Understanding how the Australian labour market works, with examples from the COVID-19 era - A talk to the 2025 Australian Labour Market Research Workshop

JEFF BORLAND *Department of Economics, University of Melbourne*

Abstract

This article is about how to understand labour markets. I suggest an overall framework for doing analysis of labour markets and give a detailed treatment of several elements: asking questions, how to think about the causes of changes in labour market outcomes and using descriptive information to evaluate causality. To illustrate these approaches, I use examples taken from my research on Australia's labour market in the COVID-19 era.

Acknowledgement: This article is based on the keynote presentation at the 2025 Australian Labour Market Research Workshop. It draws on two research papers in progress, Borland (2025a, 2025b). I am grateful to Phil Lewis for the invitation to present the talk, to participants at the Workshop for helpful feedback, and to Matt Cowgill, Bjorn Jarvis and Peter Lake for very helpful comments on a draft version of the article. All opinions and any errors are my own.

Introduction

In this talk I am going to change the pace. So far in the workshop we have heard lots of interesting and informative presentations on what is happening in Australia's labour market. I want instead to move back a step and make my focus the methods we use to get that knowledge about what is happening. Specifically, my central theme will be how we take the tools of economics and apply them to answer questions about the operation of labour markets.

My talk is in two main parts. First, I will introduce an overall framework for how to analyse labour markets. Second, I will address in more detail several elements from that framework. To illustrate my ideas, I am going to use content from my recent research on the Australian labour market during COVID-19.

Analysis of labour markets encompasses a range of purposes, types of questions and techniques. My concern in this talk is going to be with a category of analysis that I call 'policy analysis'. This category is purely my invention. But the features I attribute to the category, shown in Table 1, do I feel, represent a distinct stream of analysis.

What I have in mind is analysis to answer a question that is policy relevant at that time and is about a specific labour market. Policy relevant can mean questions that are directly about policy, or questions about what is happening in the labour market that are important background for policy making. My talk is going to mainly involve the latter type of question. Examples might be wanting to know what has happened to Australia's labour force participation rate in the past decade, understanding why nominal wage growth was so low in Australia in the 2010s, or evaluating whether there is currently a shortage of teachers in individual states and territories in Australia. It is my impression that it is this type of background question which is most prevalent in policy analysis of labour markets.

Policy analysis also brings a distinct approach to how it goes about answering these questions. That approach is primarily empirical, using descriptive statistics and perhaps simple regression analysis, integrated with a synthesis of findings from existing studies. Think of, as examples, an RBA *Bulletin* article, a Productivity Commission research report, or a Treasury *Round-Up* article. Government naturally is a major source of policy analysis of labour markets in Australia, along with researchers at think-tanks, academics and market economists.

Obviously, there are alternative approaches to labour market analysis. In Table 1, I compare the features of policy analysis against two other categories that might be identified, which I label as 'market analysis' and 'academic research'. My emphasis in distinguishing between the categories is on the questions asked and how analysis of labour markets is done, rather than who does the analysis. (For example, it is common these days for economists working in government to be doing what I classify as 'academic research'.)

Table 1. Types of labour market analysis

	Market analysis	Policy analysis	'Academic' research
Objective	Comment on narrow topic.	Answer question of policy relevance at specific time for specific country.	Original contribution of broad relevance to knowledge on labour markets.
Types of questions	Commentary on today's labour force figures. Description of trend in labour productivity.	What has happened to labour force participation rate in Australia in the past decade? Why is wage growth in Australia relatively low?	How does technology affect labour markets? How do childcare policies affect labour supply?
How analysis is undertaken	Presentation plus basic analysis of descriptive statistics.	Synthesis of existing knowledge plus new empirical research based mainly on descriptive statistics.	Original empirical (and theoretical) research.
Data used	Headline statistics	Detailed data underlying headline statistics plus some use of unit-record data.	Unit-record individual and firm-level data. Value placed on use of new data sources.
Output and audience	Short research note -> Clients, other market economists, other members of own organisation, media.	Report (of varying length) -> Other economists in own or other government departments and agencies, Ministers and staff. Sometimes external audience of professional economists through external publications. Media, business, researchers in other social sciences.	Working paper/Journal article -> Primarily other academic economists.
Causality	Usually not concerned with causal impacts.	Seek to argue for causal impact using existing knowledge and through careful interpretation of descriptive statistics and regression analysis.	Use advanced empirical techniques for identifying causal impact.
Time period	Needs to be done quickly to fit the purpose – usually just a couple of hours.	Could range from a week to several months, depending on scale of project.	Usually needs to be conducted over an extended period, with the original phase of research taking several months or longer.

To conclude this introduction, and summarise where I have got up to so far, my talk is going to:

- be concerned with questions about the labour market that arise in doing policy relevant analysis.
- suggest a structured approach to answering those questions.
- be about the practice of answering the questions – that is, how best to create bridges from technique to application.

Given this focus, I suspect the main audience for my talk is early career public policy and labour market analysts, as well as students intending on that career. At the same time, hopefully aspects of what I say will also be of wider value, for anyone undertaking labour market analysis.

The overall framework for labour market analysis

What is the 'structured approach' for labour market analysis that I am proposing? Figure 1 provides an overview. Essentially, there are three main stages in the approach, which I classify as: (1) Foundation; (2) Investigation; and (3) Coming to an answer.

Stage 1 – The foundation

The starting point is to know the question you want answered and the labour market you are interested in. Having a well-formed question is essential. Without that there won't be any purpose or structure to your analysis. You may be given the exact question you will need to answer, or it may come from initially being asked to consider issues related to a broad labour market related topic, in which case getting to the question will be part of your task of analysis, and is likely to take some time. Defining the labour market you are analysing is also a must. Sometimes your question will imply directly the labour market you are interested in. But if not, this needs to be thought about explicitly. Is your interest in labour market outcomes for the whole of Australia, for an individual industry such as mining or manufacturing, or for a specific group of workers in a specific region, such as hairdressers in Victoria? Whatever the answer, how you define the labour market will set the scope of your analysis, suggest the types of influences on labour market outcomes you will need to consider, and likely also the data you should use in any empirical analysis.

Figure 1. A structured approach to answering questions about the labour market

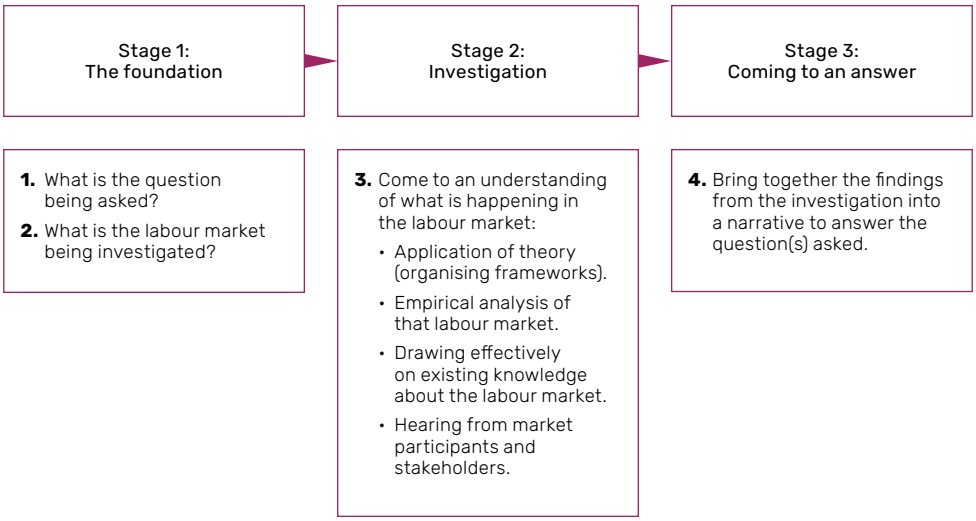

Stage 2 – Investigation

The next step is to undertake an investigation directed at answering your question. This stage involves creating and bringing to bear on the question of a range of different types of information. This information can be derived from the application of economic theory, your own empirical analysis, the findings from existing studies that relate to the question, and via stakeholders and market participants.

For the information you collect through the investigation to have the greatest chance of leading you to a correct answer to your question, you will need a solid knowledge of (i) general ways of thinking about how the labour market operates, what I call 'frameworks for organising thinking about the labour market' and (ii) how to do empirical analysis.

Theories of how labour markets operate are many and diverse. But when it comes to answering policy analysis type questions, I would argue that those theories can be distilled to four core 'ways of thinking' about labour markets:

- Representing labour market outcomes as the **equilibrium** from interactions of the forces of demand, supply and institutions.
- Recognising the role of **drivers**. Having a taxonomy of the main channels through which demand, supply and institutions determine labour market outcomes.
- How changes in labour market outcomes imply **adjustment**. Representing changes in labour market outcomes as occurring through a process of adjustment, where that adjustment can potentially occur in multiple ways.

- The **human element**. How behaviour and outcomes in a labour market depend specifically on the fact that it is the labour of humans that is being traded.

By 'how to do empirical analysis', at a high level I mean a knowledge of empirical techniques relevant to answering policy analysis type questions about labour markets – including skills in measuring, describing, and assessing causality. At the coalface of doing empirical analysis, this translates into a plethora of specific skills. These skills range from what might be seen as the very basic, such as being able to correctly interpret the meaning of variables and generate descriptive statistics, through to the advanced, including the ability to manage large data sets and apply sophisticated econometric techniques.

Stage 3 – Coming to an answer

Ultimately, you will need to use the information you collect in your investigation to come to a judgement on the answer to your question. To do that will require you to collate and interpret your information, and to use that as the basis for your answer. The final step is to present your answer and supporting evidence to whoever has asked you to do the analysis, in a format they have asked for or that you choose.

For the rest of the talk, I will draw from this suggested structure to describe three key elements of doing labour market analysis. These are:

a. Begin with a question and keep asking questions
b. Know the drivers and follow the consequences, and
c. How to judge causality

a. Begin with a question and keep asking questions

My claim is that there are four types of questions that get asked about labour markets. The first two types, with which I am mainly concerned in this talk, are what I call 'describe' and 'explain' questions. *Describe* questions ask for a description of an outcome or outcomes occurring in a labour market (or markets). *Explain* questions ask for the reason(s) that an outcome has occurred. In addition, there are two other types of question that might be asked. *Policy* questions are about whether a situation has arisen where government intervention in the operation of a labour market can improve society's wellbeing, and what that intervention should be. *Prediction* questions ask what will happen in a labour market in the future. Table 2 gives examples of each type of question.

Table 2. Analysis of labour markets – The types of questions that get asked

Type of question	Definition	Example
a. Describe	Questions asking about what outcomes are occurring in a labour market (or markets):	
	What outcome(s) are occurring at a point in time?	What is the proportion of the labour force in Australia is currently unemployed?
	Whether a difference exists in some outcome between groups of workers in the labour market or between workers in different labour markets at the point in time?	In the 2021 Australian Census, what was the average hourly wage of workers with a bachelor's degree compared to workers who did not complete high school?
	Has an outcome changed over time?	What has happened over the past year in Australia to the proportion of the workforce who move between jobs?
	Is there a difference in changes over time?	What has happened over the past year to the labour force participation rate in Australia compared to the United States?
b. Explain	Why has an outcome happened?	Could, for example, be applied to the findings from any of the 'Describe' questions – such as: What explains why the average hourly wage of workers with a bachelor's degree is higher than for workers with no post-school qualification? What explains why the proportion of the workforce who moved between jobs has increased over the past year?
c. Specify implications for policy	Is a policy response needed to address an outcome that has occurred? If so, what should that policy response be?	Could, for example, be applied to the findings from any of the 'Describe' questions – such as: Is the current rate of unemployment too high or too low? If so, what policy response is needed to alter the rate?
d. Predict	What outcome(s) will occur in future time periods?	Will the rate of unemployment in Australia increase? In which Australian regional labour markets will there be the largest impact from decarbonisation?

Any analysis must, at some stage, arrive at a question to answer. Because it is only when you get to a question that your analysis can have direction. That is the sense in which you should always 'begin with a question'. Having in mind the four types of questions can assist you in doing this. Needing to put your question into one of the categories will make you think exactly about what it is you want to know and allow you to be more precise in expressing your question. Identifying the type of question you are asking can also help to point you in the direction of the empirical methods you will need to use.

The original question is your main question, and the one that you always need to come back to. But doing analysis is almost never as simple as answering just one question. Usually, in the process of seeking to answer the question you commenced with, extra questions will come up, that you will feel you need to answer in order to be confident about or to fully understand the answer you give to the main question.

For a *describe* question, that follow-up question might be: 'Will I find the same outcome if I use a different data source, or if I compare with a different country, or if the weighting I have used to create my key variable is done differently?'

For an *explain* question, it might be that having found that your original hypothesis to explain an outcome does not work out, you need to go on to consider other potential explanations. Or even if your original hypothesis does work out, every cause has its own cause, so you may want to go to the next stage of investigating what is behind the initial causal factor you investigated.

In this way, understanding labour market outcomes most often requires you to form and answer a sequence of questions. In fact, in most analyses that you undertake it will be the case that you could go on asking extra questions forever. The issue is knowing when to stop, realising when you have reached the point where you know enough to answer your main question with sufficient detail and rigour.

There is no single right way to decide on the extra questions that should follow from the main question you begin with. But there are guiding principles you can follow. *First*, the objective in asking more questions is to add to your knowledge. So, you should try to think of follow-up questions that are going to add most to being able to answer your original question. *Second*, the type of questions that are likely to add usefully to your knowledge almost always relate to the context of your analysis: the specific question you began with, the data and variables you are using, what theory says, how policy might be applied, and what is already known. *Third*, generally there will be multiple approaches to asking extra questions that will add to your knowledge that are relevant to the main question you begin with. That makes them all potentially worth pursuing, and choosing among them will come down to your judgement.

Here is an illustration of labour market analysis as a sequence of questions, drawn from my study of the Australian labour market during the COVID-19 era. The issue of labour market tightness has been of significant interest to policymakers from around late 2021. A standard measure of tightness is the vacancy rate, and no doubt a question that many labour market analysts have been asked in the last several years is: what has happened to job vacancies in Australia during the COVID-19 era? The answer, shown in Figure 2, is straightforward: the vacancy rate has skyrocketed. Not only that, but even in late 2024 it remained well above historical levels.

Figure 2. Job vacancy rate, Australia, May 1979 to November 2024 (quarterly; sa)

Notes: Calculated as (Vacancies)/(Vacancies + Employment).
Source: (i) Vacancies – ABS, Job Vacancies Australia, Table 1: (ii) Employment - ABS, Labour Force Australia, Table 1.

Knowing that the vacancy rate hit such heights and remained elevated leads naturally to a next question: Why?' To be able to answer that question, it is necessary to think of possible explanations to be investigated. Figure 3 shows some of the different ways this could be done.

You might try to look at the extent to which the higher vacancy rate has been due to growth in labour demand or to restrictions on labour supply, such as lower migration inflows. You might study changes in the vacancy rate by industry or occupation, with the idea that those patterns would suggest potential explanations. Or you might analyse the Beveridge curve relation between vacancies and unemployment to understand to what extent the higher vacancy rate reflects cyclical or structural influences. It is a matter of judgement what to do.

In my analysis, I chose to ask as the next question to what degree the high vacancy rate at the end of 2024 was due to: (i) a higher rate of creation of new job vacancies or (ii) vacancies being filled at a slower rate. I did this because I anticipated that answering that question would directly feed into a further question, that could get me close to having a sufficiently precise answer to the question of why the vacancy rate remained so high. That was borne out as I did the analysis.

Having discovered that the high vacancy rate at the end of 2024 was mainly due to vacancies being filled more slowly, I was able to use that finding to ask as a follow-up question: What were the relative roles of the rapid pace of employment growth in Australia and labour market mismatch in explaining the high vacancy rate? And with the

answer to that question in hand, I was able to go on to also ask why there might be higher mismatch in the labour market.

Figure 3: Labour market analysis as a sequence of questions

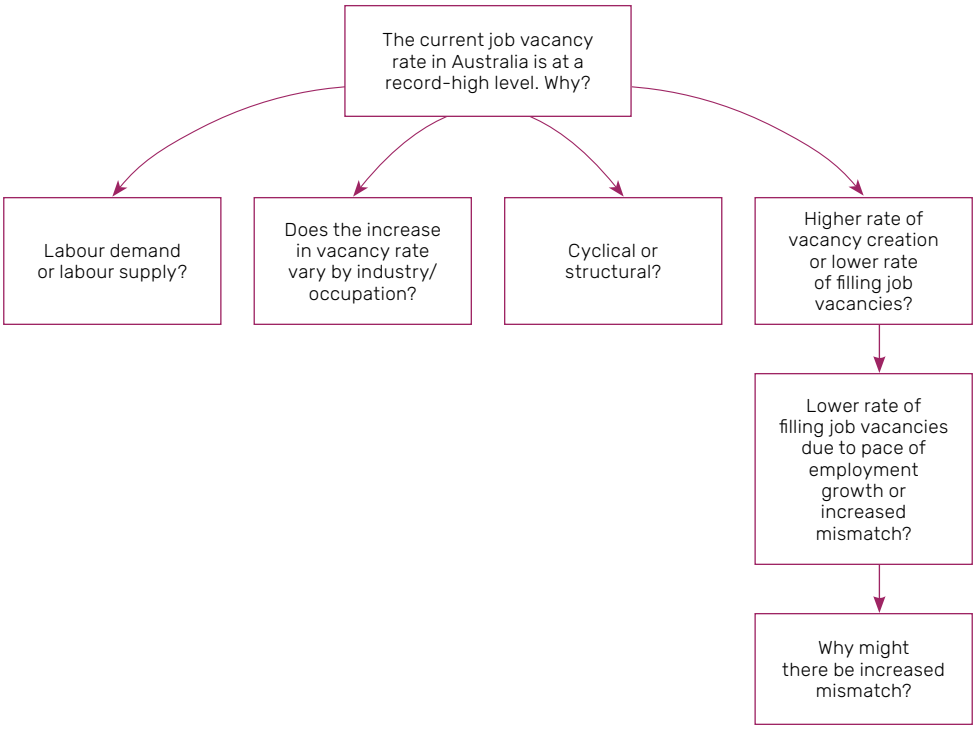

I am not suggesting that my sequence of questions turned out to be more informative than alternative ways in which the 'why' question might be asked. What is important though is that I had an argument to justify how my line of questioning would add to my knowledge in a way that was useful for understanding the elevated vacancy rate at the end of 2024. That made it worthwhile to pursue the line of questioning. No doubt fully understanding why the vacancy rate remains so high also requires pursuing alternative approaches to asking additional questions.

b. Know the drivers and follow the consequences

Economists bring a standard approach to thinking about how changes in market outcomes occur. For the analysis of labour markets, this way of thinking can be translated as:

- Labour market outcomes, such as wages and employment, are causally determined by a set of 'drivers'. Those drivers consist of the main forces that affect labour demand and labour supply, together with institutions and policies.
- It follows that changes in the drivers give rise to changes in labour market outcomes. Or put the other way around, labour market outcomes adjust in response to changes in the drivers in a market.

All this really means of course is that economists use models to understand markets, and when you change exogenous variables in a model, you get different predictions for the endogenous variables. However, recognising that this is how we think does, I would argue, have an important implication for doing labour market analysis: If you are going to analyse labour markets, it is valuable to carry round in your head (i) a list of the potential **drivers** operating on any labour market and (ii) a list of potential labour market **outcomes** – that is, the ways that adjustments can occur in a labour market.

I will now spend some time discussing each of these 'lists'.

Drivers

Different drivers cause different causal impacts on labour market outcomes. For that reason, to understand adjustment in labour market outcomes, you must be able to precisely identify the drivers that have changed. This requires being aware of the set of potential drivers, which is a lengthy list.

Tables 3a and 3b present my summary of major drivers of labour market outcomes. The tables classify drivers between those that directly affect labour demand and/or labour supply, and those that derive from economic policy. Both tables also further divide each category according to the duration over which the driver impacts on the labour market.

Table 3a. Drivers that directly affect labour demand and/or labour supply

Time horizon	Examples of type of driver	Main features
Long run/Tectonic	Technological change Climate change Demography Social norms Globalisation Educational profile of workforce	Slowly evolving and continuous incremental permanent impacts over a long time period
Long run/One-shot	Change in labour supply preferences (e.g., from learning new information about benefits of working from home)	Large permanent impact that commences at time of change in driver
Medium run	Mining boom COVID-19 Extended extreme weather event (e.g., drought over several years) War/international conflict	Large temporary impact concentrated in a time period of several years
Short run	Short-term extreme weather events (e.g., cyclones) Regular seasonal cycle Strike activity	Large temporary impact concentrated in a limited time period (e.g., a few months)

Table 3b. Policy drivers

Type of policy	Main features
Policies subject to ongoing reform to address long-run drivers (e.g., climate change policy; education and training policy)	Slowly evolving and continuous incremental permanent impacts extending over a long time period
Monetary and fiscal policy to manage business cycle	Large temporary impact concentrated at time of policy implementation
One-off policy reforms that affect labour market e.g., reforms to IR system; minimum wage; immigration policy; childcare policy; income support payments; public sector employment policy	Large permanent impact concentrated at time of introduction of policy

Outcomes and adjustment

Understanding labour market outcomes and adjustment requires being able to characterise the set of labour market outcomes that might alter when a driver changes, knowing what are the causal mechanisms that connect a change in a driver to a labour market outcome and appreciating the linkages that can exist between labour market outcomes.

I will now say something more about each of these aspects.

The set of labour market outcomes

Once you start thinking about what might be considered as a labour market outcome, you quickly realise it is a broad set. Table 4 lists a set of potential outcomes, taken from the excellent overview of adjustment in labour markets by Richard Blandy and Sue Richardson (1982).

Table 4. Labour market outcomes – Margins of adjustment when drivers change

Average pay level
Monetary and fiscal policy to manage business cycle
Relative pay (e.g. between new hires/incumbent workers; by occupation; between firms)
Total employment
Average hours worked
Entry requirements/Hiring standards
Rates of promotion within organisations
Characteristics of jobs
Rate of worker turnover/Worker tenure
Extent of occupational/industry mobility
Incidence of self-employment
Labour force participation
Worker training and acquisition of skills
Job search effort
Rates of flow between labour force states/Rate of job-to-job transitions

Source: Blandy and Richardson (1982)

Blandy and Richardson's list suggests that a thorough analysis of the impact of a change in a driver will often require investigating an extensive set of labour market outcomes. As a case study, we can look at the array of ways in which labour markets might adjust to an increase in the minimum wage, as summarised in recent work by Arindrajit Dube and Attila Lindner (2024). For a long time, research on the minimum wage pretty much exclusively looked for effects on employment. That has changed in the past decade, with a much broader potential set of modes of adjustment now being considered. The modes of adjustment reviewed by Dube and Lindner are: Changes in total employment, the composition of employment by full-time/part-time status, job amenities, worker turnover, worker productivity, the composition of workforce by skill level and labour force participation are the main margins of adjustment considered (as well as non-labour market outcomes such as firm entry and exit, output prices and profits).

Causal mechanisms

The way we think about labour markets needs to be soundly based. If we are going to suggest a linkage between a change in a driver and adjustment in a labour market outcome, we want there to be a causal basis for that association. To believe a causal relation exists will depend on both theory and empirical analysis.

Let's continue with the example of an increase in the minimum wage. Critical to thinking it worthwhile to study the modes of adjustment suggested by Dube and Lindner is that the economic theories of labour demand by a profit maximising firm and/or labour supply by a utility maximising household provide a basis for believing there is a potential causal relation between an increase in minimum wage rate and those outcomes. We know, for example, that reducing the quality of job amenities, hiring higher productivity workers or substituting workers being paid the minimum wage for other types of workers or capital, are ways that a profit maximising firm facing a requirement to pay a higher minimum wage might adjust.

Theory therefore provides the basis for identifying and investigating each potential causal relation from an increase in the minimum wage. To believe that a causal relation actually does exist with any of the outcomes suggested, it is necessary to go further and undertake empirical analysis that finds evidence in support. One example where this has happened is from quantitative studies of the impact of minimum wages on worker turnover, that almost universally finds a significant negative causal relation (Dube and Linder, 2024, Table 5).

Linkages between labour market outcomes

It is important to appreciate that changes in labour market outcomes may themselves be linked, often a further essential element to understanding adjustment. Continuing again with the minimum wage example, an increase in the minimum wage may cause firms to substitute away from workers paid the minimum wage towards other workers. If that substitution effect is widespread enough, the increase in the demand for the labour of the other workers may be sufficient to cause an increase in their wages. That is, the increase in minimum wage rate is linked to an increase in the pay of higher wage workers. Appreciating the breadth of the potential linkages between labour market outcomes, which I refer to as 'seeing the labour market whole', is vital for understanding how adjustment happens.

To illustrate the importance in labour market analysis of knowing the drivers and being able to follow adjustment, I will use as my example the initial impact of COVID-19 on Australia's labour market.

The spread of COVID-19 within Australia in March 2020, combined with growing awareness of rapidly developing health disasters in other countries, was the background to its immediate impact – a 'shutdown' of economic activities, caused by government-imposed restrictions on travel and mobility and prohibitions on activities that risked transmission of COVID-19, together with voluntary withdrawal by consumers from some activities. It was this shutdown of economic activities that was the main driver of the impact of COVID-19 on Australia's labour market in early 2020.

Identifying that shutdown was the main driver of what happened in the labour market – and this is the critical point I want to make – then allows us to trace out how employment outcomes adjusted. Figure 4 presents a flow-chart summarising this adjustment, focusing on the impact of shutdown on labour demand. (Breunig et al., 2024 identify labour demand as the main channel of transmission to employment outcomes).

Figure 4. Understanding the initial impact of COVID-19 on employment

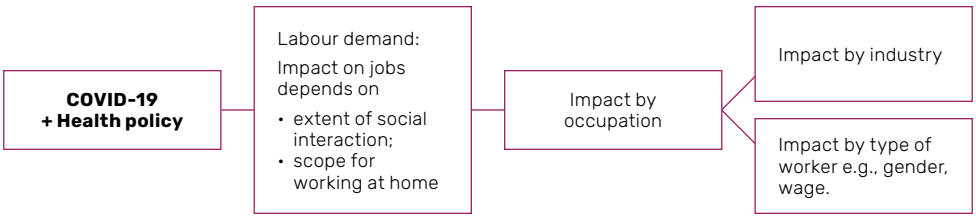

The start of the flow chart recognises that the most direct immediate impact of COVID-19 was on labour demand at the job-level. Restrictions imposed on travel and mobility implied a concentrated negative impact on labour demand in jobs where there was not a capacity to perform that job at home. That negative impact was intensified in jobs with a high level of social contact, which were in addition adversely affected by government prohibitions and voluntary withdrawal by consumers. Table 5 demonstrates how this impact played out in the pattern of decreases in employment by type of job in the earliest phase of the pandemic, using 4-digit occupations to represent jobs.

Table 5. Job-level impacts of COVID downturn by extent of social contact and scope to work at home

Type of job	Change in persons employed, Feb to May 2020 (%)	Examples of jobs
Low scope to work at home plus high social contact	-13.3	Early childhood teachers; physiotherapists; general practitioners; bricklayers; chefs; hairdressers; waiters; mail sorters
Low scope to work at home plus low social contact	-6.8	Engineering managers; land valuers; surveyors; safety inspectors; motor mechanics
High scope to work at home plus high social contact	-0.2	Primary school teachers; university lecturers and tutors; call centre workers
High scope to work at home plus low social contact	-1.7	Graphic and web designers; urban and regional planners; software and applications programmers

Source: Borland (2025a)

This impact on jobs translated into an impact on the distribution of employment between broader occupation categories, the next stage in the flow chart. Differences in how types of jobs (4-digit occupations classified on the basis of scope to work at home and extent of social contact) are distributed across broader occupation categories (1-digit occupations) meant that the shutdown due to COVID-19 differentially affected employment in the broader categories. The most adversely affected 1-digit occupations were community and personal service workers, salespersons and labourers with professionals being least affected.

In turn, because the occupational composition of employment differs between industries and across workers by demographic and job characteristics, an unequal impact of the pandemic between occupations also implied unequal impacts on employment on those other dimensions. These are the final stages in the flow chart.

The impact on employment by worker characteristics is shown in Table 6. Females were more negatively affected than males, reflecting their greater concentration in the most adversely affected occupations. For the same reason, young workers were much more adversely impacted than mature age or older workers. Casual employees were more negatively affected than permanent employees. This likely reflected the larger share of casual employment in the most affected occupations, as well as casual employees with less than 12 months tenure being ineligible for the Job Keeper program.

Table 6: Impact on employment, By worker characteristics, Early 2020

		Change in employment	Share of employment in 3 occupations (1-digit) most negatively affected by COVID-19 (%)
1]	*Sex*	*Change in monthly hours – March to May 2020 (%)*	*Feb. 2020*
	Males	-8.7	23.4
	Females	-12.5	28.9
2]	*Age*	*Change in weekly hours – March to May 2020 (%)*	*Feb. 2020*
	15-24 years	-23.4	57.9
	25 plus years	-8.1	24.5
3]	*Employee status*	*Change in weekly hours – February to May 2020 (%)*	*Feb. 2020*
	Permanent	-6.1	22.3
	Casual	-27.0	46.2

Source: Borland (2025a)

Two important themes about drivers and adjustment are illustrated from the example of the immediate impact of COVID-19 that I have just worked through:

First, it demonstrates how knowing the driver(s) is critical for tracing out adjustment. In this example, identifying the driver as the shutdown of economic activities

allows us to see a direct causal link from onset of COVID-19 to the pattern of adjustment in employment outcomes.

Second, it shows the importance of 'seeing the labour market whole'. Recognising the links between alternative labour market classifications enables us to appreciate how the initial direct impact of COVID-19 on the composition of employment by job type would have implications for the composition of employment by occupation, industry and between workers with different demographic and job characteristics. Those impacts on the different types of employment outcomes were simultaneous, since each outcome is just a different way of classifying employment at the same point in time. Hence, in this case 'following adjustment' is about specifying a logical structure for understanding a set of changes that happened simultaneously. In other cases, 'following adjustment' may be more about setting out a path of adjustment in outcomes that occurs over time with, for example, a change in a driver causing a change in a labour market outcome, which itself causes a change in another outcome.

c. How to judge causality

Establishing the causal relations that exist between drivers and outcomes in a labour market is at the core of labour market analysis. It is central to answering *explain* and *policy* questions and is also relevant to *predict* questions. These days, the mention of causal relations immediately brings to mind 'program evaluation' techniques such as randomised controlled trials, difference-in-difference, regression discontinuity and instrumental variables techniques. Increased application of these methods underlies what is often referred to as the 'causal revolution' in micro-econometrics and has brought a significant increase in the rigour with which causal relations can be identified. Seminal contributions were made by David Card (1990), and Josh Angrist and Alan Krueger (1991). (For a recent review see Garg and Fetzer, 2025.)

The problem for our purposes, however, is that for many (maybe most) policy analysis type questions we want to answer about the labour market, program evaluation techniques cannot help us at all, or will only get us so far. This is because for many questions it is likely that no (or only limited) relevant empirical evidence from program evaluation techniques will be able to be generated. This could reflect the nature of the question not being amenable to those techniques, or data limitations or time constraints that prevent you using those techniques. When this happens, descriptive information, facts about the labour market derived from descriptive statistics and simple regression analysis, become the primary source of empirical evidence that can be applied.

So, how should we proceed if we want to determine whether a causal relation exists between a potential driver and a labour market outcome, and we need to use descriptive information to do that? Suppose for example that we have been asked to analyse whether increased opportunities for female participation in higher education have had a causal impact on female labour force participation, and to determine the magnitude of any such impact. How can we go about answering this question?

We need to begin by collecting the information on which we will rely to answer the question. I would suggest several strategies for doing this:

First, begin with theory. Establishing that there is a plausible theory for a causal relation between the variables being analysed is an essential step. Without a theory for a causal relation that is consistent with our conception of how humans behave and interact, it is impossible to fully believe that such a relation exists. Meeting the standard of plausibility will depend on whether aspects of the theory such as the assumptions made and the channels through which causality is depicted as operating, themselves seem plausible. But while theory is therefore a necessary condition for believing a causal relation exists, it is rarely (if ever) a sufficient condition. Economists are notably inventive when it comes to thinking up plausible explanations for causal relations between variables, so we would usually like confirmation from data. Not only this, but theory also has limits. While it may give us confidence a causal relation exists between variables, it may not, for example, tell us anything about the strength of that relation.

So, **second**, we need to move on to look at data. Specifically, we want to do analysis that I am assuming needs to be undertaken using exclusively descriptive information that is available and relevant to the causal relation we are investigating.

The important point I want to make is that when we use that descriptive information in a directed way, it can inform us about causality. We are all familiar with the expression 'correlation is not causality'. But my own feeling is that this blanket rejection of using correlation to understand causality goes too far. Carefully directed analysis of correlation can in fact be valuable.

What do I mean by 'carefully directed'? I have in mind using descriptive information to investigate for causal relations: *first*, by applying what is known about the channels through which a causal relation would be expected to operate; and *second*, by applying simple principles underlying causality.

To illustrate the first point on causal channels, we can use my hypothetical investigation into whether female participation in higher education is causally related to female labour force participation. My argument is that it is possible to do better than, for example, just looking at a correlation over time between national-level aggregate outcomes for those two variables.

We could instead begin by asking: Which females have increased access to higher education and for which courses? Once we know that information, we can work out what it implies about the groups of females who we should observe to have higher labour force participation, and which types of jobs they should be doing. For example, where increased access to higher education is focused on females from low socioeconomic status groups, any direct causal impact on labour force participation should be for the same group. Or if increased access to higher education is for groups doing specific courses, any causal impact should primarily be reflected in higher participation by females working in occupations that those courses qualify them for. Looking more precisely at the relation between female higher education and labour force participation in the way I am suggesting may not exactly identify a causal relation, but it certainly gives more confidence of a causal relation than looking at correlation in national aggregates.

The underlying principle here is to use logic and/or theory to develop more detailed ways to test for causal relations using descriptive data. As Robert Abelson (1995, p.183) writes: 'A powerful strategy is to spell out the details of the causal mechanism... and then test the consequences this mechanism would entail.'

There is another set of questions relating to causal implications that could also be asked. This is to ask what else we might expect to observe happen, if there is a causal relation. For my example this would involve asking: Are there other outcomes we would expect to see if there is a causal relation between female participation in higher education and female labour force participation? One such implication could be for wages. If increased female labour force participation is being driven by increased higher education participation, we should expect to observe an increase in average female wages.

My second point is that whether descriptive statistics are consistent with basic principles underlying causality can be informative for supporting the existence of a causal relation. My hypothetical example can again illustrate. If education participation is truly driving labour force participation, the former must happen before the latter. That is, there is a timing element that can be part of our test of causality – as Austin Bradford Hill (1965) puts it in his list of checks for causal relations, the horse must come before the cart. There are other similar checks relying on principles of causality we could also apply. Other things equal, if a causal relation exists, we would likely expect that a larger increase in female education participation should give rise to a larger increase in female labour force participation. Similarly, we would expect that there should be a consistent impact across time for different cohorts of females for whom there is increased access to higher education.

One final point about doing your own empirical analysis to assess causality using descriptive statistics. My focus in this discussion has been on how descriptive statistics can be used to identify and therefore test causal relations. Other factors, though, must also come into play when you judge the credibility and relevance to your question of the findings from the analysis you do – such as data quality and how close the available data get you to your question, as well as the statistical and economic significance of what you find.

The **third** strategy is to use existing evidence effectively. It is likely that there will be existing studies with findings that relate to the question about a causal relation you want to answer. It is important to incorporate that evidence into your analysis – depending on your judgement as to its relevance and credibility.

The **fourth** strategy is to be willing to cast a wide net in your investigation, beyond using quantitative methods or to looking at the literature. Talking to stakeholders, people involved in the labour market you are studying, is an essential complement to the other strategies. As a pioneer of the field of labour economics, Clark Kerr, wrote (1988, p.38): 'Experience something of reality first hand, as well as second hand through statistics.'

There is a reason why an important part of any Productivity Commission inquiry is to talk to market participants and allow for public submissions, why Treasury and the Reserve Bank have business liaison programs, and why Jobs and Skills Australia has

such an extensive structure of engagement with business, unions and educators (on this latter example, see Dawkins, 2024). It is the market participants who can tell you about institutional details, about the aspects of outcomes that cannot be captured in data, and about their own motivations and those of other stakeholders.

Many of my own examples of when it is been critical to hear from stakeholders involve understanding the impact of labour market policies. Twenty years ago, in a team analysing the Mutual Obligation Initiative, we only understood why the proportion of income support recipients doing the program was much less than expected, once we talked to the Centrelink case managers in charge of administering the program. More recently, as part of the Economic Inclusion Advisory Committee, I have learned that it is simply impossible to understand the inadequacy of the current level of JobSeeker payment without hearing from people receiving the payment about their lives.

Using information to answer the question

Applying these four strategies I have suggested – looking at theory, using data, reviewing existing evidence and talking to market participants – you will have the information which you now need to use to decide if a causal relation exists. How should you do that? What you cannot do is to plug the information into an algorithm or mathematical formula that will tell you the answer. Instead, it is going to be up to you to come to a judgement, justified by your interpretation of the evidence.

There is no single correct way to come to a judgement. But I do think there are practices to follow that will make for a judgment that best reflects and is justified by your evidence.

These practices are for you to:

- Make a complete and fair minded summary of the available evidence as it relates to the question about causality you are investigating.
- Weight each piece of evidence according to its importance for answering your question (where importance depends on credibility and relevance).
- Give an integrated, logical and as coherent as possible account of what the separate pieces of evidence imply for the answer to your question (accepting that often it won't be possible to entirely resolve inconsistencies). While you will ultimately need to be able to give this account in words, graphical methods can be a useful tool for mapping out the set of evidence and assessing causal relationships; see for example, Hernan and Robins, 2024, chapter 6.
- Present your answer, which should follow directly from your account of the evidence, always being willing to indicate your degree of uncertainty and what extra evidence might be valuable.

I think it is this sort of approach that John Kay and Mervyn King (2020, p.410) have in mind with their concept of narrative reasoning:

> 'Narrative reasoning is the most powerful mechanism for organising our imperfect knowledge. Understanding the complex world is a matter of constructing the best explanation – a narrative account – from a myriad of little details and the knowledge of context derived from personal experience and the experiences of others.'

To finish up on this topic, and to illustrate the role that the directed use of descriptive information can play in thinking about causality, I will again use an example from my study of the Australian labour market during COVID-19, this time about the 'great resignation'.

A rising job quit rate and apparent increase in the incidence of retirement in the United States in 2020 was interpreted as revealing a fundamental change in attitudes to work – and to portend the global spread of mass resignation and retirement. The enthusiasm for a 'great resignation' spread to Australia, and once the Australian Bureau of Statistics (ABS) released its *Participation, Job Search and Mobility* survey data for the year ending February 2022, showing a jump in the rate of job leaving, mania took hold – with headlines such as 'The Great Resignation hits Australia…' and 'Evidence of Great Resignation Emerges'. But did the rise in job leaving really represent a change in attitudes to work?

To test this causal proposition, let's start with theory. The theory of labour supply tells us that a change in attitudes to work is indeed a possible explanation for the rise in job leaving. If (some) people increase their preference for leisure over consumption, we would expect to observe a higher proportion of workers leaving their jobs. However, there is an alternative perspective, from the theory of labour demand. Stronger economic conditions mean a higher rate of new job creation. New jobs are filled in part by individuals coming from outside the labour force, but also by current workers. The rate of current workers leaving their jobs (to take new jobs) increases when labour market conditions improve, as was happening in Australia in 2021. Hence, theory also suggests a different explanation; and it is worth noting, an explanation that is supported by extensive existing empirical evidence on procyclicality of job leaving – for Australia, see D'Arcy et al. (2012, p.5). Even with regard to labour supply, an extra consideration is that a strong motivation for moving out of work that existed in United States, to avoid becoming ill with COVID-19, should not have affected behaviour as strongly in Australia.

Having alternative theories means we need to move on to empirical analysis, using descriptive information. We can begin with the aggregate measure of job leaving, shown in Figure 5. Looking at the series it is easy to see the jump in year to February 2022, and then also to February 2023. But what is apparent as well is that the rate of job leaving had been well below norm in the year to February 2021, and in the year to February 2024 largely reversed from the higher rates in preceding years. In assessing the size of jump it is also important to be aware of the strength of growth in labour demand in Australia at that time.

What should we conclude from this? I would argue that the aggregate evidence is not consistent with a change in labour supply (great resignation) explaining the

increase in job leaving rate. First, it seems that the jump in job leaving in the years to February 2022 and 2023 job leaving can be explained by more regular features of the labour market: delayed catch-up to job leaving being below-norm in the year to February 2021 and strong employment growth in 2021 and 2022. Second, the turnaround in job leaving in the year to February 2024 seems to contradict the idea of a permanent change in preferences for working or at least shows it to have been a limited phenomenon.

Figure 5. Annual rate of job leaving, Australia 2015 to 2024

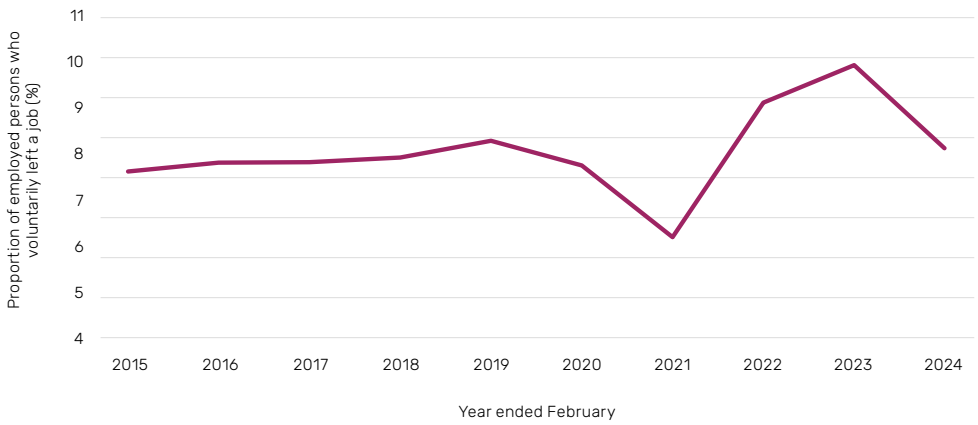

Source: (a) Employed persons who left a job in last 12 months - ABS, Job Mobility; (b) Number of persons who held a job in last 12 months – ABS, Participation, Job Search and Mobility Microdata, Table 1.

But we don't just have to rely on aggregate data, it is also possible to do a set of more directed analyses:

First, current outcomes can be compared against the longer history of job leaving, to give extra context. Doing this reveals that the type of down and up in job leaving observed during the COVID-19 era has been a common feature going back to 1980s and supports the idea that the job leaving rate rises during periods of strong growth in labour demand.

Second, an implication of the theory supporting the great resignation was that job leavers would move out of the labour force, rather than to another job. Hence, if the great resignation had been the powerful force claimed, a decrease in the labour force participation (LFP) rate should have occurred. But contrary to that prediction, we observe that the LFP rate, in aggregate and for all age groups, rose during the period in which the great resignation was supposed to be happening. In December, prior to COVID-19, Australia's LFP rate had been 65.7 per cent, but three years later, in December 2022, at the time when the impact of great resignation was supposed to be peaking, the LFP rate had grown to 66.6 per cent.

Third, it is possible to look specifically at motivations for job leaving. Job leaving due to the great resignation should have been associated with more workers quitting for reasons of wanting to do other activities, rather than, for example, their current job ending. However, using data from the ABS Labour Force Survey on workers' expectations of remaining with their employer over the next 12 months, shown in Figure 6, it can be seen that the slight increase in expectations of job leaving during 2021 and 2022 was mainly due to a larger proportion of workers expecting to quit to change jobs or seek other employment.

Figure 6. Proportion of persons in employment who expect to be with current employer/job in 12 months, by reason, February 2019 to February 2023

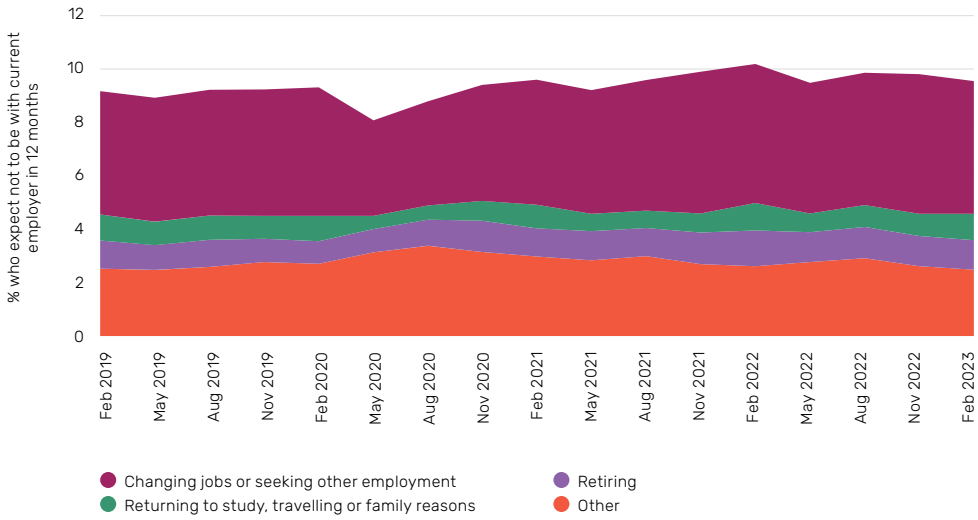

Source: ABS, Labour Force Australia – Detailed, Table 17.

What I have just presented is a combination of theory and set of descriptive analyses relating to whether the great resignation was real; that is, due to a change in labour supply. The judgement I come to from looking at that evidence is that the rise in job leaving in 2021 and 2022 was due to the cycle in labour demand, and not to a change in labour supply. That is, no great resignation.

How did I come to this judgement? It is the volume and consistency of the evidence: the movement in aggregate job leaving being consistent with the labour demand explanation, and the failure to find evidence supporting alternative implications of the great resignation (such as a decrease in labour force participation) or of the causal channel through which the great resignation would be expected to operate (such as leaving to study or travel).

Conclusion

This has been a presentation about method. Of course, in the end, the method of analysis is just the means to what really matters: the extra knowledge about the operation of the Australian labour market that our analysis brings. Still, the quality of what we learn depends hugely on the quality of the methods we apply – and there is nothing automatic about choosing those methods. That is why I think it is worthwhile for us every now and then to step back and try to say what our methods are, and what constitutes good practice. This must be more than a listing of the available technical tools, such as econometric estimators and our theoretical models of the labour market. Rather, it has to be about the ways we can or should apply those technical tools (together with our other skills).

We have a question, how are we to go about answering it? Hopefully, my talk will have given some insights into approaching that task.

References

Abelson, Robert (1995), *Statistics as a Principled Argument* (Lawrence Erlbaum Associates).

Angrist, Joshua and Alan Krueger (1991), 'Does compulsory school attendance affect schooling and earnings?', *Quarterly Journal of Economics* 106(4): 979-1014.

Blandy, Richard and Sue Richardson (1982), 'Introduction' and 'Conclusion', pages 1-22 and 261-83 in R. Blandy and S. Richardson (eds) *How Labour Markets Work: Case Studies in Adjustment* (Longman Cheshire).

Borland, Jeff (2025a), 'Decline to boom, and back again: The Australian labour market in the COVID-19 era', mimeo, University of Melbourne.

Borland, Jeff (2025b), 'How to read a labour market', mimeo, University of Melbourne.

Bradford-Hill, Austin (1965), 'The environment and disease: association or causation?', *Proceedings of the Royal Society of Medicine*, 58: 295-300.

Breunig, Robert, Wei Cheng, Laura Montenovo, Kyoung Hoon Lee, Bruce A. Weinberg and Yinjunjie Zhang (2024), 'Cross-country analysis of labor markets during the COVID-19 pandemic', National Bureau of Economic Research, Working Paper no. 33029.

Card, David (1990), 'The impact of the Mariel Boatlift on the Miami labor market', *Industrial and Labor Relations Review* 43(2): 245-257.

D'Arcy, Patrick, Linus Gustafsson, Christine Lewis and Trent Wiltshire (2012), 'Labour market turnover and mobility', *RBA Bulletin*, December, 1-12.

Dawkins, Peter (2024), 'The value of policy relevant research in education, training, jobs and skills', *Australian Journal of Labour Economics*, 26(2): 114-31.

Dube, Arindrajit and Attila Lindner (2024), 'Minimum wages in the twenty-first century', National Bureau of Economic Research, Working Paper no.32878.

Garg, Prashant and Thiemo Fetzer (2025), 'Causal claims in economics'; accessed at: https://arxiv.org/pdf/2501.06873

Hernan, Miguel and James Robins (2024), *Causal Inference, What If?* (Chapman and Hall).

Kay, John and Mervyn King (2020), *Radical Uncertainty* (The Bridge Street Press).

Kerr, Clark (1988), 'The neoclassical revisionists', in *How Labor Markets Work* by Bruce Kaufman (ed.) (Lexington Books).