DEVELOPMENT POLICY CENTRE

Discussion Paper 116

August 2025

Series ISSN 2206-303X

Do seasonal work visas suppress the earnings of incumbent farm workers?

Toan Nguyen and Ryan Edwards

Abstract

This paper uses comprehensive tax and visa records to examine the impact of seasonal work visas on the earnings of incumbent farm workers in the context of Australia's Seasonal Worker Program (SWP). Using a triple-difference approach to exploit differences in earnings growth across different low-skilled occupations within localities, we find no convincing evidence that SWP visas reduced farm worker earnings. Our main null result is robust to alternative specifications and approaches, including regional and national (skill-cell) difference-in-differences, and is consistent with the long-term population decline in agricultural regions and a growing inability or unwillingness of the domestic workforce to fill short-term agricultural roles.

The Development Policy Centre is part of Crawford School of Public Policy at The Australian National University. We undertake research and promote discussion on Australian aid, Papua New Guinea and the Pacific and global development policy.

Australian National University

Do seasonal work visas suppress the earnings of incumbent farm workers?

Toan Nguyen and Ryan Edwards

Toan is a Research Fellow at the Development Policy Centre.

Ryan Edwards is the Deputy Director of the Development Policy Centre.

Nguyen, T. Edwards, R. 2025. Do seasonal work visas suppress the earnings of incumbent farm workers?, *Development Policy Centre Discussion Paper 116,* Crawford School of Public Policy, The Australian National University, Canberra.

Thanks to Bob Breunig, Paul Burke, Alfredo Paloyo, Stephen Howes and seminar participants at Monash University, the University of Sydney, the University of the South Pacific, the Australian Public Finance Workshop at the Australian National University and the University of Melbourne for helpful comments and discussions. This research was supported by the Pacific Research Program, with funding from the Department of Foreign Affairs and Trade, received ANU Human Ethics approval under protocol 2022/491 and was conducted in the Australian Bureau of Statistics DataLab under approved project 2022023. Views and errors are the authors' own.

For more information on the Development Policy Centre, visit **devpolicy.crawford.anu.edu.au**

1 Introduction

A central question in labour economics concerns the impact of immigration on the earnings of incumbent workers (Card, 2001; Borjas, 2003; Dustmann, Schönberg, and Stuhler, 2016). Much of the existing research focuses on the impacts of long-term (permanent) migration, yet the findings remain inconclusive (Dustmann et al., 2016). Some studies find that immigration negatively affects native workers' labour market outcomes (Borjas, 2003; Llull, 2018a). Others find minimal impacts (Card, 1990, 2001) and some find effects only for particular groups, which depend on the margins along which the labour market adjusts (Dustmann, Frattini, and Preston, 2013; Dustmann, Schönberg, and Stuhler, 2017; Foged and Peri, 2016; Llull, 2018b; Manacorda, Manning, and Wadsworth, 2012; Ottaviano and Peri, 2012; Peri and Yasenov, 2019). Within this burgeoning literature, evidence on the impacts of seasonal work visas, which account for a significant share of international labour mobility, including in the United States, Canada, and New Zealand (Martin, 2025; Curtain, 2025; Bedford and Bedford, 2025), remains incredibly thin. These short-term, non-resident, temporary work visas which aim to fill narrow labour market gaps, where native labour supply is perceived to be low, are central in modern policy debates on immigration, but very poorly understood (Clemens, 2022). Here we provide, to the best of our knowledge, the first quantitative assessment of any seasonal visa program globally on native earnings.

This paper measures the impact of Australia's Seasonal Worker Program (SWP) on incumbent farm workers' earnings and employment outcomes.[1] The SWP was introduced in July 2012 as a development initiative to provide temporary job opportunities for Pacific Islanders and supply labour for Australia's agricultural sector. As Australia's primary temporary visa program focused on seasonal agricultural and a high-regulated scheme widely regarded as global best practice (Curtain, 2025; Doyle and Howes, 2015; Martin, 2025; Bedford and Bedford, 2025), the SWP offers an important case study to test whether seasonal work visas depress

[1]This includes all taxpayers who worked in the horticulture sector, including both Australian-born and foreign-born workers. However, our baseline results are based on those who remained in the region in both 2010 and 2017, so short-term migrants are excluded.

the earnings of some of the lowest paid workers in the economy, or whether domestic labour supply for these roles is sufficiently low or inflexible enough for no such consequences.

We find no compelling evidence that the SWP reduced the earnings of incumbent farm workers relative to other low-skilled workers in the same communities. Similar patterns are observed using less restrictive comparisons across regions or occupations. We also find limited evidence of any adverse employment impacts in host regions, suggesting either seasonal absorption or other margins of adjustment. Seasonal work visas are often premised on labour shortages and focus on particular sectors, occupations, and regions where domestic labour supply is presumed to be low and insufficient to meet seasonal needs. Descriptive analysis exploring historical trends in these labour markets reveals lower population growth compared to regions with fewer SWP workers, a higher concentration of agricultural activity, and declining domestic labour supply across the agricultural sector well before the SWP was introduced. We cautiously conclude that the SWP is likely to be alleviating regular preexisting seasonal labour shortages, with little consequence for the earnings or employment prospects of incumbent farm workers.

We offer two main contributions. First, we add new evidence to the nascent body of research in labour economics on the labour market impacts of seasonal work visas, a topic that remains chronically underexplored relative to its immense global economic scale and centrality in modern immigration policy debates. Clemens, Lewis, and Postel (2018) examine the labour market consequences of the 1964 termination of the U.S. Bracero program, which had allowed Mexican seasonal agricultural workers to work in the United States. Excluding Braceros did not result in higher wages or substantially increased employment for U.S. farm workers. More recently, Clemens (2022) finds that restricting foreign seasonal farm workers in the U.S. during the Great Recession also had virtually no impact on native employment. Here, we take the establishment and growth, rather than restriction or termination, of a similar modern program, and leverage remarkably precise administrative data to identify its impacts on the earnings of the most exposed domestic workers.

Our second contribution is more practical. Our findings suggest that uncapped, employer-sponsored migration programs, when carefully designed to support sectors with genuine and persistent labour shortages, do not necessarily harm or displace existing workers. The Australian SWP serves as a clear example of how well-designed, demand-driven migration can support structural adjustment without negatively affecting incumbent workers in declining regions (Duncan, Harris, Mavisakalyan, and Nguyen, 2020).

We use rich administrative data and a triple differences design to estimate the impact of Australia's SWP on existing farmworker earnings and employment. Drawing on the Personal Level Integrated Data Asset (PLIDA), which links tax, employment, visa, and other administrative records across government agencies, we observe earnings for the universe of individual taxpayers from fiscal years 2010–11 to 2017–18. Leveraging tax and visa records together, we can precisely identify the locations and occupations of individuals most directly exposed to the recent growth in SWP workers. Our primary identification strategy is a triple-difference approach exploiting variation across three dimensions: time, geography and occupation. By comparing the earnings trajectories of similar workers across treated and control regions and occupations, before and after the program's announcement and expansion, we are able to isolate the causal effect of the program on domestic workers most likely to be affected. Individual fixed effects remove any person-specific, time-invariant unobservable confounders and offer additional precision. We supplement our main triple differences approach with evidence from traditional regional and occupational difference-in-difference designs, and a regional event study tracing year on year dynamics before and after SWP establishment.

The next section introduces the SWP and situates our approach within competing empirical strategies for measuring the impacts of immigration. Section 3 introduces the Australian administrative data, details our estimation and identification strategy, and provides some preliminary descriptive evidence. Section 4 presents our main null results on earnings, key robustness checks, and consonant findings from alternative workhorse approaches. Section 5 explores potential explanations for the null result and Section 6 concludes.

2 Background

2.1 Australia's Seasonal Worker Program

The SWP was formally established in 2012 with two main objectives: (1) to contribute to economic development in partner Pacific countries, and (2) to provide benefits to Australian employers and the domestic economy (Parliament, 2016; for Global Development, 2025). The program followed a three-year pilot—the Pacific Seasonal Worker Pilot Scheme (PSWPS)—launched in 2008, with the first workers arriving in 2009 (Gibson and McKenzie, 2011). The pilot, modelled on New Zealand's successful Recognised Seasonal Employer (RSE) scheme introduced in 2007, received positive evaluations and led to the full rollout of the SWP (Doyle and Howes, 2015; Curtain, Dornan, Doyle, and Howes, 2016; World Bank, 2018). While Gibson and McKenzie (2011), World Bank (2018), and Gibson and McKenzie (2014) examined the development impacts of the SWP and RSE in sending countries, there remains a lack of empirical research assessing the effects of the RSE and SWP on incumbent farm workers in Australia and New Zealand, which is our contribution here.

The SWP is employer-sponsored, sector-specific, and demand-driven—meaning employers can recruit as many eligible workers as they need, conditional on government approval. Workers are recruited from ten Pacific countries and Timor-Leste and are tied to a single employer, location, and occupation during their stay. To recruit workers, employers must first be vetted and approved as Approved Employers by the Department of Employment and Workplace Relations. Once approved, employers bear unique pastoral care obligations, including covering travel costs, arranging accommodation and transport, and ensuring access to worker protections such as union briefings and Fair Work Ombudsman oversight. These responsibilities exceed those under comparable temporary visa categories, such as the Working Holiday Maker program (Zhao, Binks, Kruger, Xia, and Stenekes, 2018).

Participation in the scheme is highly attractive to workers: income gains relative to home-country earnings (the "place premium") are estimated at between three- and tenfold (depending on the sending country) and sending countries typically all suffer from weak domestic demand and limited formal employment opportunities (Doan, Dornan, and Edwards, 2023). Every participating country has vastly more interested workers than available places, making employer demand the only real constraint on program size. In July 2018, the Pacific Labour Scheme (PLS) was launched as a longer-stay complement to the SWP and the two were merged into the Pacific Australia Labour Mobility (PALM) scheme in 2020, with separate short (SWP) and long (PLS) streams. However, our focus here is the SWP and the introduction of the PLS, the PALM merger, and COVID-19 pandemic border closures and policy changes would further complicate our empirical comparisons, so we focus on the 2009–18 period. Throughout the paper, our data are reported by financial year, where, for example, we refer to fiscal year 2009–10 as 2009 but it covers July 2009 to June 2010.

Figure 1: Annual SWP Visas

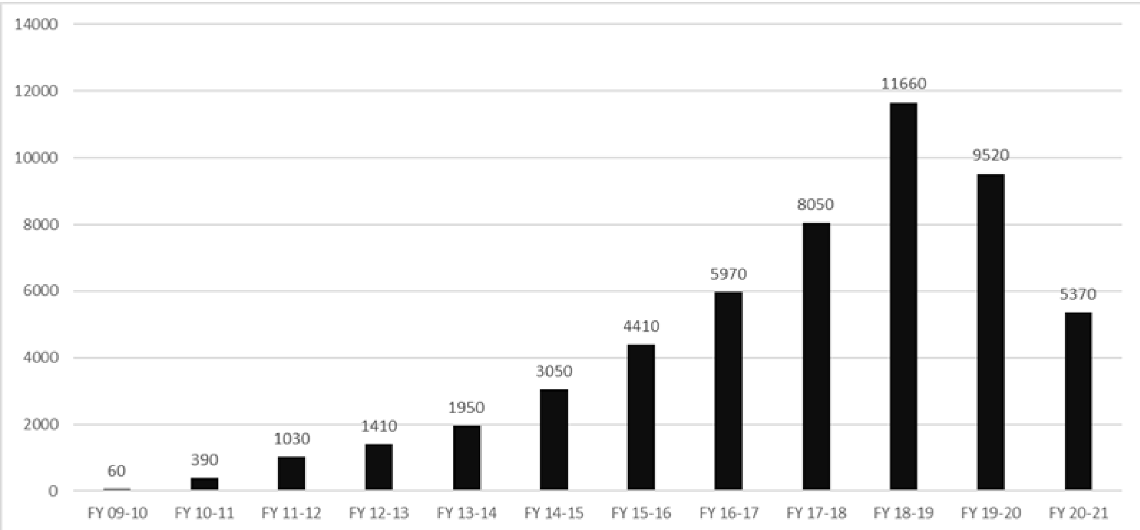

Notes: Data are drawn from the ABS PLIDA Visa module, provided by the Department of Home Affairs, and represent the number of visas issued in each fiscal year.

From its initial pilot through to COVID-19 pandemic disruptions, the SWP experienced slow but steady growth over the 2010s. Figure 1 shows the number of SWP visas granted each year from 2009 to 2020. Only 60 visas were issued in 2009, during the pilot period, but after the formal introduction of the SWP in July 2012, visa numbers jumped to 1,950 in 2013 and increased steadily to 11,660 visas in 2018. With the onset of COVID-19 in January 2020 and major internal and international travel restrictions (borders closed in April 2020), 9,520 SWP visas were issued in FY2019–20 and 5,370 in FY2020–21. However, many SWP workers already in the country had their visas extended due to border closures and food supply concerns, indeed these were among the only new travellers allowed, so in these years visas issued do not accurately reflect the numbers of active SWP workers in Australia.[2]

Figure 2: SWP Visas By Occupation

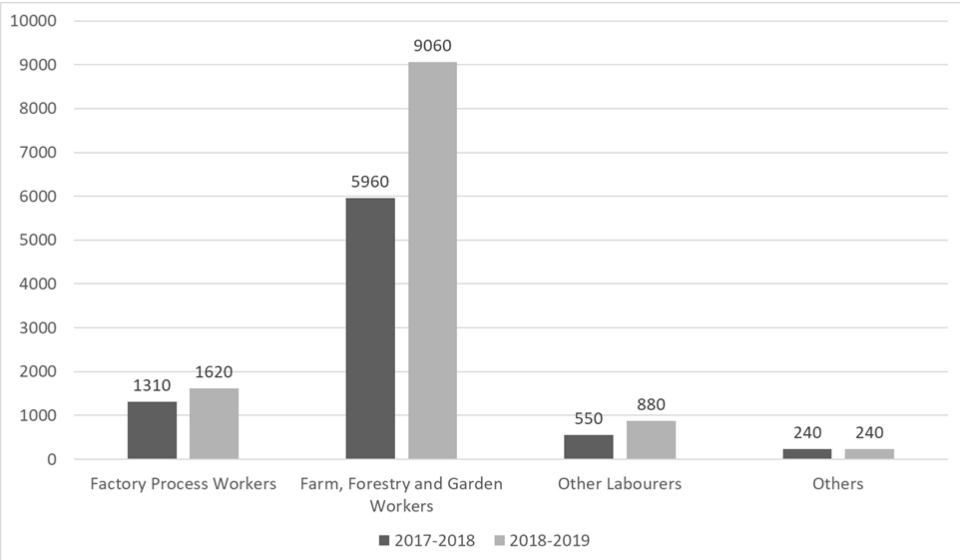

Notes: Data are drawn from the ABS PLIDA Visa module, provided by the Department of Home Affairs. Visa and Tax modules do not include occupational information prior to November 2017 for SWP workers. We assume the occupational composition of SWP visa holders remained relatively stable between 2012 and 2017, which is consistent with qualitative evidence gathered from SWP employers and other sources.

[2]PALM-short (formerly SWP) numbers have recently stablised, with just under 15,000 short-stay workers in June 2025, according to the Australian Department of Employment and Workplace Relations.

Figure 2 plots the number of SWP visas issued in 2017 and 2018 at the Australian and New Zealand Standard Classification of Occupations (ANZSCO) two digit level) (ABS, 2023). Farm, Forestry, and Garden Workers received by far the largest share in 2017 and 2018, surging from 5,960 to 9,060. By contrast, Factory Process Workers received 1,310 visas in 2017 and 1,620 in 2018, and Other Labourers were granted 550 and 880 visas.[3] The “Others” category remained constant at 240 each year. Overall, Figure 2 highlights the central role of agriculture and related sectors in driving demand for seasonal migrant labour under the SWP, and the relative stability of SWP worker occupation between years.

Statistical Area Level 4 (SA4) regions are our main geographical unit of interest in this paper. Australia is divided into 108 SA4 regions, together encompassing all of Australia without gap or overlap (ABS, 2024a).[4] SA4s typically have between 100,000 and 500,000 people, and are designed to reflect the demand (where people reside) and supply (where they are employed) side of local labour markets (Deutscher, 2020) (Kucuk and Ulubasoglu, 2024).

Figure 3 plots the absolute number of SWP visa holders in each SA4 in 2017−18. Some SA4s hosted over 1,000 SWP workers while others had fewer than 10. However, any potential impacts on incumbent workers depend on both the size of the local workforce and the scale of SWP participation in the area.

[3]Factory processing workers here are almost all working in agricultural processing, slightly after the harvest phase which is picked up in the farm, forestry, and garden workers occupation code.

[4]The 108 SA4s includes 19 special-purpose codes with few residents.

Figure 3: SWP Worker Count By SA4

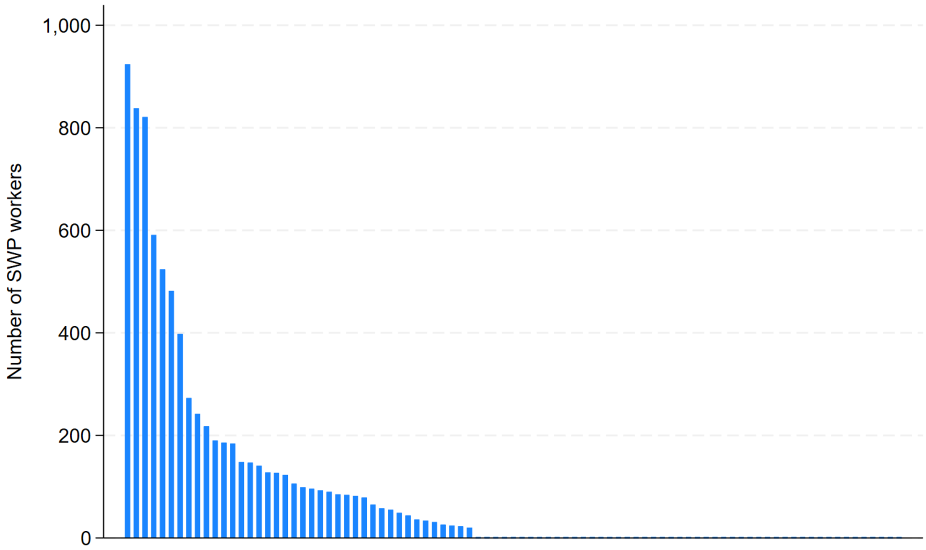

Notes: Data are drawn from the ABS PLIDA Visa module, provided by the Department of Home Affairs, and SA4 boundaries taken from the ABS. The vertical axis represents the number of SWP visa in 2018 domiciled in each SA4, and each vertical bar is one SA4. All SA4s are included, the the right side of the graph, with zero, are typically urban SA4s, which we exclude from the subsequent analysis.

Figure 4: SWP Worker Share By SA4

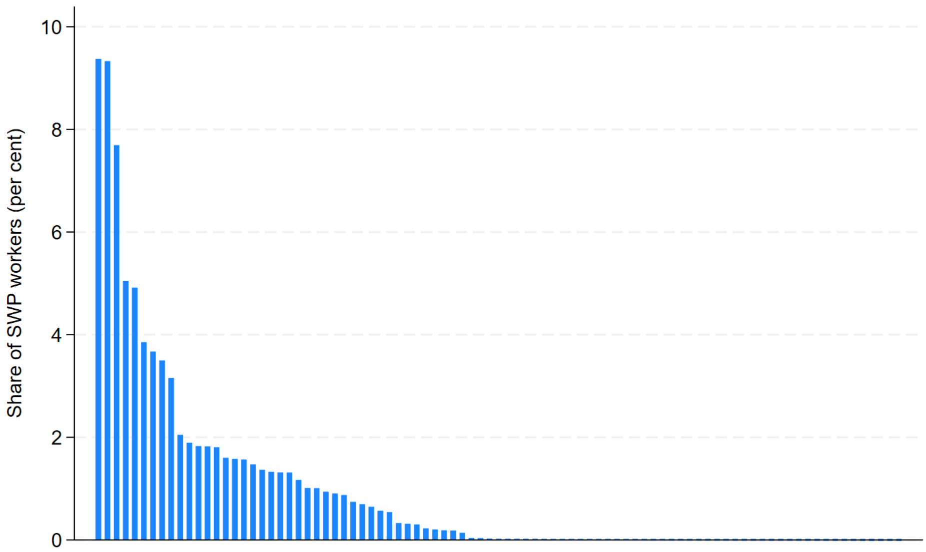

Notes: Data are drawn from the ABS PLIDA Visa module, provided by the Department of Home Affairs, and SA4 boundaries taken from the ABS. The vertical axis represents the number of SWP visa in 2018 domiciled in each SA4 divided by the number of labourers in the same SA4 at baseline, specifically in the 2011 census. Each vertical bar is one SA4. All SA4s are included, the the right side of the graph, with zero, are typically urban SA4s, which we exclude from the subsequent analysis.

Figure 4 expresses these numbers as a share of workers classified as ANZSCO 1-digit code "8 Labourers" in the baseline period, 2011, in each SA4.[5] We focus on this group because they are most likely to be substitutable with SWP workers. Figure 4 again shows a highly skewed distributions: while most SA4s had low SWP intensity, SWP workers account for more than 7% of the labourer workforce in three key growing regions.

Figure 5: Spatial Distribution of SWP Intensity

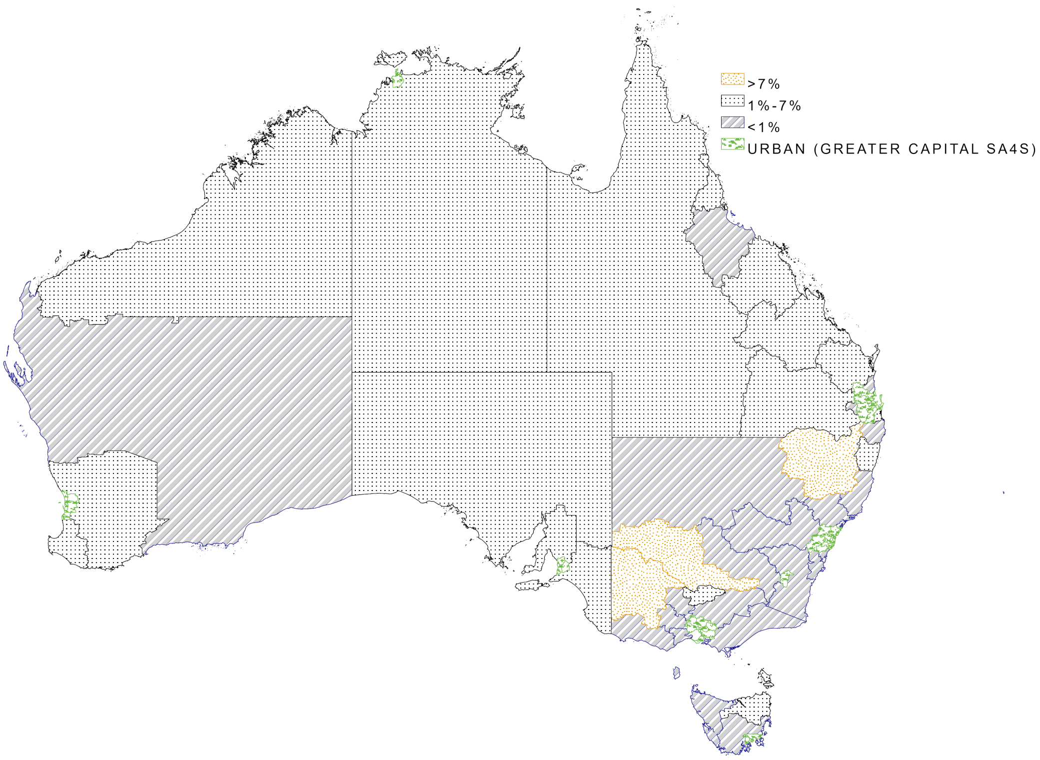

Notes: Data are drawn from the ABS PLIDA Visa module, provided by the Department of Home Affairs, and SA4 boundaries taken from the ABS. SA4s are categorised based on the number of SWP visa in 2018 domiciled in each SA4 divided by the number of labourers in the same SA4 at baseline, specifically in the 2011 census.

[5]We use 2011 as the reference year for measuring the size of the local "labourer" workforce because it predates the establishment of the SWP in 2012. According to ABS (2023), "labourers perform a variety of routine and repetitive physical tasks using hand and power tools, and machines either as an individual or as part of a team". Throughout this paper, we use the terms "local workers," "local labourers," "incumbent workers," "workers," and "labourers" interchangeably to refer to this group.

Figure 5 shows the spatial distribution SWP exposure intensity across Australia's highly heterogeneous SA4 regions. SWP visas are mostly used in rural regions with highly seasonal agricultural sectors, and not in metropolitan regions or the dry desert (outback) centre. The three regions where SWP workers accounted for more than 7 percent of the local labourer workforce in 2017–18 (one in Northern New South Wales and two along the western portion of the New South Wales-Victoria border) are particularly focused on horticulture and supply a significant share of Australian fruits and vegetables.

Table 1: Pre-SWP Baseline Characteristics

	Low Exposure	High Exposure
Average Age (years)	39.50	40.04
Average Population Size (persons)	197,081	144,645
Density (persons per km^2)	5.17	1.58
Female (%)	50.78	50.55
Bachelor Degree or higher (%)	11.82	8.30
Agriculture, forestry and fishing employment (%)	4.52	14.67
Employed Full-time (%)	35.46	35.99
Employed Part-time (%)	19.05	17.95
Unemployed (%)	3.47	3.32
Not in Labour Forces (%)	38.17	38.93
Total Temporary visa (persons)	2,790	1,803
Temporary Work (Skilled) visa (persons)	516	364
Working Holiday Maker (WHM) visa (persons)	320	488
Number of SA4s	23	3

Notes: Data are sourced from Australian Population Census 2011, except for data on temporary visas in the bottom panel which are sourced from Australian Population Census 2016. Regions with low SWP exposure are defined as those where SWP workers account for less than 1% of local labourers, while regions with high SWP exposure are those where SWP workers account for at least 7% of the local labour force.

Table 1 presents a comparison of key demographic and labour market characteristics between high-(over 7%) and low-exposure (under 1%) SA4s using pre-SWP data from the 2011 Australian Population Census, and the earliest available On average, high-exposure SA4s have a slightly older population, with a mean age of 40.04 years compared to 39.50 years in the low-exposure group. They also have an average population size of approximately 144,645 and a population density of just 1.58 persons per square kilometre, compared to 197,081 people and 5.17 persons per square kilometre in low-exposure SA4s. Women make up roughly half of the population in each, but only 8.3% of the population in high-exposure SA4s hold a Bachelor's degree or higher, compared to 11.82% in low-exposure areas. High-exposure regions have a substantially higher proportion of their workforce engaged in agriculture, forestry, and fishing (14.67% versus 4.52%), reflecting their agrarian economic base.

Crucially, Table 1 also shows that labour force indicators were relatively similar across the two groups at baseline. The proportion of people employed full-time is marginally higher in high-exposure areas (35.99%) than in low-exposure areas (35.46%), while part-time employment is slightly lower (17.95% vs 19.05%). Unemployment rates are, on average, 3.32% in high-exposure regions and 3.47% in low-exposure regions and the share of the population not in the labour force is around 38–39% for both groups.

The bottom panel of Table 1 shows the extent of temporary visa holders in both groups of SA4s. Low-exposure SA4s had a total of 2,790 temporary visa holders in 2016, compared to 1,803 in high-exposure SA4s.[6] Within this group, 516 people in low-exposure SA4s held Temporary Work (Skilled) visas, slightly higher than the 364 in high-exposure areas. The number of Working Holiday Maker (WHM) visa holders was higher in high-exposure SA4s (488) than in low-exposure regions (320), reflecting their greater reliance on agriculture. The WHM visa serves as an alternative source of temporary farm labour to the SWP, although it mostly is covers other purposes (e.g. traveling and cultural exchange) so is not our focus here.

[6]Temporary visa data are only available from the 2016 census products.

Figure 6: Weekly Income of Labourers

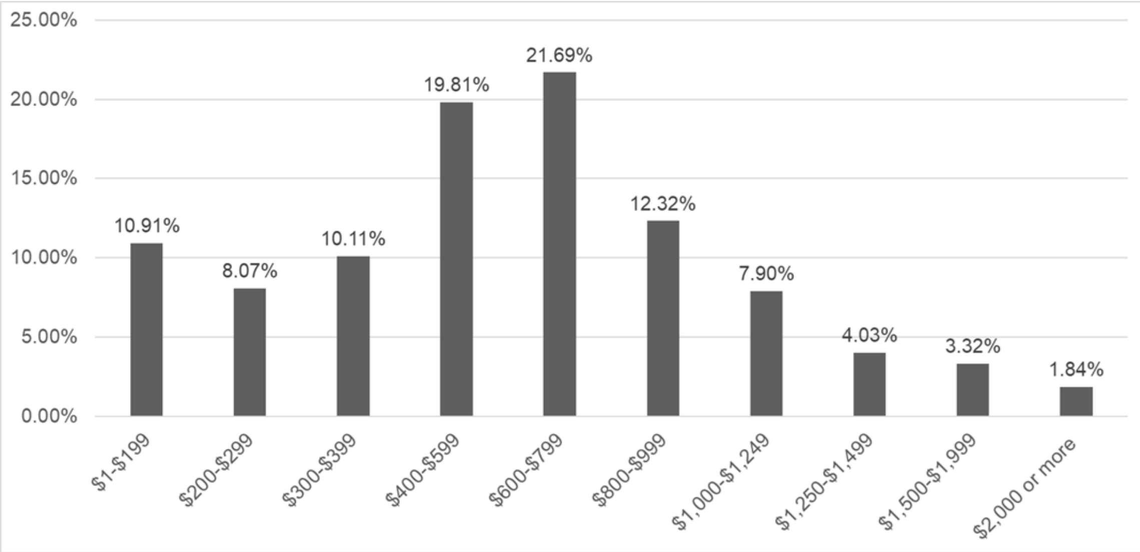

Notes: Data are sourced from the 2011 Australian Population Census and cover all individuals classified as labourers (ANZSCO 1-digit code 8). Weekly income is self-reported and includes both employees and the self-employed. The full-time minimum wage at the time (approximately $589 per week) falls within the $600–$799 bracket, which contains the largest share of workers. Those earning below this threshold—nearly half the sample—likely include part-time, casual, or informally employed workers. This distribution reflects variation in hours worked, employment arrangements, and compliance with labour regulations, all of which affect total earnings outcomes.

One potential concern, when studying wage impacts in a setting like Australia with industry-specific wage agreements, is that there may be little in "Labourers" incomes to exploit, since they mostly earn around minimum wages. Figure 6 presents the distribution of weekly income for workers classified as "Labourers" (ANZSCO 1-digit code 8) based on data from the 2011 Census. The chart reveals significant variation in income levels. The largest share of labourers (21.69%) earn between $600 and $799 per week—the band where full-time minimum wage workers sit—followed closely by 19.81% earning $400 to $599. Smaller proportions earn either lower or higher amounts, with 10.91% earning as little as $1–$199 per week and 1.84% earning $2,000 or more. This variation among labourers earnings likely stems from significant differences in working hours and employment arrangements (piece rates), and our focus on total earnings in tax returns captures total effects inclusive of these margins.

2.2 Estimating the labour market impacts of immigration

To examine how immigration affects the wages and employment of native workers, economists have often utilized three main empirical strategies: regional comparisons, skill-cell models, and individual-level panel designs. Each approach captures different margins of adjustment and involves important trade-offs (Dustmann et al., 2016). Regional (or spatial correlation) approaches compare labour market outcomes across cities/regions with varying immigrant inflows, capturing total effects at the regional level and including local capital and labour supply responses (Card, 1990; Altonji and Card, 1991; Dustmann, Fabbri, and Preston, 2005; Dustmann et al., 2017, 2013; Peri and Yasenov, 2019). These models are vulnerable to bias if immigrants disproportionately settle in booming areas, or if native workers respond by migrating elsewhere.

Skill-cell models, by contrast, define labour markets based on education and experience groups at the national level (Borjas, 2003; Aydemir and Borjas, 2007; Llull, 2018b; Breunig, Deutscher, and To, 2017). This mitigates regional sorting concerns, but assumes fixed labour supply within each skill cell—often overlooking native workers' ability to upgrade skills, switch occupations, or adjust participation. This approach estimates immigration's impact by comparing wage changes between inexperienced and experienced workers across skill levels, but its interpretation becomes problematic when wage rigidities differ between groups. Both regional and skill-cell models can miss important dynamic responses by incumbent workers.

Some studies combine these methods. For example, Card (2001), Glitz (2012), and Monras (2020), by comparing wage changes across cells defined jointly by education and region. This mixed approach often yields less negative wage effects of immigration on native workers than purely national skill-cell models. Other studies adopt structural approaches, such as (Ottaviano and Peri, 2012; Manacorda et al., 2012; Piyapromdee, 2021; Llull, 2018b), which estimate economic model parameters to predict immigration's wage effects. While these studies generally find modest positive impacts on native wages and wage losses for earlier immigrants, their results rely on assumptions about the substitutability between immigrants and natives within skill groups. Such assumptions can be biased if immigrant downgrading and misclassifications occur, potentially distorting estimates of both relative and total wage effects. Both of these approaches, as difference-in-difference type exercises of a “reduced form” nature, consider only relative differences and impacts, assume no spillovers between the designated groups and preclude any conclusions on aggregate or general equilibrium impacts beyond the region or skill-cell of interest.

Recent advances in data availability have also enabled a shift toward individual-level panel designs (Dustmann et al., 2016). These approaches, using linked administrative data, allow researchers to control for unobserved individual heterogeneity and directly track native behavioural responses—such as migration, changes in labour force status, and occupational mobility. The use of population rather than sample data also allows for greater precision in estimation, reducing concerns over statistical power and the treatment of different analytical units (e.g. weighting regions or skill-cells). A recent example of this nature is Foged and Peri (2016). By following individuals, they see that low-skilled natives tend to upgrade their occupations when exposed to immigration, rather than experience wage suppression, which may in turn be masked by design choices. Here, our preferred approach combines the strengths of several of these methods, exploiting differences in earnings growth trajectories between similar occupations, in the spirit of a skill-cell approach, across different regions, in the spirit of regional approaches, while leveraging our linked administrative data to (a) account for any individual-specific unobservable heterogeneity and (b) precisely track the earnings of long-term labourers in exposed regions, who are likely to be some of the lowest paid and least mobile workers in the Australian economy.

3 Empirical strategy

3.1 Data

This paper draws on PLIDA, the Australian Government's central platform for linking and accessing individual-level administrative data, formerly known as the Multi-Agency Data Integration Platform (MADIP). PLIDA integrates information from core government agencies and provides longitudinal data on the entire Australian population from 2000 onward. For our purposes here, it includes visa, travel, income and tax records, employment information, and demographic data (including census records) over time (ABS, 2024b).

We construct our main analytical dataset in three main steps. Our first step is collecting the universe of the Australian Taxation Office's (ATO) Individual Tax Returns from fiscal year 2009–10 to 2018–19, which allows us to track the annual earnings, location, and occupation (at the 2-digit ANZSCO level) of every taxpayer in Australia in every year. This person level unbalanced (i.e. when someone does not file a tax return) panel contains our main outcomes of interest, allows us to track individuals over time and exploit individual fixed effects, and is the base dataset we build on and narrow down for our analysis. We use data only from the 2009–19 period for three reasons. First, gathering data from 2009 allows us to examine pre-trends, and provides a sufficient window following the Global Financial Crisis. Second, we take 2018–19 as our final year to exclude pandemic disruptions and recent changes to the program, allowing us to assess medium-term impacts and adjustment over a period of relative policy and economic stability (i.e. normal times and typical labour market conditions). Third, computational tractability.

Our second step involves incorporating a special Visa and Travellers' Module, from the Department of Home Affairs (which provides detailed information on SWP visa holders), constructing precise measures of SWP exposure for each taxpayer, and refining our comparisons group in important ways. The SWP is an employer-sponsored visa which ties workers to a single Approved Employer, fixing their occupation and location. This stability means workers generally cannot self sort across locations and roles, and allows us to precisely define treatment exposure at the region-by-occupation cell level.[7]

For regional exposure, SA4s are classified as highly exposed (i.e. treated) if SWP workers account for at least 7% of labourers in that SA4 in the baseline period (3 SA4s) and unexposed if SWP workers make up less than 1% of the labourer workforce (22 SA4s).[8] We follow a "donut hole" approach and exclude (a)"low dosage" SA4s where SWP workers represent between 1–7% of labourers (18 SA4s) to make treatment-control comparisons clearer and any potential impacts easier to discern (see Figure 5), and (b) the 45 metropolitan SA4s.[9]

For occupational exposure, we define exposed occupations as those falling under ANZSCO 2-digit code 84 (Farm, Forestry, and Garden Workers) (ABS, 2023) and occupation 83 (Factory Process Workers) and use the other labourer occupations as our comparison group of similar low-skilled but less exposed occupations. These other labourer occupations within the ANZSCO 1-digit code 8 category, are 81 Cleaners and Laundry Workers, 82 Construction and Mining Labourers, 83 Factory Process Workers, and 85 Fast Food Cooks and Food Preparation Assistants. All other occupations are excluded and we are left with the population of labourers in 25 SA4s, three being the most exposed to the SWP and the rest virtually unexposed.

Our third and final step in constructing the dataset used for our main estimates is keeping only data from two years: 2010–11 (pre-treatment) and 2017–18 (post-treatment). The two-period panel and resulting "long difference" approach allows us to identify medium-term impacts after the demand-driven program had scaled up to its pre-pandemic peak and the scale of participation in local labour markets was economically significant enough to plausibly detect any potential effects.[10]

[7]For labour hire companies that move SWP workers between sites, these locations are generally within the same region, and we are able to identify their location in the data.

[8]Kucuk and Ulubasoglu (2024) use similar data and regional treatment classifications, with four SA4s in the city of Brisbane.

[9]Metropolitan SA4s are excluded because (a) the SWP was designed for regional areas and had minimal uptake in cities, (b) including urban areas, where more people live, would disproportionately weight city labour markets and dilute the relevance of the more appropriate comparison group and, most importantly, (c) cities and rural areas have fundamentally different labour markets, dynamics, and secular trends.

[10]Recall from Section 2 that uptake was limited immediately after introduction, and only by fiscal year 2017–18 did SWP participation become meaningfully large.

3.2 Estimation and identification

Our triple-difference design leverages the fact that SWP worker occupations and locations are typically fixed over time. By exploiting variation over *time* (pre- and post-SWP expansion), *region* (SA4s with high vs. low SWP intensity), and *occupation* (SWP-occupations vs. other low-skill jobs), we are able to precisely identify the impacts of the SWP on those incumbent low-skilled workers most likely to be affected most exposed incumbent workers relative to other low-skilled workers in the same area. Our main specification is:

$$\begin{aligned} Y_{iojt} &= \delta_i + \mu_o + \gamma_j + \lambda_t + \beta \cdot \text{Treat}_j \cdot \text{Treat}_o \cdot \text{Post}_t \\ &\quad + \alpha_1 \cdot \text{Treat}_j \cdot \text{Treat}_o + \alpha_2 \cdot \text{Treat}_o \cdot \text{Post}_t + \alpha_3 \cdot \text{Treat}_j \cdot \text{Post}_t + X_{it} + \epsilon_{iojt} \end{aligned} \tag{1}$$

where Y_{iojt} denotes *log real annual earnings from wages and salaries*, as reported in tax returns, for individual (labourer) i, in occupation o, residing in SA4 region j, in year t. Treat_j,Treat_o, and Post_t are dichotomous variables equal to 1 for (a) labourers in high SWP intensity regions (0 otherwise), (b) labourers working in SWP occupations, and (c) fiscal year 2017–18 (reflecting the two-period long difference setup and capturing common secular trends), and pairwise interaction terms (which flexibly control for two-way differences over time between occupations, regions, and treatment statuses) for each are included throughout. δ_i, μ_o, and γ_j denote individual, occupation, and region fixed effects, accounting for time-invariant worker heterogeneity, occupation-specific earning levels, and location-specific unobservables. X_{it} includes age (years) and age squared. Standard errors are adjusted for arbitrary heterogeneity and clustered at the *individual* level throughout.

The coefficient of interest, β on the triple interaction ($\text{Treat}_j \cdot \text{Treat}_o \cdot \text{Post}_t$), offers a causal interpretation if there are no time-varying omitted variables within SA4s systematically affecting the differences in wage trends between SWP versus other labourers in treatment rather than control SA4s (or, control rather than treatment, which leads to the same problem). While we have no reason to believe this to be the case, individual fixed effects throughout capture any further unobservable heterogeneity across individuals, and we complement our main triple difference estimates with more general difference-in-difference approaches. These approaches, while standard workhorse models in the literature, have less credible identification assumptions but allow for potential adjustments across occupations (within SA4s) and locations (within occupations) to give a broader picture.

A few further comments on interpretation are in order. First, our long difference approach means we are not estimating short-term impacts, but medium-term cumulative impacts accumulating between the announcement of the SWP in 2012 and its pre-pandemic peak in 2018–19 in the same regions. Second, our treatment definitions mean we are focusing on the most extreme cases of potential direct impacts: incumbent workers in the same locality, working in the same jobs, in the most exposed localities. Finally, likely any difference in difference type approach, our focus on impacts relative to similar workers in the same locality, assuming no spillovers across locations or localities, precludes any conclusions on aggregate or general equilibrium impacts.

Figure 7: Labourers' Earnings Trends, By Region Treatment Status

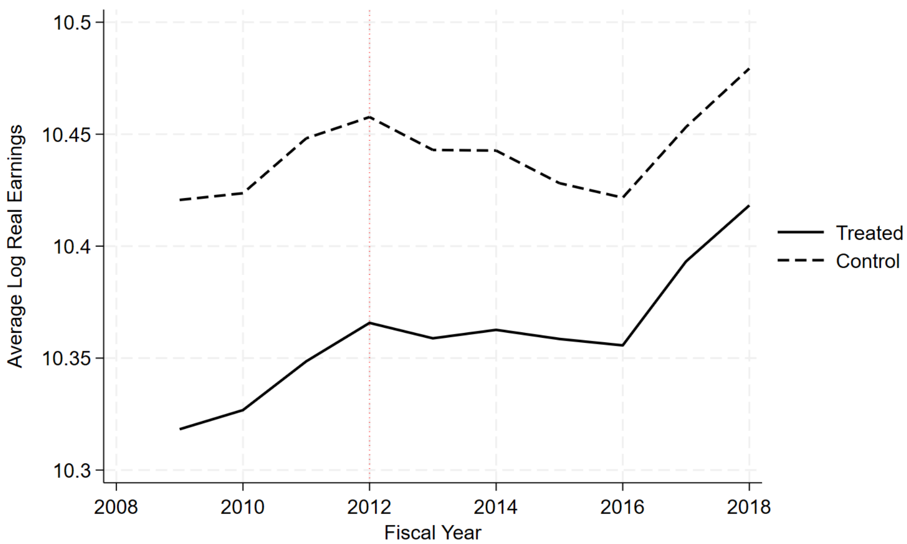

Notes: This figure plots annual average worker earnings for all labourers (i.e. an unbalanced panel) in study regions split by those in treatment and control regions. Greater capital regions and those with 1–7% of their local labourer workforce comprised of SWP workers in 2018 are excluded throughout. The vertical dotted red line indicates the establishment of the SWP. Data are calculated from ATO Individual Tax Returns in PLIDA and the DataLab.

3.3 Visual evidence

Our primary outcome of interest is annual pre-tax earnings from wages and salaries.[11] Figure 7 plots the evolution of log real earnings in constant 2011 Australian dollars (i.e. with 2011 as the base year) for labourers in treated and control SA4s from fiscal year 2009 to fiscal year 2018. Treated areas, those relatively more exposed to the SWP over time, consistently exhibit lower log real earnings than control areas. From 2009–11, log real earnings increased modestly from 10.32 to 10.35 in treated SA4s and from 10.42 to 10.45 in control SA4s. In other words, both groups followed similar trends and differences were small, 0.1 in each period, before the introduction of the SWP in 2012. The SWP period (2012–18) reveals slightly faster wage growth in SA4 regions and a further narrowing of the earnings gap.[12]

Figure 8: Labourers' Earnings Trends, By Occupational Treatment Status

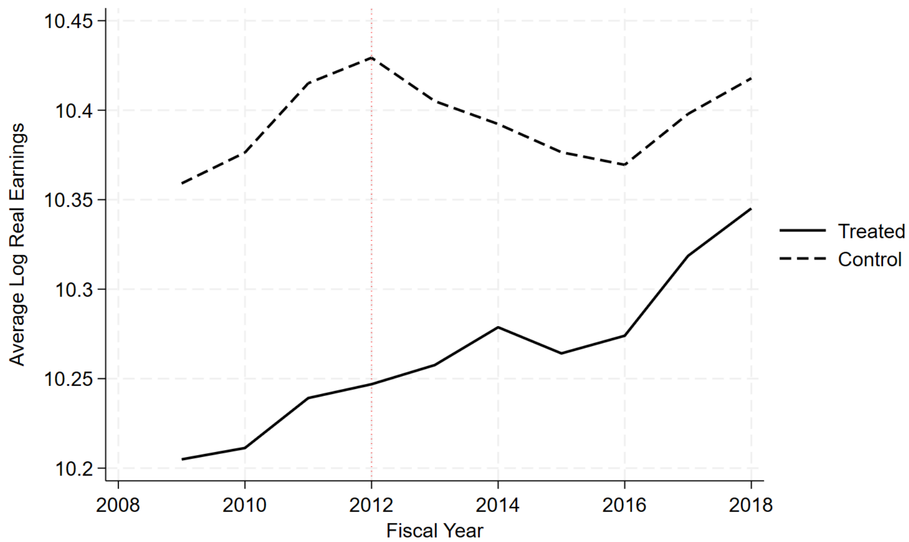

Notes: This figure plots annual average worker earnings for all labourers (i.e. an unbalanced panel) nationally split by those in treatment ANSCO code 84, Farm, Forestry and Garden Workers) and control (ANZSCO groups 81 Cleaners and Laundry Workers, 82 Construction and Mining Labourers, and 85 Fast Food Cooks and Food Preparation Assistants) occupations, excluding those in greater capital regions. The vertical dotted red line indicates the establishment of the SWP. Data are calculated from ATO Individual Tax Returns in PLIDA and the DataLab.

[11] Earnings are measured before the deduction of taxes. We use the terms wages and earnings interchangeably throughout.

[12] We show log earnings as this is our primary outcome of interest, but in 2009, labourers in treated SA4s earned an average of approximately $30,280, compared to $33,545 in control SA4s. By 2018, these had risen to $33,464 and $35,574.

Figure 8 shows the same by occupational grouping, where treated workers are Farm, Forestry and Garden Workers (ANZSCO 84) and all remaining labourers, except those in greater capital regions, form the control (81 Cleaners and Laundry Workers, 82 Construction and Mining Labourers and 85 Fast Food Cooks and Food Preparation Assistants). Again, pre-period earnings for the two groups are laregly parallel and the differences small (around 0.15), rising modestly from 10.20 in 2009 to 10.24 in 2011 in treated occupations and from 10.36 to 10.42 in control occupations. The level differences likely capture differences in occupation-specific characteristics or labour demand, although the number of workers in two groups also followed similar trends.[13] After the parallel trends in the pre-treatment period, mean earnings in SWP sectors, at the national-level, experienced slightly stronger earnings growth while other labourers' earnings fell slightly in real terms. These patterns motivate combining both types of exposure to focus on occupation differences within regions.

4 Main results on earnings

Our main result is that the introduction and expansion of the SWP does not appear to have any statistically discernible impact on the earnings of local, low-skilled workers in the most affected occupations and regions. This overall null finding of no adverse impacts for incumbent labourers is robust across a range of triple difference specifications, as well as less demanding regional and national skill cell difference-in-difference designs.

Table 2 presents the main triple difference results, comparing labourers in SWP sectors against similar labourers in the same locality. Our preferred specification is the most saturated specification using a perfectly balanced panel in Column 6, but we present estimates without such conditioning and with all observations first for completeness.

[13]The number of workers in treated occupations remained relatively stable over this period, ranging from 64,726 in 2009 to 66,272 in 2011, while control occupations employed around 158,000 to 162,000 workers.

Column 1 of Table 2 begins with the baseline specification and every tax-paying labourer in study regions. Time, SA4 and occupational fixed effects are included throughout. The coefficient of interest, on the triple interaction term ($\mathtt{Treat}_o \cdot \mathtt{Treat}_j \cdot \mathtt{Post}_t$), is negative, small (-0.02), and not statistically different from zero. Column 2 adds age and aged squared as covariates and the point estimate is virtually unchanged. Using the full unbalanced panel means these estimates capture not only labourers who remain in the sample across the two years but also new entrants such as recent migrants, younger workers, or workers switching into the ANZSCO 1 digit 8 sector, as well as leavers who exit the sample by 2017, including retirees, workers who migrate out of the region, or those who change industries. Such net effects capture the direct impact of SWP exposure on incumbents plus any indirect effects through changes in workforce composition, such as selective in-migration of higher-earning workers or the exit of lower-earning workers.

Column 3 of Table 2 restricts the estimation sample to the balanced panel of labourers who remain in the same occupations and localities from 2010–17. As the least mobile and most-exposed group, these are precisely the incumbent workers we are most interested in (cf. those who move to opportunity or higher-paying occupations, or are only in the sector transiently in their youth). Here, the coefficient flips sign to 0.012 and the standard error is almost twice the size (expected given the smaller group). Column 4 adds controls and the point estimate almost halves. Column 5 replaces the controls with individual fixed effects, removing all time-invariant sources of heterogeneity across individuals, and the coefficient is virtually identical to that in Column 3 with the balanced panel and no controls, except more precisely estimated. The variation absorbed by the individual FEs reduces precision but is unlikely to bias point estimates. Our preferred specification in Column 6 adds the two age covariates back in to the balanced panel specification with all fixed effects, finding a remarkably small coefficient of 0.0018 which, again, is not statistically discernible from zero.

Table 2: Impact of the SWP on Log Real Earnings—Triple Differences

	Log of Real Annual Earnings					
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathrm{Treat}_o \cdot \mathrm{Treat}_j \cdot \mathrm{Post}_t$	-0.0208	-0.0244	0.0123	0.0068	0.0137	0.0018
	(0.024)	(0.024)	(0.040)	(0.040)	(0.030)	(0.030)
$\mathrm{Treat}_o \cdot \mathrm{Treat}_j$	0.0973	0.0981	0.0516	0.0584	-0.0022	-0.0074
	(0.018)	(0.018)	(0.029)	(0.028)	(0.059)	(0.058)
$\mathrm{Treat}_o \cdot \mathrm{Post}_t$	0.0890	0.0931	0.0227	-0.0046	0.0144	-0.0299
	(0.010)	(0.010)	(0.017)	(0.017)	(0.013)	(0.013)
$\mathrm{Treat}_j \cdot \mathrm{Post}_t$	0.0310	0.0369	0.0027	0.0069	0.0059	0.0153
	(0.014)	(0.014)	(0.022)	(0.022)	(0.017)	(0.017)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
SA4 FE	Yes	Yes	Yes	Yes	Yes	Yes
Occupational FE	Yes	Yes	Yes	Yes	Yes	Yes
Individual FE	No	No	No	No	Yes	Yes
Additional controls	No	Yes	No	Yes	No	Yes
Balanced Panel	No	No	Yes	Yes	Yes	Yes
Observations	335,721	335,721	80,674	80,674	80,674	80,674

Notes: Dependent variable is log of real annual earnings from wages and salaries. A two period panel is used throughout: 2010–11 and 2017–18. Columns (1)–(2) use an unbalanced panel; Columns (3)–(6) use a balanced panel. All regressions include year, SA4, and 2-digit occupation fixed effects. Additional controls (where included) are age and age squared. The data are restricted to individuals classified as labourers (ANZSCO 1-digit code = 8). Standard errors are clustered at the individual level. *** p 0.01, ** p 0.05, * p 0.1.

As with any null results, the crucial question is then what size of an effect, over the 2010–17 period, we can confidently rule out. With the standard error between 0.024 and 0.04, the minimum detectable effect we can confidently rule out is between 0.067 and 0.11, which annualised over 8 years is around 1 percent faster wage growth in SWP occupations vis a vis other similar occupations. In the preferred specification (Column 6), we can confidently rule out the possibility that any direct impact, positive or negative, on workers' earnings was cumulatively more than 8 percent (or more than one percent per year) since the establishment and growth of the SWP. While these minimum detectable effects are generally smaller than negative effects found in some earlier studies, we are unable to rule out smaller effects.

4.1 Robustness of the main result

The 7% threshold used to define SWP-exposed regions in our main analysis is somewhat arbitrary, capturing a smaller group of the most exposed workers rather than a larger group including slightly less exposed workers (i.e. receiving, on average, a lower "dosage" of the treatment).[14] The exposure to the SWP for labourers in occupation 83 (Factory Process Workers) may also be limited, since most SWP participants are farm workers.

To assess the sensitivity of our results to these two important design choices, in Table 3 we present estimates excluding occupation 83 from the sample and varying the treatment threshold for exposed regions.[15] All columns use the main triple difference specification and data in the two years 2010 and 2017. Columns 1 and 2 of Table 3 show the specifications in Columns 5 and 6 of Table 2 except excluding factory process workers. The point estimates are of the same sign and slightly larger, but statistically insignificant at conventional levels. Columns 3 and 4 use a 5% threshold to define treated regions and Columns 5 and 6 adopt a more inclusive 3% cutoff. All results are consistent with our main null results, and notably

[14]Clearly, if there is no evidence of any impact in the regions with the highest intensity, there is no reason to expect to see any effect as exposure is effectively diluted. Rather, this exercise more explores whether it is these particular SA4s driving the null rather than including others which might be considered appropriate under alternative treatment definitions or with a view to use more of the available data.

[15]We also conduct the analysis without excluding occupation 83. Results are similar to those in Table 3 and excluded for brevity.

the coefficient on the triple interaction term becomes much smaller in the most expansive treatment definition, with or without covariates. Together, these estimates suggest that our main null result is unlikely to be driven by neither the occupation sample or the high exposure threshold used to select the treatment regions.

Table 3: Impacts on Log Real Earnings—Triple Difference Sensitivity Analysis

	Threshold = 7%		Threshold = 5%		Threshold = 3%	
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathsf{Treat}_o \cdot \mathsf{Treat}_j \cdot \mathsf{Post}_t$	0.0293	0.0135	0.0362	0.0180	0.0016	0.0033
	(0.032)	(0.031)	(0.029)	(0.028)	(0.023)	(0.023)
$\mathsf{Treat}_o \cdot \mathsf{Treat}_j$	-0.0510	-0.0578	-0.0705	-0.0713	0.0139	0.0086
	(0.072)	(0.071)	(0.067)	(0.066)	(0.052)	(0.051)
$\mathsf{Treat}_o \cdot \mathsf{Post}_t$	0.0227*	-0.0272**	0.0229*	-0.0271**	0.0254*	-0.0246*
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
$\mathsf{Treat}_j \cdot \mathsf{Post}_t$	-0.0063	0.0074	-0.0194	-0.0063	0.0123	0.0085
	(0.020)	(0.019)	(0.017)	(0.017)	(0.012)	(0.012)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
SA4 FE	Yes	Yes	Yes	Yes	Yes	Yes
Occupational FE	Yes	Yes	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Additional controls	No	Yes	No	Yes	No	Yes
Observations	58,546	58,546	61,150	61,150	75,204	75,204

Notes: Dependent variable is log of real annual wage and salary income. A two period balanced of labourers is used throughout, 2010–11 and 2017–18. Columns (1)–(2) use a 7% threshold to define treated regions; Columns (3)–(4) use a 5% threshold; Columns (5)–(6) use a 3% threshold. All regressions include year, SA4, and 2-digit occupation fixed effects. Additional controls (where included) are age, age squared, and occupation dummies. Standard errors are clustered at the individual level. *** p 0.01, ** p 0.05, * p 0.1.

4.2 Alternative approaches

4.2.1 Regional difference in differences

While our main triple difference approach allows precise targeting of potentially affected groups and arguably the most credible identification assumption, it assumes workers are immobile across narrowly defined job categories. Comparisons using such fixed "cells" of incumbent workers may be misleading if workers respond to immigration by changing occupations or industries (Dustmann et al., 2016). To address this concern, we estimate a simplified difference-in-differences specification based only on regional exposure (i.e. removing the occupation dimension). This approach allows for the possibility that incumbent labourers may have adjusted to the SWP not by exiting the labour market or suffering wage losses, but by shifting occupations (within the same region), and captures aggregate net impacts at the regional level, regardless of whether these adjustments occurred within or across occupations. This approach rests on a standard parallel trends assumption (recall the parallel pre-trends in Section 2) and no time-varying omitted variables systematically correlated with changes labourers earnings trends and SWP uptake.

Table 4 presents the medium-term impacts, from a simple 2x2 difference in difference approach from 2010–11 to 2017–18, of the SWP on average labourer earnings in highly exposed SA4s relative to other rural SA4s unexposed to the program. Across all specifications, we again find no evidence of any adverse impacts on incumbent labourer wages. Rather, including the subset of workers who are constantly in the data from start to finish, for example because they enter (e.g. young people) or exit the labour force or move to other regions, reveals small positive wage impacts driven by this transient group, something worth exploring in future research.

Table 4: Impact of the SWP on Log Real Earnings—Regional Difference in Differences

	Log of Real Annual Wage and Salary Income								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\text{Treat}_j \cdot \text{Post}_t$	0.0404***	0.0384***	0.0350***	0.0039	0.0055	0.0045	0.0153	0.0091	0.0091
	(0.010)	(0.009)	(0.010)	(0.015)	(0.015)	(0.015)	(0.011)	(0.011)	(0.011)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SA4 FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Individual FE	No	No	No	No	No	No	Yes	Yes	Yes
Additional controls	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
State $\times$ Time FE	No	No	Yes	No	No	Yes	No	No	Yes
Balanced Panel	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	427,272	427,272	427,272	116,574	116,574	116,574	116,574	116,574	116,574

Notes: Each column presents a separate regression on log of real annual earnings. The treatment indicator equals 1 for individuals residing in SA4s defined as SWP-exposed, and 0 otherwise. A two period panel is used throughout, where the post-treatment period is defined as 2017—18 and the pre-treatment period is 2010–11. The data are restricted to individuals classified as labourers (ANZSCO 1-digit code = 8). Additional controls refer to age and age squared. Standard errors clustered at the SA4 level in parentheses. *** p 0.01, ** p 0.05, * p 0.1.

Table 4 again begins with the full, unbalanced panel of all labourers observed in either the baseline (2010) or endline year (2017), regardless of whether they appear in both years. Column 1 is a simple two-by-two (treated-untreated and pre-post) long difference estimate, with only the treatment and "post" indicators and no additional covariates. Column 1 shows that average labourer earnings in SWP regions grew around 4 percent faster from 2010–11 to 2017–18. The coefficient is precisely estimated and statistically significant at conventional levels. Column 2 adds individual covariates (age and age squared) and Column 3 add state-by-time fixed effects, allowing different state trends and restricting our comparisons to across treatment and control SA4s within each state. The coefficients are slightly smaller, at 0.038 and 0.035, and statistically significant at conventional levels.

Columns 4–9 of Table 4 focus on the balanced panel of individuals observed in both 2010 and 2017, which we are able to precisely track over time and include individual fixed effects to control for time-invariant heterogeneity across labourers. Columns 4–6 present results without individual fixed effects but progressively adding individual covariates then state-by-year fixed effects. The estimated treatment effect is small, between 0.004 and 0.006, and standard error also small at 0.015 throughout. Columns 7–9 present balanced panel estimates including individual fixed effects. Column 7 includes no covariates and finds a point estimate of 0.015, a very small magnitude over eight years. Column 8 adds individual covariates and our preferred saturated specification in Column 9 also includes state-by-time fixed effects, capturing differential regional trends driven by state policies or macroeconomic shocks. Adding covariates reduces the magnitude slightly but the estimates with and without the state-by-time fixed effects are similar (both 0.0091 and 0.011) and not statistically distinguishable from zero.

4.2.2 Regional event study estimates

The key concern with the estimates presented in Table 4 is that treated areas may have been on different wage trajectories before the SWP. To more formally examine these pre-period differences and trace earnings dynamics over time, we extend our main two period panel to include annual data each year from 2009 to 2018 and estimate the following event study version of the regional approach:

$$Y_{ijt} = \delta_i + \gamma_j + \lambda_t + \sum_{k=2009, k \neq 2011}^{2018} \beta_k \cdot \text{Treat}_j \cdot \text{Year}_t + X_{it} + \epsilon_{ijt} \quad (2)$$

Here, β_k captures the period-specific differences in outcomes between treated and control regions in year k, relative to the omitted base year (2011) and conditional on individual covariates (X_{it}) and individual (δ_i), region (γ_j), and year fixed effects λ_t. Like the two-period regional approach, the identifying assumption is that treated and control regions would have experienced parallel trends in earnings in the absence of the SWP.

Figure 9 plots the estimated β_k coefficients. The pre-treatment period shows no discernible differences, with all coefficients close to zero and statistically insignificant at conventional levels (the blue line indicates the 95 percent confidence interval). Throughout the post-treatment period, the point estimates remain small and, with the exception of 2016–17, not statistically different from zero. The isolate and modest 2 percent increase relative to the base year observed in 2016–17, while statistically different from zero, is not statistically different from the other point estimates. Overall, that all estimates are centred near zero and have overlapping confidence intervals provides some additional reassurance that our main null results are unlikely to be an artifact of any underlying pre-trends.

Figure 9: Dynamic Impacts—Regional Event Study

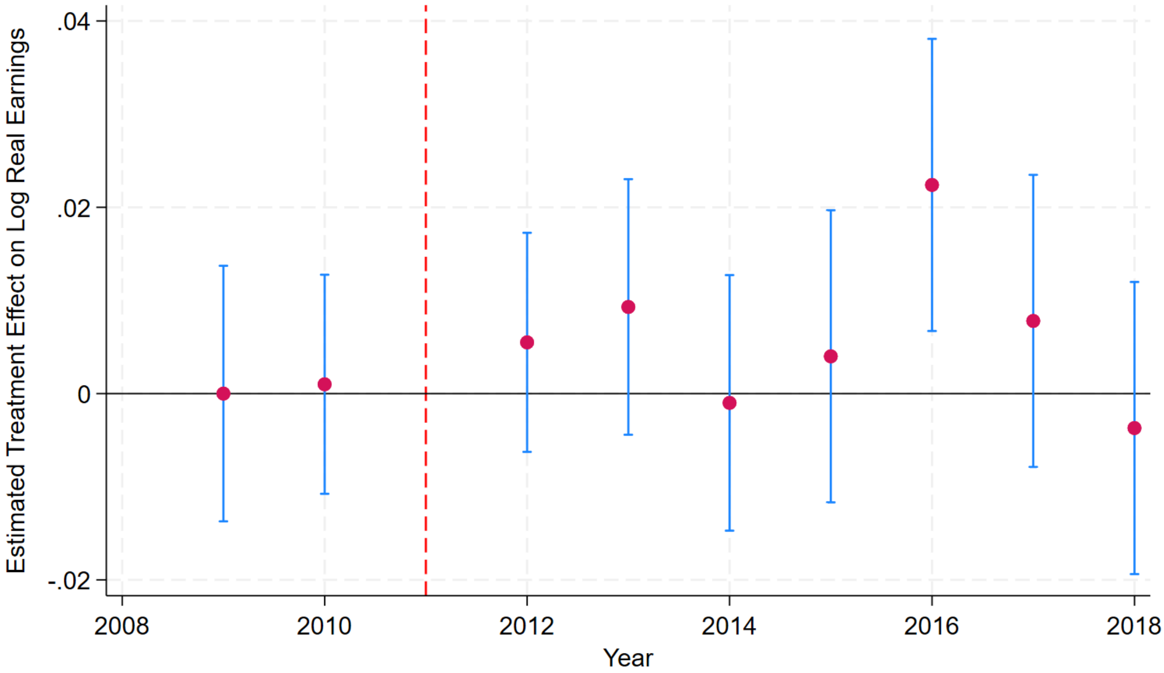

Notes: The Figure plots the year-specific treatment effects from an event study including individual, year, and region, fixed effects and age covariates and using annual data from the balanced panel of labourers (ANZSCO 1-digit code = 8), in each year from 2009 to 2018, in only treatment and control SA4s. Vertical blue lines indicate 95 percent confidence intervals and the vertical dotted redline indicates the announcement of the SWP.

4.2.3 National skill-cell difference in differences

Just as the regional approach allows people to adjust to immigration by changing occupations within regions, exploiting only the occupational dimension of the data allows people to adjust across locations—a crucial adjustment mechanism, which in settings of relatively flexible labour markets, is well known to bias spatial estimates downwards and understate the true effects of immigration (Borjas, 2003). Such "skill-cell" approaches (a) define labour markets by national-level skill groups, such as occupation, rather than geographic location, (b) better reflect competition in national labour markets, where such arbitrage is less feasible, and (c) capture the common native-immigrant substitution that tends to occur within these skill cells across locations (Dustmann et al., 2016).[16]

[16] This logic is echoed in structural and reduced-form models by Llull (2018b,a), who show how equilibrium wage effects of immigration depend on within-group elasticities of substitution and aggregate demand responses. Ottaviano and Peri (2012) and Manacorda et al. (2012) show that immigration affects the wage structure primarily within narrowly defined skill (occupation) cells, not necessarily across regions.

We compare earnings in the labourer occupation most exposed to the SWP against similar labourer occupations with the equation:

$$Y_{iot} = \delta_i + \mu_o + \lambda_t + \beta \cdot \text{Treat}_o \cdot \text{Post}_t + X_{it} + \epsilon_{iot} \tag{3}$$

where Treat_o is equal to one for workers whose ANZSCO 2-digit code is 84 (Farm, Forestry and Garden Workers) and 0 for remaining labourers, and the remainder is as in Equation 1. The estimation sample is a two period panel (2010 and 2017), excluding SWP workers and workers in greater capital regions, and including only individuals classified as labourers (ANZSCO code 1 digit 8).[17] The coefficient of interest β captures the effect of the SWP on average worker earnings in SWP occupations nationally, relative to workers in similar low-skilled occupations in rural areas, under the identification assumption that the two occupational groupings would have remained on parallel trends but for the introduction of the SWP.

Column 1 of Table 5 begins with the simplest version of Equation 3, using the unbalanced panel of all labourers and including no additional covariates, finding that average farm, forestry, and garden labourer earnings grew over ten percent faster than the earnings of labourers in regional areas. Column 2 adds age and age squared as covariates, reducing the point estimate slightly. Column 3 restricts the balanced panel of labourers present in both years. The sign reverses and the point estimate is close to zero and statistically insignificant. Column 4 swaps the two control variables for individual fixed effects.[18] The sign returns to positive but the point estimate is small, at 0.007, and statistically insignificant at conventional levels. The null findings for the balanced panel of long-term labourers starkly contrast to the wage growth for workers not part of this group, likely reflecting these workers greater mobility, while highlighting the importance of accounting for individual-level heterogeneity and using a

[17] Individuals in occupation code ANSCO 83 (Factory Process Workers) have limited exposure so are excluded.

[18] As younger workers are more likely to be employed in SWP-exposed occupations and typically experience faster wage growth due to life-cycle effects, controlling for age and age squared may mechanically absorb some of the post-treatment wage dynamics, biasing the treatment estimate downward, especially in individual fixed effects regressions. We omit X_{it} when including individual fixed effects to capture within-worker changes while avoiding potential collinearity introduced by age adjustments.

balanced panel to obtain more reliable estimates. Nonetheless, taken together, a national skill cell approach clearly provides no convincing evidence of any adverse impacts for workers in SWP occupations.

Table 5: Impacts of the SWP on Log Real Earnings—Occupation Difference in Differences

	Log of Real Earnings			
	(1)	(2)	(3)	(4)
$\mathrm{Treat}_o \cdot \mathrm{Post}_t$	0.1069***	0.0987***	-0.0029	0.0066
	(0.005)	(0.005)	(0.009)	(0.007)
Time FE	Yes	Yes	Yes	Yes
Occupational FE	Yes	Yes	Yes	Yes
Individual FE	No	No	No	Yes
Additional Controls	No	Yes	Yes	No
Balanced Panel	No	No	Yes	Yes
Observations	1,398,875	1,398,875	320,550	320,550
R^2	0.085	0.133	0.128	0.714

Notes: This table presents DD estimates of the impact of SWP exposure on log real annual earnings, using an occupational-level specification. The variable $\mathrm{Treat}_o \cdot \mathrm{Post}_t$ equals 1 for individuals in treated occupations (ANZSCO 84) in 2017. Additional controls (where included) comprise age and age squared. The balanced panel specification restricts the sample to individuals observed in both 2009 and 2017. Standard errors are clustered at the individual level. *** p 0.01, ** p 0.05, * p 0.1.

5 Explaining the overall null result

The remainder of the paper explores two potential explanations for the overall null impact we observe of the SWP on incumbent labourer earnings. First, we examine extensive margin adjustment—whether overall employment declined to accommodate the SWP workers—using the regional difference in difference design. Second, with evidence of any potential employment impacts in hand, we review the long-term trends in domestic labour supply in these sectors and regions.

5.1 Impacts on local employment

We explore whether local labour markets instead adjust on the extensive margin of employment by examining the relationship between the SWP and the probability of employment for residents in exposed SA4s. Conventional employment status is not observed in our tax record data, so we define employment as a binary indicator equal to one if any wage or salary income is reported in a given year.[19] As people not reporting income do not have an observable occupation, we broaden our focus beyond only the labourers subset of taxpayers to all individuals who have ever filed a tax return in our study regions.

Figure 10 plots the proportion of taxpayers reporting no wage or salary income in regions most exposed to the SWP against relatively unexposed regions. The levels are similar, at under 16 percent in the control group and over 17 percent in the treated group.[20] The ratios track each other remarkably closely in the pre-SWP period (from 2009–12), providing no evidence of any problematic pre-trends for a difference in difference comparison.[21]

[19] The key limitation here, however, is that individuals not reporting wage or salary income, regardless of their actual economic activity, appear as “unemployed” in our data.

[20] Note the y axis starting at 0.13 makes these differences look larger than they are, and that a one percent difference in this ratio likely translates to a much smaller difference in the unemployment rate estimated using traditional approaches.

[21] Between 2009 and 2011, the share of non-wage taxpayers in treated regions declined from 17.3% to 15.9%, while in control regions it fell from 15.6% to 14.5%.

Figure 10: Trends in Zero-Wage Income Reporting Among Taxpayers

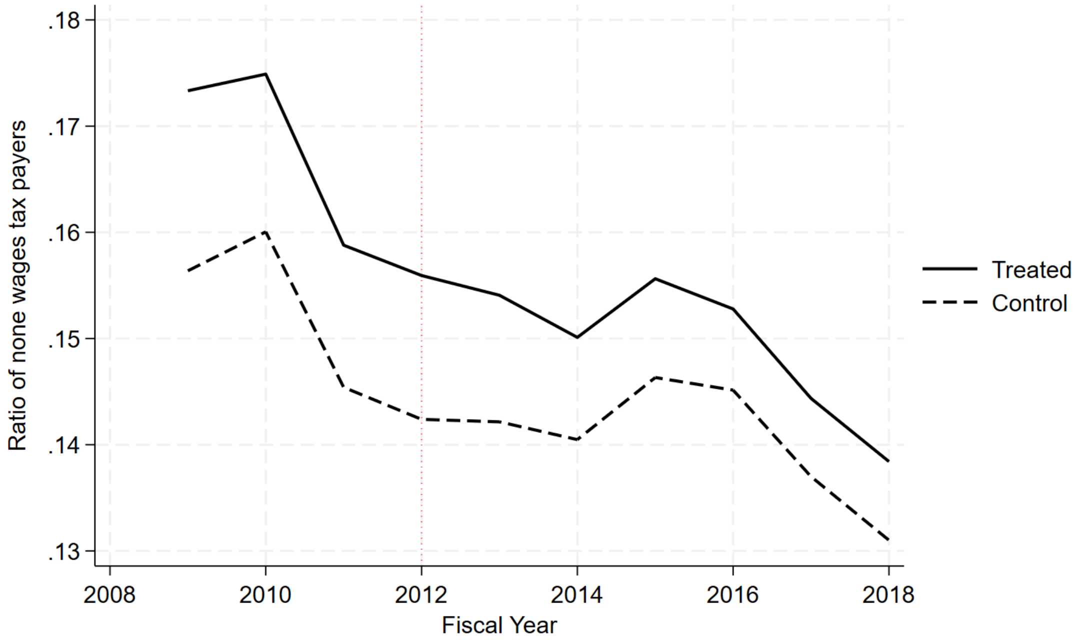

Notes: We excludes greater capital regions. Data are sourced from individual tax returns. Ratio of none wages tax payers = none wages tax payers/(total none wages tax payers + tax payers who record positive wages and salaries)

To more formally test for any differences during the SWP period, relate regional employment to SWP exposure with the long difference specification:

$$Y_{ijt} = \delta_i + \gamma_j + \lambda_t + \beta \cdot \mathrm{Treat}_j \cdot \mathrm{Post}_t + X_{it} + \epsilon_{ijt} \tag{4}$$

where Y_{ijt} is the binary employment indicator for individual i in region j and year t, Treat_j and Post_t are defined as same as in Equation 1, X_{it} includes age and age squared, standard errors are clustered at the *individual* level, and estimation is by logit.

Table 6 presents the employment results, reported as log-odds rather than changes in probability. Column 1 begins with all individual taxpayers. When converted to average marginal effects, the Column 1 estimate corresponds to roughly a 0.22 percentage point increase in employment probability in treated SA4s, assuming a mean employment rate of 90%, an economically small but statistically significant effect consistent with the trends in Figure 10. Column 2 uses a balanced panel of individuals in tax data in both years, effectively removing the oldest and the youngest workers. Here, the coefficient flips in sign, is small and not statistically distinguishable from zero. Column 3 adds individual fixed effects, which removes people whose employment status does not change and captures worker-specific heterogeneity. The estimated coefficient is negative at -0.03 but not statistically significant at conventional levels. The point estimate in Column 3, however, suggests that on the margin, among the movers only, people tend to slightly more likely be moving into unemployment. Overall, the three specifications do not provide any strong evidence that domestic workers in SWP regions adjusted to the SWP its null wage effects through dis-employment.

Table 6: Impacts on Regional Employment

	(1)	(2)	(3)
$\mathsf{Treat}_j \cdot \mathsf{Post}_t$	-0.0322	-0.0117	0.0247**
	(0.019)	(0.013)	(0.010)
Year FE	Yes	Yes	Yes
SA4 FE	Yes	Yes	Yes
Individual FE	Yes	No	No
Balanced panel	Yes	Yes	No
Observations	268,216	2,081,246	3,775,297

Notes: The dependent variable is a binary indicator equal to 1 if the individual had any wage/salary income. All Columns show results from logit models. Standard errors are clustered at the individual level. *** p 0.01, ** p 0.05.

5.2 Domestic labour supply in rural regions

Our analysis so far indicates that the introduction of the SWP does not appear to have had any adverse effect on incumbent farm workers' earnings, and that this effect is not explained by dis-employment in SWP regions. How did the labour market absorb SWP workers? Temporary seasonal work visas, like the SWP, aim to fill quite narrow labour market gaps, gaps believed to exist because native labour supply for these particular jobs is low (i.e. persistent labour shortages in specific regions and sectors of the economy). The SWP is an uncapped, employer-sponsored program, meaning that employers recruit workers based on their own labour demands. If employers primarily use the program to fill genuine labour gaps, rather than replace existing workers, then impacts on both earnings and employment for domestic workers would be limited.

Figure 11: Initial Agricultural Employment and SWP Intensity

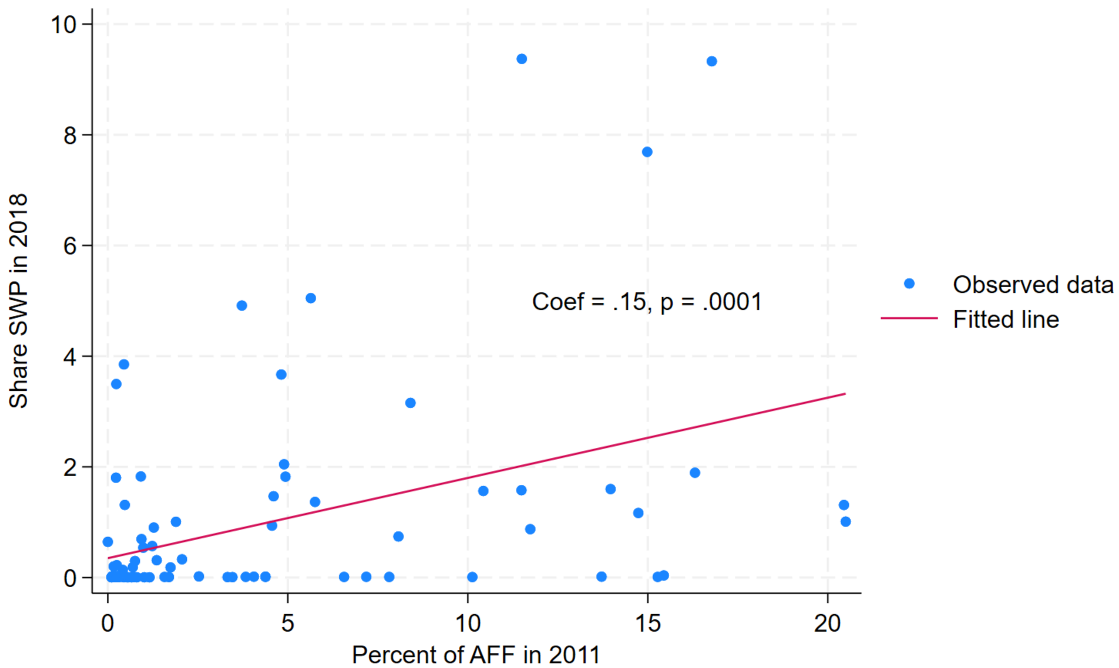

Notes: Data are sourced from Home Affairs Visa data (ABS, 2024b) and 2011 Census (ABS, 2025a)

To explore this possibility, we first look in more detail (than in Figure 5) at how SWP intensity is distributed across regions. Figure 11 plots, by SA4 area, using all SA4s nationwide, the percent of the local labour force employed in Agriculture, Forestry, and Fishing (AFF) at baseline (before the SWP, using data from the 2011 Census) against the SWP share of the local labourer workforce (measured as a proportion of all 2011 ANZSCO 1-digit code "8 Labourers" in each SA4 region) in 2018, seven years on. An area with one percent more of its workforce in AFF in 2011 had, on average, an 0.15 percent point higher SWP worker share in 2018. Regions initially more reliant on the AFF employment were also more to have SWP workers constitute more of their local labourer workforce in the following years.

Figure 12: Slower Population Growth in Agricultural Regions

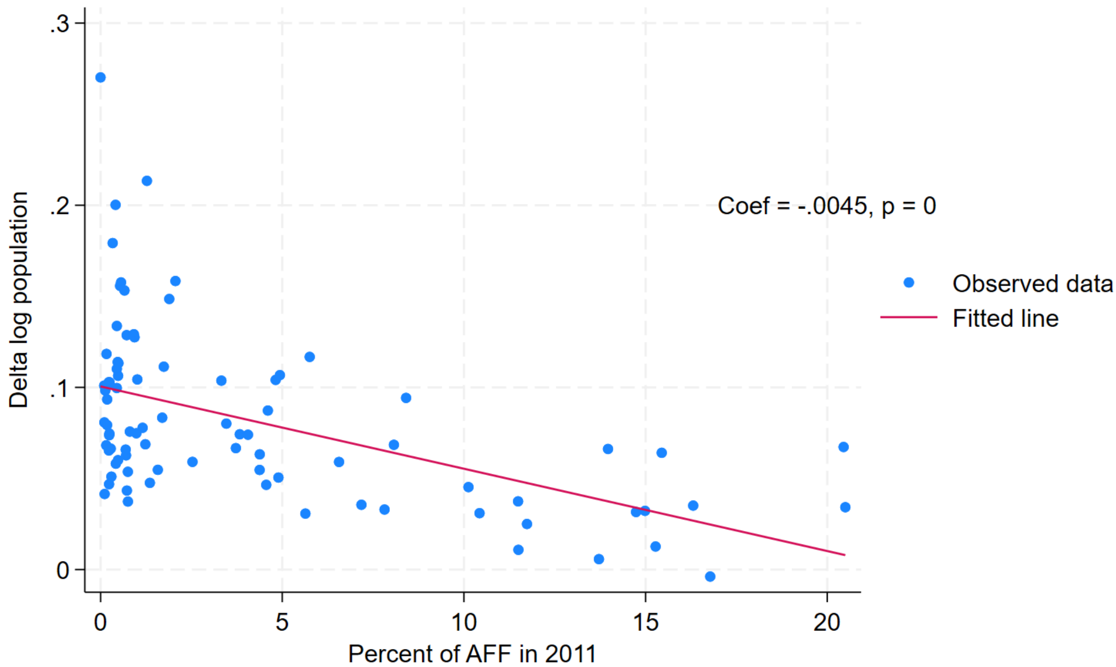

Notes: Data are sourced from ABS 2011 Census Population (ABS, 2025a) and ABS regional population estimates (ABS, 2025b).

Looking back, Figure 11 retains the same X axis and replaces the SWP share on the Y axis with the change in the logarithm of SA4 population, as measured in the the 2006 and 2011 censuses. Here, one percentage point higher reliance on the AFF sector, in 2011, corresponds to approximately 0.45 percent slower population growth in the five years before. All but one SA4 with more than 5 percent of its workforce in agriculture in 2011 sit below 0.1 and areas with the fastest population growth mostly have little to no agricultural workforce, highlighting Australia's long-term urbanisation trend and the slower population growth, from much lower levels, in agriculture-dependent regions.

Table 7: Change in Number of Workers by Occupation Group (2006–2011)

Occupation Group	2006	2011	Change (%)
Farm workers ANZSCO 84	96,090	88,045	−8.37%
Labourers ANZSCO 8	952,524	947,612	−0.52%
All Others	7,986,048	8,921,699	+11.72%

Notes: Data are sourced from the 2006 and 2011 Australian censuses by ABS Table Builder. (ABS, 2025a)

The urbanisation trends and structural changes in employment implied by Figure 12 are also clear in the national level data. Table 7 shows the changes in the occupational composition of the Australian workforce, nationally, between 2006 and 2011. During this period, employment in ANZSCO 84 (Farm Workers), where most SWP workers end up working in the following years, declined by more than 8%. At the same time, total employment in all other occupations increased by nearly 12%, highlighting clear structural change in the domestic labour market: fewer domestic workers are employed as farm workers and areas where there is demand for farm workers have fewer workers. Notably, employment in ANZSCO 8 (Labourers), a broader category that includes both farm workers and other manual or low-skilled occupations, such as construction labourers, cleaners, and factory hands, remained virtually flat, decreasing slightly by 0.52%. Thus, the farm worker (ANZSCO 84) share of the domestic labourer workforce (ANZSCO 8) was in relative decline before the SWP.

This stagnation in the labourer category, alongside a sharp decline in the farm worker subgroup, suggests that agricultural roles were becoming increasingly difficult to fill even as the broader labour market was expanding. These patterns support the view that the SWP distribution was in line with regional labour shortages in agriculture. Rather than substituting existing farm workers, the SWP supplemented a shrinking workforce in sectors where domestic supply was failing to meet employer demand. This interpretation is consistent with Duncan et al. (2020), who find that employer-sponsored visa programs can help address regional labour shortages, particularly those arising from commodity-driven economic booms.

6 Conclusion

In this paper, we measured the impact of Australia's SWP, a key program designed to address labour shortages in regional and agricultural sectors, on the earnings of incumbent farm workers. To our knowledge, this is the first quantitative assessment of any seasonal visa program globally on native wages. Using rich, linked administrative data covering the universe of taxpayers and visa holders, we isolate the effect of the program on the most exposed domestic workers with a triple difference design exploiting variation in SWP exposure across regions, occupations and time. We complement these precise comparisons with more relaxed difference in difference approaches providing relative impacts at the regional level, or at the national level across occupations.

Overall, we find no compelling evidence that the SWP reduced incumbent farm worker earnings, either using our preferred precise triple difference comparisons within localities, or broader comparisons across regions or occupations. We also find no evidence that SWP workers and these null earnings impacts were absorbed through dis-employment in host communities, and instead interpret our findings in the context of the long-term structural change and increased urbanisation. In the decade leading up to the introduction of the SWP, fewer domestic workers were employed as farm workers, and areas where there was demand for farm workers progressively had relatively fewer workers. Together, it appears the SWP has been serving as a complementary mechanism to support industries and regions facing persistent challenges in attracting labourers, especially those that only need workers on demand and for part of the year.

We offer three main caveats on our analysis. First, our analysis relies on tax declaration data submitted by individuals to the government. While this provides comprehensive population coverage, the data contain only annual income figures and do not include information on hours worked, hourly wage rates, and whether workers were paid fixed wages or piece rates. Naturally, the data do not capture unreported cash-in-hand payments and other forms of non-compliance or illegality. Thus, we are unable to determine whether changes in earnings are driven by changes in hourly wages, hours worked, or other margins. For example, the absence of any observed SWP effect on earnings could reflect a situation where all workers are paid at or near the minimum wage, with variation arising mainly from hours worked, a dimension we cannot evaluate given data constraints.

Second, the SWP is a demand-driven visa program, whose uptake is clearly non-random among participating employers and across regions or occupations. While the SWP was not motivated by initial labour market conditions (cf. a smaller development initiative), its announcement largely unexpected, and participants cannot selective sort across locations or sectors, omitted variable bias cannot be ruled out, as is the case with any non-randomised study. Future research exploiting genuinely random variation, as in Clemens (2024), to understand study these issues and adjustment mechanisms would be valuable.

Third and finally, the SWP began as a relatively small program and today its equivalent, the PALM scheme, still accounts for only around 1 percent of temporary visa holders in Australia. It may be the case that the absence of any systematic impacts here could be related to the program not being a large labour supply shock, as in other studies, but rather a purposeful program that grew slowly enough such that marginal adjustments occurred in real time, year on year, in ways masked by our long difference and not discernible in the event study due to the relatively low "dosage" in initial years. Our analysis of the impacts of the SWP should not be conflated with potential impacts from other types of immigration, certainly not larger inflows, those that are not initiated by the employer based in their specific needs, those concentrated in cities, and so forth, although these are important areas for future research and rich administrative data we have introduced in this paper provides plentiful opportunities of this nature.

References

ABS (2023). Anzsco, australian and new zealand standard classification of occupations. https://www.abs.gov.au/ausstats/abs@.nsf/0/8B1F5DDDD46033ABCA2575DF002DA75E?opendocument. Accessed: May 17, 2023.

ABS (2024a). Abs asgs. https://www.abs.gov.au/statistics/standards/australian-statistical-geography-standard-asgs-edition-3/jul2021-jun2026/main-structure-and-greater-capital-city-statistical-areas/statistical-area-level-4. Accessed: May 17, 2023.

ABS (2024b). Person level integrated data asset. https://www.abs.gov.au/about/data-services/data-integration/integrated-data/person-level-integrated-data-asset-plida. Accessed: May 17, 2024.

ABS (2025a). Abs census of population and housing. https://www.abs.gov.au/census. Accessed: May 17, 2023.

ABS (2025b). Abs regional population. https://www.abs.gov.au/statistics/people/population/regional-population. Accessed: May 17, 2023.

Altonji, J. G. and D. Card (1991). The effects of immigration on the labor market outcomes of less-skilled natives. In *Immigration, Trade, and the Labor Market*, pp. 201–234. University of Chicago Press.

Aydemir, A. and G. J. Borjas (2007). Cross-country variation in the impact of international migration: Canada, Mexico, and The United States. *Journal of the European Economic Association 5*(4), 663–708.

Bedford, C. and R. Bedford (2025). New zealands recognised seasonal employer scheme: Pathways and prospects. *Asia & the Pacific Policy Studies 12*(3), e70034.

Borjas, G. J. (2003). The labor demand curve is downward sloping: Reexamining the impact of immigration on the labor market. *the Quarterly Journal of Economics 118*(4), 1335–1374.

Breunig, R., N. Deutscher, and H. T. To (2017). The relationship between immigration to australia and the labour market outcomes of australian-born workers. *Economic Record 93*(301), 255–276.

Card, D. (1990). The impact of the Mariel boatlift on the Miami labor market. *Ilr Review 43*(2), 245–257.

Card, D. (2001). Immigrant inflows, native outflows, and the local labor market impacts of higher immigration. *Journal of Labor Economics 19*(1), 22–64.

Clemens, M. A. (2022). The effect of seasonal work visas on native employment: Evidence from US farm work in the great recession. *Review of International Economics 30*(5), 1348–1374.

Clemens, M. A., E. G. Lewis, and H. M. Postel (2018). Immigration restrictions as active labor market policy: Evidence from the Mexican Bracero exclusion. *American Economic Review 108*(6), 1468–1487.

Curtain, R. (2025). Australias seasonal worker program: Working out ways to manage risk. *Asia & the Pacific Policy Studies 12*(3), e70033.

Curtain, R., M. Dornan, J. Doyle, and S. Howes (2016). Pacific possible: Labour mobility – the ten billion dollar prize. Report, World Bank, Washington, DC.

Deutscher, N. (2020). Place, peers, and the teenage years: long-run neighborhood effects in australia. *American Economic Journal: Applied Economics 12*(2), 220–249.

Doan, D., M. Dornan, and R. Edwards (2023). The gains and pains of working away from home: The case of pacific temporary migrant workers in Australia and New Zealand. *World Bank and Development Policy Centre of the Australian National University, Canberra.*

Doyle, J. and S. Howes (2015). Australia's seasonal worker program: Demand-side constraints and suggested reforms. Technical report, The World Bank.

Duncan, A., M. N. Harris, A. Mavisakalyan, and T. Nguyen (2020). Migration flows in commodity cycles: Assessing the role of migration policies. *European Economic Review 127*, 103458.

Dustmann, C., F. Fabbri, and I. Preston (2005). The impact of immigration on the British labour market. *the Economic Journal 115*(507), F324–F341.

Dustmann, C., T. Frattini, and I. P. Preston (2013). The effect of immigration along the distribution of wages. *Review of Economic Studies 80*(1), 145–173.

Dustmann, C., U. Schönberg, and J. Stuhler (2016). The impact of immigration: Why do studies reach such different results? *Journal of Economic Perspectives 30*(4), 31–56.

Dustmann, C., U. Schönberg, and J. Stuhler (2017). Labor supply shocks, native wages, and the adjustment of local employment. *the Quarterly Journal of Economics 132*(1), 435–483.

Foged, M. and G. Peri (2016). Immigrants' effect on native workers: New analysis on longitudinal data. *American Economic Journal: Applied Economics 8*(2), 1–34.

for Global Development, C. (2025). Seasonal worker programme. https://gsp.cgdev.org/legalpathway/seasonal-worker-programme-swp/. Accessed: May 17, 2025.

Gibson, J. and D. McKenzie (2011). Australia's pswps: Development impacts in the first two years. *Asia Pacific Viewpoint 52*(3), 361–370.

Gibson, J. and D. McKenzie (2014). The development impact of a best practice seasonal worker policy. *Review of Economics and Statistics 96*(2), 229–243.

Glitz, A. (2012). The labor market impact of immigration: A quasi-experiment exploiting immigrant location rules in Germany. *Journal of Labor Economics 30*(1), 175–213.

Kucuk, M. and M. Ulubasoglu (2024). Paying income tax after a natural disaster. *Journal of Environmental Economics and Management 128*, 103044.

Llull, J. (2018a). The effect of immigration on wages: Exploiting exogenous variation at the national level. *Journal of Human Resources 53*(3), 608–662.

Llull, J. (2018b). Immigration, wages, and education: A labour market equilibrium structural model. *the Review of Economic Studies 85*(3), 1852–1896.

Manacorda, M., A. Manning, and J. Wadsworth (2012). The impact of immigration on the structure of wages: theory and evidence from Britain. *Journal of the European Economic Association 10*(1), 120–151.

Martin, P. (2025). Farm guest workers: Us experience. *Asia & the Pacific Policy Studies 12*(3), e70035.

Monras, J. (2020). Immigration and wage dynamics: Evidence from the Mexican Peso crisis. *Journal of Political Economy 128*(8), 3017–3089.

Ottaviano, G. I. and G. Peri (2012). Rethinking the effect of immigration on wages. *Journal of the European Economic Association 10*(1), 152–197.

Parliament (2016). Seasonal change: Inquiry into the seasonal worker programme. https://www.aph.gov.au/Parliamentary_Business/Committees/Joint/Migration/Seasonal_Worker_Programme/Report. Accessed: May 17, 2025.

Peri, G. and V. Yasenov (2019). The labor market effects of a refugee wave: Synthetic control method meets the Mariel boatlift. *Journal of Human Resources 54*(2), 267–309.

Piyapromdee, S. (2021). The impact of immigration on wages, internal migration, and welfare. *the Review of Economic Studies 88*(1), 406–453.

World Bank (2018). Maximizing the development impacts from temporary migration: Recommendations for australia's seasonal worker programme. Report, World Bank, Washington, DC. DOI not available.

Zhao, S., B. Binks, H. Kruger, C. Xia, and N. Stenekes (2018, February). What difference does labour choice make to farm productivity and profitability in the australian horticulture industry? a comparison between seasonal workers and working holiday makers. Research Report 18.1, Australian Bureau of Agricultural and Resource Economics and Sciences (ABARES), Canberra.