e61 MICRO NOTE 14 | 6 November 2023

PRODUCTIVITY IN MOTION: THE ROLE OF JOB SWITCHING

Jack Buckley

International research suggests that job switching can boost aggregate productivity by moving workers to more productive firms (Albagli et al., 2021), but direct evidence on this mechanism in Australia is scarce.

This note sheds new light on the possible contribution of job switching to productivity growth in Australia.

- Workers who switch jobs move to firms that are (on average) 13.1% more productive than the firms they leave. This differential is even more pronounced for young workers.
- But the average productivity gap between 'origin' and 'destination' firms has more than halved since the mid-2000s, which could suggest a decline in productivity enhancing reallocation.

Australia's productivity slowdown over recent decades has coincided with a decline in job switching (Andrews & Hansell, 2021). Job switching could boost aggregate productivity via many channels. One possible channel involves workers moving to firms with higher average productivity, and workers becoming more productive as a result. This could occur if the higher productivity of a worker's new firm is due to better management, firm-specific know-how or more and better use of capital.

We combine data on individual workers' actual job switches with firm-level estimates of productivity (see Appendix A.1). We then estimate the difference between the average labour productivity of the firms that workers leave and the average labour productivity of the firms they join.[1] The size of this 'productivity gap' indicates whether workers are, on average, moving to more productive firms.[2] We find that:

1. **The productivity gap between origin and destination firms is large.** The average productivity of the firms workers move to is 13.1% higher than the average productivity of the firms they leave.
2. **The productivity gap is more pronounced for young workers.** Workers aged 25-34 move to firms that are (on average) 14.6% more productive, while workers aged 45-54 move to firms that are 9.3% more productive (Figure 1 panel B).
3. **The average productivity gap has fallen over time** from 14.9% in the mid-2000s, to 6.7% in the late 2010s (Figure 1 panel A). The difference in the average productivity percentile (rank) of a worker's origin and destination firm has also fallen (Figure A.1 panel A), suggesting that changes in the firm-level productivity distribution are not driving the result.

Figure 1: Average productivity gap by productivity growth cycle

Difference betwen mean labour productivity of origin and destination firms*

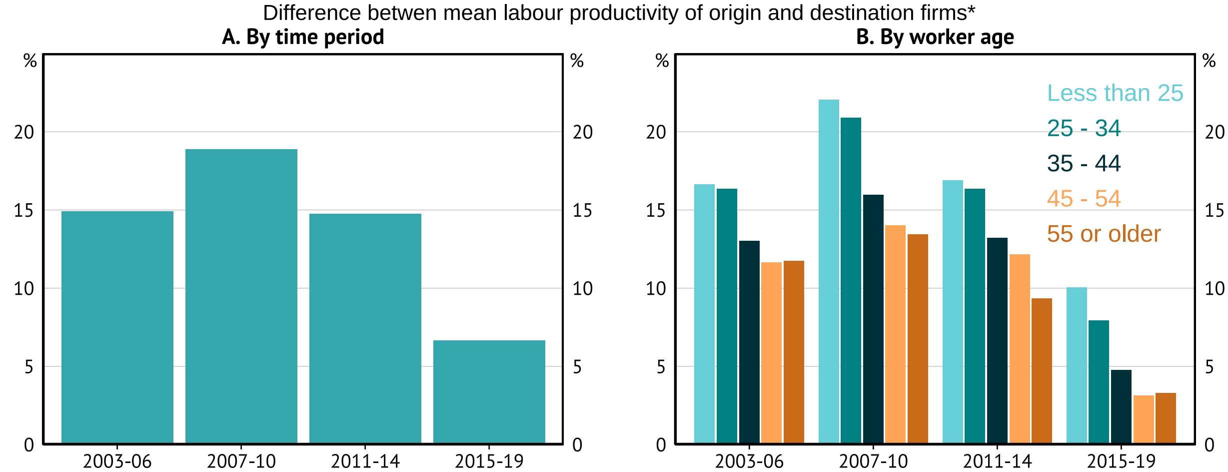

* Productivity is measured as (lagged) labour productivity (Value add/FTE). Productivity estimates are demeaned by year-industry (ANZSIC 2-digit) to account for differences in capital intensity between industries. Productivity gaps are the difference between the mean productivity of origin and destination firms.
Sources: ABS; e61

1 Labour productivity is defined as (lagged) output per worker (value added/FTE). Productivity is used interchangeably with labour productivity throughout.
2 Our underlying assumption is that firms with higher average productivity also have higher marginal productivity (the productivity of an additional worker). While the two are unlikely to be equal, with adjustment frictions average productivity will be more reflective of marginal productivity (See Appendix A.2).

What underpins the average productivity gap between origin and destination firms?

The average gap between origin and destination firms is driven by a small majority (53.6%) of switches that move workers to more productive firms. The share of job switches that involve moving to a higher productivity firm has fallen over time, from 54.2% between 2003 and 2006 to 52.8% between 2015 and 2019 (Figure A.1 panel A).

How is the average productivity ‘gain’ so large when so many workers move to less productive firms?

The average positive productivity gap between destination and origin firms is 23.7% larger (in absolute value terms) than the average negative gap associated with moves to less productive firms. This means the ‘gains’ from the 53.6% of moves to more productive firms more than outweigh the ‘losses’ from the remaining 46.4% of moves to less productive firms.

Why are so many workers moving to less productive firms?

People switch jobs for a range of reasons. Many of these reasons are likely unrelated to firm productivity, such as personal fulfilment and a better work-life balance (NAB, 2022). Of workers moving to less productive firms:

- 22.2% are changing occupation (ANZSCO major group), potentially signifying a career change. This is more common for younger workers with the rate of occupation switches 10.6 ppts higher for workers aged 25-34 than workers aged 45-54.
- 10.7% are joining a young firm (established less than five years ago) that is growing in size (headcount). These firms often have significant growth potential and may offer opportunities for career advancement (Haltiwanger et al., 2016).
- 2.8% are leaving a capital city to live in a regional (non-capital) area, potentially signifying a lifestyle change. For workers aged 25-34 this figure is even higher (3.4%) and could be related to the millennial exodus from cities (Dwyer, 2023).[3]

Why might young workers be more likely to gravitate to a higher productivity firm?

One reason may be that young workers are more likely to need to change jobs in order to find the one that is the ‘right fit’ for their skills (Adams et al., 2022; Andrews et al., 2022). Young workers are also more likely to move to a higher paying job (Wong, 2023), and higher paying jobs are more likely to be found at more productive firms (Carlsson et al., 2016).

Could job switching affect productivity in other ways?

Reallocating labour to more productive firms could help boost aggregate productivity, particularly where the sources of higher productivity are firm-specific and workers have similar capabilities. But where workers differ in their capabilities, at least two other channels come into play:

1. **Diffusion of new technologies and ideas.** When workers join a new firm they often bring with them new ideas, new technologies or better ways of doing things. This can make their new firm more productive.
2. **Better job matching.** Changing jobs can help workers find a role that better matches their skills. Workers in better matched jobs tend to be more productive (Coraggio et al., 2022).

Policy implications

1. **Declining labour mobility could have contributed to Australia’s aggregate productivity slowdown.** Our work shows that workers tend to move to firms that are (on average) more productive. But this tendency appears to have weakened.
2. **Making it easier for workers to switch jobs could help boost productivity.** Australian workers currently face a number of barriers when looking to move to a more productive firm. Some of these barriers have gotten worse in recent years. Non-compete clauses (NCCs), for instance, have become more prevalent over the last 15 years and are now a default option in many employment contracts (Andrews & Jarvis, 2023). Removing or limiting the use of NCCs would help remove one source of friction behind the decline in job mobility. Examples of other policies that could reduce frictions include:
   - reforming occupational licensing
   - replacing stamp duty with a land tax.

3 Compared to workers moving to more productive firms, workers moving to less productive firms are 1.6% less likely to be changing occupation, but 29.3% more likely to be joining a young growing firm, and 7.1% more likely to be moving to a regional area from a capital city.

References

Adams, N., Andrews, D., Auer, J., Dwyer, E., Hayward, Z., Kaplan, G., La Cava, G., Michielsen, A., & Nolan, M. (2022). *Better Harnessing Australia's Talent: Five Facts for the Summit*, e61 Institute.

Albagli, E., Canales, M., Syverson, C., Tapia, M., & Wlasiuk, J. (2021). *Productivity Growth and Workers' Job Transitions: Evidence from Censal Microdata*, Becker Friedman Institute. https://bfi.uchicago.edu/working-paper/productivity-growth-and-workers-job-transitions-evidence-from-censal-microdata/

Andrews, D., Auer, J., Elias, M., La Cava, G., & Mercier, L. (2022). *The Effect of the COVID-19 Recession on the Youth Labour Market in Australia* (tech. rep.). e61 Institute. https://e61.in/the-effect-of-the-covid-19-recession-on-the-youth-labour-market-in-australia/

Andrews, D., & Hansell, D. (2021). Productivity-Enhancing Labour Reallocation in Australia. *Economic Record*, *97*(317), 157–169. https://doi.org/10.1111/1475-4932.12601

Andrews, D., & Jarvis, B. (2023). *The ghosts of employers past: how prevalent are non-compete clauses in Australia?*, e61 Institute.

Carlsson, M., Messina, J., & Skans, O. N. (2016). Wage Adjustment and Productivity Shocks. *Economic Journal*, *126*(595). https://doi.org/10.1111/ecoj.12214

Coraggio, L., Pagano, M., Scognamiglio, A., & Tåg, J. (2022). JAQ of All Trades: Job Mismatch, Firm Productivity and Managerial Quality. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4069721

Decker, R. A., Haltiwanger, J., Jarmin, R. S., & Miranda, J. (2020). Changing Business Dynamism and Productivity: Shocks versus Responsiveness. *American Economic Review*, *110*(12). https://doi.org/10.1257/AER.20190680

Dwyer, E. (2023). *Sydney's Millennial Exodus*, e61 Institute.

Haltiwanger, J., Jarmin, R. S., & Kulick, R. B. (2016). High Growth Young Firms: Contribution to Job, Output and Productivity Growth. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.2866566

Multi-Agency Data Integration Project (MADIP). (2003 - 2019). *MADIP Modular Product, ABS DataLab. Findings based on use of MADIP data.*

NAB. (2022). *Australia's Great Resignation - Is it happening, which sectors are exposed & what are the key drivers?*, National Australia Bank. https://business.nab.com.au/wp-content/uploads/2022/02/NAB-Behavioural-Insights-Australias-Great-Resignation.pdf

Wong, A. (2023). *Climbing the wage ladder: Linking job mobility and wages*, e61 Institute.

A.1. Data construction

We construct a longitudinal linked employee-employer dataset (L-LEED) covering the period from 2003 to 2019. The L-LEED is created by linking workers (the ABS MADIP dataset) to firms (the ABS BLADE dataset) using pay-as-you-go (PAYG) tax records from the ATO (Multi-Agency Data Integration Project (MADIP), 2003 - 2019). This dataset covers all employees who work at a firm who submits a Business Income Tax return.

In the MADIP and PAYG data we use the following information about employees: age, gender, annual earnings at each job held during the financial year, occupation, unique individual and firm identifiers, and the industry (ANZSIC 4-digit) of their employer. We obtain information on firm turnover, cost of goods sold, derived FTE headcounts and other variables used to estimate our alternate productivity measures from the BLADE data.

From the full L-LEED we apply the following filters to arrive at the final dataset we use in our analysis:

- We exclude the non-market sector of the economy as productivity is difficult to measure there. We also test the sensitivity of our results to excluding the volatile mining and agriculture sectors and finance and insurance sectors where productivity can be difficult to measure. Our main results are largely unaffected by the inclusion or exclusion of these industries.
- We exclude worker-firm matches that are not a worker's main job. Main jobs are defined as a worker's highest paying job in a particular year. In order to look at main job to main job switches, we include all jobs that were either a worker's main job last year or their main job the following financial year.

Defining job switching

In our analysis of job switching we focus on job-to-job transitions. A job-to-job transition occurs when a worker changes jobs without going through a period without paid employment. However, due to the unreliable nature of employment start and end dates in the PAYG data we am not able to precisely identify the date a job-to-job transition (J2Js) occurs.

To address this issue, we focus our analysis on J2Js where a worker leaves a job that was their main job last year and starts a job that becomes their main job next year. We exclude observations where a worker changes jobs more than once within a particular year or holds more than two jobs during the year. This ensures that we correctly identify the origin and destination firm in each transition. Defining J2Js in this way means that we will inevitably capture some JNJ transitions (where a worker moves from employment to unemployment to employment again). We test the sensitivity of our results to this by dropping workers who earn less than the full time minimum wage within a given year. This exclusion has a very small positive effect on the average productivity gains in our main analysis, but the trends remain the same.

Measuring productivity

The main measure of labour productivity we use in our analysis is value add per FTE worker, where value added = turnover - cost of goods sold. We also test the sensitivity of our results to alternate measures of productivity, including gross output (turnover) per FTE worker. FTE is imputed by the ABS using information on firm headcount and worker wages.

In our headline analysis, we make the following adjustments to the raw productivity estimates that are common in the literature (see, for example, Albagli et al., 2021):

- Winsorize the top and bottom 2.5% of the firm productivity distribution to avoid extreme values distorting the results.
- Demean by year-industry (2-digit ANZSIC) to account for differences in capital intensity across industries. The broad trends in the main results are unaffected when using data that has not been demeaned, but the average productivity gains increase when the data is not demeaned.
- We take a one year lag of each firm's productivity to ensure that the job-to-job transition is not affecting our estimates.

Finally, to measure the mean productivity gap, we take the mean productivity of origin firms and compare it to the mean productivity of destination firms (dropping any observations where either the origin or destination firm's productivity cannot be measured). We do this to prevent the results being distorted by transitions that involve a small absolute productivity gap, but a large percentage difference in productivity (e.g. when a worker leaves a firm whose value add is close to zero). Taking the mean percentage gap results in implausibly large estimates. An alternate approach used in the literature is to measure productivity in log form (i.e. log(VA/FTE)) and take the mean log difference of switches. This approach produces very similar results to our main analysis (if log differences are assumed to be approximately equal to the percentage change, which is only true for small percentage changes).

A.2. Average and marginal productivity gains

Our key underlying assumption is that firms with higher average productivity also have higher marginal productivity (the productivity of an additional worker). Under this assumption, if a worker moves to a firm with higher average productivity this will have a positive effect on aggregate productivity.

Is this assumption likely to be true?

Economic theory suggests that marginal products (in our case the marginal productivity of an additional worker) will be equalised across firms in a market. But this is only true in a frictionless market. When adjustment frictions exist marginal products are less likely to be equalised. In our case the presence of labour market frictions, such as search costs for firms and switching costs for workers, means that the average productivity of firms will be more reflective of their marginal productivity (Decker et al., 2020). Further, as Decker et al. (2020) show, under certain assumptions (constant returns to scale and perfect competition), moving resources to a firm with higher average productivity will always increase aggregate productivity, because the marginal revenue product of a firm does not change with scale but only varies with their total factor productivity.

A.3. Additional analysis

The charts in Figure A.1 present additional analysis of the variation in the productivity gains from job switching. Panel A shows that the fraction of all job switches that move workers to more productive firms is only marginally higher than the fraction that move workers in the opposite direction. It also highlights that the share of transitions that move workers to more productive firms has fallen over time.

Panel B shows that the size of productivity gains also varies based on the productivity of the firm a worker is leaving. Workers at less productive firms – those in the bottom half of the firm productivity distribution – on average move to more productive firms. Workers at more productive firms – the top half of the distribution – on average move to less productive firms (Figure A.1 panel B). This pattern of reallocation is closer to a random reshuffling of workers (orange line, Figure A.1 panel B), than a steady progression of workers up the firm productivity ladder (red line, Figure A.1 panel B).

Figure A.1: Share of transitions to more productive firms*

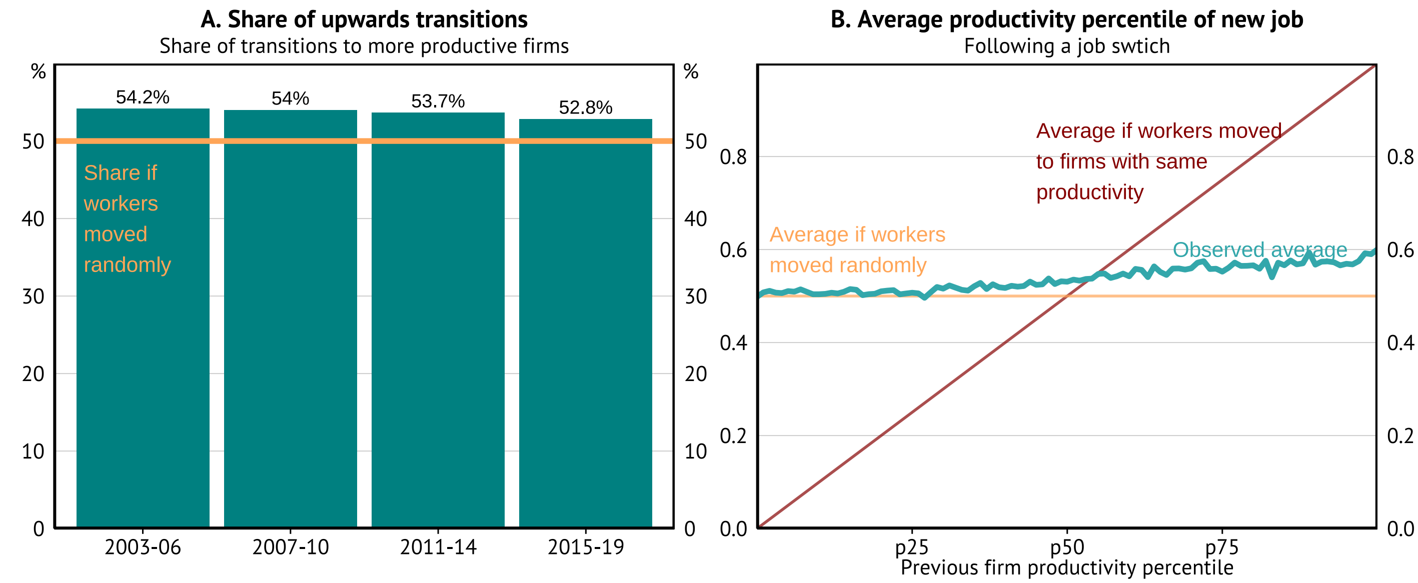

* Productivity is measured as (lagged) labour productivity (VA/FTE).
Sources: ABS; e61

A.4. Sensitivity to alternate measures of productivity

In Figure A.2 we test the sensitivity of our main results (Figure 1 panel A) to alternate measures of firm productivity, including:

(a) Difference between the productivity percentile of the origin and destination firm within their industries.

(b) Total factor productivity (TFP), estimated following the approach suggested by Ackerberg, Caves and Frazer (2015) using Stata's *prodest* package.

(c) Labour productivity measured as turnover/FTE (TR/FTE) rather than value add/FTE (VA/FTE).

Figure A.2: Alternate measures of firm productivity

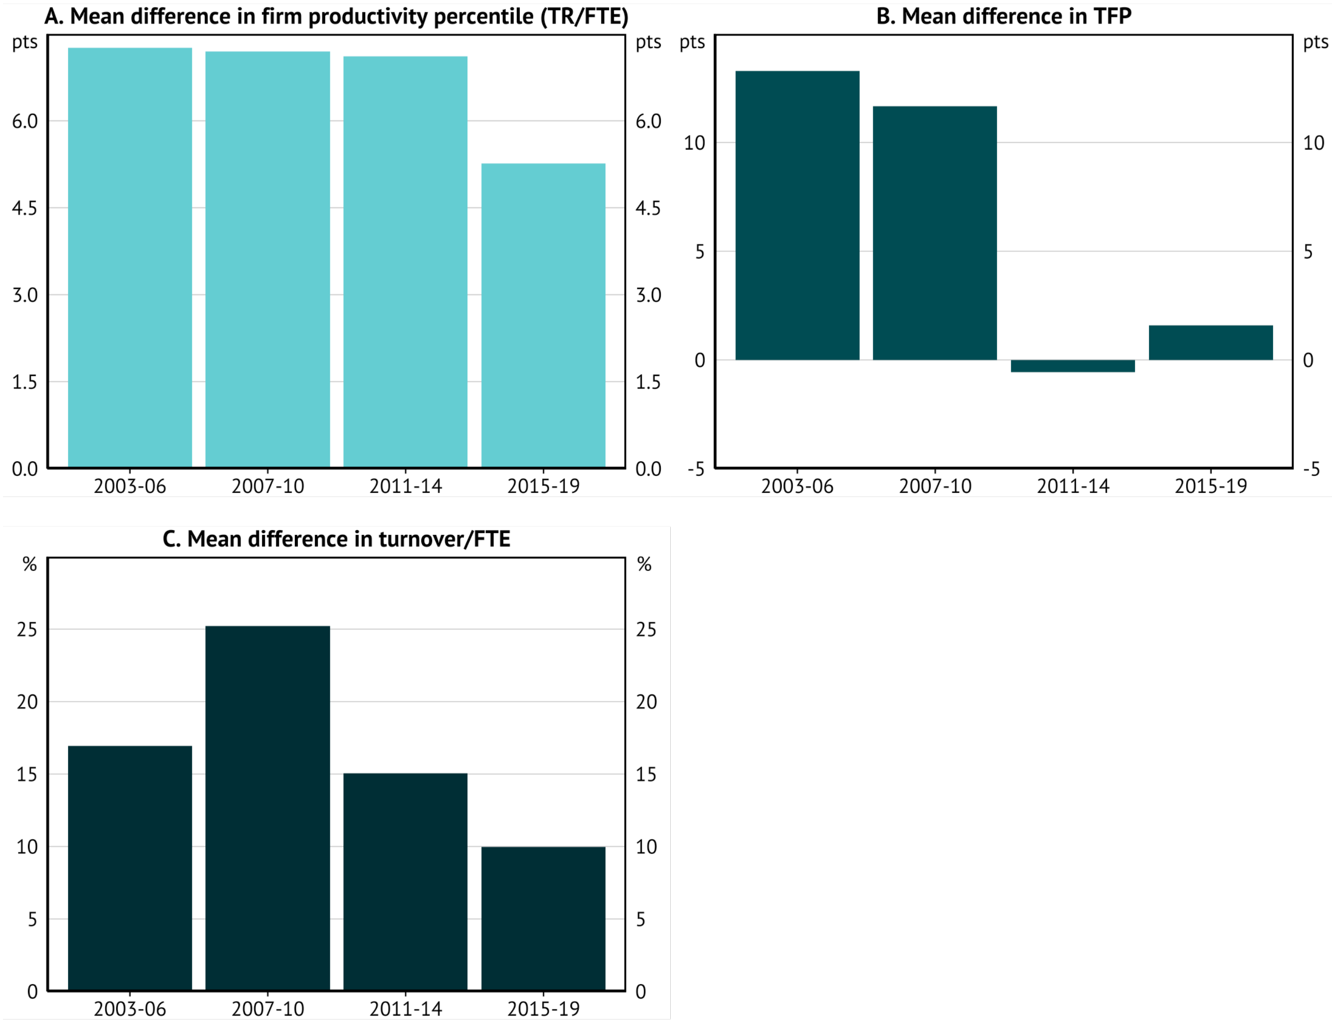

* Productivity gains are the difference between the mean productivity of origin and destination firms.
Sources: ABS; e61 analysis

The results using these alternate measures of productivity are similar to our main analysis. We still find that the average productivity of destination firms is higher than the average productivity of origin firms (with the exception of the TFP results post GFC) and that this difference has fallen in recent years relative to the early 2000s. Notably, we find that the average difference in productivity percentiles has declined recently, suggesting that our results are not simply due to a narrowing of the overall productivity distribution.