



The employment effects of innovation

Alex Coad, Rekha Rao

► To cite this version:

| Alex Coad, Rekha Rao. The employment effects of innovation. 2007. halshs-00175042

HAL Id: halshs-00175042

<https://shs.hal.science/halshs-00175042>


Submitted on 26 Sep 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



The employment effects of innovation

Alex COAD, Rekha RAO

2007.36



CENTRE NATIONAL
DE LA RECHERCHE
SCIENTIFIQUE

THE EMPLOYMENT EFFECTS OF INNOVATION ^{*}

Alex Coad ^{a b †}

Rekha Rao ^c

^a Max Planck Institute of Economics, Jena, Germany

^b Centre d'Economie de la Sorbonne, Equipe MATISSE, Univ. Paris 1 – CNRS

^c LEM, Sant'Anna School of Advanced Studies, Pisa, Italy.

July 9, 2007

Abstract

The issue of technological unemployment receives perennial popular attention. Although there are previous empirical investigations that have focused on the relationship between innovation and employment, the originality of our approach lies in our choice of method. We focus on four 2-digit manufacturing industries that are known for their high patenting activity. We then use Principal Components Analysis to generate a firm- and year-specific ‘innovativeness’ index by extracting the common variance in a firm’s patenting and R&D expenditure histories. To begin with, we explore the heterogeneity of firms by using semi-parametric quantile regression. Whilst some firms may reduce employment levels after innovating, others increase employment. We then move on to a weighted least squares (WLS) analysis, which explicitly takes into account the different job-creating potential of firms of different sizes. As a result, we focus on the effect of innovation on total number of jobs, whereas previous studies have focused on the effect of innovation on firm behavior. Indeed, previous studies have typically taken the firm as the unit of analysis, implicitly weighting each firm equally according to the principle of ‘one firm equals one observation’. Our results suggest that firm-level innovative activity leads to employment creation that may have been underestimated in previous studies.

L’INNOVATION DES ENTREPRISES ET CROISSANCE DE L’EMPLOI

Résumé: Nous présentons une analyse de la relation entre l’innovation et l’emploi dont l’originalité repose sur la méthodologie statistique. Nous nous concentrons sur quatre industries manufacturières qui ont une forte propension à déposer des brevets. Ensuite, nous utilisons l’analyse en composantes principales pour élaborer une indice d’innovation. Dans un premier temps, nous explorons l’hétérogénéité des firmes en appliquant la régression par quantiles. Alors que l’innovation peut entraîner dans quelques entreprises des réductions d’emplois, dans d’autres, elle est à l’origine d’une augmentation des emplois. Nous procédons à une analyse par les moindres carrés pondérés (‘weighted least squares’, WLS) qui prend en compte les différentes capacités de création d’emploi d’entreprises de tailles variées. Nous nous concentrons donc bien sur l’effet de l’innovation sur le nombre total d’emplois, alors

^{*}An earlier draft of this paper won the ‘best junior paper award’ at the DRUID summer conference 2007. We thank Hans Gersbach, Karin Hoisl, Camilla Lenzi, Pierre Mohnen, and Christian Seiser, as well as participants at the DRUID summer conference 2007 and the Monte Verità conference on “The Economics of Technology Policy” (July 2007) for helpful comments and discussions. The usual caveat applies.

[†]*Corresponding Author:* Alex Coad, Max Planck Institute of Economics, Evolutionary Economics Group, Kahlaische Strasse 10, D-07745 Jena, Germany. *Phone:* +49 3641 686822. *Fax:* +49 3641 686868. *E-mail:* coad@econ.mpg.de

que les études précédentes se sont attachées à l'effet de l'innovation sur le comportement des entreprises. Nos résultats suggèrent que la contribution de l'innovation sur la création d'emplois a été sous-estimée dans les études antérieures.

JEL codes: L25, O33, J01

Keywords: Technological Unemployment, Innovation, Firm Growth, Weighted Least Squares, Aggregation, Quantile Regression

Mots clés: Chômage technologique, Innovation, Croissance des firmes, Weighted Least Squares, Aggrégation, Régression par quantile

1 Introduction

Whilst firm-level innovation can be expected to have a positive influence on the growth of a firm's sales, the overall effect on employment growth is *a priori* ambiguous. Innovation is often associated with increases in productivity that lower the amount of labour required for the production of goods and services. In this way, an innovating firm may change the composition of its productive resources, to the profit of machines and at the expense of employment. As a result, the general public has often expressed concern that technological progress may bring about the 'end of work' by replacing men with machines. Economists, on the other hand, are usually more optimistic.

To begin with, theoretical discussions have found it useful to decompose innovation into product and process innovation. Product innovations are often associated with employment gain, because the new products create new demand (although it is possible that they might replace existing products). Process innovations, on the other hand, often increase productivity by reducing the labour requirement in manufacturing processes (e.g. via the introduction of robots (Fleck, 1984)). Thus, process innovations are often suspected of bringing about 'technological unemployment'. The issue becomes even more complicated, however, when we consider that there are not only direct effects of innovation on employment, but also a great many indirect effects operating through various 'substitution channels'. For example, the introduction of a labour-saving production process may lead to an immediate and localized reduction in employees inside the plant (the 'direct effect'), but it may lead to positive employment changes elsewhere in the economy via an increased demand for new machines, a decrease in prices, and increase in incomes, an increase in new investments, or a decrease in wages (for an introduction to the various 'substitution channels', see Spiezia and Vivarelli, 2000). As a result, the overall effect of innovation on employment needs to be investigated empirically. Although Van Reenen recently lamented the "dearth of microeconomic studies on the effect of innovation on employment" (Van Reenen, 1997: 256), the situation has improved over the last decade.

Research into technological unemployment has been undertaken in different ways and at various levels of aggregation. The results emerging from different studies are far from harmonious though – "[e]mpirical work on the effect of innovations on employment growth yields very mixed results" (Niefert 2005:9). Doms et al. (1995) analyse survey data on US manufacturing establishments, and observe that the use of advanced manufacturing technology (which would correspond to process innovation) has a positive effect on employment. At the firm-level of analysis, Hall (1987) observes that employment growth is related positively and significantly to R&D intensity in the case of large US manufacturing firms. Similarly, Greenhalgh et al. (2001) observe that R&D intensity and also the number of patent publications

have a positive effect on employment for British firms. Nevertheless, Evangelista and Savona (2003) observe a negative overall effect of innovation on employment in the Italian services sector. When the distinction is made between product and process innovation, the former is usually linked to employment creation whereas the consequences of the latter are not as clear-cut. Evidence presented in Brouwer et al. (1993) reveals a small positive employment effect of product-related R&D although the combined effect of innovation is imprecisely defined. Relatedly, work by Van Reenen (1997) on listed UK manufacturing firms and Smolny (1998) for West German manufacturing firms shows a positive effect on employment for product innovations. Smolny also finds a positive employment effect of process innovations, whereas Van Reenen’s analysis yields insignificant results. Harrison et al. (2005) consider the relationship between innovation and employment growth in four European countries (France, Italy, the UK and Germany) using data for 1998 and 2000 on firms in the manufacturing and services industries. Whilst product innovations are consistently associated with employment growth, process innovation appears to have a negative effect on employment, although the authors acknowledge that this latter result may be attenuated (or even reversed) through compensation effects. To summarize, therefore, we can consider that product innovations generally have a positive impact on employment, whilst the role of process innovations is more ambiguous (Hall et al., 2006).

We must emphasize, however, that investigations at the level of the firm do not allow us to infer the aggregated and cumulative effect of innovation on ‘total jobs’ – this is because datasets are composed of firms of different sizes which need to be weighted accordingly. Previous research in this area, however, has implicitly given equal weights to firms, by treating each firm as one ‘observation’ in a larger database. These studies can shed light on the effect of innovation on employment decisions in the ‘average firm’, but they do not yield conclusions on the total employment effects of innovation, for society as a whole.¹

We have strong theoretical motivations for suspecting that the relationship between innovation and employment is not invariant over the firm size distribution. For example, it may be the case that larger firms are more prone to introduce labour-saving process innovations, whereas smaller firms are often associated with product innovations. In this way, innovation in larger firms may be associated with job destruction whereas the innovative activity of small firms would be associated with job creation. On the other hand, smaller firms have less restrictive hiring-and-firing regulations, and so innovation may lead to reductions in employment that are more frequent in smaller firms than in their larger counterparts. Although there may be a relationship between the size of a firm and the employment effects of innovation, however,

¹Note however that Evangelista and Savona (2002) estimate the effect of firm-level innovation on total employment by attributing observation-specific weights to firms. Nonetheless, their analysis is rather limited because their employment growth variable is a qualitative survey response instead of a quantitative growth rate.

we consider the sign and magnitude to be an empirical question. Our empirical framework enables us to evaluate the effect of innovation on total employment by attributing weights to firms of different sizes.

“Linking more explicitly the evidence on the patterns of innovation with what is known about firms growth and other aspects of corporate performance – both at the empirical and at the theoretical level – is a hard but urgent challenge for future research” (Cefis and Orsenigo, 2001:1157). We are now in a position to rise to this challenge. In Section 2 we discuss the methodology, focusing in particular on the shortcomings of using either patent counts or R&D figures individually as proxies for innovativeness. We describe how we use Principal Component Analysis to extract a synthetic ‘innovativeness’ index from patent and R&D data. Section 3 describes how we matched the Compustat database to the NBER innovation database, and we describe how we created our synthetic ‘innovativeness’ index. Indeed, we have made efforts to obtain the best possible observations for firm-level innovative activity. Whilst our database does not allow any formal distinction between ‘product’ and ‘process’ innovation, however, we do not consider this to be a fatal caveat for the purposes of this investigation. Section 4 contains the semi-parametric quantile regression analysis, where we can observe how the influence of innovation on employment change varies across the conditional growth rate distribution. We then move on to the parametric analysis in Section 5. In particular, we compare the estimates obtained from conventional regressions (OLS and FE) with those obtained from weighted least squares (WLS). We observe that WLS estimation consistently yields a slightly more positive (although never statistically significant) estimate than other techniques, which suggests that previous studies may have underestimated the total employment gains from innovation. Section 6 concludes.

2 Methodology – How can we measure innovativeness?

Activities related to innovation within a company can include research and development; acquisition of machinery, equipment and other external technology; industrial design; and training and marketing linked to technological advances. These are not necessarily identified as such in company accounts, so quantification of related costs is one of the main difficulties encountered during the innovation studies. Each of the above mentioned activities has some effect on the growth of the firm, but the singular and cumulative effect of each of these activities is hard to quantify. Data on innovation per se has thus been hard to find (Van Reenen, 1997). Also, some sectors innovate extensively, some don’t innovate in a tractable manner, and the same is the case with organizational innovations, which are hard to quantify in terms of impact on the overall growth of the firms. However, we believe that no firm can survive without at least some degree of innovation.

We use two indicators for innovation in a firm: first, the patents applied for by a firm and second, the amount of R&D undertaken. Cohen et al. (2000) suggest that no industry relies exclusively on patents, yet the authors go on to suggest that the patents may add sufficient value at the margin when used with other appropriation mechanisms. Although patent data has drawbacks, patent statistics provide unique information for the analysis of the process of technical change (Griliches, 1990). We can use patent data to access the patterns of innovation activity across fields (or sectors) and nations. The number of patents can be used as an indicator of inventive as well as innovative activity, but it has its limitations. One of the major disadvantage of patents as an indicator is that not all inventions and innovations are patented (or indeed ‘patentable’). Some companies – including a number of smaller firms – tend to find the process of patenting expensive or too slow and implement alternative measures such as secrecy or copyright to protect their innovations (Archibugi, 1992; Arundel and Kabla, 1998). Another bias in the study using patenting can arise from the fact that not all patented inventions become innovations. The actual economic value of patents is highly skewed, and most of the value is concentrated in a very small percentage of the total (OECD, 1994). Furthermore, another caveat of using patent data is that we may underestimate innovation occurring in large firms, because these typically have a lower propensity to patent (Dosi, 1988). The reason we use patent data in our study is that, despite the problems mentioned above, patents would reflect the continuous developments within technology. We complement the patent data with R&D data. R&D can be considered as an input into the production of inventions, and patents as outputs of the inventive process. R&D data may lead us to systematically underestimate the amount of innovation in smaller firms, however, because these often innovate on a more informal basis outside of the R&D lab (Dosi, 1988). For some of the analysis we consider the R&D stock and also the patent stock, since the past investments in R&D as well as the past applications of patents have an impact not only on the future values of R&D and patents, but also on firm growth. Hall (2004) suggests that the past history of R&D spending is a good indicator of the firms technological position.

Taken individually, each of these indicators for firm-level innovativeness has its drawbacks. Each indicator on its own provides useful information on a firm’s innovativeness, but also idiosyncratic variance that may be unrelated to a firm’s innovativeness. One particular feature pointed out by Griliches (1990) is that, although patent data and R&D data are often chosen to individually represent the same phenomenon, there exists a major statistical discrepancy in that there is typically a great randomness in patent series, whereas R&D values are much more smoothed. Figure 1 shows that the variable of interest (i.e. ΔK – additions to economically valuable knowledge) is measured with noise if one takes either innovative input (such as R&D expenditure or R&D employment) or innovative output (such as patent statistics). In order to remove this noise, one needs to collect information on both innovative input and




Figure 1: The Knowledge ‘Production Function’: A Simplified Path Analysis Diagram (based on Griliches 1990:1671)

output, and to extract the common variance whilst discarding the idiosyncratic variance of each individual proxy that includes noise, measurement error, and specific variation. In this study, we believe we have obtained useful data on a firm’s innovativeness by considering both innovative input and innovative output simultaneously in a synthetic variable.² Principal Component Analysis (PCA) is appropriate here as it allows us here to summarize the information provided by several indicators of innovativeness into a composite index, by extracting the common variance from correlated variables whilst separating it from the specific and error variance associated with each individual variable (Hair et al., 1998). We are not the only ones to apply PCA to studies into firm-level innovation however – this technique has also been used by Lanjouw and Schankerman (2004) to develop a composite index of ‘patent quality’ using multiple characteristics of patents (such as the number of citations, patent family size and patent claims).

Another criticism of previous studies is that they have lumped together firms from all manufacturing sectors – even though innovation regimes (and indeed appropriability regimes) vary dramatically across industries. In this study, we focus on specific 2-digit sectors that have been hand-picked according to their intensive patenting and R&D activity. However, even within

²Following Griliches (1990), we consider here that patent counts can be used as a measure of innovative *output*, although this is not entirely uncontroversial. Patents have a highly skew value distribution and many patents are practically worthless. As a result, patent numbers have limitations as a measure of innovative output – some authors would even prefer to consider raw patent counts to be indicators of innovative *input*.

these sectors, there is significant heterogeneity between firms, and using standard regression techniques to make inferences about the average firm may mask important phenomena. Using quantile regression techniques, we investigate the relationship between innovativeness and growth at a range of points of the conditional growth rate distribution. We observe three types of relationship between innovation and employment. First, most firms do not experience much employment change in any given year, and what little change they have appears to be largely idiosyncratic and not strongly related to innovative activity. Second, for those firms that grow the fastest, we observe that innovation seems to be strongly positively associated with increases in employment. Third, for those firms that are rapidly shedding workers, this is strongly associated with innovative activity. We note that this heterogeneity of the response of employment change to innovation cannot be detected if we focus on conventional regression estimators that estimate ‘the average effect for the average firm’.

We only consider certain specific sectors, and not the whole of manufacturing. This way we are not affected by aggregation effects; we are grouping together firms that can plausibly be compared to each other. We are particularly interested in looking at the growth of firms classified under ‘complex’ technology classes. We base our classification of firms on the typology put forward by Hall (2004) and Cohen et al. (2000). The authors define ‘complex product’³ industries as those industries where each product relies on many patents held by a number of other firms and the ‘discrete product’ industries as those industries where each product relies on only a few patents and where the importance of patents for appropriability has traditionally been higher.⁴ We chose four sectors that can be classified under the ‘complex products’ class. The two digit SIC codes that match the ‘complex technology’ sectors are 35, 36, 37, and 38.⁵ By choosing these sectors that are characterised by high patenting and high R&D expenditure, we hope that we will be able to get the best possible quantitative observations for firm-level innovation.

³During our discussion, we will use the terms ‘products’ and ‘technology’ interchangeably to indicate generally the same idea.

⁴It would have been interesting to include ‘discrete technology’ sectors in our study, but unfortunately we did not have a comparable number of observations for these sectors. This remains a challenge for future work.

⁵The ‘complex technology’ sectors that we consider are SIC 35 (industrial and commercial machinery and computer equipment), SIC 36 (electronic and other electrical equipment and components, except computer equipment), SIC 37 (transportation equipment) and SIC 38 (measuring, analyzing and controlling instruments; photographic, medical and optical goods; watches and clocks).

3 Database description

3.1 Database

We create an original database by matching the NBER patent database with the Compustat file database, and this section is devoted to describing the creation of the sample which we will use in our analysis.

The patent data has been obtained from the NBER database (Hall et al., 2001b), and we have used the updates available on Bronwyn Hall’s website⁶ to obtain data until 2002. The NBER database comprises detailed information on almost 3 416 957 U.S. utility patents in the USPTO’s TAF database granted during the period 1963 to December 2002 and all citations made to these patents between 1975 and 2002. A firm’s patenting history is analysed over the whole period represented by the NBER patent database. The initial sample of firms was obtained from the Compustat⁷ database for the aforementioned sectors comprising ‘complex product’ sectors. These firms were then matched with the firm data files from the NBER patent database and we found all the firms⁸ that have patents. The final sample thus contains both patenters and non-patenters.

The NBER database has patent data for over 60 years and the Compustat database has firms’ financial data for over 50 years, giving us a rather rich information set. As Van Reenen (1997) mentions, the development of longitudinal databases of technologies and firms is a major task for those seriously concerned with the dynamic effect of innovation on firm growth. Hence, having developed this longitudinal dataset, we feel that we will be able to thoroughly investigate whether innovation drives sales growth at the firm-level.

Table 1 shows some descriptive statistics of the sample before and after cleaning. Initially using the Compustat database, we obtain a total of 4274 firms which belong to the SICs 35-38

⁶See <http://elsa.berkeley.edu/~bhhall/bhdata.html>

⁷Compustat has the largest set of fundamental and market data representing 90% of the world’s market capitalization. Use of this database could indicate that we have oversampled the Fortune 500 firms. Being included in the Compustat database means that the number of shareholders in the firm was large enough for the firm to command sufficient investor interest to be followed by Standard and Poor’s Compustat, which basically means that the firm is required to file 10-Ks to the Securities and Exchange Commission on a regular basis. It does not necessarily mean that the firm has gone through an IPO. Most of them are listed on NASDAQ or the NYSE.

⁸The patent ownership information (obtained from the above mentioned sources) reflects ownership at the time of patent grant and does not include subsequent changes in ownership. Also attempts have been made to combine data based on subsidiary relationships. However, where possible, spelling variations and variations based on name changes have been merged into a single name. While every effort is made to accurately identify all organizational entities and report data by a single organizational name, achievement of a totally clean record is not expected, particularly in view of the many variations which may occur in corporate identifications. Also, the NBER database does not cumulatively assign the patents obtained by the subsidiaries to the parents, and we have taken this limitation into account and have subsequently tried to cumulate the patents obtained by the subsidiaries towards the patent count of the parent. Thus we have attempted to create an original database that gives complete firm-level patent information.

Table 1: Summary statistics before and after data-cleaning (SIC's 35-38 only)

	sample before cleaning $n=4274$ firms		sample used $n=1920$ firms	
	mean	std. dev.	mean	std. dev.
Total Sales	1028.156	6775.733	1178.169	7046.354
Patent applications	6.387029	45.35137	9.267999	56.86579
R&D expenditure	58.6939	363.402	57.06897	351.7708
Total Employees	8.5980	40.1245	10.21704	46.0217

Table 2: Firm size distribution in SIC35-38, 1963-1998

No. of Employees		SIC 35	SIC 36	SIC 37	SIC 38
≤ 250	Mean	0.104332	0.112787	0.101951	0.095604
	Std. Dev	0.069861	0.06795	0.071023	0.070228
	obs	2570	3196	266	3667
$> 250 \ \& \ \leq 500$	Mean	0.371858	0.36585	0.375686	0.36347
	Std. Dev	0.071885	0.071307	0.075044	0.069809
	obs	969	1347	204	879
$> 500 \ \& \ \leq 5000$	Mean	1.802632	1.684009	2.091483	1.641482
	Std. Dev	1.187339	1.109291	1.174895	1.094161
	obs	3317	2941	937	2018
> 5000	Mean	33.91514	43.34083	91.30289	25.02034
	Std. Dev	50.19058	64.77395	165.8062	28.9475
	obs	1729	1322	1312	935

Note: employee numbers given in thousands.

and this sample consists of both innovating and non-innovating firms. These firms were then matched to the NBER database. After this initial match, we further matched the year-wise firm data to the year-wise patents applied by the respective firms (in the case of innovating firms) and finally, we excluded firms that had less than 7 consecutive years of good data. Thus, we have an unbalanced panel of 1920 firms belonging to 4 different sectors. Since we intend to take into account sectoral effects of innovation, we will proceed on a sector by sector basis, to have (ideally) 4 comparable results for 4 different sectors.

3.2 Summary statistics and the ‘innovativeness’ index

Table 2 provides some insights into the firm size distribution for each of the four sectors. We can observe a certain degree of heterogeneity between the sectors, with SIC 37 (Transportation equipment) containing relatively large proportion of large firms.




Figure 2: Number of patents per year. SIC 35: Machinery & Computer Equipment, SIC 36: Electric/Electronic Equipment, SIC 37: Transportation Equipment, SIC 38: Measuring Instruments.

Figure 2 shows the number of patents per year in our final database. For some of the sectors there appears to be a strong structural break at the beginning of the 1980s which may well be due to changes in patent regulations (see Hall (2004) for a discussion). Table 3 presents the firm-wise distribution of patents, which is noticeably right-skewed. We find that 46% of the firms in our sample have no patents. Thus the intersection of the two datasets gave us 1028 patenting firms who had taken out at least one patent between 1963 and 1998, and 892 firms that had no patents during this period. The total number of patents taken out by this group over the entire period was 317 116, where the entire period for the NBER database represented years 1963 to 2002, and we have used 274 964 of these patents in our analysis i.e. representing about 87% of the total patents ever taken out at the US Patent Office by the firms in our sample. Though the NBER database provides the data on patents applied for from 1963 till 2002, it contains information only on the granted patents and hence we might see some bias towards the firms that have applied in the end period covered by the database due the lags faced between application and the grant of the patents. Hence to avoid this truncation bias (on the right) we consider the patents only till 1997 so as to allow for a 5-year gap between application and grant of the patent.⁹ Concerning R&D, 1867 of the 1920 firms

⁹The gap between application and grant of a patent has been referred to by many authors, among others Bloom and Van Reenen (2002) who mention a lag of two years between application and grant, and Hall *et al.* (2001a) who state that 95% of the patents that are eventually granted are granted within 3 years of

Table 3: The Distribution of Firms by Total Patents, 1963-1998 (SIC's 35-38 only)

	0 or more	1 or more	10 or more	25 or more	100 or more	250 or more	1000 or more
Firms	1920	1028	641	435	195	119	53

Table 4: Contemporaneous rank correlations between Patents and R&D expenditure

	SIC 35	SIC 36	SIC 37	SIC 38
ρ	0.4277	0.4560	0.4326	0.4591
p -value	0.0000	0.0000	0.0000	0.0000
Obs.	8533	8751	2696	7475

report positive R&D expenditure.

Table 4 shows that patent numbers are well correlated with (deflated) R&D expenditure, albeit without controlling for firm size. To take this into account, we take employees as a measure of firm size and scale down the R&D and patents measures.¹⁰ Table 5 reports the rank correlations between firm-level patent intensity and R&D intensity. For each of the sectors we observe positive and highly significant rank correlations, which nonetheless take values lower than 0.25. These results would thus appear to be consistent with the idea that, even within industries, patent and R&D statistics do contain large amounts of idiosyncratic variance and that either of these variables taken individually would be a rather noisy proxy for ‘innovativeness’.¹¹ Indeed, as discussed in Section 2, these two variables are quite different not only in terms of statistical properties (patent statistics are much more skewed and less persistent than R&D statistics) but also in terms of economic significance. However, they both yield valuable information on firm-level innovativeness.

application. However, we allow for a five-year gap here because it has been suggested that this gap has become longer in recent years.

¹⁰We also investigate the robustness of our results by scaling down a firm’s R&D and patents by its sales instead of its employees, and obtain similar results. For a brief discussion, see the Appendix (Section A).

¹¹Further evidence of the discrepancies between patent statistics and R&D statistics is presented in the regression results in Tables 5 and 6 of Coad and Rao (2006a).

Table 5: Contemporaneous rank correlations between ‘patent intensity’ (patents/employees) and ‘R&D intensity’ (R&D/employees)

	SIC 35	SIC 36	SIC 37	SIC 38
ρ	0.1631	0.2321	0.2248	0.1990
p -value	0.0000	0.0000	0.0000	0.0000
Obs.	7906	8119	2505	6935

Table 6: Extracting the ‘innovativeness’ index used for the quantile regressions - Principal Component Analysis results (first component only, unrotated)

	SIC 35	SIC 36	SIC 37	SIC 38
R&D / Employees	0.4189	0.3934	0.4136	0.4146
Patents / Employees	0.3734	0.3983	0.3806	0.3953
R&D stock / Employees ($\delta=15\%$)	0.4399	0.4106	0.4208	0.3994
Pat. stock / Employees ($\delta=15\%$)	0.3845	0.4161	0.4027	0.4147
R&D stock / Employees ($\delta=30\%$)	0.4394	0.4125	0.4248	0.4074
Pat. stock / Employees ($\delta=30\%$)	0.3882	0.4181	0.4054	0.4175
Prop ⁿ Variance explained	0.5297	0.6913	0.5040	0.8338
No. Obs.	7271	7477	2323	6394

Our synthetic ‘innovativeness’ index is created by extracting the common variance from a series of related variables: both patent intensity and R&D intensity at time t , and also the actualized stocks of patents and R&D. These stock variables are calculated using the conventional amortization rate of 15%, and also at the rate of 30% since we suspect that the 15% rate may be too low (Hall and Oriani, 2006). Information on the factor loadings is shown in Table 6. We consider that the summary ‘innovativeness’ variable is a satisfactory indicator of firm-level innovativeness in all the sectors under analysis because it loads reasonably well with the stock variables and explains between 50% to 83% of the total variance. Our composite variable has worked well in previous studies (e.g. Coad and Rao 2006a,b,c) and in this study we find that it works reasonably well. Nevertheless, we check the robustness of our results in the Appendix (Section B) by taking either a firm’s R&D stock or its patent stock as alternative indicators of ‘innovativeness’.

An advantage of this composite index is that a lot of information on a firm’s innovative activity can be summarized into one variable (this will be especially useful in the following graphs). A disadvantage is that the units have no ready interpretation (unlike ‘one patent’ or ‘\$1 million of R&D expenditure’). In this study, however, we are less concerned with the quantitative point estimates than with the qualitative variation in the importance of innovation over the conditional growth rate distribution (i.e. the ‘shape’ of the graphs).

4 Semi-parametric analysis

In this section we use semi-parametric quantile regression techniques to explore the heterogeneity between firms with regards to their innovation and employment behavior. We begin with an introduction to quantile regression before presenting the results.




Figure 3: The (annual) employment growth rates distribution for our four two-digit sectors.

4.1 An Introduction to Quantile Regression

Standard least squares regression techniques provide summary point estimates that calculate the average effect of the independent variables on the ‘average firm’. However, this focus on the average firm may hide important features of the underlying relationship. As Mosteller and Tukey explain in an oft-cited passage: “What the regression curve does is give a grand summary for the averages of the distributions corresponding to the set of x ’s. We could go further and compute several regression curves corresponding to the various percentage points of the distributions and thus get a more complete picture of the set. Ordinarily this is not done, and so regression often gives a rather incomplete picture. Just as the mean gives an incomplete picture of a single distribution, so the regression curve gives a correspondingly incomplete picture for a set of distributions” (Mosteller and Tukey, 1977:266). Quantile regression techniques can therefore help us obtain a more complete picture of the underlying relationship between innovation and employment growth.

In our case, estimation of linear models by quantile regression may be preferable to the usual regression methods for a number of reasons. First of all, we know that the standard least-squares assumption of normally distributed errors does not hold for our database because growth rates follow an exponential rather than a Gaussian distribution. The heavy-tailed nature of the growth rates distribution is illustrated in Figure 3 (see also Stanley et al. (1996) and Bottazzi and Secchi (2003) for the growth rates distribution of Compustat firms). Whilst

the optimal properties of standard regression estimators are not robust to modest departures from normality, quantile regression results are characteristically robust to outliers and heavy-tailed distributions. In fact, the quantile regression solution $\hat{\beta}_\theta$ is invariant to outliers of the dependent variable that tend to $\pm \infty$ (Buchinsky, 1994). Another advantage is that, while conventional regressions focus on the mean, quantile regressions are able to describe the entire conditional distribution of the dependent variable. In the context of this study, high growth firms are of interest in their own right, we don't want to dismiss them as outliers, but on the contrary we believe it would be worthwhile to study them in detail. This can be done by calculating coefficient estimates at various quantiles of the conditional distribution. Finally, a quantile regression approach avoids the restrictive assumption that the error terms are identically distributed at all points of the conditional distribution. Relaxing this assumption allows us to acknowledge firm heterogeneity and consider the possibility that estimated slope parameters vary at different quantiles of the conditional growth rate distribution.

The quantile regression model, first introduced in Koenker and Bassett's (1978) seminal contribution, can be written as:

$$y_{it} = x'_{it}\beta_\theta + u_{\theta it} \quad \text{with} \quad \text{Quant}_\theta(y_{it}|x_{it}) = x'_{it}\beta_\theta \quad (1)$$

where y_{it} is the dependent variable, x is a vector of regressors, β is the vector of parameters to be estimated, and u is a vector of residuals. $Q_\theta(y_{it}|x_{it})$ denotes the θ^{th} conditional quantile of y_{it} given x_{it} . The θ^{th} regression quantile, $0 < \theta < 1$, solves the following problem:

$$\min_{\beta} \frac{1}{n} \left\{ \sum_{i,t: y_{it} \geq x'_{it}\beta} \theta |y_{it} - x'_{it}\beta| + \sum_{i,t: y_{it} < x'_{it}\beta} (1 - \theta) |y_{it} - x'_{it}\beta| \right\} = \min_{\beta} \frac{1}{n} \sum_{i=1}^n \rho_\theta u_{\theta it} \quad (2)$$

where $\rho_\theta(\cdot)$, which is known as the 'check function', is defined as:

$$\rho_\theta(u_{\theta it}) = \begin{cases} \theta u_{\theta it} & \text{if } u_{\theta it} \geq 0 \\ (\theta - 1)u_{\theta it} & \text{if } u_{\theta it} < 0 \end{cases} \quad (3)$$

Equation (2) is then solved by linear programming methods. As one increases θ continuously from 0 to 1, one traces the entire conditional distribution of y , conditional on x (Buchinsky, 1998). More on quantile regression techniques can be found in the surveys by Buchinsky (1998) and Koenker and Hallock (2001); for some applications see the special issue of *Empirical Economics* (Vol. 26 (3), 2001).

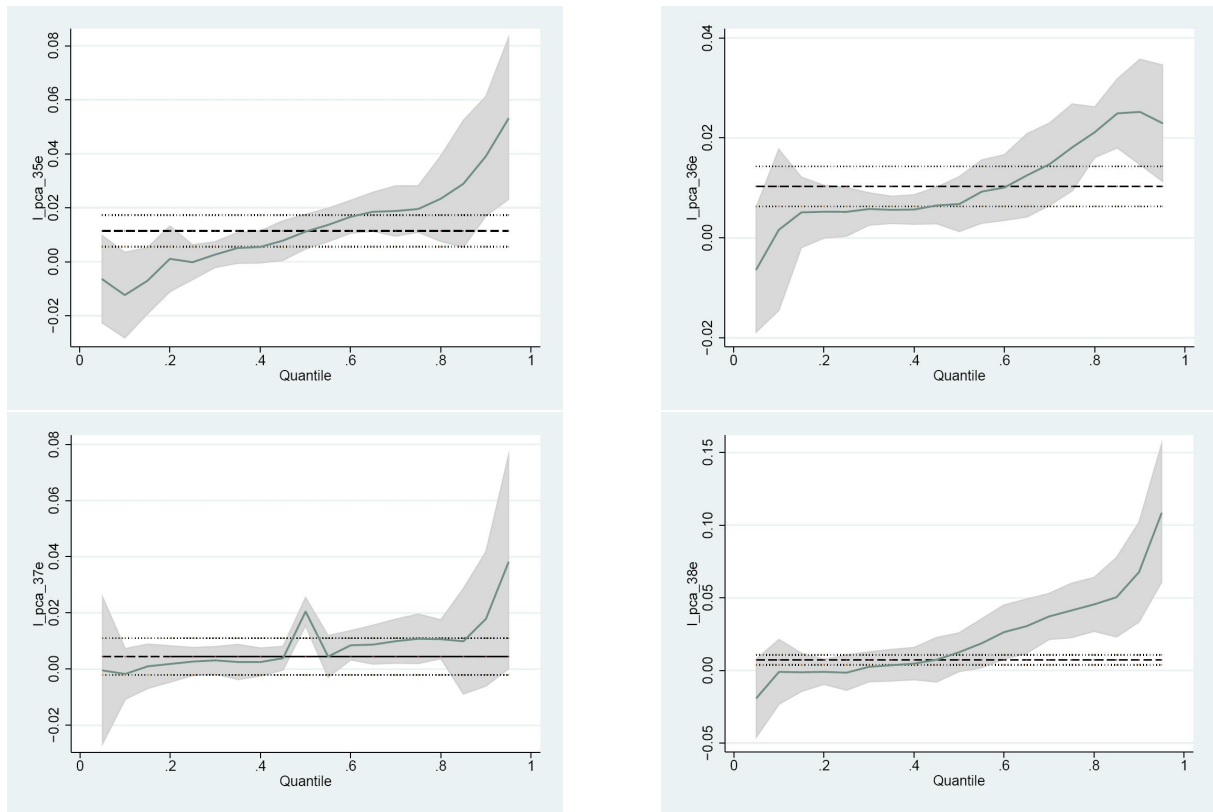


Figure 4: Variation in the coefficient on ‘innovativeness’ (i.e. β_1 in Equation (4)) over the conditional quantiles. Confidence intervals extend to 2 standard errors in either direction. Horizontal lines represent OLS estimates with 95% confidence intervals. SIC 35: Machinery & Computer Equipment (top-left), SIC 36: Electric/Electronic Equipment (top-right), SIC 37: Transportation Equipment (bottom-left), SIC 38: Measuring Instruments (bottom-right). Graphs made using the ‘grqreg’ Stata module (Azevedo, 2004).

4.2 Quantile regression results

We now apply quantile regression to estimate the following linear regression model:

$$GROWTH_{i,t} = \alpha + \beta_1 INN_{i,t-1} + \beta_2 CONTROL_{i,t-1} + y_t + \epsilon_{i,t} \quad (4)$$

where *INN* is the ‘innovativeness’ variable for firm *i* at time *t*. *CONTROL* includes all of the control variables that may potentially influence a firm’s employment growth,¹² namely, lagged growth, lagged size and 3-digit industry dummies. We also control for common macroeconomic shocks by including year dummies (y_t).

The regression results for each of the four 2-digit sectors can be seen in Figure 4 (see also Table 7). We observe considerable variation in the regression coefficient over the conditional quantiles. At the upper quantiles, the coefficient is observed to increase. This means that innovation has a strong positive impact on employment for those firms that have the fastest employment growth. At the lower quantiles, however, the coefficient on our ‘innovativeness’ variable often becomes negative (although not statistically significant), which indicates that innovation is associated with job destruction for those firms that are losing the most jobs.

To sum up, it may be useful to distinguish between three groups of firms. First of all, the ‘average firm’ stays at roughly the same size. Such firms do not change their employment levels by much, and furthermore innovation seems to have little effect on their employment decisions. This is indicated by the fact that the coefficient on ‘innovativeness’ is close to zero at the median quantile. The second group consists of those fast-growing firms that are experiencing the largest increases in employment. For these firms, innovation has a strong positive effect on employment. The third group contains firms that are losing the most jobs. In this case, increases in firm-level innovative activity are associated with subsequent reductions in employment. This could be due to two effects, however. On the one hand, it could be due to innovation leading to a reduction in the required labour inputs (this effect is the *bona fide* ‘technological unemployment’ argument). On the other hand, it could be because some firms are unsuccessful in their attempts at innovation. This is the ‘tried and failed’ category of innovators described in Freel (2000) and discussed in Coad and Rao (2006b). We suspect that both of these effects are present for this third group of firms.

In the Appendix (Section B), we check the robustness of our results by using alternative (cruder) measures of firm-level innovation. These measures are 3-year R&D and patent stocks, depreciated at the conventional rate of 15%. As expected, these two variables taken on their own are less clear-cut than our preferred composite ‘innovativeness’ variable. Broadly speaking, however, the results from this exercise appear to support our main results presented in this section.

¹²For a survey of firm growth, see Coad (2007).

We have thus observed that in some cases innovation is associated with employment creation whilst in other cases it is associated with job destruction. It is of interest to see if these two categories are correlated with a firm's size. For example, we could suspect that the latter category corresponds to the largest firms who are more likely to introduce process innovations. Previous studies have not been able to test this hypothesis because they implicitly attribute equal weights to firms of different sizes. We argue that this approach is flawed, however, given that larger firms have a greater impact on the absolute number of jobs, because of their large size. We investigate this issue in the next section.

5 Parametric analysis

We begin this section with a brief introduction to the weighted least squares estimator, and then apply it to our dataset.

5.1 An Introduction to Weighted Least Squares

“As Mosteller and Tukey (1977, p346) suggested, the action of assigning “different weights to different observations, either for objective reasons or as a matter of judgement” in order to recognize “some observations as ‘better’ or ‘stronger’ than others” has an extensive history.”

Willett and Singer (1988:236)

Consider the regression equation:

$$y_i = \beta x_i + \epsilon_i \tag{5}$$

The OLS regression solution seeks to minimize the sum of the squared residuals, i.e.:

$$\min Q = \sum_{i=1}^n (y_i - \beta x_i)^2 \equiv \sum_{i=1}^n (\epsilon_i)^2 \tag{6}$$

Implicit in the basic OLS solution is that the observations are treated as equally important, being given equal weights. Weighted Least Squares, however, attributes weights w_i to specific observations that determine how much each observation influences the final parameter estimates:

$$\min Q = \sum_{i=1}^n w_i (y_i - \beta x_i)^2 \tag{7}$$

It follows that WLS estimators are functions of the weights w_i .

Although WLS can be used in situations where observations are attributed different levels of ‘importance’, it is most often used for dealing with heteroskedasticity. In the context of this study, the weight w_i corresponds to the firm’s size, measured in terms of employees.

5.2 Regression results

We estimate equation (4) using conventional estimators such as OLS and the Fixed-Effect estimator, as well as the Weighted Least Squares estimator: the results are presented in Table 7.

The main feature of the regression results is that the coefficients obtained from the standard OLS and fixed-effect (FE) estimators are positive (though not always significant) for each of the four sectors. Furthermore, we observe that the R^2 coefficients are rather low, always lower than 7%. The standard interpretation of these results would be that, if anything, innovation seems to be positively associated with subsequent employment growth. However, our preferred interpretation of these results is informed by the quantile regression analysis presented in the preceding section. We observed that, for the fast-growing firms, innovation is positively associated with employment, whilst increases in innovation may also be associated with job destruction for those firms shedding the most jobs. This heterogeneity is indeed masked by standard regression techniques that focus on ‘the average effect for the average firm’.

We also observe that the WLS coefficient estimates are in most cases higher than the results obtained by either OLS or FE. This evidence hints that innovation in large firms is more likely to be associated with employment creation than innovation in small firms. This is an interesting finding given that larger firms have a greater potential for large increases in the absolute number of new jobs. In addition, this result is perhaps surprising given that large firms are usually associated with process innovations (see for example Klepper, 1996) and process innovations, in turn, are usually classified as labour-saving.

6 Conclusion

Our main results are twofold.

Our first main result emerges when we apply semi-parametric quantile regressions to explore the relationship between innovation and employment growth. We observe three categories of firms. First, most firms do not grow by much, and what little they do grow seems to be unrelated to innovation. Second, those firms that experience rapid employment growth owe a large amount of this to their previous attempts at innovation. Third, for those firms that are shedding the most jobs, increases in innovative activity seem to be associated with job destruction. The distinction between these three categories is effectively masked whenever

Table 7: Regression estimation of equation (4). Quantile regression estimates obtained using 1000 bootstrap replications.

	Quantile regression					OLS	FE	WLS
	10%	25%	50%	75%	90%			
SIC 35								
β_1	-0.0124	-0.0003	0.0111	0.0196	0.0392	0.0114	0.0195	0.0261
Std. Error	0.0082	0.0043	0.0035	0.0046	0.0124	0.0035	0.0057	0.0086
t -stat	-1.51	-0.06	3.11	4.24	3.16	3.23	3.42	3.03
R^2 within							0.0424	
R^2 between							0.0001	
R^2 overall	0.0723	0.0582	0.0479	0.0735	0.0894	0.0599	0.0273	0.1989
obs (groups)							601	
obs	6682	6682	6682	6682	6682	6682	6682	6682
SIC 36								
β_1	0.002	0.0052	0.0067	0.0179	0.0255	0.0103	0.0153	0.0282
Std. Error	0.0063	0.002	0.0024	0.004	0.0044	0.0024	0.0056	0.0057
t -stat	0.32	2.58	2.8	4.5	5.73	4.29	2.75	4.96
R^2 within							0.043	
R^2 between							0.0005	
R^2 overall	0.0429	0.041	0.0361	0.0487	0.048	0.0479	0.018	0.1427
obs (groups)							614	
obs	6891	6891	6891	6891	6891	6891	6891	6891
SIC 37								
β_1	-0.0017	0.0024	0.0038	0.0096	0.0179	0.0043	0.0149	0.0149
Std. Error	0.0176	0.0031	0.0026	0.0035	0.0114	0.005	0.008	0.0056
t -stat	-0.1	0.75	1.47	2.72	1.56	0.85	1.86	2.67
R^2 within							0.0548	
R^2 between							0.0048	
R^2 overall	0.1036	0.0787	0.065	0.0617	0.0716	0.0685	0.0261	0.2417
obs (groups)							178	
obs	2154	2154	2154	2154	2154	2154	2154	2154
SIC 38								
β_1	-0.0011	-0.0008	0.0125	0.041	0.0688	0.0073	0.0136	0.0044
Std. Error	0.011	0.0058	0.0086	0.0099	0.0191	0.0074	0.0107	0.0068
t -stat	-0.1	-0.14	1.45	4.13	3.6	0.99	1.27	0.65
R^2 within							0.0329	
R^2 between							0.0884	
R^2 overall	0.0459	0.0288	0.028	0.0507	0.065	0.0284	0.0049	0.1247
obs (groups)							527	
obs	5870	5870	5870	5870	5870	5870	5870	5870

Note: The fact that the WLS R^2 is higher than the OLS or FE R^2 may simply be a spurious statistical result (Willett and Singer, 1988).

conventional parametric regressions are used, because these latter focus on ‘the average effect for the average firm’ and are relatively insensitive to heterogeneity between firms.

Our second main result is observed when we investigate whether the relationship between innovation and employment varies with firm size. Our previous observations on the heterogeneity of firm behavior vis-à-vis innovation and employment effectively fuelled such suspicions. Our results indicate that, if anything, innovative activity in large firms is more positively associated with employment growth than innovative activity undertaken by their smaller counterparts.

We should mention the limitations of our results that are brought on by the specificities of our dataset. In the US, the labour market is more fluid than in other countries, and this may reduce the generality of our results. Furthermore, we focus only on high-tech manufacturing sectors. Although this particular sectoral focus allows us to get relatively accurate measures of firm-level innovation, it reduces the scope of our analysis. It may be the case that the relationship between innovation and unemployment is different for other sectors of the economy.

APPENDICES

A Scaling down according to firm size

There are at least two ways of scaling down indicators of innovative activity according to firm size (Small and Swann, 1993). The first, and perhaps most common, way is to use a firm's sales as an indicator of its size. The second involves scaling down according to a firm's employment.

Our analysis in this paper uses the second approach, which some authors have nonetheless identified as the preferable method.¹³ However, we also investigate the robustness of our analysis by scaling down according to total sales. Table 8 contains the corresponding results for the generation of our composite 'innovativeness' indicator. We observe that this indicator does not appear to perform as well when we scale down innovative activity by a firm's total sales (compare the results here with those in Coad and Rao (2006b,c)). We nonetheless pursue the analysis using this indicator, and we obtain similar results (see Table 9).

B Alternative measures of innovative activity

In this section we verify the robustness of the quantile regression results presented in Section 4 by using simpler and cruder measures of firm-level innovative activity.

We now estimate the following linear regression model:

$$GROWTH_{i,t} = \alpha + \gamma_1 INNOV_{i,t-1} + \beta_2 GROWTH_{i,t-1} + \beta_3 SIZE_{i,t-1} + \beta_4 IND_{i,t} + y_t + \epsilon_{i,t} \quad (8)$$

where $INNOV_{i,t-1}$ refers to either a firm's 3-year stock of R&D intensity (i.e. R&D / Sales) or patent intensity (Patents / Sales); the conventional depreciation rate of 15% has been used for both of these variables. The results are presented in Figures 5 and 6. Broadly speaking, these results offer support to our earlier analysis. In general, we observe that the coefficient is close to zero at the median quantile. The coefficient decreases at the very lowest quantiles, often taking on a negative coefficient. In contrast, the coefficient becomes increasingly positive at the upper quantiles.

¹³Scherer (1965) discusses the possibility of scale measurement errors entering into various firm-level data. Although he is unable to verify the hierarchy of these errors, he speculates that the measurement problems are likely to be larger for assets followed by sales and (to a lesser extent) employment (Scherer 1965: 259).

Table 8: Extracting the ‘innovativeness’ index used for the quantile regressions – Principal Component Analysis results (first component only, unrotated)

	SIC 35	SIC 36	SIC 37	SIC 38
R&D / Sales	0.1631	0.1351	0.3076	0.0302
Patents / Sales	0.2669	0.1239	0.4294	0.1614
R&D stock / Sales ($\delta=15\%$)	0.4628	0.4945	0.3530	0.4645
Pat. stock / Sales ($\delta=15\%$)	0.4840	0.4958	0.4830	0.5199
R&D stock / Sales ($\delta=30\%$)	0.4659	0.4888	0.3540	0.4653
Pat. stock / Sales ($\delta=30\%$)	0.4865	0.4870	0.4877	0.5200
Prop ⁿ Variance explained	0.5031	0.6155	0.4752	0.3762
No. Obs.	7858	8079	2559	6940




Figure 5: Variation in the coefficient on a firm’s 3-year R&D stock (i.e. γ_1 in Equation (8)) over the conditional quantiles. Confidence intervals extend to 2 standard errors in either direction. Horizontal lines represent OLS estimates with 95% confidence intervals. SIC 35: Machinery & Computer Equipment (top-left), SIC 36: Electric/Electronic Equipment (top-right), SIC 37: Transportation Equipment (bottom-left), SIC 38: Measuring Instruments (bottom-right). Graphs made using the ‘grqreg’ Stata module (Azevedo, 2004).

Table 9: Regression estimation of equation (4). Note that the quantile regression SEs have not been bootstrapped here.

	quantile regression					OLS	FE	WLS
	10%	25%	50%	75%	90%			
SIC 35								
β_1	-0.03868	-0.00390	0.00225	0.00114	0.01109	-0.00127	0.00276	0.02374
Std. Error	0.00154	0.00142	0.00066	0.00066	0.00145	0.00346	0.00494	0.01522
t -stat	-25.10	-2.74	3.40	1.71	7.64	-0.37	0.56	1.56
R^2 within							0.0401	
R^2 between							0.0006	
R^2 overall	0.06880	0.05350	0.04020	0.06060	0.07350	0.05360	0.02260	0.16560
obs (groups)							661	
obs	7273	7273	7273	7273	7273	7273	7273	7273
SIC 36								
β_1	-0.01129	-0.00984	-0.00063	0.01048	0.01895	-0.00240	-0.00235	0.02051
Std. Error	0.00178	0.00091	0.00049	0.00085	0.00123	0.00276	0.00279	0.01403
t -stat	-6.33	-10.84	-1.29	12.28	15.45	-0.87	-0.84	1.46
R^2 within							0.0352	
R^2 between							0.0050	
R^2 overall	0.04190	0.0368	0.0332	0.0436	0.0400	0.0408	0.0135	0.1333
obs (groups)							614	
obs	7495	7495	7495	7495	7495	7495	7495	7495
SIC 37								
β_1	0.00479	0.00148	0.00304	0.00673	0.02609	0.00658	0.00362	0.02215
Std. Error	0.00376	0.00129	0.00118	0.00165	0.00269	0.00379	0.00490	0.01533
t -stat	1.28	1.14	2.57	4.06	9.72	1.74	0.74	1.44
R^2 within							0.0588	
R^2 between							0.0046	
R^2 overall	0.0841	0.0750	0.0594	0.0529	0.0661	0.0659	0.0233	0.2213
obs (groups)							178	
obs	2389	2389	2389	2389	2389	2389	2389	2389
SIC 38								
β_1	0.00038	-0.00128	0.00540	0.00347	0.00121	0.00173	0.00164	0.00174
Std. Error	0.00110	0.00061	0.00063	0.00117	0.00122	0.00268	0.00333	0.00276
t -stat	0.35	-2.11	8.52	2.97	0.99	0.64	0.49	0.63
R^2 within							0.0261	
R^2 between							0.0502	
R^2 overall	0.04030	0.02570	0.02560	0.03900	0.03940	0.02510	0.00610	0.11840
obs (groups)							527	
obs	6421	6421	6421	6421	6421	6421	6421	6421

Note: The fact that the WLS R^2 is higher than the OLS or FE R^2 may simply be a spurious statistical result (Willett and Singer, 1988).

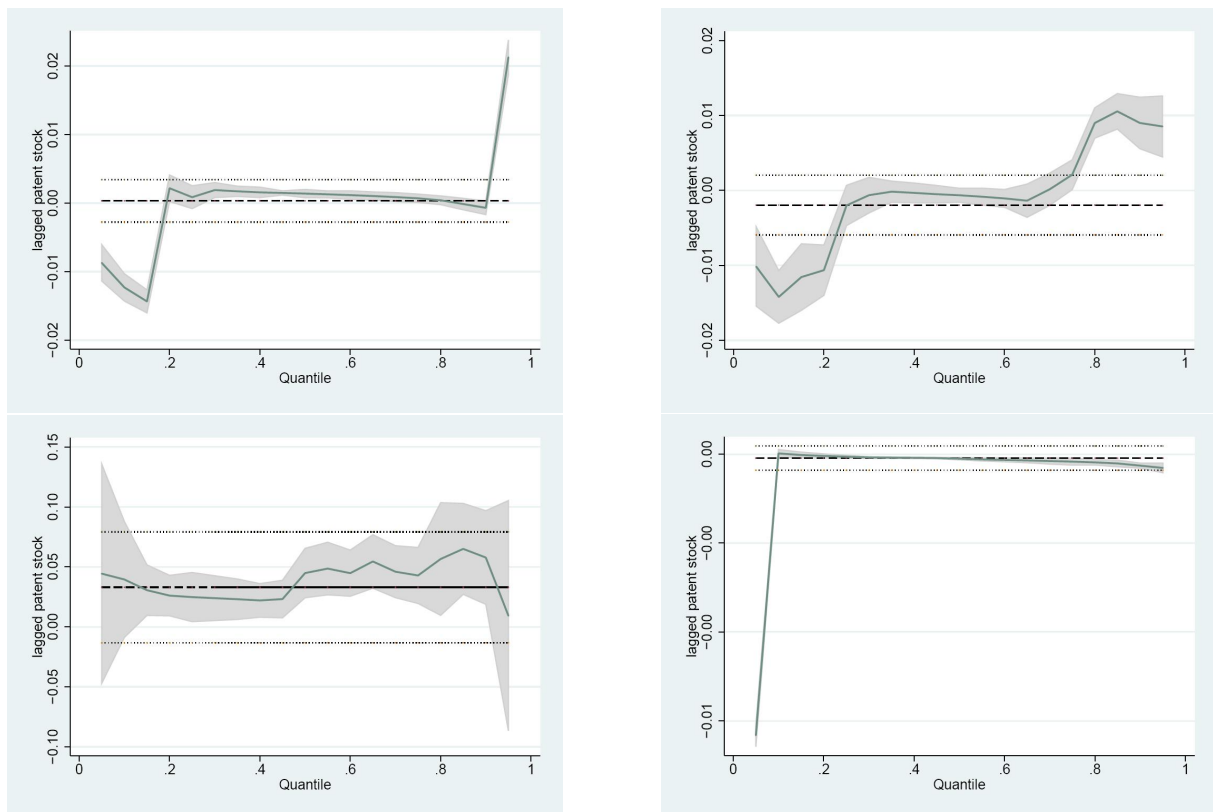


Figure 6: Variation in the coefficient on a firm's 3-year patent stock (i.e. γ_1 in Equation (8)) over the conditional quantiles. Confidence intervals extend to 2 standard errors in either direction. Horizontal lines represent OLS estimates with 95% confidence intervals. SIC 35: Machinery & Computer Equipment (top-left), SIC 36: Electric/Electronic Equipment (top-right), SIC 37: Transportation Equipment (bottom-left), SIC 38: Measuring Instruments (bottom-right). Graphs made using the 'grqreg' Stata module (Azevedo, 2004).

References

- Archibugi, D., 1992. Patenting as an indicator of technological innovation: a review. *Science and Public Policy* 19 (6), 357-368.
- Arundel, A., Kabla, I., 1998. What percentage of innovations are patented? Empirical estimates for European firms. *Research Policy* 27, 127-141.
- Azevedo, J.P.W., 2004. *grqreg*: Stata module to graph the coefficients of a quantile regression. Boston College Department of Economics.
- Bloom, N., Van Reenen, J., 2002. Patents, Real Options and Firm Performance. *Economic Journal* 112, C97-C116.
- Bottazzi, G., Coad, A., Jacoby, N., Secchi, A., 2005. Corporate Growth and Industrial Dynamics: Evidence from French Manufacturing. Pisa, Sant'Anna School of Advanced Studies, LEM Working Paper Series 2005/21.
- Bottazzi, G., Secchi, A., 2003. Common Properties and Sectoral Specificities in the Dynamics of U. S. Manufacturing Companies. *Review of Industrial Organization* 23, 217-232.
- Buchinsky, M., 1994. Changes in the U.S. Wage Structure 1963-1987: Application of Quantile Regression. *Econometrica* 62, 405-458.
- Buchinsky, M., 1998. Recent Advances in Quantile Regression Models: A Practical Guide for Empirical Research. *Journal of Human Resources* 33 (1), 88-126.
- Cefis, E., Orsenigo, L., 2001. The persistence of innovative activities: A cross-countries and cross-sectors comparative analysis. *Research Policy* 30, 1139-1158.
- Coad, A., (2007). 'Firm Growth: A Survey', Papers on Economics and Evolution 2007-03, Max Planck Institute of Economics, Evolutionary Economics Group, Jena, Germany.
- Coad, A., Rao, R., (2006a). Innovation and Firm Growth in 'Complex Technology' Sectors: A Quantile Regression Approach. *Cahiers de la Maison des Sciences Economiques* No. 06050 (Série Rouge), Université Paris 1 Panthéon-Sorbonne, France.
- Coad, A. and R. Rao (2006b), 'Innovation and Firm Growth in High-Tech Sectors: A Quantile Regression Approach', Pisa, Sant'Anna School of Advanced Studies, LEM Working Paper Series 2006/18.
- Coad, A., Rao, R., (2006c), 'Innovation and market value: a quantile regression analysis' *Economics Bulletin* 15 (13), 1-10.

- Cohen, W.M., Nelson, R.R., Walsh, J.P., 2000. Protecting their intellectual assets: Appropriability conditions and why US manufacturing firms patent (or not). NBER working paper 7552.
- Dosi, G., 1988. Sources, Procedures, and Microeconomic Effects of Innovation. *Journal of Economic Literature* 26 (3), 1120-1171.
- Evangelista, R., Savona, M., 2002. The Impact of Innovation on Employment in Services: evidence from Italy. *International Review of Applied Economics* 16 (3), 309-318.
- Evangelista, R., Savona, M., 2003. Innovation, Employment and Skills in Services: Firm and Sectoral Evidence. *Structural Change and Economic Dynamics* 14 (4), 449-74.
- Fleck, J., 1984. The Adoption of Robots in Industry. *Physics in Technology* 15, 4-11.
- Geroski, P.A., 2000. The growth of firms in theory and practice, in: Foss, N., Mahnke, V., (Eds.), *New Directions in Economic Strategy Research*. Oxford University Press, Oxford
- Greenhalgh, C., Longland, M., Bosworth, D., 2001. Technological Activity and Employment in a Panel of UK Firms. *Scottish Journal of Political Economy* 48 (3), 260-282.
- Griliches, Z., 1990. Patent Statistics as Economic Indicators: A Survey. *Journal of Economic Literature* 28, 1661-1707.
- Hall, B.H., 2004. Exploring the Patent Explosion. *Journal of Technology Transfer* 30 (1-2), 35-48.
- Hall, B.H., Jaffe, A., Trajtenberg, M., 2001a. Market Value and Patent Citations: A First Look. Paper E01-304, University of California, Berkeley.
- Hall, B.H., Jaffe, A., Trajtenberg, M., 2001b. The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools. NBER Working Paper 8498.
- Hall, B. H., Lotti, F., Mairesse J., (2006) 'Employment, Innovation, and Productivity: Evidence from Italian Microdata', mimeo, December 5.
- Hall, B.H., Oriani, R., 2006. Does the market value R&D investment by European firms? Evidence from a panel of manufacturing firms in France, Germany, and Italy. *International Journal of Industrial Organization*, in press.
- Hair, J., Anderson, R., Tatham, R., Black, W., 1998. *Multivariate Data Analysis: Fifth Edition*. Prentice Hall, Upper Saddle River, NJ.

- Harrison, R., Jaumandreu, J., Mairesse, J., Peters, B., (2005), ‘Does innovation stimulate employment? A firm-level analysis using comparable micro data on four European countries’, mimeo, Universidad Carlos III Madrid, February.
- Klepper, S., 1996. Entry, Exit, Growth, and Innovation over the Product Life Cycle. *American Economic Review* 86 (3), 562-583.
- Koenker, R., Bassett, G., 1978. Regression Quantiles. *Econometrica* 46, 33-50.
- Koenker, R., Hallock, K.F., 2001. Quantile Regression. *Journal of Economic Perspectives* 15 (4), 143-156.
- Lanjouw, J., Schankerman, M., 2004. Patent Quality and Research Productivity: Measuring Innovation with Multiple Indicators. *Economic Journal* 114, 441-465.
- Mansfield, E., Rapoport, J., Romeo, A., Villani, E., Wagner, S., Husic, F., 1977. *The Production and Application of New Industrial Technology*. Norton, New York.
- Mosteller, F., Tukey, J., 1977. *Data Analysis and Regression*. Addison-Wesley, Reading, MA.
- Niefert, M., 2005. Patenting Behavior and Employment Growth in German Start-up Firms: A Panel Data Analysis. Discussion Paper No 05-03, ZEW Centre for European Economic Research, Mannheim.
- OECD, 1994. Using Patent Data as Science and Technology Indicators. Patent Manual 1994, Paris.
- Scherer, F.M., 1965. Size of Firm, Oligopoly, and Research: A Comment. *Canadian Journal of Economics and Political Science / Revue canadienne d'Economie et de Science politique*, 31 (2), 256-266.
- Small, I., Swann, P., 1993. R&D performance of UK Companies. *Business Strategy Review*, 4 (3), 41-51.
- Stanley, M.H.R., Amaral, L.A.N., Buldyrev, S.V., Havlin, S., Leschhorn, H., Maass, P., Salinger, M.A., Stanley, H.E., 1996. Scaling behavior in the growth of companies. *Nature* 379, 804-806.
- Van Reenen, J., 1997. Employment and Technological Innovation: Evidence from UK Manufacturing Firms. *Journal of Labor Economics* 15 (2), 255-284.
- Willett, J.B., Singer, J.D., 1988. Another Cautionary Note About R^2 : Its Use in Weighted Least-Squares Regression Analysis. *The American Statistician* 42 (3), 236-238.