Organisational Economics Workshop in Australia

Markets, Contracts and Organizations Conference, 29 June 2023

Michael Brennan, Chair, Productivity Commission

Thank you for having me here. It's a great pleasure to be a part of the Markets, Contract and Organisation Conference for 2023.

I acknowledge the traditional owners of this land – the Ngunnawal people – and pay respect to elders past and present. It's also particularly pleasing to be part of a conference under the auspices of ANU economics.

It's about 30 years (give or take) since I was here as an undergraduate. I must confess that at the time, having studied intermediate microeconomics, I thought I understood pretty well the basics of the theory of the firm. That is to say, I knew about production functions, cost curves and factor returns. Then an older and wiser economist asked me: "Have you ever thought about why firms exist? Why do we even have them?" I pondered that and drew a blank. He calmly suggested that I go and read Coase. Which I did. And, intellectually at least, life has never been the same.

The version of Coase's *The nature of the firm* that I laid my hands on was actually part of a volume compiled by Oliver Williamson and Sidney Winter, based on a conference in 1987 to commemorate the 50th anniversary of that original paper. It included a paper by Oliver Hart about incomplete contracts and residual rights of control.

You will be familiar with that paper, but you might not recall that it concluded with an assessment of the emerging literature on organisations and the theory of the firm, with the statement that "there is every reason to be excited about the next fifty years."

I trust the conference is living up to that optimistic forecast. No pressure. Unlike almost all of you, I am not an academic. I am a professional policy adviser. So I will come to some policy implications – some areas where the contract/transaction cost/organisational approach to economics can make a contribution to Australia's future policy challenges.

But first, some (non-scholarly) reflections on why it's a good thing for all economists and policy professionals to read *The nature of the firm*, *The problem of social cost* and other elements of the organisational/transaction cost/property rights canon.

The first reason is that it reminds us that, so often, the biggest discoveries in economics are hiding in plain sight. The idea that an organisational hierarchy is an alternative – a substitute for contract-based market exchange – is both breathtakingly simple and incredibly powerful, whether the cause is transaction costs, bounded rationality and incomplete contracts, asset specificity or other things.

As Oliver Williamson put it, Ronald Coase "uncovered a serious lapse in the accepted textbook theory of firm and market organization". The key point being that the lapse occurred not because firms were exotic or unfamiliar or because we couldn't observe firms operating in the real economy, but precisely because they were so ubiquitous.

The textbook theory (like my undergraduate self) had assumed the existence of firms – albeit not in the usual sense of economic assumptions. Economists normally make assumptions consciously in the cause of parsimony – to simplify a model and strip it back to the core elements under investigation. This was different, more an unconscious assumption – taking the reality of the firm as given because it was such an unexceptional part of the economic landscape.

A second virtue of this strand of literature is that it focuses on bargaining and exchange processes as much as it does on final or equilibrium outcomes. That is in contrast to the general equilibrium framework – where all mutually agreeable trades have occurred by assumption – and we are interested in the properties of the stylised equilibrium (e.g. is it a social optimum?) or a comparison between two equilibria in the face of an external shock.

Here we think harder about how those trades actually occur; what the world looks like before two or more (boundedly rational) people see an opportunity to improve their lot, and then, how they go about it – shaped by their incentives, imperfect information and the rules of the game. It's economics in the active voice.

And importantly, such interactions can involve *setting* the rules of the game – assigning initial property rights and determining the institutional structure within which bargaining occurs.

Reading Armen Alchian or Harold Demsetz, for example, one is reminded that some things are logically prior to market exchange – like the precise configuration of property rights, which themselves can be shaped by a 'bargaining process' of sorts.

Economists tend to take market exchange as given, thereby assuming the presence and shape of property rights. But property rights, like firms, emerge in a society to try and solve externality and other problems. They are shaped and they evolve over time, and they are backed up by an institutional apparatus – some collective process, which can include courts or the state – to provide enforcement and predictability.

The question of how those collective structures are designed, come about and change over time is both important and fascinating. That was the spirit of Coase's article on the economics of lighthouses or Elinor Ostrom's work on the management of common resources. Both contributions took the possibility of market failure (as conventionally understood in the microeconomics textbook) as a starting point, not the inevitable conclusion.

The real story is how people and groups work to overcome a social problem that they themselves can often see and intuitively understand. That brings in another point about the nature of progress or advancement in economics as a quasi-science.

The point is this: communities wrestled with public good provision long before Arthur Pigou or Paul Samuelson; they developed rules for governing common resources for centuries before economists defined them as such; the common law tort of negligence had evolved long before Coase wrote *The problem of social cost*; and organisational hierarchies existed before *The nature of the firm*.

Economists didn't invent any of these things; it's even a judgment call whether economists can be said to have 'discovered' them. They existed as a human reality. But economists certainly mapped them by describing them and systematically analysing observed patterns of social behaviour. Often that involved drawing out the latent logic within these observed group behaviours and evolved structures. Just as much of our individual knowledge is tacit, arguably the same can be said for some of our collective knowledge. Part of the role of economics has been to (imperfectly) codify that tacit collective knowledge.

Unsurprisingly, Coase and Ostrom both had a highly empirical approach, but the empirics consisted of a painstakingly detailed understanding of how real-world systems operate. They were case studies. It is one thing to ask why we have firms – that we have lighthouses is perhaps a more mysterious fact. It is certainly a more contingent fact. Lighthouses are almost the canonical public good – non-rival and non-excludable. That lighthouses were constructed at all is a highly context specific and path-dependent fact.

Trying to distil, generalise or codify those lessons (to the extent we can) is a very important contribution. For one thing, it allows us to transfer institutional knowledge to new settings. For example, we can fast-track the evolution of property rights in the face of technological change in areas such as data or carbon abatement.

In a rapidly changing economy, which has been the pattern for the last 250 years of human history, we cannot always rely on slowly evolving home grown institutions, even though we need to be mindful of those when we propose reforms. So, the knowledge we develop and systematise from one context can guide reform in other areas.

That is the hopeful side of the story. There is also a cautious one. Once we better understand the latent logic of existing institutions, we might be more discerning about proposals to reform them. As with GK Chesterton's metaphor of the seemingly useless fence, it is only once we understand its true rationale that we can make a case for its removal.

In that spirit of hope and caution, economists should naturally think about where their insights and expertise can be deployed in ways relevant to public policy. Let me sketch out just one bit of relevant real world policy context for the Markets Contracts and Organisations agenda. A few months back the Productivity Commission completed the five-yearly Productivity Review – a document (or series of documents) entitled "Advancing Prosperity".

We put considerable emphasis on the growing 'non-market' services sector: areas like health, social assistance, education, public administration and community safety. This audience will readily recognise that the term 'non-market economy' is a misnomer. All these services are delivered in some sort of market-like setting, it's just that the 'markets' in question operate with a complicated set of rules and incentives. The disciplines of price competition, firm entry and exit, and reward for innovation (such as through expanded market share) are either lacking or highly attenuated.

The MCO acronym is a useful framework to think about these services, but I will reverse the order. There are many services delivered via organisations – government owned and run entities. Others – in Australia, a significant share – are delivered by private or not-for-profit entities via contract with government. And some are delivered in quasi-markets, with subsidised or voucherised funding from government and a heavy dose of government regulation.

Whether we have the mix right, the 'make or buy' decision, and whether within these categories we have the right design – of organisations, contracts or markets – is a key question for policy, and it could have major bearing on whether we can capture future productivity gains in these non-market services. Arguably Australia and New Zealand have had the greatest success in the reform of organisations, particularly those delivering utility-like services in water, energy, rail and telecommunications.

As Gary Sturgess has pointed out, in both countries (including in the Australian states), the reform of government business enterprises from the 1980s onwards drew on property rights theory. In essence, this reform program reflected clear thinking about the boundaries of the firm – creating some contractual clarity between government and a separate business entity, with the efficiencies of a hierarchy within that entity. And what apportionment of the rights of control at different levels that could maximise the efficiency of the service outcome.

I think that chapter of organisational reform, and its benefits, are well understood. Of course, not every aspect was perfect and further gains could be made in public sector 'firm'-level productivity. But in other cases – and in a growing share of the economy – Australian governments are making significant use of the 'C' and the 'M' – direct contracting with non-government providers or service delivery via a quasi-market.

In some ways, Australia and New Zealand have made quite significant – even radical – use of these structures to deliver services. The stakes here are high, in terms of value for money – the fiscal costs can be large, and so too the impact on people's lives for good or ill. Arguably there has been less cross-pollination between theoretical insights and policy practice in this area.

There could be some gains from trade between policy makers and researchers. There have been some fairly successful models of contracting in Australia (sometimes with in-house providers) in areas like public transport and corrections.

When done well, contracting has helped to clarify the obligations of parties, and protect front-line service deliverers from the barrage of potentially conflicting instructions that can come with public sector hierarchy. Periodic review and re-contracting can create a contestability effect.

In human services – where a great deal is delivered on behalf of government by non-government providers – it is less clear that the lessons of optimal contracting are fully developed and deployed. It is not clear that we are systematically striking the balance between ex ante clarity (via specification of clear terms in a contract) and subsequent discretion (via an ongoing relationship to deal with issues as they arise).

In many cases, we fund a considerable volume of human services via what are essentially grant programs. In 2019, the Productivity Commission was asked to complete a study on the spending by various governments on child and family services in the Northern Territory.

Much of that spending took the form of grants to non-government organisations, largely to fund services to Aboriginal and Torres Strait Islander communities. It found that two tiers of government were spending around $538 million per annum via nine agencies and 700 different grants to 500 entities. Nearly one quarter of those grants were for amounts less than $100,000. The median duration of a grant was two to four years, with 43% of grants operating for less than two years.

Many organisations were cobbling together multiple grants from different programs to fund their day-to-day operations. For example, one service provider received around $20 million from 46 grant programs across different levels of government and different departments, and a local council was receiving $10 million from 24 different programs.

A constant complaint is the level of prescriptive reporting requirements for these multiple, and often very small and short-term grants. Multiple departments in the NT and Commonwealth Governments were funding similar services – we identified three different departments (two Federal and one NT) delivering family violence programs. Feedback from community groups on the ground was that co-ordination was lacking. One pointed to the 'proliferation of new service providers in the region, with the support of government funding, often funded in roles that duplicate existing programs'. This is just a stark example of a fairly familiar pattern.

In contrast to my earlier reference to Elinor Ostrom's work, it's hard to believe that the patterns we observe here reflect some sense of socially evolved optimality. In fact, if the point of good contracting and institutional design in the public sector is to achieve an invisible hand of sorts – channelling private incentives toward the public good – then the situation we confronted seemed to be doing the opposite: taking a considerable stock of altruistic motivation and channelling it into poor public outcomes.

Some of these issues can be addressed through better contracting – longer terms, less prescriptive reporting, more use of a single funding arrangement and a better balance between up front specification and ex post adaptability.

The harder thing is to address the coordination problem. There are clear complementarities in human service delivery – in housing, employment assistance, family support, and drug and alcohol services. We often talk about wrap-around services or joined-up government. The point is that multiple small interventions work best when co-ordinated.

In the economic literature on firms, the presence of complementarities usually militates in favour of clustering those activities, e.g. within an organisational hierarchy like a firm. But in this case, it won't necessarily be a good argument for direct government delivery, since governments are notorious for their siloed and uncoordinated approach to service delivery.

The benefits of co-ordination and complementarity might have to be sought via other structures. The use of social impact bonds is one way this has been tried. Leaving aside the 'bond' element (the idea of external financing), the principle of paying for a set of outcomes can give the service deliverer an incentive to identify and deliver the right mix of interventions.

But in practice it seems as though governments that have gone down this path have found the contractual complexity quite a challenge. Beyond direct contracting, there are other instances where public services are privately delivered via a quasi-market, often with some form of voucherised public subsidy.

Medicare funding for GPs and out of hospital specialist services fits into that category, as does the contestable provision of vocational education and training. The NDIS can also be seen in this way. Aged care and centre-based childcare likewise. In these cases, market design is crucial – and hard.

Small weaknesses in the design of incentives can have outsized impacts on service quality, fiscal cost or the incentive to innovate. In some cases, we operate a hybrid between a contractual model and a market model, such as in vocational training, where in most states, public subsidies can follow student choice but are still governed via a funding agreement between training organisations and government.

There can then be a lack of clarity as to whether government is acting primarily as a funder and contract manager, or as a regulator and market steward. In judging any of these arrangements, it is important to put their imperfections in context. No perfect system exists – the choice and design of organisations, contracts and markets for delivering public services involves trade-offs. As usual in economics, we are constrained by the feasible set and we have to choose which incentive properties matter most at the margin.

Harold Demsetz warned against the Nirvana complex – comparing real world systems to an imagined ideal of perfection. Many of our quasi-market approaches have had their problems, but are still better than the system they replaced. However, the growing size of the non-market sector tells us that the salience of these issues is increasing. Getting the structure of these systems 'right' (or moving them in the right direction) will have big implications for our fiscal position and overall quality of life.

In closing, there seems a big opportunity from better linking academic research in organisational economics to the practical policy challenges we face, much as was evidenced by the approach to government business reform in the past.

And there might be benefits that flow the other way too. As with my earlier reference to the earthy, concrete empiricism that I attributed to Coase and Ostrom, there are interesting real world patterns out there to be studied.

Noting too that, in the spirit of lighthouses and common resources, we often find real world examples of individuals and groups overcoming the incentive problems caused by rigid funding models or siloed service delivery or other factors.

In the Productivity Commission's case study work on innovation in the primary health system, we found amazing examples of practitioners and officials swimming against the tide of powerful funding and governance incentives to deliver consumer centred, integrated and data-informed preventative care.

Distilling those lessons, giving them some more systematic theoretical rigour to improve actual outcomes … as well as (as I put it earlier) identifying the latent logic of observed real world solutions would be a great joint contribution by academics and policy makers.

Perhaps a worthy ambition for the next 50 years.

I wish you well for the remainder of the conference.