AUSTRALIAN MINIMUM LEGAL DRINKING AGE & EFFECTS LATER-IN-LIFE

Thomas Denford

November 2025

School of Economics
University of Sydney

A thesis submitted in partial fulfilment of the requirements for the degree of Bachelor of Economics (Honours)

Supervised by Dr. Rebecca McKibbin

STATEMENT OF ORIGINALITY

I hereby declare that this submission is my own work and to the best of my knowledge it contains no material previously published or written by another person. Nor does it contain any material which has been accepted for the award of any other degree or diploma at the University of Sydney or at any other educational institution, except where due acknowledgment is made in this thesis.

Any contributions made to the research by others with whom I have had the benefit of working at the University of Sydney is explicitly acknowledged.

I also declare that the intellectual content of this study is the product of my own work and research, except to the extent that assistance from others in the project's conception and design is acknowledged.

Thomas Denford

7 November 2025

ACKNOWLEDGEMENTS

I need to sincerely thank my incredible supervisor, Dr. Rebecca McKibbin, for her hard work and guidance throughout the year. From coming up with the initial ideas, to keeping my many minor concerns at bay, safe to say this thesis would not exist if not for your selflessness in helping me.

Special thanks to my wonderful family and Honours homies also. In the immortal words of Tina Turner, they were "simply the best."

I acknowledge the use of ChatGPT's GPT-4o and GPT-5 models, to write Stata code, develop my empirical strategy, and refine the academic language of my own work. This included providing my thoughts, ideas and complete sentences, and asking for suggestions on how to develop and improve upon them.

ABSTRACT

This thesis examines the long-run effects of Australia's 1970s minimum legal drinking age (MLDA) reforms, which aligned all states to a uniform age of 18. While alcohol's harms on adolescents are well-established, less is known about its level of persistence decades on. Using a staggered synthetic difference-in-differences (SDID) framework and HILDA data, I estimate the impact of exposure to MLDA reform on a range of outcomes at ages 55 to 63. Large and highly significant effects are found for annual income (negative) and the prevalence of long-term health conditions (positive). However, a likely violation of the identifying assumption known as 'parallel trends'—driven by structural differences between treated and control states—means these estimates cannot be interpreted as causal. A novel contribution of this thesis is its focus on the long-run outcomes of MLDA policy in an Australian context, in contrast to existing literature which has largely examined short-run effects. Most importantly, however, the study highlights the challenges of using retrospective data to assess historical policy impacts, underscoring the need for caution in formulating research questions.

1 Introduction

It is well-established that alcohol has adverse effects on the health and development of adolescents (Tetteh-Quarshie & Risher, 2023). Today, alcohol remains the leading individual risk factor for death and disability in 15- to 24-year-olds globally (Bowden et al., 2024). The minimum legal drinking age (MLDA) is recognised as the key policy for mitigating alcohol-related harms among young people, as it establishes the age at which alcohol is legally accessible. My thesis is interested in understanding the connection between the MLDA and long-run outcomes.

In the 1970s, a series of state-level reforms aligned Australia's MLDA to a uniform age of 18, previously varying across jurisdictions. My aim is to investigate whether these changes have influenced cohort outcomes three to four decades after exposure.

The relationship between early-life shocks and later-life outcomes has been widely studied, with significant findings across a range of research areas (Smith, 2015; Almond et al., 2018). The underlying MLDA theory is that by restricting legal access, alcohol consumption can be eliminated, or at least significantly reduced, thereby limiting the associated harms. Although most have had their 'first drink' before reaching legal age (Carpenter & Dobkin, 2009), there is significant evidence the MLDA does influence alcohol consumption (Yörük & Yörük, 2013; Ahammer et al., 2022). As such, if early-life shocks can influence behaviour in the short-run, their impact on habit formation may in turn affect outcomes later-in-life.

To conduct my study, I use a staggered, synthetic differences-in-differences (SDID) framework to exploit the 1970s reforms, and use HILDA (Household, Income and Labour Dynamics in Australia) survey as my primary data source for identifying and outcome variables.

In addition to evidence of the alcohol mechanism, the core identification assumption of SDID is that the treatment and control groups exhibit 'parallel trends' prior to reform. Initial tests of the MLDA's effect on teenage alcohol consumption could not be done, while later-life consumption appears to be lower in states who had MLDA reform. Critically, however, 'parallel trends assumption' is likely violated. It is probable that my results

are driven by confounding factors, namely existing imbalances between the economically larger 'control' states and the 'treated' states where these reforms took place.

Although this raises concern, I present results with a strong caveat that they should not be interpreted as causal, but instead to provide a comprehensive record of what was derived. I find strong reductions in real annual income between $9500 and $13000 (around one-fifth of the current median), and an increase in the prevalence of long-term health conditions of between 7.6 and 15.2 percentage points at ages 55 to 63 for those affected by a MLDA reform. Though highly significant estimates are not found for household net worth and university completion, the magnitude of these results is staggering.

Accordingly, a key contribution of this paper is the study of long-run outcomes, since much of the literature is focused on short-run outcomes. However, this paper's main contribution is its principal limitation: the absence of valid data prevents a direct test of the 1970s MLDA reforms. Its value in novelty is constrained by its age.

The remainder of this thesis is organised as follows. Section 2 explains the background of Australia's minimum legal drinking age laws and motivates policy context in the 1970s. I conduct a literature review on MLDA policies and their impacts on short-run and long-run outcomes in Section 3. Section 4 introduces my data: the HILDA survey. Section 5 outlines my staggered, synthetic 'differences-in differences' identification strategy. My mechanism and main results are presented in Section 6. And finally, Section 7 outlines my concluding remarks.

2 Background

Today, Australia's minimum legal drinking age sits at 18 nationwide. However, this was not always the case, where prior to 1974, the MLDA varied across states.

As of 1966, half of Australia's states and territories had an MLDA of 18 — New South Wales (NSW), Victoria (VIC), the Northern Territory (NT), and the Australian Capital Territory (ACT) — while the remaining states had an MLDA of 21 — Queensland (QLD), Western Australia (WA), South Australia (SA) and Tasmania (TAS).

Over the course of the next 8 years, all states that had an MLDA of 21 would lower it to 18. I restrict my analysis to the major reforms in WA, SA, and QLD, and compare outcomes to NSW and VIC, whose laws did not change during this period. This selection prioritises reforms with high potential significance and ensures sufficient data availability for testing.

As such, Table 1 summarises all major MLDA policy changes, with the dates at which each state's MLDA was lowered to or set at 18. WA and QLD lowered the legal age to 18 in a single reform — the former in 1970 and the latter in 1974, while SA did it in two stages — progressing to 20 in 1968, then 18 in 1971. As mentioned, I will focus on SA's larger 1971 reform.

Table 1: MLDA reform across Australia (1970s).

State	Date	Reform
Western Australia	1 July 1970	$21 \rightarrow 18$
South Australia	8 April 1971	$20 \rightarrow 18$
Queensland	18 February 1974	$21 \rightarrow 18$

These legislative changes were done in line with pushes internationally (namely the US) to extend rights to eligible servicemen during the Vietnam War. Slogans such as 'Old enough to fight, old enough to vote' were commonly used by proponents advocating for a uniform 'age of majority' at which all adult rights would be granted (Toomey at al., 2009).

As such, this adoption of a national MLDA was not health-related—if it was the case, the MLDA would likely have been raised. Furthermore, it did not coincide with other alcohol consumption policies, being merely a political act to extend freedoms to young people. Thus, it seems plausible that these MLDA policies were generally exogenous. In that case, assuming the parallel trends assumption holds, any result found would be attributable to the lowering of the MLDA, not some other concurrent policy or confounding variable.

These policies concentrate on 18 to 20-year-olds, the cohorts whose legal access to alcohol were subject to MLDA reform. Hence, the long-run indicators I study reflect the impact of greater alcohol access during adolescence on present-day outcomes. I will use the term 'adolescents' and 'teenagers' interchangeably to refer to this group.

3 Literature review

There has been extensive literature examining the minimum legal drinking age and its aggregate effects. Primarily this has been focused on immediate or near-immediate effects of policy, i.e. in the short-run. The discussion covers: (i) studies that speak to the MLDA mechanism, (ii) short-run effects for motivation, and (iii) long-run effects.

My contributions serve to enhance the literature on Australian alcohol and MLDA policy, as well as their effects on long-run outcomes — both areas where evidence remains limited.

3.1 Mechanism

3.1.1 Theory

It is important to establish that MLDA policy could have driven any long-run effects. A plausible mechanism can be visualised in Figure 1.

Figure 1: Diagram of causal mechanism.

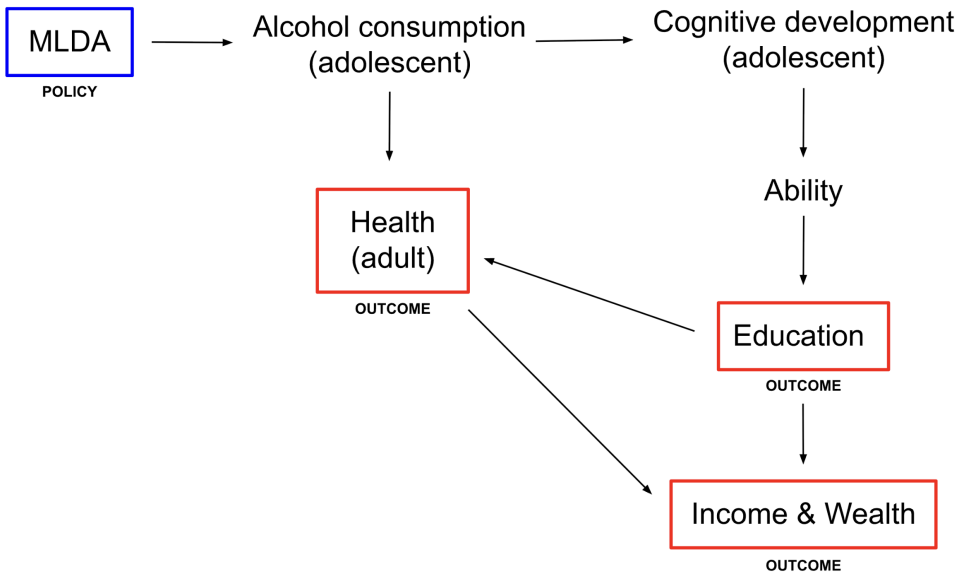

In theory, a MLDA reform, which controls legal access of alcohol for adolescents, would impact adolescent alcohol consumption. Any change in alcohol consumption would likely affect health outcomes, namely cognitive development in adolescence and health as adults.

Any effect on adult health could manifest directly in an individual's earnings (income and wealth) if their health limits their work potential. Alternatively, alternate cognitive development could directly impact an individual's inherent 'ability,' and influence their accumulation of education, income, and wealth as such.

Available evidence from trends in 1970s Australia does indicate that the MLDA reforms had a notable impact on drinking behaviour. Consumption of alcohol in Australia peaked in the 1970s, with 13.1 per capita litres drunk in 1974-75 (Australian Bureau of Statistics, 2011). This extended to young people—from available data, it is known that year 10 students experienced highs in consumption (Roche et al., 2007) and as such, this likely extended to older adolescents.

From observation research, it appears most were aware of the risks of alcohol but shared little concern, highlighted in the apparent prevalence of 'drink driving' (NFSA Films, 2017). Official guidelines on 'safe' drinking behaviour, like those produced by the National Health and Medical Research Council today, did not exist. Instead, only outdated, non-rigorous advice such as Anstie's limit—which corresponds to approximately 3.5 standard drinks daily today—served as guides (Babor et al., 1987). This would suggest the attitudes of Australians towards drinking in the 1970s were much more liberal, and an increase in alcohol consumption in conjunction with the MLDA would have been likely.

3.1.2 Alcohol consumption

In continuation, it is necessary to examine whether the MLDA is expected to impact alcohol consumption generally. The literature does overwhelmingly support a strong relationship between the two.

Carpenter and Dobkin in their seminal 2011 paper found that allowing 18 to 20-year-olds to drink (i.e. having an MLDA of 18 rather than 21) significantly increases their overall drinking participation and 'binge' drinking (5 or more drinks in one sitting). This was in line with results found by other authors who shared similar approaches, namely Dee (1999) and Miron and Tetelbaum (2009).

Furthermore, both Norberg et al. (2009) and Plunk et al. (2013) find strong links

between the MLDA and later-life alcohol dependence, with lower MLDAs increasing substance use disorder and binge drinking respectively. The presence of a link between the MLDA and alcohol consumption provides critical motivation for the presence of further effects.

3.1.3 Risks of alcohol

Next, alcohol does exert significant health effects, especially for adolescents. Upon consumption, alcohol impairs cognitive function, weakening an individual's perception, memory and decision-making (Lees et al., 2020).

Adolescents are particularly vulnerable to alcohol risks because brain development is ongoing. This limits their ability to assess risks, increases susceptibility to permanent cognitive impairment, and heightens the potential for long-run dependence (Tetteh-Quarshie & Risher, 2023). Numerous studies have affirmed a strong association between binge drinking as a teenager and poor neurocognitive performance (Jacobus & Tapert, 2013).

3.1.4 Venue access

An additional component of the MLDA is that it grants access to adult venues such as bars and clubs. Consequently, some observed outcomes associated with the MLDA may stem from access to these environments rather than from alcohol consumption itself.

Generally, studies find that 'on-premise' sales (bar or restaurant) increases crime and fatalities relative to 'off-premise' sales (liquor stores) (Baughman et al., 2001; Mair et al., 2015).

3.2 Short-run effects

Effects associated with the MLDA can be split into two categories: behaviours and outcomes. Finding impacts on both in the short-run is important for establishing the likelihood of long-run effects. Firstly, it should be expected that if MLDA policies shape behaviour, e.g. increase drug use or participation in crime, the associated effects or consequences of this behaviour would manifest later-in-life. Second, demonstrating short-term

policy effects on outcomes would provide evidence that longer term effects are plausible. Finding that the MLDA affects traffic fatalities, for example, provides evidence that the policy mechanism is sound.

Thus, although my primary interest lies in long-run effects, short-run evidence should motivate the analysis.

3.2.1 Drug & tobacco use

Observing whether 'gateway effects' exist, i.e. whether alcohol is a complement or substitute to drugs and tobacco, is essential to understanding MLDA-driven behaviour. Overall, a few studies identify a link to alcohol consumption, providing some evidence for the 'stepping-stone' theory, but more evidence to a substitute effect.

In terms of drug use, the literature differentiates between 'soft' and 'hard,' where effects vary. Deza (2015) uses a regression discontinuity design at the US MLDA of 21 and finds an increase in all measures of alcohol consumption, but a decrease in 'hard' drug use (cocaine, crack or heroin), suggesting a substitution effect. Similarly, Krauss et al. (2015) find a higher MLDA could increase the risk of 'hard' illicit drug use in adulthood, but these estimates are only moderately significant.

There is mixed evidence on the link between the MLDA and marijuana use. Crost and Guerrero (2012) find marijuana usage decreases sharply at the US drinking age. A follow-up study by Crost and Rees (2013) find no evidence of such behaviour, though they note their finding of no significance may be due to less valid data. Yörük and Yörük (2013) also find no significant effects on marijuana (or smoking) at the MLDA.

Furthermore, Dee (1999) finds an increase in the MLDA reduces smoking participation between 3% and 5%. Though their specification using cigarette taxes does not hold due to imprecise estimation, they conclude that there is evidence that cigarette and alcohol are substitutes.

3.2.2 Crime

Engagement in criminal activity represents another behaviour likely influenced by MLDA policies, given alcohol's impact on cognitive decision-making with increased risk-taking behaviour (George et al., 2005).

Carpenter and Dobkin (2015) is a landmark paper on crime and the MLDA. They find a statistically significant, discontinuous increase in the number of arrests at age 21 (in the US where the MLDA is 21). For more violent crimes, such as assault and robbery, they find a substantial increase, while less violent property and drug possession crimes see only a small increase. They note that these are arrest rates and not crime rates, where alcohol could potentially produce sloppier criminals rather than increasing criminal behaviour.

Smith and Burvell (1987) look at Australia, exploiting the same 1970s MLDA reforms in Queensland and South Australia as this study, and find an increase in male juvenile crime between 20% and 25%. Particularly high increases were found for burglary and larceny.

3.2.3 Traffic fatalities

In a shift to short-run outcomes, a huge chunk of MLDA literature encompasses regression discontinuity papers dedicated to observing effects on traffic fatalities. In medical terms, alcohol impairs driving ability by disrupting the brain's higher-order functions first—including coordination and sensory processing (Ogden & Moskowitz, 2004).

Toomey and Wagenaar (2002) collate the 'higher quality' papers in this area, where the majority—58%—find significant evidence a higher MLDA decreases traffic fatalities.

A large number of Australia-focused papers are found in this area and find similar effects. Jiang et al. (2015) uses the 1970s MLDA policy changes in Queensland and Western Australia as an alternative specification to their RBT-focused model, finding a lower MLDA increased traffic fatalities. Finally, Smith and Burvill (1986) found lowering the drinking age in three Australian states adversely affected traffic safety, with a notable increase in driver and motorcyclist injuries in South Australia. Alternatively, Lindo et al. (2016) use a regression discontinuity design around NSW's MLDA of 18 and find no

evidence of effects on traffic fatalities.

3.2.4 Adolescent mortality rates

Adverse effects on mortality rates are also associated with MLDA policies. Plunk et al. (2016) find individuals subject to a lower MLDA in the US experience an increased risk of death from liver disease and some cancers. Tran et al. (2022) exploit the increase in the MLDA in Lithuania from 18 to 20, finding negative but insignificant effects on mortality rates among 18 and 19-year-olds when confounding factors are included.

3.3 Long-run effects

The literature on long-run effects driven by the MLDA, most relevant to this study, is sparse, underscoring the value of my contribution. Some papers have looked at education and labour market outcomes of younger adults, but none to my knowledge have looked at income, wealth or health, or later-life outcomes.

3.3.1 Education & labour market

Primarily, long-run MLDA-centred papers have focused on educational attainment, due to its reliability as a metric and high data availability, given the proximity in age of college graduates to the primary MLDA target group of 18 to 21-year-olds.

Cook and Moore (1993) find that a higher MLDA leads to a higher probability of college graduation, suggesting that educational attainment is inversely related to alcohol consumption. Lindo et al. (2013) use a regression discontinuity at age 21 and find university grades fall once students are subject to legal alcohol access.

One of the most interesting studies in the MLDA literature is Yörük (2015). They use a regression discontinuity design around the MLDA and find a 1 hour reduction in weekly working hours for 21-year-olds with legal alcohol access, though no effect on wages.

3.3.2 Alternate policies

By contrast, there are numerous interesting papers that look at similar non-MLDA policy effects on long-run outcomes, with most focusing on health.

Many have been found to induce long-run effects. While not absolute evidence of potential MLDA effects, this literature does help demonstrate that early policy interventions can have enduring ramifications.

Wherry and Meyer (2016) and Brown et al. (2019) investigate the long-term impact of US Medicaid insurance. They find that additional eligibility in childhood leads to respective declines in mortality rates and increases in income from age 23 onwards. Bailey et al. (2023) leverage the staggered rollout of the U.S. food stamps program and find greater access to social resources improves human capital, self-sufficiency and quality of neighbourhood residence in adulthood.

As another health-adjacent policy, it should be reasonable to expect the MLDA to produce similar effects.

Hence, to surmise, MLDA policies are found to be determinants of a number of short-run and long-run behaviours and outcomes. These results suggest long-term harms should be an important consideration for policymakers, as the literature demonstrates they are generally large and significant.

4 Data

I utilise the 'restricted' Household, Income and Labour Dynamics in Australia (HILDA) survey as my primary data source. Beginning in 2001, it is a longitudinal survey spanning 23 years, designed to capture a representative sample of Australian households. I use all waves (1-23) for estimation.

For identification, I use date of birth and state where their highest schooling grade was completed. The former determines an individual's age, and the latter proxies which state they lived in during their adolescent years. I assign treatment according to these variables, as they establish what MLDA an individual was subject to as a teenager. 'Treated' individuals are those who were subject to a new low MLDA of 18 as a result of 1970s reform, in either QLD, SA, or WA.

My final sample includes individuals born between 1945 and 1959, and who completed high school in 'treated' state units – QLD, SA, and WA – and 'control' state units – NSW and VIC. Summary statistics are presented in Table 2.

The dataset includes several behavioural variables (e.g., drinking, smoking) and outcome measures (e.g., income, wealth, health, and education). While HILDA does contain data on crime and gambling, these variables were excluded due to insufficient observations, as they are only available in the later waves.

There are no missing observations for most variables, with some notable exceptions. Wealth is surveyed every 4 waves in HILDA, and so appears for a quarter of the total observations. Individuals who did not complete HILDA's 'Self-Completion Questionnaire' (SCQ) did not have their drinking or smoking habits recorded. Moreover, only those recorded as 'active' drinkers were asked to input their average drinks per sitting, which makes up the 'binge drinking' variable.

Overall, the data appear plausible. To note, standard deviation is large, indicating substantial dispersion and presence of outliers, though this is as expected.

4.1 Sample selection

I choose the sample based on the variables available in HILDA.

Table 2: Summary statistics.

	Mean (SD)	Median	N
Key outcomes			
Annual income (2023–24 AUD	59 763 (60 986)	48 295	52 868
Household net worth (2023 Q2 AUD)	1 729 189 (2 133 888)	1 176 576	13 487
Long-term health conditions (%)	0.3849		52 859
University completion (%)	0.1564		52 868
Behavioural outcomes			
Drinker (%)	0.8351		48 287
Binge drinking (%)	0.1241		40 237
Smoker (%)	0.1467		48 216
Identifying variables			
Treated state (%)	0.3426		52 868
Treated state * Post (%)	0.2110		52 868
Year of birth	1952.57	1953	52 868
Treated units			
QLD (%)	0.1455		7 692
SA (%)	0.1244		6 579
WA (%)	0.0727		3 844
Control units			
NSW (%)	0.3619		19 131
VIC (%)	0.2955		15 622
Observations			52 868

Firstly, I exclude all individuals who do not have a valid date of birth, or state of highest schooling grade completion observation. About 26.7% of those born within the sample period are missing the latter. This reflects individuals either missing the survey waves when the question was asked (10, 14, 17), or not completing high school in Australia. Neither should impact the validity of my results.

Given the staggered adoption of a MLDA of 18, restricting the sample to birth cohorts 1945 to 1959 ensures a minimum of four pre-policy and six post-policy years of observations for each treated state. Though this range will result in unbalanced estimation periods – 8 pre & 6 post in QLD, 6 pre & 8 post in SA, and 4 pre & 10 post in WA – a balanced panel is required for my empirical strategy.

A consequence of this sample specification is that I am limited by the age range at which I can observe my variables of interest. By virtue of the HILDA interviews taking place between 2000 and 2021, the only ages that can be observed with a full panel are 55 to 63. While this does not compromise the goal of my study, testing long-run outcomes, it does limit its ability to estimate lifecycle dynamics. These outcomes are only observed towards the end of an individual's working life. For older cohorts, an alternative 20th century Australian panel dataset would be required - something that, to my knowledge, does not exist. To measure younger cohorts, newer HILDA waves could be used as they release.

To determine the age of participants, I use their date of interview, in conjunction with their date of birth. For the 13.7% of in-sample individuals whose interview date was not recorded, I impute a default of the 1st of September for their wave's corresponding interview year. This choice assumption was made given most HILDA interviews take place in August or September (64.5% of in-sample recordings). To ensure individuals are not duplicated at any age (since individuals can be interviewed for multiple waves within one non-calendar year), I increase the age of any older duplicate observations by one year and continue cycling through until there are no duplicates.

The result is a total of 52 868 observations.

One concern with my proxy for 'state of residence as an adolescent' - the state where

their highest schooling grade was completed - is interstate migration. If an individual moved states after leaving high school, then their outcomes would be attributed to the wrong state laws. This is a particular concern for early school leavers, of which compose a significant proportion of the sample - 52% completed up to a Year 10 equivalent or lower.

Data from the late-1970s indicates interstate migration rises at ages 19 and 20 following a trough in the late teens, yet remains below the national (all-age) average of roughly 1% per annum (Bell et al., 2018). Flows were strongly directed towards Queensland—and, to a lesser extent, Western Australia—with net outflows from New South Wales and Victoria and comparatively little from South Australia (Australian Bureau of Statistics, 2022).

If MLDA reform did worsen long-run outcomes, these observed migration patterns would induce upward (towards zero) bias: many treated or 'exposed' individuals would be misclassified into control states, compressing the treated–control gap. However, with these migration estimates small, the bias is unlikely to be of significance.

4.2 Key variables

My key variables can be separated into behaviours and outcomes of interest.

Alcohol and cigarette consumption act as 'behaviours,' i.e. variables that test the mechanism for long-run effects. HILDA provides granular data on an individual's drinking habits, specifically their consumption per day, and average drinks per sitting, and tobacco use, specifically their cigarette consumption per day. I convert these into binary indicators of whether an individual 'actively' drinks or smokes, and whether they 'binge' drink on average (5 or more drinks per sitting).

My key variables of interest can be split into four areas: education, health, income, and wealth. I choose a leading indicator from each area to create a complete, but straightforward picture of general long-run effects. These are: university completion (education), long-term condition development (health), annual income (income) and household net worth (wealth). I use binary indicators for the former two variables, and inflation-adjusted numerical and log-transformed indicators for the latter two.

Household net worth specifically requires two qualifiers for estimation. First, only

a household's highest income earner in a given period is counted. This avoids double counting, and inherently assumes a positive correlation between income and wealth accumulation. Second, since wealth data is limited in HILDA to every four waves, I measure household net worth within an age 'cluster.' This is to ensure there is a sufficient number of observations, comparable to the other key outcomes. To do so, I group observations into five-year bands around each target age (e.g., 53–57 for age 55). For each individual, I then take household net worth at the oldest age available within the age band. While most individuals have only one observation per band, this method ensures that, for those with multiple, wealth is consistently captured towards the top of the range.

I also utilise the Consumer Price Index (CPI) series from the Australian Bureau of Statistics (2025) to deflate income and wealth to 2023 Australian dollars.

I focus much of my analysis on outcomes measured between 55 and 60. This is to best avoid any endogeneity associated with retirement decisions.

Moreover, specifications with Western Australian (WA) observations are only provided for transparency. Given some WA year-of-birth clusters are very small (e.g. 4-8 observations) and contain some major outliers, specifications that include WA are unlikely to be truly representative.

5 Empirical strategy

My empirical strategy employs a staggered, synthetic differences-in-differences (SDID) framework (Arkhangelsky et al., 2021).

The SDID extends the conventional difference-in-differences framework, working to create synthetic control groups when no single control unit offers a valid counterfactual.It computes optimal time and unit weights for the control units to best match the pre-intervention trajectory of each treated unit. This should support causal inference as it creates a credible counterfactual to be used as the baseline for comparison in the analysis.

The staggered component accommodates non-uniform treatment timing, allowing each unit to receive treatment at a distinct point in time. However, since a balanced panel is required, there are uneven pre- and post-reform periods across treated units (as highlighted in Section 4). As such, a 15-year panel is chosen, and includes eligible individuals born between 1945 and 1959.

The SDID model I estimate can be specified as follows:

$$Y_{st} = \alpha + \tau \cdot (T_s \cdot P_{st}) + \epsilon_{st}$$

where:

- Y_{st} : outcome — for cohort in state s of birth-year t
- T_s : treatment — 1 if state s had MLDA reform
- P_{st} : post-treatment — 1 if cohort of birth-year t turned 21 (or 20) after state s had MLDA reform
- ϵ_{st}: error term

The coefficient of interest is τ, and the estimand, the average treatment effect on treated (ATT), is denoted as $\hat{\tau}$. The full expression for $\hat{\tau}$ is outlined in Appendix C.

The ATT, $\hat{\tau}$, can be found when the interaction term $T_s \cdot P_{st}$ equals one. In event time, these are cohorts of the same birth-year, residing in a reform state, who turned 21 (or 20) after their state had MLDA reform. This is the age they reach their state's pre-reform MLDA (21 in QLD and WA, 20 in SA), where all cohorts had legal access

to alcohol regardless of birth-year. This coefficient can be interpreted as the effect of a MLDA reform on a particular outcome, e.g. annual income or probability of attaining a university degree. Since the MLDA reforms I compare are not uniform in magnitude, this is notionally the effect of some 'large' drinking age reform.

This can be illustrated with an example. Suppose we want to estimate whether being allowed to drink alcohol at a younger age affects income later in life. The "treatment" is new exposure to a lower legal drinking age, and the outcome is income measured at age 55. The SDID compares the observed average income of cohorts newly eligible to drink with the trajectory of a synthetic control group composed of individuals who had always been eligible, weighted to match pre-exposure income trends. The resulting estimate captures the intent-to-treat effect of new drinking-age exposure on later-life income, measured relative to what would have been observed for the same cohort absent the policy change.

As such, treatment 'exposure' builds up over the first few event-time periods, given how old an individual was at the time of reform. For example, the Queensland cohort who turned 21 in 1974 would receive anywhere between 1 and 317 days of 'additional' legal access (given QLD's reform date was 18 February 1974), whereas cohorts who turned 21 in 1978 would receive all three 'additional' years. Thus, the SDID framework which captures the average treatment effect on treated is essential to capturing dynamic effects.

To note, this framework does not adjust the pre-policy trends of treated units. As such, if treated units display pre-trends that are uneven or volatile, there is little that can be done to smooth them. This framework offers one of the most robust ways to account for non-parallel pre-trends while preserving the capacity for causal inference.

Furthermore, covariates (whether as controls or fixed effects) are excluded from my model for three main reasons: (i) their inclusion may introduce bias into the weighting procedure, as the relationship between covariates and outcomes likely evolves differently over time across treatment and control groups; (ii) they may absorb variation essential for reliable SDID estimation of the average treatment effect on the treated (ATT); and (iii) several typical covariates are already used as outcome variables in my analysis (Kranz, 2022; Clarke et al., 2024).

6 Results

The results of the synthetic difference-in-differences regressions and corresponding event study plots are shown for each behavioural and outcome variable here. All are estimated using `sdid` and `sdid_event` packages on Stata, developed by Clarke et al. (2023) and Ciccia (2024) respectively.

I find limited evidence of a long-run MLDA effect on my behavioural indicators of drinking and smoking habits. For my key variables of interest, I do a highly significant long-run effects on annual income and the prevalence of long-term health conditions. Annual income is significantly lower for cohorts who experienced 1970s MLDA reform, and long-term health conditions are significantly more common. No robust effects are observed for household net worth or university completion. However, event-study reviews suggest the parallel trends assumption may not hold, despite the construction of a synthetic counterfactual.

My results are presented in two parts: behavioural outcomes and main outcomes. Behavioural outcomes are examined first, as they help explain how early exposure can generate long-run effects through the development of persistent habits that influence my key variables of interest.

6.1 Behavioural results

I begin by examining long-run effects on drinking and smoking habits and find no consistent evidence of a link. With some negative estimates, it may provide evidence that the mechanism from the MLDA to drinking habits is not as direct as expected.

Accordingly, though the mechanism of the MLDA's effect on alcohol (or cigarette) consumption for adolescents cannot be tested, I can test for the MLDA's effect on later-life habits.

In theory, a lower minimum legal drinking age which grants legal alcohol access to adolescents should lead to increased alcohol consumption among adolescents. As such, if dependence is shaped in adolescence, cohorts whose adolescent years occurred after MLDA reform in the affected states (QLD, SA, WA) would be expected to be relatively

'heavier' drinkers. This is because they would have received up to three 'additional' years of access to alcohol relative to older cohorts.

Tables 3 and 4 lists the regression results for the synthetic differences-in-differences models on the present-day drinking behaviour of individuals exposed to MLDA reform as a teenager.

Table 3: SDID regression on alcohol use (participation).

	Age = 56		Age = 58		Age = 60	
	(1)	(2)	(3)	(4)	(5)	(6)
ATT	-0.0339	-0.0083	-0.0201	-0.0321***	-0.0952***	-0.0425***
	(0.0353)	(0.0325)	(0.0261)	(0.0108)	(0.0410)	(0.0130)
Includes WA	✓	X	✓	X	✓	X
N	2196	2033	2266	2102	2289	2126

Notes: The dependent variable is a binary indicator for current alcohol consumption: 1 if an individual currently drinks; 0 otherwise. ATT, the average treatment effect on the treated, is the estimand. Robust standard errors are obtained by bootstrapping (2000 replications) and are reported in parentheses. Coefficients significantly different from zero are denoted: * 10%, ** 5%, and *** 1%.

Table 3 depicts my estimates for alcohol use, the extensive margin of drinking. Interestingly, most of my results predict that a lower MLDA reduced the expected proportion of adults in the treatment states who drink. In particular, the older specifications in columns (4) and (6) (age 58 and 60) are highly significant. Specifically, column (4) suggests that exposure to a MLDA change sees the probability of an individual in QLD, or SA drinking at age 58 fall by 3.21 percentage points (pp), on average, ceteris paribus. Assuming MLDA reform elevated youth consumption, this finding points to a life-cycle smoothing effect whereby individuals offset early increases with later reductions in alcohol use.

To visualise these results and examine the validity of the parallel trends assumption, Figure 2 plots event studies for alcohol use at age 58, with QLD and SA as the respective treated units. These should show how the reform's effects evolve overtime.

For QLD (left panel), there are two significant (negative) lags, at τ = -3 and τ = 2, i.e. the cohorts who turned 20 three years before, and two years following the law change, respectively. While the significant coefficient at lag τ = 2 intimates a smaller number

Figure 2: SDID event studies on alcohol use (participation).

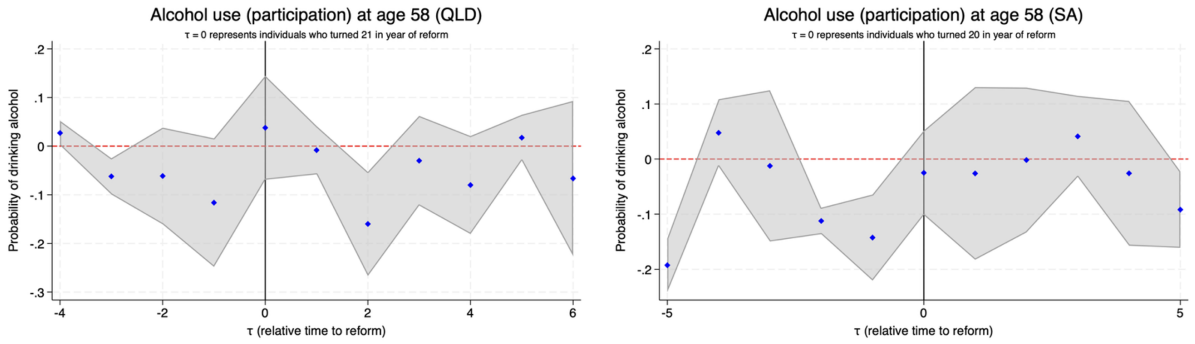

Notes: Left panel displays QLD as sole treatment state; right panel displays SA as sole treatment state. The dependent variable is a binary indicator of alcohol consumption (1 if an individual currently drinks; 0 otherwise). ATT, the average treatment effect on the treated, is the estimand. Pre-treatment coefficients and confidence intervals are obtained via placebo inference (bootstrapped to 2000 replications). QLD and SA are graphed separately because the placebo requires the number of control units (NSW, VIC) to exceed treated units.

of drinkers among Queenslanders born in 1955, in-line with the Table 2's regression estimates, a significant estimate at lag τ = -3 means significant pre-reform trends cannot be ruled out. Specifically, a downward pre-trend could be present.

SA (right panel) paints a more complex picture. There are three significant (negative) lags, and they all come prior to reform, at τ = -5, -2 and -1. It would almost certainly appear that there were downward pre-reform trends in SA for alcohol participation relative to control states. As such, the parallel trends assumption required for differences-in-differences estimation is probably violated, which would render these estimates invalid. Unfortunately, there are limited changes in specification that could address this issue, given most would also violate an identifying assumption necessary for causal inference. The interpretation and implications of this finding are discussed in detail with the main results, where similar concerns arise.

Table 4 provides estimates for prevalence of 'binge' drinking at middle-age. As a supplement to alcohol use, this indicator helps us in understanding the magnitude of individual-level alcohol consumption, the intensive margin.

Under all specifications, the ATT of binge drinking is highly insignificant, and has no clear sign, though it is more positive for older observations. Given the magnitude of coefficients match that of the 'alcohol use' indicator, they do appear to be of regular size.

Table 4: SDID regression on binge drinking.

	Age = 56		Age = 58		Age = 60	
	(1)	(2)	(3)	(4)	(5)	(6)
ATT	-0.0271	-0.0868	0.0188	-0.0314	0.0365	0.0614
	(0.0683)	(0.0855)	(0.0602)	(0.0765)	(0.0374)	(0.0463)
Includes WA	✓	X	✓	X	✓	X
N	1852	1710	1893	1749	1910	1765

Notes: The dependent variable is a binary indicator of 'binge' alcohol consumption (given an individual currently drinks): 1 if an individual has ≥ 5 drinks on average per occasion, 0 otherwise. ATT, the average treatment effect on the treated, is the estimand. Robust standard errors are obtained by bootstrapping (2000 replications) and are reported in parentheses. Coefficients that are significantly different from zero are denoted: * 10%, ** 5%, and *** 1%.

These results suggest MLDA reform had little effect on the probability of binge drinking in later-life. Though, measurement error associated with misreporting likely attenuates these estimates and limits its interpretability.

Thus, the results for drinking participation and binge drinking indicate no clear long-run effect on alcohol consumption. Of note, the coefficient for the youngest specification (age 56) is not significant and nearly positive, indicating the predicted negative effect may not hold for younger age groups (e.g. in youth). This indicates that an MLDA-induced increase in adolescent alcohol consumption is at least plausible.

For any potential complementary effects, I evaluate the effect of the MLDA reforms on smoking behaviour later-in-life. Table 5 provides these estimates.

Table 5: SDID regressions on smoking (participation).

	Age = 56		Age = 58		Age = 60	
	(1)	(2)	(3)	(4)	(5)	(6)
ATT	0.0255	-0.0020	0.0640	0.0459	0.1269**	0.0765
	(0.0302)	(0.0348)	(0.0408)	(0.0433)	(0.0514)	(0.0524)
Includes WA	✓	X	✓	X	✓	X
N	2194	2034	2257	2094	2291	2127

Notes: The dependent variable is a binary indicator for cigarette use: 1 if an individual currently smokes; 0 otherwise. ATT, the average treatment effect on the treated, is the estimand. Robust standard errors are obtained by bootstrapping (2000 replications) and are reported in parentheses. Coefficients significantly different from zero are denoted: * 10%, ** 5%, and *** 1%.

In most specifications, I find a positive, but insignificant effect of the MLDA reform

on the likelihood of smoking. Older specifications are again more positive and closer to significance, but not meaningfully.

Figure B1 of Appendix B plots the event studies for smoking at age 60. For SA, there are four successive significant positive pre-reform lags (τ = -4, -3, -2 and -1), which together, indicate a strong downward pre-trend for cigarette consumption in SA that exceeded that of NSW and VIC. Hence, the upward post-reform trend, and positive ATT, cannot be extrapolated on, as the parallel trends assumption is almost certainly violated.

As such, the coefficients for smoking participation are unreliable. Although QLD's estimates do not individually violate the parallel trends assumption, the overall ATT estimate for smoking is undermined by the inclusion of South Australia, whose violation of the assumption invalidates half of the treatment group.

Thus, whether the MLDA exerts long-run effects through changes in drinking behaviour remains uncertain. By statistical significance alone, it would appear the MLDA reduced alcohol participation in old-age, though the event studies highlight the identifying assumption is not satisfied, and thus, this cannot be asserted. Similarly, there is limited evidence in favour of effects on smoking behaviour, in-line with Deza (2015).

Ultimately, these behavioural tests provide only supplementary analysis of dependence effects. My main results for the outcomes of interest constitute the primary focus of this paper and may, alternatively, provide the grounds for robust inference.

6.2 Main results

My main results for long-run outcomes of interest yield several statistically significant findings, most notably for income and health, where estimates are large and robust. Evidence of effects on wealth is also present, albeit to a lesser extent. However, the likelihood that these effects are truly causal remains uncertain due to substantial concerns regarding the validity of the parallel trends assumption, which underpins identification. The SDID regressions on my key, long-run variables of interest are provided in Tables 6, 7, 8 and 9.

I first observe real annual income as an indicator of career success. I measure income at the lower end of the age window (55-61), to avoid any shocks associated with retirement. Table 6 provides these estimates of MLDA reform effects on annual income. To conserve table space, specifications that include Western Australia are reported in Table A1, found in Appendix A.

Table 6: SDID regressions on annual income.

	Age = 55		Age = 57		Age = 59	
	(1) Level	(2) Log	(3) Level	(4) Log	(5) Level	(6) Log
ATT	-11 924.2***	-0.093*	-9575.9***	-0.112	-13 307.7**	-0.051
	(876.1)	(0.056)	(2607.8)	(0.086)	(5423.2)	(0.116)
Includes WA	X	X	X	X	X	X
N	2085	2034	2174	2113	2239	2181

Notes: The dependent variables are annual income in 2023–24 Australian dollars, and log annual income where specified. ATT, the average treatment effect on the treated, is the estimand. Robust standard errors are obtained by bootstrapping (2000 replications) and are reported in parentheses. Coefficients significantly different from zero are denoted: * 10%, ** 5%, and *** 1%.

Per column (1), new exposure to a lower minimum legal drinking age is estimated to reduce an individual's annual income at age 55 by an average of $11 924, ceteris paribus. This is significant at the 1% level, and estimates of similar magnitude are found for at ages 57, and 59. Compared to the median in the sample of $48 300, this is also quite large, nearly a quarter of the size.

This result is not robust to alternate specifications, however. With the natural logarithm of income as the dependent variable, age 55 income maintains moderate significance, but age 57 and 59 income do not. This suggests that there may be some outliers in the dataset, likely high-income earners in control units, that pull the raw estimates significantly negative.

Thus, there is some evidence that new legal access to alcohol as a teenager had a meaningful effect on annual income in later-life. As theorised, the direction of this effect is negative, suggesting reform to a lower MLDA results in worse long-run outcomes for subject cohorts.

To assess these findings, I observe state-specific event studies for annual income of QLD and SA, found in Figure 4.

Figure 3: SDID event studies on annual income at age 55.

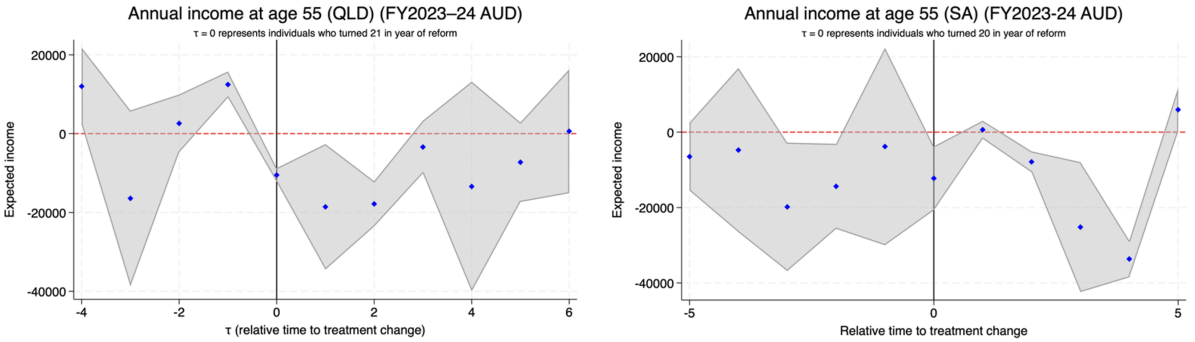

Notes: Left panel displays QLD as sole treatment state; right panel displays SA as sole treatment state. The dependent variable is annual income in 2023–24 Australian dollars. ATT, the average treatment effect on the treated, is the estimand. Pre-treatment coefficients and confidence intervals are obtained via placebo inference (bootstrapped to 2000 replications). QLD and SA are graphed separately because the placebo requires the number of control units (NSW, VIC) to exceed treated units.

While the event studies do highlight some significant (negative) post-reform lags for both states in-line with the estimated ATT — at $\tau = 1$ and 2 for QLD, and $\tau = 2$, 3, and 4 for SA — it does also depict multiple significant pre-reform lagged effects. There is no clear level difference or permanent trend however, with mostly noisy, non-monotonic differences between the treatment and control groups.

On its own, this would provide little evidence of systematic bias, where parallel trends could be considered 'plausible.' However, considering the low statistical power of my model, with small 'year of birth' clusters and few treated and control units, the confidence intervals are likely large, which would falsely rule out additional significant lags (Roth, 2022). In this case, an upward trend in both QLD and SA could have been plausibly ruled out.

As such, by visual check, a parallel trends assumption violation is probable, meaning any causal interpretations should be cautioned. This aligns with my takeaways from the event studies on drinking and smoking participation.

This is a critical mark on this paper. The key identifying assumption of differences-in-differences models cannot be met. Uncontrolled volatility in the treated group is leading

to a likely violation of the parallel trends assumption, implying that observed significance may simply reflect confounding differences.

In my models, the SDID estimator attempts to mitigate non-parallel pre-reform trends by learning unit and time weights that align donors to treated units. This smooths the trends of the control group, but for my data, it is the treatment group whose trends are volatile. Thus, pre-period fit is imperfect. The volatility in the outcomes of treated units stems from having very few observations within each birth cohort cluster, with counts between 10 and 15 common. This makes estimation prone to outliers and ungeneralisable to the population. By contrast, the control states (NSW, VIC) are highly populous, yielding cluster-level samples typically double or triple those of the treated states. Exact cluster counts for annual income at age 55 are provided in Table A2, found in Appendix A.

This pattern emerges under alternative age and log specifications also.

Event studies for real annual income at age 57 are pictured in Figure B2 of Appendix B highlight significant pre-reform trends in both QLD and SA. In QLD, there appears to just be a positive level difference, however in SA, there is some visual evidence of a downward pre-trend. With greater statistical power, these estimates would yield even stronger statistical significance. As such, the same conclusion is reached: the parallel-trends assumption is likely violated.

Switching the dependent variable to log income produces similar findings. Figure 5 displays these event studies at age 55, where a potential upward trend in QLD (left panel) rules out much likelihood of parallel trends. The volatility and unreliability of the data can be seen in the variation of the standard errors, with the confidence intervals at lags at τ = -4 and -3 in the SA plot (right panel) impossibly small.

Unfortunately, few credible alternatives to a differences-in-differences approach are available. Standard regression designs do not provide a credible counterfactual, any instrumental variable will likely be collinear with time, and regression discontinuity designs rely on a sharp treatment cutoff which MLDA reform does not supply. Synthetic control methods are most similar but suffer the same limitations as DID, only smoothing the

Figure 4: SDID event studies on log annual income at age 55.

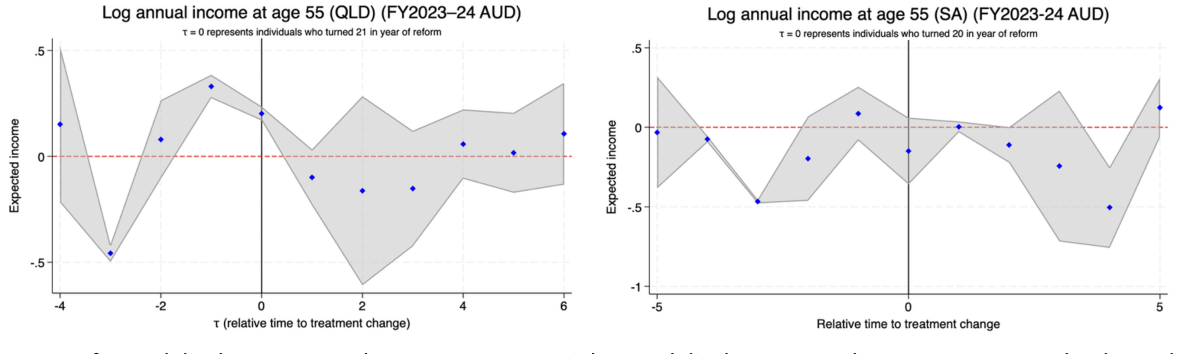

Notes: Left panel displays QLD as sole treatment state; right panel displays SA as sole treatment state. The dependent variable is log annual income in 2023–24 Australian dollars. ATT, the average treatment effect on the treated, is the estimand. Pre-treatment coefficients and confidence intervals are obtained via placebo inference (bootstrapped to 2000 replications). QLD and SA are graphed separately because the placebo requires the number of control units (NSW, VIC) to exceed treated units.

trends of the control group in addition to a lack of time weighting.

Thus, I cannot report valid estimates for annual income. Similarly and regrettably, the estimates of long-term health conditions, household net worth, and university completion are also not credible.

Next, I use household net worth to approximate lifetime wealth accumulation. I only measure it at two age specifications, 55 and 60, due to the limits of HILDA's wealth data. Table 7 reports these estimates.

Table 7: SDID regressions on household net worth.

	Age = 55			Age = 60		
	(1) Level	(2) Level	(3) Log	(4) Level	(5) Level	(6) Log
ATT	456 967*	142 517***	0.3168	121 819	245 023	0.7234
	(267 179)	(28 295)	(0.2763)	(345 217)	(546 447)	(0.4973)
Includes WA	✓	X	X	✓	X	X
N	1686	1562	1533	1814	1686	1649

Notes: The dependent variables are household net worth in 2023 Q2 Australian dollars, and log household net worth where specified. ATT, the average treatment effect on the treated, is the estimand. Each age specification includes the oldest observation within a five-year band around the specified age. Robust standard errors are obtained by bootstrapping (2000 replications) and are reported in parentheses. Coefficients that are significantly different from zero are denoted: * 10%, ** 5%, and *** 1%.

From column (2), the main specification is large, positive and highly significant, implying cohorts who turned 21 (or 20) after MLDA reform in treated states were able to

accumulate an additional $142 000 in wealth by age 55 approximately, on average, ceteris paribus. This is bizarre, given it is the opposite result to the negative income estimates. However, this result is not robust, either at age 60, or when the natural log is taken, implying it is likely imprecisely estimated.

Familiar issues arise once the event studies are observed, however, pictured in Figure B3 of Appendix B. With significant outliers prior to reform and considerable variability, there is limited support for the parallel trends assumption. Given the universal data limitations, this result would very likely hold for age 60 also.

However, since the parallel trends assumption may just potentially hold for QLD, I estimate QLD wealth. These estimates are provided in Table A3 of Appendix A. With highly significant estimates in both directions for different age specifications, there appears to be no clear MLDA-driven wealth effect, though these interpretations must be cautioned.

Thus, there is minimal evidence of a MLDA-driven effect on wealth.

I then estimate the MLDA's effect on the prevalence of long-term health conditions, my key measure of physical health. These estimates are displayed in Table 8.

Table 8: SDID regressions on prevalence of long-term health conditions.

	Age = 55		Age = 60		Age = 63	
	(1)	(2)	(3)	(4)	(5)	(6)
ATT	0.0255	0.0952***	0.1176***	0.1523***	0.0576**	0.0755***
	(0.0602)	(0.0141)	(0.0316)	(0.0309)	(0.0255)	(0.0050)
Includes WA	✓	X	✓	X	✓	X
N	2253	2085	2432	2258	2418	2244

Notes: The dependent variable is a binary indicator of a long-term health condition (1 if an individual reports a condition, 0 otherwise). ATT—the average treatment effect on the treated—is the estimand. Robust standard errors are obtained by bootstrapping (2,000 replications) and are reported in parentheses. Coefficients significantly different from zero are denoted: * 10%, ** 5%, and *** 1%.

There appears to be a significantly positive effect on health condition prevalence driven by the MLDA, with highly significant coefficients for all age specifications. By column (4), cohorts who gained legal access to alcohol after MLDA reform in treated states are expected to be 15.23 pp more likely to have developed a serious health condition by age 60, on average, ceteris paribus. This is quite stark, but in the expected direction. Interestingly, if the MLDA concurrently increased mortality rates, this estimate particularly

may be downward biased, as individuals who experience the worst health outcome – death – are not counted.

A review of the event studies in Figure 6 displays a lack of parallel trends, however. The QLD plot (left panel) displays an positive pre-trend outlier, and two additional near-significant lags (which would very likely be significant with more precise estimation), while the SA plot (right panel) appears to exhibit a downward trend.

Figure 5: SDID event studies on long-term health conditions at age 60.

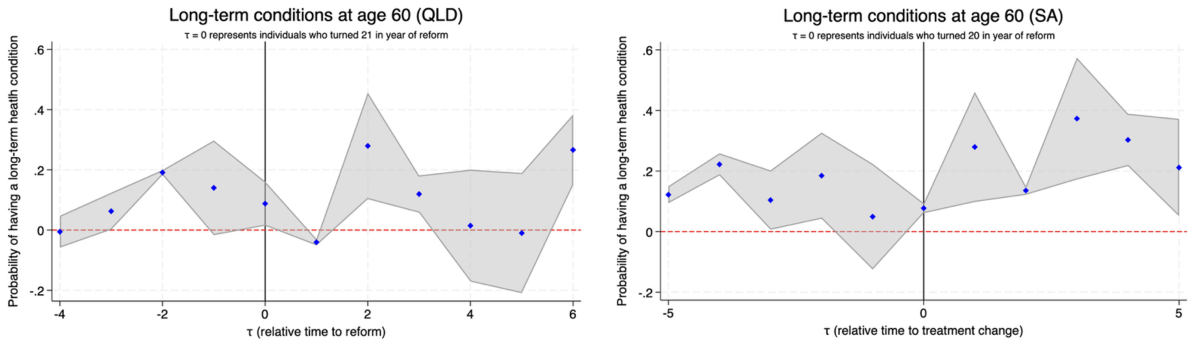

Notes: Left panel displays QLD as sole treatment state; right panel displays SA as sole treatment state. The dependent variable is a binary indicator of a long-term health condition: 1 if an individual reports a condition, 0 otherwise. ATT, the average treatment effect on the treated, is the estimand. Pre-treatment coefficients and confidence intervals are obtained via placebo inference (bootstrapped to 2000 replications). QLD and SA are graphed separately because the placebo requires the number of control units (NSW, VIC) to exceed treated units.

Thus, these health estimates should also be interpreted with significant caution.

Finally, I observe my measure of educational attainment: university completion. Table 9 reports these estimates.

All coefficients are negative and of regular size. Column (2) suggests cohorts exposed to a new MLDA of 18 in treated states were, on average, 4.88 pp less likely to have completed university at age 55, ceteris paribus. However, with the age 55 specification containing the only coefficient of marginal significance, it appears there is no robust effect of the MLDA on educational attainment.

As such, these synthetic differences-in-differences specifications appear to have failed to produce any valid causal estimates of the MLDA's effect on long-run outcomes. The significant coefficients I do estimate likely arise due to confounding factors. Namely,

Table 9: SDID regressions on university completion.

	Age = 55		Age = 60		Age = 63	
	(1)	(2)	(3)	(4)	(5)	(6)
ATT	-0.0005	-0.0488*	-0.0044	-0.0566	-0.0070	-0.0635
	(0.0551)	(0.0279)	(0.0591)	(0.0409)	(0.0638)	(0.0559)
Includes WA	✓	X	✓	X	✓	X
N	2253	2085	2432	2258	2419	2245

Notes: The dependent variable is a binary indicator of a university completion: 1 if an individual has a university degree, 0 otherwise. ATT, the average treatment effect on the treated, is the estimand. Robust standard errors are obtained by bootstrapping (2,000 replications) and are reported in parentheses. Coefficients significantly different from zero are denoted: * 10%, ** 5%, and *** 1%.

persistent economic differences between the smaller states who experienced reform and the larger 'control' states would explain much of this perceived 'gap,' or treatment effect.

If NSW and VIC were wealthier and had better healthcare on average compared to QLD and SA in the 1970s, which is widely known to be true, then a poor causal instrument would, though imprecisely, produce a negative treatment effect equal to the size of these differences.

This fact would explain the highly significant estimates for annual income and the prevalence of long-term health conditions, and should explain much of the results of this study.

7 Concluding remarks

My thesis finds limited evidence that the MLDA produces any long-run effect. Despite highly significant estimates for annual income and the prevalence of long-term health conditions, both of which point to a negative outcomes, a very likely invalidation of my identification strategy means causal interpretations cannot be made. Even with a synthetic counterfactual constructed using optimally weighted units, the resulting estimates lack credibility primarily due to small sample size. The 'parallel trends assumption' cannot be met, and as such, confounding differences would appear to explain much of the perceived effect.

As such, though this study fails to measure anything of significance, its questions retain much importance for policy. Continued research into the long-run effects of health shocks and the effectiveness of current health policies will remain as vital as ever to the development of economies and people globally. It is my hope that research on alcohol policy continues, ultimately enabling a clearer understanding of whether a MLDA of 18 is truly appropriate in Australia.

References

Ahammer, A., Bauernschuster, S., Halla, M., & Lachenmaier, H. (2022). Minimum legal drinking age and the social gradient in binge drinking. *Journal of Health Economics, 81*, 102571. https://doi.org/10.1016/j.jhealeco.2021.102571

Almond, D., Currie, J., & Duque, V. (2018). Childhood circumstances and adult outcomes: Act II. *Journal of Economic Literature, 56*(4), 1360–1446. https://doi.org/10.1257/jel.20171164

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., & Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review, 111*(12), 4088–4118. https://doi.org/10.1257/aer.20190159

Australian Bureau of Statistics. (2011). Apparent consumption of alcohol: Extended time series, 1944–45 to 2008–09 (Catalogue No. 4307.0.55.002). https://www.abs.gov.au/ausstats/abs@.nsf/Lookup/4307.0.55.002main+features31944-45+to+2008-09

Australian Bureau of Statistics. (2022). *Migration within Australia* [Presentation slides]. https://www.abs.gov.au/system/files/documents/4a8c0805128406c21e65414a486d332a/Internal%20migration%20slides_for%20web.pdf

Australian Bureau of Statistics. (2025). *Consumer Price Index, Australia: Latest release*. Australian Bureau of Statistics. https://www.abs.gov.au/statistics/economy/price-indexes-and-inflation/consumer-price-index-australia/latest-release

Babor, T. F., Kranzler, H. R., & Lauerman, R. J. (1987). Social drinking as a health and psychosocial risk factor: Anstie's limit revisited. In M. Galanter (Ed.), Recent developments in alcoholism: Volume 5. Memory deficits, sociology of treatment, ion channels, early problem drinking (pp. 373–396). Springer. https://doi.org/10.1007/978-1-4899-1684-6_16

Bailey, M. J., Hoynes, H., Rossin-Slater, M., & Walker, R. (2024). Is the social safety net a long-term investment? Large-scale evidence from the food stamps program. *The Review of Economic Studies, 91*(3), 1291–1330. https://doi.org/10.1093/restud

/rdad063

Baughman, R., Conlin, M., Dickert-Conlin, S., & Pepper, J. (2001). Slippery when wet: The effects of local alcohol access laws on highway safety. *Journal of Health Economics, 20*(6), 1089–1096. https://doi.org/10.1016/S0167-6296(01)00103-5

Bell, M., Wilson, T., Charles-Edwards, E., & Ueffing, P. (2018). Australia: The long-run decline in internal migration intensities. In T. Champion, T. Cooke, & I. Shuttleworth (Eds.), *Internal migration in the developed world: Are we becoming less mobile?* (1st ed., pp. 147–172). Routledge. https://doi.org/10.4324/9781315589282-7

Bowden, J. A., Bartram, A., Harrison, N. J., Norris, C. A., Kim, S., Pettigrew, S., Olver, I., Jenkinson, R., Bowshall, M., & Miller, C. (2024). Australian parents' attitudes, perceptions and supply of alcohol to adolescents: A national cross-sectional survey. *Health Promotion International, 39*(6), daae173. https://doi.org/10.1093/heapro/daae173

Brown, D. W., Kowalski, A. E., & Lurie, I. Z. (2020). Long-term impacts of childhood Medicaid expansions on outcomes in adulthood. *The Review of Economic Studies, 87*(2), 792–821. https://doi.org/10.1093/restud/rdz039

Carpenter, C., & Dobkin, C. (2009). The effect of alcohol consumption on mortality: Regression discontinuity evidence from the minimum drinking age. *American Economic Journal: Applied Economics, 1*(1), 164–182. https://doi.org/10.1257/app.1.1.164

Carpenter, C., & Dobkin, C. (2011). The minimum legal drinking age and public health. *Journal of Economic Perspectives, 25*(2), 133–156. https://doi.org/10.1257/jep.25.2.133

Carpenter, C., & Dobkin, C. (2015). The minimum legal drinking age and crime. *The Review of Economics and Statistics, 97*(2), 521–524. https://doi.org/10.1162/REST_a_00489

Ciccia, G. (2024). *sdid_event* [arXiv preprint 2407.09565]. arXiv. https://arxiv.org/abs/2407.09565

Clarke, D., Pailañir, D., Athey, S., & Imbens, G. (2023). *On synthetic*

difference-in-differences and related estimation methods (IZA Discussion Paper No. 15907). IZA Institute of Labor Economics. https://docs.iza.org/dp15907.pdf

Clarke, D., Pailañir, D., Athey, S., & Imbens, G. (2024). On synthetic difference-in-differences and related estimation methods in Stata. *The Stata Journal: Promoting Communications on Statistics and Stata, 24*(4), 557–598. https://doi.org/10.1177/1536867X241297914

Cook, P. J., & Moore, M. J. (1993). Drinking and schooling. *Journal of Health Economics, 12*(4), 411–429. https://doi.org/10.1016/0167-6296(93)90003-W

Crost, B., & Guerrero, S. (2012). The effect of alcohol availability on marijuana use: Evidence from the minimum legal drinking age. *Journal of Health Economics, 31*(1), 112–121. https://doi.org/10.1016/j.jhealeco.2011.12.005

Crost, B., & Rees, D. I. (2013). The minimum legal drinking age and marijuana use: New estimates from the NLSY97. *Journal of Health Economics, 32*(2), 474–476. https://doi.org/10.1016/j.jhealeco.2012.09.008

Dee, T. S. (1999). The complementarity of teen smoking and drinking. *Journal of Health Economics, 18*(6), 769–793. https://doi.org/10.1016/S0167-6296(99)00018-1

Deza, M. (2015). The effects of alcohol on the consumption of hard drugs: regression discontinuity evidence from the National Longitudinal Study of Youth, 1997. *Health Economics, 24*(4), 419–438. https://doi.org/10.1002/hec.3027

George, S., Rogers, R. D., & Duka, T. (2005). The acute effect of alcohol on decision making in social drinkers. *Psychopharmacology, 182*(1), 160–169. https://doi.org/10.1007/s00213-005-0057-9

Jacobus, J., & Tapert, S. F. (2013). Neurotoxic effects of alcohol in adolescence. *Annual Review of Clinical Psychology, 9*, 703–721. https://doi.org/10.1146/annurev-clinpsy-050212-185610

Jiang, H., Livingston, M., & Room, R. (2015). Alcohol consumption and fatal injuries in Australia before and after major traffic safety initiatives: A time series analysis. *Alcoholism: Clinical and Experimental Research, 39*(1), 175–183. https://doi.org/10.1111/acer.12609

Kranz, S. (2022). *Synthetic difference-in-differences with time-varying covariates* [Working paper]. GitHub. https://github.com/skranz/xsynthdid/blob/main/paper/synthdid_with_covariates.pdf

Krauss, M. J., Cavazos-Rehg, P. A., Agrawal, A., Bierut, L. J., & Grucza, R. A. (2015). Long-term effects of minimum legal drinking age laws on marijuana and other illicit drug use in adulthood. *Drug and Alcohol Dependence, 149*, 173–179. https://doi.org/10.1016/j.drugalcdep.2015.01.043

Lees, B., Meredith, L. R., Kirkland, A. E., Bryant, B. E., & Squeglia, L. M. (2020). Effect of alcohol use on the adolescent brain and behavior. *Pharmacology, Biochemistry, and Behavior, 192*, 172906. https://doi.org/10.1016/j.pbb.2020.172906

Lindo, J. M., Siminski, P., & Yerokhin, O. (2016). Breaking the link between legal access to alcohol and motor vehicle accidents: Evidence from New South Wales. *Health Economics, 25*(7), 908–928. https://doi.org/10.1002/hec.3197

Lindo, J. M., Swensen, I. D., & Waddell, G. R. (2013). Alcohol and student performance: Estimating the effect of legal access. *Journal of Health Economics, 32*(1), 22–32. https://doi.org/10.1016/j.jhealeco.2012.09.009

Mair, C., Lipperman-Kreda, S., Gruenewald, P. J., Bersamin, M., & Grube, J. W. (2015). Adolescent drinking risks associated with specific drinking contexts. *Alcoholism: Clinical and Experimental Research, 39*, 1705–1711. https://doi.org/10.1111/acer.12806

Miron, J. A., & Tetelbaum, E. (2009). Does the minimum legal drinking age save lives? *Economic Inquiry, 47*(2), 317–336. https://doi.org/10.1111/j.1465-7295.2008.00179.x

NFSA Films. (2017, November 28). Drinking driving surviving [Video]. YouTube. https://www.youtube.com/watch?v=GW-rn9zpVUo

Norberg, K. E., Bierut, L. J., Grucza, R. A., & Wang, J. C. (2009). Long-term effects of minimum legal drinking age laws on past-year alcohol and drug use disorders. *Alcoholism: Clinical and Experimental Research, 33*(12), 2180–2190. https://doi.org/10.1111/j.1530-0277.2009.01056.x

Ogden, E. J. D., & Moskowitz, H. (2004). Effects of alcohol and other drugs on driver performance. *Traffic Injury Prevention, 5*(3), 185–198. https://doi.org/10.1080/15389580490465201

Plunk, A. D., Cavazos-Rehg, P. A., Bierut, L. J., & Grucza, R. A. (2013). The persistent effects of minimum legal drinking age laws on drinking patterns later in life. *Alcoholism: Clinical and Experimental Research, 38*(6), 1623–1629. https://doi.org/10.1111/j.1530-0277.2012.01945.x

Plunk, A. D., Krauss, M. J., Syed-Mohammed, H., Hur, M., Cavazos-Rehg, P. A., Bierut, L. J., & Grucza, R. A. (2016). The impact of the minimum legal drinking age on alcohol-related chronic disease mortality. *Alcoholism: Clinical and Experimental Research, 40*(8), 1761–1768. https://doi.org/10.1111/acer.13123

Roche, A. M., Bywood, P. T., Borlagdan, J., Lunnay, B., Freeman, T., Lawton, L., Tovell, A., & Nicholas, R. (2007). Young people and alcohol: The role of cultural influences. National Centre for Education and Training on Addiction (NCETA), Flinders University. Published by DrinkWise Australia Ltd.

Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights, 4*(3), 305–322. https://doi.org/10.1257/aeri.20210236

Smith, J. P. (2015). Economic shocks, early life circumstances and later life outcomes: Introduction. *The Economic Journal, 125*(588), F306–F310. https://doi.org/10.1111/ecoj.12280

Smith, D. I., & Burvill, P. W. (1986). Effect on traffic safety of lowering the drinking age in three Australian states. *Journal of Drug Issues, 16*(2), 183–198. https://doi.org/10.1177/002204268601600206

Smith, D. I., & Burvill, P. W. (1987). Effect on juvenile crime of lowering the drinking age in three Australian states. *British Journal of Addiction, 82*(2), 181–188. https://doi.org/10.1111/j.1360-0443.1987.tb01458.x

Tetteh-Quarshie, S., & Risher, M.-L. (2023). Adolescent brain maturation and the neuropathological effects of binge drinking: A critical review. *Frontiers in*

Neuroscience, 16, 1040049. https://doi.org/10.3389/fnins.2022.1040049

Toomey, T. L., & Wagenaar, A. C. (2002). Effects of minimum drinking age laws: Review and analyses of the literature from 1960 to 2000. *Journal of Studies on Alcohol, Supplement*, (14), 206–225. https://doi.org/10.15288/jsas.2002.s14.206

Toomey, T. L., Nelson, T. F., & Lenk, K. M. (2009). The age-21 minimum legal drinking age: A case study linking past and current debates. *Addiction, 104*, 1958–1965. https://doi.org/10.1111/j.1360-0443.2009.02742.x

Toomey, T. L., Rosenfeld, C., & Wagenaar, A. C. (1996). The minimum legal drinking age: History, effectiveness, and ongoing debate. *Alcohol Health & Research World, 20*(4), 213–218. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6876521/

Tran, A., Jiang, H., Lange, S., Livingston, M., Manthey, J., Neufeld, M., Room, R., Štelemėkas, M., Telksnys, T., Petkevičienė, J., Radišauskas, R., & Rehm, J. (2022). The impact of increasing the minimum legal drinking age from 18 to 20 years in Lithuania on all-cause mortality in young adults: An interrupted time-series analysis. *Alcohol and Alcoholism, 57*(4), 513–519. https://doi.org/10.1093/alcalc/agab076

Wherry, L. R., & Meyer, B. D. (2016). Saving teens: Using a policy discontinuity to estimate the effects of Medicaid eligibility. *Journal of Human Resources, 51*(3), 556–588. https://doi.org/10.3368/jhr.51.3.0913-5918R1

World Health Organization. (2011, August 10). Action needed to reduce health impact of harmful alcohol use. https://www.who.int/news/item/10-08-2011-action-needed-to-reduce-health-impact-of-harmful-alcohol-use

Yörük, C. E. (2015). The effect of alcohol consumption on labor market outcomes of young adults: Evidence from minimum legal drinking age laws. *The B.E. Journal of Economic Analysis & Policy, 15*(3), 1297–1324. https://doi.org/10.1515/bejeap-2014-0104

Yörük, B. K., & Yörük, C. (2013). The impact of minimum legal drinking age laws on alcohol consumption, smoking, and marijuana use revisited. *Journal of Health Economics, 32*(2), 477–479. https://doi.org/10.1016/j.jhealeco.2012.09.007

Appendix

A Additional tables

Table A1: SDID regressions on annual income (WA).

	Age = 55	Age = 57	Age = 59
	(1)	(2)	(3)
ATT	7460.78	-1750.44	1032.80
	(13 904.11)	(10 149.48)	(10 893.39)
Includes WA	✓	✓	✓
N	2253	2347	2416

Notes. The dependent variables are annual income in 2023–24 Australian dollars. ATT, the average treatment effect on the treated, is the estimand. Robust standard errors are obtained by bootstrapping (2,000 replications) and are reported in parentheses. Coefficients significantly different from zero are denoted: * 10%, ** 5%, and *** 1%.

Table A2: Year of birth clusters for annual income at age 55.

YOB	1945	1946	1947	1948	1949	1950	1951	1952
NSW	33	37	37	42	38	51	46	40
VIC	29	30	38	46	39	36	27	44
QLD	9	18	19	18	12	11	15	21
SA	12	8	16	13	13	16	21	18
YOB	1953	1954	1955	1956	1957	1958	1959	
NSW	44	64	60	82	75	85	83	
VIC	34	48	59	59	55	56	72	
QLD	28	23	22	33	30	30	29	
SA	20	21	31	24	21	20	24	

Table A3: SDID regressions on household net worth (QLD).

	Age = 55		Age = 60	
	(1) Level	(2) Log	(4) Level	(5) Log
ATT	182,231.10***	0.6339***	-681,657.10***	-0.1838
	(38,574.37)	(0.0572)	(213,102.1)	(0.1397)
N	1362	1341	1463	1431

Notes. The dependent variables are household net worth in 2023 Q2 Australian dollars, and log household net worth where specified. ATT, the average treatment effect on the treated, is the estimand. Each age specification includes the oldest observation within a five-year band around the specified age. Robust standard errors are obtained via placebo inference (bootstrapped to 2000 replications) and are provided in parentheses. Coefficients significantly different from zero are denoted: * 10%, ** 5%, and *** 1%.

B Additional figures

Figure B1: SDID event studies on smoking at age 60 (participation).

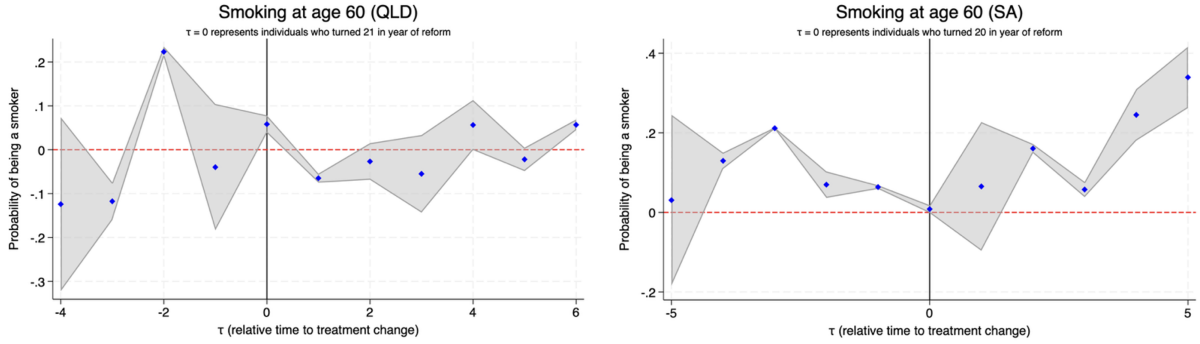

Notes: Left panel displays QLD as sole treatment state; right panel displays SA as sole treatment state. The dependent variable is a binary indicator for cigarette use: 1 if an individual currently smokes; 0 otherwise. ATT, the average treatment effect on the treated, is the estimand. Pre-treatment coefficients and confidence intervals are obtained via placebo inference (bootstrapped to 2000 replications). QLD and SA are graphed separately because the placebo requires the number of control units (NSW, VIC) to exceed treated units.

Figure B2: SDID event studies on annual income at age 57.

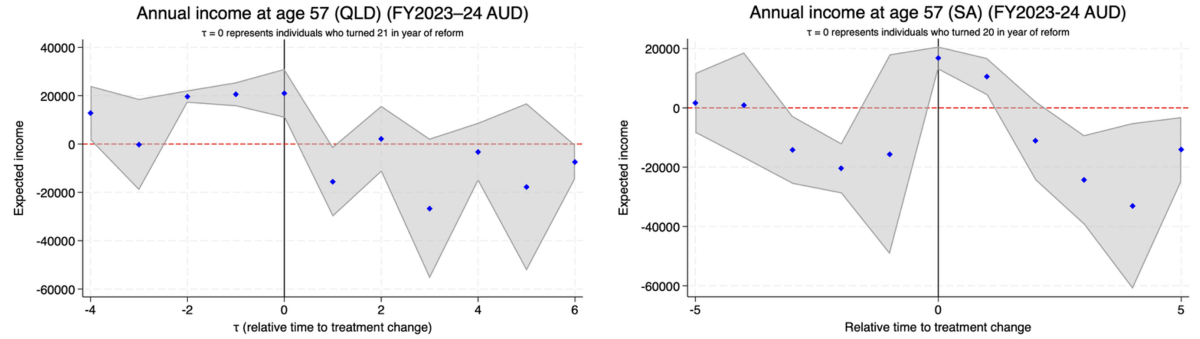

Notes: Left panel displays QLD as sole treatment state; right panel displays SA as sole treatment state. The dependent variable is annual income in 2023–24 Australian dollars. ATT, the average treatment effect on the treated, is the estimand. Pre-treatment coefficients and confidence intervals are obtained via placebo inference (bootstrapped to 2000 replications). QLD and SA are graphed separately because the placebo requires the number of control units (NSW, VIC) to exceed treated units.

Figure B3: SDID event studies on household net worth at age 55.

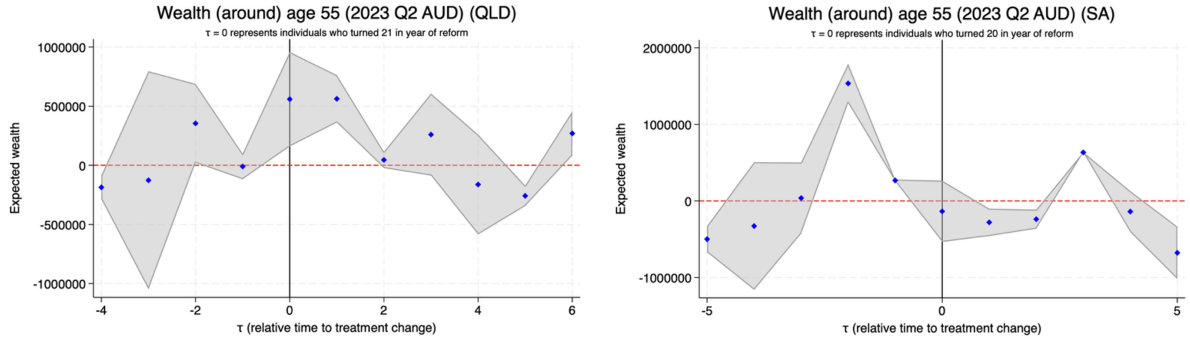

Notes: Left panel displays QLD as sole treatment state; right panel displays SA as sole treatment state. The dependent variable is household net worth in 2023 Q2 Australian dollars. ATT, the average treatment effect on the treated, is the estimand. Pre-treatment coefficients and confidence intervals are obtained via placebo inference (bootstrapped to 2000 replications). QLD and SA are graphed separately because the placebo requires the number of control units (NSW, VIC) to exceed treated units.

C Additional equations

The estimand, the average treatment effect on treated (ATT), denoted as $\hat{\tau}$, can be expressed as:

$$\hat{\tau} = \underbrace{\frac{1}{N_1} \sum_{s \in \mathcal{T}} \frac{1}{T_{1,s}} \sum_t P_{st}\, Y_{st}}_{\text{treated–post mean}} - \underbrace{\sum_{c \in \mathcal{C}} \hat{\omega}_c \, \frac{1}{T_1} \sum_t \bar{P}_t \, Y_{ct}}_{\text{synthetic counterfactual}} .$$

where:

- $\mathcal{T}$ / $\mathcal{C}$: sets of treated units and control units; $N_1 = |\mathcal{T}|$.
- $T_1 = \sum_t \bar{P}_t$: average post length across treated units; $T_{1,s} = \sum_t P_{st}$: number of post periods for treated unit s.
- $\bar{P}_t = \frac{1}{N_1} \sum_{s \in \mathcal{T}} P_{st}$: average treated post-indicator at time t.
- $\hat{\omega}_c \geq 0, \ \sum_{c \in \mathcal{C}} \hat{\omega}_c = 1$: control-unit weights.

The estimation of the treated post-mean uses weights to simply average outcomes across cohorts, whereas the synthetic counterfactual utilises the control-unit weights $\hat{\omega}_c$ optimised for SDID. Time-unit weights are used to learn the control-unit weights, but average out to one, meaning they do not appear in the ATT estimation.