Economics of Working from Home

Committee for Economic Development of Australia (CEDA) panel discussion on 12 October 2021

Michael Brennan, Chair

Thanks for having me here to discuss our recent research on *Working from Home*.

Working from home is a highly topical issue, because so many have been forced to try it, and so much has been written – from differing perspectives – about whether it will or should continue.

It is also fascinating from an economic perspective for a number of reasons.

First, although our work emphasised the role of the 'forced experiment' brought about by COVID, it is important to remember that there is an underlying relative price effect that has made remote working more feasible and attractive.

As I never tire of reminding people, the central workplace is a relatively recent phenomenon in human history, and it was aided by the rapid falls in the real cost of moving people from place to place.

For much of the 20th century, transport became faster, cheaper and better. Efforts at futurism reflected that too – I have talked elsewhere about the US Postmaster General in the 1950s who predicted the advent of rocket mail.

George Jetson – that futuristic figure who first appeared in the 1960s – drove around in a space car in a space city. But he still drove to work.

It is a reminder that in our long-range predictions we so often get it wrong because we extrapolate along a single dimension – an observation made by Australian tech investor Rodney Brooks.

For the last three decades, moving people has not become much faster or much easier, particularly with congestion in major cities. Cars are cheaper, but the time cost of the commute is going up. Because of its service element, underlying public transport costs also tend to rise faster than overall inflation.

But the cost of computing, information technology and communications has plummeted.

In much of the debate over the pros and cons of remote work, we often forget that there is a fundamental technological shift – some underlying tectonic forces at work.

Second, it is interesting that despite this relative cost shift, the level of working from home hardly budged for the last 20 years – HILDA suggests 8 per cent of employees did it some of the time; we estimate that translated into about 2 per cent of hours worked; on Census day around 5 per cent of workers did not commute to work.

Those levels have stayed fairly constant since 2000.

And then the pandemic came along and changed everything, not least of which was the perception of working from home.

It seems to contradict the traditional micro-economic assumption of rational profit maximising firms and utility maximising households, all operating on the basis of known, fixed preferences and complete knowledge.

There shouldn't be $100 bills on the sidewalk, and yet prior to 2020 there was clearly some untapped potential for more employees to work effectively from home and for employers to allow it. Why didn't they?

Economists have made a few observations about this:

- Before the pandemic, the pay-offs for a firm to experiment with work from home were asymmetric: if it went well, others might copy. If it was a disaster, then that would be, well, disastrous. So it was safer not to.

- There is something akin to multiple equilibria when it comes to perceptions of working from home. Before COVID, work from home had a stigma – was it working from home, or shirking from home? In that world, industrious, productive employees might shy away from seeking remote work because of what it might signal about their work ethic. So the stigma can become self-reinforcing (a stable equilibrium). COVID removed the stigma by requiring large numbers to work from home irrespective of their relative work ethic.

- The assumption of profit maximisation with complete information has always been a very imperfect representation of how business decisions are made, particularly in the context of new technologies. This is an observation made by Armen Alchian in his 1950 classic Uncertainty, Evolution and Economic Theory.

The point about experimentation is that it leads to learning. The forced experiment of COVID was no exception. It appears that firms and workers learned two things:

1. They learned something about where they sat in the productivity distribution of remote work across firm/workers pairings

2. They learned that the mean of that distribution was a bit higher / better / more productive than they would have thought in 2019.

For many, that was a double whammy, creating a strong upside surprise.

Another reflection is that the forced experiment of the pandemic was just the beginning. As restrictions ease, we will observe a second wave of experimentation.

This time, instead of everyone doing the same thing, firms and workers will be trying out almost endless different and individualised approaches and learning from that experience. Fully remote, fully centralised and all the possible versions of the hybrid model.

The learning from this second experiment comes in two forms – the conscious learning that comes from firms identifying what works and doesn't work for them, as well as the unconscious learning that happens at the level of the economy as a whole, as successful models are rewarded by survival and expansion while unsuccessful models fall by the wayside.

Despite the traditional economic affection for representative agent models (like the single representative firm or household – more on that later) the truth is that the economy relies on the sheer variety of models put forward in this second wave of experimentation. This variety facilitates the learning process – precisely because of the uncertainties and lack of knowledge faced by workers and firms.

The Productivity Commission is an evidence-based organisation, but to say much of interest about work from home one has to supplement evidence with judgement. The evidence just isn't there yet, at least not in a definitive form.

We estimate that 35 per cent of jobs can be done from home. Few of those will be done exclusively from home, as they are in occupations and sectors that benefit from in person contact to stimulate collaboration and creativity.

So we sketched out a couple of scenarios: if all those who could work from home did so 2 days a week on average (1 day for part timers), we would see the number of hours worked from home rise to around 13 per cent.

Perhaps a more plausible assumption would be that only half of those who can work from home do so any of the time – again 2 days per week and 1 day for part timers. That would imply more like 7 per cent of hours worked from home.

It is still a big change – if for example the number of hours worked went from 2 per cent to 7 per cent and the number of people working from home some of the time rose to about 1 in 5.

It would represent a stark increase in a short space of time – much more rapid than other big changes like the rise of female participation or the rise of service sector employment, both of which have been dramatic but relatively gradual.

Will it hurt productivity?

Most of the public and media debate about the productivity of working from home is focused on the experience of individuals – often on the observation that some people are not particularly productive working at home.

By contrast, we were trying to answer a different question: namely, what happens at the economy wide level if an extra (say) 5, or even 10, percentage points of hours worked is worked from home?

In taking a sanguine view on that, in effect we are judging that the extra hours of remote work are coming disproportionately from the right-hand side of the distribution (those who do it better than average), or at worst a random draw. In other words, the extra hours worked from home are likely to be in those firms and by those workers and on those days where it can be done without diminution of productivity (and possibly with increased productivity).

And that can be true even if some individuals are less productive. The average effect is likely to be neutral or positive for a number of reasons: first, bosses have to agree to their employees working from home and they are unlikely to agree to a productivity-sapping arrangement. Also, the first day at home is likely to be more productive than the third, fourth or fifth (and an incremental 5 percentage points of remote work must come disproportionately from first days rather than fourth or fifth days). Finally, employees who want most to work from home have a very strong incentive to find employers and jobs where this can be done without loss of productivity and wages.

We are similarly positive about cities. Again, there is the tendency for some to extrapolate on a single dimension – more remote work means less work in the CBD, office vacancies will rise, cafes and restaurants will suffer, people will move further away from city centres and these trends will all become self-reinforcing.

In theory all those things are possible if the shock is big enough. Through history and across the world, cities have hollowed out or at least taken a long time to recover from a big shock.

Our judgement is that the shock here is not big enough. And for smaller shocks, the second-round effects tend to mitigate, rather than amplify, the initial impact. So, for example, if office vacancies did increase, we would expect rents to fall. That makes it more attractive for new businesses to locate centrally.

If people commute less often, and trade this off against a longer trip when they do, then it's possible that more jobs concentrate in the CBD rather than less.

These things are all speculative, but the broad point is that there is significant adaptive capacity in our cities – they are complex systems and they adjust on multiple margins.

One policy implication is that the case for more flexible regulation of land use gets stronger.

Planning is naturally controversial because of strong community interest in issues of built form. But we have over-planned and over-regulated land use. In a world of significant uncertainty, we still have planning systems that purport to determine – at very high levels of resolution – what forms of economic activity can occur where.

Most jurisdictions in Australia have multiple different commercial and business zones for different types of activity.

But do this thought experiment. Imagine someone sitting in their home office, ordering some groceries online between meetings and then ordering some new office equipment for the home set-up.

All those land uses – office, residential, supermarkets and large format retail – have separate zonings in the physical world of land use planning. But surely in light of the dramatic changes we are seeing through remote work and online retail, we have to be open to new, novel combinations and co-locations of these things.

The boundary between retail and industrial zonings (e.g. fulfillment centres) is already becoming blurred – except in planning schemes, where retail and industrial zonings are quite distinct.

The workplace relations system is another policy framework that has been tested by remote work – not just because the regulation of hours of work is based on presence in a physical workplace. During the pandemic, the system showed itself to be quite adaptive in giving businesses and workers the flexibility to undertake work outside the 'standard workday' or require workers to perform different duties.

Many of the occupations and sectors most likely to work from home have less contact with the formal workplace relations (WR) system – they tend not to be award reliant, are paid well above statutory minima and often work flexible hours. Many of them would be on formal enterprise agreements, however.

The challenge for the WR system is that the preferences for, and feasibility of, working from home is quite individual, whereas the system is largely based on collective entitlements – the *National Employment Standards* (NES) enshrine mandatory community-wide minima; awards operate across industries and

collective agreements across a business or worksite. It is very hard for any of these instruments to enshrine a substantive right to work from home because a collective, one-size-fits all entitlement is unlikely to be appropriate.

So firm policies and individual arrangements are most likely to be the mechanisms used to give effect to work from home. The NES and some awards do enshrine certain procedural entitlements – like the right to request remote work for certain classes of employees – but these stop short of substantive entitlements.

A contextual issue is that – in general – workers appear to desire the ability to work from home more than bosses want to extend it, largely because the main benefit is the commute avoided, and this is hard for the employer to monetise.

So employer representatives tend to voice some caution about work from home – both its productivity impact and the implications for areas like work health and safety. But many commentators also worry about work from home from an employee perspective – particularly the concern as to whether it might lead to more contracting and insecure work.

So if we imagine the WR system as being there, in part, to protect workers against harms, we first have to work out which harms matter more – not enough working from home or too much of it. The system can do both, but realistically only through procedural rules, like the right to request and the obligation on employers to consult about major workplace changes, rather than providing a substantive entitlement to work from home or to work from a central location.

Following on from the publication of our report a few weeks ago, we have also published on our website what we describe as a ‘simple model’ of working from home. We used this model as scaffolding – not so much to tell us ‘the answer’ but to help us think through how the bargain between employer and employee might work out and what variables would shape it.

I guess it’s simple in two ways: one, that it makes a lot of simplifying assumptions – a single representative firm, a single representative individual, both maximising profit and utility with well-defined production and utility functions. Both are price takers. The usual things.

Second, it doesn’t rely on overly complex mathematical methods. It’s a standard ‘constrained optimisation’.

But a quick look at the paper will show that it’s not that simple. There are complicated looking equations everywhere. The functional form and the first order conditions get quite long and complicated because we have to add new terms to try and capture what is going on when we add working from home to a more standard model.

To start from the beginning, a standard model of the household's decision to work involves them trading off consumption (the fruits of paid work) and leisure – two goods which increase utility. The assumption is that no one would want to work except for the pay – the consumption it makes possible. So households choose a level of labour (to pay for consumption) and leisure.

To capture working from home, we add two types of labour – work in the office and work at home – into the utility function. Both forms of labour provide some benefit to the household: working from the office has the benefit of sociability; working from home provides some flexibility (e.g. to combine work with other domestic priorities).

This has two effects. One is that it means labour is now entering the utility function explicitly. It is no longer just the opposite of leisure. This implicitly brings in the notion that workers are getting some positive utility from work, over and above the consumption it makes possible. That is contrary to the standard economic assumption, but as an approximation of the real world it's not completely crazy. Many people do derive considerable intrinsic benefit from various forms of labour.

But it does mean that when it comes time to put (hypothetical) parameters into the model, you have to make sure the utility from these two types of labour isn't too large, or your household becomes a complete workaholic. It has to be the case that consumption and leisure still dominate the overall labour supply choice.

Nonetheless, what the model shows is that the ability to work from home does generally increase the household's labour supply. Much of that is because of the avoidance of the commute, which provides more time for workers to allocate between leisure and work.

Then there is another subtlety, which is the second effect of bringing two types of labour into the utility function. It is that the two types of labour in the utility function have a degree of complementarity. It is as though workers prefer variety rather than working all the time in one setting. This effect is small, compared with the role of the commute (partly because of the illustrative parameter values we have assigned in our simulation) but it does mean that when work from home becomes an option, there is a small labour supply effect purely because workers get higher marginal utility from an extra hour of work now that they are splitting their work across two locations rather than one. So in our simulation, you notice there is a (small) positive labour supply effect even for a worker with a hypothetical zero commute.

Again, this seems like a quirk of the model but not entirely crazy. Isn't there something in the attraction to the hybrid model that comes from having a change of scene – that would translate into working longer on my first day at home than I would on my fifth day in the office? And it is consistent with survey evidence that shows that most workers want a hybrid of working from home and the office.

As I mentioned before, the model is based on representative agents, maximising known objective functions and taking wages as given. And the model is static. We observe the change in behaviour once work from home becomes possible by comparing outcomes in one model specification to those in another. It is a comparative static exercise.

These seem like significant limitations, given we are in a world that is changing and adapting to a new way of working in the face of great uncertainty. But it can still yield insights into the forces that might shape the adjustment.

For example, one obvious point borne out in the model is that it involves two types of labour (at home and at the central location) but usually there will only be one wage. Assuming the employer cannot pay differential wages for different locations of work, they are constrained in their choice as to how much working from home they will want to offer – there is a unique point where the marginal product of labour is the same for each type of labour.

It is highly unlikely that this point will match the relative preference of the household. This creates a natural tension in the relationship. You can imagine three ways in which it might be resolved:

- The employer might find other ways of improving the relative attractiveness of working in the central location to shape worker preferences, such as through ways to foster greater social interaction in the office.

- Labour will reallocate over time, so that workers with a high preference for working from home seek out employers and workplaces where working from home is more productive.

- Even if workers stay put, there is incentive for them to demonstrate higher productivity when working at home, to shift their employer's preferences into greater alignment with their own.

All of these things reflect the possible ways in which a market economy can adapt to a change like this. It reminds us that there are huge incentives to share the gains that this technology makes possible and to get better at using it over time. This is part of the reason we are cautiously positive about the overall impact of working from home across the economy as a whole.