Information Technology and Intangible Output:
The Impact of IT Investment on Innovation Productivity

by

Landon A. Kleis

B.Com., Queen's University, 1996

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE IN BUSINESS ADMINISTRATION

in

THE FACULTY OF GRADUATE STUDIES

Sauder School of Business
Management Information Systems Division

We accept this thesis as conforming to the required standard

The University of British Columbia

June 2004

Library Authorization

Landon Kleis
Name of Author *(please print)*

24/06/2004
Date (dd/mm/yyyy)

Signature

Title of Thesis: Information Technology and Intangible Output: The Impact of IT Investment on Innovation Productivity

Degree: MScB Year: 2004

Department of Sauder School of Business
The University of British Columbia
Vancouver, BC Canada

Abstract

Research investigating the contribution of information technology (IT) to firm productivity is moving beyond the question of whether IT has an impact, to how that impact is created. Traditional measures of firm output, such as value added, profit and market value, may be less useful in this context. However, the use of intangible outputs holds some promise in assessing how IT creates value within the firm. An important and information-intensive intangible output is innovation, which can be supported through the application of information technology. A patent production function is specified using R&D capital and IT capital as inputs, and citation-weighted patent output as an index of the overall inventive output of the firm. A panel of 262 large U.S. firms, from the years 1987 to 1996, is analyzed using OLS regression. Results indicate that, at the margin, the effect of IT capital on patent output is negative. This implies that increasing the general level of IT investment in a firm cannot be assumed to automatically improve the productivity of formal R&D. Recent IT productivity research suggests this may be due to the role of unmeasured complimentary investments in making IT effective.

Table of Contents

List of Tables

Acknowledgements

I owe a debt of gratitude to my supervisor, Paul Chwelos, for providing patience, inspiration and levity. I would also like to sincerely thank Ron Ramirez for his thoughtful suggestions. I thank both gentlemen along with Kenneth Kraemer for sharing their data. This work also benefited from comments received at WISE 2003.

I am also grateful to my family for their wisdom and support throughout my lifetime and in particular during this program. Finally, I would like to thank Marie Rekkas for stylistic guidance and econometric advice.

Landon Kleis

1. Introduction

Until the widespread adoption of personal computers in the late 1980's, information technology was generally seen by economists and managers simply as capital equipment. Computer systems were employed only by large companies to create transaction processing systems that shared several attributes of manufacturing assembly lines: standardized inputs and a narrow range of outputs; carefully defined operating procedures; and a large capital investment for hardware, data centres, and specialized training. Yet with the advent of the PC, the local area network (LAN), and improved software, IT has become a general-purpose technology. These breakthrough technologies, such as the steam engine or the electric motor, have the power to transform business models, operations and whole economies in radical ways, for three key reasons: they have a wide variety of uses; price declines facilitate the discovery of even more applications; and network externalities create increasing value as more of the technology is adopted (Bresnahan and Trajtenberg, 1995). The data-gathering and communications-enabling powers of IT infrastructure make possible new organizational capabilities and management structures at the micro level, while promising productivity improvements and even mitigation of the business cycle at the macro level (DeLong and Summers, 2001). The prominent and growing levels of IT investment attracted considerable attention from researchers and the popular press, who were surprised by initial calculations that showed IT had a disproportionately small impact on productivity. This "productivity paradox" spawned a significant body of research, which is summarized in several review articles (Brynjolfsson, 1993; Brynjolfsson and Hitt, 1996, 2003; Kohli and Devaraj, 2003).

Although there are a few exceptions, most of the recent research concurs that IT has a positive impact upon firm productivity. However, this may be viewed as only the first stage of research into the subject, because most studies call for further investigation into the specific ways in which IT is used to create value. New studies have asked the questions "when, where and how" IT creates value, rather than "whether" value is created at a higher level (Thatcher and Oliver, 2001). An interesting addition to this line of inquiry is the work of Mittal and Nault (2004), which showed that IT can have both direct and indirect effects. While the direct impact of IT may manifest in its substitution for other factors of production, the indirect impact occurs by making the other factors more efficient than they would be in the absence of IT.

The overwhelming proportion of IT productivity research at the firm level has investigated the impact of IT capital investment upon traditional "post-production" measures of firm outputs; value added, profit, and market value are the most common. IT capital is typically defined as a tangible asset: it has a physical embodiment that can be counted. Recently, researchers have considered the intangible aspects of IT, such as the role played by complementary organizational capital in the effectiveness of IT investments (Brynjolfsson et al., 2002). However, the question of whether IT capital has an impact on intangible output—defined here as the production or accumulation of non-physical, non-financial assets—has been neglected. Both innovation and new product development are key examples of intangible outputs with a close relationship to IT. Innovation has long been the subject of economic research, much of which uses patents as an index of the overall inventive output of the firm. Given the importance of innovation to firms and economies alike, this is a crucial area of research.

Innovative activity can be supported through the application of information technologies. Knowledge management systems, intranets, and group support systems, for example, can be used explicitly in the new product development process (Nambisan, 2003). At a lower level of analysis, IT supports innovation by reducing the costs of information collection, management, and dissemination. This stimulates the knowledge creation process by enabling more efficient collaboration within firms, within research communities, and between business partners (Malone and Crowston, 1994; Gurbaxani and Whang, 1991). Thus, it is reasonable to expect a relationship exists between information technology and the knowledge creation process (Lee and Choi, 2003).

This research will investigate the contribution of IT to the productivity of formal R&D-based innovation in large U.S. firms, using citation-weighted patents as an output indicator. Although this places restrictions on the types and sources of innovation being investigated, it represents a substantial portion[1] of quantifiable innovation in the U.S. between 1987 and 1996. A modified form of the Knowledge Production Function, incorporating IT capital, is presented, and operationalized using the Cobb-Douglas and translog functional forms. Data obtained from the Computer Intelligence database, Standard and Poor's Compustat, and the NBER patent citations data file (Hall et al., 2001) are analyzed using ordinary least squares regression. The results indicate that, for a given level of R&D capital, increasing IT capital has a negative effect on the expected output of citation-weighted patents. Although somewhat surprising amid the recent IT productivity research, this finding contributes to our understanding of how IT impacts

[1] Hall et al. (2001) note that between 72% and 77% of U.S. patents are granted to corporations during the sample period of their data (1963-1999). Of these, between 40% and 50% were matched to publicly-traded U.S. corporations for the period 1987-1999.

business value. There are a number of possible explanations and opportunities for future research.

2. Literature Review

IT Productivity

IT productivity research often considers IT capital as an economic input to production. This approach is especially prevalent in macroeconomic studies, which view IT as a form of capital that fits into a conventional production function. It is also true of many firm-level studies that investigate the nature of this new kind of capital investment in the context of production. As such, IT capital is almost universally treated as an independent variable, while the dependent variable may take the form of tangible outputs such as value added or sales growth.

Researchers typically define IT capital to include numerous forms of computer hardware, such as PCs, mainframes, and networks for data communication. Hardware is a preferred measure since in most situations it is included in a firm's assets on the balance sheet. Software, on the other hand, tends to be expensed and in many firms may even be developed internally, making it difficult to produce estimates of its worth.

Rapid advances in hardware technology, in terms of speed, capability and availability, coupled with steady or declining nominal prices, are a well-known phenomenon (DeLong and Summers, 2001). This trend has fulfilled a decades-old prediction that has come to be known as Moore's law, which posits that the circuit density on a semiconductor memory wafer can be doubled every 18 months (Moore, 1965). The popular interpretation of this "law" is that the power of microprocessors will follow this trend. Empirical research has found this to be generally true, and in some cases somewhat conservative. Yet there is a disadvantage to the temptation to measure the economic potential of IT through purely technical metrics such as clock speed or chip density. Although this approach has intuitive and even historical appeal, the complexities of implementing IT frustrate simplistic analyses. Doubling the speed with which a CPU processes instructions does not necessarily translate into a doubling of the overall speed of that computer system, nor of the productivity of work performed on that computer. The efficiency of the CPU is bound by the other subsystems with which it integrates (memory, storage, and networking), and the ability of the user (via software) to continue feeding it instructions (Chwelos, 2003).

The same may be said for considering conventional economic input measures exclusively, such as IT capital. As Brynjolfsson et al. (2002) so aptly put it, "A computer that is integrated with complementary organizational assets should be significantly more valuable to a business than a computer in a box on the loading dock." The authors argue that the effective use of IT requires complementary organizational assets such as decentralized decision making, IT and management skills and procedures to capture and use data in a more profitable way. As such, the true cost of a successful IT implementation may cost as much as ten times the initial technology investment. Therefore, they argue, a firm's stock market valuation is superior to output as a measure of the productive value of the combination of tangible IT capital and the associated intangible assets that determine its effectiveness.

Three key findings have emerged from the IT productivity literature. First, IT capital investments on the whole have generally positive returns (Brynjolfsson and Yang, 1996). In several cases where negative returns were found, other researchers have subsequently argued that errors in measurement and price index assumptions have the potential to cause significant distortion in the results (Dewan and Min, 1997; Barua and Lee, 1997). Repeated studies of the same data but with different deflators for IT prices have shown positive results for IT productivity. Second, IT productivity results are more pronounced over the long term, rather than in short-term time frames (Brynjolfsson and Hitt, 2003). This result is thought to be due to the lags required to implement the necessary complementary investments and business operations restructuring needed to take advantage of IT capital investments. Third, individual firms can vary significantly in their ability to create value from IT capital investments (Brynjolfsson et al., 2002).

A common characteristic of the extant research in this area is the use of traditional measures of firm outputs to gauge productivity. The most common are profit, sales, ROI, value added and market value (Kohli and Devaraj, 2003). This works well in the production function methodology, where optimum levels of inputs can be determined in order to maximize an output objective.. However, as some researchers have pointed out, the impact of IT capital investment upon production may not be immediate (Brynjolfsson, 1993; Thatcher and Oliver, 2001). Firms may invest in IT to facilitate organizational changes, over the medium term, to gain strategic advantage over competitors. This in turn may lead to supranormal profits over the long term (Porter, 2001). IT can facilitate a differentiation strategy, such as mass customization or service quality improvements. Alternatively, a firm may use IT to vertically restructure, using IT to reduce communication and co-ordination costs among market-based "electronic hierarchies" (Malone et al., 1987). Furthermore, even when the objective is to maximize profits, some firms may take a longer-term approach to this goal. There is evidence that productivity gains arising from IT investments may require years of complimentary investments and organization learning to effect significant changes in traditional measures of output (Brynjolfsson and Hitt, 2003). Intangible outputs, such as new product development and patents, may have a more direct relationship with the firm's chosen levels of inputs. Since IT plays a large role in these pursuits, this new area of productivity research is potentially fertile.

Innovation, R&D, and Patents

The contribution of innovation to long-term economic growth has been the subject of economic research for almost 50 years (Abramowitz, 1956; Solow, 1956). A multitude of studies have demonstrated this link at both the economy and firm levels (see Jones, 1995 for a useful review). Yet the investigation into how innovation is accomplished, and how to measure it, remains far from complete. Both of these issues have been complicated by the widespread use of IT. The process of innovation has been changed by IT as much as accounting or logistics, while the outputs of innovation are in many cases less quantifiable than in the past.

It is important to distinguish between the concepts of innovation, R&D, and patenting. Innovation, the process of developing new products and services, takes place in both formal and informal contexts, and often goes unrecorded by official agencies. The research and development laboratory is but one source of innovation. Dosi (1998) adds three others: the ability of the firm to "learn by doing"; informal knowledge gathering, such as publications, technical associations and personnel transfers; and embodied innovation adopted through the use of new types of capital inputs. Patents are commonly associated with innovation and R&D processes, but are only one outcome of innovation, and are not guaranteed at that. As Griliches (1990) summarized, "not all inventions are patentable, not all inventions are patented, and the inventions that are patented differ greatly in … the magnitude of inventive output associated with them."

Information technology has contributed to all types of innovation. Advances in computer modelling and simulation have improved formal research productivity in fields such as pharmaceuticals, automotives and chemicals. The Internet has hastened the flow of, and access to, scientific information such as electronic versions of scholarly journals and public research databases, both of which were cited by R&D managers as being used in their R&D processes in a recent survey of French firms (Kremp and Mairesse, 2004). "Learning by doing" can be supported by computerized engineering analysis of production processes, leading to adjustments and improvements with significant results (Hatch and Mowery, 1998). Small-scale innovation and experimentation occurs frequently in the service sector due to the existence of customizable tools for knowledge workers, improving business processes through the use of macros, programming languages, and interoperable systems and data formats.

It is difficult to estimate even the relative share of informal and formal modes of innovation within firms or industries. Denison (1985) estimated that R&D accounted for only 20 per cent of all technical progress, although other researchers (Rosenberg, 1985; Nelson, 1986) concluded that R&D was the dominant organizational form for technological search. So-called "learning by doing" or "learning by using" innovation has been researched extensively (Rosenberg, 1982) and can be more difficult to measure empirically. In general, the phenomenon refers to the role of experience in improving the productivity of the manufacturing process over time (in contrast to pre-production innovations produced inside the R&D laboratory). Often no budget exists for this sort of innovation, so it can appear as a mysterious source of productivity improvement, such as in the case of the "Horndal effect," where a Swedish steel factory experienced an annual 2% increase in output without any capital investment for 15 years (Arrow, 1962). Learning by doing is also cited as a source of first-mover advantage in research-intensive industries such as computers and semiconductors (Hatch and Mowery, 1998; Grabowski, 2002).

There is also a great deal of variation among industries in the propensity to do research and the modes in which it is pursued (Bound et al., 1984). This diversity is explained by the extent of opportunities for innovation under the technological paradigms for each field at any given time (Dosi, 1998). Scherer (1986) found 42.5 percent of the variance in patenting rates could be attributed to differences across industries. R&D productivity research typically accounts for this by using industry dummy variables.

Researchers who study innovation have spent considerable effort finding a suitable way to measure both inputs and outputs of innovation. Formal innovation is more commonly studied, since a specific R&D program has quantifiable inputs (labour and capital) being utilized to produce, it is hoped, quantifiable innovations (patents or products). However, there are numerous difficulties with these measures. For inputs, it can be difficult to obtain data from companies wishing to protect the size and nature of their R&D program. Even if the input shares and prices are known, it is difficult to assess the contribution of one year's research toward the innovation outputs from that year. Knowledge, the main conceptual input to the innovation process, accumulates (and depreciates) over time, just as some innovations may take many years to produce (Griliches, 1979). Output measurement also presents a host of challenges, some of which are discussed below.

A major contribution to the understanding of innovation in an economic context is the Knowledge Production Function (KPF), credited to Pakes and Griliches (1984). The authors proposed a model based on a Cobb-Douglas production function,

$$\dot{k}_{i,t} = a_i + \gamma t + \sum_{\tau=0}^{5} \theta_\tau RD_{i,t-\tau} + u_{i,t} \ ,$$

where $\dot{k}_{i,t}$, the output of knowledge generation activity for firm i at time t, is determined by the firm's current and past five years of research investment $RD_{i,t-\tau}$. Both $\dot{k}_{i,t}$ and $RD_{i,t-\tau}$ are expressed in log terms. The level of investment is assumed to be determined exogenously. Firm-specific differences in research productivity are captured by the a_i term, while $u_{i,t}$ represents the stochastic component.

Since a firm's knowledge stock is inherently unobservable, Pakes and Griliches operationalized the model using patents as a proxy for $\dot{k}_{i,t}$ (sometimes called the Patent Production Function or PPF). A patent represents the formal disclosure of an idea or process that has passed some standard of novelty and possesses an expected economic value to its owner (Griliches, 1990). Similarly, research and development expense represents the formal commitment of a firm's resources (such as scientists and engineers, labs and prototypes) to knowledge-generating activities. Using patents as an output indicator, and lagged R&D expenditures as the investment input, Pakes and Griliches found a statistically significant relationship between R&D expenditures and patent output at both the firm and industry levels. For cross-sectional estimates, the median R-squared is about 0.9, while for time-series estimates it is not as strong (about 0.3), but nevertheless positive and significant.

Other researchers have commonly employed patents as a proxy for additions to the knowledge stock or knowledge output (see Griliches, 1990, for a thorough review of the literature). Bound et al. (1984) investigate the nature of the relationship between patents and R&D, and are especially interested in the question of returns to scale. Depending on the estimation method being used, returns can be increasing (ordinary least squares, negative binomial) or decreasing (Poisson, nonlinear least squares) with scale. Hall et al. (1986) investigate the question of whether lagged R&D influences patent outcomes. They conclude that lag effects are difficult to estimate because R&D expenditure patterns are highly autocorrelated within the firm. As such, there is little difference between the sum of estimated coefficients of a series of lagged R&D and the coefficient where only contemporaneous R&D is estimated (Hall and Ziedonis, 2001).

An interesting addition to this research is the work of Kortum and Lerner (1998), who use the PPF framework to investigate the efficacy of venture capital in procuring patents. By adding a new kind of capital input to the model, it is possible to consider the relative effectiveness of R&D and venture capital. They found that the impact of venture capital on patent output is 6 times greater than R&D, and that this type of capital, in general is responsible for a significant portion of innovation in the U.S.

The use of patents as a measure of innovation presents both positive and negative aspects to the researcher, which must be considered in the context of one's investigation to determine suitability. For many researchers, the appeal of patents as an output measure stems from their availability as the quantitative representation of an idea that has been sufficiently developed to pass the scrutiny of a patent agency, and which is evidence of some expectation of positive utility for the patent-holder. Patent data for the U.S. are readily available in computerized format as far back as 1963, and are richly endowed with information about the innovation itself, the patent owner ("assignee"), citations of relevant prior patents, and the field into which the patent has been classified. The data are also plentiful: the U.S. Patents and Trademark Office (USPTO) granted over 3 million patents between 1963 and 1999, and the number of patent applications per year has doubled between 1992 and 2003 (USPTO, 2003). Patents are assigned a field among over 120,000 technological classifications, as determined by the USPTO.

Unfortunately there are numerous caveats to the use and interpretation of patents as a measure of innovation output. The first issue is, with respect to R&D expenditures, whether or not a patent represents an input or an output. There is some debate as to the direction of causality between these two concepts. In order to be patentable, an innovation requires both basic research and some amount of development to prove that it can be a commercially-viable product. Yet once a patent is granted, further development is often required to determine whether the product would be economically viable for its owner. In this context, obtaining a patent usually requires future R&D spending by the assignee.

When comparing patent output across firms and years, Griliches (1990) identified two key issues: quality and classification. Clearly, the value of patents is subject to tremendous variation, even in a controlled group such as within a single year's patents for a given company. This variance makes it difficult to compare the sum of patent output across firms, industries or years. Classification becomes a problem when attempting to determine characteristics of particular fields or groups. One difficulty arises because of the tendency for the USPTO to reorganize its classification system periodically, and the constant addition of new categories to deal with new types of innovations. It is important to be able to group like categories of patents together for analysis because the divergence of patenting rates across fields, and their respective trends over time, must be recognized when using patent statistics.

Finally, there are two additional issues to consider when interpreting patent outputs: patents do not represent all of a firm's innovative output, and among the innovations that are patentable, there are a number of mitigating factors in the decision to apply for a patent. Factors affecting patent-seeking include: economic conditions; changes to rules about what can be patented; and the patent-infringement litigation climate (Griliches, 1990; Hall and Ziedonis, 2001). Competitive factors also play a role, such as the relative advantage of patenting versus other forms of intellectual property protection for the firm's products and markets. For example, if time-to-market constraints are present, it may be more prudent to develop and market the new product quickly, rather than spend several months preparing a patent application.

Some researchers have attempted to address the difficulties with raw patent counts as the dependent variable. The causality problem may be addressed by testing for endogeneity in the model. The variability in the quality or value of patents has been addressed to a large extent by using citations received by other patents as a proxy for economic value (Hall et al., 2001). Existing patents are cited by both the patent applicant and the patent examiner, so determining each patent's accumulated (forward) citations from subsequent patents yields a reasonably objective measure of its relative importance.

The authors also propose a classification system and a method for removing unwanted year and year-field interaction effects, while leaving field effects intact. This is elaborated in §4. The extent to which patents capture all innovation is addressed, first, by noting that patents are used as an indicator of the level of innovative activity among firms, rather than measuring the entire innovative output. Second, changes in patenting trends due to economic, legal and competitive conditions are captured by using year indicator variables.

Integrating the Two Areas of Research

Although IT productivity and innovation research have several themes in common, researchers have brought them together on only a few occasions. Greenan et al. (2001) examined the relationship between specialized skills utilized in IT and R&D and the productivity outputs of firms. They found positive productivity returns to skill-related investments in these areas, although only in the cross-sectional dimension. Kremp and Mairesse (2004) found that knowledge management policies (associated with the use of IT and the Internet) had a significant impact on firms' innovation and patenting performance. Given the knowledge-intensive nature of innovation, it is logical to further investigate the relationship between IT and innovation.

3. Model

A production function is simply an analytical framework for characterizing the nature of one or more inputs in relation to a measure of output. One very common framework is the Cobb-Douglas production function, which typically sets the dependent variable as some measure of output (y), such as sales or profit, while the independent variables consist of two inputs, such as labour (L) and physical capital (K):

$$y = A(L^{\alpha} K^{\beta}).$$

The output elasticity of each input is simply the exponent associated with the input. Additional input variables, such as materials, may be added. The Cobb-Douglas form has interesting properties: it is a constant elasticity model, and the elasticity of substitution between capital and labour is equal to 1. An attractive property of this production function is that it is linear in (natural) logs:

$$\ln y = \ln A + \alpha \ln L + \beta \ln K .$$

To test the impact of IT upon the productivity of innovation requires an empirical model that addresses the nature of innovation itself. The development of such a model is fraught with difficulty, due to the variety of modes of innovation (formal; informal; embedded), and its manifestations (intermediate outputs such as process innovations; traditional post-production outputs; increased consumer surplus).

This research will focus on formal innovation activities as embodied in the R&D spending reported in firms' financial statements. Accounting rules allow firms to expense research, defined as the "planned search or critical investigation aimed at discovery of new knowledge" which is specifically directed at a new or improved output (Oliver, 2003, pg. 46). Development is defined as transforming "research findings or other knowledge into a plan or design," which can include prototyping and building and operating pilot plants. Declared R&D expense also accounts for in-process research assets and intangibles purchased from other companies.

Formal innovation depends upon the generation of new knowledge, in both the research and development phases. As such, we require some way of relating the inputs of this process (of which R&D is a formal component) to its outputs (which could be represented by any number of indicators).

Knowledge Production Function

Pakes and Griliches (1984) characterize innovation as the production of new knowledge. The statistical model they propose uses patents as an index of the output of "inventive activity," and R&D expenditures as an input. This is operationalized as a Cobb-Douglas production function, which lends itself to econometric analysis. Simply put, innovation is a form of production, with an input (R&D) and an output (patents). Other researchers (Pardey, 1989; Trajtenberg, 1990; Hall & Mairesse, 1995) have used this model to explore the differences in research productivity between firms in different industries and countries.

Existing research in innovation economics does not specify a role for firm IT spending in the KPF. However, as introduced earlier, it is reasonable to expect information technology to have an influence on firm innovation efforts. The innovation process relies heavily on knowledge generation, acquisition, organization, and application (Dosi, 1998). Information technologies can assist these activities, both directly and indirectly. Direct assistance occurs through IT capital investments targeted at firms' R&D programs, for example, high-powered workstations, servers, and connectivity to other labs and resources. Information technology infrastructure investments assist research indirectly by providing the firm as a whole the potential to collect and analyze information that impacts the innovation process, for example by analyzing product returns and warranty service records to determine ways to enhance the product. Knowledge Management Systems can contribute to the organization's ability to absorb and utilize new (external) information, and transform internal knowledge into new products and processes (Adams and Lamont, 2003). Klein, Gee and Jones (1998) note that effective IT support for firm-wide knowledge gathering is necessary for the proper application of certain modes of R&D managerial decision-making.

To incorporate IT capital into the KPF, the method of Kortum and Lerner (1998) is adopted, whereby a new form of capital augments the Pakes and Griliches (1984) model. In this case, the introduction of the IT capital variable, *IT*, will allow the impact of IT capital on the effectiveness of R&D in the patent-generating process to be estimated:

$$\begin{aligned} PATCIT_{i,t} = \alpha + \beta_1 RDSTOCK_{i,t} + \beta_2 IT_{i,t} + \beta_3 EMPS_{i,t} + \beta_4 EMPS_{i,t}^2 \\ + \gamma_t + \eta_i + \varepsilon_{i,t}\ , \end{aligned} \tag{1}$$

where $PATCIT_{i,t}$ represents the logarithm of the number of patent citations received for firm i at time t, $RDSTOCK_{i,t}$ is the logarithm of R&D expressed as a stock variable (i.e., R&D capital rather than yearly expenditure), and employees and employees squared ($EMPS_{i,t}$ and $EMPS_{i,t}^2$) control for firm size (discussed in the next section). The γ_t term captures unobservable time effects (operationalized as year dummy variables), η_i captures firm fixed effects, and $\varepsilon_{i,t}$ is the stochastic component for firm i at time t.

The use of stock measures for both R&D and IT eliminates the need for lagged terms as this information is inherently included in the model, and is consistent with much of the literature investigating the productivity of R&D with respect to firm output (Hall and Mairesse, 1995).

In addition to the somewhat restrictive Cobb-Douglas model (2), a more general model is specified in the translog form:

$$
\begin{aligned}
PATCIT_{i,t} = {} & \alpha + \beta_1 RDSTOCK_{i,t} + \beta_2 RDSTOCK_{i,t}^2 + \beta_3 IT_{i,t} + \beta_4 IT_{i,t}^2 \\
& + \beta_5 RDSTOCK_{i,t} IT_{i,t} + \beta_6 EMPS_{i,t} + \beta_7 EMPS_{i,t}^2 \\
& + \gamma_t + \eta_i + \varepsilon_{i,t} \quad .
\end{aligned} \tag{2}
$$

As the Cobb-Douglas model in (1) is a special case of the translog specification in (2), testing the validity of the former is straightforward.

Control Variables

Prior research has established a number of controls that should be employed when using patent (or patent citation) data. First, one must control for year effects. The number of patents applied for in any year can be influenced by notable patent litigation events or changes to patent regulations. Further, the number of patents granted or cited in a given year is influenced by the productivity of the USPTO, which fluctuates according to available manpower and the investigation requirements of each cohort of patents.

Second, there is the issue of firm size, or scale effects in R&D. Two arguments come to bear on this matter. On one hand, large firms possess a scale advantage because their internal legal departments allow for more efficient preparation of patent applications (Scherer, 1965; Lerner, 1995). On the other, Griliches (1990) found smaller firms (having less than 1,000 employees or spending less than two million dollars per year on R&D) in his sample of publicly-traded companies obtained more patents in proportion to their size than larger ones. He offered three possible explanations: stock markets are highly selective in allowing entry to small public companies, and as a result such firms may possess a number of patents upon entry; perhaps smaller firms needed a patent to enter the market; and the likely practice of doing more informal R&D than larger companies, thus increasing the patents-to-R&D expenditure ratio. To accommodate these competing ideas, firm size in this research is controlled by using two variables to measure the number of employees in the firm: employees, and the square of employee. This allows for nonlinear or quadratic effects.

The third factor to control for is industry. Patent statistics show clear differences across industries in their propensity to patent (see Griliches, 1990). This variation arises from the different stages in technological evolution, general industry maturity, level of competition and the degree to which patents are seen as effective methods for securing intellectual property rights. The pharmaceutical industry, for example, places higher importance on patenting than other methods such as time to market and trade secrets (Grabowski, 2002). However, since equations (1) and (2) already contain terms that control for firm effects, industry controls cannot be added. Instead, the data is split into industry sub-samples (see Table 3), and the results are reported separately.

4. Data

This empirical investigation employs three sources of data. First, the "Computer Intelligence" (CI) database provides data on the IT capital stock in large American firms from 1987 to 1999[2]. This data becomes the basis for the observations in this research, since the IT capital figures can be linked to the other two data sets using standard Compustat identifiers. Second, individual firm data (R&D, employees, assets) are obtained from Standard and Poor's Compustat Industrial Annual. Third, patent citation data are obtained from an extensive and detailed data set of U.S. patents granted from

[2] Originally collected by Computer Intelligence, this database became known as ZD Market Intelligence, and ultimately, Harte Hanks CI Technology Database.

1975 to 1999, which accompanies the paper "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools" (Hall et al., 2001).

IT Capital

The CI database consists of twelve years of annual IT data for large U.S. companies from 1987 to 1999. Overall, 1,818 companies are included in the database but only about 450 of these appear in all years. Through a combination of surveys and physical site visits, an appraisal of the quantity of systems is recorded for approximately 850 companies per year. The data include the number of mainframes, minicomputers, PCs, network nodes, CPU power and hard disk storage capacity. Upon aggregating these quantities, the purchase value for the firm's total IT capital stock, CPURCH, is computed.

The database requires three modifications to be useful in this context. The first is to address the change in methodology for the calculation of CPURCH, beginning in 1995. Until 1994, CPURCH included the value of all IT systems in the firm; after that, the variable represents only the processor value (i.e., networks, storage, and PCs are excluded). Chwelos et al. (2004) propose a modification which estimates a constant version of CPURCH by decomposing the original measure into the values assigned to mainframes, minicomputers and PCs over the sample years. Since there was no change in the methodology for computing these prices, the resulting estimate of CPURCH_E is not subject to the shift in 1995. The authors demonstrate the validity of this method by noting the comparable results produced by alternative approaches.

Second, in order to obtain accurate estimates of the productivity of IT capital, the dollar amounts must be expressed in real, rather than nominal, terms. Chwelos et al. propose a two-pronged approach: one price index is calculated for PCs, while another is calculated for all other IT hardware. This addresses the issue of the more rapid decline in the quality-adjusted price of personal computers, calculated by Berndt and Rappaport (2001) to average -28.1% in the years of the CI database. Other IT hardware is deflated by the BEA price index for computers and peripherals, which averaged -14.7% per year over the period.

Finally, some attention must be paid to the categorization of each firm into industry groups for the purposes of analyzing innovation productivity by industry. Industry classification is obtained from the Compustat database, using the Standard Industrial Classification (SIC) code. In order to keep the number of industries manageable, 7 industry classifications were produced using the two-digit level in most cases (see Table 3). The industries are further classified as either "manufacturing" or "services" for the purpose of broader comparisons. The vast majority of patents are granted to manufacturing firms, which are segregated in my results. However, since some companies' classifications changed over the time period, adjustments were made to place them in a single industry so as not to confound the impact of a particular firm on multiple industries.

R&D and Firm Financial Data

Publicly-traded firms in the U.S. are required to disclose research and development expense in their annual income statements (Lev and Sougiannis, 1996). The reported R&D expense is governed almost entirely by SFAS 2 and 86 (Oliver, 2003). This figure includes the cost of labour and capital used in discovering new product or process innovations and developing these into viable commercial products (excluding government-sponsored research). Thus, there is an issue with double-counting of IT capital employed in R&D, since that capital will be counted twice: once in the CI data, and again in the R&D stock data. The implications of this issue are discussed in the Conclusion (§6). Standard and Poor's Compustat database provides historical company financial statement information for all the years in this study.

Using financial statements as a source for R&D expenditure data presents a number of challenges. Principally, one must consider that the number found in a firm's annual report is offered in the context of financial reporting regulations. This means that recognizable R&D expense, at best, includes no less and no more than that which meets the definitions used by accounting professionals. Many firms have no formal R&D program, yet some of these have one or more patents. Like most accounting regulations, SFAS 2 stipulates that costs be expensed or capitalized and depreciated if they are substantially devoted to the R&D function. There are many instances in which one could imagine this test to be subjective, the most obvious of which are large systems such as mainframes or networks which serve numerous company functions. As a result, there may be some potential for managerial manipulation in recognizing expenses as being R&D-related. Another challenge is the frequently spotty nature of the R&D expense history for a given firm, resulting from a discontinued or divested research program, a shift in strategic focus, or the acquisition of other companies which themselves have no work-in-process R&D. Any of these situations may correspond, in a given year, with successful patent applications due to applications previously submitted that year. This could produce outliers in the data.

In order to make the R&D measures comparable to the IT capital measures, the R&D values need to be adjusted to constant dollar amounts and then aggregated into a stock figure. Although a more finely-tuned price index that accounted for labour and capital inputs to R&D would be desirable, the National Science Foundation uses gross domestic product (GDP) implicit price deflators to convert R&D expenditures to constant (1992) dollars (National Science Foundation, 2003). In this research, the GDP deflator was then used to rescale the amounts to constant 1993 dollars to facilitate matching to IT spending data.

Annual R&D expense amounts constitute an input expressed as a "flow," rather than a stock. The stock measure for R&D is calculated with the perpetual inventory model employed by Hall (1990):

$$K_t = (1-\delta)K_{t-1} + R_t,$$

where K_t is the R&D capital stock at the end of period t, and R_t is the R&D expenditure during the year, in real dollars. Griliches and Mairesse (1984) found the choice of δ, the depreciation rate, makes little difference to production function estimates. In this case, the Hall method is used with a δ of 15 percent. The first year of R&D capital is obtained by using an initial "seed" value of R&D expenditure in the first year and dividing it by the sum of the depreciation rate δ and a pre-sample growth rate of 8 percent per year, as per Hall. The figure used for this purpose and for the yearly additions to the R&D capital stock is found by taking the average of the R&D expenditures in that year and the prior year. This is done to approximate the level of R&D that the firm was using throughout the year, rather than its state at year-end as reported in the financial statements.

One disadvantage of this method is that any missing observations for yearly R&D require interpolation. In the data used in this research, there are 25 companies for which this is the case. It may also be possible to increase the number of observations by including firms that do not report R&D for any of the years in the sample, despite having IT and patent data available. These observations are dropped from the main sample, although an additional regression is run where they are included. Some researchers have addressed the problem of having no R&D statistics for a large number of companies, in which case they substituted zeroes (Ziedonis and Hall, 2001; Brynjolfsson et al., 2002) and marked these observations with dummy variables. In the current dataset, doing this would add only 98 observations, which interestingly account for 1703 patents, although 522 of these were the result of single observation (North American Philips Corporation in 1992).

Patents and Patent Citations

Patents have been used as indicators of knowledge output in empirical studies since the 1960's (Schmookler, 1966). However, the shortcomings of simple patent counts have affected researchers' abilities to draw conclusions, since patents vary significantly in value. Recent investigations in patent statistics have shown promise in addressing this challenge by considering the citations (of other patents) made by each patent, and by inverting this data, determining the number of citations each patent has received (subject to temporal limitations).

Previous work by Trajtenberg (1990) shows that the number of citations received by a patent is correlated with its economic value. Hall, Jaffe and Trajtenberg's research applies this finding to the available (post-1975) computerized data at the USPTO, generating a large database of patents, citations made and received, and where possible, links to the Compustat identifiers for the patent-holder.

Another key piece of data provided is the field to which each patent belongs. This is to be distinguished from the industry groups to which firms belong. The field to which a patent belongs is derived from the USPTO classification of patents into similar groups, for example chemical process patents or steam engine patents. Since the USPTO has over 120,000 classifications, the authors created two simpler schemes: one of six fields and another more granular one of 36 fields. A company in any SIC industry might hold patents in any or all of the fields; there is no automatic relationship or connection between a company's industry and the fields in which it holds patents.

While the number of citations received by a particular patent may serve as a measure of its quality, the authors caution that patent citations cannot reasonably be compared across years or fields due to differences in both propensity to patent and propensity to cite. As Hall et al. (2001) point out (p. 27):

1. the average number of citations **received** by patents in their first 5 years has been rising over time;
2. the average number of citations **made** per patent has been rising over time; and
3. the observed citation-lag distributions for older cohorts have fatter "tails" than those of more recent cohorts.

In addition, since the citation lag distributions tend to be quite long (20-50 years), there is a serious truncation problem as one approaches 1999, the last year in the data. This leads the authors to recommend a three-year "safety lag" so as not to use biased data where not enough time has passed to allow the application process to complete and citations to begin. This imposes a limit upon the data available for analysis in this research: 1996 becomes the last valid year for which citation data may be used.

As mentioned earlier, when patents are grouped according to field classifications, it becomes clear that there are different propensities to patent and propensities to cite across the patent-producing fields, and across years. There are two methods proposed by Hall et al. (2001) to adjust patent citations to remove these effects. The first ("fixed-effects approach") is relatively simple to explain and calculate because it involves only the aggregate year and field strata. The second ("quasi-structural approach") has a better ability to extract signal from noise, but is more complex (and computationally expensive) since it requires the individual adjustment of each patent based on all of its particular citations.

The fixed-effects approach removes these trends by determining the average number of citations received for all patents in each year and each field. By dividing the citations received by any given patent by its year/field cohort average, the trends are removed and the patent's quantity of citations received may be compared to that of other patents. This method removes all year, field, and year-field effects. Hall et al. (2001) caution that it may not be reasonable to remove field effects, since there are systematically different rates in propensity to cite across fields. On the other hand, they also point out that one could construe these differences as artefacts (due to administrative differences over the years at the patent office) and should therefore be removed. Due to this confusion, they also present a method to remove the year and year-field effects but leave the field effects intact. This is achieved by dividing the year-field means by the overall means for each field. The disadvantage to this approach is that it may remove variance components that may be real.

In order to make patent data useful in the corporate context, each patent must be linked, if possible, to the company that owns it. Although the USPTO records this information ("assignee" and "assignee code"), it is not available in a format that facilitates automated linking to conventional sources of financial data, such as Compustat. Only 55-65% of patents between 1987 and 1996 were matched by Hall et al. (2001). This is due to two factors. First, the proportion of patents held by U.S. firms (and therefore eligible to match to Compustat) is roughly 70-80% (Hall et al. (2001), Table 3). Second, the authors used the list of Compustat companies as of 1989. As a result, there is some potential for measurement error relating to the matching of patents obtained by companies which did not yet exist in 1989, or existing companies later acquired by other companies.

Another factor to consider is the year used to link the patent to its inputs in the PPF. The USPTO records both the application year and the granting year. Often these can be 2-3 years apart, although some patents are granted in the application year. Hall et al. (2001) point out that using application year makes more sense for production functions, since this is when the resources devoted to preparing the patent application are known to be utilized. The granting year would be more useful to models that view patents as an input to some other variable, such as firm value. This research will therefore adopt the application year convention.

Linking the Data

To form a single body of data, the Compustat, NBER patent citation and CI data sets must be linked by both year and company identifier. The temporal scope of the data is limited by the years for which all three datasets were available or valid. The first year constraint is imposed by the CI data (1987), while the last year cutoff is imposed by the patent data's forward citation lag constraint (1996). The three datasets were linked using the U.S. standard CUSIP identifier. The CI data is used as the base set, to which the other two are linked. Out of 8,327 possible observations in the CI data, a total of 3,068 links are possible with the R&D data, while 2,491 links are possible from the CI data to the patent citation data. The intersection of the three data sets yields 1,809 observations in an unbalanced panel.[3]

However, this number of observations imposes the assumption that, should a company receive no patents or citations in a given year, the R&D and IT capital stocks in that year somehow do not count, and are discarded. This is a problem which should be addressed. Unfortunately there is no way (short of examining several hundred thousand patent records and attempting to match their assignees) to determine from the patent data whether a firm, for a given year, had no successful patent applications, or rather that the patents it was awarded were not linked to the company name. Since Hall et al. (2001) use a 1989 listing of companies, it is conceivable that any new or restructured company in any of the other years in this sample could fail to link to the patent data. A prudent method of generating valid zero-patent years is to fill in any gaps that exist between other valid years. For example, if a match is produced from a company to patents awarded in 1990 and 1992, but not in 1991, it is reasonable to assume that the company had no successful patent applications in 1991. This technique adds another 110 observations. Some manual adjustments to CUSIPs that failed to match due to notational differences

[3] The largest balanced panel that could be constructed, with full rank, would consist of 1320 observations for 132 firms.

add another 59 observations, bringing the total to 1,978. Unfortunately, employees data is missing for 16 observations so the final total is 1,962 (261 firms).

5. Analysis and Results

Summary Statistics

The available data consist of 262 firms with a total of 1,962 observations (Table 1). The firms are large, with a median of 14,492 employees and assets of nearly 2 billion dollars. Firms obtain 18 patents per year at the median, with 62 citations per year. This sample includes 53.4% of the patents matched to Compustat company identifiers by Hall et al. (2001) for the years 1987 to 1996, but only 18.46% of patents issued to corporations during this sample period.

As one would expect, there is significant variance in the amounts these firms spend on IT and R&D. The median IT capital amount, across all firms and years (in constant 1993 dollars), is approximately $14 million, with a standard deviation of $79 million. The mean R&D capital amount is approximately $276 million, with a standard deviation of $3.7 billion. The skew found in these variables is addressed largely by taking the natural log, as discussed in the model. Some of the variables have a minimum of zero. To avoid losing observations when the natural log is taken, a small arbitrary quantity is added to these variables.

The panel, as mentioned earlier, is unbalanced, with approximately half (132) of the firms reporting data in all 10 years (see Table 2). The remaining 130 firms have from 1 to 9 observations, with a median of 5 observations. The data is heavily concentrated in the manufacturing sector (250 of 262 firms), for the reason that firms in the service sector tend to engage in R&D and patenting less frequently, if at all. Although beyond the time frame of this sample, the 1998 change to patent law, allowing for the granting of business process patents (Surowiecki, 2003), may begin to increase the service sector's share.

Initial Regressions

The Cobb-Douglas and translog specifications are estimated, as given in (1) and (2), using OLS regression. Table 4 reports the results of these regressions. The dependent variable in all cases is (log) patent citations received, adjusted to remove year and year-field effects. The coefficient estimates in bold text are significant at the 5% level, with p-values shown in parentheses below the parameter estimates. The controls used in each regression consist of two variables representing the number of employees, as reported in the CI database: employees (EMPS) and employees squared ($EMPS^2$).

In column 1, a base specification is used to test whether the data supports the relationship established in the PPF literature. As expected, the parameter estimate for R&D is positive and significant at the 1% level. Columns 2 through 4 present the results of a fixed-effects panel data estimation. Column 2 shows the results for the Cobb-Douglas specification, which reports a change in the sign of the R&D estimate and a positive but statistically insignificant estimate for IT. Column 3 presents the Cobb-Douglas, augmented by the interaction term RDSTOCK*IT. The interaction term was found to be significant, but the main effects were not. The translog specification in column 4 finds the IT estimates to be negative, while the R&D and interaction terms are positive.

To determine whether the PPF is of the Cobb-Douglas or translog form, two tests are performed to establish whether the squared and interaction term estimates are jointly equal to 0. In the first test, under the null, the PPF is Cobb-Douglas, while under the alternative it is translog. The result from this test suggests that it is the translog form that best fits the production function (p=0.001). A second test is conducted to verify the joint significance of the squared terms. The result indicates that, jointly, RDSTOCK^2 and IT^2 do contribute to the explanatory power of the model once R&D, IT and RDSTOCK*IT are included in the model (p=0.001). This confirms that the PPF is best characterized as a translog production function. The Cobb-Douglas form, augmented by the interaction term, is presented for comparison purposes in column 3.

Subsequent to regression analysis, it is prudent to test whether the data meets assumptions of normality, homoskedasticity and uncorrelated errors. A histogram plot of residuals appears bell-shaped, conditioning on a number of firms reporting 0 patents. Further, an examination of the distribution of standardized residuals for outliers suggests their bias to be inconsequential. A total of 92 observations were found outside 2 standard deviations of the mean standardized residual. However, if the distribution is normal, one would expect 5% of observations (in this case, 98 data points) to be outside these bounds. Therefore, the data appear to be normally distributed. To detect heteroskedasticity, a test was conducted to determine if the variance of the error terms was linearly related to the expected value of the dependent variable. Under the null, the error terms are homoskedastic. Using an auxiliary regression, the product of the number of observations and the regression R-squared fails to reject the null at the 1% level using a chi-square distribution (with one degree of freedom). Finally, to test if autocorrelation influences the result of the panel regression, an AR1 panel regression was compared. Autocorrelation is marginally present, but the results from the regression are not sensitive to the correction.

Output Elasticities

In order to interpret the results of a translog estimation, the partial output elasticities are evaluated at the median, since they do not have a constant effect on expected patent citations. In Table 4, the partial output elasticities are reported below the regression results (p-values are reported in parentheses). For the translog model, a 1% increase in R&D stock (for a given level of IT capital) would yield a 0.9854% (p=0.0001)[4] increase in patent citations. A 1% increase in IT capital (for a given level of R&D stock) would yield a 2.2356% (p=0.0001) *decrease* in patent citations. Thus, while the effect of R&D capital upon patents is reasonably consistent with the extant literature, the impact of increasing the IT stock appears to be negative. The analysis also identifies decreasing returns to scale in the patent production function because the sum of the output elasticities is less than one.

Endogeneity

This research thus far assumes patents, as mentioned earlier, are the output of the innovation process as measured by R&D inputs. There is, however, general agreement among researchers that the relationship between R&D and patents is at least mildly reciprocal, since successful patents require additional development funds to become successful products. For the same reasons, IT may also be endogenous since innovation may come from and stimulate further spending on the firm's IT infrastructure. This leads one to suspect that endogeneity may be a problem in the proposed models.

---

[4] The standard error used to compute the p-value is computed analytically.

It is possible to test for endogeneity using the Durbin-Wu-Hausman (DWH) test. The test determines if the OLS estimator is biased when compared to a two-stage least squares (2SLS) estimation using an instrumental variable. The criteria for a good instrument are that it is highly correlated with the endogenous regressor and that it is uncorrelated with the model error term. A common source of instrumental variables tends to be lagged versions of the regressor. In this case, the lagged R&D stock is suitable: the correlation coefficient between current and lagged R&D stock is 0.9016. Interestingly, the DWH test fails to reject that R&D is exogenous (p=0.5978).

While it is possible that there is a reciprocal and contemporaneous relationship between patents and IT, there are several reasons to expect is not likely to be a measurable effect. First, the measure of IT capital is firm-wide in nature, so increases in R&D-related IT would be relative to all other activities of the firm. Second, IT capital responds slowly to changes in annual IT spending due to the accumulated capital stock of previous years. Nevertheless, the DWH test found that IT capital, instrumented with lagged IT capital (which share a correlation of 0.9658), failed to reject exogeneity (p=0.5071). As such, further specifications using instrumental variables are unnecessary.

Further Investigation

In order to test for bias due to short-term membership in the panel,

Table 5 presents the results from the three model specifications (translog, Cobb-Douglas and Cobb-Douglas with interaction) for a balanced panel of 132 firms. The results are very similar to those produced by the full sample. Again, the squared and interaction terms of the translog model are found to contribute jointly to the model, and the partial output elasticities are similar in sign and magnitude.

Although firm effects were included in the model specification, it is prudent to determine if their presence affects the results in a meaningful way. The model was estimated leaving out firm effects, but using the Huber/White/Sandwich estimator of variance. Results are provided in Table 6. Post-estimation tests of significance fail to reject the translog model. The partial output elasticities again show a positive effect for R&D and a negative one for IT, although both are smaller in magnitude than the specification that includes firm fixed effects. This indicates that the firm fixed effects have some overall impact, which is consistent with the literature that contends the successful implementation of IT is dependent upon other organizational factors.

One concern raised by the omission of firms with missing R&D expenditure data, in some or all sample years, is that this may bias the results. To address this concern, a larger sample was created by including these observations. The results are presented in Table 7. Although the sample is larger (N=2,648; 413 firms), the results are qualitatively similar. The partial output elasticity of IT is negative (-0.3950) and statistically significant (p=0.0020). This indicates that the results are not sensitive to missing R&D data.

Restricting the sample by industry is a natural avenue for further investigation. With firm fixed effects included in the model specification, industry fixed effects cannot be employed simultaneously. Instead, the sample may be restricted to industry groups if there are sufficient observations in these new panels to produce satisfactory results. Since 250 firms belong to the four manufacturing industry classifications, only these four groups were isolated. Table 8 presents the partial output elasticities. Post-estimation testing rejected both Translog and interaction models for non-durable manufacturing (column 1) and process manufacturing (column 4). The partial output elasticity estimates for non-durable manufacturing found both R&D and IT have negative and statistically significant impact on expected patent citation. Process manufacturing estimates were not statistically significant at the 5% level, but exhibit a positive sign on R&D and a negative one on IT. Durable manufacturing and high-technology manufacturing tested for the Cobb-Douglas with interaction form. For the former, the effect of increasing either of R&D and IT had a negative impact on expected patent citation, while for the high-technology manufacturing industry, the signs and magnitude of the estimates were similar to those in the unrestricted sample.

6. Conclusion

The results from the model presented here report that there is a negative marginal impact of overall IT investment upon the productivity of R&D, as measured by citation-weighted patent output. The results are robust to panel characteristics (balanced or unbalanced), the presence of firm effects, and industry classification at the 2-digit SIC level. In no case does IT exhibit a positive partial output elasticity. At best, it is slightly less negative for durable and non-durable manufacturing, although for these industries, R&D yields a negative partial output elasticity as well.

In the wake of recent IT productivity literature, these results are surprising, given the information-intensive nature of innovation. One interpretation of the results is that firms choose IT investment levels not to maximize patent output levels, but to maximize profits or facilitate a new business strategy, as discussed earlier. This idea can be seen by examining the estimated production function more carefully:

$$\begin{aligned}PAT\dot{C}IT &= 0.2967RDSTOCK + 0.0059RDSTOCK^2 - 1.5118IT + 0.0526IT^2 \\ &- 0.0231RDSTOCK \times IT + \dots .\end{aligned}$$

The marginal product of R&D capital with respect to patents is evaluated at the median level of IT (expressed in logs):

$$\frac{\partial PATCIT}{\partial RDSTOCK} = 0.2967 + 2(0.0059)(5.6235) - 0.0231(16.4441) = 0.3608$$

For firms in competitive markets, the price of an input is equal to its marginal product multiplied by marginal revenue. Therefore, firms are implicitly expecting marginal revenue for each patent citation, in terms of R&D stock, to be $2,771,618 (R&D is expressed in millions of dollars). The marginal product of IT capital with respect to patents is evaluated at the median level of R&D:

$$\frac{\partial PATCIT}{\partial IT} = -1.5118 + 2(0.0526)(16.4441) - 0.0231(5.6235) = -1.3001$$

Thus, in terms of IT capital, the implicit marginal revenue of a patent citation is $-769. First-order conditions for maximization of the production function require ratio of these two derivatives to be equal to the ratio of their prices. Since the ratios are not equal, it appears that the first-order conditions for maximizing output are not being met. Because the marginal product for IT is negative, for a given level of R&D, the median firm is over-invested in IT capital .

The primary implication of this result is that increasing the general level of IT investment in a firm cannot be assumed to automatically improve the productivity of formal R&D. What seems likely is that the measure of IT investment used here is inadequate to capture the nuances of how IT assists innovation. Two firms may have the same levels of R&D and IT capital, yet may choose very different implementations. Because the data do not reveal the extent of IT capital devoted to R&D, it is difficult to characterize this investment. Further, given the role of complementary investments in making effective use of IT, there can be significant latitude in the innovation efficiencies of intangible outputs, even among firms with ostensibly comparable levels of IT and R&D capital.

Returning to one of the major caveats of using patents as indicators of innovative output, one must exercise caution in assuming that the results reject the contribution of IT to innovation as a whole. Patent statistics may be useful as an index of output, but do not tell the entire story of a firm's innovative capacity. Another consideration in this regard is that IT capital may be more helpful for creating innovations that are not patentable (or were not, during the period under investigation). Dosi (1998) notes that patents tend to be more effective protection for product patents, while lead times and learning curves are more effective for process innovations.

Two changes in the structure of corporate R&D over the past number of years may further help to explain the lack of positive IT impact in the patent production function specified. One phenomenon is the departure of firms from doing in-house research, substituting the purchase of R&D embedded in products or licensing arrangements from other firms. For example, pharmaceutical companies often license new molecules from biotech firms (the latter do not appear in this sample); other manufacturing industries exhibit similar patterns of R&D outsourcing (Love and Roper, 2001). Another confounding phenomenon may be the strategy of some high-technology firms to grow through continuous acquisition of smaller, innovating companies. In such cases, the IT capital and in-process R&D of the acquired firms would appear on the parent firm's financial statements, while the accumulated R&D knowledge stock would not. As a result, the IT capital of the parent company would be large relative to its R&D capital. Further, if the acquired company held patents arising from its R&D investment, the patent data used in this research would not have been updated to indicate the patents' new owner, thus reducing the estimated patent productivity of R&D and IT for the acquiring firm.

Limitations and Future Research

The principal limitation to this research is the inability to discern the proportion of IT capital invested in formal innovation. Although the emphasis is on detecting the impact of overall IT investment upon patent output, the fact that some of this investment is also being counted in the R&D expenditure means that determining the impact upon patent production is bound to be confused among the inputs. The double-counting problem was encountered by Kortum and Lerner (1998), who acknowledged that some venture capital would also be counted in firms' R&D expenditures. They noted early in their paper that this made it less likely an impact of venture capital would be found on R&D productivity, although their findings were nevertheless significant. Hall and Mairesse (1995) also noted the large impact of double-counting on R&D productivity measurement due to R&D-related labour being included in their overall firm labour variable. However, since their data was derived from a government survey and therefore much more detailed, they were able to employ a correction to their data by estimating the proportion of firm labour devoted to R&D. This remedy is not available for the data used in the present research.

Future research in this topic area depends chiefly upon the availability of new and more detailed data. Although Hall and Mairesse (1995) note that the choice of R&D deflator matters little when applied across all industries, it may be that shifts in demand for research skills in certain industries may have an impact if the deflator is constructed in a more comprehensive fashion (taking into account shifting prices in R&D labour and capital). The double-counting issue identified above presents a challenge which could be overcome if the proportion of IT spent in R&D endeavours was known. A limitation also exists in patenting data in the pre-1998 period, when business process patents were not yet approved. This structural break may be a fruitful area for study once sufficient time has passed for patents and citations to have accrued in volumes adequate for econometric analysis. Finally, it would be worth investigating other methods of measuring innovation, such as new product introductions, in order to capture more of the total innovation of the firm.

References

Abramowitz, M., "Resources and Output Trends in the United States Since 1870", *American Economic Review, 46*, 2 (1956), pp. 5-23.

Adams, G., and Lamont, B., "Knowledge management systems and developing sustainable competitive advantage," *Journal of Knowledge Management, 7,* 2 (2003), pp. 142-154.

Arrow, K., "The Economic Implications of Learning by Doing," *Review of Economic Studies*, *29*, 3, (June 1962), pp. 155-173.

Baily, M., Chakrabarti, A., and Levin, R., "Innovation and Productivity in U.S. Industry" (in Reports), *Brookings Papers on Economic Activity*, *1985*, 2 (1985), pp. 609-639.

Barua, A., and Lee, B., "The Information Technology Productivity Paradox Revisited: A Theoretical and Empirical Investigation in the Manufacturing Sector," *The International Journal of Flexible Manufacturing Systems*, 9 (1997), pp. 145-166.

Berndt, E.R., and Rappaport, N.J., "Price and Quality of Desktop and Mobile Personal Computers: A Quarter-Century Historical Overview," *American Economic Review*, *91,* 2, (2001), pp. 268-273.

Bound, John, Cummins, C., Griliches, Z., Hall, B.H., and Jaffe, A., "Who Does R&D and Who Patents?" In *R&D, Patents and Productivity*, edited by Zvi Griliches. Chicago: University of Chicago Press, 1984, pp. 21-55.

Bresnahan, T., and Trajtenberg, M., "General Purpose Technologies: 'Engines of Growth'?", *Journal of Econometrics,* 65, (1995), pp. 83-108.

Brynjolfsson, E., "The Productivity Paradox of Information Technology," *Communications of the ACM*, *36*, 12, (1993), pp. 67-77.

Brynjolfsson, E. and Hitt, L., "Paradox Lost? Firm Level Evidence of High Returns to Information Systems Spending," *Management Science*, *42*, 4, (1996), pp. 541-558.

Brynjolfsson, E., and Hitt, L., "Computing Productivity: Firm Level Evidence," *Review of Economics and Statistics*, *85,* 4, (2003), pp. 793-808.

Brynjolfsson, E., Hitt, L., and Yang, S., "Intangible Assets: Computers and Organizational Capital," *Brookings Papers on Economic Activity: Macroeconomics, 1*, 2002, pp. 137-199.

Bryrijolfsson, E., and Yang, S., "Information Technology and Productivity: A Review of the Literature," *Advances in Computers*, 43, (1996), pp. 179-214.

Chwelos, P. "Approaches to Performance Measurement in Hedonic Analysis: Price Indexes for Laptop Computers in the 1990's," *Economics of Innovation and New Technology, 12,* 3, (2003), pp. 199-224.

Chwelos, P., Ramirez, R. and Kraemer, K., "Substitute or Complement? The Role of Information Technology Capital in Firm-Level Production 1987-1999," working paper (2004).

DeLong, B., and Summers, L., "The 'New Economy': Background, Historical Perspective, Questions, and Speculations," *Economic Review – Federal Reserve Bank of Kansas City, 86,*4, (2001), pp. 29-59.

Denison, E., *Trends in American Economic Growth, 1929–1982,* Washington, D.C: Brookings Institution, 1985.

Dewan, S., and Min, C., "The Substitution of Information Technology for Other Factors of Production: A Firm Level Analysis," *Management Science, 42,* 12, (1997), pp. 1660-1675.

Dosi, G., "Sources, Procedures and Microeconomic Effects of Innovation," *Journal of Economic Literature*, *26* (September 1998), pp. 1120-1171.

Grabowski, H., "Patents, Innovation and Access to New Pharmaceuticals," *Journal of International Economic Law, 5,* 4, (December 2002), pp. 849-860.

Greenan, N., Mairesse, J. and Topiol-Bensaid, A., "Information Technology and Research and Development Impacts on Productivity and Skills: Looking for Correlations on French Firm Level Data," NBER Working Paper 8075, (2001).

Griliches, Z., "Issues in Assessing the Contribution of Research and Development to Productivity Growth," *The Bell Journal of Economics*, *10*, 1 (Spring, 1979), pp. 92-116.

Griliches, Z., "Patent Statistics as Economic Indicators: A Survey," *Journal of Economic Literature*, *28* (December 1990), pp. 1661-1707.

Griliches, Z. and Mairesse, J., "Productivity and R&D at the Firm Level." In *R&D, Patents and Productivity,* edited by Zvi Griliches. Chicago: University of Chicago Press, (1984), pp. 339-374.

Gurbaxani, V. and Whang, S., "The Impact of Information Systems on Organizations and Markets." *Communications of the ACM*, *34*, 1 (1991), pp. 61-73.

Hall, B.H., "The Manufacturing Sector Master File: 1959-1987," NBER Working Paper No. 3366 (May 1990).

Hall, B.H., Griliches, Z. and Hausman J.A., "Patents and R&D: Is There a Lag?" *International Economics Review, 27*, 2 (1986), pp. 265-83.

Hall, B.H., Jaffe A.B., and Trajtenberg, M., "The NBER Patent Citations Data File: Lessons, Insights and Methodological Tools." NBER working paper 8498, (2001).

Hall, B.H., and Mairesse, J., "Exploring the relationship between R&D and productivity in French manufacturing firms," *Journal of Econometrics,* 65 (1995), pp. 263-293.

Hall, B.H., and Ziedonis, R.H., "The Effects of Strengthening Patent Rights on Firms Engaged in Cumulative Innovation: Insights from the Semiconductor Industry," in Libecap, Gary (ed.), Entrepreneurial Inputs and Outcomes: New Studies of Entrepreneurship in the United States, Vol. 13 of Advances in the Study of Entrepreneurship, Innovation, and Economic Growth, Amsterdam: Elsevier Science, 2001.

Hatch, N.W., and Mowery, D.C., "Process Innovation and Learning By Doing in Semiconductor Manufacturing," *Management Science, 44,* 11 (November 1998), pp. 1461-1477.

Jones, C. I., "R&D-Based Models of Economic Growth," *Journal of Political Economy, 103,* 4, (1995), pp. 759-784.

Klein, J., Gee, D., and Jones, H. "Analysing clusters of skills in R&D—core competencies, metaphors, visualization, and the role of IT," *R&D Management, 28,* 1, (January 1998), pp.37-43.

Kohli, R., and Devaraj, S., "Measuring Information Technology Payoff: A Meta-Analysis of Structural Variables in Firm-Level Empirical Research," *Information Systems Research, 14,* 2, (2003), pp. 127-145.

Kortum, S. and Lerner, J., "Does Venture Capital Spur Innovation?" NBER Working Paper 6846, (1998).

Kremp, E., and Mairesse, J., "Knowledge Management, Innovation and Productivity: A Firm Level Exploration Based on French Manufacturing CIS3 Data," NBER Working Paper 10237, (2004).

Lee, H. and Choi, B., "Knowledge Management Enablers, Processes, and Organizational Performance: An Integrative View and Empirical Examination," *Journal of Management Information Systems*, *20*, 1 (Summer 2003), pp. 179-228.

Lerner, J., "Patenting in the Shadow of Competitors," *Journal of Law and Economics*, (October 1995), pp. 463-495.

Lev, B., and Sougiannis, T., "The capitalization, amortization and value-relevance of R&D," *Journal of Accounting and Economics*, 21, (1996), pp. 107-138.

Love, J., and Roper, S., "Outsourcing in the innovation process: Locational and strategic determinants," *Papers in Regional Science*, 80, (2001), pp. 317-336.

Malone, T., Yates, J., and Benjamin, R. "Electronic Markets and Electronic Hierarchies," *Communications of the ACM, 30*, 6, (1987), pp. 484-497.

Malone, T., and Crowston, K., "The Interdisciplinary Study of Coordination," *ACM Computing Surveys*, *26*, 1 (1994), pp. 87-119.

Moore, G.E., "Cramming more components onto integrated circuits," *Electronics, 38,* 8, (April 19, 1965).

Nambisan, S., "Information Systems as a Reference Discipline for New Product Development," *MIS Quarterly*, *27*, 1 (March 2003), pp. 1-18.

National Science Foundation "IRIS" database, Table A-1 (R&D GDP deflator), retrieved September 2003 from http://www.nsf.gov/sbe/srs/iris/start.htm.

Nelson, R.R., "Institutions Generating and Diffusing New Technology." In *Frontiers in innovation diffusion*. Eds.: Fabio Arcangeli, Paul David and Giovanni Dosi. Oxford: Oxford U. Press, 1986.

Oliver, Joseph R., "Accounting and tax treatment of R&D: An update," *The CPA Journal, 73,* 7, (July 2003), pp. 46-49.

Pakes, A. and Griliches, Z. "Patents and R&D at the Firm Level: A First Look," in Griliches, Z., ed. *R&D, Patents, and Productivity*. Chicago, IL: University of Chicago Press, 1984, pp. 55-72.

Pardey, Philip G., "The Agricultural Knowledge Production Function: An Empirical Look," *The Review of Economics and Statistics*, *71*, 3, (August 1989), pp. 453-461.

Porter, M.E., "Strategy and the Internet," *Harvard Business Review,* March 2001, pp. 63-78.

Rosenberg, Nathan, *Inside the black box.* Cambridge: Cambridge U. Press, 1982.

Rosenberg, Nathan, "The Commercial Exploitation of Science by American Industry," in *The uneasy alliance: Managing the productivity-technology dilemma.* Ed.: Kim B. Clark, Robert H. Hayes and Christopher Lorenz. Cambridge, MA: Harvard Business School Press, 1985.

Scherer, F.M. "Firm Size, Market Structure, Opportunity, and the Output of Patented Inventions," *American Economic Review*, 55 (1965), pp 1097-1125.

Scherer, F.M., *Innovation and growth. Schumpeterian perspectives.* Cambridge: MIT Press, 1986.

Schmookler, J., *Invention and economic growth.* Cambridge: Harvard U. Press, 1966.

Solow, R.M., "A Contribution to the Theory of Economic Growth," *Quarterly Journal of Economics*, *70,* 1, (1956), pp. 65-94.

Surowiecki, J., "Patent Bending," *The New Yorker*, July 14 and 21, (2003). Retrieved April, 2004 from http://www.newyorker.com/printable/?talk/030714a_talk_surowiecki.

Thatcher, M.E., and Oliver, J.R., "The impact of technology investments on a firm's production efficiency, product quality, and productivity." *Journal of Management Information Systems, 18,* 2 (Fall 2001), pp. 17-45.

Trajtenberg, Manuel, "A penny for your quotes: patent citations and the value of innovations," *RAND Journal of Economics, 21,* 1, (Spring 1990), pp. 172-187.

USPTO, "Performance and accountability report: Fiscal Year 2003." Retrieved May, 2004 from http://www.uspto.gov/web/offices/com/annual/2003/index.html.

Tables

Table 1: Descriptive Statistics

Variable	Units	Median	Std. Dev.	Min	Max
IT	thousands	13,854.23	79,367.60	9.18	1,160,282.00
RDSTOCK	millions	276.86	3,694.68	0.00	39,612.78
EMP	units	14,492.50	73,608.05	1,500.00	876,000.00
Patents Received	units	18.00	177.70	0.00	2,405.00
Citations Received	units	62.00	1,026.81	0.00	12,795.00
Citations Received (adjusted)	units	59.38	918.65	0.00	14,919.47
Assets	thousands	1,958,550.00	21,700,000.00	0.00	252,000,000.00

Table 2: Panel Characteristics

Years	Observations	Firms
10	1320	132
9	189	21
8	192	24
7	49	7
6	54	9
5	45	9
4	32	8
3	33	11
2	46	23
1	18	18

Table 3: Industry Classifications

Industry	Sector	SIC	Firms
Agriculture, Forestry, Fishing	[none]	01xx-09xx	5
Non-Durable Manufacturing	Manufacturing	20xx-23xx, 27xx	19
Durable Manufacturing	Manufacturing	24xx-25xx, 30xx-35xx, 39xx	77
Process Manufacturing	Manufacturing	26xx, 28xx-29xx	73
High-tech Manufacturing	Manufacturing	36xx-38xx, 3571	81
Wholesale Trade	Services	50xx-51xx	3
Services	Services	70xx-89xx	4

Table 4: Regression Results

	Base	**Cobb-Douglas**	**Cobb-Douglas with Interaction**	**Translog**
Column	(1)	(2)	(3)	(4)
ln RDSTOCK	**0.0001** (0.0001)	**-0.0438** (0.0250)	0.0816 (0.2137)	0.2967 (0.2177)
ln $RDSTOCK^2$				**0.0059** (0.0031)
ln IT		0.0728 (0.1630)	0.1121 (0.0847)	-1.5118 (0.3348)
ln IT^2				**0.0526** (0.0105)
ln RDSTOCK* IT			**-0.0076** (0.0129)	**-0.0231** (0.0132)
N	1962			
Firms	261			
Controls	EMPS, $EMPS^2$, year			
R-squared (overall)	0.1806	0.2034	0.1802	0.2346
$\eta_{PatCit,RDSTOCK}$		**-0.0438** (0.0250)	**-0.0436** (0.0251)	**0.9854** (0.0001)
$\eta_{PatCit,IT}$		0.0728 (0.1630)	0.0692 (0.1871)	**-2.2356** (0.0001)

Notes for all results tables:

1. The column containing the translog model is shaded, corresponding to tests indicating it best fits the data.
2. Coefficients in bold are significant at the 5% level.
3. *p*-values appear in parentheses below the coefficient.

Table 5: Balanced Panel Results

	Cobb-Douglas	**Cobb-Douglas with Interaction**	**Translog**
Column	(1)	(2)	(3)
ln RDSTOCK	-0.0272 (0.3410)	0.1084 (0.2859)	0.2756 (0.3068)
ln $RDSTOCK^2$			**0.0222** (0.0129)
ln IT	-0.0065 (0.9280)	0.0383 (0.1187)	-1.5831 (0.4803)
ln IT^2			**0.0536** (0.015)
ln RDSTOCK* IT		**-0.0082** (0.0172)	**-0.0333** (0.0183)
N	1320		
Firms	132		
Controls	EMPS, $EMPS^2$, year		
R-squared (overall)	0.0087	0.0196	0.0682
$\eta_{PatCit,RDSTOCK}$	-0.0272 (0.3410)	-0.0293 (0.3092)	**1.2855** (0.0001)
$\eta_{PatCit,IT}$	-0.0065 (0.9280)	-0.0104 (0.8869)	**-2.5712** (0.0001)

Table 6: Results without firm fixed effects

	Cobb-Douglas	Cobb-Douglas with Interaction	Translog
Column	(1)	(2)	(3)
ln RDSTOCK	-0.0354 (0.5560)	-0.8259 (0.5759)	-0.5949 (0.5389)
ln $RDSTOCK^2$			**-0.0045** (0.0114)
ln IT	**0.9393** (0.0001)	0.6840 (0.1877)	-1.3948 (0.8137)
ln IT^2			**0.0660** (0.0261)
ln RDSTOCK* IT		**0.0477** (0.0340)	**0.0357** (0.0334)
N	1962		
Firms	261		
Controls	EMPS, $EMPS^2$, year		
R-squared (overall)	0.3718	0.3742	0.3803
$\eta_{PatCit,RDSTOCK}$	-0.0354 (0.5560)	-0.0419 (0.5195)	0.0748 (0.2128)
$\eta_{PatCit,IT}$	**0.9393** (0.0001)	**0.9524** (0)	**-0.2489** (0.0147)

Table 7: Results including missing R&D observations

	Cobb-Douglas	C-D+i	Translog
Column	(1)	(2)	(3)
R&D	**0.1206** 0.0420	-0.7060 (0.2147)	-0.1788 (0.296)
R&D * R&D			**0.0407** (0.0269)
IT	-0.1161 0.3270	-0.2943 (0.1261)	-1.5509 (0.8298)
IT * IT			**0.0286** (0.0162)
RD * IT		**0.0526** (0.0131)	**0.0447** (0.0206)
N	2648		
Firms	413		
Controls	EMPS, $EMPS^2$, year		
R-squared (overall)	0.2795	0.2726	0.3758
$h_{PatCit,RD}$	**-0.0438** (0.025)	**0.1586** (0.008)	**0.9572** (0.0022)
$h_{PatCit,IT}$	0.0728 (0.163)	0.0016 (0.9894)	**-0.3950** (0.0020)

Table 8: Output elasticities of manufacturing industries

	Non-durable Manufacturing	Durable Manufacturing	High-Tech Manufacturing	Process Manufacturing
Column	(1)	(2)	(3)	(4)
Model form	Cobb-Douglas	Cobb-Douglas with Interaction	Translog	Cobb-Douglas
$\eta_{PATCIT,RD}$	**-.3361** (0.008)	**-.0753** (0.013)	**0.5191** (0.001)	.0627 (0.086)
$\eta_{PATCIT,IT}$	**-.8311** (0.013)	**-.2760** (0.018)	**-1.5807** (0.001)	-.1410 (0.096)
N	145	575	620	536
Firms	19	77	80	73
R-sq	0.0003	0.0801	0.2471	0.2277