NET WORK LISER

Discussion Paper Series

IZA DP No. 18708

June 2026

Policy or Pandemic? Explaining the Rise in Young Adults' Private Health Insurance Take-Up in Australia

Nathan Kettlewell
University of Technology Sydney
and IZA@LISER

Tiannan Gao
University of Technology Sydney

IZA DP No. 18708 June 2026

Policy or Pandemic? Explaining the Rise in Young Adults' Private Health Insurance Take-Up in Australia

Abstract

This study uses administrative data to estimate how an age-based discount policy introduced in 2019 affected the take-up of private health insurance in Australia before and during the COVID-19 pandemic. The policy provided a perpetual premium discount of 2%–10% for people who purchased insurance before age 30. Applying a synthetic difference-in-differences model, we find that the policy had only a small and positive effect on take-up prior to the pandemic, which was a period of declining insurance membership. However, there was a significant increase in take-up in the treatment group in the period during the COVID-19 pandemic, when the overall demand for private health insurance was increasing. We explore the extent to which the larger treatment effect post-COVID is likely to be due to the discount policy, rather than other age-specific demand factors relevant to the pandemic. We conclude that the relative increase in demand among the young is predominantly driven by factors other than the discount policy.

JEL classification

I13, I18

Keywords

health insurance, premium discount, COVID-19, synthetic difference-in-differences

Corresponding author

Nathan Kettlewell
Nathan.Kettlewell@uts.edu.au

1 Introduction

In countries with dual public-private health systems, subsidies for health insurance are widely used to stimulate demand and address adverse selection (Colombo and Tapay, 2004). Several studies investigate how premium subsidies affect demand generally in different institutional settings, often leveraging natural experiments in the ways such policies are implemented (e.g., Cheng, 2014; Finkelstein, 2002; Frean et al., 2017; Gruber and Washington, 2005; Hinde, 2017; King and Mossialos, 2005; Nicolás and Vera-Hernández, 2008; Rodríguez and Stoyanova, 2008). Studies typically find that health insurance demand is inelastic. For example, Finkelstein (2002) estimates the price elasticity for supplementary health insurance in Canada to be -0.5. In Australia, estimates range from -0.35 to -0.5 (Duckett et al., 2019).

In this paper, we evaluate the effect of a price discount policy for private health insurance (PHI) implemented in Australia, a country with a dual public-private insurance system. A unique feature of this setting is that the policy specifically targets young adults only (ages 18-29 years), allowing us to learn about the price-responsiveness of this group specifically within a natural experiment design. While not always the primary target of health insurance subsidies, and often overlooked in health insurance subsidy research, young adults are a particularly important group. Their participation in the insurance pool is critical for reducing adverse selection and ensuring market stability. Indeed, the discount policy was introduced precisely in response to growing concerns about declining PHI membership among younger Australians, which threatened the viability of the PHI system. This can be seen in Figure 1a – starting in 2015, there was a decline in PHI membership that only reversed at the onset of the COVID-19 pandemic. This decline reignited concerns from the 1990s about an adverse selection "death spiral" (Duckett et al., 2019; Duckett, 2020; Zhang, 2020). Demographic trends further illustrate the challenge of attracting younger people, as the average age of insured individuals has been increasing annually, with the 20–29 age group persistently exhibiting the lowest coverage rates (Figure 1b).

Age-based discounts represent a controlled step away from broad community rating, a cornerstone of PHI in Australia (with some exceptions discussed later). They allow insurers to offer a premium discount for those aged 18-29 years of 2%–10% (further details in Section 2.1). The policy commenced in April 2019 during the downward trend in memberships. This trend reversed in Q2 2020, coinciding with the onset of the COVID-19 pandemic. While the precise drivers of this reversal remain unclear, it is likely that increased concerns about healthcare access and public system capacity constraints

Figure 1: Coverage Trend

(a) Australian PHI Coverage

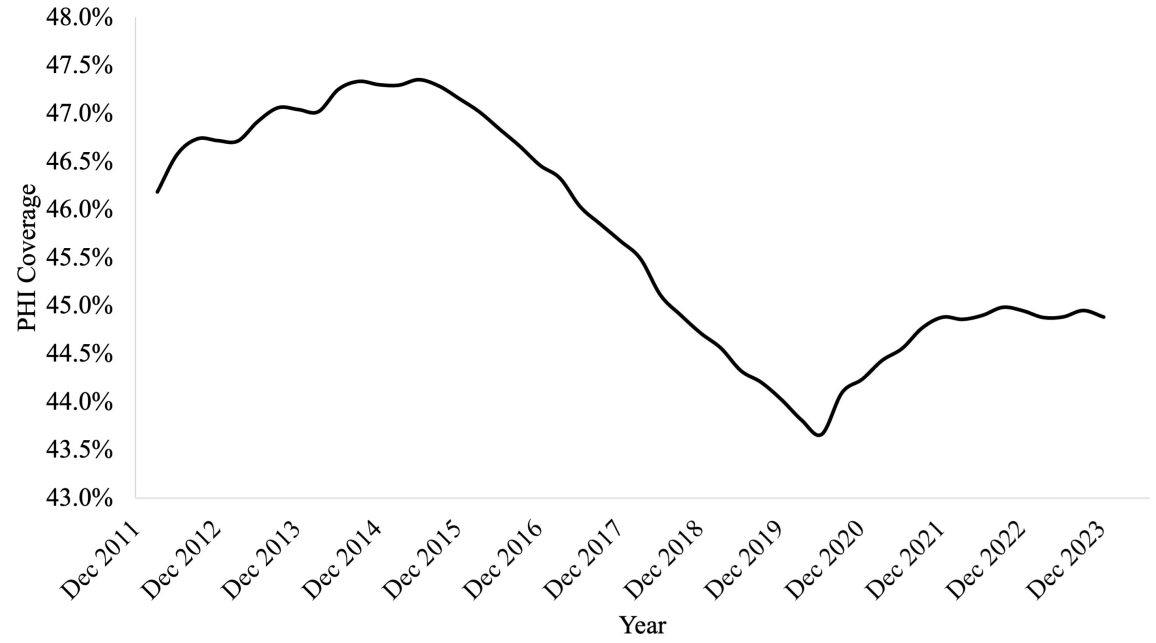

(b) Trends in PHI coverage by age group
(insured persons as % of population)

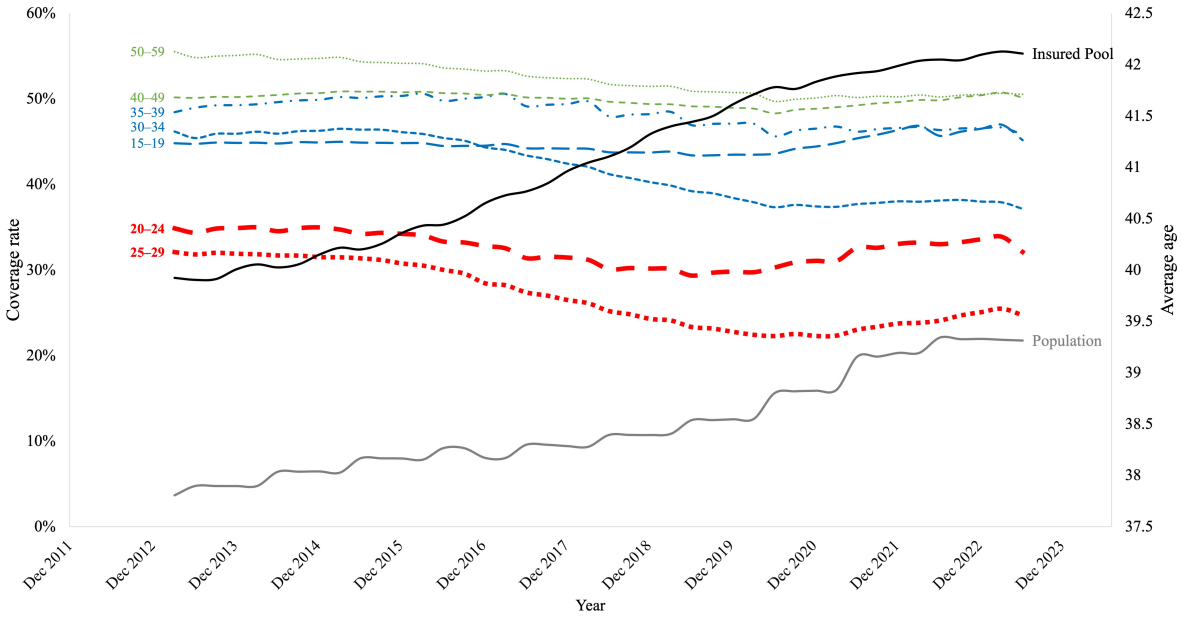

Source: (a) Coverage rate is from the APRA *Quarterly private health insurance statistics* issue released on 28 February 2024 (Australian Prudential Regulation Authority, 2024b). (b) The number of insured persons by age cohort is from APRA *Quarterly private health insurance statistics* issue released on 28 February 2024 (Australian Prudential Regulation Authority, 2024b). The population size by age cohort is from the Australian Bureau of Statistics *Australian Demographic Statistics* (Australian Bureau of Statistics, 2020a). The average age is based on a weighted average of the mid-points for each age group.

contributed to rising PHI demand (Higginson et al., 2020; Koce, 2022; Walker, 2020).

We evaluate age-based discounts for those aged 20–24 and 25–29 by using adminis-

trative data and a synthetic difference-in-differences design (Arkhangelsky et al., 2021). This method compares people in the target age groups to weighted counterfactuals of people in other age groups. A key empirical challenge arises from the fact that the treatment groups (at least for people in the 25–29 years age group) experienced the steepest decline in PHI coverage prior to policy implementation, violating the parallel trends assumption required for DID estimation. To deal with this, we apply a rotation correction to the pre-intervention period, following a similar approach proposed by Freyaldenhoven et al. (2021). We demonstrate that, after a linear trend adjustment, the treatment and control groups exhibit parallel trajectories prior to the intervention, strengthening the validity of our identification strategy.

Our findings indicate that the age-based discount policy had a modest effect on take-up of PHI in the three quarters before the onset of the COVID-19 pandemic. In the period during the pandemic, the estimated policy effect is larger. Through a series of exploratory analyses we investigate whether this is likely to be a 'true' treatment effect (e.g., reflecting delayed response to the policy and/or a positive interaction with heightened overall demand during the pandemic) or instead an age-specific shift in demand due to the pandemic itself. We conclude that the majority of the treatment effect we estimate in the pandemic period is due to the pandemic itself, rather than the age-discount policy, implying that the small recovery in coverage among young Australians is due to external influences rather than government incentive policy.

Our paper contributes to the international literature on the effectiveness of incentive policies on demand for private insurance (Bilgrami et al., 2021; Courtemanche et al., 2017; Finkelstein, 2002; Frean et al., 2017; Kettlewell et al., 2018; Kettlewell and Zhang, 2024; Liu and Zhang, 2023; Marenzi et al., 2021; Palangkaraya and Yong, 2005; Stavrunova and Yerokhin, 2014; Thomasson, 2003). Identifying cost-effective strategies to enhance private insurance uptake is particularly important in contexts where private insurance is heavily regulated (e.g., community rating) since increased private coverage can alleviate pressures on public healthcare systems, and governments may face pressure to bolster demand for private insurance.

In the Australian context, our paper adds to the evidence-base on the many incentive policies that exist for PHI. Closest to the age-based discount policy we evaluate is Lifetime Health Cover (LHC) loading, which acts as a kind of antipode to the discount policy by allowing insurers to charge an additional 2% premium loading for each year uninsured after age 31. Ellis and Savage (2008) attribute a large increase in PHI membership to the introduction of LHC and its associated advertising in 2000, although Kettlewell and Zhang (2024) find it only increases take-up at age 31 by 4-6% in more recent years.

The most controversial incentive policy has been premium subsidies, which comprise 6.5% of federal health spending in 2020-21 (Australian Institute of Health and Welfare., 2022). Cheng (2014) simulated the effect of removing the subsidy and found it would reduce government expenditure since spending on the rebate exceeds reduced spending on public hospital care. Kettlewell et al. (2018) and Liu and Zhang (2023) evaluate the effect of the higher subsidy for older people and find nil and small effects on take-up respectively. The age-based discount is an attractive alternative to the subsidy program because it is targeted and its cost is borne by insurers, rather than taxpayers. Our findings indicate that the program may have helped insurers attract younger individuals into the insurance pool and back-of-the-envelope calculations suggest it is plausible the policy improved insurers' bottom lines.

Our paper is organised as follows. Section 2 provides the institutional background for health insurance in Australia. In Section 3 we discuss our data and identification strategy. In Section 4, we discuss our results, conduct robustness checks and disentangle COVID-19 effects. In Section 5 we provide some back-of-the-envelope calculations for the effect of the discount policy on insurers. Section 6 concludes and discusses policy implications.

2 Institutional background

All Australians are entitled to free public hospital care and heavily subsidised primary and specialist care through a system of universal cover known as Medicare, which is funded through general revenue. Despite this, around 45% of people have chosen to purchase PHI with hospital cover.[1] PHI is essentially a duplicate of Medicare – holding PHI means that if a person chooses to receive treatment in a private (rather than public) hospital, they can claim against their policy (patients without PHI can pay out-of-pocket, but in practice, this is rare due to high costs). PHI's primary benefit is that it gives patients faster access to treatment in the case of elective surgery (like joint replacements). Patients using the public system are subject to priority waiting lists for such procedures, whereas patients using the private system are typically treated much faster. Additionally, private patients can sometimes select their specialists, and may prefer the amenities in private hospitals.

[1] Many people also purchase PHI that covers out-of-hospital expenses that are not covered by Medicare, such as dental, optometry and physiotherapy. For the purpose of this paper, PHI is taken to mean cover for private hospital care.

PHI premiums adhere to community rating, with uniform pricing for the same policy within each state and territory regardless of personal characteristics, including age (ignoring discount and penalty policies), gender, and prior health conditions. This means young people with low risk pay above-actuarially fair rates, while older (high-risk) people pay below. Hit-and-run purchasing behaviour is managed through mandatory no-claim periods for new purchasers.

2.1 Age-based discount policy

As discussed earlier, from around 2015 until the onset of the COVID-19 pandemic, PHI membership was declining, accompanied by a steady increase in the average age of insureds. Introduced on 1 April 2019, age-based discounts aimed to arrest this trend by making health insurance more affordable for young individuals and incentivising them to obtain coverage earlier in life. This initiative enables insurers to potentially reduce premiums by up to 10% for people aged 18 to 29 who hold PHI. The discount is applicable annually until the policyholder reaches the age of 41, after which it reduces by 2% each year until it is fully phased out, provided the policyholder continuously maintains their cover (Commonwealth Ombudsman., 2021). For example, a 25-year-old paying $1,500 annually for PHI could save $150 per year with this discount until he or she turns 41. The size of the discount varies depending on the age at which the person first purchases PHI. If they purchase between ages 18-25, a 10% discount is granted, then the discount steps down in 2% intervals for purchases between ages 26-29 (i.e., 26 = 8%, 27 = 6%, 28 = 4%, 29 = 2%). This discount applies to both new policyholders and those who purchased before 1 April 2019. However, the discount is not available to those covered as a dependent under a family or single-parent policy, and it is averaged between the individual discounts of the two adults in cases of a couple or family policy. It is also not mandatory – insurers retain the discretion to offer it or not, although in practice, this has been widely implemented (discussed further below).

Since its introduction, the total number of people receiving discounts in the target age groups has remained fairly stable (Figure 2a), although the total number receiving discounts has grown steadily due to people retaining their discount if they maintain cover into their 30s. Most people receiving discounts are receiving a discount rate of 10% (indicating they purchased insurance when aged 18-25 years) (Figure 2b). This cohort has steadily grown, with the proportion rising from 30% initially to 45% by Q4 of 2023 (Figure 2b).

Figure 2: Trend in people receiving age-based discount over time

(a) Number of persons by age

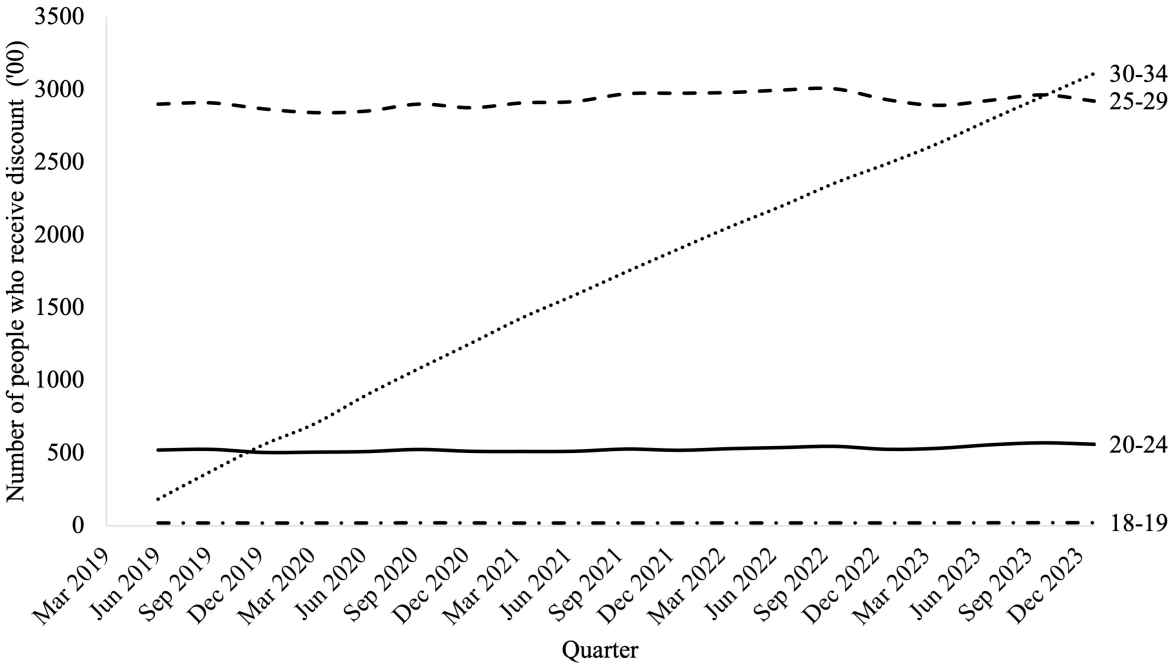

(b) Trend in discount rate received

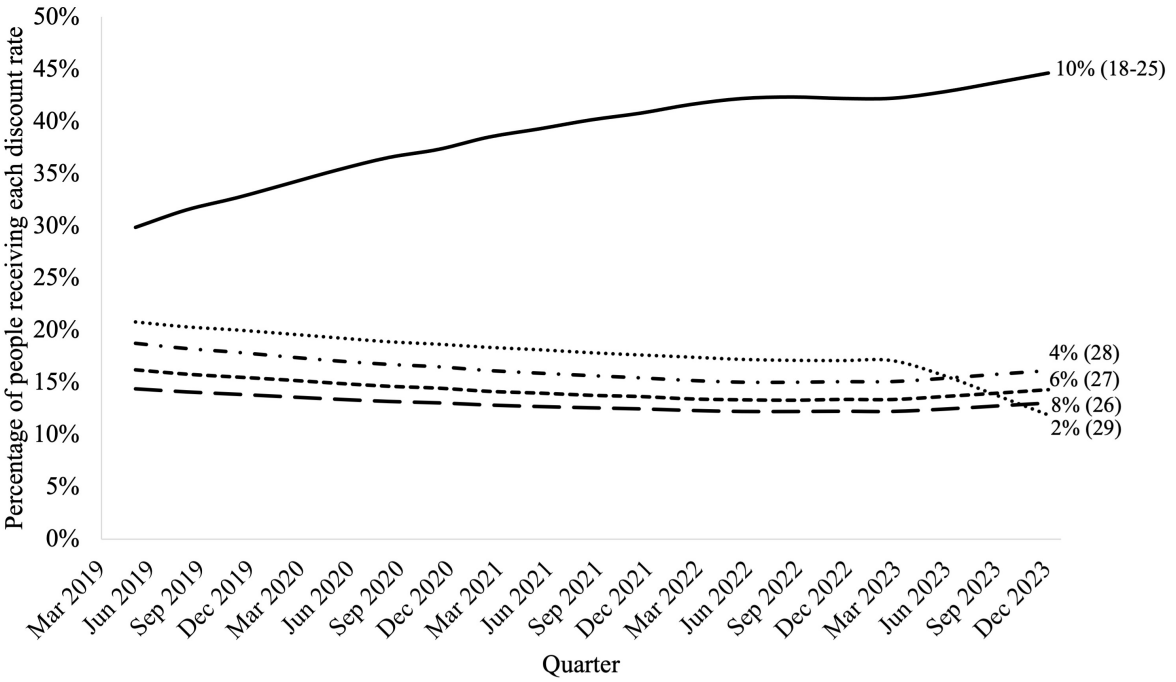

Source: Data on the number of people receiving discounts is from *Private health insurance reform data quarterly trends report* released on 15 April 2024 (Department of Health and Aged Care., 2024). Percentage in Figure 2b is calculated by the number of people with a certain discount (i.e. 10%) divided by the total number of persons who receive an age-based discount.

As a percentage, approximately 60–70% of insured persons aged 25–29 receive discounts and 10% of people aged 20–24 years (Figure 3). There are two main reasons why these numbers are not 100%. First, some people aged 20–29 are covered under their guardian's PHI plan as a dependent.[2] This is particularly the case for those in the 20–24 age bracket. The second reason is that not all insurers offer discounts since the policy is voluntary. However, the fact that the majority of people aged 25–29 are receiving a discount indicates that most insurers have adopted this policy.

Figure 3: Age-based discount by age cohort
(Percentage of insured people receiving discount by age group)

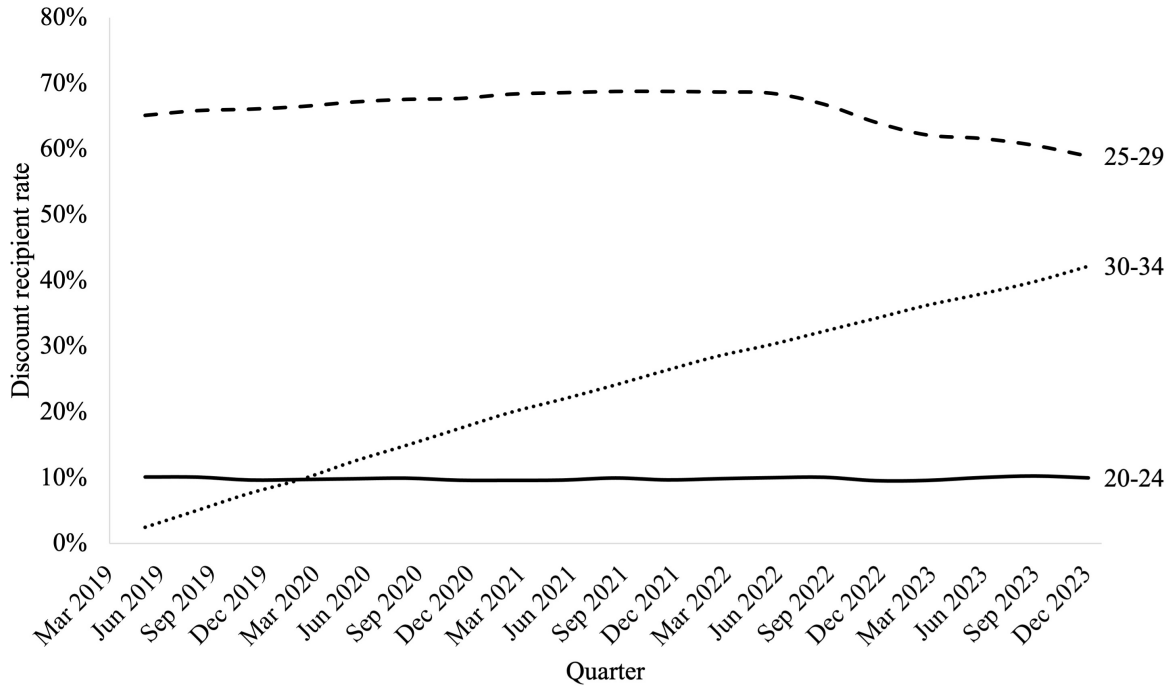

Source: The number of persons who receive a discount is the number of people receiving discounts from *Private health insurance reform data quarterly trends report* released on 15 April 2024 (Department of Health and Aged Care., 2024). The number of insured population by age cohort is from APRA *Quarterly private health insurance statistics* issue released on 28 February 2024 (Australian Prudential Regulation Authority, 2024b). Age-based discounts for 18–29 years are applied to age groups 20–24 years, 25–29 years and 30–34 years.

In 2021, a set of reforms expanded the legal definition of a dependent, most importantly by raising the maximum age from 24 to 31 years (though insurers can choose to

[2] The current definition of a dependent includes: (1) Dependent child: 0–17 years old and does not have a partner; (2) Dependent student: 18–31 years old, does not have a partner, is receiving a full-time education and as defined in the rules of their private health insurer; (3) Dependent non-student: 18–31 years old, does not have a partner, is not receiving a full-time education and as defined in the rules of their private health insurer; (4) Conditional dependent non-student: a dependent non-student that has their own general treatment policy with the same insurer covering them for hospital cover; (5) Non-classified dependent person: 18–24 years old, does not have a partner, is not a dependent student or non-student by the rules of their private health insurer; (6) Dependent person with a disability: 18 years old and over, may have a partner and is a participant in the National Disability Insurance Scheme or is a 'person with a disability' as defined by the rules of their private health insurer.

implement their own age range for dependent students and non-students from 18 years up to 31 years). The effect of this change is evident in Figure 3, where from around Q2 2022 there is a declining share of people aged 25–29 years receiving discounts (the time that most insurers actually implemented the measure). Some young adults who would purchase PHI independently and qualify for the discount policy instead remain covered as dependents. The impact of the expanded definition of a dependent can also be seen in Figure 4, which shows that the fraction of policies that include cover for dependents has risen slightly since the second half of 2022. Given this reform confounds the age-based discount policy, our evaluation separates treatment effects prior to and after this policy change (in effect, the treatment after this change is discounts plus expanded dependency eligibility).

Figure 4: The percentage of insured people with policies that cover dependents

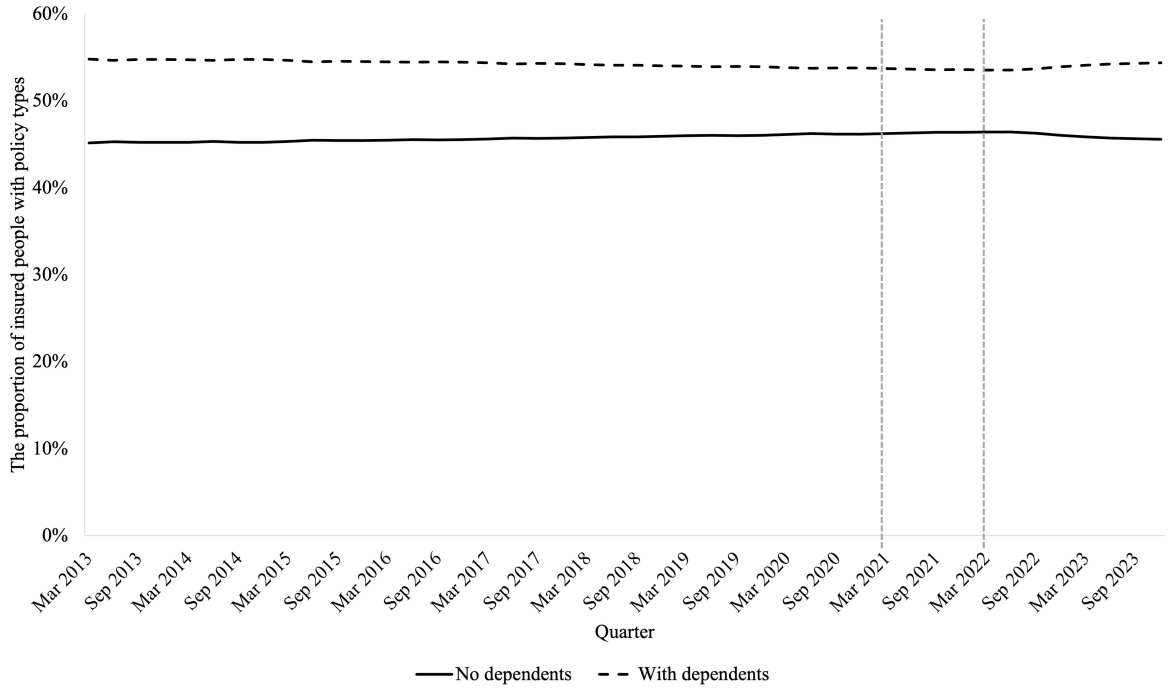

Notes: APRA dataset only includes 6 categories: single, family, single parent, couple, 2+ persons no adults and 3+ adults. The figure groups single, couple and 3+ adults into "No dependents"; and groups family and single parent into "With dependents"; 2+ persons no adults is a policy type for children only, which is not a relevant group to our evaluation. Vertical dot line for Q1 2020 represents that COVID-19 starting point, and Q1 2022 is approximately when the definition of a dependent changed.
Source: The data are from APRA *Quarterly private health insurance statistics* issue released on 28 February 2024 (Australian Prudential Regulation Authority, 2024b).

To give a rough sense of the incentive effects of discounts, Figure 5 illustrates (i) discount amounts and (ii) the disparity between average hospital benefits and premiums for insured individuals aged 20–24 and 25–29 annually from 2013 to 2022. The average net benefits of PHI are strongly negative for both age groups and have broadly increased over time, with some recent fluctuations during the COVID-19 pandemic likely associated with delayed surgeries in 2020. The premium discounts offered are comparatively small in contrast to the significant gap between expected benefits and premiums for this

demographic. Consequently, the discounts, while not trivial, still keep premiums well above actuarially fair amounts for this demographic. This can explain why insurers have been happy to apply this policy, but may also limit the effectiveness of discounts in incentivising early uptake.

To close this section, it is worth briefly mentioning other incentive policies for PHI in Australia, which contribute to the high rate of PHI coverage and operate alongside the discount policy. First, there is a means-tested premium rebate (subsidy) that sits on top of age-based discounts. The rebate amounts change each year, but the maximum is 24.6% in 2024.[3] Second, there is a tax penalty mandate, known as the Medicare levy surcharge, that applies a tax penalty for higher income earners without PHI. Third, there is an age-based penalty, known as Lifetime Health Cover loading, which allows insurers to charge an additional 2% premium loading for each year a person was uninsured after age 31. Importantly, these policies did not change during the evaluation period for our study.

2.2 COVID-19 pandemic

The first reported case of COVID-19 in Australia was 25 January 2020, and on 22 March 2020 the country initiated a national "lockdown" to curb the spread. The impact of the pandemic was highly disparate across Australia and varied over time. Apart from those aged 85 and older, the rate of confirmed COVID-19 cases was higher for those aged 20–24 and 25–29 in 2020 (168 and 183 per 100,000 people, respectively) than other age groups, which can be attributed to the associated restrictions on movement and disruptions to the labour markets (Australian Institute of Health and Welfare., 2021).

From an evaluation perspective, the pandemic is both a challenge and an opportunity. On the one hand, since there was no change in the discount policy during the pandemic, it will be informative to see whether the policy was differentially effective during a period of "health emergency", which saw many Australians entering the PHI market. On the other hand, differences in the way the pandemic affected, for example, employment and earnings across age groups, mean that we need to be cautious about attributing effects solely to the discount policy. We revisit this distinction later.

[3] This is increased by 5 and 10 percentage points for those aged 65 and 70 years and older respectively.

Figure 5: Net benefits and discount amount on hospital treatment

(a) Age 20–24

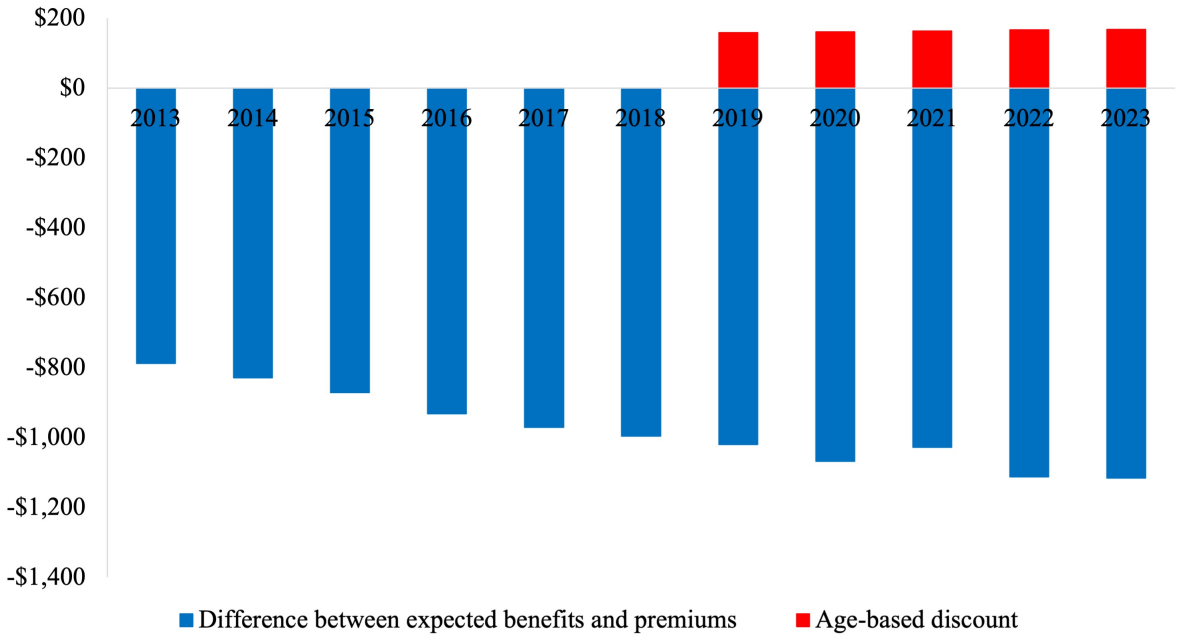

(b) Age 25–29

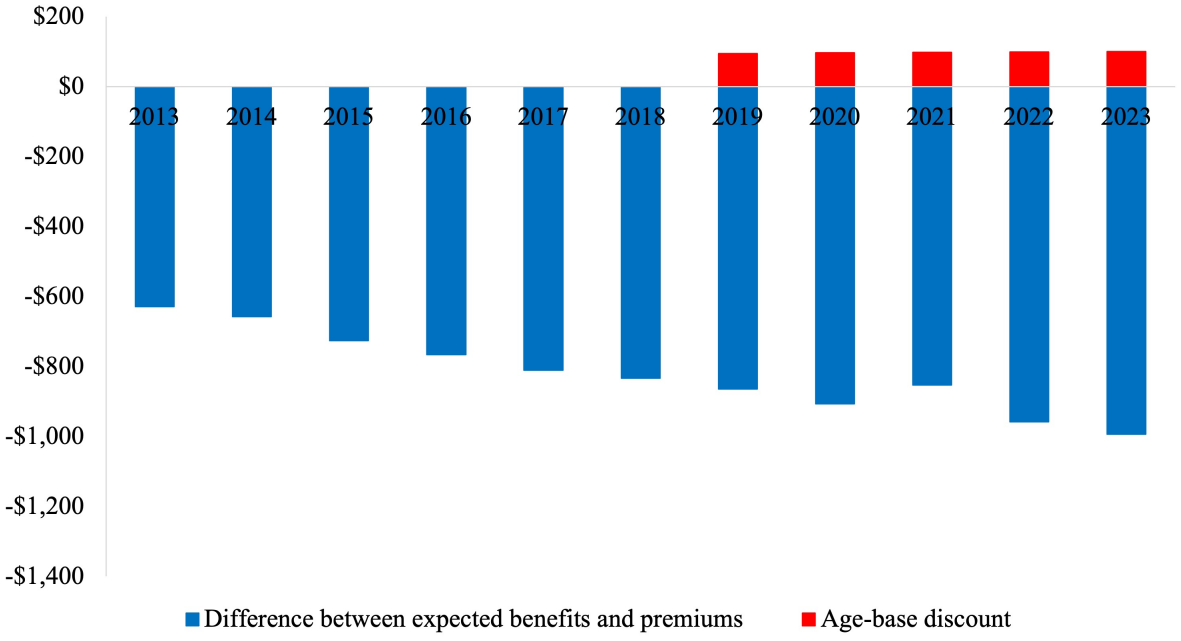

Source: The premium for hospital coverage is assuming people in those age groups pay the same average premium as all population, from APRA *Operations of private health insurers annual report* data, which is $\frac{Premium\ revenue\ in\ hospital\ treatment}{Insured\ people\ in\ hospital\ treatment}$ (Australian Prudential Regulation Authority, 2024a). The benefits for hospital treatment is from APRA *Quarterly private health insurance statistics* issue released on 28 February 2024 (Australian Prudential Regulation Authority, 2024b).

3 Data and Methods

3.1 Data

Our main datasets are the Quarterly Private Health Insurance Statistics published by the Australian Prudential Regulation Authority (APRA) and Australian Demographic Statistics published by the Australian Bureau of Statistics (ABS). The APRA data comprises quarterly information submitted by all private health insurers in Australia on memberships and other details (Australian Prudential Regulation Authority, 2024b). One downside of the APRA data is that membership statistics are aggregated into five year age groups and contain only limited details on members. However, a major advantage of these data is that they cover the entire population. The data are also public and likely to be extremely accurate. The ABS data provides quarterly estimates of population size by age group which allows us to derive the rate of insurance for each age group (number of insured persons/population size by age) in each quarter.[4]

3.2 Methodology

Our identification strategy exploits the temporal aspect of our data and relies on the fact that the age-based discount affected the incentives to insure for younger people but not for older people. More specifically, we classify two treatment age groups – those aged 20–24 years and those aged 25–29 years.[5] Those aged 20–24 qualify for a discount of 10%; those aged 25–29 will receive a discount of 2%–10%, depending on at which age they first insure. Because not everyone in our treatment age groups is entitled to discounts, our estimates can be considered intention-to-treat.

In settings like ours, the traditional difference-in-differences (DID) method (Angrist and Pischke, 2009; Ashenfelter and Card, 1984; Bertrand et al., 2004) and the synthetic control (SC) method (Abadie et al., 2010; Abadie et al., 2015; Abadie and Gardeazabal, 2003; Abadie and L'hour, 2021) are commonly used to link observed data with counter-

[4] Quarterly estimated resident population is calculated by taking the population estimate at the start of the quarter and adding natural increase (births minus deaths), net overseas migration and (in the case of state/territory populations) net interstate migration. These calculations are performed separately for each age-cohort and sex (Australian Bureau of Statistics, 2020b). The dataset is linked to ABS Data Explorer (BETA) by searching Quarterly Estimated Resident Population: https://explore.data.abs.gov.au/.

[5] Discounts are also available for those aged 18–19 years. However, since APRA aggregate membership statistics into 5-year age groups, we only observe outcomes for the cohort 15-19, which mostly comprises dependent children who are not eligible for the discount.

factuals. The DID method is most suitable where there are numerous units treated by the policy, comparing differences in trends between treated and pre-specified control units and relying on the assumption of parallel trends between those groups. The SC method is helpful when only one (or a few) units are treated, selecting a weighted average of control units based on pre-trends and (optionally) other predictors to serve as the counterfactual.

Recently, an approach that generalises DID and SC has been proposed, known as synthetic difference-in-differences (SDID) (Arkhangelsky et al., 2021; Clarke et al., 2023). SDID integrates the strengths of both DID and SC methods. For example, it allows treatment and control units to exhibit trends at different levels before the reform (like DID), but the counterfactual is a weighted average of control units (like SC). From a practical perspective, the features of SDID make it particularly attractive for our setting, and this method is likely to be more efficient than alternatives (Arkhangelsky et al., 2021). SDID optimises the creation of matched control units to relax the need for strict parallel trend assumptions, but unlike SC does not require treatment units to fall within the "convex hull" of the control units (i.e., for the levels of the outcome variable (PHI coverage) not to be at the upper or lower extreme). This assumption is not met in our setting since those aged 20–29 always have the lowest rates of insurance. Meanwhile, SDID avoids the problematic temptation in DID to specify the control unit(s) based on pre-testing of parallel trends (Roth, 2022). Inference in DID settings with only one or a few treated units is challenging because of the need to cluster errors at the unit level (Bertrand et al., 2004). In addition, SDID lends itself to an intuitive randomisation inference procedure based on the distribution of placebo treatment effect estimates,[6] but has the potential to be more efficient because it relies on weaker criteria for constructing a counterfactual. For all these reasons, we use SDID to isolate the causal effect of the policy in our analysis.

Adopting the notation in Arkhangelsky et al. (2021), the SDID estimator is given by:

$$(\hat{\tau}^{sdid}, \hat{\mu}, \hat{\alpha}, \hat{\beta}) = \underset{\tau,\mu,\alpha,\beta}{\operatorname{argmin}} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} \left(Y_{it} - \mu - \alpha_i - \beta_t - D_{it}\tau\right)^2 \hat{\omega}_i^{sdid} \hat{\lambda}_t^{sdid} \right\} \tag{1}$$

In equation (1), Y_{it} is the PHI coverage rate for the treated unit i at quarter t.

[6] The main idea of such placebo test is to consider the behaviour of SC estimation when replacing the unit that was exposed to the treatment with different units that were not exposed, which is closely connected to permutation tests in randomisation inference (Arkhangelsky et al., 2021; Abadie et al., 2010, 2015).

$N = 13$ (including one treated group, either 20–24 or 25–29), with control age groups spanning 30–34 to 85-89.[7,8] $t \ldots T$ is Q1 2017, Q2 2017, ..., Q2 2023. α_i and β_t are unit- and time-fixed effects, respectively. D_{it} is a binary indicator that takes the value 1 if observation i is treated by time t and 0 for unit i if it is untreated at time t. Our target parameter τ represents the causal effect of the treatment, and the results are in Appendix Table A1.

What separates SDID from DID and SC is the inclusion of all two-way fixed effects, unit and time weights in equation (1). The presence of unit-fixed effects allows for any constant differences between treatment and control units, because SDID only seeks to match treated and control units on pre-treatment trends. $\hat{\omega}^{sdid}$ are unit weights (shared with the SC estimator) that are selected by the SDID algorithm to ensure treated units and controls roughly follow parallel trends prior to intervention (Arkhangelsky et al., 2021). The selection of time weights $\hat{\lambda}^{sdid}$ (omitted in SC estimator) helps to place more weight on pre-treatment periods that are more similar to post-treatment periods, in the sense of finding a constant difference between the post-treatment average of each control unit and pre-treatment weighted averages across all selected controls (Clarke et al., 2023). The selected optimal weights in our application are shown in Appendix Table A2.

Equation (1) assumes a time-invariant treatment effect (τ), but in practice, the treatment effect might vary over time. In particular, two significant events exist in the post-treatment period: the COVID-19 pandemic (Q2 2020-) and the expansion of the definition of a dependent (Q2 2022-). This motivates estimating a panel event-study version of our SDID model, which also allows us to explore whether there are any differences between treated and control units prior to the intervention. This is also helpful for evaluating parallel trends. Specifically, for each period, omitting i subscripts, this estimator can be achieved by calculating $(Y_t^{treated} - Y_t^{control}) - (Y_{baseline}^{treated} - Y_{baseline}^{control})$ for each t, in our case setting the period prior to the treatment as the baseline. Clarke et al. (2023) discusses how confidence intervals for each estimate can be constructed using a block bootstrap procedure.

We estimate the SDID model using the 'sdid' package for Stata (Clarke et al., 2023).

[7] We exclude older age groups since they are much smaller and more volatile in their trends. We also exclude age groups aged under 20 since they are primarily dependent children whose insurance status is tied to parents represented in other age groups.

[8] Those in the 30–34 age group are a potentially contaminated control unit since over time, some people who join the insurance pool at age 25–29 will age into this category. Given only a few years of post-treatment data, and this being only one control unit, we expect any bias from this to be small (and in the direction of zero), so retain this age group as a control in our main analysis. In the sensitivity analysis presented later we show that excluding this group has little effect on our estimates.

For inference, we use the randomisation inference approach based on placebo estimates described in Algorithm 4 of Arkhangelsky et al. (2021).

3.3 Graphical analysis

Before our formal analysis, we start by comparing trends in PHI coverage rates between those in our treatment groups and those aged 35–39, who are the most proximate age group to our treatment ages without being adjacent, and intuitively might be the most reasonable control group in a classic DID design (see Figure 6).

Figure 6: Trends in PHI coverage: treatment groups versus ages 35–39

(a) Treated: 20–24

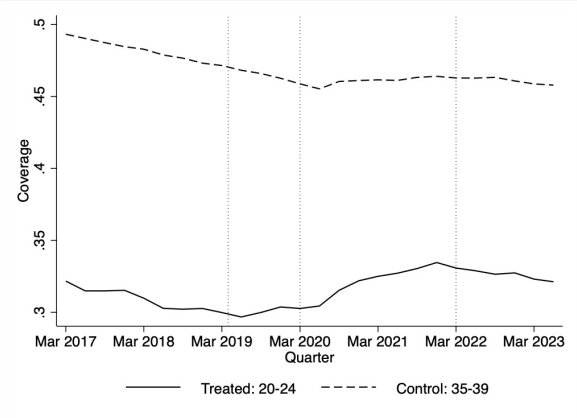

(b) Treated: 25–29

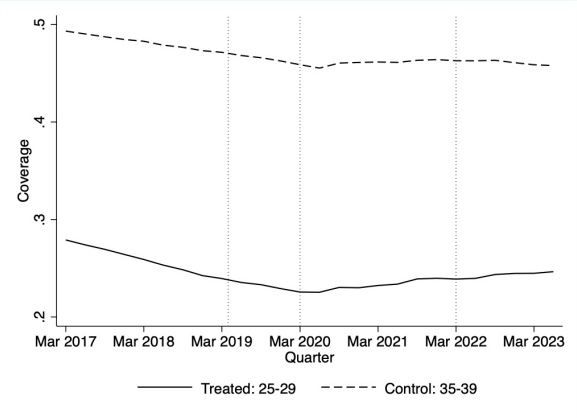

Notes: Vertical dotted lines mark three key events: the introduction of the age-based discount in April 2019, the onset of COVID-19 after Q1 2020, and approximate definition of a dependent changed after Q1 2022.

Figure 6 shows weak evidence of an increase in PHI coverage for the 20–24 years group in the period between treatment and COVID-19, and then a larger effect in the period after the pandemic. There is potentially a similar pattern for the 25–29 years group, although the trends in the pre-treatment period are not parallel, highlighting that traditional DID is insufficient in this setting.

4 Results

4.1 Baseline estimates

Table 1 presents SDID estimates for our two treatment groups, considering different post-treatment periods. Graphs illustrating the trends for both the treated and synthetic

control groups (using the full post-treatment period specification) are in Figure 7.

Table 1: Baseline SDID estimates

Sample window	20–24	25–29	Obs.
2017Q1 – 2020Q1	0.0050	-0.0004	169
	(0.0034)	(0.0036)	
2017Q1 – 2022Q1	0.0210***	0.0036	273
	(0.0071)	(0.0072)	
2017Q1 – 2023Q2	0.0232**	0.0067	338
	(0.0100)	(0.0088)	

The first sample window (2017Q1 – 2020Q1) ends before the onset of COVID-19. The second sample window (2017Q1 – 2022Q2) ends before the dependent definition change around Q1 2022. The final specification uses the full sample (2017Q1 – 2023Q2). Standard errors are reported in parentheses. ***, **, * = statistically different from zero at the 1%, 5%, and 10% level, respectively. Observations are at the cohort-quarter level, and standard errors are based on a placebo inference procedure (Arkhangelsky et al., 2021).

Figure 7: Trends in coverage over time for treated age and the relevant weighted average of control ages

(a) Treated: 20–24

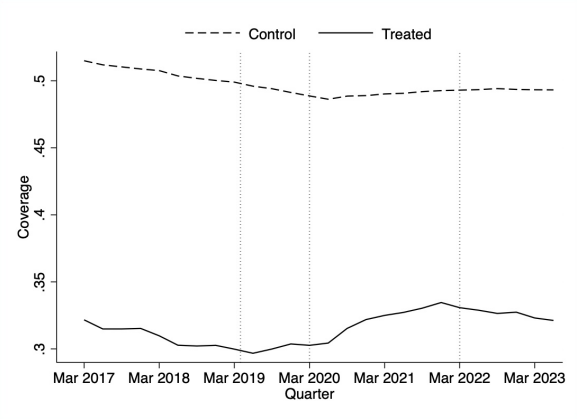

(b) Treated: 25–29

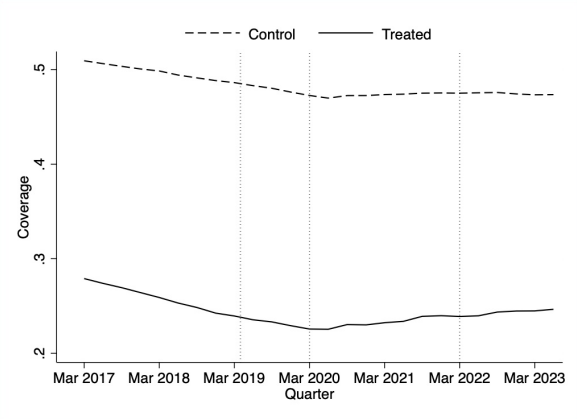

Notes: Vertical dotted lines mark three key events: the introduction of the age-based discount in April 2019, the onset of COVID-19 after Q1 2020, and approximate definition of a dependent changed after Q1 2022. The shaded area represents the time weights assigned to the counterfactual trend for the treated group in the SDID estimator.

There is no evidence that PHI coverage increased in either treatment group in the period prior to COVID-19. The point estimates are small in absolute value (0.50 percentage points (ppts) for 20–24 and -0.04 ppts for 25–29) and not significantly different from zero. When we extend the period to include after the pandemic, we estimate that those in the 20–24 treatment group were 2.3 ppts more likely to insure, while we continue to find no statistically significant evidence the policy increased PHI coverage for those aged 25–29. These estimates suggest the discount policy had limited (negligible) short-term effectiveness, with larger policy impacts emerging in the post-pandemic period.

4.2 Panel event study estimates

Next, we investigate dynamic treatment effects more formally by presenting panel event study estimates in Figure 8, which allows us to examine how PHI take-up evolved over time following the introduction of the policy. Our results suggest the parallel trends assumption does not hold for those in the 25–29 age group.

Figure 8: Dynamic Treatment Effect in SDID Model

(a) Treated: 20–24

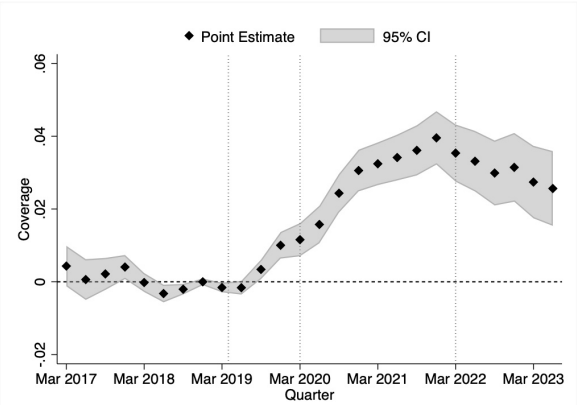

(b) Treated: 25–29

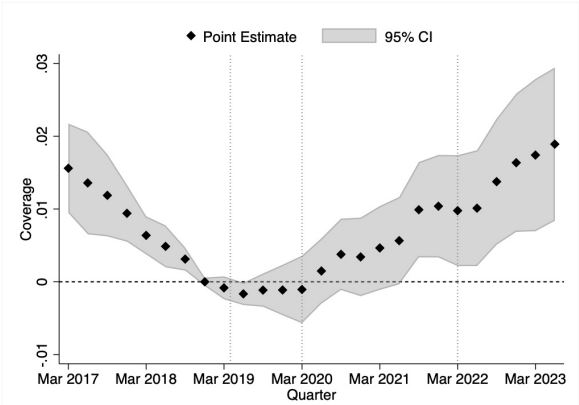

Notes: Vertical dotted lines mark three key events: the introduction of the age-based discount in April 2019, the onset of COVID-19 after Q1 2020, and approximate definition of a dependent changed after Q1 2022. Standard errors using the placebo method are calculated using bootstrap with 1000 replications. Confidence intervals are built by bootstrap estimation predictions using only the treated units to estimate the noise level.

In Figure 8a, we can see that the policy has a gradual effect on PHI take-up for those aged 20–24, reaching 4 ppts around December 2021 and gradually declining but remaining above 2 ppts afterwards. Estimates are all close to zero in the pre-treatment periods, which bodes well for the parallel trends assumption. Interestingly, while the standard SDID analysis estimated a small positive but statistically insignificant average effect across the post-treatment period to Mar 2020 (prior to the COVID-19 pandemic), the panel event study estimates suggest there was a significant effect but only in the last two quarters, with an effect size of approximately 1 ppt at the onset of the pandemic.

The results in Figure 8b for the 25–29 age group, in contrast to those in Figure 8a, show a clear, close to linear, downward trend in the pre-treatment period, casting doubt on the parallel trends assumption being satisfied. This result is unsurprising and explained by the fact that the 25–29 age group was experiencing the steepest decline in PHI demand prior to the policy, so a weighted average of other age groups is unlikely to follow the same path.

To deal with this, we propose adjusting the values of estimates and assume that parallel trends hold conditional on a linear trend difference between treatment and con-

Figure 9: Dynamic treatment effect with adjustment for ages 25–29

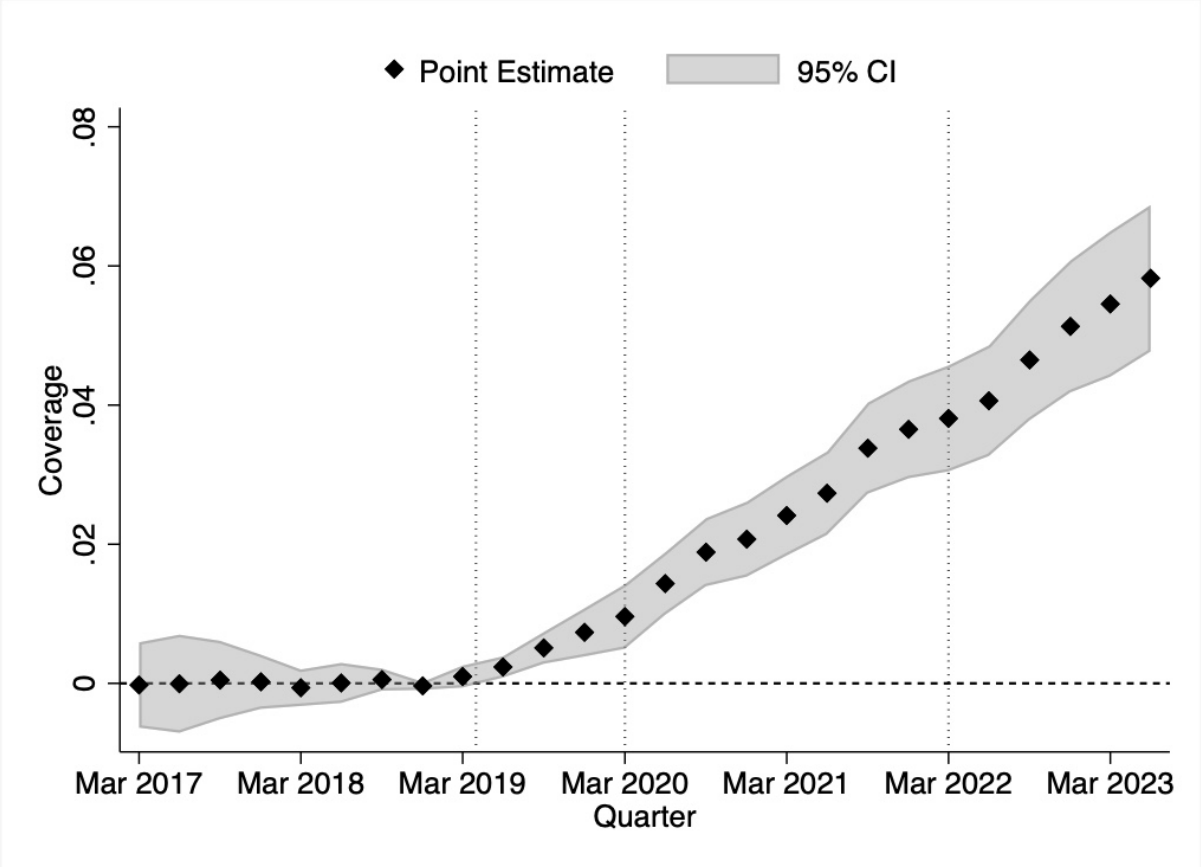

Notes: Vertical dotted lines mark three key events: the introduction of the age-based discount in April 2019, the onset of COVID-19 after Q1 2020, and approximate definition of a dependent changed after Q1 2022. Standard errors using the placebo method are calculated using bootstrap with 1000 replications. Confidence intervals are built by bootstrap estimation predictions using only the treated units to estimate the noise level.

trol in the pre-treatment period, as described in Freyaldenhoven et al. (2021) (see also Dobkin et al. (2018) for a similar application). The fact that the pre-treatment estimates are strongly linear supports this approach. To operationalise this, we first fit a linear regression through the pre-treatment estimates and then subtract the predictions from that regression line from the unadjusted estimates (in both the pre- and post-treatment periods). In effect, this rotates the event study plot to make the pre-treatment period estimates close to zero.

In Figure 9, we present the adjusted panel event study estimates for the 25–29 age group. Adjusting is not necessary for the 20–24 age group since the pre-trends are basically flat already, and so we do not impose additional assumptions for this group. Now, pre-trends for the 25–29 age group are much flatter, in fact, the estimates for each period are very close to zero. After controlling for the initial downward trend, the results suggest the policy did increase PHI take-up, with the treatment effect being small prior to the pandemic, but gradually increasing over the period, eventually reaching 6 ppts in

June 2023.

4.3 Robustness

To ensure the robustness of our findings, we conduct a series of sensitivity analyses. We find that key results hold consistently across all robustness checks. All figures presented in this section for the 25–29 age group use the trend adjustment described above.

4.3.1 Windows of pre-treatment periods

The selection of windows of pre-treatment periods in our analysis is somewhat arbitrary. In our main results, we only include data from 2017 onward since trends further in the past may poorly reflect trends today. However, including data from five years earlier (2013 onward) does not materially affect our estimates (see Appendix Figure A1), which suggests the results are not unduly sensitive to pre-treatment window selection.

4.3.2 Synthetic control selection

Since those aged 30–34 can become treated over time (i.e., people who initially lock-in discounts age into this age group), their inclusion in the donor pool for the SDID model may bias our estimates. To address this concern, we re-estimate our analysis after excluding this age group from the donor pool (Appendix Figure A2). This does not materially alter our findings, indicating that potential contamination from the 30–34 age bracket does not drive our results.

4.3.3 Interstate migration

Interstate migration may affect PHI take-up patterns, potentially influencing the magnitude of estimates. To control for this, we apply the SDID algorithm to residuals obtained after controlling for the fraction of insured people in age group i living in each Australian state and territory (see Clarke et al., 2023).[9] Our results remain qualitatively unchanged, as shown in Appendix Figure A3.

[9] There are eight states and territories in Australia.

4.3.4 Placebo test

Placebo tests are commonly employed in DID designs (Abadie and Gardeazabal, 2003; Bertrand et al., 2004; Kleven et al., 2019; Huntington-Klein, 2021). This is often done by acting as if the policy intervention occurred at an earlier period and estimating the model on this 'placebo' treatment. In our SDID setting, this is useful to help rule out the possibility that the estimation strategy systematically picks up some unusual dynamics in the trajectories of different age groups. If we estimate nil effects at placebo treatment dates, that can improve our confidence in the main results.

To operationalise our test, we act as if the discount policy was introduced in April 2016, April 2017 and April 2018. We estimate panel event-study SDID treatment effects for these placebo-intervention points. For each age bracket, we generate three comparative estimates corresponding to the three placebo periods. For consistency we keep the length of the pre-treatment period the same as in our main results for each placebo estimate, and make the same rotation adjustment for the 25–29 age group. The estimates are shown in Figure 10.

For the 20–24 age group, we often see strong downward pre-trends for the placebo years. We do not see any evidence of a systematic increase in coverage that coincides with any placebo thresholds, including if we rotate the estimates as we do for the 25–29 age group (see Appendix Figure A4). Nonetheless, we interpret the fact that pre-trends in the placebo years are different from those closer to the actual policy introduction as implying that we should not put too much weight on this analysis in the case of the 20–24 age group.

For the 25–29 age group, the evidence is clearer, with no sharp changes in trajectory at any of the placebo dates, and if anything placebo estimates are negative (but generally not different from zero).

4.4 Disentangling incentive and COVID-19 effects

Our results suggest much larger treatment effects in the period after COVID-19, which may be because there was a moderator effect between the pandemic and the discount policy, or it simply took time for people to respond to the incentive (delayed behavioural reaction), or the pandemic itself may have had different effects on PHI demand for different age groups. In this section we attempt to distinguish between these explanations.

Figure 10: Comparison dynamic SDID estimates in placebo tests with adjustment

(a) Treated: 20–24

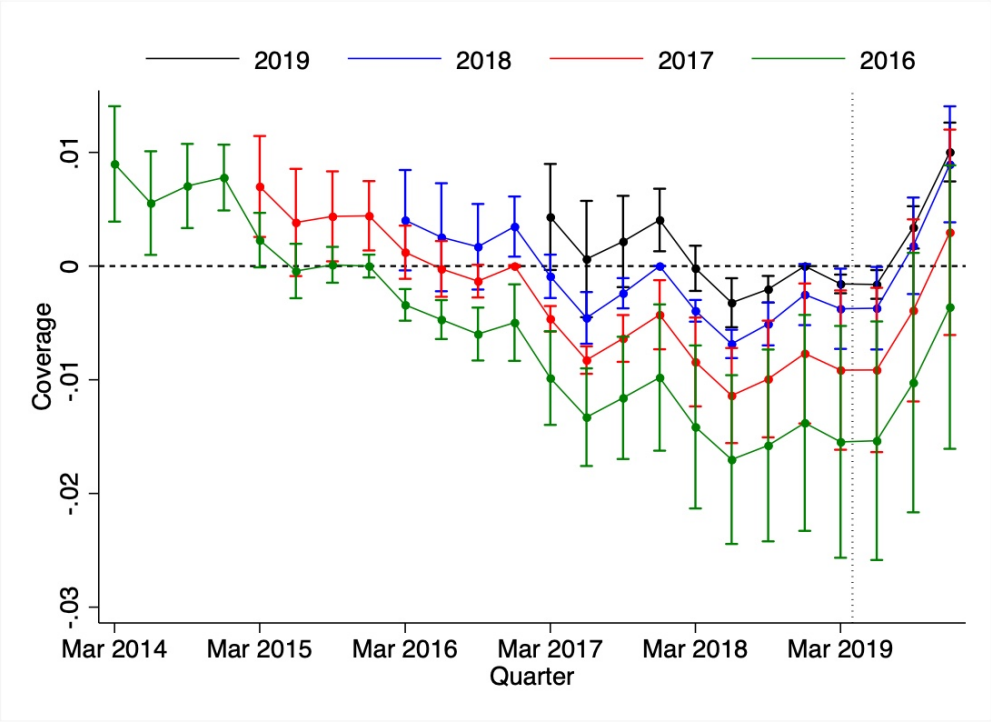

(b) Treated: 25–29

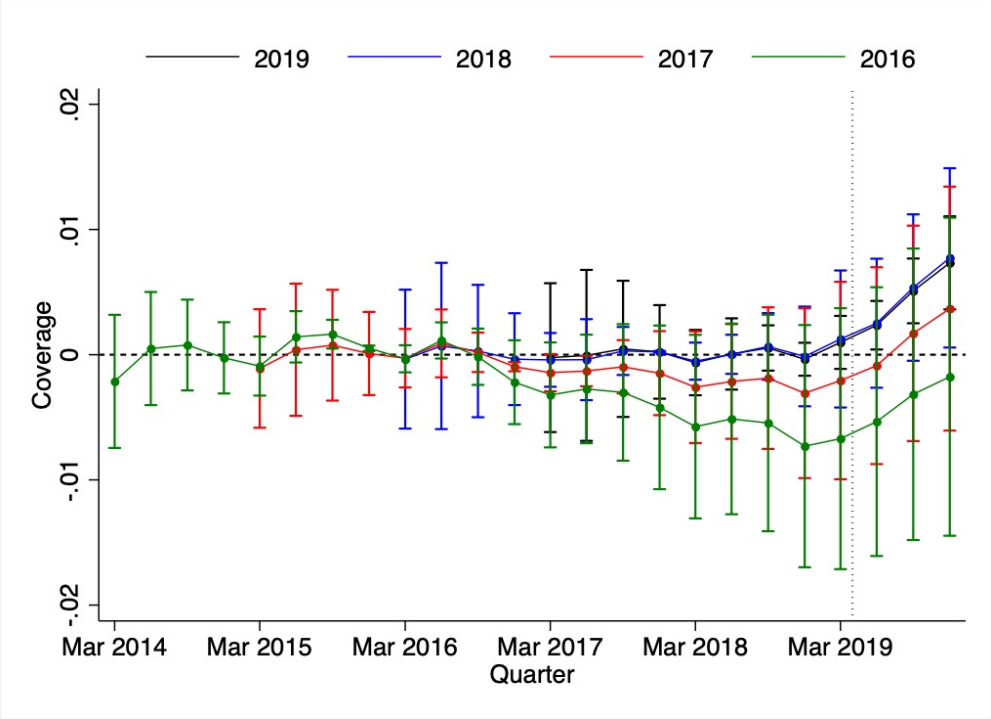

Notes: Vertical dot line for April 2019 represents when the age-based discount was introduced. Standard errors using the placebo method are calculated using bootstrap with 1000 replications. Confidence intervals are built by bootstrap estimation predictions using only the treated units to estimate the noise level.

Conceptually, some arguments can be made to suggest the effect of COVID-19 is at least in part moderation (or at least not a confounder). First, the pandemic saw an increase in PHI enrolment broadly across the population (Figure 1). With it, more people across all age groups interacted with insurers, and subsequently it is likely that awareness of the discount would have increased, spurring demand among the young. Another argument is that the treatment effect seems to have already begun rising in the year before the pandemic and continues on a similar trajectory during the pandemic, suggesting that COVID-19 reinforces or exaggerates rather than fundamentally alters the policy's effect. Next, to the extent there were economic differences in the experience of COVID-19 and associated lockdowns, these appear to have been borne disproportionately by the young, who experienced larger drops in employment at the outset of the pandemic than other age groups (Australian Institute of Health and Welfare., 2021). If concerns about the effect of the pandemic on the public healthcare system are the main reason to induce people to buy PHI, presumably the highest healthcare users (older individuals) would be driven most strongly by this, rather than younger people.

While a moderation effect is plausible, there are also dissenting arguments. For example, when demand for PHI picked up, a much larger fraction of younger people were uninsured than older people, and the effect may have been larger due to this lower baseline. Additionally, while younger people purchasing PHI for the first time in 2020 were encouraged by age discounts, older people were *discouraged* by Lifetime Health Cover loading (LHC). Recall that this policy allows insurers to charge an additional 2% premium loading for people who insure for the first time after age 31, with 2 ppts added for each year uninsured from 31, meaning someone aged 50 years and insuring for the first time would pay a 40% loading, while someone aged 25 would receive a 10% discount. This pricing structure may have discouraged PHI uptake among older cohorts while amplifying the relative effect of age-based discounts for younger individuals.

Additionally, COVID-19 severely affected international migration in Australia, which is important since younger people are over-represented among immigrants. New migrants are unlikely to buy PHI, so a reduction in immigrants will increase the coverage rate for the age group, as non-insured young migrants entered the denominator of the coverage rate calculation. Finally, the expanded age of dependency policy may explain some effects from 2022. Recall that the age of dependency reform allowed more young adults to remain on family PHI policies instead of purchasing individual coverage. Since dependents can be added without cost to family PHI plans, it is likely this policy saw some young people (in the 25–29 age group) become insured who would have otherwise not purchased individual cover for themselves.

We now investigate each of these possibilities empirically.

4.4.1 LHC loading effect

We test whether the discouragement effect of LHC loading for our older control age groups drives our results in two ways. First, we act as if people aged 30–34 and 35–39 are the treatment group and use age groups 40–89 as controls. Second, we use only the most proximate age group with no policy contamination – those aged 35–39 – as the sole control group and estimate dynamic treatment effects using classic event study DID, and compare to our main results.

The logic for our first test is as follows. Assume the supply curve for PHI is perfectly elastic at the common market price and demand is downward sloping. There are two groups, no LHC and liable for LHC, and they have the same demand curve and start from the same equilibrium quantity of PHI, and both experience the same shift in demand. For the latter group, the supply curve effectively shifts up at the same time as the shift in demand, which means that the new equilibrium is at a lower quantity than for the former group. This suggests testing whether coverage increased by more for age groups subject to no or low LHC compared to those subject to high LHC after the pandemic as a way to test whether this effect is present in our estimates. We can only causally identify an LHC loading effect if we assume there are no other age-related reasons why demand would shift disproportionately for different age groups. This is obviously strong, but nonetheless, if we see an effect for these younger but untreated groups, that would bring into question the attribution of the post-COVID effects, regardless of mechanism.

We test for an LHC loading effect by setting 30–34 and 35–39 (done separately) as the treatment group and 40–89 as controls, and then estimate our panel event study SDID model. Those in the 30–34 age group are subject to a maximum 8% loading, and those in the 35–39 group 18%, while LHC is higher in the control groups, up to a maximum of 70% for those 65 or older.[10] We set the treatment date as April 2019 to see whether there is an effect prior to the pandemic (this is a kind of placebo test). The results are in Figure 11.

The estimates in the period after the discount policy but prior to the pandemic are small and not significantly different from zero, which provides some reassurance for the positive effects we estimate in this period being due to the discount policy. However,

[10] Because some people aged 30–34 become eligible for discounts over time, estimates for this group should be interpreted with caution as they may partly reflect a response to the discount policy.

Figure 11: Dynamic Treatment Effect in SDID Model with LHC effect

(a) Treated: 30–34

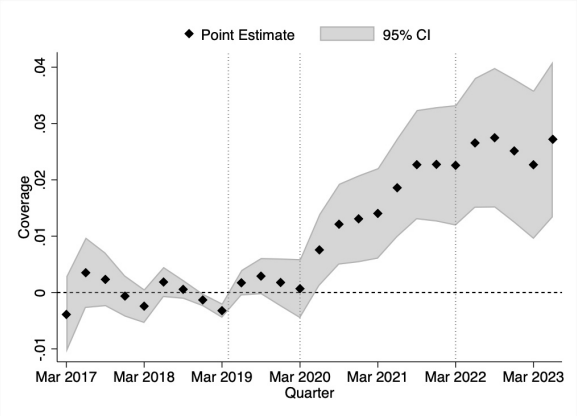

(b) Treated: 35–39

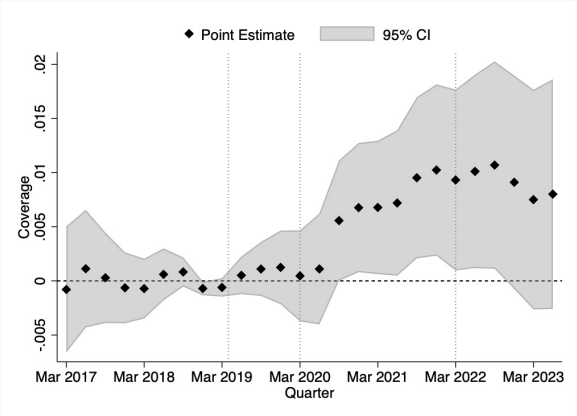

Notes: The same rotation adjustment maintained for the 25–29 years treatment group is applied to the 30–34 and 35–39 placebo treatment groups, since we observed similar linear pre-trends. Vertical dotted lines mark three key events: the introduction of the age-based discount in April 2019, the onset of COVID-19 after Q1 2020, and approximate definition of a dependent changed after Q1 2022. Standard errors using the placebo method are calculated using bootstrap with 1000 replications. Confidence intervals are built by bootstrap estimation predictions using only the treated units to estimate the noise level.

we find that there was stronger growth in the younger age cohorts after COVID-19 than older cohorts. The magnitudes are lower than our estimates for the actual treatment age groups, but nonetheless, bring into question the elevated treatment effect estimates post-COVID-19 being solely due to the discount policy.

Next, we use 35–39 as the *only* control group[11] and compute $(Y_t^{treated} - Y_t^{control}) - (Y_{baseline}^{treated} - Y_{baseline}^{control})$ to map out the treatment effect at each quarter, comparing this to our SDID estimates. March 2019 is set as the baseline. Figure 12 compares the SDID and DID estimators. Despite the fact that our previous analysis suggested we may be overestimating effects due to LHC, the SDID and DID estimates are very similar.

Overall, the evidence is mixed as to whether the effects we estimate post-COVID-19 are due to older age groups having a smaller demand response to the pandemic, potentially due to LHC. But there is some evidence this could be a factor.

4.4.2 Migration effect

One of most significant issues during the COVID period was migration disruptions due to border closures and associated lockdown actions.[12] This reduced the large number

[11] We use 35–39 since this is the most proximate age group that was not directly affected by the discount policy during the analysis period.

[12] Starting from early March 2020, Australian Home Affairs successively implemented inward travel restrictions on foreign nationals entering Australia from Iran, South Korea, and Italy to any country, with restrictions gradually easing from late April. Meanwhile, outward travel restrictions were also imposed on Australians travelling abroad (Australian Institute of Health and Welfare., 2021;

Figure 12: Comparison in SDID and DID with LHC effect

(a) Treated: 20–24

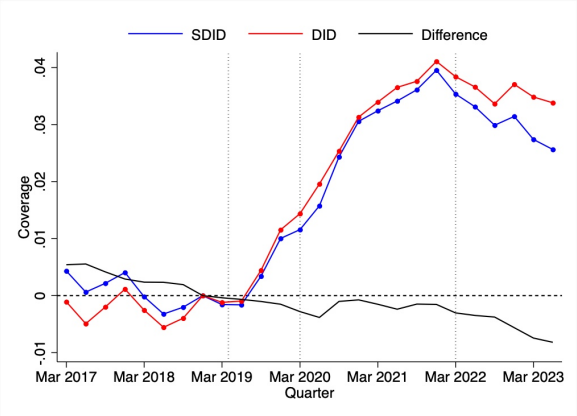

(b) Treated: 25–29

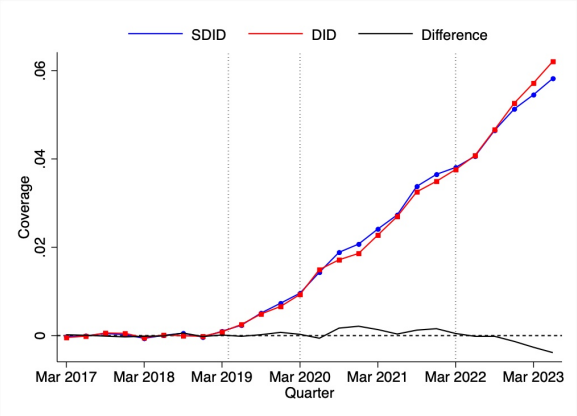

Notes: The rotation adjustment was only maintained for the 25–29 years treatment group. Vertical dotted lines mark three key events: the introduction of the age-based discount in April 2019, the onset of COVID-19 after Q1 2020, and approximate definition of a dependent changed after Q1 2022.

of younger migrants regularly coming into the country – the usual proportion of net overseas migrants aged 20–24 and 25–29 being around 25% and 15%, respectively (Australian Bureau of Statistics, 2024).[13] Since young migrants are unlikely to purchase PHI (due to having to purchase separate overseas visitors health cover), their absence from the population may have artificially increased the rate of coverage in the younger age groups relative to older ones.

To eliminate the effect of migration, we apply bounding analysis, in the same spirit as Lee bounds (Lee, 2005). We estimate the effect by simulating extreme (counterfactual) migration patterns under the assumption that COVID-19 never occurred. We use an ARIMA model to forecast the number of arrivals and departures of migrants by extrapolating pre-pandemic trends.[14] This means migration flows under a no-pandemic scenario

Auditor-General Report No.12 2021–22, 2022).

[13] The number of arrivals and departures of migrants comes from Australian Bureau of Statistics (2024). The detailed quarterly age-specific net overseas migration data are extracted from the ABS TableBuilder platform using the 'Net Overseas Migration' dataset. Access to TableBuilder is generally available at https://www.abs.gov.au/statistics/microdata-tablebuilder. The custom table includes 5-year 'Age' in rows, and 'Direction of migration' crossed with quarterly 'Reference period' (from 2013 onward) in columns. See: https://www.abs.gov.au/statistics/microdata-tablebuilder/available-microdata-tablebuilder/net-overseas-migration.

[14] Creating 'Net Overseas Migration' dataset via ABS TableBuilder platform by constructing a quarterly age-specific series, then estimating separate ARIMA models for arrivals and departures. The models are fitted using data from Q1 2013 to Q4 2019 (exclusive) to dynamically generate out-of-sample forecasts for Q4 2019 to Q4 2021 (exclusive), with the first-difference series depending on the first lag of the error term only, equivalently a random walk with MA(1) errors (optionally with drift) (equivalently to ARIMA(0,1,1)). Then we combine the observed data (prior to Q3 2019 and after Q1 2022) with the predicted values to form a complete series for use in SDID estimation. One limitation of the migration data is that people who enter/leave multiple times in a quarter can be counted multiple times.

Figure 13: Comparison in with and without migration effect in SDID

(a) Treated: 20–24

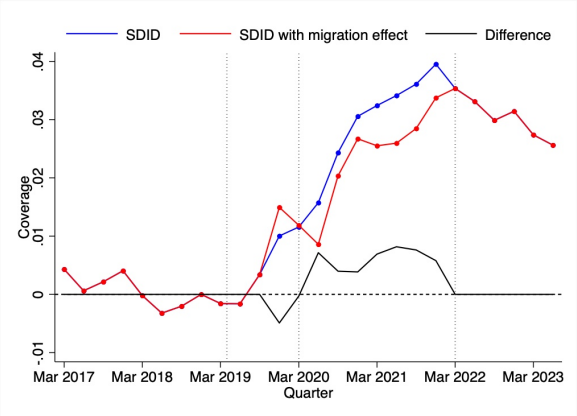

(b) Treated: 25–29

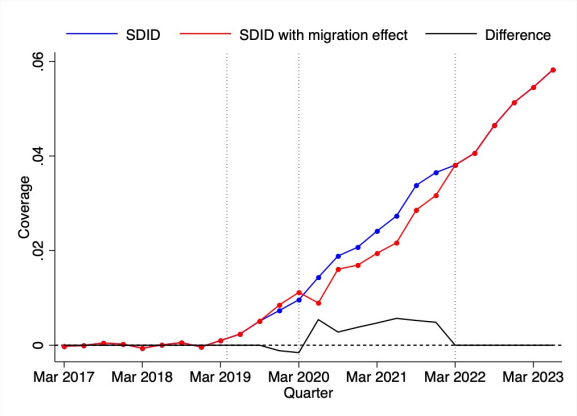

Notes: Vertical dotted lines mark three key events: the introduction of the age-based discount in April 2019, the onset of COVID-19 after Q1 2020, and approximate definition of a dependent changed after Q1 2022. The rotation adjustment was only maintained for the 25–29 years treatment group.

and preserving the original migration trajectory during the COVID period. We then 'add' the predicted net-migrants to the treatment age populations (20–24 and 25–29) and make the extreme-case assumption that none of the migrants into Australia would have purchased PHI, while all of the people leaving would have held PHI. We recalculate the new total number of insured individuals, new population size, and new coverage rate per quarter and re-estimate the SDID models. The results are presented in Figure 13.

While the treatment effect is (mechanically) smaller during the COVID period in our analysis, the difference to our baseline estimates is small. This suggests that migration disruptions only explain a small amount (if any) of the elevated treatment effect estimates during the pandemic.

4.4.3 Changed age of dependency effect and children living with parents

Recall that since around mid-2022, insurers started adopting an expanded definition of a 'dependent' for the purposes of family PHI policies, extending eligibility to some 25–29 year olds. Dependents can join their family's plan at no additional cost. According to Figure 4, there is a small but measurable increase in the proportion of insured individuals covered under policies "with dependents" at the time of this change, namely a 1.56% increase in the "with dependents" policy proportion (from 53.56 to 54.40 ppts) between 2022 and 2023. Whether this represents a jump in the number of people in this age group joining PHI or simply dropping own cover to join their family's plan, is unclear. However, it is plausible that some of the increase in the treatment effect for this age group after

mid 2022 – which is absent for the 20–24 group – is due to the expanded dependency definition.

For the 20–24 age group, a related issue is that people in this age group may have decided to move in with their parents, or delayed moving out of their family home, during the pandemic. Indeed, one survey of young people found that 5% of surveyed 20-21 year olds moved back to their family home during the pandemic, and 91% of those cited COVID-19 restrictions as the reason (Evans-Whipp and Prattley, 2023). If these people were then able to join their parents' insurance, that could have driven up the rate of coverage for this age group. Unfortunately it is difficult to test for this with our data, but the fact that there is little change in rate of policies with dependents between 2020 and 2022 (Figure 4) suggests this effect may have been small. On the other hand, the drop in the treatment effect from around 4 ppts at the start of 2022 to below 3 ppts towards the end of 2023 could indicate this was partly responsible for the large effects during the lockdown period, with the treatment effect going down once people resumed normal living arrangements.

4.4.4 Maximum bound test

In this subsection, we pivot our attention away from directly explaining the drivers of our estimates and instead focus on whether their magnitude can reasonably be attributed to the discount incentive alone.

Australian Government Department of Health, Disability and Ageing (2024) report the precise number of people receiving discounts in each age group. This provides a natural upper limit to how large our treatment effect estimates can reasonably be. For example, if our estimates imply that the policy induces more people into PHI than are receiving discounts, this would suggest confounding factors are driving part of the estimated effect. If the estimated number of people induced by the discount policy remains below the maximum bound, the treatment effect is plausible.

In Figure 14, the black solid line shows the maximum bound value (the total number of people who received the discount policy). It is zero in March 2019, before the policy commenced, then sharply increases in June 2019, since young people who already had PHI automatically received the discount, as well as those who joined the pool in that period. In the following periods, the bound value is relatively stable for both treatment groups. One may infer from this that the policy did not induce anyone new into the market and was therefore ineffective. However, there are two important reasons why

Figure 14: The number of people induced in SDID and 'Maximum bound'

(a) Treated: 20–24

(b) Treated: 25–29

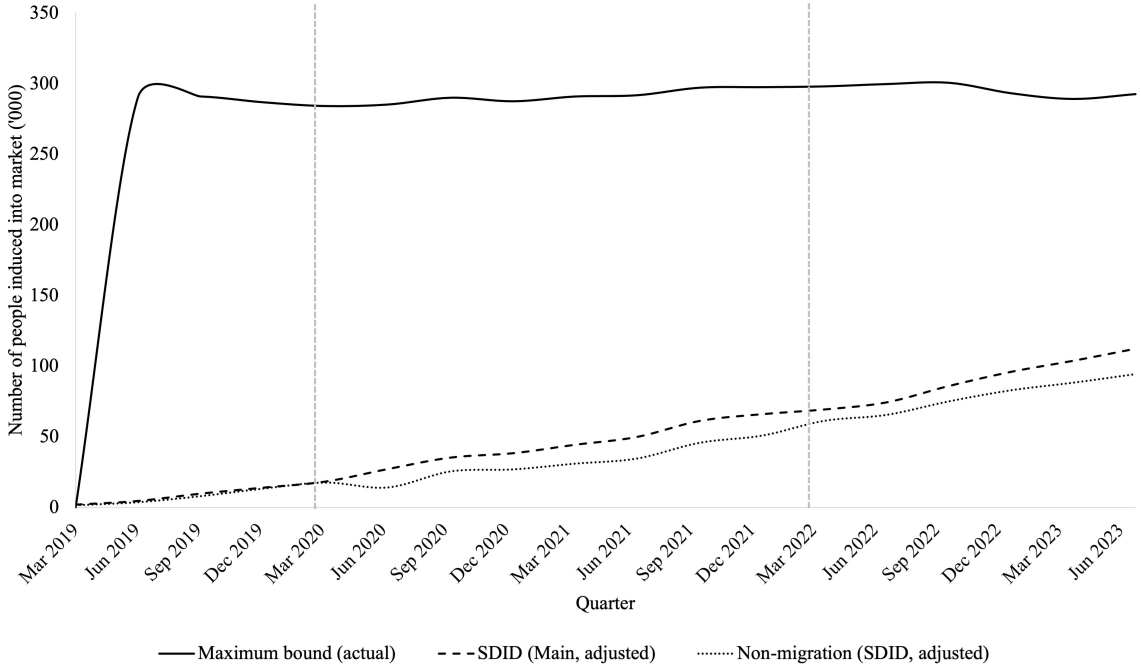

Notes: The number of individuals induced by the discount policy is calculated by multiplying the relevant SDID estimate by the corresponding population size. The lines (solid, dashed and dotted) represent, respectively, the maximum bound (actual), the original SDID estimates, and the SDID estimates excluding the migration effect. The dashed and dotted lines in the age 25–29 correspond to the adjusted SDID estimates and the adjusted SDID estimates excluding the migration effect. Vertical dot line for Q1 2020 represents that COVID-19 starting point, and Q1 2022 is approximately when the definition of a dependent changed.

this interpretation is inappropriate. First, recall that coverage was declining for these treatment groups when the policy was introduced (see Figure 1), and a stable bound curve is consistent with a reversal (or slowing) of that trend. Second, in each period the number of people receiving discounts changes due to natural age changes, with some people moving from younger age pools to older age pools. This is why it is important to estimate a model of insurance demand with a thoughtfully constructed counterfactual, as we have attempted to do.

Now, we compare the maximum bound to the estimated number of persons induced by the discount effect implied by our SDID estimates. We obtain estimates for the number of people induced to insure by multiplying the treatment effects from our panel event study models by the number of people insured. Figure 14a shows that the positive effect before the COVID-19 pandemic is 'plausible' for the 20–24 age in all cases. However, the effect size of the discount for this treatment group is questionably large post-COVID, even exceeding the maximum bound in some periods. This is our strongest evidence that the effects in the COVID-19 period are not solely due to the discount policy.

For the 25–29 group, the estimates are all well within the maximum bound. However, this is unsurprising since the maximum bound is naturally much higher for this age group because, unlike the 20–24 age group, the majority of these people are independently insured (i.e., not covered as a dependent on a family policy) and so a lot more of them are receiving discounts.

A natural complement to the maximum bound analysis is to consider what the estimates imply about the price elasticity of demand. We cannot estimate a pure price elasticity because the discount policy is not a one-off price drop. Purchasing PHI while young allows people to lock-in the discount rate until age 41, so forward looking individuals may bring forward their entry into the market. As a result, implied elasticities should be treated as upper bounds. The implied arc elasticity as of Sept 2019 for 20–24 is 1.2 and as of Dec 2019 is 3.9.[15] These figures are large but arguably reasonable. First, the pre-policy rate of coverage among 20–24 year olds, who are not covered as dependents, is extremely low – around 4%. With such a low baseline, small shifts in demand imply large elasticities. Second, the dynamic incentives to lock-in discount rates mean these figures should be larger than the pure price elasticity. Other research has also found that young people are more responsive than older cohorts to premiums in Australia (Stavrunova and Yerokhin, 2014; Liu and Zhang, 2023). From 2020 (during the pandemic), implied elas-

[15] We estimate arc elasticities as $\frac{\frac{Q_1-Q_0}{(Q_1+Q_0)/2}}{\%Discount}$ where Q_1 is the maximum bound and Q_0 is the difference between the maximum bound and the estimated number of people induced to insure. $\%Discount$ is the average discount received, which for the 20–24 age group is 10% and for 25–29 is 5.5%.

ticities range up to above 30, further evidencing our conclusion that effects post-COVID cannot be fully attributed to the discount policy.

For the 25–29 treatment group, the implied elasticities are smaller – 0.6 and 0.9 in Sept 2019 and Dec 2019 respectively – perhaps due to the higher baseline. Effects may also be relatively smaller because the long-term discount was less and because this group already had an incentive to insure before their 31st birthday due to LHC. The implied elasticities are much (questionably) larger post-COVID, reaching 8.6 by June 2023.

5 Cost-effectiveness analysis

Has the age-based discount policy improved the financial position of insurers? One rationale behind incentivising younger people to purchase PHI is that they reduce the degree of adverse selection and therefore put downward pressure on premiums. This would only be the case though if the policy is self-financing for insurers. Discounts will not be self-financing if the amount spent on discounts (which includes windfall gains to those who would have insured anyway) exceeds the net income received per induced person (premium received minus benefits paid).

We perform simple back-of-the-envelope calculations to gauge whether the policy has improved the position of insurers. We use information from the APRA Operations of Private Health Insurers Annual Report (Australian Prudential Regulation Authority, 2024a) on premiums, membership and cost of discounts by age group, and population from ABS. The insurers' quarterly extra cost is equal to the quarterly benefits of people who are induced by the discount, which is measured by the total quarterly cost of induced people and the number of people induced by the discount each quarter.[16] The insurers' net income is equal to the difference between quarterly premiums and the cost of hospital treatment paid by the insurer. Results are in Figure 15. We only focus on estimates up to the onset of the COVID-19 pandemic since we do not believe effects after this period can reasonably be attributed to the discount alone, and ignore second-order effects (e.g., lower premiums overall due to more favourable selection of insurers).

The trends in insurers' cost move in tandem with the discount effect. Focusing on Dec 2019, for the 20–24 age group the additional revenue received ($1437.41) exceeds the additional costs borne ($527.51), implying the discount provided a positive financial

[16] (1) The total quarterly cost of induced people is the number of people receiving discounts multiplied by the premium of insured people times the average discount. (2) The number of people induced by discount each quarter is obtained from SDID estimates.

Figure 15: Average cost and net revenue per induced person for insurers

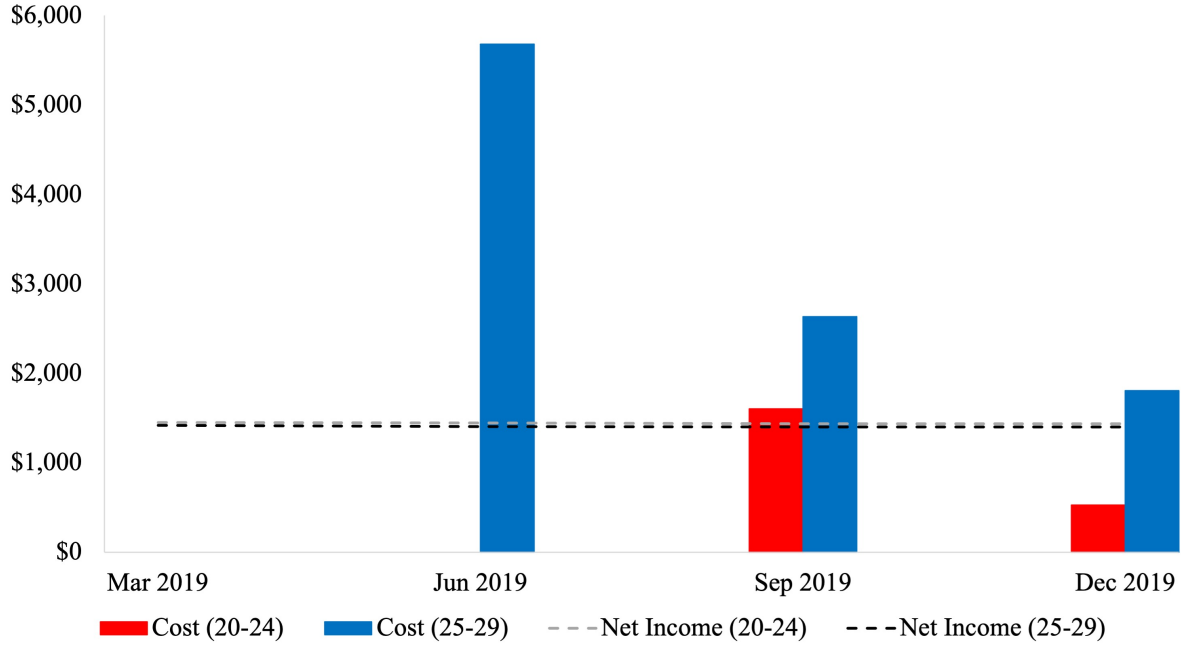

Notes: (1) Net revenue per insured $= \frac{\text{Hospital treatment premium revenue}}{\text{Number of Hospital treatment insured persons}} - \frac{\text{Hospital treatment total benefits}}{\text{Hospital treatment insured persons}}$
(2) Cost per induced $= \max\left(0, \frac{\text{Total cost of the discount}}{\text{Number of person induced to insure}}\right)$, where total cost of the discount equals the number of persons who receive a discount multiplied by the average premium and the average discount rate, and the number of person induced to insure equals the population in the treated age group multiplied by the estimated treatment effect. Estimates are bounded below at zero because negative treatment effects imply economically meaningless negative inducement costs.
Source: The yearly hospital treatment premium revenue and yearly number of Hospital treatment insured persons are both from APRA *Operations of Private Health Insurers Annual Report* data (Australian Prudential Regulation Authority, 2024a). The hospital treatment total benefits is from APRA *Quarterly private health insurance statistics* issue released on 28 February 2024 (Australian Prudential Regulation Authority, 2024b). The number of persons who receive a discount is the number of people receiving discounts from *Private health insurance reform data quarterly trends report* released on 15 April 2024 (Department of Health and Aged Care., 2024). The average premium is calculated by $\frac{\text{total premium revenue}}{\text{total number of insured persons}}$. The average discount rate for 20–24 is 10%, and average discount for 25–29 is set to 5.5% based on our calculations for how many people are receiving each rate of discount. The population of treated age group is from the Australian Bureau of Statistics *Australian Demographic Statistics* (Australian Bureau of Statistics, 2020a).

return for insurers. This is consistent with the large relative price response from this group. For the 25–29 group, the average additional revenue ($1398.18) does not exceed the additional costs ($1808.17). However, bear in mind this is a static analysis, based on relatively short-term impacts, and does not consider whether those induced to insure are more favourably selected than the existing insurance pool. Overall, the analysis suggests that if insurers were able to charge lower premiums to young people less restrictively, as some commentators have advocated (Paolucci et al., 2019), it is plausible that they would be willing to do so.

6 Conclusion

This study evaluates a policy aimed at young individuals (aged 20–29) that grants them a premium discount for purchasing PHI. Our evidence suggests that the policy generated only a modest positive effect in the period leading up to the COVID-19 pan-

demic, and then coincided with much larger estimated effects post-COVID. However, the observed larger effect after the pandemic does not appear to be attributable to the policy itself, but rather to a heterogeneous response to the pandemic by age group. This indicates that external shocks can reshape behavioural responses in ways that obscure or disrupt the underlying policy effects.

Our findings imply that while young people are indeed responsive to price incentives, leading to a modest increase in insurance take-up, the policy does not materially alter the demographic composition of the insured pool. Since adverse selection is largely driven by the disproportionately high participation of older (higher-risk) people, the modest inflow of younger and healthier entrants achieved under the current discount is insufficient to offset this structural imbalance. In other words, the magnitude of the current discount policy is not sufficient enough to fundamentally address the ageing insurance pool.

Has the discount policy been worthwhile? The pre-COVID evidence seems to provide the most reliable estimate of the true effect of the policy. Based on what we observed in this period, it is plausible that the policy induced some additional insurance take-up among young people, and that the premium revenue from those young people would have been sufficient to cover the extra costs to insurers. This suggests that the discount may have potentially exerted downward pressure on premiums, albeit a likely very modest contribution.

The COVID-19 disruption prevents us from reliably estimating the medium- or longer-term impact of the policy at this stage. Nevertheless, the apparent disproportionate demand response from younger people during the pandemic is a noteworthy finding in and of itself. Exploring this further would be an interesting direction for future research. Understanding why younger and older cohorts might respond differently to a generalised demand shock – whether this be due to policy settings, beliefs or something else – could help to improve predictions about behaviour in health insurance markets.

Recall that the age-based discount sits within a suite of incentives designed to prop up demand for PHI in Australia. It is evident that the discount policy has not dramatically reshaped the Australian PHI market. However, it has plausibly increased demand for PHI (by a small margin) without being punitive to insurers or consumers, or directly costing taxpayers. In contrast, premium rebates for older people are an expensive lever with limited long-term gains (Zhang et al., 2023; Kettlewell et al., 2018; Liu and Zhang, 2023), and LHC and MLS increase demand through punishments. There is ongoing controversy around government intervention in PHI in Australia (Duckett et al., 2019), and more broadly about the appropriate design of regulatory incentives in supplementary health

insurance markets. We hope the evidence in this paper can help to shape the continued debate.

References

Abadie, A., Diamond, A., and Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American statistical Association*, 105(490):493–505.

Abadie, A., Diamond, A., and Hainmueller, J. (2015). Comparative politics and the synthetic control method. *American Journal of Political Science*, 59(2):495–510.

Abadie, A. and Gardeazabal, J. (2003). The economic costs of conflict: A case study of the basque country. *American economic review*, 93(1):113–132.

Abadie, A. and L'hour, J. (2021). A penalized synthetic control estimator for disaggregated data. *Journal of the American Statistical Association*, 116(536):1817–1834.

Angrist, J. D. and Pischke, J.-S. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review*, 111(12):4088–4118.

Ashenfelter, O. C. and Card, D. (1984). Using the longitudinal structure of earnings to estimate the effect of training programs.

Auditor-General Report No.12 2021–22 (2022). Management of international travel restrictions during covid-19.

Australian Bureau of Statistics (2020a). National, state and territory population. `https://www.abs.gov.au/statistics/people/population/national-state-and-territory-population`.

Australian Bureau of Statistics (2020b). National, state and territory population methodology. `https://www.abs.gov.au/methodologies/national-state-and-territory-population-methodology`.

Australian Bureau of Statistics (2024). Overseas migration. `https://www.abs.gov.au/statistics/people/population/overseas-migration`.

Australian Government Department of Health, Disability and Ageing (2024). Private health insurance reform data quarterly trends report. `https://www.health.gov.au/resources/publications/private-health-insurance-reform-data-quarterly-trends-report`.

Australian Institute of Health and Welfare. (2021). Covid-19 and the impact on young people. https://www.aihw.gov.au/reports/children-youth/covid-19-and-young-people.

Australian Institute of Health and Welfare. (2022). Health expenditure australia 2020-21. https://www.aihw.gov.au/reports/health-welfare-expenditure/health-expenditure-australia-2020-21.

Australian Prudential Regulation Authority (2024a). Operations of private health insurers annual report. https://www.apra.gov.au/operations-of-private-health-insurers-annual-report.

Australian Prudential Regulation Authority (2024b). Quarterly private health insurance statistics. https://www.apra.gov.au/quarterly-private-health-insurance-statistics.

Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *The Quarterly journal of economics*, 119(1):249–275.

Bilgrami, A., Cutler, H., Sinha, K., and Cheng, Z. (2021). The impact of means-tested premium rebates and tax penalties on the demand for private hospital cover in australia. *Economic Record*, 97(317):170–211.

Cheng, T. C. (2014). Measuring the effects of reducing subsidies for private insurance on public expenditure for health care. *Journal of Health Economics*, 33:159–179.

Clarke, D., Pailañir, D., Athey, S., and Imbens, G. (2023). Synthetic difference in differences estimation. *arXiv preprint arXiv:2301.11859*.

Colombo, F. and Tapay, N. (2004). Private health insurance in oecd countries: the benefits and costs for individuals and health systems.

Commonwealth Ombudsman. (2021). Age-based discount. https://www.privatehealth.gov.au/health_insurance/surcharges_incentives/\discount_age.htm.

Courtemanche, C., Marton, J., Ukert, B., Yelowitz, A., and Zapata, D. (2017). Early impacts of the affordable care act on health insurance coverage in medicaid expansion and non-expansion states. *Journal of Policy Analysis and Management*, 36(1):178–210.

Department of Health and Aged Care. (2024). Private health insurance reform data quarterly trends report. https://www.health.gov.au/resources/publications/p

rivate-health-insurance-reform-data-quarterly-trends-report?language=en.

Dobkin, C., Finkelstein, A., Kluender, R., and Notowidigdo, M. J. (2018). The economic consequences of hospital admissions. *American Economic Review*, 108(2):308–352.

Duckett, S. (2020). Young people dropping private health hurts insurers most, not public hospitals.

Duckett, S., Cowgill, M., and Nemet, K. (2019). Saving private health 2: Making private health insurance viable.

Ellis, R. P. and Savage, E. (2008). Run for cover now or later? the impact of premiums, threats and deadlines on private health insurance in australia. *International journal of health care finance and economics*, 8:257–277.

Evans-Whipp, T. and Prattley, J. (2023). Young adults returning to live with parents during covid-19.

Finkelstein, A. (2002). The effect of tax subsidies to employer-provided supplementary health insurance: evidence from canada. *Journal of Public Economics*, 84(3):305–339.

Frean, M., Gruber, J., and Sommers, B. D. (2017). Premium subsidies, the mandate, and medicaid expansion: Coverage effects of the affordable care act. *Journal of health economics*, 53:72–86.

Freyaldenhoven, S., Hansen, C., Pérez, J. P., and Shapiro, J. M. (2021). Visualization, identification, and estimation in the linear panel event-study design. Technical report, National Bureau of Economic Research.

Gruber, J. and Washington, E. (2005). Subsidies to employee health insurance premiums and the health insurance market. *Journal of Health Economics*, 24(2):253–276.

Higginson, S., Milovanovic, K., Gillespie, J., Matthews, A., Williams, C., Wall, L., Moy, N., Hinwood, M., Melia, A., and Paolucci, F. (2020). Covid-19: The need for an australian economic pandemic response plan. *Health Policy and Technology*, 9(4):488–502.

Hinde, J. M. (2017). Incentive (less)? the effectiveness of tax credits and cost-sharing subsidies in the affordable care act. *American Journal of Health Economics*, 3(3):346–369.

Huntington-Klein, N. (2021). *The effect: An introduction to research design and causality.* Chapman and Hall/CRC.

Kettlewell, N., Stavrunova, O., and Yerokhin, O. (2018). Premium subsidies and demand for private health insurance: results from a regression discontinuity design. *Applied Economics Letters*, 25(2):96–101.

Kettlewell, N. and Zhang, Y. (2024). Age penalties and take-up of private health insurance. *Health Economics*, 33(4):636–651.

King, D. and Mossialos, E. (2005). The determinants of private medical insurance prevalence in england, 1997–2000. *Health services research*, 40(1):195–212.

Kleven, H., Landais, C., and Søgaard, J. E. (2019). Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics*, 11(4):181–209.

Koce, M. (2022). More young australians choose members health funds. `https://membershealth.com.au/publications/more-young-australians-choose-members-health-funds`.

Lee, D. S. (2005). Training, wages, and sample selection: Estimating sharp bounds on treatment effects.

Liu, J. and Zhang, Y. (2023). Elderly responses to private health insurance incentives: Evidence from australia. *Health Economics*, 32(12):2730–2744.

Marenzi, A., Rizzi, D., and Zanette, M. (2021). Incentives for voluntary health insurance in a national health system: evidence from italy. *Health Policy*, 125(6):685–692.

Nicolás, Á. L. and Vera-Hernández, M. (2008). Are tax subsidies for private medical insurance self-financing? evidence from a microsimulation model. *Journal of Health Economics*, 27(5):1285–1298.

Palangkaraya, A. and Yong, J. (2005). Effects of recent carrot-and-stick policy initiatives on private health insurance coverage in australia. *Economic Record*, 81(254):262–272.

Paolucci, F., Melia, A., and Henriquez, J. H. (2019). Private health insurance premiums should be based on age and health status. The Conversation.

Rodríguez, M. and Stoyanova, A. (2008). Changes in the demand for private medical insurance following a shift in tax incentives. *Health economics*, 17(2):185–202.

Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, 4(3):305–22.

Stavrunova, O. and Yerokhin, O. (2014). Tax incentives and the demand for private health insurance. *Journal of health economics*, 34:121–130.

Thomasson, M. A. (2003). The importance of group coverage: How tax policy shaped us health insurance. *American Economic Review*, 93(4):1373–1384.

Walker, B. (2020). The effects of covid-19 on the australian health system and phi. `https://www.actuaries.digital/2020/03/20/the-effects-of-covid-19-on-the-australian-health-system-and-phi`.

Zhang, Y. (2020). Private hospital insurance premiums should vary by age. *MJA Insight+*, 2020(October 5).

Zhang, Y., Liu, J., and Kettlewell, N. (2023). Who really benefits from private health insurance rebates? not people who need cover the most.

Appendix A - Table and Figure

Table A1: Dynamic SDID estimator

	20–24		**25–29**	
$\hat{\tau}$	**original**	**rotated**	**original**	**rotated**
Mar 2017	0.0043	0.0010	0.0156	-0.0002
Jun 2017	0.0006	-0.0020	0.0136	-0.0000
Sept 2017	0.0022	0.0003	0.0119	0.0005
Dec 2017	0.0041	0.0030	0.0094	0.0002
Mar 2018	-0.0002	-0.0005	0.0064	-0.0006
Jun 2018	-0.0032	-0.0028	0.0049	0.0001
Sept 2018	-0.0020	-0.0009	0.0031	0.0005
Dec 2018	0	0.0019	0	-0.0004
Mar 2019	-0.0016	0.0011	-0.0008	0.0010
Jun 2019	-0.0016	0.0018	-0.0017	0.0024
Sept 2019	0.0034	0.0076	-0.0011	0.0051
Dec 2019	0.0100	0.0150	-0.0011	0.0073
Mar 2020	0.0116	0.0173	-0.0011	0.0096
Jun 2020	0.0157	0.0222	0.0015	0.0143
Sept 2020	0.0243	0.0316	0.0038	0.0189
Dec 2020	0.0306	0.0386	0.0034	0.0207
Mar 2021	0.0324	0.0412	0.0046	0.0241
Jun 2021	0.0341	0.0437	0.0056	0.0273
Sept 2021	0.0361	0.0464	0.0099	0.0338
Dec 2021	0.0395	0.0506	0.0104	0.0365
Mar 2022	0.0354	0.0472	0.0098	0.0381
Jun 2022	0.0331	0.0457	0.0101	0.0406
Sept 2022	0.0299	0.0432	0.0138	0.0465
Dec 2022	0.0314	0.0455	0.0164	0.0513
Mar 2023	0.0274	0.0422	0.0174	0.0545
Jun 2023	0.0256	0.0412	0.0189	0.0582
Treatment group	✓	✓		
			✓	✓
Observations	338		338	

Table A2: Optimal SDID unit and time weights selection

	$\hat{w}^{sdid}$			$\hat{\lambda}^{sdid}$	
Unit	**20–24**	**25–29**	Time	**20–24**	**25–29**
30–34	0.1252	0.2246	Mar 2017	0	0
35–39	0.0984	0.1287	Jun 2017	0	0
40–44	0.0742	0.0632	Sept 2017	0	0
45–49	0.0840	0.0841	Dec 2017	0	0
50–54	0.0942	0.1142	Mar 2018	0	0
55–59	0.1038	0.1385	Jun 2018	0	0
60–64	0.1082	0.1398	Sept 2018	0	0
65–69	0.0863	0.0720	Dec 2018	1	1
70–74	0.0773	0.0349			
75–79	0.0491	0			
80–84	0.0489	0			
85–89	0.0504	0			
Treatment group	✓			✓	
		✓			✓

Figure A1: SDID Estimates (Robustness Check)
— windows of pre-treatment period

(a) Treated: 20–24

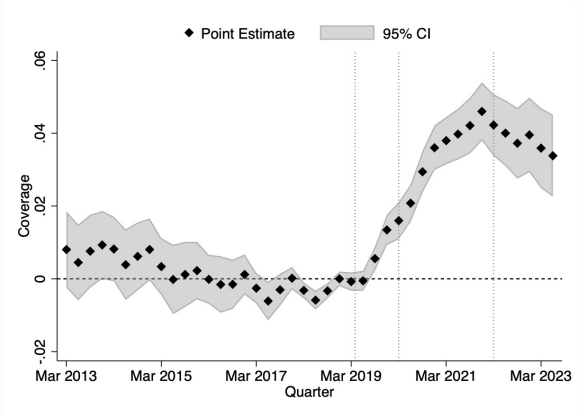

(b) Treated: 25–29

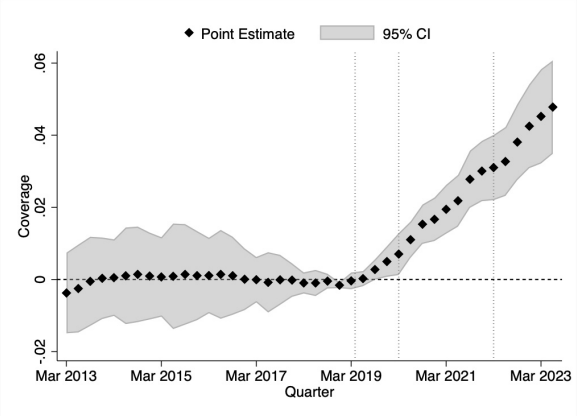

Notes: The rotation adjustment was only maintained for the 25–29 years treatment group. Vertical dotted lines mark three key events: the introduction of the age-based discount in April 2019, the onset of COVID-19 after Q1 2020, and approximate definition of a dependent changed after Q1 2022. Standard errors using the placebo method are calculated using bootstrap with 1000 replications. Confidence intervals are built by bootstrap estimation predictions using only the treated units to estimate the noise level.

Figure A2: SDID Estimates (Robustness Check)
— dropping age 30–34 from donor pool

(a) Treated: 20–24

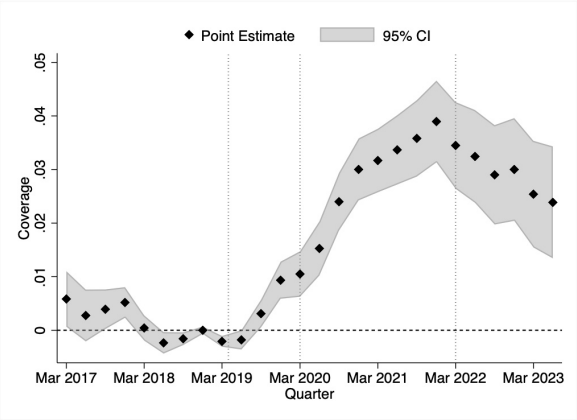

(b) Treated: 25–29

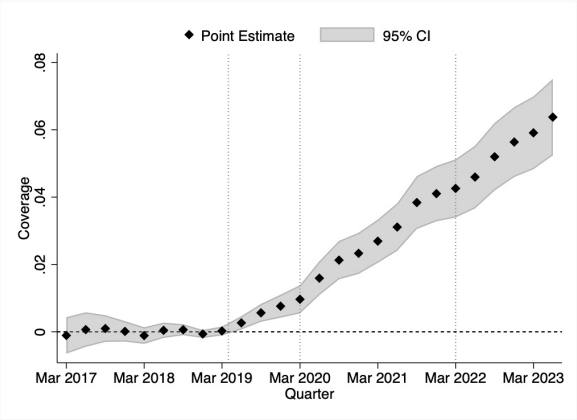

Notes: The rotation adjustment was only maintained for the 25–29 years treatment group. Vertical dotted lines mark three key events: the introduction of the age-based discount in April 2019, the onset of COVID-19 after Q1 2020, and approximate definition of a dependent changed after Q1 2022. Standard errors using the placebo method are calculated using bootstrap with 1000 replications. Confidence intervals are built by bootstrap estimation predictions using only the treated units to estimate the noise level.

Figure A3: SDID Estimates (Robustness Check)
— Interstate migration

(a) Treated: 20–24

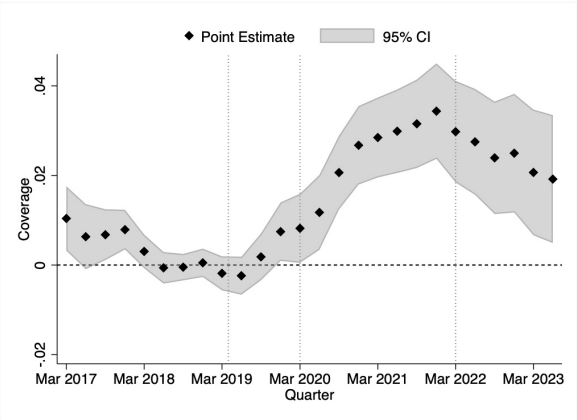

(b) Treated: 25–29

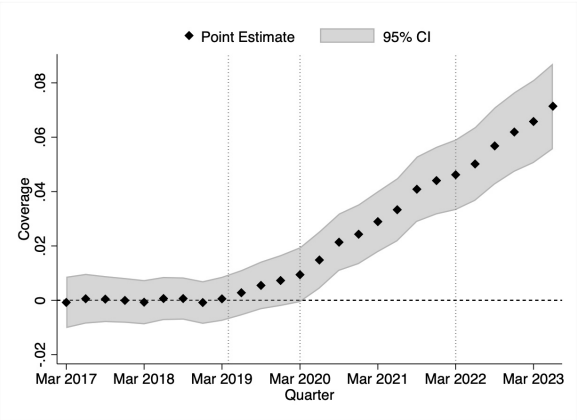

Notes: The rotation adjustment was only maintained for the 25–29 years treatment group. Vertical dotted lines mark three key events: the introduction of the age-based discount in April 2019, the onset of COVID-19 after Q1 2020, and approximate definition of a dependent changed after Q1 2022. Standard errors using the placebo method are calculated using bootstrap with 1000 replications. Confidence intervals are built by bootstrap estimation predictions using only the treated units to estimate the noise level.

Figure A4: Dynamic Treatment Effect in Placebo tests (different intervention time)

2016

(a) 20–24 (b) 20–24 (adjusted) (c) 25–29 (adjusted)

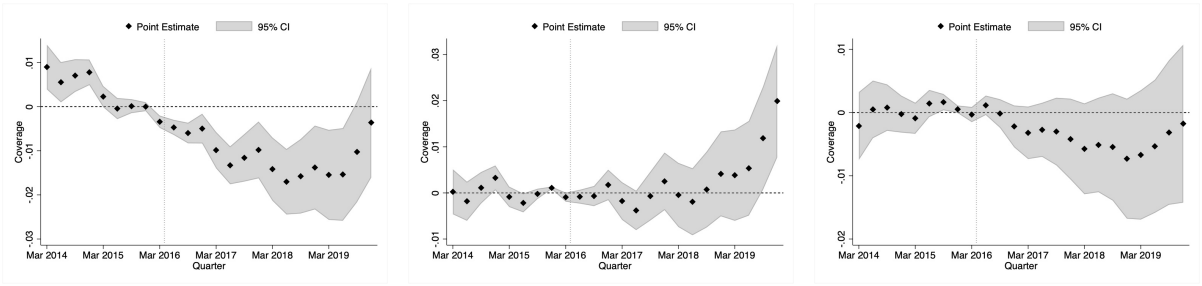

2017

(d) 20–24 (e) 20–24 (adjusted) (f) 25–29 (adjusted)

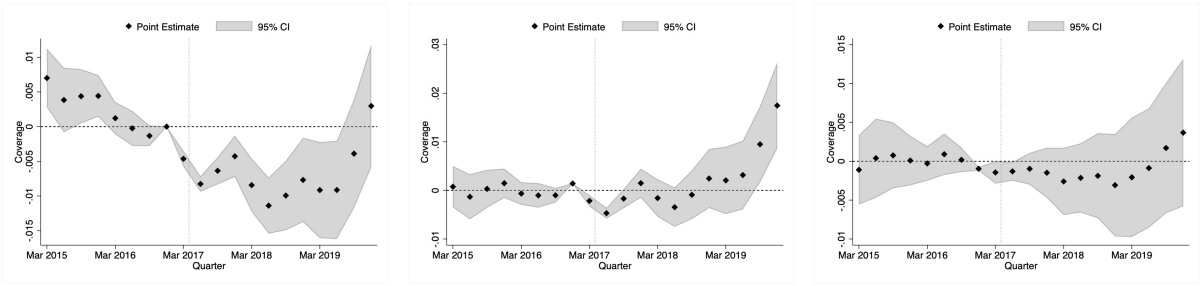

2018

(g) 20–24 (h) 20–24 (adjusted) (i) 25–29 (adjusted)

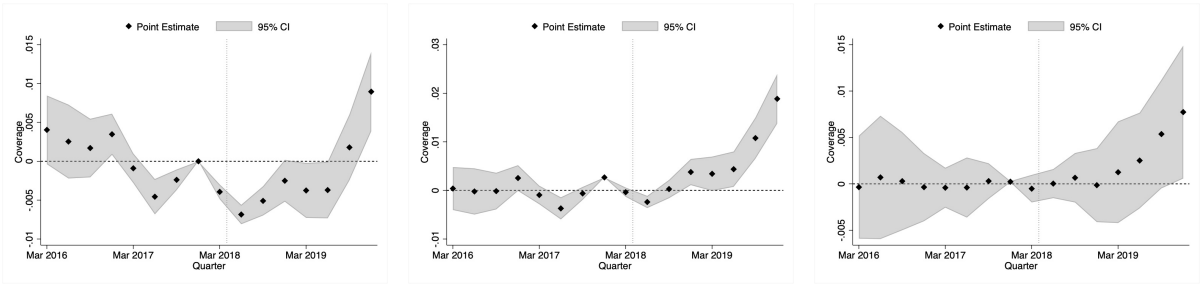

Notes: Vertical dot line for pseudo-treatment test time point. Standard errors using the placebo method are calculated using bootstrap with 1000 replications. Confidence intervals are built by bootstrap estimation predictions using only the treated units to estimate the noise level.