Wage Growth in Australia: Lessons from Longitudinal Microdata

Dan Andrews, Nathan Deutscher, Jonathan Hambur and David Hansell[1,2]

2019-04

Date created: July 2019

Wage Growth in Australia: Lessons from Longitudinal Microdata
Dan Andrews, Nathan Deutscher, Jonathan Hambur and David Hansell
2019-04
July 2019

ABSTRACT

This paper uses novel microdata sources spanning 2001-02 to 2015-16 to explore the structural drivers of wage growth in Australia, with a view to better understanding recent weak wage growth – a phenomenon observed across a range of advanced economies. Controlling for a range of cyclical and other factors, we show that part of firms' idiosyncratic productivity growth tends to be passed-through to workers in the form of higher wage growth, consistent with the idea that firms share rents with their workers. The size of this pass-through is around the midpoint of leading international estimates, and appears to decline modestly after 2012-13, when aggregate wage growth begins to slow. We then discuss a range of possible mechanisms for this modest decline, including the transition from the mining boom, globalisation, changing shock processes, declining labour market fluidity and uneven technology diffusion. Given that our dataset does not cover the past few years, however, it is not clear whether lower pass-through of productivity to wages has persisted. In this respect, our results represent a tentative first step in examining the microeconomic drivers of a macroeconomic phenomenon with particular reference to wages. We provide a proof-of-principle demonstration of the value of longitudinal microdata and empirical techniques to such investigations, and suggest further analysis into the factors behind declining labour market fluidity would be a fruitful exercise.

1. INTRODUCTION

Across a variety of measures, wage growth has been slow in Australia, and in many other advanced economies, over recent years (Treasury 2017; Arsov and Evans 2018). A range of analyses suggest that wage growth has been lower than can be explained by measures of labour market spare capacity, inflation expectations and output price inflation, and the usual lags between these variables and wage growth (e.g. Cassidy 2019). This has led to suggestions that deeper structural factors may be at play, including changes in competitive dynamics owing to globalisation and technological change (Lowe 2018). But systematic evidence on the contribution of these structural factors in Australia is scarce and there is growing recognition that the answers to these questions may lie at the micro rather than the macro level (Weir 2018).

Accordingly, this paper uses novel microdata sources to explore the structural sources of wage growth in Australia. More specifically, we are interested in understanding that modest portion of weakness in wage growth that cannot be explained by historical relationships with labour market slack, labour productivity and inflation. Conceptually, we follow Summers and Stansbury (2017) and note that recent wage growth could have been weaker than expected either because there has been a recent and potentially transient change in either the relationship between wages and productivity; or because the relationship remains the same but something else has weighed on wages. Since the latter explanation tends to encompass a number of cyclical aspects – such as mismeasurement of labour market slack – we focus our efforts on identifying the underlying structural link between wages and productivity. Empirically, credible identification depends on the ability to link workers to firms – through the use of matched firm-worker data – in order to control for unobserved worker and firm characteristics, and the increasing tendency of high-ability workers to sort into high productivity firms. We thus follow a recent literature that seeks to understand aggregate wage outcomes in terms of the important links that exist between the dispersion in firm-level productivity and worker-level wages within narrowly defined sectors (Card, Heining and Kline 2013; Bagger et al 2014).

We first document an important link between individuals' wages and firm-level productivity over the period 2001-02 to 2015-16, consistent with the idea that firms share rents with their workers. Such a link does not arise in neoclassical models of wages, where all firms pay the same wage, but is an empirical regularity in the international literature and is consistent with leading models of job search. Our estimates – which are the first for Australia – imply that an idiosyncratic shock to firm-level productivity of 10 per cent is associated with an increase in wages of around 1 per cent, controlling for a battery of time-varying cyclical shocks and assortative matching between high-ability workers and high productivity firms. This is around the midpoint of the range of leading international estimates, and highlights the important link between firm productivity and wages, consistent with recent analysis using firm level data (Treasury 2017).

We also find evidence that this pass-through may have declined modestly in the period after 2012-13, when aggregate wage growth slows. Drawing aggregate implications from these structural micro-level relationships is complex, reflecting the many assumptions required to extrapolate with any confidence from estimates from the micro -level to the aggregate economy. Indeed, this reflects the more general tension that can at times exist between the internal and external validity of results in empirical work (Roe and Just 2009).

A back-of-the-envelope estimate nonetheless provides useful context, suggesting that growth in the aggregate wage price index could have been around 0.15 percentage points higher annually after 2012-13, had an aggregate shift in pass-through – equivalent to the same magnitude as estimated from micro-data over this period – not occurred. This corresponds to around one-third of the unexplained weakness in wages growth observed in recent years. The change in the relationship between workers' wages and

firm-level productivity thus provides a new mechanism to understand the meaningful but modest portion of the weakness in aggregate wage growth over recent years that cannot be explained by historical relationships. Given that our dataset for this part of the analysis concludes in 2014-15, however, it is not clear whether lower pass-through of productivity to wages has persisted.

We then indirectly explore and discuss some potential mechanisms for this apparent modest, and potentially transient, shift in the relationship between wages and firm-level productivity. We start by discussing some Australia-specific explanations, such as the notion that there might be a real wage 'overhang' following the end of the mining boom, or that the shock processes affecting firms might have changed. Reflecting the fact that wage growth has been unexpectedly low in a number of advanced economies over the past few years, we then go on to explore in more detail two explanations that are more 'global' in nature.

The first is the decline in labour market fluidity, which has been documented for the United States (US) and which we document for Australia. To the extent that it implies fewer outside options for workers, lower labour market fluidity may have reduced workers' confidence and power in negotiations, and thus lowered the scope for rent sharing. Equally, lower labour market fluidity could reflect decreased feelings of job security amongst workers due to globalisation and technological advancement (Lowe 2018). Consistent with both of these, we discuss recent evidence which shows that higher job switching rates are associated with higher wage growth at the local labour market level in Australia, even after controlling for a range of cyclical and demographic factors. Moreover, even workers who remain with their incumbent employer appear to benefit from more fluid labour markets, which provides further evidence that the decline in job switching rates – to the extent it is structural – could be related to the shift in rent-sharing.

Second, we explore the conjecture that given the rising productivity gap between frontier and laggard firms observed internationally (Andrews, Criscuolo and Gal 2015, 2016), laggard firms may try to cut costs to compete, leading their wages to be less responsive to productivity (Weir 2018). While this explanation may be relevant for a number of OECD countries, it does not appear to fit the Australian data. This is not to say that technology does not influence wages and productivity. Instead, the two necessary ingredients for the technology diffusion argument – namely noticeable increases firm-level productivity dispersion, and growing wage inequality driven by differences in average wages paid between firms – are less evident in Australia.

Finally, even if the shift in the relationship between wages and firm-level productivity had not occurred, aggregate wage growth may still have been affected by changing patterns of within-industry labour reallocation. Indeed, the period of low aggregate wage growth has coincided with a slowdown in the extent to which labour is reallocated from less productive to more productive firms. Given high productivity firms pay higher wages, slowing reallocation implies that fewer high paying jobs are being created than otherwise. We consider this compositional effect separate from the relationship between wages and productivity, as it implies both lower aggregate productivity and wage growth.

Overall, what emerges is not a single unified explanation of slower aggregate wage growth, but rather a pattern of evidence that lends credence to the idea that structural factors may have weighed on wage growth in Australia up until at least 2014-15. On the one hand, it is consistent with the idea that the response of wages to improving labour market conditions has been slower and more muted than in past cycles (MYEFO, 2018). This suggests a need to continue to look beyond these old relationships when considering the outlook for wages growth. On the other hand, to the extent that we have not explicitly identified the drivers of the reduced pass-through of productivity to wages, it is not clear whether this effect will persist into the medium-term and may well turn out to be a transient phenomenon. A key priority for future work will be to understand these drivers, their origins and potential responses.

The remainder of the paper is structured as follows. The next section reviews what we know about low wage growth in Australia and presents some new evidence illustrating the broad-based nature of the wage slowdown at the firm and worker level. Section 3 then examines the empirical link between individual wages and firm-level labour productivity, while Section 4 indirectly explores some potential mechanisms behind the observed changes in the relationship between wages and productivity. Section 5 examines the implications of changing patterns of labour reallocation for aggregate wages and then Section 6 discusses some implications for future research.

2. LOW WAGE GROWTH IN AUSTRALIA

Wage growth has been low in recent years, relative to history. Moreover, it has generally been somewhat lower than would have been predicted based on its historical relationship with labour market slack, labour productivity, inflation expectations, and other key theoretical determinants (Treasury 2017; Bishop and Cassidy 2017; Arsov and Evans 2018). For example, the lower rate of inflation, or inflation expectations, in recent years would be expected to lower nominal wages growth one-for-one. But wages growth has been lower than expected even after accounting for this, and other changes in relevant determinants. In this regard, the RBA's Phillips curve model has over-estimated actual wage growth (based on the WPI) by around ½ percentage points per annum over past 5 years (Figure 1; Cassidy 2019) – and it is this unexplained portion of wages weakness that we are particularly interested in.

Figure 1: Contribution to Wage Growth – RBA Model
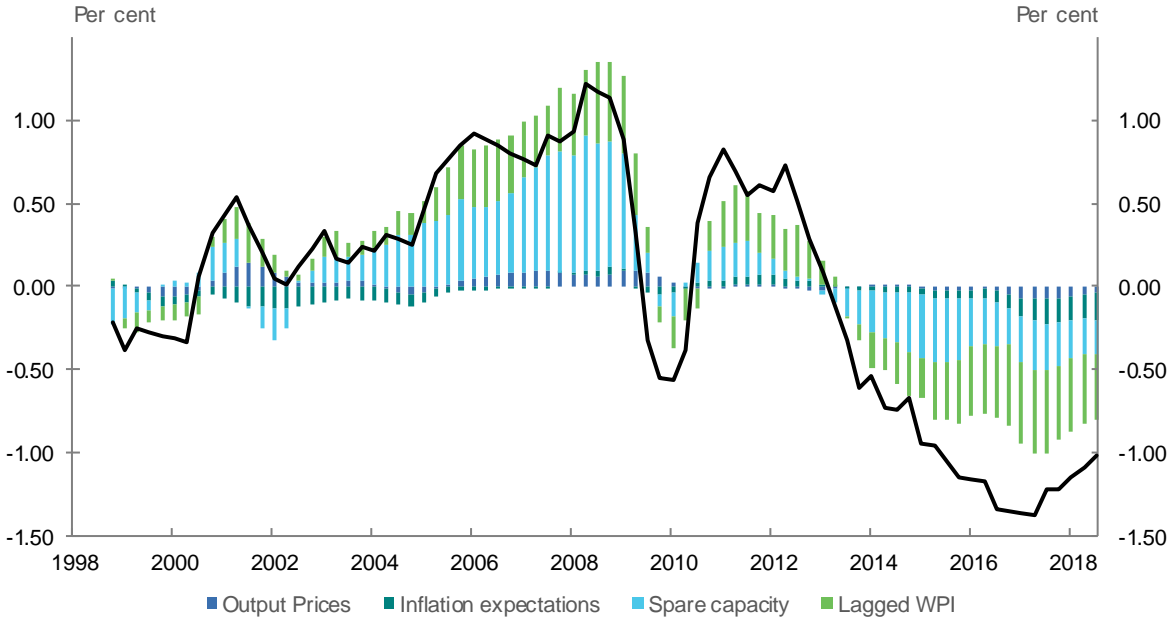
Notes: Wage growth measure is year-ended private WPI growth; all series shown as year-ended deviations from own mean.
Sources: ABS; RBA

Relatedly, there has also been a marked divergence between wages and other cyclical measures that tend to co-move closely with wages, such as measures of capacity constraints (Figure 2 Panel A). Perhaps unsurprisingly then, actual wage growth has persistently been below that forecast by the RBA since 2011 (Figure 2, Panel B).[3]

---

3   This graph is only illustrative, as some portion of the forecast miss will also reflect errors in the forecasts for the variables that help to determine wage growth.

A number of papers have tried to explain recent low wage growth. Most of these analyses have been done using economy-wide macro-data. While such analyses are useful, they necessarily have some shortcomings. As the period we are considering is relatively short, it can be difficult to identify any structural change, and also to differentiate between different explanations for the change. Moreover, economy-wide data are not particularly useful if we are trying to identify the effects of compositional change.

Figure 2: The Slowdown in Aggregate Wage Growth in Australia

Panel A: Constraints on Output and Wage Growth   Panel B: RBA Wage Price Index Forecasts
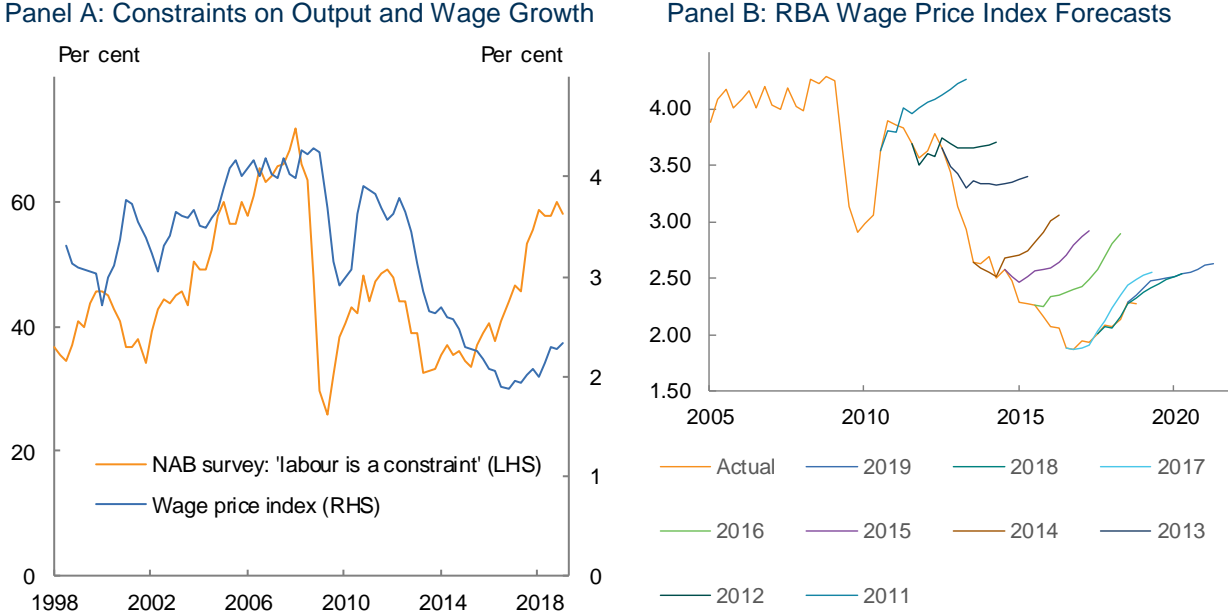
Notes: In left panel, wage growth measure is year-ended growth in WPI. NAB survey measure shows share of firms reporting availability of suitable labour as a minor or significant constraint on output. In right panel, forecasts are from February Statement of Monetary Policy.
Sources: ABS; NAB; RBA

Digging deeper, recent analysis using microdata suggests that the slowdown in wage growth is relatively broad based and has been experienced across income, education, age and occupation categories (Treasury 2017). Moreover, the slowdown in wage growth is apparent even when following trends in wage growth for individual firms (Figure 3), with average wage growth within firms around 2 percentage points lower in 2015-16, relative to 2010-11. We also see a similar decline when we track wage growth for individual workers, or fixed firm-worker relationships – that is, just looking at variation in the wages of a worker that remains at one firm (Appendix B, Figure B1). This set of results rules out a number of compositional explanations. For example, the slowdown is not driven by the changing nature of workers (for example, skills or age) or firms (for example, industry), at least not entirely.

Figure 3: Broad-based slowdown in wages
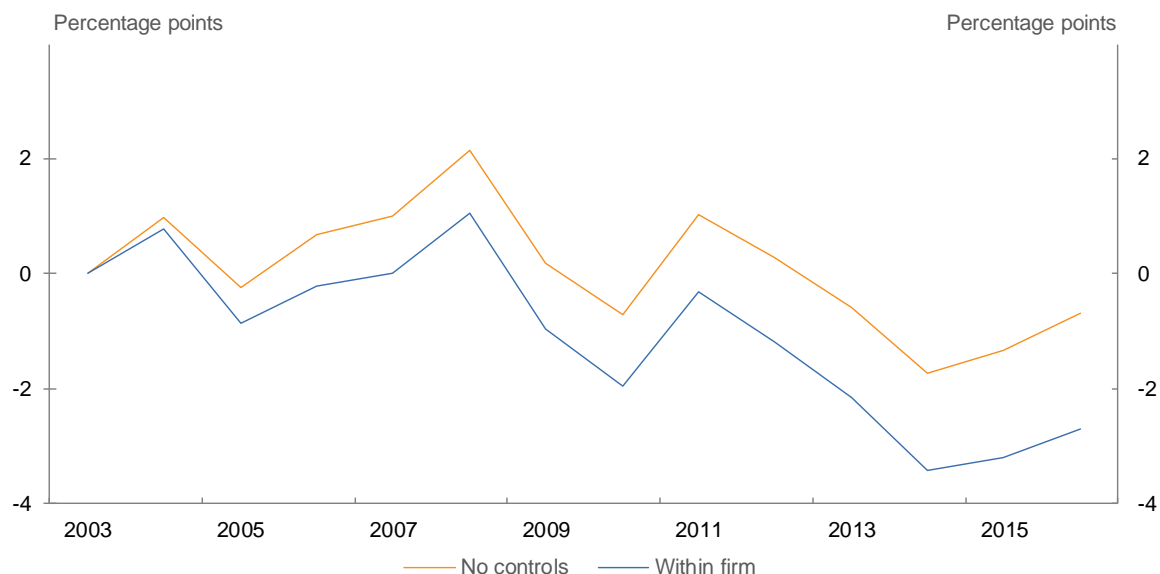
Note: Presents coefficients from a regression of change in log annual wages on financial year dummies and either: no controls; or firm fixed effects. Values represent deviation of nominal wages growth from 2003 level.
Source: Treasury analysis of de-identified tax data.

Another advantage of microdata is that they can make it easier to abstract from cyclical factors – or even certain forms of mismeasurement, for example in inflation expectations or labour market slack – and to identify structural factors. This is because we can include time-varying controls at the state and industry level that allow us to abstract from these aggregate influences.

These points highlight the potential benefits of using microdata to try to better understand the recent unexplained weakness in wage growth. There have been relatively few such analyses to date, and these have tended to focus on firm-level data. But firm-level data is only likely to get us so far. Numerous papers have documented the role of worker heterogeneity in explaining heterogeneous wage outcomes across firms (for example, Song et al 2019; Card et al 2016). Failing to account for worker differences can lead to incorrect inference, reflecting the fact that highly skilled workers will tend to raise firm productivity, and given high productivity workers may tend to congregate at certain firms.

As such, this paper makes use of de-identified matched firm-worker panel data that allows us to account for worker heterogeneity when trying to shed light on the factors affecting aggregate wage growth in Australia. This is the first paper to use such a database in Australia.

3. MICRO EVIDENCE ON WAGES AND PRODUCTIVITY

One useful framework for thinking about the slowdown in wage growth is that proposed in Summers and Stansbury (2017). Broadly, we can think of two potential sets of explanations for the fact that wage growth has been weaker than expected: there has been a shift in the relationship between productivity and wage growth; or, the wages-productivity relationship has remained unchanged but something else has been weighing on wage growth.

Many of the more structural explanations posited to date can broadly be interpreted as a change in the relationship between wage and productivity growth. For example, as suggested by Foster and Guttmann (2018), if workers feel less secure in their jobs this is likely to lower their wage bargaining power and mean that more of the gains from increases in productivity and profits are likely to flow to employers instead of workers. This would lead to a weakening of the relationship between wages and productivity.

Increasing integration of local and global labour markets might have similar effects (Lowe 2017). These contrast to some of the more cyclical explanations, such as mismeasured labour market slack, that will weigh on wage growth, but will not necessarily affect the relationship between wages and productivity, particularly at the micro level.

As we are interested in identifying whether there are structural explanations for recent weak wage growth, we explore whether there is evidence of a shift in the relationship between wage and firm productivity growth in recent years.

3.1 Wages, productivity and rent sharing

To consider whether the link between productivity and wages has changed in the past few years, we draw on the rent-sharing literature. In the presence of labour market frictions, any job-worker match will generate some economic surplus, or 'rent'. This rent will be split between the worker and the firm, and this split will be a function of the worker's bargaining power.

Within these models there is a link between productivity and wages, as worker productivity affects the amount of rent available to split. In particular, wages usually take on the reduced-form:

$$log(wage_i) = \gamma log(LP_i) + (1 - \gamma)b_i$$

Where LP_i is the productivity of the worker and b_i is a reservation wage, which captures the outside option of the worker. The parameter γ captures the worker's degree of bargaining power.

As such, to consider the relationship between firm productivity and worker wages, we estimate the following equation:

$$log(wage_{ijt}) = \propto + \beta_1 log(LP_{jt}) + \delta X' + \varepsilon_{ijt}$$

The vector X' has a number of controls, including various fixed effects. It includes a number of time fixed effects that account for time-varying state- and industry-specific shocks. This implies that our estimates give information on the pass-through of shocks to idiosyncratic firm productivity onto wages – that is, they net out the effects of average changes in these variables at the state and industry level. This will ensure that our results are not unduly affected by cyclical factors, such as conditions and wages in the rest of the industry, which are likely to be a key determinant of the reservation wage. One potential cost, however, is that it is more difficult to infer aggregate implications (see below). The regression also includes a variety of firm and worker fixed effects, which are included to control for time-invariant factors, such as a worker's work ethic or a firm's inherent innovativeness.

A large number of papers have estimated equations of this form to examine whether there is a link between firm- or industry-level productivity and wages. Or, more precisely, they try to examine whether shocks to firm- or industry-level productivity flow through to wages, while trying to avoid any endogenous changes in both wages and productivity due to, for example, demand shocks. These papers are well summarized in Card et al (2016) and Manning (2011). Estimates from more recent papers, which exploit both firm-level productivity data and worker-level wage data, generally find estimates of γ ranging from 0.05 to 0.15.

While the relationship between wages and productivity is of interest, the motivation of this paper is to examine whether this relationship has changed. To do so, we also estimate the following equation:

$$\Delta log(wage_{ijt}) = \propto + \beta_1 \Delta log(LP_{jt}) + \beta_2 [\Delta log(LP_{jt}) * D_{12/13}^{Post}] + \delta X' + \varepsilon_{ijt}$$

Where $D_{12/13}^{Post}$ is a dummy that takes on the value one for 2012-13 onwards, and zero otherwise. Therefore, our coefficient of interest is β_2, which identifies whether there has been a change in the relationship between firm-level productivity and wages over the recent period where unexplained weakness in aggregate wage growth emerges. We choose 2012-13 to coincide with period of unexplained weakness in wages growth per the RBA's Philips curve model, though some basic testing suggests our results are generally robust to the exact choice of year.

We prefer a differences specification when trying to identify a change in the relationship, rather than the levels specification that is more common in the literature and that we use for identifying rent sharing. This is because the differences specification is likely to be more suited to capturing a change in the relationship between wages and productivity. For example, if the change reflects a shift in worker bargaining power, we would expect it to affect the sharing of future rents, but not necessarily past rents, making the differences specification more appropriate. More generally, if we are transitioning to a new and different relationship between the levels of wages and productivity, the transition is likely to take a number of years, and can be better accommodated by the differences specification. The differences specification can also help to control for variation in worker quality (Card et al 2016), though this is not overly important for our application as we have access to employee-level data.[4]

At the macro level, the relationship between wages and productivity will reflect not only the idiosyncratic firm productivity shocks that we capture, but also common productivity shocks at the national, state and industry levels that are absorbed through our fixed effects structure. Despite these apparent costs, we include fixed effects since our aim to credibly identify the causal relationship between productivity and wages, and any change therein. To the extent that this aids causal identification, if we find any evidence of a change in the relationship between firm-level productivity and wages, then it will likely be suggestive of a change in the aggregate relationship. This will be the case even if the responsiveness of wages to common and idiosyncratic shocks differs somewhat (as has been demonstrated in Card et al (2016)).

Our measure of wages is the gross wages received by the worker in their main job (that is, the job that pays them the most). As demonstrated in Appendix A, this measure tracks official aggregate wage measures (such as the WPI and AENA) reasonably well in trend terms, although some non-trivial differences arise due to definitional issues. Our measure of labour productivity is the firm's real value added, divided by the number of employees at the firm during the year. We construct the measure of real value added by deflating nominal value added by an industry-specific value-added deflator.[5] We only have data on value-added up to 2014-15, so this limits the sample for this part of the analysis.

It would be preferable to measure both wages and productivity on an hours basis, as this might be more robust to changes along the intensive hours margin (for example, Bishop, Gustafsson and Plumb 2016), which could be particularly important given the growth of the 'gig economy'. But, so long as both are measured on the same basis, any such adjustment is unlikely to substantially affect our results.

3.2 Is there a relationship between wages and firm productivity in Australia?

Table 1 outlines the results from five specifications with different fixed effects structures. In all specifications there is a positive and significant relationship between firm productivity and worker wages over the period 2001/02-2014/15. The estimates are around 0.1, which implies that an unexpected 10 per cent increase in firm productivity leads to a 1 per cent increase in worker wages. This is around

---

4   To show that our results do not reflect some inherent difference between the level and difference specifications, we estimate the basic model with no interaction in differences for comparison. The results are very similar to the levels specification. See Appendix B.

5   More details are available in Appendix A.

the midpoint of the leading estimates from the international literature. For context, a firm that has productivity one standard deviation above the industry mean will have wages that are about 10 per cent higher.[6] The fact that we find a relationship between firms' productivity and wages is consistent with Treasury (2017), which finds that more productive firms tend to pay higher wages.

Table 1: The link between wages and firm productivity
Dependent variable: Log wages of the individual

	(1)	(2)	(3)	(4)	(5)
Productivity	0.244***	0.096***	0.111***	0.077***	0.075***
	(0.004)	(0.002)	(0.002)	(0.002)	(0.001)
Common controls	X	X	X	X	X
Fixed effects					
Firm		X		X	
Worker			X	X	
Firm#Worker					X
R^2	0.311	0.413	0.739	0.757	0.791
Adjusted R^2	0.311	0.403	0.673	0.689	0.721
N	27,325,184	27,325,184	27,325,184	27,325,184	27,325,184

Note: All regressions include a set of common controls covering worker characteristics (sex, a quadratic in age and residential location (Statistical Area 4)), firm characteristics (size and industry division), and state- and division-wide time-varying factors (state and industry division separately interacted with financial year). Errors are clustered at the firm-level.

It is worth noting that, while the pass-through of firm productivity to wages seems low, in a neoclassical model it will be zero, as all firms pay the same wage. Moreover, we are only capturing the short-run pass-through and any nominal rigidity in wages, as might occur if individuals sign multi-year contracts, will necessarily lower the measured pass-through. Such rigidities could also be important if firms expect some of the productivity growth to be retraced (that is, firm-specific productivity does not have a unit root). In fact, Card et al (2016) finds some evidence that the pass-through of industry-level productivity to wages is stronger than the pass-through of firm-level productivity. So the micro-level productivity-wage nexus can coexist with the longer-run one-to-one association present in macro studies.

As we move from Columns 1 to 5, we incorporate increasingly demanding micro fixed effects structures and the explanatory power of the models increase accordingly. In particular, the inclusion of worker fixed effects in Column 3 leads to a particularly large increase in explanatory power, which is consistent with numerous papers that find that variation in worker quality explains a large share of variation in wages (for example, Song et al (2019); Card et al (2016)). This highlights the importance of human capital and skills in the determination of wages.

At the same time, the coefficient on productivity tends to decline when we include worker fixed effects, compared to the case with no fixed effects. This suggests that failing to account for worker effects leads to an upward bias in the coefficient, consistent with Dobbleaere and Mairesse (2018). The bias reflects the fact that firms that employ higher productivity workers are likely to have both higher wages and productivity. Indeed, the inclusion of interacted worker-firm fixed effects in Column 5 allows us to control for the assortative matching of high-ability workers to high productivity firms. Given our aim is to isolate the pass-through of exogenous productivity 'shocks' to wages, it is crucial to abstract from these compositional effects. This highlights the importance of having linked employer-employee data.

---

6   The standard deviation of labour productivity within an industry is approximately 1 log unit.

3.3 Has the relationship between wages and firm productivity changed?

Table 2 shows the results from the regression that looks for changes in the relationship between firm productivity and wages. Again, as we move from Column 1 to 5 we incorporate an increasingly burdensome fixed effects structure. As the model is in differences, firm and worker fixed effects are now allowing for differential trends in wages.[7]

Table 2: The changing in the link between wage growth and firm productivity growth
Dependent variable: Change in log wages of the individual

	(1)	(2)	(3)	(4)	(5)
ΔProductivity	0.121***	0.118***	0.105***	0.100***	0.097***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)
ΔProductivity#post13	-0.016***	-0.015***	-0.031***	-0.027***	-0.024***
	(0.004)	(0.004)	(0.005)	(0.005)	(0.004)
Common controls	X	X	X	X	X
Fixed effects					
Firm		X		X	
Worker			X	X	
Firm#Worker					X
R^2	0.036	0.063	0.264	0.283	0.325
Adjusted R^2	0.036	0.041	0.013	0.008	0.046
N	27,325,184	27,325,184	27,325,184	27,325,184	27,325,184

Note: All regressions include a set of common controls covering worker characteristics (sex, a quadratic in age and residential location (Statistical Area 4)), firm characteristics (size and industry division), and state- and division-wide time-varying factors (state and industry division separately interacted with financial year). All are interacted with the post13 dummy. Errors are clustered at the firm-level.

Turning now to our variable of interest (that is, productivity#post13), the negative coefficient on the interaction term suggests that there has been some weakening in the relationship between productivity and wages in recent years. Our estimates suggest that the relationship weakened by 10-20 per cent, depending on the specification.

Overall, the results indicate that the relationship between wages and productivity has declined modestly in recent years. While the results specifically relate to the relationship between idiosyncratic firm productivity and wages (as discussed above), they are suggestive of a broader weakening in the relationship. This lends credence to the view that structural factors may have weighed on wage growth in recent years. We explore this in more detail in the next section. Nevertheless, given our dataset concludes in 2014-15 it is not clear whether lower pass-through of productivity to wages has persisted.

To try to quantify these effects, we run a simple counterfactual. We assume that prior to 2012-13 there was a (long-run) one-to-one relationship between wages and productivity at the aggregate level.[8] We also assume that the 10 per cent weakening in the relationship evident at the micro level can be

---

[7]   Freestone (2018) finds no systematic heterogeneity in wage growth, which would suggest that these fixed effects are unnecessary. Nevertheless, some of these micro-level fixed effects appear to improve the model fit, particularly the inclusion of firm fixed effects. This is even more evident in the non-interacted differences specification in Appendix B.

[8]   This assumption is consistent with the fact that labour's share of income has been broadly flat for a number of decades.

generalized to the macro level.[9] Our counterfactual then asks, what would wage growth have looked like if the relationship had not weakened, given labour productivity growth of around 1½ per cent?

The red line in Figure 4 plots this counterfactual wage growth series. [10] In this case, aggregate nominal wage growth might have been about 0.15 percentage points higher in each year since 2012-13, eliminating around 1/3 of the unexplained weakness in nominal wage growth. The results are obviously only illustrative, but they do suggest that structural forces could potentially account for a non-trivial share of the unexplained weakness observed in recent years.

Figure 4: A modest drag on aggregate wage growth
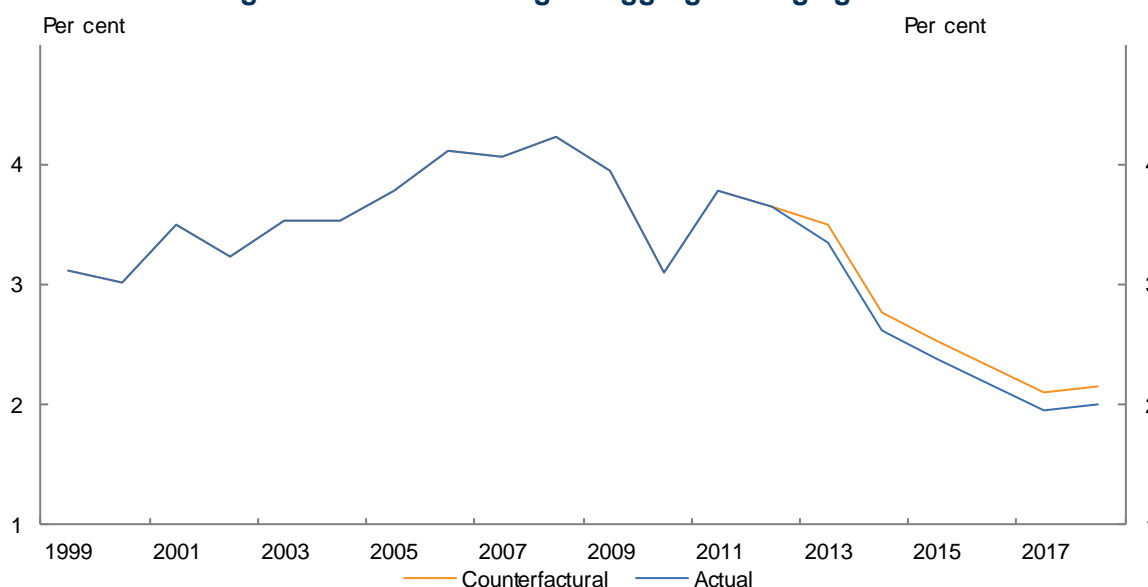
Notes: Wage growth measure is year-ended growth in WPI. Counterfactual assumes a one-to-one historical relationship between labour productivity and wages, and that the relationship has weakened by 10 per cent. Source: ABS; Authors' calculations

4. POTENTIAL EXPLANATIONS FOR THE CHANGING RELATIONSHIP

There are numerous potential explanations for a recent, potentially transient, shift in the relationship between productivity and wages, and for unexpectedly lower wage growth more generally.[11] Some of these are more Australia specific and we briefly discuss a few of these in Section 4.1.

Nevertheless, as discussed in Arsov and Evans (2018), wage growth has been unexpectedly low in a number of advanced economies over the past few years. This would suggest that these phenomena are unlikely to be explained solely by domestic factors.

As such, we then go on to explore some explanations that are more global in nature. One explanation is declining labour market fluidity, which has been documented in the United States. Declines in fluidity could reflect increasing distortions and adjustment costs, and so could be the cause of lower wage growth and the change in the relationship between firm productivity and wages we observe in the data. But it

---

9   This represents a lower bound based on our estimates and thus reflects a conservative stance on the effect of micro-level decoupling between wages and productivity.

10   WPI is quality adjusted, so changes in productivity may be less likely to affect this wages measure. However, this graph is just intended as illustrative.

11   We return to some other potential explanations in our discussion of future work in Section 6.

could also be symptomatic of factors such as technological progress and globalisation that make workers feel less secure and therefore less willing to change jobs.

Section 4.2 then explores that the role of the growing productivity gap between the productivity of 'leaders' and 'laggards' (Andrews, Criscuolo and Gal 2015, 2016) in explaining the change in the relationship between firm productivity and wages, and low wage growth (see Weir 2018).[12]

4.1 Australia-specific explanations

A number of potential explanations for any change in the relationship between firm-level productivity and wages relate specifically to Australia. One prominent example is that wages grew faster than productivity during the mining boom, and now need to grow more slowly than productivity for a period.[13]

Assessing this real wage overhang explanation is complicated by the fact that the relevant wage for consumers and firms differs. Consumers care about how much they can buy with their wage, and so the relevant deflator is the Consumer Price Index (CPI). The nominal wage deflated by the CPI is often referred to as the real consumer wage. In contrast, firms care about how much their labor inputs cost relative to what they can sell their goods for, so the relevant deflator is the Producer Price Index (PPI). The nominal wage deflated by the PPI is often referred to as the real producer wage. As documented by Treasury (2017), the real consumer wage grew faster than productivity during the mining boom, while the real producer wage tracked labour productivity fairly closely.

If the weakness in wages that has been observed in recent years does reflect some form of real wage overhang, it could potentially account for the decline in rent-sharing. In particular, if firms were handing out an 'abnormally' large share of productivity gains to workers during the boom, and are now trying to 'claw' this back by giving out a lower share, this could account for the decrease in the rent sharing coefficient. That said, we might expect any such change to be more evident when looking at the relationship between the aggregate or industry-wide component of productivity and wages, rather than between the firm-specific component and wages.

Another potential explanation for the observed decline in rent-sharing that specifically relates to our study is a change in the economy's shock process. In order to isolate firm-specific productivity shocks, and to abstract from demand or other shocks that could affect firm's measured productivity and wages, many rent sharing studies instrument their productivity variables. We choose not to do so for two reasons. First, we do not have access to any instruments that we feel are valid.[14] Second, to the extent that firm-specific demand shocks raise firm's profits, for example by raising the prices of their goods, we might still be interested in how they are passed on to wages.

One concern, though, is that the pass-through of increased profits from productivity and demand shocks could differ if, for example, the persistence of the two shocks differs. Firm-specific productivity shocks might be more persistent than firm-specific demand shocks, and as such the pass-through of these shocks could be higher. This would be consistent with van Reenen (1996), who finds a much larger estimate of rent sharing when he instruments using a measure of patents. If this is the case, our estimated coefficient will be an average of the two pass-throughs, weighted by the relative importance of the two shocks.

---

12   A related, econometric, explanation could be that a growing productivity divergence might introduce an omitted variable bias, given we only have firm-level productivity in our regression. The bias is likely to be negative, as workers with higher wage growth have more understated productivity growth.

13   See, for example, the Q&A to Phil Lowe's Address to the AFR Business Summit on 6 March 2019, as reported in Kehoe (2019).

14   Many papers use lags of the right hand side variables, or measures of industry profit and productivity. Such instruments are only valid under certain, fairly restrictive assumptions.

Moreover, if the relative importance of the two shocks has changed in recent years, this could lead to a change in the estimated coefficient.

While this is possible, there is little evidence of such a change in the relative importance of firm-specific demand and productivity shocks. For one, if productivity shocks had become less important, we might expect to see a broad-based decline in the dispersion of productivity amongst firms in recent years. As discussed in Section 4.3, there is little evidence of such a broad-based decline. That said, Campbell et al (2019) do find some evidence of a decline for certain industries, so this could be worth exploring in future work.

Similar arguments could be made regarding efficiency wages: firms paying higher wages to promote greater productivity. This has the potential to lead to reverse causality, and therefore bias our estimates of the relationship between firm idiosyncratic productivity and worker wages. However, it is unlikely to affect our estimates of the change in this relationship, unless the propensity to use efficiency wages has changed.

4.2 Declining labour market fluidity

A number of papers have documented the fact that labour market fluidity has declined in the US (for example, Davis and Haltiwanger, 2014). Moreover, a nascent literature from the United States suggests that lower job switching rates are associated with lower wage growth (Faberman and Justiniano 2015; Karahan et al 2017; Moscarini and Postel-Vinay 2017). This association has been found in a variety of settings, and is robust to different definitions of local labour markets – such as demarcated by geography (states) or demography (sex, race, age and skills) – and to using both aggregate- and individual-level wages. The job switching rate emerges as a more important correlate of wages than the unemployment to employment transition rate, a common cyclical indicator of labour market strength.

A more fluid labour market can benefit workers' wages in two broad ways. First, it may lift labour productivity by facilitating moves from low to high productivity firms, as examined in Section 5, and ensuring workers' skills and firms' needs are well matched. Second, it may improve the bargaining position of workers. If a worker has more credible outside options in more fluid labour markets, then they will be better placed to negotiate higher wages with their employer, and therefore to earn a higher share of the rent from the match. As such, declines in labour market fluidity associated with, for example, increasing adjustment costs, could weigh on aggregate wages and help to explain the micro decoupling evident in the data.

At the same time though, lower labour market fluidity could also be a symptom of other, broader factors. A number of academics and policymakers have argued that factors such as increasingly globalised labour market and technological innovation could make workers feel less secure in their jobs (see: Lowe 2018; Foster and Guttmann 2018). For example, cross-country evidence shows that those sectors that experienced a larger increase in global value chain integration recorded lower wages growth, especially when low-wage countries are integrated in supply chains (Andrews, Gal and Witheridge 2018). Indeed, workers that feel less secure might be less willing to change job, particularly if they worry that, as the newest employees, they might be the first to be fired in any downturn. Moreover, to the extent that globalisation benefits the largest firms (for example, Autor et al 2017), and could therefore weigh on the entry and growth rates of new firms that drive much of the labour market turnover, it could directly affect labour market fluidity (Davis and Haltiwanger 2014).

Consistent with the US evidence, Australian labour markets appear to have become less fluid in recent years. This is demonstrated in Figure 5, which shows declines in excess job reallocation rates (Panel A) – job churn in excess of what is required to accommodate net job creation (see Davis and Haltiwanger (2014) for details) – and job-to-job switching rates (Panel B). Abstracting from cyclical fluctuations, both appear to have trended down over the sample.[15]

Figure 5: Labour market fluidity has declined

Panel A: Excess job reallocation rates   Panel B: Job switching rate
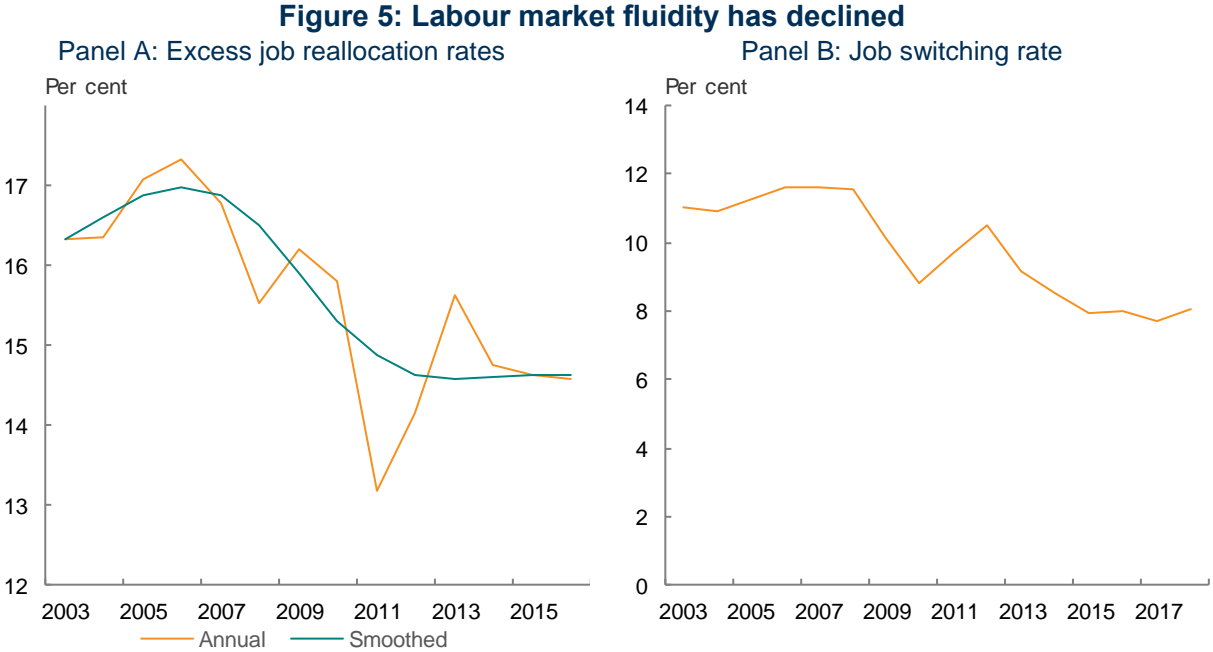
Notes: In Panel A, excess job reallocation is calculated as (JC+JD)/|(JC-JD)| where JC refers to job creation rates and JD refers to job destruction rates. The smoothed series utilises STATA's robust nonlinear smoothing command "smooth" (option = 4235eh). The job switching rate is calculated as the number of workers that switch jobs, divided by the number that have jobs at the start and end of the year.
Source: Treasury analysis of de-identified tax data.

Could the decline in job switching rates be associated with less upward pressure on Australian wages?

Deutscher (2019) examines the effects of job switching on wages in Australia and finds that local labour markets that have higher job switching rates have higher wage growth. This result remains after controlling for demographic influences and a range of time-varying shocks, including the unemployment to employment transition rate. A one percentage point increase in the rate at which workers switch jobs is associated with a ½ percentage point increase in growth in average wages. While the association is stronger for those that switch jobs, who may be in the best position to negotiate wages, it remains for those that stay in their job. This provides evidence that it is not only the act of switching jobs that affects worker wages, but also the potential to switch jobs, as it provides workers with an outside option and therefore a stronger bargaining position. These associations also hold when estimated using employee-level data, which allows for a rich set of controls.

The association between job switching and wages is economically meaningful, with plausible variation in switching rates being associated with wage growth worth hundreds of dollars a year to the typical worker. As an illustration of this, Figure 6 shows the predicted effect of a 2 percentage point decrease in the job-switching rate – which roughly corresponds to the decline in the aggregate switching rate observed in the sample we analyse – for a worker on the mean wage.[16] For the full sample, the predicted effect is

---

15   Ideally we would extend the sample, as the early part of our sample coincides with the mining boom, when job switching rates may have been unusually high. This is something we hope to do in future work.

16   The sample we analyse differs from that in Figure 5; for example, it only includes those for whom we have location data.

a decline in wage growth of around 1 per cent (or $600) a year. This reflects a predicted effect of around -3 per cent (or $1600) a year for job switchers and around -0.6 per cent (or $300) a year for job stayers.

Figure 6: Predicted impact on wages of a 2 percentage point decline in the job-switching rate in a local labour market

Panel A: Dollar decrease (2016 dollars)	Panel B: Percentage point decrease
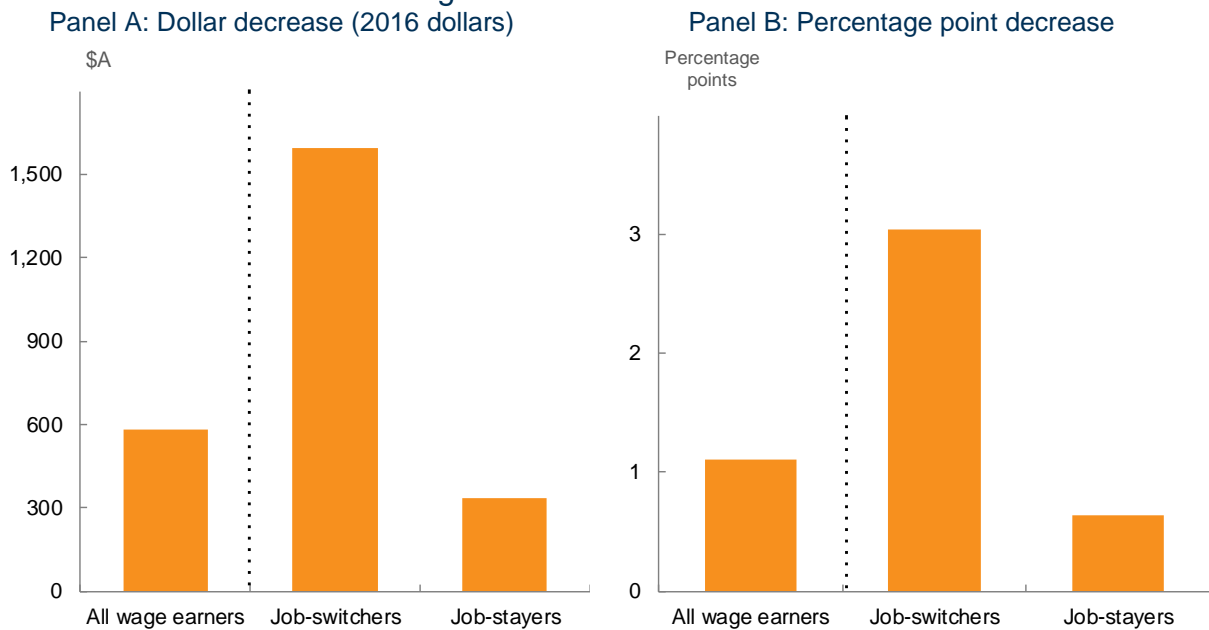
Notes: Based on coefficients from an OLS regression of the change in the natural logarithm of average wages in an SA4 in a given year on the job-to-job and unemployment-to-job transition rates and a set of control variables. The control variables include year and SA4 fixed effects, state interacted with year, and a set of SA4 time-varying controls encompassing a linear trend, the unemployment rate, average number of weeks searching for a job, share of workers in each of 19 one-digit industry classifications, and the average age and proportion female of the relevant sample. The regression is run on three samples: all wage earners; job-switchers; and job-stayers.
Source: Treasury analysis of de-identified tax data.

The results of this exercise suggest that labour market fluidity may be an important factor in the evolution of wages – benefitting not just those who switch but those who choose to stay in their job. Given the trend decline in labour switching rates has been coincident with the unexplained weakness in wages, it seems likely that the two are related. Moreover, to the extent that lower job switching rates imply fewer outside options for workers, or reflect declines in workers' sense of security more generally, it could help to explain the modest decline in rent-sharing.

As with the majority of this paper's analysis, these findings rely on rich underlying microdata following workers and firms over time. A logical next step would be to directly examine whether the decline in rent-sharing was more pronounced in local labour markets where job switching declined more. As more comprehensive and finely-grained data become available, further light will be able to be shone on the role of job switching on productivity, and on individual careers and wages. Such data will also help to identify the underlying causes of the decline in job switching rates.

4.3 Uneven technology diffusion

Another factor that could be contributing to weak wage growth in Australia is an increasingly uneven diffusion of technology that leads to a growing divergence between the productivity of 'leaders' and 'laggards'. A growing divergence in productivity outcomes is likely to be associated with a growing divergence between the wages paid by high and low productivity firms, and therefore increasing wage inequality, particularly if the technologies require highly skilled workers (for example, Caselli 1999).

As well as leading to greater wage inequality, these dynamics could also weigh on aggregate wage growth, as discussed in Weir (2018) and Lowe (2018). In order to compete less productive firms will need to try to cut costs, and might do so by paying lower wages. This would place downward pressure on aggregate wage growth, and would lead to a further divergence between the wages paid by high and low productivity firms, beyond what would be implied by the differing marginal productivities of labour at these firms. As this might make the wages of low productivity firms less responsive to productivity changes, it could also help to explain the decline in rent-sharing.

There is some tentative evidence for such dynamics overseas. For example, a number of papers have documented a growing gap between the productivity of leader and laggard firms, such as Andrews et al (2015, 2016) for a range of OECD countries, and Haldane (2018) for the United Kingdom. A number of papers have also documented that the dispersion of wages paid by firms has increased over time (for example, Card, Heining and Kilne 2013; Song et al 2019). Moreover, Berlinegieri, Blanchenay and Criscuolo (2017) find some evidence that the relationship between productivity and wages is weaker for low productivity firms, and that much of the increase in the dispersion of wages between firms reflects weakness in wage growth for firms at the bottom of the wages distribution. Both of these finding are consistent with low productivity firms having to lower wages to compete with their highly productive competitors.

However, the evidence for Australia is far less compelling. In particular, there is little evidence of rising productivity dispersion in general, or a growing divergence between the productivity of leaders and laggards in Australia. This is evident in Figure 7 Panel A, which is constructed by calculating the ratio of value-added labour productivity for firms at the top of the industry productivity distribution to all other firms in an industry, and then taking the median across industries.[17] These findings are consistent with the findings in Treasury (2017) and Campbell et al (2019) and are robust to different industry definitions, weighting and trimming rules, divergence measures, to focusing on the top 5 percentile instead of the top 10 percentile, and to using estimates of multi-factor productivity.

There is also little evidence of a growing divergence between wages paid by firms at the top and bottom of the productivity distribution. Despite some volatility, the ratio of wages paid by the most productive firms to those paid by other firms has been broadly stable over the sample (Figure 7, Panel B).[18]

---

17   Industries are defined at the 4-digit ANZSIC industry level. The analysis uses data from the BLADE database. Value-added productivity is calculated as the ratio of value-added to full time equivalent (FTE) employees. We add one to FTE employees to account for owner-managers who affect output, but not employment statistics.

18   The large spike in the mid-2000s appears to reflect a number of small firms. If we weight firms by their employment shares, which may lessen measurement error and may be more relevant for aggregate wage outcomes, the spike is not evident and the ratio is far less volatile.

Figure 7: Key components of technology diffusion-based explanations

Panel A: Ratio of productivity for most productive firms to all other firms has not risen

Panel B: Ratio of wages paid by most productive firms to all other firms has not risen
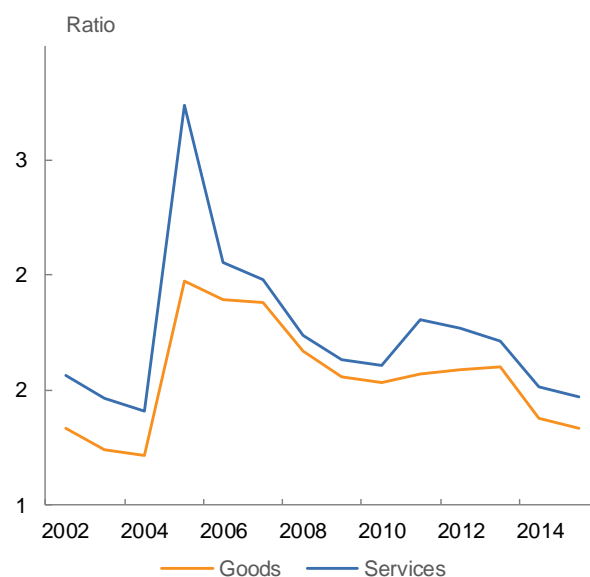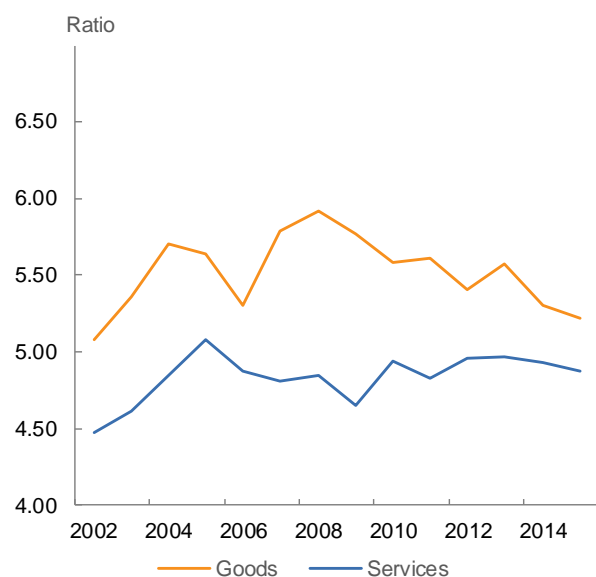
Notes: Ratios calculated as average productivity or wages for top 10 percentile of firms in a 4-digit industry, over average for rest of industry. Goods sector includes manufacturing, construction, wholesale trade, retail trade and transport; services sector includes accommodation, information media & technology, rental services & real estate, professional services and administrative services. In left panel, one added to all firm FTE to account for owner managers.
Source: ABS BLADE database

So far we have focused on firm-level data, but employee-level data may provide additional insights as we can examine individual wage outcomes, rather than the firms' average wages. Using our de-identified matched firm-worker panel data we follow Song et al (2019) and decompose the total variation in individuals' wages into two components: a between-firm component that captures variation in average wages across firms; and a within-firm component that captures variation in individuals' wages within firms (see Box 1). If uneven technology diffusion was important, we would expect to see a substantial increase in the between-firm component. While there has been some increase, it has been fairly small at only around 0.03 log points, and has only accounted for around half of the overall increase in wage dispersion (Figure 8). In contrast, in the US over the same sample the between component increased by around 0.1 log points, and accounted for the majority increase wage dispersion over the sample (Song et al 2019).

Figure 8: Decomposing wage dispersion
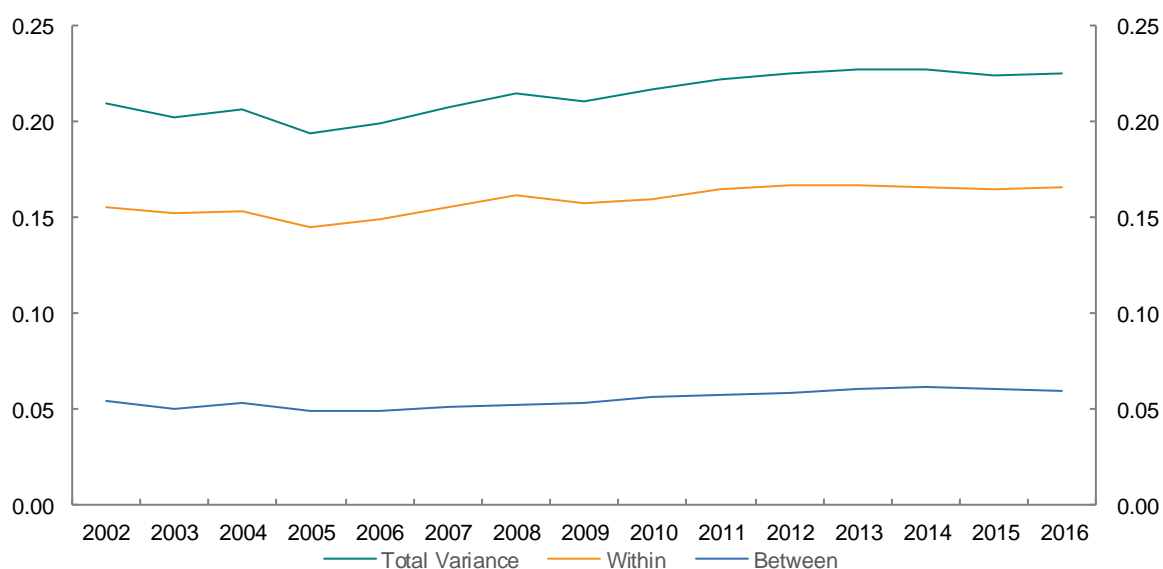
Notes: See Box 1
Source: Treasury analysis of de-identified tax data.

Box 1: Wage Variance Decomposition

In order to compare Australian and US trends in earnings equality, we adopt the approach of Song et al (2019) and decompose the cross-sectional variance of log earnings into within and between components. That is, we let $w_t^{i,j}$ be log wages of employee i at firm j in period t. Thus,

$$w_t^{i,j} \equiv \overline{w}_t^j + \left(w_t^{i,j} - \overline{w}_t^j\right)$$

Where $\overline{w}_t^j$ is the mean wage in firm j and so

$$var\left(w_t^{i,j}\right) = var\left(\overline{w}_t^j\right) + var\left(w_t^{i,j}\middle| i \in j\right)$$

The first term on the right-hand side is the *between-firm* dispersion as measured by the variance of mean firm wages. The second term on the right-hand side is the within-firm dispersion term as measured by the deviation of individual earnings from the within-firm mean.

In order to make our estimates comparable with those of Song et al (2098), we weight the within-term by each firm's employment share and limit our sample to full-time employees[19] in companies with a headcount of 20 or more.

---

5. PRODUCTIVITY, WAGES AND REALLOCATION

Throughout the paper we have noted the potential for lower market dynamism to weigh on both aggregate productivity and wage growth. For example, lower job switching rates could lead to less reallocation of workers from unproductive to productive firms, and therefore to slower productivity and

---

19   Our data only contain annual earnings. We proxy full-time workers as those earning above the annualised minimum weekly wage; for example, in 2010 this was $29,624 per year.

wage growth. A natural question then is, what evidence is there of decreasing market dynamism and efficient reallocation?

Declines in firm entry rates over recent years are suggestive of a decline in business dynamism and efficient reallocation. Andrews and Hansell (2019) look at this more explicitly, using the approach proposed by Decker, Haltiwanger, Jarmin and Miranda (2018). This approach exploits a prediction from canonical models of firm dynamics: conditional on initial size, more productive firms are likely to grow more quickly than less productive firms, and less productive firms are more likely to fail.

The approach entails estimating the responsiveness of employment growth to lagged labour productivity at the firm level, where the simplest form of the model is:

$$g_{i,t+1}^{E} = \beta_0 + \beta_1 LP_{it} + X_{it}'\theta + \varepsilon_{it+1}$$

Where future growth ($g_{i,t+1}^{E}$) is a function of current (labour) productivity (LP) of firm i at time t, expressed in relative terms as deviation from the industry average, and where the vector X' includes controls for firm size plus a range of cyclical controls.[20]

To identify how productivity-enhancing reallocation is evolving in a structural sense, the model controls for cyclical influences in three key ways. First, the firm-level labour productivity is expressed as a deviation from the industry average at the four-digit ANZSIC06 level, which is akin to controlling for time-varying industry-specific shocks. Second, the impact of common time-varying shocks (that is, the national business cycle) is absorbed by the inclusion of time fixed effects. Finally, the inclusion of the state unemployment rate directly controls for state-level business cycles and, through its interaction with the labour productivity term, sweeps out the impact of state-level shocks on the responsiveness of employment growth to firm productivity.

Turning to the results, the estimated coefficient on LP (that is, β_1) is positive and significant at the 1 per cent level, suggesting that on average over the sample period, more productive firms are more likely to expand and less productive firms are more likely to contract. This suggests that resources are flowing in the right direction over time and thus that reallocation provides a positive contribution to aggregate productivity growth. Importantly, evidence emerges to suggest that this reallocation process may have weakened around the same time that aggregate wage growth slowed. In this regard, the authors interact the (lagged) firm productivity term with:

- First, a linear time trend and its square, to account for non-linearity in the trend. This exercise returns a positive coefficient on LP*Trend coupled with a negative coefficient on LP*Trend[2]. To illustrate this effect, Figure 9 Panel A shows the difference in employment growth between a high productivity firm – that is, a firm one standard deviation above the industry mean – and a low productivity firm – that is, a firm one standard deviation below the industry mean. Over the first few years of the sample, the employment growth differential between a high and low productivity firm rises slightly but then declines from around 2010 to reach around 7 per cent by 2016.

- Second, a piecewise dummy that takes the value of 1 in the post-2011-12 period. This interaction term is negative and statistically significant and implies that the difference in employment growth

---

20   The model defines LP as real revenue per FTE in order to avoid indirectly censoring units with negative value-added but using real value-added per FTE measure does not change the story. While future analysis will look to exploit estimates of multi-factor productivity (MFP), the estimated responsiveness of firm employment growth to (lagged) productivity in Decker et al (2018) is insensitive to the choice of LP or MFP.

between a high and low productivity firm fell by 1½ percentage points in the low wage growth period (Figure 9, Panel B).[21]

Given higher productivity firms pay higher wages, the decline in the rate at which labour is reallocated to high productivity firms implies that there will be fewer high paying jobs than there would be otherwise. Put differently, a slowdown in the pace of labour reallocation might have resulted in a compositional effect that imposes a drag on the dynamics of aggregate wages and productivity growth. Finally, to the extent that the decline in the ability of the economy to reallocate labour to high productivity firms reflects a rise in adjustment frictions, this could partly account for why more firms are reporting it more difficult to find suitable workers (Figure 1, Panel A).

Figure 9: Reallocation of workers from low to high productivity firms has slowed
Difference in employment growth between a high and low productivity firm
Panel A: Quadratic time trends   Panel B: Piecewise specification
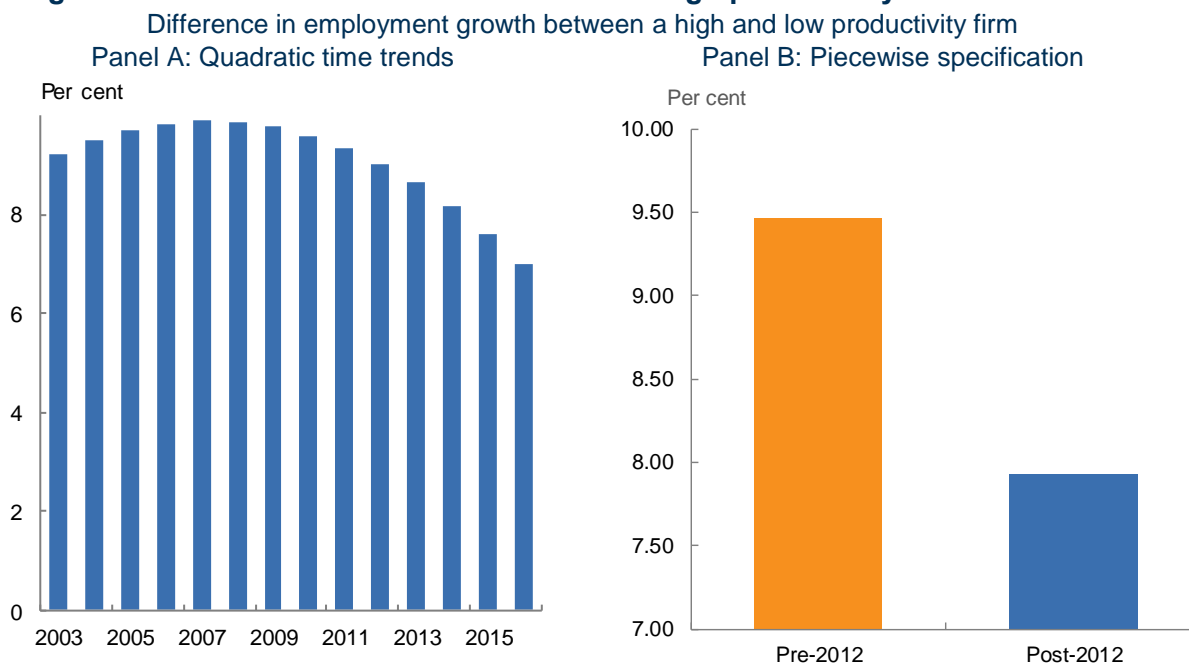
Notes: High (low) productivity firms are defined by being one standard deviation above (below) the industry mean labour productivity.
Source: Andrews and Hansell (2019).

6. DISCUSSION AND FUTURE WORK

Our paper exploits novel microdata sources to assess the relevance of a range of potential structural barriers to higher wage growth in Australia. In doing so, we provide a proof-of-principle demonstration of the value of longitudinal microdata, and related empirical techniques, to such investigations. In particular, our use of de-identified matched employer-employee data allow us to employ state-of-the-art high-dimensional fixed effects modelling to account for worker heterogeneity, and assortative matching between high ability workers and high productivity firms. Consistent with the international literature, our results highlight the importance of accounting for worker heterogeneity in explaining wage outcomes, and in making inference about their determinants.

Moreover, our analysis provides useful insights into low wage growth in Australia. This is a topic of some international interest given recent weak wage growth across a range of advanced economics, and is of domestic interest given the importance of wages at both an individual and aggregate level.

---

[21]   This result is robust to using alternative break points, such as 2012/13.

Overall, we find evidence that structural forces may have weighed on wage growth in recent years, and that these forces have seen a modest shift in the relationship between wage and productivity growth. Given that our dataset concludes in 2014/15, however, it is not clear whether lower pass-through of productivity to wages has persisted. While we cannot provide conclusive evidence on what these forces are, our results do help to guide future research by ruling out some, and providing tentative evidence in support of others.

In particular, our results suggest that declining labour market fluidity and market dynamism are potentially relevant factors to understand changes in the relationship between firm productivity and wage growth, and weak wage growth more generally. As such, any explanation for weak wage growth probably needs to also be able to explain changing patterns of market dynamism.

The implications of these findings are obviously highly dependent on the cause of lower wage growth and dynamism. For example, if firm-specific human capital has become more important – due to the rise in intangible capital – then the cost of job separation from both the perspective of the worker and firm will be higher, implying less need for job switching (Cairo and Cajner 2018). To the extent that this affects the bargaining position of workers and firms similarly, there may be limited ability or benefit to respond to these trends. Conversely, if these concurrent phenomena instead reflect increasing adjustment frictions in the economy (Decker et al., 2018), then there may be a scope to improve outcomes. Thus, identifying the reasons for the decline in market dynamism is a high priority for future research.

It is important not to overstate the significance of these results for wage growth going forwards. To date, the degree of 'unexplained' weakness in wage growth has been only moderate: potentially around ½ percentage point. Moreover, given we have not explicitly identified the drivers of the reduced pass-through of productivity to wages, it is not clear whether this effect will persist into the medium term. In particular, if the structural forces reflect a shift to a new 'equilibrium' relationship between the *level* of productivity and wages, the effects on wage *growth* should eventually diminish and thus be less relevant in the medium term.

In either case, if the public and private sectors are aware of these forces, they can continue to incorporate them into their forecasts. For example RBA (2018) notes:

"The central forecast for a moderate increase in wages growth over the next couple of years assumes that some of these factors will continue to weigh on wages growth for a while yet."

A greater understanding of the nature of the structural weakness will help to facilitate this.

Our results also provide further empirical support for the relationship between productivity and wages by providing the first empirical evidence on firm- and worker-level rent sharing for Australia. These results confirm international evidence of modest but not insignificant sharing of idiosyncratic firm-level productivity gains with workers (Card, Heining and Kline 2013).

Of course, there are other potential explanations for the observed decline in rent-sharing that we have not explored in this paper, some of which can only be tested with matched firm-worker panel data. These could be considered in future work.

For example, research has also linked the increasingly globalised nature of the labour market, and the rise of China, to lower worker bargaining power (for example, Autor, Dorn and Hanson 2016). Future work could examine this, particularly if we can couple the location and industry data in our de-identified matched firm-worker panel data with detailed data on imports. An examination of the link between wage bargaining patterns and rent-sharing would also be possible with the incorporation of micro-level wage setting data into our de-identified matched firm-worker panel database (see Guertzgen 2009).

Finally, it would be interesting to dig deeper into the relationship between wages and market dynamism. A logical next step would be to directly examine whether the shift in the relationship between wages and firm-level productivity was more pronounced in local labour markets where job switching declined more. Furthermore, younger workers are generally more likely to switch jobs, and job switching represents a key channel for them to increase their wages (for example, Haltiwanger, Hyatt and McEntrafer, 2017). As such, they could be more adversely affected by lower labour market fluidity. Finally, it would be interesting to explore how the results vary across high or low wage earners or for firms in regions with larger declines in labour mobility. This could give additional clues regarding the structural sources of wage growth in Australia.

REFERENCES

Andrews D, C Criscuolo and P Gal (2015), 'Frontier Firms, Technology Diffusion and Public Policy: Micro Evidence from OECD Countries', OECD Productivity Working Papers No 2

Andrews D, C. Criscuolo and P Gal (2016), 'The Best versus the Rest: The Global Productivity Slowdown, Divergence across Firms and the Role of Public Policy', OECD Productivity Working Papers, No. 5

Andrews D, P Gal and W Witherridge (2018), 'A Genie in a Bottle? Globalisation, Competition and Inflation', OECD Economics Department Working Paper No. 1462

Andrews D and D Hansell (2019), 'Productivity-enhancing Labour Reallocation in Australia', Treasury Working Paper (forthcoming)

Arsov I and R Evans (2018), 'Wage Growth in Advanced Economies', RBA *Bulletin*, March 2018

Australian Treasury (2017), 'Analysis of Wage Growth', Treasury Working Paper, November 2017

Autor DH, D Dorn and GH Hanson (2016), 'The China Shock: Learning from Labor Market Adjustments to Large Changes in Trade', NBER Working Paper No 21906, January 2016

Autor DH, D Dorn, L Katz, C Patterson and J Van Reenen (2017), 'The Fall of the Labour Share and the Rise of Superstar Firms', NBER Working Paper No. 23396, May 2017

Bagger, J, BJ Christensen and DT Mortenson (2014), 'Productivity and Wage Dispersion: Heterogeneity or Misallocation?' Mimeo.

Benmelech E, N Bergman and H Kim (2018), 'Strong Employers and Weak Employees: How Does Employer Concentration Affect Wages?'**,** NBER Working Paper No 24307, February 2018

Berlingieri G, P Blachenay and C Criscuolo (2017), 'The Great Divergence(s)', OECD Science, Technology and Innovation Policy Papers No. 39, May 2017

Bishop J and N Cassidy (2017), 'Insights into Low Wage Growth in Australia', RBA *Bulletin*, March 2017

Bishop J and I Chan (2019**),** 'Is Declining Union Membership Contributing to Low Wages Growth?', RBA Research Discussion Paper No 2019-02

Bishop J, L Gustafsson and M Plumb (2016), 'Jobs or Hours? Cyclical Labour Market Adjustments in Australia', RBA Research Discussion Paper No 2016-06

Campbell S, T Nguyen, A Sibelle and F Soriano (2019), 'Measuring Productivity Dispersion in Selected Australian Industries', Treasury-ABS Working Paper No. 2019-02, February 2019

Card D, AR Cardoso, J Heining and P Kline (2016), 'Firms and Labour Market Inequality: Evidence and Some Theory', NBER Working Paper No. 22850, November 2016

Card D, J Heining and P Kline (2013), 'Workplace Heterogeneity and the Rise of West German Wage Inequality', *The Quarterly Journal of Economics*, 128(1), pp 967-1015

Caselli F (1999), 'Technological Revolutions'**,** *American Economic Review*, 89(1), pp 78-102

Cassidy (2019), 'Low Wage Growth: An Introduction', Proceedings from RBA Conference on Low Wage Growth, Sydney, 4-5 April 2019

Davis SJ and J Haltiwanger (2014), 'Labor Market Fluidity and Economic Performance', NBER Working Paper No. 20479, September 2014

Decker RA, JC Haltiwanger, RS Jarmin and J Marinda (2018), 'Changing Business Dynamism and Productivity: Shocks vs Responsiveness', NBER Working Paper No 24236, January 2018

Deutscher N (2019), 'Job-to-job Transitions and the Wages of Australian Workers', Treasury Working Paper (forthcoming)

Dobbelaere S and J Mairesse (2018), 'Comparing Micro-evidence on Rent Sharing from Two Different Econometric Models', *Labour Economics*, 52, pp 18-26

Faberman RJ and A Justiniano (2015), 'Job Switching and Wage Growth', Chicago Fed Letter No. 337

Freestone O (2018), 'The Drivers of Life-Cycle Wage Inequality in Australia', *Economic Record*, 94(307), pp 424-444

Foster J and R Guttmann (2018), 'Perceptions of Job Security in Australia', RBA *Bulletin*, March 2018

Guertzgen N (2009), 'Rent-sharing and Collective Bargaining Coverage: Evidence from Linked Employer-Employee Data', *The Scandinavian Journal of Economics*, 111(2), pp 323-349

Haldane AG, 'The UK's Productivity Problem: Hubs No Spokes**,** Speech to Academy of Social Sciences Annual Lecture, London, 28 June 2018

Haltiwanger J, H Hyatt and E McEntrafer (2017), 'Who Moves up the Job Ladder?', NBER Working Paper No 23693, August 2017

Hansell D and B Rafi (2018), 'Firm-level Analysis Using the ABS' Business Longitudinal Analysis Data Environment (BLADE)', *Australian Economic Review*, 51(1), pp 132-138

Karahan F, R Michaels R, B Pugsley, A Şahin and R Schuh (2017), 'Do Job-to-Job Transitions Drive Wage Fluctuations over the Business Cycle?', *American Economic Review*, 107(5), pp 353-57

Lowe P (2017), 'The Labour Market and Monetary Policy', Address to the Anika Foundation Luncheon, Sydney, 26 July 2017

Lowe P (2018), 'Productivity, Wages and Prosperity', Address to Australian Industry Group*,* Melbourne, 13 June 2018

Manning A (2011), 'Imperfect Competition in the Labor Market', in Card D and O Ashenfelter (*eds*), *Handbook of Labor Economics: Volume 4b,* Elsevier, Amsterdam, pp 973-1041

Moscarini G and Postel-Vinay F (2017), 'The Relative Power of Employment-to-Employment Reallocation and Unemployment Exits in Predicting Wage Growth', *American Economic Review*, 107(5), pp 364-68

RBA (2018), 'Economic Outlook', RBA *Statement on Monetary Policy*, November

Roe, BE, and DR Just (2009), 'Internal and external validity in economics research: Tradeoffs between experiments, field experiments, natural experiments, and field data', *American Journal of Agricultural Economics* 91(5), pp 1266-1271

Song J, DJ Price, F Guvenen, N Bloom and T von Wachter (2019), 'Firming Up Inequality', *The Quarterly Journal of Economics*, 134(1), pp 1-50

Stansbury AM and LH Summers (2017), 'Productivity and Pay: Is the Link Broken?', NBER Working Paper No 24165

Van Reenen J (1996), 'The Creation and Capture of Rents: Wages and Innovation in a Panel of U.K. Companies', *The Quarterly Journal of Economics*, 111(1), pp 195-226

Weir G (2018), 'Wages Growth Puzzles and Technology', RBA Research Discussion Paper No 2018-10

APPENDIX A: MICRODATA SOURCES

Business Longitudinal Analysis Data Environment (BLADE)

The firm-level data used in this paper draws heavily on the Business Longitudinal Analysis Data Environment (BLADE), compiled by the Australian Bureau of Statistics. BLADE consists of administrative data from the Australian Taxation Office (ATO) matched with ABS-produced survey microdata, such as the Business Characteristics Survey.

In particular the analysis on productivity reallocation and average wages draws on two key administrative data sources. First, measures of output and intermediate use are sourced from Business Activity Statements (BAS). While each annualised BAS dataset contains more than two million records, only around 800,000 or more units report formal employment via annual payment summaries (APS).[22] We also draw on this payment summary information in combination with the industry of origin of the employing business to estimate full-time equivalent employees – our measure of labour input.[23] For the BLADE data analysis in Section 4.2, we construct value-added as sales less non-capital expenditure, which capture intermediate inputs used by the firm in producing their output.

The main demographic information in BLADE consists of industry and institutional sector codes, which permit us to focus on industries and businesses that are privately-owned. The ABS adds an additional year of administrative and survey data to BLADE each year. As of writing, the latest available year was 2016 and so our sample – unless specified otherwise – encompasses the period 2002 to 2016.[24]

As is standard in the literature, we confine our analysis to focus on the market sector (that is, we exclude utilities; education, public administration and safety and health) and exclude finance and insurance where productivity is notoriously difficult to measure. Our final firm-level dataset for analysis covers around 75 per cent of gross value added and 84 per cent of compensation of employees in the market sector over 2002-2016.

Matched Firm-Worker Panel

In addition to the firm-level data results from BLADE, the Australian Treasury receives de-identified Australian Taxation Office (ATO) for the purposes of revenue analysis.[25] These data encompass annual employee earnings for all Pay-As-You-Go taxpayers and detailed balance sheet information on incorporated entities. For this paper and other revenue-related analysis, we have created a panel of de-identified matched firm-worker data. In this database, we define firm value added as a measure of gross operating surplus (essentially profits) plus compensation of employee. We use this definition, instead of the definition used in BLADE, due to differences the reporting of intermediate inputs across tax forms.

Since these data encompass all compliant entities within the Australian economy, it is difficult to provide summary statistics since the cross-sectional and longitudinal scope of each slice of the data is determined by the analytical question at hand. Currently, our data cover 2001-02 to 2014-15 (2015-16 for some of

---

22   It is important to note that the ABS does not permit access to person-level data. All information from the payment summaries has been aggregated to the business unit.
23   For more information see Hansell, Nguyen and Soriano (2015).
24   The Australian fiscal year ends in June. For simplicity, we refer to the last year of the fiscal year, and so the 2001-02 fiscal year becomes 2002.
25   Note the de-identified data contains no information capable of directly identifying any individual or business.

the analysis in this paper) but is limited to those employed by incorporated entities. For example, the following table contains summary statistics on annual individual wages over the sample period.[26]

Fiscal Year	N	Mean	Median	25th Percentile	75th Percentile	SD
2002	2,936,914	27,906	21,923	6,406	39,200	39,128
2003	3,329,139	28,153	21,854	6,319	39,906	35,602
2004	3,583,533	28,654	21,130	5,995	40,352	48,154
2005	3,698,581	29,324	21,742	6,137	41,874	38,451
2006	3,845,696	30,836	22,935	6,586	43,798	41,075
2007	3,919,153	32,575	23,783	6,803	45,778	46,912
2008	4,091,544	34,703	25,017	7,264	48,402	49,378
2009	3,951,879	36,897	27,684	8,042	51,384	50,740
2010	3,865,464	38,110	28,324	8,424	52,844	51,355
2011	3,993,307	40,373	29,889	8,962	55,599	55,253
2012	4,060,880	42,232	31,199	9,454	58,322	54,264
2013	4,122,927	44,246	32,677	10,000	60,706	57,359
2014	4,079,778	45,604	33,959	10,600	62,352	60,297
2015	4,243,596	46,524	34,662	11,050	63,585	61,976

For companies this corresponds to approximately 739,000 distinct entities with the following attributes.

Variable	N	Mean	Median	25th Percentile	75th Percentile	SD
Employment	4,048,349	16.02	4	2	10	360.9
Value-added	4,048,349	1,599,857	133,941	46,743	374,935	116,000,000

Central to the questions of whether wage growth has slowed and by how much is how one defines wages. The most cited measures of are the wage price index (WPI), total compensation or average compensation of employees (AENA).[27]

When comparing these measures of wage growth to our micro-data, it is important to consider what exactly is captured in these measures. The WPI tracks the compensation paid to given occupations, and so is not affected by compositional shifts, such as workers leaving high paying mining jobs at the end of the mining boom, and taking up other, lower paid, jobs. In contrast, AENA tracks the aggregate compensation paid to all workers in the economy, and is affected by such compositional shifts. As such, it has tended to be below the WPI in recent years.

Our microdata, on the other hand, track individual workers. The data will be affected by compositional change, similar to AENA, as a worker shifting from a high paying job to a low paying job will be recorded as having low wage growth (assuming we are not focusing on worker-employer matches). In this sense, it is somewhat akin to the AENA measure. However, we cannot calculate wage growth for an individual that transitions from work to unemployment, or unemployment to work, which might make our measure less volatile.

---

26   Note that to make our regressions more tractable, we limit each individual to their main income source for the year. In the (rare) event of a tie, we use random assignment.

27   AENA is obtained by dividing total compensation of employees by the total number of employees (see introduction to ABS catalogue number 6105 (2005).

Notwithstanding these differences, when we compare the relative rates of change for these three series, it shows that the microdata we exploit in this paper exhibits similar trends over the sample period (Figure A1). Indeed, the relative rate of change for our sample is often in the middle between the WPI and AENA.

Figure A1: Wage Growth Measures
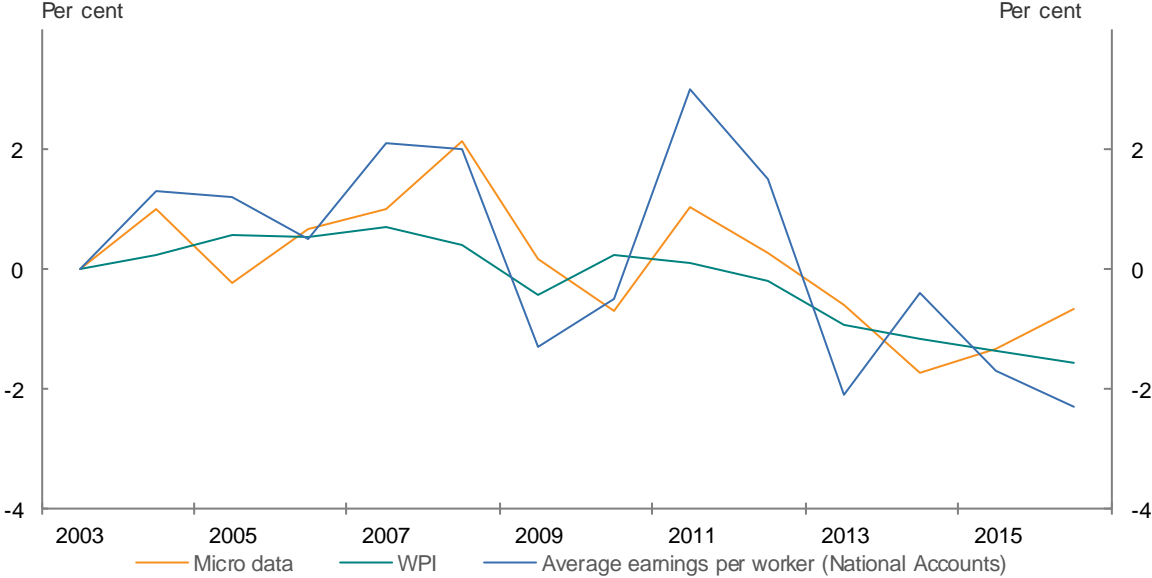
Notes: Micro data series show unweighted average of individual workers' wage growth.
Sources: ABS; Treasury analysis of de-identified tax data

Appendix B: Additional results

Figure B1 demonstrates the broad-based decline in wages growth over recent years. It plots the differences in the coefficients on financial year dummies for 2010-11 and 2015-16, from regressions of the change in log annual wages on financial year dummies and certain micro fixed effects. The model with no fixed effects simply captures the change in workers' unweighted average wage growth, while the other specifications capture average wage growth for individual workers, firms, or worker firm matches (that is, a given worker at a given firm). In all cases, the decline in wage growth over the period is around 2 percentage points. This provides further evidence that the decline in wage growth has been broad-based, and has not reflected any compositional shifts between different types of workers, firms, or the matches therewith.

Figure B1: A broad-based decline in wage growth since 2010-11
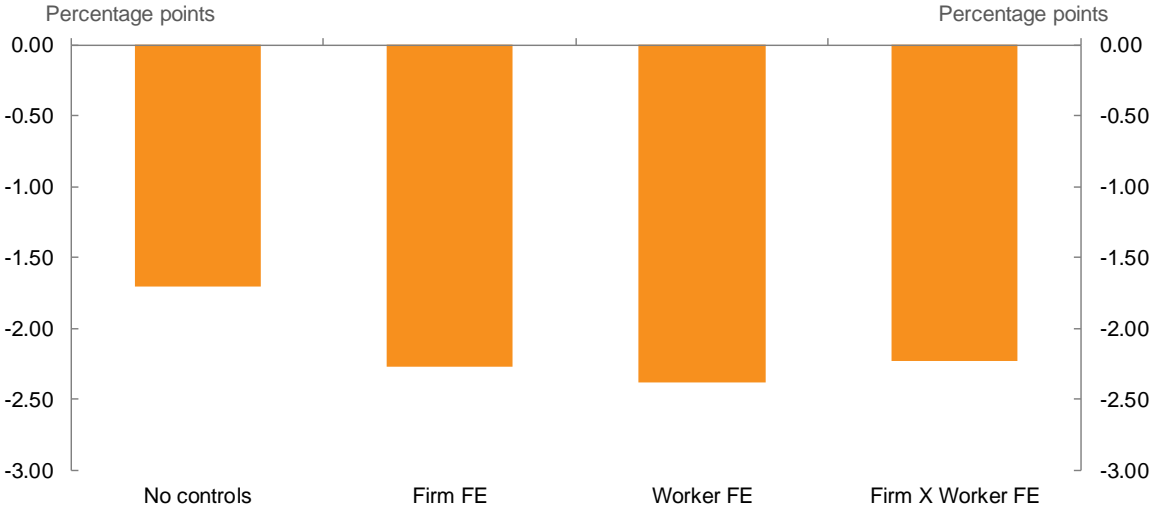
Note: Presents the fall in nominal wage growth between 2010-11 and 2015-16, estimated as the difference in the coefficients from a regression of change in log annual wages on financial year dummies and either: no controls; firm fixed effects; worker fixed effects; and firm-worker pairing fixed effects. The latter two specifications also include controls for a quadratic in worker age.
Sources: Treasury calculations based on confidential taxation data.

Table B1 shows the results from the changes version of our model examining the relationship between firm wages and productivity. Consistent with the levels specification, the coefficient on productivity tends to be around 0.1. This provides additional robustness to the results from our level model. Moreover, the fact that the level and changes specifications are quite similar suggest that it there is not some inherent difference between the levels and changes specifications that is driving our finding that the relationship between firm wages and productivity has weakened.

Table B1: The changing link between wages growth and firm productivity growth
Dependent variable: Change in the log wages of the individual

	(1)	(2)	(3)	(4)	(5)
ΔProductivity	0.119***	0.116***	0.105***	0.101***	0.097***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)
Common controls	X	X	X	X	X
Fixed effects					
Firm		X		X	
Worker			X	X	
Firm#Worker					X
R^2	0.036	0.058	0.231	0.252	0.303
Adjusted R^2	0.036	0.042	0.037	0.043	0.072
N	27,325,184	27,325,184	27,325,184	27,325,184	27,325,184

Note: All regressions include a set of common controls covering worker characteristics (sex, a quadratic in age and residential location (Statistical Area 4)), firm characteristics (size and industry division), and cyclical effects (state and industry division separately interacted with financial year). Errors are clustered at the firm-level.