



RESERVE BANK OF AUSTRALIA

Speech

On Lags

Luci Ellis Assistant Governor (Economic)

Sir Leslie Melville Memorial Lecture, Australian National University

Canberra – 17 August 2018



I'd like to thank the Australian National University for providing the opportunity to speak to you tonight. ANU is one of my three *alma maters*, and I'm pleased to be able to engage once again with the university community. It's important that central banks and the academic community collaborate. We can learn a lot from each other.

I'm conscious that this talk was originally scheduled to take place almost exactly a year ago, but I had to cancel due to illness. That seems especially pertinent on this occasion, because my topic tonight is lags. By choosing this topic, I hope to encourage those of you who engage in economic analysis of any kind, and the broader academic community, to think carefully about time, and in particular, about how the passage of time is built into your analysis.

I want to talk about lags: where they come from; how they might change; and what they might be telling us.

This is more than just an academic point about what our models should look like, though. Understanding lags is central to day-to-day analysis and interpretation of economic developments. As such, they are central to our work in economic policy.

Before I get to that, I'd like to make some brief remarks about the current situation. We released the Statement on Monetary Policy last week, and earlier today the Governor and some others of us appeared before the House of Representatives Standing Committee on Economics. I don't have anything new to add but thought it would be worth summarising the main messages from the Statement and the testimony.

Unemployment is still a bit higher than our estimates of the rate consistent with full employment. Inflation is still a bit low. It is therefore appropriate that monetary policy should be supporting the economy right now.

But the Australian economy is growing a bit above trend, which has helped the progress on getting unemployment down and seeing inflation return closer to the midpoint of the Bank's target. The world economy is continuing to expand at a good rate and that is also helping. We expect further progress on both the inflation and unemployment fronts over the period ahead. If things turn out broadly as we expect, it is therefore more likely that the next move in the cash rate will be up than down.

That progress on getting unemployment down and inflation up is likely to be only gradual, however. So any increase is still likely to be some time away. That's been the story for a while, and these remain the key messages about the policy outlook, as the Governor explained earlier today.

Of course, there are a number of risks that need to be kept in mind. Trade tensions have increased in recent months, which could worsen the investment climate as well as directly limiting global trade and growth. Conditions in a few emerging markets have worsened for country-specific reasons. The Chinese authorities are balancing the need to support growth with the need to reduce financial risk. It is hard to predict how that balance plays out in their macro policy settings in the short term, and how this will affect growth in China. The outcome will be directly relevant to the prices paid for Australia's bulk commodity exports, as will be global energy demand more generally. It can't be ruled out that certain commodity prices, and thus Australia's terms of trade, remain higher for longer than we expect. That would be a positive income boost if it happened.

Another risk to the upside is that the US economy is currently seeing significant fiscal stimulus at a time when there is already not much spare capacity. So it can't be ruled out that global growth and inflation turn out to be a bit stronger than we currently expect. In those circumstances, policy rates in the United States might rise a bit faster than some people expect.

Domestically, we are alert to the risks that high household debt could pose for household spending, especially if other negative events should occur. We are also mindful that, as the labour market tightens, wages growth could be slow to pick up in earnest. This has been the pattern in some other countries where unemployment rates are already very low. Among the questions we have been grappling with lately are: how low the unemployment rate needs to go before wages growth picks up more emphatically; how quickly wages growth will respond to tighter conditions; and how much it will pick up. Fundamentally these are questions about distinguishing between structural change and lags. This brings me to my topic tonight.

What Lags Are

A lag occurs in any instance where time passes between when an activity is initiated and when that activity has its impact. Almost all economic phenomena involve lags. Not all of these lags are significant, of course. The time taken to make a hamburger doesn't matter that much for economic outcomes. The time taken to build a house is rather more important. But even lags that are short individually can impose significant costs if they occur frequently. For example, tender times – the time needed to pay and, where relevant, receive change – are a non-trivial share of the cost of making payments (Stewart *et al*/2014). Waiting time is a cost, and can affect people's choices.

In thinking about lags, I have found it useful to categorise them into three types, which I will discuss in turn:

- Process lags
- Stock–flow lags, and
- Learning lags.

Process lags

A *process lag* is the time taken for any readily defined production process. Think of laying a brick or baking a cake. It involves a more or less predictable delay, and its outcome is unambiguous. The brick gets laid or it doesn't. (The cake might or might not turn out as you expect, but you have a cake.)

One important but underappreciated example of a process lag is the time required to make a decision. By this I mainly mean the formal processes and governance that organisations place around decision-making. Some of this might involve information gathering or other analysis, to know what decision to make. But even with all the facts in front of you, there are lags involved in any formal decision process. Documents must be prepared, meetings convened, review processes completed. Good governance takes resources; it also takes time.

Decision lags underlie an aspect of economic policymaking that is often taught to undergraduates. They account for an important difference between monetary and fiscal policies as tools for macroeconomic demand management. Monetary policy changes, we were all taught, can be enacted more quickly. The decision lag was shorter. But that change might take longer to affect the economy. The 'long and variable lags' of monetary policy are those after the decision is taken and enacted. Fiscal policy changes can have more immediate effect, but the decision lag is longer. Parliament must vote on the government's budget. Other decision processes might be needed before funds are finally disbursed or spent. These can also take some time.

Many other decision lags can be relevant to our analysis of the economy. Consider the approval processes for building projects, or for loans. Most companies have formal decision processes for allocating budget to particular investment projects, based on a written business case. Tender processes for awarding contracts and application processes for hiring workers involve lags arising from the search process, which is a process lag in itself. But in the end, there's a decision, which also involves a lag.

It should be clear by now that time is an input into almost every economic activity. Certainly anything that involves a labour input involves a time input. Often we oversimplify or ignore the resulting process lags in our models, for example by thinking of everything in discrete time with quite long discrete periods. But those process lags are there.

The existence of process lags has a number of implications. These go beyond the standard lag assumptions in our models' production functions. One is that for every production process with a material process lag, there is an 'in-process' state. Some of these provide important, additional

information on the state of the economy. The stock of construction projects approved or commenced, but not completed, tells us a lot about near-term future construction activity. The number of job vacancies – the in-progress state between deciding to hire and actually hiring – tells us a lot about the state of the labour market.

A second important implication of process lags is that hog cycles can arise. It takes time to fatten a hog, or build a building, or ship something from far away. In the interim, prices can fluctuate significantly. Too many economic models forget about these dynamics. But the short-term fluctuations can matter a lot, on our way to the long-term.

Stock-flow lags

The second kind of lag – stock-flow lags – also produces hog cycles, for slightly different reasons. These lags are the inherent implication of past activity having long-lived consequences. They can be seen in the property market, in the shipping fleet – indeed in most classes of capital goods. Because these products are long-lived, the increment to the stock in any one year is small. And because of that, the response of the stock is usually quite sluggish. Price cycles can therefore arise. I have spoken about these kinds of hog cycles before in the context of property markets (Ellis 2015).

More generally, stock-flow lags by definition have long-lived consequences. Past decisions affect outcomes far into the future. Because of this, they create an option value of waiting. Economic theory has long allowed for this in models of investment, for example. Vintage capital, s - S pricing and many other mechanisms all arise from this type of lag. I don't have much to add to that already rich literature tonight. But I would reiterate a point I've made before (Ellis 2014): if you aren't taking stocks and flows seriously, you probably don't have a realistic model of the economy.

Learning lags

The third type of lag that I identify, I've loosely termed '*learning lags*'. This is the time taken to effect an intellectual change in a person or group of people. Think of the time taken to draw a conclusion, develop a skill, or detect a structural change. This type of lag is similar to the process lags I have already spoken about. But it has a number of important differences. I think these differences are central to many of the economic uncertainties that central banks must grapple with.

Firstly, the time taken is inherently uncertain. People's learning styles and development rates differ. And some people are just more decisive than others.

Secondly, and perhaps more importantly, the input to the process can be noisy and uncertain. Think of the example of detecting and concluding that a structural change has occurred in the economy. That's a signal extraction problem. How long it takes before you can conclude that a change has occurred depends on many things: the noisiness of the data, the signal extraction technologies you have available – even the strength of your priors versus your willingness to change your view. That can differ depending on individual psychology or organisational culture. It also means that the results of a learning process might not be replicable across people, time or circumstances.

Thirdly, and related to the last point: the output of the process can be uncertain. You don't know what conclusion you will draw until you've drawn it. You don't know how much skill you will build up until you have done the training.

Building on all that, we can also see that *micro learning lags lead to macro learning lags*. Micro learning lags are when individual actors in the economy take uncertain amounts of time to learn about changes in their circumstances, before they can respond to those changes. Because of that, noise and uncertainty arise in the lags we see in macroeconomic developments. So it takes time to learn about changes in the economy – the macro learning lags. These macro learning lags could be as simple as the time it takes to be confident that growth has picked up, for example. Was it a change in trend or just one quarter's noise? This isn't always easy to know. Detecting structural change in the economy is even more difficult. It depends on the variability of those learning-and-response processes at the micro level.

One might think that the presence of variable lags would make observers more reluctant to conclude the economic structures have changed. In fact, it seems the opposite is true. Too often we see confident pronouncements that things are now different. In fact, often those observers just needed to be a little more patient, before the normal patterns reasserted themselves.

A typical example has occurred after most of the recessions in industrialised countries over the last several decades. Compared with earlier recessions in the post-war period, those since the early 1990s have tended to involve more financial sector distress. The uncertainty this engenders – or perhaps simply the balance sheet distress – changes the nature of the recovery.

In particular, it tends to make the subsequent pick-up in employment quite drawn out. We therefore start to see claims of “jobless recoveries” or “skills mismatch”. It is often assumed that there has been a structural change in the economy to which workers are not able to adapt. But in reality the labour market does start to pick up, eventually. It just takes longer than some of these observers seem to have expected. The economic structure and skills requirements aren't changing in any unusual way. It was just that the lags had lengthened. Or maybe, they hadn't lengthened. People just forgot that the lags are there.

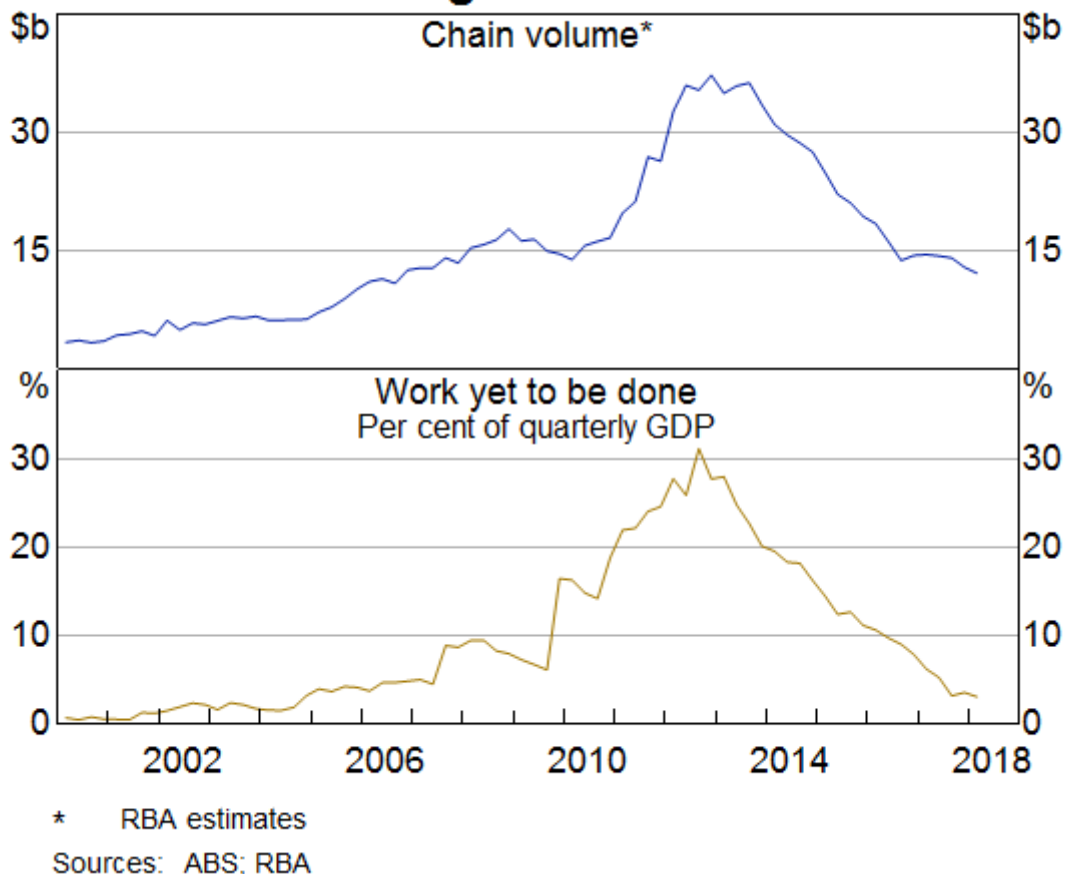
Three Important Lags in Australia

I'd like to illustrate these different types of lags with some examples that have been important to economic developments in Australia in recent years.

Perhaps the best example of a *process lag* is the multiyear boom in mining investment. Largely spurred by the increased demand for commodities in China, earlier this century mining firms decided to expand production capacity for coal, iron ore and LNG here in Australia. It would have taken some time for these firms to make those decisions. More importantly, it took several years for each of these projects to be built.

The effect of the resulting process lags is clear in this graph (Graph 1). The lags involved in decision-making and project preparation meant that this investment took a while to ramp up. And the process lag to completion of these projects has taken many years; for LNG, there is still a bit more to go.

Graph 1 Mining Investment

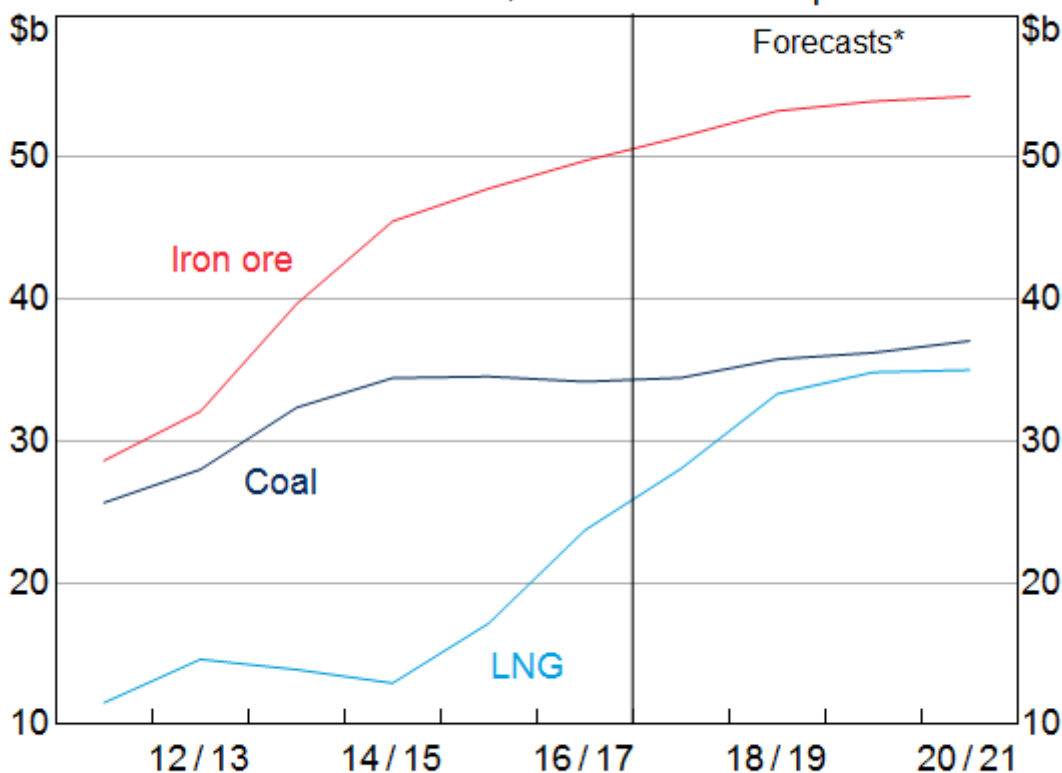


As a result, the ramp-up of resource exports flowing from this decision was also quite lagged (Graph 2). For coal and iron ore, the process is more or less complete. But for LNG, again, some of the new production capacity is still to come on line.

Graph 2

Resource Export Volumes

Selected commodities, constant 2015/16 prices



* Based on DOIIIS projections

Sources: ABS; DOIIIS; RBA

More recently, there have been announcements of new mining projects. Most of these are 'sustaining investment', that is, the investment you need to do to keep production at its current level (Jenner *et al*/2018). However, there are a couple of coal projects underway and one planned LNG expansion that involve increases in productive capacity. These projects have in part been a response to stronger demand – and thus higher prices – for these key commodities than some observers might have expected a couple of years ago. But again, it takes a while before this addition to capacity comes on line.

It is typical to use property as an example of a *stock–flow lag*. The flow of new construction is small relative to the stock already existing. That flow supply is determined in part by the price, and thus the profit to be made from the construction project. But it is the balance of supply of and demand for the *stock* that determines prices. Even without the process lags involved in construction, it would therefore take a while for any supply response to bring prices down.

I'd like to emphasise a different example tonight, though, because it highlights how simple lags can sometimes be misinterpreted as something else. It also shows how the existence of stock–flow lags means that cyclical developments can have long-term consequences. That example is long-term unemployment.


Data on unemployment durations first became regularly available in the early 1990s recession. You can see how, at the beginning of this graph, the rate of people unemployed for more than three

months but less than a year ('medium-term unemployment'), and the long-term unemployment rate – people unemployed for more than a year – were still rising even after short-term unemployment peaked. Both these rates start rising around the time that unemployment overall starts rising. This is because the stock of unemployed starts to rise partly because the flow back into employment slows down. This affects people who are already unemployed. But unemployment also rises because some previously employed people lose their jobs. This is especially relevant in recessions when many workers are laid off. So of course there is a process lag: it takes more than a year before someone who previously had a job becomes long-term unemployed, just by definition. This is why long-term unemployment continues to rise even after the layoffs decline to more normal rates. You can also see this effect during the smaller increases in overall unemployment in the late 1990s and in 2001.

A recovering economy, solid employment growth, but rising long-term unemployment: it would have been tempting to misinterpret this as a structural problem. Indeed, many did so at the time. But that probably is not what was going on. What this shows is the effect on a stock – unemployment – when the economic decision – to hire or lay off workers – is a flow in and out of that stock.

Graph 3

Unemployment Rates*



* Seasonally adjusted by RBA

Sources: ABS; RBA

Once unemployment starts to decline, the stock–flow lag becomes more relevant. Consider that, if the flow out of unemployment back into employment declines for a period, a stock of unemployed people can build up. Even if that flow returns to more normal levels in the subsequent recovery, it can take a long time for that stock to return to low levels. That is exactly what we saw after the

1990s recession. It took 15 years from the peak, before long-term unemployment finally troughed, at a bit above ½ per cent of the labour force.

Again, this should not be interpreted as a structural deterioration in the performance of the labour market. It just took a long time to work down that stock of unemployment.

More recently, the end of the mining investment boom in 2013 induced a period of slower growth and softer labour market conditions. As a result, both medium-term and long-term unemployment increased for a time. They are currently a little below their peaks in 2015, but only a little.

As an aside, you might have noticed the sharp move down in the latest data for medium-term unemployment. One should never make too much of one month's figures. It's hard to know whether it will persist, or whether it's a result of the vagaries of our seasonal adjustment process. ^[1] And it's a good example of the signal extraction problem that gives rise to learning lags in economic policymaking.

An important lesson here is that, most of the time in an expansion, growth at so-called 'potential' is not enough. There is a stock of unemployment that needs to be worked down. That means that an above-average flow out of unemployment back into employment is needed. And that is one reason why macro policy might need to be supportive of growth for an extended period during an expansion.

A natural question to ask is why one wouldn't seek to run the economy even faster, to open that spigot back into paid employment even wider and get long-term unemployment down sooner. It's a legitimate question. In essence it depends on how quickly unemployed people can be matched with vacant jobs. It is possible that there are limits to the efficiency of the labour market's ability to do that matching. Job search involves process lags, such as the window of time that an employer chooses to receive applications, as well as the decision lag involved in selection. So beyond some point, running the economy faster might not open the spigot that much wider.

There is also the question of how much faster the economy could be induced to run at all. For the economies most affected by the financial crisis (by which I *don't* mean Australia), that might also have had its limits. Banks' credit supply had been weakened and existing spare capacity reduced firms' need to invest. At the same time, during this period, it was not always the case that the various arms of macroeconomic policy in these countries were all pulling in the same direction.

In the Australian context, as the Bank has explained on several occasions, there have instead been concerns regarding a different stock–flow dynamic. One key way expansionary monetary policy works is by getting people to borrow more and spend. Lower interest rates also reduce the servicing burden on existing debt, which can free up some cash to spend more. Overall, though, we would ordinarily expect debt to rise when monetary policy is boosting growth. When household debt is high, it poses some risks. In particular, it could make household spending and well-being less resilient to shocks, including individual-specific outcomes.

So the question then arises: what is the appropriate balance between the risk of higher debt and the benefits of the hoped-for faster decline in unemployment? In part, this depends on whether a

particular level of interest rates induces households to keep increasing debt indefinitely, or whether debt instead accumulates up to some point, or some relativity to income, and no further. The experience of recent decades, and the arithmetic of debt-servicing burdens, suggests the latter. [\[2\]](#) But it is hard to be certain.

Another question to ask when weighing up this balance is whether the question of how low unemployment can go depends on how fast you are getting there. It's usually assumed that there is a level of unemployment, below which wages growth starts to pick up. This is sometimes called the 'full employment' rate of unemployment. A more technically correct, but less attractive, name is the 'non-accelerating inflation rate of unemployment', or NAIRU. [\[3\]](#) In some other countries, unemployment has fallen well below previous estimates of the NAIRU before wages growth actually started to pick up. As a result, more recent NAIRU estimates have tended to be revised down.

That brings me to two examples of a *learning lag* that I consider to be particularly relevant to the current situation. The first is the simple one that it might just take a while before firms start to realise how hard it is to find suitable labour, and decide to offer higher wages. So it could be that it takes a little while for wages to pick up once unemployment has fallen below the NAIRU. If so, it would mean that NAIRU hasn't really moved in these countries, and the revisions we see are just an artefact of the way this unobservable variable is measured.

If that is what is going on, then this mechanism would also be relevant in Australia. As the Bank has explained on several occasions recently, although there appears to still be a deal of spare capacity in the labour market, many firms are reporting that they are finding it hard to recruit suitably skilled workers. Some of these issues are industry-specific. For example, we hear from contacts in our liaison program that it is particularly difficult to recruit certain kinds of specialists in construction and IT, as well as local workers into hospitality roles in regional areas, where foreign workers on short-term visas had previously been more easily recruited. But so far, wages growth is not picking up much even in these areas, despite scattered reports of high wage offers in some cases. It could be that, after a period of realisation and learning, we will see wages growth pick up more emphatically in these areas.

But there is another kind of learning lag that might also be driving the outcomes we see. It might be that employers are learning that their definition of 'suitable labour' can change. After a long period where there is plenty of unemployed labour, employers get used to being able to pick and choose. As the labour market begins to tighten, they have to start looking at people who are less credentialed or who have less prior experience. And they have to start looking at the way they organise the work in their firm, and maybe investing in productivity-enhancing innovation. And if they do that, suddenly previously 'unsuitable' workers start looking a lot more 'suitable'. Productivity isn't something an individual has handed to them at birth. Different people have different skills and experience which affects their productivity in a given context. But the context of how work is organised and supported with capital is crucially important.

If this is true, then a period where the labour market tightens gradually, and not in every area at once, could induce firms to realise that workers are more 'suitable' than they previously assumed. And if they realise that before they start resorting to paying higher wages, the result is that

unemployment can fall further before wages start to rise. In other words, the NAIRU falls. And if that is what is happening, it means that learning lags can induce path-dependence in the macroeconomy, in a particularly consequential way.

How Lags Change and Why That Can Matter

Most lags can change in duration. Some are more stable than others. Policymakers need to be alert to which lags can shift, and be ready to detect those shifts. It should be obvious that technological change, by increasing productivity, often reduces process lags. Changes in rates of obsolescence or physical depreciation can change stock-flow lags. And as technology makes more information available more frequently, we might expect that learning lags could speed up. This particular implication is not guaranteed, though. Sometimes more information just adds noise. I shall talk more about that in a moment.

Some lags, especially decision lags, depend on legal and social processes. So they can change, if those responsible for those processes desire it. A good example here is the streamlining of building approval processes in some jurisdictions over recent years. Decision processes can also slow down, if more and more rigorous governance is being demanded, whether by external parties or the decision-makers themselves.

Whatever the reason for it, a change in the length of lags can have huge economic implications. For the hog cycle examples, price dynamics can change noticeably. And where a shorter lag comes from increased productivity and therefore lower cost, we can see marked changes in relative prices.

There can also be big social implications from a change in lags. Probably the best example is the enormous decline in the cost of recording and transmitting information over the centuries. Ideas, good and bad, can now reach the whole world in seconds. In past centuries, each new copy of a text was done by hand and could take months. The result is more information to more people, not all of it better information.

What it Means for Monetary Policy

I've already referenced the long and variable lags of monetary policy. I would like to draw out three more general implications of lags for monetary policy making.

First, most forecasts leverage known process lags. For example, we use building approvals to help forecast residential construction. Stock-flow lags are also important, but in my view these are often underappreciated in many forecasting models.

Second, the lags from monetary policy to macroeconomic effects are primarily learning lags. That is why they are variable, as well as long. People take time to realise that demand has increased enough that it's worth hiring a new employee. Likewise people take time to decide that demand has increased enough to be worth changing their prices. The time taken will depend on many factors. For example, the lags to changing prices are determined by things such as the pricing models prevailing in that industry, menu costs, and the level of competition, which determines the cost of getting a pricing decision 'wrong'.

Third, both the learning lags involved in noticing stronger demand and the decision lags involved in responding to stronger demand depend in part on the stories we tell ourselves – our narratives and beliefs. That's because these help determine how readily we update our hypotheses, and change our minds.

This raises the question of whether better macroeconomic outcomes can be achieved by telling ourselves better stories – more accurate stories, and stories that are more consistent with each other. I believe this supports the case for fulsome and frequent communication by the central bank. Many researchers have focused on the role of central banks' public commitments to a nominal anchor such as an inflation target, to help coordinate expectations around that target. Others have focused on central bank communication as a signal of future central bank actions, which other economic actors can respond to.

I can't help thinking that the case goes beyond that. When a central bank explains publicly how it interprets economic developments, it creates a public good. Putting out a view gives everyone something to compare with their own expectations. So by sharing the results of its efforts publicly, perhaps central banks can help everyone learn a little faster, just as we learn from comparing our views to those of other observers.

To conclude, lags pervade economic activities. They matter. As researchers, and as policymakers, we need to understand when lags are important, and develop techniques to detect when those lags are changing. The passage of time needs to be taken seriously in economic inquiry. So I'll take no more of yours tonight, and thank you for listening.

Bibliography

Ellis L (2014), '[Future Directions in Financial Stability Analysis: Learning from Others, Learning from the Past](#)', Address to the Paul Woolley Centre for the Study of Capital Market Dysfunctionality Conference, Sydney, 9 October.

Ellis L (2015), '[Property Markets and Financial Stability: What We Know So Far](#)', University of New South Wales (UNSW) Real Estate Symposium 2015, Sydney, 8 September.

Jenner K, A Walker, C Close and T Saunders (2018), '[Mining Investment Beyond the Boom](#)', RBA *Reserve Bank Bulletin*, March.

RBA (2003), '[Household Debt: What the Data Show](#)', *Reserve Bank Bulletin*, March, pp 1–11.

Stewart C, I Chan, C Ossolinski, D Halperin and P Ryan (2014), '[The Evolution of Payment Costs in Australia](#)', RBA Research Discussion Paper No 14.

Endnotes

[1] The ABS only publishes seasonally adjusted data for some categories of unemployment by duration. The figures shown here have been seasonally adjusted by RBA staff.

[2] This is the implication of RBA (2003).

[\[3\]](#) Strictly speaking, it's the non-accelerating wages growth rate of unemployment, but 'NAWGRU' has not caught on.

© Reserve Bank of Australia, 2001–2018. All rights reserved.