From Micro to Macro: Using Administrative and Other Microdata Sources to Answer Macroeconomic Questions

Jonathan Hambur

A thesis submitted for
Doctor of Philosophy
The Australian National University

Crawford School of Public Policy

June 2026

Disclaimer

I hereby declare that the work undertaken in this thesis has been conducted by me, and that I have made significant contribution to each of four Chapters in this thesis. I undertook parts of this research while employed by the Australian Treasury and Reserve Bank of Australia. Some of these papers have been published, or had early versions released as working papers, as noted below. No part of this work has been submitted to this or any other university as part of a degree.

Chapter 1: Does Monetary Policy Impact Innovation? Evidence from Australian Administrative Data, with Omer Majeed and Robert Breunig

I initiated the project idea, designed the econometric approach and tests, ran the macro-data regressions, and refined and ran the final specifications of the microdata analysis. I drafted the majority of the first draft of the paper and provided editing and oversight of other's contributions.

This paper has been published in the Journal of Macroeconomics

Chapter 2: Adoption of Emerging Digital General-purpose Technologies: Determinants and Effects, with Kim Nguyen

I initiated the project, designed the econometric approach, did the initial text analysis, designed the basic event study analysis, and undertook the WoolDridge extensions. I contributed to the first draft of the paper and took forward further revisions and editing.

Published as RBA Research Discussion Paper 2023-10

Chapter 3: Shifts in Australian Price-setting Behaviour Around Large Shoch, with Matt Fink

I secured the data, initiated the project, designed the sales algorithm, undertook the hazard rate regressions and the DSGE modelling exercises. I drafted first drafts of the DSGE and hazard rate sections and edited and refined the rest of the paper.

Working paper published as ANU CAMA working paper 07/2026

Chapter 4: The effect of labour market concentration and declining entry rates on wages: Evidence from Australia

I was the sole author of this paper

Published as RBA Research Discussion Paper 2023-02

Acknowledgements

I would like to acknowledge the assistance of Paula Drew in copy editing of all four chapters of this thesis. Her assistance in picking up my numerous typos and formatting mistakes was very much appreciated. I would like to thank my supervisors Robert Breunig, Begoña Dominguez and Yixao Zhou for their input, assistance and encouragement. I would like to thank my examiners, Aniket Baksy, Erwan Gautier and Christian Gillitzer. And I would like to thank the Reserve Bank of Australia and Australian Treasury where I was employed while I undertook much of the research. I would also like to acknowledge staff at the ABS for their work in building, curating, and clearing the microdata assets I have used in this thesis. Without their tireless work I would not have been able to complete this thesis.

Finally, I would like to thank my family and my wife for their ongoing love and support while I was working on this thesis. This thesis is dedicated to my sister Michelle, who passed during the preparation of this thesis.

Overview

This thesis presents four chapters using Australian microdata sources and techniques to explore important macroeconomic issues.

The first two chapters focus on productivity, the key driver of living standards over the long-term. The first of these explores the extent to which monetary policy, or demand shocks more generally, can influence innovative activity and productivity over the medium-term. While this has been explored for the US, it is yet to be considered for a small open technology importer, nor using broad measures of innovation that capture technology adoption, rather than just frontier innovation. Combining macroeconomic and firm-level measures of innovation with monetary policy shocks in a local projections framework, this work finds that contractionary policy can influence innovation. However, the effects are shorter-lived than previously found in the US and more heterogenous, with demand and financial constraint channels playing an important role in this heterogeneity. These findings reinforce the importance of macro-stabilisation policies to avoid medium-term productivity scarring

The second chapter examines the adoption of emerging digital technologies: Cloud Computing and AI/Machine Learning. Understanding barriers and enablers to profitable adoption of such technologies is crucial if we are to reap their benefits over coming years. Using private sector data, along with natural language processing, this Chapter builds a unique merged dataset of firm-level adoption, profits, hiring and management capabilities. Applying event study approaches, it finds that the average firm does not benefit from adoption in terms of profitability. But firms whose Board have relevant skills do. They are also the firms most likely to increase hiring of skilled workers around adoption. So, this work highlights the important role of skills, particularly management skills, in the profitable adoption of emerging technologies, which has implications for the focus of skills policies.

The third chapter focuses on inflation, and whether changes in the frequency of firm-level price changes can account for the pick-up in inflation post-COVID. Using firm-level web-scraped prices microdata linked to firm tax data, the Chapter shows that firms began resetting prices more frequently in response to the large inflationary shocks post-COVID. This is counter to standard macroeconomic models, which assume constant price-setting frequencies, and suggests nonlinearities in inflation dynamics following large shocks. Using a workhorse policy DSGE model with constant price-setting frequencies, the Chapter shows that failing to account for these changes in price-setting leads to a significant underprediction of recent inflation outcomes. So central banks either need to invest in new frameworks or add judgement during such periods.

The final chapter explores the labour market, and whether growing monopsony power caused slow wages growth in Australia over the 2010s. Using linked employer-employee administrative data it demonstrates that employer concentration is associated with lower wages. While labour markets did not become more concentrated, the effect of given levels of concentration appears to have increased. A key driver was the decline in firm entry, which meant there were fewer young growing firms trying to poach workers from incumbents. This work highlights the important links between firm dynamics and competition, and labour market dynamism, power and wages. It therefore highlights how policies aimed at increasing firm dynamism can potentially support worker bargaining power.

Table of Contents

List of Figures (main text)

List of Tables (main text)

1. Chapter 1: Does Monetary Policy Impact Innovation? Evidence from Australian Administrative Data

Omer Majeed*, Jonathan Hambur** and Robert Breunig*

*Crawford School of Public Policy, Australian National University (ANU)
**Economic Research Department, Reserve Bank of Australia (RBA), Crawford School of Public Policy, Australian National University (ANU)

Abstract

We examine whether monetary policy affects innovative activity and productivity in Australia. As a small open economy that primarily imports and adopts existing technology, research on Australia complements previous research on the United States. While contractionary policy reduces aggregate research and development spending, and in turn productivity, the effects appear more short-lived compared to the United States. When using survey measures of innovation that capture adoption, we find heterogenous responses. Small firms decrease innovation in response to contractionary monetary policy shocks, whereas large firms increase innovation. This heterogeneity may reflect differing exposures to the demand and financial constraint channels of monetary policy. United States monetary policy affects Australian firms' innovation, at least in part by influencing global economic conditions.

JEL Classification Numbers: E32; E52; G32; O30

Keywords: Innovation; Monetary Policy; Firm-level data

We are grateful to Matt Read, Beth Webster, Benjamin Wong, Petr Sedláček, John Simon and Tim Taylor and seminar participants at the Tax and Transfer Policy Institute (ANU) and Reserve Bank of Australia (RBA) for their comments on the paper. We thank Chloe Heiniger with help in preparing the manuscript. We thank an editor and two anonymous referees for comments which helped to improve the paper. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the RBA.

This paper was published in the Journal of Macroeconmics, Volume 86 in December 2025.

1.1 Introduction

Productivity growth is the key determinant of living standards over the long term. Many factors can influence how quicky productivity grows, including technological change, competition, skilled labour, regulation, trade and tax policy (Aghion and Howitt 2008). But economists have traditionally assumed that productivity is unaffected by 'cyclical' factors such as current economic conditions and monetary policy, at least over the medium term.

More recently, though, there is a growing literature questioning this assumption. One stream of this literature argues that weaker economic conditions, particularly economic downturns, can lead to slower rates of innovation and technology adoption by businesses, which can in turn lead to persistent changes in productivity and economic output. These hysteresis effects are considered in Stadler (1990), Comin and Gertler (2006), Anzoátegui et al. (2019), Bianchi, Kung and Morales (2019) and Amador (2022) among others.

In a similar vein, contractionary monetary policy, which weakens economic conditions, can slow innovation and technology adoption, and therefore have persistent effects on productivity (Moran and Queralto 2018; Jordà, Singh and Taylor 2024; Ma 2023; Ma and Zimmerman 2023).[1] While the effects of expansionary and contractionary monetary policy are likely to cancel out over a cycle, such a finding highlights the potential for medium-run economic scarring to occur if policy is constrained by the effective lower bound on interest rates and therefore cannot offset economic downturns.

This has important implications for macroeconomic policy, but there is little evidence on these questions outside the United States (US). Effects could differ substantially in a small open economy that mostly imports innovation, like Australia, compared to a large economy that pushes the technological frontier. Moreover, for a small open economy there is scope for spillovers from US policy via economic conditions, or through changes in the technology frontier itself.

In addition to its US focus, the evidence to date has mostly focused on narrow measures of innovative activity like Research and Development (R&D) spending and patenting. These measures likely miss a large amount of innovative activity for countries such as Australia, where adoption of existing technologies and processes drives aggregate productivity growth (OECD 2015; Comin and Mestieri 2018; Majeed et al. 2021; Argente et al. 2020). Moreover, adoption of technologies features heavily in theoretical models of the effect of monetary policy on productivity (Moran and Queralto 2018). However, the adoption channel has not been empirically explored to date.

Our paper makes several contributions to the literature. First, we examine these questions in a small open economy that largely imports overseas technologies. We find that contractionary domestic monetary policy has little effect on patenting, but a moderate effect on R&D spending and trademarks. However, this effect appears shorter-lived compared to US studies, perhaps indicating that policy is working by influencing shorter-term decisions around adopting and adapting existing technologies, rather than long-term investment in developing new cutting-edge technologies.

1 Baqaee, Farhi and Sangani (2024) propose a different mechanism whereby contractionary shocks lead to a reallocation of resources towards lower productivity firms.

Second, we explore broader measures of innovation that capture adoption of existing technologies. We find that contractionary monetary policy has limited effects on innovation in aggregate. However, this reflects offsetting dynamics for small and large firms, who decrease and increase innovation, respectively, following a contractionary monetary policy shock. This suggests that the effects of policy on adoption may be more complex and heterogeneous than previously identified.

Building off this finding, our third contribution is to explore the mechanisms through which policy could affect innovation, to shed light on this heterogeneity. We find that domestic monetary policy has a larger effect on innovation amongst domestically-focused firms than on exporting firms. This is consistent with the demand channel of monetary policy playing an important role: contractionary policy weakens domestic conditions and therefore incentives to invest. We also find that contractionary policy leads to an increase in the number of firms reporting that a lack of funds is hampering their investment, consistent with policy influencing innovation by making financial conditions more restrictive. This channel appears more important for small and medium firms. We also find that small firms are more likely to report that the costs of innovation are a barrier after a monetary policy shock, whereas there is no effect for large firms. Together, these results help to explain the differing effects of policy across firm sizes.

Our final contribution is to document the fact that US monetary policy appears to spillover and affect innovation and adoption in a small open economy like Australia. While part of this may reflect the effects that US policy has on the technology frontier, at least part also appears to reflect effects on global economic conditions. This is supported by the fact that Australian exporting firms decrease their innovation by more than domestically-oriented firms following a contractionary US monetary policy shock. This builds on the existing broader literature on spillovers from US policy (e.g. Georgiadis 2016; Kearns, Schrimpf and Xia 2022).

Overall, our results suggest that monetary policy can affect innovation and therefore productivity, though they paint a more complex and heterogenous picture than previous work. This reflects both our focus on a small, technology-importing open economy and on broader measures of innovation adoption. From a policy standpoint, the results reinforce the importance of using macro stabilisation policy to avoid sharp economic downturns, given the potential for medium-run economic scarring, whilst also highlighting the costs of having such policy constrained (Benigno and Fornaro 2018; Ikeda and Kurozumi 2019; Barlevy 2004; Garga and Singh 2021). But they also suggest that the costs and trade-offs could differ somewhat for a technology importer, compared to an economy at the technology frontier.

This paper proceeds as follows. We first discuss the relevant literature on the determinants of innovation in Section 1.2 before giving an overview of our data and methodology in Section 1.3. We then look at the effects of monetary policy shocks on aggregate innovative activity – patents, trademarks and R&D – and firm-level innovation in Section 1.4, before exploring the channels though which monetary policy may affect innovation in Section 1.5. We consider the aggregate effect of changes in R&D spending on productivity in Section 1.6 before concluding.

1.2 Literature review

Until recently, the prevailing wisdom was that current macro-economic conditions and monetary policy have no effect on productivity over the medium term. For monetary policy, this assumption is reflected in its so called 'neutrality'. However, a growing number of papers have begun to question this assumption, contributing to the broader literature on hysteresis. Much of this literature focuses on medium-run employment effects, due for example to labor market scarring for workers (e.g., Blanchard and Summers 1986; Ball 2009; Andrews et al. 2020), or changes in the nature and number of start-ups (Sedláček and Sterk 2021; Ouyang 2009; Davis and Haltiwanger 2024).

More recently the scarring literature has examined the medium-run implications for productivity, building on the existing evidence that economic and financial conditions can influence the amount of innovative activity undertaken (e.g., Brown, Fazzari and Petersen 2009; Ouyang 2011; Huber 2018; Webster and Jensen 2009). Early papers in this productivity hysteresis literature focused on economic downturns, and whether they could contribute to lower productivity in the medium-run by weighing on the amount of innovation and technology adoption (Stadler 1990; Comin and Gertler 2006; Anzoátegui et al. 2019; Bianchi, Kung and Morales 2019; Barlevy 2007).[2] Focusing mostly on R&D spending as their measure of innovation, these papers find evidence that innovation declines during downturns and that this can have medium-run effects on productivity and output. Different papers propose different mechanisms for the decline in innovation, including reduced incentives due to lower demand (Comin and Gertler 2006; Anzoátegui et al. 2019; Barlevy 2007) or tighter credit market conditions (Bianchi, Kung and Morales 2019; Queralto 2020). Some papers have also argued that innovation will increase during downturns, as weaker economic activity can free up resources and make innovation cheaper (Aghion et al. 2012).

Recent papers have begun to explore whether monetary policy can influence innovation and technology adoption, and through this have medium-run effects on productivity and economic output. Jordà, Singh and Taylor (2024), using cross-country data, provide empirical evidence that contractionary monetary shocks lead to lower productivity growth over the medium-term. They build a New Keynesian model with endogenous TFP growth, to motivate their empirical findings.

Moran and Queralto (2018) provide more evidence on the mechanisms for the US. Using a Vector Auto Regression (VAR) model, they find that contractionary monetary policy lowers innovative activity, as measured by R&D spending. In turn, lower R&D spending leads to slower productivity growth and therefore lower economic output. They build these mechanisms into a New Keynesian model where weaker economic conditions lower the returns to, and thus investment in, innovation and adoption. Monetary policy can thus influence productivity and economic output in the medium-term. Their model indicates that the medium-run effects of policy on productivity and output can be sizeable and should therefore be of first order importance to policymakers. They estimate that US output and productivity would have been permanently 2 per cent higher had monetary policy not been constrained by the effective zero lower bound following the Global Financial Crisis. Similarly, US productivity and output would have been permanently 1 per cent higher if the tightening of monetary policy from 2016 had been more gradual.

Ma (2023) focuses on patenting as a measure of innovative activity, again for the US. He finds that the value and number of patents rises following an expansionary monetary policy shock, which in turn can contribute to higher productivity. He finds that firms with higher liquidity are more responsive than firms with lower liquidity, suggesting that financing constraints play an important

[2] In a related literature, Sedláček and Ignaszak (2021) consider the role that firm-level demand can play in incentivising innovation, and, in turn, in affecting aggregate growth.

role. In a heterogenous firm model, he shows how this financing constraint mechanism can contribute to longer-lived impacts of monetary policy shocks on productivity and output.

Ma and Zimmerman (2023) focus on R&D spending, patenting and venture capital funding, and find similar results to the above papers for the US. Similar to our paper, they explore the channels through which monetary policy can affect innovative activity. They find that R&D and patenting are more responsive to monetary policy shocks in cyclical industries, suggesting monetary policy affects innovation by influencing demand conditions – the demand channel. They also find evidence that venture capital funding falls following a contractionary monetary policy shock, which in turn is likely to limit the amount of funds available for innovative activity – the financing constraint channel. Applying standard multipliers from innovative activity to output, they suggest that a 100 basis point contractionary shock could lead to output that is 1 per cent lower 5 years after the shock.

Amador (2022) focuses on cross-country measures of the take-up of general-purpose technologies, like electrification. He finds that contractionary policy leads to slower diffusion of these technologies, which can weigh on output in the medium-term.

Our paper also touches on several other literatures. This includes the broader literature examining the heterogenous effects of monetary policy on firms. Most relevant for our paper, a number of these papers have examined the role of credit constraints, firm age and firm size in explaining this heterogeneity. Durante et al. (2022) examines the role of firm age and leverage. Jeenas (2024) focuses on liquidity and finds that less liquid firms respond more strongly, consistent with monetary policy easing constraints on firms. Similarly, Cloyne et al. (2023) assume non-dividend-paying young firms are financially constrained and find that such firms are more responsive to monetary policy. In contrast, Ottonello and Winberry (2020) find that less leveraged and less risky firms respond more to monetary policy shocks as it is less costly for them to gain additional financing. Nolan, Hambur and Vermeulen (2025) find that firms that claim to be more financially constrained are more responsive to policy, particularly along the intensive margin of investment.

Second, our paper touches on the broader literature exploring the relationship between economic cycles and innovation. For example, Manso, Balsmeier and Fleming (2023) document that firms tend to shift from exploration to exploitation as they move from contraction to expansion. This highlights the potential value in exploring broad measures of innovation as we do below.

Finally, our paper touches on the large literature exploring global spillovers from US monetary policy. This literature focuses on the effects of US policy on interest rates, investment, cross-border lending and capital flows, or output in other countries (e.g., Bruno and Shin 2015; Georgiadis 2016; Iacoviello and Navarro 2019; Kearns, Schrimpf and Xia 2022; Arbatli-Saxegaard et al. 2025). We extend this to consider spillovers from US monetary policy to Australian firms' innovative activity.

1.3 Data and Methodology

We do not explicitly develop our own theoretical model for the analysis but the channels of monetary policy that we consider are those of models developed in other papers such as Aghion et al. (2012), Moran and Queralto (2018) and Ma (2023). In the Moran and Queralto (2018) model, innovation adoption is an important element. Innovators spend on R&D to develop new products. Adopters then expend resources to adopt the new innovation, which they can then feed into production and sales. While some R&D spending will capture adoption rather than innovation, adoption activity is much broader and is the mechanism whereby new innovation diffuses into the economy. They only test the effect of monetary policy on the initial innovation step (i.e. R&D spending). Motivated by their work, we look to examine the effects on a broader range of innovation measures, in particular ones that will capture adoption.

1.3.1 Data

We use a combination of survey data linked to firm-level administrative data.

We focus on four different measures of innovative activity. The first three are narrow measures of innovative activity and the fourth is a broad, survey-based measure. All four have been considered in previous literature.

- The (log) flow of new patents filed in Australia by Australian residents from IP Australia's IPLORD database (IP Australia, 2023). This is a fairly narrow measure of innovative activity that is more likely to capture the creation of 'new-to-world' innovation and which has previously been linked to economic growth (Atun, Harvey and Wild 2007).

- The (log) flow of new trademarks filed in Australia by Australian residents from IP Australia's IPLORD database. This is a slightly broader measure of innovation that will likely capture 'new-to-firm' innovation (Mendonça, Pereira and Godinho 2004; Claes 2005).

- Aggregate (log) R&D spending from Australian Bureau of Statistics (ABS) National Accounts. R&D has been linked to both higher novelty of innovation and adoption of innovation (Majeed and Breunig 2022; D'Este, Amara and Olmos-Peñuela 2016). This is a slightly broader measure of innovative activity that will capture spending on the creation of 'new-to-world' innovation, as well as spending to adopt innovation (e.g. Cohen and Levinthal 1989).

All three of these variables are observed quarterly from March 1994 through December 2019.

All three measures have limitations. Patents may be a poor measure of innovation for some industries. They can reflect anti-competitive behaviour rather than innovation (Argente et al., 2020). Trademarks can reflect innovation but are often filed by bad actors trying to usurp valuable property rights (von Graevenitz, 2013). R&D spending focuses on inputs rather than outputs. Innovation can be done without patenting or trademarking. There is an extensive literature canvassing these weaknesses. For a small number of examples, see Reeb and Zhao (2020), Nappi and Kelly (2021), Glaeser and Lang (2024), Stundziene et al. (2024) and references therein.

Our fourth measure of innovative activity is a survey measure of innovation collected in the ABS Business Characteristics Survey (BCS); see Australian Bureau of Statistics (2023). This is a broad measure of innovation based on the Oslo Manual (2018), the OECD benchmark for innovation measurement. It is widely used by statistical agencies around the world (OECD, 2018). Firms are asked whether they introduced new or significantly improved: goods or services; operational processes; organisational/managerial processes; or marketing methods.[3] This measure includes adoption of existing technologies or processes, which is an important mechanism in models such as that of Moran and Queralto (2018), but one that has not been explored empirically in relation to monetary policy shocks. Gault (2013) describes the extensive testing that has been conducted on this measure to confirm that it accurately captures firm innovative activity.

Adoption is particularly important in small open economies which tend to be importers of innovation. For example, in Australia 8 per cent of firms report introducing a new-to-world good or service and only 1.2 per cent report a new-to-world marketing innovation. The vast majority of firms (from 75 to 92 per cent depending on the type of innovation) who report innovative activity report that it is new to the business only. See Australian Bureau of Statistics (2016-17).

Previous Australian research found this self-reported innovation measure to be associated with firm productivity (Palangkaraya, Spurling and Webster, 2015). We find a positive association between current innovation and future productivity in our data, suggesting this self-reported measure captures real activity (Appendix Table A1). Nonetheless, there could be some systematic differences in how firms interpret and report innovation. For example, large firms may have more formalized innovation tracking systems, potentially creating systematic measurement differences across firm sizes. Benchimol, El-Shagi and Saadon (2022) highlight how respondent characteristics can influence forecasts among experts. Similarly, firms with different characteristics might have different approaches to forming expectations about future economic conditions that influence innovation decisions. We discuss this further in interpreting the results below.

The survey data we use come from the ABS' BCS, a census of firms with more than 300 employees and a stratified random sample for firms with less than 300 employees. It captures around 8,000 firms each year from 2005/06 to 2019/20. Response rates are very high as firms are legally required to respond to this survey which is conducted by the national statistics agency. Stratification by industry and business size is implemented to produce data that are representative of Australian businesses. The ABS do not provide sample weights so we use unweighted data. In our results across all firms, large firms will be somewhat overweighted. For this reason, and to explore heterogeneity, we provide pooled estimates and results split by firm size.

Firms with fewer than 300 employees are included on a rotating 5-year basis. This could create some biases in our results, particularly if small and large firms respond differently to shocks as the share of large firms in our sample will be larger as we consider longer time horizons. Taking any given year as a base period, only 2/5th of small firms will still be in the sample in 4 years' time. Splitting the sample into small and large firms helps to understand any potential bias created when we pool all firms together. As a robustness test, we also estimated the models including firms in the regression when there is a shock at time t and where we observe the firm at horizons $t-1$ and $t+3$. This produces a set of estimates which are unaffected by the rotating panel since we have removed small firms that are affected by shocks in the middle or the end of their spell in the data. These results are discussed in the robustness section below.

Another potential bias could come from attrition. If monetary policy shocks cause firms to exit, and if the remaining firms were more likely to be innovative, it may appear that the share of firms innovating has increased, but this apparent result would merely be reflecting a positive

[3] The exact wording of the question is included in Appendix A.

survivorship bias. To avoid this issue, our main regressions focus on firms with at least 5 observations. These are firms who do not exit during their period in the sample. Our firm-level results thus measure the intensive margin effect of monetary policy on firm innovation only. Our key results are unaffected if we include all firms. Moreover, we find that exit in response to a monetary policy shock is not related to prior innovative behavior or intentions. See section 4.3 for further discussion of these robustness checks. In keeping with standard practice, we exclude micro-businesses by removing all firms with one full time equivalent (FTE) employee or less.

The firm-level BCS data are linked to the ABS Business Longitudinal Analysis Data Environment (BLADE), which contain demographic and tax data from administrative sources. This allows us to model innovation at the firm level accounting for firm-level covariates and to explore heterogenous effects across firm types. Descriptive statistics for the firm-level sample are included in Appendix Tables A2 and A3. We provide descriptive statistics for all firms and for firms split by firm size, exporting status and foreign ownership.

1.3.2 Methodology

Monetary Policy Shocks

An inherent difficulty in examining the effect of monetary policy on innovation is that the official policy rate will be endogenous. That is, innovation activity and monetary policy are co-determined by other factors. For example, the central bank is likely to raise rates if it expects economic activity and inflation to increase. But improvements in economic activity might also spur further innovative activity. Thus, it might appear that higher interest rates lead to more innovation because both are moving with economic conditions.

To address this endogeneity issue, we use a monetary policy shock measure developed in Beckers (2020). This is a Romer and Romer (2004) (RR) style approach, in that it measures shocks as divergences of the observed policy rate from what would be expected based on an estimated policy reaction function. It builds on Caldara and Herbst (2019), who extend the RR approach to also account for responses to financial market variables. RR-style approaches are common in the literature, for example used in Cloyne and Hürtgen (2016), Tenreyro and Thwaites (2016) and Barnichon and Mesters (2020) and Holm, Paul and Tischbirek (2021).

Beckers (2020) estimates an augmented Taylor rule that includes a forecast for economic conditions and several indicators of domestic and foreign financial conditions (e.g. bond spreads, option-implied volatility). Shocks are constructed as the deviation of the actual policy rate from that implied by the rule. This approach removes the endogenous component of monetary policy by purging changes in the policy rate of the central bank's systematic response to its own forecasts for inflation and activity, which in turn should reflect all available information about domestic and foreign economic outcomes. This includes US monetary policy and other factors, e.g. expected activity, inflation, etc., that might drive firm innovation. We use the measure constructed by Beckers (2020).

This strategy will control for firms adjusting their innovation strategies in anticipation of monetary policy effects unless firms are reacting to things that affect both their innovation and monetary policy but that are not part of the central bank's forecast. This seems unlikely to us.[4]

The Beckers (2020) measure is constructed by estimating a forward-looking Taylor-type rule, augmented to account for responses to financial market variables, of the following form:

$$\Delta cr_t = \alpha + \beta_r cr_{t-1} + \boldsymbol{Y}^{fc}_{t+h|t}\beta_y + \boldsymbol{CS}_t\beta_{cs} + \varepsilon_t$$

Where Δcr_t is the change in the policy rate at the central bank meeting at time *t*, cr_{t-1} is the level at the prior meeting, $\boldsymbol{Y}^{fc}_{t+h|t}$ are the central bank's *h*-quarter ahead macroeconomic forecasts, and $\boldsymbol{CS}_t$ is a set of money or credit market spreads or other financial indicators available at the time. The residual ε_t is the monetary policy shock. Beckers (2020) preferred model includes the two-quarter-ahead forecasts for GDP growth, inflation, revisions to forecasts, the nowcast of unemployment, domestic money-market spreads, domestic business lending spreads, and US credit spreads. The measure purges market-anticipated changes in the policy rate at the meeting, though this has little impact.

We adopt this measure as our preferred shock measure because previous research found that it is able to overcome the price puzzle in Australian data: that contractionary monetary policy is often estimated to raise prices. This is not the case for the simpler RR-style measure that does not account for responses to financial market variables (Bishop and Tulip 2017). It is also not the case for measures based on high-frequency changes in bond yields over short windows around announcements (Hambur and Haque 2024). We consider these monetary policy shock measures and two other alternatives in section 1.4.3.

For the regressions with macroeconomic variables, we use quarterly observations of the shocks. For the annual firm-level regressions, we sum the shocks in the relevant year. See Figure 1.1.

[4] If we estimate the impact of monetary policy shocks at time *t* on innovation at time *t-2* or *t-3* we find no statistically significant impact. We look at cumulative change since period *t-1*, so we cannot use *t-1* to test anticipation. These results are available from the authors upon request.

Figure 1.1: Beckers Monetary Policy Shock Measure

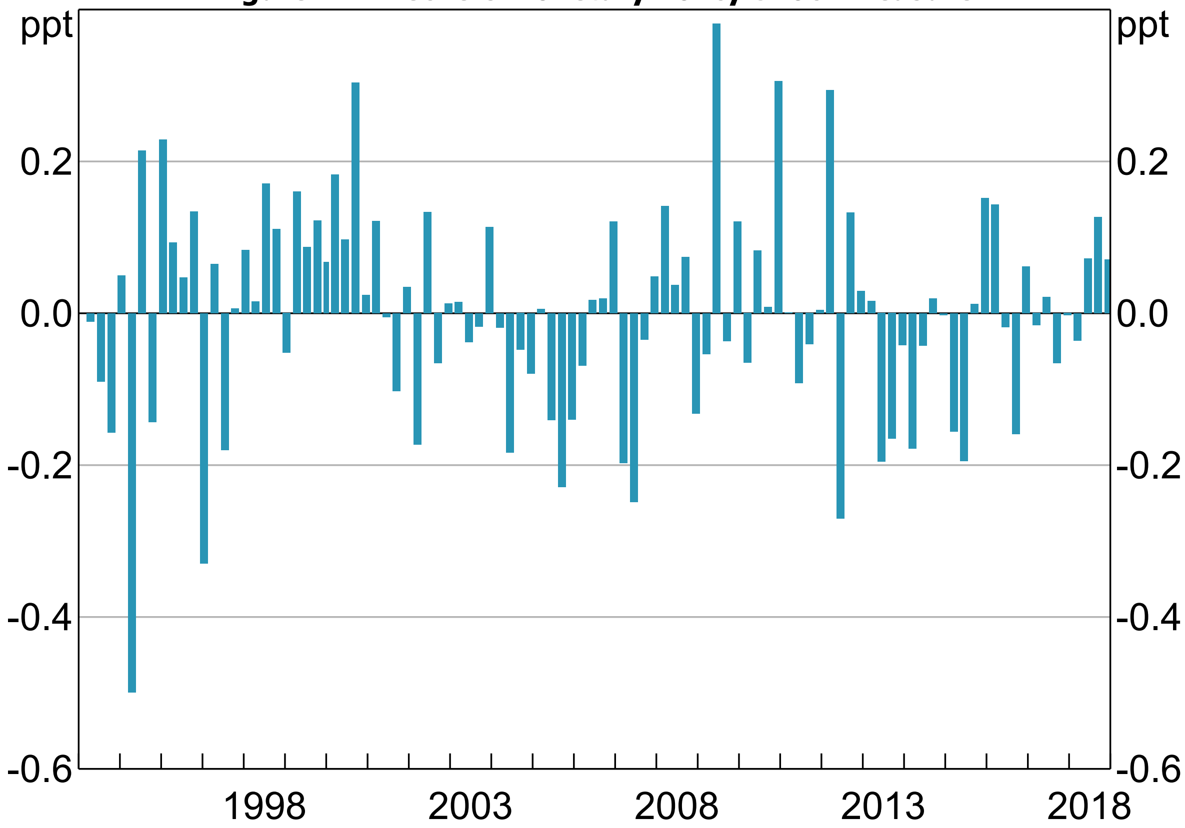

Note: Figure shows our preferred monetary policy shock measure estimated from equation (1). This is the "unanticipated BT-CS measure" from Beckers (2020). Quarterly measure.
Source: Beckers (2020)

Our paper, in using the shocks from Beckers (2020), applies a specific form of monetary policy rule, rather than considering what may be best. As such, we abstract from related and important issues around differences between the monetary policy rules that small open economies follow and the ones they should follow, as discussed in Benchimol (2024). That paper shows the important role for limited information in the transmission of monetary policy which may help to explain part of the heterogeneity in firm responses that we find.

Estimation

We employ a local projection regression (Jordà 2005) to trace out the effect of a monetary policy shock at time *t* on our measures of innovative activity over a number of different time horizons *h*. We estimate separate regressions for each time horizon *h*. This is a common approach in the literature (e.g. Ma 2023; Jordà, Singh and Taylor 2024). Our regressions take the form:

$$Inn_{i,t+h} = \beta_h shock_t + \sum_{j=1}^{J} \alpha_{h,j} Inn_{i,t-j} + \sum_{j=1}^{J} \gamma_{h,j} Controls_{i,t-j} + \sum_{j=1}^{J} \theta_{h,j} shock_{t-j} + v_{i,t+h}$$

where $h \geq 0$, denotes the time horizon, *t* denotes time (years in firm-level regressions and quarters in aggregate outcome regressions) and *i* indexes the firm (for firm-level regressions). We estimate the equation using aggregate variables which we can represent identically to the above, simply by dropping the *i* subscripts. $Inn_{i,t+h}$ is our measure of innovation. This is either (log) R&D,

(log) patents or (log) trademarks at the aggregate level; or, for the firm level, a 0/1 indicator of whether a firm innovates. $shock_t$ is our continuous measure of monetary policy shocks described above. Typically, local projections control for macroeconomic variables to improve efficiency. We control for standard aggregate measures: Gross Domestic Product (GDP), the Consumer Price Index (CPI) and the trade-weighted index (TWI) exchange rate. For our firm-level regressions we also control for firm-level variables: (log) employment, turnover growth, (log) capital expenditure, industry and age (a dummy if a firm is more than four years old). With the exception of age and industry, we do not include contemporaneous controls, only lagged variables. This is to avoid imposing the implicit assumption that monetary policy does not affect contemporaneous outcomes (Ramey 2016).

In our preferred specification of the macro-level regressions that we present below (in Figure 2, for example) we include 8 quarterly lags of the control variables, four lags of the shock variable and four lags of the dependent variable in the aggregate-level regressions. If we use more than one year of lagged shocks, the sample decreases significantly. If we use an automatic lag selection procedure for the control variables, such as minimising the Akaike Information Criteria (AIC), we get different preferences for different horizons of lags in different models and subsamples. As the results don't change much and the differences in AIC are small, we use two years in all the results we report.

For the firm-level regressions (see Table 1, for example), we include one annual lag of all the firm-level and macro controls. Age is included contemporaneously since it is uncorrelated with the shock. In the robustness checks, we provide results without any macroeconomic controls to show that this does not impact our results. We cluster standard errors at the period level, reflecting that the key variable of interest varies across time but not across firms. As a result, we have a small number of clusters, which can bias standard errors downwards. To address this, we assess significance based on a *t*-distribution with *T-n* degrees of freedom, where *T* is the sample length and *n* is the number of coefficients on variables that do not vary across firms, see Cameron and Miller (2015). We allow for serial correlation in the unobservables ($v_{i,t+h}$), though we still include the lagged dependent variable to allow for interpretation of the regressions as the change in probability of innovating relative to before the shock.[5]

1.4 Results

1.4.1 Aggregate innovation measures

In this section, we examine how monetary policy shocks affect patents, trademarks and R&D activity. We focus on the coefficients of the shock variable in equation (2) over different time horizons and trace them out in Figure 2. For consistency, all the results of this paper are based on a 100 basis point contractionary shock.

5 In the firm-level regression results that we report, we do not include a lag of the shock. The results are robust to including one lag of the shock, but we prefer the more parsimonious specification.

Figure 1.2: Effect of Monetary Policy Shock on Aggregate Innovation Measures

100 basis point shock

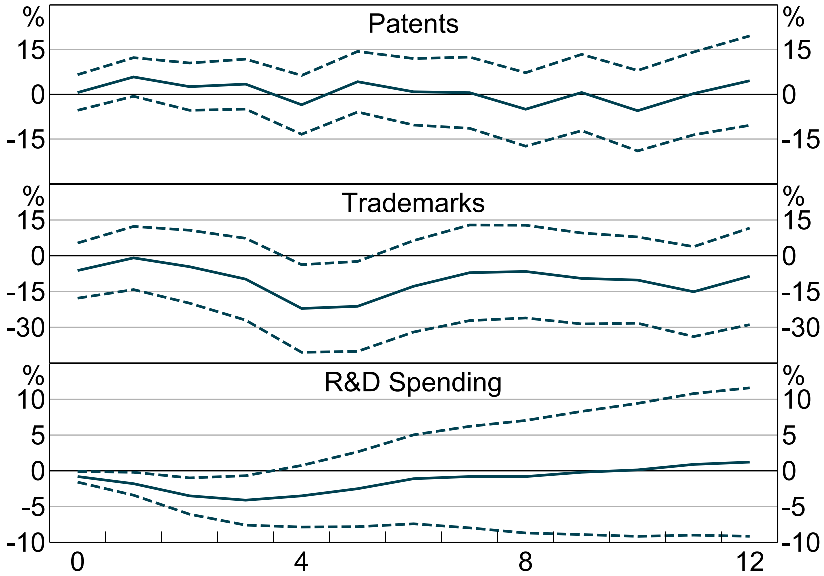

Note: Results from local projection model of aggregate innovation metrics on monetary policy shock. Results show percentage change in flow of patents, trademarks or R&D spending. Trademark sample excludes 1994-1995 due to break in time series. Patents and trademarks include only those with Australian filers. Dashed lines show 90 per cent confidence intervals, with Huber/White standard errors. Models have 8 quarterly lags of growth in GDP trade-weighted exchange rate index and CPI growth, and four lags of shocks and lagged dependent variable. Results robust to other specifications of lags. Data are quarterly.

Unlike the US (Ma 2023), monetary policy shocks have no effect on the number of patents filed (Figure 1.2; top panel). This is consistent with Australia being an importer of new technologies, rather than a producer and Australia having fewer patents filed compared to the US. This suggests that monetary policy may have less of an effect on the technology frontier in an adopting country and therefore a less persistent effect on productivity.

Nevertheless, domestic monetary policy has a significant influence on broader measures of innovation. In the year following a 100 basis point contractionary shock, R&D spending declines by almost 5 per cent (bottom panel), while the number of trademarks falls by over 15 per cent one-to-two years after the shock (middle panel).[6] The response of R&D spending is somewhat larger than documented in Moran and Queralto (2018), though it is also shorter lived and the peak effect is earlier – they document a 0.8 per cent decline, with the peak effect around 4 years after the shock.[7]

6 While the 5 per cent fall in R&D spending may seem large, it is not unreasonable in the context of this series – over the sample, the growth rate of R&D spending has ranged between around negative 8 per cent and 22 per cent, with a standard deviation of around 7.5 per cent. This unconditional standard deviation is much larger than the standard deviation of the structural R&D shock from the small VAR discussed later, suggesting much of the volatility in R&D spending reflects other factors such as aggregate demand. Moreover, as shown in Beckers, the identified shock is quite persistent, with a half-life of around 2 years.

7 Results are very similar if we instead estimate a simple 5-variable VAR with four lags where the shock is ordered first (see Appendix Figure B1), as would be expected given the equivalence results outlined in Plagborg-Møller and Wolf (2021). The effects are slightly longer-lived if we include only two lags, but this appears to reflect bias coming from truncation of the VAR order. We also experimented with a VAR identified using timing restrictions as in Moran and Queralto (2018). This provided no evidence of a significant effect on R&D spending, likely reflecting the strong (and probably inaccurate) identification assumption.

Our results are closer in magnitude to those of Ma and Zimmerman (2023), who also use local projection, and find that aggregate investment in intellectual property falls by around 1 per cent one to two years after the shock, and that listed company R&D spending falls by around 3 per cent, with a peak effect two to three years after the shock.

Overall, our effects are shorter lived compared to US findings. While it is difficult to draw too strong a conclusion due to the inherently different nature of the identified shocks (particularly their persistence), this suggests that domestic monetary policy in Australia has shorter-lived effects on innovation. This is not surprising if monetary policy is mainly influencing Australian firms' decisions around adopting and adapting overseas technologies. These investments may have shorter horizons compared to developing cutting-edge technologies, so policy has shorter-lived effects on associated spending. Similarly, R&D spending appears to have shorter-lived effects on productivity in Australia, compared to in the US (see section 6 below).

Notably, similar regressions using US monetary policy shocks show no significant effect on these narrower measures of innovation. This is in contrast to the impact of US monetary policy shocks on the broader measure of innovation and adoption explored below.

Firm-level innovation and adoption measures

As discussed above, much of the literature to date has focused on relatively narrow measures of innovation that are unlikely to fully capture the broader effects of monetary policy on adoption of existing technologies by firms. This is the case even though technology adoption is a crucial mechanism in models such as that of Moran and Queralto (2018). Moreover, it is likely to be the core channel through which monetary policy affects productivity and innovation in small open economies that tends to import new technologies.

Focusing on our broader survey measure of innovation we find that a contractionary 100 basis point monetary policy shock has relatively little effect on the overall share of firms innovating (Table 1.1, top panel). But this average result hides very different outcomes for small and large firms. The share of small and medium enterprises (SMEs; less than 200 employees) innovating falls by 6 percentage points the year after a 100 basis point contractionary monetary policy shock (Table 1.1, middle panel). This equates to around 52,000 fewer firms innovating.

In contrast, the share of larger firms (more than 200 employees) innovating increases by 5 percentage points two years after the shock, equating to around 200 firms (Table 1.1, last panel). This latter finding may seem surprising in the context of the existing literature, but it is consistent with some of the broader innovation literature that argues that innovation may be countercyclical as firms may have an incentive to undertake longer run investments in innovation when input costs are relatively cheap and when demand is weak (Aghion et al. 2012). Larger firms may be better placed to take advantage of countercyclical conditions.

We present the results from separate regressions for small and large firms. The differences between small and large firm responses that we discuss are statistically significant. We assess these statistical differences based on a single pooled regression with the effect of all control variables allowed to differ by sub-group.

For context, from 2005/06 to 2019/20 the share of firms innovating rose consistently, increasing by around 7 percentage points by the end of the sample (Department of Industry, Science, Energy and Resources 2024). Given this, monetary policy appears to have meaningful effects on the share of firms innovating.[8] Moreover, the heterogeneity suggests that firms may be differentially exposed to different monetary policy channels. We explore this in section 5.

Table 1.1: Effect of 100 basis point shock on share of firms innovating

By firm size

	Year 0	Year 1	Year 2	Year 3
All firms				
Effect	-0.76	-2.76*	1.29	7.06**
(s.e.)	(1.65)	(1.47)	(2.09)	(2.67)
R^2	*0.20*	*0.14*	*0.11*	*0.09*
Observations	*45,053*	*35,635*	*26,302*	*17,536*
SMEs				
Effect	-2.85	-6.13***	-2.45	1.39
(s.e.)	(1.61)	(0.99)	(1.83)	(1.54)
R^2	*0.19*	*0.14*	*0.12*	*0.10*
Observations	*29,551*	*22,185*	*14,616*	*7,291*
Large firms				
Effect	3.47	2.70	5.02*	9.80*
(s.e.)	(2.64)	(2.65)	(2.35)	(4.46)
R^2	*0.19*	*0.11*	*0.07*	*0.06*
Observations	*15,502*	*13,450*	*11,686*	*10,245*
Difference by size significant at	5 per cent level	1 per cent level	1 per cent level	10 per cent level

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggested by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. Regressions include age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth capital expenditure and employment, and one lag of the dependent variable.

Source: Authors' calculations

Overall, these results suggest that the effects of monetary policy on adoption may be more complex and heterogenous than the effects of more 'leading-edge' innovation, which previous work had consistently found was negatively affected by contractionary policy (Moran and Queralto 2018, Ma 2023; Ma and Zimmermann 2023). Given the important role that adoption plays in models in this literature, and in productivity growth more generally, this is an important finding.

Robustness

We test to see if our firm-level results differ substantially when we use other monetary policy shock measures. We consider four measures:

- The change in the policy rate (cash rate) itself.

8 The response may seem large relative to the variation in the number of innovating firms. But it is important to keep in mind that a 100 basis point monetary policy shock is larger than any shock which has occurred in the sample.

- A version of the Beckers (2020) shock that does not purge market expectations, which we call 'Beckers'.

- A simple Romer and Romer-style shock where the policy reaction function includes only economic variables and excludes the financial market variables used in Beckers (2020). This is taken from Bishop and Tulip (2017). We refer to this as 'BT'.

- A measure based on high-frequency changes in bond yields based on a 90-minute window around announcements (Hambur and Haque 2024), which we call 'levels shock'. This measure is orthogonalized by regressing it against economic data which are available at the time of the shock (labor force information, inflation, etc.). Failure to do this produces a `shock' variable that can be predicted from contemporaneous data, violating the basic principle of a shock being something that is not predictable.

The results are provided in Appendix Tables B1-B4. Using the Beckers variable gives qualitatively and quantitatively similar results to our preferred estimates from section 1.4.2, consistent with the high correlation between the two shock measures. The results with the BT variable are also qualitatively similar, with SMEs decreasing their innovation and large firms increasing it. Using the levels shock or the changes in the cash rate itself, we find that monetary policy has no statistically significant effect on innovation. The cash rate is a poor `shock' measure as it reflects both anticipated and unanticipated changes in monetary policy. The `levels' shock has significant shortcomings discussed in Hambur and Haque (2024) and continues to exhibit the price puzzle. [9] The Beckers shocks remain our preferred estimates.

Appendix Table B5 provides estimates where we only include firms in the regression when there is a shock at time *t* and where we observe the firm at horizons *t-1* and *t+3.* This produces a set of estimates which are unaffected by the rotating panel since we have removed small firms that are affected by shocks in the middle or the end of their spell in the data. Appendix Table B6 shows the results if we include firms that appear in fewer than 5 years in the data. Both tables show that our results are robust to our choice of sample restrictions and the differential sampling of small and large firms. Appendix Table B7 shows results excluding the macroeconomic controls. Inclusion or exclusion of macroeconomic controls does not affect our results. This provides further evidence that choice of lag length (including zero lags) and controls does not affect our results.

Sample selection does not appear to be driving our results, which focus on the intensive margin, and whether monetary policy affects innovation of surviving firms. But another margin could be firm exit, with contractionary monetary policy disproportionately causing innovative firms to exit. While this would not affect our intensive margin results, this is still a potentially interesting margin to explore. To directly test this, we estimate the probability of firm exit in response to a monetary policy shock controlling for observable firm characteristics and macroeconomic conditions. We test whether the exit probability differs by previous innovation activity and by whether the firm indicated that they were focused on innovation at the time of the shock. Mostly we find no differences between innovators and non-innovators (see Appendix Table B8). There is a small

[9] Even after orthogonalization, the `levels' shock continues to exhibit the price puzzle, making it unsuitable for this exercise. If we use the levels shocks without orthogonalization, we find a negative effect of monetary policy shocks on innovation. These results are available from the authors upon request. The cash rate itself exhibits the price puzzle.

difference at *t+1*, with innovators more likely to exit, but the difference is only statistically significant at the 10 per cent level. We find precise zeros for all other time horizons and for the difference between firms who are focused on innovation and those who are not. There does not appear to be an extensive exit margin effect of monetary policy on innovation in Australia.

1.5 Through which channels does monetary policy affect innovation?

As discussed in the literature review, there are several channels through which monetary policy and economic conditions could lower the amount of innovative activity in the economy. One is the demand channel: contractionary monetary policy may weaken aggregate demand and therefore lower the returns to innovation. Another is the credit constraint channel: monetary policy may lead to tighter credit conditions and make cash-flow constraints more binding for some firms. In turn this may lessen their ability to fund and undertake innovative activity.

It is difficult to perfectly separate these different channels and to precisely identify the effect of each one. We follow other papers (Jeenas, 2024; Ottonello and Winberry, 2020; Cloyne et al., 2023) in examining the importance of these channels by comparing outcomes for firms that we would expect to be more exposed to the channels to outcomes for firms that we would expect to be less exposed. This provides evidence that is suggestive of the roles for different channels rather than a sharp causal identification of these channels.

1.5.1 Demand channel

Aggregate demand in the economy is an important channel for monetary policy. As a contractionary monetary policy shock lowers aggregate demand in the economy, it also lowers the potential future profits of a firm. As the likelihood of future profits is the key reason for firms to innovate (Aghion and Howitt 2008), the probability of firms innovating can decrease as aggregate demand falls.

To consider the importance of the demand channel, we examine outcomes for exporting and non-exporting firms. Exporting firms are likely to be less exposed to domestic conditions, as their sales are not determined exclusively by the domestic economy. We interact our shock variable with a dummy variable that takes value one if a firm ever exported while in the sample and zero otherwise. We use this to trace out the effect of a shock separately for exporters and non-exporters.

Consistent with the demand channel playing an important role, the negative effect of contractionary monetary policy on innovation is less evident for exporters (Table 1.2). This pattern provides suggestive evidence that the demand channel may be a key mechanism through which monetary policy affects firm-level innovation. These results are similar for large firms and SMEs (Appendix Tables B9 and B10), though there is only a statistically significant difference between exporters and non-exporters for the large firms.[10] This suggests that the results are not simply a reflection of the fact that exporters tend to be larger.

[10] Small sample sizes (about 2,000 firms at horizon *t+3)* for exporting SMEs may impact our ability to detect differences for that group.

Table 1.2: Effect of 100 basis point shock on share of firms innovating

By exporter status

	Year 0	Year 1	Year 2	Year 3
Exporter				
Effect	0.24	1.60	4.29	8.83*
(s.e.)	(2.12)	(1.96)	(2.49)	(4.01)
R^2	*0.19*	*0.12*	*0.09*	*0.06*
Observations	*17,974*	*14,734*	*11,672*	*8,904*
Not exporter				
Effect	-1.19	-5.58***	-1.14	4.59**
(s.e.)	(1.78)	(1.43)	(1.92)	(1.93)
R^2	*0.19*	*0.12*	*0.10*	*0.10*
Observations	*27,079*	*20,901*	*14,630*	*8,632*
Difference by export status significant at:	-	1 per cent level	1 per cent level	-

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. Regressions include age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth, capital expenditure and employment, and one lag of the dependent variable.

Source: Authors' calculations

It is important to note that there could be other differences between exporting and non-exporting firms that affect these results. For example, exporting firms tend to be more productive and may have better management. To isolate the demand effects, we re-estimate the regression using a measure of US monetary policy shocks taken from Choi, Willems and Yoo (2024).[11] If the previous results reflected the demand channel, rather than other inherent differences between exporters and non-exporters, we might expect exporters to respond more strongly to overseas monetary policy shocks as they are more directly exposed to foreign demand than are non-exporting firms. This is the 'trade channel' of policy spill overs (Arbatli-Saxegaard et al. 2025). This exercise is also valuable as it provides (to our knowledge) the first evidence on spillovers from US monetary policy onto innovation and adoption in another country.

The first thing to note is that contractionary US monetary policy shocks appear to lower the share of firms innovating in Australia (Table 1.3). The effects appear larger than for domestic shocks. However, it is difficult to compare the magnitudes given the two shocks are inherently different: they are identified using different techniques in separate country contexts and differ both in terms of the nature of the policy shock and the Taylor rule of the local central bank.

[11] This measure is based on high-frequency changes in bond yields.

Table 1.3: Effect of 100 basis point US shock on share of firms innovating

By export status

	Year 0	Year 1	Year 2	Year 3
All firms				
Effect	-5.73	-11.02***	-8.16	-7.59
(s.e.)	(5.40)	(3.20)	(5.17)	(8.88)
R^2	*0.20*	*0.14*	*0.11*	*0.09*
Observations	*45,053*	*35,635*	*26,302*	*17,536*
Exporters				
Effect	-7.38	-10.86*	-9.76	-11.28
(s.e.)	(6.97)	(5.37)	(6.29)	(10.17)
R^2	*0.19*	*0.12*	*0.09*	*0.06*
Observations	*17,974*	*14,734*	*11,672*	*8,904*
Non-exporters				
Effect	-4.53	-10.61**	-6.71	-6.00
(s.e.)	(5.04)	(3.79)	(4.95)	(5.89)
R^2	*0.19*	*0.12*	*0.10*	*0.09*
Observations	*27,079*	*20,901*	*14,630*	*8,632*
Difference by exporter status significant at	-	-	-	-

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. Regressions include age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth, capital expenditure and employment, and one lag of the dependent variable.

Source: Authors' calculations

The second key finding is that the response for exporters appears larger than for non-exporters, which is the opposite of the findings when we considered a shock to Australian monetary policy. This is observed for both SMEs and large firms (Appendix Tables B11 and B12) and is consistent with other work focusing on the response of firm-level investment and sales in foreign countries following US policy shocks (Arbatli-Saxegaard et al. 2025). While the differences are not statistically significant, the greater sensitivity of exporters to foreign shocks, and of non-exporters to domestic shocks, provides further evidence of the importance of the demand channel.

One concern may be that we are still not directly isolating the demand channel. In particular, we may be picking up differential effects on credit constraints and the credit constraint channel (explored further below). For example, US monetary policy shocks may influence global capital markets, and this may have a larger effect on access to credit for exporters.

To account for this, we control for the change from *t-1* to *t+h* in firms' self-reported measure of whether lack of access to funds is hampering innovation (discussed more below). This will directly control for the effect that the shock had on credit constraints. The approach is similar to Holm, Paul and Tischbirek (2021) who decompose the effects of monetary policy on consumption into the direct and indirect (general equilibrium) channels by controlling for changes in income.

As evident in Table 1.4, the key results are unchanged. Exporters are more responsive to US shocks (if still not statistically significantly so), and non-exporters to domestic shocks. Overall, these results provide support for the importance of the demand channel.

Table 1.4: Effect of 100 basis point shocks on share of firms innovating

By exporter status, controlling for financing constraints

	Year 0	Year 1	Year 2	Year 3
Domestic shocks				
Exporter				
Effect	-0.65	-0.13	2.90	8.97*
(s.e.)	(2.19)	(1.90)	(2.23)	(4.24)
R^2	*0.20*	*0.13*	*0.10*	*0.07*
Observations	*16,450*	*13,290*	*10,380*	*7,680*
Not exporter				
Effect	-2.60	-6.81***	-2.46	2.68
(s.e.)	(1.92)	(1.46)	(1.38)	(1.97)
R^2	*0.20*	*0.14*	*0.12*	*0.10*
Observations	*25,600*	*29,450*	*13,410*	*7,550*
Difference by export status significant at:	-	1 per cent level	5 per cent level	-
US shocks				
	Year 0	Year 1	Year 2	Year 3
Exporter				
Effect	-10.71	-15.75***	-13.81**	-21.09*
(s.e.)	(5.97)	(3.74)	(5.12)	(9.53)
R^2	*0.20*	*0.13*	*0.10*	*0.07*
Observations	*16,450*	*13,290*	*10,380*	*7,680*
Not exporter				
Effect	-7.61	-11.47**	-9.94**	-14.40**
(s.e.)	(4.42)	(4.48)	(4.05)	(5.23)
R^2	*0.20*	*0.14*	*0.12*	*0.10*
Observations	*25,600*	*29,450*	*13,410*	*7,550*
Difference by export status significant at:	-	-	-	-

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. Regressions include age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth, capital expenditure and employment, and one lag of the dependent variable. Also include change in declared financial constraints from *t-1* to *t+h.*

Source: Authors' calculations

1.5.2 The credit constraint channel

The credit constraint channel of monetary policy captures any way in which contractionary monetary policy (or weaker economic conditions) makes it harder for a firm to finance innovative activity. This can reflect: tighter aggregate credit supply; lower asset prices which reduce borrowers' collateral value and therefore borrowing capacity – the 'financial accelerator channel' (Bernanke, Gertler and Gilchrist 1999); or lower liquidity due to lower revenue or higher interest payments making existing financing constraints more binding (e.g. Jeenas 2024). We are considering a broad concept of the credit constraint channel that includes more than simply the supply or availability of credit.

To examine the credit constraint channel directly, we explore an additional question in the BCS. This question asks whether a lack of access to additional funds is hampering the business's ability to innovate. We examine whether monetary policy shocks lead to a change in the share of firms reporting that access to additional funds is limiting their innovative activity.[12]

Consistent with the credit constraint channel being important, contractionary monetary policy shocks lead to an increase in the likelihood that firms report that lack of funds is significantly hampering their ability to undertake innovation (Table 1.5). This is almost entirely driven by SMEs, which is consistent with the evidence that SMEs have a larger decline in innovation following a monetary policy shock. To put the results in context, over the sample, around 17 per cent of SMEs note that a lack of funds is hampering their innovation (compared to 8 per cent for large firms). So, a 100 basis point shock (which is very large historically) would cause the share of SMEs reporting that a lack of funds are hampering innovation to increase by around 15-20 per cent. These findings are also consistent with the broader literature that finds that SMEs are more likely to be credit or cash constrained than larger firms — for example Mancusi and Vezzulli (2010) and Bakhtiari et al. (2020).

Cloyne et al. (2023) find that firms that are prone to financial constraints are more responsive to monetary policy. Our results suggest that contractionary policy influences innovation by tightening constraints and we find that small firms are more responsive. If small firms are most likely to face constraints, our results are in line Cloyne et al. (2023). Our results also align with Manso, Balsmeier and Fleming (2023) and Acemoglu et al. (2018) who examine the relationship between financial constraints and firms' innovation strategies. They emphasise the difference between radical and incremental innovation. The difference in response of small and large firms could be driven by their engaging in different types of innovation.

12 The survey question asks firms about what factors are hampering their innovative activity. As well as a lack of access to additional funds, other factors are: lack of skilled workers; cost of development; regulations and compliance; uncertain demand for new goods/services. Firms can select more than one factor.

Table 1.5: Effect of 100 basis point shock on share of firms reporting lack of funds hampering innovation

By firm size

	Year 0	Year 1	Year 2	Year 3
All firms				
Effect	1.40**	2.06***	1.44	2.37***
(s.e.)	(0.63)	(0.63)	(0.99)	(0.54)
R^2	*0.28*	*0.32*	*0.35*	*0.39*
Observations	*42,138*	*32,901*	*23,826*	*15,245*
SMEs				
Effect	2.02*	2.92***	2.45*	3.67***
(s.e.)	(1.00)	(0.70)	(1.30)	(0.65)
R^2	*0.28*	*0.32*	*0.34*	*0.37*
Observations	*28216*	*20,966*	*13,518*	*6,360*
Large firms				
Effect	0.16	0.63	0.25	1.10
(s.e.)	(0.49)	(1.41)	(1.19)	(0.83)
R^2	*0.28*	*0.34*	*0.39*	*0.42*
Observations	*13,922*	*11,935*	*10,308*	*8,885*
Difference by size significant at	-	-	10 per cent level	10 per cent level

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. Regressions include age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth, capital expenditure and employment, and one lag of the dependent variable.

Source: Authors' calculations

1.5.3 The cost channel

As discussed, another channel through which contractionary monetary policy could lead to increased innovation and adoption of technologies is by freeing up resources and lowering the cost of innovating (Aghion et al. 2012). To explore this, we can exploit survey questions around whether "the cost of developing or introducing innovation" is hampering innovation. This is an imperfect measure of the cost of innovation. Lower demand might reduce the benefit of innovating and firms may report this as an increase in the relative cost of innovation. Nevertheless, it does represent a somewhat direct way to try to identify the channel.

Table 1.6 shows that following a contractionary policy shock, the share of small firms reporting that costs are a barrier to innovation rises by around 4 percentage points. In contrast, there is no significant evidence of an increase for larger firms.

The results appear contradictory to the cost channel as small firms report *higher* costs of innovating post-monetary policy shock, not lower costs. If the cost channel were important, the Aghion et al. (2012) argument would suggest that the share of firms reporting costs as a barrier should decrease, not increase, following a contractionary shock. The fact that SMEs are more likely to say that costs are hampering their innovation following a contractionary shock may be reflecting lower demand or higher financing costs. The difference in outcomes between SMEs and large firms is interesting and reinforce our findings elsewhere that small and large firms respond quite differently to monetary policy shocks, but these results are probably explained by the lack of our ability to precisely identify the costs of innovation, per se. Overall we find no evidence in support of the cost channel.

Table 1.6: Effect of 100 basis point shock on share of firms reporting "the cost of developing or introducing innovation is hampering innovation"

By firm size

	Year 0	Year 1	Year 2	Year 3
All firms				
Effect	1.23	1.12	2.75*	2.36*
(s.e.)	(0.79)	(1.14)	(1.14)	(1.20)
R^2	*0.33*	*0.36*	*0.38*	*0.41*
Observations	*44,087*	*34,859*	*25,686*	*17,049*
SMEs				
Effect	1.22*	1.72	4.51**	3.75*
(s.e.)	(0.57)	(1.30)	(1.95)	(1.67)
R^2	*0.33*	*0.36*	*0.37*	*0.39*
Observations	*29,072*	*21,823*	*14,355*	*7,131*
Large firms				
Effect	1.24	0.46	0.38	0.96
(s.e.)	(1.75)	(2.14)	(1.38)	(1.35)
R^2	*0.32*	*0.38*	*0.40*	*0.42*
Observations	*15,015*	*13,029*	*11,331*	*9,918*
Difference by size significant at	-	-	5 per cent level	-

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. Regressions include age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth, capital expenditure and employment, and one lag of the dependent variable.
Source: Authors' calculations

1.6 Macroeconomic effects on productivity

Thus far we have focused on measures of innovation rather than productivity. Moran and Queralto (2018) demonstrate in their model and empirically that shocks to R&D spending have an effect on productivity in the medium-run. They show this focusing on the US and using a cross-country panel (that includes Australia). However, given the vastly different structures of the Australian and US economies, it is worth examining whether the results for the US also hold for Australia.

To examine if their results hold for Australia, we reproduce the small VAR model used in Moran and Queralto (2018). We estimate a small, three variable VAR with annual data on GDP, total-factor productivity from Bergeaud, Cette and Lecat (2016), and R&D spending from 1988 to 2019. As in Moran and Queralto (2018) we examine the effect of R&D shocks, which are identified by ordering R&D last in a Cholesky decomposition. This implies that R&D cannot affect TFP contemporaneously, consistent with the fact that it generally takes time for R&D expenditure to result in new technologies. While this identification assumption may be strong, it is the same as used in Moran and Queralto (2018) and it should highlight whether similar economic patterns and mechanisms are evident in the US and Australia.

Taking this approach, we find that an R&D shock leads to a persistent increase in TFP that peaks around 5 years after the shock (Figure 1.3). The magnitudes of the responses are somewhat larger than Moran and Queralto (2018) report, with an equivalent size increase in R&D (4 per cent) leading to a 1.6 per cent increase in TFP, compared to 0.4 per cent increase in Moran and Queralto (2018). In the Australian data, the volatility of the R&D shock is an order of magnitude smaller (around ½ per cent, compared to 4 per cent in Moran and Queralto (2018)), suggesting that such a large shock would be extremely unusual.[13] Our results are somewhat more in line with Ma and Zimmerman (2023). Drawing on other estimates, they suggest that a 100 basis point shock would lower TFP and output by 0.5 to 1 per cent, whereas our estimates put this at around 1.6 per cent.

Interestingly, the response is also less long-lived compared to Moran and Queralto (2018): the peak occurs around 5 years after the shock for Australia, compared with around 8 years in the US. One possible explanation for this difference is that Australia imports innovation and that its R&D spending is more targeted at adapting overseas innovation. If so, and if the investment horizon on adaption is shorter than on developing frontier technologies, declines in R&D spending may be expected to have shorter-lived effects on productivity in Australia. Moreover, while adoption of technologies declines, the actual stock of global knowledge is unaffected by lower R&D spending in Australia, allowing firms to catch up more quickly.

This is speculative and there are many differences between the US and Australia. Future work could examine whether shorter-lived dynamics feature in other small, technology importing countries.

[13] We experimented with looking directly at the effects of monetary policy shocks on productivity. We were unable to identify any effects. This could be due to the lack of quarterly data on total-factor productivity, resulting in our estimates being based on a small number of annual observations and highly aggregated shocks.

Figure 1.3: Effect of an R&D Shock on Total Factor Productivity
Aggregate data, one standard deviation R&D shock

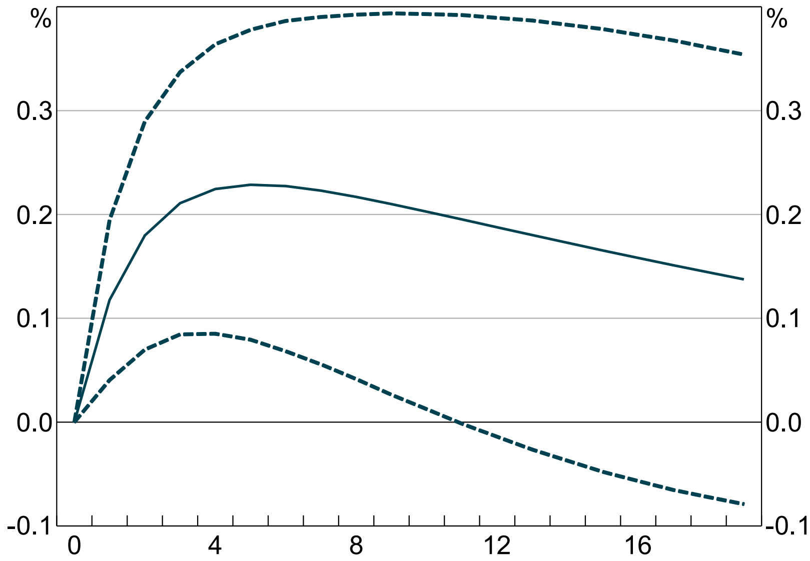

Note: Figure shows per cent response of TFP to an R&D spending shock from a VAR containing the log levels of real GDP, real R&D spending, and TFP. Response based on Cholesky decomposition with R&D spending ordered last. VAR(1) model. Dashed lines are the 90 per cent confidence intervals. Data are annual.

1.7 Conclusion

There is a small but growing literature arguing that macroeconomic conditions and monetary policy shocks can have medium-run effects on productivity and therefore living standards and long-run economic activity. This has important implications for macroeconomic policy. While such effects are likely to cancel out over a cycle, the potential for medium-run productivity scarring increases the importance of macro-stabilisation policy. It also increases the costs associated with such policies being constrained, such as at the effective zero lower bound on interest rates. And it potentially alters the trade-offs a central bank faces when stabilising inflation and activity in the face of supply shocks, provided that inflation expectations can be kept anchored.

We provide the first evidence on the effects of monetary policy on innovation for a small open economy that imports innovation rather than creating it. Consistent with the overseas literature, we find evidence for Australia that contractionary monetary policy shocks are associated with a decline in R&D spending, and that declines in R&D spending are associated with lower levels of productivity. However, the effects appear to be shorter-lived compared to the US.

We also provide the first evidence on the effects of monetary policy on broader measures of innovation that better capture adoption of existing technologies. Such adoption has been shown to be important empirically for small open economies and is a key channel in many macroeconomic models. We find substantial heterogeneity in the response of firms. SMEs are less likely to innovate and adopt after a contractionary monetary policy shock, but large firms are more likely to do so. This suggests that effects of policy on the adoption of innovation may be more complex and heterogeneous than the effects of monetary policy on patenting and R&D spending.

These results can potentially be explained by SMEs and large firms being differentially exposed to the channels through which monetary policy can affect innovative activity (see Table). Exporting firms (which tend to be larger) react less to contractionary monetary policy shocks than domestically-focused firms. SMEs are also more likely to have their innovation hampered by a lack of funds following contractionary monetary policy shocks, providing evidence for the financial constraint channel of monetary policy.

Table 1.7: Table of Channels

Channel	Evidence
Demand channel	Yes: exporters more responsive to US shocks, non-exporters to domestic shocks
Credit constraint channel	Yes: contractionary shocks raise the share of firms reporting constraints. More notable for small firms
Cost channel	No evidence

Still, there could be other reasons that small and large firms respond differently. They may have different approaches to forming expectations about future economic conditions that influence innovation decisions, akin to the effect of different characteristics on forecasting documents in Benchimol, El-Shagi and Saadon (2022). Limited information, shown to be important in Benchimol (2024), may impact the transmission of monetary policy differently for small and large firms. Small and large firms may also interpret and report innovation differently.

Small and large firms might be doing different types of innovation and may thus respond differently to monetary policy shocks. Manso, Balsmeier and Fleming (2023) and Acemoglu et al. (2018) showed that firms pursuing different innovation strategies (radical versus incremental) respond differently to financial constraints.

Our findings on the credit constraint channel are somewhat consistent with Cloyne et al. (2023) who find that firms that are prone to financial constraints are more responsive to monetary policy. Our results suggest that contractionary policy influences innovation by tightening constraints, and we might expect that those firms who are prone to constraints are likely to be most affected. This contrasts with Ottonello and Winberry (2020) who argue that less risky firms are more responsive to monetary policy (if we equate riskiness with likelihood of credit constraints binding).

A novel aspect of our paper is considering the effects of US monetary policy shocks on the innovative activity of Australian firms. We find important spillovers from US monetary policy shocks onto the innovative activity of Australian firms. This effect is larger for Australian firms who export, reinforcing the importance of the demand channel of monetary policy.

Overall, our results suggest that domestic and foreign monetary policy can affect innovative activity in a small open economy. However, they paint a more complex picture relative to the existing literature, with significant heterogeneity in effects by firm size and export activity and less persistent impacts. If our results can be generalized to small, open economies where adoption of innovation is important, the effects of monetary policy and economic fluctuations on innovation and productivity may be smaller compared to a frontier economy, lessening, though not removing the scope for medium-term scarring.

Appendix A: Additional summary statistics and data description

A.1 Innovation metric

Our survey measure of innovation is based on the definition from the Oslo Manual (2018):

An innovation is a new or improved product or process (or combination thereof) that differs significantly from the unit's previous products or processes and that has been made available to potential users (product) or brought into use by the unit (process).

This measure is gathered every second year in the Business Characteristics Survey conducted by the Australian Bureau of Statistics. Firms are first instructed that `` A **new good or service** means any **good or service** or combination of these which is **new to this business**. Its characteristics or intended uses differ significantly from those previously produced/offered by this business." They are then asked if they developed a new good, a new service or none of the above. (Firms can choose both Good and Service.) [Bold is reproduced from the questionnaire.]

Firms are then asked "Were any of the **new or significantly improved** goods or services introduced by this business:

- New to the world?
- New to Australia but not new to the world?
- New to the industry within Australia, but not new to Australia or the world?
- New to this business only (i.e. none of the above)"

Firms are then asked an identical question about `new or significantly improved processes' (split into `operational' and `organisational/managerial') and `marketing methods'.

This sequencing of questions is promoted by the Oslo Manual as being the best way to obtain this information from firms. A recent version of the questionnaire is available at: https://consult.abs.gov.au/industry-statistics/business-characteristics-surveys-consultation/supporting_documents/1BCC1%2020222023%20%20Sample%20only.pdf

To demonstrate the link between this self-reported measure of innovation and productivity, we use the same local projections framework described in the paper at the firm level, but tracing out the effect of innovating on firms' future productivity, measured as the ratio of value-added to full-time equivalent employees. While not causal, we do see that innovating is associated with stronger productivity growth, particularly for smaller firms where variation between innovating and not innovating year-to-year is much more common. This reinforces findings in previous work in Australia (Palangkaraya et al., 2015 and Majeed and Breunig, 2022) and in Europe (van Leeuwen, 2002) that this innovation measure is linked to firm performance and productivity.

Table A1 shows the correlation between self-reported innovation at time t and future firm-level productivity (at time $t+h$), controlling for age, industry, (lag) GDP growth, (lag) inflation, (lag) growth in the exchange rate, (lag) turnover growth, (lag) capital expenditure and (lag) employment, and one lag of the dependent variable.

Table A1: Effect of innovating at time *t* on future firm productivity

By firm size

	Year 0	Year 1	Year 2	Year 3
All firms				
Effect	0.44	1.46**	1.19*	1.17***
(s.e.)	(0.37)	(0.59)	(0.55)	(0.42)
R^2	0.73	0.67	0.62	0.59
Observations	84,060	76,900	72,170	64,510
SMEs				
Effect	0.62	2.72***	2.29**	1.75***
(s.e.)	(0.46)	(0.55)	(0.53)	(0.58)
R^2	0.7	0.64	0.58	0.55
Observations	63,530	56,750	52,460	46,170
Large firms				
Effect	-0.90	-2.07	-1.92	-0.53
(s.e.)	(0.64)	(1.14)	(1.18)	(0.86)
R^2	*0.19*	*0.11*	*0.07*	*0.06*
Observations	*15,502*	*13,450*	*11,686*	*10,245*

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. All regressions include controls for industry, (lag) GDP growth, (lag) inflation, (lag) growth in the exchange rate, (lag) turnover growth, (lag) capital expenditure) and (lag) employment, and lag of the dependent variable.

Source: Authors' calculations

A.2 Summary statistics

Table A2: Summary Statistics for Firm-Level Data

Full sample

Variables	Mean	Median	Share
Employment (FTE)	293	14	
Sales growth %	6	3	
Age	25	18	
Share innovating %			52
Share of firms >5 years old %			81
Share of firm finance access hampers innovation %			15
Share of SME %			78
Total observations '000s			103

Source: ABS; Authors' calculations

Table A3: Summary Statistics for Firm-Level Data

	Mean	Median	Share	Mean	Median	Share
		SMEs			**Large firms**	
Employment (FTE)	22	8		1,246	569	
Sales growth %	6	3		8	4	
Age	21	15		44	31	
Share innovating %			48			63
Share of firms >5 years old %			78			90
Share of firm finance access hampers innovation %			17			8
Total observations '000s			80			20
		Non-exporters			**Exporters**	
Employment (FTE)	143	9		624	79	
Sales growth %	7	3		5	3	
Age	22	15		34	24	
Share innovating %			47			62
Share of firms >5 years old %			79			89
Share of firm finance access hampers innovation %			16			14
Total observations '000s			70			32
		Domestically owned			**Foreign owned**	
Employment (FTE)	136	10		906	355	
Sales growth %	6	3		6	4	
Age	23	16		35	23	
Share innovating %			50			60
Share of firms >5 years old %			81			85
Share of firm finance access hampers innovation %			17			10
Total observations '000s			81			21

Source: ABS; Authors' calculations

Appendix B: Additional regression results

Figure B1: Effect of Monetary Policy Shock on R&D Spending in VAR Model

100 basis point shock, quarters

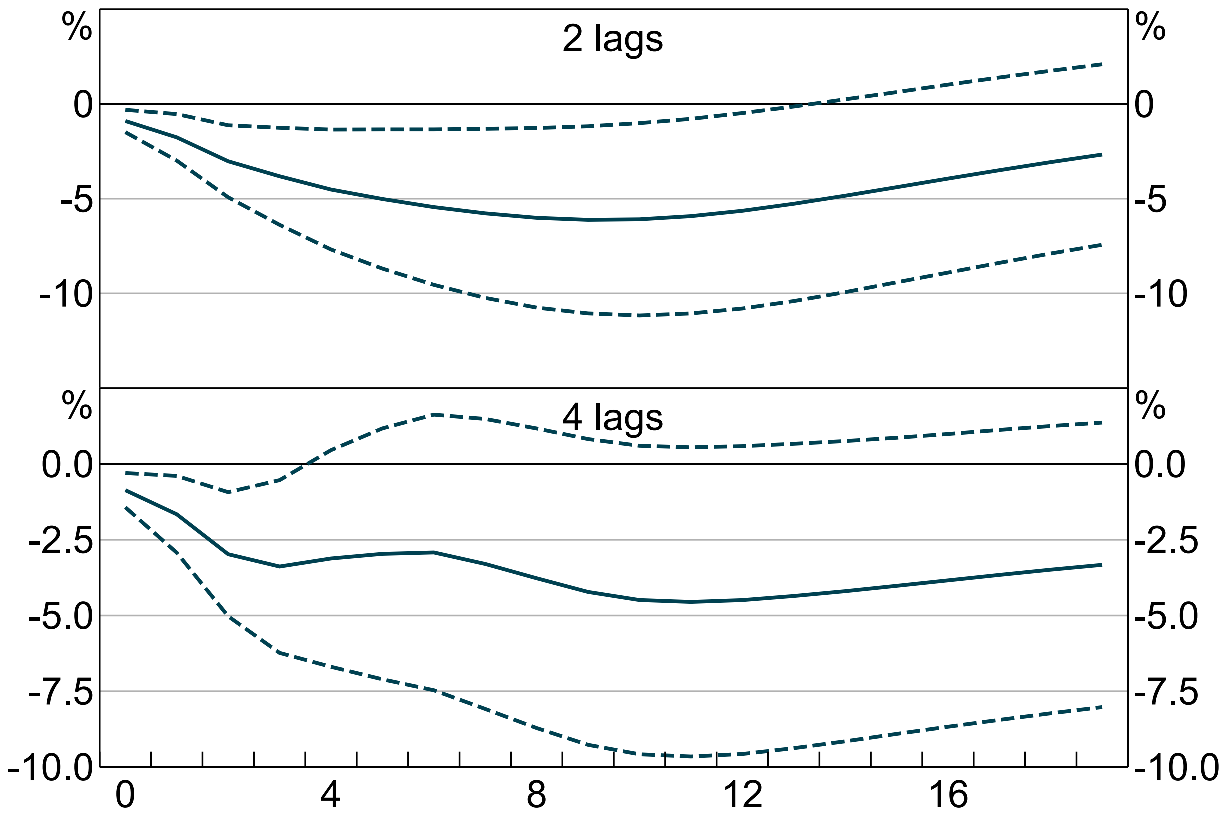

Note: Figure shows per cent response of R&D spending to monetary policy shock in a VAR containing the log levels of the CPI, real GDP, real R&D spending, the trade-weighted exchange rate index, the level of the cash rate, and the Beckers (2020) shock. Response based on a Cholesky decomposition with the shock measure ordered first. Top panel is VAR(2) model. Bottom panel is VAR(4) model. Dashed lines are 90 per cent confidence intervals.

Table B1: Effect of 100 basis point shock on share of firms innovating

Robustness check of Table 1 using alternative measure of monetary policy shock:
Cash Rate Change
All firms and by firm size

	Year 0	Year 1	Year 2	Year 3
All firms				
Effect	-0.49	0.28	0.25	0.64
(s.e.)	(0.32)	(0.38)	(0.28)	(0.62)
R^2	0.20	0.14	0.11	0.09
Observations	45,050	35,640	26,300	17,540
SMEs				
Effect	-0.45	0.14	0.04	-0.03
(s.e.)	(0.41)	(0.49)	(0.36)	(0.86)
R^2	0.19	0.13	0.12	0.1
Observations	29,550	22,190	14,620	7,290
Large firms				
Effect	-0.51	0.50	0.45	1.15
(s.e.)	(0.43)	(0.52)	(0.4)	(0.82)
R^2	0.19	0.11	0.07	0.06
Observations	15,500	13,450	11,690	10,250

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. All regressions include controls for age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth, capital expenditure and employment, and one lag of the dependent variable.

Source: Authors' calculations

Table B2: Effect of 100 basis point shock on share of firms innovating

Robustness check of Table 1 using alternative measure of monetary policy shock:
Beckers shock without purging market expectations
All firms and by firm size

	Year 0	Year 1	Year 2	Year 3
All firms				
Effect	-4.38**	-3.52**	-0.03	5.51**
(s.e.)	(1.49)	(1.55)	(2.06)	(2.26)
R^2	*0.05*	*0.05*	*0.05*	*0.04*
Observations	*45,151*	*35,705*	*26,342*	*17,545*
SMEs				
Effect	-5.60***	-6.44***	-3.87*	-0.93
(s.e.)	(1.44)	(1.13)	(1.84)	(1.67)
R^2	*0.05*	*0.05*	*0.05*	*0.05*
Observations	*29,648*	*22,255*	*14,656*	*7,300*
Large firms				
Effect	-1.37	1.01	3.84	10.48**
(s.e.)	(3.01)	(2.65)	(2.64)	(3.99)
R^2	*0.03*	*0.02*	*0.03*	*0.03*
Observations	*15,503*	*13,450*	*11,686*	*10,245*

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. Regressions include age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth, capital expenditure and employment, and one lag of the dependent variable.

Source: Authors' calculations

Table B3: Effect of 100 basis point shock on share of firms innovating

Robustness check of Table 1 using alternative measure of monetary policy shock:
Bishop and Tulip (2017) shock (see section 4.3 of the main paper)
All firms and by firm size

	Year 0	Year 1	Year 2	Year 3
All firms				
Effect	-1.75	-0.83	0.19	2.28
(s.e.)	(1.09)	(1.40)	(0.97)	(1.69)
R^2	*0.05*	*0.05*	*0.05*	*0.04*
Observations	*45151*	*35,705*	*26,342*	*17,545*
SMEs				
Effect	-2.51	-2.16	-1.51	-2.10
(s.e.)	(1.50)	(1.45)	(1.22)	(1.43)
R^2	*0.05*	*0.05*	*0.05*	*0.05*
Observations	*29,648*	*22,255*	*14,656*	*7,300*
Large firms				
Effect	-0.16	0.47	2.02*	4.94*
(s.e.)	1.18	1.90	1.09	2.46
R^2	*0.03*	*0.02*	*0.03*	*0.03*
Observations	*15,503*	*13,450*	*11,686*	*10,245*

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. Regressions include age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth, capital expenditure and employment, and one lag of the dependent variable.

Source: Authors' calculations

Table B4: Effect of 100 basis point shock on share of firms innovating
Robustness check of Table 1 using alternative measure of monetary policy shock:
Levels shock (see section 4.3 of main paper)
All firms and by firm size

	Year 0	Year 1	Year 2	Year 3
All firms				
Effect	0.08	-0.15	-0.09	-0.60
(s.e.)	(0.4)	(0.41)	(0.27)	(0.56)
R^2	0.05	0.05	0.05	0.04
Observations	45,050	35,640	26,300	17,540
SMEs				
Effect	-0.18	-0.05	-0.05	0.16
(s.e.)	(0.52)	(0.49)	(0.34)	(0.75)
R^2	0.05	0.13	0.12	0.1
Observations	29,650	22,260	14,660	7,300
Large firms				
Effect	0.50	-0.30	-0.10	-1.00
(s.e.)	(0.41)	(0.51)	(0.42)	(0.78)
R^2	0.03	0.02	0.03	0.03
Observations	15,500	13,450	11,690	10,250

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. Regressions include age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth, capital expenditure and employment, and one lag of the dependent variable.
Source: Authors' calculations

Table B5: Effect of 100 basis point shock on share of firms innovating

Robustness check of Table 1: keeping only firms observed at time *t-1* and *t+3* for shocks at time *t*

All firms and by firm size

	Year 0	Year 1	Year 2	Year 3
All firms				
Effect	0.86	-0.26	2.34	7.06**
(s.e.)	(2.52)	(1.49)	(2.36)	(2.64)
R^2	*0.21*	*0.13*	*0.11*	*0.09*
Observations	*17,528*	*16,770*	*16,694*	*17,536*
SMEs				
Effect	-8.52**	-8.16***	-3.31	1.39
(s.e.)	(0.96)	(0.88)	(2.89)	(1.54)
R^2	*0.19*	*0.13*	*0.13*	*0.10*
Observations	*7,284*	*7,078*	*7,054*	*7,291*
Large firms				
Effect	6.48**	4.65	5.58**	9.80*
(s.e.)	(2.84)	(2.80)	(2.15)	(4.45)
R^2	*0.20*	*0.10*	*0.07*	*0.57*
Observations	*10,244*	*9,692*	*9,640*	*10,245*

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. Regressions include age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth, capital expenditure and employment, and one lag of the dependent variable. We only include firms in the regression when there is a shock at time *t* and where we observe the firm at horizons *t-1* and *t+3.* This produces a set of estimates which are unaffected by the rotating panel since we have removed small firms that are affected by shocks in the middle or the end of their spell in the data.

Source: Authors' calculations

Table B6: Effect of 100 basis point shock on share of firms innovating

Robustness check of Table 1: including firms with fewer than 5 observations

All firms and by firm size

	Year 0	Year 1	Year 2	Year 3
All firms				
Effect	-0.75	-3.07*	1.28	7.04**
(s.e.)	(1.89)	(1.50)	(1.95)	(2.74)
R^2	*0.20*	*0.14*	*0.11*	*0.09*
Observations	*58,686*	*41,610*	*28,296*	*17,761*
SMEs				
Effect	-2.23	-6.11***	-2.23	1.33
(s.e.)	(1.99)	(1.26)	(1.50)	(1.64)
R^2	*0.19*	*0.13*	*0.12*	*0.10*
Observations	*41,842*	*27,602*	*16,464*	*7,485*
Large firms				
Effect	2.98	2.49	5.09*	9.91*
(s.e.)	(2.55)	(2.75)	(2.47)	(4.42)
R^2	*0.19*	*0.10*	*0.08*	*0.06*
Observations	*16,844*	*14,008*	*11,832*	*10,276*

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. Regressions include age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth, capital expenditure and employment, and one lag of the dependent variable. No requirement to report for 5 years is imposed.

Source: Authors' calculations

Table B7: Effect of 100 basis point shock on share of firms innovating

Robustness check of Table 1: model re-estimated dropping all macroeconomic control variables

All firms and by firm size

	Year 0	Year 1	Year 2	Year 3
All firms				
Effect	0.14	-1.57	2.72	6.51**
(s.e.)	(1.54)	(1.45)	(3.05)	(2.7)
R^2	0.2	0.13	0.11	0.08
Observations	*45,053*	*35,635*	*26,302*	*17,536*
SMEs				
Effect	-2.68	-5.21***	-0.47	0.68
(s.e.)	(1.89)	(0.91)	(2.73)	(2.25)
R^2	0.19	0.14	0.12	0.1
Observations	*29,551*	*22,185*	*14,616*	*7,291*
Large firms				
Effect	5.56	4.41	6.8	9.77**
(s.e.)	(3.29)	(2.79)	(4.23)	(3.86)
R^2	0.19	0.1	0.07	0.05
Observations	*15,502*	*13,450*	*11,686*	*10,245*

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. Regressions include age and lagged controls for industry, turnover growth, capital expenditure and employment, and one lag of the dependent variable.

Source: Authors' calculations

Table B8: Effect of 100 basis point shock on firm exit probability

By innovation status, difference in effect by innovation status

	Year 0	Year 1	Year 2	Year 3
Ever innovated relative to never				
Effect	-0.04	1.04*	-0.10	-0.26
(s.e.)	(0.37)	(0.52)	(0.30)	(0.43)
R^2	*0.00*	*0.03*	*0.00*	*0.00*
Observations	*84,440*	*73,220*	*64,080*	*55,560*
Ever focused on innovation relative to never				
Effect	-0.041	0.10	-0.33	0.29
(s.e.)	(0.38)	(0.11)	(1.29)	(0.92)
R^2	*0.00*	*0.00*	*0.00*	*0.00*
Observations	*81,610*	*70,680*	*61,820*	*53,450*

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. Regressions from binary indicator of future exit on shocks. Include age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth, capital expenditure and employment, and one lag of the dependent variable. Coefficients show difference between groups.

Source: Authors' calculations

Table B9: Effect of 100 basis point shock on share of firms innovating

Robustness check of Table 2: Estimated only using SMEs

By exporter status

	Year 0	Year 1	Year 2	Year 3
Exporter				
Effect	-3.84*	-3.29	-0.16	4.07
(s.e.)	(1.98)	(1.94)	(3.92)	(3.26)
R^2	*0.19*	*0.14*	*0.13*	*0.10*
Observations	*7,733*	*5,788*	*3,847*	*2,009*
Not exporter				
Effect	-2.24	-6.90***	-2.85	0.99
(s.e.)	(1.60)	(1.21)	(1.88)	(1.34)
R^2	*0.19*	*0.13*	*0.10*	*0.09*
Observations	*21,818*	*16,397*	*10,769*	*5,282*
Difference by exporter status significant at	-	-	-	-

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. Regressions include age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth, capital expenditure and employment, and one lag of the dependent variable.

Source: Authors' calculations

Table B10: Effect of 100 basis point shock on share of firms innovating

Robustness check of Table 2: Estimated only using Large firms

By exporter status

	Year 0	Year 1	Year 2	Year 3
Exporter				
Effect	3.47	4.76	5.91*	9.73*
(s.e.)	(2.76)	(3.06)	(2.62)	(4.92)
R^2	*0.20*	*0.11*	*0.08*	*0.06*
Observations	*10,241*	*8,946*	*7,825*	*6,895*
Not exporter				
Effect	3.65	-1.49	3.12	9.79*
(s.e.)	(2.98)	(3.30)	(2.63)	(4.52)
R^2	*0.18*	*0.09*	*0.07*	*0.06*
Observations	*5,261*	*4,504*	*3,861*	*3,350*
Difference by exporter status significant at	-	10 per cent level	-	-

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. Regressions include age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth, capital expenditure and employment, and one lag of the dependent variable.

Source: Authors' calculations

Table B11: Effect of 100 basis point US shock on share of firms innovating

Robustness check of Table 3: Estimated only using SMEs

By export status

	Year 0	Year 1	Year 2	Year 3
All firms				
Effect	-7.00	-10.33**	-9.85*	-21.19
(s.e.)	(4.52)	(3.80)	(5.22)	(16.05)
R^2	*0.19*	*0.14*	*0.12*	*0.10*
Observations	*29,551*	*22,185*	*14,616*	*7,291*
Exporters				
Effect	-11.44*	-11.76	-18.81**	-16.62
(s.e.)	(5.61)	(6.85)	(6.92)	(19.25)
R^2	*0.19*	*0.14*	*0.13*	*0.10*
Observations	*7,733*	*5,788*	*3,847*	*2,009*
Non-exporters				
Effect	-5.48	-9.59**	-6.93	-25.36
(s.e.)	(4.59)	(4.58)	(5.41)	(18.18)
R^2	*0.19*	*0.12*	*0.10*	*0.09*
Observations	*21,818*	*16,397*	*10,769*	*5,282*
Difference by exporter status significant at	-	-	-	-

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with t-n degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where t is sample length and n is number of coefficients. Regressions include age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth, capital expenditure and employment, and one lag of the dependent variable.

Source: Authors' calculations

Table B12: Effect of 100 basis point US shock on share of firms innovating

Robustness check of Table 3: Estimated only using Large firms

By export status

	Year 0	Year 1	Year 2	Year 3
All firms				
Effect	-3.66	-11.49*	-5.59	-9.62
(s.e.)	(9.19)	(6.05)	(5.71)	(8.97)
R^2	0.19	0.11	0.07	0.06
Observations	15,502	13,450	11,686	10,245
Exporters				
Effect	-5.36	-10.66	-5.90	-12.94
(s.e.)	(9.22)	(6.44)	(7.54)	(10.87)
R^2	*0.20*	*0.11*	*0.08*	*0.06*
Observations	*10,241*	*8,946*	*7,825*	*6,895*
Non-exporters				
Effect	-0.14	-14.17	-5.39	-2.67
(s.e.)	(10.20)	(7.93)	(4.51)	(7.54)
R^2	*0.18*	*0.09*	*0.07*	*0.06*
Observations	*5,261*	*4,504*	*3,861*	*3,350*
Difference by exporter status significant at	-	-	-	-

Note: Standard errors in parentheses, ***P < 0.01; **P < 0.05; *P < 0.1. Standard errors clustered at an annual level. Significance assessed using T-distribution with *t-n* degrees of freedom as suggests by Cameron and Miller (2015) to account for small number of clusters, where *t* is sample length and *n* is number of coefficients. Regressions include age and lagged controls for industry, GDP growth, inflation, growth in the exchange rate, turnover growth capital, expenditure and employment, and one lag of the dependent variable.

Source: Authors' calculations

References

Acemoglu D, U Akcigit, H Alp, N Bloom and W Kerr (2018), "Innovation, reallocation, and growth", *American Economic Review*, 108(11), pp 3450-3491.

Aghion P, and PW Howitt (2008), *The Economics of Growth.* The MIT press.

Aghion P, N Bermna, L Eymard, P Askenazy and G Cette (2012), "Credit constraints and the cyclicality of R&D investment: Evidence from France", *Journal of the European Economic Association*, 10(5), pp 1001-1024.

Amador S (2022), 'Hysteresis, Endogenous Growth, and Monetary Policy', Working Paper No. 348, University of California Davis Department of Economics.

Andrews D, N Deutscher, J Hambur and D Hansell (2020), 'The Career Effects of Labour Market Conditions at Entry', Treasury Working Paper No. 2020-01, Australian Government: The Treasury,

Anzoátegui D, D Comin, M Gertler and J Martinez (2019), 'Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence', *American Economic Journal: Macroeconomics*, 11(3), pp 67-110.

Arbatli-Saxegaard EC, M First, D Furceri and J Verrier (2025), 'U.S. Monetary Policy Shock Spillovers: Evidence from Firm-level Data', *Journal of International Money and Finance,* 158, 103402.

Argente D, S Baslandze, D Hanley and S Moreira (2020), 'Patents to Products: Innovation, Product Creation, and Firm Dynamics', Federal Reserve Bank of Atlanta Working Paper No. 2020-4,

Atun R, I Harvey and J Wild (2007), 'Innovation, Patents and Economic Growth', *International Journal of Innovation Management,* 11(2), pp 279-297.

Australian Bureau of Statistics (2017), 'Innovation in Australian Business', ABS Website https://www.abs.gov.au/statistics/industry/technology-and-innovation/innovation-australian-business/2016-17. Accessed 11 July 2025.

Australian Bureau of Statistics (2023), 'Business Characteristics Survey', ABS Website https://www.abs.gov.au/participate-survey/business-survey/business-characteristics-survey. Accessed 11 July 2025.

Bakhtiari S, R Breunig, L Magnani and J Zhang (2020), 'Financial Constraints and Small and Medium Enterprises: A Review', *Economic Record*, 96(315), pp 506–523.

Ball LM (2009), 'Hysteresis in Unemployment: Old and New Evidence', National Bureau of Economic Research Working Paper No. 14818, March.

Baqaee D, E Farhi and K Sangani (2024), 'The Supply-side Effects of Monetary Policy', *Journal of Political Economy,* 132(4), pp 1065-1112.

Barlevy G (2004), 'The Cost of Business Cycles Under Endogenous Growth', *American Economic Review*, 94(4), pp 964-990. https://doi.org/10.1257/000282042002615.

Barlevy G (2007), 'On the Cyclicality of Research and Development', *American Economic Review*, 97(4) pp 1131-1164.

Barnichon P and G Mesters (2020), 'Identifying Modern Macro Equations with Old Shocks', *The Quarterly Journal of Economics*, 135(4), pp 2255-2298.

Beckers B (2020), 'Credit Spreads, Monetary Policy and the Price Puzzle' RBA Research Discussion Paper 2020-01.

Benchimol J (2024), **'**Central bank objectives, monetary policy rules, and limited information', *Journal of Macroeconomics*, 80, 103604.

Benchimol J, M El-Shagi and Y Saadon (2022), 'Do expert experience and characteristics affect inflation forecasts?', *Journal of Economic Behavior & Organization*, 201, pp 205-226.

Benigno G, and L Fornaro (2018), 'Stagnation Traps', *The Review of Economic Studies*, 85(3) pp 1425-1470.

Bergeaud A, G Cette and R Lecat (2016), 'Productivity Trends in Advanced Countries between 1890 and 2012', *Review of Income and Wealth*, 62(3), pp 420–444.

Bernanke BS, M Gertler and S Gilchrist (1999), 'The financial accelerator in a quantitative business cycle framework', in: J. B. Taylor and M. Woodford (ed.), Handbook of Macroeconomics, edition 1, volume 1, chapter 21, pages 1341-1393, Elsevier.

Bianchi F, H Kung and G Morales (2019), 'Growth, slowdowns, and recoveries', *Journal of Monetary Economics*, 101 pp 47-63.

Bishop J and P Tulip (2017), 'Anticipatory Monetary Policy and the 'Price Puzzle'.' RBA Research Discussion Paper 2017-02.

Blanchard OJ and LH Summers (1986), 'Hysteresis and the European Unemployment Problem', *NBER Macroeconomics Annual*, 1, pp 15-78.

Brown JR, SM Fazzari and BC Petersen (2009), 'Financing Innovation and Growth: Cash Flow, External Equity, and the 1990s R&D Boom', *The Journal of Finance*, 64(1), pp 151-185.

Bruno V and HS Shin (2015), 'Capital flows and the risk-taking channel of monetary policy', *Journal of Monetary Economics*, 71, pp 119-132.

Caldara D and E Herbst (2019), 'Monetary Policy, Real Activity, and Credit Spreads: Evidence from Bayesian Proxy SVARs', *American Economic Journal: Macroeconomics*, 11(1) pp 157-192.

Cameron AC and DL Miller (2015), 'A Practitioner's Guide to Cluster-Robust Inference', *The Journal of Human Resources*, 50(2), pp 317-372.

Choi S T Willems and SY Yoo (2024), 'Revisiting the Monetary Transmission Mechanism Through an Industry-Level Differential Approach', *Journal of Monetary Economics*, 145, p 2-17.

Cloyne J and P Hürtgen (2016), 'The macroeconomic effects of monetary policy: a new measure for the United Kingdom', *American Economic Journal: Macroeconomics,* 8(4), pp 75-102.

Cloyne J, C Ferreira, M Froemel and P Surico (2023), 'Monetary Policy, Corporate Finance, and Investment.' *Journal of the European Economic Association,* 21(6), pp 2586-2634.

Cohen WM and DA Levinthal (1989), 'Innovation and learning: the two faces of R&D', *The Economic Journal*, 99(397), pp 569-596.

Comin D and M Gertler (2006), 'Medium-term Business Cycles', *American Economic Review*, 96(3), pp 523-551.

Comin D and M Mestieri (2018), 'If technology has arrived everywhere, why has income diverged?' *American Economic Journal: Macroeconomics*, 10(3), pp 137-178.

Davis SJ and JC Haltiwanger (2024), 'Dynamism diminished: The role of housing markets and credit conditions', *American Economic Journal: Macroeconomics*, 16(2), pp29-61.

Department of Industry, Science, Energy and Resources (2024), 'Australian Innovation Statistics'*,* 18 October 2024. Canberra: Australian Innovation Statistics. https://www.industry.gov.au/publications/australian-innovation-statistics.

D'Este P, N Amara and J Olmos-Peñuela (2016), 'Fostering Novelty While Reducing Failure: Balancing the Twin Challenges of Product Innovation.' *Technological Forecasting and Social Change,* 113, pp 280–292.

Durante E, A Ferrando and P Vermeulen (2022), 'Monetary policy, investment and firm heterogeneity.' *European Economic Review,* 148(104251).

Garga V and SR Singh (2021), 'Output Hystersis and Optimal Monetary Policy.' *Journal of Monetary Economics*, 117, pp 871-886.

Gault F (2013), 'The Oslo Manual,' chapter 2 in: Fred Gault (ed.), *Handbook of Innovation Indicators and Measurement*, pages 41-59. Edward Elgar Publishing.

Georgiadis G (2016), 'Determinants of global spilloers from US monetary policy.' *Journal of Internal Money and Finance*, 67 ,pp 41-61.

Glaeser S and M Lang (2024), 'Measuring innovation and navigating its unique information issues: A review of the accounting literature on innovation.' *Journal of Accounting and Economics*, 78(2-3), 101720.

Hambur J and Q Haque (2024), 'Can We Use High-Frequency Yield Data to Better Understand the Effects of Monetary Policy and Its Communication? Yes and No!' *Economic Record*, 100(328), pp 3-43.

Holm MB, P Paul and A Tischbirek (2021), 'The Transmission of Monetary Policy Under the Microscope.' *Journal of Political Economy*, 129(10), 28612904.

Huber K (2018), 'Disentangling the Effects of a Banking Crisis: Evidence from German Firm and Countries.' *American Economic Review*, 103(3), pp 868-898.

Iacoviello M and G Navarro (2019), 'Foreign effects of higher U.S. interest rates.' *Journal of International Money and Finance*, 95, pp 232-250.

Ikeda D and T Kurozumi (2019), 'Slow Post-financial Crisis Recovery and Monetary Policy.' *American Economic Journal: Macroeconomics,* 11(4), pp 82-112.

IP Australia (2023), 'Intellectual Property Longitudinal Research Data (IPLORD).' Last modified August 30, 2023. Accessed December 19, 2023. https://aut.ac.nz.libguides.com/c.php?g=685064&p=4840642.

Jeenas P (2024), 'Firm Balance Sheet Liquidity, Monetary Policy Shocks, and Investment Dynamics', Barcelona School of Economics Working Paper 1409, October 2023 (revised May 2024).

Jordà O (2005), 'Estimation and Inference of Impulse Responses by Local Projections.' *American Economic Review,* 95(1), pp 161-82.

Jordà O SR Singh and AM Taylor (2024), 'The Long-Run Effects of Monetary Policy', *The Review of Economics and Statistics,* doi: https://doi.org/10.1162/rest_a_01527.

Kearns J, A Schrimpf and FD Xia (2022), 'Explaining Monetary Policy Spillovers: The Matrix Reloaded', Research Discussion Paper RDP 2019-03, Reserve Bank of Australia.

Ma C (2023), 'Firm liquidity and the innovation channel of monetary policy.' Available at SSRN: https://doi.org/10.2139/ssrn.4555571.

Ma Y and K Zimmermann (2023), 'Monetary Policy and Innovation', National Bureau of Economic Research Working Paper No. 31698, September.

Majeed O and R Breunig (2022), 'Determinants of innovation novelty: evidence from Australian administrative data.' *Economics of Innovation and New Technology*, 32(8), pp 1249-1273.

Majeed O, A Balaguer, D Hansell, L Hendrickson, A Latcham and T Satherley (2021), 'What Drives High Growth? Characteristics of Australian Firms.' *Economic Record,* 97(318), pp 350-364.

Malmberg C (2005), 'Trademark Statistics as Innovation Indicator? – A Micro Study', Paper in Innovation Studies No. 2005/17, Lund University. https://portal.research.lu.se/en/publications/trademark-statistics-as-innovation-indicators-a-micro-study.

Mancusi ML and A Vezzulli (2010), 'R&D, Innovation and Liquidity Constraints', KITeS Working Paper No. 030, ITeS, Centre for Knowledge, Internationalization and Technology Studies, Universita' Bocconi, Milano, Italy, revised May 2010.

Manso G, B Balsmeier and L Fleming (2023), 'Heterogeneous innovation over the business cycle.' *Review of Economics and Statistics*, 105(5), pp 1224-1236.

Mendonça S, TS Pereira and MM Godinho (2004), 'Trademarks as an indicator of innovation and industrial change.' *Research Policy,* 33(9), pp 1385-1404.

Moran P and A Queralto (2018), 'Innovation, productivity, and monetary policy.' *Journal of Monetary Economics*, 93, pp 24-41.

Nappi V and K Kelly (2021), 'Measuring knowledge management in the innovation process: a systematic literature review', *International Journal of Knowledge Management Studies,* 12(2), pp 161-182.

Nolan G, J Hambur and P Vermeulen (2025), 'How does Monetary Policy Affect Business Investment? Evidence from Australia.' CAMA Working Paper 9/2025.

OECD (2015), 'The Future of Productivity'. OECD Publishing. Available at: OECD-2015-The-future-of-productivity-book.pdf.

OECD (2018), 'Oslo Manual 2018: Guidelines for Collecting and Interpreting Innovation Data', Fourth Edition. Paris: OECD Publishing. In series *The Measurement of Scientific, Technological and Innovation Activities*. Available at : https://www.oecd.org/science/oslo-manual-2018-9789264304604-en.htm.

Ottonello P and T Winberry (2020), 'Financial Heterogeneity and the Investment Channel of Monetary Policy', *Econometrica*, 88(6), pp 2473–2502.

Ouyang M (2009), 'The Scarring Effect of Recessions.' *Journal of Monetary Economics*, 56(2), pp 184-199.

Ouyang M (2011), 'On the Cyclicality of R&D.' *The Review of Economics and Statistics*, 93(2), pp 542–553.

Palangkaraya Am T Spurling and E Webster (2015), 'Does Innovation Make (SME) Firms More Prodcutive', in Angus Moore and John Simon (Eds.), *RBA Annual Conferencce 2015 – Small Business Conditions and Finance*.

Plagborg-Møller M and CK Wolf (2021), 'Local Projections and VARs Estimate the Same Impulse Responses.' *Econometrica* 89(2) pp 955-980.

Queralto A (2020), 'A model of slow recoveries from financial crises.' *Journal of Monetary Economics,* 114, pp 1-25.

Ramey VA (2016), 'Macroeconomic Shocks and Their Propagation', Chapter 2 in (eds) JB Taylor and H Uhlig, *Handbook of Macroeconomics,* Vol. 2, pp 71-162.

Reeb DM and W Zhao (2020), 'Patents Do Not Measure Innovation Success.' *Critical Finance Review* 9(1-2), pp 157-199.

Romer CD and DH Romer (2004), 'A New Measure of Monetary Shocks: Derivation and Implications.' *The American Economic Review,* 94(4), pp 1055–1084.

Sedláček P and M Ignaszak (2021), 'Productivity, Profitability and Growth', CEPR Discussion Paper No. 16205.

Sedláček P and V Sterk (2017), 'The Growth Potential of Startups over the Business Cycle.' *American Economic Review*, 107(10), pp 3182-3210.

Stadler GW (1990), 'Business Cycle Models with Endogenous Technology.' *The American Economic Review*, 80(4), pp 763-778.

Stundziene A, V Pilinkiene, M Vilkas, A Grybauskas and M Lukauskas (2024), 'The challenge of measuring innovation types: A systematic literature review.' *Journal of Innovation & Knowledge*, 9(4), 100620

Tenreyro S and G Thwaites (2016), 'Pushing on a string: US monetary policy is less powerful in recessions.' *American Economic Journal: Macroeconomics* 8(4), pp 43-74.

van Leeuwen G (2002), 'Linking Innovation to Productivity Growth Using Two Waves of the Community Innovation Survey.' *OECD Science, Technology and Industry Working Papers*, No. 2002/08, OECD Publishing, Paris.

von Graevenitz G (2013), 'Trademark cluttering—evidence from EU enlargment.' *Oxford Economic Papers*, 721, pp 721-745.

Webster E and P Jensen (2009), 'Macroeconomic Conditions and Successful Commercialization.' *Cambridge Journal of Economics,* 35(1), pp 123-44.

2. Chapter 2: Adoption of Emerging Digital General-purpose Technologies: Determinants and Effects

Jonathan Hambur* and Kim Nguyen**

*Reserve Bank of Australia and Australian National University
** Moody's Analytics

Abstract

This paper examines whether having Board members with relevant experience and skills is associated with more profitable adoption of cloud computing and artificial intelligence/machine learning technologies. Using text-based measures of adoption constructed from Australian listed companies' reports and event study techniques, we show that firms with Board members who have relevant technological backgrounds are more likely to adopt profitably. These same firms also hire more skilled workers around adoption, suggesting they are making larger investments in complementary skills and processes to facilitate ethe profitable adoption. These findings highlight the critical role of management capabilities in the profitable adoption of emerging digital GPTs.

JEL Classification Numbers: J2, O32, O33

Keywords: technology adoption, productivity, management capability, skills

Thanks to John Simon, Callan Windsor, Flavio Calvino, Min Zhu and seminar participants at the RBA, Australian Conference of Economists, Curtin University, Swinburne University, University of Queensland and the OECD. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Reserve Bank of Australia. The authors are responsible for any errors.

2.1 Introduction

Emerging digital general-purpose technologies (GPT) such as cloud computing and artificial intelligence/machine learning (AI/ML) have the potential to revolutionise economies and drive faster growth in productivity and living standards. However, there remains significant uncertainty as to how quickly these benefits will accrue. As with many GPTs before, the key determinant won't just be how quickly firms adopt the technologies. Rather, it will be how quickly they learn to use them effectively, make complementary investments, and shape their broader production processes to fully leverage the technologies (e.g. Agrawal, Gans and Goldfarb 2023; Brynjolfsson, Rock and Syverson 2019). For example, Acemoglu (2025) finds that the benefits of AI will be only modest over the next decade if there is only shallow adoption of these technologies in the form of task augmentation.

The observation that not all technology adoption is equal is also mirrored in the emerging empirical literature on whether adoption of these digital GPTs raises profitability and productivity. For example, while some studies have documented that adoption of cloud computing technology can be associated with higher profitability (Zheng *et al* 2023; Jin and Bai 2022), others have found it to be associated with losses and ultimately abandonment of the technology (Makhlouf 2020; McMillan *et al* 2022; West 2022). Similarly, some papers have found that adoption of AI/ML is associated increased profitability (Babina *et al* 2024), while other have found negative or nil effects (McElheran *et al* 2025). Notably, Calvino and Fontanelli (2026) and Brynjolfsson, Jin and McElheran (2021) both find that adoption is only associated with gains when accompanied by other investments.

To better understand how quickly the benefits of these technologies are likely to accrue, as well as how businesses and policymakers can speed this process, we need to have a clearer picture of what determines whether a firm can profitably and productively adopt these technologies. One factors that has received very little focus to date is management capabilities. This is despite the fact that management quality has long been shown to be an important determinant of firm profitability and productivity, innovation, and absorbative capability (e.g. Scur at al 2021; Bloom et al 2019). And the fact that management capability is likely to be key in any firm decisions around complementary investments and changing processes to best leverage digital GPT.

This paper looks to fill this notable gap in the literature by examining the role of management skills in the profitable adoption of technology. To do so we construct a new database capturing adoption of cloud computing and AI/ML, two emerging digital GPT, using text from earnings calls and annual reports of Australian listed firms from the early 2010s to 2022. This allows us to identify the timing of adoption and trace out its effects. We combine this with data on Board members' skills to understand the impact of management skills on profitable adoption. And finally, we also merge in data on hiring to better understand the mechanisms though which management skills may facilitate profitable adoption, as well as the interaction between worker and management skills more generally. In combining these three sources, which as far as we are aware we are the first to do, we can gather a rich understanding of the importance of management skills in profitable adoption of digital GPT, the role of complementary investments, and the effects of adoption more generally.

Australia represents an interesting country in which to explore these issues. It is a small open economy that tends to import technology rather than being at the frontier. And it ranks in the middle of the pack in terms of developed country management capability rankings (Bloom, Sadun and Van Reenen 2017). As such, it is likely to be representative of what diffusion looks like in a broad range of technology-importing countries, and what barriers exist in terms of the availability of highly skilled managers and workers. This compares to most studies to date that focus on larger, technology-developing countries such as the US and Western Europe.

To preview the results, we find that firms with directors with strong technical backgrounds (experience with these GPTs) are far more likely to adopt the technologies and to increase their profitability post-adoption. This contrasts with other firms, where there is no evidence that adoption increases profitability. Moreover, firms with directors with relevant technical backgrounds are also more likely to hire workers with GPT-related skills around adoption. This may indicate that the adoptions involve greater complementary investments, including in adapting internal tools. The fact that such investments are more likely to lead to increased profitability is consistent with theoretical literature, and recent empirical evidence in Calvino and Fontanelli (2026). Our findings importantly link this to management capabilities, highlighting that they appear to be important in facilitating and motivating deeper and more transformative.

These findings have several implications for firms and our view on the productivity benefits of AI. First, they reinforce existing findings that not all adoption of digital GPTs is the same, and that complementary investments and adjustments may be needed for firms to leverage them. Second, they highlight that having management, and workers, with the right technical skills may be important in this process. This suggest that effective diffusion of the technologies could be uneven and slow. It also highlights an important message for policymakers trying to speed the process of productive adoption: investments in education and training to related to digital GPTs needs to consider management capabilities, not just worker capabilities.

The rest of the paper is laid out as follows. We start with a brief review of the literature in Section 2.2, before discussing our data in Section 2.3. Section 2.4 outlines the methodology, before section 2.5 shows our results. Section 2.6 then concludes.

2.2 Literature

Our paper builds on several streams of literature. Most closely related to our paper is the small strand of the literature focusing on the importance of management skills in adoption of emerging digital GPT technologies. McEhleran *et al* (2024) explores how adoption differs based on owner characteristics, finding that AI use was more prevalent amongst firms with more educated owners in the US. Similarly, Nicoletti, von Rueden and Andrews (2020) finds that adoption is lower in countries and industries in Europe where managerial experience is lower. Most closely related to our work is Alekseeva *et al* (2020), who identify adoption of technology using job advertisements. They find that firms that post advertisements for managers are more likely to see firm growth and investment. We take a different approach to identifying adoption and management skills that allows us to better isolate adoption from hiring, which we show do not always go together.

Our paper also relates to the literature on the role of worker skills in adoption. For example, Babina *et al* (2023) finds evidence that firms with more highly-educated and STEM workers invest more in AI in the US, and that increasing AI investment tends to be associated with increased use of high-skilled staff. Similarly, Harrigan Reshef and Toubal (2021) find that the prevalence of 'techies' in a firm is an important determinant in adoption and within industry firm polarization. And Fonatnelli *et al* (2025) find that French AI adopters with few IT-skilled staff experience more volatile productivity post adoption. In terms of hiring, Calvino and Fontanelli (2023b) find AI adoption for French firms is positively associated with investments in human capital. We similarly find evidence that adoption is associated with increased demand for skilled workers, but that this is somewhat dependent on the nature of management skills.

Also relevant are recent papers examining the effects of adoption on productivity, profitability and sales. For example, Babina *et al* (2024) find evidence that US firms that invest in AI have strong growth in sales and market valuations, particularly for larger firms. Czarnitzki, Fenandez and Rammer (2023) find a positive association between AI use and productivity in German survey data, while Acemoglu et al (2023) argues that similair associations in US data are driven by selection. And Zheng et al (2023) and Jin and Bai (2022) both document that firm profitability tends to rise following adoption of cloud technologies in the US, whereas Destefano, Kneller and Timmis (2025) find such evidence mainly for new firms.

Several papers have also highlighted the conditions for such gains. McElheran *et al* (2025) find that industrial AI use tends to lead to lower profits in the short run in the US, consistent with a J-curve of costly adjustment before profits. This is particularly notable for older businesses who often struggle to maintain certain practices related to monitoring performance indicators. Most relevant for some of our findings, Calvino and Fontanelli (2026) find that, adoption is only associated with higher productivity for firms that are developing their own AI solutions, not taking off-the-shelf ones. And Brynjolfsson et al (2021) find that adopting predictive technologies only raises revenue for those firms who make investments in IT capital, educated workers, or have well designed workplaces. All three highlight the importance of complementary investments and skills if firms are to reap the benefits of adoption. We build on such studies by identifying management skills as an important predictor of gains, and for complementary investments.

Our paper also relates to the broader literature documenting the pace of adoption of digital GPT. This includes a few papers for the US, including Bloom *et al* (2021) and Alekseeva *et al* (2021), as well a number for European countries (e.g. Calvino and Fontanelli 2023a; Calvino *et al* 2022). We expand the literature to explore a small-open technology-importing country in the form of Australia.

Finally, our paper also touches on the literature on the importance of management quality in firm productivity, profitability and innovation. As well as the more general work linking management quality to productivity (e.g. Scur at al 2021; Bloom et al 2019), management quality has been shown to have implications for innovation and adoption. For example, Bloom, Sadun and Van Reene (2021) show that firms with better management gain more productivity benefits from productivity investments. Similarly, with Nguyen (Forthcoming) shows that more trusting CEOs tend to promote more innovation, and Hsieh et al (2022) shows that companies with Boards with STEM skills tend to be more innovative. He latter two papers show that these benefits extend to the upper echelons of management, similair to our findings.

2.3 Data and summary statistics

2.3.1 Adoption data

Measuring adoption directly is challenging. Various sources of data have been explored in the literature, including surveys (McMillan *et al* 2022; Calvino and Fontanelli 2023; Czarnitzki, Fenandez and Rammer 2023) job advertisements (Alekseeva *et al* 2021; Acemoglu, *et al* 2022), patents (da Sillva Marioni, Rincon-Aznar and Venturini 2024), and information on firms' websites (Calvino *et al* 2022). Each of these sources has advantages and disadvantages. For example, surveys and websites can be representative of the broader business population, but often have limited longitudinal information. Hiring and patent data may have more longitudinal information, but potentially only capture a narrow and indirect slice of the adoption process.

We take an alternative approach by using the text of annual reports and earnings calls of listed firms, obtained via Refinitiv. This provides a potentially broad measure of adoption with good longitudinal information, though the sample will necessarily be skewed towards larger firms. Similar approaches have been applied in the literature to measure other metrics such as firms' climate change exposures (Sautner *et al* 2023) and firms' debt covenant exposures (Nguyen 2022).

We look for mentions of words related to each of the technologies of interest in the text. The list of words we use is taken from Bloom *et al* (2021), who identify word pairings, or 'bigrams', that appear to capture technological advancements that have influenced businesses, based on patents and US firms' earnings calls. They then group these bigrams into technology clusters. While they identify several clusters, we focus on two for this paper: AI/ML and cloud computing. See Appendix A for a list of the bigrams contained in each technology cluster.

For the paper we focus on these two technologies for two reasons. First, both represent emerging GPT, in that they are 'generic' technologies that have the potential to become pervasive and transform business processes, and that it will take time from initial introduction for these technologies to mature and for complementary investments to be effectively utilised (Etro 2010; Bayrak, Conley and Wilkie 2011; Crafts 2021). Second, these technologies are inherently interlinked, with increased computing power provided by cloud computing helping to fuel AI/ML, and many cloud technologies incorporating AI/ML (Loucks 2018).

We identify usage of one of these technologies based on the inclusion of one of these bigrams in the text, and identify adoption as the first reference to the technology. Rather than simply taking any mentions of the terms as evidence of technology use, we apply several extra filters to ensure that we are identifying adoption or use of the technology. First, we remove pages with three or more proper nouns, which we identify using in-built parts of speech identification functions in Python's Natural Language Toolkit. This allows us to remove Board member biography pages, which often have references to a director's experience.[14] We also require the annual report or earnings call to have at least two references to the technology. Having applied these filters, we

14 Initial analysis of the unfiltered data indicated that a very large share of false positives came from these biographies. We also used other cut-offs, such as the 90th percentile of the number of proper nouns within each page in the document. Results are similar. Our approach may lead to some false negatives, for example if Amazon Web services is identified as a proper noun. That said, most of the results are robust to including these pages so this does not appear to be a major issue.

then spot check 10 reports or calls each year and find no false positives. As a further robustness, we also consider the results focusing only on annual reports and removing earnings calls where there may be more references to 'hot topics'. Doing so does not change the results.

While we are confident that we have removed all false positives, we cannot rule out the possibility of false negatives where firms are using the technology but not reporting it, or not reporting it in a way that our method identifies. Assuming these false negatives occur randomly, and we are not consistently missing certain types of adoption or certain types of firms' adoption, the false negatives will tend to be 'noise' biasing our analysis towards finding no results. That said, we would expect our approach to mainly pick up substantive implementations of the technology, rather than, for example, simply shifting to a cloud version of word processing software. As such, our results should be interpreted as relating to substantive implementations of these technologies.

Given we identify adoption based on mentions in company reports, it is possible that adoption occurs before firms enter our sample. This could occur for two reasons. First, adoption occurs before our sample begins. This is unlikely given our sample goes back to the early 2000s. Second, firms adopt before listing on the stock exchange entering our sample. Such firms would either be recorded as adopting later than they actually did, if the technology is mentioned, or never adopting. Both will tend to bias us to wards finding no effects of adoption. However, there is no reason to argue this bias would differentially affect the sub-samples of firms with and without Board members with certain skills.

We end the sample in 2022, before the launch of ChatGPT and pick up in the popularity of Large Language Models. While this may limit the ability to assess the effects of recent advancements and adoption, it allows us to avoid measurement error that may come from 'AI-washing', where firms report using AI/ML technologies as a marketing tool. Moreover, recent surveys have highlighted that much of the recent use of AI technologies in Australia is very shallow (Fernando, McLoughlin and Ratnayake 2025). So by focusing on these measures in the pre-ChatGPT era, we should be able to better identify more significant technology adoption.

Appendix A shows the patterns of adoption over time. The key takeaways are the adoption in Australia has lagged patterns in the US, as documented in Bloom *et al* (2021) for example. The outcomes have also been highly heterogenous across sectors, highlighting the importance of including sector-level controls in the analysis.

2.3.2 Firm-level financials

We merge our adoption information with firm financial information Morningstar.[15] Table 2.1 displays the median values of various financial statistics for firms outside of the IT sector (which we abstract from for most of the paper) in 2022 as an example, comparing those adopting in that year to those who are yet to adopt.[16]

15 Note that firms will leave our sample if they delist. This could cause some sample selection bias, if those firms that exit following adoption are doing particularly poorly, biasing us towards finding a positive impact on profitability.

16 As such, those firms that previously adopted are not included. We take this approach to think about the characteristics of firms when they adopt, rather than after adoption, given adoption could be associated with firm growth. That said, the patterns are similar if we take the simpler approach of just looking at firms that ever adopted, versus those that have not (Table B1).

Table 2.1: Descriptive Statistics of Firm Characteristics

By whether a firm adopted GPT in 2022, median

	Adopters in 2022	Firms yet to adopt in 2022
Return on assets (%)	0.8	–3.9
Revenue ($m)	108.3	1.6
Assets ($m)	413.9	42.9
Cash ratio (%)	0.1	4.2
Debt ($m)	54.0	0.7
Gearing ratio	0.4	0.1
Number of observations	70	928

Notes: Firms in the IT sector are excluded, as are firms that have adopted prior to 2022. Return on assets is measured as EBITDA/assets*100, cash ratio is measured as cash/assets*100 and gearing ratio is measured as debt/equity*100.

Sources: Authors' calculations; Morningstar; Refinitiv.

Generally, adopters are larger in terms of both revenue and assets. This finding is consistent with several papers focusing on AI use (e.g Calvino *et al* 2022; Calvino and Fontanelli 2023; McEhleran *et al* 2024), who have tended to attribute the finding to scale advantages, possibly due to high fixed costs in AI adoption. They also appear less liquid based on the raw summary statistics. However, once we account for industry controls, higher liquidity is found to be associated with adoption (Table C1), consistent with findings for AI usage overseas (Alekseeva *et al* 2021; Babina *et al* 2024). This suggests that financing constraints could be a relevant barrier to adoption.[17]

2.3.3 Board of Directors characteristics

To consider the role of the Board of Directors, we use information from S&P Capital IQ. Specifically, this database has information and biographies on the past and present Board members at Australian listed companies. Using a snapshot of the Board members as at March 2023, we derive variables for each individual, covering: age; gender; whether they have a Master of Business Administration degree; whether their study was in a STEM-related field[18]; whether they were previously a Board member or a 'key professional' in the IT industry; whether they have experience with GPT, as indicated by reference to any of the technology-related words in their biography. For each firm, we create metrics capturing the share of Board members with certain characteristics, based on the position held as of March 2023.

Table 2.3 provides a snapshot of the data, splitting the sample into those firms adopting technologies in 2022 and those that have not adopted to that point, and again focusing on non-IT firms. On average, adopting firms' boards are more likely to have a female member, an individual with an MBA or STEM degree and previous Board positions in IT firms. Notably though, the correlation with female representation and education appears to largely reflect other factors such as industry and firm size. Once accounting for these, female representation and education are no longer a significant predictors of adoption (Table C4).

17 We run several robustness tests which do not change the results. These include using a probit model and taking a non-parametric approach by replacing the linear explanatory variables with dummy variables showing which decile of the industry-level distribution the firm was in (see Appendix C). The decile model indicates that the largest firms in each industry (i.e. top two deciles) are far more likely to adopt than other firms.

18 A STEM degree is a degree in the fields of science, technology, engineering or mathematics.

Table 2.3: Descriptive Statistics of whether at Least One Member of the Board of Directors have a Particular Characteristic

By whether a firm adopted in 2022

	Adopters in 2022	Firms yet to adopt in 2022
Female (%)	49	36
STEM degree (%)	39	28
MBA degree (%)	42	25
Experience with GPT (%)	25	10
Experience in IT industry (%)	9	2
Age (average, years)	60	60
Number of observations	67	787

Notes: Based on March 2023 snapshot of Board members. Firms in the IT sector are excluded, as are firms that have adopted prior to 2022. The number of observations that have information about age of Board members is smaller (41 adopters and 423 non-adopters).

Sources: Authors' calculations; Morningstar; Refinitiv; S&P Capital IQ.

Most relevant to the rest of our analysis, adopters are also much likely to have a director with experience in emerging digital GPT, based on the relevant keywords appearing in their biography.[19] One explanation for this finding is that having Board members with these skills contributes to the decision to adopt. Another is that firms bring on Board members with such skills because they intend to, or have adopted, these technologies. So these skills facilitate adoption, but do not lead to adoption. Given the lack of longitudinal information, we cannot separately test these two explanations. However, both suggest that having directors with technology-related knowledge is important for adoption.

It is important to reiterate that these findings should be thought of as correlations, rather than causal. For example, it could be that firms with strong governance have more diverse boards, or boards with more diverse skills, and these same firms may also be more likely to adopt technologies. Still, they provide some evidence that having a mix of skills and backgrounds is associated with greater adoption. And the provide some initial indication of potentially important drivers and barriers to explore.

2.3.4 Hiring data

As a final dataset, we also incorporate information on firms' job advertisements and whether they mention GPT, which we take as an indicator of trying to bring in those skills. We use the database constructed by Bahar and Lane (2022), which uses Lightcast job ad data from 2012 to create indicators identifying whether firms are advertising for GPT-related skills. We match these to our database using a combination of ASX tickers and company names, along with fuzzy matching. Similar data were used by Bloom *et al* (2021), as well as several papers looking at AI use (e.g. Alekseeva *et al* 2021; Calvino *et al* 2022; Babina *et al* 2023, 2024).

To provide some initial insights, we consider whether firms that adopt are more likely to place job advertisements for workers with these skills, running a simple regression of the following form:

19 One concern with this metric is that references to the technology might reflect the adoption of technology by the firm in question. However, spot checks of the data suggest this is not the case, and that references tend to relate to past work or educational background.

$$Hire\,a\,skilled\,worker\,by\,2022_i = \beta * Adopt\,by\,2022_i + \gamma_{ind} + \log\left(Assets_{i,2022}\right) + \varepsilon_i$$

Our outcome variable $Hire\,a\,skilled\,worker\,by\,2022_i$ takes on the value 1 if the firm has put a job ad looking for GPT-skilled workers during the sample, and 0 if none of their job ads asked for these skills. We regress this on an indicator for whether or not the firm has adopted a GPT by 2022. Again we control for industry and size to capture other factors that might explain outcomes.

For the sub-sample of firms for whom we can observe job advertisements, those that have adopted a GPT are far more likely to have advertised for GPT skills at some stage, with the probability being around 16 percentage points higher than for non-adopters once we control for firm characteristics (Table 2.5). This suggests that firms increase their demand for these skills when trying to use these technologies.

Table 2.5: Relationship between Hiring Skilled Workers and Technology Adoption

Linear probability regression of ever advertised for GPT skills on ever adopted

	With no controls	With industry controls	With industry and size controls
Adopt	0.372***	0.397***	0.162**
	(0.079)	(0.083)	(0.070)
Number of observations	215	215	215

Notes: Excludes firms in the IT industry. Only includes firms that have job ads captured in the Bahar and Lane (2022) sample. *, **, *** indicate significance at the 10, 5 and 1 per cent level, respectively. Standard errors are shown in parentheses and are clustered at the industry level. Size control accounts for firm's largest ever (log) assets.

Sources: Authors' calculations; Bahar and Lane (2022) using Lightcast data; Morningstar; Refinitiv.

2.4 Methodology

We employ a panel event study framework to estimate the effect of GPT adoption on firm-level outcomes. The framework allows us to analyse changes in firm outcomes before and after the adoption of GPT, accounting for the fact that firms adopt GPT at different points in time which can create issues around overlapping treatment periods in difference-in-difference models (Goodman-Bacon 2019). The approach is also similar in nature to the dynamic analysis undertaken in Babina *et al* (2024) for AI adoption.

The variable $Adopt_i$ indicates the period when the technology was first referenced by firm i. The outcome of interest is denoted as y_{it}, and the panel event study specification is as depicted in the equations:

$$y_{i,t} = \alpha + \sum\nolimits_{2\le j\le J} \beta_j \left(Lag\ j\right)_{i,t} + \sum\nolimits_{1\le k\le K} \delta_k \left(Lead\ k\right)_{i,t} + \Gamma * \log\left(Assets_{i,t}\right) + \mu_i + \theta_{s,t} + \epsilon_{i,t}$$

where:

$$(Lagj)_{i,t} = 1\{t = Adopt_i - j\} \text{ for } j \in \{1,\dots,J\}$$

$$(Leadk)_{i,t} = 1\{t = Adopt_i + k\} \text{ for } k \in \{1,\dots,K\}$$

The adoption event's lags and leads are defined as binary variables signifying that a specific firm was a given number of periods away from the adoption event. The coefficients of interest are the betas related to the lags and deltas for the leads. We focus on up to four years before adoption and three years after.[20] While a longer post-adoption window could be appropriate if these investments have very long payoff windows, we are constrained by the sample period available.

We control for firm fixed effects in μ_i and industry-time fixed effects $\theta_{s,t}$. The former allows us to control for time-invariant unobserved heterogeneity at the firm level, while the latter accounts for industry-level patterns in adoption, demand conditions, tax incentives and reporting standards. We also control for firm size, measured as log assets, given the earlier findings around its importance for adoption. We cluster standard errors at the firm level to account for correlation of outcomes within firms over time.

For the analysis, our reference or base period, relative to which all other results are interpreted, is two years before adoption. We choose this rather than the identified year of adoption as there is some evidence that firms may be investing or changing their behaviour ahead of adoption.

The key assumption made in these regressions is parallel trends. In absence of the treatment (adoption) the treated and control firms would have behaved similarly. One way to assess this is to explore the pre-trends. If the firms were following a similair trajectory up until adoption, this provides some evidence that such an assumption is reasonable. Controlling for size and industry, to of the key predictors of adoption, should help with this.

Nevertheless, our results should still be thought of as correlations, rather than as strictly causal. Ultimately firms are selecting into adoption (treatment), rather than it being allocated randomly, and they may do so in expectation of some future outcome. It may be that, notwithstanding the similair pre-trends, their selection into treatment reflects something about their future prospects. Such a bias could go either way: they may select in to take advantage of good opportunities, or they may do so to avoid an impending downturn. In either case, this does suggest some caution should be taken in interpreting the results as causal.

Finally, we note that there is an emerging literature highlighting the biases that can come from using event study approaches when the effects of treatment vary across treatment cohorts (see Sun and Abrahams 2021 and Roth 2023 for discussions). To address this, we apply the imputation method proposed by Wooldridge (Forthcoming) for all our main regressions, using the never treated group as out main comparison.[21] These results are shown in Appendix J.

2.5 Results

2.5.1 Effect of adoption on profitability

We start our analysis focusing on non-IT firms over the full sample period. This includes any adoption event occurring from 2013 to 2022, and any balance sheet data back to 2009 to account

20 We restrict our attention to the four periods prior to adoption only as there is high volatility and noise in the more distant history.

21 This is done using the jwdid package in Stata. I do not include log assets as a control in this case given difficulties including time-varying covariates in this model.

for the pre-adoption behaviour of firms. Profitability generally remains broadly stable following adoption, showing neither statistically nor economically significant differences compared to the pre-adoption period (Figure 2.1).

Figure 2.1: Return on Assets around GPT Adoption

Full sample

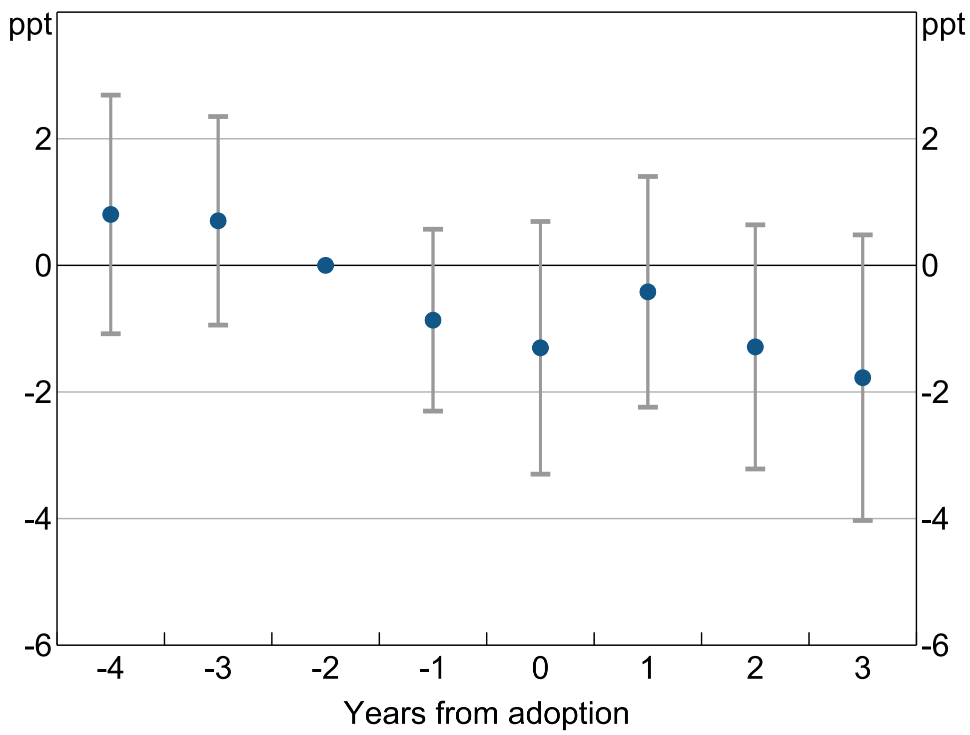

Notes: Dots represent point estimates. Whiskers indicate 90 per cent confidence intervals.

Sources: Authors' calculations; Morningstar; Refinitiv.

There is a large literature documenting that the diffusion of GPT tends to be slow because of the need to create complementary inventions, or make organisation changes, that facilitate profitable adoption of the GPT (e.g. Agrawal *et al* 2023). In the case of cloud computing, this could reflect the developments in cloud-based programs or cloud services. For AI/ML, this could include the development of simple machine learning models and packages.

As such, it is worth examining whether profitability evolves differently for early and late adopters.[22] Figure 2.2 shows that there is a significant downturn in profitability in the years following the adoption event for early adopters, defined as those adopting from 2013 to 2016. Conversely, profitability of late adopters (adopting from 2017 to 2022) seems to remain broadly stable post-adoption.[23] To test whether the difference in post-adoption outcomes for early and late adopters is significant we stack the two groups in a staggered difference-in-difference-in-difference model. We do find a marginally significant difference in outcomes, suggesting that late adopters had better post-adoption outcomes than early adopters (Appendix F). These findings also align with Hajikhani

22 Early adopters tend to be smaller but have similar profitability pre-adoption to late adopters (see Table B1). They are also more likely to be in certain sectors, as shown in Figure 2.

23 For this regression we begin the sample in 2014 (in line with the lags shown) and allow firm adopting pre-2012 to enter as additional controls treating them as never adopters to increase the sample size. But results are robust to excluding them.

et al (2025) who find less evidence of productivity increases following adoption for early Finnish adopters.

Figure 2.2: Return on Assets around GPT Adoption

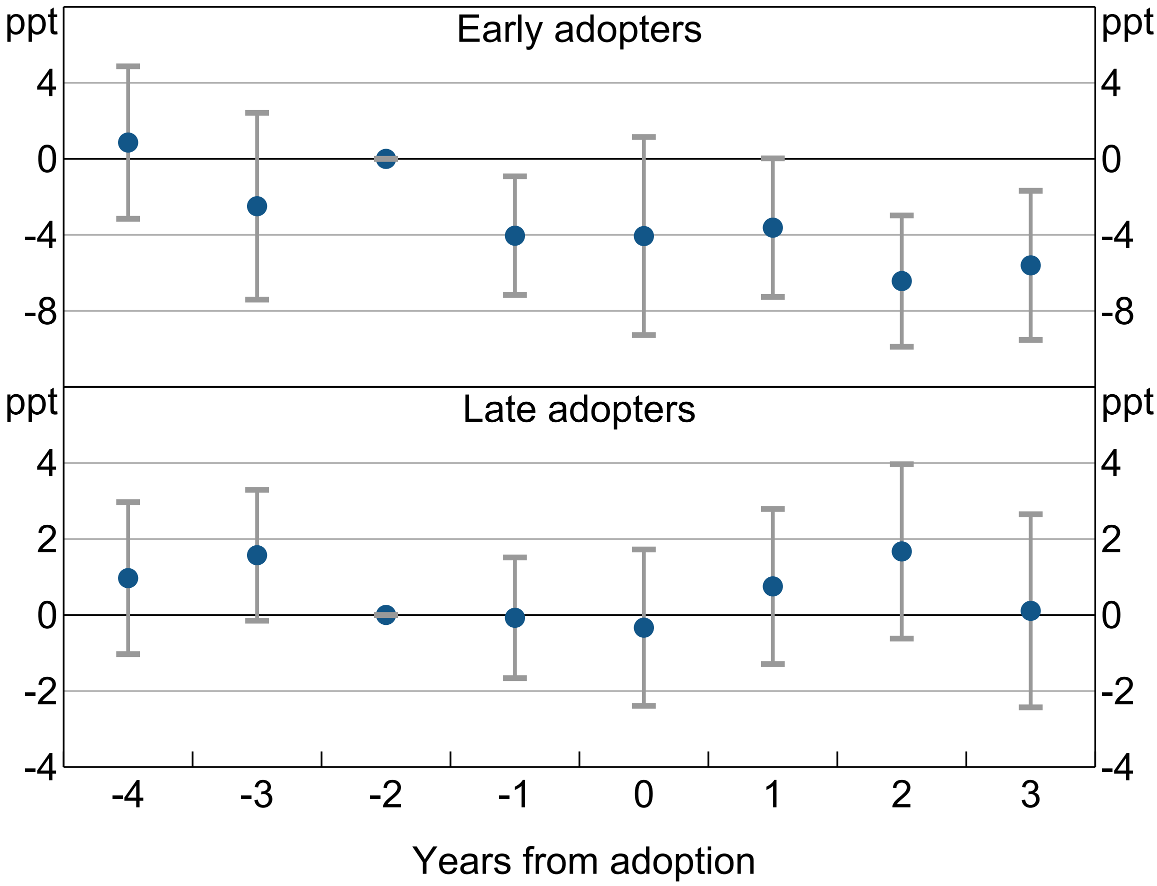

Notes: Dots represent point estimates. Whiskers indicate 90 per cent confidence intervals.
Sources: Authors' calculations; Morningstar; Refinitiv.

Importantly, there is no significant trend in outcomes in the pre-adoption period across all samples, providing some evidence that the findings do not reflect differing prospects and trajectories for the firms. That said, there is some volatility in profitability in the years pre-adoption. However, the results are robust to changing the exact period definitions, for example by lumping all pre-and post-adoption periods together (Appendix E), or testing their joint difference from zero.

It is also possible that there are some other differences between early and late adopters that are driving the differences in their outcomes. One is that changes in accounting standards could potentially play a role. Around 2019 and 2020 there was a clarification in accounting standards indicating that most cloud computing expenses should be expensed, though a small share could be recorded as investment (Department of Finance 2022). This is likely to have led more firms to record GPT-related outlays as expenses, which should weigh on their profitability. More generally, the adoption of cloud technology could lead firms to shift from capital expenditure with subsequent depreciation expenses, to operating expenses, making it look like firms' profitability is falling post-adoption when profitability is measured using EBITDA (as we do).

We take two simple approaches to considering these issues. The first is to deduct capital purchases from profitability. The second is to account for depreciation expenses by measuring profitability as EBIT (earnings before interest and tax). The results remain broadly unchanged, though the measure including capital expenditure is noisier (Appendix H).[24]

One further concern might be the inclusion of the COVID-19 pandemic in the late adopters sample. It may be that those firms that had adopted cloud or AI/ML technologies were better able to navigate the COVID-19 shock, either because they were better managed or because the technologies helped navigate the shock (consistent with the large spike in adoption). As a simple test, we ended our sample in 2020. Doing so leads to, if anything, stronger evidence that late adopters increase profitability post-adoption in terms of the magnitude of the effect, though the response is not significant (Appendix I).[25]

The change in the effects over time particularly highlights some of the issues recently documented around using event studies with varying cohort effects. Nevertheless, the results are very similair when we use the Wooldridge estimator (Appendix J).

Taken together these results suggest that profitable adoption has become easier over time, consistent with further developments in these technologies.[26] If this is the case, it suggests that there may be more scope for effective adoption of these technologies going forward, which could improve productivity. That said, there still appears to be a large degree of variation in outcomes, as reflected in the wide error bands. This suggests other factors may dictate whether adoption is profitable, which we explore below.

2.5.2 The role of Board of Directors

In this section, we delve into the characteristics of the Board of Directors that correlate with the profitable adoption of technologies. We conduct the same event study looking at profitability post-adoption, but split the sample based on whether the Board has members with certain characteristics or not. We focus on late adopters (post-2016), given this is when the bulk of adoption occurs. Focusing on the most recent adoptions also should help to limit issues around the cross-sectional snapshot of directors becoming less accurate over time.

Focusing first on whether the Board has some members with a background in GPT, there is evidence of an increase in profitability for adopting firms whose current Board has some degree of GPT expertise (Figure 2.3). This is not the case for boards with no members with GPT experience.

---

24 Another approach considered was to use coarsened exact matching, to match our firms across several observables. We considered this approach, and while some of the patterns were similar, the small sample size meant that the results were very volatile and imprecisely estimated. Moreover, given the pre-trends are quite flat in our main regressions, and given the main concerns around causality likely arise around selection on unobservables, we have not pursued this approach further.

25 The difference in outcomes between firms with and without GPT skills on their board is also evident, though the former no longer shows a significant increase. Removing mining firms, which may have very different technology needs, also does not change the results (Figure I1).

26 We can't rule out the possibility that the nature of the early and late adoptions differed substantially. For example, early adoptions may have been larger investments, more experimental and more disruptive. They may also have had positive spillovers or greater upside risk in terms of skills formation, learning by doing or pushing the knowledge frontier forwards. Nevertheless, the evidence suggests that more recent adoptions were easier to turn into profit (or at least to avoid a loss), even if this simply reflected the potentially less ambitious nature of these adoptions.

The difference in outcomes is statistically significant (Appendix F) and robust to period definitions (Appendix E), and again tests of pre-trends find no evidence of trends prior to adoption.[27] The difference in outcomes appears to be driven by a reduction in operating expenses for firms with a GPT-experienced Board, while revenue remains stable post-adoption for all firms (Appendix D). This contrasts to Babina *et al* (2024) who find that AI adoption leads to increased sales, but relatively few cost savings.

Figure 2.3: Return on Assets around GPT Adoption by Board Skills

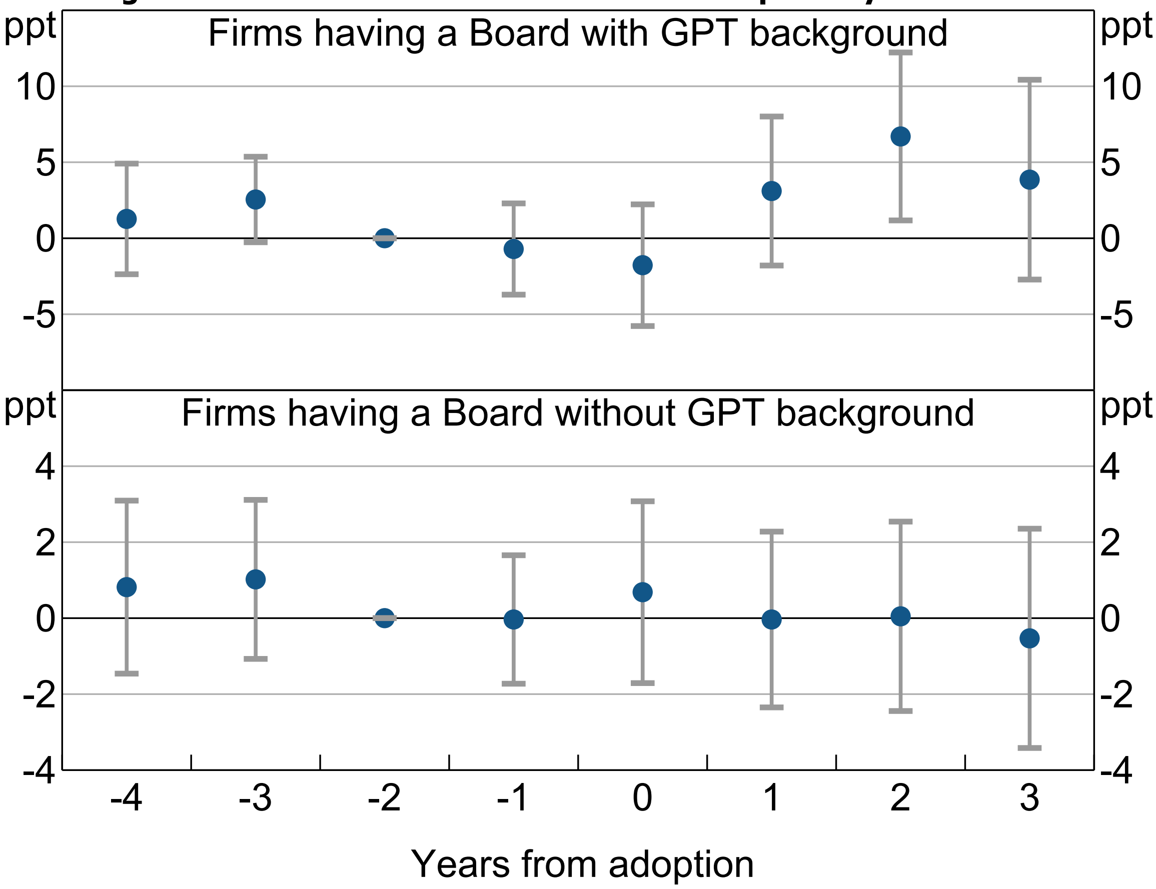

Notes: Late adopter sample. Dots represent point estimates. Whiskers indicate 90 per cent confidence intervals.
Sources: Authors' calculations; Morningstar; Refinitiv; S&P Capital IQ.

These results provide some evidence that boards with relevant experience may help profitable adoptions to GPT. Such a finding is somewhat consistent with evidence from Alekseeva *et al* (2021) and Calvino *et al* (2022) that having managers with AI skills is highly valued in terms of wage premia or firm productivity, respectively. It is also consistent with Fontanelli *et al* (2025) finding that firms without such skills have more volatile productivity post adoption.

One concern with this finding may be that Board members add mentions of GPT to their biographies if their company profitably adopted the technology. However, spot checks of the biography data suggest that this is not a key driver, and that we are mainly picking up the directors' educational background and previous work experience. Again, this result is also robust to using the methodology suggested by Wooldridge (Forthcoming) (Appendix J).

27 If we extend the pre-adoption window back further the evidence of pre-trends remains weak, though there is some volatility 5 years before adoption. Accumulating all periods from four years ahead of adoptions together provides no significant evidence of pre-trend for those with GPT-skilled Board members, though there is some for other firms.

None of the other Board characteristics were associated with differing profitability post-adoption. While there was some evidence for female representation, this findings was not robust to using the Woolridge estimator.

2.5.3 Adoption and Hiring

One natural question is, why does having a Board with relevant skills promote profitable, rather than non-profitable, adoption? Does it appear that these adoptions are potentially accompanied by broader changes to business processes and staffing, consistent with the benefits of these technologies mainly accruing when such accompanying investments are made?

To consider this we look at how firm advertisements for relevant skills change around adoption events, using the dataset created in Bahar and Lane (2022). This allows us to consider whether the firms are investing in complementary skills at the same time. This may be indicative of making larger investments in surrounding processes, and more generally of more substantive investments in these technologies, given hiring of such workers is a common measure of AI adoption used in the literature (e.g. Alekseeva *et al* 2021).

To do so we apply the same event study as above, with the outcome variable of interest being a dummy variable that takes on the value 1 if any of the job advertisements the firm posted during the year were for GPT, and 0 if none were asking for these skills.[28] As such, we look at the effect of adoption on the probability that, conditional on placing a job advertisement, the advertisement was for staff with GPT skills.

We do see a pick-up in hiring post-adoption (Figure 2.4). However, this is only evident for firms with Board members with GPT experience. As noted, these are also the same firms that have the most evidence of increased profitability post-adoption. Other firms do not show any change in hiring of GPT-skilled workers around adoption. Again this is robust to using the Woolridge (Forthcoming) estimator.

While certainly not definitive, this provides some tentative evidence that hiring skilled workers is an important step in the profitable adoption of some of these GPT. Such a finding is also consistent with overseas findings that firms with more skilled workers are more likely to invest in AI/ML in the form of further hiring of skilled workers (eg. Babina *et al* 2023). This highlights how digital GPT adoption may be associated with increased demand for associated skills, and the importance of having a skilled workforce plays in the effective adoption of GPT.[29]

More generally, it also provides some evidence that these firms are making more substantial complementary investments in technologies and processes as part of the adoption. The fact that these firms are also those making gains from adoption suggests those investments are important in getting the most out of the technologies, consistent with the existing literature. And that having management with relevant skills makes that deeper adoption and complementary investment more likely.

---

28 Those for whom we cannot see a job advertisement in the year are not considered.

29 We explored whether firms that have placed job advertisements for GPT-skilled workers at some stage are more likely to increase profitability post-adoption. There is no evidence of a significant difference. In part this may reflect a lack of information on the stock of workers. For example, firms may already have skilled workers, and not need to hire. Further work could look to incorporate information on the stock of workers to better consider this question.

Figure 2.4: Hire Tech Skills around GPT Adoption

By Board skills

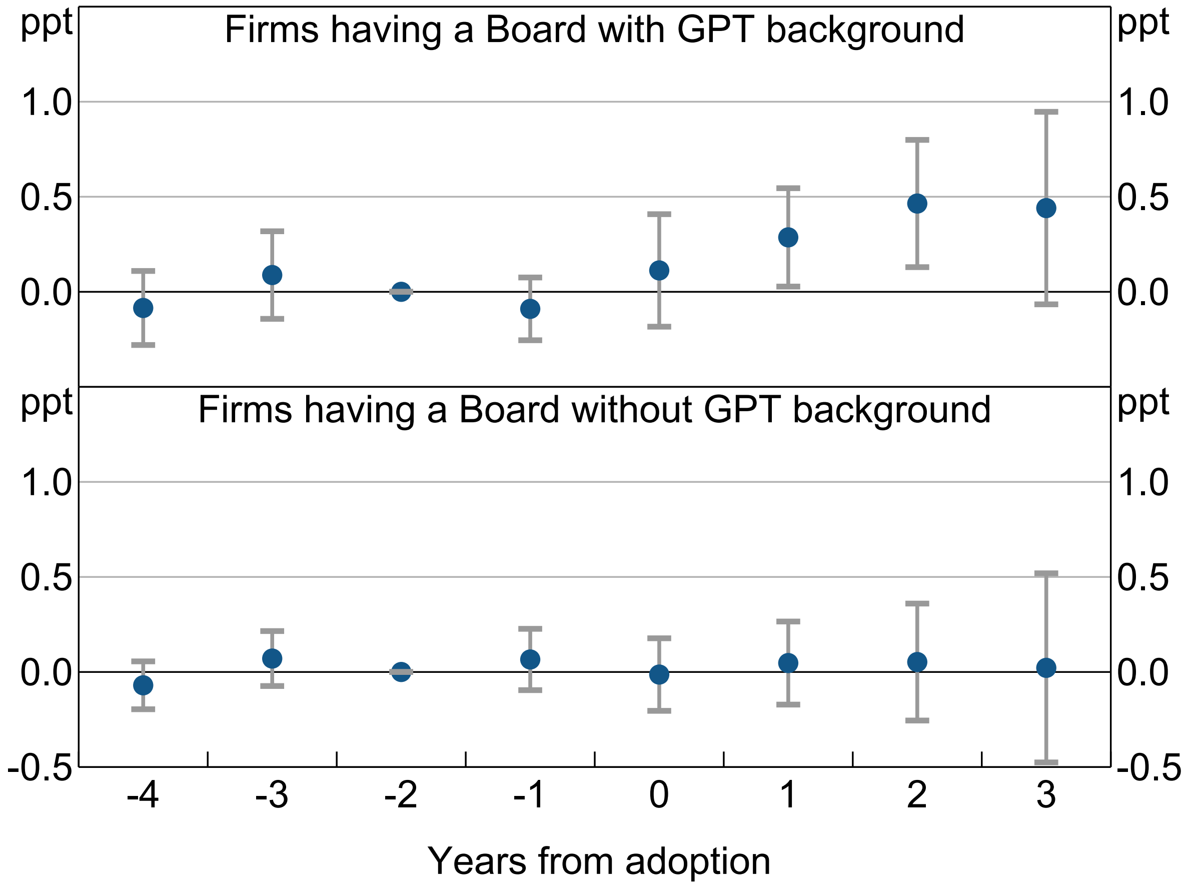

Notes: Late adopter sample. Dots represent point estimates. Whiskers indicate 90 per cent confidence intervals.

Sources: Authors' calculations; Bahar and Lane (2022) using Lightcast data; Morningstar; Refinitiv; S&P Capital IQ.

2.6 Conclusion

Emerging digital GPT have the potential to revolutionise the economy and help turnaround slow productivity growth outcomes observed in many countries over recent decades. But the pace of these gains will depend on how quickly firms can effectively adopt and adapt these technologies. Understanding what factors predict profitable adoption and deeper integration of these technologies into production will be important in understanding and facilitating this process.

Our results help us to better understand this process. Leveraging a unique dataset derived of annual reports, Board characteristics and hiring data for listed firms in Australia we show that not all adoption is equal. Many firms show no increase in productivity following adoption. Those firms that do profitably adopt are far more likely to have Board members with relevant skills, highlighting the importance of management skills in profitable adoption. Moreover, they are far more likely to invest in skilled workers alongside the adoption, consistent with the adoption coming alongside complementary investments.

These results have several implications for the prospects for productivity from digital GPTs. First, they reinforce findings elsewhere that complementary investments will be crucial if we are to get the most out of these technologies. This may mean that the gains will be slower than some anticipate. Second, they highlight the importance of skilled management in the profitable adoption of these technologies. While this is intuitive, given management will play an important role in redesigning processes and making decisions around complementary investments, to date the role of technically skilled managers in profitable adoption has been of limited focus. This potentially represent a significant barrier to profitable and productive adoption in many countries, particularly smaller technology-importing counters, given such skills are in limited supply. This highlights the importance for policymakers of building a skilled and dynamic labour force that extends to the management level.

While the increasing ease of use of these technologies may overcome some of these barriers, firms will still need to make decisions around changes to operational processes to make best use of the technologies. As such, the important role of management in profitable adoption is unlikely to disappear. In fact, it could become even more important as the scope of these technologies becomes even wider and the decisions become more multi-faceted. Nevertheless, a valuable direction for future analysis would be to explore similair dynamics using the period post the emergence of ChatGPT and other large language models using similair approaches, but with differing ways of identifying adoption given the likelihood of false positives and ‘AI-washing’ in firm reports.

Appendix A: Adoption metrics

Table A1: Keywords Used to Identify Technologies
From Bloom *et al* (2011)

Cloud computing	Artificial intelligence and machine learning
paas, cloud infrastructure, distributed cloud, cloud provider, cloud offerings, cloud service, cloud applications, community cloud, private cloud, public cloud, cloud deployments, cloud environments, cloud management, cloud services, cloud security, enterprise class, iass, hybrid cloud, cloud platform, cloud hosting, personal cloud, enterprise network, cloud computing, cloud based, saas, cloud storage, enterprise applications, cloud solution, enterprise cloud, cloud deployment	neural network, deep learning, language processing, machine learning, machine intelligence, natural language, artificial intelligence, ai technology, supervised learning, learning algorithms, unsupervised learning, reinforcement learning, ai machine

Appendix B: Trends and Additional Descriptive Statistics

Figure B1: Share of Firms Reporting GPT

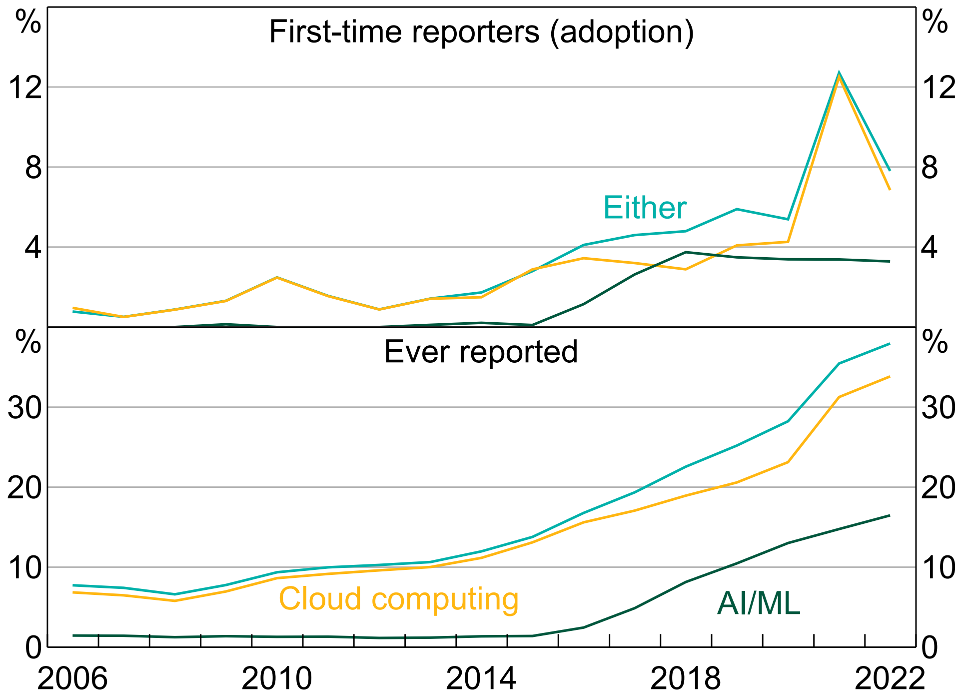

Notes: Figure shows the share of firms reporting a technology for the first time in the top panel, and share that have ever reported in the bottom panel.

Sources: Authors' calculations; Refinitiv.

Figure B2: Share of Firms Ever Reporting GPT

By sector

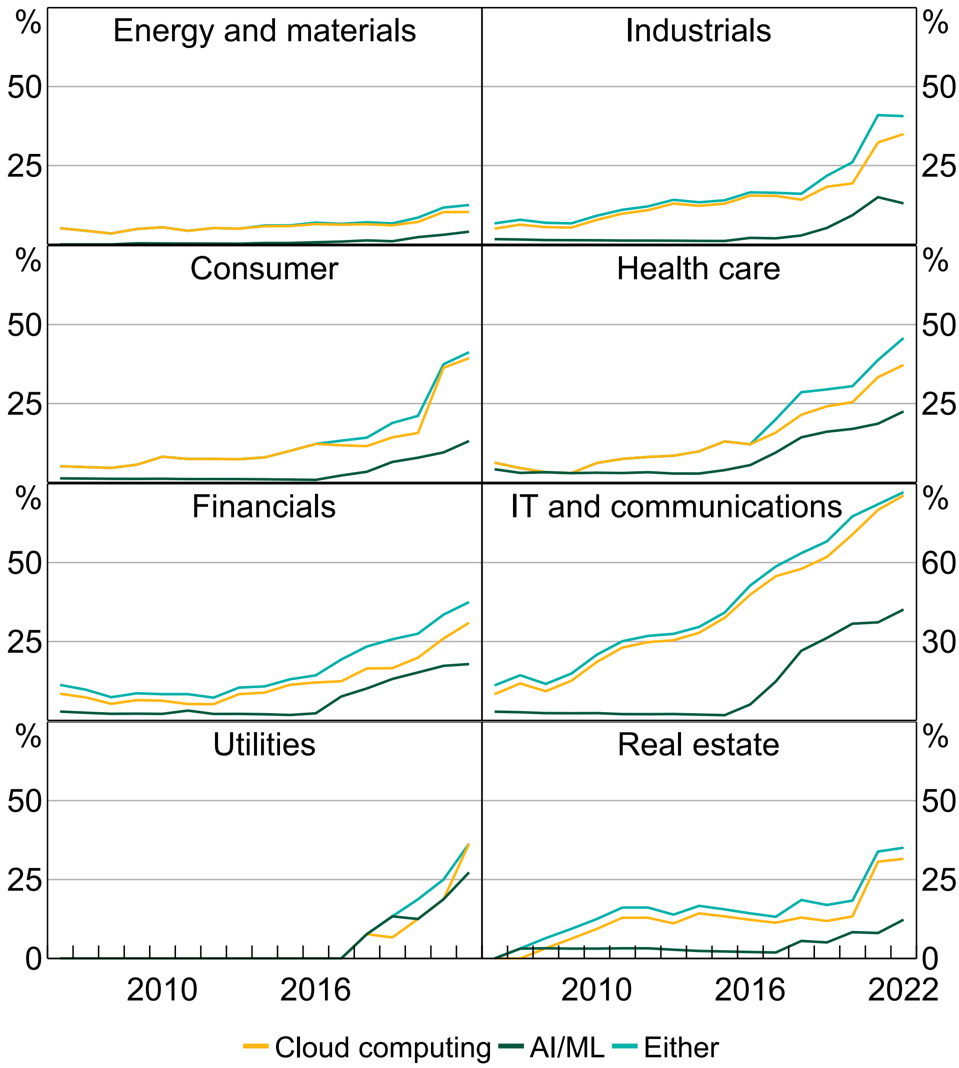

Notes: Figure shows the share of firms that have even reported a technology split by sector.

Sources: Authors' calculations; Refinitiv.

Table B1: Descriptive Statistics of Firm Characteristics

Prior to adoption, median

	Early adopters	Late adopters
Return on assets (%)	4.9	4.2
Revenue ($m)	32.6	98.0
Assets ($m)	83.7	274.4
Liquid assets ($m)	0.5	0.3
Debt ($m)	5.4	37.0
Gearing ratio	0.3	0.4
Number of observations	374	3,456

Notes: Firms in the IT sector are excluded, as are firms that have adopted prior to the collection year. Return on assets is measured as EBITDA/assets*100 and gearing ratio is measured as debt/equity*100. The number of observations is at firm-year level.

Sources: Authors' calculations; Morningstar; Refinitiv.

Appendix C: Robustness Checks for Adoption Determinants

C.1 Adoption and firm characteristics regressions

To account for differences across industries we consider a regression of the following form, where we lag each the factor by two periods.[30]

$$Adopt_{i,t} = \beta * Factors_{i,t-2} + \gamma_{ind,t} + \varepsilon_{i,t}$$

where $Adopt_{i,t}$, takes on the value 1 if firm i adopts in year t, and 0 otherwise. $Factors_{i,t-2}$ includes several variables of interest, including balance sheet metrics and firm age, lagged by two years. We account for industry-by-time fixed effects to abstract from adoption patterns in each industry ($\gamma_{ind,t}$).

We focus on firms yet to adopt by that date, so we estimate what characteristics predict adoption among those firms that are yet to adopt these technologies. The regression covers the period from 2013 to 2022, though the results are similar if we focus on more recent adopters (e.g 2017 onwards). We also consider versions using probit regressions (Table C2) and non-parametric measures of size (Table C3).

Table C1: Determinants of Adoption – Firm Balance Sheet Metrics

Estimates from linear probability regression of adoption on firm characteristics

Assets (log)	Return on assets	Positive cash flow	Cash ratio	Labour share of expenses	Age
0.012***	–0.012	0.018***	0.018*	0.013**	–0.016
(0.001)	(0.009)	(0.007)	(0.010)	(0.006)	(0.017)

Notes: Firms in the IT sector are excluded, as are firms that have adopted previously. There are approximately 5,300 observations. Includes controls for industry*time effects. *, **, *** indicate significance at the 10, 5 and 1 per cent level, respectively. Standard errors are shown in parentheses and are clustered at the firm level. Return on assets is measured as EBITDA/assets*100 and cash ratio is measured as cash/assets*100.

Sources: Authors' calculations; Morningstar; Refinitiv.

30 We choose two periods given results from our event study analysis suggest firms may already be undertaking investments in the year before they reference the technology adoption.

Table C2: Determinants of Adoption – Firm Balance Sheet Metrics

Probit regression of adoption on firm characteristics

Assets (log)	Return on assets	Positive cash flow	Cash ratio	Labour share of expenses	Age
0.142***	0.240	0.301***	0.246	0.257*	–0.174
(0.017)	(0.003)	(0.107)	(0.257)	(0.001)	(0.232)

Notes: Firms in the IT sector are excluded, as are firms that have adopted previously. There are approximately 3,650 observations. Includes controls for industry*time effects. *, **, *** indicate significance at the 10, 5 and 1 per cent level, respectively. Standard errors are shown in parentheses and are clustered at the firm level. Return on assets is measured as EBITDA/assets*100 and cash ratio is measured as cash/assets*100.

Sources: Authors' calculations; Morningstar; Refinitiv.

Table C3: Determinants of Adoption – Firm Balance Sheet Metrics

Linear regression of adoption on firm size decile indicators

Lag 2	Assets (log)	Standard errors
2nd decile	0.004	0.009
3rd decile	–0.001	0.008
4th decile	0.012	0.009
5th decile	0.016*	0.009
6th decile	0.019*	0.009
7th decile	0.018*	0.009
8th decile	0.010	0.009
9th decile	0.055***	0.011
10th decile	0.081***	0.010

Notes: Firms in the IT sector are excluded, as are firms that have adopted previously. There are 5,283 observations. *, **, *** indicate significance at the 10, 5 and 1 per cent level, respectively. Standard errors are clustered at the firm level. Deciles defined by industry and year.

Sources: Authors' calculations; Morningstar; Refinitiv.

C.2 Adoption and Board Characteristics

Again, we also explore these patterns using a regression framework that allows us to account for covariates. Due to the cross-sectional nature of our board data, we estimate the model as a cross-sectional regression, where adoption is measured as ever having adopted by 2022:

$$Adopt\,by\,2022_i = \beta * Board\,Characteristics_i + \gamma_{ind} + \log\left(Assets_{i,2022}\right) + \varepsilon_i$$

where $Adopt\,by\,2022_i$ is a dummy variable that takes on the value 1 if the firm has adopted a GPT by 2022 (inclusive). Again, we include industry controls to abstract from differences in adoption across industry and just focus of the effect of board characteristics. The variable $Board\,Characteristics_i$ is a variable that takes on the value 1 if the firm has a director with the characteristic of interest. We include all characteristics together in a multivariate regression. Finally, given the importance of size for adoption, and the possibility that larger firms are more likely to have directors with certain characteristics, we also control for the size of the firm, measured by their (log) assets in 2022.

Table 4: Determinants of Adoption – Board of Directors' Characteristics

Estimates from linear probability regression of whether a firm has adopted by 2022 on whether Board has a member with a particular characteristic

	Without size control	With size control
Female	0.076**	0.004
	(0.026)	(0.023)
STEM degree	0.014	–0.005
	(0.040)	(0.034)
MBA degree	0.044	0.001
	(0.030)	(0.030)
Experience in IT industry	0.250**	0.316***
	(0.091)	(0.076)
Experience with GPT	0.113**	0.086**
	(0.042)	(0.043)
Number of observations	1,251	1,249

Notes: Excludes firms in the IT industry. Adoption metric is whether a firm has adopted by 2022. All regressions control for industry*time effects. All explanatory variables are dummies indicating whether the firm has any Board member (as of March 2023) with the relevant characteristic. *, **, *** indicate significance at the 10, 5 and 1 per cent level, respectively. Standard errors are shown in parentheses and are clustered at the industry level.

Sources: Authors' calculations; Morningstar; Refinitiv; S&P Capital IQ.

Appendix D: Operating Expenses and Revenue

Figure D1: Operating Expenses around GPT Adoption

As a share of assets

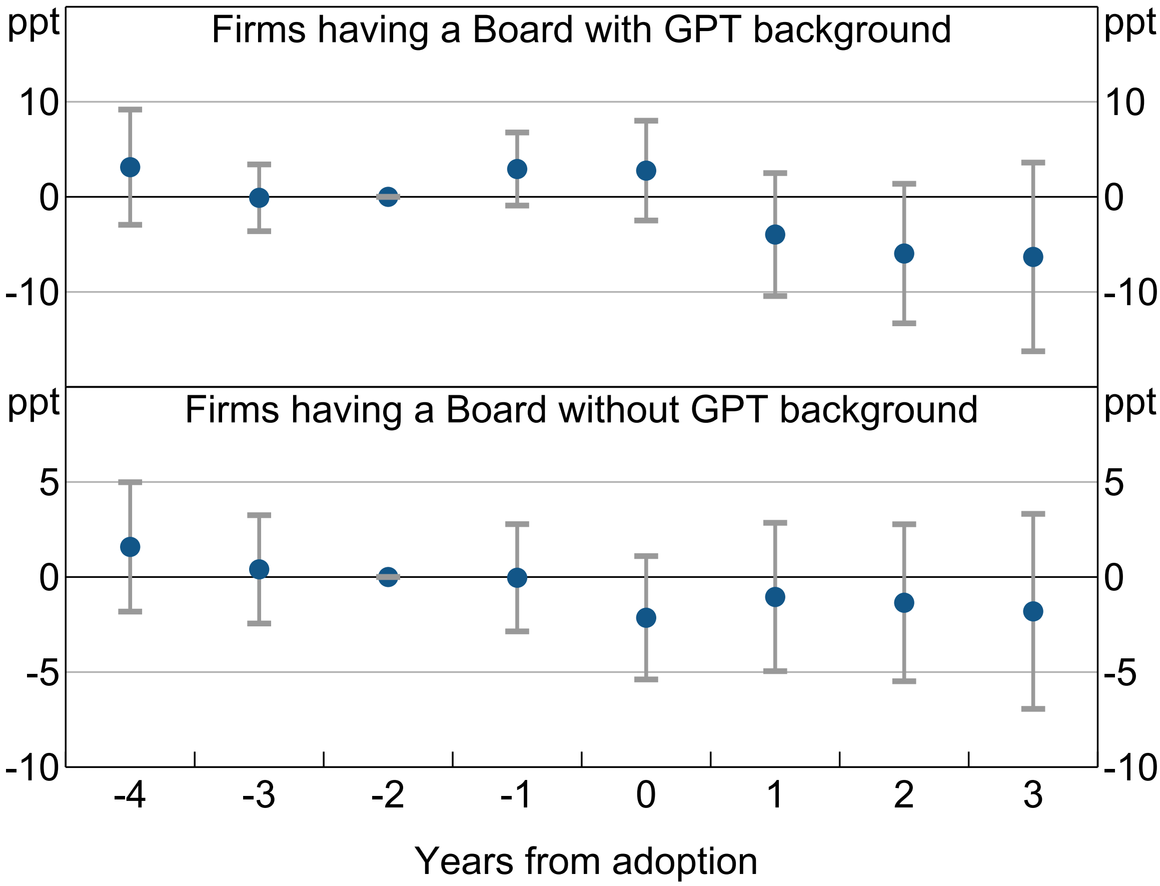

Notes: Late adopter sample. Dots represent point estimates. Whiskers indicate 90 per cent confidence intervals.

Sources: Authors' calculations; Morningstar; Refinitiv; S&P Capital IQ.

Figure D2: Revenue around GPT Adoption

As a share of assets

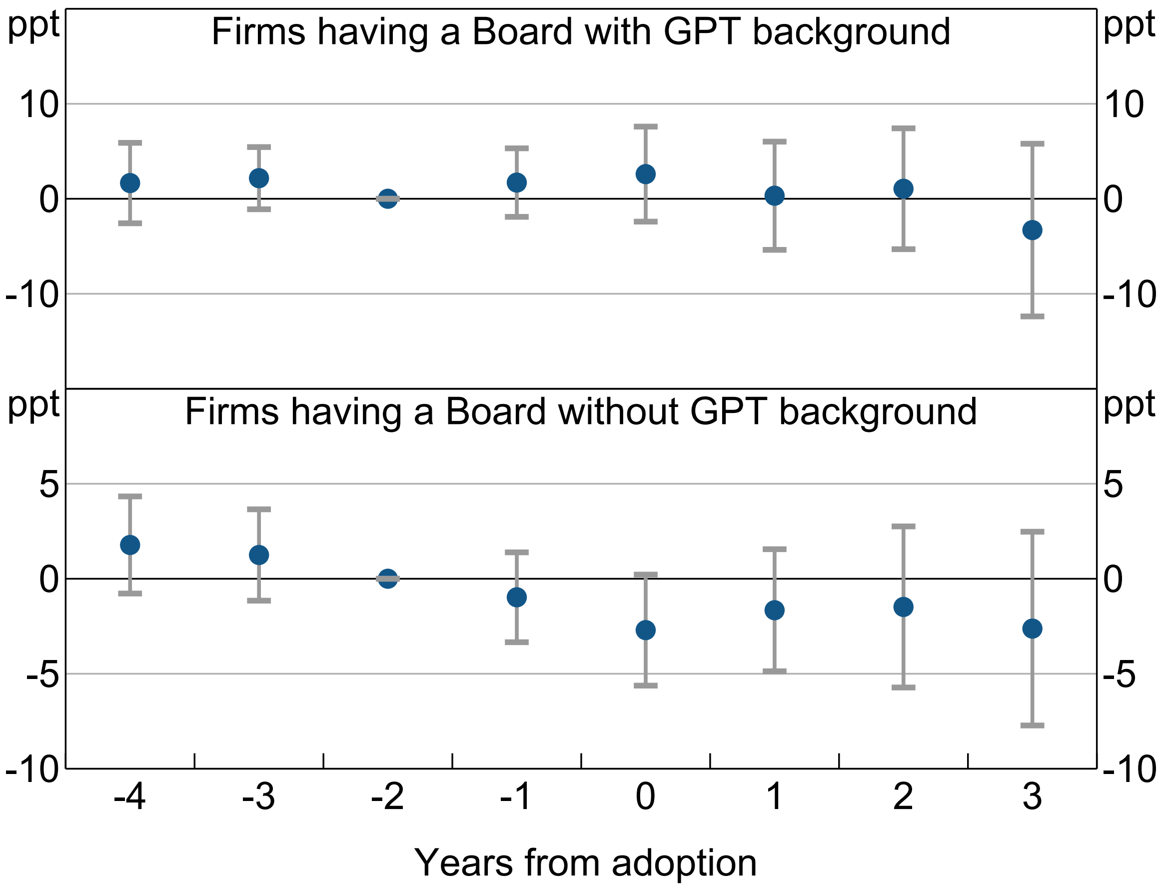

Notes: Late adopter sample. Dots represent point estimates. Whiskers indicate 90 per cent confidence intervals.

Sources: Authors' calculations; Morningstar; Refinitiv; S&P Capital IQ.

Appendix E: Robustness Checks Using Accumulated Lags and Leads

While the event study set-up with multiple leads and lags allows us to examine the timing of adoption effects on profitability, it is useful to know if, taken together, there are significant changes in profitability post-adoption from the pre-adoption periods. To this end, we accumulate all the leads and lags into binary variables indicating pre-adoption and post-adoption, as follows:

$$y_{i,t} = \alpha + \beta_J \left(Lag\right)_{i,t} + \sum_j \beta_j \left(Treat_j\right)_{i,t} + \beta_K \left(Lead\right)_{i,t} + X'_{i,t}\Gamma + \mu_i + \theta_s + \varepsilon_{i,t}$$

where:

$$\left(Lag\right)_{i,t} = 1\{t \leq Adopt_i - 2\} \text{ indicating pre-adoption}$$

$$\left(Treat_j\right)_{i,t} = 1\{t = Adopt_i - j\} \text{ for } j \in \{0,1\} \text{ indicating adoption periods}$$

$$\left(Lead\right)_{i,t} = 1\{t \geq Adopt_i + 1\} \text{ indicating post-adoption}$$

We estimate the equation for the same sub-samples as explored previously. Results show significant evidence of an increase in profitability in the joint post-adoption periods for late adopters with GPT experience, though the increase is not significant for those with female representation on the Board.

Figure E1: Return on Assets around GPT Adoption

Accumulated lags and leads

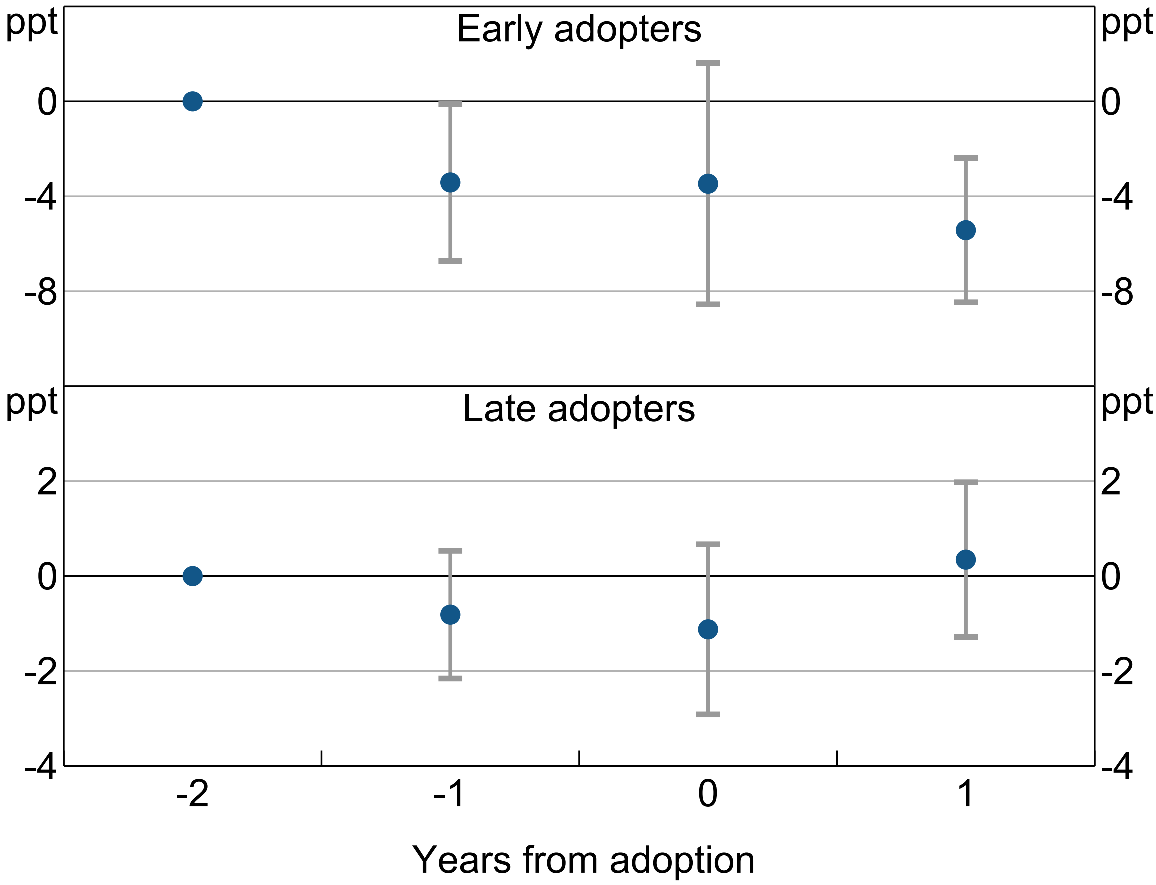

Notes: Dots represent point estimates. Whiskers indicate 90 per cent confidence intervals.

Sources: Authors' calculations; Morningstar; Refinitiv.

Figure E2: Return on Assets around GPT Adoption

Accumulated lags and leads

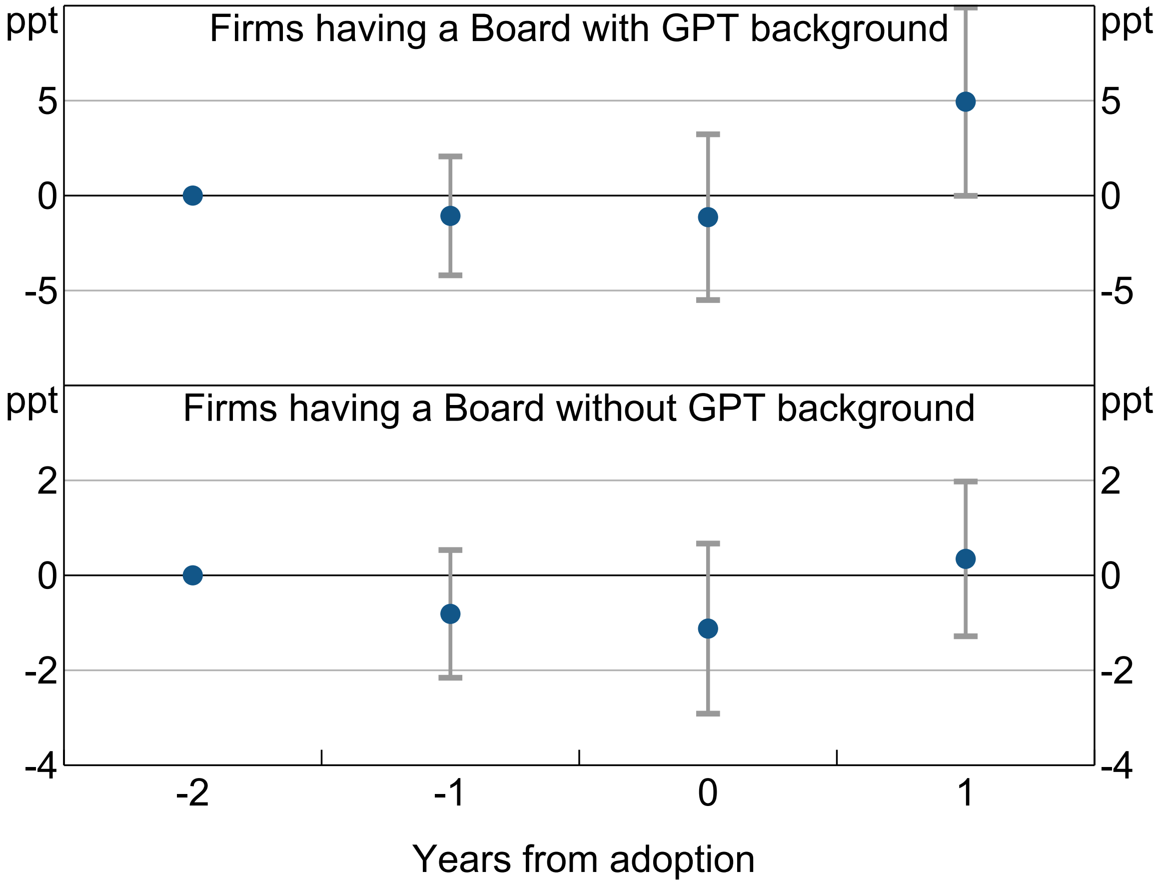

Notes: Late adopter sample. Dots represent point estimates. Whiskers indicate 90 per cent confidence intervals.

Sources: Authors' calculations; Morningstar; Refinitiv; S&P Capital IQ.

Appendix F: Robustness Checks Using Staggered Difference-in-difference-in-difference

We also employ a staggered difference-in-difference-in-difference framework where we interact the treatment variable and the group split. This approach is very similar to the event studies but allows for easier testing of differences in outcomes between groups.

$$
\begin{aligned}
y_{i,t} &= \alpha + \sum_h \beta_h 1\{Treat = h\} + \sum_g \gamma_g 1\{Group = g\} + \sum_h \sum_g \omega_{h,g} 1\{Treat = h\} * 1\{Group = g\} \\
&\quad + \mu_i + X'_{i,t}\Gamma + \theta_s + \sum_g \eta_g 1\{Group = g\} * \left(X'_{i,t} + \theta_s \right) + \epsilon_{i,t}
\end{aligned}
$$

where:

$$
h = \begin{cases} 0 \text{ if } t \le Adopt_i - 2 \\ 1 \text{if } t = Adopt_i \text{ OR } t = Adopt_i - 1 \\ 2 \text{if } t = Adopt_i + 1 \\ 3 \text{if } t = Adopt_i + 2 \\ 4 \text{if } t = Adopt_i + 3 \end{cases}
$$

Firms are split into groups g based on their characteristics of interest (having a Board member with GPT background, having a female Board member, early or late or non-adopters). Our coefficients of interest are $\omega_{h,g}$ which trace the difference in profitability of firms in group g between post-adoption and pre-adoption periods, relative to the control group. We start the GPT and female regressions in 2015 and the late versus early adoption regression in 2011 given we do not model longer lags so have less need for prior information.[31]

Results shown in Table F1 confirm the event study findings that suggest the role of GPT background and female representation on the Board for firm profitability following adoption. Similarly, there is a significant difference between outcomes for early and late adopters post-adoption.

31 Results are quite similar if we match the earlier specification periods exactly. However, we preferred this sample given the pooling of past periods in this specification.

Table F1: Difference in Post-adoption Outcomes by Firm, Board and Period

Using staggered difference-in-difference-in difference, relative to alternate cohort

	Firms with GPT experience, relative to none	Late adopters, relative to early
At adoption	−1.798	3.880
	(2.142)	(3.223)
Post-adoption		
One year	3.430	4.496
	(3.334)	(3.101)
Two years	7.098*	5.606*
	(3.951)	(3.304)
Three years plus	1.671	5.924
	(4.890)	(3.187)
Number of observations	7,555	11,598

Notes: Firms in the IT sector are excluded. *, **, *** indicate significance at the 10, 5 and 1 per cent level, respectively. Standard errors are shown in parentheses and are clustered at the firm level. First two columns focus on late adopters.

Sources: Authors' calculations; Morningstar; Refinitiv; S&P Capital IQ.

Appendix G: Robustness Checks Using Firm Size

Figure G1: Return on Assets around GPT Adoption by Firm Size

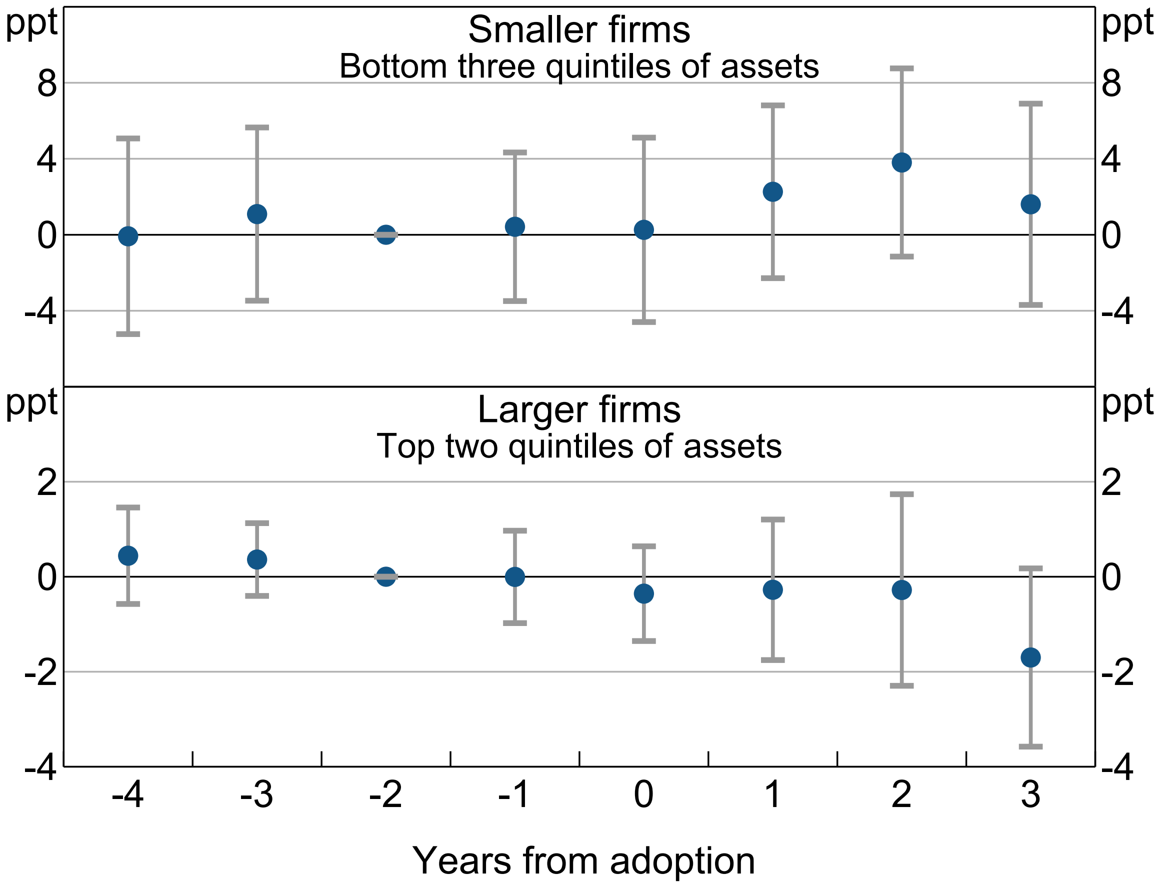

Notes: Late adopter sample. Dots represent point estimates. Whiskers indicate 90 per cent confidence intervals.

Sources: Authors' calculations; Morningstar; Refinitiv.

Appendix H: Robustness to Other Profitability Measures

Figure H1: Return on Assets around GPT Adoption

Net of capital expenditure

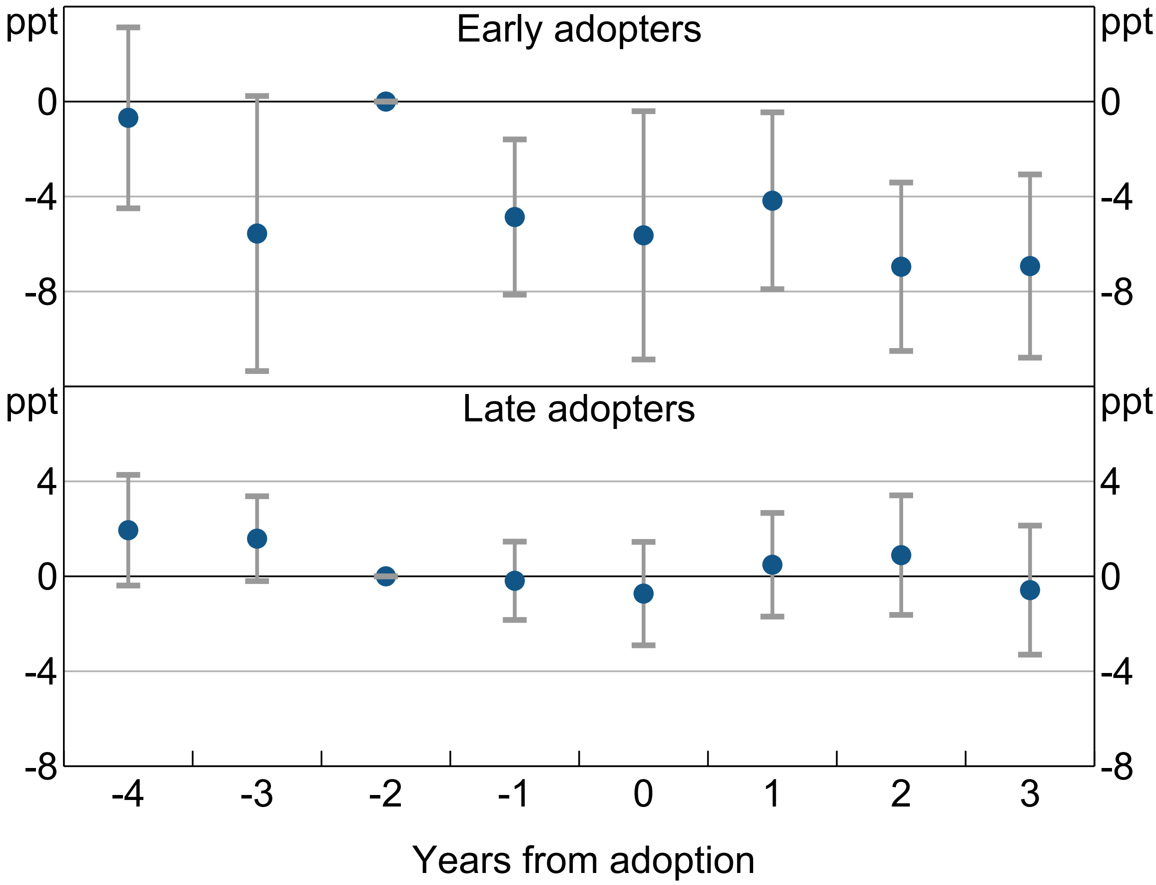

Notes: Dots represent point estimates. Whiskers indicate 90 per cent confidence intervals.

Sources: Authors' calculations; Morningstar; Refinitiv.

Figure H2: EBIT around GPT Adoption

As a share of assets

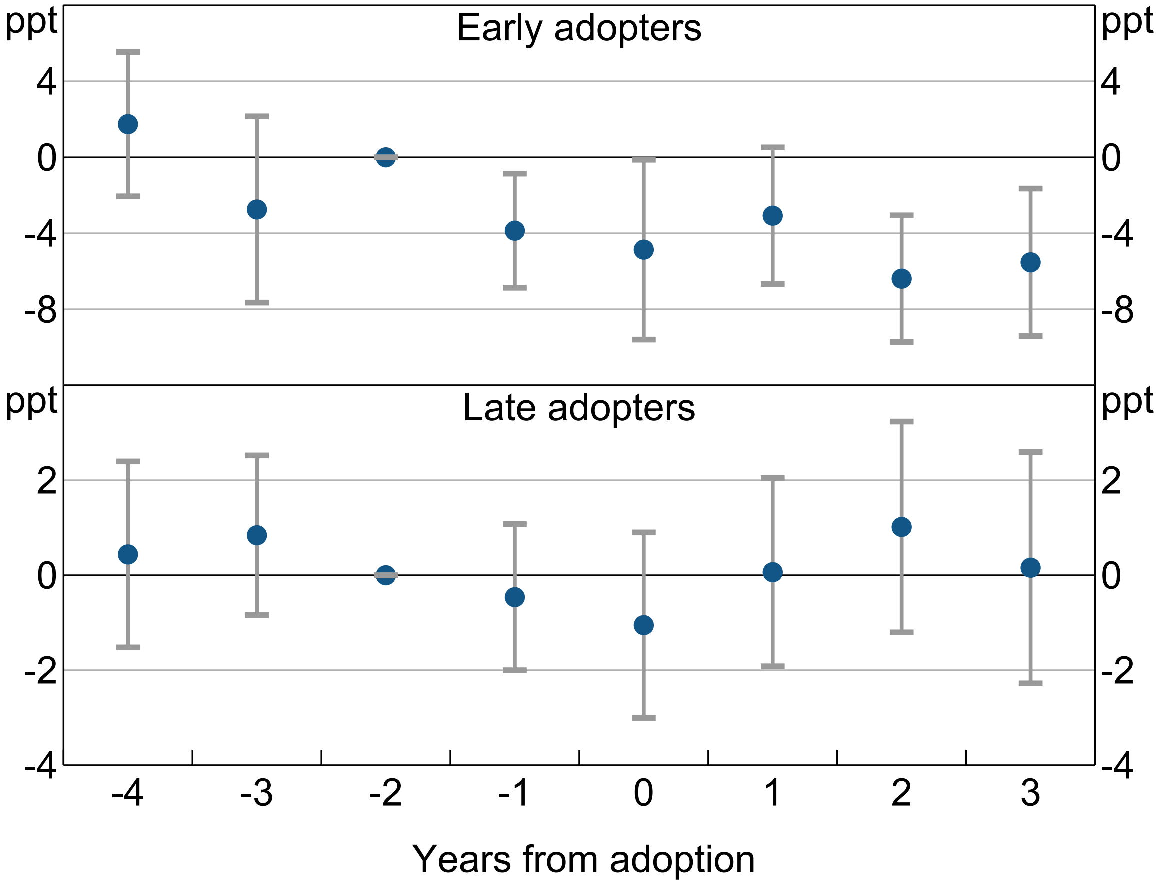

Notes: Dots represent point estimates. Whiskers indicate 90 per cent confidence intervals.

Sources: Authors' calculations; Morningstar; Refinitiv.

Appendix I: Other Robustness Checks

Figure I1: Return on Assets around GPT Adoption

For non-mining firms

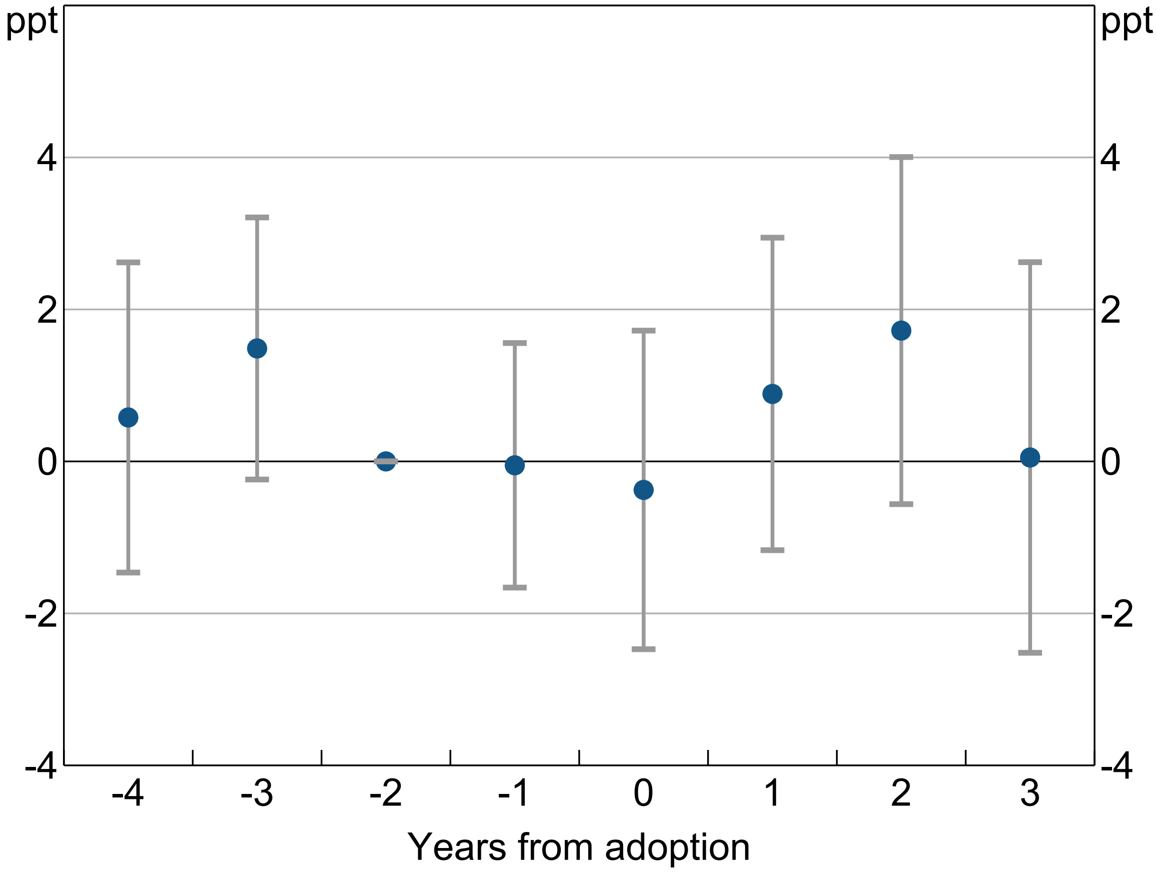

Notes: Late adopter sample. Dots represent point estimates. Whiskers indicate 90 per cent confidence intervals.

Sources: Authors' calculations; Morningstar; Refinitiv.

Figure I2: Return on Assets around GPT Adoption

Removing 2021 onwards from sample

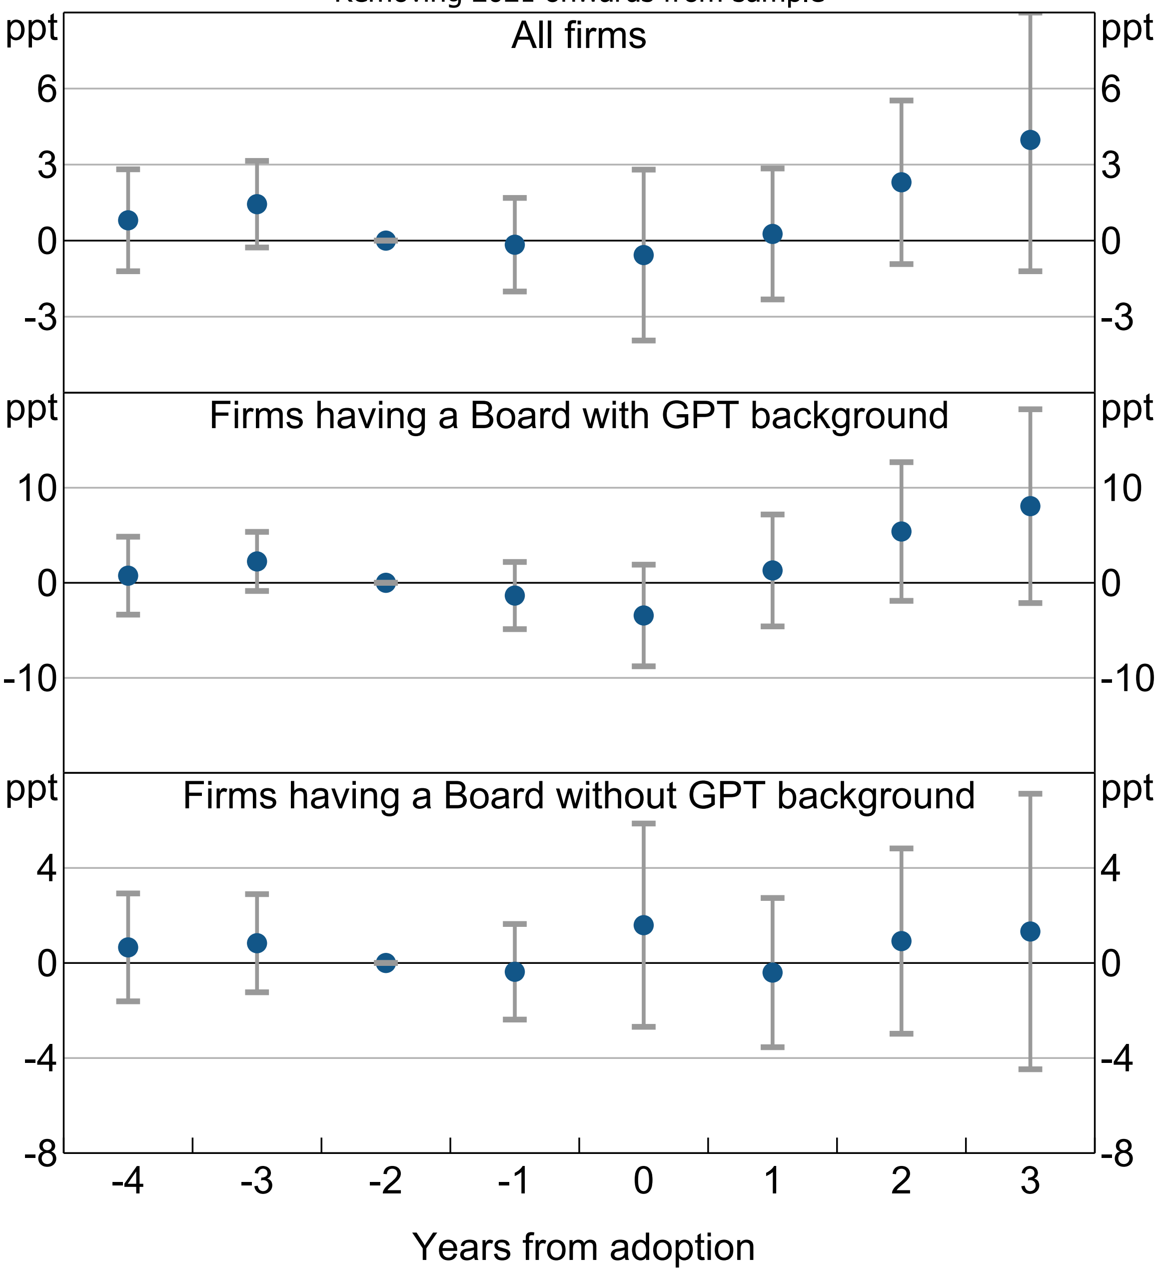

Notes:	Late adopter sample. Dots represent point estimates. Whiskers indicate 90 per cent confidence intervals.
Sources:	Authors' calculations; Morningstar; Refinitiv; S&P Capital IQ.

Appendix J: Wooldridge estimator

Figure J1: Return on Assets around GPT Adoption

Wooldridge estimator

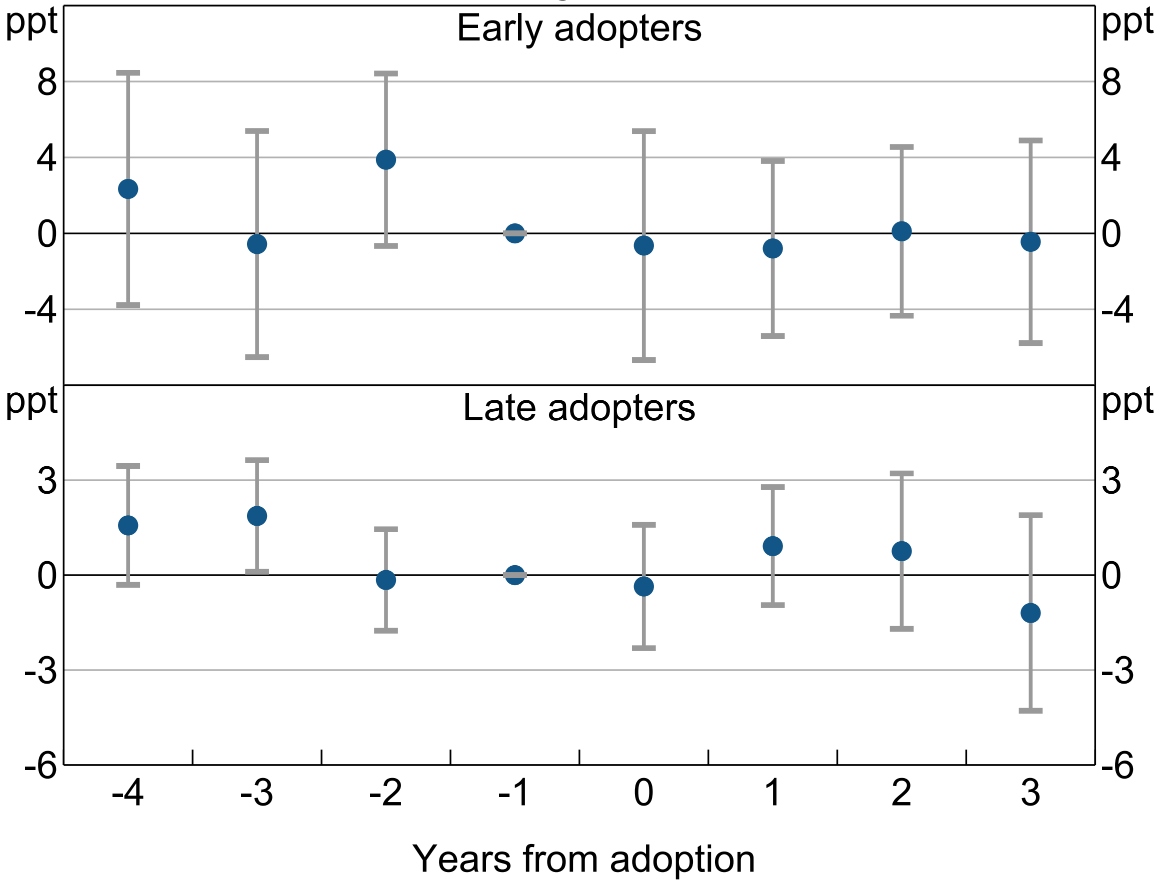

Notes: Dots represent point estimates. Whiskers indicate 90 per cent confidence intervals.

Sources: Authors' calculations; Morningstar; Refinitiv; S&P Capital IQ.

Figure J2: Return on Assets around GPT Adoption by Board Skills

Wooldridge estimator

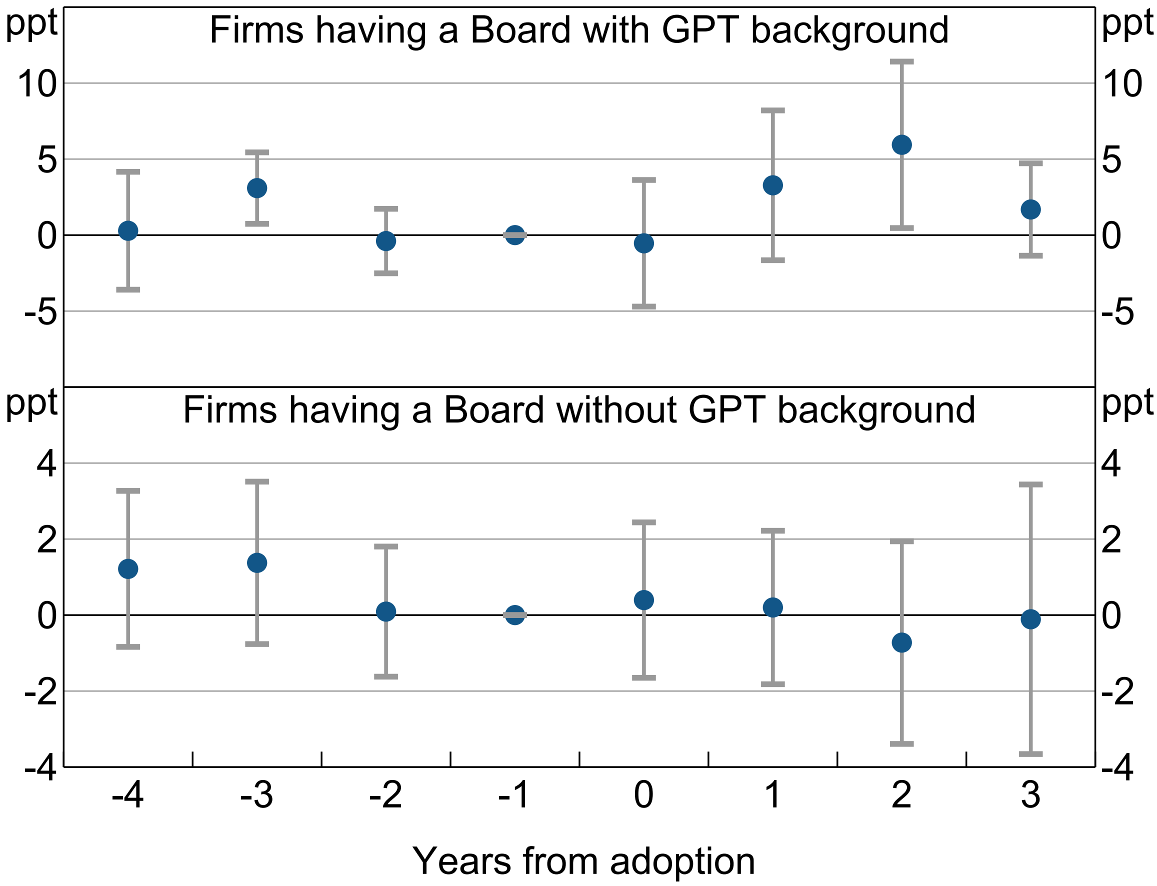

Notes: Late adopter sample. Dots represent point estimates. Whiskers indicate 90 per cent confidence intervals.

Sources: Authors' calculations; Morningstar; Refinitiv; S&P Capital IQ.

Figure J3: Hire Tech Skills around GPT Adoption

By Board skills, Wooldridge estimator

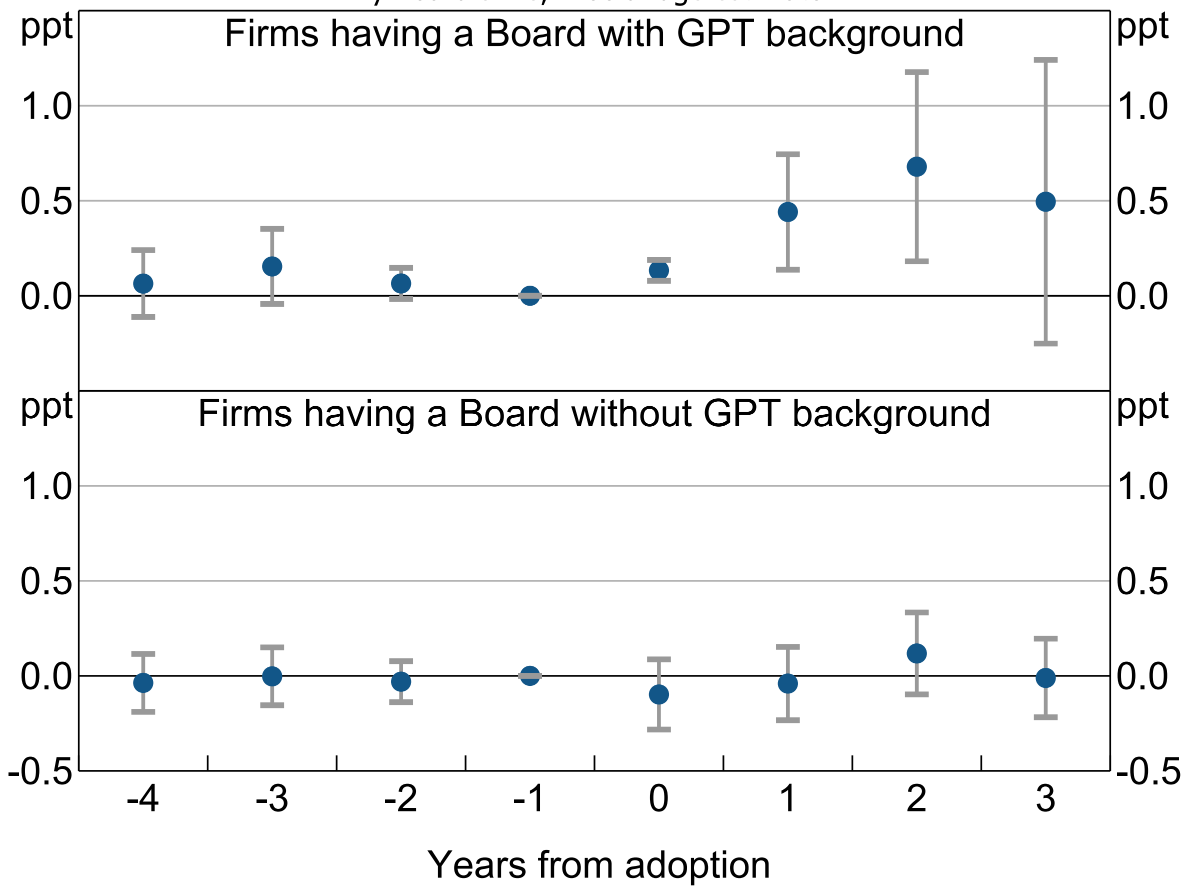

Notes: Late adopter sample. Dots represent point estimates. Whiskers indicate 90 per cent confidence intervals.

Sources: Authors' calculations; Bahar and Lane (2022) using Lightcast data; Morningstar; Refinitiv; S&P Capital IQ.

References

Acemoglu D (2025), 'The Simple Macroeconomics of AI', *Economic Policy*, 40(121) pp 13-58.

Acemoglu D, G Anderson, D Beede, C Buffington, E Childress, E Dinersoz, L Foster, N Goldschlag, J Haltiwanger, Z Kroff, P Restrepo and N Zolas (2023), 'Advanced Technology Adoption: Selection or Causal Effects?', *AEA Papers and Proceedings*, 113, pp 210-2014

Acemoglu D, D Autor, J Hazell and P Restrepo (2022), 'Artificial Intelligence and Jobs: Evidence from Online Vacancies', *Journal of Labor Economics,* 40(51), pp 294-340.

Agrawal A, JS Gans and A Goldfarb (2023), 'Similarities and Differences in the Adoption of General Purpose Technologies', in S Basu, L Eldridge, J Haltiwanger and E Strassner (eds), *Technology, Productivity, and Economic Growth*, Studies in Income and Wealth, University of Chicago Press, Chicago.

Alekseeva L, M Giné, S Samila and B Taska (2020), 'AI Adoption and Firm Performance: Management versus IT', Available at SRN: https://ssrn.com/abstract=3677237 or http://dx.doi.org/10.2139/ssrn.3677237

Alekseeva L, J Azar, M Giné, S Samila and B Taska (2021), 'The Demand for AI Skills in the Labor Market', *Labour Economics*, 71, Article 102002.

Andrews D, J Hambur, D Hansell and A Wheeler (2022), 'Reaching for the Stars: Australian Firms and the Global Productivity Frontier', Australian Treasury Working Paper No 2022-01.

Babina T, A Fedyk, A He and J Hodson (2023), 'Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition', NBER Working Paper No 31325.

Babina T, A Fedyk, A He and J Hodson (2024), 'Artificial Intelligence, Firm Growth, and Product Innovation', *Journal of Financial Economics*, 151, Article 103745.

Bahar E and O Lane (2022), 'How Dispersed are New Technologies in the Australian Job Market?', *Treasury Round Up*, October, pp 21–27.

Bayrak E, JP Conley and S Wilkie (2011), 'The Economics of Cloud Computing', *The Korean Economic* Review, 27(2), pp 203–230.

Bloom M, E Brynjolfsson, L Foster, R Jarmin, M Patnaik, I Saporta-Ektsen and J Van Reenen (2019), 'What Drives Differences in Management Practices', *American Economic Review*, 109(5), pp 1648-1683.

Bloom N, TA Hassan, A Kalyani, J Lerner and A Tahoun (2021), 'The Diffusion of Disruptive Technologies', Centre for Economic Performance, CEP Discussion Paper No 1798.

Bloom N, R Sadun and J Van Reenen (2012), 'Americans do IT Better: US Multinationals and the Productivity Miracle', *American Economic Review*, 102(1), pp 167-201.

Bloom N, R Sadun and J Van Reenen (2017), 'Management as Technology', NBER Working Paper No 22327, October.

Burgess J and J Connell (2020), 'New Technology and Work: Exploring the Challenges', *The Economic and Labour Relations Review*, 31(3), pp 310–323.

Brynjolfsson E, D Rock and C Syverson (2019), 'Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics', in *The Economics of Artificial Intelligence: An Agenda*, eds. A Agarwal, J Gans, and A Goldfarb, 23-60. Chicago: University of Chicago Press.

Brynjolfsson E, W Jin and K McElheran (2021), 'The power of prediction: predictive analytics, workplace complements, and business performance', *Business Economics*, 56, pp 217-239.

Calvino F and L Fontanelli (2023a), 'A Portrait of AI Adopters across Countries: Firm Characteristics, Assets' Complementarities and Productivity', OECD Science, Technology and Industry Working Papers No 2023/02.

Calvino F and L Fontanelli (2023b), 'Artificial intelligence, complementary assets and productivity: evidence from French firms', LEM Working Paper No. 2023/35..

Calvino F and L Fontanelli (2026), 'AI Users Are Not All Alike: The Characteristics of French Firms Buying and Developing AI', *Resaerch Policy*, 55(5).

Calvino F, L Samek, M Squicciarini and C Morris (2022), 'Identifying and Characterising AI Adopters: A Novel Approach on Big Data', OECD Science, Technology and Industry Working Papers No 2022/06.

Crafts N (2021), 'Artificial Intelligence as a General-purpose Technology: An Historical Perspective', *Oxford Review of Economic Policy,* 37(3), pp 521–536.

Czarnitzki D, GP Fernandez and C Rammer (2023), 'Artificial intelligence and firm-level productivity', *Journal of Economic Behaviour and Organization,* 211, pp 188-205.

Da Silva L, A Rincon-Aznar and F Venturini (2024), 'Productivity performance, distance to frontier and AI innovation: Firm-level evidence from Europe', *Journal of Economic Behaviour and Organization,* 228, pp 1-25.

DeStefano D, Kneller R and Timmis J (2025), 'Cloud Computing and Firm Growth', *The Review of Economics and Statistics*, 107(6), pp 1638-1651.

Department of Finance (2022), 'Cloud Computing Arrangements', accessed April 2023. Available at <https://www.finance.gov.au/publications/resource-management-guides/accounting-internally-developed-software-and-cloud-computing-arrangements-rmg-109/cloud-computing-arrangements>.

DISER (Department of Industry, Science, Energy and Resources) (2021), 'Innovation Activity', *Australian Innovation System Monitor*, accessed April 2023. Available at <https://www.industry.gov.au/sites/default/files/minisite/static/e809cbb0-a803-4827-a45b-51598ba272b2/australian-innovation-system-monitor/business-innovation/innovation-activity/index.html>.

***Economist, The* (2023),** 'Chief Executives Cannot Shut up about AI', *The Economist* online, 3 June, accessed June 2023. Available at <https://www.economist.com/business/2023/06/01/chief-executives-cannot-shut-up-about-ai>.

EmadEldeen, R, AF Elbayoumi, MAK Basuony and EKA Mohamed (2021), 'The Effect of the Board Diversity on Firm Performance: An Empirical Study on the UK', *Corporate Ownership & Control*, 18(3, Special Issue), pp 337–347.

Etro F (2010), 'The Economics of Cloud Computing', VoxEU Column, 25 February, accessed April 2023. Available at https://cepr.org/voxeu/columns/economics-cloud-computing.

Fernando J, K McLoughlin and R Ratnayake (2025), 'Technology Investment and AI: What are Firms Telling US?', RBA *Bulletin*, November.

Fontanelli L, Gueini M, Miniaci R and A Secchi (2025), 'Predictive AI and productivity growth dynamics: evidence from French firms', Fondazione Eni Enrico Mattei Working paper 11.2025.

Goodman-Bacon A (2019), 'So You've Been Told to Do My Difference-in-differences Thing: A Guide', Unpublished manuscript, Vanderbilt University.

Gordini N and E Rancati (2017), 'Gender Diversity in the Italian Boardroom and Firm Financial Performance', *Management Research Review*, 40(1), pp 75–94.

Hajikhani A, A Jager A, T Schubert and A Suominuen (2025), 'Return of the Solow-paradox in AI?: AI-adoption and firm productivity', *Mimeo.* Available at *https://doi.org/10.21203/rs.3.rs-7434983/v1*

Harrigan J, A Reshef and F Toubal (2021), 'The March of the Techies: Job Polarization Within and Between Firms', *Resaerch Policy*, 50.

Hsieh TS, JB Kim, RR Wang and Z Wang (2022), 'Educate to innovate: STEM directors and corporate innovation', *Journal of Business Resaerch*, 138, pp 229-238.

Jin W and J Bai (2022), 'Cloud Adoption and Firm Performance: Evidence from Labour Demand', Available at SSRN https://dx.doi.org/10.2139/ssrn.4082436

Loucks J (2018), 'Artificial Intelligence: From Expert-only to Everywhere', in 'Technology, Media, and Telecommunications Predictions 2019', Deloitte Insights, pp 14–23.

Makhlouf R (2020), 'Cloudy Transaction Costs: A Dive into Cloud Computing Economics', *Journal of Cloud Computing*, 9, Article 1.

McElheran K, JF Li, E Brynjolfsson, Z Kroff, E Dinlersoz, L Foster and N Zolas (2025), 'AI adoption in America: Who, what, and where', *Journal of Economics & Management Strategy*, 33(2), pp 375-415.

McElheran K, MJ Yang, E Brynjolfsson, Z Kroff (2025), 'The Rise of Industrial AI in America: Microfoundatons of the Productivity J-curve(s)', Available at SSRN https://dx.doi.org/10.2139/ssrn.5036270

McMillan H, T Murray, C de Fontenay and R Lattimore (2022), 'Head in the Cloud: Firm Performance and Cloud Service', Paper presented at the 2022 ABS and RBA Joint Conference on 'Economic Implications of the Digital Economy', 9–10 March, Sydney.

Nguyen K (2022), 'The Real Effects of Debt Covenants: Evidence from Australia', RBA Research Discussion Paper No 2022-05.

Nguyen KT (Forthcoming), 'Trust and Innovation witin the Firm: Evidence from Matched CEO-Firm Data', *Quarterly Journal of Economics*.

Nicoletti G, C von Rueden and D Andrews (2020), 'Digital technology diffusion; A matter of capabilities, incentives or both?', *European Economic Review*, *128,* pp 1-37.

Productivity Commission (2023), *5-year Productivity Inquiry: Innovation for the 98%*, Vol 5, Inquiry Report No 100, Productivity Commission, Canberra.

Roth J, PHC Sant'Anna, A Bilinski and J Poe (2023), 'What trending in difference-in-differences? A synthesis of the recent economic literature', *Journal of Econometrics*, 235(2), pp 2218-2244.

Sautner Z, L van Lent, G Vilkov and R Zhang (2023), 'Firm-level Climate Change Exposure', *Journal of Finance*, 78(3), pp 1449–1498.

Scur D, R Sadun, J Van Reenen, R Lemo and N Bloom (2021), 'The World Management Survey at 18: lessons and the way forward', *Oxford Review of Economic Policy,* 37(2), pp 231-258.

Sun L and S Abraham (2021), 'Estimating dynamic treatment effects in event studies with heterogenous treatment effects', *Journal of Econometrics*, 225(2), pp 175-199

West D (2022), 'Weighing up the Cost of Disconnected Digital Business Experiences', MYOB online blog, 9 June, accessed April 2023. Available at <https://www.myob.com/au/blog/digital-disconnection-costs-to-business/>.

Wilson M, R Agarwal, W Li and C Bajada (2023), *Dynamic Capabilities: How Australian Firms Can Survive and Thrive in Uncertain Times*, Committee for Economic Development of Australia, Melbourne.

Wooldridge J (Forthcoming), 'Two-way fixed effects, the two-way Mundlak regression, and difference-indifferences estimators', *Empirical Economics*

Zheng M, R Huang, X Wang and X Li (2023), 'Do firms adopting cloud computing technology exhibit higher future performance? A textual analysis approach', *International Review of Financial Analysis*, 90, pp 1-17.

3. Chapter 3: Shifts in Australian Price-setting Behaviour Around Large Shock

Matthew Fink* and Jonathan Hambur**

* Economic Research Department, Reserve Bank of Australia (RBA)**Economic Research Department, Reserve Bank of Australia (RBA), Crawford School of Public Policy, Australian National University (ANU)

Abstract

The sharp rise in inflation following the COVID-19 pandemic has renewed interest in how firms adjust prices after large economic shocks and the implications for modelling inflation and setting monetary policy. Using a large dataset of web-scraped Australian retail prices, we document an increase in the frequency of price changes in 2022 and 2023, alongside strong goods price inflation. We incorporate these microdata-based estimates of price-setting frequency into the Reserve Bank of Australia's dynamic stochastic general equilibrium model to assess their macroeconomic implications. We find that failing to account for higher rates of price adjustment during the high-inflation period leads to inflation forecasts that are up to 1.2 percentage points too low, even when the underlying shocks are known. The increase in the frequency of price resets also steepens the Phillips curve, reducing the policy trade-off between inflation and output. Given knowledge of this change in price-setting behaviour, a hypothetical central bank with unchanged preferences would tend to raise interest rates more aggressively than in a scenario where price rigidity was stable. Our findings highlight the importance of accounting for shifts in price-setting behaviour when interpreting inflation and setting monetary policy.

JEL Classification Numbers: E31; E52

Keywords: Inflation; pries microdata; price setting; optimal policy

Thanks to Erwan Gautier, Christopher Gibbs and Matthias Klein for their helpful feedback on this work, and to seminar participants at the RBA, University of Melbourne, University of Queensland, University of Sydney, Australian National University and the e61/UNSW Macroeconomic Policy Workshop. Thanks to Omer Majeed for making significant contributions to the initial setup and exploration of our dataset.

3.1 Introduction

The sharp rise in inflation in 2022 and 2023 has renewed interest in the microeconomic foundations of price-setting behaviour, particularly how firms respond to large economic shocks. A growing body of international empirical research suggests that large supply shocks may increase the frequency of price adjustments, reducing the degree of price rigidity. In such periods, firms appear to pass through cost changes more rapidly than in normal times. This challenges a common assumption in standard macroeconomic models that the frequency of firms' price changes is invariant to economic conditions. Under this assumption, the pass-through of cost shocks to final prices is stable and the inflationary effects of large shocks are simply scaled-up versions of small ones. If price setting instead becomes more flexible after large shocks, inflation may respond more strongly and faster than conventional frameworks predict.

This has direct implications for policymakers. Central banks and other macroeconomic policy institutions typically rely on models that assume stable price-setting behaviour when forecasting inflation and calibrating policy. If this assumption breaks down during large shocks – such as those seen during and after the COVID-19 pandemic, or potentially in the context of future geopolitical or climate-related disruptions – forecasts may be systematically biased and policy responses miscalibrated. Quantifying how price-setting behaviour changes in such episodes is therefore critical for improving inflation forecasts and guiding policy decisions. It also informs how price rigidity should be modelled more generally, which matters for a broader set of questions, including how costly inflation is for the economy and households.

In this paper, we examine how price rigidity in Australia changed between 2018 and 2023, using a novel dataset of item-level prices collected from the websites of around 60 large retailers. We focus particularly on the inflation surge in 2022 and 2023 (Graph 3.1), documenting shifts in the frequency of price changes during this time. These findings are incorporated into the Reserve Bank's benchmark dynamic stochastic general equilibrium (DSGE) model to assess implications for inflation dynamics, policy transmission and policy trade-offs.

Graph 3.1: Consumer Price Inflation

Year-ended

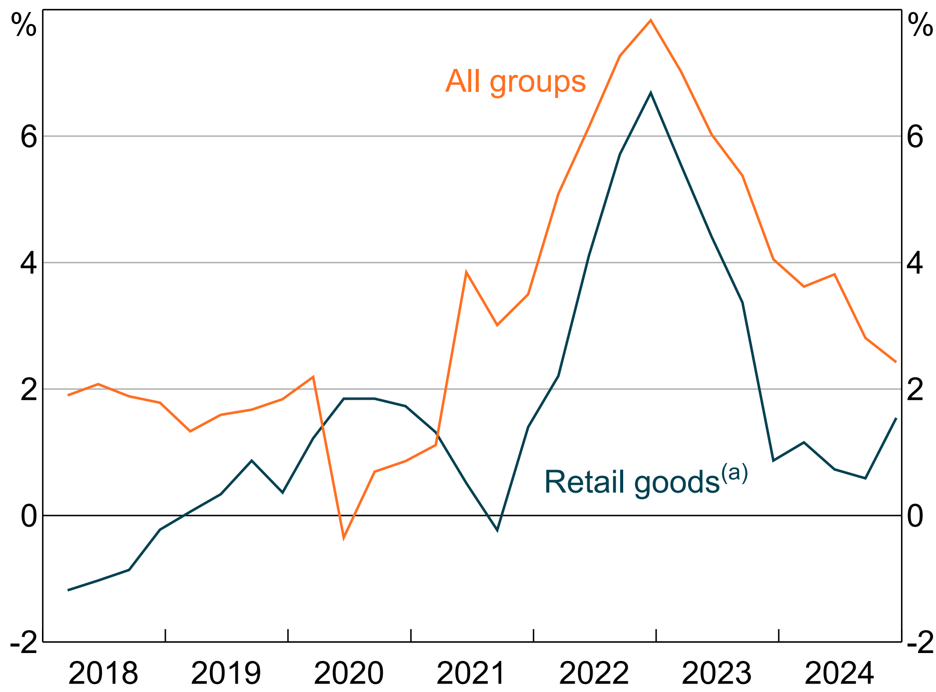

Note: (a) ABS CPI goods component excluding items not in web-scraped dataset: food & non-alcoholic beverages, tobacco, new dwellings, motor vehicles and automotive fuel.

Sources: ABS; Authors' calculations; RBA.

Our results suggest that price-setting behaviour became significantly less flexible at the peak of COVID-19 economic uncertainty than it was before COVID-19. But as inflation surged in 2022 and 2023, prices became more flexible, with firms adjusting prices more frequently than they did before the pandemic. This 'state-dependent' rather than 'time-dependent' pricing – in which firms adjust prices in response to economic conditions rather than on a fixed schedule – implies that large shocks may have disproportionately larger effects on inflation (i.e. that the inflationary effects of a shock are nonlinear in the size of the shock).

We then look to explore how important these changes could be, leavarging existing tools available to central banks. In particular, using the RBA's DSGE model, we find that failing to account for changes in price-setting frequency during the 2022–2023 inflation surge would have led to a substantial underprediction of inflation. Our estimates suggest that the decline in rigidity could explain up to 1.2 percentage points of inflation in this period, meaning that changes in price-setting behaviour accounted for a sizeable share of the increase in prices. To better capture these dynamics when the economy faces large shocks in the future, forecasters should consider using models that allow for shifts in price-setting frequency, such as 'menu cost' models, or apply explicit judgement when using standard models that assume frequency is constant.

Monetary policymakers typically face a trade-off between lowering high inflation and supporting economic growth and employment, particularly in the wake of a supply shock, because reducing inflation in these circumstances requires curbing demand and activity to create slack and ease price pressures. In assessing the implications of our findings for monetary policy, we find that reduced price rigidity lessens the central bank's trade-off between stabilising inflation and real activity during large supply shocks. In this environment, policymakers can lower inflation with fewer real economic costs, suggesting a stronger focus on inflation stabilisation is warranted compared to a situation in which rigidity did not change.

We quantify these effects in the post-COVID high-inflation period, using a so-called 'optimal control' exercise. This exercise uses the RBA's DSGE to calculate the policy path that would have minimised deviations of inflation and unemployment from target, given the shocks that hit the economy, the structure of the model, and some set of simple fixed preferences for trading off inflation and unemployment in the face of a supply shock. We compare these policy paths with and without the change in rigidity.

Relative to a scenario where rigidity remained unchanged, this exercise calls for policy that raises interest rates more aggressively – by up to around 40 basis points more than otherwise – followed by a sharper subsequent easing. Importantly, this does not imply that actual rates should have been higher than they were at the time. Our work compares these 'optimal' policy paths under different assumptions about price rigidity, rather than against the realised cash rate path. Moreover, the exercise uses a specific set of policy preferences around the trade-off between inflation and unemployment, which may differ from those used by policymakers during the high-inflation episode in 2022 and 2023. And the exercise relies on the structure of the model, which by its nature is a simplification of reality. Rather than prescribing a specific policy stance, our findings suggest that the observed decline in price rigidity would have made a front-loaded tightening more appealing for a central bank having to trade off inflation and unemployment, relative to the case where rigidity had not changed.

Beyond shedding light on this high-inflation episode, our paper contributes to the broader economic literature in several ways. First, it adds to the growing international empirical evidence on variability in price-setting frequency and nonlinear inflation dynamics. Our dataset covers a broader range of products than much of the existing literature using web-scraped data, including many discretionary and non-perishable items that may be less flexibly priced in normal times. Second, we are the first in this literature to apply formal duration modelling approaches to estimate price-setting frequency. These offer advantages over simpler 'share of prices changing' methods. Finally, we contribute to the relatively sparse literature on price-setting dynamics in Australia, focusing on a more recent period than previous studies and using web-scraped Australian data for the first time.

The objective of our analysis is to assess whether Australian firms' price-setting behaviour shifted over time within our dataset, and if so to explore the policy implications of those shifts. The remainder of the paper proceeds as follows. Section 3.2 reviews the existing literature in this area. Section 3.3 introduces our dataset. Section 3.4 introduces our preferred rigidity measure, price duration, and presents a time-varying measure estimated using a regression-based survival analysis approach that maps to the well-known Calvo parameter. Sections 3.5 and 3.6 explore the implications of changing price rigidity for inflation dynamics and monetary policy. Section 3.7 concludes with a discussion.

3.2 Literature review

Modern macroeconomic models of the business cycle typically incorporate price rigidity to explain how fluctuations in nominal economic variables affect the real economy over time. A widely used framework is the Calvo pricing model, which assumes that only a fixed proportion of firms can adjust their prices in each period (Calvo 1983). This *frequency* of price changes—often referred to as the 'Calvo parameter'—is usually treated as constant. This results in 'time-dependent' pricing, where firms' decisions to change prices depend solely on time and not on economic conditions. Consequently, nominal rigidity is also constant over time, producing linear relationships between the size of a given economic shock and its inflationary impact.

A substantial empirical literature has emerged over time to assess price-setting dynamics, to which this paper contributes. Early studies using CPI and scanner data found that prices were more flexible than assumed in standard macroeconomic models (e.g. Nakamura and Steinsson 2008; Klenow and Malin 2010). More recent work has leveraged web-scraped data, which offers advantages such as higher frequency, broader item coverage, and fewer measurement distortions compared to official statistics. Pioneering studies by Cavallo and Rigobon (2016) and Cavallo (2018) explore these advantages and demonstrate that web-scraped data can yield more accurate estimates of price-setting behaviour.

A general finding across this literature is that the frequency of price changes varies with economic conditions, undermining the fixed rigidity assumption in the Calvo framework. Several studies have used this evidence to motivate and calibrate models with 'state-dependent' pricing, which better match observed economic behaviour (e.g. Alvarez et al. 2016). These models often imply nonlinear dynamics, where large shocks lead to increased rates of price-resetting and so have larger and faster effects on inflation.[32]

Recent papers have explored the importance of pricing dynamics in state-dependent models, particularly in the context of the pandemic (Dedola et al. 2023; Auclert et al. 2024; Cavallo, Lippi and Miyahara 2024; Blanco et al. 2024a, 2024b, 2025; Gautier et al. 2025). These studies find that price rigidity tends to decline during large shocks, resulting in stronger inflation responses than predicted by time-dependent models. For example, Auclert et al. (2024) show that while time- and state-dependent models yield similar inflation dynamics for small shocks in the euro area, state-dependent models produce markedly stronger responses when shocks are large.

The incorporation of state-dependent pricing into models also has important implications for monetary policy trade-offs. Karadi et al. (2024), for instance, show that the trade-off between lowering inflation and preserving real economic activity in the US is reduced during large shocks, as falling price rigidity steepens the Phillips curve. This implies that central banks should respond more aggressively to inflationary supply shocks than they would if price rigidity did not decrease.

32 Rotemberg pricing is another form of non-state dependent pricing commonly used in macro models, where rigidity arises non-linearly because firms face quadratic price adjustment costs (Rotemberg 1982). In log-linearised models, it results in identical predictions of pass-through as Calvo models, despite the different pricing mechanism. Non-linear Philips curves may arise through other mechanisms, such as the specific form of the demand function (Harding, Lindé and Trabandt 2022), and via labour market frictions (Benigno and Eggertsson 2023; Schmitt-Grohé and Uribe 2025).

Our paper contributes to the growing literature using prices microdata to examine price-setting shifts during and immediately after the COVID-19 period. These include Rudolf and Seiler (2022) for Switzerland, Montag and Villar (2023) for the US, Henkel et al. (2023) for euro area countries, Klein et al. (2024) for Sweden, and Bilyk et al. (2024) for Canada. Across these studies, the early pandemic phase generally saw little change in the frequency of price adjustments. In contrast, studies covering the subsequent high-inflation period consistently document a strong increase in the frequency of price changes, driven mainly by more frequent price increases. This pattern suggests a shift toward more flexible pricing in response to large shocks, consistent with state-dependent pricing behaviour.

Finally, our work builds on the limited literature examining price-setting dynamics in Australia. Cagliarini et al. (2010), using RBA firm survey data from 2000 to 2006, found that prices were less rigid than typically assumed, implying slower inflation pass-through in standard models. Sutton (2017) used CPI microdata from 2001 to 2017 and identified moderate variation in price stickiness over time, with only modest increases in flexibility around the Global Financial Crisis. Our paper extends these analyses, using a different data source—web-scraped prices—and focusing on a more recent period of heightened economic volatility with an emphasis on policy implications.

3.3 Data

3.3.1 Data structure and treatment

The most commonly used types of price microdata in empirical research include survey data underlying consumer and producer price indices, scanner data sourced from point-of-sale transactions and prices collected from firms' websites ('web-scraped' prices). A discussion of the strengths and limitations of these types of prices in a research context can be found in Cavallo (2018). The present paper uses a web-scraped dataset collected by the Australian Bureau of Statistics (ABS), which captures advertised prices for every item sold on up to 61 large Australian retailers' websites. Prices were collected every few days between 2016 and 2023, with the number of firms sampled increasing over time. The dataset becomes sufficiently large to be reliable from 2018 and in total has around 549 million price observations for 13 million items. While data collection occurred frequently across the dataset (every few days), the timing of collections is irregular.

The dataset covers a broad range of retail items, such as apparel, homewares, and electronic goods, which together represent around 25 per cent of the Australian Consumer Price Index (CPI) basket.[33] Retailers and items are de-identified, though their industry classification is known. Two key attributes of sampled firms are that: (i) they derive substantial revenue from both online and brick-and-mortar stores; and (ii) they account for a large share of total sales turnover in their respective industry subdivision — around 30 per cent across sampled firms in each quarter on average. These characteristics support two assumptions underpinning our analysis: that online

33 Items in the goods component of the CPI that are notably not part of the web-scraped dataset include food and non-alcoholic beverages, tobacco, new dwellings, motor vehicles and automotive fuel.

advertised prices generally reflect in-store prices; and that price-setting behaviour observed in the sample can be generalised to the broader retail sector.[34]

Compared to some recent studies of price-setting behaviour—such as Auclert et al. (2024) and Cavallo, Lippi and Miyahara (2024)—our dataset includes a broader range of products, particularly discretionary and durable goods. These types of goods should exhibit slower-evolving pricing and longer item lifecycles. In this sense, our data is more comparable to that used in Dedola et al. (2023) and Blanco et al. (2024a, 2024b, 2025), which cover a wider set of products and services. This allows us to explore how retail firms with pricing that is typically less flexible respond to large economic shocks, and what implications this has for inflation.

The collection method used by the ABS for our dataset ensures that the full pricing lifecycle of each item is captured—from the first day an item is listed for sale (or slightly after subject to scraping timing) to the day it is removed. This typically results in a downward-sloping price path for items over time. Full lifecycle measurement differs from CPI-based microdata, where items are substituted periodically to control for quality changes within product categories. We are unable to make such adjustments due to the lack of information on the nature of each item in the dataset. Consequently, our dataset is likely to contain a higher share of price decreases than CPI data, making it unsuitable for calculating CPI-like average inflation rate measures but well-suited for measuring the breadth of offered prices across the products firms sell and shifts in price-setting behaviour.

There is ongoing debate about whether changes in advertised (sales-inclusive) or 'regular' (non-sale) prices are more relevant for characterising price rigidity. While some researchers argue that only regular price changes reflect firms' responses to macroeconomic shocks or cost pressures, others contend that both regular and sale prices may respond to economic fundamentals, and that what ultimately matters for inflation dynamics are the prices at which consumers transact.[35] This latter view aligns with the construction of the Australian CPI, which incorporates all advertised prices, including sale prices.

Rather than taking a firm position, we present rigidity results for both advertised and regular prices. While the ABS dataset contains both, regular prices are only available from 2020 onward. To avoid relying on a short sample and ensure consistent classification of sales across firms and over time, we develop an algorithm to identify sales based on pricing patterns.[36] Specifically, we identify 'V-shaped' sales, where a price drops and returns to an equal or higher level within three months, and 'clearance' sales, where a price falls by at least 25 per cent and exits the dataset within a month. We use these imputed regular and sale price series in our analysis rather than the reported regular price data, though we find that the reported and imputed regular price data generally align well over time (Graphs A1 and A2).

---

34 These assumptions are standard in the web-scraped prices microdata literature, as well as in statistical agencies that collect online prices as a part of compiling official price indices. For example, see sections 2.31-2.32 of ABS (2018). A further discussion of principles for selecting retail firms for web-scraping-based research can be found in Cavallo and Rigobon (2016).

35 For contrasting views on the relevance of sales pricing to inflation dynamics and monetary policy analysis see Guimaraes and Sheedy (2011) and Kryvtsov and Vincent (2021).

36 This builds on the approach described in the Supplemental Appendix to Gautier et al. (2024).

Importantly, while we will show that levels of measured price rigidity in our dataset may differ substantially depending on whether sales are included, *changes* in rigidity over time in our dataset are similar across the two price types. This means that questions around the inclusion of sales do not materially affect our conclusions about the impact of shifts in price rigidity on inflation dynamics and policy decision-making.

Our ability to explore heterogeneity within the dataset is limited. Information on the nature of each product is not currently provided, limiting the ability to split by item types. And while we have data on the firm industry, the sample number of firms limits our ability to split the sample by retailer type while still maintaining firm anonymity.

A further description of our data, cleaning methodology and sales algorithm is provided in Appendix A.

3.3.2 Summary statistics

To validate the data, we first provide some basic statistics and visualisations. Graph 3.2 shows the distribution of the log size of price changes in the data over the full sample period. Advertised prices exhibit symmetrical peaks around large changes such as ±25, ±30, and ±40 per cent, alongside a central concentration of smaller changes between -5 and +10 per cent (Graph 3.2, left panel). Disaggregating into regular and sale prices reveals that most regular price changes fall between -5 and +10 per cent (centre panel), while sale price changes are typically larger, clustering around common discount rates such as '25 per cent off' (right panel).[37] The symmetry in sale price changes reflects items going on sale temporarily and subsequently returning to full price. As noted above, there is a relatively large tail of negative price changes, compared to what we might see if using CPI microdata. This reflects that the web-scraped prices capture the full lifecycle from when items go on sale to when they are retired.

37 See Appendix B for charts showing how the distribution of the size of price changes has evolved over time.

Graph 3.2: Distribution of Log Price Changes by Size
2018–2023

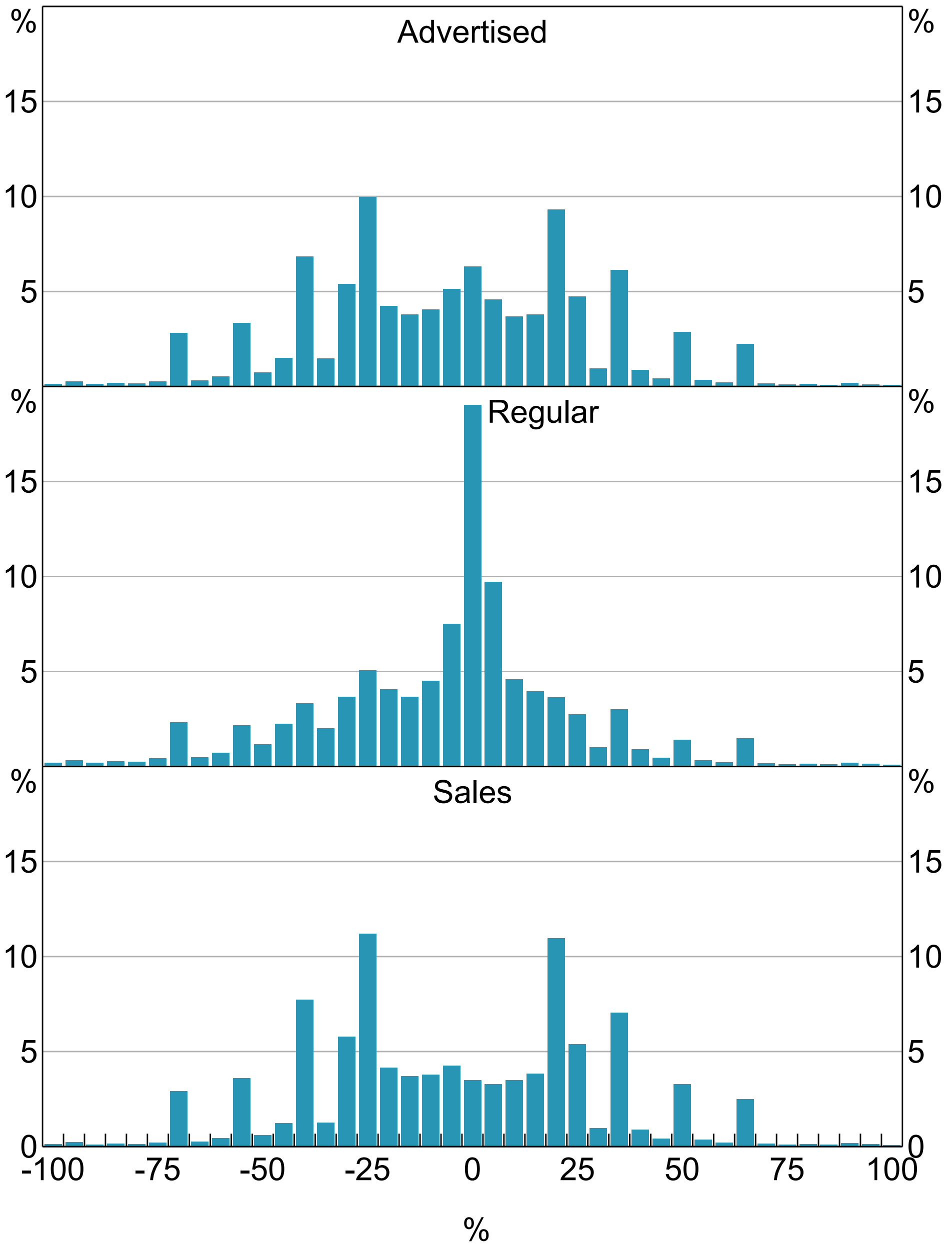

Notes: Item-weighted. x-axes are segmented in 5 percentage point intervals, where, for example, the zero bucket is changes between 0 and 5 per cent. Regular and sales prices are imputed using an algorithm that identifies sales patterns (see Appendix A).

Sources: ABS; Authors' calculations.

As a further validation, we look for evidence that the high inflation period for goods is captured in our data. Given the nature of the data, rather than trying to recreate a CPI growth rate or focus on the size of price changes, we follow a number of other papers and look at the share of prices increasing and decreasing over each month (e.g. Bilyk et al. 2024; Klein et al. 2024). Specifically, we take the final price observation for each item in our data in each month and calculate the difference in the price compared to the end of the previous month. We then take the separate share of price increases and decreases out of all prices in each month, applying a 12-month moving average to the resulting series to account for the highly seasonal nature of price-setting. Subtracting decreases from increases provides a 'relative frequency' measure that can be used as a gauge for inflationary pressure. This approach abstracts from issues around lifecycle pricing, as well as intra-month variation that is less relevant for observing the trend inflationary impulse (but may still be relevant in assessing price rigidity).

We find that the share of price increases rose notably over 2022 and 2023, while the share of price decreases declined (Graph 3.3).[38] This means that the relative frequency of price increases rose substantially, in line with the increase in retail goods price inflation (Graph 3.4). More generally, this metric has a close association with retail good price inflation, picking up slightly in 2020, before declining in 2021 alongside inflation. These results suggest that the dataset does effectively capture the inflationary impulse seen in the official Australian CPI in 2022 and 2023.

Graph 3.3: Frequency of Price Changes

Percentage of prices that change each month

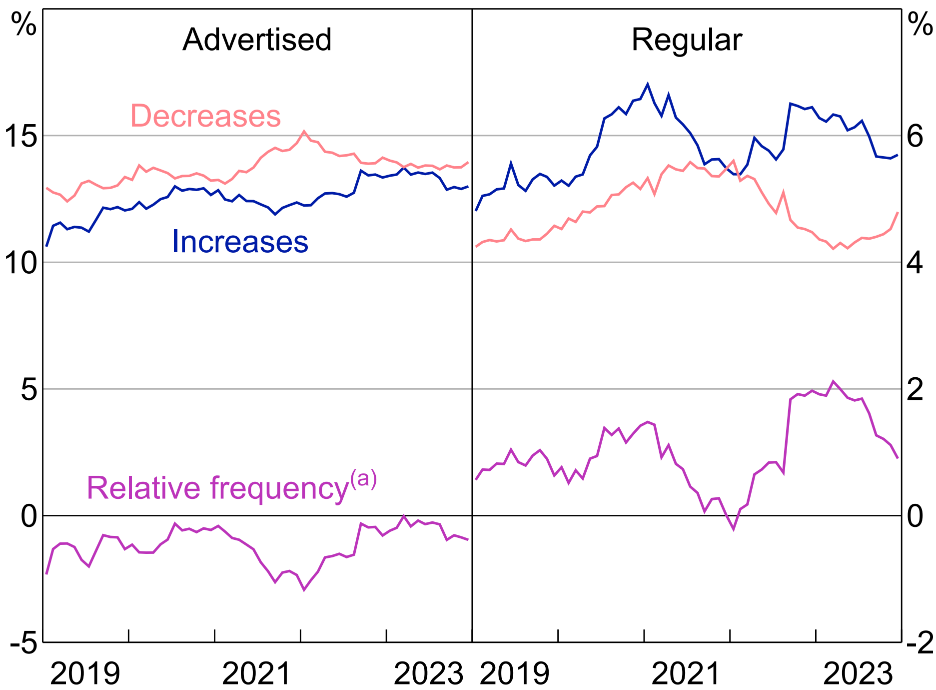

Notes: 12-month trailing averages. All series are averaged across firms using CPI weights.
(a) Relative frequency is calculated as the percentage point share of all end-of-month item prices that constitute an increase compared to the previous month less the share that constitute a decrease.

Sources: ABS; Authors' calculations.

38 In general, there tend to be more price decreases than increases for advertised retail prices, and more increases than decreases for regular retail prices. This reflects that advertised prices include clearance sales while these are excluded from regular prices.

Graph 3.4: Relative Frequency of Price Changes versus Inflation

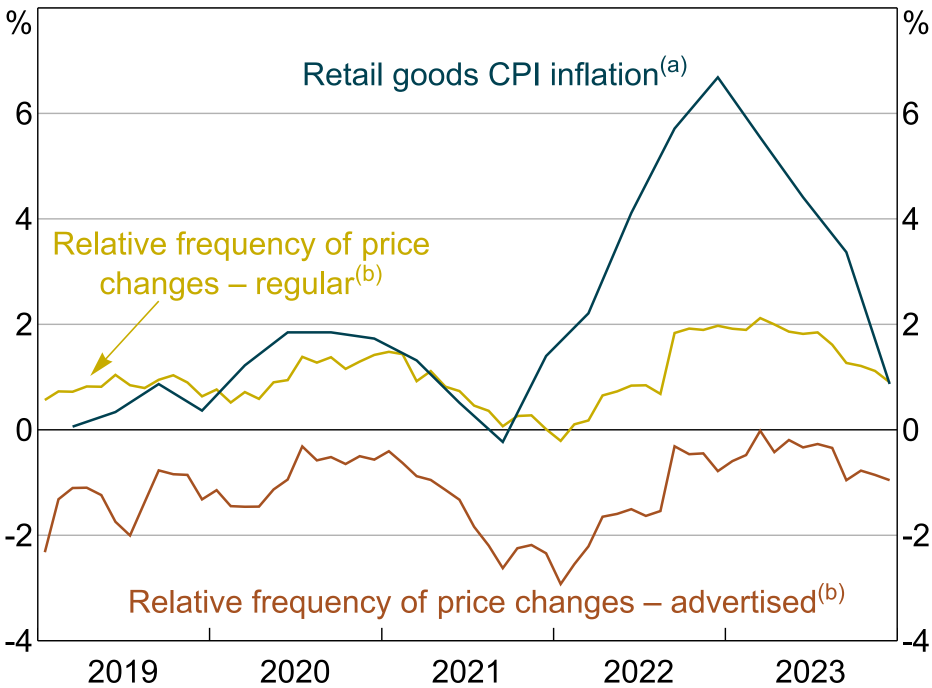

Notes: (a) Year-ended growth in the ABS CPI goods component excluding items not in web-scraped dataset: food & non-alcoholic beverages, tobacco, new dwellings, motor vehicles and automotive fuel.
(b) Relative frequency is calculated as the percentage point share of all end-of-month prices that constitute an increase compared to the previous month less the share that constitute a decrease. Shares of increases and decreases are calculated as 12-month trailing averages. Averaged across firms in the web-scraped dataset using CPI weights.

Sources: ABS; Authors' calculations.

3.4 Measuring price rigidity

Several ways of measuring price rigidity have been proposed in the literature. One is by simply looking at the share of prices that change in a month. However, this does not account for cases where prices may change multiple times in a month, which would abstract away from the richness of our high-frequency price data.[39] Moreover, such measures do not account for how long a price has remained unchanged before adjusting—an important dimension of rigidity. Another approach is to look at the kurtosis of the distribution of price changes, which have been shown to provide a sufficient statistic for parameterising macro models with state-dependent pricing mechanisms (Alvarez et al. 2016). However, kurtosis measures can be heavily affected by sample changes and noise.

Instead, we focus on measures of price duration – how many days since the price last changed. These overcome some of the above limitations and are well suited to our dataset. Moreover, they have a direct mapping to objects of interest in standard macroeconomic models of the business cycle, such as Calvo parameters.

39 Using such a measure provides less evidence of a change in rigidity, though this is not our preferred metric for the above noted reasons. An alternative would be the number of price changes each month, but due to irregular sampling frequency such measures may be biased.

3.4.1 Simple price duration measures

We start by showing the average number of days since an item last recorded a price change (conditional on a price change having occurred). Following a period of relative stability through mid-2021, advertised and regular price durations rose as goods inflation rose, with longer intervals between changes on average (Graph 3.5, top panel). These results seem counterintuitive. However, they are largely explained by compositional effects. As inflation picked up, many prices in our dataset that had remained unchanged for extended periods—sometimes several years—began to adjust. The entry of these long-lived prices into the sample of changing prices increased the mean duration, even though firms were likely changing prices *more* frequently overall.

Graph 3.5: Price Duration

Number of days since items last recorded a price change

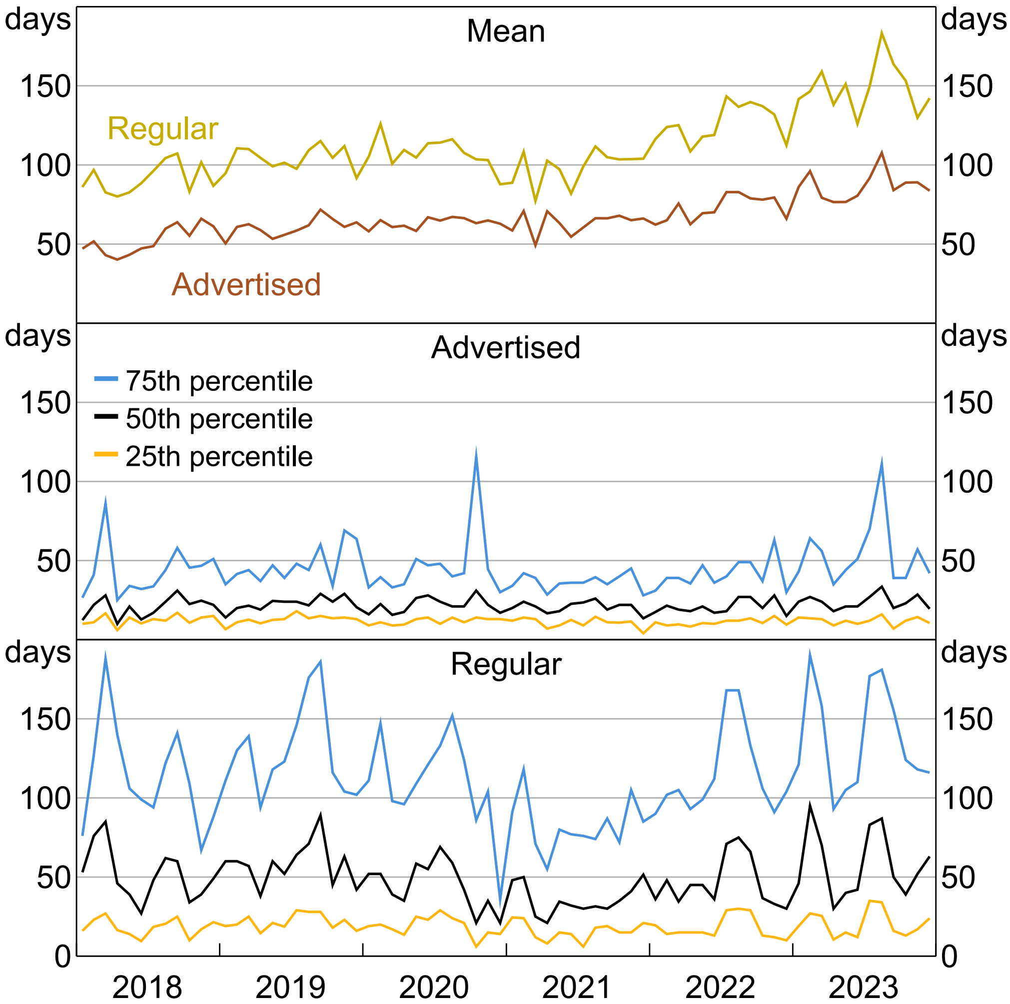

Note: Item-weighted.

Sources: ABS; Authors' calculations.

To see this, we can look at the distribution of price durations. The rise in average duration was concentrated in the upper tail of the distribution (for example, the 75th percentile and above), while the median remained relatively stable (Graph 3.5, middle and bottom panels). The fact that long-stable prices began adjusting more often provides some initial evidence that price price-setting frequency increased during the high inflation period.

3.4.2 Survival analysis methodology

To explore this more formally we use a parametric survival model to estimate the probability that a price remains unchanged over time. This approach has several advantages. First, survival analysis incorporates the full set of price spells—including prices that do not change—making it less sensitive to compositional shifts in the sample. Second, the regression-based framework allows us to formally test whether changes in rigidity over time are statistically significant. Third, survival models can control for firm-level heterogeneity and censoring (e.g. items entering or exiting the sample mid-spell).

We use a parametric survival model with an exponential survival function, which assumes a constant hazard rate—i.e. the probability of a price persisting is independent of information about how long it has been since that price last changed.[40] This assumption aligns with the Calvo pricing model commonly used in DSGE frameworks, allowing us (through a simple conversion) to interpret the estimated survival rates of prices as Calvo probabilities.

We define a 'price spell' as the number of days between when a price is set and when it next changes. The model is estimated via maximum likelihood over the full set of price spells, accounting for two key features of the data[41]:

1. Firm fixed effects: These control for persistent differences in pricing behaviour across retailers, helping isolate within-firm changes over time and abstract from changes in the composition of firms in the sample.

2. Censoring: To address left censoring (e.g. unknown price history before an item enters the sample), we exclude the first observed spell for each item. The model also handles right censoring, where items exit the sample without a price change through standard approaches.

In our baseline specification, price survival probabilities are allowed to vary by the quarter or year in which a price spell begins.[42] The estimation takes the following form:

---

40 The results of our analysis were robust to an alternative assumption that the probability distribution underlying the function was a Weibull distribution.

41 To account for data collection gaps that could be misinterpreted as long price spells, we also tested a specification that excludes price spells with gaps longer than 30 days between observations. While this adjustment affected the level of estimated rigidity, it did not materially alter the trend over time.

42 We considered allowing baseline hazard rates to vary by item, but this proved computationally intensive given the size of the dataset. We also tested a specification in which price spells were split across quarters, allowing hazard rates to reflect all periods spanned rather than just the start date. However, this approach produced volatile estimates that were sensitive to future data updates. For example, the addition of new data for 2024 could retroactively alter estimates for 2022 by changing the interpretation of spells spanning both years. Our preferred specification—anchoring hazard rates to the quarter or year in which a spell begins—offers greater stability and interpretability.

$$S_{ir}(t|\boldsymbol{x}_j) = \exp(-\lambda_r \exp(\boldsymbol{x}'_j\boldsymbol{\beta})\, t)$$

Where:

- $S_{ir}(t|\boldsymbol{x}_i)$ is the survival function giving a probability that a price remains unchanged up to time t for item *i* and retailer *r* conditional on the year or quarter $\boldsymbol{x}_j$,
- λ_r is a retailer-specific baseline hazard (rigidity) rate capturing firm fixed effects; and
- $\boldsymbol{\beta}$ is a vector of coefficients on time dummies that map to a period-specific survival rate across firms.

The survival function gives the probability that a price does not change up to a given point in time, for each year or quarter in which underlying prices were observed. To make these results comparable to the Calvo parameter commonly used in macroeconomic models—which is expressed as the probability of a price remaining unchanged over one quarter—our results are scaled to reflect the probability of prices remaining unchanged after 90 days (i.e. we set t=90).[43]

Standard errors are not clustered. But the results are robust to allowing for heteroskedasticity across, or autocorrelation within, items.

3.4.3 Survival analysis results

We estimate price rigidity separately for advertised and regular prices. As expected, regular prices are more rigid, reflecting the high prevalence of temporary sales in advertised prices. On average, the probability that a regular price remains unchanged after one quarter is 51 per cent, compared with 15 per cent for advertised prices (Graphs 3.6).[44] These estimates are broadly consistent with empirical evidence from Sutton (2017), though they imply greater flexibility than standard macroeconomic models—an observation common in the international literature.[45]

---

43 We perform the analysis across all firms and products. We are limited from exploring heterogeneity across firm industries by privacy requirements. Future analysis could be done by CPI component if item-level categorisations are made available by the data provider.

44 The results are robust to an alternative specification that assumes a Weibull distribution.

45 Importantly, differences between microdata estimates and DSGE model calibrations do not imply that macro models are incorrect. Calvo parameters in DSGE models are designed to capture aggregate price stickiness, which may reflect broader economic frictions and abstract from firm-level heterogeneity. For further discussion, see Maćkowiak and Smets (2008) and Cagliarini, Robinson and Tran (2010).

Graph 3.6 Share of Retail Prices Unchanged after One Quarter

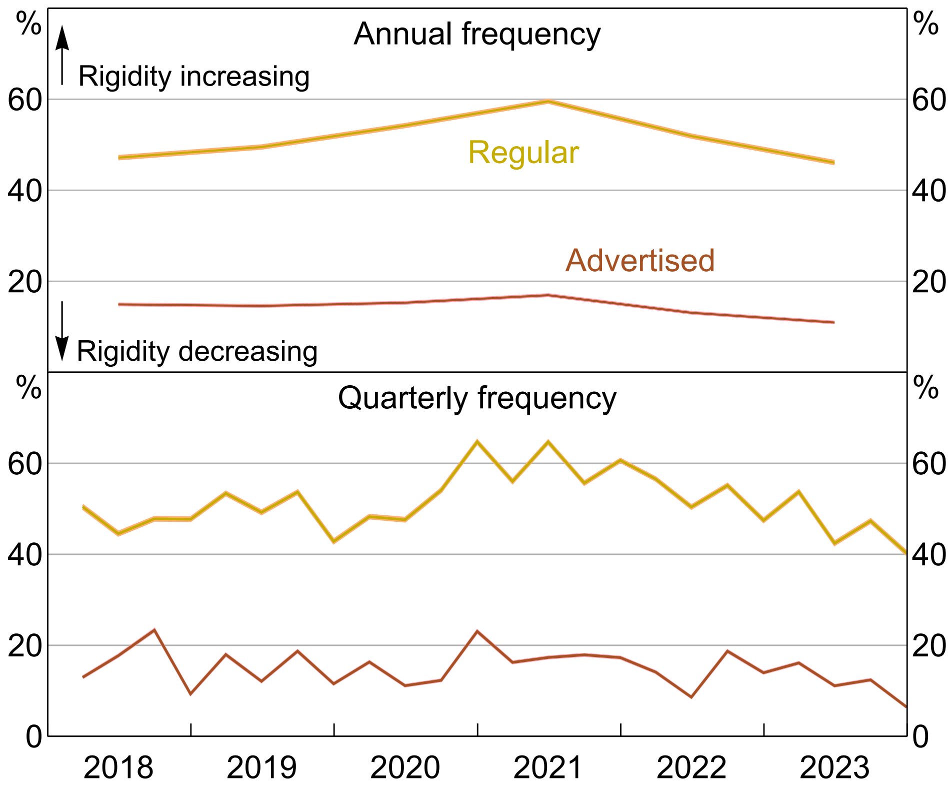

Notes: Item-weighted. Shading shows 95 per cent confidence intervals.
Sources: ABS; Authors' calculations.

We also find meaningful variation in rigidity over time. Regular prices became more rigid and less likely to change during the early pandemic period, with the probability of remaining unchanged rising by around 10 percentage points between 2019 and 2021. This was less evident for advertised prices, suggesting that firms were, in relative terms, increasingly using sales as a mechanism for price flexibility. This is consistent with heightened uncertainty evident during this period, which may have made firms relatively less willing to make more permanent price change prices.[46]

From 2022 onward, as general retail goods inflation accelerated, both regular and advertised prices became significantly more flexible: by the end of the sample the share of advertised and regular prices remaining unchanged over a quarter was 5 and 10 percentage points lower than pre-COVID. This suggests that when faced with large inflationary shocks and rising input costs from 2022, firms started changing prices more regularly, including by making more permanent changes.[47] This is consistent with international evidence that price rigidity tends to decline during high inflation periods.[48] Moreover, it is broadly in line with comparable estimates – for example

46 A theory of price rigidity based on firms' uncertainty about the competitive environment is developed in Ilut et al. (2020).

47 Our data includes some products (such as clothes and electronics) whose prices would typically be expected to decline steadily over their lifespan. It is possible that upward pressure on input costs for these products might result in firms slowing the frequency of price decreases, rather than increasing the frequency of changes. This would give the appearance that price rigidity was increasing alongside rising inflation, though the lower frequency of price changes would support rising inflation. To the extent that this is a factor, we expect that it would make the decline in rigidity in our dataset look more modest than it was.

48 See Cavallo, Lippi and Miyahara (2024) and Gagliardone et al. (2025).

Gautier et al. (2025) found that the share of regular prices not changing over a month and quarter in European CPI data fell by around 5 and 10 percentage points, respectively, over a similar period (the pickup is slightly later in Australia, consistent with the later increase in Australian inflation).

3.5 How shifts in price rigidity affect inflation dynamics

Having established in Section 4 that firms' price rigidity changed materially over the pandemic period, we assess whether these changes had meaningful implications for inflation dynamics. Specifically, we ask: if price rigidity declined between 2019 and 2023, as our estimates suggest, how would that have affected the transmission of shocks to inflation?

There are several ways to explore this question. One approach would be to build a new macroeconomic model calibrated to our microdata estimates. Another would be to incorporate our results into the Reserve Bank's existing benchmark DSGE model, which is used for policy analysis. We adopt the latter approach, as it allows us to assess the implications of changing price rigidity within a framework familiar to policymakers. This helps bridge the gap between micro-level evidence and benchmark policy tools, and is closer in spirit ot how policymakers may adapt their existing tools to address these challenges. Our exercise is similar in spirit to Gelain and Lopez (2024), who adapt the Federal Reserve's linear model frameworks to better capture inflation dynamics around the pandemic.

While the key elements of the model are outlined in Gibbs et al. (2021), we briefly recap them here. The DSGE is a multisector open economy model with a tradeable, non-tradeable, housing, mining and imports sectors. Each of these sectors is monopolistically competitive and faces equivalently Calvo or Rotemberg prices frictions (given the model is log-linearised around a zero inflation steady state). Goods are combined by a final goods firm into a composite good used for investment and consumption. Markets are complete and bonds are in zero net supply.

We focus on variation in the Calvo parameter—the probability that a firm does not adjust its price in a given period—which we estimated in Section 4 using survival analysis. Because the level of our estimated Calvo parameters differs from that used in the baseline DSGE model, we adjust the slope of the Phillips curve using the change in rigidity between 2019 and the peak of the inflation surge in 2023. This embeds an assumption that 'normal' price rigidity in our data is best represented by the observed rigidity probability immediately prior to the start of the COVID-19 pandemic. We think this is broadly reasonable given the extended period of price stability prior to 2020. An alternative approach would be to test the effect of the shift in rigidity from its highest point in 2021 (immediately prior to the inflation surge) and its lowest point in the data during the surge in 2023. However, this would be less in line with the intent of comparing price rigidity at the peak of an inflationary episode to 'normal' rigidity embedded in our benchmark policy model.

3.5.1 Adjusting the slope of the Philips curve

The RBA's DSGE model uses a Rotemberg-type price adjustment mechanism. To incorporate our Calvo-based rigidity estimates, we exploit the well-known equivalence between Rotemberg and Calvo pricing under a first-order approximation around the zero-inflation steady state.[49] This allows us to adjust the Phillips curve slope in the model using our estimated Calvo parameters.

49 See proof and discussion in Rotemberg (1987) and Roberts (1995).

Specifically, we exploit the fact that the slope of the Philips curve in the model can be expressed as:

$$\pi_t = \gamma * MC_t + \delta * E(\pi_{t+1}) + \epsilon_t$$

where:

- π_t is inflation,
- MC_t is marginal costs,
- ϵ_t is a cost-push shock,
- γ is the slope of the Phillips curve, defined as:

$$\gamma = (1 - \beta * \theta) * \frac{(1 - \theta)}{\theta} * \frac{1}{1 + \varphi\beta}$$

Here, β is the discount factor, φ is the degree of indexation and θ is the Calvo parameter (the share of prices that remain unchanged in each period). Given the direct estimates of γ in the Bank's model and the parameters β and φ, we can back out θ. We can then adjust $\hat{\theta}$ using our microdata estimates to create an alternate Philips curve slope $\hat{\gamma}$.[50]

One natural question is how to map our estimated Calvo parameters into the model, given that there are levels differences between the estimated slopes in our empirical data and those embedded in the various sectors of the DSGE model. We test two approaches: one based on the absolute change in our estimate of rigidity (percentage points), and another based on the proportional change (per cent). These serve as lower and upper bounds for our analysis (Table 3.6).

Table 3.6: Range in Observed Retail Price Rigidity[(a)]

Rigidity defined as the probability a price remains unchanged after one quarter

	Advertised prices	Regular prices
2019	14.6	49.5
2023	10.9	46.1
Difference (ppt)[(a)]	-3.7	-3.4
Difference (%)[(a)]	-25.3	-6.9

Note: (a) Lower bound = percentage point change; upper bound = proportional change.

Sources: ABS; RBA

Rather than having a single Philips curve, the DSGE model has a separate Philips curve for each of five sectors: resources, tradables, non-tradables, imports and housing. We uniformly apply the observed variations in price rigidity between 2019 and 2023 to the slope of each curve; i.e. we treat the observed shift in retail price rigidity as representative of a shift in economy-wide rigidity. On the face of it this may seem a somewhat strong assumption. However, it is unclear *a priori*

50 For the analysis, we abstract from the indexation and assume this is unchanged.

whether changes in other sectors would be larger or smaller. For example, if we were just to apply the change to the tradables sector this would imply no changes in rigidity in the rest of the economy, which could conceivably be a stronger assumption than uniformly applying the shift observed in our available data across industries.

3.5.2 Illustrating the role of price rigidity in the DSGE framework

To illustrate the impact of price rigidity on inflation dynamics, we simulate a one per cent foreign inflation (cost-push) shock in the DSGE model, using our estimated range of Calvo parameters for regular prices. We find that lower price rigidity leads to a stronger and faster inflation response. Specifically, the initial pass-through of the shock to year-ended inflation is higher by between 0.03 and 0.25 percentage points (Graph 3.7, panel 1). The long-run impact on the price level also increases, rising from around 0.2 per cent under baseline rigidity to as much as 0.35 per cent under lower rigidity—an increase of nearly 75 per cent (panel 3).[51]

Graph 3.7: Response to a Cost-Push Shock under Alternative Price Rigidity Assumptions

Regular prices

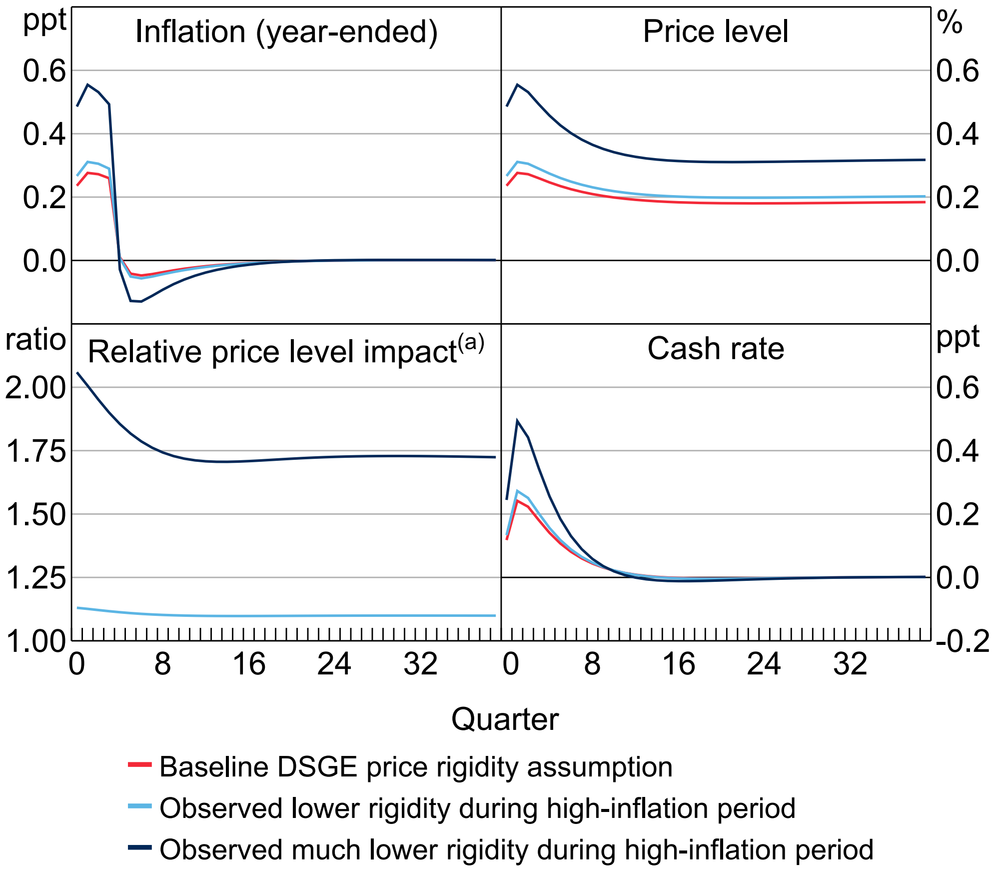

Notes: 100 basis point positive shock in the price of foreign imports.
(a) Ratio of the price level impact from the shock using alternative rigidity assumption over impact using baseline rigidity.

Sources: ABS; Authors' calculations.

---

51 We also tested the impact of a positive one per cent domestic consumption (demand) shock on the price level under alternative price rigidities; in general, the results are similar to those following the cost shock, though of somewhat smaller magnitude. For example, accounting for the decline in price rigidity with a consumption shock results in a price level estimate around 2 to 20 per cent higher in the long run.

Given the model's Taylor rule, this stronger inflation response translates into a more aggressive policy rate path. The cash rate rises by between 3 and 25 basis points more in the first year compared to the baseline. Results are similar—but somewhat larger—when using the range of Calvo parameters estimated for advertised prices, which include temporary sales and tend to be more flexible (Graph C1).

3.5.3 Forecasting inflation with time-varying price rigidity

The previous subsection illustrated how changes in price rigidity affect the transmission of a single shock. To assess the practical relevance of these changes, we now ask: if price rigidity declined between 2019 and 2023 as our estimates suggest, how much would this have affected inflation forecasts during the post-pandemic surge?

To answer this, we conduct a scenario forecasting exercise using the DSGE model. We treat our microdata-based estimates of shifts in price rigidity as true and adjust the slope of the Phillips curve accordingly. We then use the model to recover the set of shocks that must have hit the economy to generate observed outcomes. Finally, we simulate the model forward from the start of 2022 using these estimated shocks but revert the Phillips curve slope to its baseline value. Comparing the resulting inflation forecasts to actual outcomes allows us to isolate the impact of misspecified price rigidity. This approach is similar to De Fiore et al. (2023), who use it to evaluate alternative monetary policy frameworks during the post-pandemic inflation period.[52] The exercise can be thought of as asking: how wrong would forecasters have been if they had known all the shocks that were going to hit the economy, but didn't account for the fact that price rigidity had changed?

More precisely, the DSGE model can be represented in its VAR form:

$$Y_t = \boldsymbol{\rho} Y_{t-1} + \boldsymbol{\vartheta} \varepsilon_t$$

where:

- Y_t is the vector of endogenous variables,
- $\boldsymbol{\rho}$ is a matrix linking past values to current outcomes based on the model parameters,
- ε_t are structural economic shocks,
- $\boldsymbol{\vartheta}$ is a matrix capturing the transmission of shocks based on the model parameters.

Changing the Phillips curve slope (via γ) alters both $\boldsymbol{\rho}$ and $\boldsymbol{\vartheta}$**.** Using our lower rigidity estimates, we denote the adjusted matrices as $\widehat{\boldsymbol{\rho}}$ and $\widehat{\boldsymbol{\vartheta}}$, and the recovered shocks as $\hat{\varepsilon}_t$. The model then becomes:

52 It is also somewhat similair to the decomposition in Lane (2024), where European inflation outcomes are decomposed into those reflecting incorrect assumptions about exogenous variables like energy prices, and those that cannot be explained by incorrect assumptions.

$$Y_t = \widehat{\boldsymbol{\rho}} Y_{t-1} + \widehat{\boldsymbol{\vartheta}} \hat{\varepsilon}_t$$

We take these to be the true structure of the economy, and so $\hat{\varepsilon}_t$ are the true structural shocks that were hitting the economy during the period.

With these 'true' shocks in hand, we can then revert to the original model structure from 2022 onwards and calculate what forecasters might have expected inflation to be if they knew the shocks hitting the economy but had the wrong model structure and rigidity. The counterfactual conditions would be:

$$\widehat{Y}_t = \boldsymbol{\rho} \widehat{Y}_{t-1} + \boldsymbol{\vartheta} \hat{\varepsilon}_t$$

The difference between observed and counterfactual outcomes, $Y_t - \widehat{Y}_t$, gives an estimate of the forecast error attributable to misspecified price rigidity.[53]

Results

Graph 3.8 show the results. The red line shows actual year-ended inflation, while the blue lines show model-predicted inflation using the upper and lower bounds of our estimated rigidity decline. Even assuming perfect knowledge of the shocks, failing to account for the decline in price rigidity would have led to under-predicting inflation by between 0.41 and 1.33 percentage points one year ahead. The peak of inflation would also appear slightly more persistent and delayed. The magnitude of these effects if broadly in line with findings in Gautier et al. (2025), which argues that failing to account for changes in price-setting frequency in Europe during the post-pandemic high inflation would have led to inflation predictions being one percentage point lower.

---

53 One limitation of this approach is that the recovered shocks prior to 2022 are based on lower rigidity assumptions, whereas our results in Section 5 suggest that rigidity may have been higher during the early pandemic. This could introduce minor bias due to autoregressive shock dynamics, but we expect the impact to be second order.

Graph 3.8: Inflation Outcomes under Various Price Rigidity Assumptions

Regular prices, year-ended

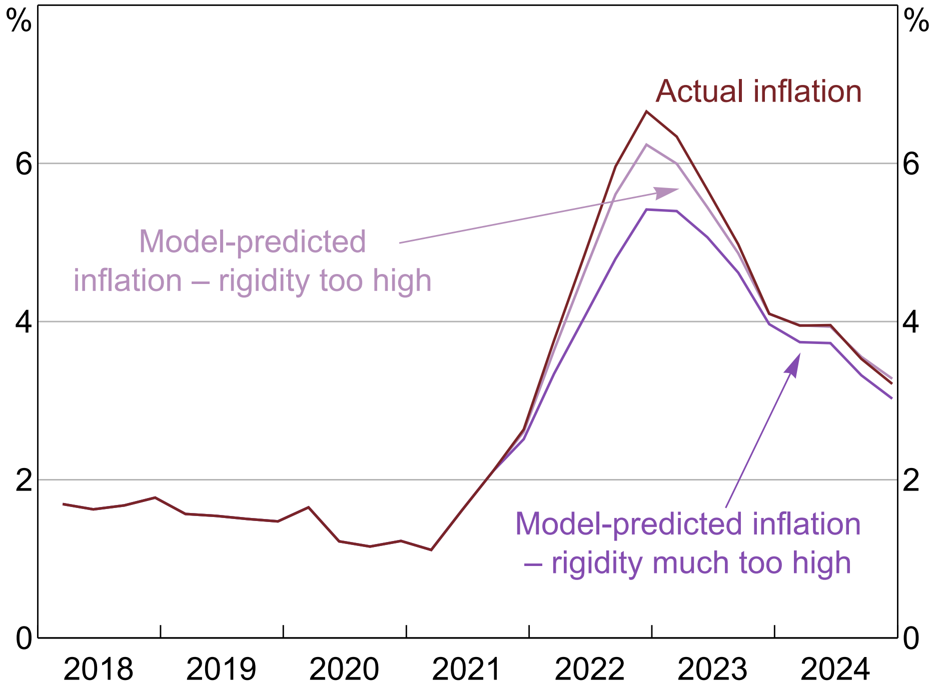

Note: Actual inflation rate versus predicted path of inflation if price rigidity in the model was overestimated.

Sources: ABS; Authors' calculations.

These findings suggest that ignoring time variation in price rigidity—particularly during periods of large cost shocks—can lead to significant forecast errors. While our analysis is based on a DSGE framework, the implications are broader. Many forecasting models, including single-equation Phillips curves commonly used in public and private sector institutions, assume a fixed Phillips curve slope. Our results highlight the importance of allowing for time-varying price-setting frequencies when inflation dynamics are shifting rapidly and demonstrate that this was critically important during the recent inflationary episode.

As a supplementary exercise, we also tested whether varying the assumed degree of price rigidity affects the model's decomposition of inflation into supply and demand shocks. The results suggest minimal sensitivity to this assumption. See Appendix C for details.

3.6 How shifts in price rigidity impact monetary policy

Changes in price-setting behaviour can have important implications for monetary policy. If price rigidity varies over time, then the transmission of monetary policy—and the trade-offs faced by policymakers, particularly during supply-side shocks—may also change.

In this section, we explore the implications of changes price rigidity for monetary policy transmission, trade-offs and strategy. We conduct two exercises. First, we compare undertake a so-called optimal control exercise, calculating the policy paths that would have minimised (weighted) deviations of inflation and unemployment from target during the 2022–2023 high-inflation period under baseline model rigidity and under the lower rigidity observed in that period. Second, we examine how the optimal weights (those that minimise weighted deviations of inflation and unemployment) in a simple Taylor rule change as rigidity declines. This latter exercise is especially relevant in the context of supply-side shocks, where inflation and output move in opposite directions, creating a sharper policy trade-off. In contrast, demand-driven shocks typically move inflation and output in the same direction, making the implications of changing rigidity for those trade-offs less pronounced.

3.6.1 Transmission of monetary shocks under varying rigidity assumptions

To motivate the analysis, we first simulate the impact of a 100 basis point monetary policy shock under baseline and lower rigidity assumptions. The results show that when prices are more flexible, the disinflationary effect of a given rate increase is stronger (Graph 3.9). Crucially, this stronger inflation response is achieved without a larger decline in output. A steeper Phillips curve means that a given reduction in real activity translates into a larger fall in inflation. As a result, the trade-off between stabilising inflation and supporting real output is reduced.

Graph 3.9: Response to a Cash Rate Shock under Alternative Price Rigidity Assumptions

Regular prices

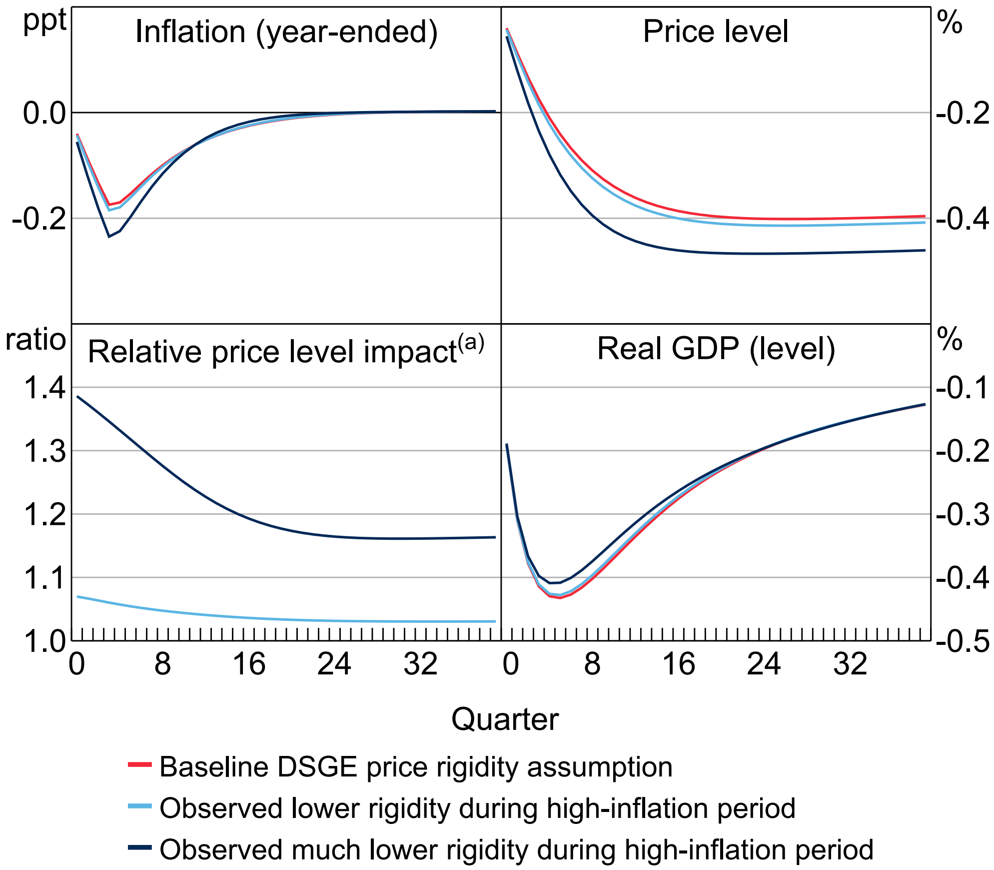

Notes: 100 basis point positive shock in the cash rate.
(a) Ratio of the price level impact from the shock using alternative rigidity assumptions over impact using baseline rigidity.

Sources: ABS; Authors' calculations.

3.6.2 Optimal control exercises under varying price rigidity assumptions

We now examine how shifts in price rigidity effect monetary policy trade-offs and strategy. To do this, we implement an optimal control exercise using the DSGE model. Unlike a fixed policy rule approach (e.g. estimating a Taylor rule), optimal control calculates the path for the cash rate that minimises a standard policy loss function, given the shocks that hit the economy. This is a common method used by policymakers (e.g. in Lowe and Ellis 1997 and Adolfson et al. 2011).

Methodology

For our analysis we assume that the loss function places equal weight on stabilising inflation, output (as a proxy for unemployment) and interest rate smoothing. Formally:

$$Loss = \sum_{j=1}^{N}\sum_{i=0}^{T} \beta^i * \omega_j * (a_{j,t+i} - a_{j,t+i}^{target})^2$$

where:

- a_j are the policy-relevant variables (inflation, output gap, and change in the nominal interest rate),
- a_j^{target} are their respective 'targets' (2.5 per cent for inflation, 0 for output gap and interest rate changes)[54],
- ω_j are the weights assigned to each variable,
- β is the discount factor (set to 0.9996),
- The summation covers the period from March 2022 to December 2024.

Inflation and interest rate changes are each assigned a weight of 1. The output gap is weighted at 1/64, which equates a one percentage point deviation in inflation to a one percentage point unemployment gap, based on Okun's law. This implies equal weighting across inflation, unemployment, and interest rate volatility in the loss function.

We focus on unanticipated monetary policy shocks, solving the path of unanticipated shocks that minimizes the loss function. In this sense we are not solving for a full commitment solution. We also do not compare the paths coming from these exercises to the actual policy response in the period as this would not be a like-for-like comparison. The RBA Board may have used a different loss function, and the model has access to ex post outcomes that were unknown to policymakers in real time. Instead, we compare optimal paths under different rigidity assumptions to draw general lessons about how monetary policy should respond when price-setting behaviour changes, ceteris paribus.

---

54 More precisely the target for inflation is the average rate, which is closer to 2.6 per cent.

Results

Graph 3.10 shows the policy paths under baseline and lower rigidity assumptions. When prices are less rigid, the optimal response is to raise interest rates more aggressively in the face of inflationary shocks. The cash rate peaks between 10 and 40 basis points higher than under baseline rigidity within the first year, resulting in substantially lower inflation.[55]

Graph 3.10: Cash Rate and Economic Outcomes from Optimal Control Exercise

Relative to paths with baseline rigidity

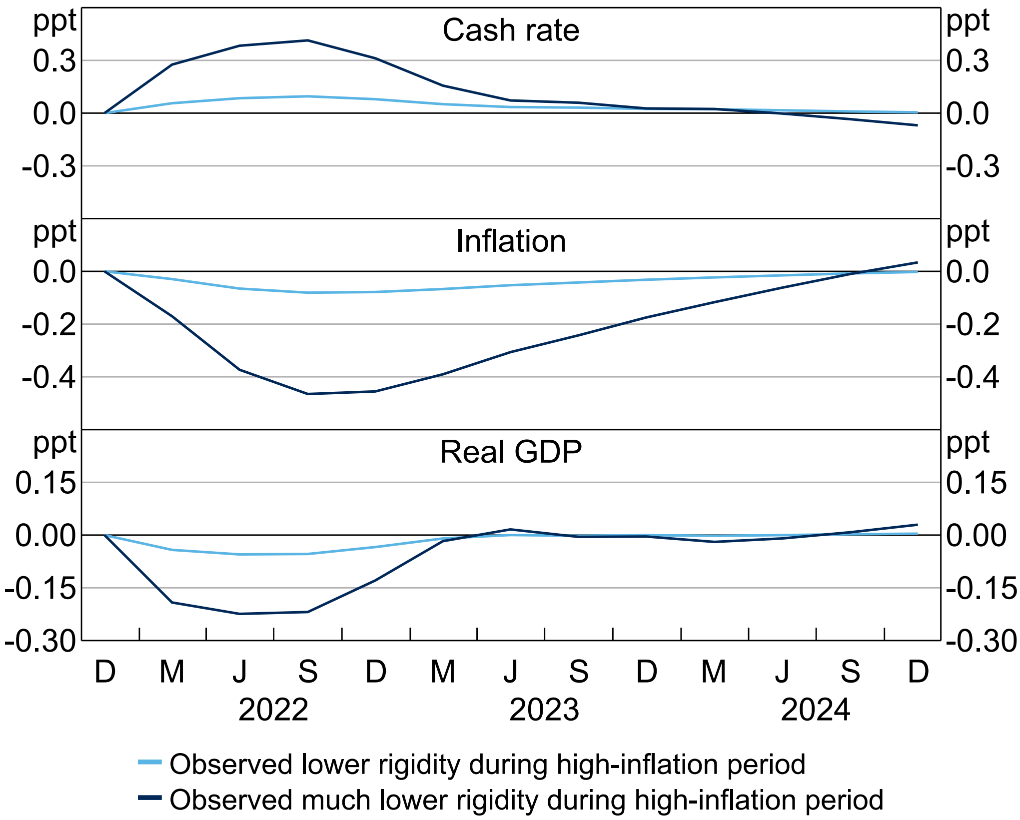

Notes: Cash rate path from optimal control exercise and associated inflation and real GDP outcomes under alternative assumptions about price rigidity. Results are shown relative to those obtained using the baseline price rigidity estimate embedded in the model. Inflation and real GDP growth are annualised growth rates.

Sources: ABS; Authors' calculations.

While output growth also declines slightly more than under the baseline, the trade-off is more favourable: the additional disinflation is achieved with only a modest additional fall in output. This reflects the steeper Phillips curve under lower rigidity, which allows inflation to be stabilised more effectively for a given change in real activity.

55 A decline in price rigidity also implies a lower welfare cost of inflation, which could justify placing less weight on inflation in the policy loss function. We tested a variant of the exercise where the inflation weight was reduced by 30 per cent, consistent with the decline in the social cost of inflation implied by lower rigidity in simple DSGE models (e.g. Blanchard and Galí 2008). While this adjustment brought the optimal policy path closer to the baseline, we caution against overinterpreting the result. The RBA's DSGE model is more complex than the stylised frameworks used to derive such adjustments, and the calibration of the weight change may not be directly transferable.

Overall, our findings suggest that when price-setting frequencies pick up—particularly during supply-side inflationary shocks—monetary policy can be tightened more aggressively to achieve faster disinflation with limited additional cost to output and unemployment. These results are consistent with those described in Karadi et al. (2024), who conduct a similar exercise for the 2022-23 US inflation surge using a model with state-dependent menu cost rigidities.

3.6.3 Optimal simple rules under different rigidity assumptions

While optimal control exercises provide a full policy path under a particular scenario, central banks often rely on simple policy rules – such as the Taylor rule – for communication, forecasting, and operational guidance. The Taylor rule prescribes a policy interest rate response given deviations in inflation and output from their targets (Taylor 1993). If price rigidity changes over time, and so the trade-off between inflation and unemployment changes, this could alter how aggressively policymakers want to respond to inflation and output fluctuations, and therefore the preferred weights in the rule.

To explore this, we assess how the coefficients in a standard Taylor rule hat minimise a loss function vary under different assumptions about price rigidity. This helps quantify how monetary policy trade-offs evolve when firms adjust prices more flexibly, particularly during supply-side shocks.

Methodology

The Taylor rule specification embedded in the RBA's DSGE model takes the following form:

$$i_t = \rho i_{t-1} + (1-\rho)\left[r^* + \phi_\pi(\pi_t - \pi^*) + \phi_y \Delta y_t\right] + \varepsilon_t$$

Where:

- i_t is the nominal interest rate,
- i_{t-1} is the lagged interest rate,
- r^* is the neutral real interest rate,
- π_t is the two quarter average inflation rate,
- π^* is the inflation target,
- Δy_t is the two quarter average GDP growth rate,
- ϕ_π and ϕ_y are the policy response coefficients to inflation and output, respectively,
- ρ is the interest rate smoothing parameter,
- ε_t is a monetary policy shock.

To estimate how the values of ϕ_π and ϕ_y that minimise losses vary with price rigidity, we simulate the DSGE model under different Calvo parameters (as estimated in Section 4). For each rigidity setting, we solve for the Taylor rule coefficients that minimise the same policy loss function used in Section 6.2, subject to the model's structure and shock variances. In this sense this exercise is thinking about what general rule would be optimal on average over all time under the different rigidities.[56]

Results

Under the baseline rigidity setting in the model, the loss-minimising weight on inflation is approximately 1.8 times the weight on output. When rigidity is lowered in line with observed changes during the post-pandemic inflation, this ratio increases to between 1.9 and 2.4 — a rise of 5 to 33 per cent.[57] These results are consistent with our earlier findings. When prices are more flexible, inflation responds more to changes in activity and the Phillips curve becomes steeper. As a result, placing greater emphasis on inflation in the policy rule is preferred. Our estimated ratios provide a simple metric for thinking about how the central bank's policy considerations might shift during periods of elevated inflation and lower price rigidity.

3.7 Conclusion

Our analysis shows that the large supply shocks of the late pandemic period coincided with a decline in price rigidity, as firms became more willing to adjust prices. This shift contributed meaningfully to the sharp rise in inflation in 2022 and 2023.

The observed tendency for price rigidity to decline during large shocks has important implications for macroeconomic modelling and policy. First, inflation may accelerate more rapidly than predicted by standard linear models, requiring forecasters – including central banks – to account for more flexible price-setting behaviour. Second, as rigidity falls, the Phillips curve steepens, reducing the trade-off between inflation and output. This means that, all else equal, central banks facing large supply-side inflationary shocks can raise interest rates more aggressively and achieve inflation objectives more quickly, with limited additional costs to economic activity.

Applying these insights in real time is challenging. Prices microdata are not currently available on a timely basis in many countries, including Australia, and the nature of shocks during turbulent periods can be difficult to identify. Policymakers must therefore rely on informed judgement to assess whether firms are adjusting prices more frequently. Business liaison and business surveys may offer valuable real-time signals in these periods.

Accurate identification of the type of shock hitting the economy is also critical. The policy implications we describe for periods of lower price rigidity are most relevant for supply-side shocks, where inflation and output move in opposite directions. In contrast, demand shocks typically push both inflation and output in the same direction, making the policy trade-off less

56 One weakness of the approach is that the times that we face lower rigidities are exactly those times where the shocks hitting the economy are extreme. That said, the exercise still provides a general sense of the change in trade-offs.

57 Note that both coefficients decline in absolute terms, particularly the output coefficient. This reflects the fact that, under lower rigidity, monetary policy has a stronger effect on real interest rates. Given this, demand shocks — which do not involve a trade-off between inflation and output — can be addressed with smaller interest rate adjustments.

relevant. This underscores the importance of robust shock identification, as discussed in Beckers, Hambur and Williams (2023).

One way to improve forecasting accuracy under shifting price-setting behaviour is to expand the set of models used in policy analysis. Models with state-dependent pricing allow firms' responsiveness to vary with the size of economic shocks and can be calibrated using empirical moments from microdata – such as those estimated in this paper.

In the absence of timely data or formal models that explicitly capture changing rigidity, policymakers must lean more heavily on informed assessment. In these instances, the findings in this paper offer an intuitive evidence-based guide to how price-setting behaviour is likely to evolve during large supply shocks and how these changes affect inflation dynamics and the design of monetary policy responses.

Appendix A Data, cleaning and sales algorithm

A.1 Data description

The raw data consists of prices for retail goods listed on the websites of major Australian retailers. These were collected by the Australian Bureau of Statistics (ABS) every few days between the start of 2016 and end of 2023. In general, the collection process involved taking the price of every item listed on a given retailer's website. The ABS performed some additional cleaning to reduce the likelihood that price collections were made for duplicates of the same item; however in some instances it is possible that very similar items are recorded separately.[58]

Price collection dates vary by firm, and the frequency of price collections for a given firm may be irregular. For example, all prices for the items of a given firm might be collected on one date, then two days later, then four days later, then three days later, etc. The median frequency of price collections per item across firms is around 6 per month; this varies over time but generally increases over the sample (Graph A1).[59] Because of this, our main analysis of price rigidity uses a survival analysis method that accounts for bias from the irregular timing and frequency of price collections.

Graph A1: Median Number of Price Observations per Item
CPI-weighted across firms, monthly

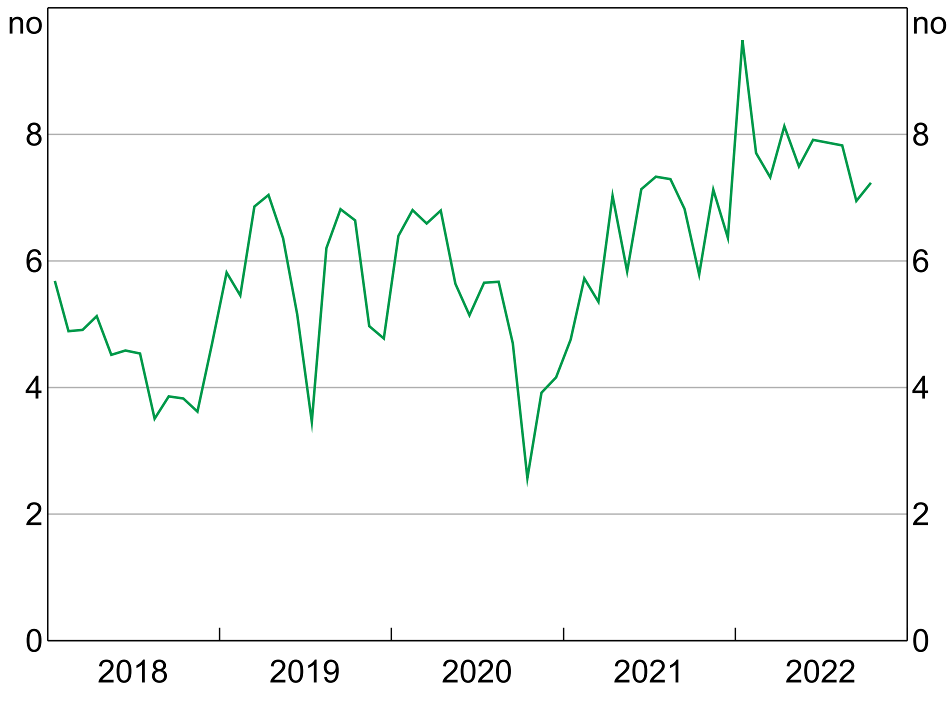

Sources: ABS; Authors' calculations.

58 An example might be two t-shirts in the same style that are different colours.

59 These are likely to be lower bound estimates for a given item over the sample period, as in an practice items will regularly drop out of the sample as they are retired.

Data fields available for each retailer include a unique item number that is consistent across collection dates, the date of each price observation and up to three types of prices: 'advertised' prices, which are the current retail price of the item offered to consumers on that date (this may or may not be a discounted price); 'full' prices, which are the undiscounted price of the item; and 'discounted' prices, which are the sale price of the item if a sale is active (Table A1). Advertised prices are available from the start of the dataset, while full and discounted prices are only available from 2020. The recording of full and discount prices in the raw data reflects the classification of these price types by the retailers themselves as per their website. This means that there may be inconsistency in how sales are recorded across firms. For this reason, we standardise the definition of regular and sales prices across firms using the sales algorithm described in Section A.3.

Table A1: Sample visualisation of uncleaned web-scraped prices data

Indicative

Retailer number	Retailer group	Item number	Date	Advertised price	Full price	Discounted price
5	clothing	1	01/01/2020	20		
5	clothing	1	03/01/2020	15		
5	clothing	1	06/01/2020	15	20	15
5	clothing	2	01/01/2020	10		
5	clothing	2	03/01/2020	12		
5	clothing	2	06/01/2020	12		

Source: RBA

The presence and total number of firms captured in the dataset varies over time, with the number of firms being too small to provide comfort about statistical validity of the sample prior to 2018. Between 2018 and 2023, the average number of firms present in each month is 53 and the total number of unique firms is 61 (Graph A2). Each firm in the dataset is deidentified, although a 'retailer group' variable corresponding to the broad retail category of each firm is provided. The set of retailer categories is: alcohol, automotive, clothing, cosmetics, department store, electrical, eyewear, footwear, furniture, hardware, homewares, jewellery, magazines and pharmaceuticals.

Graph A2: Number of Unique Firms in Prices Data
Monthly

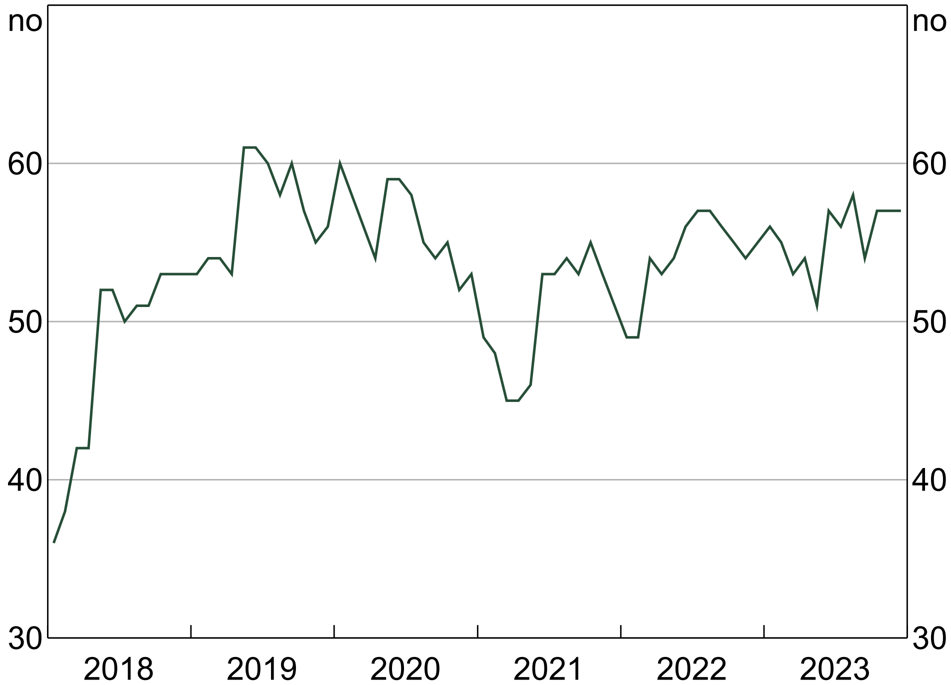

Sources: ABS; Authors' calculations.

Apart from the general category of the retailer and a deidentified product code, no information about the nature of a given sampled product is available in the dataset. This means that analysis at a more granular 'product type' level is not possible. That said, the ABS has collected information about the product type of items and may make this available in some form in the future.

ABS confidentiality requirements restrict the amount of information that can be shared about the pricing behaviour of individual firms present in the data.

Where appropriate, measures constructed as a part of our analysis are weighted using weights derived from mapping retailer categories to their approximate category matches in the Australian Consumer Price Index. This mapping is shown in Table A2. Where several retailers are part of the same category, we split their category weight evenly amongst them.

Table A7: CPI weights

Retailer group	Nearest analogous ABS CPI expenditure groupings	CPI weight used in analysis[(a)]
Alcohol	Alcoholic beverages	5.0
Automotive	Spare parts and accessories for motor vehicles	0.8
Clothing	Garments	2.0
Cosmetics	Personal care products	0.9
Department	Combination of: household textiles; household appliances, utensils and tools; clothing and footwear	5.3
Electrical	Combination of: major household appliances; small electric household appliances; audio, visual and computing equipment	1.9
Eyewear	Accessories	0.7
Footwear	Footwear	0.5
Furniture	Furniture	1.4
Hardware	Combination of: tools and equipment for house and garden; maintenance and repair of the dwelling	2.5
Homewares	Combination of: household appliances, utensils and tools; household textiles	2.0
Jewellery	Accessories	0.7
Magazines	Newspapers, magazines and stationery	0.4
Pharmaceuticals	Pharmaceutical products	1.0

Notes: (a) Based on average annual weights for mapped CPI expenditure groupings between 2018 and 2023.

Source: ABS; RBA

A.2 Data cleaning

Price changes are measured between two consecutive observations in the dataset and are calculated in log difference terms. For the survival analysis, we drop price observations that may be stale, which we define as when the distance between two price observations is 30 days or more. As noted in Section 4.2, we also drop the first price spell for each item to avoid the possibility of left censoring. Further cleaning of our imputed regular prices series is described in the next section. Apart from these measures, we include all price observations in the data set.

A.3 Sales algorithm

Firm-reported regular and sales prices are available in our dataset from late 2020. To extend these series back to the start of 2018 and ensure a consistent definition of sales pricing across firms, we develop an algorithm that identifies temporary and clearance sales based on observed pricing patterns. This allows us to construct imputed series for regular and sales prices over the full sample period.

Our approach follows the methodology introduced by Nakamura and Steinsson (2008) and further developed by Gautier et al. (2023). The algorithm classifies price changes as either regular or sale-related by identifying specific patterns in the evolution of advertised prices over time:

- Temporary (or 'V-shaped') sales are flagged when the advertised price of an item falls and then returns to an equal or higher level within three months. To ensure the return to full price for all items can be observed, we exclude prices that fall within the final three months of a firm's data.

- A temporary sale flag is removed if the post-sale price is followed by another price decline within 14 days. This helps to avoid misclassifying volatile pricing behaviour as sales activity.

- Clearance sales are identified when an advertised price falls by at least 25 per cent and the item exits the dataset within one month. Prices within one month of the end of a firm's data are excluded from the data to ensure the exit condition can be verified.[60]

To assess the performance of the algorithm, we compare our imputed regular and sales prices series to ABS-recorded full and sales prices series available from late 2020 onwards (Graphs A3 and A4). The imputed and ABS-recorded series generally align well in level and trend. However, at disaggregated levels, the ABS-recorded data suggest some inconsistencies in how firms distinguish between regular and sales prices, likely reflecting variation in how retailers label discounts on their websites or how the web-scraper captured this information.

A key advantage of our algorithm is that it applies a uniform definition of sales across all firms. By focusing on the underlying structure of price changes — rather than relying on retailer-provided labelling — the algorithm ensures that sales are identified consistently, regardless of how they are attributed in the source data.

60 Adjusting the time windows used to identify temporary and clearance-style sales had only minor effects on results, and the overall pattern of identified sales remained stable across these tests.

Graph A3: Comparison of ABS-recorded and Imputed Regular Prices Data

Price duration

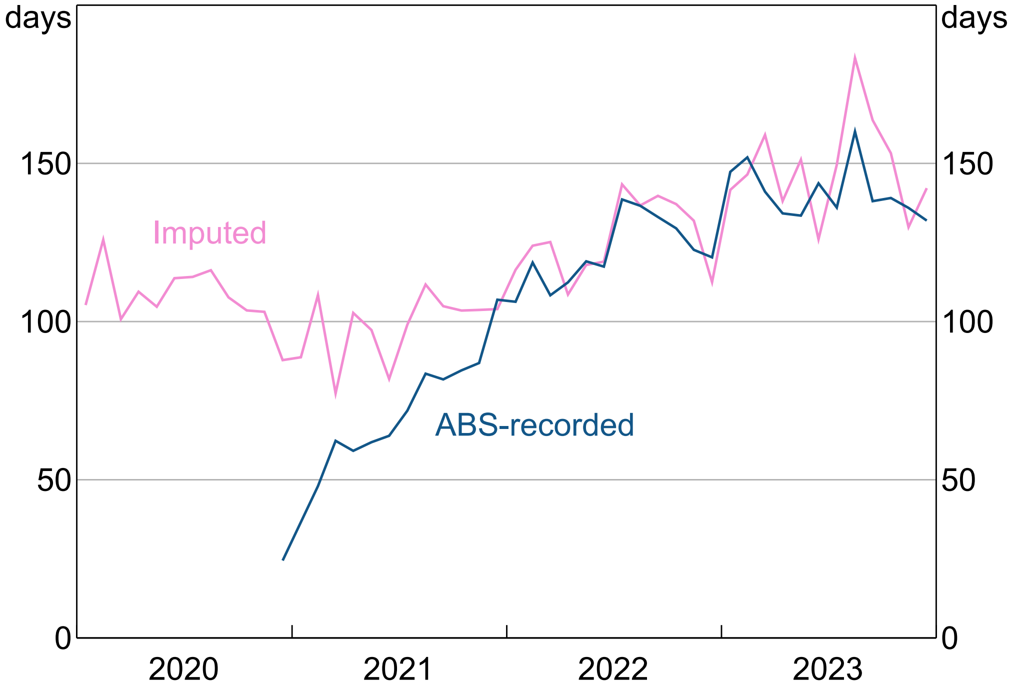

Notes: Item-weighted. Average number of days since prices last changed, conditional on a change having occurred.

Sources: ABS; Authors' calculations.

Graph A4: Comparison of ABS-recorded and Imputed Prices Series

High-inflation period, 2022–2023

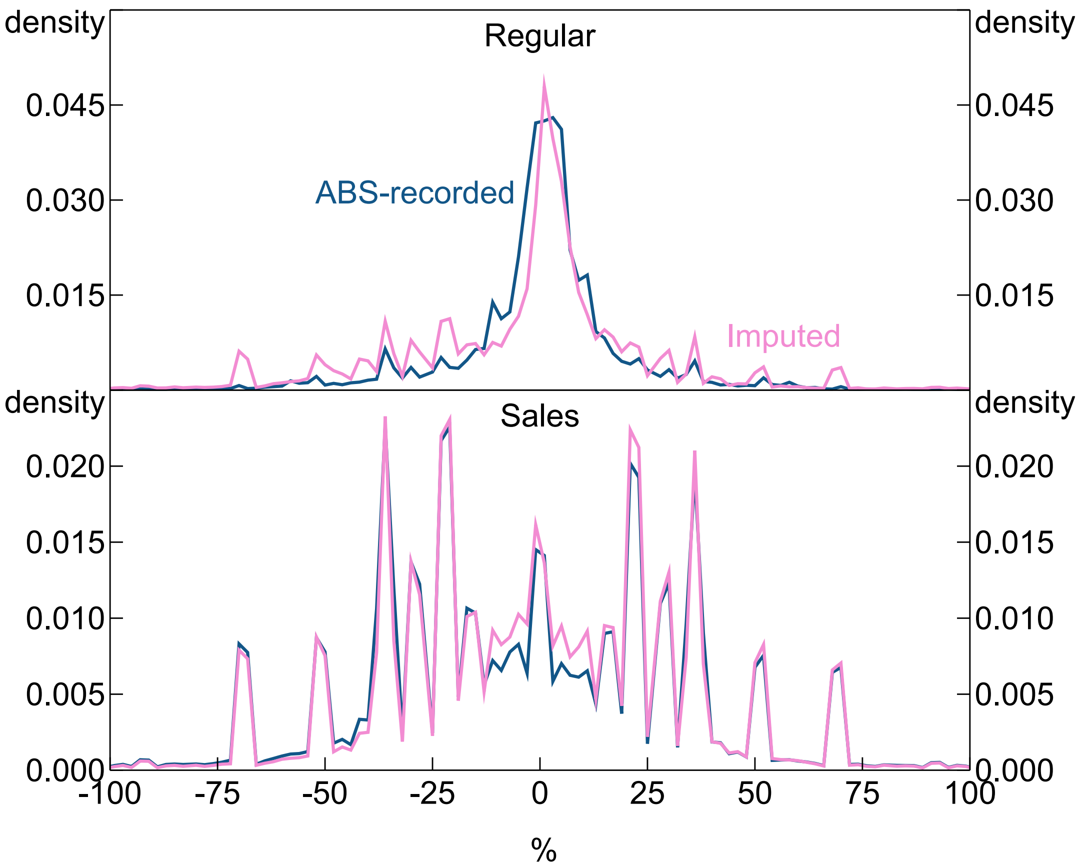

Notes: Item-weighted distribution of log price changes by size estimated by kernel density. The density is evaluated on the x-axes in 2 percentage point intervals. For example, the value plotted around zero corresponds to price changes in the interval –1 to 1 per cent.

Sources: ABS; Authors' calculations.

Appendix B Additional analysis of the size of price changes

We document how the distribution of price changes changed over time. At the onset of the pandemic, the distribution of both advertised and regular price changes became more dispersed, with increased weight in the tails—indicating a higher incidence of large price increases and decreases Graphs (B1 and B2).[61] This is reflected in a notable rise in the kurtosis of price changes around 2021 in both price measures (Table B). The increase in kurtosis suggests a shift toward more extreme price adjustments, consistent with international findings from the early pandemic period (e.g. Montag and Villar 2023). These changes coincided with a sharp drop in demand and heightened economic uncertainty.

Graph B1: Distribution of Advertised Price Changes by Size

Item-weighted, kernel densities

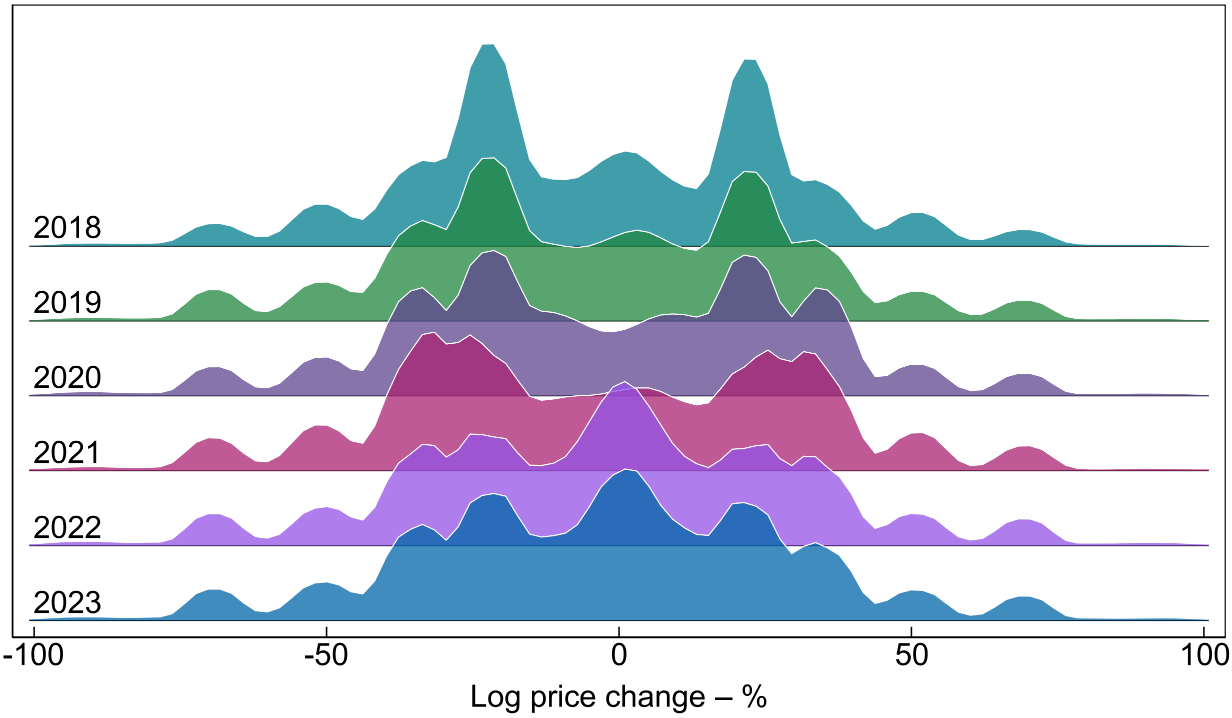

Sources: ABS; Authors' calculations.

61 As a robustness test, we ran versions of the item-weighted yearly kernel distributions holding the firm sample fixed between 2019 and 2023. The measured shape of distributions was not significantly different, nor was the evolution of changes in distributions over time. That said, there was some variation in the height of various peaks in the distribution away from the centre. This suggests the need for caution in interpretating sales-related peaks; i.e. the relative height of the peaks is affected by compositional change in which firms are sampled over time.

Graph B2: Distribution of Regular Price Changes by Size
Item-weighted, kernel densities

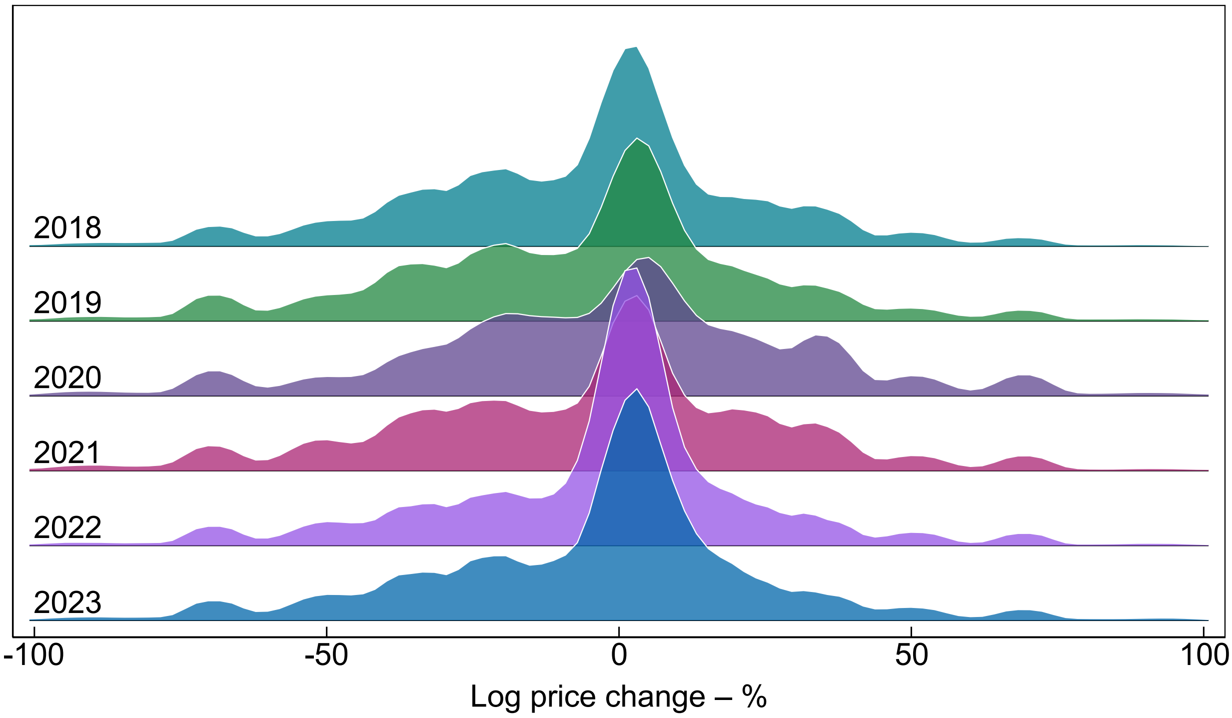

Note: Imputed regular prices.
Sources: ABS; Authors' calculations

Table B1: Mean kurtosis of the size of price changes

	Advertised Item-weighted	Advertised Firm-weighted[(a)]	Regular Item-weighted	Regular Firm-weighted[(a)]
2018	10	12	11	21
2019	8	10	10	20
2020	7	11	9	17
2021	15	18	48	22
2022	6	14	11	27
2023	7	12	12	18

Note: (a) Cross-firm CPI-weighted mean calculated using a fixed effects regression to control for in-sample volatility in firm composition.
Sources: ABS; RBA

During the high-inflation phase, the distribution of regular price changes became more centralised, with a higher frequency of smaller adjustments and a decline in kurtosis. This pattern is consistent with theoretical and empirical work showing that higher inflation reduces the relative cost of frequent price changes, leading firms to adjust prices more often but by smaller amounts.[62] However, interpreting these shifts is not straightforward. While visual inspection suggests a greater share of small price changes during the high-inflation period, kurtosis levels are broadly similar to pre-pandemic values. This makes it difficult to draw firm conclusions about whether price rigidity was ultimately higher or lower than before the pandemic. Moreover, changes in the nature of the shocks driving marginal costs will also influence the distribution of price changes.

This highlights some of the challenges in using distributional statistics to infer changes in price-setting behaviour. In particular, estimates of kurtosis are highly sensitive to changes in sample composition and outlier values. For example, combining subgroups with different variances—even if each has the same kurtosis—can inflate the overall measure. While our firm-weighted metrics attempt to control for compositional shifts by focusing on within-firm changes, variation in the number of items per firm can still affect results. More broadly, measurement error—whether from prices being captured inaccurately or irregular data collection—can bias kurtosis estimates. So while our distributional measures offer useful insights, they should be interpreted with caution and complemented by more robust methods.

62 For theory, see Alvarez et al. (2022) and Cavallo, Lippi and Miyahara (2024). For overseas empirical evidence, refer to the recent studies cited in our literature review.

Appendix C Survival analysis regression results

The tables below include coefficients from the survival analysis that correspond to the beta term in equation 1. Transformation of a given coefficient (β) to a cross-sample quarter frequency Calvo parameter equivalent (θ) is achieved via:

$$\theta = \exp\left(-\exp(\beta) * \left(\frac{365.25}{4}\right)\right)$$

Table C1: Survival Analysis Results

Annual

	Advertised prices	Regular prices
2018	–3.871***	–4.800***
	[–3.866, –3.876]	[–4.791, –4.809]
2019	–3.859***	–4.866***
	[–3.855, –3.864]	[–4.857, –4.875]
2020	–3.884***	–5.005***
	[–3.879, –3.888]	[–4.996, –5.014]
2021	–3.940***	–5.171***
	[–3.935, –3.945]	[–5.163, –5.180]
2022	–3.804***	–4.936***
	[–3.800, –3.809]	[–4.927, –4.945]
2023	–3.720***	–4.769***
	[–3.715, –3.725]	[–4.760, –4.778]
No of observations	49,916,224	13,139,762

Notes: *** denotes statistical significance at the 1 per cent level. Coefficients are computed by adding estimated year dummy coefficient to base year coefficient and testing the difference from zero using a Wald test. All standalone coefficients are statistically significantly different from zero. Square brackets show lower and upper bound 95 per cent confidence intervals, which are based on the linear combination of base year and dummy year coefficients and take account of covariance between these parameters.

Sources: ABS; Authors' calculations.

Table C2: Survival Analysis Results

Quarterly (*continued next page*)

	Advertised prices	Regular prices
2018:Q1	–3.800***	–4.894***
	[–3.794 , –3.805]	[–4.883 , –4.904]
2018:Q2	–3.966***	–4.726***
	[–3.961 , –3.972]	[–4.716 , –4.737]
2018:Q3	–4.139***	–4.818***
	[–4.134 , –4.145]	[–4.808 , –4.828]
2018:Q4	–3.650***	–4.815***
	[–3.645 , –3.655]	[–4.806 , –4.825]

Table C2: Survival Analysis Results
Quarterly (*continued*)

	Advertised prices	Regular prices
2019:Q1	–3.974***	–4.980***
	[–3.969 , –3.979]	[–4.970 , –4.990]
2019:Q2	–3.766***	–4.858***
	[–3.761 , –3.772]	[–4.848 , –4.868]
2019:Q3	–3.998***	–4.988***
	[–3.993 , –4.003]	[–4.978 , –4.998]
2019:Q4	–3.745***	–4.679***
	[–3.740 , –3.750]	[–4.670 , –4.689]
2020:Q1	–3.920***	–4.831***
	[–3.915 , –3.925]	[–4.822 , –4.841]
2020:Q2	–3.728***	–4.812***
	[–3.723 , –3.733]	[–4.802 , –4.822]
2020:Q3	–3.774***	–5.001***
	[–3.769 , –3.780]	[–4.991 , –5.011]
2020:Q4	–4.131***	–5.346***
	[–4.126 , –4.136]	[–5.337 , –5.356]
2021:Q1	–3.917***	–5.060***
	[–3.912 , –3.922]	[–5.050 , –5.070]
2021:Q2	–3.953***	–5.345***
	[–3.948 , –3.958]	[–5.336 , –5.354]
2021:Q3	–3.972***	–5.048***
	[–3.967 , –3.977]	[–5.039 , –5.058]
2021:Q4	–3.952***	–5.207***
	[–3.947 , –3.957]	[–5.197 , –5.216]
2022:Q1	–3.843***	–5.076***
	[–3.838 , –3.848]	[–5.067 , –5.085]
2022:Q2	–3.617***	–4.892***
	[–3.612 , –3.622]	[–4.883 , –4.901]
2022:Q3	–3.998***	–5.033***
	[–3.993 , –4.003]	[–5.024 , –5.042]
2022:Q4	–3.837***	–4.808***
	[–3.832 , –3.842]	[–4.799 , –4.817]
2023:Q1	–3.913***	–4.990***
	[–3.908 , –3.918]	[–4.981 , –4.999]
2023:Q2	–3.727***	–4.668***
	[–3.722 , –3.732]	[–4.659 , –4.677]
2023:Q3	–3.779***	–4.804***
	[–3.774 , –3.784]	[–4.795 , –4.814]
2023:Q4	–3.504***	–4.608***
	[–3.499 , –3.509]	[–4.599 , –4.618]
N	49,916,224	13,139,762

Notes: *** denotes statistical significance at the 1 per cent level. Coefficients are computed by adding estimated year dummy coefficient to base year coefficient and testing the difference from zero using a Wald test. All standalone coefficients are statistically significantly different from zero. Square brackets show lower and upper bound 95 per cent confidence intervals, which are based on the linear combination of base year and dummy year coefficients and take account of covariance between these parameters.

Sources: ABS; Authors' calculations.

Appendix D Supplementary exhibits

Graph D1: Response to a Cost-Push Shock under Alternative Price Rigidity Assumptions

Advertised prices

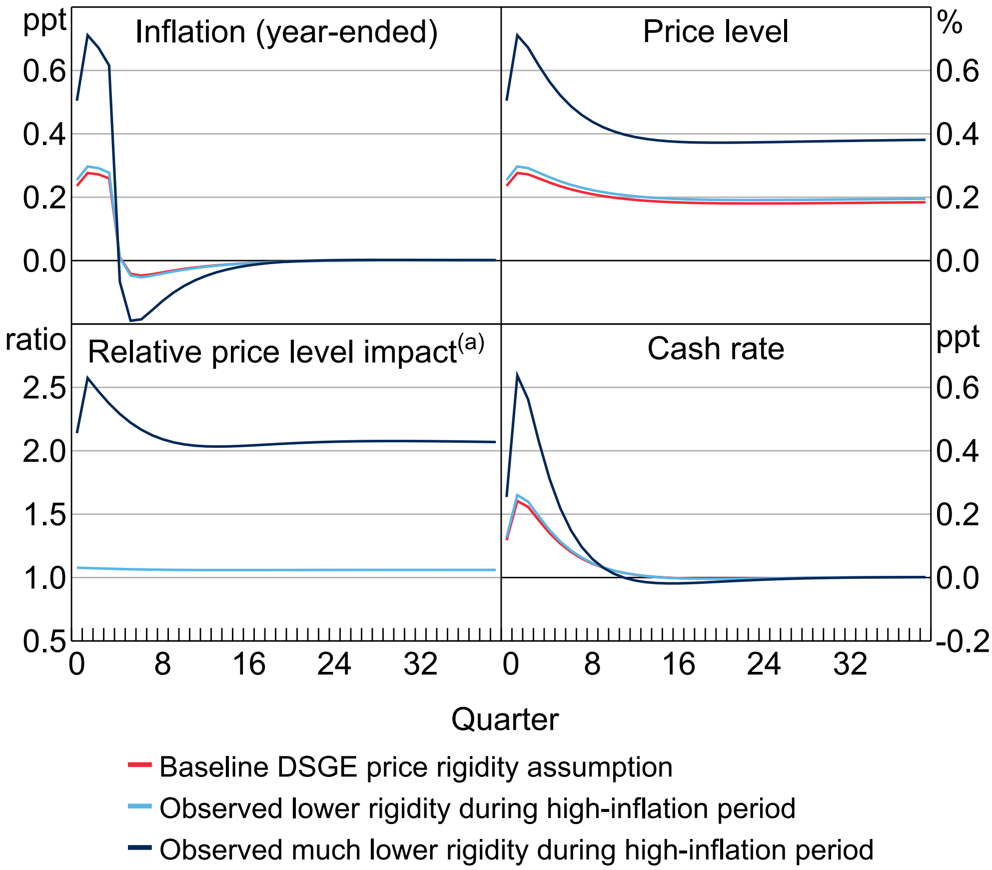

Notes: 100 basis point positive shock in the price of foreign imports.
(a) Ratio of the price level impact from the shock using alternative rigidity assumption over impact using baseline rigidity.

Sources: ABS; Authors' calculations.

Appendix E Impact of shifts in price rigidity on supply-demand decomposition

DSGE models can be used to estimate the relative contributions of supply and demand factors to inflationary shocks. This distinction is important because the effectiveness of monetary policy depends on the nature of the underlying shocks. Building on the approach in Beckers, Hambur and Williams (2023), we tested whether different assumptions about price rigidity influence the model's decomposition of deviations in trimmed mean inflation from its steady-state value (assumed to be 2.5 per cent, the midpoint of the RBA's target band).

Graph E1 presents the results. While the decomposition varies slightly across specifications, the differences are minor in magnitude. This suggests that the model's interpretation of the balance between supply and demand shocks is not materially affected by reasonable changes in the assumed degree of price rigidity.

Graph E1: Inflation Decomposition

Trimmed mean, year-ended, deviation from 2.5 per cent

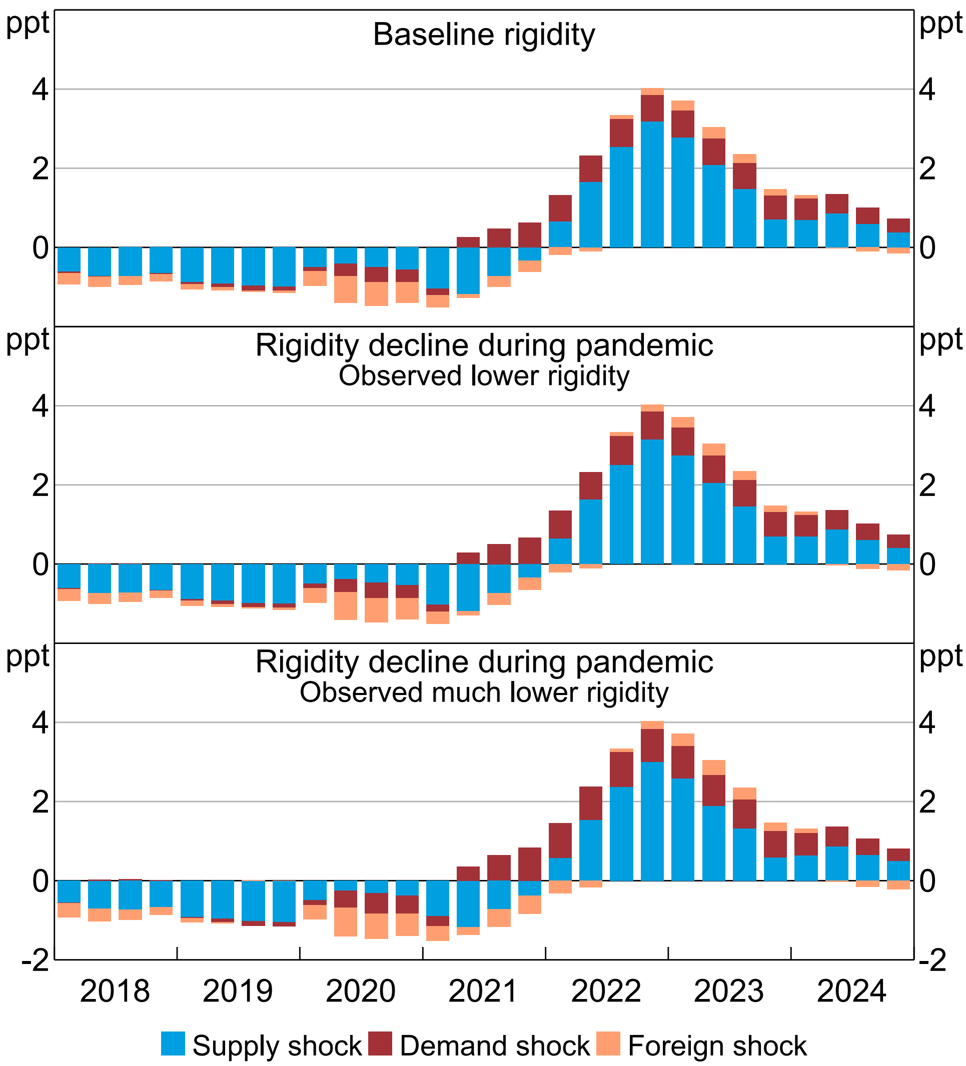

Notes: Supply shocks include mark-up, technology and labour supply shocks, and will capture some overseas shocks. Demand shocks include monetary policy, investment and consumption shocks. Foreign shocks include foreign demand and cost shocks.

Sources: ABS; Authors' calculations.

References

Adolfson M, S Laséen, J Lindé and LEO Svensson (2011), 'Optimal monetary policy in an operational medium-sized DSGE model', *International Finance Discussion Paper 1023, Board of Governors of the Federal Reserve System*.

Alvarez F, H Le Bihan and F Lippi (2016), 'The real effects of monetary shocks in sticky price models: a sufficient statistic approach', *American Economic Review*, 106(10), pp 2817–2851.

Auclert A, R Rigato, M Rognlie and L Straub (2024), 'New pricing models, same old Phillips curves?', *The Quarterly Journal of Economics*, 139(1), pp 121–186.

Australian Bureau of Statistics (2018), 'A guide to the Consumer Price Index: 17th series', *https://www.abs.gov.au/statistics/economy/price-indexes-and-inflation/guide-consumer-price-index-17th-series/latest-release*.

Beckers B, J Hambur and T Williams (2023), 'Estimating the relative contributions of supply and demand drivers to inflation in Australia', *RBA Bulletin*.

Benigno P and GB Eggertsson (2023), 'It's Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve', *NBER Working Paper 31197*.

Bils M and PJ Klenow (2004), 'Some evidence on the importance of sticky prices', *Journal of Political Economy*, 112(5).

Bilyk O, M Khan and O Kostyshyna (2024), 'Pricing behaviour and inflation during the COVID-19 pandemic: Insights from consumer prices microdata', *Bank of Canada Staff Analytical Note 2024-6*.

Blanchard O and J Gali (2010), 'Labor markets and monetary policy: a New Keynesian model with unemployment', *American Economic Journal: Macroeconomics*, 2(2), pp 1–30.

Blanco A, C Boar, CJ Jones and V Midrigan (2024), 'Nonlinear dynamics in menu cost economies? Evidence from U.S. data', *NBER Working Paper Series No 32748*.

Blanco A, C Boar, CJ Jones and V Midrigan (2024), 'Non-linear inflation dynamics in menu cost economies', *NBER Working Paper Series No 32094*.

Blanco A, C Boar, CJ Jones and V Midrigan (2025), 'The inflation accelerator', *NBER Working Paper Series No. 32531*.

Cagliarini A, T Robinson and A Tran (2010), 'Reconciling microeconomic and macroeconomic estimates of price stickiness', *RBA Research Discussion Paper 2010-01*.

Calvo GA (1983), 'Staggered prices in a utility-maximizing framework', *Journal of Monetary Economics*, 12(3), pp 383–398.

Cavallo A (2018), 'Scraped data and sticky prices', *Review of Economics and Statistics*, 100(1), pp 105–119.

Cavallo A and R Rigobon (2016), 'The billion prices project: using online prices for measurement and research', *Journal of Economic Perspectives*, 30(2), pp 151–178.

Cavallo A, F Lippi and K Miyahara (2024), 'Large shocks travel fast', *American Economic Review: Insights*, 6(4), pp 558–574.

De Fiore F, B Mojon, D Rees and D Sandri (2023), 'Monetary policy frameworks away from the ELB', *BIS Working Paper 1156*.

Dedola L, E Gautier, A Nakov, S Santoro, E De Veirman, L Henkel and B Fagandini (2023), 'Some implications of micro price-setting evidence for inflation dynamics and monetary transmission', *European Central Bank Occasional Paper Series No 321*.

Gagliardone L, M Gertler, S Lenzu and J Tielens (2025), 'Micro and macro cost-price dynamics in normal times and during inflation surges', *NBER Working Paper 33478 February*.

Gautier E, C Conflitti, D Enderle, L Fadejeva, A Grimaud, E Gutiérrez, V Jouvanceau, J-O Menz, A Paulus, P Petroulas, P Roldan-Blanco and E Wieland (2025), 'Consumer Price Stickiness in the Euro Area During an Inflation Surge', *Working paper*.

Gautier E, C Conflitti, RP Faber, B Fabo, L Fadejeva, V Jouvanceau, J-O Menz, T Messner, P Petroulas, P Roldan-Blanco, F Rumler, S Santoro, E Wieland and H Zimmer (2024), 'New Facts on Consumer Price Rigidity in the Euro Area', *American Economic Journal: Macroeconomics*, 16(4), pp 386–431.

Gelain P and P Lopez (2024), 'Understanding post-pandemic surprises in inflation and the labor market', *Federal Reserve Bank of Cleveland Economic Commentary*.

Gibbs CG, J Hambur and G Nodari (2021), 'Housing and commodity investment booms in a small open economy', *Economic Record*, 97(317), pp 212–242.

Guimaraes B and KD Sheedy (2011), 'Sales and Monetary Policy', *American Economic Review*, 101(2), pp 844–76.

Harding M, J Lindé and M Trabandt (2022), 'Resolving the missing deflation puzzle', *Journal of Monetary Economics*, 126, pp 15–34.

Henkel L and others (2023), 'Price setting during the coronavirus (COVID-19) pandemic', *European Central Bank Occasional Paper Series No 234*.

Ilut C, R Valchev and N Vincent (2020), 'Paralyzed by fear: rigid and discrete pricing under demand uncertainty', *Econometrica*, 88(5), pp 1899–1938.

Karadi P, A Nakov, GN Barrau, E Pasten and D Thaler (2024), 'Strike while the iron is hot: optimal monetary policy with a nonlinear Phillips curve', *BIS Working Papers No 1203*.

Klein M, K Strömberg and O Tysklind (2024), 'Inflation dynamics in the high inflation period: insights from microdata', *Sveriges Riksbank Economic Commentary NR 14*.

Klenow PJ and BA Malin (2010), 'Microeconomic evidence on price-setting', *Chapter 6 in Handbook of Monetary Economics*, Volume 3, pp 231–284.

Kryvtsov O and N Vincent (2021), 'The cyclicality of sales and aggregate price flexibility', *Review of Economic Studies*, 88(1), pp 334–377.

Lane PR (2024), 'The 2021-2022 inflation surges and the monetary policy response through the lens of macroeconomic models', *Speech at the SUERF Marjolin Lecture hosted by the Banca d'Italia November 18*.

Lowe P and L Ellis (1997), 'The smoothing of official interest rates', *RBA Annual Conference Volume, in Philip Lowe (ed.), Monetary Policy and Inflation Targeting, Reserve Bank of Australia*.

Maćkowiak B and F Smets (2008), 'On implications of micro price data for macro models', *European Central Bank Working Paper Series No 960*.

Montag H and D Villar (2022), 'Price-setting during the COVID era', *US Bureau of Labour Statistics Office of Survey Methods Research*, Economic Working Paper WP-547.

Nakamura E and J Steinsson (2008), 'Five facts about prices: a reevaluation of menu cost models', *The Quarterly Journal of Economics*, 123(4), pp 1415–1464.

Roberts JM (1995), 'New Keynesian economics and the Phillips curve', *Journal of Money, Credit and Banking*, 27(4), pp 975–984.

Rotemberg JJ (1982), 'Sticky prices in the United States', *Journal of Political Economy*, 90(6), pp 1187–1211.

Rotemberg JJ (1987), 'The New Keynesian microfoundations', *NBER Macroeconomics Annual*, 2, pp 69–104.

Rudolf B and P Seiler (2022), 'Price setting before and during the pandemic: evidence from Swiss consumer prices', *European Central Bank Working Paper Series*.

Schmitt-Grohé S and M Uribe (2025), 'Heterogeneous Downward Nominal Wage Rigidity: Foundations of a Nonlinear Phillips Curve', *Working paper*.

Sutton M (2017), 'The average size and proportion of price changes in the CPI', *Consumer Price Index, Australia, (September), 12-21*.

Taylor JB (1993), 'Discretion versus policy rules in practice', *Carnegie-Rochester Conference Series on Public Policy*, 39, pp 195–214.

4. Chapter 4: The effect of labour market concentration and declining entry rates on wages: Evidence from Australia

Jonathan Hambur

Economic Research Department, Reserve Bank of Australia (RBA), Crawford School of Public Policy, Australian National University (ANU)

Abstract

This paper documents the degree of labour market concentration and its effects on wages in Australia using administrative tax data. As in other countries, greater concentration within labour markets leads to larger wage markdowns. Moreover, declining firm entry and dynamism appears to have increase the impact of given levels of concentration, as it has meant there are fewer young, growing firms trying to poach from incumbents, raising the market power of incumbents. Quantitatively, declining entry rates account for a moderate portion of the weak wage growth in Australia, highlighting the importance of firm dynamism in labour market power.

JEL Classification Numbers: C23, C55, D22, D43, E24, J30

Keywords: wages, market power, labour markets, concentration

This work was completed while working in Australian Treasury's Macroeconomic Group and the Reserve Bank of Australia's Economic Research Department. I would like to thank Cyrille Schwellnus, Dan Andrews, Jeff Borland, Sephorah Mangin, and seminar participants at the RBA, RBA Quantitative Macro Workshop, OECD, Treasury, Labour Econometrics Workshop and TTPI Public Finance Conference for their helpful comments. An earlier version of this paper was released as RBA RDP 2023-02 'Did Labour Market Concentration Lower Wages Growth Pre-COVID?'.

4.1 Introduction

Across a variety of measures, wage growth was slow in Australia, and in many other advanced economies, in the years leading up to COVID-19. The impetus to understand weak wages growth is clear. Wages growth plays a crucial role in inflation dynamics, in determining government revenue, and in determining the welfare of individuals. An understanding of whether and what structural factors were weighing on wages is crucial to understanding future wage growth outcomes, and to potentially creating policy to counter the structural factors. To the extent that these structural factors remain, they could continue to weigh on wages in the post-COVID period.

One explanation that has received attention in recent years is the possibility of decreasing competition and increasing concentration amongst employers (e.g. OECD 2020). All else equal, where there are fewer employers competing for workers, workers have fewer outside options to leverage in wage negotiations. As such, employers tend to have more power to set lower wages: this is referred to as monopsony power.

The focus on employer concentration and market power in labour markets as a potential explanation for low wages growth is unsurprising given the documented increase in concentration in product markets (Hambur 2023 for Australia, Autor et al (2020) for the US and Bajgar et al (2019) for Europe and North America). Moreover, increases in employer concentration could help to explain other economic phenomena witnessed over the past decade, such as declines in measures of job-switching, which were evident both in Australia (Deutscher 2019) and overseas (Moscarini and Postel-Vinay 2017). If there are fewer employers to switch between, there is likely to be less job switching, all else equal.

This paper contributes to the growing literature documenting the effects of labour market concentration on wages empirically. I examine the effects using a comprehensive de-identified linked employer-employee dataset (LEED) based on administrative tax data. This allows me to look at the effect of firm-level market shares and local labour market concentration on wages, for the entire Australian economy, while accounting for other factors such as firm productivity, and number of plants.

The paper makes a few contributions to the literature. Most notably it explores the role that firm-dynamism can play in mitigating the effects of labour market concentration. In particular, I document the important role that firm entry plays in driving job-switching in Australia, and how the rate of firm entry influences the effect that a given level of concentration has a larger effect on wages. In doing so I bring together the existing literature on firm dynamism and labour dynamism (e.g. Bilal et a 2021), with the literature on labour market concentration and labour market power, highlighting how policies that support business dynamism can support wages growth not only by supporting productivity growth, but also worker bargaining power.

Second, I provide the first evidence on whether labour market concentration has affected wages growth for Australia. This is valuable for policy purposes. But beyond this Australian evidence is also generally interesting given its unique labour market institutions, which include differing minimum wages for different occupations, alongside enterprise and individual-level bargaining. These factors, particularly the differentiated minimum wage, have the potential to somewhat offset the effects of labour market concentration (e.g. Berger et al 2022), so Australia makes an interesting case study.

Third, I contribute to the ongoing debate around whether models of monopsonistic competition or oligopsony are more appropriate by testing both types of models in a common context.

In terms of my key findings, I find that wages tend to be lower in more concentrated markets (all else equal), consistent with economic theory and overseas findings. But even within markets there is a great deal of variation, with larger firms tending to exert more market power and therefore setting lower wages (after accounting for differences in productivity). This suggests that models of oligopsony may be more appropriate than models of monopsony, consistent with recent papers such as Berger et al (2022).

Second, while concentration has been broadly unchanged since the mid-2000s, the impact of concentration on wages appears to have increased over time. For any given level of concentration, its impact on wages has more than doubled, compared to the mid-2000s. So despite the fact that concentration remained broadly unchanged over the period, concentration may still have weighed on wages growth pre-COVID. Simple back-of-the-envelope partial equilibrium estimates suggest that wages were a little under 1 per cent lower on average from 2011–2015 than they would have been had the impact of concentration not increased.

Finally, I document the important role played by declining firm dynamism in these dynamics. I document the fact that firm entry is strongly associated with stronger outside options for workers and increased job-to-job switching, particularly from incumbent to young firms, creating more competition for workers. Consistent with this, in industries with lower entry rates, a given level of concentration has a larger effect on wages. This finding is robust to measuring entry rates as a firm or plant level, or to using churn rate (entry plus exit rates). It is also robust to controlling for other factors that have been shown to influence the effect of concentration over time, such as economic conditions, industry-switching rates, and union membership.

The next section of the paper provides an overview of the international literature on labour market concentration, monopsony and wages. Section 4.3 discusses the data used in the analysis. Section 4.4 documents labour market concentration in Australia, and how it has changed over time. Section 4.5 examines the relationship between wages and labour market concentration, while Section 4.6 examines explanations for the increasing impact of concentration. Section 4.7 then concludes.

4.2 Literature review and framework

This paper links to the large literature examining monopsony power and its implications for wages, employment, and optimal policymaking with respect to, for example, minimum wages. The topic has received a lot of focus in recent years in response to slow wages growth, and evidence of increasing inequality and declining labour shares in some advanced economies.[63]

Most closely related to this paper is the fast-growing literature exploring the role of concentration in labour market power. In thinking about the impact of concentration on monopsony power,

[63] For a more detailed literature review, see Manning (2020) and Azar and Mzrinescu (2024).

search-and-match type models are useful for intuition. In these models, workers search for job offers to gain or change jobs, while firms make job offers and search for workers. When a worker (n) and a firm match (i), the wage ($w_{n,i,t}$) is set to split the 'surplus' created by the job based on their bargaining power. This surplus is the difference between the revenue generated by the job ($p_{i,t}$), and the worker's outside options ($oo_{i,t}$), which includes unemployment (and associated benefit payments) and the chance finding other jobs ($X_{i,t}$).

$$w_{n,i,t} = \beta(X_{n,i,t}) * p_{n,i,t} + \left(1 - \beta(X_{n,i,t})\right) * oo_{n,t}(Concentration_{i,t}; X_{n,i,t})$$

As discussed in Jarosch, Nimczik and Sorkin (2024) and Schubert, Stansbury and Taska (2024), if there are only a small number of potential other employers the chances of getting an offer from one of them will be lower. This will lower the value of workers outside options, and therefore their ability to bargain for a higher wage.[64]

Numerous papers have demonstrated this link empirically. They find that increases in labour market concentration are associated with lower wages, all else equal. This has been documented for the US (Benmelech, Bergman and Kim 2022), UK (Abel et al 2018), Austria (Jarosch et al 2024) and in cross-country analysis (OECD 2021). These papers have varied in the data used to construct measures concentration, creating concentration metrics using employment stocks (Rinz 2018), job advertisements (Azar et al 2020) and new hires (Marinescu et al 2021). Work has also explored the extent to which affects come through new hires, or incumbent workers, with Bassanini, Batut and Caroli (2023) showing that both are important.

Most closely related to this paper is the portion of the literature trying to explore changes in the impact of concentration on wages over time. Empirically, Benmelech et al (2022) finds that concentration weighed more heavily on wages in the US in the period 2008 to 2016, compared to earlier periods. They link this to declining union membership and therefore lower offsetting collective worker power, finding that the impact of concentration on wages was smaller (in absolute terms) where union enrolment rates were higher. Abel et al (2018) find similar results for collective bargaining. Finally, Schubert et al (2024) documents the fact that employer concentration is less impactful in industries or occupations with greater occupational mobility, as workers can more easily change occupations, making their set of outside options larger.

This paper also relates somewhat to the part of the literature accounting for heterogeneity amongst firms within markets. While early papers focused on monopsonistic models, Berger et al (2021) recently extended the focus to oligopsonistic models, highlighting the potential for differing labour supply elasticities, and therefore wages (or wage markdowns), amongst firms with different market shares. Benmelech et al (2022) and Jarosch et al (2024) outline models with similar results. An important implication is that the entire distribution of market shares may be important in considering the effect on wages, and that aggregate concentration metrics could mask relevant changes. Moreover, when markets are granular, the emergence of new emerging firms who provide new additional outside options may be particularly important.

Finally, this paper also relates to the literature examining the relationship between firm dynamism and labour market dynamism. Most relevant, Bilal et al (2022) build an integrated model of firm

64 Berger et al (2022) derive a general equilibrium model with finite firms and inelastic labour supply across and within firms, which leads to similar conclusions.

dynamics and growth, and labour market matching. They show that in such a model new firms tend to grow by poaching workers. As such, shocks that lower rates of firm entry lead to slower rates of poaching and fewer vacancies, hence lowering labour market dynamism. Such dynamics align with empirical regularities identified by Haltiwanger et al (2018), where firms enter, grow, and poach workers from established firms, providing important competitive dynamics for incumbents.

4.3 Data and measurement

4.3.1 Data sample and market definition

For the analysis, I use administrative tax data on workers and firms in Australia. This dataset links the near universe of employees and their wages to employers (via annual Pay-As-You-Go PAYG statements). As such, it has demographic information on the workers (e.g. location), as well as detailed information on the firm, including its industry, starting date, revenue, expenses and employment. The dataset is highly representative, and the number of workers covered in each industry lines up well with ABS employment releases (see Appendix Table A1).

For most of the analysis I focus on the period 2004/5–2015/16. While data are available for a longer sample, for this period job-level location data are available, making for more accurate measures of plant size. I focus on the market sector, removing the Public Administration, Health and Education industries. This is because assessing market power in these sectors, where government and administered prices play a large role, is less sensible.

The literature tends to focus on concentration within 'local' labour markets. These are defined as the intersection of a location, and an industry or occupation. Accounting for location is important given it can be hard or costly for workers to move between areas. Similarly, many jobs have occupation- or industry-specific skills, preventing workers from moving easily between occupations or industries.

For locations, I use working zones constructed in Department of Infrastructure, Regional Development and Cities (2016). These are constructed using Census data on the home and work locations of workers and are designed to capture local areas within which people tend to both work and live. This makes them an ideal location definition. The results are robust to using alternate geographies such as SA4, or SA4 and Greater Capital City regions (Appendix A and B).

For industry, I use ANZSIC 2006 3-digit industries. I use industry rather than occupation data as it is available on a job-level, while occupation data are only reported annually on an individual level, and can be slow to be updated (Hathorne and Breunig 2022). Moreover, other analysis has shown that people may be slow to update their occupations in their tax reporting. The results of the analysis are robust to using occupation data, as well as less granular ANZSIC categories (Appendix A and B), consistent with findings in other papers (e.g. Handwerker and Dey 2018). However, given the lower quality of the occupation data, these results are robustness rather than the baseline. [65]

65 An alternative would be to use data driven industry groupings based on flows as in Schubert et al (2024). I explore a variation of this in Appendix C.

Based on these definitions, I have around 190 industries, 290 working zones. This creates around 55,000 possible permutations, of which only around half exist, leaving me with 25,000 local labour markets per year.

My unit of observation is a plant. This is the intersection of a firm and a locality.[66] So, if a large chain has stores in 2 different working zones, these will be considered 2 different plants. Based on this dataset I find that around half of all plants belong to multi-plant firms – i.e. firms with workers, and so operations, in multiple locations.

My measure of concentration is the Herfindahl-Hirschman Index (HHI), defined as:

$$HHI_{market,t} = \sum_{plant \in Market} \left(\frac{Plant\ headcount\ or\ wagebill_{\ plant,t}}{Market\ headcount\ or\ wagebill_{\ t}} \right)^2$$

This is a standard measure of concentration used in the literature, and can be derived theoretically as the relevant object. The HHI is bounded from 0 to 1. Low levels of the HHI indicate low levels of concentration (lots of small firms). High levels of the HHI indicate a highly concentrated market, with an HHI of 1 showing that there is only one employer.

Most papers focus on employment-based measure of concentration. However, Berger et al (2021) argue that wage-based measures are more relevant. For example, firms with lots of part time workers, or high turnover, may be measured as being too larger when employee-based measures are used. But wage-based measures are less affected by such biases. As such, I analyse both employee- and wage-based concentration measures, but mainly focus on employment-based measure for comparability to the literature. Results are robust to this choice (see Appendix A and B for wage-based concentration results).

As discussed below, in the regression analysis I control for productivity. Unfortunately, this can only be modelled at the firm-, rather than plant-level. For these regressions I focus on companies and exclude 2015/16 due to data constraints. I measure productivity as the ratio of value-added to employees. Value-added is defined as income less all expenses other than depreciation, interest expenses and wage expenses (i.e. these expenses are added back to the value-added measure). Wage rates and hours are not available, only wage payments. As such, wages are measured on a per-employee basis.

4.3.2 Summary statistics

Table 4.8 provides summary statistics on the degree of concentration in Australian local labour markets. Assessing whether labour markets are highly concentrated is difficult given the lack of clear guidance on what constitutes a 'concentrated' market. One benchmark that can be applied is whether the HHI is above 0.25. This benchmark is used by some competition regulators in product markets (e.g. US Department of Justice 2010), though whether it is appropriate for labour markets is unclear. On this metric around 58 per cent of markets are concentrated. However, these market

---

[66] Where firms have complex tax reporting structures I collapse them down to a single reporting unit nationally. But I still allow for different plants based on different geographic locations of workers in that single reporting unit.

account for only around 14 per cent of workers, as smaller markets are, unsurprisingly, more concentrated.

Table 4.8: Labour Market Concentration Metrics

Summary statistics

	Unweighted	Employment-weighted
25th percentile	0.153	0.011
Median	0.298	0.031
Mean	0.346	0.103
75th percentile	0.500	0.133
Share with HHI above 0.25	58%	14%

Note: Summary statistics for working zone by 3-digit ANZSIC industry local labour markets

While it can be difficult to compare metrics across countries due to different industry and locality definitions, Australian markets appear less concentrated than those in Europe. In the sample of European countries examined in OECD (2021) markets with an HHI above 0.25 account for around 20 per cent of workers, on average. Australian markets similarly appear less concentrated than US markets, based on evidence from Rinz (2018) (Figure 4.1), though are potentially more concentrated than UK markets (Abel et al 2018).

Figure 4.1: Employment-weighted Average Employment HHI

By sector, 2012

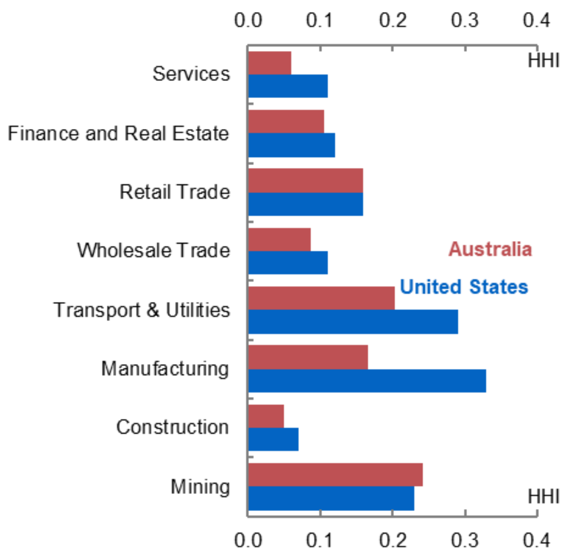

Note: Services include Accommodation & Hospitality, Information, Media & Technology, Professional Services, Administrative Services, Arts & Recreation, and Other Services. Interpretation of the results for mining might be affected by fly-in-fly-out work. Working zone by 3-digit ANZSIC industry local labour markets.

Sources: Author's calculations; Rinz (2018)

These average levels hide a great deal of heterogeneity, as shown in Table 4.8. Part of this heterogeneity reflects difference across industries. As in the US, industries such as mining, manufacturing, transport and utilities, and retail trade tend to be more concentrated. In contrast, services, including accommodation and hospitality, tend to be less concentrated, though this could in part reflect the fact that industry classifications tend to be less detailed in the services sector.[67]

Much of the variation also reflects geographic differences. Markets in major cities tend to be less concentrated (Figure 4.2). Unsurprisingly, regional and rural areas tend to have more concentrated labour markets

Figure 4.2: Unweighted Average Employment HHI

By remoteness, 2005-2016 average

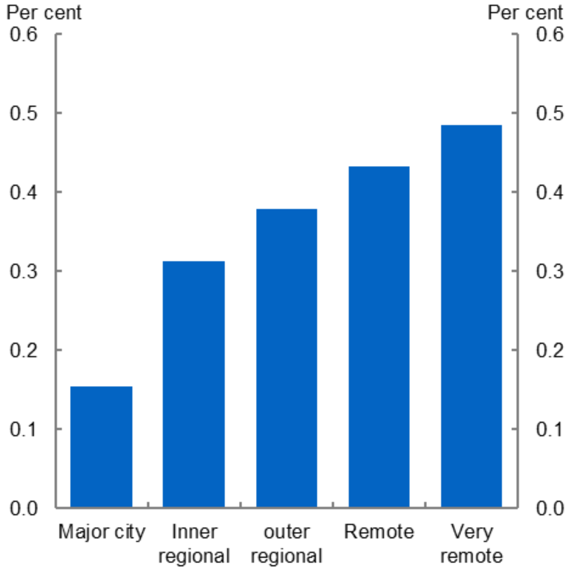

Note: Remoteness based on ABS remoteness structures. Working zone by 3-digit ANZSIC industry local labour markets.

Overall, local labour market concentration has not changed in Australia over the sample. Figure 3 plots the average level of concentration. Using an employment-weighted average of the individual labour markets shows some decline over the period. However, this reflects a compositional shift with workers moving to urban centres, which tend to be less concentrated. The unweighted average has been broadly unchanged (Figure 4.3), while the distribution has been similarly stable (Figure 4.4). Taken together, this shows that concentration has been broadly unchanged over the period.

67 Some caution should be taken in interpreting the mining results, given the use of fly-in-fly-out (FIFO) workers. Results for construction may also be affected by the prevalence of contractors, who will not be recorded as employees.

Figure 4.3: Average Employment HHI Over Time

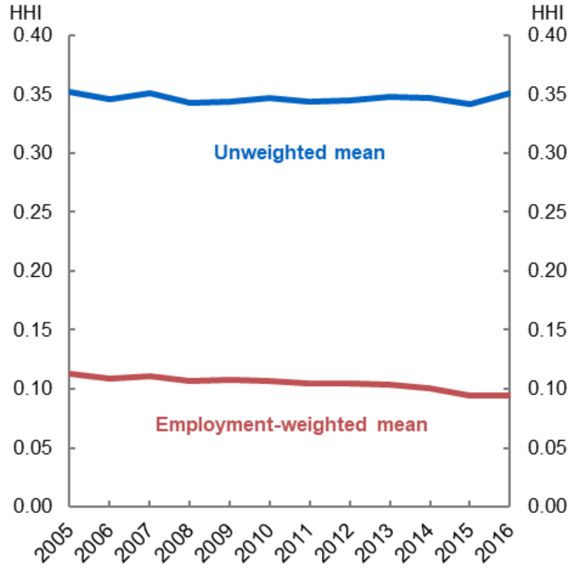

Note: Working zone by 3-digit ANZSIC industry local labour markets.

Figure 4.4: Distribution of Employment HHI Over Time

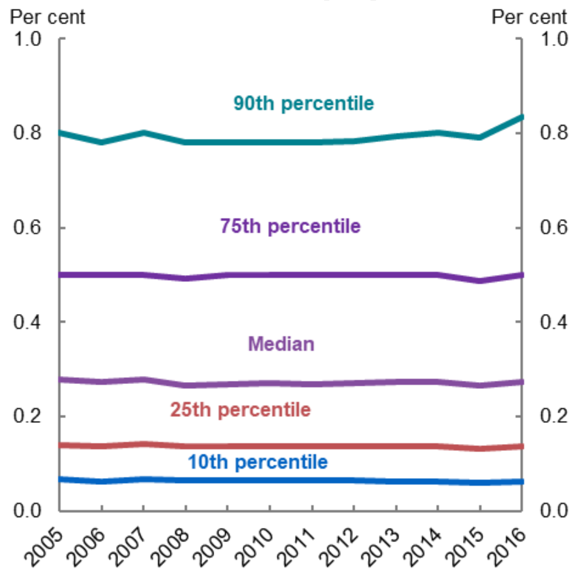

Note: Working zone by 3-digit ANZSIC industry local labour markets.

4.4 Wages and Concentration

To provide some initial evidence on the relationship between concentration and wages, Figure 4.5 plots changes in the HHI index over the sample for labour markets against their average annual real wage growth. Real wage growth has tended to be higher in local labour markets where the HHI has declined, providing some initial evidence that local labour market concentration is associated with lower wages.

Figure 4.5: Employment HHI Versus Wages Growth

By local labour market, changes from 2004/5 to 2015/16

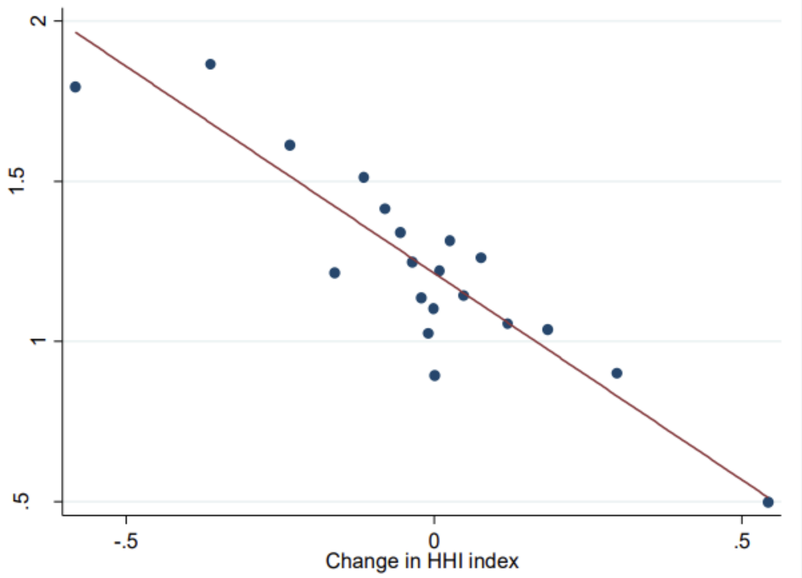

Note: Plots changes in HHI in local labour markets (x-axis) from 2004/5 to 2015/16, against the average growth in their real wages over the same period. Dots each represent the average of a number of individual local labour markets. Line is line of best fit. Working zone by 3-digit ANZSIC industry definition of local labour markets.

4.4.1 Model specification and identification

To test this more formally, I now turn to a regression framework. In specifying the regression, it is important to think about how wages are set.

As discussed in Manning (2020), in many models of labour market power the firm optimally sets its wages some fraction below its marginal revenue product of labour (MRPL), which is closely linked to labour productivity. The extent of this markdown will be a function of the firms' market power (the elasticity of its labour supply curve $1/\varepsilon_{i,t}$). As such, these models predict that the wage (for the market or firm) can be written as: [68]

$$w_{i,t} = MRPL_{i,t}\left(\frac{1}{\varepsilon_{i,t}+1}\right)$$

$$\ln(w_{i,t}) = \ln(MRPL_{i,t}) - \ln(\varepsilon_{i,t}+1)$$

In our case, the market power is summarised by the measure of concentration. The degree of market power could be market-specific, or firm/plant-specific, depending on the model.

The main model I estimate assumes that the market structure is one of monopsonistic competition. Wages are set as markdown on the MRPL, with this mark-down, and so degree of market power, being the same for all firms in the market (m) and being related to the market concentration. This is the standard regression framework used in most papers (e.g. Rinz 2018; Azar 2020)

[68] Search and match models such as Schubert et al (2024) also often reduce to such a characterisation.

$$\ln (w_{m,t}) = \gamma * \ln (Prod_{m,t}) + X_{m,t} + \beta * Concen_{m,t} + \epsilon_{m,t}$$

In each case, the coefficient of interest is β. For market-level regressions, $Prod_{m,t}$ and $w_{m,t}$ are weighted averages (though this makes little difference).

I also consider versions of the model assuming oligsopsony allowing the degree of market power to vary with plant size. In this case the mark-down differs based on the plant's (i) labour market power, which is related to its market share. This oligopsonistic framework is more in line with recent papers such as Berger et al (2022).

$$\ln (w_{i,t}) = \gamma * \ln (Prod_{i,t}) + X_{i,t} + \beta * Plant\ Market\ Share_{i,t} + \epsilon_{i,t}$$

For both models X_t contains several controls, including industry by time and location by time effects, and market size. These controls help me to strip out factors affecting wages for all firms or markets, such as local or industry demand conditions and prices, and focus in on changes in wages and concentration that reflect changes in market power. I also include market fixed effects, which ensure that I am focusing on changes over time, and abstracting from confounding factors that affect the level of concentration and wages.

Ultimately, what the regressions are trying to capture any change in the markdown over MRPL (i.e. labour market power) that reflects a change in concentration. For example, entry of a new competitor that lowers concentration and market power, and so leads to a smaller markdown and higher wages. Or equally a firm closure or merger that raises concentration and market power, and leads to a larger markdown and lower wages. Or similair changes for the firm in the firm-level regressions.

This leaves two key sources of bias in the regressions. The first is classical measurement error in the concentration measure, and in particular variation over time that does not relate to changes in market power (due to the inclusion of fixed effects). This will lead to an attenuation bias.

The second, and more concerning is endogeneity stemming from other shocks that drive both concentration and wages in the local market, but do not relate to changes in the markdown. One natural set of shocks are industry- or local-level economic shocks. For example, increases in concentration in the manufacturing sector that coincide with the general shrinking of the manufacturing sector over time that are also associated with weaker wages in that sector. Or similarly general economic changes in a local area, such as weak conditions in locations that were previously reliant on manufacturing. Inclusion of industry-time and location-time fixed remove these trends, ensuring that identification will be coming from changes in concentration in the local market.

Still there may be other local shocks that influence wages and concentration without influencing the markdown. For example, consider a firm-specific demand shock that pushes up its prices and MRPL, and causes it to gain market share. As I include a measure of productivity this concern is lessened somewhat – any shock that influences wages through influencing MRPL should be accounted for, with only shocks that influence the markdown remaining.[69]

Nevertheless, as productivity measurement is imperfect, I also estimate a version of the monopsony model where I use an instrument for concentration in the local market. Specifically, I use the average of the (inverse of the) number of plants in the industry in other locations, with each location being weighted by its share of industry employment. This approach is used in several papers (e.g. Azar et al 2020). The use of the number of firms, rather than the HHI, should lessen the impact of shocks to individual firms, while focusing on national changes allows us to abstract from local shocks that may influence concentration and wages. That said, this approach makes it difficult to include industry by time effects, and so is not the preferred specification.[70]

One final concern is that the assumption of constant passthrough from productivity to wages, as captured by the above models may not be accurate. In fact, in some models imply that changes in concentration can influence the passthrough rate directly, rather than just the markdown. As such, as a robustness I use a passthrough specification of the regression. This leads to similair results (Appendix C).

4.4.2 Baseline results

Table 4.9 shows the results from the monopsony model. Running the model at the local labour market level (*Local Labour Market Models*), we can see that increases in concentration are associated with lower wages, consistent with expectations. The coefficient is larger using the instrumental variable. This suggests that changes in the local HHI driven by national-level variation in concentration have a larger impact on wages, and provides additional support to the baseline findings.

69 A concern working the other way could be that there is a mechanical relationship between market share and productivity, meaning productivity is a bad control, especially for the firm-level regressions. This concern must be weighed up against other sources of endogeneity from tis exclusion.

70 Results are also robust to using the lagged HHI to abstract from contemporaneous shocks. See Appendix B. Another approach would be to try to instrument using some 'direct' measure of changes in competition, such as mergers, as done in Benmelech et al (2022). Recently constructed databases of mergers may facilitate such approaches in the future (Win et al 2024).

Table 4.9: Monopsony and Oligopsony models

Employment concentration

	Local Labour Market Models		Plant-level Models		
	Monopsony	Monopsony – IV	Monopsony	Monopsony – IV	Oligopsony
Productivity	0.077***	0.070***	0.153***	0.153***	0.158***
	(0.001)	(0.002)	(0.004)	(0.001)	(0.004)
Concentration - HHI	-0.122***	-0.542**	-0.005	-0.021**	
	(0.008)	(0.253)	(0.011)	(0.009)	
Concentration – plant market share					-0.103***
					(0.027)
R^2	0.67	0.009	0.781	0.042	0.801
Observations	302,000	302,000	3,660,000	3,660,000	3,560,000
Fixed effects					
Market	Y	Y	Y	Y	N
Time*Industry	Y	N	Y	N	N
Time*Location	Y	Y	Y	Y	N
Time*Market	N	N	N	N	Y
Plant	N	N	Y	Y	Y

Note: All models contain controls for market size (number of workers in the market), and plant-level models account for number of plants in the parent firm.

The final three columns show results run at the plant-level. Doing so using the aggregate HHI there is less evidence of a negative relationship with wages. However, if we replace the aggregate HHI with the firm market share there is a strong negative relationship. This suggests that Australian labour markets are better represented as oligopsonistic, with larger firms within local labour markets having larger wage markdowns. This is even the case once including time-by-market fixed effects, so accounting for the average markdown in the market.

Figure 4.6 provides a sense of the magnitude of these results. Focusing first on the regressions using market concentration, wages are on average 4½ per cent lower in a highly concentrated industry (75th percentile HHI of 0.5), compared to a low concentration industry (25th percentile HHI 0.15). But gaps don't just exist across markets. Even within markets there is a great deal of difference between high and low market share firms. High market share firms pay wages that are 2 per cent lower (once accounting for differences in productivity and other factors.

Figure 4.6: Gap in Average Wages and Rent-sharing

Gap between high and low productivity firms and markets

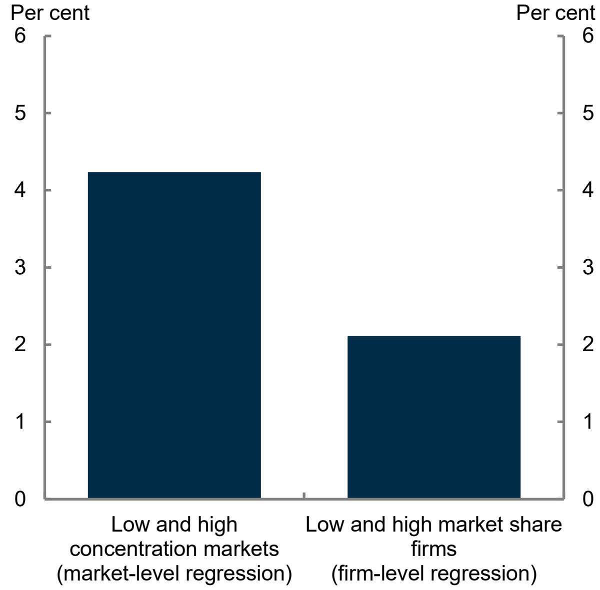

Note: Low and high concentration markets are the 25th and 75th percentile, respectively. Low market share is zero, high market share is 20 per cent of market (95th percentile). Results from headcount models, either market-level or firm-level, as indicated. Red dashes show 2 standard deviation range of estimates.

One concern with the results above is that firms may set wages at a firm-level, not a plant level. As such, the inclusion of multi-plant firms may bias the results towards zero. We do find some evidence for this, with the estimates of the coefficient on the HHI being more negative when we focus only on single-plant firms (though it is still negative and significant for multi-plant firms – Appendix Table B8). Given the growth of multiple-plant firms, this is likely to bias the coefficient on the HHI towards zero over the sample, working against our findings in the next section.

4.4.3 Does the effect of concentration change over time?

As noted above, the degree of labour market concentration has been unchanged over the sample. At first blush, this would suggest that concentration cannot explain low wages growth pre-COVID. However, papers such as Benmelech et al (2022) have found evidence that the impact of given levels of concentration can change over time.

To examine this, I run the baseline market-level regression, but allowing all coefficients to vary across 3 financial year sub-samples: 2005–2007; 2008–2010 (GFC); 2011–2015 (post-GFC). This will allow me to consider whether the impacts of HHI on wages have increased, even if concentration has remained unchanged.[71]

[71] These results are robust to changing the period definition. Note market fixed effects are not allowed to differ across samples, so variation in wages and HHI across sample periods is allowed to affect the results. So we are in some senses identifying based both on whether changes in HHIs in the later samples had a larger impact on wages, as well as whether markdowns became relatively larger in more concentrated markets across the periods. Allowing the market fixed effects to vary and focusing on the former would be a stricter test, but is not feasible due to the short sample period. Moreover, the latter aspect is also of relevance, particularly in thinking about the macro impact of the change.

Table 4.10 shows the results of the sub-sample regression. The impact of concentration more than doubled from the early period to the 2011–2015 period.[72]

Table 4.10: Sub-sample results

Employment concentration

	Baseline	Sub-samples	National unemployment	State unemployment
Concentration	-0.122***	-0.067***	-0.123***	-0.122***
	(0.008)	(0.020)	(0.015)	(0.015)
Concentration*GFC		-0.076***		
		(0.021)		
Concentration*Post GFC		-0.079***		
		(0.022)		
Concentration*Unemployment			-0.359	-0.035
			(1.557)	(1.106)
R^2	0.67	0.67	0.67	0.67
Observations	302,000	302 ,000	302 ,000	302 ,000

Note: All regressions include a set of common controls covering market characteristics (number of workers in the market), as well as market, time*industry and time*location fixed effects. For GFC models, all controls other than market FE are interacted with period dummies. Pre-GFC is 2005–2007. GFC is 2008–2010. Post GFC is 2011 to 2015. For unemployment models, interacting unemployment with productivity does not change the results. Errors clustered at the local market level.

One potential explanation could be softer labour markets. In many search-and-match models a tighter labour market, leads to greater bargaining power for works given the greater chance of finding another jobs, which could lower the impact of concentration. However, this does not appear to be the case, as there is no significant evidence that the impacts of concentration on wages are stronger when the labour market is weaker, as measured by the national or state-level unemployment rate (Columns 3 and 4).

To put the sub-sample result in perspective, Figure 4.7 shows the estimated wage gap between a high and low concentration markets between the early and recent periods. This gap is around to 2 per cent in the early period, but 5 per cent in the more recent period.

72 Excluding the mining sector does not impact the results. As such, it does not appear that the inclusion of the mining boom in the initial period is directly accounting for the results.

Figure 4.7: Difference in Average Wages between High and Low Concentration Markets

By period

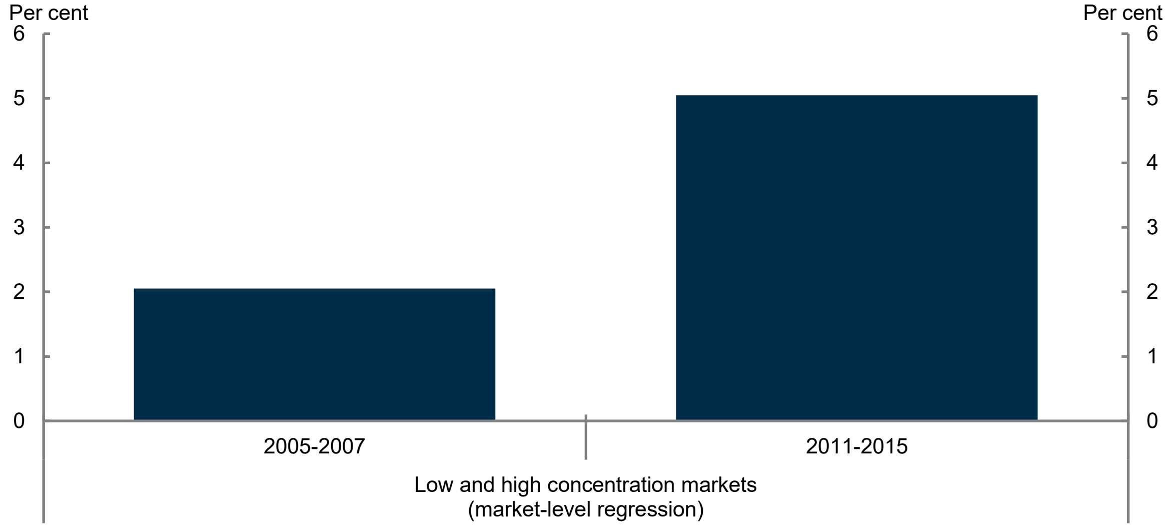

Note: Low and high concentration markets are the 25th and 75th percentile, respectively. Results from headcount models estimated at the market level shown over sub-samples in Table 4.

To get a sense of the macroeconomic impacts I use a simple back of the envelope calculation. The national employment-weighted average HHI is around 0.1. Multiplying this by the change in the HHI coefficient between periods (0.79) suggests that nationally wages were around 0.80 per cent lower than they would have been on average from 2011–2015, compared to the case whether the coefficient did not change. To put this in context, wages growth towheads the end of this period was around 2.5 per cent.

Obviously, the calculation is very simple and does not account for heterogeneity in outcomes or general equilibrium mechanisms. In particular, Berger et al (2022) show that the reduced form parameters estimated in regressions such as the above will not equate to the true structural parameters due to general equilibrium effects. So static calculations such as these should be interpreted with some caution. Still it does suggest that the increasing impact of concentration has had a substantial effect on aggregate wages.

4.5 Firm entry, concentration and labour market power

4.5.1 Motivating evidence and mechanism

There is a growing literature linking firm dynamism and labour market dynamisms. In particular, Bilal et al 2022 builds a model where firms enter and attract staff away from incumbent firms, meaning that declines in firm entry can be associated with lower job-to-job switching and few outside options for workers. These models coincide with empirical evidence in papers such as Haltiwanger et al 2018, where young growing firms tend to attract workers way form incumbent firms, leading to greater labour market dynamism.

Such dynamics are evident in Australia. Much of the decline in job-to-job switching in Australia has reflected less workers switching from established mature firms to young firms. And this has coincided with declines in firm entry rates (Figure 4.8).

These patterns are also evident within local areas. Figure 4.9 shows that there is a strong positive correlation between firm entry rates in a local area and the rate of job-to-job switching. This is the case even controlling for time-invariant differences across geographies and the overall business cycle. Moreover, this relationship is only evident when focusing on switching between mature and young firms (Figure 4.10). There is no relationship between firm entry and mobility of workers between mature incumbent firms (Figure 4.11).

Figure 4.8: Firm Entry Rates and Rate of Job-switching

By firm age

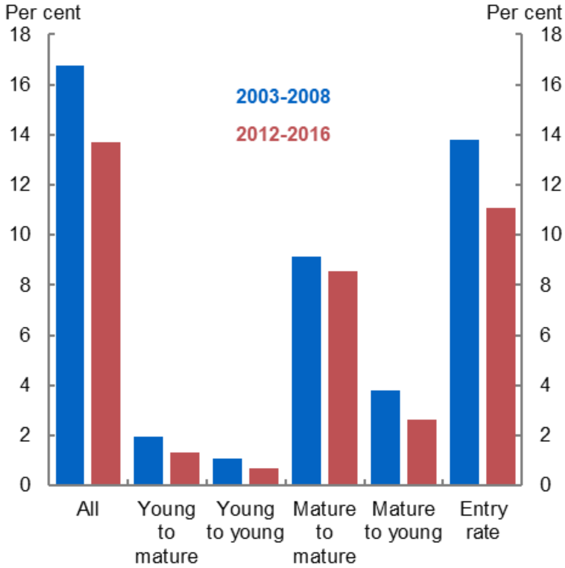

Note: Young firms are 0–5, mature firms are older than 5 years.

Figure 4.9: Firm Entry Rates and Job-to-Job Switching Rate

By SA4 region, over 2001/2 to 2015//16

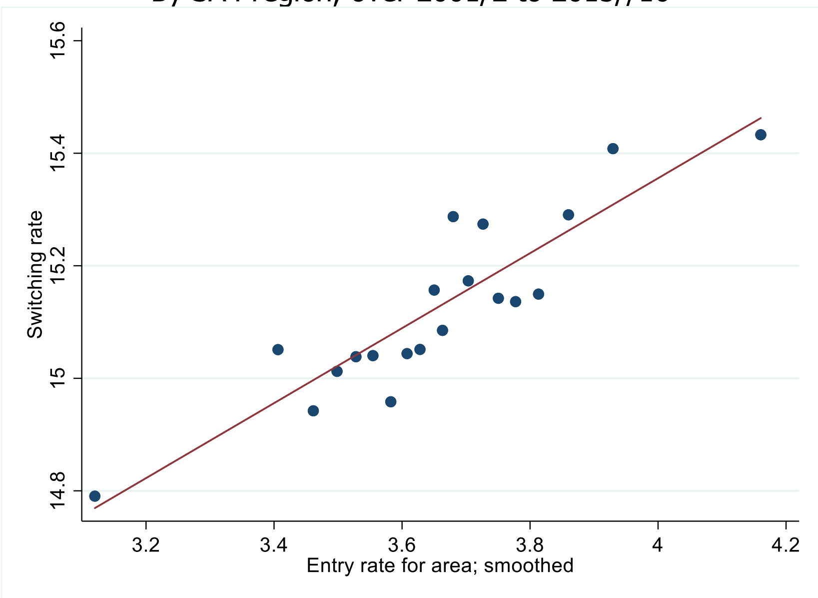

Note: Entry rates for SA4s or greater capital city areas versus the job-to-job switching rate. This is shown having residualised with respect to time effects, area effects, and the employment share of young firms. Approximately 780 SA4*year combinations based on 120 million workers. Entry rate calculated as entries, divided by sum of previous and current number of firms (smoothed).

Figure 4.10: Entry Rate and Young-to-mature Firm Job-to-Job Switching Rate

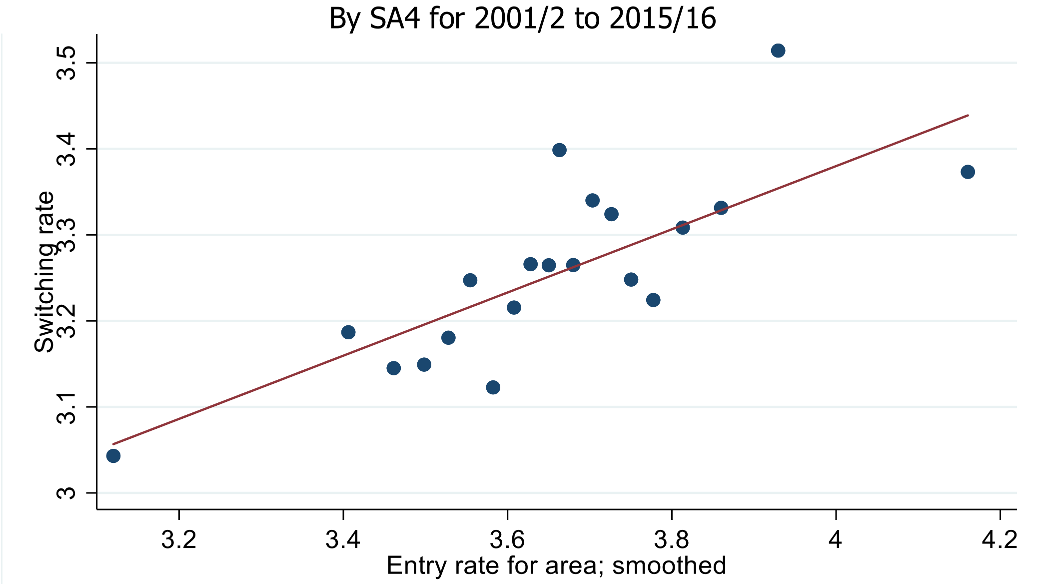

Note: Plots entry rates for SA4s or greater capital city areas versus the job-to-job switching rate. This is shown having residualised with respect to time effects, area effects, and the employment share of young firms. Approximately 780 SA4*year combinations based on 120 million workers. Entry rate calculated as entries, divided by sum of previous and current number of firms (smoothed). Young firms are 0–5, mature firms are older than 5 years.
Source: Author's calculations

Figure 4.11: Entry Rate and Mature-to-mature Firm Job-to-Job Switching Rate

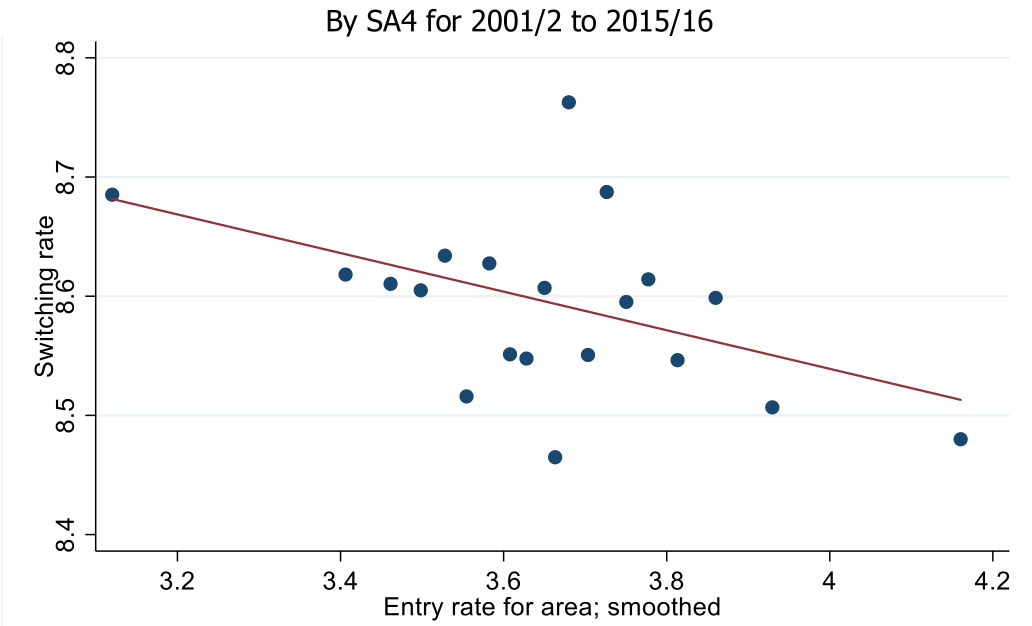

Note: Plots entry rates for SA4s or greater capital city areas versus the job-to-job switching rate. This is shown having residualised with respect to time effects, area effects, and the employment share of young firms. Approximately 780 SA4*year combinations based on 120 million workers. Entry rate calculated as entries, divided by sum of previous and current number of firms (smoothed). Young firms are 0–5, mature firms are older than 5 years.
Source: Author's calculations

Taken together, these results suggest that firm entry represents an additional source of competition for labour for incumbents, with entrants poaching staff as they grow and so supporting job switching. Entrants could therefore broaden the set of outside option for workers, and so raise their bargaining power. As such, declining firm entry and dynamism could make concentration amongst incumbents more impactful.

To make this more concrete, take for example the framework outlines in Jarosch et al (2024). In this framework the equilibrium relationship between the wages and concentration is (Equation 9):

$$\overline{w} = 1 - (1 - \alpha) * \frac{1 - \beta(1 - \delta)}{1 - \beta(1 - \lambda\alpha[1 - C] - \delta[1 - \tau C]}$$

Where $\overline{w}$ is the average wage, α is the worker share of income in a standard model, β is the discount factor, δ is exogenous separation rate, τ ar the higher order effects of concentration through potentially meeting a previous employer, and C is the degree of concentration. We can see that the effect of concentration on the wage is smaller where the term λ, the job finding rate, is lower.

While in many cases we would think of this object as a cyclical variable, structural changes in the number of vacancies would also affect the job finding rate. In particular, as shown in Bilal et al (2022), shocks that lessen the rate of firm entry lead to lower rates of vacancies and poaching, as there are fewer new firms looking to grow and poach. So falling rates of firm entry could make concentration more impactful in this framework.

One final point to highlight is that, while firm entry and market concentration will naturally be linked, given levels of concentration can be supported by different entry rates. Concentration reflects the combination of firm entry, growth and exit, so a given level of concentration could exist with higher or lower entry rates, depending on these other factors. For example, an HHI of 0.25 could be supported with four equal size static firms, or with firm entry, growth and exit.

4.5.2 Regression framework

To examine the role of declining firm entry, I extend the baseline market-level monopsony model to incorporate firm entry rates. Specifically:

$$\ln(w_{i,t}) = \gamma * Prod_{i,t} + X_{i,t} + \beta * \ln(Concen_{i,t}) + \delta * \ln(Concen_{i,t}) * Entry_{ind,t} + \epsilon_{i,t}$$

$Entry_{ind,t}$ is the national rate of entry in the relevant ANZSIC 3-digit industry, demeaned across time. I consider entry rates both constructed at a firm-level, but also at a plant level (i.e. firms entering into a new geography).

I focus on the national, rather than local rate of entry, to avoid endogeneity between local conditions and entry rates, as well as any mechanical feedback from changes in local entry rates onto concentration. It also helps to limit noise introduced by very small markets, where rates can be volatile. Entry rates are demeaned for each industry. As such I am focusing on changes in switching rates and abstracting from structural differences across industries that might affect entry, concentration and wages.

The coefficient of interest in this case is δ. This captures whether a given level of concentration has a larger or smaller effect on wages in industries where entry rates have risen. If concentration amongst incumbents is less impactful in industries with rising entry, the coefficient would be positive.

Table 4.11 shows the results of the regression with entry rates constructed using the entries of firms, and entries of plants (i.e. recording a firm opening a plant in a new locality as an entry). The results consistently show that, where industry entry rates are higher, local labour market concentration is less impactful.[73] That is, having a high level of firm entry seems to provide an additional source of competition, lessening the market power associated with concentration amongst incumbents.

Table 4.11: Firm Entry Rate Models

Employment concentration

	Baseline	Plant-based entry rate	Firm-based entry rate
Concentration	-0.122***	-0.110***	-0.111***
	(0.008)	(0.016)	(0.015)
Concentration*Entry Rate		0.800***	0.540***
		(0.228)	(0.204)
R^2	0.670	0.683	0.686
Observations	302,000	274,000	274,000

Note: All regressions include a set of common controls covering firm characteristics (number of plants in the firm) and market characteristics (number of workers in the market), as well as plant fixed effect and time*market fixed effects. Errors clustered at the local market level.

These results are robust to various choices. This includes using plant- or firm-based entry rates, using churn rates (entry plus exit), and using previous number of firms or the average number of firms across t-1 and t as the denominator. It is also robust to accounting for other factors that have been found to influence the affect of concentration in the literature, like union membership rates and cross-industry switching rates and economic conditions (see Appendix C).

To provide a sense of the macro importance of declining entry rates I do a similar calculation as for the sub-sample regressions. On average there was around a 5 percentage point decline in the entry rate of plants and firms across the sample. Applying this to the coefficients from Columns 2 and 3 of Table 4.11 and using the employment-weighted average HHI (0.1), this suggests that wages would have been around 0.25–0.50 per cent higher in 2015, had entry rates not declined.

Again, the estimate should not be interpreted too precisely as it is a simple, partial equilibrium estimate. But these results do suggest that declining firm entry and dynamism had a substantial negative impact on worker's wages by lowering their bargaining power and increasing the power of incumbent firms.

It is important to note here that while declining firm entry appears to be a proximate cause of the increasing impact of concentration on wages, in turn it could reflect other factors. For example, it may be that access to finance became harder post-GFC, and this lowered firm entry, which in turn

[73] These results are robust to using different methods for calculating entry rates that are more robust to outliers, and using measures of churn, rather than entry (results available on request)

influenced the effect concentration on wages. So while we interpret our findings as pointing to the proximate mechanism behind the increasing impact of concentration, what caused entry rates to decline is unknown.

4.6 Conclusion

The impact of labour market concentration and power on wages has been an area of increasing interest for academics and policymakers across the globe. This paper contributes to the growing body of evidence documenting the affects of labour market concentration on wages.

Most notably, it has highlighted the important role that firm entry and dynamism plays in offsetting the effects of labour market concentration. Where entry rates are higher, there are more young firms looking to grow and poach workers from incumbents, and this appears to offset some of the effects of concentrated labour markets wages. This highlights how policies that improve economic and firm dynamism can be pro-worker, as well as supporting aggregate productivity growth.

The paper also highlights the important role that labour market concentration and declining economic dynamism has played in weakening worker bargaining power and wages in Australia. While labour market concentration did not increase, the increasing impact of labour market concentration may have weighed on wages growth pre-COVID, subtracting a bit under 1 per cent from wages on average from 2011–2015. Much of this reflected declining firm entry and dynamism, as it has meant that incumbents faced less competition from new entrants, allowing them to exert more market power for given concentration levels.

Appendix A: Additional Charts and Tables

Figure A2: Employment by Industry

LEED versus LLFS, 2001/2 to 2015/16

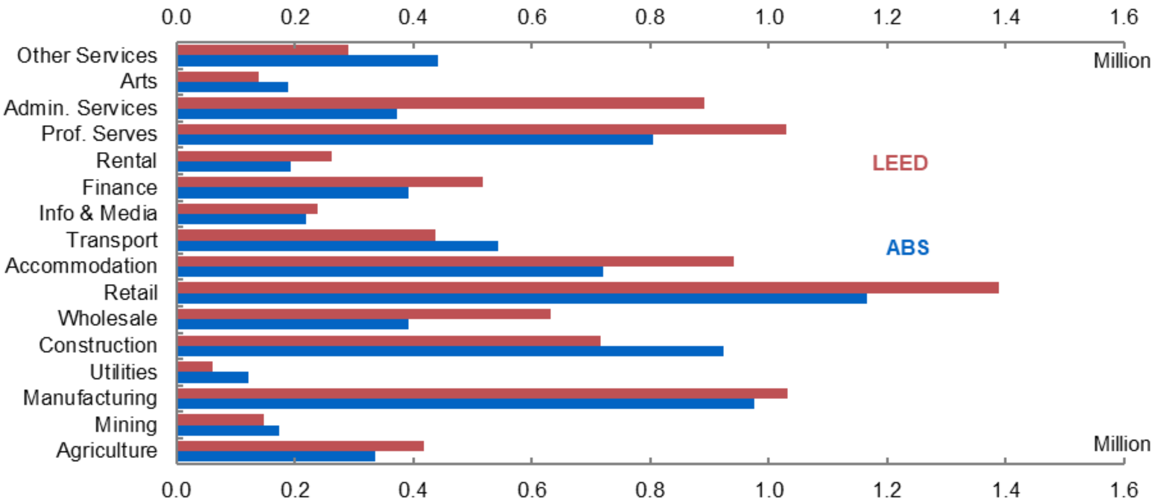

Notes: ABS data is from Labour Force Survey detailed and is on a person basis. LEED is from administrative tax data and is displayed on a job basis. As such, sectors with high rates of multiple-job holdings, like Retail and Accommodation, are larger in LEED. LEED also excludes self-employed, which may be more important in Construction. The higher LEED employment in Administrative Services reflects differences in how employment in labour hire and employment services are captured in tax data (via the service providers), compared to in labour statistics (via the end employer).

Source: Author's calculations; ABS

Table A1: Labour Market Concentration Metrics

Summary statistics, wage-based measures

	Unweighted	Employment-weighted
25th percentile	0.186	0.011
Median	0.363	0.031
Mean	0.401	0.103
75th percentile	0.584	0.133

Note: Summary statistics for working zone by 3-digit ANZSIC industry local labour markets

Figure A3: Unweighted Average Employment HHI

By working zone 2005/6 to 2015/16

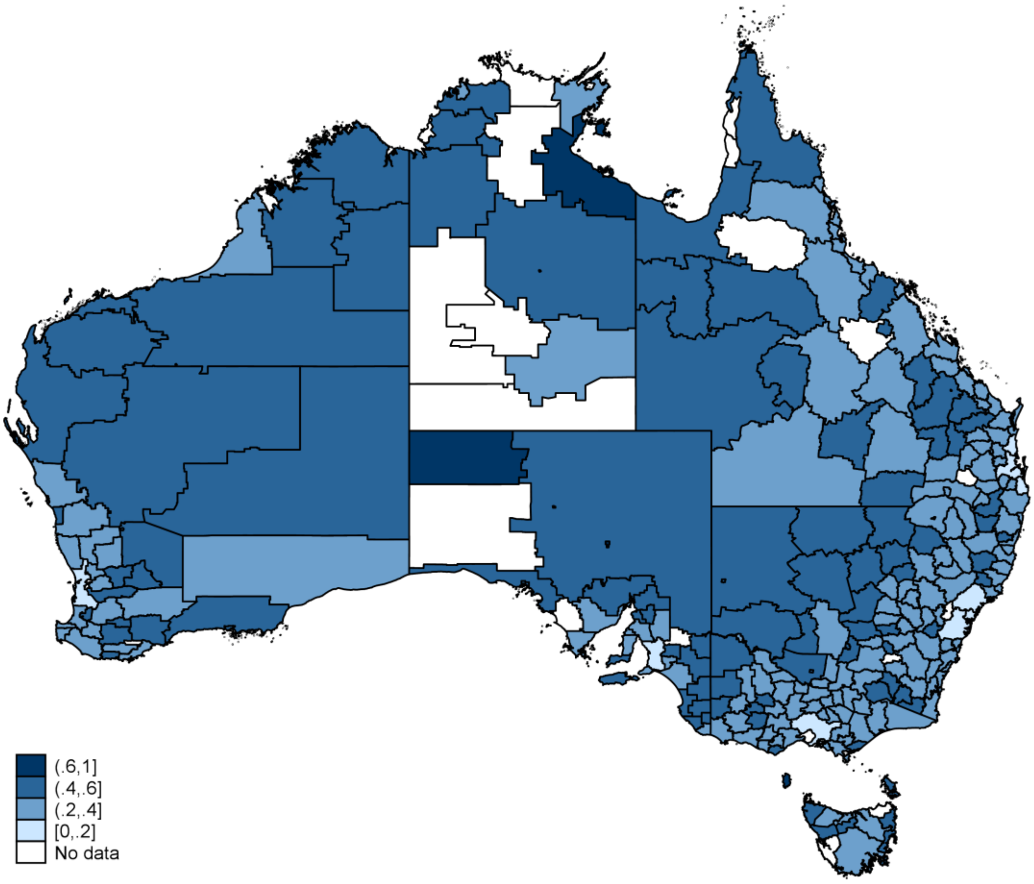

Figure A4: Average Wage HHI

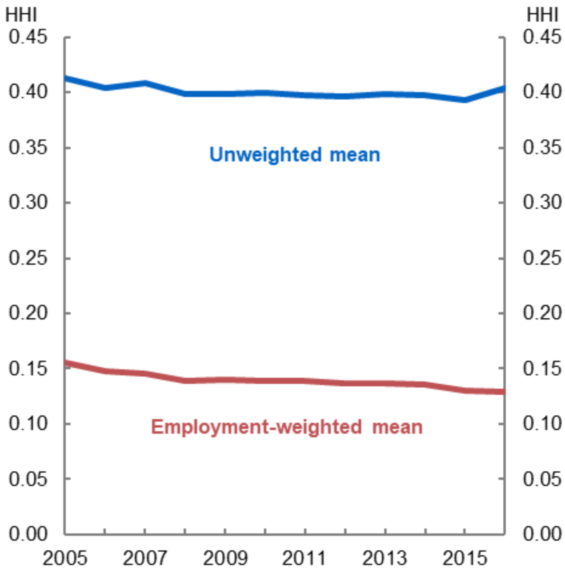

Figure A5: Average Employment HHI

ANZSIC 2-digit industry based markets

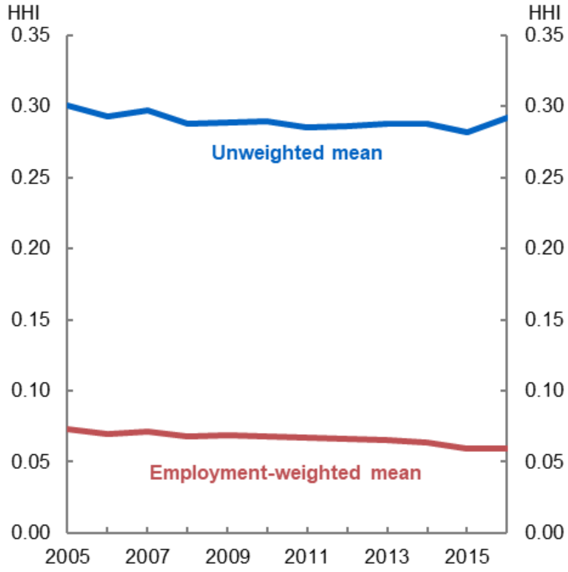

Figure A6: Average Employment HHI

SA4/GCCSA based markets

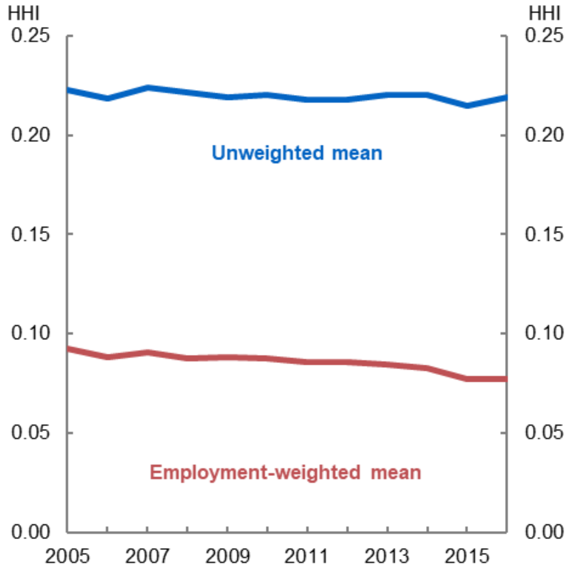

Figure A7: Average Employment HHI

ANZSIC 2-digit industry by ANZSCO 2-digit occupations based markets

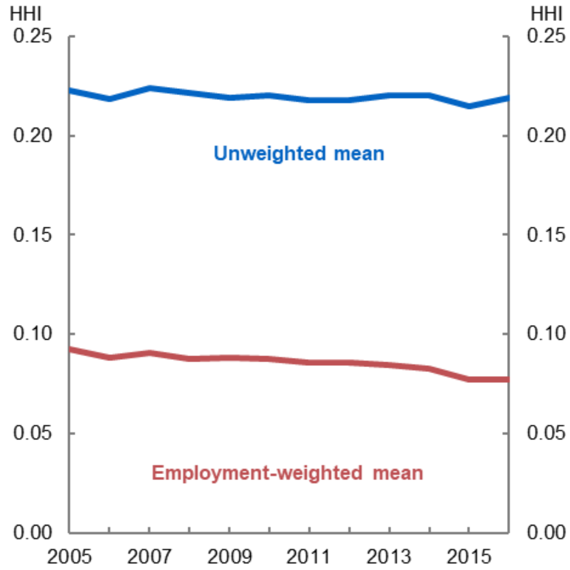

Appendix B: Regression Robustness Tests

Table B1: Robustness – Local Labour Market Models

Fixed effects, headcount

	Headcount			Wages		
	None	Local	Time trends	None	Local	Time trends
Concentration (HHI)	-0.162***	-0.153***	-0.122***	-0.238***	-0.359***	-0.311***
	(0.012)	-0.015	(0.014)	(0.012)	(0.015)	-0.025
Productivity	0.175***	0.087***	0.077***	0.174***	0.085***	0.075***
	(0.002)	(0.002)	(0.002)	(0.004)	(0.002)	(0.002)
R^2	0.107	0.631	0.671	0.110	0.633	0.672
N	306,000	302,000	302,000	306,000	302,000	302,000
Fixed effects						
Local market	N	Y	Y	N	Y	Y
Time*Industry	N	N	Y	N	N	Y
Time*Location	N	N	Y	N	N	Y

Note: All regressions include a set of common controls covering market characteristics (number of workers in the market). Errors clusters at the local market level.

Table B2: Robustness – Rent-sharing Models

Fixed effects, headcount

	None	Local	Ind area time	Plant	Local time	All
Productivity	0.371***	0.318***	0.316***	0.164***	0.324***	0.170***
	(0.005)	(0.05)	(0.05)	(0.004)	(0.05)	(0.05)
Local Concentration (HHI)* Productivity	-0.207***	-0.233***	-0.239***	-0.122***	-0.286***	-0.174***
	(0.016)	(0.014)	(0.015)	(0.011)	(0.017)	(0.017)
Firm Concentration (HHI)* Productivity	-0.121***	-0.201***	-0.203***	-0.008	-0.322***	-0.322***
	(0.010)	(0.008)	(0.008)	(0.007)	(0.010)	(0.010)
R^2	0.211	0.326	0.344	0.781	0.367	0.797
N	3,970,000	3,970,000	3,970,000	3,660,000	3,860,000	3,560,000
Fixed effects						
Local market	N	Y	Y	Y	Y	Y
Time*Industry	N	N	Y	Y	N	Y
Time*Location	N	N	Y	Y	N	Y
Time*Market	N	N	N	N	Y	Y
Plant	N	N	N	Y	N	Y

Note: All regressions include a set of common controls covering market characteristics (number of workers in the market). Errors clusters at the local market level.

Table B3: Robustness – Local Labour Market Models

Coverage, headcount

	Baseline	No small markets (<10 plants)	Lagged HHI	No rural	No primary industries
Concentration (HHI)	-0.122***	-0.135***	-0.087***	-0.131***	-0.144***
	(0.014)	-0.028	(0.013)	(0.015)	(0.016)
Productivity	0.077***	0.031***	0.077***	0.076***	0.073***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
R^2	0.670	0.856	0.67	0.670	0.655
N	302,000	117,000	302,000	277,000	261,000
Fixed effects					
Local market	Y	Y	Y	Y	Y
Time*Industry	Y	Y	Y	Y	Y
Time*Location	Y	Y	Y	Y	Y

Note: All regressions include a set of common controls covering market characteristics (number of workers in the market). Errors clusters at the local market level.

Table B4: Robustness – Local Labour Market Models
Concentration metrics, headcount

	Baseline	Top 4 share	Top 10 share	Log HHI
Concentration (HHI)	-0.122***	-0.122***	0.014	-0.146***
	(0.014)	(0.018)	(0.026)	(0.005)
Productivity	0.077***	0.078***	0.077***	0.077***
	(0.002)	(0.002)	(0.002)	(0.002)
R^2	0.670	0.670	0.670	0.670
N	302,000	302,000	302,000	277,000
Fixed effects				
Local market	Y	Y	Y	Y
Time*Industry	Y	Y	Y	Y
Time*Location	Y	Y	Y	Y

Note: All regressions include a set of common controls covering market characteristics (number of workers in the market. Errors clusters at the local market level.

Table B5: Robustness – Local Labour Market Models
Market definitions, headcount

Industry/occupation	ANZSIC 3-digit	ANZSIC 2-digit	ANZSIC 3-digit	ANZSIC 2-digt and ANZSCO 2-digit
Location	Working Zone	Working Zone	SA4/Greater capital area	Working Zone
Concentration (HHI)	-0.122***	-0.115***	-0.073***	-0.810***
	(0.014)	(0.020)	(0.033)	(0.025)
Productivity	0.077***	0.064***	0.049***	0.062***
	(0.002)	(0.002)	(0.003)	(0.002)
R^2	0.670	0.677	0.779	0.672
N	302,000	165,000	91,000	1,740,000
Fixed effects				
Local market	Y	Y	Y	Y
Time*Industry	Y	Y	Y	Y
Time*Location	Y	Y	Y	Y

Note: All regressions include a set of common controls covering market characteristics (number of workers in the market. Errors clusters at the local market level.

Table B6: Robustness – Wage models

Baseline models

	Local Labour Market models		Plant-level model		
	Labour market	Labour market Over time	Plant monopsony	Plant oligopsony	Plant rent-sharing
Market Concentration (HHI)	-0.311***	-0.246***	-0.279***		
	(0.008)	(0.023)	(0.011)		
Market Concentration GFC (HHI)		-0.085***			
		(0.021)			
Market Concentration post GFC (HHI)		-0.094***			
		(0.021)			
Local Concentration – (plant share)				3.300***	
				(0.032)	
Local Concentration (plant share)* Productivity					-0.117***
					(0.015)
Firm Concentration (plant share) * Productivity					-0.259***
					(0.01)
R^2	0.672	0.67	0.781	0.801	0.802
N	302,000	302,000	3,660,000	3,560,000	3,560,000
Fixed effects					
Local market	Y	Y	Y	N	N
Time*Industry	Y	N	Y	N	N
Time*Location	Y	Y	Y	N	N
Time*market	N	N	N	Y	Y
Plant	N	N	Y	Y	Y

Note: All regressions include a set of common controls covering market characteristics (number of workers in the market) Plant-level model contain control for number of plants in parent firm. Errors clusters at the local market level.

Table B7: Robustness – Wage Models

Mechanism models

	Entry models			Union models (2005–2013)	
	Baseline	Plant-based entry rate	Firm-based entry rate	Baseline	Union
Concentration	-0.122***	-0.110***	-0.111***	-0.315***	-0.322***
	(0.008)	(0.016)	(0.015)	(0.016)	(0.016)
Concentration * Entry Rate		0.540***	0.800***		
		(0.204)	(0.228)		
Concentration * Union coverage rate					0.010**
					(0.003)
R^2	0.67	0.683	0.686	0.985	0.701
N	302,000	274,000	274,000	247,000	217,000
Fixed effects					
Local market	Y	Y	Y	Y	Y
Time*Industry	Y	Y	Y	Y	Y
Time*Location	Y	Y	Y	Y	Y

Note: All regressions include a set of common controls covering market characteristics (number of workers in the market) Plant-level model contain control for number of plants in parent firm. Errors clusters at the local market level.

Table B8: Robustness – Single and Multi-plant Models

Employment concentration

	Baseline	Single plants only	Firm-based entry rate
Productivity	0.158***	0.214***	0.078***
	(0.004)	(0.005)	(0.001)
Firm share	-0.103***	-0.300***	-0.148***
	(0.027)	(0.108)	(0.030)
R^2	0.801	0.774	0.824
Observations	3,560,000	1,680,000	1,740,000

Note: All regressions include a set of common controls covering firm characteristics (number of plants in the firm) and market characteristics (number of workers in the market), as well as plant fixed effect and time*market fixed effects. Errors clustered at the local market level.

Appendix C: Additional regression frameworks

C.1 Rent-sharing models

A number of recent models have also argued that the degree of market concentration can affect the pass-through of firm-level productivity shocks onto wages. In oligopolistic models such as by Berger et al (2022) and Benmelech et al (2022), increases in productivity, which make the firm larger, move it onto a steeper part of the labour supply curve. As such the markdown increases, leading wages to increase by less than productivity, meaning the pass-through to wages is incomplete. Moreover, the degree of pass-through tends to decline as firms become larger, and as such pass-through becomes a function of firms' market shares. Such a specification also naturally comes out of some search-and-match based models of the labour market (e.g. Jarosch et al 2024).

As such, I also consider a rent-sharing specification as robustness:

- **Rent sharing**: The pass-through of productivity is a function of the labour market power, and therefore concentration. This could be plant-level or market-level concentration:

$$\ln(w_{i,t}) = X_{i,t} + \gamma * \ln(Prod_{i,t}) + \beta * Concen_{i,t} * \ln(Prod_{i,t}) + \epsilon_{i,t}$$

$X_{i,t}$ contains similar controls to the main models. But in this case I can also include local market by time fixed effects, given the focus is on the interaction between concentration and productivity.

Table C shows the results from the rent-sharing model. In all cases, pass-through from productivity to wages declines when concentration rises. Of particular interest, Column 3 shows the results when accounting both for market- and firm-level concentration. It shows that in more concentrated markets pass-through tends to be lower. But even within these markets there is variation, with larger firms having even lower pass-through.

Table C1: Rent-sharing regression

Employment concentration	HHI	Firm Share	HHI and Firm Share
Productivity	0.170***	0.161***	0.170***
	(0.005)	(0.004)	(0.004)
Productivity* Local Concentration (HHI)	-0.194***		-0.174***
	(0.016)		(0.017)
Productivity* Plant Concentration (plant share)		-0.163***	-0.074***
		(0.011)	(0.01)
R^2	0.797	0.796	0.797
Observations	3,560,000	3,560,000	3,560,000

Note: All regressions include a set of common controls covering firm characteristics (number of plants in the firm) and market characteristics (number of workers in the market), as well as plant fixed effect and time*market fixed effects. Errors clustered at the local market level.

C.2 Exploring declining occupational mobility

One explanation for an increased affect of concentration could be declining occupational or industry mobility. Much of the work on labour market concentration effectively assumes that workers cannot move between local labour markets due to geographic or skill barriers. However, in practice workers can move between geographies, industries or occupations, even if it is costly or difficult. If it becomes easier to move between these markets, and so the set of potential outside options becomes larger, the effect of concentration within the market might become smaller.

For example, in Berger et al (2021)'s model they allow for differentiation in supply of labour across, as well as within industries. When differentiation across industries declines, consistent with easier mobility across industries, the impact of differentiation and concentration within industries is lessened.

Schubert et al (2021) focus on this issue empirically. They derive a model where workers can potentially move to other industries in their locality. They show that in this case wages will be impacted by the 'outside options' that other occupations in their locality represent. Failing to include these outside options can bias the estimated impact of the HHI on wages. Similarly, they show that the impact of labour market concentration on wages tends to be smaller in industries with more cross-industry mobility, and the outside options become more important.

Declining occupational mobility and increasing frictions in moving between industries and occupations may be relevant in the Australian case. For example, Figure C1 shows that a moderate portion of the decline in job-to-job switching reflects fewer people changing occupations. This might suggest increasing frictions between labour markets. Moreover, failing to account for these dynamics could potentially impact the regressions exploring the how entry rates influence the effects of concentration.

Figure C1: Job-to-job Switching Rate

By type, over time

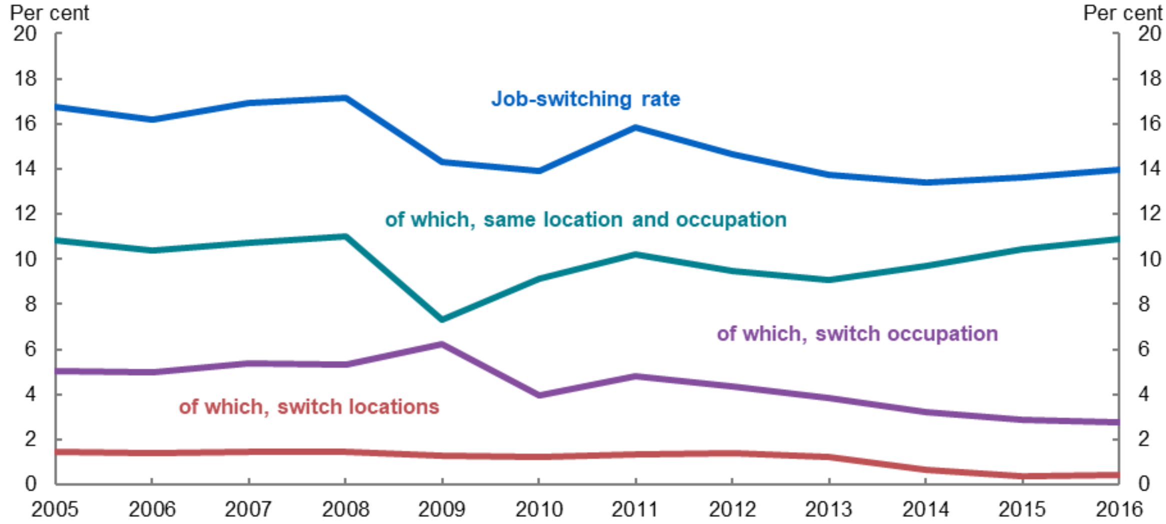

Note: Job-to-job switching as defined in Deutscher (2019). Location defined as SA4 or greater capital city area. Occupation is 2-digit ANZSCO.

To consider this, I include the measure of outside options suggested by Schubert et al (2021) and examine the impact on the coefficient on the HHI in the full sample, and in the sub-sample regressions.

The outside-option variable is constructed as a weighted average of the wages for firms in all other industries in the local area. The weights are a function of 2 factors: the share of employment for each of these industries in the local area (excluding the base industry), to capture how 'important' each industry is in the local area; and the national rate of switching between the base and the other industry (measured as a share of job-switchers only), to capture the ease with which people can move from the base industry to the alternative industry.[74]

The results of the regression are contained in Table C2. Focusing on Column 2, we can see that having strong outside options leads to higher wages. Moreover, the effect of concentration within a market is lower in this regression, suggesting that failing to account for these options can cause us to overstate the effects of the HHI.

[74] Note that due to the inclusion of location by time fixed effects in the regression, we are identifying whether wages increase more in industries where wages in closely linked industries rose relatively more. This makes it difficult to use this set-up to understand macroeconomic implications, rather than simply examining the mechanism. For example, taken at face value, if the outside option becomes less important, this could help workers with poor outside options.

Table C2: Outside Options Models

Employment concentration

	Baseline	Outside Options	Outside Options – Sub-sample
Concentration	-0.122***	-0.104***	-0.047**
	(0.008)	(0.014)	(0.020)
Outside options		0.040***	0.050***
		(0.004)	(0.004)
Concentration*GFC			-0.076***
			(0.021)
Concentration*Post-GFC			-0.083***
			(0.021)
Outside options *GFC			-0.010***
			(0.004)
Outside options *Post-GFC			-0.020***
			(0.004)
R^2	0.670	0.671	0.671
Observations	302,000	302,000	302,000

Note: All regressions include a set of common controls covering firm characteristics (number of plants in the firm) and market characteristics (number of workers in the market), as well as plant fixed effect and time*market fixed effects. All controls other than firm FE are interacted with period dummies. Pre-GFC is 2005–2007. GFC is 2008–2010. Post GFC is 2011 to 2015. Errors clustered at the local market level.

Focusing on the sub-sample in Column 3, we see outside options becoming less impactful over time. This is consistent with what we might expect to see if it became harder to move between industries. This suggests that declining occupational mobility may have weighed on wages growth by limiting workers' ability to leverage options outside their industries in wage negotiations.

Still, this regression provides no evidence that changes in the importance of outside options and occupational mobility have increased the impact of concentration. The coefficient on the HHI declines by a similar amount over the sample as in the model without outside options, so changes in the importance of the omitted outside options variable was not driving the decline.

To test more directly the impact of declining occupational mobility on the role of concentration, I extend the baseline market-level monopsony model to incorporate industry switching rates. Specifically:

$$\ln(w_{i,t}) = \gamma * Prod_{i,t} + X_{i,t} + \beta * \ln(Concen_{i,t}) + \delta * \ln(Concen_{i,t}) * Switch_{ind,t} + \theta * \ln(OO_{i,t}) + \rho * \ln(OO_{i,t}) * Switch_{ind,t} + \epsilon_{i,t}$$

$Switch_{ind,t}$ is the national rate of industry switching in the relevant ANZSIC 3-digit industry. This is either expressed as a share of all job switchers (*leave share*), or as a share of all job switchers and stayers (*leave rate)* focusing on worker's main jobs.

I focus on the national, rather than local rate of switching, and demean switch rates by industry.

The coefficients of interest in this case are δ and ρ. If concentration amongst incumbents is less impactful in industries with more occupational switching, the coefficient δ would be positive. Similarly, if outside options are more important where switching is more prevalent ρ would be positive.

Table C3: Leave Share Models

Employment concentration

	OO model	OO model with leave share	OO model with leave rate
Concentration	-0.104***	-0.104***	-0.104***
	(0.014)	(0.014)	-0.011
Outside options	0.040***	0.039***	0.040***
	(0.004)	(0.004)	(0.004)
Concentration*Leave share		-0.288*	-0.076***
		(0.151)	(0.021)
Outside options *Leave share		0.086***	-0.083***
		(0.032)	(0.021)
Concentration*Leave rate			0.035
			(0.324)
Outside options *Leave rate			0.228***
			(0.066)
R^2	0.671	0.671	0.671
Observations	302,000	299,000	302,000

Note: All regressions include a set of common controls covering firm characteristics (number of plants in the firm) and market characteristics (number of workers in the market), as well as plant fixed effect and time*market fixed effects. All controls other than firm FE are interacted with period dummies. Pre-GFC is 2005–2007. GFC is 2008–2010. Post GFC is 2011 to 2015. Errors clustered at the local market level.

Using both the leave share and rate, we see no evidence that concentration is less impactful where there is more industry switching (Table). However, there is some evidence that outside options are more important in industries with higher switching rates.

Taken together, these results provide some evidence that decreasing occupation and industry mobility may have weighed on wages pre-COVID by limiting the role of outside options. However, the direct evidence with respect to its role in increasing the impact of concentration is limited. As such, it suggests that failing to account for these dynamics is not impacting the main entry rate regressions.

C.3 Declining Union Membership

Benmelech et al (2022) argue that declining union membership rates could account for the increased impact of concentration on wages in the US. They argue that unions provide an offset to monopsony power. With lower union membership there is less of an offset, and so concentration could be more impactful. Failing to account for these dynamics may be important, particularly if there is some relationship between entry rates and union membership that biases my main results.

From a timing standpoint it seems somewhat unlikely that declining union coverage can explain the 'sudden' weakness in wages growth in the 2010s, given union coverage has been declining for decades. And Bishop and Chan (2019) found no evidence that declining union membership led to lower wage outcomes in Enterprise Bargaining agreements in Australia, as the share of agreements with union involvement remained unchanged.

Nevertheless, it is still worth examining the interaction between union coverage, labour market concentration and wages, to provide a fuller assessment of the interaction between union coverage and firm market power. To do so, I run a similar specification to the above, and to that used in Benmelech et al (2022):

$$w_{i,t} = \gamma * Prod_{i,t} + X_{i,t} + \beta * Concen_{i,t} + \delta * Concen_{i,t} * Union\ share_{div,t} + \epsilon_{i,t}$$

$Union\ share_{div,t}$ is the share of employees in each ANZSIC division that are part of a union. This is taken from the ABS Employee Earning, Benefit and Trade Union Membership release for 2006–2013. Unfortunately, more granular industry-level data were not available. Again, rates were demeaned to focus on changes in union membership, rather than the levels.

Table C4 shows the results. There is a significant positive coefficient on δ, indicating that the impact of market concentration on wages is smaller (in absolute terms) when union membership rates are higher. This does suggest that declining unionisation rates may explain part of the increased impact of concentration on wages. However, the evidence is somewhat weaker compared to the entry rate regressions, and after accounting for entry rates union coverage no longer plays a statistically significant role in influencing the impact of concentration (Column 3). This suggest that declining entry rates are a more important driver of the increased effect of concentration.

Table C4: Union Share Models

Employment concentration

	Baseline	Union	Union and firm-based entry rate
Concentration	-0.133***	-0.144***	-0.127***
	(0.016)	(0.016)	(0.018)
Concentration*Union Rate		0.009**	0.005
		(0.004)	(0.004)
Concentration*Entry Rate		0.939***	0.939***
		(0.254)	(0.254)
R^2	0.67	0.70	0.70
Observations	247,000	219,000	219,000

Note: All regressions include a set of common controls covering firm characteristics (number of plants in the firm) and market characteristics (number of workers in the market), as well as plant fixed effect and time*market fixed effects. All controls other than firm FE are interacted with period dummies. Regressions for shorter sample of 2006-2013. Errors clustered at the local market level.

One related explanation not explored in detail due to data constraints is changes in the share of workers on different pay-setting mechanisms. To the extent that pay is set nationally for a firm or industry, local concentration may be less impactful. This does not necessarily mean that wages will be higher, but the relationship between local concentration and wages may be weaker.

Unfortunately, data on pay-setting shares by ANZSIC division are only available for three years across my sample: 2008, 2012 and 2014. Based on this limited data, there is cross-sectional evidence that the HHI is more impactful in divisions with higher use of individual agreements. However, it is possible that this is driven by some common factors affecting both, such as level of education required for staff. Focusing within sectors, there is no evidence that the HHI has become relatively more impactful where the individual agreements usage increased.

With additional data this channel could be assessed more directly. However, given wages in individual agreements must be at least as high as award wages, and given that individual agreement usage increased up to 2008, but then fell, increased use of individual awards seems unlikely to be behind the results in this paper

References

Abel W, S Tenreyro and G Thwaites (2018), 'Monopsony in the UK', Centre for Economic Policy Research Discussion Paper DP13265.

Azar J, I Marinescu, M Steinbaum and B Taska (2020), 'Concentration in US Labor Markets: Evidence from Online Vacancy Data', *Labour Economics*, 66, Article 101886.

Azar J and I Marinescu (2024), 'Monopsony Power in the Labour Market: From Theory to Policy', *Annual Review of Economics,* 16, pp491-518.

Bassanini A, C Batut and E Caroli, 'Labor Market Concentration and Wages: Incumbents Versus New Hires', *Labour Economics*, 81, April, pp 1-13.

Benmelech E, NK Bergman and H Kim (2022), 'Strong Employers and Weak Employees: How Does Employer Concentration Affect Wages?', *The Journal of Human Resources*, 57(S), pp S200–S250.

Berger D, K Herkenhoff and S Mongey (2021), 'Labor Market Power', *American Economic Review*, 112(4), pp 1147-93.

Bilal A, N Engbom, S Mongey and GL Violante (2019), 'Firm and Worker Dynamics in a Frictional Labor Market', *Econometrica*, 90(4), pp 1425-1462.

BITRE (Bureau of Infrastructure, Transport and Regional Economics) (2018), 'Working Zones 2016', BITRE Information Sheet 98.

Deutscher N (2019), 'Job-to-Job Transitions and the Wages of Australian Workers', Treasury Working Paper No 2019-07.

Hambur J (2023), 'Product Market Competition and Its Implications for the Australian Economy', *Economic Record*, 99(324), pp 32-57.

Handwerker EW and M Dey (2022), 'Some Facts about Concentrated Labor Markets in the United States', U.S. Bureau of Labor Statistics BLS Working Paper No 550.

Hathorne C and R Breunig (2022), 'Occupational Mobility in the ALife Data: How Reliable are Occupational Patterns from Administrative Australian Tax Records?', *Economic Papers*, 41(4), pp 297-324.

Jarosch G, JS Nimczik and I Sorkin (2024), 'Granular Search, Market Structure, and Wages', *The Review of Economic Studies*, 91(6), pp 3569-3607..

Manning A (2021), 'Monopsony in Labor Markets: A Review', *ILR Review*, 74(1), pp 3–26.

Marinescu I, I Ouss and L-D Pape (2021), 'Wages, Hires, and Labor Market Concentration', *Journal of Economic Behavior & Organization*, 184, pp 506–605.

Moscarini G and F Postel-Vinay (2017), 'The Relative Power of Employment-to-Employment Reallocation and Unemployment Exits in Predicting Wage Growth', *The American Economic Review*, 107(5), pp 364–368.

OECD (Organisation for Economic Co-operation and Development) (2020), *Competition in Labour Markets*, Background Note, OECD, Paris. Available at <https://www.oecd.org/daf/competition/competition-concerns-in-labour-markets.htm>.

OECD (2021), *The Role of Firms in Wage Inequality: Policy Lessons from a Large Scale Cross-Country Study*, OECD Publishing Press, Paris.

Rinz K (2018), 'Labor Market Concentration, Earnings Inequality, and Earnings Mobility', U.S. Census Bureau Center for Administrative Records Research and Applications, CARRA Working Paper 2018-10.

Schubert G, A Stansbury and B Taska (2024), 'Employer Concentration and Outside Options', Available at SSRN: https://ssrn.com/abstract=3599454 or http://dx.doi.org/10.2139/ssrn.3599454.

U.S. Department of Justice and Federal Trade Commission (2010), 'Horizontal Merger Guidelines'. Available at <https://www.justice.gov/atr/merger-enforcement>.

Win NN, J Hambur, D Hansell and N Crighton (2025), 'Tracking Mergers and Acquisition in Australian Administrative Data', Australian Treasury Competition Review Working Paper, September 2025.

5. Conclusion of the thesis

This thesis has demonstrated the power of administrative and other integrated microdata sources in better understanding important microeconomic dynamics, and their implications for the macroeconomy and macroeconomic policy in Australia. As such, it highlights the value of continuing to invest in such data, as well as the capabilities to leverage it, to improve economic forecasting and policymaking, and continue to build our understanding of the macroeconomy.

Through the first two chapters the thesis provided important insights into the past and future drivers of productivity in Australia, and how economic policy can influence productivity growth. These insights are crucial given productivity growth is the key driver of long-term living standards and has been slowing over recent decades.

Chapter 1 highlighted how macroeconomic conditions and monetary policy shocks can have medium-run effects on productivity, and therefore activity, in a small open technology-importing country like Australia. That said the effects appear more muted, compared to those in a country at the technology frontier like the US, potentially as domestic conditions do little to influence the global productivity frontier. So, domestic cyclical factors are likely to have more limited medium-term implications for productivity in a small open economy like Australia.

Specifically, while contractionary monetary policy shocks are associated with a decline in R&D spending in Australia, and declines in R&D spending are associated with lower levels of productivity, the effects appear to be shorter-lived compared to the US. Moreover, the effects of monetary policy on broader measures of innovation that better capture adoption of existing technologies (a crucial channel in aggregate productivity growth) are more complex and heterogeneous. SMEs are less likely to innovate and adopt after a contractionary monetary policy shock, but large firms are more likely to do so. This appears to reflect differing exposures to the channels through which monetary policy can affect innovation, particularly the demand channel and the financial constraint channel, both of which appear important. Finally, US monetary policy is also an important driver of Australian firms' innovation and adoption, likely reflecting a combination of the aforementioned channels, as well as its impacts on the global frontier.

Chapter 2 took a more forward-looking lens on productivity by examining the effects of, and barriers and enablers to, adoption of emerging transformative digital technologies: Cloud Computing and AI/Machine Learning. Applying event study approaches, it found that the average firm does not increase profitability post-adoption. But firms whose Board have relevant skills do. They are also the firms most likely to hire more skilled worked around adoption. So, this work highlights the important role of skills, particularly management skills, in making the most of these technologies going forward. This has natural implications for skills-related policies, suggesting they should focus both on the use of these technologies, but also on helping managers support implementation and integration into firm processes.

Pulling the insights of these two chapters together, they suggest that cyclical dynamics and policies have some implications for productivity growth. But that structural policies are likely to be more important, such as those related to building a skilled workforce.

The final two chapters use microdata to consider the extent to which certain firm dynamics that are abstracted away from in macroeconomic models can explain macroeconomic outcomes in Australia. As such, they provide insights into where policymakers and forecasters may need to invest to improve their frameworks and models they use.

Chapter 3 focused on inflation, and whether changes in the frequency of firm-level price changes can account for the pick-up in inflation post-COVID. The chapter shows that firms began resetting prices more frequently in response to the large inflationary shocks post-COVID, which is counter to standard macroeconomic models that assume constant price-setting frequencies. Incorporating these dynamics into a workhorse policy DSGE, we showed the failure of standard models to account for these changes can potentially account for a significant portion of the underprediction of inflation in Australia in the post-COVID period.

The key policy insight is that, when faced with a large inflationary shock, standard frameworks used by policymakers may break down. As such they either need to invest in alternative frameworks or apply judgement to their forecasts and policy actions during such times (though doing so in real-time is inherently challenging). In either case, the investments made in Australian firm-level prices microdata as part of this project will be an important facilitator.

Chapter 4 explores the labour market, and whether increases in monopsony power caused slow wages growth in Australia over the 2010s. While labour markets did not become more concentrated, the effect of given levels of concentration became larger. The increasing impact of concentration subtracted a bit under 1 per cent from wages on average from 2011–2015. A key driver of the increased impact of concentration appears to have been the decline in firm entry: fewer young firms were entering and trying to grow by poaching workers from incumbents, and this reinforced the effects of concentrated labour markets. This highlights the important links between firm dynamics and competition, and worker bargaining power, which are absent from most economic models. It also highlights how policies that improve economic and firm dynamism can be pro-worker, as well as supporting aggregate productivity growth.

Taken together, these two chapters highlight how standard economic models often abstract from dynamics that are important in understanding macroeconomic outcomes over the short- and medium-run. While no one model or framework can capture all these important dynamics, gathering evidence from microdata provides can help us to understand which ones are likely to be important, and when.

Data Disclaimer

The results of these studies are based, in part, on data supplied to the ABS under the Taxation Administration Act 1953, A New Tax System (Australian Business Number) Act 1999, Australian Border Force Act 2015, Social Security (Administration) Act 1999, A New Tax System (Family Assistance) (Administration) Act 1999, Paid Parental Leave Act 2010 and/or the Student Assistance Act 1973. Such data may only used for the purpose of administering the Census and Statistics Act 1905 or performance of functions of the ABS as set out in section 6 of the Australian Bureau of Statistics Act 1975. No individual information collected under the Census and Statistics Act 1905 is provided back to custodians for administrative or regulatory purposes. Any discussion of data limitations or weaknesses is in the context of using the data for statistical purposes and is not related to the ability of the data to support the Australian Taxation Office, Australian Business Register, Department of Social Services and/or Department of Home Affairs' core operational requirements.

Legislative requirements to ensure privacy and secrecy of these data have been followed. For access to PLIDA and/or BLADE data under Section 16A of the ABS Act 1975 or enabled by section 15 of the Census and Statistics (Information Release and Access) Determination 2018, source data are de-identified and so data about specific individuals has not been viewed in conducting this analysis. In accordance with the Census and Statistics Act 1905, results have been treated where necessary to ensure that they are not likely to enable identification of a particular person or organisation.