

Session 4

| | |
|---|-----|
| Keynote address: Professor Bronwyn Hall, University of California at Berkeley | 122 |
| Discussant: Dr Simon Wakeman, New Zealand Productivity Commission | 132 |
| Discussion | 139 |

Patrick Nolan: Welcome back from afternoon tea. You've got an exciting afternoon of presentations lined up.

I'm thrilled to have Professor Bronwyn Hall and Dr Simon Wakeman presenting the next session, where they will get into a little bit more detail around the relationship between innovation and productivity and particularly draw on some of the firm-level data. So it should be a very good session. A few times I've mentioned the potential of linked micro-data, so this is where we actually start to see some of the fruits of that agenda.

Professor Bronwyn Hall's an Emeritus Professor at the University of California, Berkeley. She is also a Research Associate of the National Bureau of Economic Research and the Institute for Fiscal Studies in London and a Fellow of NIESR which is the National Institute of Economic and Social Research, also in London. She currently serves as an Associate Editor of the Economics of Innovation and New Technology and of Industrial and Corporate Change and she's also a member of several advisory boards.

Simon is a Principal Advisor at the Productivity Commission where he leads the Commission's research on the relationship between innovation and the productivity of New Zealand firms. Prior to joining the Commission he was an Associate Professor at the European School of Management and Technology where he taught courses in Business Strategy and Capturing Value from Innovation. So I'm very much looking forward to the presentations.

I'll pass over to Bronwyn. Bronwyn, thank you.



Keynote address:

Professor Bronwyn Hall, University of California at Berkeley

Bronwyn Hall: I want to thank all the organisers, especially Patrick, for not only inviting me to this wonderful country – I'm really happy to be here – but also making the trip so smooth and so forth. It really has been a good experience.

Simon and I have cleverly divided the topic. I'm not talking about New Zealand and he's talking about New Zealand, but I think that will work out since he knows more about it than I do.

This presentation is an overview of what we've learned about R&D innovation and productivity at the firm level. It's going to focus on what people have learned from the Community Innovation Survey which was pioneered in Europe but which is done in many countries, including this one. They don't call it the community innovation survey, but they do a survey which has questions which are similar to the ones on the Community Innovation Survey. And the US is a laggard country, but they've finally begun to do an innovation survey as part of their R&D survey a couple of years ago. Well, at this point, probably five or six years ago.

I want to talk a little about a framework for interpreting some of the results that come from the measurement of the relationship between research and development,

innovation and productivity, and mention some of the problems. As usual when you do empirical research you discover some problems and so I'm going to mention how we might improve things in the future a little bit. I want to emphasise that this analysis that I'm reviewing here, and it's partly for time reasons, is really only the micro firm-level analysis rather than the economy-wide results of R&D and innovation. The fact that we believe that there are spill-overs already tells you that economy-wide results might be different from just adding up the firms'. And that's an interesting topic. It's just I couldn't do everything. I had some slides at one point, but I've just had to shorten things.

So what are the mechanisms that we think connect innovation and productivity?

So what are the mechanisms that we think connect innovation and productivity?

The one that we focus on most of the time is improvements within existing firms, namely, creation of new goods and services and process and organisational innovation which lead to efficiency gains. Now at the economy-wide level you want to keep in mind that big mechanisms, which, among other people, Eric Bartelsman has studied, are the entry of more efficient firms and the exit of less efficient firms. That has good economy-wide effects. The entry of firms with new technologies are on the frontier, leaving aside the question of efficiency. And I'm not focused on those set of topics.

So to summarise my view of what we know in this area – by the way I know it's fun to photograph slides and I think you should, but I guarantee you these will go on somebody's website, on mine if nobody else's, so you won't have to photograph every slide. I just saw somebody photographing, so I figured maybe I should save them the trouble – we know a great deal about the contribution of R&D and innovation to firm-level productivity as conventionally measured. There's been a lot of research using data from a lot of countries



and that's a subject which has been studied for – I always have to do arithmetic here – probably the mid-50s – I mean econometrically been studied from the mid-50s to today. That makes 60 years.

We know something about the contribution of entry and exit to aggregate productivity growth from work. Again, I mentioned Eric and all the work with US data also. We know something about the contribution of R&D to quality improvement and therefore to productivity growth that comes from lower prices. The simplest example, of course, is lower prices for the computers that are sitting on everybody's table. When I started computing I worked on a room-sized computer that was nowhere near as powerful as the little laptop I have sitting there. I started computing in 1963 and in those days, you know, they filled up rooms.

We know much less about some topics we'd like to know something about, which is the contribution of R&D and innovation to welfare and poorly measured but important outputs, namely, health, environmental quality and things of that kind that are basically not priced...

We know much less about some topics we'd like to know something about, which is the contribution of R&D and innovation to welfare and poorly measured but important outputs, namely, health, environmental quality and things of that kind that are basically not priced, either because they're totally public goods or because we subsidise them at the government level and price signals are not very good, as in the case of health. And I think this got mentioned earlier by Adam, among

other people. It's one of the reasons why even if you accept the BERD goal, you still have this problem, is it producing the GDP or the welfare change that you really wanted? And without measuring these other outputs it's hard to say that the quality of life has improved.

...even if you accept the BERD goal, you still have this problem, is it producing the GDP or the welfare change that you really wanted? And without measuring these other outputs it's hard to say that the quality of life has improved.

 **Donal Curtin** @donal_curtin: Bronwyn Hall from Berkeley is giving #innovatenz a v good summary of what we do & don't know about effects of R&D/innovation on productivity

...labour market regulation is very important in influencing how effective innovation can be in a given country.

We don't have full aggregate growth models that incorporate behaviour of individual firms that let us look at things in detail, such as regulatory impacts which is really quite important. When you look across countries, regulation has a big influence, not

just regulation of goods, but labour market regulation is very important in influencing how effective innovation can be in a given country.

I want to talk about something that was mentioned earlier which is the distinction between research and development and innovation, because it's important. That's why the Survey got started and why these surveys of innovation got started in the first place – was a perception that R&D wasn't the whole story. It's especially not the whole story if you're in a service economy. The really important innovations in the service economy are frequently not – they come from other people's R&D, but they don't come from your own R&D.

... R&D doers are more likely to be innovators, but there are plenty of innovators that don't do R&D...

The first thing is – I'd show you Italian firms, just because that was the data I had – and the first thing to say is, yes, R&D doers are more likely to be innovators, but there are plenty of innovators that don't do R&D, based on these surveys. It's especially true in the service sector, partly because they rely on purchased technology to a great extent – customer management software, this sort of thing. So that means that just focusing on R&D might not be the same thing as looking at innovation.

...what do firms spend on innovation?

The second thing is that when you go and ask, "what do firms spend on innovation?" – I'm now giving you the UK numbers on these figures and I believe Simon has some numbers for New Zealand – this is the average breakdown across firms, not the totals, on the spending on innovative activities.

The first thing to note is the category Acquisition of machinery and computer hardware/software associated with innovation. Okay, I didn't put that on the slide, but this is not just I went out and bought new machinery. This is, whatever I bought was associated with the introduction of new processes in my firm – is the largest category. Even in manufacturing. R&D is smaller, so the real spending is this equipment spending. It's even more important in services and "Other" – I forget what's in "Other" – not much, it's mostly services, maybe transportation or something. I mean this is economy-wide, but there are certain sectors, primary sectors are excluded here. I'm not including those.

...internal R&D, even in manufacturing, averages a quarter of total innovation spending.

The other thing to note is you have these other things which become much harder to measure as you go down. Marketing expenses associated with the introduction of new products, processes, training expenses associated with the introduction of new product and processes or processes design expense again. Buying external knowledge or spending money. Buying external knowledge is distinct from spending money on R&D externally. It's more like I took a patent licence. That's why it's a different category. Those are all much smaller, but the main message here is that internal R&D, even in manufacturing, averages a quarter of total innovation spending.

So now you go back and say, "well how have people measured innovation in the past?" You all probably, or many people in this room, know that mostly, most of the work has been done using R&D flows or historical stocks, meaning depreciated accumulated – past R&D depreciated – as proxies for innovation input. In manufacturing that's not a terrible idea. We have a survey on that – what the results are.

There's a smaller literature that uses patents as a proxy for intermediate innovation output. There are weaknesses, having to do with, they don't cover all innovation, especially patents, in earlier periods.

...[do] the direct measures of innovation improve our measurement of the returns to innovation via higher productivity.

So the reason for going to the Community Innovation Survey is to see whether the direct measures of innovation improve our measurement of the returns to innovation via higher productivity. It's a question of how does that fit in the productivity relationship. We know that R&D fits reasonably well for the firms that do R&D. That's not all the firms, and we'd like to see if we could expand our coverage. So just to review what's in the surveys on innovation.

The first thing and the most important question, and the question that is used most often for analysis, is they ask the firm, "did you introduce a new product during the last three years that was new to the firm or new to the market?" In New Zealand's case, I believe they also ask whether it was new to New Zealand or new to the global market. In Canada, also, they ask that question I think, but most countries don't make that distinction.

Sometimes they ask about organisational innovation separately from process innovation. And another question which turns out to be a useful question is they ask firms, "what's the share of sales during the past three years that came from products that were new to the firm?" And that question turns out, based on firm interviews, to be one they know how to answer, in general. Even my tiny firm knows the answer to that question. I mean I don't have it any more, but I had it for 30 years. Because they track sales that way, so it's, it's actually a question that they know they can ... because they track sales by product and so they know when they introduce a new product they have a new item in the register.

The second thing that you get in these surveys is data on productivity employment and what you don't get generally is capital expenditures, which means they're good at measuring some form of broad labour productivity, but they're not going to be terribly good in many cases at measuring TFP – which is okay. In most developed countries that's not a serious problem.

There's a survey by Mairesse and Mohnen in the Handbook of the Economics of Innovation that goes into a lot of detail about these measures and what countries have them and how they vary and all sorts of things like that.

When I get to the framework here, what I want to try to explain is this issue that has to do with price and quantity, because it affects how you interpret the results and I think it's actually potentially quite interesting. At the firm level, when we take productivity, labour productivity, which is usually measured as sales per employee or TFP based on employees and capital – by sales, I mean net revenue or net sales – the typical pattern is that you're not deflating the revenue by a price index that really captures the quality of the product that the firm is selling. You're deflating, generally, by an industry-wide price index and if this firm is producing much higher quality products than another firm in the industry, what you're doing is you're not accounting for quality change by this firm. That's not a bad thing. I just want to be clear that what we're measuring with revenue is price times quantity.

...[these surveys are] good at measuring some form of broad labour productivity, but they're not going to be terribly good... at measuring TFP...

...my interpretation of what the innovation is, is affected by whether it's price or quantity. It might make me more efficient... [or] people are much more willing to buy my product because it's much more higher quality... increasing price or increasing quantity, both will increase my revenue.

Why do I care? I care, because my interpretation of what the innovation is, is affected by whether it's price or quantity. It might make me more efficient, innovation, which means that I can produce more output per unit of input. Another possibility, and in fact the one that is dominant, is that what it causes is demand curves to shift out, which is to say people are much more willing to buy my product because it's much more higher quality, therefore I can charge a higher price. I might be producing the same quantity but I can charge a higher price. That will increase my revenue also. Increasing price or increasing quantity, both will increase my revenue. But as an econometrician, my interpretation of what happened will be affected by how I had deflated to begin with. I'm not doing this with equations, because I think that makes sense, but it makes the explanation much more wordy.

So if you assume that innovation can affect both price and quantity and you've got the product of the two and if you assume some fairly innocuous assumptions – which is, the first one, in perfect competition, which basically characterises almost all sectors, and then you assume that the goal of process innovation is to reduce cost, to make the firm more efficient and the goal of product innovation is typically to shift the demand curve out, which is to say

consumers have higher willingness to pay for the same good, the higher willingness to pay for the computer or whatever because it's much better. They like to pay a lot for the iPhone 6 because it's a whole lot better than the iPhone 4s that they have, or it's a higher quality good for the same price which typically characterises, in fact ICT. I've bought laptops every two or three years for the last 15 years and for a long time I was just paying the same price every time, but the computer I was getting was a whole lot better every time. So if you assume that the product is really pushing demand and process is really pushing efficiency, and you can actually derive what the implications are going to be.

...they actually built a supply and demand model and they were able to show that ... process innovation spending did look like it was going towards efficiency and product innovation spending did look like it was pushing demand out.

I just want to take a little detour and tell you that there's work in progress. I just talked to Emile Petrone in Minnesota and they're still working on this, but unfortunately the paper that's out is dated 2011 and they're expanding to other data. But this is the only case that I have seen where they actually had R&D spending at the product and process level, so the firm could tell them I spent this much on product and this much on process innovation – because most firms can't tell you that. Of course that's with error, right. Obviously not every firm's going to be able to make that split, but still what they did was they actually built a supply and demand model and they were able to show that indeed, as we would expect, process innovation spending did look like it was going towards efficiency and

product innovation spending did look like it was pushing demand out. In other words, there was quality improvement and so forth.

So they were able to identify this contribution separately. So I throw that out as support for the notion that the two can be considered somewhat distinct, but not completely because they tend to go together. The reason why this matters – why does it matter? It matters because product innovation, you can easily show is unambiguously going to increase revenue productivity and labour demand. Both labour demand and the firm will want more workers and it will increase its revenue productivity.

And in many economies there's concern that innovative activity may be so good that what it does is reduce labour demand... That's of concern to policy makers because their biggest concern is making sure there are enough jobs.

Process innovation increases revenue productivity and labour demand only if the demand is elastic, and even in this case it gets dampened. And I throw this out because the data show that this is the case. It hints at this being the case. And in many economies there's concern that innovative activity may be so good that what it does is reduce labour demand, okay. That's of concern to policy makers because their biggest concern is making sure there are enough jobs. There's a larger debate going on right now about digital technologies and labour but I won't get into that.

The allocation of where you think the innovation is going between price and quantity depends on the type of price deflator you're going to use – whether you deflate by the quality adjusted price or whether you deflate by just the GDP deflator or something like that. But that's just an allocation. From the firm perspective, the firm just cares about the revenue, okay, not about the quantity or the price. We might be interested in the quantity and the price, but they're concerned about the revenue.

So to back up here a little, what do the data say? The survey that I have here – which I'm not going to go into in detail because you don't really want to look at all of these numbers; I'll summarise what it says – comes from a large collection of papers that have used a standard model, which is a cross-sectional model, which means that it has the limitations of a cross-sectional model, namely if there's anything left out that's correlated with what's going on inside the model you may have some upward or downward biases to some of the coefficients. Unfortunately, doing more than this is very difficult using innovation surveys because innovation surveys in most countries are based on a sample of firms and the sample is redrawn every time the survey is done in order to reduce respondent burden. The consequence of that is it's very hard to look anything long-term at an individual firm level, which is what you'd really like to do, and you can't really do it because of the way the surveys are designed.

It's a simple model where you essentially say, is the firm going to do R&D and how much is it going to do, and you have equations to describe that behaviour. And that R&D is going to produce some innovations, could produce process and product or it could produce increased sales of new products – whatever innovation measure you choose, you choose. And then you say well the production function will include these innovation outcomes in the production function but we will instrument them – we will predict them. We'll say this is what the firm is expected to get as an innovation, given how much R&D it's doing, or not doing.

One feature of this model – Adam mentioned the fact that firms, especially smaller firms,

are often doing R&D and not measuring it – and one of the features of this model is that you predict R&D for those smaller firms from the first equation. So you predict if this firm did not report how much R&D it was doing but we’re going to assume that it did the amount you would predict given its size, its industry, all of its characteristics – exports, foreign ownership, the things that Beth mentioned. And then we’ll use that as their R&D. You don’t want to know about the econometrics.

...it could also be that process does indeed contribute to productivity, just not as much as product innovation.

It’s been estimated for 20+ countries. It does confirm high rates of return to R&D that have been found in a lot of the earlier studies. Innovation output statistics are a lot noisier than R&D and the reality is that innovation doesn’t add much to the model if you have R&D. You basically don’t need it. The trick is that there are a lot of firms you don’t have R&D for, so you still might want to do things this way. What I did a couple of years ago was summarise the results, and the main thing that I want to point to in this slide is that if you include both product and process innovation in the regression, which is the top half of the slide, the top section of the slide down to that line, basically you get nothing for process, or negative, which is really what you had predicted from thinking about it. The bottom three are where they are included separately. When you include process separately and you don’t include product you get something from process innovation. It does increase productivity. One of the reasons for that, of course, is that process and product are often done together, but it could also be that process does indeed contribute to productivity, just not as much as product innovation.

So the TFP on innovation sales, the elasticities range from meaning, – Point 4 is basically Point 2 – if my innovative sales share doubles, that’s 20% on productivity and productivity goes up 20%, so that’s pretty high. What you also find is that the R&D-intensive and the high technology firms have higher elasticities. That is entirely predicted if you think rates of return to R&D are equalised across sectors, and this is a point I cannot emphasise enough because I’ve had people come and tell me over and over, “oh it’s a high rate of return – we should invest more”.

No. If you have an industry where there are high shares of R&D you will have high rates of return and industries with high share of R&Ds will have higher elasticities. Industries with low shares of R&D will have lower elasticities. That will translate into roughly equal rates of return for the additional R&D dollar, and I won’t do the algebra. So, let’s skip that. I’m just conscious, if I have 15 minutes, is that right? No, I don’t, because I’m supposed to finish in half an hour. I just want to make sure I get through this. I’d hate to rush.

We have done some research where we tried to look at the employment impacts of innovation. The idea is we want to decompose employment change and given data limitations this is the decomposition we can do. We can say there’s an industry productivity trend in old products. There’s growth due to process innovation in old products. There’s growth due to output growth of old products. There’s growth due to product innovation. This is now across Europe, for which I have data, and that shows that the average industry-specific employment growth, without innovation, is negative in all countries. It shows up here right away. The process innovation in old products is pretty positive in these European countries for employment growth and product innovation is, on average, equally positive but it has a much lower impact for Italy. It has a higher impact in France and Germany, lower in Spain.

I want to emphasise – by the way, this is 1998-2000. This is the period we can compare. This is a terrible period to compare, okay. Does everybody know what happened in 1999 to 2000?

Male Participant: Y2K.

Bronwyn Hall: No, not Y2K, much worse, much worse. The Euro was introduced. It's very clear from the Italian accounts that they didn't adjust quickly. There are a lot of bad numbers in the Italian data, because you had to adjust your entire system of accounts to go from Lira to Euros, so that was a problem. But I have this as the only period for which I have comparable data. The other interesting thing to note here is non-innovators don't have any employment growth, period. They're barely on the graph. And that's not relative to anything – that's just they don't have any. So that's a summary, but let's skip that.

...does innovation spending tell us something different than R&D spending[?]

Let's just go straight on to innovation spending. The idea here was, does innovation spending tell us something different than R&D spending, because for the UK I actually had innovation spending. And this graph shows you for the larger European countries, of the 28 European countries, with the key exception of two countries who apparently don't like to share the data with the rest of the EU – namely Germany and the UK. How about that. But they have everybody else here and it goes from Poland all the way down to Cyprus, which is a bit of an exception, and what it shows you is that the blue is intramural R&D. The red is extramural R&D. The green is machinery equipment software and the purple, which is small, is the other external knowledge acquisition. They don't all measure all of the things that I showed you before.

But the interesting thing here is the variation. Across countries there's real variation in shares here. Poland is all about new machinery equipment and software coming in from the Germans so they can run their plants in Poland. Basically, it's what that's about. And you see

it in the Czech Republic, another one. The next one is the Czech Republic with the big green, and the next one after that is Hungary. So we know what happened there. This is 2010 – there was a lot of stuff coming in.

Compared to R&D, innovation spending is more strongly associated with information from suppliers and innovation to meet environmental or health and safety standards, which basically is an equipment story.

The UK evidence, I used total innovation spending. Twice as many, or nearly twice as many, firms have innovation spending as have R&D spending. The median innovation spending is five times the median R&D, as opposed to that four that you saw before. Compared to R&D, innovation spending is more strongly associated with information from suppliers and innovation to meet environmental or health and safety standards, which basically is an equipment story. I should have mentioned that what I did with the innovation spending, was substitute it for R&D in the first part of that model and run the whole model again. It's less strongly associated with exporting and collaboration with other firms and information for consumers. That is, it is more process than product. That makes sense, right. It's a better predictor of innovation probability. However, doubling innovation spending has the same impact on TFP as doubling R&D and I haven't fully figured out why that should be the case. It's interesting.

...innovation dummies are a serious measurement problem.

So my conclusion, and I didn't emphasise this – I'm anxious to make sure that we have time for discussion – one thing I didn't emphasise so I should say more about it: innovation dummies are a serious measurement problem. If you think about asking a firm like IBM, did you introduce a new product in the last three years – that is a stupid question. And so you get people writing papers saying, "large firms innovate more". Well, they don't innovate more – it's just if you ask them the dummy question, of course they introduced a new product. If they didn't they hopefully exited the market. Whereas a small firm may not introduce new products very often but it may be perfectly viable, and if you look at the data, you'll see that. The shares of firms that innovate go climbing up as you increase firm size, but that isn't really informative to you, because that's something you knew a priori and it doesn't really capture how innovative the firm is. The share of sales due to new products does capture it. It actually turns out to be a much better measure. So that's the first thing.

The problem with the share of sales due to new products is that we don't have a corresponding measure for process innovation...

And the only country that has ever dared to do this is Germany, and they did once ask how much cost reduction they got from their process innovation and they got answers, and it's in Bettina Peters' thesis, but that's the only time it's been done. My understanding is that people understand this is an issue. That's a much

harder question to answer. That's the reason why they don't ask it now.

Share of sales to new products is something firms track. Cost reduction from process innovation may or may not be tracked by firms and so there's ongoing work. The OECD is basically exploring how they can improve this level of this question.

But the reason I bring this up is because when you stick process and product innovation in the same productivity regression, because they're both measured with incredible error – they're both very noisy indicators – you essentially will get typically equal – you'll get positive on the one and negative on the other and it doesn't really mean that that's what the story is. It doesn't mean that one type is good for productivity and the other type is bad for productivity. What it means is that they're measuring the same underlying thing, which is innovation, because they are correlated, although not perfectly, and that they're both very poor measures and one of them is a little poorer than the other one, and that's the one with the negative coefficient. So I just bring this up because I've seen it so often in many results and it sometimes gets misinterpreted.

So I'll stop there and go on to Simon.
[Applause]



Discussant:

Dr Simon Wakeman, New Zealand Productivity Commission

Simon Wakeman: Thank you. Bronwyn has outlined to you some of the challenges of doing research in this area on innovation and productivity and so it's my job now to translate this for the New Zealand evidence and understand what we can learn from the research we've been doing through the Productivity Hub on these issues.

...while New Zealand's close to the top of the OECD in terms of scientific research, we're much less effective in getting commercial value out of that research.

So to start off, the international evidence suggests that while New Zealand's close to the top of the OECD in terms of scientific research, we're much less effective in getting commercial value out of that research.

This chart shows you the average ranking out of 40 OECD member and affiliated countries in a range of categories, ordered from the upstream science end to the commercialisation end. And it shows that as we move further down the chain, the worse New Zealand looks.

One thing that the international evidence highlights is that New Zealand businesses invest relatively little in R&D. This is a fact that has drawn a lot of attention in recent years and has already been covered earlier today, particularly by Adam in his talk.

But one thing that he highlighted is that the research shows us some things that are correlated with this, or potential explanations for this, but doesn't give us a smoking gun, doesn't tell us why New Zealand firms invest relatively little in R&D or, by extension, innovation. Possibly the most obvious potential explanation is that New Zealand firms may earn relatively little from investing in innovation, but apart from some anecdotal evidence, it's not something we have a very good handle on.

Through the Productivity Hub and our collaboration with Motu we've been trying to generate some quantitative evidence on the returns to innovation for New Zealand firms that hopefully will help us to shed light on what drives this relatively low investment in innovation by New Zealand firms. For this we use Statistics New Zealand's longitudinal business database which integrates data from a number of different sources, including financial data from tax records and financial surveys, firm responses to the innovation module of the Business Operations Survey, as well as patent and trademark filings. However I'm also required in using this data to tell you two things: first, that this information has been confidentialised or anonymised and access controlled to protect the identity of the firms and their data, and second, that the results do not represent official statistics.

 **Donal Curtin** @donal_curtin:
And now @NZprocom's Simon Wakeman is giving us the empirical innovation/R&D story on NZ data #InnovateNZ

...we can count 40% of New Zealand firms in any year that are innovating.

The LBD contains a range of measures of innovation, including both inputs and outputs of the innovation process, and in this data we have a whole range of results about how many New Zealand firms engage in innovation. On the narrowest measure, only a fifth of a percent, are actually filing patents. However, when we think very broadly, introducing any innovation of any type new to the firm, we can count 40% of New Zealand firms in any year that are innovating. And the LBD data also contains some measures of innovation intensity within the firm, including the one that Bronwyn just mentioned, the share of returns that come from new products.

Only a quarter to a third of firms that engaged in innovative activity are also engaged in R&D...

It's only when we [look at]...innovation that's new to the world, either filing patents or introducing products and services new to the world, that we actually find that more than half the firms that innovate are engaged in R&D.

One thing that this data highlights is that R&D is only one of the inputs into innovation. Only a quarter to a third of firms that engaged in innovative activity are also engaged in R&D. And this is true even when we talk about specific types of innovation – product innovation or process innovation. It's only when we get into the very narrow definitions of innovation, the innovation that's new to the world, either filing patents or introducing products and services new to the world, that we actually find that more than half the firms that innovate are engaged in R&D.

There's also data within the Business Operations Survey on the makeup of the innovation spending or product development spending, and we see that R&D expenditure is a big contributor in some industries, but not in all. So looking across the sectors we see that in primary and manufacturing sectors R&D expenditure makes up about 80% of the innovation or product development spending, although when we break this down into specific types of industries, we see that this is driven primarily by Machinery & Equipment and Food & Beverage with other manufacturing industries being much lower R&D intensity, while R&D expenditure only makes up a smaller fraction of innovation spending in the services – one-third



in the total services sector and less than 10% in some services sectors such as wholesale trade, retail trade and in the hospitality industries. In those industries, marketing or market research drives the bulk of our spending.

This highlights one of the weaknesses of using R&D as a measure of innovation and pushes us to look more at some of these output measures of innovation. In the results that follow I'm going to concentrate on three of these output measures. They are measures of product innovation, process innovation and organisational innovation, partly because I think these are the broadest definitions of innovation. Also, they're the ones that seem to show the strongest results.

Bronwyn recommended the use of the sales from new products innovation, but I've also looked at this as a measure, but find that although we get similar results to what we get on product innovation, the results are very noisy. It's perhaps because the way it's measured in New Zealand is quite crude. The spending on the sales from new goods and services is grouped into some broad categories and we get a lot of variance around our results.

This chart shows you the results from a series of regressions where I've regressed the output measure, one of the various measures of firm output or firm performance, on the individual innovation indicators and doing them separately, one by one in separate regressions. The error bars on these charts show the 95% confidence intervals and the size of the bar is the difference between innovating and non-innovating firms. So we have employment gross output or revenue, which you know what those mean. The productivity estimates used here, particularly the multi-factor productivity estimates here, what Bronwyn referred to as TFP – or we usually refer to as MFP – are derived by the approach developed and tested for New Zealand by Richard Fabling and Dave Maré, and if you want to know more about this it's described in a recent working paper that Motu produced for the Productivity Hub.

...firms engaging in any type of innovation are around 75% larger, both in terms of employment and annual gross output. However, firms that innovate do not appear to be any more productive than firms that do not innovate.

We see in these baseline results that firms engaging in any type of innovation are around 75% larger, both in terms of employment and annual gross output. However, firms that innovate do not appear to be any more productive than firms that do not innovate. There's some evidence that organisational innovators have around 20% higher productivity, and you see that the error bar here is not overlapping with the zero axis, so that this is statistically significant. However, when we look at multi-factor productivity, which captures more the knowledge-based capital, it's not significantly different. And for product innovations the differences aren't statistically different and for process innovation we see no differences at all.

However, as Bronwyn has just outlined, measuring these relationships is not easy and there are a number of issues we need to consider. For a start, using revenue or gross output instead of physical output to measure productivity distorts the relationship between innovation and productivity, especially with respect to process innovation, and so we have to be careful in our interpretation. Also, because multi-factor productivity is the residual from a production function, it depends on the accurate measurement of a number of different inputs and it's likely to be quite noisy. So the way the literature that Bronwyn described deals with this is to use a multi-stage model that instruments for innovation with factors

that predict whether a firm innovates, such as R&D. However, most of the prior literature is in this cross-sectional context and one of the problems with this is that you cannot identify whether the relationships you observe are driven by innovation itself or other factors that you would expect to be correlated with innovation but you can't control for.

So in this research I take a different approach, which is to look at change in firm outputs or productivity over subsequent years. This controls for the unobserved firm characteristics that we think might be correlated with innovation that also drive productivity. I also control for those firm characteristics that we do observe, such as age, size, industry, whether the firm is exporting, etc.

Finally, to deal with the measurement error in MFP, productivity here is calculated as a two-year moving average. So, using the year before and the actual year in which it's calculated.

...firms that are introducing innovation grow faster than firms which are not, and after three years the innovating firms are on average between 5 and 10% larger in terms of output.

These are results in terms of revenue or gross output – we see that firms that are introducing innovation grow faster than firms which are not, and after three years the innovating firms are on average between 5 and 10% larger in terms of output. Product innovation and organisational innovation show stronger results than process innovation, although there's still a statistically significant difference in the first year in terms of process innovation.

...only firms introducing product innovations and organisational innovations are clearly doing better, with around 3% higher growth after 3 years.

However, to emphasise, this is only growth in terms of total output. When we look at productivity, so output for inputs, we see that only firms introducing product innovations and organisational innovations are clearly doing better, with around 3% higher growth after 3 years. There does not appear to be any difference between firms that introduce process innovations and firms that do not, although as Bronwyn has pointed out, process innovation could show up in terms of lower prices rather than higher demand, and so may not be reflected in the relationships that we observe.

...it's firms in the 5 to 10 year age group that see the most improvement after innovation... somewhere between 6% and 10% improvement over 3 years, relative to firms that do not innovate.

Now I look at how these results vary by firm characteristics. For this analysis I've interacted the innovation variable, it's a separate variable in each regression, with the firm characteristics. And I'm also focusing here on the three-year change, as that appears to be the period over which we saw the returns being clearest. And from these results we see that it's firms in the 5 to 10 year age group that see the most improvement after innovation. Innovating firms in this age group see somewhere between 6%

and 10% improvement over 3 years, relative to firms that do not innovate.

We also know from other research that this is the age where there's a lot of firms that either make or break, survive and grow or drop out. It also seems to be a stage at which whether the firm is innovating or not seems to make the most difference.

...whether a firm is engaged in R&D does not appear to make a difference to whether the firm benefits from organisational innovation.

Our data also allows us to separate firms by where the innovation comes from. We see that firms that are engaged in R&D are much more likely to show improvement following product innovation than those that are not. However, we know that only about a third of product innovators are engaged in R&D, but it seems to be this subset of firms that benefit most from doing product innovation. However, whether a firm is engaged in R&D does not appear to make a difference to whether the firm benefits from organisational innovation. So R&D is very relevant when we're talking about product innovation. It does not seem to be very relevant when we're talking about other types, particularly about organisational innovation.

Looking at other places in which the innovator might get its ideas from, we see that those that get their ideas from more formal sources such as universities, Crown Research Institutes and conferences do better following product innovation. However, those that get their ideas from their business environment, especially their staff, their customers and other businesses, show more improvement following organisational innovation.

...[innovating firms] that get their ideas from more formal sources such as universities, Crown Research Institutes and conferences do better following product innovation. However, those that get their ideas from their business environment, especially their staff, their customers and other businesses, show more improvement following organisational innovation.

Finally, we see that those firms with larger potential markets are likely to do better if they are innovating. This chart shows the three-year change in productivity for innovating firms by the extent to which the product that they produce is tradeable across the country, across regions within New Zealand. So firms with the least tradeable products, such as your local coffee shop, do not appear to improve much depending on whether they innovate or not. By comparison, those in industries where the products are traded across multiple regions or across the whole country do much better if they are innovating.

...firms with international connections, exporting, foreign-owned or with overseas investments themselves, do better if they're innovating than firms that are just focused domestically.

Looking more broadly, we see that those firms with international connections – exporting, foreign-owned or with overseas investments themselves – do better if they're innovating than firms that are just focused domestically. This is consistent with other research that shows that market size limits innovation. However, it highlights that for those firms that have the ability to expand their markets beyond their local community or even beyond New Zealand, innovation does appear to make a difference.

...it's not clear... whether innovation actually leads to productivity benefits for New Zealand firms, given what they invest in it.

At the start of this presentation I asked whether one of the main reasons that New Zealand do not invest very much in research and development or innovation is because the benefits from innovation are low. We don't have a full answer to this question yet, but the results are suggestive. We see that product innovations and organisational innovations show a higher productivity growth. The difference is only 3 to 5% percent over 3 years, although when you think of the average productivity growth of New Zealand firms being only 0.1%, this could be significant. However, this does not take into account the costs of innovation and the risk it creates for firm survival. And it's not clear, therefore, whether innovation actually leads to productivity benefits for New Zealand firms, given what they invest in it.

However, there are specific categories in which innovating firms do much better than those that are not. Among young firms, especially that 5 to 10 year age group, those that innovate are clearly doing better than non-innovating firms. At the same time we see that firms that are engaged in R&D or get their ideas from universities or more traditional sources, do better out of product innovation. However, firms that get their ideas from the sources in their business environment show improvement following organisational innovation.

...it shows the benefits of encouraging firms to engage in R&D and connecting with universities and so on if we want to create more product innovation. However, making sure they're connected with their business environment also matters, particularly given that organisational innovation appears to be one of the equal drivers of productivity growth.

From a policy perspective, it shows the benefits of encouraging firms to engage in R&D and connecting with universities and so on if we want to create more product innovation. However, making sure they're connected with their business environment also matters, particularly given that organisational innovation appears to be one of the equal drivers of productivity growth.

The other thing that appears to make a difference is the size of the market, both domestically and internationally. In some cases the size of the markets is exogenously determined by the characteristics of the product, something the firm can't do much about, but in other cases it is. The firms can reach out, they can become more connected both domestically and internationally.

...hopefully a better understanding of the relationship between innovation and productivity and particularly where innovating firms are more likely to experience productivity improvements, will help us shape our innovation policies to focus on where we can make a difference, and that's the next challenge.

Before I finish, I should emphasise that these results we see are merely correlations; that we don't have a lever we can pull and test whether that innovating makes a difference to productivity. And so we can't say explicitly whether there's a causal link between innovation and productivity. Nevertheless, hopefully a better understanding of the relationship between innovation and productivity, and particularly where innovating firms are more likely to experience productivity improvements, will help us shape our innovation policies to focus on where we can make a difference, and that's the next challenge. [Applause]

Discussion



Patrick Nolan: Great. Well, thank you, Simon and Bronwyn. Again two very good presentations. Bronwyn pointed out how we have to always be incredibly cautious about the data, particularly the data coming from some of the business surveys, which was interesting because I think it's also a challenge, I guess, for the work that you've been doing with the LBD, Simon, and I think other work you've prepared has shown actually how hard it is to measure innovation and exactly what is meant by innovation. I thought it was also fascinating, the importance of the product innovation relative to organisational and process innovations, which I thought was really interesting context for some of the earlier discussions that we had today – we were really driving towards management matters and some of these sorts of other factors. I guess that's a counter to that as well. So any questions?



Beth Webster: Thank you both for the presentation. There's some great graphs in there which I will borrow one day.

Just this issue of how to measure innovation as an investment. Bronwyn, you talked about it, and so did you Simon. It's a really vexed one, because until we know who's got the best rate of return, we really need a better measure. And we're not there yet, as you clearly illustrate. The R&D's flawed. It's a great money measure, but it's too narrow. The dummy measures are just dummies, but the share of sales due to new products is a mixture between outcomes and activity. So I don't know if that really gets it for me either. We really need a "how much did you spend on innovation", which is R&D, plus more.



if you introduce a new product in the last three years, that doesn't necessarily mean success, but at least it means you innovated. It could be that that product fails.

Bronwyn Hall: Well, I would distinguish those two things. In fact, when I estimated the model with innovation spending, I replaced R&D with innovation spending, not the innovation output. So I view the share of innovative sales as really an output measure, a success measure. And I view the innovation dummy as a measure of output of innovation. It is true that if you introduce a new product in the last three years, that doesn't necessarily mean success, but at least it means you innovated. It could be that that product fails.

Measuring something other than R&D is a good thing.

Whereas R&D is clearly an input and measuring the inputs is not a bad thing, because the input is the thing that the firm decides on and that's the thing you can tweak with your incentives. So you actually want both measures. So I distinguish among them, but I do think that, especially in service economies which is what most of the people above 20 on this chart are also service economies. The majority of the sectors, the service sectors are much bigger than anything else they have. Measuring something other than R&D is a good thing. I actually had a question for Simon, which is just a clarifying question which is, where was capital equipment spending in the innovation spending you had? Was it left out or was it, was it in "Other"? Because that's an interesting question, given how big it is in the UK.

● ● ●

Simon Wakeman: Well, thanks for both questions. To Beth's question, I haven't looked at the results with innovation spending. It was actually only Bronwyn suggesting that as a potential alternative to R&D that I discovered that we actually have this data on innovation spending in New Zealand. It doesn't seem to be the same definition as the UK. As far as I know, "Other" doesn't capture the capital equipment spending. That's more prototyping commercialisation costs and so on, so maybe that would be left out as a type of innovation spending. But at least it captures the marketing thing. So I'd like to redo the analysis with innovation spending.

Measuring innovation is an endless challenge. I think we don't have any perfect measure. We have patents, we have R&D spending, we have some of these innovation output measures. The best we can do perhaps is to look at our results across these multiple measures and see if we're getting consistent results, and if we're not getting consistent results, then asking why. But if we are getting consistent results, then we have a consistent story we can tell.

● ● ●

Patrick Nolan: Okay, Donal and then Eric Bartelsman.

● ● ●

Donal Curtin: A couple of questions for Simon. First of all, the low results are even on some of the measures in terms of the impact on MFP. Were you surprised by that? Do you think it's very robust? And if you do think it's really what's going on, it's a bit of an issue for us all, because we're at a Productivity Symposium and all this innovation and R&D doesn't seem to be contributing any.



Simon Wakeman: I don't think we actually had any significant negative results. Correct me if I'm wrong, but we did have some very low results on some of the indicators. With the process innovation, as Bronwyn pointed out, you have to be quite careful in interpreting those because process innovation may not necessarily show up in terms of higher revenue productivity for firms. But it suggests that either firms are being rational and not investing in innovation because it doesn't pay off, or there are other forms of things that they're engaging in that are not innovation that are helping to grow the firm. And so even without innovating, these firms are already growing. I can't dissect what is actually going on there, but there are potential explanations for it.



Maybe what we need is actually an Oxford comma between innovative and productive Kiwi firms, because if actually what we're showing is the link between innovation and productivity is a bit more complicated than we may have assumed.

Patrick Nolan: Eric. While the microphone's going there, I can't help myself but just pick up on Donal's question. Maybe what we need is actually an Oxford comma between innovative and productive Kiwi firms, because if actually what we're showing is the link between innovation and productivity is a bit more complicated than we may have assumed. Maybe they're completely separate things, or relatively separate or there's actually more we need to unpack. Eric.



Eric Bartelsman: I'll actually talk about that in the next session. I really liked listening to these conversations and presentations, and we've got all this data so we can take a step back, we can be the scientist looking at everything that's happening at a macro level. And we have in our mind a model where there are firms making decisions. I'd like to hear more about the firms making decisions, and I know Bronwyn's actually spent a lot of her life talking to R&D managers and innovative managers of firms. If the calculations are hard for us to make with all the data on what the rate of return is, how is a manager making this decision ex-ante going to take a stab at what the rate of return might be and how do they make those decisions?



Bronwyn Hall: Maybe there's a manager in the room who will actually answer this seriously, because I really haven't talked enough to be an expert on the topic.

The first rule of thumb is to look over at the guy next to you in your industry and ask what his R&D intensity is, and then that's how you choose yours.

The first rule of thumb is to look over at the guy next to you in your industry and ask what his R&D intensity is, and then that's how you choose yours. [Laughter] And that's about as good as any indicator, because in fact he faces all the same uncertainties that you face. So it's a question of pooling information across members of the industry, all of whom are more or less seeing the same things that you're seeing, but maybe processing them differently. Quite seriously, they do actually do that – they benchmark to a great extent.

I do think, however, that individually their attempts to calculate rates of return ex post which might give you some info, and these guys, the MBAs, when they go to HBS [Harvard Business School] or wherever, I've seen the textbooks; the textbooks spend an awful lot of time teaching them how to compute internal rates of return. I presume that some of them actually try to do it, but it's clearly expected and what you realise is going to be quite different from expected.

However, I want to underline that's one of the reasons why I like R&D spending as a variable and innovation spending as a variable, is because it contains within it something about what the firm thought. It's the choice variable, and innovation outcomes are so uncertain in some ways that the firm clearly has influence on them, but doesn't really control them. And so they don't tell you as much about what the firm thought, because they're full of other pieces of information. They're both good, but they're very different animals. But maybe somebody here knows more about this topic.



Patrick Nolan: I don't know if you want to come in quickly, Simon, and then we've got David and then Richard. It's not compulsory.



Simon Wakeman: Okay, just continue with the questions.



Patrick Nolan: Okay, yeah over there.



David Hargreaves: Hi there, David Hargreaves from the Reserve Bank, so of course my question's about finance. I am interested in whether you've looked at all, Simon, at whether the sources of finance available to firms affect their decisions about whether to innovate or not? And I know from Bronwyn Hall's work, that's something that Bronwyn's worked on, so she might want to comment as well.



Simon Wakeman: The answer to mine. I haven't looked specifically in this research about decisions about whether to innovate or not. This is really looking more at the relationship between innovation and productivity. In previous work I've done I've looked at what are the challenges that New Zealand firms face in commercialising innovation. And I'd say that's not whether to innovate but whether to take your product to a larger market.

Finance does play a role, particularly at the level of going international, which for New Zealand firms can often be quite early, but also very expensive... being able to get the right personnel that are able to take your product overseas in terms of management capability, in terms of marketing ability and so on, can be as much of a challenge as getting the finance.

Finance does play a role, particularly at the level of going international, which for New Zealand firms can often be quite early, but also very expensive. It's having the funding to be able to take your product internationally. But there are other things that can be just as challenging. Being able to get the right personnel that are able to take your product overseas in terms of management capability, in terms of marketing ability and so on, can be as much of a challenge as getting the finance.



Patrick Nolan: Bronwyn, did you want to come in on that as well? Access to capital.



Bronwyn Hall: Oh, I'm sorry. I thought the question was directed at Simon, because the question was about the New Zealand firms, and I really can't comment on whether they have difficulties attracting finance. I guess I don't even know whether that question about access to finance is even asked on the survey. That's a question which is commonly asked on the Community Innovation Survey.



Simon Wakeman: We do have information on our survey about some of the barriers to innovating, but I haven't looked at those as part of this.



Bronwyn Hall: Okay. In the future one must look at them.



Patrick Nolan: Richard.



Richard Fabling: Hi, Richard Fabling from Richard Fabling. [Laughter] An observation and a question. The observation is, when we first put in place the Business Operation Survey and we looked at the Community Innovation Surveys, a lot of the reasons why we don't have some of the questions that are in the Community Innovation Survey is because we did a lot of cognitive testing here in New Zealand and we decided that some of the questions were either unanswerable by the firms that we were asking the questions of, or it imposed a massive respondent load on them. So we made a decision to not force them through that load and get potentially low quality responses.

My question is, and in the morning sessions we talked a lot about intangible assets and stocks of intangible capital and things like that, and the analysis that I'm seeing in this session I think, partly because it's been driven by cross-sectional studies, looks at annual investment in R&D. And is that really what we want to look at the returns to? I think of R&D as an investment in a stock, and a stock has a return over a number of years. Is that the right way to think about this and does it make sense to put R&D expenditure in one year and into a regression with productivity on the left-hand side or is that R&D expenditure in that one year acting as a proxy.



Bronwyn Hall: You're right. Of course, we don't have the data.



Richard Fabling: But we do have the data.



Bronwyn Hall: Yeah, we don't have it typically – we typically don't have the data. However, that's okay, because if you work in logs and you're working with established firms which most people are, because they're basically restricting the sample to something that is reasonably well-measured, it doesn't make any difference basically. Because in logs – I've done it with different depreciation rates – the depreciation rate doesn't enter in logs, because fundamentally R&D for most firms doesn't change a lot from year to year.

So when you take the log, you can write down the capital formulation for the R&D and if it's R over depreciation plus growth rate, if that's stable, that goes into the constant, and so the coefficient you get on R&D is the same, whether you use the stock or whether you use the flow. In a world in which R&D is fluctuating a lot and there are a lot of new entrants that's not going to be as true. But given cross-section results, you're looking across big size ranges – it doesn't really matter what you do. R&D, the flow, is going to be just as good a proxy as R&D, the stock. And the problem – the depreciation rate becomes a bit of an issue. So that doesn't worry me so much.

I'm interested in your comments about the Innovation Survey because he showed all the questions that we've been using, but I agree that there are many questions on the survey that I wouldn't want to go near, because I can tell that they're unanswerable. I'm curious what you dropped.



Richard Fabling: We don't collect a lot of dollar information, like the dollars spent on innovation.



Bronwyn Hall: Where does he get his innovation spending numbers from?



Richard Fabling: Spend, it's 5 to 10% of sales.



Simon Wakeman: Yeah, so we collect the variable but it's zero to 10, 10-20, 20-30 and the rest.



Bronwyn Hall: Oh yes, but, but most of the innovation surveys, the share of sales, they actually collect ranges. You're not the only ones. They collect 0-5, 5-10, this sort of thing. I was going to explain to you later, you actually have to transform the variable and you probably didn't. There's a problem with the distribution of the variable because it has a distribution like that. If you just stick it in, you won't get anything. So I was going to go through that later.



Patrick Nolan: Well, thank you. We're, we're right out of time. Thank you Bronwyn and Simon for your discussion. I think, if anything, you've highlighted the importance of being aware of the measurement issues and really being explicit. But this is an exciting area to be working in and I think in New Zealand we're quite excited about what we're able to do now with the micro-data, so I hope that we can continue these conversations. So if you could just join me in thanking Simon and Bronwyn. [Applause]