The impact of R&D grants on the performance of New Zealand firms

Research Note 2017/5

October 2017

Author: Simon Wakeman

New Zealand Productivity Commission Research Note 2017/5: The impact of R&D grants on the performance of New Zealand firms

Date: October 2017

Author: Simon Wakeman

JEL classification: O31, O38

ISBN: 978-0-478-44039-3

Acknowledgements: I would like to thank Sarah Holden, Kirsten Keene, Sara McFall, Scott Wombwell, Donal Krouse, Alice Feng, and Chris Barnett for working together on the project that led to this paper and the various insights on the R&D grants programme and the data that they shared with me in the process; Guanyu Zheng for his assistance with the econometric work and understanding of the LBD; Trinh Le for sharing her code for propensity-score matching; members of the Microdata Access Team at Statistics New Zealand for support in using the data; Richard Fabling and Dave Mare for sharing their firm-level estimates of productivity; Lisa Meehan, Lynda Sanderson, and Simon Rae for comments on an earlier version of this work; and especially Paul Conway for his guidance and numerous rounds of feedback on this paper throughout the process. I retain responsibility for any errors.

Disclaimer: The opinions, findings, recommendations, and conclusions expressed in this paper are those of the author, not Statistics New Zealand or the New Zealand Productivity Commission.

The results in this paper have been created for research purposes from the Integrated Data Infrastructure (IDI) managed by Statistics New Zealand and are not official statistics. Access to the anonymised data used in this study was provided by Statistics New Zealand in accordance with security and confidentiality provisions of the Statistics Act 1975. Only people authorised by the Statistics Act 1975 are allowed to see data about a particular person, household, business or organisation and the results in this paper have been confidentialised to protect these groups from identification. Careful consideration has been given to the privacy, security and confidentiality issues associated with using administrative and survey data in the IDI. Further detail can be found in the privacy impact assessment for the IDI available from www.stats.govt.nz.

The results are based in part on tax data supplied by Inland Revenue to Statistics New Zealand under the Tax Administration Act 1994. This tax data must be used only for statistical purposes, and no individual information may be published or disclosed in any other form, or provided to Inland Revenue for administrative or regulatory purposes. Any person who has had access to the unit-record data has certified that they have been shown, have read, and have understood section 81 of the Tax Administration Act 1994, which relates to secrecy. Statistics NZ confidentiality protocols were applied to the data sourced from the Ministry of Business, Innovation and Employment. Any discussion of data limitations or weaknesses is in the context of using the IDI for statistical purposes, and is not related to the data's ability to support these government agencies' core operational requirements.

Information on the Productivity Commission can be found on www.productivity.govt.nz or by contacting +64 4 903 5150.

Abstract

This paper presents the results from an evaluation of the impact of New Zealand Government R&D grants on the performance of New Zealand firms. The analysis uses information on grants and firm performance from 2004 to 2012 available in Statistics New Zealand's Longitudinal Business Database. Grant recipients are matched to non-recipients that had a similar propensity to receive a grant. The performance of these two groups of firms is then compared across a range of measures including R&D spending, innovation activity, employment, output, and productivity growth. As the available data precedes the creation of Callaghan Innovation in 2013, this paper does not directly evaluate the performance of Callaghan Innovation's R&D grants programme. Nevertheless it illustrates the range of insights that can be gained by assessing the impact of government programmes using the LBD, and highlights both the strengths and limitations of using the LBD to evaluate the impact of government interventions.

Contents

Tables

Figures

1 Introduction

Investment in knowledge-based capital (KBC) is a major driver of economic performance (OECD, 2013). Although relatively little is known about how much investment in KBC drives economic performance in the New Zealand context, de Serres, Yashiro, and Boulhol (2014) use low business expenditure on R&D (BERD) to argue that weak investment in KBC could account for up to 40% of New Zealand's productivity gap relative to the OECD average. Partly in response to this belief, New Zealand's innovation policy in recent years has had a strong focus on raising BERD, primarily through direct grants to firms that invest in R&D.

This paper evaluates the impact of R&D grants on the performance of New Zealand firms. It draws on data in the Longitudinal Business Database (LBD), which is compiled by Statistics New Zealand from a range of administrative sources including company tax filings, responses to business surveys, and records of government assistance. The paper does not directly evaluate the performance of the R&D grant scheme run by Callaghan Innovation as the necessary data is not yet available.[1] Nevertheless, this work does provide an example of how the impact of receiving an innovation grant on firm performance can be assessed. More broadly, it also illustrates the potential for using LBD to assess the impact of government interventions at the firm level.

An important issue in evaluating the impact of the R&D grants scheme is untangling the difference in firm outcomes that occurs as the result of receiving the grant from that which would occur without it. To address this issue, this paper uses the propensity-score matching method to select a set of non-recipients that are not systematically different from grant recipients, and compare the performance outcomes of the two groups.

This paper builds on two earlier evaluations of the New Zealand Government's R&D assistance programme. A paper by researchers in the Evaluation and Research teams of the Ministry of Economic Development (2011) examined the impact of R&D grants awarded between 2002 and 2008 on the economic performance of firms up to four years after the grant.[2] These researchers split grants into "capability building" and project grants, and applied a propensity-score matching approach similar to the one used in this paper. They found that firms which received capability-building assistance – generally smaller grants – experienced higher growth in employment and sales in the year after the grant (by 6 and 8 percentage points respectively), and higher growth in multifactor productivity (MFP) four years after the grant (by 13 percentage points over four years). However, counterintuitively, they found no impact from project grants – the larger grants – on any of their measures of firm performance.[3] Pooling the capability and project grants together, they found that productivity growth for small firms was 20 percentage points higher 4 years after the grants but was no different for large firms or for firms already engaged in R&D.

Jaffe and Le (2015) evaluated the impact of the New Zealand Government's R&D grants programme on the innovation outcomes of recipient firms. Like MED (2011) and this paper, they focused on the series of grant programmes that existed prior to Callaghan Innovation and used a propensity-score matching approach. They found that receiving a grant increases the probability of a firm introducing both new-to-the-world and new-to-the-firm goods & services by around 10 percentage points and introducing a new process by 5 percentage points. Receiving a grant also increases the probability of patenting from 1 percent to 2 percent. They did not find any differential effects between firms based on size. However, they examined only the impact on innovation outcomes – patent filings with the Intellectual Property Office of New Zealand (IPONZ) and innovation activity reported in the Business Operations Survey (BOS) – and did not look at the impact on measures of firm performance.

[1] The latest available data is from 2012 while Callaghan Innovation was created in 2013.

[2] The paper does not clearly state the time period of which it is able to measure the impact of the R&D grants. However, due to the time lag in data being made available to researchers, it would appear that the latest data is for 2008.

[3] The result is counterintuitive as the project grants were generally larger amounts and therefore it was expected would have a greater impact.

2 Empirical analysis

2.1 Data sources

The data used in this analysis is drawn from the Longitudinal Business Database (LBD), which is compiled by Statistics New Zealand from a range of administrative data sources and surveys. The LBD includes information on government assistance (including R&D grants), self-reported measures of R&D expenditure and innovation (e.g., introducing new goods & services) collected via the R&D Survey and the Business Operations Survey (BOS), patent and trademark registration data from the Intellectual Property Office of New Zealand (IPONZ), and financial data. In addition, it also includes derived estimates of employment, capital stock, intermediate inputs, and gross output that were constructed by Fabling and Maré (2015). Figure 2.1 summarises the sources used for this analysis.

Figure 2.1 Data sources feeding into the Longitudinal Business Database

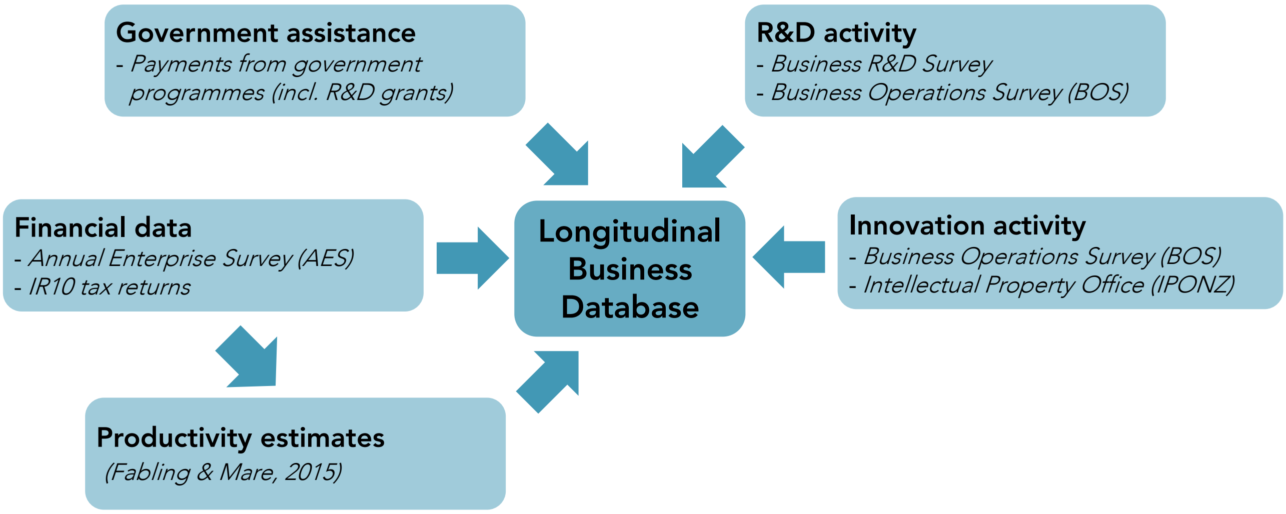

The availability of data used in the study depends on the specific characteristics of the respective datasets. For instance, financial data is only available if the firm either filed an IR10 form with the Inland Revenue Department (IRD) or responded to the Annual Enterprise Survey.[4] The availability of productivity estimates depends on data for all the various components of the production function (described below) being available from at least one of these sources for a given year.

As Table 2.1 shows, financial data is available for around 50% of all firms, but the data necessary to estimate productivity is only available for about 40% of all firms. Nevertheless, as Project grant recipients tend to be larger firms, the data necessary to estimate productivity is more likely to be available for grant recipients (for around 60% of the subset). Similarly, while only a very small proportion of all firms respond to the R&D Survey or the Business Operations Survey in any given year, a much larger proportion of Project grant recipients are covered by one or both of these surveys.

[4] All New Zealand-resident firms must declare their key financial data to IRD, but they may choose either to file an IR10 or to submit complete copies of their annual accounts. Only the IR10 data is available in the LBD. In general, smaller firms are more likely to file an IR10 form while larger firms are more likely to file their annual accounts. Meanwhile, Statistics New Zealand surveys a subset of firms in the Annual Enterprise Survey, but samples a greater proportion of larger firms so financial data for these firms is more likely to be available through the AES.

Table 2.1 Data coverage of all firms vs grant recipients by data source

	All firms in the LBD dataset		Firms with IR10 or AES data		Firms with productivity estimates		Firms with IPO data		Firms with R&D Survey data		Firms with BOS data	
Year	All	Project grant recipients (%)	All (%)	Project grant recipients (%)	All (%)	Project grant recipients (%)	All (%)	Project grant recipients (%)	All (%)	Project grant recipients (%)	All (%)	Project grant recipients (%)
2005	485022	405	49.6	60.7	37.9	58.5	0.5	17.0	-	-	1.4	20.7
2006	496545	387	49.1	62.8	37.3	59.7	0.5	18.6	0.5	58.9	1.1	17.8
2007	507762	426	48.8	59.9	36.7	56.3	0.5	21.8	-	-	1.2	19.0
2008	514962	348	49.6	56.9	36.9	53.4	0.5	19.8	0.5	57.8	1.2	17.2
2009	511119	261	49.7	62.1	36.5	58.6	0.5	23.0	-	-	1.2	17.2
2010	501741	174	49.8	58.6	36.0	56.9	0.5	25.9	0.5	58.6	1.2	24.1
2011	495504	213	50.1	60.6	36.1	54.9	-	-	-	-	1.2	21.1
2012	445440	237	51.5	60.8	37.7	55.7	-	-	0.6	62.0	1.2	21.5

Notes: This table shows the number of firms in the full LBD sample in year against the number of firms with the data used to generate measures used in the analysis. Observation counts rounded to base 3.

Receiving an R&D grant

The New Zealand Government has provided R&D assistance to firms through a range of schemes, and the names and granting agencies for these programmes have changed over time. Under Callaghan Innovation, created in 2013, the grant types have been simplified into three broad programmes: Growth, Project, and Student. Table 2.2 presents summary information on the number and total amount of grants of each type awarded in 2015.

Table 2.2 R&D grant types under Callaghan Innovation in 2015

Type	Description	Value (#) in 2015
Growth	Covers 20% of R&D costs up to $5 million a year. They are available to businesses that invest over 1.5% of turnover in R&D	$134,927,861 (85)
Project	Covers up to 50% of R&D costs and are awarded primarily to businesses undertaking research for the first time	$24,114,907 (302)
Student	Help businesses access undergraduate and postgraduate students who can assist with R&D projects	$6,402,790 (280)

Notes: Growth grants are available to all firms that meet the eligibility requirements. Project grants and Student grants are at the discretion of Callaghan, based on their judgement against a set of criteria.[5]

Source: Information taken from Callaghan Innovation 2015 Annual Report.

As described above, a lack of productivity data after 2012 precludes an evaluation of the grants administered by Callaghan Innovation. Instead, this Research Note evaluates the set of grants that existed in the period from 2005 to 2012. During this time, a range of grant types were administered by different agencies under different schemes. To make these broadly comparable to the set of grants that now exist under Callaghan Innovation, Figure 2.2 shows a timeline of grant types that existed from 1994 to 2012, with the colours representing the corresponding grant type under Callaghan Innovation.

[5] The Growth Grant criteria require that to be eligible a business must:

- Have had a minimum of $300,000 in eligible R&D expenditure sourced from non-government funds in each of the last two financial years
- Have had eligible R&D expenditure of at least 1.5% of their revenue in the last two financial years.
- Meet financial and management due diligence requirements sufficient to justify three years of funding
- Provide an R&D plan:
  - Suitable to assess progress in the business R&D programme; and
  - An estimate of R&D expenditures over the next three years.

The Project Grant Judgement Criteria fall under the following headings: Private Investment returns; Pathway to market (commercial outcomes); Ability to deliver (R&D outputs); Develop R&D Programme; Grant Impact; and Benefits Outside the Business.

Figure 2.2 Classification of pre-Callaghan R&D grant types into Callaghan types

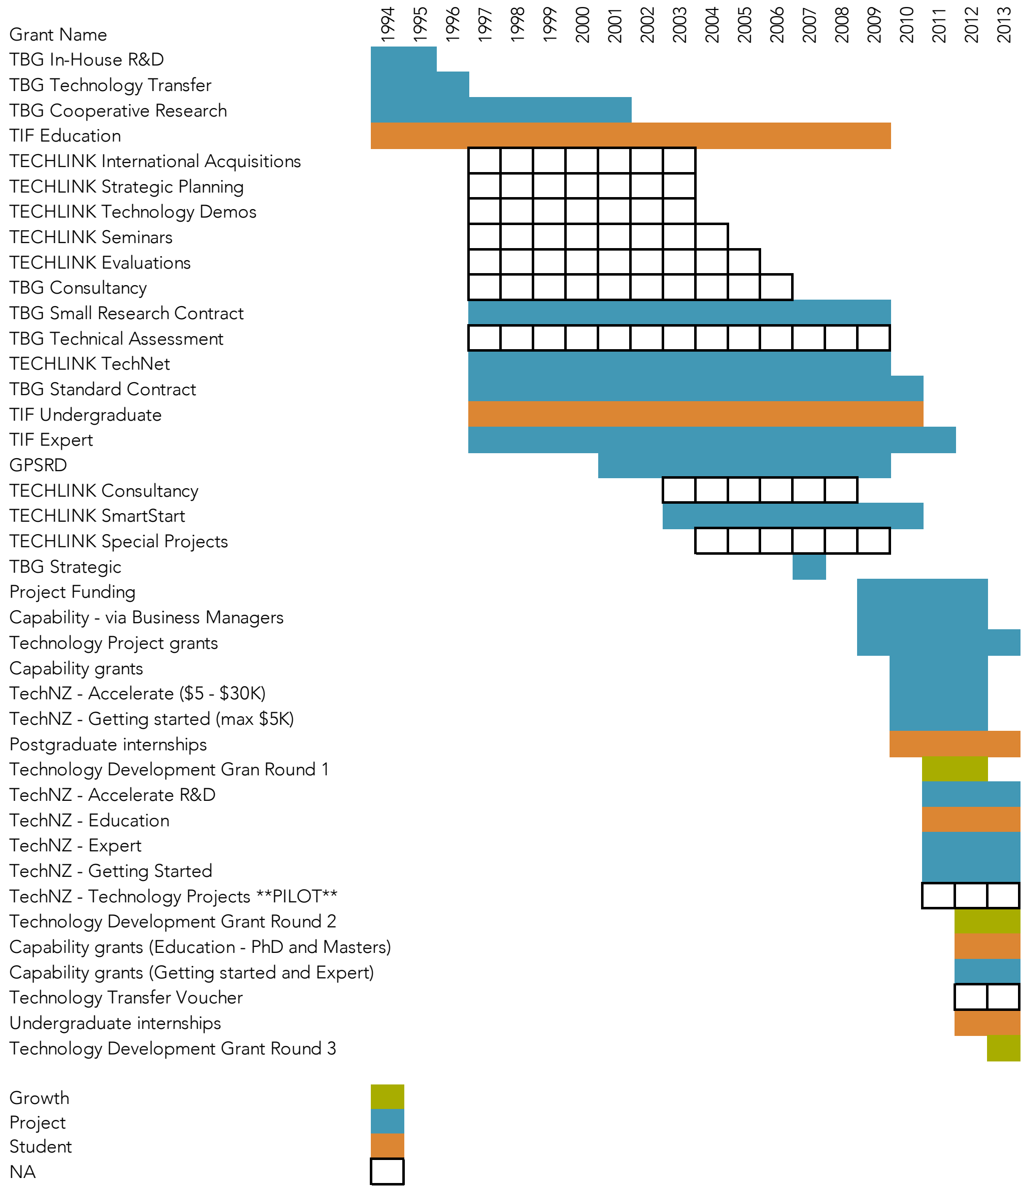

Notes: Based on information in Feng (2016).

The analysis that follows only uses grant types corresponding to Callaghan Innovation's Project grant in which the total amount of Project grants received in a given year was at least $10,000. It is easier to construct a set of otherwise-comparable firms that did not receive a Project grant and hence to construct a counterfactual set of non-recipients than for Growth-grant recipients. The granting agency has some discretion in whether to award a Project grant so there are likely to be eligible firms that did not receive a Project grant. As long as there are no systematic differences to the Project grant recipients, these firms can form a good comparison set. By contrast, all eligible firms are automatically entitled to receive a Growth grant, so if a firm did not receive a Growth grant then either it chose not to apply or the firm was not eligible or for some reason. Whichever the reasons, the non-recipient is unlikely to be a good comparison for a Growth-grant recipient. Meanwhile, Project grants are more substantial than Student grants and therefore more likely to impact on firm performance.

Figure 2.3 plots the number of recipients and total value of Project grants paid out from 2003 to 2012. It shows that both the number and the value of Project grants fluctuated significantly over the time. In particular, there was a large increase in awards for both the TBG Standard Contract and the TBG Strategic grants in 2007. Some of the fluctuation may also be due to scheme changes that reduced the availability of Project grants.

Figure 2.3 Number and amounts of pre-Callaghan Project grants 2003-2012

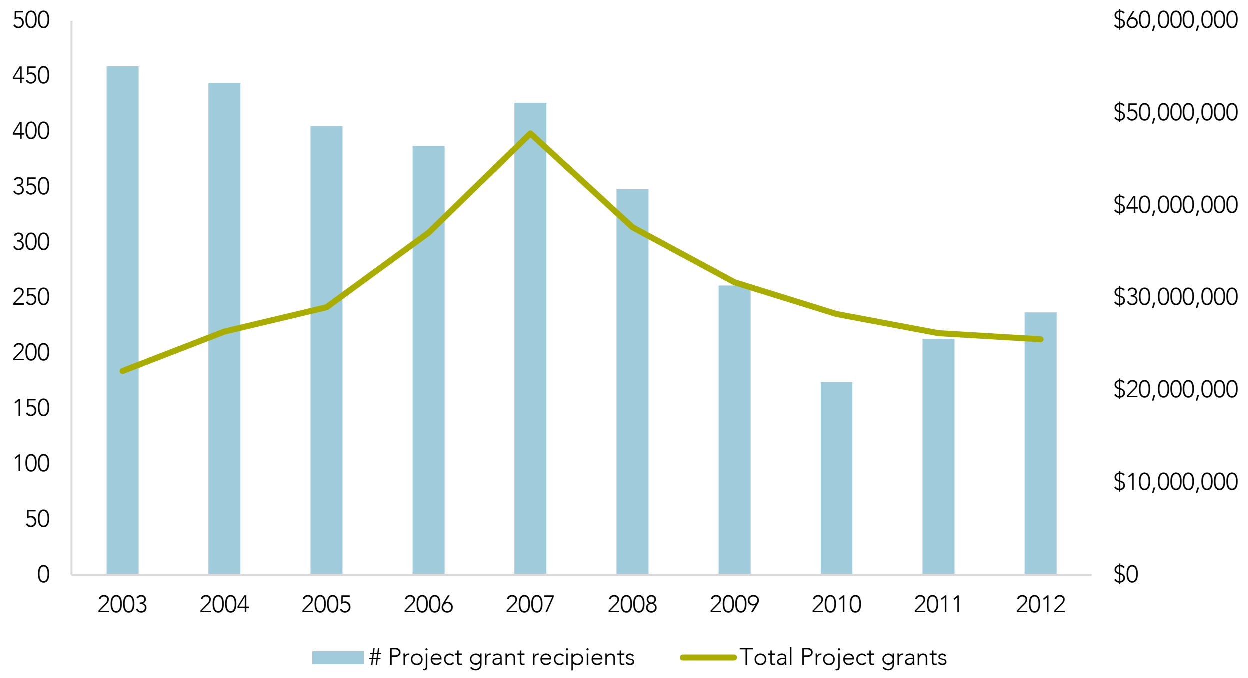

Notes: Chart shows number of recipients (against left axis) and total value (against right axis) for grants paid out from 2003 to 2012. Information is constructed from data on grant recipients matched to the Longitudinal Business Frame so excludes data on grants that were not matched to identifiable firms.

2.2 Constructing measures of firm performance

Various measures are used to gauge firm performance following the receipt of an R&D grant: changes in R&D expenditure, patenting activity, survival, various types of innovation activity, change in employment, change in output, and change in productivity (both labour and multifactor) are all used in the analysis.

Information on R&D expenditure from responses to the R&D Survey and/or the BOS Survey is used to construct a measure of the change in R&D expenditure over time, from the year prior to receiving the grant until 1 to 4 years after the grant year for the subset of firms for which data is available.[6] However, as neither survey is explicitly longitudinal, the data necessary to measure changes in R&D expenditure over a 3-year period is only available for a fraction of firms: around 25% of grant recipients and a much smaller fraction of non-recipients. More importantly, it is unclear if this sample is representative, given that Statistics New Zealand is more likely to collect data on larger firms.

R&D expenditure is zero for a large number of firms, and it is common for firms to have a positive value in one year and zero in another. To reduce the distortion caused by zero values, the change is calculated relative to the average level of R&D expenditure (RD) before and after. That is, change is measured using the following formula: $\Delta RD = \dfrac{RD_{t+n} - RD_{t-1}}{\frac{1}{2}(RD_{t+n} + RD_{t-1})}$, which constrains ΔRD to between -2 and 2.

To measure patenting activity and innovation output of various types, indicators are constructed for whether a firm filed a patent and/or whether it reported innovation activity of a specific type (conditional on having responded to the BOS survey) in a specific year.

To measure firm productivity performance, the analysis applies the method described by Fabling and Maré (2015) for estimating firm-level measures of multifactor productivity (*MFP*). More specifically, MFP

[6] The data from the year prior to grant is used so that the grant amount is not included the baseline calculation. The data is taken from the R&D survey if available, and from the BOS if not. To avoid any inconsistencies in measurement of R&D expenditure across the surveys, the data on the levels is for the start and the end of the period is drawn from the same survey. This requires that the firms are surveyed at least twice.

is the residual from estimating a trans-log production function with gross output as the dependent variable and firm fixed effects, with each industry estimated separately. However, in contrast to Fabling and Maré, the production function specification does not include year fixed effects, which allows MFP to vary across the 12-year period for which data is available and makes it possible to compare MFP levels (i.e., to measure changes) across time.

Formally the specification of the production function is:

$$ln\left(GO_{it}\right) = \alpha_j + \sum_r \beta_r ln\left(X_{it}^r\right) + \sum_r \sum_{s \neq r} \delta_{rs} ln\left(X_{it}^r\right) ln\left(X_{it}^s\right) + \gamma_i + \varepsilon \text{ for each industry } j \in J$$

where $r, s \in \{L, K, M\}$

GO_{it} is firm i's gross output in year t

X_{it}^L is firm i's level of employment in year t

X_{it}^K is firm i's capital stock in year t

X_{it}^M is intermediate inputs firm i uses in year t

γ_i is a fixed effect for firm i

J is the set of industries

The value of MFP for firm i in year t is calculated by:

$$\ln(MFP_{it}) = ln\left(GO_{it}\right) - \sum_r \beta_r ln\left(X_{it}^r\right) - \sum_r \sum_{s \neq r} \delta_{rs} ln\left(X_{it}^r\right) ln\left(X_{it}^s\right)$$

where MFP_{it} is multi-factor productivity of firm i in year t

As the various measures of firm performance, and particularly *MFP*, are measured with error, the two-year moving average is used in the analysis. This reduces variation that may be caused by measurement error in any particular year.

The analysis uses information from various sources to construct a set of firm characteristics that are used to predict the propensity to receive a grant. This includes indicators of the firm's age, its rolling mean employment (*RME*), its *MFP* (as above), its capital-labour ratio ($\frac{K}{L}$), whether it is part of a business group, whether it is foreign owned, whether it is engaged in exporting, and its primary industry, using information from the core LBD.

Information from a combination of the R&D survey, BOS, and the IR10 form are used to create an indicator of whether the firm engaged in R&D activity in any particular year.[7] Information from the government assistance programme (GAP) database is used to construct indicators of whether a firm received a R&D grant or another form of government assistance in the prior 3 years. In order to allow comparison to the results generated by Jaffe and Le (2015), the set of variables that were used in that paper are constructed for only firms that responded to the BOS.

It is important to note some issues in the data on R&D recipients available in the LBD. In particular, the descriptive statistics in the raw dataset indicate that around 20% of firms that supposedly received R&D grants have no employees. As the criteria for Project grants require that the recipient have the capability to deliver on the technical aspects and to commercialise the results, it is unlikely that a grant would be awarded to a firm with no employees. Instead, one possibility is the recipient is part of a business group and the information on employees (and probably also the financial information, including R&D expenditure) was recorded under another member of the group. A second explanation is that that the grant information was simply matched to the wrong firm in the LBD.[8]

Observations with no employees are dropped from the sample to eliminate the obvious errors. However, there may be some cases in which the grant information was matched to a unit of a business

---

The IR10 form contains a line item for R&D expenditure. However, because of the way the income & expense statement in that form is constructed, it implicitly excludes spending on salaries & wages and physical equipment (Fabling, 2008). Therefore the quantitative measure is not comparable to the measures of R&D expenditure in the R&D survey and the BOS. Nevertheless, if a firm reports positive R&D expenditure on the IR10 form it is an indicator that it was engaged in R&D activity in that year.

[8] The grant information was matched to the LBD by name and GST number (if available), but not all matches may have been correct.

group that had some employees but which only reports part of that group's financial data. In that case the information on the firm's performance after receipt of the grant will be incorrect.[9] Hence the results should still be treated with caution.

2.3 Descriptive statistics

Table 2.3 shows the mean values of the variables for all firms with available data in the LBD alongside the mean values for R&D active firms in the same industries as grant recipients and just the set of grant recipients. It shows that the Project-grant recipients differ substantially, not only from the full dataset but also from R&D active firms in the same industries. This highlights the need for caution in interpreting the raw (unmatched) results, and the importance of constructing a reasonable counterfactual.

In any case, Figure 2.4 shows various aspects of firm performance of grant recipients relative to the performance of non-recipients to provide a point of comparison with the matched results in the following section. The charts show the average values for both grant recipients and non-recipients separately as well as the difference of means – grant recipients minus non-recipients (with the level of statistical significance indicated by *'s).

The results are estimated by regressing the various measures of firm performance on an indicator of whether the firm received a grant or not. For the continuous performance measures (change in R&D expenditure, etc.) the estimates were derived using ordinary least squares (OLS), whereas for the binary measures (patenting, innovation activity, and survival) estimates were derived using a logit model. The sample includes all firm-year observations with data on firm performance and whether a grant was received.

The results show that the change in R&D expenditure of grant recipients is flat to declining, while for non-recipients it is growing over time, and that the difference after 4 years is statistically significant. However, on all the other measures grant recipients perform better than non-recipients. Grant recipients are significantly more likely to patent in subsequent years, more likely to survive, have higher rates of all types of innovation, and grow significantly faster.

That said, no attempt has been made to account for systematic differences between grant recipients and non-recipients that may explain these differences in performance. Hence these results should not be interpreted as reflecting any causal impact of receiving a grant on firm performance.

[9] This error may be partly ameliorated because economic performance is measured relative to a baseline that has also been understated. Moreover, the propensity-score matching process takes into account the information on employment and productivity, so a grant recipient with understated financial data will be matched to a smaller firm than it would if it had the correct data. Nevertheless, if there are a substantial number of incorrectly matched firms still remaining in the sample then the resulting statistics will not give a true depiction of the grant recipient's actual performance.

Table 2.3 Mean values of key variables

	All economically active firms	R&D active firms in same industries as grant recipients	Project grant recipients
Economic output			
Multi-factor productivity	0.01	-0.07	-0.15
Labour productivity	$183,960	$170,967	$91,874
Gross output	$1,075,457	$11,543,549	$21,320,006
Employment	3.63	27.40	49.99
Sales (as per GST filing)	$858,276	$10,094,374	$15,921,853
Sales (combined)	$602,051	$9,013,105	$12,249,689
Innovation output			
Any innovation new to the firm	43.3%	71.9%	76.1%
Good/service new to firm	24.8%	61.8%***	72.6%***
new to the world	4.0%	19.5%***	39.3%***
new to New Zealand	6.1%	17.1%***	14.3%*
new to the firm	14.8%	25.2%***	17.9%***
Operational process new to firm	22.8%	43.3%***	42.9%
Organisational process new to firm	27.9%	46.1%***	40.5%***
Marketing method new to firm	24.0%	40.2%***	40.5%
Filed patent in year	0.0%	0.7%***	4.2%***
# patents filed	0.00	0.01***	0.08***
R&D expenditure			
Reported positive R&D spending	98.9%	100.0%***	59.4%***
R&D expenditure (as per BOS)	$83,468	$676,303***	$1,509,618***
R&D expenditure (as per R&D survey)	$314,123	$780,313***	$1,159,959***
Total expenditure	$658,178	$10,527,943***	$13,379,739***
Firm characteristics			
Age	10.19	11.24***	11.67***
0-5 years	32.4%	29.9%***	31.6%***
10-20 years	28.6%	27.3%***	26.9%
20-50 years	13.0%	14.1%***	12.4%***
5-10 years	25.5%	27.0%***	26.4%
50+ years	0.6%	1.7%***	2.6%***
Employment			
0-20 emp	39.9%	61.8%***	37.7%***
20-50 emp	59.2%	29.7%***	48.7%***
50-100 emp	0.7%	5.4%***	8.3%***
100+ emp	0.2%	3.1%***	5.4%***
Exporter (as per LBD)	2.5%	21.9%***	46.5%***
Foreign owned (as per LBD)	0.8%	4.5%***	6.5%***
Received R&D grant in prior 3 years	0.3%	8.3%***	71.3%***
Received non-R&D assistance in prior 3 years	0.4%	9.0%***	40.5%***
Belongs to a business group	2.4%	9.5%***	23.3%***
State-Owned Enterprise	0.0%	0.2%***	0.3%***
Predominant sector (as per ANZSIC 2006)			
Primary	21.8%	8.6%***	4.5%***
Manufacturing	5.9%	21.6%***	40.4%***
Services	72.3%	69.8%***	55.2%***
Predominant industry (as per ANZSIC 2006)			
Agriculture, forestry & fishing	21.7%	8.5%***	4.3%***
Mining	0.1%	0.1%	0.2%
Manufacturing	5.9%	21.6%***	40.4%***
Electricity, gas, water & waste services	0.2%	0.2%	0.5%***
Construction	22.5%	6.9%***	1.7%***
Wholesale trade	4.6%	9.3%***	10.5%***
Retail trade & accommodation	11.8%	6.2%***	1.3%***
Transport, postal & warehousing	4.5%	1.0%***	0.4%***
Information media & telecommunications	1.2%	6.2%***	1.1%***
Financial & insurance services	2.0%	1.0%***	1.5%
Rental, hiring, and real estate services	3.3%	2.4%***	1.2%
Professional & administrative services	16.5%	29.2%***	36.1%
Personal, household & other services	0.1%	0.0%	0.0%
Arts, recreation, and other services	5.6%	7.4%	0.9%

Notes: This table shows count and mean for a range of variables for all firms that were (1) economically active; (2) R&D active and in same industries as a Project grant recipient; and (3) received a Project grant. Sample includes all firms with available data in given year. Counts are rounded to base 3. Proportions for binary and categorical variables are calculated using sum and counts rounded to base 3. Asterisks against the means in columns 2 and 3 indicate that the difference in means relative to column 1 and 2 (respectively) is statistically significant, where *** = p<0.01, ** = p<0.05, * = p<0.1.

Figure 2.4 Performance of grant recipients vs non-recipients (unmatched)

Panel A: R&D expenditure, patent filing, and survival

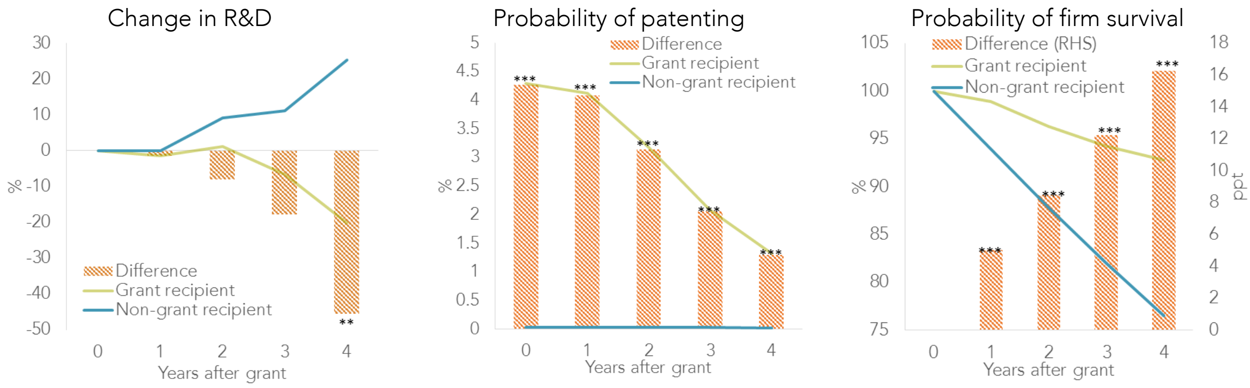

Panel B: Innovation activity

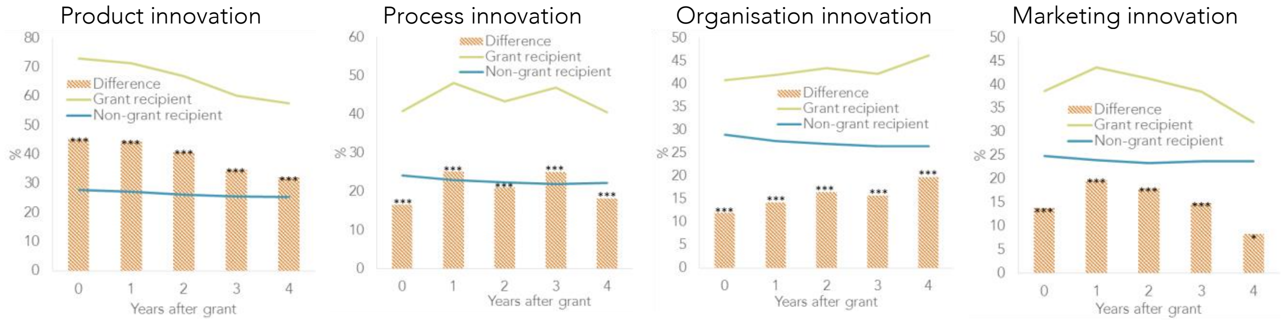

Panel C: Economic performance

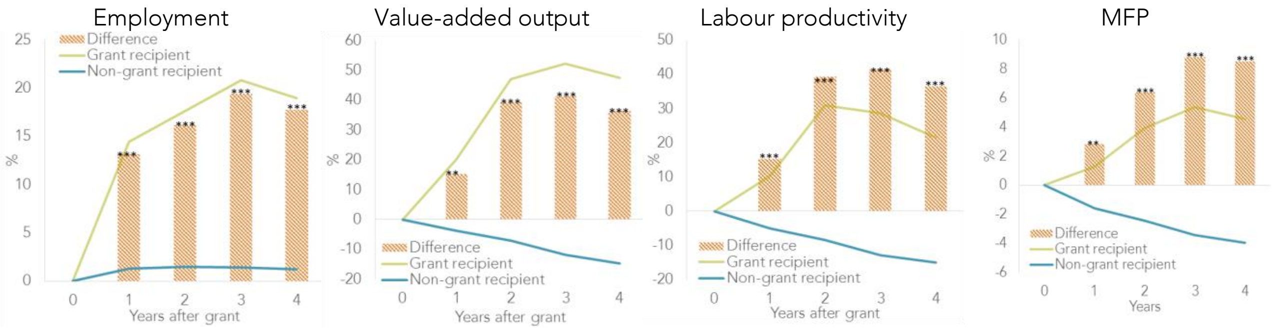

Notes: Figure contains a series of charts showing the performance of grant recipients relative non-recipients. The lines show the average values of performance measures for both grant recipients and non-recipients separately while the bars show the difference between the two. The values are estimated from a regression of the performance measure on an indicator of whether the firm received a grant. For the binary measures the regression is estimated using a logit model and for the continuous measures using an OLS model. The regression sample includes all observations with available data on the performance measure and grant recipient in a given year. Asterisks indicate whether the difference is statistically significant, where *** = p<0.01, ** = p<0.05, * = p<0.1.

2.4 Constructing the counterfactual

Choice of method

The objective of this paper is to evaluate the impact of receiving an R&D grant on firm performance – that is, to compare the outcome if a firm receives government assistance to the outcome if it does not. However, the "counterfactual" – the outcome if a grant recipient had not received assistance – does not exist in the real world, so it is necessary to approximate it using the outcome(s) of firms that did not receive grant assistance.

The challenge in constructing the counterfactual from the outcomes of non-recipients is to include as many firms as possible without including firms that are systematically different from R&D grant recipients. The ideal is when the "treatment" – receiving a grant – is randomly assigned across firms, as in a randomised controlled trial. This is clearly not possible when using an evaluation constructed after the fact, but it provides a benchmark against which to compare alternatives.

One approach is to use details of the grant process to choose non-recipients that would have received a grant but for some (almost) random circumstance. There are a range of reasons a firm does not receive a grant, including it was not eligible, it applied but was judged not good enough, or it was not aware that a grant was available and hence did not apply. When there are bright-line criteria for determining whether a firm is eligible for a grant, as exists for Growth grants under the Callaghan Innovation schema,[10] then it may be possible to identify firms that were (just) outside the cut-off but are otherwise not systematically different from grant recipients. Similarly, if grant recipients are selected using a score-based selection process it is possible to use non-recipients (just) below the cut-off core.[11]

Another approach is to match non-recipients to recipients based on characteristics (e.g., size, age, industry, R&D activity) that are important factors in whether a firm receives a grant. For instance, Figure 2.5 shows the probability of receiving a grant by industry, from only 1 percent for firms in Information Media and Telecommunications to at least 15% for firms in the Manufacturing industry. Hence the firm's industry might be an appropriate characteristic to use for matching.

Figure 2.5 Grant recipients as proportion of all R&D active firms by industry

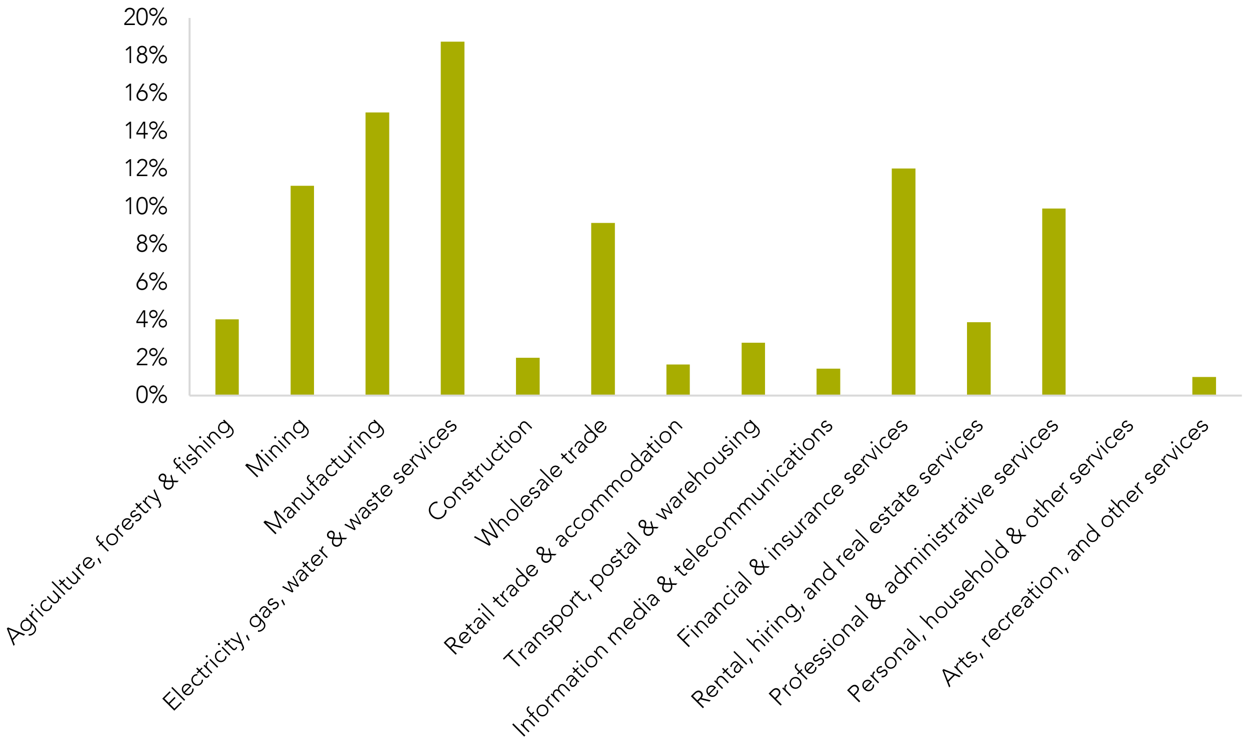

However, using characteristics that do not definitely determine the eligibility/propensity to receive a grant may wrongly exclude firms that are otherwise comparable. Moreover, individual characteristics on their own are unlikely to definitely determine whether a firm receives a grant. Instead, in most cases this will depend on a (non-linear) combination of multiple factors. Hence it may make sense to use a combination of firm characteristics to determine the likelihood of receiving a grant, and then choose non-recipients with similar "propensity" as grant recipients. This method is known as propensity-score matching (Rosenbaum & Rubin, 1983).

[10] To be eligible for a Growth grant a firm must have R&D expenditure of $300,000 over past 2 years.

[11] In the econometrics literature, a quasi-experimental approach using a cut-off score or threshold is commonly referred to as a "regression discontinuity design" (Angrist & Pischke, 2009).

In the first stage of the propensity-score matching approach, a combination of all (potentially) relevant pre-treatment characteristics are used to estimate a propensity score for each firm-year observation. Then in the second stage the treated and untreated observations with a similar probability of receiving treatment (i.e., similar propensity scores) are matched to each other. Observations that do not have a close enough match – both treated and untreated – are excluded from the analysis.

This method assumes that any systematic differences between the treated and untreated groups can be removed by conditioning on the propensity score. In particular, it assumes that after conditioning on the propensity score there are no characteristics that influence both the probability of treatment and potential outcomes – that is, that assignment to treatment is essentially random. This assumption will be violated if the propensity score estimation includes any variables that are themselves affected by treatment. Therefore it is important to use only characteristics generated prior to the treatment period.

The propensity-score matching approach also does not account for unobserved characteristics that affect both the likelihood of treatment and potential outcomes. Another alternative would be use an instrumental variables (IV) approach. In theory, this would remove the non-random component of selection into treatment and hence make it possible to interpret the remaining effect as causal. However, in practice identifying a valid instrument is difficult.

Propensity-score matching approach

The analysis that follows uses the propensity-score matching approach. As touched on above, this involves the following steps:

1. Estimate the propensity score using the combination of pre-treatment variables that best removes any systematic difference between treated and untreated observations after matching.
2. Choose the algorithm for matching treated and untreated observations, trading off benefits of having a larger number of observations with inaccuracy introduced by using bad matches.
3. Estimate the average treatment effect on treated observations using the dataset of matched observations.

Caliendo and Kopeinig (2008) provide guidance on the choice of variables for estimating the propensity score. The main objective is to remove any systematic difference between the treated and untreated observations after matching. Therefore, although it is important to obtain the best possible prediction of selection into treatment all else being equal, it is even more important not to introduce any variables that are affected by treatment. Moreover, variables that on their own describe systematic differences between treated and untreated observations should be used to stratify the observations before matching (rather than being used to estimate the propensity score).

Table 2.4 shows the results from a logit model of receiving a grant regressed on various combinations of variables that predict the propensity to receive a grant. The initial choice of variables was based on the specification used in Ministry of Economic Development (2011), but then several changes in the choice of variables were made to improve on the predictive power of the estimation. In particular, following Jaffe and Le (2015), the regression includes a variable that captures non-government R&D assistance in the previous three years. It also includes an indicator of whether the firm received other R&D assistance in the previous 3 years.[12]

The specification containing the optimised choice of variables, first without industry or year dummies, is shown in column 1. Dummies to capture industry, year, and their interaction are then progressively added in columns 2-5. For comparison, the results from using the specifications in Ministry of Economic Development (2011) and Jaffe and Le (2015) are shown in columns (6) and (7) respectively. For the purposes of comparing the different specifications, the correction prediction rate (Heckman, Ichimura, Smith, & Todd, 1998) is reported for the whole sample and for grant recipients only.

[12] Jaffe and Le (2015) also included a number of other firm characteristics in their specification of the propensity score, including the (self-reported) level of competition that firm faces in its market, the ease of access to capital, and the level of skill in the labour market. However, as these variables are taken from the BOS, they are only available for a subsample of the firms used in this analysis and hence not included here.

The table shows that the only variable in the optimised specification that is not significant is *log(MFP$_{t-1}$)*. However, this variable is significant when included without *log(MFP$_{t-2}$)* or when the *change in log(MFP)* variable is included instead of *log(MFP$_{t-2}$)*. As such, this variable is kept in the regression so that the coefficient on *log(MFP$_{t-2}$)* represents the *change in log(MFP) from t-2 to t-1.*

The specifications with just industry dummies and with both industry and year dummies result in the highest correction prediction rates for grant recipients. The analysis that follows uses the propensity score predicted from the specification with both industry and year dummies (column 4), and stratifies the matching by industry and year.

The two histograms in Figure 2.6 show the predicted propensity scores for grant recipients (in Panel A) vs non-recipients (in Panel B), where a propensity score of close to zero means that the model predicts there is close to zero probability that the firm will receive a grant (given its characteristics). The charts demonstrate that, although the majority of observations have propensity scores close to zero, the propensity scores stretch across the range up to one. Nevertheless, the propensity scores vary widely by industry/year. This means that, after stratifying by industry & year, there will be a number of treated observations that do not have untreated observations with a similar propensity score.

Figure 2.7 plots the propensity scores for grant recipients by industry and year. The solid and hollow dots indicate those treated observations that are and are not matched (respectively) after imposing a requirement that the propensity score of the untreated observation be within 0.01 (or 1 percent) of the propensity score of a treated one.

More generally, restricting matches to within a tighter range means that more of the treated observations with higher propensity scores are dropped from the sample used for analysis. Figure 2.8 shows the number of grant recipients that are matched within a propensity score range of 0.1, 0.01, and 0.001 respectively.

The matching algorithm determines which untreated observations are to be used for the counterfactual. Table 2.5 lists the alternative matching algorithms considered for this analysis.[13] The most straightforward matching estimator uses only the *k* observations with the closest propensity score (known as "nearest neighbour" matching). However, as shown above, there are many treated observations for which the propensity score of the nearest untreated observation is very different. The caliper method restricts the choice of nearest neighbours to untreated observations within a pre-determined propensity-score range of the treated observation. The radius method also restricts to untreated observations within the range, but uses all of those observations regardless of how far away they are. The kernel method is a compromise between the two: it includes all untreated observations within the range but weights the observations by the inverse of the distance between the propensity scores of the treated and the untreated observation.

As described above, in choosing the matching algorithm the trade-off is between the benefits of a larger number of observations versus the inaccuracy of bad matches. The analysis that follows is based on the caliper method with observations matched within a range of 0.01. Also, following the advice in Caliendo and Kopeinig (2008), matching is done "with replacement" (i.e., untreated observations may be matched to more than one treated observations) and restricted to observations in which the propensity-score range of treated and untreated observations overlap (i.e., lie on the common support).[14]

[13] Caliendo and Kopeinig (2008) also describe several further matching algorithms, including stratification and interval matching, local linear matching, and weighting on the propensity score.

[14] In robustness checks observations are matched within different ranges and alternative matching algorithms.

Table 2.4 Coefficients from logit model of receiving a R&D grant

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		industry dummies	year dummies	year & industry dummies	year-industry dummies	MED (2011)	Jaffe & Le (2015)
Age	-0.0308***	-0.0332***	-0.0305***	-0.0320***	-0.0336***		
	(0.0088)	(0.0089)	(0.0089)	(0.0091)	(0.0093)		
Age squared	0.0003***	0.0003***	0.0003***	0.0003***	0.0004***		
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)		
log(Age)						-0.3206***	-0.0333
						(0.0484)	(0.1080)
log(RME_{t-1})	1.0385***	1.0824***	1.0540***	1.0974***	1.0942***	0.9416***	0.1179*
	(0.1326)	(0.1391)	(0.1338)	(0.1398)	(0.1400)	(0.0901)	(0.0711)
log(RME_{t-2})	-0.8427***	-0.8286***	-0.8593***	-0.8413***	-0.8304***		
	(0.1322)	(0.1392)	(0.1333)	(0.1396)	(0.1398)		
log(RME, squared)						-0.0934***	
						(0.0155)	
Change in log(RME)						0.4423***	
						(0.1015)	
log(MFPt-1)	-0.0294	0.0076	-0.0195	0.0184	0.0221	-0.6289***	
	(0.0938)	(0.0975)	(0.0937)	(0.0965)	(0.0985)	(0.0452)	
log(MFPt-2)	-0.2261***	-0.2177**	-0.2268***	-0.2203**	-0.2230**		
	(0.0828)	(0.0861)	(0.0835)	(0.0865)	(0.0867)		
Change in log(MFPt)						0.3730***	
						(0.0526)	
log(Capital-labour $ratio_{t-1}$)	0.3257***	0.4106***	0.3411***	0.4270***	0.4265***	0.1234***	
	(0.0866)	(0.0892)	(0.0871)	(0.0897)	(0.0904)	(0.0399)	
log(Capital-labour $ratio_{t-2}$)	-0.2604***	-0.2386***	-0.2694***	-0.2465***	-0.2380***		
	(0.0776)	(0.0783)	(0.0781)	(0.0789)	(0.0797)		
Change in log(capital-labour ratio)						0.0595	
						(0.0793)	
Belongs to a business group						0.4623***	0.2708
						(0.1116)	(0.1677)
Exporter	0.6252***	0.5483***	0.6071***	0.5275***	0.5332***	1.6143***	1.1926***
	(0.1381)	(0.1384)	(0.1389)	(0.1394)	(0.1397)	(0.1377)	(0.2456)
Exporter of manufactured goods						-0.4381**	
						(0.1738)	
Goods manufacturer						0.7645***	
						(0.1456)	
Foreign owned	-0.3856**	-0.4796***	-0.3901**	-0.4753***	-0.4991***	-0.5611***	0.0309
	(0.1693)	(0.1562)	(0.1697)	(0.1559)	(0.1604)	(0.1464)	(0.1730)
Reported positive R&D spending in	1.0625***	0.9700***	1.0866***	0.9991***	1.0170***	2.7221***	
prior year1	(0.1026)	(0.1000)	(0.1033)	(0.1008)	(0.1028)	(0.0908)	
Received R&D grant in prior 3 years	4.9653***	4.4779***	4.9688***	4.4577***	4.4698***		
	(0.1592)	(0.1566)	(0.1589)	(0.1563)	(0.1559)		
Received non-R&D government	0.9816***	0.9356***	1.0129***	0.9754***	0.9665***		1.8195***
assistance in prior 3/5 years	(0.1203)	(0.1193)	(0.1213)	(0.1206)	(0.1214)		(0.1720)
Industry dummies	N	Y	N	Y	Y	N	N
Year dummies	N	N	Y	Y	Y	N	N
Industry-year dummies	N	N	N	N	Y	N	N
Firm characteristics (from BOS)	N	N	N	N	N	N	Y
Constant	-8.8857***	-10.9191***	-8.8759***	-10.9374***	-11.2508***	-9.0367***	-6.9843***
	(0.4459)	(0.7211)	(0.4580)	(0.7296)	(1.2393)	(0.3890)	(0.4332)
# grant recipients	1035	1032	1032	1032	1032	1326	237
# non-grant recipients	789243	763467	789246	763467	529995	918075	17787
Pseudo R2	0.524	0.540	0.527	0.543	0.531	0.316	0.223
Correct prediction rate	0.986	0.977	0.986	0.976	0.977	0.903	0.798
Correct prediction rate (recipients)	0.880	0.912	0.883	0.910	0.902	0.838	0.772

Notes: This table shows the coefficients from a series of logit regressions of an indicator of whether the firm received an R&D grant in a given year on a set of predictors. Following Jaffe and Le (2015), the regression in column (7) includes the indicator of non-R&D assistance in the previous 5 years.

Figure 2.6 Predicted propensity scores for grant recipients vs non-recipients

Panel A: Grant recipients

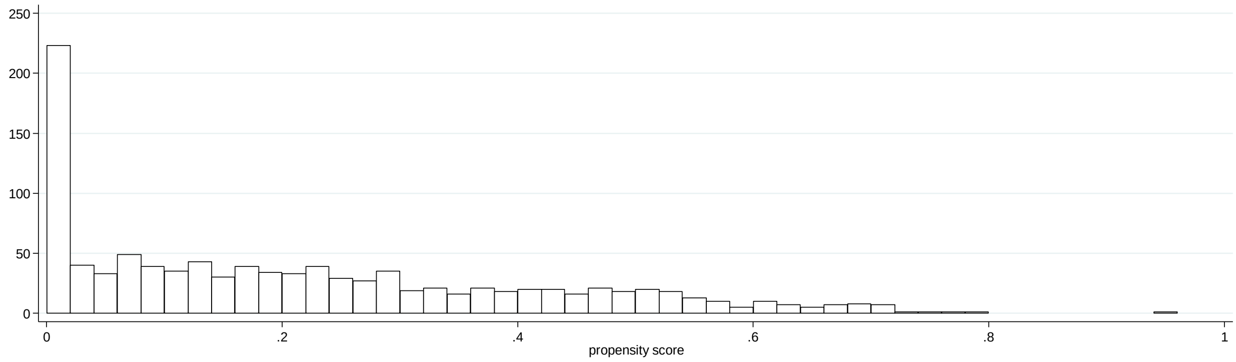

Panel B: Non-recipients

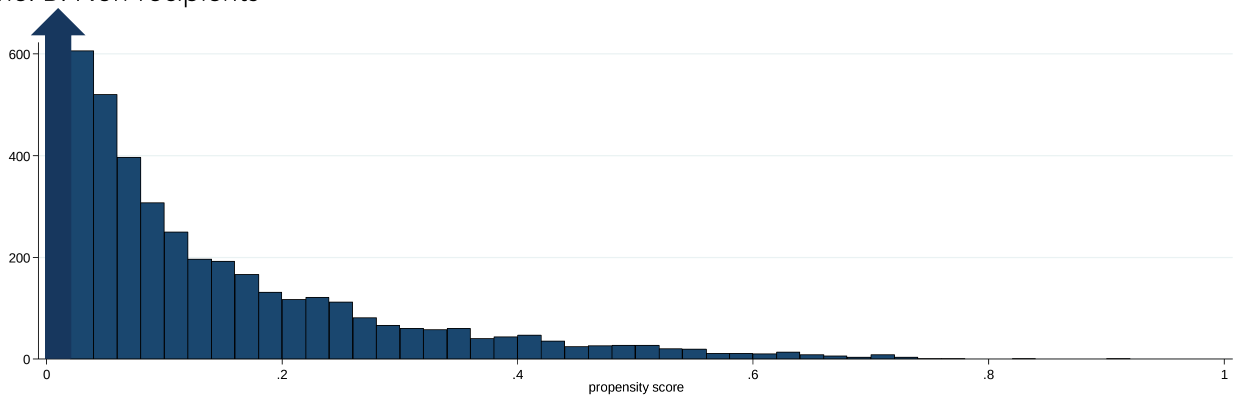

Notes: Figure shows histogram of propensity scores for grant recipients (Panel A) and non-recipients (Panel B), predicted from the regression in reported in Table 2.4. The height of the bars represents counts of treated/untreated firms, except for the first bar in Panel B where count is outside range of y-axis.

Figure 2.7 Propensity scores for grant recipients by industry & year (matched with PS±0.01)

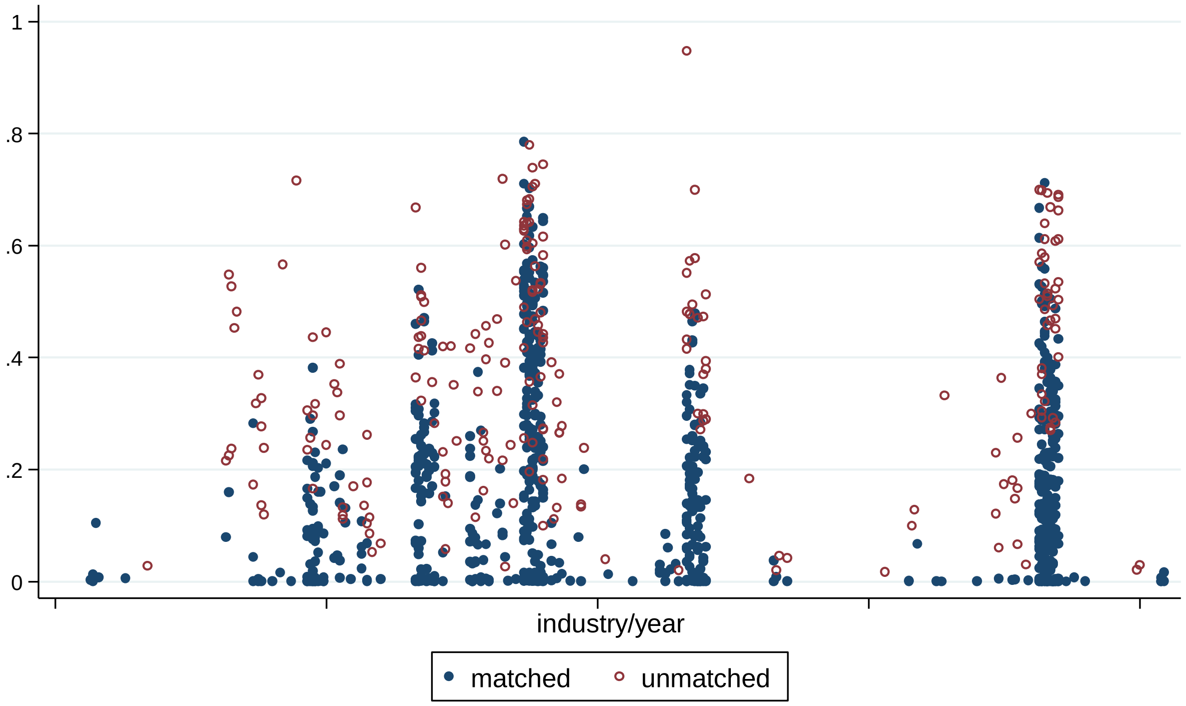

Notes: Figure plots the propensity scores for grant recipients predicted from the regression in Table 2.4 (up the y-axis) by industry and year (along the x-axis). The solid and hollow dots indicate those treated observations that are matched and unmatched (respectively) to an untreated observation with PS±0.01.

Figure 2.8 Propensity scores for grant recipients matched within different ranges

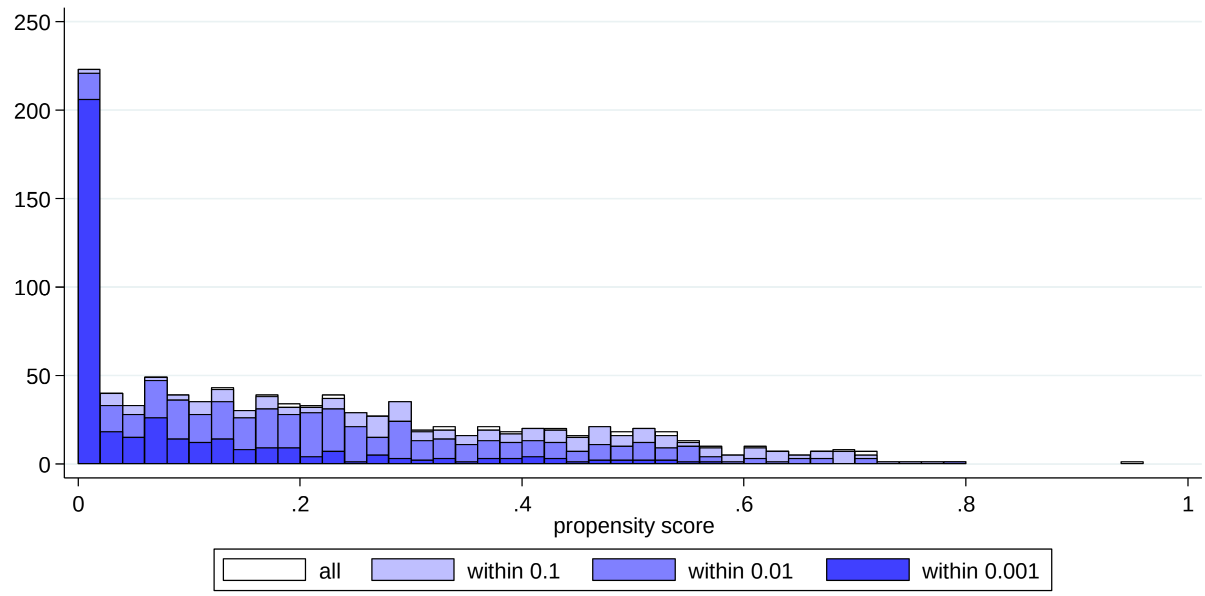

Notes: Figure shows a set of histograms overlaid on top of each other, starting with all grant recipients, and adding those that are matched to a non-recipient with PS±0.1, ±0.01, and ±0.001 respectively.

Table 2.5 Alternative matching algorithms

Algorithm	Condition
Nearest neighbour	Match each grant recipient to *k* non-recipients with closest PS
Caliper	Match each grant recipient to *k* non-recipients with closest PS *within range (e.g., 0.01)*
Radius	Match each grant recipient to *all* non-recipients with PS within range
Kernel	Match each grant recipient to all non-recipients with PS with range *and weight non-recipients based on difference vs recipient's PS*

Constructing the matched dataset implies choosing one or more untreated observations for each treated firm. However, the set of observations for which performance information is available varies for the different output measures and over different time horizons (one year, two years, etc.). This means that the dataset of observations available for matching also varies across these dimensions.

If the observations are re-matched for every variation, then the set of other firms for which information is available could affect the results (particularly when using the nearest-neighbour/caliper and kernel approaches). Instead, for the measures of firm performance calculated as a change from a base year, the match is done using observations that have information for the output measure in the base year.

For the measures created from BOS variables (i.e., innovation output) for which the sample changes every year, all observations in the BOS sample are matched (i.e., with information in any year). For the change measures, there will be some attrition of observations when data is not available in later years (most likely because the firm ceases being economically active). If the attrition of treated and matched controls is related to treatment then this may affect the results, and the results need to be interpreted in light of this. By contrast, as the BOS sample is random and representative of firms in the BOS population in any year, the attrition should not affect the results on the BOS measures.

Having chosen the matched dataset, it is important to verify that treated and untreated observations are "balanced" – that is, there are no systematic differences between them. Panel A of Table 2.6 presents the results of a t-test of differences of means on the variables used to estimate the propensity score, both before and after matching. It shows that the significant differences between treated and control firms in the full sample are removed once the observations have been matched. Moreover, Panel B shows that overall the bias between the treated and control groups has been effectively removed after matching.

Table 2.6 Balance of treated and control samples before and after matching

Panel A: Individual variables

Variable		Mean		%bias	%reduction in abs(bias)	t-test		V(T)/V(C)
		Treated	Control			t	p>abs(t)	
Age	Unmatched	17.14	12.97	31.10		13.36	0.00***	2.59+
	Matched	16.28	16.45	-1.30	95.70	-0.24	0.81	1.13
log(RME_{t-1})	Unmatched	2.76	0.57	150.60		70.77	0.00***	3.29+
	Matched	2.57	2.64	-5.00	96.70	-0.82	0.41	1.02
log(RME_{t-2})	Unmatched	2.63	0.55	138.30		67.58	0.00***	3.64+
	Matched	2.45	2.51	-4.40	96.80	-0.71	0.48	1.01
log(MFP_{t-1})	Unmatched	-0.09	0.04	-19.40		-6.33	0.00***	1.06
	Matched	-0.05	-0.04	-0.80	95.80	-0.19	0.85	1.17+
log(MFP_{t-2})	Unmatched	-0.12	0.04	-23.30		-7.92	0.00***	1.23+
	Matched	-0.06	-0.08	2.90	87.40	0.64	0.52	0.79+
log(Capital-labour ratio$_{t-1}$)	Unmatched	9.66	9.34	31.00		8.55	0.00***	0.47+
	Matched	9.61	9.60	1.30	95.80	0.32	0.75	0.96
log(Capital-labour ratio$_{t-2}$)	Unmatched	9.63	9.31	29.90		8.47	0.00***	0.56+
	Matched	9.58	9.59	-0.50	98.40	-0.11	0.91	1.07
Exporter	Unmatched	0.64	0.05	155.80		82.91	0.00***	
	Matched	0.59	0.62	-6.70	95.70	-1.02	0.31	
Foreign owned	Unmatched	0.10	0.01	41.10		28.79	0.00***	
	Matched	0.10	0.10	-1.00	97.70	-0.15	0.88	
Reported positive R&D spending in prior year	Unmatched	0.56	0.02	146.20		114.99	0.00***	
	Matched	0.49	0.48	2.30	98.40	0.34	0.73	
Received R&D grant in prior 3 years	Unmatched	0.80	0.01	276.50		336.73	0.00***	
	Matched	0.74	0.74	1.30	99.50	0.17	0.87	
Received non-R&D government assistance in prior 3 years	Unmatched	0.53	0.01	144.60		174.30	0.00***	
	Matched	0.47	0.45	4.60	96.80	0.67	0.50	
	+ if variance ratio outside [0.89; 1.13] for U and [0.87; 1.15] for M							

Panel B: Overall measures

Sample	Ps R^2	LR chi^2	p>chi^2	Med. Bias	Mean Bias	B	R	%Var
Unmatched	0.542	8511.1	0	41.1	93.8	337.0*	11.47*	88
Matched	0.002	4.54	0.972	1.3	2.6	10.6	1.08	25
* if B>25%, R outside [0.5; 2]								

Notes: Panel A shows the means and variance of the variables used to predict the propensity score for the full sample before matching and for the matched sample. Means of the binary variables are generated using sum and counts rounded to base 3 in order to protect confidentialty. Panel B shows a set of summary statistics based on the comparisons between the treated and control groups before and after matching.

2.5 Estimating the impact of grants on performance

Figure 2.9 shows the same set of charts as in Figure 2.4 above, with the sample restricted to the matched observations (using a caliper of 0.01). The results show that the R&D expenditure of grant recipients is not significantly different from that of non-recipients in the years after receiving the grant. However, consistent with Jaffe and Le (2015), grant recipients are more likely to file patents and to introduce new products but not to engage in process innovation. Nevertheless, they are more likely to engage in a marketing innovation in the year following the grant. They are more likely to be economically active 3-4 years after grant, and have higher employment and labour productivity growth. However, they have significantly lower MFP growth than the control group in the first year.

Figure 2.9 Performance of grant recipients vs non-recipients (matched)

Panel A: R&D expenditure, patent filing, and survival

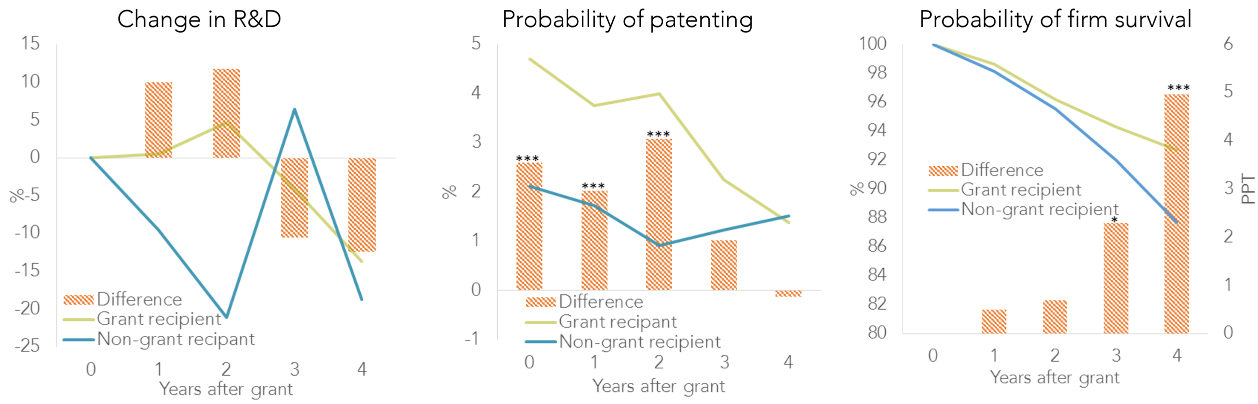

Panel B: Innovation

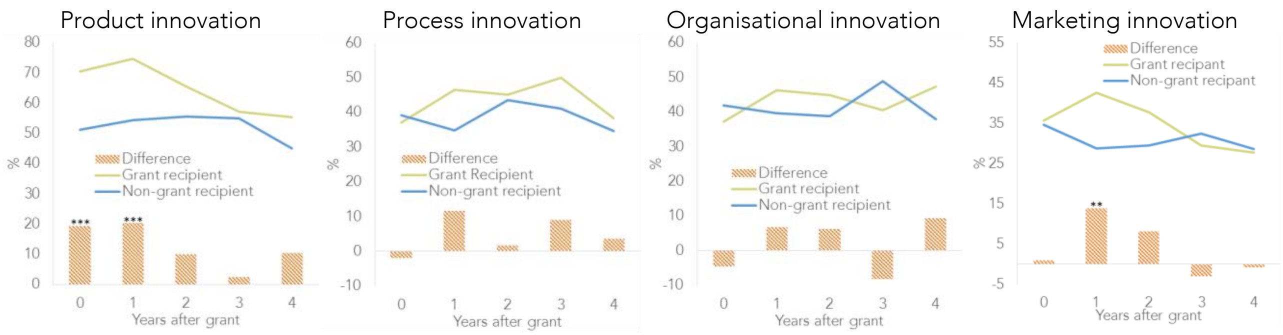

Panel C: Economic performance

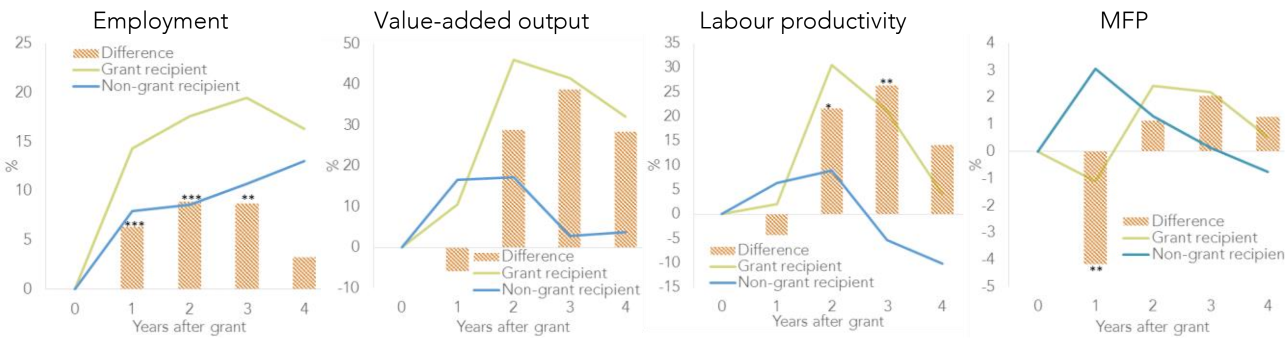

Notes: Figure contains a series of charts reflecting the performance of grant recipients relative a control group of non-recipients. Sample includes observations matched to 5 nearest neighbours within a caliper of 0.01. Other details same as Figure 2.4.

As the results on R&D expenditure and MFP are surprising, the following sub-sections examined these results in more detail.

R&D expenditure

Figure 2.10 shows the change in R&D expenditure for grant recipients and the control group by year in which the firm received the grant. It reveals some variation in the impact of receiving an R&D grant on R&D expenditure, but across all years there is a wide margin of error around the estimates. Only firms that received grants in 2006 show a significantly higher R&D expenditure and only in the first year after the grant.

Figure 2.10 Change in R&D expenditure for grant recipients vs controls by year

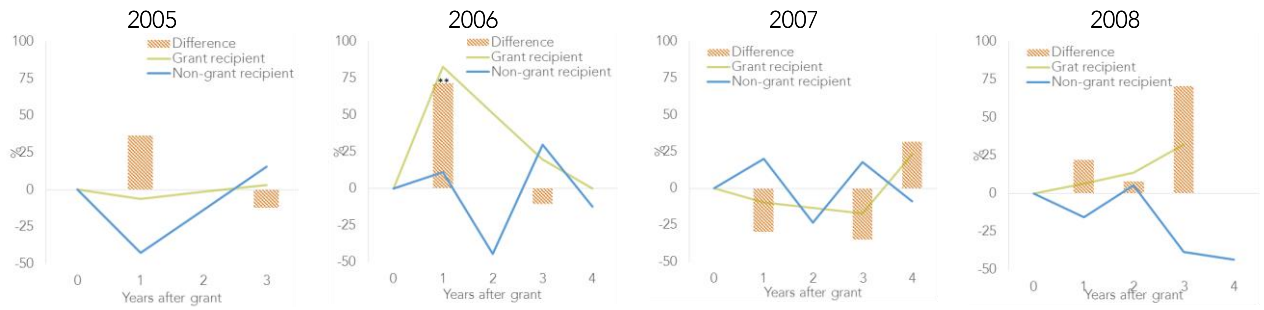

Notes: Figure contains a series of line charts showing the change in R&D expenditure of grant recipients relative a control group of non-recipients in each grant year. Sample includes observations matched to 5 nearest neighbours within a caliper of 0.01. Other details same as Figure 2.4.

Figure 2.11 plots the kernel density distribution of the 3-year change in R&D expenditure (with the results on all years pooled together) for grant recipient and non-recipients separately. The vertical lines represent the weighted mean of the two samples. It shows that grant recipients are more likely than non-recipients to have positive R&D expenditure growth, but a large proportion have negative growth and that on average the mean R&D expenditure growth over three years for grant recipients is below that for non-recipients (consistent with Figure 2 Panel A).

Figure 2.11 Kernel density of 3-year change in R&D expenditure for grant recipient vs controls

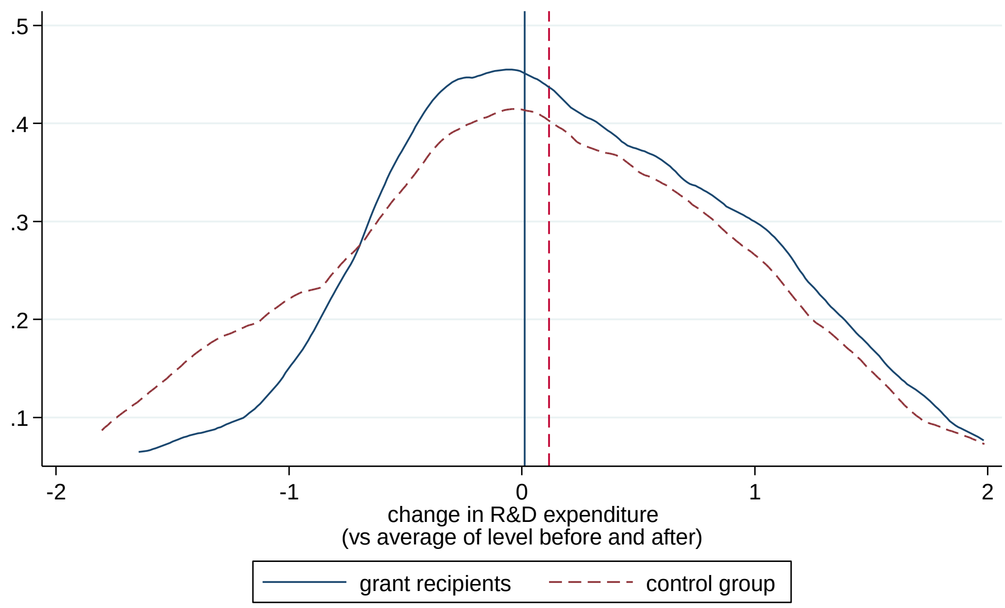

Notes: Figure contains kernel density plots of 3-year change in R&D expenditure for grant recipient (solid blue) and control group (dash red). Top and bottom 5% of distribution not shown. Vertical lines represent the weighted mean of each sample. Sample includes observations matched to 5 nearest neighbours within a caliper of 0.01. Other details same as Figure 2.4.

Multi-factor productivity

Figure 2.12 shows the change in MFP for grant recipients and the control group by year of grant. It shows that MFP growth has fluctuated for both grant recipients and non-recipients over the period, and in some years (particularly 2006, 2007 & 2009) MFP growth of grant recipients is higher while in other years (2005, 2008 & 2010) the control group had faster productivity growth. The only grant years for which recipients and the control group grew at significantly different rates are 2005 and 2010, when the control group grew 10 and 20 percentage points faster (respectively) in the first year than great recipients.

Figure 2.12 Change in MFP for grant recipients vs controls by year

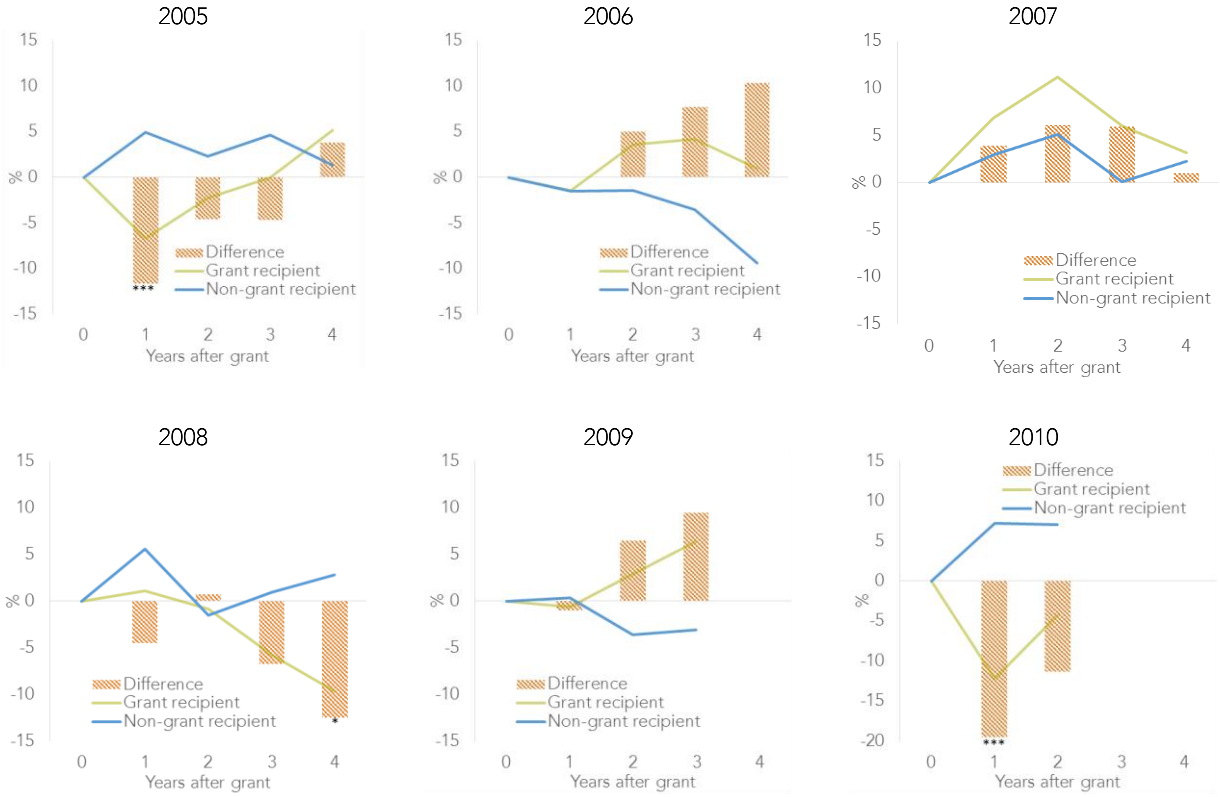

Notes: Figure contains a series of line charts showing the change MFP of grant recipients relative to a control group of non-recipients in each grant year. Sample includes observations matched to 5 nearest neighbours within a caliper of 0.01. Other details same as Figure 2.4.

The kernel density plots of the 3-year change for grant recipients vs the control group, shown in Figure 2.13, reveals that productivity growth of grant recipients has a wider variance than for the control group, and the mode is negative. By contrast the mode is positive for the control group, although for a large proportion of the firms the 3-year change is negative.[15]

[15] The kernel density plot of the control group has several local peaks (i.e., appears more volatile). This indicates there is clustering around several different growth rates (both positive and negative).

Figure 2.13 Kernel density of 3-year change in MFP for grant recipients vs controls

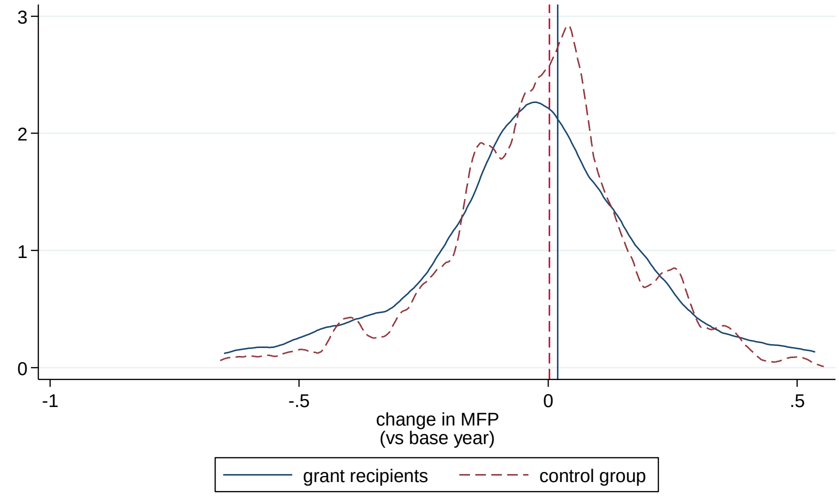

Notes: Figure contains kernel density plots of 3-year change MFP for grant recipient (solid blue) and control group (dash red). Top and bottom 5% of distribution not shown. Vertical lines represent the weighted mean of each sample. Sample includes observations matched to 5 nearest neighbours within a caliper of 0.01. Other details same as Figure 2.4.

3 Discussion & conclusion

Overall, the impact of receiving a Project grant between 2004 and 2012 on firm performance appears to be mixed. Project grant recipients are spending no more on R&D in subsequent years than the control group. Moreover, relative to the full sample of (unmatched) non-recipients, the R&D expenditure of grant recipients is declining. However, consistent with Jaffe and Le (2015), grant recipients are more likely to patent and to introduce new products (but not to engage in process innovation). In addition, the results show that Project grant recipients are more likely to innovate in their marketing approach in the first year after the grant.

Taken together, the results suggest that after receiving a Project grant the recipients direct their attention away from R&D and toward product development and commercialisation.

The prescribed objective of the current R&D grant programme administered by Callaghan Innovation is to increase business investment in research & development to support long-term economic growth.[16] One of the criteria for awarding a Project grant is the extent to which the grant will "contribute to the development of a more stable and substantial New Zealand-based R&D programme within the business". If the series of grant programmes that existed prior to Callaghan had a similar objective, then evidence that the R&D programme does not increase future R&D expenditure relative to the counterfactual would be concerning.

Of course, it may be inappropriate to ascribe the objectives of the current system to the myriad of Project-like grant programmes that existed prior to Callaghan Innovation. It would certainly be wrong to ascribe to Callaghan Innovation's grant programme the outcomes that were achieved – or not achieved – prior to its creation. One of the main purposes for creating of Callaghan was to clarify the objectives of the New Zealand Government support for R&D and to bring the complementary activities together so they could work more effectively. Moreover, the Project grant – to which the grants examined in this analysis are meant to be equivalent – is only one of the channels for achieving the overall objective of increasing business expenditure on R&D, and it is only a relatively minor share (around 15%) of the total budget. A proper evaluation would take into account the suite of grant types that are available.

The results also show that in the 2-3 years after receiving a Project grant, recipients experience faster employment and labour productivity growth than non-recipients. Moreover, grant recipients are more

[16] See: https://www.callaghaninnovation.govt.nz/sites/all/files/guidelines-ministerial-direction-explained.pdf

likely to survive over the following 3-4 year period. However, consistent with Ministry of Economic Development (2011), receiving a Project grant does not have a positive impact on MFP.

The lack of evidence of an effect from R&D grants on MFP is surprising. The grants appear to have a positive impact on innovation, and this flows through to higher employment and labour productivity, but not to MFP. Traditionally, MFP is taken to represent a firm's "technology" – the portion of output that is not explained by its inputs – and MFP growth to represent technological change. As productivity growth is the primary driver of long-term economic growth, a necessary condition for R&D grant to support long-term economic growth is that it increase productivity.

However, other work (Wakeman & Conway, 2017) on New Zealand firms has also struggled to find a link between innovation and the productivity growth. That research finds that, in general, firms engaged in innovation – and specifically those that are also engaged in R&D – do not experience higher productivity growth compared to non-innovators. There is a connection for certain types of firms (e.g., start-ups, manufacturing firms, and firms with international connections), but not over all firms. The finding that there is an impact of R&D grants on innovation but not on productivity is consistent with this result.

As the discussion above indicates, whether the R&D grant programme is judged to be a success depends on the programme's objectives. Therefore it is important to be clear in setting up a government programme about what are its objectives. Under Callaghan Innovation the prescribed objective is to increase R&D expenditure, although whether that is true specifically for Project grants is less clear. This means it should be possible to use the same methods used in this paper to evaluate whether the R&D programme has been a success.

However, for this to happen the data necessary to conduct an evaluation needs to be available. At present not all of the data necessary to evaluate Callaghan Innovation's R&D programme is being collected. For instance, reliable R&D expenditure information is only collected for a subset of firms through the R&D and Business Operations Surveys, and not even for all R&D grant recipients. Efforts by Callaghan Innovation to collect this information from its customers going forward (including for a number of years after the grant has been paid) are to be encouraged. It would also be helpful to include a question on R&D expenditure into the Annual Enterprise Survey to capture this information for a larger set of non-recipients.

At the same time, it is necessary to improve the information on R&D grants (or government assistance more generally) available in the LBD. Right now this information is only merged into the LBD on an ad hoc basis, with a multiple-year interval between the last two times this was done. To be able to evaluate government programmes in as close to real time as possible, this information needs to be added to the LBD on a regular basis. It would also be helpful to improve the matching between recipients and the firms' LBD records. Using the New Zealand Business Number should facilitate this.

This research note has illustrated the potential for using the LBD to evaluate the impact of New Zealand Government assistance programmes. The rich data in the LBD makes it possible to measure outcomes using a range of variables with a sufficiently large number of firms to construct a counterfactual control group. However, it has also highlighted some of the limitations of this, including the time delay in data becoming available in the LBD, the lack of data on some variables, and imperfect matching of recipients to the Longitudinal Business Frame. Addressing these issues will make the LBD a much more valuable tool for monitoring the impact of government interventions on firms.

References

Angrist, J. D., & Pischke, J.-S. (2009). *Mostly harmless econometrics: an empiricist's companion*. Princeton, NJ: Princeton University Press.

Caliendo, M., & Kopeinig, S. (2008). Some Practical Guidance for the Implementation of Propensity Score Matching. *Journal of Economic Surveys, 22*(1), 31-72. doi:10.1111/j.1467-6419.2007.00527.x

de Serres, A., Yashiro, N., & Boulhol, H. (2014, April 2014). *An International Perspective on the New Zealand Productivity Paradox*. New Zealand Productivity Commission Working Paper, (2014/01). Wellington, New Zealand.

Fabling, R. (2008). *Will a BERD model fly? Estimating aggregate R&D expenditure using a micro model.*

Fabling, R., & Maré, D. C. (2015, Sep). *Production function estimation using New Zealand's Longitudinal Business Database*. Motu Economic and Public Policy Research Working Paper, (15-15).

Feng, A. (2016). *R&D grants: 2005 – 2009 - A look at descriptive trends*

Heckman, J., Ichimura, H., Smith, J., & Todd, P. (1998). Characterizing Selection Bias Using Experimental Data. *Econometrica, 66*(5), 1017-1098. doi:10.2307/2999630

Jaffe, A., & Le, T. (2015). *The impact of R&D subsidy on on innovation: a study of New Zealand firms*. Motu Working Papers, (WP 15-08).

Ministry of Economic Development. (2011). *Evaluation of the Impacts of Cross-Vote Government Assistance on Firm Performance, Stage 2: Impacts of Direct Financial Support for R&D*. Ministry of Economic Development. Wellington.

Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika, 70*(1), 41-55. doi:10.1093/biomet/70.1.41

Wakeman, S., & Conway, P. (2017). *Innovation and the performance of New Zealand firms*. New Zealand Productivity Commission working paper. Wellington, New Zealand.