GPTs, it is useful to focus on the process of technology diffusion. Previous analyses of the technology diffusion process have offered a number of new questions (Crafts 2004; Atack et al 2008; Basu et al. 2008; Rosenberg and Trajtenberg 2004; Kim 2005; Klaus and Esteban Rossi-Hansberg 2008). Accordingly, we will examine evidence of knowledge spillovers running across invention sectors as a part of our analysis.

3. OUR PATENT ANALYSIS

To complete our empirical analyses, we use US patent data from 1975 to 2006. We do not believe that limiting our data to US patented inventions presents a barrier to generalizing our results. First, the United States remains one of the world's largest markets, and firms with a global marketing strategy will generally patent in Europe, Japan, and the US. Second, although a substantial share of global innovative activity in emerging technology fields takes place outside the US, it is nevertheless unlikely that any

[13] This view is not without its critics. Meyer (2006) suggests that nano-materials support the value chain rather than constituting the initial element because of the linking function these materials play. Based on a cluster analysis of more than 5,400 patent classifications, he offers another candidate nano-industry structure: Measurement-focused, materials in composites and coatings, pharmaceuticals/chemicals, and electronics/devices, with instrumentation serving as a connecting and enabling technology.

limitation of the geographical composition of innovation activity significantly affects our statistical analysis. This latter point is buttressed by our use of both the USPTO and the World Intellectual Property Organization (WIPO) patent classification systems in our analysis.

In our analysis, we employ patent citations. These citations are patent references that newly-granted patent documents include as in indication of “prior art” for the focal invention. If a focal patent B cites back to some earlier-issued patent A, we know that an individual involved in the invention, the prosecution, or the examination of the B patent believed that the earlier-issued patent A described a critical piece of knowledge upon which the invention specified in B built. Although citations have commonly been used as a source of information for measuring knowledge spillovers, there are three caveats that researchers ought to bear in mind. First, the person adding patent A to the B document may be the inventor, the inventor’s patent agent or lawyer, or it may be the patent examiner at the USPTO (Alcacer and Gittleman, 2006). Accordingly, it is difficult to draw a direct causal link between the piece of knowledge embodied in document B and the inventor of the invention described in patent A. That said, in the aggregate these citation patterns are useful for tracking the development of a technology over time. Second, non-patented innovations may draw from technical knowledge described in patent documents. Third, inventors may draw important information from reading the relevant scientific literature. These last two observations imply an underestimation of the spillover effects, whereas the first one would bias the result in the opposite direction.[14]

Given a sufficiently long time span, patents may also be meaningfully categorized according to the citations that they receive ex post, often called their “forward citations.” In the above example, patent B would be a forward citation to patent A. These forward citations, however, tend to develop slowly, given that the mean lag between patent application and patent grant is approximately 3 years (Graham, 2004). The count of citations received by a given patent provides a proxy for the “importance” or “technological impact” of a patent. Forward citations have also been shown to correlate strongly with the

[14] Recently, Adams and Clemmons (2006) sought to control some of these problems by complementing citations with firms’ R&D expenditure.

market value of the patent (Harhoff, et al 2005) and with the market value of firms holding the patents (Hall et al 2004).

Figure 1 about here

3.1 Patent Growth Trends

Figure 1 offers a broad view of some of the data we use. It displays the time series for overall US patenting, as well as patenting in the Combustion Engine (CE), Nanotechnology and Software arts. We are interested in CE as a comparison, primarily because it is a mature technology and was closely tied to one particular industry—automobiles—during the 20th century.[15] Software is an important component of information technologies, a class of technologies which have been hypothesized to be a GPT. While information technologies have been diffusing since at least the 1950s, the growth spurt of software-related technologies—and the patenting thereof—tend to be of more recent vintage (Graham and Mowery, 2003). Therefore, at a given point in time, these three technologies would likely be at different phases in their life-cycles. Given their specific trajectories, we expect that the knowledge spillovers to other technologies generated by CE should be constant (or even falling) over our study period (1975 to 2006), while those attributable to IT should be increasing although at a lower rate than those associated with nanotechnology.

The "software" technology time series presented in Figure 1 is based on US patents assigned to classes 707, 709 and 711. These classes roughly correspond to international patent class G06, the primary "software" class identified in Graham and Mowery (2004). Consistent with our assumptions, we see growth over time in the number of software inventions being patented in the United States. Moreover, this growth is occurring at a greater rate relative to both "all patenting" and combustion-engine patenting (Figure 1). Combustion engine patents are included in US class 123 and 60, which roughly correspond to F02 in the WIPO classification (see Graham, 2006).

[15] In 1859 Etienne Lenoir brought forth a motor fired by a mixture of gas and air. Lenoir's prototype consumed too much gas to be commercially competitive. Nevertheless it furnished the pattern, and from then on a large number of engineers and tinkers devoted themselves to the problem. The crucial conceptual contribution was made in 1862 by Beau de Rochas whose four-stroke cycle has since become the standard. But no one put this principle into effect until N.A. Otto combined it in 1876 with pre-compression of the charge to produce the first practical gas engine. The Otto 'silent' engine swept the market. This form of internal combustion offered important advantages over steam. It was more efficient, cleaner, and the nature of the fuel was such that feeding was easier to automate.

Figure 2 about here

We utilize three different but overlapping classifications of nanotechnology patents in our analyses. The relationship among these three definitions is summarized in Figure 2. Corresponding to the data plotted as "nanotechnology" patents in Figure 1, we use the experimental classification applied at the USPTO, patent class 977. From 1975-2005, a total of 4,216 patents have been assigned by the USPTO to this experimental class.[16] The trend for these nanotechnology patents shows growth in inventors seeking patents from 1975-2005, with a growth rate at least as high as that shown in software for most of the 1990s, although the growth rate in nanotechnology appears to have slowed during the period 2002-2005. While we use the USPTO nanotechnology definition exclusively for this trend analysis, elsewhere in our paper we employ comparisons between the USPTO "nanotechnology" classification and a key-word definition created at the Georgia Tech School of Public Policy (the use of which produces 9,707 patent matches between 1975-2005) and another experimental classification created by the European Patent Office, Y01N (matching 10,148 US-issued patents during 1975-2005).

Figure 1 demonstrates that the USPTO issued about 72,000 patents in 1975 and approximately 175,000 in 2006, implying an annual growth rate of 2.9%.[17] As one would expect, a smaller growth rate is shown in the patenting of combustion-engine inventions over the same time period: 2.6%. Conversely, the two emerging technologies of software and nanotechnology display a relatively high growth: 15% and 17% respectively (due to the paucity of patenting prior to the mid-1980s, we limited the calculation for nanotechnology to the period 1986-2006).

3.2 Technology Concentration

Measured by the international patent classification (IPC), there are approximately 250 three-digit WIPO technology categories of which about 90% have assigned at least 10 patents during 1976-2006. For computational convenience we restrict our attention to this group of patents, which amounts to 2,626,821

[16] The USPTO assigns each "nanotechnology" patent to at least one permanent non-experimental class, and also assigns each patent in parallel to the experimental class 977.
[17] Coincidentally, this figure is close to the annual growth rate of per capita income, which for the same period is about one percentage point lower.

US patents distributed across 226 3-digit IPC groupings. This same stock of patents is distributed across 424 three-digit USPTO patent classes. Since we are interesting in investigating the trail of "knowledge spillovers" across technological fields, we conduct several analyses to gain insights into the distribution of patents across technology categories.

Figure 3 about here

The statistical literature has elaborated several measures to compute the concentration of an attribute (see the excellent review by Atkinson, 1970). Some of these measures have become popular in recent political economy studies that investigate the relationship between the concentration of income, land, or wealth and the process of development. To understand—in Rosenberg's words—the 'direction' of technological change driven by emerging technologies, we borrow from this literature and use the Lorenz curve, the Gini index, and a Quintiles representation of patents across different patent classes (Table 1, panels A and B). Figure 3 contains a Lorenz curve representing patent-class concentration during 1975-2002. A value of x on the horizontal axis is the fraction of the WIPO technology classes ranked in increasing order with respect to the number of patents. The vertical axis reports the percentage of the all patents accounted for by the x fraction of WIPO categories. For instance, the point (0.9, 0.3) indicates that 90 percent of the WIPO classes with the lowest patent frequency account for 30 percent of all patents. Hence, as the curve diverges from the forty-five degree line, the patents in the distribution are increasingly concentrated in an increasingly small number of classes.

Table 1, Panels A and B about here

As another concentration measure, we present a Gini index, reported in the second column of Table 1, Panel B. The Gini index measures the area between the Lorenz curve and the straight 45-degree line, normalized in the zero-one interval. The Gini index approaches zero as the Lorenz curve approaches the forty-five degree line, and goes to one when the Lorenz curve approaches the horizontal axis. A visual inspection of the Lorenz curve for all patents in our sample reveals a high degree of concentration among a handful of patent classes. This feature is confirmed by a Quintile representation of the data in Table 1, Panel A, which shows that the top quintile accounts for 85 percent of the overall patents in our sample.

The first row of Table 1, Panel A also reports the concentration at the top 10^{th}, 5^{th}, and 1^{st} percentile, which account for 68, 50, and 20 percent, respectively.

In order to gain further insights into the dynamics of the concentration indices, we split our observation period into three sub-periods: 1975-82, 1983-93, and 1994-2002.[18] The Gini index and percentile data are reported in the remaining rows of Tables 1, Panel B. All measures suggest an increasing degree of concentration as we move forward in time. The Gini index rises 0.77 to 0.83 from the first period to the second, while the top quintile increases from 0.8 to 0.88 and the top 1 percentile from 0.18 to 0.25 across the two periods.

In sum, our analyses of the patent data show a great deal of concentration among a handful of 'star' technological classes. Furthermore, this concentration has become more and more pronounced during the period for which we have data (1975-2002). In Section 7 below, we clarify the extent to which these two observations affect the main analysis of our paper.

4. INNOVATION SPAWNING

One of the characteristics of a GPT as defined in Bresnahan and Trajtenberg (1992) is that the candidate technology should induce further innovation in other sectors. In Jovanivic and Rousseau (2005) this feature is evaluated by examining the dynamism of patenting activity, and by observing entry and exit in the stock market, this latter indicator intended to proxy for the replacement of an obsolete GPT with a new one. They found that patenting in each of the IT and electricity eras was more intense than in the decades separating the two periods. Our disaggregated data on citations allow us to assess the hypothesis that this pattern will hold at a different level. The phenomenon that we want to measure is the extent to which nano-inventions have been used as an input for inventions in other technological fields. While we examined some anecdotes of nano-materials as inputs in Section 2 above, these isolated cases do not help us to determine whether the "spawning" phenomenon is a general characteristic of nanotechnologies.

Using our comprehensive data, we build a 'knowledge dissemination curve' with the patent citation data. Our method employs patent citations, exploiting this information to infer "knowledge

[18] Data limitations necessitate us restricting the coverage to no later than 2002.

flows" from inventions assigned into one patent technology class to later-in-time inventions embodied in patents assigned to a different technology class. As an example, imagine that we are interested in measuring the spillover from nano-inventions to a given technological class *N*. For the sake of simplicity assume in period t that 100 patents are granted, and that a fraction *r* of these patents cite backward to at least one nanotechnology patent. The variable *r* is our proxy for the R&D externalities generated by nanotechnology, calculated by employing cross-technology class citations on issued patents.

We compute this index for each of the USPTO's 424 different patent classes during three different eras for three distinct technologies (as defined by 3-digit patent classifications): Combustion engines (CE), information technology (IT), and nanotechnology. We selected IT and nanotechnologies because the former is a GPT candidate technology, and we are attempting to discover whether the latter shares characteristics with the former. For comparison purposes, we also chose CE technology because it is closely identified with a single industry (automotive), and has not been cited as a GPT. Both CE and IT have a long technology pedigree, and thus we were able to choose three periods for analysis that permitted us a sufficient "forward citation window" in which to have a more complete analysis, settling on the cohorts of 1986-89, 1992-95, and 1998-2001 (spanning 16 years). Due to the late emergence of sufficient numbers of nanotechnology patents, we settled on the cohorts 1990-94, 1996-99, and 2002-05 for an analysis of our nanotechnology patents (also spanning 16 years).

To conduct our analysis, we obtained a time series of the index *r* consisting of three points, each of which is meant to represent the flow of knowledge during a period for a particular candidate technology. Since it would be cumbersome to represent over 400 time-series on a plot, we chose to depict *r* by percentile for the top 5% of the distribution, and by deciles for the remaining part of the overall distribution. In other words, in every period the 424 USPTO classes are ranked in increasing order relative to the index *r*, and only the values of *r* associated with the PTO class located at the first, second, et seq. deciles or at the 95^{th}, 96^{th} et seq. percentiles are plotted. The outcome of this procedure is what we call a 'knowledge dissemination curves' (henceforth KDC) for a given technology. In Figures 4 and 5, the

KDC represents nanotechnology, whereas Figures 6 and 7 are associated with information technologies, and Figures 8 and 9 are associated with combustion engines.

Figures 4, 5, 6, 7, 8, and 9 about here

Figures 8 and 9 show that the 98th and 99th percentiles of the combustion engine KDCs correspond to approximately 0.2 and 0.25 and show a mild positive time trend. The corresponding KDCs for nanotechnology in Figures 4 and 5 begin below 0.02 at the leftmost (earliest) point and rise to 0.07 and 0.09 in the most recent period. It does not surprise us that citations are greater in a mature technology than in an emerging one.

The interesting novelty is the rapidity at which the gap closes between these two sets of KDCs. Should these time trends remain stable, the figures suggest that the gap at the 98th and 99th percentile would be eliminated in a bit more than one decade. A quick inspection of the lower percentiles reveals the difference between the two families of KDCs is significantly smaller, suggesting that catching-up may require an even shorter period of time.

The comparison between IT (Figures 6 and 7) and nanotechnology KDCs at the 98th and 99th percentile shows a much wider gap. The information technology KDCs reach a level of about 0.8 and 0.5, respectively. Given these trends, more than half a century would be required for the nanotechnology KDCs to reach the same level as demonstrated for IT. But it is not a given that slope of nanotechnology KDCs will remain constant: Indeed, trends may increase sharply and follow the typical S-shape of many dissemination curves.

Perhaps the most surprising result of these KDCs is that the 90th percentile of IT and nanotechnology KDCs reach the same level in the final years of our examination period. In fact, IT "knowledge spillovers" appear to be more concentrated in a handful of patent classes than those characterizing either combustion engines or nanotechnologies. This observation may be consistent with another type of dissemination curve employed in Jovanovic and Rousseau (2005), who used as a variable the percentage of IT capital investment across industries. They found that the dissemination curves tend to flatten very quickly as one moves down on the percentiles. Conversely, they found greater regularity in the dissemination of electricity. Our data suggest that similar patterns may hold for knowledge diffusion.

In brief, our dissemination curves suggest that the knowledge embodied in nanotechnology is spreading more evenly across other patent classes than either the knowledge embodied in information technologies or combustion engines. We do find, however, that the intensity of spillovers from nanotechnology is not as high as in the other two technologies we examined. Further analysis is required to understand the extent to which our results depend on the fact that nanotechnology is in an early phase of its technology life-cycle, or whether this feature that will persist.

5. ON KNOWLEDGE CONVERGENCE

One interesting feature of our KDCs is that they generally have a positive trend. We observe this positive trend for CE patents (which represent a mature technology that was primarily adopted in one—admittedly large—industry), for information technology (a well-established GPT technology), and for nanotechnology (an emerging GPT candidate). This finding is inconsistent with our initial hypothesis that a mature technology should show a flat or perhaps declining pattern of "knowledge spillovers."

While it is difficult to generalize from one mature technology, the finding raises the question: Is there a common force that drives dissemination curves in a positive direction? While it is tempting to interpret these patterns as evidence of an increasing level of "knowledge spillovers" across all technological fields over time (Grodal and Thoma 2008), we have too little evidence to make any such heroic contention in this article. We do, however, stress that the debate on cross-pollination of ideas is relevant for the patent literature that bases the GPT test on generality indices, for these may be inflated by a phenomenon of knowledge convergence.

Table 2 about here

Our data do, however, allow us to compute the opposite of cross-pollination indicators. For a given number of WIPO patent classes, we calculate the ratio between the number of same-class citations and the overall citations in that focal WIPO patent class (for a list of these classes, see Table 2). As that ratio increases, the flow of knowledge (as indicated by patent citations) from other technological areas decreases. While the level of the ratio is difficult to interpret, its variation over time gives a clear indication of whether a given patent class is building more or less from other patent classes. A reduction in the ratio for the majority of patent classes would suggest a “convergence” across different areas of research and development. Figure 10 plots the ratio just described for two periods: 1985-87 (horizontal axis) and 1992-94 (vertical axis). The chart is comparative across time periods: If a patent class lies on the 45 degree line, that position indicates there is no change in the relative “importance” of knowledge (as measured by citations) that flows from other fields. If a class lies below the 45 degree line, such position indicates that the focal technology class is drawing relatively more from other fields.

Figures 10 and 11 about here

An examination of Figure 10 does not suggest to us a clear pattern of convergence from the second half of the 1980s to the early 1990s. In fact, the number of technology classes falling below the 45 degree-line are only slightly more numerous that those falling above it. Accordingly, we repeat this graphical analysis during a more recent period, depicting the results in Figure 11. By plotting 1992-94 is on the horizontal axis and 1999-2001 on the vertical one, we are able to demonstrate that in more recent years, the majority of patent classes plotted below the 45-degree line. Such placement suggests that during the 1990s, a shift occurred where the majority of technology classes were benefiting more from knowledge inputs derived from different technological fields. Our analysis thus supports the hypothesis that “knowledge spillovers” between technologies have accelerated during the 1990s, lending some credence to the “convergence” hypothesis.

6. GENERALITY SCORES

In order to test for the “pervasiveness” of a GPT, Hall and Trajtenberg (2004) suggested using a focal patent’s forward citations to generate a “generality score.” The generality score is a species of the Herfindahl-Hirschman concentration Index (HHI) of the patent classes assigned to the focal patent's forward citations. This resulting measure of "pervasiveness" is defined by the formula

$$Generality_i = 1 - \sum_{j}^{n_i} s_{ij}^2$$

where S_{ij} = share of patent i's forward citations in class j. The theory underlying the use of this measure is that it captures information about the extent to which the focal patent is being applied in a wide range of technologies—the so called "pervasiveness" of a patented invention. As a patent's "Generality" score increases, that patent is being cited by patents in a broader range technology classes, and we therefore infer that the patent is being applied more broadly across other technologies.

If we examine the "Generality" scores across all patents in a particular technology (in nanotechnology for instance) and compare these against the scores for patents in other technologies, we may infer something about the "pervasiveness" of the technology's application throughout the economy, at least as compared with other technologies. Obviously, any measure built in this way will be very sensitive to right truncation. In the study of new and embryonic technologies in which the patent record is slowly developing, the absence of a sufficient "forward" time window will pose great difficulties in calculating a useful "Generality" index for individual patents, and by extension entire patented technology areas. The trends depicted in Figure 1 demonstrate that, in the emerging nanotechnologies, substantial numbers of patents began to issue from the USPTO in the 1990s, thus giving us a sufficiently long "forward window" to develop credible "Generality" scores on the earliest patents issued in this new technology space.

Table 3 about here

Table 3 reports generality scores for several categories of patents, including all patents issued in years 1976-2005 as well as two definitions of IT patents, and three definitions of nanotechnology patents. The first IT definition is derived from Hall (2004) and is a USPTO-class definition of "computing and communication" information technology patents.[19] The second definition is more specific to computers,

[19] By this definition, ICT patents are those assigned to US classes 178, 333, 340, 342, 343, 358, 367, 370, 371, 375, 379, 385, 455, 704, 341, 364, 380, 382, 395, 700, 701, 702, 706, 708, 709, 712, 713, 714, 715, 717, 345, 347, 349, 710, 360, 365, 369, 707, 711, 703, 705, 725, and 902.

and particularly to software, and is derived from Graham and Mowery (2005).[20] For simplicity, we conducted our analysis on patents in 5-year cohorts, 1976-2005. Table 3 also discloses the number of patents issued within each sample in the cohort, and also the number of "Cited" patents, the only patents for which the calculation of Generality scores is possible.[21]

Several conclusions can be drawn from our examination of Table 3. First, our results show a substantial truncation effect: The generality scores for patents issued after the year 1995 demonstrate deflation, likely due to the increasingly sparse numbers of citations received as the "forward citation window" becomes increasingly short. More crucially, both the IT patent samples show generality scores significantly higher than those for aggregate US patenting, year on year. Tests for significance in the difference of means confirm these differences in each 5-year cohort at the 99% confidence level. It is interesting to note that while the generality scores for the patents selected according to the "computing and communication" definition applied in Hall (2004) (Sample 2) are higher than overall patenting (Sample 1), those selected by the "computer / software" definition offered by Graham and Mowery (2005) (Sample 3) are significantly higher than both. This observation is borne out by tests for significance reported in Table 3 showing between-sample differences significant at the 99% confidence level within each cohort. The fact that both these IT sampling methods selected patents with significantly higher "generality" scores year on year than in the overall population suggests that these technologies are being adopted comparatively widely in the economy, and thus are examples of 'pervasive' technologies. It is accordingly strong evidence for the existence of a GPT in these information technologies.

We extend this same analysis to nanotechnology patents, employing for the first time we are aware three different definitions of "nanotechnology" patents. Table 3 produces the results of our analysis and the generality scores during 1976-2005 for patents selected by reference to the USPTO's

---

[20] This software patent definition is also derived from the US classes, including classes 364, 395, 700, 701, 702, 706, 708, 709, 712, 713, 714, 715, and 717.

[21] A "cited" patent is a patent that has been cited by some other patent. Uncited patents have had no citations (through 2006, the final year in which data is available to us), and thus for these it is impossible to calculate a generality score.

experimental class 977 (Sample 4), a keyword method innovated at the Georgia Institute of Technology's Public Policy School (Sample 5), and the EPO's experimental class Y01N (Sample 6). The patents in Sample 6 are US-issued patents, connected through priority and family information to equivalent EPO-issued patents assigned to class Y01N. We find consistent results across all three definitions.

Table 4 about here

Generality scores for 1976-2005 irrespective of the definitional scheme we employ demonstrate that "nanotechnology" patents appear to share the characteristic of "pervasiveness" with IT patents. First, under all three definitions, nanotechnology patents are more "general" than patents as a whole (see Table 3 and between-sample tests of significance). Moreover, nanotechnology patents are by and large just as "general" as the patents selected to represent IT technologies (see Table 3 and between-sample tests of significance), and in fact appear more general.

In Table 4, we present the results of an analysis of patents in which the union of all the patents classified as "nanotechnology" under our three definitions (the union of samples 4, 5, and 6) is compared with those classified as "information technology" (Sample 2). Table 4 reiterates the results of Table 3, Sample 2 for comparative purposes. Tests for significance show that the combined nanotechnology sample (Sample 7) shows mean generality scores within each 5-year cohort significantly different from all patents, and also significantly higher than "information technology" patents in all cohorts, with each difference being significant above the 99% confidence interval.

The generality scores presented in Table 3 and Table 4 present strong evidence of the "pervasiveness" of nanotechnology patents, providing us with evidence that these technologies exhibit one of the necessary characteristics of a GPT. Not only do we show significantly higher generality scores year on year regardless of which of three different, although overlapping (see Figure 2), selection criteria we employ, but the generality scores we find for nanotechnology patents compare favorably with those of IT patents, a technology commonly considered a GPT. The trend for nanotechnology patents is strong and consistent: Nanotechnology patents offer evidence of the "pervasiveness" of this emerging technology.

7. DISCUSSION

One serious limitation in using citation data to infer the flow of knowledge spillovers in a particular class of innovations is that both patenting activities and innovation activities may be only weakly correlated. If that is the case, one would tend to overestimate the pace of innovation in a field where for technical, legal, or economic reasons the tendency to patent an innovation is unusually pronounced relative to other fields. For instance, it may be that the remarkable rise of electronic-related patenting since the 1970s may have outpaced the 'actual' underlying rate of IT innovation. This issue has been raised in the literature, and some partial corrections have been offered. Adams and Clemens (2006) for instance uncover the flow of basic research, the patentability of which is notoriously poor, by using data on R&D spending and on scientific publications.

We are left with a question after our analysis, however. To what extent are the conclusions of the previous three sections affected by the possible differential degrees of patenting across technological classes? We will argue that it is unlikely that the generality indices, or our convergence results, are affected by such a bias, but that our knowledge dissemination curves likely are affected. In addition we suggest that the position of the 'actual' IT-knowledge dissemination curves are likely to be below the ones we plotted (in Fig. 6 and 7), but in all likelihood not below the ones plotted for combustion engines (Fig. 8 and 9).

The generality index measures the degree of technology-class dispersion in forward citations of any given patent. By construction this measure does not depend on the absolute number of citations received by a patent. Imagine that there are only two classes of patents, *i* and *j*, and we are interested in computing the generality index of patents A and B, each of which are cited 100 and 200 times by later-issuing patents, respectively. Assume that patent A is cited 20 and 80 times in class i and j, and that patent B is cited 40 and 160 times (the order between *i* and *j* does not matter). In either case the generality index would be identical (in this case, 0.42).

The suggestion that there has been a tendency for knowledge to converge in the second part of the 1990s (and that no such a pattern was identifiable for the first part of the decade) was based on a 'first-difference' argument, which is therefore immune from any concern over the number of patents being issued and assigned into any particular patent class. Specifically, convergence was inferred from observing that for most of the technological classes in the second part of the 1990s, the percentage of intra-class citation was lower than in the first part of the decade. In terms of Figure 11, this observation suggests that most technology classes lay below the 45 degree line. In truth, the wide variation of points along the 45 degree line–aligning roughly from 0.3 to 0.8 on each axis–may be in fact due to the different degree of patentability across technological fields (but that information was not used for the convergence result).

Such is not the case when we consider the Knowledge Dissemination Curves (KDC) discussed in Section 4. We observed that the information technology KDC (at the top percentile) was approximately an order of magnitude higher than was the corresponding nanotechnology KDC. It remains for us to comment upon what part of this gap, if any, can be explained by some unusually high degree of patentability among IT innovations. Our answer consists of two parts. First we present data that support the contention that IT is not a field showing unusually intense patenting activity. Secondly, we propose a way to make an educated guess about the patentability bias.

Table 5 about here

First, we point to evidence reported in the Carnegie-Mellon Survey on appropriability conditions in the US (Cohen, et al, 2000). The CMS presents evidence that patenting in the information technologies has not outpaced the overall population rate. Their Table A1 reports that the product- and process-innovation patenting rate in electronic components, semiconductors, and computing was lower than the average of all surveyed industries in the mid-1990s.

Another piece of information is derived from Hall and Trajtenberg (2004) who base their analysis upon a small sample of the most cited US patents since 1972. A distribution of the top 20 of these patents is contained in their Table 6, showing that only one such patent is associated with IT, fewer than for non-IT electrical (6), chemical (6), mechanical (3), and drugs (3) patents. While limited, this observation is supported by data that we present in Table 5. Table 5 records the top seven WIPO patent classes into which issued US patents have been assigned since the mid-1970s, accounting for about 1/3 of all US patenting during this time period. Although two of these classes (G01 and G06) are associated with IT inventions, neither of these two are ranked at the very top. The number of patents assigned into these categories is of similar magnitude to those assigned into chemistry, electricity, drugs and medical equipment, and packing.

The fact that IT appears from this evidence to not be patented or cited more than other technological fields does not rule out the potential bias of easy-patentability. Such a characteristic may affect all the top patent classes simultaneously. But growth accounting studies provide valuable numerical information to clarify the issue. Several authors have argued that the post-1995 productivity revival in the US and in other advanced countries is in great part attributable not so much to the spread of computers but to productivity gains attributable to the IT industry.[22] There are no studies that point to advances in drugs or chemicals or packaging as being the main driver of the remarkable surge seen in productivity. Hence, if productivity is correlated with the "value" of new ideas, it is reasonable to conclude that ideas in the IT area are not disproportionately patented relative to other technology classes.

Yet, we are willing to accept that some bias towards excess patentability may well exist. Fortunately, our data give us a reliable upper-bound of what such a bias may be. An inspection of Figures 10 and 11 reveals that the lowest values of the intra-class citation ratio are associated with patent class B32 (layered products) and B05 (spraying apparatus), each with a value of about 0.4. Presumably neither of these two classes is related to information technologies—and neither are listed in common definitional schemes for IT. In fact, the electronic classes with the highest self-citation ratios are G03, G06, and G11, each scoring roughly between 0.65 and 8. These are all "electronics" classes and, because of the high self-citation rate refer to like-classed patents about twice as often as do patents assigned in the other non-electronics categories.

[22] Robert Gordon is perhaps the strongest advocate of this view. See for instance his 2000 and 2003 NBER working papers.

We believe accordingly that it is likely that part of the difference between these two groups is accounted for by a genuine difference in the value of the underlying knowledge. Moreover, we believe that individuals inventing new spraying apparatus (class B05) are more likely to incorporate IT-related knowledge (whether fixed in a patent or not) than is an IT-engineer likely to use knowledge developed for spraying apparatus in a new device or piece of software. If by a rule of thumb we say conservatively that only half of the difference is due to the patentability bias, then it is fair to conclude that the position of the "real" IT Knowledge Dissemination Curve is at most 25% lower than the ones we depict in Figures 6 and 7. We can therefore suggest with confidence, even after such a downward correction, a substantial gap between the IT and nanotechnology KDCs would remain.

8. CONCLUSION

In this article, we have examined the existing literature, and applied where appropriate its teaching on GPTs to the emerging nanotechnology arena.. As stated, GPTs drive the expansion of the technological frontier in modern economies, and eventually improve standards of living in the long run. In this context, we asked ourselves whether evidence is now sufficiently strong to argue that nanotechnology ought to be admitted to the club of GPTs. Any such determination has important policy ramifications, especially since GPTs are believed to be the prime-movers of long-run productivity waves.

In contrast to most known examples of GPTs found in the literature, such as steam engines, electricity, and ICTs, the specialized literature we surveyed does not suggest an obvious aspect of nanotechnology that can be labeled as a 'generic function,', such as the rotary motion for motors or the transistorized binary logic for microelectronics. However, by using steel as an historical example, we argued that this feature should not be an exclusion factor, for it does not appear particularly relevant and it is not readily testable with economic data.

We instead offer contributions in methods, data, and results to the current debate over GPTs generally, and nanotechnology in particular. We test for two of the three main defining features of a GPT: "pervasiveness" and a likelihood to foster downstream innovations. Following the lead of Hall and

Trajtenberg (2004) we employ data on US patents and patent citations to build "generality" indices, finding evidence of a consistent and strong "pervasiveness" in nanotechnology innovations.

Moreover, in order to give a quantitative content to the concept of innovation "spawning," we exploited the alternative classification possibilities offered by the USPTO, WIPO and the EPO in nanotechnology patenting. We employed these data to quantify the intensity and direction of "knowledge spillovers" flowing from patented nanotechnology inventions to patented inventions in other fields, and compared these with the spillovers we observed for information technologies and combustion-engine inventions. By applying a methodological advance—namely the 'knowledge dissemination curve'—we obtained evidence that nanotechnology "knowledge spillovers" appear to be more uniformly distributed across technological classes, are less intensive, and have a much more pronounced time trend than those obtained for our other focal technologies. In other words, nanotechnology appears to be following an S-shaped technology development pattern, and to be positioned somewhere prior to the inflexion point. We leave it to further research to verify whether this 'uniformity' feature of nanotechnology across technological field will persist in the subsequent phases of diffusion. If it does, and the amount of "knowledge spillover" continues to rise apace, we will likely before long see the results in some hard economic data, such as in the productivity and investment numbers.

REFERENCES

Adams, J.D, and J.R. Clemmons (2006). "Science and Industry: Tracing the Flow of Basic through Manufacturing and Trade." NBER Working Paper n. 12459.

Aghion, P. and P. Howitt (1992). "A Model of Growth through Creative Destruction." Econometrica 60: 323-351.

Aghion, P. and P. Howitt (1998). On the Macroeconomic Effects of Major Technological Change. Annales d'Economie et Statistique, Vol. 49/50 pp: 53-55.

Arnold, W. 1995, "The SIA Lithography Roadmap," Microlithography World," Winter, pp. 7-11

Alcacer, J. and M. Gittelman. "Patent Citations as a Measure of Knowledge Flows: The Influence of Examiner Citations." Review of Economics and Statistics 88, no. 4 (November 2006): 774-779.

Atack, J., F. Bateman, and R. Margo (2008). "Steam Power, Establishment Size and Labor Productivity Growth in Nineteenth Century American Manufacturing." Explorations in Economic History, forthcoming.

Atkinson, A. B.; (1970). On the Measurement of Inequality. Journal of Economic Theory, v. 2, iss. 3, pp. 244-63

Barro, R. J. and X. Sala-i-Martin (2004). Economic Growth, 2nd ed., McGraw_Hill, New York.

Basu, S. and J.G. Fernald (2008). 'Information and Communications Technology as a General Purpose Technology: Evidence from U.S. Industry Data"

Federal Reserve Bank of San Francisco Economic Review, 2008, pp. 1-15

Bresnahan, T. F., and M. Trajtenberg (1992). "General Purpose Technologies: Engine of Growth?" NBER Working Paper 4148.

Bresnahan, T. F., and M. Trajtenberg (1995). "General Purpose Technologies: Engine of Growth?" Journal of Econometrics, 65, pp: 83-108.

Bunnell, T. (2003). Malaysia, Modernity and the Multimedia Corridor: A Critical Geography of Intelligent Landscapes. Routledge Pacific Rim Geographies.

Cohen, W.M., R.R. Nelson, and J.P. Walsh (2000). "Protecting their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or not)." NBER working paper 7552.

Crafts, N. (2004). "Steam as a General Purpose Technology: A Growth Accounting Perspective." Economic Journal, 114, pp: 338-351.

David, P.A.(1990) "The Dynamo and the Computer: A Historical Perspective on the Modern Productivity Paradox." American Economic Review, Papers and Proceedings 80, n. 2, pp: 355-361.

Devine, W.D., Jr. (1983). "From Shafts to Wires: Historical Perspective on Electrification." Journal of Economic History 43, n.2 pp: 347-372.

Glaeser, E.L. H.D. Kallal, J.A. Scheinkam, and Andrei Shleifer (1992). “Growth in Cities.” Journal of Political Economy 100, n. 6, pp: 1126-1152.

Gordon, R. (2000). Does the "New Economy" Measure up to the Great Inventions of the Past? NBER working paper 7833.

Gordon, R. (2003). Hi-Tech Innovation and Productivity Growth: Does Supply Create its own Demand? NBER working paper 9437.

Graham, S.J.H. (2004). "Continuation, Complementarity, and Capturing Value: Three Studies Exploring Firms’ Complementary Uses of Appropriability Mechanisms in Technological Innovation" University of California, Berkeley Dissertation, Spring.

Graham, S.J.H. (2006). “The Determinants of Patentees' Use of 'Continuation' Applications in the United States Patent and Trademark Office, 1980-99,,” in Intellectual Property Rights: Innovation, Governance and the Institutional Environment, B. Andersen, ed., Edward Elgar Publishers.

Graham, S.J.H. and D.C. Mowery (2003). "Intellectual Property Protection in the Software Industry," in Patents in the Knowledge-based Economy: Proceedings of the Science, Technology and Economic Policy Board, W. Cohen and S. Merrill, eds., National Academies Press.

Graham, S.J.H. and D.C. Mowery (2004). “Submarines in Software? Continuations in U.S. Software Patenting in the 1980s and 1990s,” Economics of Innovation and New Technology (EINT), Vol. 13(5), pp. 417-442.

Graham, S.J.H. and D.C. Mowery (2005). “Software Patents: Good News or Bad News?” in Intellectual Property Rights in Frontier Industries: Software and Biotechnology, Robert Hahn, ed., AEI-Brookings Joint Center.

Grodal, S. and T. Grid (2008). Cross-Pollination in Science and Technology: Concept Mobility in the Nanobiotechnology Field. Manuscript presented at the NBER Conference on Emerging Industries: Nanotechnology and NanoIndicators, May 2008.

Desmet, K. and E. Rossi-Hansberg (2007). Spatial Growth and Industry Age. NBER Working Paper no. 13302.

Hall, B.H. and M. Trajtenberg (2004), "Uncovering GPTS with Patent Data," NBER Working Paper No. 10901, Cambridge, MA: National Bureau of Economic Research.

Hall, B.H. (2004), “Exploring the Patent Explosion,” NBER Working Paper, Cambridge, MA: National Bureau of Economic Research.

Hall, B.H., A. Jaffe and M. Trajtenberg (2005). “Market Value and Patent Citations,” Rand Journal of Economics vol. 36, pp. 16-38.

Harhoff, D., Sherer, M. & Vopel, K. (2003). Citations, Family Size and, Opposition and the Value of Patent Rights-Evidence from Germany. Research Policy, 32, 1343-1363.

Harriott, L., 2000, “Limits of Lithography,” Working Paper, Bell Labs-Lucent Technologies.

Helpman, E. and M. Trajtenberg (1998). “A Time to Sow and a Time to Reap: Growth Based on General Purpose Technologies,” in General Purpose Technologies and Economic Growth, pp. 55-84, edited by E. Helpman. Cambridge and London: MIT Press, 1998.

Homburg, E. , A.S. Travis and H.G. Schröter (1998). The Chemical Industry in Europe, 1850-1914: Industrial Growth, Pollution, and Professionalization, Kluwer Academic Publishers.

Jovanovic, B. and R. Rob (1990). “Long Waves and Short Wave: Growth through Intensive and Extensive Search”, Econometrica, 58,pp. 1391-1409.

Jovanovic, B. and P.L. Rousseau. (2005) “General Purpose Technologies.” In Handbook of Economic Growth, Volume 1B, edited by Philippe Aghion and Steven N. Durlauf. Chapter 18. Amsterdam, the Netherlands: North Holland.

Kanellos, M., 2005, AMD’s 65 nano question solved: it’s a second half thing, CNET News.com, November 2. http://news.com.com/2061-10791_3-5929097.html?tag=nl, 2006, Intel shows test chips made on future processes, CNET News.com, January 25.

Kim, S. (2005). “Industrialization and Urbanization: Did the Steam Engine Contribute to the Growth of Cities in the United States?” Explorations in Economic History, 42, vol. 4, pp: 586-598.

Klenow, Peter, J. and Andres Rodriguez-Clare (2005). “Externalities and Growth.” In Handbook of Economic Growth, Volume 1A, edited by Philippe Aghion and Steven N. Durlauf. Chapter 11. Amsterdam, the Netherlands: North Holland.

Kostoff, R., J. Stump, D. Johnson, J. Murday, C. Lau, W. Tolles, 2006, The structure and infrastructure of global nanotechnology literature, Journal of Nanoparticle Research, 8 (3-4), pp. 301-321.

Landes, D. S. (1969). The Unbound Prometheus: Technological Change and Industrial Development in Western Europe from 1759 to the Present. Cambridge University Press.

Lo, S. and D. Sutthiphisal (2008). “Crossover Inventions and Knowledge Diffusion of General Purpose Technologies: Evidence from the Electrical Technology.” NBER Working Paper 14043.

Lipsey, R.G., C.T. Bekar , and K. I. Carlaw .(1998). “What Requires and Explanation,” in General Purpose Technologies and Economic Growth, pp. 15-54, edited by E. Helpman. Cambridge and London: MIT Press, 1998.

Lipsey, R.G., K.I. Carlaw and C.T. Bekar. (2006). Economic Transformations: General Purpose Technologies and Long-Term Economic Growth. Oxford University Press.

Liu, M. and S. Chen (2003). International R&D Deployment and Locational Advantage: A Case Study of Taiwan. NBER Working Paper No. W10169.

Lucas, R.E. Jr. (1988) “On The Mechanics of Economic Development.” Journal of Monetary Economics 22, n.1, pp: 3-42.

Lucas, R.E. Jr. (2002) Lectures on Economic Growth, Harvard University Press, Cambridge.

Lux Research, 2006, The Nanotech Report: 4th Edition, Volume 1, New York: Lux Research.

Meyer, M., 2006, What Do We Know About Innovation in Nanotechnology? Some Propositions about an Emerging Field between Hype and Path-Dependency, Paper presented at the 2006 Technology Transfer Society Conference, September 27-29, Atlanta, Georgia.

Moser, P. and T. Nicholas. (2004) “Was Electricity a General Purpose Technology?” Evidence from Historical Patent Citations.” American Economic Review, Papers and Proceedings 94, n.2, pp: 388-394.

Moser, P. (2007). Why Don't Inventors Patent? NBER Working Paper No. 13294.

Palmberg, C. and T. Nikulainen, 2006, Industrial Renewal and Growth Through Nanotechnology? Helsinki: ETLA, The Research Institute of the Finnish Economy.

Porter, A.L., M. Meyer, and I. Rafols (2008). The Cognitive Geography of Nanotechnologies: Location and Knowledge Flows of Nano-Research in the Map of Science. Mimeo.

Roco, M., 2004, Nanoscale Science and Engineering: Unifying and Transforming Tools, AIChE Journal, 50 (5), pp. 890-897.

Roco, M., 2005, International Perspective on Government Nanotechnology Funding in 2005, Journal of Nanoparticle Research 7 (6), pp. 707-712.

Romer, P.M. (1990). "Endogenous Technological Change." Journal of Political Economy, 98, pp: S71-S102.

Romer, P.M. (1986) "Increasing Returns and Long-run Growth." Journal of Political Economy, 94, n 5, pp: 1002-1037.

Rosenberg, N. (1976), "Technological Change in the Machine Tool Industry, 1840-1910," In Perspectives in Technology, Cambridge: Cambridge University Press, pp: 9-31.

Rosenberg, N., and M. Trajtenberg (2004). "General-Purpose Technology at Work: The Corliss Steam Engine in the Late-Nineteenth-Century United States." The Journal of Economic History, 64, n. 1, pp: 1-39.

Shea, C., 2005, "Future management research directions in nanotechnology: A case study," Journal of Engineering and Technology Management, 22, 185-2000.

Tannenbaum, R., 2005, "Advances in Materials and Self-Assembly," in Advanced Technology and the Future of Manufacturing, Proceedings, Atlanta, Georgia: Georgia Institute of Technology.

Youtie, J. , M. Iacopetta, and S. Graham (2008). Assessing the Nature of Nanotechnology: Can We Uncover an Emerging General Purpose Technology? Journal of Technology Transfer, vol. 33, n. 3, pp: 315-329.

TABLES

Table 1: Distribution of Patents and Concentration of Patent Classes

Panel A: Distribution of Patents by Quintiles (Q)

Period	Q1	Q2	Q3	Q4	Q5
1975-2002	0.0002	0.0005	0.0233	0.1239	0.8521
1975-82	0.0001	0.0005	0.0393	0.1590	0.8011
1983-93	0.0002	0.0006	0.0246	0.1353	0.8393
1994-02	0.0002	0.0005	0.0187	0.0992	0.8814

Panel B: Concentration of Technologies, by Patent class

Period	Gini Index	Top 10%	Top 5%	Top 1%
1975-2002	0.8087	0.6857	0.5040	0.2011
1975-82	0.7740	0.6280	0.4546	0.1793
1983-93	0.7993	0.6674	0.4859	0.1908
1994-02	0.8318	0.7291	0.5428	0.2474

Source: Authors' Elaboration

Table 2: Relevant International Patent Classifications (WIPO)

A01	Agriculture; forestry; animal husbandry; hunting; trapping; fishing
A47	Furniture; domestic articles or appliances; coffee mills; spice mills, suction cleaners in general
A61	Medical or veterinary science; hygiene
A63	Life-saving; fire-fighting
B01	Physical or chemical processes or apparatus in general
B05	Spraying or atomising in general; applying liquids or other fluent materials to surfaces, in general
B23	Machine tools; metal-working not otherwise provided for
B29	Working of plastics; working of substances in a plastic state in general
B32	Layered products
B41	Printing; lining machines; typewriters; stamps
B60	Vehicles in general
B65	Conveying; packing; storing; handling thin or filamentary material
C07	Organic chemistry
C08	Organic macromolecular compounds; their preparation or chemical working-up; compositions based thereon
C12	Biochemistry; beer; spirits; wine; vinegar; microbiology; enzymology; mutation or genetic engineering
E04	Building
F01	Machines or engines in general; engine plants in general; steam engines
F02	Combustion engines; hot-gas or combustion-product engine plants
F16	Engineering elements or units; general measures for producing and maintaining effective functioning of machines or installations; thermal insulation in general
G01	Measuring; testing
G02	Optics
G03	Photography; cinematography; analogous techniques using waves other than optical waves; electrography; holography
G06	computing; calculating; counting
G09	Educating; cryptography; display; advertising; seals
G11	Information storage
H01	Basic electric elements
H02	Generation, conversion, or distribution of electric power
H03	Basic electronic circuitry
H04	Electric communication technique
H05	Electric techniques not otherwise provided for

Source: http://www.wipo.int/classifications/ipc/ipc8/?lang=en

Table 3: Generality Scores for Patents Issued by Year, 1976-2005, overall and in various sectors

	Sample 1: All patents					Sample 2: Information Technology patents					Between Samples	
Cohort	Generality (mean)	St.Dev.	Cited	All patents	t-test: btw cohort	Generality (mean)	St.Dev.	Cited	All patent	t-test: btw cohort	t-test: S1 - S2	t-test: S1 - S3
1976-80	0.405357	0.279369	293035	316783	0.0000	0.456188	0.269047	19409	20225	0.0000	0.0000	0.0000
1981-85	0.41226	0.277641	301180	322407	0.0000	0.48092	0.259162	23109	23852	0.0303	0.0000	0.0000
1986-90	0.41804	0.275195	395962	421473	0.0000	0.485446	0.25329	42279	43234	0.5425	0.0000	0.0000
1991-95	0.412795	0.277261	464325	499663	0.0000	0.484461	0.256121	60405	61951	0.0000	0.0000	0.0000
1996-00	0.372309	0.283336	604677	684756	0.0000	0.446299	0.270599	121222	127582	0.0000	0.0000	0.0000
2001-05	0.233065	0.275062	460523	817341	- -	0.283895	0.286438	122106	184769	- -	0.0000	0.0000

	Sample 3: Computer software patents					Sample 4: US 977 nanotech patents					Between Samples	
Cohort	Generality (mean)	St.Dev.	Cited	All patents	t-test: btw cohort	Generality (mean)	St.Dev.	Cited	All patent	t-test: btw cohort	t-test: S2 - S3	t-test: S1 - S4
1976-80	0.592044	0.228478	2082	2119	0.0327	0.656823	0.151735	5	5	0.7510	0.0000	0.0441
1981-85	0.605242	0.211009	3111	3152	0.0000	0.681349	0.144172	14	14	0.2668	0.0000	0.0003
1986-90	0.575746	0.225413	6426	6529	0.0000	0.613884	0.209561	45	49	0.1808	0.0000	0.0000
1991-95	0.556469	0.241636	9750	9976	0.0000	0.560693	0.248397	221	221	0.0034	0.0000	0.0000
1996-00	0.534064	0.253507	21889	22861	0.0000	0.499059	0.281373	749	803	0.0000	0.0000	0.0000
2001-05	0.348714	0.299487	24109	36332	- -	0.333545	0.315544	1068	1559	- -	0.0000	0.0000

	Sample 5: Nanotech keyword patents					Sample 6: EPO Y01N-linked nano patents					Between Samples	
Cohort	Generality (mean)	St.Dev.	Cited	All patents	t-test: btw cohort	Generality (mean)	St.Dev.	Cited	All patent	t-test: btw cohort	t-test: S1 - S5	t-test: S1 - S6
1976-80	0.6795	0.165136	4	4	0.7559	0.577711	0.243431	118	124	0.2156	0.0497	0.0000
1981-85	0.703633	0.130376	16	16	0.0309	0.540681	0.25855	183	187	0.0654	0.0000	0.0000
1986-90	0.57483	0.235737	324	334	0.2394	0.576272	0.231173	832	847	0.3352	0.0000	0.0000
1991-95	0.557434	0.24835	1971	2017	0.0000	0.566877	0.239334	2042	2088	0.0000	0.0000	0.0000
1996-00	0.500142	0.272762	3726	3965	0.0000	0.495359	0.267494	3536	3761	0.0000	0.0000	0.0000
2001-05	0.320756	0.307005	4111	6675	- -	0.39954	0.30062	2445	3095	- -	0.0000	0.0000

Table 4: Generality Scores for Patents, 1976-2005, Information Technology and Nanotechnology

	Sample 2: Information Technology patents					**Sample 7: All nanotechnology patents**					Between Samples	
Cohort	Generality (mean)	St.Dev.	Cited	All patents	t-test: btw cohort	Generality (mean)	St.Dev.	Cited	All patents	t-test: btw cohort	t-test: S1 - S7	t-test: S2 - S7
1976-80	0.456188	0.269047	19409	20225	0.0000	0.580927	0.24048	123	129	0.2917	0.0000	0.0000
1981-85	0.48092	0.259162	23109	23852	0.0303	0.550678	0.254489	197	201	0.2990	0.0000	0.0000
1986-90	0.485446	0.25329	42279	43234	0.5425	0.569862	0.235505	1090	1114	0.0459	0.0000	0.0000
1991-95	0.484461	0.256121	60405	61951	0.0000	0.552957	0.247004	3538	3627	0.0000	0.0000	0.0000
1996-00	0.446299	0.270599	121222	127582	0.0000	0.489086	0.270682	6513	6950	0.0000	0.0000	0.0000
2001-05	0.283895	0.286438	122106	184769	- -	0.337251	0.303696	5988	9058	- -	0.0000	0.0000

Table 5: Top WIPO Patent Classes, assigned to US Patents 1975-2002

WIPO Patent Class	Description	Frequency	Position on Lorenz Curve
H01	Basic electric elements	199,902	1.00
A61	Medical or veterinary science; hygiene	185,990	0.92
G01	Measuring; testing	141,258	0.85
G06	computing; calculating; counting	116,985	0.80
H04	Electric communication technique	112,402	0.75
C07	Organic chemistry	109,662	0.71
B65	Conveying; packing; storing; handling thin material	94,912	0.67

Source: Authors' Elaboration

Table 6: Field Distribution of the 20 most forward-cited patents, 1975-99

Electrical	6
Chemicals	6
Mechanical	3
Drugs and Medical Instruments	3
Computing	1
Others	1

Source: Authors Elaboration based on Hall and Trajtenberg (2004), table 2.

FIGURES

Figure 1: Growth in Issued Patents, 1975-2006

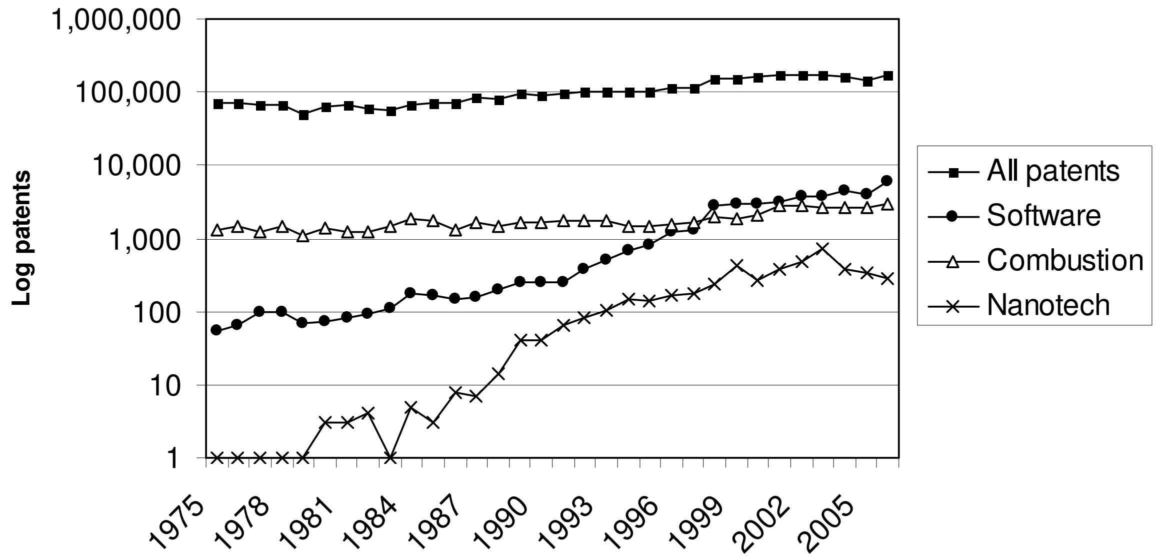

Fig. 2: Relationship among three Nanotechnology Patent Datasets

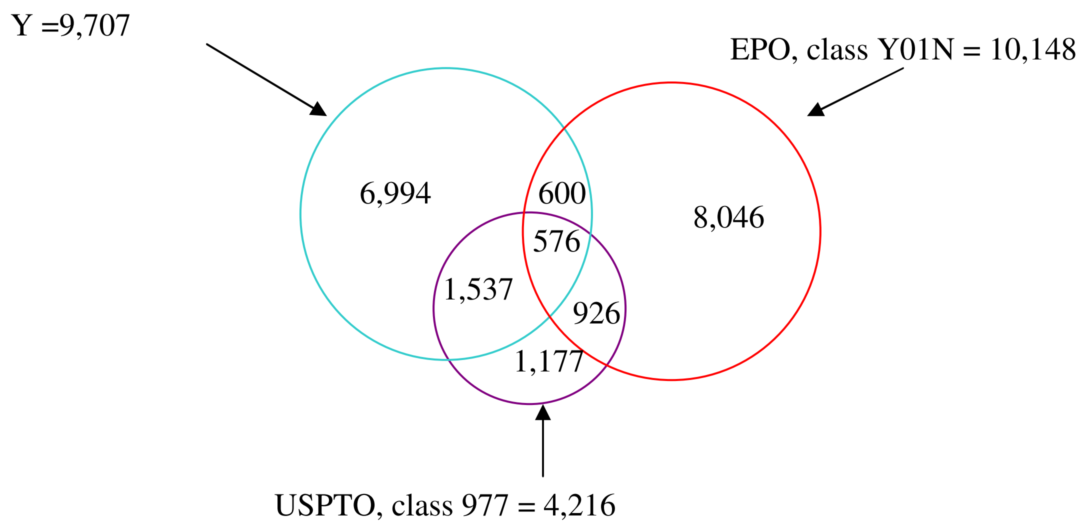

Note: Each circle represents the set of a different Nanotechnology Patent Dataset. The rightmost set represents patents in the family of those classed in the experimental class Y01N of the European Patent Office (EPO). The bottommost set is based on patents issued in class 977 of the US Patent Office (USPO), whereas the leftmost set is based on patents elaborated by Youtie (Y). The patents represented here were issued over the period1976-2005.

Figure 3: Lorenz Curve: 1975-2002
Number of Patents: 2.6 Million
Number of (IPO) classes: 226

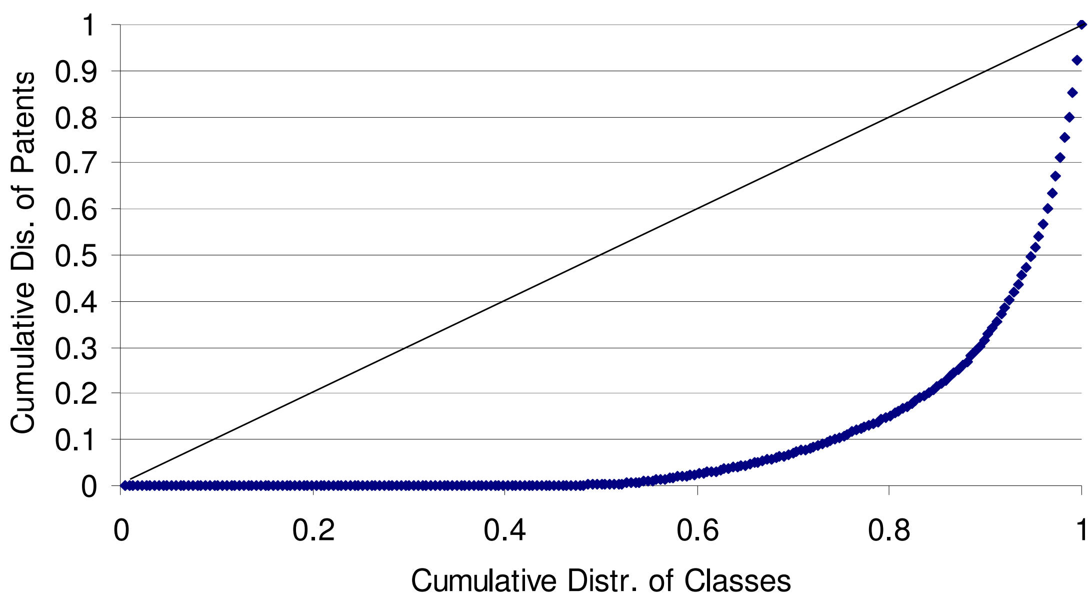

Figure 4: Dissemination of NANO-Knowledge by perc. at the top 5%

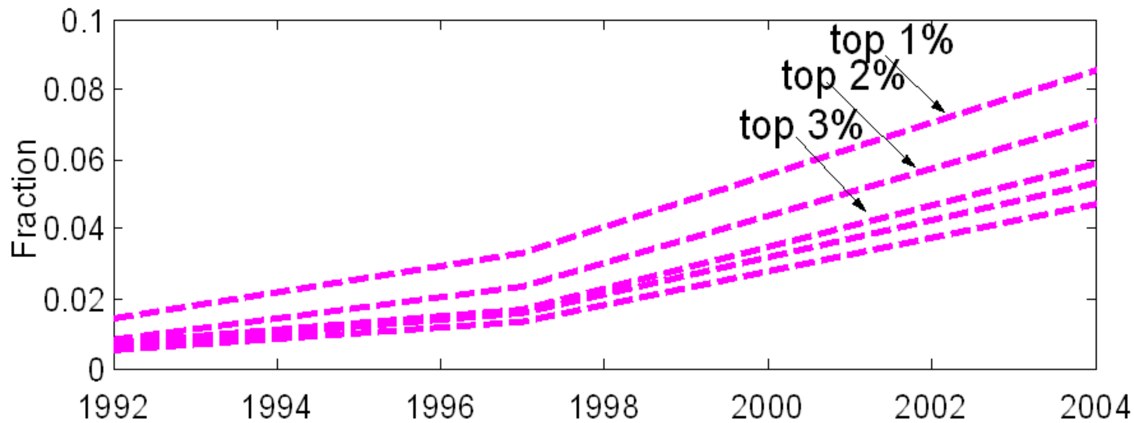

Figure 5: Dissemination of NANO-Knowledge by deciles

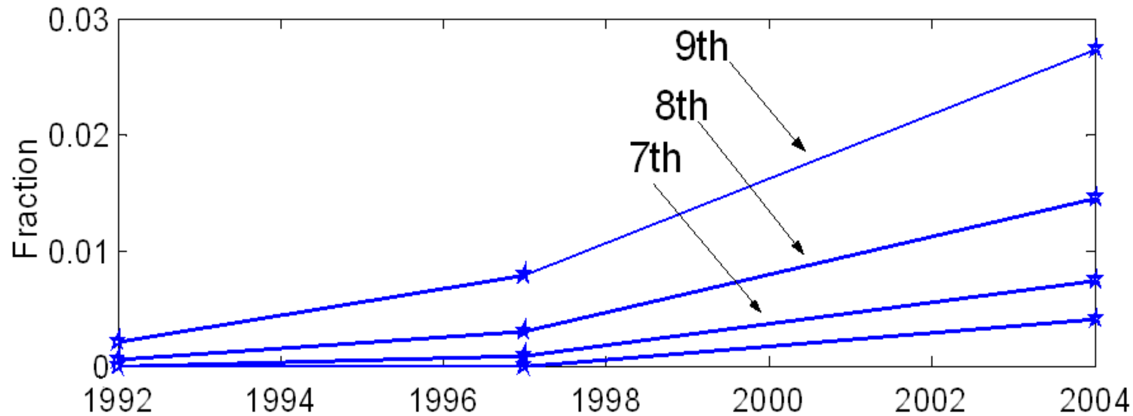

Source: Authors' elaboration based on US Patents data.

Note: The top dashed line of Figure 4 displays a ratio: the number of patents in a focal class that cite to at least one nanotechnology-classed patent divided by the total number of patents classified in the focal class, here ranked at the 99th percentile for the number of nano-patent citations. The second line from the top shows the same numerical information but for the technology patent class that occupies the 98th percentile, and so forth. Figure 5 is similar in construction to Figure 4, except that the position is indicated by deciles instead of percentiles.

Figure 6: Dissemination of IT-Knowledge by percentile at the top 5%

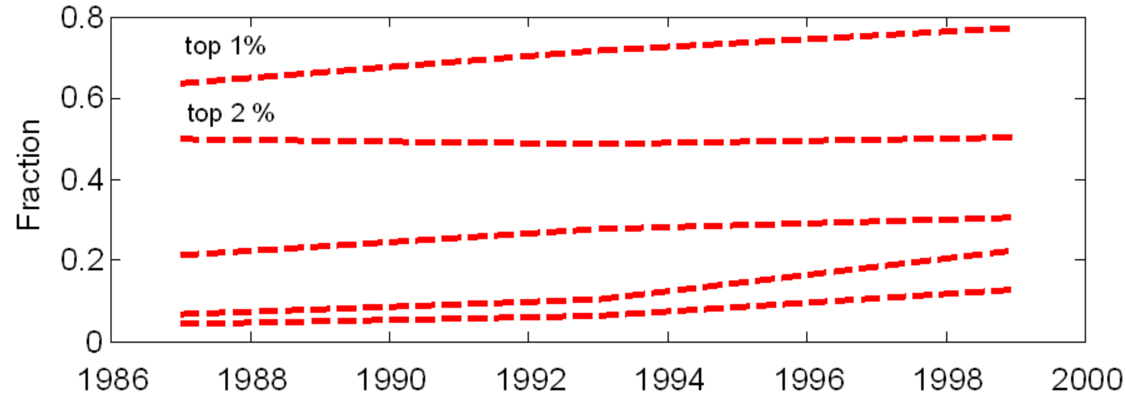

Figure 7: Dissemination of IT-Knowledge by deciles

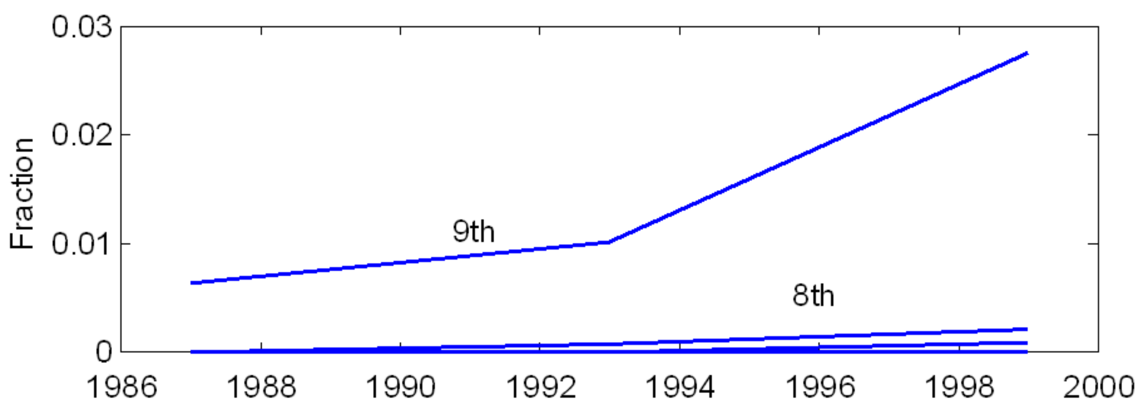

Source: Authors' elaboration based on US Patents data.
Note: Figures 6 and 7 are constructed identically to Figures 4 and 5, respectively, except that patent citations here refer to information technology classes instead of nanotechnology classification.

Figure 8: Dissemination of CE-Knowledge by perc. at the top 5%

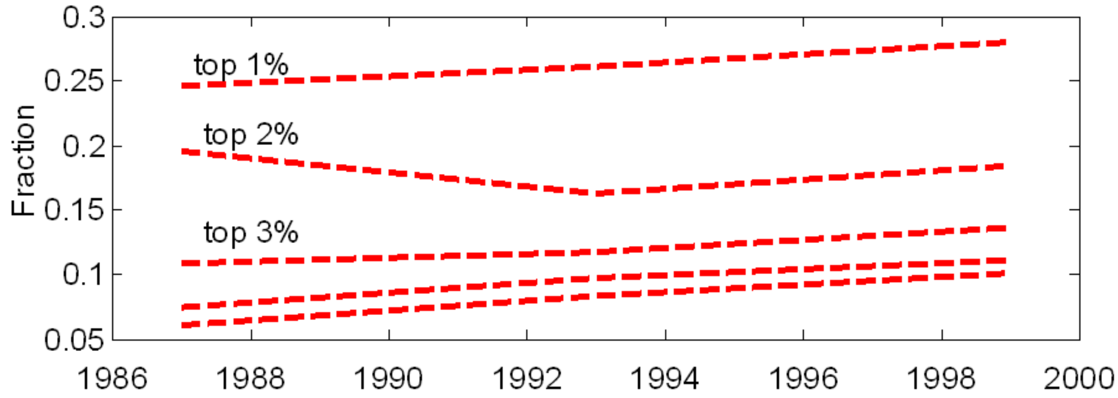

Figure 9: Dissemination of CE-Knowledge by deciles

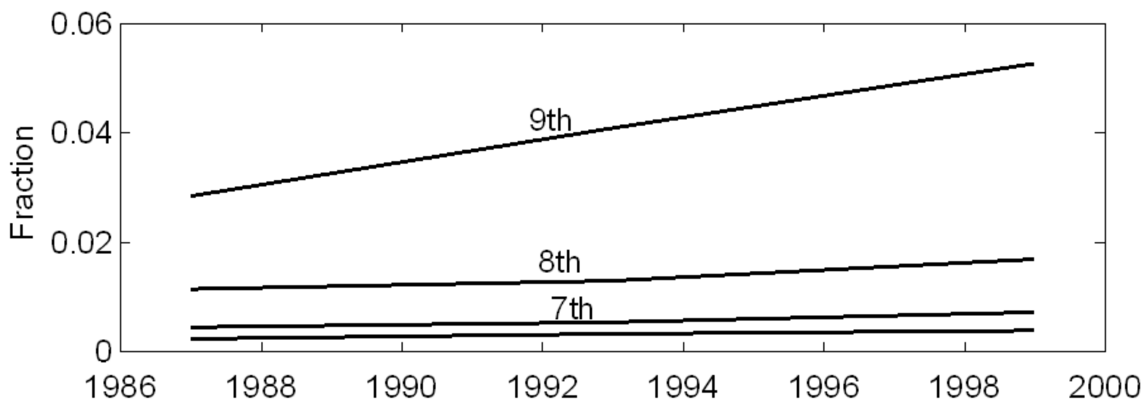

Source: Authors' elaboration based on US Patents data.
Note: Figures 8 and 9 are constructed identically to Figures 4 and 5, respectively, except that patent citations here refer to combustion engine technology classes instead of nanotechnology classification.

Figure 10: Patterns of Self-Citations: 1985-87 versus 1992-94

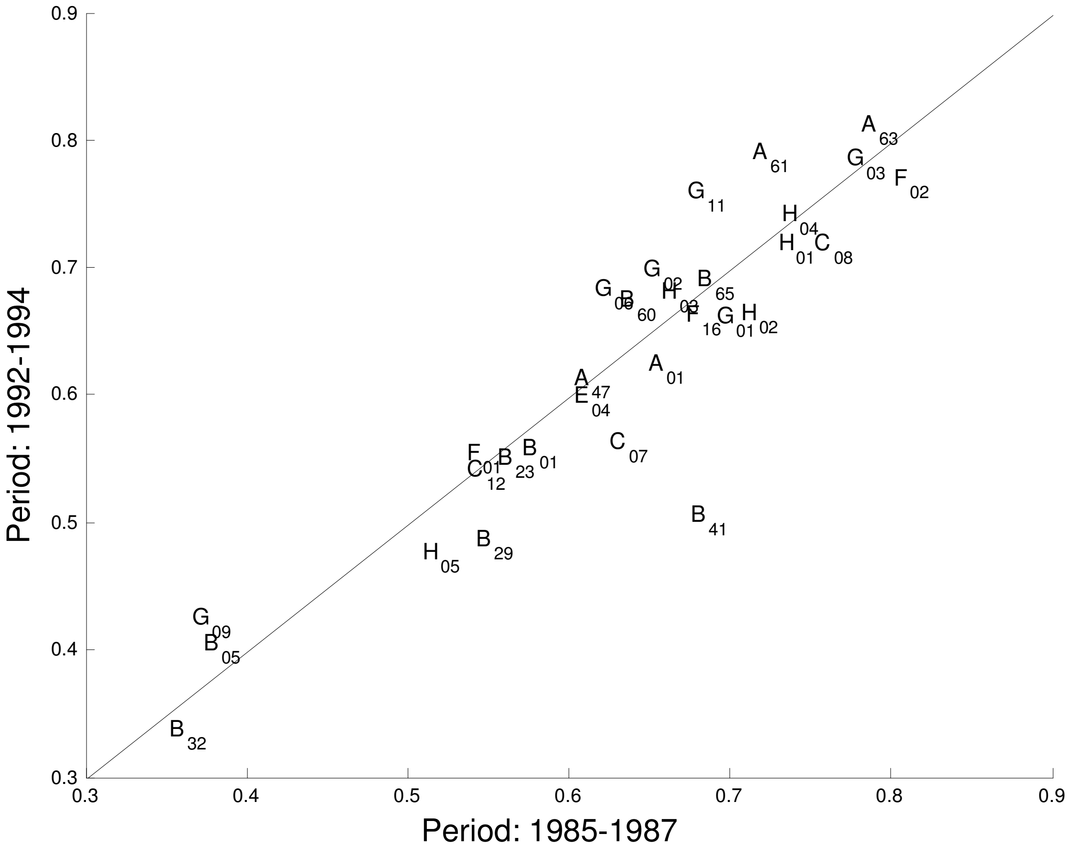

Figure 11: Patterns of Self-Citations: 1992-94 versus 1999-2001

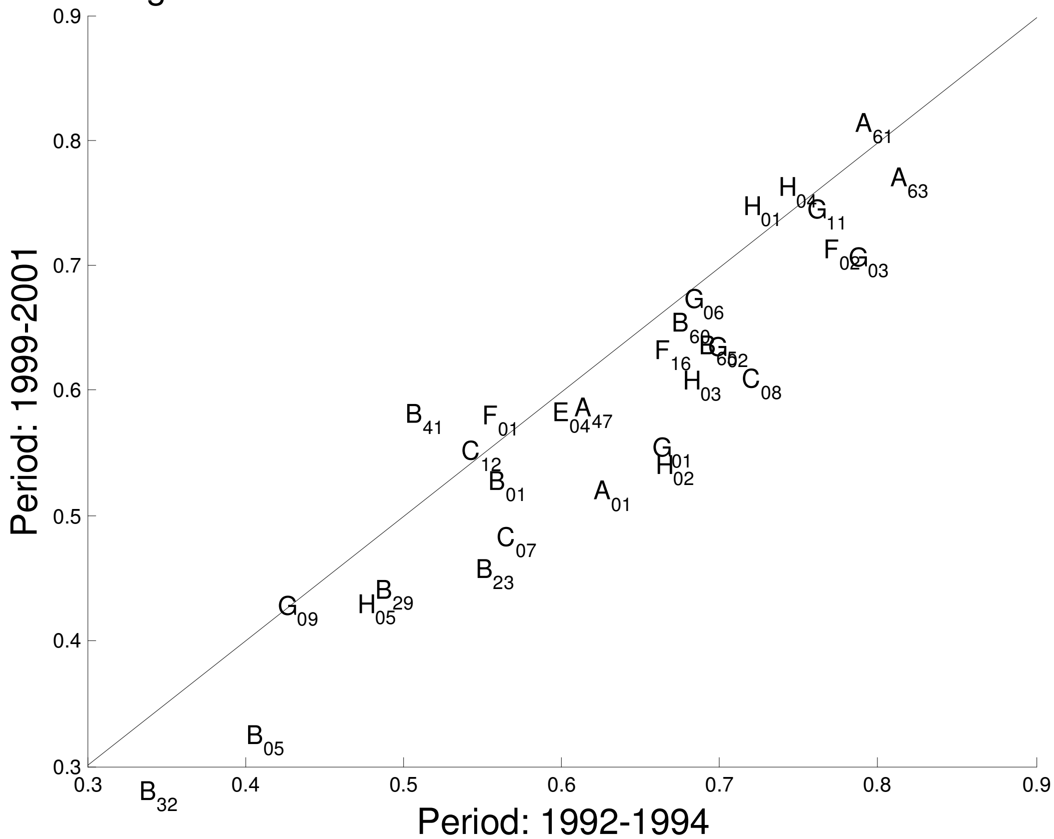

Source: Authors' elaboration based on USPTO dataset. A description of the codes can be found in Table 2.