Comments on "The Use of Randomized Field Experiments to Assess the Effectiveness of Government Programs to Support Business Innovation: A Feasibility Study"

Kenneth J. McKenzie

Department of Economics

University of Calgary

Presented at the Industry Canada Roundtable Discussion of Field Experiments and Government Programs to Support Business Innovation

January 15, 2009

Ottawa, Ontario

I should state from the outset that I don't have any direct experience in field experiments, though I am generally familiar with the basic ideas.

I would like to begin by congratulating Industry Canada for considering the implementation of a field experiment in this area. I am not aware of any government initiated field experiments involving corporations of the type being considered here. I think it shows real vision. Well done.

I sense that the motivation for this study is a feeling on the part of the government, and perhaps Industry Canada in particular, that we "seem to be doing everything right" in terms of providing incentives for R&D and innovation, but that it doesn't seem to be paying dividends in terms of Canada's performance internationally. Ron Parker alluded to this in his opening remarks.

And taking a fairly narrow view of "doing everything right" – and I stress that I think it is a narrow view – this does indeed seem to be the case. We have very generous tax credits for R&D – among the most generous in the world, particularly when we take account of provincial programs which are added on top of the federal SR&ED. We also have generous grant programs, again at both the federal and provincial level. Yet Canada tends to lag other developed countries in most measures of R&D and innovation.

As I have argued elsewhere, and some of you are familiar with this, I think we need to take a somewhat broader view of the policy landscape and its interaction with R&D and innovation, and look not just at programs targeted at R&D. For example, I have argued that the overall production tax regime, which affects how the "fruits" of R&D – new products and processes – are taxed, can be important. Some econometric work that I have undertaken bears this out. I will return later to this idea that we have think in broader terms about this issue.

However, it is clearly important to understand as best we can the impact of programs specifically targeted at R&D, and in particular whether they actually lead to an increase in activity, and if so how much of an increase.

As mentioned in the discussion paper, the vast bulk of the studies on this question done to date are econometric in nature. A typical approach would be to use a panel data set of, say, cross country

data and regress R&D expenditures on various control variables as well as some measure of the after-tax user cost of R&D. Other approaches use micro, firm level data, and take a time series approach. As discussed in the paper, most studies of this type tend to find that the long run elasticity of R&D with respect to its after-tax price is about unity – so a 10% decrease in the price of R&D, for example from a tax credit, leads to a 10% increase in R&D expenditures (and note well I say EXPENDITURES on R&D; I will come back to this). The short run elasticity is lower, around 0.2.

This suggests that R&D is quite responsive to R&D in the long run, and somewhat less so in the short run.

Moreover, as was discussed in the paper, while I think that some of these econometric studies are quite good, there are always problems with econometric papers in this (and other) areas, most particularly relating to endogeneity problems and controlling for other influences ("holding all of your ceterises perbus"). Another issue is that they tell you something about the average response of firms to subsidies (after controlling for as many things as they can, often using fixed effects), which can be misleading when considering specific policies applied in a particular country or sector. For example, many of the cross-country studies "mix together", into an imperfect summary measure typically related to the cost of R&D capital, very different regimes (incremental, absolute credits, etc.)

There is also dearth, however, of studies specific to Canada.

This, then, is where controlled field experiments enter the picture – the paper refers to them as the "gold standard", and, under the right conditions, I believe they can be. As I said, to my knowledge there has been no field experiment of the type considered here.

So let me now turn to the field experiments considered in the paper. My remarks will focus on two main themes – the first relates to **timing**, in a couple of different ways that I will expand upon; the second relates to the choice of the population from which the program and control groups are randomly chosen, and what I will call **"population self-selection"**. I will also address other issues at the end.

I view my role as one of being a somewhat cynical devil's advocate. Indeed, Doug Tattri mentioned some of the problems that I will discuss both in the paper and in his remarks.

Timing

The first issue relating to timing involves the **lead time** the program groups need to effectively incorporate the higher tax credit rates or grants into their plans. This is mentioned briefly in the paper, but I think it needs to be thought about carefully. R&D programs typically involve multi-year investments and commitments, and companies need sufficient lead time to incorporate the new policy environment into their plans. In Experiment I, for example, I understand that firms in the program group will be told ahead of time that they have been randomly chosen to be eligible for an enhanced (higher) SR&ED rate for the next 3 years. I think the firms need to be informed of this well in advance of the implementation of the higher credit rate so that they have time to adjust their spending plans accordingly.

A second issue relating to timing is, I think, potentially very important and it is one that I would like to particularly emphasize. All of the proposed experiments involve a **TEMPORARY** increase in assistance – a temporary increase in the SR&ED credit rate, a temporary increase in the grant amount. Focussing my remarks on the tax credit experiment, a problem here is that firms may simply adjust the *timing* of their R&D expenditures away from the future and into the years when the temporary enhanced tax credit is available. This is an issue that is relevant for any temporary, non-permanent tax credit. Indeed, temporary credits are often introduced with precisely this inter-temporal shifting in mind (i.e., to shift investment from the future into the present). To the extent that this shifting occurs, all that may happen is that the timing of R&D will be accelerated (higher now, lower later) with no real increase in R&D. This may lead to an upward bias in the estimate of the sensitivity to R&D expenditures to the tax credit measured by the field experiment. A similar problem exists for the grant experiment. **In my mind this is a potentially serious problem, one that I'm not sure how to address (other than lying, and telling them that the increased credit is permanent – I suspect you aren't allowed to do this)**. We can discuss this, but it may render the experimental design questionable.

At the very least, more attention should be paid to examining the time profile of R&D over the experiment horizon. Another approach might be to have two program groups – one with, say, a three year horizon and one with five, and then examine differences in the time profile between the two groups vis-à-vis the control group.

Population Self Selection

The proposed population from which the program and control groups will be chosen in the "experimental recruitment stage" are the set of CCPCs that have **already taken advantage** of the tax credit or grant program in the previous year. As such, the population has already been determined by **self-selection**. These are firms that have already self-selected into the population and revealed one or both of two things: they are R&D oriented; they are responsive to government programs to encourage R&D.

As discussed in the paper, by choosing the two groups from this self-selected population you can perhaps think of the experiments as measuring the responsiveness of firms to R&D subsidies along the intensive margin, rather than the extensive margin. Perhaps and fair enough.

But I am again concerned about potential biases and in particular about using the results of the study to form conjectures about the reactions of a broader, non-self-selected population.

Is the elasticity measured along the intensive margin more or less than the extensive margin? If we are, as suggested, just measuring the intensive margin and missing the extensive margin, will this introduce a downward bias?

Does the fact that these firms have already revealed their "sensitivity" to government programs suggest that the results will be biased towards finding a big impact?

What about the fact that these are already R&D oriented firms?

Other Issues

A minor point is that I was a bit confused about how, precisely, the grant experiment would work. A similar concern was expressed by Gary Burtless in his remarks. For example, on page 35 it is indicated that "the funding organization could fund 50 per cent of accepted control-group proposals but 60 per cent of accepted program-group proposals." Is this referring to the proportion of companies whose programs are accepted or to the proportion of dollars asked for? Do the grants actually work in this way – i.e., to the programs have target percentages expressed in this way?

Another issue that should be kept in mind when interpreting and designing the experiments is that they tell you something about the response of firms to the specific program or experiment that is implemented. Responses to different programs may be, well, different. And, responses may be different in different institutional and economic settings (which needs to be kept in mind when using results from other countries and programs). If econometric studies suffer from the problem of averaging over these institutional differences (after controlling for as many of them as you can), experimental studies suffer from the problem that they are too specific to those characteristics.

I would like to conclude with a couple of broader points that relate to big picture question motivating the study – how can we increase R&D activity and innovation in Canada.

The outcome that will be measured in this study is the sensitivity of R&D **expenditures** to an increase in different types of subsidies. A potential issue with all of these studies – econometric and field experiment alike – is what I call the **"p times q"** problem.

R&D expenditures = the price of a unit of R&D (p) times the quantity of R&D undertaken (q).

What we are really interested in is q – the **quantity** of R&D undertaken. But we can't measure this. We also can't really measure p, the price of a unit of R&D. What we do instead is measure p times q, or **R&D expenditures**. It is possible, of course, that an increase in the subsidy rate leads to an increase in p rather than an increase in q. This would be the case, for example, if the inputs used to produce R&D are relatively inelastic in supply. Thus, we might measure an increase in R&D **expenditures** in response to a subsidy, but there may be no increase in **real R&D**.

A paper by Austin Goolsbee of the University of Chicago in the late 1990's makes this point. The paper is called "Does Government R&D Policy Mainly Benefit Scientists and Engineers?". It argues that there is a potential problem with government efforts to increase inventive activity. The majority of R&D spending is actually just salary payments for R&D workers. Their labor supply, however, is quite inelastic so when the government funds R&D, a significant fraction of the increased spending goes directly into higher wages. Using US data on wages of scientific personnel, the paper shows that government R&D spending raises wages significantly, particularly for scientists related to defense such as physicists and aeronautical engineers. Because of the higher wages, Goolsbee argues that conventional estimates of the effectiveness of R&D policy may be 30 to 50% too high. The results also imply that by altering the wages of scientists and engineers even for firms not receiving federal support, government funding directly crowds out private inventive activity.

Now this is a problem with ALL studies of the impact of government subsidies on R&D, not just field experiments, but I think it bears keeping in mind. Also, it suggests that focussing on targeted R&D subsidies as a way of encouraging inventive activity may be problematic.

Actually, the experimental approach might actually help with this issue. It might be argued that given the relatively small group of participants in the experiment, the "p times q" problem is small because p won't really change that much (and can reasonably be treated as fixed), and thus you are actually measuring what is close to a real R&D effect. This may be true, but when the program is expanded to a broader population this assumption is more questionable. It may be that an enhanced credit, for example, works in a small group where p is fixed, but not in a bigger population where p might increase.

Another paper by Paul Romer of Stanford University in 2000 relates to Goolsbee's in the sense that it also suggests that the supply of R&D workers (scientists and engineers) is rather inelastic. The main argument of this paper is that the government R&D policies may have been largely ineffective as they have essentially relied on demand-side instruments and incentives, without recognizing the fact that the supply of R&D workers is quite unresponsive to changes in demand. The policy implication of this is that any successful policy must not discard the supply side of the R&D workers market. That is, the Government must provide incentives to the increase of the number of under and postgraduate degrees awarded in the natural science and engineering fields, along with the subsidies given to innovative firms, so that the stimulus received by the demand side of the R&D workers market is accompanied by a similar stimulus of the supply side.

And, to close the circle, this relates indirectly to a point I made at the beginning of my comments regarding some of my own work regarding the impact of the production tax regime on R&D activity. That is, we need to consider the entire policy landscape as it relates to R&D, not just narrow policies and subsidies targeted at lowering the cost of R&D. The entire tax and fiscal system, and our policies in support of post-secondary education, all bear on the issue. Rather than fixating on R&D credits and grants, perhaps that is where we should be looking for the answer to the question of why we "seem to be doing everything right" without the results to show for it.