Innovation and the productivity challenge in the public sector

Talk given at the inaugural *Policy Reflection Forum* at the Department of Communications

Wednesday 5 March 2014

Dr. Mark Matthews[1]
HC Coombs Policy Forum
A Joint Australian Government – Australian National University initiative

We have heard a lot about innovation in the public sector over recent years. How useful has this focus actually been for the efficiency and effectiveness of policy design, service delivery and indeed enhancing public value?   My talk today will consider some of the key dimensions of the innovation agenda, and how it does, and does not, relate to the public sector productivity challenge that we face in Australia's distinctive federal system.

I'll consider: (a) the advantages and the disadvantages of attempts to articulate an innovation agenda within government derived from the significantly different risk environment found in the private sector, especially given the fact that government is the uncertainty and risk manager of last resort and must therefore act in a risk-averse manner in many key areas of responsibility, and (b) the disconnect between the advocacy of innovation in the public sector (as a 'nice to have' objective) and the major 'need to have' challenge of increasing public sector productivity over the long-term as we grapple with the fiscal consequences of an ageing economy (it is not just the population that ages, but also our infrastructure and institutions).

I will then consider what could be done to reduce that disconnect, specifically;

- the productivity gains that stem from fostering a greater tolerance for uncertainty and ambiguity in governments' approach to information requirements (and hence the ICT costs necessary to support government) – especially given that tolerating uncertainty signals the existence of trust and reciprocity, and fosters greater self-reliance amongst those supported by government;

---

[1] Contact: mark.matthews@anu.edu.au

- lessons from the HC Coombs Policy Forum's experimental work carried out in partnership with a state government aimed at adapting the structured hypothesis testing methods used by the US security intelligence community as a faster and cheaper method for monitoring and evaluating government expenditure, analysing and calibrating information from multiple sources, and indeed as a way of designing policies themselves as hypothesis tests (which facilitates learning and waste reduction);

- the more general advantages of balancing the current dominating emphasis on 'evidence-based policymaking' with 'intelligence-based policymaking' delivered by the use of these formal hypothesis-testing methods – an approach far better suited to making timely decisions under conditions of unavoidable uncertainty and ambiguity;

The productivity challenge in the public sector

I will start by stressing the importance of re-connecting innovation and productivity growth. On the basis of statistical analyses of past decades, the Treasury's Intergenerational report assumes future annual average productivity growth of 1.6%. The public sector is roughly one-third of the economy. These are major challenges for the public sector. In order to head off some serious problems, a large chunk of the economy must increase its productivity in a steady way and by a significant amount. Yet, we can't even measure the productivity of the public sector itself effectively at present. What are the consequences if this rate of productivity growth does not eventuate? How well positioned are we to judge these consequences if we don't have a good sense of past rates of public sector productivity growth and the relative importance of different drivers. Perhaps we need an old fashioned growth accounting framework for the public sector that relates outputs growth to the growth of different inputs and estimates the 'residual'?

Firstly, I want to take you briefly through a little intellectual history concerning the study and understanding of science, technology and innovation in the private sector. This is important in order to explain the strengths and the weaknesses of the way in which the innovation agenda has recently been adopted in the public sector.

The interest in innovation in the private sector originated, in part, in a reaction against the difficulties faced by neo-classical economics in explaining technological advance. If one assumes a world of perfect information and a state of

equilibrium in which markets are operating in a stable manner, then technological advances must be treated as externally originating deviations from these equilibrium conditions — processes of disruption to which the economic system must respond and adapt.

The finding from the early growth accounting studies that long-run productivity growth had a large 'residual' element that could not be explained by increases in the standard factors of production (capital and labour, etc.) stimulated a large and productive line of investigation that eventually led to the 'innovation studies' work that is currently informing thinking on public sector innovation. As innovation studies has evolved it has moved away from economic theory and econometrics and toward more managerial approaches — with a particular (and useful) emphasis on documenting and understanding real practices in business.

Inevitably, this emphasis on how businesses *do* innovation in practice leads to a focus on how firms accumulate and exploit proprietary knowledge and capabilities: how they seek to exploit intangible assets that their competitors do not have. The emphasis is on *differences* between firms' capabilities — on how innovation drives markets in such a way that they are in continual evolution — rarely in states of equilibrium. It should be of little surprise that the management of uncertainty and risk feature strongly in this perspective on innovation.

Bring back growth accounting – but relate it to the public sector?

There has been a negative side effect for government from this pattern of evolution though. As work on 'innovation' has flourished and shifted from economics departments to business and management schools is has become a little too disconnected from our understanding of long-run productivity growth.

In the old days a lot of attention was paid to relating R&D expenditure to productivity growth. This was helped by the availability of pretty good data on R&D. We still have pretty good data, however many researchers who study innovation nowadays stress that R&D is essentially an accounting and tax break-based concept that does not reflect actual industrial realities in many sectors. We hear much less about R&D than we used to and more about innovation. One problem is that this shift in emphasis has weakened the link between measured productivity growth and

innovation – the link it asserted, the link makes intuitive sense, but we are actually rather poorly positioned nowadays to work out how future long-run productivity growth might behave – and to determine now it might react to efforts to simulate R&D and innovation investment.

Another problem is the tendency to ignore the 'inconvenient truth' – for the R&D and innovation advocates at least – that large chunks of publicly funded research expenditure have little or nothing to do (directly at least) with the generation of innovation outcomes. The so called 'linear model' that links R&D to commercial innovation (scientists invent – industry applies) may be widely debunked amongst most of the cognoscenti – but persists in government policy frameworks and the media in a zombie like manner – not properly alive but it won't die. Zombies are not good for public policy.

It is no surprise that some econometric analyses of the relationship between R&D and productivity growth in Australia have failed to find a statistically significant causal relationship – much of our R&D effort (e.g. public science) is quite legitimately about other things than productivity growth.

That said, in my view we do need to re-connect our modern understanding of innovation with its genesis in growth accounting. The use of a growth accounting framework in Treasury approaches to managing an ageing population, i.e. the three P's of participation, population and productivity, is a clear signal of the importance of closing this loop. The point is that we must close this loop with a more realistic conception of what R&D and innovation are all about.

By evolving in the manner in which it has, work on innovation studies now has the (largely unrealised) potential to converge with another stream of thinking in economics known as 'Austrian' or subjectivist economics. This stream of thinking is distinguished from neo-classical economics in some fundamental ways — and ways that are highly relevant to understanding public sector innovation.

Rather than a world of quantitative uncertainty, the Austrian economic perspectives describe a human condition in which creativity is a necessary response to qualitative uncertainty (effectively ignorance) over what the future has in store — both good and bad. In some circumstances there are no probabilities to assign to future states of the world, but rather the necessity to act *creatively* in order to generate parameters that can be assigned probabilities (and hence managed 'rationally'). The resulting competition is *inherently* a process of discovery and innovation. From this

standpoint, markets are inherently *exploratory* and *innovative* collective endeavours that operate via selection.

If we think about markets in this more analytical way — as exploratory processes and selection mechanisms — then it is easier to understand their limitations and, hence, grasp why public sector innovation is so important in helping us to manage uncertainty. Markets can cope with risk (quantifiable likelihoods) but they cannot cope with uncertainty as easily.
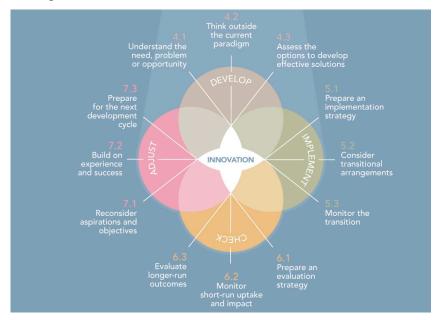
This is why governments spend vast amounts of taxpayers money on translating uncertainty into risk. Many scientific and technological inventions are driven by the fundamental human desire to transform ignorance into uncertainty and risk. There are whole rafts of imaging technologies (X-ray, ultrasound and magnetic resonance imaging, microscopes, particle accelerators, telescopes, seismic analysis, magnetic anomaly analysis, etc.) that provide us with data that we would not otherwise have access to (i.e. that translate ignorance into indications and likelihoods). Much scientific theory is concerned with translating ignorance into risk (i.e. the analysis of complex data sets in order to generate patterns of risk — such as crop planting strategies in the face of unpredictable weather patterns). In short, investments in scientific instrumentation and pattern recognition are, collectively, investments in translating ignorance into risk. We are very rarely certain of what is or may happen, particularly in complex situations such as human health, but we collectively prefer to have more information than less information to guide our decision-making.

Risk and innovation

The private sector has well established and understood mechanisms for managing the risks faced when investing in innovation. An 'appetite for risk' is a key component of corporate strategy and an influence on competitiveness. However, when it comes to fostering innovation in the public sector, it is unwise to attempt to simply transpose the private sector approach to managing the investment risks associated with innovation.

The approach we worked on with the Australian National Audit Office was to attempt to formulate a less risk-averse perspective than currently exists. We sought to do this by developing a decision support framework and by trying to convince the auditors that if a department of agency followed these guidelines but things went

wrong - then they would not be punished. That was the theory anyway. This framework is presented in Exhibit 1.

Source: ANAO (2009) *Innovation in the Public Sector: Enabling better Performance, Driving New Directions.* Australian National Audit Office *Better Practice Guide*. Canberra

Governments are special when it comes to innovation & risk taking

Governments are the uncertainty and risk managers of last resort and, nowadays, operate in a very different environment of transparency and accountability in which risk-taking of any kind can be problematic. Governments must also play the role of an 'immune system' and attempt to identify and pre-empt possible negative unintended consequences both to the nation's economy, society and environment and international relationships and obligations. Consequently, the risk-reward relationship for the public sector varies significantly from that of the private sector. Therefore, it is imperative to recognise these differences in order to enhance the effectiveness of efforts to stimulate public sector innovation.

A key difference between public sector innovation and private sector innovation is that market-based selection mechanisms play a different role in the innovation process. In the private sector, the litmus test for attempts at innovation is market success. Not all innovations prevail in the market, and indeed various other factors mean that the 'best' solutions may not become the dominant solutions. However, markets do enforce selection processes that tend to eliminate less competitive solutions. Competing firms therefore do their best to second-guess what

will prevail in the market, often applying vigorous structured decision-making processes (such as Stage-Gate methods) to weed out less promising concepts and solutions.

However, in the final analysis it is the market, and the social and cultural preferences that are reflected in markets, that will decide which innovations succeed and which do not. The academic and policy literature on how these processes work (and do not work) is well developed and full of useful insights. These insights work backwards from market processes into the R&D and demonstration stages that drive new product introduction. They also work forwards into how market processes drive the incremental innovations that continue once new products have been introduced into the market place.

In a public sector context, the relationship between innovation and markets (as selection mechanisms) is significantly different. As I've stressed before, governments deal with the uncertainties and risks that markets cannot handle. This requires innovations in what governments seek to do. But, crucially, governments cannot rely on market processes to play the critical 'weed-out' stage in the innovation process by eliminating solutions that do not align well with the preferences expressed in markets and encouraging those that do. Rather, governments need to try to mimic this aspect of the functionality of market-based selection processes without the recourse of relying on markets to actually carry out this selection process. This requires that the public sector draw heavily upon external and internal expertise to weigh-up complex risks of generating negative externalities. When there is no market-based 'shortcut' available the sheer weight of evidence that may need to be assessed poses major challenges, and raises important questions about whether 'hierarchies of evidence' are required to deal in a rational way with the sheer quantity and complexity of information available. The consequences of incorrectly judging what will and won't work when seeking to innovate are disproportionately greater for this type of public sector innovation than for private sector innovation.

Furthermore, when private sector innovation goes badly wrong (e.g. a new drug that has unforeseen and terrible side-effects) it is governments that bear the responsibility by virtue of their regulatory roles. This is why, in comparison to the private sector, public sector decision-making processes can appear cumbersome, risk-averse and time-consuming. The unintended consequences of getting it wrong are far too severe to rely on the market to correct problems — as in the private sector. The far greater complexity of what governments do generates great uncertainty over *what*

to do in response to challenges. The extraordinarily damaging potential associated with unintended consequences necessitates robust risk-averse decision-making. Governments are indeed the uncertainty and risk managers of last resort.

Unfortunately, I don't see this dimension in the various surveys of public sector innovation now taking place. Prodded on now by the OECD, these surveys collect lots of data – but data of uncertain purpose. As many junior medical doctors are taught: only order blood tests if you have some hypotheses you want to test or standard protocols to follow that will influence the decisions you have to make. In contrast, those various surveys are fishing nets with a strong process focus, with no obvious purpose beyond asserting the importance of innovation.

The public sector innovation agenda risks irrelevance

If we saw the public sector innovation agenda really pushing the boundaries of how governments act as the uncertainty and risk managers of last resort and contributing to demonstrated productivity growth I'd be unworried about this displacement issue.

The worrying thing is that I don't see this happening. What I see is the concept of public sector innovation looking for problems, and innovation being banded about as an end it itself – with self proclaimed experts on 'innovation' busily touting for business (and taxpayers dollars). This new game is disturbingly disconnected from the reality of the nasty fast and nasty slow moving problems that governments must handle – and that nobody else can handle.

Ideas and concepts have been adopted from the private sector in an attempt to frame an agenda able to address these challenges. Prototypes and experiments, incubators etc .

I think the public sector innovation agenda started to wobble off its tracks precisely because it gained 'critical mass' and started to focus on innovation as a process. A process that needs surveying, a process that needs mapping etc. This process focus has been combined with the tendency to broaden the concept of what constitutes 'innovation' to such a scope that it means almost everything and hence loses its utility. Innovation has become an over-spun and over-hyped slogan.

I have been particularly concerned at a tendency to avoid relating public sector innovation to uncertainty and risk management in the policy learning cycle –

despite my own efforts to frame things in this manner via my contribution to the *Australian National Audit Office*'s Better Practice Guide on this issue – as highlighted earlier.

Developing a better tolerance for uncertainty and ambiguity

The current dominant paradigm in public management is characterised by an intolerance for ambiguity and uncertainty. Although governments are the uncertainty and risk managers of last resort this aspect of governance sits rather uneasily against managerial systems and cultures that view ambiguity and uncertainty as dangers to be avoided. Indeed, ambiguity and uncertainty are themselves viewed as risks. This stance is especially visible in ISO31000, which frames risk as uncertainty over the achievement of clearly specified objectives.

In this paradigm, various important activities, including the manner in which output-outcome budgeting has been implemented favours *precision*: commitments to firm targets that cannot be fudged. The assumption is that a rather mechanistic stance enhances transparency and accountability and contributes to the legitimacy of governance.

This mechanistic approach can restrict the ability to learn-by-doing, share insights amongst peers and generally experiment with better ways of governing. In such a context, the 'experimentalist governance' approach promoted by Charles Sabel at the University of Columbia is refreshing because it prioritises learning, sharing insights, ideas and experiences.

A tolerance for ambiguity is central to experimental governance because it involves open-ended approaches characterised by learning under conditions of substantive uncertainty. Intended outcomes are broad, provisional and can be modified in the light of experience. Governance shifts from a command and control architecture to a distributed learning architecture. Effective public policy is not a matter of 'getting things right' ex ante, but rather establishing the conditions for learning-by-doing in an uncertain world in which objectives are rarely 'right' in the light of hindsight and unfolding, frequently unexpected, events and experiences.

Given these points, I think it would be more useful if the policy narrative shifted to a concern with how best to cope with uncertainty and ambiguity rather than risk per se. Risk is driven by the existence of uncertainty and ambiguity – and these are factors that rarely go away. As I have stressed, governments are the uncertainty and risk managers of last resort, attempting to cope with levels of uncertainty and associated risk that business and markets cannot cope with.

This means that simply adopting risk management methods developed in the private sector (a notable characteristic of the dominant 'new' public management paradigm) is not effective or especially helpful. Rather paradoxically then, whilst the private sector does have useful experience in defining and managing the appetite for risk necessary to innovate (the risk-reward relationship etc) this is not the aspect of private sector practice that has shaped approaches to risk in the public sector.

Would fostering a greater tolerance for ambiguity, uncertainty and risk give us a productivity dividend?

My *Australian National University* colleague Michael Smithson has usefully stressed that exhibiting a tolerance for uncertainty both signals and reinforces the importance of social capital. If I signal that I am willing to tolerate uncertainty about what you are doing then I signal that I trust you. This opens up a window for self-empowerment and, potentially, innovation.
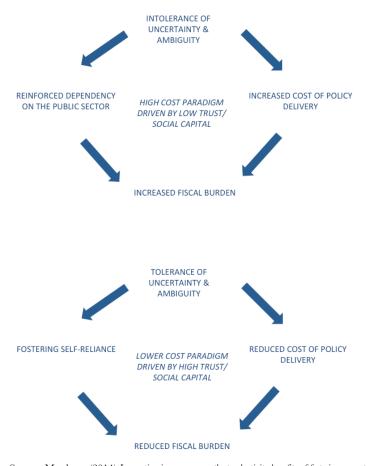
As someone who works in partnership with a range of government officials in helping to try to address challenges myself and my colleagues in the HC Coombs Policy Forum are keen to explore ways of increasing the cost-effectiveness of what governments do. I am therefore especially interested in finding ways for government to eliminate wasted effort and therefore cost.

My own past experiences in working closely with aerospace manufacturing business in the UK and to some extent the USA has left me with a lasting respect for the advantages of focusing attention on wasted work (scrap, excessive inventory, re-work, warranty claims, customer litigation etc). This focus on the outputs that one does <u>not</u> want can be a useful framework for thinking about processes in government. Whilst the difficulties in measuring outputs makes in hard to measure productivity in

the public sector it is of course far easier to measure prevailing levels of wasted effort - and reductions over time in these wasted efforts.

I suspect that if we were ever to transition to a governance paradigm in which a tolerance for uncertainty and ambiguity was much stronger then we could eliminate vast swathes of cost to taxpayers. If we have a low tolerance for uncertainty and ambiguity then we lock ourselves into a cost-escalation spiral driven by attempts to reduce ambiguity - especially if big data and associated ICT costs are involved. In a manner similar to the arms race dynamic we have to spend more and more taxpayers funds to try to re-assure ourselves that we know what our stakeholders are doing. This stance reduces empowerment and fosters a passive, compliance oriented culture.

This contrasts with the alternative of fostering a tolerance for uncertainty and ambiguity with the aim of breaking this cost-escalation cycle and, in so doing, encouraging self-reliance, innovation etc. The two contrasting dynamics are illustrated in the following two diagrams.

INTOLERANCE OF
UNCERTAINTY &
AMBIGUITY

REINFORCED DEPENDENCY
ON THE PUBLIC SECTOR

*HIGH COST PARADIGM
DRIVEN BY LOW TRUST/
SOCIAL CAPITAL*

INCREASED COST OF POLICY
DELIVERY

INCREASED FISCAL BURDEN

TOLERANCE OF
UNCERTAINTY &
AMBIGUITY

FOSTERING SELF-RELIANCE

*LOWER COST PARADIGM
DRIVEN BY HIGH TRUST/
SOCIAL CAPITAL*

REDUCED COST OF POLICY
DELIVERY

REDUCED FISCAL BURDEN

Source: Matthews (2014) *Innovation in governance: the productivity benefits of fostering a greater tolerance for uncertainty and ambiguity.* Paper given at the ANZSOG workshop on '*Twenty-first Century public management: the experimentalist alternative*'. 11-12 Feb. Crawford School of Public Policy. Available at: http://marklmatthews.com/recent-lectures-and-seminars/

It is, of course, essential to set clear transgression boundaries when being more tolerant of ambiguity and uncertainty (i.e. tests for identifying circumstances and behaviours for when I should stop trusting you). Of course, transgression boundaries can also be evolved by the general community and applied to government too – the principle is reciprocal. This pragmatic approach to accountability avoids excessive, and costly, transparency – hence marketing and PR budgets can be reduced.

Re-thinking our reliance on evidence-based policymaking

I now want to turn to another aspect of uncertainty and ambiguity – the extent to which our reliance on the concept of evidence-based policymaking is 'fit for purpose' when you work in the uncertainty and risk manager of last resort.

The concept of evidence-based policymaking is currently in vogue in many OECD governments. That approach relies on the robust analyses of data and other information to assess policy challenges and the performance of government interventions. The extent to which evidence-based analyses are able to inform *future* policy interventions rests upon the combined impact of the rate of change in pertinent conditions together with the degree of uncertainty and risk faced over potential future states of the world. In general terms, the greater the combined impact of the rate of change in pertinent conditions together with the degree of uncertainty and risk the less useful, and indeed even dangerous, a reliance on evidence-based policymaking alone.

I suggest that we would be wise to counter-balance the current emphasis on evidence-based policymaking with the concept of intelligence-based policymaking. Intelligence-based policymaking involves the adaptation of the structured competing hypothesis testing methods used in (US) intelligence tradecraft and 'weak signals' analysis for more general application in public policy.
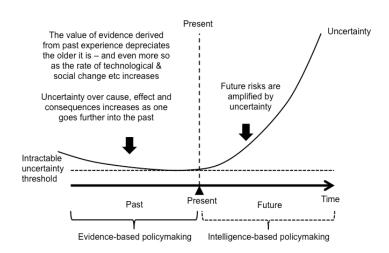
One advantage of intelligence-based policymaking is that it has been explicitly designed to handle decision-making under conditions of substantive uncertainty, ambiguity and risk – situations in which there may be no option to wait before more evidence is available before deciding what to do about a possible threat.

Another advantage of intelligence-based policymaking is that it is better positioned to handle *reflexive responses*. If I release an evidence-based assessment of a particular policy challenge (e.g. in social policy or business regulation) it is likely that

the behavior of the actors and entities whose behaviors constitutes the policy challenge may change in response to their improved understanding of what government plans to do in the future. There are many examples of this.

Given this reality, simply relying on evidence-based policy prescriptions to generate public value can be costly. The very release of the diagnosis of the problem(s) and intended solution(s) can, in effect, reduce the *Net Present Value* of the intervention. This is a familiar issue in intelligence tradecraft (e.g. counter-intelligence practice) whose broader significance for public policy is, arguably, under-valued. This aspect further reinforces the utility of intelligence-based policymaking as a counter-balance to the limitations and risks of evidence-based policymaking.

The following diagram illustrates the principle that uncertainty over understanding both the future *and* past experiences should, perhaps, be treated as a 'U' shaped curve that reaches its lowest point in the immediate past (lags in receiving and assessing information mean that there is always more uncertainty over the present than the immediate past). Although there is far greater uncertainty in regard to what the future may involve we can never be certain about our understanding of the lessons from experience to date. All other things being equal we are more uncertain the further we go into the past. This is a constraint to the net present value of evidence-based policymaking derived from historical data.

Source: Author

We have already started to make some progress in improving the cost-effectiveness of handling uncertainty and ambiguity by developing and piloting a more parsimonious approach to evaluating government interventions based on the use of structured hypothesis testing techniques – expressed as conjectures and refutations in a short tabular format.

This approach emerged from efforts to assist a state government in Australia to develop a cost-effective internal evaluation capacity that also significantly reduced the (often onerous) reporting burden on those being evaluated.

This work has adapted methods widely used by the US national security community. In essence, these methods implement the scientific method (conjectures and refutations) but within the command and control systems characteristic of that arm of government. Structured hypothesis testing, especially when it uses competing hypotheses, is especially useful because it leaves room for uncertainty and risk as core challenges, rather than attempting to treat risk management as a compliance exercise. The approach is, of course, compatible with the 'developmental evaluation' approach that makes it both legitimate and desirable to address evolving objectives in the light of experience (an approach with is I understand from recent discussions at the OECD gaining in importance outside of Australia).

Experience to date in pilot work carried out in partnership with government departments suggests that structured hypothesis-testing methods can significantly increase the speed and accuracy of two key governance activities: the analysis of evidence and monitoring & evaluation processes. As a result, a range of government departments and agencies in Australia are now expressing strong interest in the use and further development of this sort of approach. In situations where sufficient data are readily available, structured hypothesis testing can significantly reduce the cost of evaluations of government spending in comparison to conventional audit and narrative-based evaluation methods. Monitoring & evaluation findings are expressed in a tabular manner that does not require lengthy written explanation or rely on nuancing of phrases in the finalisation of conclusions. This format is illustrated in the
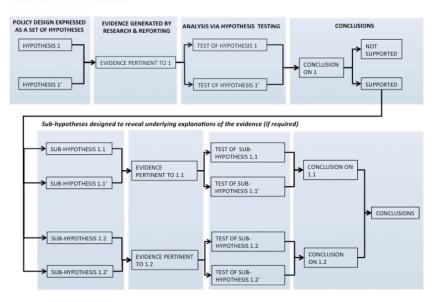
following fictional (and partial) version of a real but confidential evaluation framework.

	Propositions	Evidence	Conclusion (based on balance of probabilities)	Observations
1.	The Institute developed and demonstrated *world-class* research capacity in public policy	• The research impact (as reflected in global citations) is very high by world standards • The proportion of Institute publications with an author from the rest of the world increased to 56% from 23% over the observation period • More than half of Institute publications with an overseas author also involve authors from the USA and/or Europe – linking them to international networks	**Proposition supported**	No further evidence required
2.	The Institute attracted increased funding from Federal and State governments, the business sector and from overseas (demonstrating its increased global standing)	• The Institute attracted levels of Federal and State funding ($23M and $3M respectively) that compares well with benchmark institutes • The share of its income represented by competitive research grants had increased to over 63% by 2012 – a 12% above benchmark institutes • The levels of business funding – and as a proportion of total income – were higher than benchmark institutes	**Proposition supported** with regard to domestic sources of funding but no evidence as yet on overseas funding	No further evidence required with regard to domestic sources of funding. Clarification required on definition of business income for KPI purposes Evidence required on overseas funding
3.	The Institute reached full capacity and met its targets in terms of employment and PhDs	• The target level of employment was achieved 3 years early and it has continued to operate well above that level since then • The Institute produces more PhDs and honours students than benchmark institutes even when controlled for differences in institute incomes	**Proposition supported**	No further evidence required
4.	The PhD outputs from the Institute provided research and technical skills in public policy for Australia with associated long-term national benefits	• 78% of the PhDs in stayed in Australia and 19% went overseas • 63% gained employment in the public sector	**Proposition supported** but no evidence yet on the associated long-term national benefits	No further evidence required on PhD numbers and destinations Evidence required on associated long-term national benefits

Source: Matthews (2014) *Implementing experimental governance: Implications for the evaluation of industrial policy and for maximising learning from industrial practice.* Presentation to the OECD Expert Group on the Evaluation of Industrial Policy. 16 January. Paris. Available at: http://marklmatthews.com/recent-lectures-and-seminars/

One of the major advantages of this succinct method is cost-effectiveness: it can eliminate swathes of 're-work' in finalising often long reports as they journey up and down the government hierarchy on the winding road to completion and sign-off.
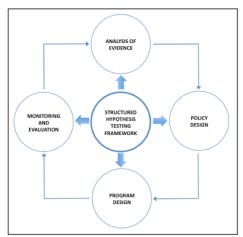
Experience to date also indicates that the ability to 'nest' hypothesis tests in evaluation and review work may have some significant productivity advantages. Diagnostic nesting allows for more parsimonious approaches to be adopted by using the test of a high-level hypothesis to determine whether or not further diagnostic scrutiny of evidence on that issue is required. This nesting is illustrated in the following diagram.

ILLUSTRATION OF HYPOTHESIS NESTING

Source: Matthews & White (2013) *Faster & smarter: using a hypothesis-testing methodology to reduce the time and cost of evaluations.* Discussion Paper prepared for consideration by the OECD Expert Panel on the Evaluation of Industrial Policy. Available on request.

The main challenge encountered to date in using this approach within government is when these diagnostic hypothesis tests need to rely on expert opinion rather than quantitative data (if this is not available). Some officials are uncomfortable if expert opinion is treated as a form of evidence.

We are now moving on to pilot this more parsimonious approach as a means of designing government interventions explicitly as competing hypotheses and, also, as a means of identifying and managing risk in policy and program design. This hypothesis-based approach is illustrated in the following diagram.

Source: Source: Matthews & White (2013) *Faster & smarter: using a hypothesis-testing methodology to reduce the time and cost of evaluations.* Discussion Paper prepared for consideration by the OECD Expert Panel on the Evaluation of Industrial Policy. Available on request.

Interestingly, when I discuss these (fairly radical ideas) with senior government officials in Australia they tend to be well received because the productivity dividend is easily grasped and understood via the reduced re-work loop dimension.

Framing this sort of approach against experimental governance, the issue that intrigues me is: can we find practical ways of fostering a greater tolerance for ambiguity in government that has the twin advantages of helping governments to be better at being the uncertainty and risk manager of last resort and also delivering a significant productivity dividend?

My guiding proposition is that a greater tolerance for ambiguity can reduce the costs of seeking spurious precision in how intervention rationales and objectives are set, programs and designed and their efficacy assessed and communicated.

Let's stop side-lining risk management

A major problem with the current 'risk averse' paradigm in governance is that risk is approached as a distinct stage in project planning, a stage in which success is defined as demonstrating awareness of a set of well-defined risks and then putting in place risk mitigation tactics that reduce these risks to acceptable levels. This amounts to an approach that says in effect "I've dealt with the risks and we can now get on with delivering". There is little scope in this approach for continuously monitoring and reacting to emerging risks. Risk is a problem – not, as it is in the private sector, a source of competitive advantage. Whilst businesses pay considerable attention to framing and managing their appetite for risk in order to innovate, vast swathes of the public sector seek, in effect, to have no appetite for risk at all (and of course in doing so they actually amplify risks).

The final point I want to make is that operationalising effective risk and uncertainty aware learning cycles of the type alluded to here requires suitable monitoring and evaluation tools. This requires risk management to be continuous discipline rather than a 'tick the boxes and forget' approach.

In an uncertain and ambiguous world it is likely that setting up competing hypotheses over risks - hypotheses designed to identify emerging risk factors and to treat risk management not as a compliance exercise but as a creative tension and debate over risk can help us to get much better at coping with, and indeed benefiting from, risk.

The ability to access monitoring & evaluation tools that facilitate 'failing early' in experimental initiatives is important here – although that fail early capacity in turn raises important political considerations as the party-political blame game can be exacerbated.

We therefore see great potential in developing an approach to risk management in public policy that is based upon the structured analysis competing hypotheses as a process of continuous monitoring. Such an approach has the potential to contribute to experimental governance because it provides a more cost-effective method for mutual learning and peer assessment based on establishing and exploiting creative tensions over uncertainty, risk and ambiguity – rather than assuming these things away by virtue of the way in which risk management is approached.

Conclusions

I now want to wrap things up by summarizing my main conclusions, framed as answers to the question: *what do we have to do to keep the public sector innovation agenda on a useful, long-term and sustainable track – a track that yields a productivity dividend?*

Firstly, we need to do much more to re-connect the public sector innovation agenda to the challenge of sustained long-term productivity growth in the public sector. This is especially important in an ageing society. This objective also poses some pretty daunting challenges in regard to measuring changes in the output of the public sector in quality adjusted terms. Some people say this is an old fashioned perspective. I suggest that we can't afford to duck this challenge.

To give you a specific example. The accuracy of the price indices used to measure real output growth, and therefore productivity growth, in the computer industry was significantly improved by the US official statisticians switching the emphasis from the (declining) cost of a computer (as a unit of hardware) to the far greater declines in the cost of a calculation – the 'capital services' provided by that computer. In other words, the key was to switch from a capital good to a capital services perspective. I see the potential for using this lesson to re-invigorate how service sector outputs are measured and used to estimate productivity growth. At present, there are major methodological challenges in measuring outputs in the service sector (of which government is part) because the only readily available output measures are input measures. Lawyers' billable hours for example. We need to focus more strongly on the value of the services delivered to users – including the

uncertainty and risk management dimensions to those services. These quality-adjusred measures are known as 'hedonic' metrics.

Secondly, we need to frame the public sector innovation agenda around governments' distinctive role as the uncertainty and risk manager of last resort. The problem has been that the risk-averse manner in which the New Public Management ethos has been developed (risk is to be avoided rather than grasped) means that this opportunity has been missed.

Thirdly, public sector innovation is a means to an end – not an end it itself. The panoply of e-government seeking to demonstrate 'innovativeness' by being visibly trendy is a concern. Twitter, blogs and the like may indeed be trendy, but they risk becoming a distraction to the core business of governments. Is this really useful – or is it just a fad that will disappear. I suspect the latter. Where is the innovation in developing more effective committee structures, budget processes, legal and regulatory frameworks. These are core functions of government that really could probably do with some innovation. Following the path of least resistance (engagement via the new media) may be new, but is it particularly useful?

Finally, we should recognize that a paradigm of escalating data acquisition, storage and analysis that signals low social capital (low trust and reciprocity) in relationships with stakeholders may be self-defeating. It locks us in to cost-escalation both directly (in ICT budgets) and indirectly in terms of stifling self-reliance amongst stakeholders. I find the idea of counter-balancing the dominating concept of evidence-based policymaking with the (potentially) emerging area of (uncertainty-aware) intelligence based policymaking compelling because it may be a faster, cheaper and more 'fit for purpose' approach to formulating policy and learning-by-doing in delivering policy. Put bluntly: a transition to intelligence-based policymaking may be the step change in public sector productivity that we are searching for – simply because it involves much lower levels of wasted person-hours within the public sector and lower risks of wasted spending on intervention designs and the monitoring & evaluation of this spending that does not align with the reality that governments are the uncertainty and risk managers of last resort.

The emerging concept of intelligence-based policymaking may provide us with a basic roadmap for starting to re-design our business processes in the public sector to reduce costly re-work loops and the consequent 'work-in-progress' inventory of backlogged paperwork being re-worked – and that's where the significant productivity dividend will lie.

The bottom line is that public sector innovation has to be about approaches able to generate cost savings over the long run – not spending more for bells and whistles and optional extras without fundamentally re-thinking the diagnostic, risk management, monitoring/evaluation and decision-making methods that are central to governance.