

The Role of the IT Revolution in Knowledge Diffusion, Innovation and Reallocation^{*}

Salomé Baslandze[†]

Abstract

What is the impact of information and communications technologies (ICT) on aggregate productivity growth and industrial reallocation? In this paper, I analyze the impact of ICT through facilitating *knowledge diffusion* in the economy. There are two opposing effects. The increased flow of ideas between firms and industries improves learning opportunities and spurs innovation. However, knowledge diffusion through ICT also results in broader accessibility of knowledge by competitors, reducing expected returns from research efforts and hence harming innovation incentives. The nature of the tradeoff between these opposing forces depends on an industry's technological characteristics, which I call *external knowledge dependence*. Industries whose innovations rely more on external knowledge benefit greatly from knowledge externalities and expand, while more self-contained industries are more affected by intensified competition and shrink. This results in the reallocation of innovation and production activities toward more externally-focused, "knowledge-hungry" industries. I develop a general equilibrium endogenous growth model featuring this mechanism. In the model, firms belonging to technologically heterogeneous industries learn from external knowledge and innovate. These firms' abilities to access external information is governed by ICT. Using NBER patent and citations data together with BEA industry-level data on ICT, I empirically validate the mechanism of the paper. Quantitative analysis from the calibrated model illustrates that it is important to account for both technological heterogeneity and the knowledge-diffusion role of ICT to explain U.S. trends in productivity growth and sectoral reallocation in recent decades. Counterfactual experiments are conducted to quantitatively assess separate channels and illustrate various growth decompositions.

JEL Classification: O3, O4

Keywords: Information and Communications Technologies, Endogenous Growth, Innovation, R&D, Industry Heterogeneity, Reallocation, Knowledge Spillovers.

^{*}I am grateful to Ufuk Akcigit, Jeremy Greenwood, and Guillermo Ordonez for their invaluable guidance. I thank Harold Cole, Cecilia Fieler, Chris Forman, Lorin Hitt, Karen Kopecky, Dirk Krueger, Iouri Manovskii, and Guido Menzio for all the useful comments. I especially would like to thank Penn Growth Reading Group members for their helpful suggestions and feedbacks. I also thank seminar participants at Penn Macro Club, Penn Macro Workshop, Federal Reserve Bank of Atlanta, and Bilkent Macroeconomics Jamboree. All errors are mine.

[†]Einaudi Institute for Economics and Finance (EIEF), Via Sallustiana 62, Rome 00187, Italy. E-mail: salome.baslandze@eief.it

1 Introduction

In 1973, when Kenbak Corporation went bankrupt after its unsuccessful sales of the world's first personal computer, Kenbak-1, few people could foresee the beginning of a new information and communications era and the structural transformation it would bring to the economy. Since then, we have witnessed a widespread penetration and a tremendous rise in the quality of information and communications technologies (ICT). Most of the economic analysis of the impact of these changes has focused on the productivity gains from ICT resulting from the associated cost reductions or complementary organizational changes. However, goods brought about by the IT revolution also allow firms to access a wider set of external information and to learn from others. This *knowledge diffusion* aspect of ICT, in turn, can have a pronounced effect on the processes through which new knowledge and productivity-improving innovation are created. In this paper, I study the impact of ICT on productivity growth and sectoral reallocation by analyzing the role of ICT as a tool for improved knowledge dissemination in the economy.

To understand the impact of knowledge diffusion brought about by ICT, we need to take into account two opposing channels. On the one hand, the increased flow of ideas between firms and industries allows for more learning of external knowledge. Because innovations often build “*on the shoulders of giants*”, this increased diffusion of knowledge makes the process of innovation more efficient and, in turn, positively affects firms' innovation incentives. On the other hand, this same process of knowledge diffusion in the economy results in greater accessibility of knowledge to both potential and existing competitors. Thus, by increasing *creative destruction* in the economy and hence reducing expected returns from research efforts, ICT can decrease the innovation incentives of firms.

At the heart of this paper is the interplay between these two forces. Depending on how these opposing forces differ across industries, ICT can lead to a sizable *sectoral reallocation* in the economy. Industries are inherently heterogeneous in the extent of external knowledge required to generate new innovations. For example, innovations in an industry such as electronic testing and measurement instruments, which is a part of the high tech sector, rely on advances in a broad set of external fields, such as nanotechnology, engineering, chemicals and others. However, innovations in the textile industry, for example, are more self-contained and do not require many external ideas for own product advancement. As a result, with the greater knowledge diffusion brought about by ICT, more externally dependent industries will benefit and expand, while more self-contained industries will be more affected by competition

and therefore will shrink.

Exploring this mechanism in a structural framework is important in order to address the main questions of the paper: What is the quantitative impact of ICT on aggregate productivity growth? What is the impact of ICT on sectoral reallocation? And how much do different facets of ICT contribute to these dynamics?

To answer these questions, I proceed in three steps. The first step is to empirically validate the mechanism at the core of the paper by reduced form facts from the NBER patent and citations data, together with the industry-level ICT data from the Bureau of Economic Analysis (BEA). I begin by constructing input-output matrices of citations data and find evidence of increased knowledge diffusion, as measured by the extent of citations to patents from external fields. Next, using patent citations information, I develop a new industry classification scheme based on the measure of *external knowledge dependence*. This measure reflects how intensively new inventions created in an industry rely on knowledge from outside. Then, consistent with the mechanism in this paper, I document that there has been a reallocation of both innovation and production activities toward sectors that are more externally dependent. In addition, I find that the data supports the relationship between increased ICT and increased competition, and ICT’s interaction with heterogeneous knowledge spillovers.

In the second part of my paper, in order to understand the interplay of opposing forces caused by knowledge diffusion and to quantify the importance of these channels for growth and observed reallocation, I build a new general equilibrium Schumpeterian growth model featuring technological and competitive spillovers from knowledge diffusion. In the model, firms belonging to technologically heterogeneous industries are engaged in research and development (R&D); they learn from external knowledge, and their ability to access this external information is governed by ICT. Improvements in ICT are associated with a widening of firms’ information sets, which increases their research efficiency. Firms belonging to more externally dependent industries experience higher technological externalities from knowledge diffusion because external learning is more important for them. As firms become more efficient at research, they increase their innovation and productivity, while also forcing other firms whose technologies they improve upon to exit. As a result, an increase in knowledge diffusion caused by ICT lets the more externally dependent industries expand, resulting in sectoral reallocation toward more “knowledge-hungry” industries. The model also features the more traditional investment-specific impact of ICT on production, which I refer to as a *direct* impact of ICT. Productivity improvements from the direct channel come about as a

result of an increase in ICT investment and ICT-specific technological progress, which are treated here as exogenous.

Lastly, I quantitatively assess the implications of the IT revolution for innovation, growth and sectoral reallocation in the U.S. economy. I calibrate the structural parameters of the model to match important features of the data in 1976-1978, before the IT revolution. I evaluate the implications of the model for the period 1976-2003 by feeding in an exogenous knowledge diffusion process due to ICT, estimated outside the model. Quantitative evaluation of the model illustrates the importance of knowledge diffusion through ICT in explaining important aspects of the U.S. experience in aggregate growth and sectoral reallocation from 1976 to 2003. In particular, through the lens of my model, the increase in knowledge diffusion and implied heterogeneous knowledge spillovers explain 74% of the increased difference between innovation intensities across the most and least externally dependent industries, and explains essentially all of the reallocation of production between these two types of industries. During this period, the U.S. economy saw an acceleration of average growth in labor productivity. This acceleration of growth is largely attributed to the rise in ICT: the model explains 76% of the increase in growth rates between the first and second halves of the sample period. Growth decompositions show that, during the initial period of the IT revolution, most of the returns came from the direct investment-specific impact of ICT. These were also sizable in the mid-1990s, but became very small afterward. The main source of growth acceleration in the second half of the sample period comes from the effect of ICT on knowledge diffusion and the consequent effect on innovation incentives and reallocation. This effect on growth is sustainable in the long run. Finally, I evaluate the importance of the competition channel induced by knowledge diffusion from ICT. This channel turns out to be quantitatively large; in the absence of a negative competition effect, productivity growth in the last sample year would increase by 28%.

Related Literature. This paper contributes to several strands of literature. First, it contributes to the extensive literature studying the impact of ICT on productivity and growth (Stiroh, 2002; Jorgenson, Ho, and Stiroh, 2003; Oliner and Sichel, 2000; Brynjolfsson and Yang, 1996; Brynjolfsson and Hitt, 1996, 2003; Tambe, Hitt, and Brynjolfsson, 2012; Dedrick, Gurbaxani, Kraemer, 2003). The prevalent channel considered in the literature is a “direct” productivity-improving channel, which entails treating ICT as one type of capital investment or analyzing complementary organizational changes (better customer-supplier linkages, electronic commerce, and others). This direct channel is deemed to have

a short-term (immediate or lagged) impact on growth unless it is sustained by constant improvements and increased investment in ICT. The “indirect” channel proposed in this paper, however, is found to be responsible for sizable long-term returns, because it changes certain fundamental characteristics of the process of innovation. This finding contrasts with some of the pessimistic views (Gordon, 2000) on the transitory productivity returns from the IT revolution.

The research line studying the impact of ICT on industry-level growth finds a great deal of heterogeneity of productivity returns from ICT (see Dedrick, Gurbaxani, Kraemer, 2003, for a comprehensive review; see also Acemoglu, Autor, Dorn, Hanson, and Price, 2014). Some industries have experienced significant and persistent productivity growth since the IT revolution, while others have seen only small and temporary gains. This paper sheds some light on these questions as it illustrates how fundamental heterogeneity in technological characteristics of industries may be important for the realized size and persistence of productivity returns from ICT.

The papers studying productivity slowdowns caused by the introduction of new technologies constitute an integral part of the literature on IT and aggregate growth. Papers by Greenwood and Yorukoglu (1998), Yorukoglu (1998) and Hornstein and Krusell (1996) emphasize the role of learning or quality mismeasurements in explaining why the initial stages of rapid investment-specific technological progress may be associated with productivity declines. In this paper, I abstract from the learning mechanism and consider the direct investment-specific channel as implying immediate returns. Finally, there are a few empirical studies that relate IT and innovation. Kleis, Chwelos, Ramirez, and Cockburn (2012) look at firm-level regressions of patents on IT and find a positive effect, especially in the mid to late 1990s. Forman and Zeebroek (2012) and Forman, Greenstein, and Goldfarb (2014) find that the Internet influenced long-distance collaborations, while Ajay and Goldfarb (2008) find that an early version of the Internet, Bitnet, facilitated collaboration among U.S. universities. In addition, Boppart and Staub (2014) show that the online accessibility of economic articles led to more follow-on research that referenced different strands of economic literature in novel ways.

Second, this paper relates to the literature on reallocation and structural transformation. This literature has mostly focused on reallocation of activities across three broadly defined sectors: manufacturing, agriculture and services (Buera and Kaboski, 2012a, 2012b; Herrendorf, Rogerson, and Valentinyi, 2014; Acemoglu and Guerrieri, 2008). In this paper, I consider the sectoral reallocation between more narrowly defined sectors. I propose a new

mechanism emphasizing the idea that differences in the underlying nature of industries' technological dependence, together with exogenous shifts in knowledge diffusion, drive changes in sectoral composition. The literature has emphasized the role of reallocation from less to more productive firms in aggregate productivity growth (in the U.S. or across countries) (Foster, Haltiwanger, and Krizan, 2001; Foster, Haltiwanger, and Syverson, 2008; Syverson, 2011; Bartelsman, Haltiwanger, and Scarpetta, 2013; Acemoglu, Akcigit, Bloom, and Kerr, 2013; Lentz and Mortensen, 2008). This paper complements the literature on the impact of reallocation. Here, reallocation occurs between industries with different external knowledge needs. A move toward unrestricted knowledge flows in the economy is accompanied by reallocation toward industries that are intrinsically more efficient. This kind of reallocation of economic activity also implies important gains in aggregate growth.

Third, this project builds on and extends the models of Schumpeterian growth (Aghion and Howitt, 1992; Klette and Kortum, 2004; Akcigit and Kerr, 2010; Aghion, Akcigit, and Howitt, 2013; Acemoglu and Akcigit, 2012). Most of the theoretical work using models of endogenous growth are studies of single sectors or multiple-sector models with the same innovation production process across sectors. Cai and Li (2012) is an exception in this regard. They consider the heterogeneity of industries in terms of their knowledge applicability. My theoretical departure relative to the previous literature is formalizing the notion of industry's technological interdependence and cross-industry heterogeneity with regard to technological dependence, as well as modeling knowledge diffusion in this context.

Last, methodologically we can draw parallels with a paper by Rajan and Zingales (1998). In order to understand the link between financial development and growth, they propose measures of an industry's dependence on external finance and observe that industries which are more in need of external finance grow faster in countries with better financial markets. Similarly to their paper, by proposing a particular theoretical mechanism for the relationship between ICT and growth, and examining cross-sector variation between industries with different technological dependence, I am able to tease out the growth impact of ICT and partially alleviate the concerns about endogeneity and model specification that are associated with reduced-form specifications of this kind.

The paper proceeds as follows. Section 2 describes the data and documents motivating empirical facts. Section 3 develops a theoretical model. I calibrate the model and present the results from a range of quantitative experiments in Section 4. Section 5 concludes.

2 Empirical Evidence

In this section, I enumerate new facts on innovation, knowledge dependence and ICT, and the well-known large increase in ICT usage in the U.S. economy over recent decades. First, I show evidence of the increase of knowledge diffusion as measured by patent citations to external fields. I postulate a new measure of an industry’s technological characteristic – external knowledge dependence – and document that there has been a sectoral reallocation toward more externally dependent industries since the IT revolution. A second set of reduced-form facts supports the mechanism of the paper. In particular, I show that higher ICT is associated with higher levels of various industry-level competition measures. I also provide evidence of heterogeneous technological spillovers for industries with varying degrees of external dependence.

2.1 Data

My main data sources are derived from the patent and citations dataset from NBER and the non-residential fixed capital dataset from BEA. The NBER Patent Database contains information on all patents granted by the United States Patent and Trademark Office (USPTO) from 1976 to 2006 (Hall, Jaffe and Trajtenberg, 2001). The data contains rich information on patent characteristics such as their technological classifications, backward citations to prior art knowledge, forward citations received from future patents, application dates and others. In addition, each patent contains information on identifiers of patent assignees, which enables analysis of firm-level patent data as well.

A feature of this dataset most important for our purpose is ability to trace knowledge flows embedded in patents through the citation record. Unfortunately, there is no available information on citations made by patents granted before 1976 because of the lack of computerized USPTO citations data, so the initial year of the sample is 1976. There is a well-known truncation problem with patent data. It comes from the fact that there is a time lag between patent application and the actual grant of the patent. Because the dataset only contains patents granted, the number of patents issued in the later years of the sample can be underestimated. A similar truncation issue is observed for the number of forward citations of patents – for patents issued in later sample periods, there is not enough time to accumulate as many forward patent citations.

To mitigate these problems, I take the following steps: Firstly, I drop the last years of the

sample and the resulting sample period becomes 1976-2003¹. Secondly, for forward citations, I use a truncation-adjustment factor from Hall, Jaffe and Trajtenberg (2001). And finally, where appropriate, regressions contain year dummies to control for time-fixed effects.

BEA data on net capital stock of private non-residential fixed assets is used for information on ICT-related variables (Stiroh, 2002). The data is given at NAICS industry levels (74 industries with 2, 3 or 4-digit NAICS aggregation) for the period of 1947-2011 by many types of assets. Among other types of assets, I identify those belonging to the information and communication category, as in Stiroh (2002). Computer hardware is classified into mainframes, PCs, printers, terminals, tape drives, storage devices and system integrators, while software is classified into prepackaged software, custom software and own-account software. In addition, communications equipment is listed as a separate category. Variables are given both in current cost and chain-type quantity indices to account for changing quality of ICT goods over time. I use asset-specific NIPA price indices over time to deflate current cost numbers². To aggregate different ICT assets into one measure of ICT stock, I use a Fisher-type aggregation for chain-type quantity indices and I use the Tornquist procedure to aggregate the price indices.

To connect the two datasets, I use various concordances that allow me to convert industry-level variables into patents' technology-class-level variables and vice versa. Appendix B contains a description of existing concordances and the details on how I extend them to my data. At the end, we are left with industry and technology-level datasets for the sample period of 1976-2003.

A period in the beginning of our sample can be loosely referred to as the time before the IT revolution. By comparing real stock of industry-level ICT capital in 1976-1978 to that in 2001-2003, we clearly see a huge rise. The ICT quantity index for a median industry increased 57 times between these two periods while the ratio of real ICT stock to real stock of equipment and structures increased 13 times. Appendix C shows the figures illustrating these changes.

Data on industry-level value added, output and employment come from the BEA.

¹The average lag between the patent application and the grant year is between 2 and 3 years.

²To address a big increase in the quality of goods over time, Gordon (1990) constructed price series with hedonic adjustments for many different assets. However, these series end in 1983. By comparing NIPA prices to Gordon's price series, Cummins and Violante (2002) extrapolated the bias in NIPA prices and extended quality-adjusted price series to 2000. They notice a large bias in NIPA prices for some assets but fortunately NIPA price indices for IT-related assets do not show a significant bias (because BEA was using hedonic adjustments for them).

2.2 Empirical Facts

Fact 1: *There has been an increase in knowledge diffusion over time. This increase is correlated with higher usage of ICT.*

A good way to study technological interconnections and the flow of ideas between different classes of technology is to utilize information embedded in patent citations. Patents describing new inventions are linked to previous patents by citing those that have impacted creation of new knowledge. Thus citations constitute good proxies for knowledge flows between different technologies. It has been a well-established practice in the patent literature to use citations as a measure of knowledge diffusion and learning (Jaffe, Trajtenberg, and Henderson, 1993)³.

To illustrate knowledge diffusion using citation information, consider a knowledge production process of a particular class of technology. In this process, inputs will include ideas generated by earlier patents from other fields. Outputs will be new knowledge in a technology class (proxied by new patents). We can then illustrate this process using input-output tables. Each row of this table corresponds to a particular technology class and shows the composition of citations made by its patents to all technology fields. I am interested in seeing how the input-output tables change over time. Because this in essence represents a comparison of two contingency tables, I standardize them to make them more comparable (Mosteller, 1968; Greenwood, Guner, Kocharkov, and Santos, 2014). This standardization takes into account that distribution of existing patents across technology fields changes over time. In essence, the process involves iteratively adjusting distributions in the tables so that eventually both have the same marginal distributions (see Appendix C for the details on this procedure).

Figure 1 illustrates resulting input-output matrices for the 1976-1978 and 2001-2003 periods. Each row represents one of the 37 technology categories⁴ and each cell depicts a

³There has been a steady increase in the number of citations over time. This may be mainly due to three factors: First, in addition to the citations listed by patent applicants, patent examiners search for additional relevant references in the records. In the 1980s, USPTO computerized its records, which may have made examiners more efficient in finding relevant prior art, so the number of citations made could have accordingly increased. Second, over time the pool of available patents for potential citation increases. This may mechanically increase the number of citations made by patents. The third factor, the focus of this paper, is the extent to which real knowledge diffusion increases over time. We need to ensure the third effect is properly separated from the other two. In order for the first two effects not to confound my analysis, I control for year of the citing patent and existing patent pool, as appropriate. In addition, patent technological classes are used to control for systematic differences in citation behavior between technological classes. I use these steps in the regression analysis in Appendix C.

⁴Here, for the sake of better visualization I use the patent classification into 37 broad technological categories (*subcat*). In Appendix C, I also depict a similar table but for more detailed technology categories




Figure 1: Citations Input-Output Matrix

Notes: The figure illustrates citations standardized input-output matrices for the periods 1976-1978 and 2001-2003. Each row represents one of the 37 technology categories and each cell depicts a share of citations given by citing technology class to corresponding cited technology class. Higher shares are depicted in darker colors. All shares in each row add up to one. Cited technology classes on the horizontal axis are ranked by citation shares received.

share of citations given by citing technology class to corresponding cited technology class. Higher shares are depicted in darker colors. All shares in each row add up to one. Cited technology classes on the horizontal axis are ranked by citation shares received. Because the number of citations to own technology class is always largest, the citing and cited classes in the leftmost cells coincide. As seen from Figure 1, technologies rely on knowledge from external classes in a significant way. External dependence also varies a great deal across technology fields – both the share of self-citations (darkness of cells in the left column) and the number of external technology fields cited vary extensively by technologies.

Comparison of input-output tables across time illustrates a clear shift toward more learning from external fields. There is more citation of new fields in patents in 2001-2003 compared to 1976-1978 (almost no white cells in later period). In addition, as can be seen from the shading of the cells, external fields became more heavily cited; during this period, the average share of self-citations decreased from 65% to 55%⁵. The increase in newly-cited fields is more drastic if we look at the similar table for finer technology categories (412 *nclass* fields). On average, 8% of other technology classes were cited by patents in the earlier period, while 35% of other technology classes were cited in the later period. The share of total citations

(*nclass*).

⁵Notice that I consider only technology classes that are consistently represented during the entire sample period and omit classes that enter the data later.

from the same technology sector decreased from 58% to 43% (see Appendix C for details).

These increasing external learning trends are correlated with ICT expansion. Using technology-class variation in real ICT capital stock and controlling for year and industry dummies along with other characteristics, we observe significant positive correlation between industry-level ICT usage and external learning. In particular, we can infer from Table 6 in Appendix C that controlling for total number of firms, patent stock and total fixed capital stock in an industry, one standard deviation increase in log ICT stock is associated with 5 percentage point increase in the percent of external citations and 20% increase in expected counts of external classes cited⁶.

Fact2: *There is large heterogeneity in external knowledge dependence among technology classes.*

To understand how an increase in overall knowledge flows in the economy affects industries, it is important to consider variability in the benefits that different industries realize from external knowledge acquisition. From observing input-output matrices in Figure 13, we can already notice differences between industries in their external knowledge dependence. Citations of patents in some industries primarily consist of citations within their own technology class, while others credit a large proportion of citations to patents from other fields, implying that those industries gain more significant knowledge from external fields.

To define a measure of technology-specific external knowledge dependence, I use a detailed industry classification breakdown into 412 technology classes (*nclass*) from patent data. For each technology class, consider external knowledge dependence as the fundamental characteristic of its innovation process, which is constant over time. Because of limited diffusion of knowledge from other classes in the early periods studied, observed citation patterns early in the sample do not reveal full information on real technology-specific external knowledge dependence. However, assuming that the diffusion process had become nearly perfect by the end of the sample⁷, we can learn about the time-invariant knowledge dependence of each technology class by looking at observed backward citation patterns of patents in 2001-2003.

Define external knowledge dependence of a technology class j as the number of different technology classes cited by all patents in a technology class j over the period 2001-2003,

⁶For the share of external citations: $1.19 \times 0.042 = 0.05$, for the number of external classes cited: $\exp(1.19 \times 0.151) = 1.197$.

⁷Evidence in support of this assumption will be provided in the calibration section of the paper.




Figure 2: Distribution of External Knowledge Dependence

Notes: Distribution of the measure of external knowledge dependence for 412 technology classes (n_{class}) from NBER patent data. External knowledge dependence of a technology class j is the number of different technology classes cited by all patents in a technology class j over the period 2001-2003, divided by the total number of technology classes. Names of some technology classes with various degrees of external knowledge dependence are given.

divided by the total number of technology classes. Formally, the measure is

$$\frac{(\# \text{ classes cited})_{\forall \text{ patent } i \in j}}{N} \quad (1)$$

where i indexes a patent, j is a technology class and N is a number of all technology classes. The higher this index, the more externally dependent is an industry⁸.

Figure 2 shows the distribution of technology classes based on their values of external knowledge dependence. Specific examples of some industries are illustrated as well. According to this measure, high-tech industries show higher levels of external knowledge dependence than do low-tech industries, which generally appear more self-contained.

⁸It may be interesting to compare this index to alternative measures of external knowledge dependence. For that, I compare the Spearman's correlations between the rankings implied by this measure and alternative measures. If a similar measure is defined based on the whole sample period, the Spearman's correlation is 0.98; if an alternative measure is defined on the sample of 1976-1978, the Spearman's correlation is 0.81. Another alternative measure to compare is $\frac{\sum_{\text{patent } i \in j} \# \text{ classes cited}_i}{N \times N_j}$ in 2001-2003 (with N_j being a number of patents in a technology class). This measure is similar to our benchmark measure but takes an average share among the patents instead of pooling all cited classes together. Spearman's correlation in this case is 0.52.

Fact3: *There has been a sectoral reallocation of innovation and production toward more externally dependent industries.*

Consider the dynamics of industries characterized by different external knowledge dependence. Figure 3 shows time trends in innovation and real activities for industries in the top and bottom 25% of the external knowledge dependence distribution. The left panel plots regression coefficients from the OLS regression of log number of patents in a technology class on year and technology class dummies. Coefficients on year dummies are reported (with a base level of 1976). We see an initial decline in innovation activities everywhere, followed by an increase in innovation, but this increase is much more pronounced and prolonged in industries that are more externally dependent.




Figure 3: Sectoral Reallocation

Notes: The left panel: OLS regression of yearly log patent stock of a technology class (n_{class}) on year dummies and industry dummies. Each dot represents coefficients on corresponding year dummies. The omitted group is 1976, so coefficients compare to the base level of zero in 1976. The bands around the point estimates represent 95% confidence intervals. The right panel: Evolution of shares of value added in total GDP over time. Data on value added for NAICS industries comes from BEA. Industries in the top 25% of the measure of external dependence are colored in red, and the bottom 25% industries, in blue. External knowledge dependence of a technology class j is the number of different technology classes cited by all patents in a technology class j over the period 2001-2003, divided by the total number of technology classes. Corresponding concordances are used to extend this classification to NAICS industries.

The right panel of Figure 3 reviews industrial activities measured by their share of value added in total GDP⁹. We observe that the share of value added in total GDP for more externally dependent industries is increasing, while it is decreasing for more self-contained industries. The trend is even more pronounced if we review a similar comparison between the top and bottom 10% of the same distribution of industries.

⁹I extend the measure of external knowledge dependence to NAICS industries using concordances, as discussed in Appendix B.

In Appendix C, I provide some robustness checks for these trends. Specifically, I show that sectoral reallocation, i.e., relative expansion of more externally dependent sectors - persists even if I 1) drop IT-related sectors that boomed over that period because of the huge technological advances, and 2) compare only industries with broadly similar periods of origin, ruling out industry maturity and/or cyclicalities as key drivers.

Fact4: *More externally dependent industries experienced higher knowledge spillovers from ICT.*

From the previous analysis, we know that knowledge diffusion increased over time and higher ICT is associated with more technological spillovers. Here, I test whether higher ICT is associated with relatively higher technological spillovers for classes of technologies that are more externally dependent. I use the shares of citations made to external classes and numbers of technology classes cited as proxies for technological spillovers and knowledge diffusion.

Table 1 shows the results from panel regressions of the two measures of technological spillovers over time as a function of ICT, and the interaction of ICT with a measure of external dependence – here, a categorical variable that corresponds to quartiles of the external knowledge dependence distribution. Regressions control for year and technology class dummies, number of citation-adjusted patents in a technology class per year, number of firms (assignees) that patent in a class, and real stock of capital from BEA data. The results show that the interaction terms of external dependence with ICT are positive and significant, meaning technological spillovers from ICT are more important for more externally dependent technology classes. Interpretation of the magnitudes is the following: all else equal, if log ICT stock increases by one standard deviation, it increases external citations share of the top-quartile industries by 10 percentage points more compared to the industries in the bottom quartile. The same change increases expected count of external classes cited for all types of industries, but by 16% more for the top-quartile industries than for the bottom-quartile ones¹⁰.

Fact5: *ICT is associated with higher levels of various measures of industry-level competition.*

I consider the relationship between ICT and various measures of competition within technology classes. The first set of measures relates to the number of entrants. Existing firms can be challenged both by incumbents operating in other technology classes and entering different classes and by completely new entrant firms. From the data, I construct the number of firms

¹⁰For the top quartile, the effect on the share of citations is $0.028 \times 3 \times 1.19 = 0.094$, while for the bottom quartile, the effect is $-0.008 \times 1.19 = -0.009$. The difference is 0.103. For the number of classes cited, I compare the effect for the top quartile, $\exp(1.19 \times 0.062) = 0.073$, to the effect for the bottom quartile, $\exp(1.19 \times (0.062 + 0.045 \times 3)) = 0.234$. The difference is 0.16.

Table 1: ICT and Technological Spillovers

	Num. external classes cited	External citations share
Log ICT	0.062*** (0.004)	-0.008** (0.004)
Log ICT $\times$ External dependence	0.045*** (0.007)	0.028*** (0.003)
Year	Yes	Yes
Technology class	Yes	Yes
Observations	11286	11286
R^2	0.89	0.88

Notes: Industry (n_{class}) $\times$ Year panel regressions. Additional controls are numbers of citation-adjusted patents in a technology class per year, number of firms (assignees) patenting in a class of technology and real capital stock from BEA data. *External dependence* is a categorical variable for quartiles of the measures of external knowledge-dependence distribution, the highest corresponding to the top 25%. The dependent variable in the first column is the number of different external technology classes cited by patents in a given class of technology per year (Poisson regression). The dependent variable in the second column is the share of external-class citations made by patents in a given technology class per year (OLS regression). Standard errors are clustered by technology classes. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$

entering a particular technology class in a year. I distinguish between first-time entrants and entrants that have patented in some other technology fields before. To see whether a firm patents in a technology class for the first time, I need to observe the firm's patenting history. Because it is hard to gain this for the early years of our sample, for this analysis only, I focus on a truncated 1979-2003 sample.

Columns 1 and 2 of Table 2 show the results of Poisson regressions of the two types of entry on ICT. In addition to industry and year effects, I also control for number of existing incumbents in a technology class (number of firms that have patented in a technology class within the last five years). Both columns show that the increase in log real ICT stock is correlated with more entry into technology classes. Results are very similar if, instead of current ICT stock, I use lagged ICT stock, thus diminishing potential concerns about reverse causality in this regression.

Another competition measure I look at is a Herfindahl index of patenting concentration, calculated as a standard Herfindahl index but based on shares of patents contributed by firms to a particular technology class within the last 5 years. A higher index for a technology class means a higher concentration of patenting in fewer firms. Column 3 of Table 2 shows that

Table 2: ICT and Competition

	Number of entrants (new)	Number of entrants (other incumbents)	Log (Herf index)
Log ICT	0.361*** (0.020)	0.550*** (0.079)	-0.389*** (0.028)
Year	Yes	Yes	Yes
Technology class	Yes	Yes	Yes
Observations	9682	9682	11286
R^2	0.76	0.74	0.86

Notes: Industry ($nclass$) $\times$ Year panel regressions of various competition measures on log real ICT capital stock. Column 1: Dependent variable is the number of new entrant firms – number of firms entering the technology class that have not patented in any other field. Column 2: Dependent variable is number of firms entering the technology class and that have patented in other fields before. Column 3: Dependent variable is Herfindahl index of patenting concentration in a technology class. Columns 1 and 2 - Poisson regressions (years: 1979-2003), Column 3 - OLS (years: 1976-2003). All regressions include year, technology class fixed effects, and number of existing incumbents (for the last five years) in a technology class. Standard errors are clustered at the technology class level. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$

higher ICT is associated with greater dispersion of patenting activities among firms.

3 Model

Motivated by new empirical observations, I develop a general equilibrium model of endogenous technical change. Firms are engaged in innovation and are technologically heterogeneous with respect to external knowledge dependence. In the model, increase in ICT is associated with broader accessibility of external knowledge. This, in turn, leads to both heterogeneous knowledge spillovers across industries and increased creative destruction that renders existing technologies obsolete. The goal is to understand the contribution of the IT revolution to innovation, sectoral reallocation and growth. This section describes the details of the model.

3.1 Preferences and Final Good Technology

Time is continuous in this economy. There is a representative household that is endowed with one unit of labor that it supplies inelastically. The household seeks to maximize its

expected discounted lifetime utility of

$$U = \int_0^{\infty} \exp(-\rho t) \log C(t) dt,$$

where ρ is the household's discount rate and $C(t)$ is consumption of a final good $Y(t)$.

The final good is produced by combining a fixed measure of intermediate goods from a unit continuum of product lines:

$$\log Y_t = \int_0^1 \log y(j, t) dj, \quad (2)$$

where $y(j, t)$ is output produced by product line j at time t .

I assume perfect competition in the market for final good production and normalize the price of the final good to 1. I denote the price of an intermediate good in sector j at time t by $p(j, t)$. Profit-maximizing final good producers choose intermediate inputs to maximize:

$$\max_{\{y(j, t)\}_{j \in [0, 1]}} \left(\exp \left\{ \int_0^1 \log y(j, t) dj \right\} - p(j, t) y(j, t) \right)$$

This maximization leads to the unit-elastic demand of

$$y(j, t) = \frac{Y(t)}{p(j, t)}. \quad (3)$$

3.2 Intermediate Goods Market

Intermediate goods in each product line (the terms product line and industry will be used interchangeably) are produced using a linear technology utilizing labor input $l(j, t)$ scaled by the time-variant firm-specific composite labor productivity:

$$y(j, t) = q(j, t) f(ICT(j, t)) l(j, t), \quad (4)$$

The composite productivity consists of two terms: $q(j, t)$ is a quality in line j and it evolves endogenously, while $f(ICT(j, t))$ is a function of ICT usage in product line j . The evolution of the first term will capture an indirect effect of ICT through its impact on quality-improving investments, while the second term is intended to capture the direct investment-specific productivity gain from ICT. The focus of the model is on the evolution of the endogenous term, $q(j, t)$. The evolution (and allocation across industries) of ICT is not the

focus of the model and is taken as exogenous.

Within a product line, assume that the same intermediate good can be produced by two firms competing a la Bertrand. Then a good is produced by the leading firm that has higher productivity. Index by i , a producing incumbent and by $-i$, the laggard firm. Then, $q_i(j, t)f(IT_i(j, t)) > q_{-i}(j, t)f(IT_{-i}(j, t))$. Limit pricing in Bertrand competition then implies a leader setting a price equal to the marginal cost of a follower¹¹. Thus, the price of an intermediate good in industry j is

$$p(j, t) = \frac{w(t)}{q_{-i}(j, t)f(IT_{-i}(j, t))}, \quad (5)$$

where $w(t)$ denotes an equilibrium wage rate in the economy.

As a result of the demand curve in equation (3) and the price in equation (5), instantaneous profit of an intermediate goods producer is

$$\Pi_i(j, t) = \left(1 - \frac{q_{-i}(j, t)}{q_i(j, t)} \frac{f(ICT_{-i}(j, t))}{f(ICT_i(j, t))}\right) Y(t) \quad (6)$$

We see that the profits of a firm are scaled with the total output in the economy (market size effect) and they depend only on the ratio of current leading technology and the existing technology base in the industry. Hence the goal of the firms is to increase profits by widening this technology gap. Given that I take the ICT process to be exogenous, it is not restricting to assume that ICT availability for the firms within the same product line is the same, i.e., $ICT_{-i}(j, t) = ICT_i(j, t)$. This does not alter the innovation decisions in an important way and it brings cleaner analytical results.

Evolution of Productivity. Product quality $q_i(j, t)$ in each product line j evolves stochastically as a result of R&D efforts of the firms (described below). If the innovation process in a product line j is successful within a small time interval Δt , it brings an improvement of the previous productivity by a fixed step size λ , with $\lambda > 1$:

$$q_i(j, t + \Delta t) = \lambda q_i(j, t)$$

Denote the productivity of a laggard firm in the industry j by $q_{-i}(j, t) = \lambda^{n_{-ij}}$ and the productivity of an incumbent by $q_i(j, t) = \lambda^{n_{ij}}$. Then denote the number of step improvements made by an incumbent relative to the laggard in this product line by $n_j(t) \equiv$

¹¹We can also interpret this structure as the pricing decision of a firm with some competitive fringe that is able to produce at some base level of technology freely accessible to everyone.

$n_{ij}(t) - n_{-ij}(t)$. I will refer to $n_j(t)$ as industry j 's technology gap. As a result of successful innovation by the leading firm, the gap in the industry increases by one: $n_j(t+\Delta t) = n_j(t)+1$. Going back to equation (6), I rewrite the incumbent's profits as

$$\Pi_i(j, t) = (1 - \lambda^{-n})Y(t) \quad (7)$$

In what follows, I drop a subscript of a leading firm i and refer to its productivity as the productivity of its product line.

3.3 Research and Development with Industry Heterogeneity

Having described the production process for intermediate goods, I next describe the process of research and development. Knowledge production takes as inputs the knowledge created in external product lines. Importantly, motivated by observations from the data, I assume that the extent of external knowledge dependence varies across industries. Here is how I model this.

Knowledge Dependence Structure. Consider a continuum of industries located on a technological circle of a unit length (see Figure 4). Each point on the circle corresponds to some product line. As noted, because of the Bertrand competition, each product line is occupied by a single incumbent. In addition, each firm will own only one product line (more detailed discussion of this follows later).

Product lines differ by their intrinsic degree of external knowledge dependence. A product line j has a product-line-specific external knowledge dependence distribution with density f_j . This density is distributed on the $(j - \frac{\delta_j}{2}, j + \frac{\delta_j}{2})$ arc of a unit circle. This means that a firm learns the ideas in the δ_j -vicinity of its own technological location on the circle. The distribution is centered around its own product line and puts smaller weights on the technologies that are further away. In this sense, one can think of points located nearby as technologies with more interdependent research problems. The area under the density f_j integrates to one.

An example of industries with different technological parameters δ is illustrated on Figure 4. Here, product line i relies more on external knowledge than does product line j . In other words, a wider span of external technologies is used in the knowledge production, ($\delta_i > \delta_j$).

Diffusion of Knowledge. In the economy, there are frictions in accessibility of external knowledge. Though the industries rely technologically on potential knowledge from the whole δ -vicinity of their location, knowledge does not perfectly disseminate in the economy and new advances in each product line are not accessible to everyone. Instead, assume there is a parameter ε_t that describes the knowledge diffusion process and that varies over time. In particular, assume that the firm at location j is able to observe innovations only on the $(j - \frac{\varepsilon_t}{2}, j + \frac{\varepsilon_t}{2})$ arc. In Figure 4, shaded blue areas correspond to the shares of accessible knowledge for two industries. Denote that area by $\bar{\mathcal{E}}(j, \varepsilon_t) = \int_{j - \frac{\varepsilon_t}{2}}^{j + \frac{\varepsilon_t}{2}} f_j dj$. Notice that it depends both on technological characteristics of the product line, δ , and the current state of knowledge diffusion, ε_t .




Figure 4: Technological Circle and External Knowledge Dependence

For example, under the assumption of a triangular distribution for external knowledge dependence, we get:

$$\bar{\mathcal{E}}(j, \varepsilon_t) = \begin{cases} \frac{2\delta_j \varepsilon_t - \varepsilon_t^2}{\delta_j^2}, & \text{if } \delta_j > \varepsilon_t \\ 1, & \text{if } \delta_j \leq \varepsilon_t \end{cases} \quad (8)$$

In Figure 4, a firm in the product line j utilizes the whole amount of necessary knowledge, while industry i is lacking a significant share of the useful knowledge that it needs in the innovation process. Diffusion parameter ε_t is exogenous and time-specific. It is assumed that ε_t increases with higher ICT, which is consistent with empirical observation from Section 2.

This distinction between ideas that are usable versus observable has some similarity with Caballero and Jaffe (1993). In their set-up, though, availability of knowledge is affected by the time dimension, while here it is affected by the cross-sectional dimension – firms first observe ideas from “closer” technologies and there are frictions in learning more “distant” ideas.

The next section describes how the areas of accessible knowledge affect the efficiency of R&D and how competition enters the model.

Research and Development. Firms undertake two types of research activities - vertical and horizontal R&D. Vertical R&D is aimed at improving the qualities of goods in one’s own product line, while horizontal R&D is directed toward the improvement of qualities in other product lines. If a firm in industry j is successful in horizontal innovation, it will end up with an innovation that improves a quality of some random product line in ε_t -vicinity of j , $(j - \frac{\varepsilon_t}{2}, j + \frac{\varepsilon_t}{2})$. This means that a firm can only improve upon the product lines that are in its information set.

The two types of research efforts are the firm’s choice variables. Denote a chosen Poisson arrival rate of horizontal and vertical innovation by x_j and z_j , respectively. Industries differ by the R&D technology. Assume that a continuum of industries is classified into a finite number of S types of industries. A firm of type $s \in \{1...S\}$ with a technology gap n needs to invest $C_{s,t}^{vert}$ resources in terms of final output to achieve x_j rate of vertical innovation, where

$$C_{s,t}^{vert} = \frac{\alpha_s}{\bar{\mathcal{E}}(s, \varepsilon_t)^\phi} x^\gamma \lambda^{-n} \quad (9)$$

and needs to invest $C_{s,t}^{horiz}$ to achieve a vertical innovation rate z_j :

$$C_{s,t}^{horiz} = \frac{\alpha_s \beta}{\bar{\mathcal{E}}(s, \varepsilon_t)^\psi} z^\gamma. \quad (10)$$

The cost efficiency of a firm of type s at time t comes from the combination of two factors. The first is a type-specific productivity α_s and the second is a current share of accessible knowledge to the firm, $\bar{\mathcal{E}}(s, \varepsilon_t)$. Though an industry may have the highest intrinsic productivity α_s , if the amount of knowledge unobservable to it is high enough, that industry may not necessarily show up as being the most cost-efficient. The degree of this trade-off is allowed to vary for vertical and horizontal innovation production functions. The parameter γ is a common curvature in cost functions. The λ^{-n_j} term in the vertical cost function

is similar to the assumption of knowledge capital in Klette and Kortum (2004) and Peters (2013): Firms with greater technological leadership have an advantage in upgrading the productivity of their own product lines. As we will see later, this assumption will also facilitate the derivation of some analytical results.

3.4 Innovation and Competition

To summarize the innovation dynamics, in the small interval of time Δt the following can happen:

$$\text{Innovation by the firm at } j \rightarrow \begin{cases} q(j, t + \Delta t) = \lambda q(j, t) & \text{with prob. } x_j \\ q(i, t + \Delta t) = \lambda q(i, t), \quad i \in (j - \frac{\varepsilon_t}{2}, j + \frac{\varepsilon_t}{2}) & \text{with prob. } z_j \end{cases}$$

If vertical innovation is successful, the incumbent firm advances its current technology gap, n , one step further. This increases the firm's mark-ups and profits (see equation (7)). On the other hand, if the firm is successful at horizontal innovation in some random product line i , it will creatively destroy the incumbent, driving it to exit. In this case, the technology gap in the i product line is re-set to one (the best technology that has been employed by a previous incumbent in i becomes public knowledge). I assume that the firm j does not keep that product line but sells it to an outside entrant for price p , which is determined by Nash bargaining with all the bargaining power given to the seller, and by the expected value of a new product line in equilibrium¹².

The distinction between the two types of innovation that firms undertake has been used in Akcigit and Kerr (2010) and Atkinson and Burstein (2010). In this set-up, it introduces competition into the model. In particular, the assumption of horizontal innovation captures the idea of creative destruction: firms from other product lines replace existing incumbents by improving upon their technologies. Denote by τ_{jt} the total competition that a firm in product line j faces from the horizontal innovations of incumbents in other product lines. Then it can be expressed as

¹²Alternatively, one could think of different stories behind this assumption: The firm innovates in a new product line and establishes a new lab that acts as a separate entity, or an existing incumbent in i , instead of exiting the economy, buys out a new innovation from a firm in j . All these stories are equivalent in terms of implications for model's dynamics. This assumption is made for tractability reasons. If the firm were allowed to acquire other product lines, it would also gradually expand its information set, making the problem intractable. Though this assumption assumes away the possibility of strategic innovations, the firm still has the full incentive for horizontal innovation as it gets the whole value of that product line.

$$\tau_{jt} = \int_{j-\frac{\varepsilon_t}{2}}^{j+\frac{\varepsilon_t}{2}} \frac{z_{it}}{\varepsilon_t} di$$

The reason for this expression is the following: the product line j faces creative destruction from all firms in ε_t -vicinity (those product lines whose information set contains j industry). The rate of competition induced by each product line i is equal to its horizontal innovation rate, z_{it} . In turn, this intensity should be divided by the measure of product lines in which that horizontal innovation may randomly be realized, ε_t .

In what follows, I assume there is an equal measure of different types of industries. In addition, to facilitate analysis, assume that product lines of different types are evenly distributed on the technological circle. This assumption on symmetry implies that any area on the circle has the same measure of product lines of different types. As a result, neither the creative destruction nor the innovation rates depend on a specific location of a product line on the technological circle. Thus, beginning with the following discussion, I drop the subscripts of the location j and $\tau_{jt} = \tau_t$.

As a result of competition and innovation dynamics, technology gaps characterizing each product line move up and down the ladder. Denote by $\mu_t(n, s)$ a measure of product lines with s type and n technology gap at time t . Clearly, $\sum_n \mu_t(n, s) = \frac{1}{S}$. Then at each point in time, equilibrium distribution over (n, s) is characterized by a set of differential equations:

$$\begin{aligned} \dot{\mu}_t(1, s) &= \tau_t(1/S - \mu_t(1, s)) - x_t(1, s)\mu_t(1, s), & \forall s \in \{1 \dots S\} \\ \dot{\mu}_t(n, s) &= -x_t(n, s)\mu_t(n, s) - \tau_t\mu_t(n, s) + x_t(n-1, s)\mu_t(n-1, s), & (11) \\ & \text{if } n \geq 2, \quad \forall s \in \{1 \dots S\} \end{aligned}$$

The interpretation of these inflow-outflow dynamics is straightforward. The outflow from the (n, s) state happens if firms in (n, s) either successfully innovate or are replaced by competitors. This explains the first two terms on the right-hand of the first differential equation. The third term shows the inflow to (n, s) state; it can be achieved by the successful vertical innovation of firms in $(n-1, s)$ state. The difference between inflow and outflow gives the change in the measure of the firms, $\dot{\mu}_t(n, s)$. The second equation has a similar interpretation.

3.5 Labor Market

Recall that there is one unit of labor supplied in each period. From equations (3), (4) and (5), the labor demand of firm in industry j is

$$l(j, t) = \frac{\frac{q_{-i}(j, t)}{q_i(j, t)}}{w_t} Y(t) = \frac{\lambda^{-n_j(t)}}{w_t} Y(t) \quad (12)$$

Hence, the labor market clearing requires

$$1 = Y_t \int \frac{\lambda^{-n_j(t)}}{w_t} dj \quad (13)$$

Consider $Q(t)$ to be an average composite quality in the economy, $\log Q(t) = \int \log q_j(t) f(ICT_j(t)) dj$. Then the resulting equilibrium wage can be expressed as (see Appendix A for the derivation).

$$w_t = Q(t) \lambda^{-\sum_{n,s} n \mu_t(n,s)} \quad (14)$$

3.6 Steady State and Value Functions

Consider the economy in stationary equilibrium, in that knowledge diffusion parameter ε , distribution and policy functions are constant, growth of aggregate composite ICT stock is constant, and the economy exhibits a constant growth rate g .

Each industry is described by n , the current technology gap in an industry, and by the area of accessible knowledge, $\bar{\mathcal{E}}(s, \varepsilon)$. The last is uniquely determined by the combination of δ and ε . Hence, the only state variables of the problem are n and s . Denote by $v(n, s)$ a value of a s -type firm with a technology gap n , normalized by output. Then, by using a household's standard Euler equation of $g = r - \rho$, the Bellman equation can be expressed as (see Appendix for the derivations):

$$\rho v(n, s) = \max_{x(n,s), z(n,s)} \left\{ \begin{aligned} & \pi(n) - \frac{\alpha_s}{\bar{\mathcal{E}}(s, \varepsilon)^\phi} x(n, s)^\gamma \lambda^{-n} \\ & - \frac{\alpha_s \beta}{\bar{\mathcal{E}}(s, \varepsilon)^\psi} z(n, s)^\gamma \\ & + x(n, s)(v(n+1, s) - v(n, s)) \\ & + z(n, s)p - \tau v(n, s) \end{aligned} \right\}, \quad (15)$$

The first term on the right-hand side is instantaneous profit, and the second and third are,

respectively, the cost of horizontal and vertical R&D. Next, with the flow rate of $x(n, s)$, vertical innovation is successful and the firm advances one step ahead; then the incremental value it is getting is $v(n + 1, s) - v(n, s)$. With the flow rate of $z(n, s)$, horizontal innovation is successful and the firm ends up innovating in an outside product line. Then the firm sells this product line to an outside entrepreneur for price p . And finally, with a flow rate of τ , creative destruction hits the firm's product line and the firm exits the market, getting an exit value of zero.

Proposition 3.1

i) The optimal R&D decisions of firms are independent of the technology gap n :

$x(n, s) = x(s)$ and $z(n, s) = z(s)$.

ii) The value function (15) of an s -type firm with a technology gap n can be expressed as

$$v(n, s) = A(s) - B(s)\lambda^{-n} \quad (16)$$

where $A(s)$ and $B(s)$ are defined by

$$A(s) = \frac{1 + z_s p - \frac{\alpha_s \beta}{\mathcal{E}(s, t)^\psi} z_s^\gamma}{\rho + \tau}, \quad \forall s \in \{1 \dots S\} \quad (17)$$

and

$$B(s) = \frac{1 + x_s^\gamma \frac{\alpha_s}{\mathcal{E}(s, t)^\phi}}{\rho + \tau + x_s \frac{\lambda - 1}{\lambda}}, \quad \forall s \in \{1 \dots S\} \quad (18)$$

where τ is the equilibrium aggregate horizontal innovation rate

$$\tau = \frac{\sum_{s=1}^I z_s}{S}, \quad (19)$$

and the price p of an expected acquired product line satisfies

$$pS = \sum_{s=1}^S A(s) - \sum_{s=1}^S \frac{B(s)}{\lambda}. \quad (20)$$

Proof See Appendix A.

The first part of the proposition states that this model displays proportionate growth, as in Gibrat's law. Intuitively, per-period profits are concave in n , while marginal cost of

innovation decreases with n , and scaling of the cost by λ^{-n} exactly balances these two effects. The second part of the proposition states that the value function is linearly separable into two parts. Those parts come from two distinct innovation opportunities that firms have. As is evident from equation (17), the first part of the value function corresponds to an option value of innovating on a new product line. Equation (18) tells us that the second part of a value function comes from the option value of improving on own product line.

Using the results from part ii) of Proposition 3.1, the stationary distribution of firms over the state space can easily be characterized as:

$$\mu(n, s) = \frac{1}{S} \left(\frac{x_s}{x_s + \tau} \right)^{n-1} \frac{\tau}{x_s + \tau} \quad (21)$$

Proposition 3.2 *Steady state growth rate can be expressed as*

$$g = \log \lambda \left(\frac{\sum_s x_s}{S} + \tau \right) + \frac{\Delta F(ICT)}{F(ICT)}, \quad (22)$$

where $F(ICT)$ stands for the direct effect of the composite ICT input: $\exp \int_0^1 \log f(ICT(j, t)) dj$

Proof See Appendix A.

As Proposition 3.2 makes clear, aggregate growth in this economy, on the one hand, comes from innovations in each product line – vertical improvements by incumbents or horizontal innovations by entrants. These innovation efforts increase productivity by step size λ . On the other hand, growth also comes from exogenous growth in ICT input.

Lastly, I summarize the equilibrium of the model as follows.

Definition (Steady-State Equilibrium) *Given the exogenous allocation of ICT capital across product lines and the parameter of the diffusion process, ε_t , an equilibrium of the economy consists of $\{x_s^*, z_s^*, \mu(n, s)^*, p^*(n, s), y^*(n, s), Y^*, w^*, \tau^*, g^*, r^*\}$ such that: (i) Aggregate output Y_t^* is given by (2); (ii) Intermediate goods prices and output $p_t^*(n, s), y^*(n, s)$ satisfy (3) and (5); (iii) Wage w^* clears the labor market in (13); (iv) Innovation decisions x_s^*, z_s^* maximize a firm's value in (15); (v) Equilibrium creative destruction τ^* is given by (19); (vi) Distribution $\mu^*(n, s)$ is given by (21); (vii) r^* satisfies the Euler equation; (viii) Growth g^* satisfies (22)*

3.7 Discussion

Here, I discuss in more detail the predictions of the model regarding the effects of increased ICT. Once the increased ICT widens the area of accessible knowledge ε_t , the cost of innovation drops, as can be seen from the specifications in (9) and (10). On the one hand, both vertical and horizontal innovation processes become more efficient (the magnitude of the gains depends on ϕ and ψ). An increase in horizontal innovations from other firms ultimately affects creative destruction, τ , faced by firms at any location on the technological circle. This reduces returns from R&D investments, thus decreasing firms innovation incentives. On the other hand, the technological spillovers to the firm depend on the degree of its external knowledge dependence. If its δ_s is high enough and ε_t is less than δ_s , the firm sees positive technological gains. However, these gains diminish with further expansion in knowledge diffusion as a marginal contribution to the area of accessible knowledge $\bar{\mathcal{E}}$ becomes lower. Once ε_t makes a firm technologically saturated, additional technological gains from that point on are zero. However, as there are other firms in the economy that still gain from additional expansion in knowledge diffusion (those with higher δ), overall competition in the economy still rises. Theoretically, this process is able to generate rich industry dynamics in that some industries may experience initial acceleration in productivity growth from ICT and negative or zero gains afterward, and others may see higher and prolonged gains. This process of sectoral reallocation continues until ICT stops expanding ε_t . Once this happens, new growth rates stabilize at a steady-state level and industries with higher intrinsic productivities α_s will be the ones growing fastest.

4 Quantitative Analysis

In order to quantitatively assess the effect of an increase in ICT on industry dynamics, growth and reallocation margins in the economy, I first calibrate the structural parameters of the model. In this section, I first lay out the solution algorithm and describe the calibration. I then show the results and quantitative experiments from the calibrated model.

4.1 Solution

Solution of the steady-state equilibrium reduces to solving a system of nonlinear equations as follows. Equations (17) and (18) solve for $A(s)$, $B(s)$, components of the value function;

first order conditions given by

$$\gamma x_s^{\gamma-1} \frac{\alpha_s}{\bar{\mathcal{E}}(s, \varepsilon_t)^\phi} - B(s) \frac{\lambda - 1}{\lambda} = 0, \quad \forall s \in \{1 \dots S\}$$

and

$$p = \gamma \frac{\alpha_s \beta}{\bar{\mathcal{E}}(s, t)^\psi} z_s^{\gamma-1}, \quad \forall s \in \{1 \dots S\}$$

solve for vertical and horizontal innovation rates x_s, z_s ; equation (19) solves for creative destruction rate τ , and equation (20) solves for the value of an acquired product line p . In fact, for any given τ and p , one could solve for industry-specific policies as functions of the area $\bar{\mathcal{E}}$, and then search for τ and p that would bring the whole system into equilibrium.

The solution outside the steady state proves not to be hard either. Consider the dynamics in which the rate of diffusion parameter ε changes exogenously over time. This model features instantaneous transition to policy functions corresponding to a new steady state after this type of exogenous change. The only variable that will not adjust instantaneously is the distribution of industries $\mu_t(n, s)$. However, because of the symmetry of the model and the results of Gibrat's law, this distribution does not affect policy functions. Hence, the system of equations outlined above will correspond to the equilibrium solution at any time, given a parameter for the knowledge diffusion process, ε_t . However, to compute allocations of some real variables over time, I need to calculate $\mu_t(n, s)$ distributions in transition. I can do so by solving the differential equations in (11). Because of this tractability, I start from solving the pre-IT-revolution steady state of the model, corresponding to the period of 1976, and then trace the evolution of the economy over time as ICT levels exogenously increase.

4.2 Calibration and the Choice of Parameters from Patent Data

In this section, I first describe a construction of type-specific distributions of knowledge dependence and an estimation of an exogenous process of knowledge diffusion from the data. Next, I turn to the calibration of structural parameters of the model.

4.2.1 External Knowledge Dependence

In the quantitative exercise, I consider the case of a triangular distribution of external knowledge dependence. Product lines are classified into S number of types according to δ . Though there is much wider industry heterogeneity in the data in terms of external knowledge dependence (see Figure 2), I distinguish between four types of external dependence in this

analysis, hence $S = 4$.

The main characteristic of a particular type of industry δ matches one to one the earlier defined measure of external knowledge dependence, as shown in (1). As in the model, this measure proxies for the span of technologies from which an industry learns. I divide the empirical distribution of external knowledge dependence (Figure 2) into four quartiles. Then δ_1 , δ_2 , δ_3 and δ_4 are assigned the average values of the measures of external knowledge dependence in each quartile, where $\delta_1 < \delta_2 < \delta_3 < \delta_4$.

4.2.2 Knowledge Diffusion

An exogenous time-varying process of knowledge diffusion is constructed from the data. The goal is to capture the effect of ICT on knowledge diffusion.

First, for each technology class in the data, define $\bar{\varepsilon}_{jt}$ to be the share of different technology classes cited by patents in j until time t ¹³. Through the lens of the model, $\bar{\varepsilon}_{jt}$ is a proxy for a current state of the span of knowledge accessibility for a technology class at a point in time. Then, to identify the influence of ICT on this variable, I estimate a linear regression of $\bar{\varepsilon}_{jt}$ on ICT-related variables and other controls:

$$\bar{\varepsilon}_{jt} = \beta_0 + \beta_1 ICT_{jt} + \beta_2 ICT_{jt}^2 + \beta_3 ICT_{jt}^{others} + controls + error_{jt}$$

In particular, I take own-class real ICT stock, ICT_{jt} , its square, ICT_{jt}^2 , and ICT_{jt}^{others} , that is, the total ICT stock in technology classes that have ever been cited by j . The latter captures the idea that diffusion of knowledge is also affected by information-related investments made by others. Other controls in the regression include number of patents in a technology class in a year, total capital stock, time trend and technology class dummies. Finally, the diffusion parameter ε_t over time is equal to the part of $\bar{\varepsilon}_{jt}$ that is attributable to the evolution of ICT, averaged across all classes j :

$$\varepsilon_t = mean_j(\hat{\beta}_0 + \hat{\beta}_1 ICT_{jt} + \hat{\beta}_2 ICT_{jt}^2 + \hat{\beta}_3 ICT_{jt}^{others})$$

Figure 5 illustrates the derived ε_t . The diffusion parameter increases sharply in the mid-1990s, when the economy saw drastic growth in investment in new and improved ICT goods. It stabilizes around 2000, as the data shows close to zero growth in real ICT stock after the

¹³One could also look at the share of different technology classes cited by patents in j at time t . This definition produces very similar results but it is reasonable to assume that $\bar{\varepsilon}_{jt}$ is weakly increasing over time, which may not necessarily be the case under this alternative specification.




Figure 5: Knowledge Diffusion over Time

collapse of the dot-com bubble.

4.2.3 Calibration

There are four types of industries, which implies estimating four parameters $\alpha_1, \alpha_2, \alpha_3, \alpha_4$. Next, the parameters from the horizontal and vertical R&D cost functions $\phi, \psi, \beta, \gamma$ and the step size of innovation λ need to be identified. I adopt the following specification of the function governing the direct productivity impact from ICT: $f(ICT) = ICT^\kappa$ ¹⁴. This adds one more parameter, κ . As a result, the model has 10 parameters in total: $\alpha_1, \alpha_2, \alpha_3, \alpha_4, \beta, \phi, \psi, \gamma, \lambda, \kappa$.

I identify the parameter κ using reduced-form regressions from the data. This parameter is responsible for the direct benefits coming from the ICT investments, such as automation of production processes, improvements in customer-supplier links and others. Empirical literature studying impact of these benefits on productivity growth has highlighted reasons why these benefits might vary across firms. These include varying management and workplace practices and complementary investments (Brynjolfsson and Hitt, 2003, Loveman, 1994, Black and Lynch, 2001). In this exercise, I abstract from these heterogeneities and focus on the average productivity impact.

I transform the production function of intermediaries from the model (equation (4)) to

¹⁴Notice that function can be scaled by any arbitrary positive number that is a free parameter. Here, the coefficient is normalized to one.

get the equation describing growth of labor productivity in a product line:

$$gLP_j = \kappa gICT_j + gq_j$$

In the model, gq_j results from innovations. In the data, I capture innovations by observing the issuance of patents. Thus, I take this equation to the data and estimate labor productivity growth as a function of ICT growth, controlling for growth in patents and other controls. As available data on ICT stock is at the industry level, I estimate the following industry-level equation:

$$gLP_{jt} = \hat{\kappa} gICT_{jt} + gPatents_{jt} + gK_{jt} + \gamma_j + \delta_t + error_{jt},$$

where $gPatents_{jt}$ is the yearly growth of citations-adjusted patent stock and gK_{jt} is the growth of total real capital stock. The regression includes controls for industry and time dummies as well. An estimate of the elasticity κ is .02 with a standard error of .009. This elasticity is comparable to the values obtained in the firm-level study by Brynjolfsson and Hitt (2003), who find that the elasticity of computers is in the range of .01 to .04.

The rest of the nine parameters are calibrated to match important moments of the data in 1976-1978. Clearly, none of the moments pins down a particular parameter exactly – instead, each moment is related to a set of parameters that jointly determine the equilibrium values of all moments. However, we can still conceptually overview identification and the role of each moment in pinning down the parameters of interest.

The scaling parameter of the horizontal R&D cost β is directly related to the amount of horizontal innovation in the economy τ . To find a counterpart for τ in the data, I first describe how I construct the variables of internal and external innovations. Using the patent data, I identify the main technology class of a firm's operation. Next, I classify the firm's patent as an innovation stemming from vertical R&D if the patent belongs to the technology class of the firm's main operation. If the patent's technology class is different, the patent is classified as a horizontal innovation. Next, in order to find a counterpart for τ , I look at the yearly share of patents in a technology class issued by firms whose main operation is in a different technology class.

Parameters of industry-specific cost efficiency $\alpha_1, \alpha_2, \alpha_3, \alpha_4$ as well as curvature parameters ϕ and ψ , can be identified using the innovation ratios of different types of industries in different periods as well as the ratio of vertical to horizontal innovation rates. As before, I assign a main technology class of operation to each firm. To obtain average patenting of firms of a type s , I look at average patenting of all firms that belong to technology classes of

type s . For example, $\frac{\text{Innov}_4}{\text{Innov}_1}_{1976}$ denotes the ratio of average patenting intensity by firms of type $s = 4$ to that of $s = 1$ type firms in 1976. In addition, I use the ratio of vertical to horizontal patenting $\frac{\text{Verical innovation}}{\text{Horizontal innovation}}_{1976}$, constructed using the average firm-level patenting intensities as well.

Firms' R&D intensity can help pin down the elasticity of the R&D cost function. In order to obtain the firm's R&D data, I match the sample of firms from the patent data to the Compustat sample of public firms. I calculate average R&D intensity as an average value of the ratio of R&D expenditures to firms' sales from the matched sample.

Parameter λ is a step size of a productivity improvement brought about by a new innovation. Average growth of labor productivity is an important moment helping us to pin down this parameter. I take information on growth of labor productivity from BEA industry-level data. I focus only on the industries that correspond to the technology classes for which I have measures of external knowledge dependence (value added of these industries accounts for 78% of GDP in 1976-2003). As I match the moments to the 1976 data, I get a low value of productivity growth, 1.31%. According to equation (22), this growth comes from two sources: direct ICT investments and endogenous quality improvements from innovations. After subtracting the growth term coming from ICT, I match the residual growth rate from the data to one generated by the model with endogenous growth.

Table 3: Moment Values

Targets	Data	Model
Aggregate Growth ₁₉₇₆	1.31%	1.30%
Average Horizontal Innovation ₁₉₇₆	0.22	0.19
R&D Intensity ₁₉₇₆	0.09	0.05
$\frac{\text{Vertical Innovation}}{\text{Horizontal Innovation}}_{1976}$	1.20	1.24
$\left[\frac{\text{Innov}_4}{\text{Innov}_1}, \frac{\text{Innov}_3}{\text{Innov}_1}, \frac{\text{Innov}_2}{\text{Innov}_1}\right]_{1976}$	[2.25, 1.15, 1.53]	[2.31, 1.09, 1.53]
$\left[\frac{\text{Innov}_4}{\text{Innov}_1}, \frac{\text{Innov}_3}{\text{Innov}_1}, \frac{\text{Innov}_2}{\text{Innov}_1}\right]_{1977}$	[2.34, 1.12, 1.68]	[2.33, 1.05, 1.61]
$\left[\frac{\text{Innov}_4}{\text{Innov}_1}, \frac{\text{Innov}_3}{\text{Innov}_1}, \frac{\text{Innov}_2}{\text{Innov}_1}\right]_{1978}$	[2.56, 1.37, 1.85]	[2.38, 1.18, 1.72]

Having found the values of targeted moments, the calibration procedure involves finding the values for unknown structural parameters that minimize the Euclidean distance between

the targeted data moments and moments generated by the model.

This procedure delivers the following match in Table 3. The model performs quite well in matching the salient features of the data in 1976-1978. Calibrated parameter values are reported in Table 4. We see that intrinsic cost efficiency of the most externally dependent industries is highest, while the least externally dependent industries have the lowest efficiency. Curvature parameters ϕ and ψ imply high elasticities of innovation in the area of accessible external knowledge; implied elasticities for vertical and horizontal R&D costs are .83 and .70, respectively. Parameter λ is slightly lower than the estimated values from the literature (Acemoglu, Akcigit, Hanley, and Kerr, 2014) but that can be explained by the fact that the targeted growth rate in this exercise is also lower. The estimate of the curvature parameter γ is in the range of [1.67, 10] provided by Kortum (1993).

Table 4: Calibrated Parameters

Parameters	Meaning	Value
$[\alpha_1, \alpha_2, \alpha_3, \alpha_4]$	Scaling of vertical R&D	[2.02, 0.5, 0.52, 0.05]
β	Scaling of horizontal R&D	6.47
λ	Step size of innovation	1.03
ϕ	Vertical R&D curvature to external knowledge	2.48
ψ	Horizontal R&D curvature to external knowledge	2.10
γ	R&D cost curvature	2.98


4.3 Quantitative Experiments

Having calibrated the parameters of the model to match the 1976-1978 data, I can now quantitatively evaluate the effect of an increase in ICT on innovation, growth and sectoral reallocation over time. The experiment is to feed the time series of exogenously estimated diffusion parameter ε_t into the model and solve for firms' innovation policies over time. I compare model-generated trends in aggregate growth and sectoral reallocation to assess how much of the trends outlined in Section 2 can be attributed to the changes brought about by the IT revolution.

4.3.1 Sectoral Reallocation

Innovation Activities. To what extent is the knowledge diffusion mechanism conceptualized in this paper responsible for the overall increase in both innovation intensities and the gap between innovation levels of different types of industries? To answer this question, I calculate aggregate innovation rates of different types of industries given the estimated values of the diffusion parameter over time. Aggregate innovation in an industry comes from both vertical innovation by incumbents and horizontal innovation by entrants. Figure 6 reports these innovation rates for the most and least externally dependent industries ($s = 1$ and $s = 4$). Reported values are normalized such that total innovation for $s = 1$ in 1976 is equal to one. The same figure illustrates the dynamics of innovation in the data. In particular, I plot the average patenting in the technology classes that comprise the bottom and top quartiles of external knowledge dependence, normalized to match innovation rates from the model in 1976.

Figure 6: Sectoral Reallocation. Innovation



Overall, the rise in ICT is responsible for a significant increase in innovation intensities over time. Consistent with the data, this increase is most pronounced in most externally dependent industries. In the data, the ratio of average patenting of top quartile industries to that of the lowest quartile increased 2.35 times from 1976-1978 to 2001-2003. The same statistic in the model equals 1.75. Hence, the model can explain 74% of the reallocation of innovation activities observed in the data. This reallocation is a result of a combination of heterogeneous knowledge spillovers and competitive effects implied by increased ICT.

Production Activities. In the empirical section, we have observed that production in the economy has become more concentrated in more externally dependent industries. How much of this reallocation can the mechanism in the paper explain?

From the model, I derive the expression for the share of total output accounted for by a particular type of industries s :

$$Share_t(s) = \frac{Q_t(s) ICT_t^\eta(s) \sum_n \lambda^{-n} \mu(n, s)}{\sum_s Q_t(s) ICT_t^\eta(s) \sum_n \lambda^{-n} \mu(n, s)}, \quad (23)$$

where $Q_t(s)$ is the average quality index for industries of type s and $ICT_t(s)$ is their t -period composite ICT stock. Distribution $\mu(n, s)$ over time follows equation (11). As can be seen, share of output of s -type industries is higher if s -type industries have higher composite quality index. At the same time, if the industries' average value of markups, n , is higher, they demand less labor and contribute less to the overall production. Clearly, share also decreases as the quality improvements by other firms increase.

Figure 7: Sectoral Reallocation. Value Added




Figure 7 illustrates the reallocation of the share of value added in GDP for industries in the top and bottom quartiles of external knowledge dependence in the data, along with its counterpart from the model. The initial levels of productivity for different types of industries in 1976 are calculated by matching their shares of total output from the model to the shares of value added in GDP from the data in 1976¹⁵. As seen, the mechanism in the model


¹⁵Notice that, in the data, the sum of value added of our industries does not equal total GDP. For consistency, I normalize the shares of value added in the data so that they add up to one.

explains essentially the entire reallocation of production toward more externally dependent industries over time.

4.3.2 Aggregate Growth. Decomposition of ICT Impact

Next, consider the model's implications for aggregate growth. Aggregate growth from the model is calculated using equation (22). Figure 8 compares growth rates implied by the model with the average growth rate of labor productivity from the industry-level data. The growth series from the data are very volatile but there is an increasing trend. The average yearly growth rate in 1990-2003 is 45% higher than the average growth rate in 1976-1990. In the model, the same statistic is equal to 34%. Thus, the model can explain 75.6% of an increase in average growth of the economy after 1990.


Figure 8: Growth of Labor Productivity. Data and Model



The aggregate growth effect from the IT revolution can be decomposed into contributions from direct and indirect channels. To this end, I compare the following counterfactuals. First, what would be the growth rate of labor productivity over time if there were no growth in ICT? This thought experiment corresponds to an economy with a constant baseline steady-state growth rate, as in 1976. In Figure 9, this is a horizontal line at a 1.3% growth rate. Next, what would be the growth rate of labor productivity over time if growth in ICT had only a direct impact? In other words, this experiment disregards the role of ICT in knowledge diffusion and its implications for firms' innovation incentives. The red area in Figure 9 depicts the contribution coming from the direct channel. Finally, the blue area is residual growth,

which corresponds to the indirect growth contribution from ICT. This contribution comes from the changes in firms' innovation behavior implied by increased knowledge diffusion from ICT.

Figure 9: Decomposition of ICT Impact



As seen in Figure 9, ICT plays a significant role in the revival of growth, and both direct and indirect channels are quantitatively important. The average yearly contribution of ICT to labor productivity growth in 1976-1990 is equal to 28.7 basis points¹⁶, while, on average, ICT contributes 78 basis points to growth in 1990-2003. In the early period, a moderate 24% of this contribution is attributed to the indirect effect. However, in the later period, on average 70% of the impact results from the indirect effect resulting from knowledge diffusion.

The direct investment-specific impact of ICT on growth is short-run and largely depends on big technological improvements in ICT-producing sectors, as well as on industry-level ICT investment decisions (that are taken as being exogenous in this paper). For example, when the growth in real ICT investments declined dramatically after the collapse of the dot-com bubble in 2000, that led to a very small investment-specific growth effect in subsequent years.

¹⁶The empirical literature studying the impact of ICT on growth does not usually find significant gains from ICT in the 80s. This can be reconciled with my findings in the following way. When I estimate the parameter κ governing the indirect effect, I take an average (contemporaneous) estimate of ICT on growth obtained on the full sample period from 1976 to 2003. However, new investments in IT goods can have a lagged effect on growth because of a required period of learning (Yorukoglu, 1998, Greenwood and Yorukoglu, 1998). Brynjolfsson and Hitt (2003) also provide evidence for a lagged effect of ICT on firms' growth.

On the other hand, the indirect contribution from ICT remains high because it operates more through the level of ICT and not through its growth. The quantitatively large indirect growth contribution from ICT is a sustainable gain in the long run. This finding contrasts with some of the pessimistic views (Gordon, 2000) on transitory nature of productivity returns from the IT revolution.

4.3.3 Competitive Spillovers

Having quantified the growth and reallocation implications of indirect effect of ICT through knowledge diffusion, in this section I look more closely at the quantitative significance of technological and competitive spillovers created by knowledge diffusion.

In line with what is observed in the data, implied competition from the model increases over time. In the model, competition plays a dual role in growth. On the one hand, it has a negative role in disincentivizing firms' innovation because of the business stealing effect. On the other hand, competition has a positive effect on growth because new entry itself increases productivity. In this exercise, I evaluate the size of the negative competition effect.

Table 5: Growth in 2003. Three scenarios.

No growth in ICT	Total impact of ICT (direct+indirect)	ICT without competitive spillovers
1.16%	2.02%	2.60%

In the counterfactual experiment, I ask what innovation rates would be if we shut down the negative competition margin in the model. That is, I re-solve the model over time, setting the average horizontal innovation rate, τ , representing competition, at its initial level in 1976. I compare the model-implied growth rates in the last sample year, 2003, under three different scenarios, with 1) no growth in ICT, 2) benchmark effect from ICT that combines both direct and indirect channels, and 3) what the effect from ICT would be without the negative competition effect.

Table 5 reports the results. The scenario with no growth in ICT would imply a growth rate of 1.16%. The benchmark model with all combined effects from ICT implies a 2.02% growth rate, which constitutes a 74% increase from the first scenario. In the third scenario, where I shut down the competition margin, implied growth further increases by 28%. This suggests that an increase in knowledge flows and enhanced technological linkages in the economy over time come with negative growth implications: they are not as large as the positive

knowledge externalities, but still significantly dampen aggregate growth. This finding can partially explain why realized productivity returns from the IT revolution are not as high as a casual observer of drastic changes in ICT over time would expect.

5 Conclusion

There has been a tremendous rise in the quality and quantity of ICT during the past four decades in the U.S. Despite extensive empirical literature, the productivity implications of ICT investments are still debated. There are interesting open questions as to 1) why the productivity returns from ICT have been so heterogeneous among industries, 2) whether growth from the IT revolution can be sustainable, and 3) why the huge productivity improvements in ICT did not translate into larger gains in aggregate economic growth. This paper studies a new channel that emphasizes the role of ICT in improving the flow of ideas in the economy. I show that this channel is quantitatively large and contributed significantly to the acceleration of growth from 1976 to 2003. I also show that 1) knowledge diffusion matched with substantial technological heterogeneity across industries explains most of the widely heterogeneous ICT returns and at least 74% of the observed sectoral reallocation in the U.S., 2) the gains from knowledge diffusion induced by ICT are sustainable in the long run and contribute at least 65 basis points to aggregate growth, and 3) along with large knowledge spillovers, knowledge diffusion through ICT also creates negative competition effects, without which growth would increase by 28%.

This paper highlights the importance of the link between knowledge diffusion, technological heterogeneity of industries and reallocation. Clearly, the differences in the environments in which businesses operate put certain types of businesses in a more advantageous position. Taking this argument a step further, we can think about whether cross-country differences in knowledge diffusion could partially explain observed cross-country differences in specialization. While this paper explores the relationship between knowledge diffusion and ICT, many other factors influence knowledge flows. One example can be labor market regulations. Economies with stricter labor market regulations and lower employee mobility will also witness less exchange of information and a lesser degree of knowledge flows. The model in this paper can serve as a platform for analyzing additional questions about knowledge diffusion, technological heterogeneity, sectoral specialization and growth.

References

- [1] ACEMOGLU, D., AND AKCIGIT, U. Intellectual Property Rights Policy, Competition and Innovation. *Journal of the European Economic Association* 10, 1 (2012), 1–42.
- [2] ACEMOGLU, D., AKCIGIT, U., BLOOM, N., AND KERR, W. R. Innovation, Reallocation and Growth. NBER Working Paper 18993, National Bureau of Economic Research, Inc, 2013.
- [3] ACEMOGLU, D., AKCIGIT, U., HANLEY, D., AND KERR, W. Transition to Clean Technology. *University of Pennsylvania mimeo* (2014).
- [4] ACEMOGLU, D., AUTOR, D., DORN, D., HANSON, G. H., AND PRICE, B. Return of the Solow Paradox? IT, Productivity, and Employment in U.S. Manufacturing. NBER Working Paper 19837, National Bureau of Economic Research, Inc, 2014.
- [5] ACEMOGLU, D., AND GUERRIERI, V. Capital Deepening and Nonbalanced Economic Growth. *Journal of Political Economy* 116, 3 (2008), 467–498.
- [6] AGHION, P., AKCIGIT, U., AND HOWITT, P. What do We Learn from Schumpeterian Growth Theory? NBER Working Paper 18824, National Bureau of Economic Research, Inc, 2013.
- [7] AGHION, P., AND HOWITT, P. A Model of Growth through Creative Destruction. *Econometrica* 60, 2 (Mar. 1992), 323–351.
- [8] AGRAWAL, A., AND GOLDFARB, A. Restructuring Research: Communication Costs and the Democratization of University Innovation. *American Economic Review* 98, 4 (Aug. 2008), 1578–1590.
- [9] AKCIGIT, U., AND KERR, W. R. Growth through Heterogeneous Innovations. NBER Working Paper 16443, National Bureau of Economic Research, Inc, 2010.
- [10] ATKESON, A., AND BURSTEIN, A. T. Innovation, Firm Dynamics, and International Trade. *Journal of Political Economy* 118, 3 (2010), 433–484.
- [11] BARTELSMAN, E., HALTIWANGER, J., AND SCARPETTA, S. Cross-Country Differences in Productivity: The Role of Allocation and Selection. *American Economic Review* 103, 1 (Feb. 2013), 305–334.

- [12] BLACK, S. E., AND LYNCH, L. M. How to Compete: The Impact of Workplace Practices and Information Technology on Productivity. *The Review of Economics and Statistics* 83, 3 (2001), 434–445.
- [13] BOPPART, T., AND STAUB, K. E. Online Accessibility of Scholarly Literature, and Academic Innovation. *Working Paper* (2014).
- [14] BRYNJOLFSSON, E., AND HITT, L. Paradox Lost? Firm-Level Evidence on the Returns to Information Systems Spending. *Management Science* 42, 4 (1996), 541–558.
- [15] BRYNJOLFSSON, E., AND HITT, L. M. Computing Productivity: Firm-Level Evidence. *The Review of Economics and Statistics* 85, 4 (2003), 793–808.
- [16] BRYNJOLFSSON, E., AND YANG, S. Information Technology and Productivity: A Review of the Literature. *Advances in Computers, Vol. 43, p.179-214* (1996).
- [17] BUERA, F. J., AND KABOSKI, J. P. The Rise of the Service Economy. *American Economic Review* 102, 6 (2012), 2540–69.
- [18] BUERA, F. J., AND KABOSKI, J. P. Scale and the Origins of Structural Change. *Journal of Economic Theory* 147, 2 (2012), 684–712.
- [19] CABALLERO, R. J., AND JAFFE, A. B. How High are the Giants’ Shoulders: An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Economic Growth. NBER chapters, National Bureau of Economic Research, Inc, 1993.
- [20] CAI, J., AND LI, N. Growth through Inter-Sectoral Knowledge Linkages. Mimeo: International Monetary Fund, 2012.
- [21] CUMMINS, J. G., AND VIOLANTE, G. L. Investment-Specific Technical Change in the US (1947-2000): Measurement and Macroeconomic Consequences. *Review of Economic Dynamics* 5, 2 (2002), 243–284.
- [22] DEDRICK, J., GURBAXANI, V., AND KRAEMER, K. L. Information Technology and Economic Performance: A Critical Review of the Empirical Evidence. *ACM Comput. Surv.* 35, 1 (Mar. 2003), 1–28.
- [23] FORMAN, C., GOLDFARB, A., AND GREENSTEIN, S. Information Technology and the Distribution of Inventive Activity. Working Paper 20036, National Bureau of Economic Research, Apr. 2014.

- [24] FORMAN, C., AND ZEEBROECK, N. v. From Wires to Partners: How the Internet Has Fostered R&D Collaborations Within Firms. *Management Science* 58, 8 (2012), 1549–1568.
- [25] FOSTER, L., HALTIWANGER, J., AND SYVERSON, C. Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability? *American Economic Review* 98, 1 (2008), 394–425.
- [26] FOSTER, L., HALTIWANGER, J. C., AND KRIZAN, C. J. Aggregate Productivity Growth. Lessons from Microeconomic Evidence. NBER chapters, National Bureau of Economic Research, Inc, 2001.
- [27] GORDON, R. J. *The Measurement of Durable Goods Prices*. University of Chicago Press.
- [28] GORDON, R. J. Does the "New Economy" Measure Up to the Great Inventions of the Past? *Journal of Economic Perspectives* 14, 4 (2000), 49–74.
- [29] GREENWOOD, J., GUNER, N., KOCHARKOV, G., AND SANTOS, C. Marry Your Like: Assortative Mating and Income Inequality. *American Economic Review* 104, 5 (2014), 348–53.
- [30] GREENWOOD, J., AND YORUKOGLU, M. 1974. Tech. rep., Carnegie-Rochester Conference Series on Public Policy, 1998.
- [31] HALL, B. H., JAFFE, A. B., AND TRAJTENBERG, M. The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools. NBER Working Paper 8498, National Bureau of Economic Research, Inc, 2001.
- [32] HERRENDORF, B., ROGERSON, R., AND VALENTINYI, K. Growth and Structural Transformation. *Handbook of economic growth*, Elsevier, 2014.
- [33] HORNSTEIN, A., AND KRUSELL, P. Can Technology Improvements Cause Productivity Slowdowns? NBER chapters, National Bureau of Economic Research, Inc, 1996.
- [34] JAFFE, A. B., TRAJTENBERG, M., AND HENDERSON, R. Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations. *The Quarterly Journal of Economics* 108, 3 (Aug. 1993), 577–598.

- [35] JORGENSEN, D., HO, M., AND STIROH, K. Growth of US Industries and Investments in Information Technology and Higher Education. *Economic Systems Research* 15, 3 (2003), 279–325.
- [36] KERR, W. R. Ethnic Scientific Communities and International Technology Diffusion. *The Review of Economics and Statistics* 90, 3 (2008), 518–537.
- [37] KLEIS, L., CHWELOS, P., RAMIREZ, R. V., AND COCKBURN, I. Information Technology and Intangible Output: The Impact of IT Investment on Innovation Productivity. *Information Systems Research* 23, 1 (2012), 42–59.
- [38] KLETTE, T. J., AND KORTUM, S. Innovating Firms and Aggregate Innovation. *Journal of Political Economy* 112, 5 (2004), 986–1018.
- [39] KORTUM, S. Equilibrium R&D and the Patent-R&D Ratio: U.S. Evidence. *American Economic Review* 83, 2 (1993), 450–57.
- [40] LENTZ, R., AND MORTENSEN, D. T. An Empirical Model of Growth Through Product Innovation. *Econometrica* 76, 6 (2008), 1317–1373.
- [41] LOVEMAN, G. W. An Assessment of the Productivity Impact of Information Technologies. *Information Technology and the Corporation of the 1990s: Research Studies*, pp. 84–110 (1994).
- [42] MOSTELLER, F. Association and Estimation in Contingency Tables. *Journal of the American Statistical Association* 63, 321 (Mar. 1968), 1–28.
- [43] OLINER, S. D., AND SICHEL, D. E. The Resurgence of Growth in the Late 1990s: Is Information Technology the Story? *Journal of Economic Perspectives* 14, 4 (2000), 3–22.
- [44] SILVERMAN, B. S. Technological Resources and the Direction of Corporate Diversification: Toward an Integration of the Resource-Based View and Transaction Cost Economics. *Management Science* 45, 8 (Aug. 1999), 1109–1124.
- [45] SINKHORN, R., AND KNOPP, P. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific Journal of Mathematics* 21, 2 (1967), 343–348.
- [46] STIROH, K. J. Information Technology and the U.S. Productivity Revival: What Do the Industry Data Say? *American Economic Review* 92, 5 (2002), 1559–1576.

- [47] SYVERSON, C. What Determines Productivity? *Journal of Economic Literature* 49, 2 (June 2011), 326–365.
- [48] TAMBE, P., HITT, L. M., AND BRYNJOLFSSON, E. The Extroverted Firm: How External Information Practices Affect Innovation and Productivity. *Management Science* 58, 5 (Jan. 2012), 843–859.
- [49] YORUKOGLU, M. The Information Technology Productivity Paradox. *Review of Economic Dynamics* 1, 2 (1998), 551–592.

Appendix

A Proofs and Derivations

Bellman Equation. The value of a s -type firm with a technology gap n at time t is equal to

$$V_t(n, s) = \max_{x_t(n, s), z_t(n, s)} \left\{ +e^{-r_{t+\Delta t}\Delta t} \left[\begin{aligned} &\left[\begin{aligned} &\Pi_t(n) - \frac{\alpha_s}{\mathcal{E}(s, t)^\phi} x_t(n, s)^\gamma \lambda^{-n} \\ &- \frac{\alpha_s \beta}{\mathcal{E}(s, t)^\psi} z_t(n, s)^\gamma \end{aligned} \right] \Delta t + o(\Delta t) \\ &(x_t(n, s)\Delta t + o(\Delta t))V_{t+\Delta t}(n+1, s) \\ &+(z_t(n, s)\Delta t + o(\Delta t))(p_{t+\Delta t} + V_{t+\Delta t}(n, s)) \\ &+(\tau_t\Delta t + o(\Delta t)) \times 0 \\ &(1 - x_t(n, s)\Delta t - z_t(n, s)\Delta t - \tau_t\Delta t - o(\Delta t))V_{t+\Delta t}(n, s) \end{aligned} \right] \right\}$$

The first part of the expression is a flow profit minus R&D cost of horizontal and vertical innovation in a time interval of Δt . $o(\Delta t)$ denotes second-order terms. Second part is a discounted continuation value after the interval has elapsed. There are different scenarios: either successful horizontal innovation happens with probability $x_t(n, s)\Delta t + o(\Delta t)$, or horizontal innovation happens with probability $z_t(n, s)\Delta t + o(\Delta t)$, or creative destruction hits the product line with probability $\tau_t\Delta t$, or with complementary probability none of those events get realized. Under these scenarios, respectively, firm gets value of $V_{t+\Delta t}(n+1, s)$ – next period value from having a larger technology gap, or $p_{t+\Delta t} + V_{t+\Delta t}(n, s)$ – price from selling external product line and then continuing with the same technology gap, n , or the value of exit which is zero, or the value of continuing with the same technology gap, n . To obtain the expression in the main text, subtract $V_t(n, s)$ from both sides, divide every term

by Δt and take the limit as $\Delta t \rightarrow 0$. We get:

$$r(t)V_t(n, s) - \dot{V}_t(n, s) = \max_{x_t(n, s), z_t(n, s)} \left\{ \begin{array}{l} \Pi_t(n) - \frac{\alpha_s}{\bar{\mathcal{E}}(s, t)^\phi} x_t(n, s)^\gamma \lambda^{-n} \\ - \frac{\alpha_s \beta}{\bar{\mathcal{E}}(s, t)^\psi} z_t(n, s)^\gamma \\ + x_t(n, s)(V_t(n+1, s) - V_t(n, s)) \\ + z_t(n, s)p_t - \tau_t V_t(n, s) \end{array} \right\} \quad (24)$$

Proof of Proposition 3.1. Conjecture that $v(n, s) = A(s) - B(s)\lambda^{-n}$. Plug into equation (15) (where $x(n, s)$, $z(n, s)$ are already taken optimally and equation (7) is used)

$$\begin{aligned} \rho A(s) - \rho B(s)\lambda^{-n} &= 1 - \lambda^{-n} - \frac{\alpha_s}{\bar{\mathcal{E}}(s, \varepsilon)^\phi} x_t(n, s)^\gamma \lambda^{-n} - \frac{\alpha_s \beta}{\bar{\mathcal{E}}(s, \varepsilon)^\psi} z(n, s)^\gamma \\ &\quad + x(n, s)B(s)\lambda^{-n} \frac{\lambda - 1}{\lambda} + z(n, s)p - \tau A(s) + \tau B(s)\lambda^{-n} \end{aligned}$$

For that identity to always hold, we need to make sure that for any n , the right-hand side is equal to the left-hand side. After grouping the terms with and without λ^{-n} term, the following should always be true:

$$\rho A(s) = 1 - \frac{\alpha_s \beta}{\bar{\mathcal{E}}(s, \varepsilon)^\psi} z(n, s)^\gamma + z(n, s)p_t - \tau A(s) \quad (25)$$

$$-\rho B(s) = -1 - \frac{\alpha_s}{\bar{\mathcal{E}}(s, \varepsilon)^\phi} x(n, s)^\gamma + x(n, s)B(s) \frac{\lambda - 1}{\lambda} + \tau B(s) \quad (26)$$

Next, I show that innovation rates are independent of n . First order conditions imply:

$$\gamma \frac{\alpha_s}{\bar{\mathcal{E}}(s, \varepsilon)^\phi} x(n, s)^{\gamma-1} = B(s) \frac{\lambda - 1}{\lambda}$$

and

$$\gamma \frac{\alpha_s \beta}{\bar{\mathcal{E}}(s, \varepsilon)^\psi} z(n, s)^{\gamma-1} = p$$

From the above FOC's, we see that $x(n, s) = x(s)$ and $z(n, s) = z(s)$. In both cases neither x nor z depends on n . Using this fact, equations (25) and (26) make it clear that both identities hold for all n . Solving for $A(s)$ and $B(s)$ gives the expressions in equations (17) and (18). This verifies the conjecture.

Derivation of Equation (14):

$$\begin{aligned}
\log Y(t) &= \int_0^1 \log[q_i(j, t) f(ICT(j, t)) l(j, t)] dj \\
&= \int_0^1 \log[q_i(j, t) f(ICT(j, t))] dj + \int_0^1 \log \frac{\lambda^{n(j, t)}}{w_t} Y(t) dj \\
&= \int_0^1 \log[q_i(j, t) f(ICT(j, t))] dj + \ln Y(t) - \log \lambda \int_0^1 n(j, t) dj - \log w_t
\end{aligned}$$

where the second row comes from using the labor demand equation in (12). After rearranging and canceling the terms, we get:

$$w_t = \exp \left[\int_0^1 \log[q_i(j, t) f(ICT(j, t))] dj - \log \lambda \int_0^1 n(j, t) dj \right]$$

Because we have a finite state space, we can rewrite $\int_0^1 n(j, t) dj = \sum_{n,s} n \mu_t(n, s)$. In addition, if we introduce the notation $\log Q(t) = \int_0^1 \log[q_i(j, t) f(ICT(j, t))] dj$, equation (14) follows:

$$w_t = Q(t) \lambda^{-\sum_{n,s} n \mu_t(n, s)}$$

Derivation of Equation (22): Using labor market clearing condition (13) and wage equation (14), we can rewrite the equation for output in the following way:

$$Y(t) = Q(t) L \frac{\lambda^{-\sum_{n,s} n \mu_t(n, s)}}{\sum_{n,s} \lambda^{-n} \mu_t(n, s)}$$

As a result, in steady state, because the distribution of product lines across the state-space is constant, growth of output g is same as growth of composite quality $Q(t)$. Outside the steady state, growth will also be affected by the changes in the distribution $\mu_t(n, s)$. If, for example, distribution over the technology gaps shifts left, the growth will also increase because of the reduced markups. To derive the steady state growth g , rewrite:

$$Q(t) = \exp \left(\int_0^1 \log[q(j, t) f(ICT(j, t))] dj \right)$$

Then the growth rate is

$$g_Q = \frac{\Delta \exp \int_0^1 \log q(j, t) dj}{\exp \int_0^1 \log q(j, t) dj} + \frac{\Delta \exp \int_0^1 \log f(ICT(j, t)) dj}{\exp \int_0^1 \log f(ICT(j, t)) dj}$$

Growth of the endogenous quality component is:

$$\frac{\Delta \exp \int_0^1 \log q(j, t) dj}{\exp \int_0^1 \log q(j, t) dj} = \lim_{\Delta t \rightarrow 0} \frac{\int_0^1 \log q(j, t + \Delta t) dj - \int_0^1 \log q(j, t) dj}{\Delta t}$$

In a unit interval Δt , quality improves by factor λ with probability of vertical innovation plus creative destruction. As a result,

$$\int_0^1 \log q(j, t + \Delta t) dj = \log \lambda \left(\frac{\sum_s x_s}{S} \Delta t + \tau \Delta t + o(\Delta t) \right) + \left(\int_0^1 \log q(j, t) dj \right)$$

After plugging this expression into the above ratio, simplifying it and taking a limit, we get:

$$g = \log \lambda \left(\frac{\sum_s x_s}{S} + \tau \right) + \frac{\Delta F(ICT)}{F(ICT)},$$

where the second term is an exogenous growth in composite ICT input and $F(ICT) = \exp \int_0^1 \log f(ICT(j, t)) dj$ (which is constant in the steady state).

B Concordances

Because the BEA data on ICT stock is given using NAICS industry classification and the USPTO issues patents according to the technology classifications, I use concordances between these two datasets. In particular, for my purpose, I need to have a concordance from 3-4 digit NAICS industries to *nclass* classification and vice versa. To map USPTO technology classification (*nclass*) to 3-digit SIC codes, I use the concordance developed by Kerr (2008). Mapping does not go one-to-one and it contains the probabilities based on the industries where new inventions from USPTO technology classes are manufactured or used. In my analysis, I use probabilities based on manufacturing-industry frequencies. I complement this concordance with SIC-NAICS concordance from BEA. This way, I convert the BEA ICT data by NAICS classification into ICT data by patent technology classification. This data is used in all industry-level regressions that use both patents and ICT data.

In other cases, when I conduct the analysis at the technology class level, I require to have the industry ICT data converted to *nclass*-level data. For this purpose, I develop an opposite concordance. In particular, I reverse an original concordance developed by Silverman (1999) which connects IPC with Canadian SIC and then complement it with other concordances from Statistics Canada (2007) to get the intermediate output of a concordance going from NAICS industry to IPC. In the final step, I connect IPC to *nclass* classification by looking at joint assignments of IPC and *nclass* classifications to patents from NBER data.

C Additional Empirical Facts

Increase in ICT. Figure 10 illustrates distributions of total real ICT stock (chain typed quantity indexes) across NAICS industries in the period of 1976-1978 and 2001-2003. Distribution strongly shifted to the right – literally all industries experienced huge increase in the levels of ICT stock starting from very low investments in late 70s. Not only levels but ICT intensities increased as well. Figure 11 shows the shift in the distribution of shares of industry’s real ICT stock in total fixed assets.




Figure 10: Total Real ICT Stock over Time and Industries

Notes: BEA data on Chain Type Quantity Index of fixed non-residential assets over 66 NAICS industries, base index year=2005. ICT stock includes communications, hardware and software: mainframes, PCs, printers, terminals, tape drives, storage devices, system integrators, prepackaged software, custom software, own account software. For the median industry, ICT quantity index increased 57 times.




Figure 11: Share of Real ICT Stock in Total Real Fixed Assets

Notes: BEA data on fixed non-residential assets over 66 NAICS industries, base index year=2005. ICT stock includes communications, hardware and software: mainframes, PCs, printers, terminals, tape drives, storage devices, system integrators, prepackaged software, custom software, own account software. Total fixed assets is the sum of equipments and structures. Both ICT stock and total fixed assets are deflated by corresponding price indexes. For median industry, ratio of two series increased 13 times.

Knowledge Diffusion Below, is the algorithm employed to standardize input-output tables described in the paper. Algorithm similar to Sinkhorn-Knopp (1976) is used. I standardize the table for 1976-1978 such that it has same marginal distributions as a table from the data in 2001-2003¹⁷. Following are the steps to achieve it:

1. Compute marginal distributions for the column (citing technology class) and the row (cited technology class) of 2001-2003 input-output table. I will call them target-row and target-column marginal distributions.
2. Start an iteration from the input-output table of 1976-1978. Compute marginal distributions for its row and column.
3. Multiply each cell by the ratio of target to existing marginal distribution of the rows. Compute a new marginal distribution for columns. Multiply each cell by the ratio of target to existing new marginal distribution of the columns. Compute new marginal distributions for columns and rows.

¹⁷Similarly, one could adjust 2001-2003 table to match marginal distributions of 1976-1978 table.

4. Compare differences between target and existing marginal distributions. If they are different, go back to Step 2.

Figure 1 in the paper is based on this procedure. The pattern is very similar if we look at the raw tables without standardization. Figure 12 illustrates this.




Figure 12: Citations Input-Output Matrix

Notes: Figure illustrates citations input-output matrices for the periods of 1976-1978 and 2001-2003. Each row represents one of the 37 technology categories and each cell depicts a share of citations which are given by citing technology class to corresponding cited technology class. Higher shares are depicted in darker colors. All shares in each row add up to one. Cited technology classes on the horizontal axis are ranked by citation shares received.

As noted in the main text, I examine a similar input-output tables for finer technology categories (*nclass*). In fact, these categories are the ones used in subsequent empirical and quantitative analysis. However, for the purpose of better visualization I first picked broader categories to illustrate an increase in external knowledge flows over time. Figure 13 depicts the relationship on the finer level. Though dark shadings are not well visible, the drastic increase in the number of new technology classes cited clearly stands out.




Figure 13: Citations Input-Output Matrix

Notes: Figure illustrates citations input-output matrices for the periods of 1976-1978 and 2001-2003. Each row represents one of the 412 technology categories and each cell depicts a share of citations which are given by citing technology class to corresponding cited technology class. Higher shares are depicted in darker colors. All shares in each row add up to one. Cited technology classes on the horizontal axis are ranked by citation shares received.

Table 6 shows a positive relationship between ICT and external learning at the technology-class level. Details are discussed in the main text.

Table 6: ICT and Knowledge Diffusion

	Num. external classes cited	External citations share
Log ICT	0.151*** (0.003)	0.042*** (0.012)
Year	Yes	Yes
Technology class	Yes	Yes
Observations	11286	11286
R^2	0.86	0.85

Notes: Industry (n_{class}) $\times$ Year panel regressions. Additional controls are number of citations-adjusted patents in a technology class per year, number of firms (assignees) patenting in a class and real fixed capital stock from the BEA data. Dependent variable in the first column is a number of different external technology classes cited by patents in a technology class per year (Poisson regression); dependent variable in the second column is a share of self-class citations among all citations made by patents of a technology class in a year (OLS regression). Standard errors clustered by technology class. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$