

**Conditions for Promoting Innovation**  
**Presentation to the Australian Payments Forum**  
**Professor Jonathan West<sup>1</sup>**  
**Monday, 22 November 2010**  
**Intercontinental Hotel, Sydney**

PROF WEST: Thank you, Ian. I'm glad you made the point that I'm not an expert in payment systems, I hope that this will not become even more obvious as I speak.

Let me begin by saying a little bit about what I will address today. What I *won't* talk about, directly, is how to innovate a payments system. What I *will* talk about is the field of study of innovation, the economics of innovation, and what we've learned about how it happens, where it happens, what retards it, and then offer some thoughts that will, I hope, help you reflect on how to promote innovation in this sector.

Innovation is one of those fields where some believe there really isn't much point in studying it all, because everybody is an expert and everybody thinks they really knows how it happens. The problem is that most of what people think they know about innovation is wrong. This is one of the most myth-prone fields of any that I know of in terms of the study of society and how it behaves. It is very prone to fads, it is very prone to instant quick fix solutions, and I think a lot of what we learn from the rigorous, scholarly, academic boring study of innovation, undermines a lot of those fads and quick fix solutions.

Perhaps the most interesting finding of the serious academic study of innovation is that, by and large, it is *not* actually a good thing. Academics, and other innovation gurus, are supposed to be all in favour of innovation, and tell people we need more innovation, we need cultures of innovation, and that we need to promote it, but what the academic study of innovation finds most obviously is that innovation is an enormously wasteful inefficient and often economic destructive process. So it is by no means the case that the promotion of innovation for its own sake is a good thing. Innovation should be undertaken when a clearly identified problem, or a clearly identified opportunity, is present, because I can guarantee you, any serious innovation effort will waste a lot of resources and be tremendously inefficient.

The desire to design efficient and economically productive innovation systems is a false goal, in my view. Nonetheless, there are many circumstances when we need to, and we should, innovate, and indeed the study of economic development shows that innovation is, in the long term, the primary driver of productive advances, which is, in turn, the primary driver of the development of prosperity. So innovation remains tremendously important, and that is true, I think, of the financial system as much any other part of the economy.

Even to me, the two most obvious characteristics of the payments system, as a field in which innovation might take place, are both suggesting that innovation in this field will be difficult, that it will be slow, and hard to achieve. These two features are

---

<sup>1</sup> This is a transcript taken directly from an oral presentation and as such may contain grammatical errors.

obviously that firstly this is a heavily regulated field in which failure is greatly to be discouraged. If your iPod fails, which they very often do, it is frustrating, it is irritating, costs you a bit of money, but it is not catastrophic. If your payments system fails, it is a catastrophe, so a lot of effort is put into regulating it to make sure it doesn't fail. Now putting a lot of effort into something to make sure it doesn't fail, discourages, obviously, experimentation. But experimentation has to be at the heart of any innovation process.

So when you have a system that is designed to discourage experimentation because failure is, or preventing failure is, the highest priority, which naturally retards innovation. The second feature of this industry that seems to me to be not conducive to innovation is, as has been pointed out, that this is a highly networked form of economic activity. It is networked in Australia, in the most obvious sense that it is, itself, a network. People have to cooperate and coordinate, and have to be able to understand each other in doing things collaboratively. Additionally, in Australia we have a financial payment system that is networked in the sense that it doesn't have, or it tries to avoid having, a single hub through which all payments move, and it tries to work through bilateral relationships amongst financial institutions.

So the Australian payments system is networked in the same sense that every financial system is networked, but it is also networked in this second sense, in such sectors where you have to get multiple players, indeed in this sector, hundreds and thousands of different institutional players to coordinate and collaborate, that naturally raises the bar to retard innovation, all things being equal.

So what I thought I would do today is offer you some thoughts about what we've learned from the study of innovation in general that may be relevant to an industry that is characterised by an aversion to failure, and therefore is heavily regulated, on the one hand, and has highly networked operations on the other hand.

So what have we learnt other than innovation is in generally very wasteful and inefficient? The first important finding of the study of innovation is a very simple one, and that is that innovation is pervasive across the economy. There is a myth that there are innovative sectors, and there are mature sectors, and that most innovation in the economy is concentrated in a few leading high technology sectors. This is not true. Innovation is pervasive across the economy at a surprisingly equal rate. Many of you will know that if you look at the Australian economy over the last twenty years, which sector of the economy has increased its productivity fastest? The answer is telecommunications, on an average of 5.5 per cent cumulative average annual productivity advanced. Not unexpected coming from a high technology sector.

What is number two? Number two, at 5.4 per cent cumulative average annual productivity advancement, is agriculture, the oldest industry in the world, even older than the famous second oldest profession in the world - a very mature sector by definition of mature sector, yet roughly neck and neck with the high technology industries in terms of productivity advanced, not in terms of absolute productivity, but in terms of productivity advanced. So there are such things as leading sectors and mature sectors. I don't think, in fact, there are such things as knowledge intensive and non knowledge intensive sectors, it is more, in a modern economy, that all sectors are

characterised by the pervasive importance of knowledge and information, and all are able to innovate.

But the second very important finding within industry sectors, is that innovation is very concentrated. That is to say, while all sectors in the economy in general do innovate, it is very much the case that within the individual industries, innovation is concentrated in a very few players. It is not any kind of normal distribution. The typical pattern is for a few players to break new ground, and for lots of others to follow them. This is a very important finding. It is kind of obvious to anyone in any industry, they know that there is a few innovative players and whole lot of followers, but government policy in general ignores this fact, and it often, in fact, flies in the face of this well known statistical fact.

Most government policy for innovation aims at supporting the laggards or the failures and encourages them not to be kind of about average, and it deliberately removes any support or backing for the leading innovators. In so doing, it retards the development of innovation. What often really matters in innovation isn't the encouragement of a general culture of innovation, where everyone has a go, but backing a few leaders, and removing obstacles for these leaders.

The third finding of the academic study of innovation, is that the typical pattern of innovation, beginning with a technical or scientific discovery, and then proceeding through a sequence to commercialisation, is not how innovation in general happens. Again most of our government policy focuses on this process, a discovery in a lab by guys in white coats, tested with the market, behind the market place, prototype it, see if you can produce it, and take it to market. That isn't typically what happens. About two thirds of innovations are actually suggested to the companies that introduce them by customers.

The overwhelmingly biggest source of innovative ideas for the company is their own customers, not their R and D labs. The R and D labs, and the research, tends to be important later in the process after the problem or opportunity has been identified, and indeed after the option for solving that problem has been identified. It is in the testing, the rigorous testing using proper scientific methods of those ideas, that the scientists and researchers are typically involved. This is a very, very important finding, because it talks about where we get the ideas: What should we work on? Where do we get the ideas for what are the real possibilities? Very rarely do they come from the industry talking to itself.

The fourth big finding of the study of innovation is that the rewards of innovation flow disproportionately to those who bear the financial risk and have the ability to appropriate the financial return. This is really unfair, because it is really often that the guys bearing the risk, and being able to appropriate the return, who are not doing very much work. The people doing the work are often the inventors, the companies, often in academic institutions. When you spend your time talking about innovation, as I do, every conference on innovation is characterised by people saying how badly rewarded they are.

Thinking about how to capture the return, and how to manage the risk in innovation, is absolutely critical to promoting it and designing a system that will work, rather than

thinking about how we produce more science, more research, more start up companies. These things tend to be the focus of discussions about how to promote innovation. Thinking about the central core problem, which is the management of the risk involved in innovation, in fact I would say the defining characteristic of innovation as an economic activity is risk, that it is a more risky activity than routine economic activity for growth. Rewards flow disproportionately to those who manage the risk and can appropriate return.

The final finding that I think is relevant to us today, and I hope when I speak, that the reasons why I have been bringing forward these findings of the field will become apparent - is that many of the capabilities and facets of activities that are required to make innovation happen and reside outside the individual firms. Firms innovate within a system, the key elements of which are often not provided by the firm itself, or in particular not provided by the firm's research and development, from its own formal R and D processes, but they exist outside the firm in clustering, in relationships with universities and researchers, in relationships with other infrastructure that are provided around them. This is, I think, true across the board in most industries and it is even true in the financial industry.

If you look at the biggest innovations in the 20th century in the financial world, certainly in the United States where I have spent most of my career, I would say the biggest three would be:

1. the development of derivatives and options, itself an outcome of an academic theory in the University of Chicago;
2. the development of venture capital, invented by Georges Doriot, a professor at the Harvard Business School, or founded capital fund gain at that University; and
3. high yield, so called junk bonds, invented by Michael Milken at the University of Wharton in his masters thesis.

So even in a field like finance, when you see academic theory, what has that really got to do with the practical world of innovation? The role of those public institutions that surround the industry is critically important, and having good relationships among those is vital to making innovation successful. All of these features of what we know about innovation suggest that opening out your perspective is at the heart of innovation, rather than closing in and talking to oneself.

When we think about the process of innovation, and the system by which innovation happens, it is clear that any such system has to perform four key functions, progressively, sequentially, but also integratively, and there are very many different ways of doing that. How each sector performs these functions defines how its innovative systems will work, and making assumptions about how they should happen is a very common source of failure.

What are those four functions? The first is, and I think the most important, is how does the sector decide which problems or opportunities to work on? How do you decide what issues are we going to take up in our innovation process? There are many techniques for doing this. There are marketing surveys, customer feedback, discussions among the industry, statistical techniques, things like total quality

management, just-in-time, these are all techniques to try and surface problems and decide which ones should be the priority.

I think the most common source of failure, waste, inefficiency, destruction of economic value in innovation processes is working on the wrong problems, working on problems that don't matter, working on problems that matter to the R and D Department, and are part of their career path and CV building, but don't really matter to the people that have to pay for the final product.

Great innovators are great at defining the right issue to work on. You're probably sick to death of hearing about how great an innovator Apple is, but the genius of that company isn't through any new kind of technological break-through, any kind of science capability that they have, or any kind of enormous patent portfolio. It is the ability to see, before any competitors, or even before their customers, what the customer will want over the horizon. No amount of interviews, customers or interviewing each other in the industry would have told you that what the customers wants is an iPod. No survey would come back with this result, because customers aren't able to conceive of such a product before they see it in operation. Why is Apple able to continually come forward with new products that are, in retrospect, obvious, and technological rather straightforward, but, in an innovation sense, break-throughs?

The current product that Apple is about launch is so damn obvious that it is astounding that no one else has really done it. What is it? It is Bluetooth-enabled head phones. The irritating thing about having head phones when you're listening to your iPhone or your iPod is those damn wires that get all entwined. Apple is about to introduce headphones that don't connect through wires, but connect through the blue tooth. There is nothing technologically difficult about that, really. There is a bit of a problem with miniaturising the interface, but why don't we have that feature already? It is actually quite a mystery. It is technically very simple, but what Apple has is a way of defining what will be really valuable for the customer, and I think once you say to customers, "Do you want headphones for your iPhone or your iPod that don't have wires" people say, "Of course." But if you say to them, "What do you want in the way of headphones?" They'll think in terms of the standard set of headphones.

This industry, I think, has a real issue, how do we define what are the right problems to work on? What really matters? What will really add value? Where are the ideas that will change the way things will happen likely to come from? I don't know the answer, but if I have one suggestion for you, and it is that I would spend more time on that topic than anything else. What is the right process for defining the priorities for innovation? And part of the answer has to be to have more than one approach going simultaneously.

The second big challenge that any innovation system will face is how do we define the set of potential solution options that we'll work on? Again the tendency is to too narrowly and too quickly focus down on a couple. We know roughly what the answer is likely to be because we all talk to each other, and that's what we've been trained to do, but the genius of effective innovators is the ability to consider other options early, and include them in the set of ideas to be tested, to be taken forward, and to bring in thought and ideas outside of the normal processes and forums for doing that.

Again there are many different ways of doing this, and of opening things out so that different choices can be considered. Again there are statistical techniques, there is brainstorming, all of these faddish and junky things, but they are all aimed at trying to open up the system so that it brings in ideas from outside.

The third big task of any innovation system is to find some way to test ideas before those options that we've defined in stage 2 are introduced. The gold standard is the formal scientific experiment. There are many other ways to test ideas: prototyping; investigative studies; computer models; customer surveys. What firms normally do is they use the method that they have available to them internally that they know how to do. But so very often those methods aren't appropriate to the type of problem they're working on, or the type of options that they've chosen for inclusion in their set of potential solutions.

The fourth task is to make a decision about what options will be taken forward, and how they will be defused among all the people that have to use them. This also is an area where failure is pervasive, and bad ways of approaching this problem are probably more typical than good ways.

The most common way you make a decision about whether to take an innovation solution forward is to have the boss decide. You work away for years on a project, on an idea and it slowly bubbles up the hierarchy until some total ignoramus who knows nothing about it, had no early input into the project, has had no training, but is the Chief Executive, or worse, the Chairman of the Board, has to make the decision.

There is a beautiful graph that has been prepared on this topic, which shows the ability of senior leadership to actually influence where the innovation path of the firm will go over time, versus where they actually spend their time. Of course the ability of the senior leadership to influence the innovation path of the organisation is greatest at the beginning in defining what problems to work on, in defining what options we're going to look at, in deciding what we are going to think about.

The least ability to influence the outcome is the final decision point, to go or not go, at which point all the work is done, the problem was defined a long time ago, and all you can do is say yes or no. If you say no, you waste all the resources that were committed to it, if you say yes, you run all the risks of it failing. Where do Chief Executives and senior leaders spend their time? At the beginning, helping to define what we're going to be work on, and thinking about how the innovation process is going to be structured, or at the end? Obviously it is at that time – at the end. So there is a diametric opposition between where the most value would be gained by the senior leadership from their input and where it actually comes.

This is a very common source of failure. The organisation gets itself on a path, producing ideas, working on ideas, investing in it, taking it forward, structuring all the reports, writing them in such a way you box in the CEO, so that they basically have to go with it. You withhold all the information that says it might not work and you're worried about. You're trying to sell it and the CEO ends up basically boxed in and says, "Okay, we'll go with it." And in their heart of hearts, everyone in the organisation knows it probably won't work, but that dynamic is set up.

Those are the four big tasks that need to be really seriously thought about in designing an innovation process. Most important, how do we decide what problems to work on? Second, what options are we going to test? Third, what techniques and methods shall we use to test them?

The worst method and technique to test them is the consultant's report. Why is that the worst possible method? Because consultants make a living out of figuring out what the organisation wants to hear and writing a report that says that. That's how you get hired as a consultant. Getting an industry expert consultant, who knows their reputation depends on saying what the organisation wants to hear, who writes a report that says, "Yes, this is a great idea, boss, go with this one." The very best is the absolutely objective scientific experiment, but there are various fields where you can't do that. And, fourth, how do you make the decision and how do you defuse it to the people that have to use it?

Let me finish up with some thoughts about highly regulated, highly networked industries. In such industries, it is very clear that you have to have an innovation system that is robust to failure, systemic failure in particular, and in which lots of people agree what you're going to do. In all of human history, we've discovered only two ways of doing this and they are, I think, largely mutually exclusive. You can either centralise those four steps and processes that have I described, you can have a brains trust who does it on behalf of the organisation, and then hands down the outcomes of their processes, or you can decentralise the process and have multiple pieces of the organisation or the system working semi-independently and combining the results.

Neither of these two approaches is right or wrong. They work in different circumstances, and indeed in the same industry you often find firms following one approach or the other effectively. However, both of them induce a typical, albeit different, set of challenges and problems. The centralised approach is typically faster. It is typically better at big systemic projects, and it leads to, commonly, problems of failure of that central organisation due to information overload or decay. The information, by the time it reaches the central brains trust, has been structured so that everybody is telling them what they want them to hear, and by the time all the information goes up to the top, and it is processed and it flows back down to the people that have to use it, it is old, it is out-of-date, and it lacks the exposure to the diversity of reality that people who have to use the results confront.

The decentralised approach, typically, runs the risk of overlapping effort or lapsed effort. You know, "I thought you were doing that. You thought I was doing it. What? We've both been working on the same thing and the thing was really important neither of us are working on" because of poor coordination. To make the decentralised approach work, it is critically important that a lot of effort is put into defining and standardising the interfaces in the system that the people who are working on problems, within an innovation process, know what are the criteria for the acceptance of the innovation.

So if you think about a personal computer, if I'm designing a hard drive, I can do that independently of designing the whole computer, but I need to know under what conditions is it going to be acceptable, even as an upgrade, to that PC. For example, I

need to know how many pins does it require. What configuration do they need? How many bits of information is the system able to handle when it exchanges information with the central processing unit? If I have a standardised set of criteria, and my innovation meets that criteria, it will be accepted, then I can work independently of the PC designer. In fact, that's how the PC should work.

The designers of the PC are largely systems, standardisers and definers and aesthetics. All the components are innovated, if you like, quite independently of the major system providers and that includes Apple. What Apple does is design the aesthetics, and design the interfaces among those components, and they design that brilliantly, but all of those pieces are outsourced. And that means Apple can draw upon all the innovations going on in the world, they can also send messages to the world they work in about what sort of innovations they are looking for by defining the interfaces and the criteria for making them acceptable.

I was thinking about what industry would be a good parallel to the one you're in, and the best I can come up with, one I happen to know very well, and I might be subject to my own biases here, but that is the US pharmaceutical industry. It's heavily regulated. It's tolerance for failure is very small. If you give people the wrong drugs and it kills them, this is highly undesirable, and regulation is designed to prevent that happening. So it is heavily regulated, every aspect of everything you do in the pharmaceutical industry is subject to some rule or regulation, and there are great screeds of text telling you everything you have to do.

It is highly networked. Innovation happens in many, many different places, and all the pieces have to fit together in the health care system for it to work. You can't just introduce a product and sell it to the consumer, particularly if it is a patented drug. You have to work with a whole other system of providers and decision makers and coordinators. The US pharmaceutical system is overwhelmingly dominant in the world of pharmaceutical innovation. Almost every significant drug comes out of the US, even when it is introduced by a European company such as Novartis, most of whose activity now is in the US.

What makes this industry so strong? Its strength is the brilliance, I think, of a coordinative mechanism, the Food and Drug Administration, and a highly decentralised system for producing innovation. The FDA defines the criteria. If you meet these rules and regulations, your product will be allowed to be introduced. The FDA lays out those rules, defines the interfaces between your drugs and the health care system. You know exactly what standards you have to meet if your innovation is going to be accepted. The FDA then backs the ability to appropriate the returns from it. It protects you if you make an innovation, so that you can make a lot of money. If you are successful in getting through all the hoops and hurdles, they will help you make money from it.

Third, even when undertaking it's own activities, the FDA relies heavily on input from all the industry expertise that is available. Mostly, the FDA employee is a government official, and it's most important activities (defining the rules and regulations, evaluating candidate's drugs, evaluating processing and production facilities), are done by people from the industry, often operating on a volunteer basis. They have the best experts in the country providing their time very cheaply, or often

pro bono to the system to make the rules as good as they possibly can, both in defining very high standards for performance, and making them sensible so that you're actually going to get drugs that work, and this isn't done with armies of government officials.

The government officials exist to coordinate the system of coordination, if you like, but the real work is by people from the industry, and their expertise is on tap, and sourced from academia. It is a system that enables the FDA to signal what sort of innovation it wants. If there are already a lot of drugs to treat an existing condition, it is going to be very hard to get approval. It is going to take a long time and the standard you have to meet is going to be very high. If there are no existing treatments, and it is a deadly, lethal condition, there are a very different standards. If you come up with candidate drug for pancreatic cancer, you will be fast tracked. The standard that you have to meet in terms of side effects will much lower, because basically if you get pancreatic cancer, the only existing treatment is a transplant, and that is very dangerous, similarly with liver cancer.

Those tweaking of the rules, the criteria, the interfaces, enables the system to guide the innovation process to channel it without directing it. So the FDA says liver cancer is an incredibly important problem to solve and we don't have good answers to it, so we'll help you get through the system. We will fast track you. We will actually finance a lot of your R and D, but you'll still have to meet these criteria for bringing your new product to market.

So I think there is a lot to be learned from looking at a system like that that taps the vast amount of expertise and initiative that exists in the research and pharmaceutical industry in the US, that is able to guide the industry to work on certain problems, and prioritise things, and yet to get, in a sense, the benefit of centralised thinking and coordination with the benefits of decentralised innovation activity. So I'll leave it there and open it for questions.