The impact of school disruptions during the COVID-19 pandemic on parental labor supply and earnings in Australia

Nicolás Salamanca[♣*], Tanya Gupta[♣], Irma Mooi-Reci[♦] and Mark Wooden[♣]

♣ Melbourne Institute of Applied Economic and Social Research, University of Melbourne, Victoria 3010, Australia.

♦ School of Social and Political Sciences, University of Melbourne, Victoria 3010, Australia.

June 2026

Accepted at the *Journal of Labor Economics*

Abstract

We estimate the impact of COVID-19 school disruptions on parental labor supply in Australia using newly collected data on local school closures, HILDA Survey data, and a difference-in-difference-in-differences design. We find no evidence that disruptions affected labor force participation or work hours among parents relative to childless adults. With 95% confidence, we can rule out reductions in labor force participation exceeding 8% for mothers and 4% for fathers. This null result holds for non-essential workers and with an estimator robust to heterogeneous and dynamic treatment effects. Generous wage subsidies, widespread remote work, and short, localized disruptions may explain our findings.

Keywords: School disruptions; labor force participation; hours worked; wages; triple-difference; COVID-19 pandemic.

JEL Codes: J22, I28, I18

---

[*] Corresponding author: n.salamanca@unimelb.edu.au. This paper uses unit-record data from the Household, Income and Labour Dynamics in Australia (HILDA) Survey. The HILDA Survey was initiated and is funded by the Australian Government Department of Social Services (DSS) and is managed by the Melbourne Institute of Applied Economic and Social Research (Melbourne Institute). The findings and views reported in this paper, however, are those of the authors and should not be attributed to either DSS or the Melbourne Institute. These data are available free of charge to researchers through the National Centre for Longitudinal Data Dataverse at the Australian Data Archive (https://dataverse.ada.edu.au/dataverse/ncld). Access is subject to approval by the DSS and is conditional on signing a license specifying terms of use. This research was supported by a US National Institutes of Health Grant (#R01AG071649: PI Lillard, with a sub-award to the University of Melbourne: PI Salamanca). We thank the Guest Editor, the Associate Editor and a reviewer for helpful comments.

I. Introduction

In response to the COVID-19 outbreak in early 2020, governments in many countries implemented policies intended to contain the spread of the virus and reduce its impact on the health services system. In Australia, governments at both the federal and state levels imposed laws and orders that restricted population movement and reduced inter-personal contact. Among the suite of restrictions implemented were business closures, stay-at-home orders, and the suspension of in-person teaching at schools. It is the latter that is the focus of this study.

In the absence of alternative arrangements and together with increased home-schooling demands, school disruptions may affect labor supply decisions, especially those of primary carers. Some parents may have been forced to reduce their work hours or withdraw from the labor force entirely because of such disruptions. Some of this negative labor supply response, however, will have been offset by government advice directing employers to allow employees to work from home wherever possible, and by increased support via the welfare system.

In Australia prior studies have examined the impact of lockdowns on various outcomes (e.g., Guven et al. 2023; Schurer et al. 2023). Schurer et al. (2023), for example, reported that the 2020 lockdown in Melbourne was associated with a significant reduction in usual weekly work hours using Sydney as comparison. None of these studies, however, attempted to isolate the role of school closures in these effects. In contrast, studies have been conducted in the US that have made use of inter-state differences in the incidence of school closures or in teaching mode. While empirical evidence is somewhat mixed, with studies focused on the first weeks of the pandemic finding little evidence of marked effects on hours worked (e.g., Heggeness 2020), subsequent research shows school disruptions lead to non-negligible reductions in parental labor supply (e.g., Amuedo-Dorantes et al. 2023; Collins et al. 2021; Garcia and Cowan 2024). Furthermore, these negative effects on labor supply have usually been found to be largest for mothers (Amuedo-Dorantes et al. 2023; Collins et al. 2021). Studies that leverage differences across states in the timing of school re-openings in both the US and Canada further support this finding (Hansen et al. 2024, Beauregard et al. 2022). Very differently, Meekes et al. (2023) focused on the difference in labor supply responses of persons employed in sectors defined by the Dutch Government as essential or non-essential, with essential workers permitted to send children to school and childcare services during periods of lockdown while non-essential workers were not. This permitted a difference-in-differences design.

Our study provides the first estimates of the impact of school disruptions (or more strictly, of the impact of the suspension of in-person teaching) on parental labor force participation and hours worked in Australia. We leverage longitudinal data from the Household, Income and Labour Dynamics in Australia (HILDA) Survey covering the period 2017 to 2021, together with uniquely detailed data on pandemic mitigation policies. This period thus includes the COVID-19 pandemic during which residents of some of Australia's largest states — Victoria, New South Wales, and the Australian Capital Territory (ACT) — were subject to highly restrictive lockdowns involving curfews and compulsory stay-at-home orders for all but essential workers. Our identification strategy exploits geographic and temporal variation in school closure mandates across greater capital city statistical areas, the narrowest geographic unit at which meaningful school disruption policies were effectively implemented. Because school disruptions and stay-at-home orders are highly correlated at this level of geographic and temporal aggregation, we use this variation in two complementary triple-difference (DDD) designs. The first compares parents of school-age children to age-matched childless adults. The second complementary design compares parents in households without and with an essential worker, where the latter group's children could continue attending school during closures.

We find no evidence that school disruptions affected parental labor force participation or hours worked. The DDD estimates are small and precisely estimated: 95 percent confidence intervals that, combined with point estimates close to zero, strongly suggest smaller effects than those documented in the US literature. We also find null, though less precisely estimated, wage effects. Our findings hold across an extensive set of subgroups — by child age, gender, marital status, education, partnership status, and closure intensity — indicating that our null main effects are unlikely to be hiding large subgroup heterogeneity.

Interestingly, Australia's proactive implementation of strict international and interstate travel bans early in the pandemic places this identifying variation in a setting with remarkably lower COVID-19 case numbers in 2020 and 2021 compared to other countries, driven by geographically clustered and mostly contained outbreaks. The low case numbers make it unlikely that fear of COVID itself drove labor force exits among parents, reducing concerns that our null results mask compositional changes from virus-related withdrawals.

The Australian null result is substantively important. It suggests that the effects of school disruptions on parental labor supply depend on the institutional context, including generous government support programs such as Australia's JobKeeper wage subsidy, the prevalence of flexible work arrangements, and the duration and geographic concentration of school disruptions. Our findings provide a useful counterpoint to the US evidence and highlight the role that policy can play in mitigating the labor market consequences of pandemic-related disruptions to schooling.

II. The Australian Response to the COVID-19 Pandemic

A. The Public Health Response

Following the arrival of the first cases of COVID-19 infection in Australia in late January 2020, the national government began imposing travel bans that were soon extended to all non-residents and non-citizens. In addition, returning Australian residents were required to enter a 14-day quarantine at a designated facility (mostly hotels). With the notable exception of a period of quarantine-free travel between Australia and New Zealand (from 19 April to 23 July 2021), these restrictions remained in place until the end of October 2021. International borders, however, did not fully re-open until 21 February 2022.

These international border closures were accompanied by internal border closures, with governments of the less populous states seeking to protect their populations from the growing number of infections in the two most populous states, Victoria and New South Wales. Such restrictions did not mean interstate travel was not feasible, but persons living in border zones required permits and negotiating hard border checkpoints, while others had to enter mandatory periods of quarantine (typically of 2 weeks duration) upon arrival at their destination.

Within states, the activity and movement of Australian residents were also heavily restricted. A newly formed 'national cabinet' comprising the Prime Minister and all state and territory leaders agreed, on 15 March 2020, to a ban on non-essential organized public gatherings of more than 500 people. This was followed by ever tighter restrictions culminating on March 24 in the decision to close or heavily restrict the operation of all non-essential businesses. Australia was now effectively in a state of partial lockdown, with residents being advised to: (i) stay at home unless shopping for essentials, exercising, seeking medical care, or traveling to work or education; (ii) work from home wherever possible; (iii) keep visitors to the home to a minimum; and (iv) not congregate in groups when outdoors.

By early May this approach seemed to be working, with new COVID-19 cases numbering fewer than 20 per day. Consequently, the national cabinet announced a plan to ease restrictions, though the rate at which restrictions were eased varied across states (Stobart and Duckett 2022). In late June 2020, however, Melbourne, the largest city in Victoria, became the epicenter of an outbreak that saw a marked increase in the number of COVID-19 cases. This prompted the Victorian state government to re-impose a range of lockdown measures, including business closures, stay-at-home orders, remote schooling, bans on all public gatherings, restricted access to family members living in aged-care facilities, and making the wearing of face masks mandatory whenever outside the home. These restrictions were not mere advice but legally enforceable requirements subject to financial penalties. In Greater Melbourne, these restrictions commenced on 9 July, and stay-at-home orders would remain in place until 27 October. In regional Victoria, stay-at-home orders commenced on 4 August and were lifted on 17 September. This, however, was to be only the first in a series of lockdowns for Victorian residents. Short snap lockdowns commenced on 13 February 2021 (for 5 days), on 28 May (for 7 days) and on 16 July (5 days). Then, on 5 August 2021, Victoria entered another prolonged lockdown that for Melbourne residents did not lift until 21 October.

Elsewhere in Australia, life seemed to be returning to some degree of normality. While many restrictions remained in place — notably capacity limits on public venues, mask-wearing mandates in certain settings, and restrictions on inter-state travel — stay-at-home orders were confined to a few short and often localized episodes (see Online Appendix Table A1).

The notable exceptions here were New South Wales and the ACT, which from the middle of 2021 would, like Victoria, also be the subject of a prolonged lockdown. Following an upsurge in the number of cases of the relatively new Delta variant, mask mandates and a range of restrictions intended to reduce social interactions were imposed on residents and businesses in Greater Sydney and major adjacent regions commencing on 23 June. Three days later, stay-at-home orders were announced. On 14 August this lockdown went state-wide (and also included the ACT), and from 23 August involved a 9 pm evening curfew for many Sydney residents. The state-wide lockdown lasted until 11 September, after which date stay-at-home orders were relaxed for residents living in local government areas without any new cases during the preceding 14 days. It was not until 11 October that lockdowns in New South Wales came to an end, and even then, only for the double vaccinated. The unvaccinated would remain under stay-at-home orders until 15 December.

B. The Economic Policy Response

During the pandemic, Australian governments significantly boosted public spending to ease financial strain on the working-age population. According to Breunig and Sainsbury (2023), approximately 6.5 million Australians, representing 42% of the working-age population, received some form of financial assistance. Chief among these support measures were: (i) the JobKeeper program, a $88 billion wage subsidy initiative aimed at helping struggling businesses retain their workforce and keep employees employed; (ii) allowing early access to private pension (superannuation) accounts, which injected an additional $38 billion into households; and (iii) the allocation of $52 billion in supplementary income payments to support various groups, including the unemployed, students, and parents.

Of most relevance to the analysis of labor supply is the JobKeeper program. Described in more detail in Borland and Hunt (2023), this scheme involved payments to eligible employers of up to $750 per week per employee. Only employers anticipating significant reductions in business turnover were eligible.[1] The scheme was in place for one year and at its peak is reported to have covered almost 30% of the workforce. The primary aim of this program was to keep employees connected to their employer, but a side consequence could have been an increased willingness on the part of employers to allow employees to work from home. Additionally, for some employees of eligible employers (those earning less than $750 per week) the JobKeeper payments should have acted to boost wage income. HILDA Survey data, however, suggests the proportion affected was relatively small. When interviewed in wave 20 (conducted roughly 6

[1] Additionally, workers employed on a casual basis who had been employed with the employer for less than one year were also not covered. The HILDA Survey indicates that in wave 19, conducted roughly 6 months prior to the introduction of JobKeeper, such workers represented about 8% of all employees.

months after the commencement of JobKeeper), only 5.4% of employees (after weighting) reported both having received JobKeeper payments and, at the previous interview in wave 19, earning less than $750 per week in their main job.

C. School Closures and Disruptions

School closures were, at least initially, controversial. The Australian Government insisted that it was safe for children to continue to physically attend school. Additionally, there were concerns about how essential workers, and especially those working in health and aged care settings, would be able to cope should their children be unable to attend school (and the associated before-and-after-school care services). Schools were thus exempt from the ban on public gatherings announced in March 2020.

School education, however, is a mainly state government responsibility and the states took a different view and moved rapidly away from the traditional in-person learning model to reliance on online learning. New South Wales, Victoria, Queensland and the ACT all mandated school closures in late March (but with exceptions for children of essential workers) and / or brought forward and extended the approaching school break between mid and late April. The other states opted for a hybrid approach where it was left to parents to decide whether they wished their child to physically attend school or to learn from home. These restrictions, however, did not last long. In South Australia and the Northern Territory, students were back in the classroom in late April. In other states the transition was slower, typically involving a period of transition and only commencing once other community restrictions were eased in early May. In Victoria, the return to the classroom was both the slowest and short-lived. On 13 July, schools in Melbourne were again required to cease in-person teaching as part of the wider lockdown imposed that month, a requirement that would, in August, be extended to all Victorian schools. The main exceptions here were Year 11 and Year 12 students, but even that exemption was removed in August. The other exception was children of essential workers, who were permitted to physically attend school but with learning still delivered online. In effect, schools were required to provide a child-minding service for essential workers. This approach would form the model followed in later lockdowns, both in Victoria and elsewhere. Disruptions to childcare services followed a different pattern, and are discussed in Online Appendix C.

An important institutional feature of the Australian school closure regime is that primary and secondary schools were nearly always closed simultaneously by state and territory governments, with few exceptions. This means that school disruptions affected parents of primary-school and secondary-school children in effectively identical ways, a feature that shapes our analyses.

A major issue of contention was the determination of who was an essential worker. In the early months of the pandemic this was largely left to schools and parents, and thus confusing. Victoria, however, eventually introduced a system (in August 2020) that required employers who were eligible to continue to keep trading to issue permits to their employees that would then provide them with access to childcare and school. This was supported by an official list of essential service providers and industries.

III. Data and Methods

A. Data Collection and Sample

The primary data source is the HILDA Survey, a household panel survey that commenced in 2001 with a nationally representative sample of Australian households residing in private dwellings, after weighting for initial non-response (Watson and Wooden 2021). The first wave comprised 13,969 participants from 7,682 households. Interviews were then sought on an annual basis with all members of those initial households aged 15 years or older, along with any persons who subsequently joined a household in which an original household member resided. A top-up sample providing another 2,153 households was added in 2011. Response rates to the HILDA Survey are relatively high, especially the annual re-interview rate, which rose from 87% in wave 2 to over 94% by wave 5 and, for the main sample at least, has remained at levels above that in every wave since (Summerfield et al. 2022, Table 8.3).

HILDA Survey data are collected via interviews and self-administered questionnaires, and each year is concentrated in August to November. Collection in 2020 and 2021 thus overlapped closely with prolonged stay-at-home orders in Victoria and New South Wales.

We then merge into these data a series of variables identifying the presence of COVID-related mitigation and containment policies, the dates on which restrictions were imposed and lifted, and the local government areas (LGAs; similar to US counties or municipalities in terms of scale and function) to which they were applied. These policy data were hand-collected as part of a large-scale project supported by the US National Institutes of Health Grant that compiled and harmonized federal, state, and local COVID-19 mitigation policies across multiple countries, covering approximately nine policy categories: restrictions on gatherings, domestic and international travel, cancellation of public events, stay-at-home orders, public information campaigns, and closures of public transportation, schools, and workplaces. The data collection effort tracked the exact date each policy was enacted, modified, and relaxed, at the most granular geographic level available, and carefully distinguished between recommended and mandated restrictions. For school disruptions specifically, the data distinguish between closures of childcare, primary, and secondary schools, and record whether essential worker exemptions applied. We also collected the essential worker lists released by each state government and constructed a concordance between these lists and four-digit Australian and New Zealand Standard Classification of Occupations (ANZSCO) industry codes to classify workers in the HILDA Survey.

For our analyses, we restrict our sample to people aged 18 to 75 years who responded at least once in survey waves 17 to 21 and who lived in a household where at least one of its members reported being in the labor force at the time of at least one of these five survey waves. Our estimation sample includes two groups: (i) parents of school-age children (aged 5–17, or younger if they have a school-age sibling) residing in the same household, and (ii) childless adults and parents of non-school-aged children, who serve as their comparison group. This comparison group is restricted to adults whose age falls within the range of ages observed among parents of school-age children, ensuring that the two groups are comparable in terms of lifecycle stage and labor market attachment. In addition, we distinguish within parents between households without and with at least one essential worker, providing a second comparison group. Essential workers were allowed to take their children to school in most states during spells of school disruption, so parents in households with an essential worker were plausibly unaffected by school disruptions.

The resulting estimation sample comprises 75,573 person-year observations from 18,729 unique persons in 11,963 unique households: 56,525 person-year observations from age-matched childless adults and 19,048 from parents of school-age children. Among parents, approximately 55 percent reside in households with at least one essential worker. Descriptive statistics are reported in Online Appendix Table A2.

B. Identifying Variation

Our main treatment variable is a binary indicator for whether there was any government-mandated school closure in a respondent's greater capital city statistical area (GCCSA, henceforth "area") and in a given calendar month. This treatment measure is constructed entirely from our collected COVID mitigation policy data, independent of the HILDA Survey.

Even though we have school disruption information at the narrower LGA and date level, we coarsen this variation at the area × month level for two reasons. First, to increase transparency and simplicity. School disruptions and closures in Australia mostly applied broadly within greater capital city regions. We therefore give up little variation with this aggregation yet force our estimates to use only meaningful changs in closures, also bypassing the need to allow for many untreated areas. Second, coarsening maximizes average geographical cell size, increasing statistical precision. The coarsening of the time dimension to the calendar month of interview follows a similar rationale. HILDA Survey's annual interview window is concentrated in August to November, so in waves 20 and 21 interviews overlap closely with the timing of prolonged lockdowns in Victoria and New South Wales. Coarsening time to calendar month eliminates some small but problematic variation for assigning school disruption timing, which facilitates event study estimates. Moreover, since the HILDA Survey conducts annual interviews and interviews are heavily clustered at the end of the year, our event studies present meaningful variation in annual rather than monthly responses to school disruptions. This increases the interpretability and precision of intertemporal effect estimates, at the cost of foregoing our ability to detect very short-run effects (which is anyway limited by the small number of interviews conducted in any one area-by-calendar-month cell).

Figure 1 illustrates the geographic and time variation in primary and secondary school disruption policies across Australian states and territories, and how this variation overlaps with HILDA Survey interview dates. Online Appendix Figure C1 illustrates similar variation in disruptions to childcare services.

In Australian, primary and secondary school closures were independently mandated by state and territory governments. Yet the correlation between monthly primary and secondary school disruptions at the area level is 0.997, and there was only limited variation in grade-specific measures, which we cannot exploit given the small samples of children in any one grade age in the HILDA Survey data. This near-perfect collinearity means that separating our sample into parents of primary- and secondary-school children or estimating grade-specific effects is not feasible. We therefore pool all parents of school-age children regardless of the age of their children. Our sample includes parents of all children aged 5–17; parents of daycare-age children (under five) are also included when they have an older, school-age sibling in the household. Disruptions to childcare services were also widespread and correlated with school disruptions (Online Appendix Figure C1), and hence we control for daycare disruptions in all specifications. Our estimates therefore capture the effect of disruptions to primary and secondary schools combined, net of any effects of concurrent daycare closures.

At the area × month level, school disruptions and stay-at-home orders are also highly correlated ($\rho = 0.871$), making it difficult to separate the labor supply effects of school disruptions from those of general lockdown restrictions in a standard difference-in-differences framework. This high correlation motivates our DDD strategy, which uses within-area-and-month variation across parent types and essential workers in households to isolate the school-disruption channel from the effects of broader stay-at-home orders. However, if there are residual factors correlated with school closures that differentially affect the labor supply of parents (or essential worker parents), our estimate of interest will be biased.

For our DDD strategy, our first comparison group consists of age-matched childless adults, who were not directly affected by school disruptions by virtue of not having children. We define childless adults as individuals aged 18–75 who do not have a school-age child (or any younger child with a school-age sibling) residing in the same household. This non-parents-vs-parents comparison has the advantage of a clean separation between affected and unaffected groups. Its limitation is that structural differences between both groups could be strongly related to labor supply, potentially affecting parallel pre-trends. To alleviate some of this issue, we restrict non-

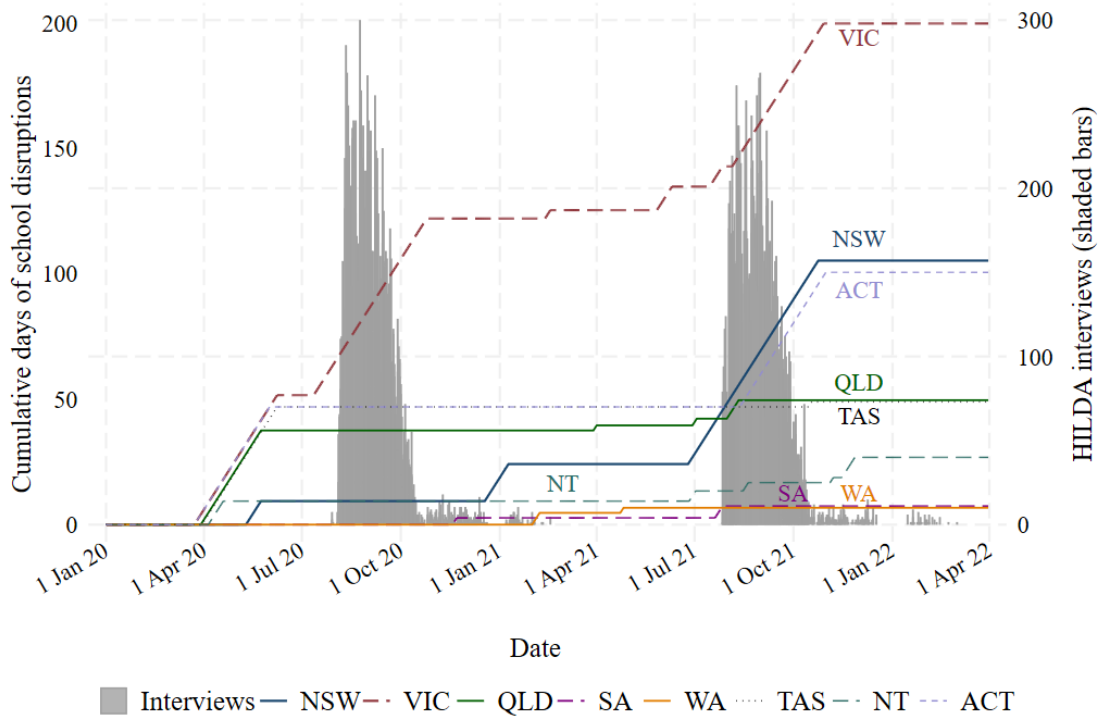

Fig. 1. Across-state variation in cumulative days of school disruption, with HILDA Survey interview dates, 2020 and 2021

In this figure, states are classified as experiencing primary and secondary school disruptions if at least one greater capital city statistical area (GCCSA) within the state is experiencing them at any given date. HILDA Survey interview dates represented by the grey columns. ACT = Australian Capital Territory; NSW = New South Wales; NT = Northern Territory; QLD = Queensland; SA = South Australia; TAS = Tasmania; VIC = Victoria; WA = Western Australia.

parents to adults whose age falls within the range observed among parents of school-age children, ensuring comparability in lifecycle stage.

Our second comparison exploits variation in the presence or absence of essential workers in the household during the pandemic. Essential workers were allowed to take their children to school in most states during spells of school disruptions, so households with at least one essential worker were plausibly unaffected by school disruptions. The no-essential-vs-essential-worker-in-household comparison has the advantage of comparing within the arguably structurally similar affected population of parents, but it comes with two key limitations. First, essential workers were also exempted from other stay-at-home restrictions beyond school disruptions, including workplace closure mandates and some movement restrictions, meaning their labor supply may have been affected differently through channels correlated with but different to school disruptions. Second, the essential worker lists released by Australian state governments were broad and inconsistently defined, and our concordance with ANZSCO industry codes likely introduces misclassification.

Neither of our two DDD designs is without issues, which is precisely why we present both side by side. Each addresses a weakness of the other: One uses a clean affected group assignment but across groups that can be less comparable, while the other uses more comparable groups but treatment differences across groups need to be more carefully interpreted. Together, these two estimators provide a more complete picture than either would alone.

C. Identification and Estimation

Our two complementary designs compare labor market outcomes in affected areas, before and after they are affected by school disruptions, drawing two different comparisons: parents of school-age children versus non-parents of school-age children, and parents without versus with an essential worker in the household. The estimand in both cases is the effect of school disruptions over and above the effect of other events — such as stay-at-home orders, daycare disruptions, and COVID case rates. The identifying assumption is also the same in both cases: that the gap in labor market outcomes between groups would have evolved in a parallel way in the absence of school disruptions. Under this assumption, any differential change in labor supply among parents (without an essential worker at home) relative to non-parents (to parents with an essential worker at home) during school disruptions can therefore be attributed to the disruption itself. We provide evidence supporting this assumption in Section IV.B.

To estimate the Average Treatment Effect on the Treated (ATET) of school disruptions, we fit the following model:

$$Y_{it} = \beta_1\left(\text{Disruption}_{at} \times \text{Group}_i\right) + \beta_2'(X_{at} \times \text{Group}_i) + \delta_{at} + \gamma_{a\times G_i} + \lambda_{t\times G_i} + \varepsilon_{it}$$

where Y_{it} is the labor market outcome of individual i at time t, Disruption_{at} is a binary indicator for whether there was a school disruption in area a in the calendar month containing t, and Group_i is a time-invariant indicator for either i) having a school-age child in the household during the pandemic, or ii) not having an essential worker in the household in the year immediately preceding the pandemic, and X_{at} is a vector of covariates that vary at the area-month level—and in our main estimation consists of daycare disruptions. The estimation samples include parents and non-parents of school-aged children during the pandemic in the former specification, and only parents in the latter. The fixed effects δ_{at} (area-by-time) absorb all common shocks to a given area in a given period — including stay-at-home orders, COVID case rates, and other area-level policies — while $\gamma_{a\times G_i}$ (area-by-group) and $\lambda_{t\times G_i}$ (time-by-group) absorb time-invariant differences between groups within areas and national trends that differ by group. These also absorb the main effects for Disruption_{at}, Group_i, and X_{at}. The DDD coefficient of interest β_1 captures the differential effect of school disruptions.

For inference, we cluster at the area level, consistent with Bertrand et al. (2004). With only 13 clusters, the standard cluster-robust variance estimator is unreliable because its finite-sample distribution can be badly skewed when clusters are few, unequal in size, or have heterogeneous within-cluster error correlations. We therefore report 95 percent confidence intervals from the wild cluster bootstrap (see e.g., Cameron et al., 2008) throughout. The wild cluster bootstrap addresses this unreliability not by estimating a variance, but by directly simulating the distribution of the test statistic under the null and inverting that distribution to recover confidence intervals.

In Panel I of Table 1, we compare parents of school-age children to age-matched childless adults — our preferred specification. In Panel II, we compare parents in households without an

Tab. 1. Effect of School Disruptions on Labor Supply by Gender

	In labor force		*Weekly hours worked*	
	Women	Men	Women	Men
	Panel I: Parents vs Non-parents			
School disruptions	-.00679	.0104	.168	-.628
	[-.066, .046]	[-.034, .1]	[-2.6, 1.4]	[-2.9, 1]
Outcome mean	0.80	0.92	33.6	42.9
Observations	23,196	21,021	17,838	18,344
	Panel II: Parents w/o vs w/ Essential Workers in HH			
School disruptions	-.0129	-.000577	.383	-2.36**
	[-.13, .1]	[-.057, .074]	[-1.2, 2.8]	[-5.9, -.45]
Outcome mean	0.74	0.93	30.5	43.6
Observations	10,690	7,986	7,600	7,214

*Triple Difference (DDD) Average Treatment Effect on the Treated estimates of experiencing a school disruption in the month of interview. Panel I identifies effects on parents relative to non-parents of school-aged children during the pandemic. Panel II identifies effects on parents without an essential worker in the household relative to parents with an essential worker in the household. Estimates use Greater City as the group dimension and month of interview as the time dimension. Square brackets show 95 percent confidence intervals based on wild cluster bootstrap inference at the Greater City level, calculated using 1,000 replications. ** denotes statistical significance at the 5 percent level based on wild cluster bootstrap inference.*

essential worker to parents in households with one, where the latter group's children could continue attending school during closures.

IV. Results

A. Main DDD Estimates

Table 1 shows that school disruptions had no detectable effect on parental labor supply in Australia.

In Panel I, we compare parents of school-age children to age-matched childless adults. Australian maternal employment, closely linked to maternal labor force participation, was around the OECD average (much higher than the USA but lower than most EU countries) before the pandemic, and childcare subsidies and flexible work arrangements were generally available — features that may have limited the scope for school disruptions to generate large labor supply responses. Consistent with this baseline, the DDD estimates are small relative to the outcome means and have 95 percent confidence intervals that include zero for both outcomes and both genders. The estimated effect on women's labor force participation is −0.007 on a mean of 0.80 (95% CI: [−0.066, 0.046]), corresponding to less than 1 percent of the mean. For men, the estimate is 0.010 on a mean of 0.92 (95% CI: [−0.034, 0.101]). The North American literature has reported reductions in maternal labor force participation in the order of 2 to 5 percentage points (e.g., Heggeness 2020; Collins et al. 2021, Beauregard et al. 2022). While our confidence intervals do not fully rule out effects at the upper end of that range — for women, we can rule out negative effects larger than about 6.6 percentage points (8.3 percent of the mean), and for men, negative effects larger than 3.4 percentage points (3.7 percent of the mean) — our point estimates are close to zero and the overall pattern of results strongly suggests that school disruptions in Australia had substantially smaller effects on labor force participation than those documented in the United States and Canada.

Effects on weekly hours worked are similarly small. The point estimates are 0.17 hours for women (mean: 33.6; 95% CI: [−2.59, 1.44]) and −0.63 hours for men (mean: 42.9; 95% CI: [−2.85, 1.01]), both with confidence intervals that exclude effects larger than about 3 hours.

In Panel II, we compare parents in households without an essential worker to parents in households with one. We find a statistically significant negative effect on men's weekly hours worked (−2.4 hours, 95% CI: [−5.94, −0.45]), though the point estimate is modest relative to the mean of 43.6 hours (5.5 percent). As noted in Section III, though, households with essential workers were differentially affected by concurrent stay-at-home restrictions, which needs to be taken into consideration when interpreting these results. For women, the Panel II estimates have confidence intervals that include zero for both outcomes. Note also that these confidence intervals are wider than in Panel I, reflecting the smaller estimation sample (parents only). For women's labor force participation, Panel II CI spans 23 percentage points (versus 11 in Panel I), meaning we cannot rule out moderate effects in this specification. For other outcomes, the loss in precision is more modest. With the exception of men's hours in Panel II, both panels paint a consistent picture of null effects on labor supply.

B. Event Study Analysis

The ATETs are useful summaries of the effects of school closures on parents but suffer from two shortcomings. First, they may be biased in the presence of heterogeneous and dynamic treatment effects. Second, they do not allow us to provide evidence on the key identifying assumption of parallel pre-trends. To address both shortcomings, we turn to the de Chaisemartin and D'Haultfoeuille (2024; hereafter DCDH) estimator, which is designed for difference-in-differences settings with staggered treatment adoption and is robust to treatment effect heterogeneity across groups and over time. The DCDH estimator (1) avoids using earlier school disruptions as controls for later disruptions by defining counterfactuals as areas that have not yet experienced any change in their school disruption status, thereby ensuring that already-disrupted areas are not used as controls (de Chaisemartin and D'Haultfoeuille, 2024, p. 3), and (2) allows for a treatment measure that can increase and decrease over time, essential for us since schools may close, reopen, close, and reopen again. However, the DCDH estimator requires separately examining parents of school-aged children and childless adults in difference-in-differences (DD) regressions rather than in fully interacted DDD regressions.

Figure 2 presents DCDH event studies for the parents-versus-non-parents comparison, with a four-panel layout — two rows (labor force participation, weekly hours worked) by two columns (women, men) — overlaying estimates for parents (in blue) and non-parents (in red). In Online Appendix Figure A1 we show analogous event studies for the essential-worker comparison. Pre-treatment estimates across groups in both designs evolve in parallel across all four panels,

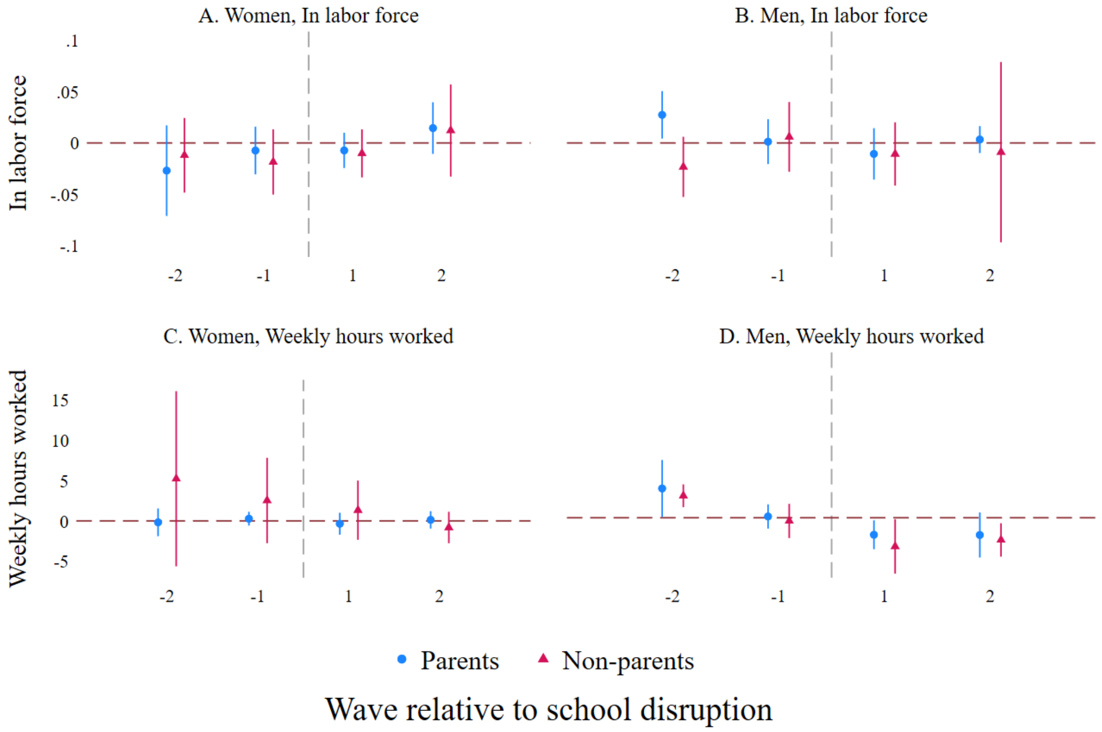

Fig. 2. Intertemporal treatment effects of school disruptions on labor supply, for parents and non-parents

de Chaisemartin and D'Haultfoeuille (2024) difference-in-differences event study estimates, estimated separately for women and men. Treatment effects relative to outcomes in the last non-disrupted HILDA wave (t=0). Effect estimates are from de Chaisemartin and D'Haultfeuille (2024), and plots are using the graphical package by Jann (2014). Blue circles: parents o school-aged children. Red circles: non-parents of school-aged children. The group dimension is greater capital city statistical area and the time dimension is HILDA Survey wave. Analytical 95 percent confidence intervals shown.

supporting the parallel trends assumption underlying our DDD designs. Both figures also suggest null intertemporal treatment effects on labor supply.

DCDH average total effect estimates by gender, partnership status, and closure intensity (Tables A8–A10) are fully consistent with the DDD patterns in Section IV.A: null effects across all specifications, suggesting limited bias from heterogeneous and dynamic treatment effects.

C. Wage Effects

We report wage effects separately in the Online Appendix since there are some parallel pre-trend violations for this outcome. The DCDH event studies in the parents-versus-non-parents design show that pre-treatment wage trends diverge for both women and men, with the divergence particularly pronounced for women (see Online Appendix Figure A2). Pre-trends hold better in the essential-versus-non-essential-worker design.

Also, ATETs of school disruption on wages, while also null, are less precisely estimated (Online Appendix Table A3). In the parents-versus-non-parents design, the estimates are small and statistically insignificant for both women (\$75, 95% CI: [−\$50, \$168], mean: \$1,312) and men (\$70, 95% CI: [−\$324, \$266], mean: \$1,850). In the non-essential-vs-essential-worker-households design, the estimate for men is negative and statistically significant (−\$405, 95% CI: [−\$890, −\$46], mean: \$2,096), the only significant estimate in the table. However, given the potential measurement error in essential worker classification discussed in Section III, and the fact that essential workers were exempted from other lockdown restrictions beyond school disruptions, we hesitate to read too much into this one significant result. It is also difficult to reconcile a negative effect on wages with the main null effects on labor supply.

D. Heterogeneity, Robustness, and Mechanisms

The null effects in Table 1 are not hiding any evidence of heterogeneity. As evidence, we first show they remain across different measures of treatment. In Table 2 we explore the effects of continuous measures of school disruption intensities within 30 days, 60 days, and 90 days of the interview date, and of the cumulative exposure over the entire sample period. Across all four intensity measures, both DDD comparisons, and both labor market outcomes, all 16 estimates have 95 percent confidence intervals that include zero. The point estimates are uniformly small across all outcomes and specification. The 95% CIs are narrow for hours worked and for the parent-vs-non-parent comparison of labor force participation, but wider for labor force participation when using the non-essential-vs-essential-worker-in-household approach.

Tab. 2. Effects by Closure Intensity, Comparing Parents to Non-Parents

	In labor force		*Weekly hours worked*	
	Parents vs Non-parents	Non-essential vs Essential	Parents vs Non-parents	Non-essential vs Essential
	Panel I: 30-day threshold			
No → Full disruption	-.000735	-.0321	-.0884	-.471
	[-.037, .028]	[-.17, .076]	[-1.9, 1.4]	[-1.9, .33]
Outcome mean	0.86	0.83	38.3	36.9
Observations	43,467	18,676	35,437	14,814
	Panel II: 60-day threshold			
No → Full disruption	.002	-.0249	-.228	-.264
	[-.038, .032]	[-.17, .07]	[-2, 1.5]	[-1.1, .38]
Outcome mean	0.86	0.83	38.3	36.9
Observations	43,467	18,676	35,437	14,814
	Panel III: 90-day threshold			
No → Full disruption	.00432	-.0181	-.149	.0954
	[-.058, .047]	[-.16, .1]	[-2.5, 2]	[-.88, .9]
Outcome mean	0.86	0.83	38.3	36.9
Observations	43,467	18,676	35,437	14,814
	Panel IV: Cumulative			
100 days disruption	.00911	-.0166	.527	1.17
	[-.012, .045]	[-.093, .11]	[-.7, 2.4]	[-.43, 2.7]
Outcome mean	0.86	0.83	38.3	36.9
Observations	43,467	18,676	35,437	14,814

*Triple Difference (DDD) Average Treatment Effect on the Treated estimates of experiencing a school disruption in the month of interview. Panel I identifies effects on parents relative to non-parents of school-aged children during the pandemic. Panel II identifies effects on parents without an essential worker in the household relative to parents with an essential worker in the household. Estimates use Greater City as the group dimension and month of interview as the time dimension. Square brackets show 95 percent confidence intervals based on wild cluster bootstrap inference at the Greater City level, calculated using 1,000 replications. ** denotes statistical significance at the 5 percent level based on wild cluster bootstrap inference.*

Table 3 then probes the robustness of our null results across subsamples. Focusing on the parents-versus-non-parents DDD design, we re-estimate effects by child age, partnered status, and education level, estimating each specification separately for women and men. All 24 estimates (six panels × two outcomes × two genders) have 95 percent confidence intervals that

Tab. 3. Effect Across Subgroups, Comparing Parents to Non-Parents

	In labor force		*Weekly hours worked*	
	Women	Men	Women	Men
	Panel I: Parents of Children Under 12 vs Childless Adults			
School disruptions	.000915	.015	.895	.422
	[-.081, .063]	[-.051, .058]	[-.61, 3.1]	[-1.1, 1.8]
Outcome mean	0.81	0.91	34.3	42.7
Observations	17,484	17,281	13,588	14,907
	Panel II: Parents of Children 12+ vs Childless Adults			
School disruptions	.0111	-.0184	-1.02	.324
	[-.036, .071]	[-.049, .013]	[-3.9, 1.5]	[-2, 3.1]
Outcome mean	0.84	0.91	35.4	42.5
Observations	15,649	15,232	12,547	13,017
	Panel III: Unpartnered - Parents vs Childless Adults			
School disruptions	.0264	.0356	.535	-.356
	[-.09, .1]	[-.069, .11]	[-3.8, 3.1]	[-7.3, 7.3]
Outcome mean	0.78	0.83	35.0	40.8
Observations	6,185	5,539	4,474	4,153
	Panel IV: Partnered - Parents vs Childless Adults			
School disruptions	.00281	-.00489	.127	-.157
	[-.033, .054]	[-.034, .024]	[-1.3, 2.8]	[-2.4, 1.1]
Outcome mean	0.81	0.95	33.2	43.5
Observations	17,011	15,482	13,364	14,191
	Panel V: Less Educated - Parents vs Childless Adults			
School disruptions	.00993	.00452	-.263	2.81
	[-.038, .071]	[-.097, .068]	[-3.8, 4.9]	[-2, 8.9]
Outcome mean	0.68	0.86	31.4	42.2
Observations	5,967	5,989	3,771	4,785
	Panel VI: Highly Educated - Parents vs Childless Adults			
School disruptions	-.00197	.0102	.275	-.326
	[-.051, .067]	[-.026, .038]	[-.62, 1.8]	[-2.2, .77]
Outcome mean	0.84	0.94	34.2	43.1
Observations	17,234	15,037	14,071	13,563

*Triple Difference (DDD) Average Treatment Effect on the Treated estimates of experiencing a school disruption in the month of interview across subsamples. Panels I and II split by child age group (straddler families excluded). Panels III and IV split by marital status. Panels V and VI split by education level (Year 12 or below vs post-secondary). In all panels the comparison group is childless adults matched on the relevant characteristic. Estimates use Greater City as the group dimension and month of interview as the time dimension. Square brackets show 95 percent confidence intervals based on wild cluster bootstrap inference at the Greater City level, calculated using 1,000 replications. ** denotes statistical significance at the 5 percent level based on wild cluster bootstrap inference.*

partnered parents, parents of children under 12 — the confidence intervals are comparable in width to those in Table 1, Panel I, providing meaningful bounds on subgroup effects. For smaller subsamples (e.g., unpartnered parents), the wider confidence intervals cannot rule out moderate effects, which is to be expected. The absence of any systematic pattern across subgroups suggests that the null results are unlikely to mask heterogeneous effects that cancel in the aggregate, at least for the subgroups where we have adequate power.

The null results further extend to every heterogeneity dimension we examine in the Online Appendix. DDD estimates by partnership status are uniformly null in the no-essential-vs-essential-worker design (Table A4, Panel II). The underlying DD estimates — decomposed by gender, partnership status, and closure intensity across all four subsamples (parents, non-parents, non-essential-worker households, essential-worker households; Tables A5–A7) — show point estimates of similar magnitude across treated and comparison groups, supporting the DDD finding that school disruptions had no incremental effect on parental labor supply beyond the general effects of stay-at-home orders.

The null results are robust to alternative specifications. When we add individual-level controls — gender, age, indigeneity, country of birth, and education — to the DDD (Table A11), the point estimates remain similar to our main specification, with confidence intervals that include zero. Given that the HILDA is an individual-level panel dataset, we also experiment with including both individual and area fixed effects in our specification. The findings, which are identified using person-specific labor supply changes within jurisdictions, are qualitatively similar to those presented above and available upon request.

V. Discussion

School disruptions during the COVID-19 pandemic did not reduce parental labor force participation or hours worked in Australia. Our DDD estimates comparing parents of school-age children to age-matched childless adults are small, precisely estimated, and statistically indistinguishable from zero for both outcomes and for both women and men. In the United States, school closures reduced maternal labor force participation by 2 to 5 percentage points (Heggeness 2020; Collins et al. 2021; Amuedo-Dorantes et al. 2023). While our confidence intervals do not fully rule out effects at the upper end of that range, our point estimates are close to zero and the overall pattern of results — null effects across both DDD designs, multiple subgroups, and alternative estimators — strongly suggests that school disruptions in Australia had substantially smaller effects on parental labor supply than those documented in North America. These null results are robust across an extensive set of specifications. They hold when we vary treatment intensity and across several groups and specifications. Wage effects are also null, though their identification is less convincing and they are less precisely estimated.

Why did school disruptions not reduce parental labor supply in Australia, when they appear to have done so in the United States and Canada? We hypothesize that three institutional features distinguishing the Australian experience may explain the null results, though we note that our data do not allow us to isolate the contribution of each channel.

First, the Australian Government's economic policy response was swift and generous. The JobKeeper Payment, introduced in March 2020, subsidized wages at a flat rate of $1,500 per fortnight for eligible employees, effectively providing an incentive for firms to retain workers and for workers to remain attached to the labor force. Additionally, the Coronavirus Supplement doubled the standard unemployment payment. These programs were explicitly designed to maintain labor force attachment during the pandemic and may have counteracted any tendency for parents to withdraw from the labor force in response to school disruptions.

Second, school disruptions in Australia were, in general, shorter and more localized than in many other countries. The average number of disrupted days was modest by international standards, and many regions — outside Victoria and New South Wales — experienced minimal disruption to in-person teaching. The brief and geographically concentrated nature of Australian school disruptions may have limited the scope for large labor supply adjustments.

Third, the rapid and widespread adoption of working from home may have enabled parents to manage childcare responsibilities without reducing their labor supply. Government advice early in the pandemic directed employers to facilitate remote work wherever possible, and many white-collar industries shifted to remote arrangements within weeks. The flexibility to work from home could have served as a buffer, allowing parents to absorb the childcare burden of school disruptions without withdrawing from employment or reducing hours.

Our study does carry several limitations. The hours worked variable in the HILDA Survey measures hours in a "usual" week, and it is unclear how respondents whose hours were temporarily disrupted during the pandemic interpreted this question. If respondents anchored on pre-pandemic usual hours rather than their actual hours during the survey reference period, our estimates of the effect on hours worked would be biased toward zero. This concern is mitigated somewhat by the fact that we also find null effects on labor force participation, which is not subject to the same measurement ambiguity. The under-representation of recent immigrants in the HILDA Survey is another limitation, since this group may have been disproportionately affected by pandemic restrictions (see e.g., Van Barneveld et al. 2020). We also note that our coarsened treatment to the area × month level, could obscure finer-grained variation in disruption timing and intensity that could have generated localized labor supply effects. Also, the HILDA Survey's annual interview window is concentrated in August to November, limiting our ability to distinguish short-run from medium-run effects in most cases. Finally, we note again the limitation to both of our DDD specifications, as discussed in Section III.

Despite these limitations, the Australian null result is substantively important. It suggests that the effects of school disruptions on parental labor supply depend on the institutional context in which disruptions occur. In countries with generous wage subsidies, widespread remote work options, and relatively short and localized school disruptions, parents may be able to absorb the childcare burden without measurable changes in labor market outcomes. The Australian experience provides a useful counterpoint to the US evidence and highlights the role that government policy can play in mitigating the labor market consequences of pandemic-related disruptions to schooling.

References

Amuedo-Dorantes, Catalina, Miriam Marcén, Marina Morales, and Almudena Sevilla. 2023. Schooling and parental labor supply: Evidence from COVID-19 school closures in the United States. *ILR Review* 76(1): 56-85.

Beauregard, Pierre-Loup, Marie Connolly, Catherine Haeck, and Timea Laura Molnár. 2022. Primary school reopenings and parental work. *Canadian Journal of Economics / Revue canadienne d'économique* 55(S1): 248-281.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics*, 119(1): 249-275.

Borland, Jeff, and Jennifer Hunt. 2023. JobKeeper: An initial assessment. *Australian Economic Review* 56(1): 109–123.

Breunig, Robert, and Tristram Sainsbury. 2023. Too much of a good thing? Australian cash transfer replacement rates during the pandemic. *Australian Economic Review* 56(1): 70–90.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2008. Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics* 90(3): 414–427.

Collins, Caitlyn, Leah Ruppanner, Liana Christin Landivar, and William J. Scarborough. 2021. The gendered consequences of a weak infrastructure of care: School reopening plans and parents' employment during the COVID-19 pandemic. *Gender & Society* 35(2): 180-193.

de Chaisemartin, Clément, and Xavier D'Haultfoeuille. 2024. Difference-in-differences estimators of intertemporal treatment effects. *Review of Economics and Statistics* (published online ahead of print version).

Garcia, Kairon Shayne D., and Benjamin W. Cowan. 2024. Childcare responsibilities and parental labor market outcomes during the COVID-19 pandemic. *Journal of Labor Research* 45(2): 153-200.

Guven, Cahit, Panagiotis Sotirakopoulos, and Aydogan Ulker. 2023. Individual labour market transitions of Australians during and after the national COVID-19 lockdown. *Applied Economics* 55(8): 853-868.

Hansen, Benjamin, Joseph J. Sabia, and Jessamyn Schaller. 2024. Schools, job flexibility, and married women's labor supply. *Journal of Human Resources* (published online ahead of print version).

Heggeness, Misty L. 2020. Estimating the immediate impact of the COVID-19 shock on parental attachment to the labor market and the double bind of mothers. *Review of Economics of the Household* 18(4): 1053-1078.

Jann, Ben. 2014. Plotting regression coefficients and other estimates. The Stata Journal, 14(4): 708-737.

Meekes, Jordy, Wolter H.J. Hassink, and Guyonne Kalb. 2023. Essential work and emergency childcare: Identifying gender differences in COVID-19 effects on labour demand and supply. *Oxford Economic Papers* 75(2): 393-417.

Salamanca, Nicolas, Tanya Gupta, Mark Wooden, Irma Mooi-Reci, 2026, "Replication Data for: The impact of school disruptions during the COVID-19 pandemic on parental labor supply and earnings in Australia", Harvard Dataverse, https://doi.org/10.7910/DVN/J2PQLN.

Schurer, Stefanie, Kadir Atalay, Nick Glozier, Esperanza Vera-Toscano, and Mark Wooden. 2023. Quantifying the human impact of Melbourne's 111-day hard lockdown experiment on the adult population. *Nature Human Behaviour* 7(10): 1652–1666.

Stobart, Anika, and Stephen Duckett. 2022. Australia's response to COVID-19. *Health Economics, Policy and Law* 17(1): 95-106.

Summerfield, Michelle, Brooke Garrard, Roopah Kamath, Ninette Macalalad, Mossamet Kamrun Nesa, Nicole Watson, Roger Wilkins, and Mark Wooden. 2022. HILDA User Manual – Release 21. Melbourne Institute: Applied Economic and Social Research, University of Melbourne. (https://melbourneinstitute.unimelb.edu.au/hilda/for-data-users/user-manuals)

Van Barneveld, Kristin, Michael Quinlan, Peter Kriesler, Anne Junor, Fran Baum, Anis Chowdhury, PN (Raja) Junankar, Stephen Clifford, Frances Flanagan, Chris F Wright, Sharon Friel, Joseph Halevi, and Al Rainnie. 2020. The COVID-19 pandemic: Lessons on building more equal and sustainable societies. *Economic and Labour Relations Review* 31(2): 133-157.

Watson, Nicole, and Mark Wooden. 2021. The Household, Income and Labour Dynamics in Australia (HILDA) Survey. *Jahrbücher für Nationalökonomie und Statistik* 241(1): 131-141.