

Kettlewell, Nathan; Zhang, Yuting

Working Paper

Financial Incentives and Private Health Insurance Demand on the Extensive and Intensive Margins

IZA Discussion Papers, No. 16248

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Kettlewell, Nathan; Zhang, Yuting (2023) : Financial Incentives and Private Health Insurance Demand on the Extensive and Intensive Margins, IZA Discussion Papers, No. 16248, Institute of Labor Economics (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/278946>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DISCUSSION PAPER SERIES

IZA DP No. 16248

**Financial Incentives and Private Health
Insurance Demand on the Extensive and
Intensive Margins**

Nathan Kettlewell
Yuting Zhang

JUNE 2023

DISCUSSION PAPER SERIES

IZA DP No. 16248

Financial Incentives and Private Health Insurance Demand on the Extensive and Intensive Margins

Nathan Kettlewell

University of Technology Sydney and IZA

Yuting Zhang

University of Melbourne

JUNE 2023

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Financial Incentives and Private Health Insurance Demand on the Extensive and Intensive Margins*

In countries with dual public and private healthcare systems, individuals are often incentivised to purchase private health insurance through subsidies and penalty. We use administrative data from Australia to study how high-income earners respond on both the intensive and extensive margins to the simultaneous withdrawal of a premium subsidy, and the increase of a tax penalty. We estimate regression discontinuity models by exploiting discontinuous changes in the penalty and subsidy rates. Our setting is particularly interesting because means testing creates different incentives at the extensive and intensive margins. Specifically, we could expect to see higher take-up of insurance coupled with downgrading to less expensive plans. We find evidence that the penalty – despite being large in value – only has a modest effect on take-up. Our results show little evidence of downgrading, which is consistent with a low price elasticity for the high-income earners we study.

JEL Classification: I13, I18, I12

Keywords: health insurance, tax penalty, regression discontinuity, Australia

Corresponding author:

Nathan Kettlewell
University of Technology Sydney
Economics Discipline Group
15 Broadway
Ultimo NSW 2007
Australia
E-mail: nathan.kettlewell@uts.edu.au

* For useful discussions and comments we thank Olena Stavrunova, Oleg Yerokhin and Jongsay Yong, as well as participants at seminars at the University of Sydney, 2022 Economics of health and wellbeing workshop, and the 2022 Australian Health Economics Society Conference. This research uses data from the Australian Taxation Office Longitudinal Information Files. All findings, opinions and conclusions are those of the authors and do not necessarily represent the views of the Australian Government or any of its agencies. Funding: This work is supported by the Australian Research Council (project ID FT200100630). Professor Zhang is the recipient of an Australian Research Council Australian Future Fellowship award funded by the Australian Government (project ID FT200100630).

1 Introduction

Many countries with dual public and private healthcare systems use financial incentives to encourage the take-up of private health insurance (PHI) either through government subsidies to those who buy, or a tax penalty to those who do not, or a mix of both ([Colombo and Tapay, 2004](#)). In Australia, since the late-1990s, the government has used both premium subsidies (known as rebates) and tax penalties (known as “Medicare Levy Surcharge”) to encourage people to buy PHI.

Currently in Australia both rebates and the Medicare Levy Surcharge (MLS) are means tested and share the same income thresholds for each tier. The means testing begins at incomes above \$AU90,000 for singles and \$AU180,000 for families. This means the policy mainly affects high income earners (in our administrative tax data, only around 18% of people are affected). There are three discrete thresholds at which the rebate (a fixed percentage of the premium) decreases, while simultaneously the MLS (a fixed percentage of income) increases (see Table 1 for precise details). The intent is for the increase in the MLS penalty (which increases the incentive to insure) to offset the reduced incentive to insure due to the lower rebate. We estimate the net effect of the MLS/rebate policy on demand for PHI, considering both the extensive (take-up) and intensive (premium spending) margins.

The Australian MLS/rebate policy is an interesting case study because it attempts to support demand for PHI in a fiscally (and equitably) mindful way. This will be increasingly important for nations dealing with ageing populations, increasing health care costs, and substantial fiscal constraints (see [Lorenzoni et al., 2019](#); [WHO, 2020](#)). Government subsidies are very commonly used to encourage people to take-up PHI, but it is very expensive. In comparison, tax penalties are not as commonly used and under studied, and if they can be used effectively together with subsidies, more countries could consider implementing them.

Like many insurance markets, Australian PHI is subject to community rating, whereby everyone can purchase insurance for the same price, with insurers sharing the risk through a risk equalization scheme. While equitable for consumers, this setting can lead to a vicious

cycle of adverse selection where healthier people, who are forced to pay substantially above actuarially fair rates for insurance, drop out, putting upward pressure on premiums, leading to further drop outs. Such a pattern was observed in Australia in the late 1990s and the MLS and rebate were introduced to address it (Smits et al., 2022). However, rebates are expensive and in 2012 the new system of means-tested incentives was introduced to better target this scheme. This was controversial at the time – the opposition government opposed the measure and cited analysis by Deloitte that predicted that over five years “1.6 million Australians would drop cover and 4.3 million would downgrade their cover”.¹ In reality, these claims turned out to be overstated, with the trend in coverage remaining fairly stable through this period, and the introduction of means testing associated with an overall small increase in take-up of PHI (Bilgrami et al., 2021). However, while the extent of downgrading was likely exaggerated, the policy does in theory create an incentive to downgrade. Our study is the first to investigate this empirically.

Another important feature of the Australian MLS/rebate policy is that it is ideally structured for empirical evaluation. The policy tier thresholds provide numerous sources of exogenous variation where the financial incentives to purchase insurance change, providing opportunities to understand how price responsive consumers are in a dual public-private system at different levels of income. The availability of large administrative data also allows us to obtain precise estimates for the effect of this policy.

To formally estimate the effect of the MLS and rebate changes, we estimate regression discontinuity design (RDD) specifications separately for each income threshold. We use a 10% random sample of all registered tax-filers in Australia from the 2017-18 financial year, which is the latest year available at the time of writing.

Intuitively, in an RDD design, we compare people with incomes just above and below the relevant threshold. We can obtain causal estimates if there are no other discontinuous changes in outcomes at the income threshold which could affect PHI coverage. In our setting,

¹Second reading speech by the shadow Minister for Health and Ageing, Peter Dutton, on 9 February 2012 regarding the *Fairer Private Health Insurance Incentives Bill 2011*.

this assumption may not hold due to sorting on the running variable in an effort to avoid the penalty. To deal with this, we identify a window around the threshold where sorting occurs and remove those observations from the analysis ([Barreca et al., 2016](#)), essentially extrapolating through that range of the running variable.

Our baseline estimates indicate that the net effect of the MLS and rebate reduction at the tier 1 threshold is an increase in PHI take-up of 1.1-1.4 percentage points (ppts) for families and 3-3.5 ppts for singles, but no robust effects at other thresholds. We also estimate a jump in premium spending of \$44-\$55 at the families tier 1 threshold and \$124-\$125 at the singles threshold, but again, no effect at other thresholds. Considering that at the tier 1 thresholds where the MLS kicks in, around 90% of families and around 70% of single people are already insured, the policy only explains a small fraction of demand. Moreover, the fact that people do not downgrade their cover, even at higher thresholds, where the withdrawal of the penalty provides a clear incentive to do so, suggests the high income earners affected by these policies are highly price inelastic.

2 Background and literature review

2.1 Australian healthcare system

Australia has universal public health insurance – Medicare, which covers free hospital treatment in public hospitals, subsidised medications and primary care doctors and specialist treatment. Like many countries with universal public health insurance, non-urgent hospital treatment has a waiting list in public hospitals. To manage this, Australia relies on a parallel private health insurance (PHI) system to encourage people to use more private hospitals.

As of March 2022, 45.1% of Australians purchased PHI to cover hospital treatment ([Australian Prudential Regulation Authority, 2022](#)). Having PHI does not preclude people from using public hospitals, and patients can choose to be public or private patients in public hospitals. While people without PHI can pay out-of-pocket for private treatment, few do so

in practice due to the high costs. By choosing to be private patients either in public hospitals or private hospitals, individuals can bypass waiting times in public hospitals and sometimes have more freedom to choose care providers.

The high PHI take-up rate is partially driven by three financial incentives implemented by the Australian government between 1997-2000: an income-based tax penalty – the MLS – introduced in 1997, PHI rebates in 1999, and Lifetime Health Cover loading (LHC) in 2000. The MLS imposes additional income tax for people who earn above a certain threshold and do not hold private hospital cover.² Rebates are government subsidies for PHI premiums for those with incomes below certain thresholds. LHC is an age-based premium penalty to encourage individuals to enrol earlier in life. Specifically, if people do not have hospital cover following their 31st birthday and decide to buy after, they have to pay a 2% loading on top of their hospital premium for every year they are aged over 30.

In the past twenty years, the Australian government has made some changes to the above three incentives. For example, age-specific premium rebates were introduced in 2005 to increase rebates for adults older than 65, and rebates became means-tested, and their growth capped in 2012. The MLS thresholds have increased over time and since 2012 there have been three different rates for the MLS aligned with the means testing of the PHI rebate. Nevertheless, these three incentives have largely maintained their core structures.

The justification for these government interventions is that, if more people buy PHI and use the private system, it may alleviate pressure on the public system. However, it is not clear how much the substitution of private care for public care has been occurring. In addition, these policies involve a significant sum of public spending: in the fiscal year of 2019-2020, the Australian government spent \$AU6.7 billion in PHI rebates to reduce premiums ([Hunt, 2021](#)) and, combined with State and Territory Government expenditure, \$AU6.1 billion to fund

²The definition of income for the purpose of the MLS is taxable income (excluding the net amount on which family trust distribution tax has been paid or any assessable first home super saver released amount), reportable fringe benefits, net investment losses, reportable super (compulsory pension fund) contributions and for couples, the share of the net income of a trust on which the trustee must pay tax and which has not been included in their taxable income. This is known as 'income for MLS purposes' in the tax system.

services provided in private hospitals ([Australian Institute of Health and Welfare, 2022](#)).

2.2 MLS and rebate policy settings and consumer incentives

Table 1 sets out the schedule of MLS and rebate rates by income and age in the 2017-18 financial year (the year corresponding to our dataset). Family threshold is applied to single parents and couples (including de facto couples) with and without children. In addition, for families with children, the income thresholds are increased by \$AU1,500 for each child after the first.³ The policy settings create a number of sharp discontinuities in the average tax rate for those who do not insure. For example, a single person earning \$AU90,000 will need to pay \$900 in additional tax if they do not have private hospital insurance.⁴

Table 1: MLS and rebates by income and age in FY2017-18


Threshold	Base tier	Tier 1	Tier 2	Tier 3
Single threshold	\$90k or less	\$90-\$105k	\$105-\$140k	\$140k +
Family threshold	\$180k or less	\$180-\$210k	\$210-\$280k	\$280k +
MLS	0%	1%	1.25%	1.5%
Rebate				
Under 65	25.934%	17.289%	8.644%	Not eligible
65-69	30.256%	21.612%	12.966%	Not eligible
70 and over	34.579%	25.934%	17.289%	Not eligible

To better understand how these policy settings affect incentives to join the private insurance pool, Figure 1 shows the premium cost after rebates and MLS liability of a single person aged under 65 years for different incomes for an entry level policy costing \$AU1,011 in 2017. This premium is equal to the value of the cheapest policy in New South Wales listed on the government policy comparison website www.privatehealth.gov.au on 19 May 2023, adjusted for growth in premiums using published figures by the Department of Health and Aged Care ([Department of Health and Aged Care, 2023](#)).

³In 2022-23, the MLS rates and policy thresholds are the same. The rebate rates are slightly lower as growth in the government’s spending on the rebate is capped at a rate of inflation, with the rebate rates adjusted to match this growth each year.

⁴In some cases the liability is less than this because MLS is only payable on taxable income, but ‘income for MLS purposes’ (see footnote 2) is used to determine the MLS tier.

Figure 1: Cost of basic PHI and MLS penalty by income



Notes: In the example, the cost of insuring is \$AU1,011 per year. Green line denotes premiums after rebates.

Since the net cost of PHI is actually less than the MLS penalty at the policy threshold, a person who only cares about maximizing income should choose to insure rather than pay the penalty, even if their willingness to pay for insurance is zero. In practice, many people choose not to buy PHI and pay the MLS instead – in our data around 10% of people who are liable. This may be due to misunderstandings about the financial implications of not insuring, or due to mistakes in predicting their income. It may also partly represent a preference to pay additional tax to fund public services, rather than to increase the revenue of private insurers, even when it is financially costly to do so.⁵

According to the figure, we should expect a large increase in take-up (extensive margin) of insurance at the tier 1 threshold, but not necessarily at the higher thresholds, because the net cost of insurance is already strongly negative when the tier 1 threshold kicks in. If people

⁵This seems to be the case for some Australians. See for example [this discussion](#) on Reddit. Failure to purchase PHI might also reflect the hassle cost involved with purchasing cover.

have not purchased PHI after reaching the tier 1 threshold, they may be highly resistant to the incentives. We also expect to see a decrease in average spending on insurance at each threshold, since the reduction in rebate acts like an increase in price. For people under 65, the increase in price at tier 1 is 11.7% of premiums, at tier 2 is 10.5% and at tier 3 is 9.5%. This increase in price affects both ‘marginal joiners’ (those who purchase PHI due to the MLS) as well as those who would purchase PHI regardless. From a policy perspective, both take-up and premium spending matter, since people who purchase less expensive insurance may have gaps in their coverage that mean they still choose to use the public system even if they are privately insured, undermining the motivation for encouraging these people to buy PHI in the first place.

2.3 Literature review on the effects of MLS and rebates

The three policy changes discussed above – the MLS, rebates and LHC – were effective in increasing PHI take-up immediately after they were introduced: the percentage of people with PHI rose from 31 per cent in 1999 to 45 per cent at the end of 2001 ([Frech et al., 2003](#)). Prior work has estimated the effects of tax and price incentives on the demand for PHI ([Frech et al., 2003](#); [Cheng, 2014](#); [Ellis and Savage, 2008](#); [Palangkaraya and Yong, 2005](#)), especially focusing on the combination of the three policies when introduced during 1997-2000. In these earlier studies, it is difficult to disentangle the independent effects of the three different policy incentives because they were implemented within such a short period of time.

Although there is an extensive literature on the effect of subsidies on the demand for health insurance internationally ([Finkelstein, 2002](#); [Finkelstein et al., 2019](#); [Frean et al., 2017](#); [Gruber and Washington, 2005](#); [King and Mossialos, 2005](#); [Nicolás and Vera-Hernández, 2008](#); [Rodríguez and Stoyanova, 2008](#)) as well as in the Australian context ([Bilgrami et al., 2021](#); [Cheng, 2014](#); [Kettlewell et al., 2018](#); [Palangkaraya and Yong, 2005](#)), the effect of financial penalty is relatively under studied and less certain ([Hackmann et al., 2015](#); [Buchmueller](#)

et al., 2021; Stavrunova and Yerokhin, 2014).

The paper closest to ours is Stavrunova and Yerokhin (2014). They used an RDD approach to study the effect of the MLS on take-up of PHI. They focused on childless singles only and used data in the financial year 2007-2008, when the MLS only had one tier with a much lower income threshold at \$AU50,000 for singles. In addition, they did not have premium information so could not study effects on the intensive margin. They concluded that the MLS increased PHI coverage by 6.5 percentage points (or 15.6%). Other related studies are Kang et al. (2015), Gong and Gao (2018) and Kang et al. (2019) which also used RDD to study the MLS, but were more focused around methodological challenges associated with RDD and used even earlier tax data. They estimated a jump in take-up ranging from 12%-22% at the threshold. Buchmueller et al. (2021) used a different approach, relying on income fluctuations over time and a large increase in the MLS threshold that occurred in 2008. They conclude that being subject to the MLS increases the probability of purchasing PHI by between 2 to 3 percent that year, but due to state dependence, the probability is estimated to increase over the following years if the individual remains liable for the MLS penalty.

None of these studies evaluated the net effect of the new means-tested MLS/rebate settings. The only study to focus on this period is Bilgrami et al. (2021). They used survey data and a differences-in-differences design to evaluate the effect of the 2012 PHI policy changes on PHI take-up, finding that the net effect of the 2012 policy changes was an increase in the PHI take-up rate by 1.5 percentage points on average. The key limitation is that high- and low-income earners are not entirely comparable, particularly since income is measured with error in surveys, it is difficult to identify individuals who fall into MLS exemption categories, and the treatment and control groups are not fixed over time.

Our study extends the existing literature in several ways. First, with the exception of Bilgrami et al. (2021), we are the first to evaluate the effect of the MLS/rebate in the post means testing period. We use administrative tax data and an RDD design, which identifies

causal effects under arguably weaker assumptions (which importantly, have simple testable implications). Second, we are the first study to consider how the incentives affect decisions at the intensive margin. This is important since, in theory, the withdrawal of premium rebate and onset of MLS penalty should lead to people purchasing less expensive insurance, a concern raised by insurers and politicians when means testing was introduced. Our study also benefits from using recently available tax data, which gives us a much larger sample than previous studies, and allows us to look at singles and families separately.

3 Data and methods

3.1 Data

Dataset. We use the Australian Taxation Office’s (ATO’s) ALife dataset, which is a 10% random sample of all registered tax-filers in Australia (see [Abhayaratna et al., 2022](#), for further details). We use data from the 2017-18 financial year, which is the latest year available in ALife at the time of writing. This year also corresponds to a period where the MLS tiers had been stable for four consecutive financial years. This is helpful for us because it implies certainty for tax-filers around the policy environment, and considerable time to adjust to these settings. Our data comprise individuals’ tax and superannuation records and includes detailed information on all their income sources, such as salary and wages, government pension and allowances, annuities and superannuation, interests, and dividends.

Private health insurers provide details to the ATO about whether an individual has insurance, during what period they were insured, how much they have paid in PHI premiums, and how much they receive as rebates for the financial year.⁶ We classify people as privately insured if they were covered by a policy expiring after the end of the 2017-18 financial year. In a small number of cases, we also classify people as insured who were not classified as

⁶Our custom release of ALife has this information from insurers; however, it is not available in the standard release. The standard release only includes an indicator for PHI status and number of days covered as self-reported in the tax return.

insured by this definition, but reported being covered by a complying PHI policy for the entire year on the MLS exemption section of the tax return. Premium expenditure is the total across all policies and does not include the PHI rebate.⁷ We also observe rebates separately. For family policies, people nominate the fraction of premium they paid (this has implications for how much rebate they will receive). Note it is not possible to link spouses in ALife, but ALife does contain information on spousal incomes.

Sample construction. We exclude those who are exempted for MLS. This includes foreign residents, members of Veterans’ Affairs and members of the defence force, blind pensioners and sickness allowance recipients. We also remove a small fraction of people (around 1%) who reported a relationship status change in the year as the MLS implications are complicated for this group.

People are classified as a family if they report having a spouse or dependent children.⁸ To account for the fact that the threshold increases by \$1,500 for each dependent child after the first one, we normalize income for people in families by deducting this additional amount (e.g., for someone with 3 dependent children we re-scale their income as $I - 3000$).

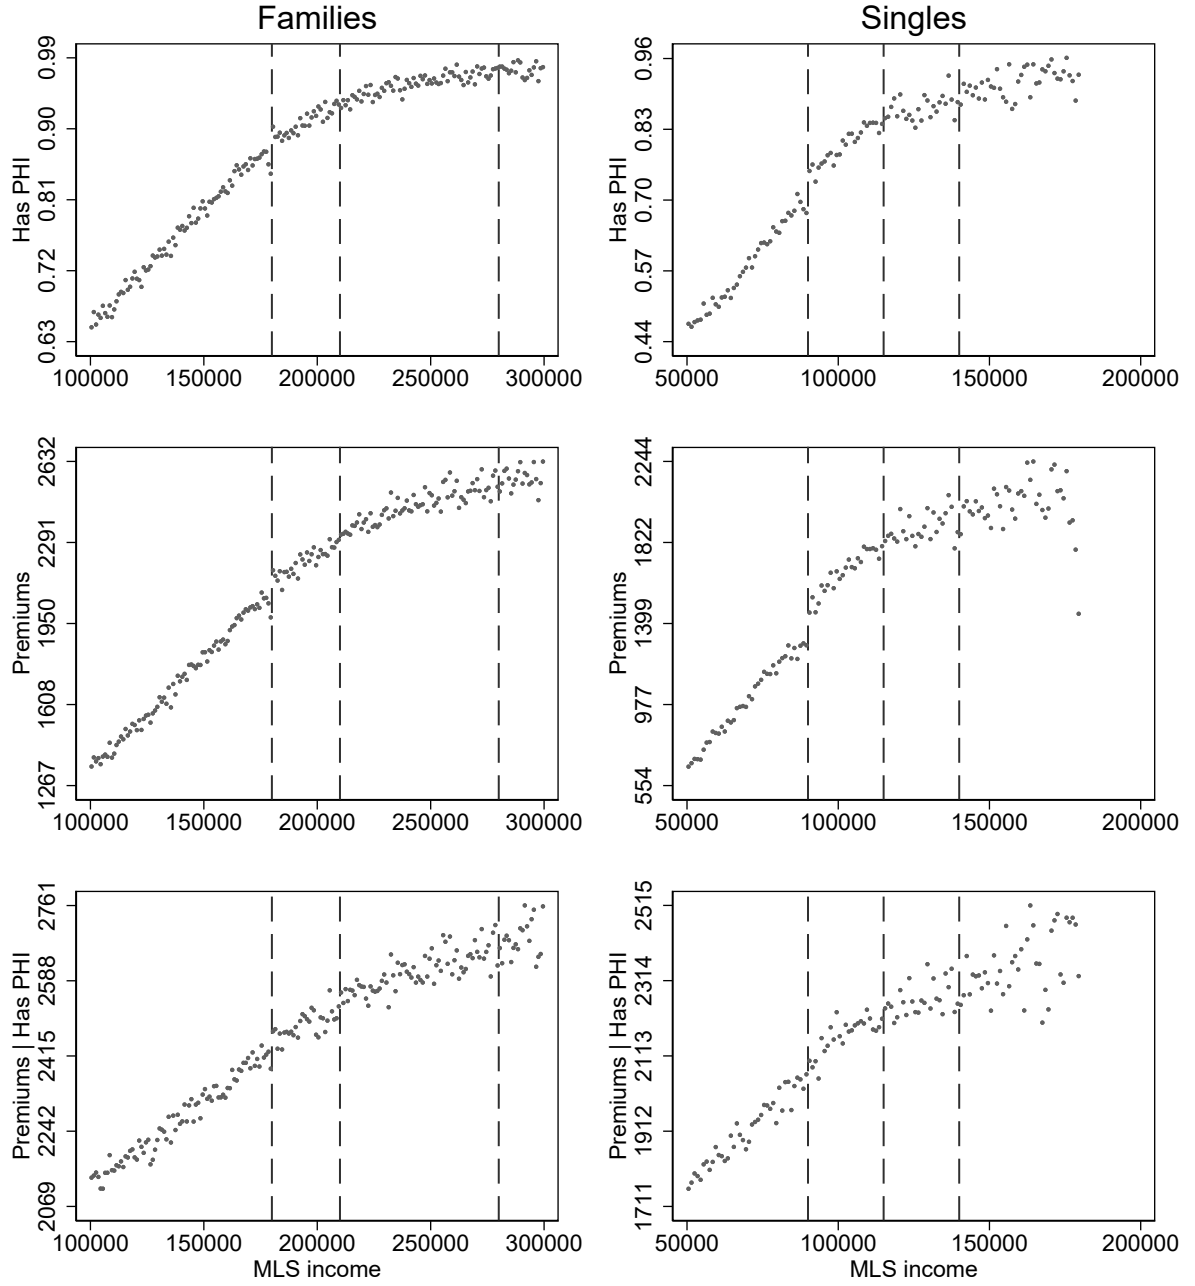
Our analysis sample includes 1,227,962 people, of which 18% have income above the MLS tier 1 threshold, 12% above tier 2 and 6% above tier 3.

Descriptive evidence. Figure 2 provides a visual overview for the relationships between income and PHI take-up, and between income and annualized premium expenditure. There is a clear positive correlation between income and insurance coverage for both singles and families – those with higher incomes are both more likely to be insured and spend more on premiums. There is also a visually clear discontinuity at the tier 1 thresholds, but not necessarily at other thresholds.

⁷Due to some suspiciously large values, we top-code this variable at the 99th percentile. This is inconsequential to our results.

⁸Due to a non-trivial number of missing values for the tax return item on dependents for Medicare Levy purposes, on advice from the ATO we use the derived variable *c_depend_child* to identify if a person has dependent children, which uses details from the current tax return as well as imputation from previous years. We drop observations where this variable is missing since we cannot accurately classify these people as single/family.

Figure 2: The relationship between income and PHI coverage



Notes: ALife 2017-18 release. Income is grouped into \$1,000 bins.

3.2 Econometric model

Estimation equation. To formally estimate the effect of the MLS and rebate changes on PHI take-up and premium expenditure, we estimate regression discontinuity design (RDD) specifications separately for each tier threshold and separately for singles and families. Our basic estimation equation is given by:

$$Y_i = \alpha_1 + \alpha_2.f(\text{income}_i - C) + \beta T_i + \alpha_3 T_i.f(\text{income}_i - C) + \epsilon_i \quad (1)$$

Y_i is the dependent variable (either an indicator for if the individual i has PHI or their total premium expenditure for the financial year), $f(\text{income}_i - C)$ is income for MLS purposes (own income plus spouse's income for those who are coupled) centered at zero at the value of the relevant MLS tier, T_i is an indicator for having income above the MLS tier being analysed and ϵ_i is a stochastic error term.

Identification. The intuition behind RDD is to compare people with incomes just above/below the relevant threshold, who are expected to be a good counterfactual for each other and differ only with respect to their policy regime. In this case, $\hat{\beta}$ is the causal effect of the policy change on Y_i for those at the MLS threshold provided the continuity assumption is satisfied ([Hahn et al., 2001](#)). This requires there are no other discontinuous changes in outcomes at the income threshold which could affect PHI coverage, an assumption that can be violated if there is sorting into treatment or control groups at the threshold. In practice, this assumption may not hold in our setting. This is because the steep penalty from the MLS creates a strong incentive for income manipulation (for example by over-stating deductions or working fewer hours) for people with incomes just above the tier 1 threshold, which has been empirically documented in other studies (e.g., [Hamilton, 2018](#); [Stavrunova and Yerokhin, 2014](#)).

In Appendix Figures A1-A2 we explore whether there is bunching in our data. As expected, there is strong visual evidence of this at the singles tier 1 threshold. Moreover, we

also see bunching just below this threshold around \$87,000, which is where the marginal tax rate changes from 32.5% to 37%.⁹ Bunching is also present at the families tier 1 threshold. There is little evidence of bunching at any of the higher tiers, likely reflecting that the incentives for sorting are much weaker at these thresholds.

We address bunching by creating ‘donut holes’ around the thresholds where we observe it (Barreca et al., 2016). Specifically, we follow Stavrunova and Yerokhin (2014) by collapsing our data into \$100 income bins¹⁰ and fitting the following specification before we estimate our RDD models:

$$D_i = \alpha_0 + \sum_{j=1}^3 \alpha_j income_i^j + \sum_{b=k}^K \delta_b bin_i + \epsilon_i \quad (2)$$

In Eq. 2, the dependent variable D_i is the count of individuals in income bin i . The model controls for a cubic in income along with a series of indicator variables for bins within the range of $[k, K]$ of the relevant threshold. We set this threshold to $\pm \$1,200$. After estimating Eq. 2, we sequentially test the statistical significance of each δ_b moving up and down from k and K .¹¹ Working up from k , we find the first income bin where the bin dummy coefficient is statistically significant ($p < 0.05$), and set that as our lower value for the donut. We perform an identical exercise working down from the upper limit. We then drop observations within this donut. Appendix Figures A1-A2 show the results of this exercise for each threshold.

RDD specification selections. To operationalize RDD we need to make several decisions about specifications. Arguably the most critical of these is how much data to use above and below the given policy threshold (the bandwidth), which involves a trade-off between bias and variance. Algorithms that aim to minimize the asymptotic mean-squared error

⁹There is also a marginal tax rate increase at \$180,000 from 37% to 45%. This creates an additional reason for bunching at that threshold. It does not, however, violate the continuity assumption since the average tax rate changes smoothly and there is no reason to expect this to be discontinuously related to demand for PHL.

¹⁰Stavrunova and Yerokhin (2014) use \$250 bins because this was the level at which their data were available. We have exact income data so are able to construct bins at a more granular level.

¹¹For singles we use all taxfilers with income above \$50,000 in the regression. For families we use all taxfilers with combined couple income above \$140,000.

(AMSE) of the boundary estimator ($\hat{\beta}$ in Eq. 1) are popular approaches in the literature (Imbens and Kalyanaraman, 2012; Calonico et al., 2014). A recent alternative proposed by Kettlewell and Siminski (2022) is to use observations away from the policy discontinuity (the placebo zone) as a training ground to inform bandwidth at the true discontinuity. They show that in many examples, this approach results in lower root MSE than the AMSE minimizing algorithms.

We use local linear RDD using the Calonico et al. (2014) and Kettlewell and Siminski (2022) approaches to bandwidth selection as baselines. We then present estimates using a range of alternative bandwidths, linear and quadratic control functions, and with and without control variables as sensitivity analyses. Sensitivity analysis is important in our context because along the range of our running variable, there are many points where there is heaping and theoretical discontinuities (e.g., at the other MLS thresholds and at tax brackets). Both the AMSE minimizing approach (Calonico et al., 2014) and the placebo zone approach (Kettlewell and Siminski, 2022) use data along the range of the running variable to inform bandwidth selection so are potentially affected by this. The AMSE approach may also result in bandwidths that exceed the distance between MLS tiers unless the observations outside those bounds are dropped, but then the selected bandwidth is necessarily interior (which may not be optimal) owing to the regularization term in the bandwidth selection algorithm. Constraints on the bandwidth are easy to impose for the placebo zone approach. Consequently, for the AMSE approach we allow the bandwidth to extend beyond other MLS tiers but for the placebo zone approach, which takes the minimum and maximum possible bandwidth as inputs, we constrain it so that the bandwidths never overlap with other MLS tiers. We use a uniform kernel since this is simple to interpret and provides a direct correspondence to the descriptive results in Figure 2 (Lee and Lemieux, 2010). Further details are in the table notes.

4 Results

4.1 Main results

Extensive margin. Our estimates for the probability of having PHI are summarized in Table 2. For both families and singles, we estimate a statistically significant jump in the probability of insurance at the tier 1 threshold. For families, the jump is estimated to be 1.1-1.4 ppts, depending on the approach used to select the bandwidth, and for singles it is estimated to be 3-3.5 ppts. We generally estimate nil effects at the other thresholds for both families and singles, and our point estimates are close to zero. The one exception is a drop in take-up at the singles tier 2 threshold, but this is only for the AMSE approach and we will show later it is not robust to alternative bandwidths.¹²

Intensive margin. Our estimates for premium spending are summarized in Table 3. In theory, the estimates could go in either direction. On one hand, the jump in take-up means that some people increase premium spending from \$0 to >\$0. On the other hand, the lower rebate means people may decide to purchase less expensive insurance. Given that we do not see evidence for a jump in take-up at tiers 2 and 3, we only expect the latter effect (reduced premium spending) at those thresholds.

Our estimates are qualitatively similar to what we observe for the extensive margin – there is a jump in premium expenditure at the tier 1 thresholds but only weak evidence for an effect at the tier 2 thresholds, with point estimates that go in an unexpected direction for families, and that we will see later are not robust to functional form decisions. We estimate a jump in premium spending of \$44-\$55 at the tier 1 family threshold and \$124-\$125 at the singles threshold.

Sensitivity. In Figures 3 and 4 we consider how sensitive our estimates are to differences

¹²This drop seems to be related to the fact the AMSE bandwidth just overlaps the tier 1 threshold, where there is a clear jump in take-up, steepening the local linear trend estimate on the left of the threshold. If we use the AMSE approach with a triangular kernel (which places more weight on observations near the threshold, potentially offsetting this bias), the estimate is closer to zero and insignificant ($\hat{\beta} = -0.01$, $SE = 0.009$).

Table 2: RDD estimates using data-driven bandwidths: extensive margin

	Tier 1		Tier 2		Tier 3	
	Families	Singles	Families	Singles	Families	Singles
(A) AMSE approach						
$\hat{\beta}$	0.0136***	0.0371***	-0.00238	-0.0192**	0.00349	0.0179
SE	(0.00387)	(0.00881)	(0.00341)	(0.00929)	(0.00416)	(0.0114)
CI LL	0.00656	0.0179	-0.00885	-0.0432	-0.00422	-0.00478
CI UL	0.0236	0.0591	0.00653	-0.00153	0.0141	0.0470
BW	2.452	1.725	3.053	1.696	2.300	2.144
Obs.	114661	47965	92665	30333	23690	12906
(B) Placebo zone approach						
	0.0114***	0.0342***	-0.000620	-0.00393	0.00139	0.0149
	(0.00356)	(0.00990)	(0.00350)	(0.0120)	(0.00374)	(0.0112)
BW	2.900	1.400	2.900	1.000	2.900	2.200
Obs.	135749	37818	87533	17284	30237	13334

Notes: Dependent variable is an indicator for if the person has PHI. Heteroskedasticity robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Bandwidths are expressed in \$0000s. (A) Local linear RDD estimates obtained using the *rdrobust* package for Stata (Calonico et al., 2017) with uniform kernel, with the robust confidence intervals reported in rows CI LL and CI UL (95% level). (B) local linear RDD estimates obtained using the *pzms* package for Stata with uniform kernel. 50 Placebo thresholds in the income range [\$50,000,\$350,000] (families) and [\$30,000,\$200,000] (singles) are used to select the preferred bandwidth (the one with the lowest RMSE across those thresholds, where the true treatment effect is assumed zero) and bandwidths between \$5,000-\$30,000 are considered in steps of \$1,000. For the singles tier 1 (2) threshold, the maximum bandwidth on the right (left) is set to \$15,000.

in bandwidth, polynomial order (linear versus quadratic) and the inclusion of controls. The controls we include are: total tax deductions; personal services income; government payments income; salary/wage income; self-employment flag; superannuation (compulsory retirement savings fund) balance; and whether they live in a rural area. The income, wealth and employment variables were selected since these might vary near the thresholds if there is sorting on the running variable (not addressed by our donut algorithm), while age, sex and region are known to predict PHI take-up.

For families, the jump in take-up and premium spending at the tier 1 threshold (Panel A) is robust to all specification choices, and is larger when smaller bandwidths are used. There is no robust evidence for a jump in take-up or spending at any other thresholds (Panels B

Table 3: RDD estimates using data-driven bandwidths: intensive margin

	Tier 1		Tier 2		Tier 3	
	Families	Singles	Families	Singles	Families	Singles
(A) AMSE approach						
$\hat{\beta}$	44.31***	125.2***	29.19**	-68.59**	6.601	7.723
SE	(11.45)	(25.20)	(14.04)	(27.09)	(21.82)	(46.44)
CI LL	15.39	62.51	2.698	-136.9	-38.17	-108.0
CI UL	66.06	170.6	65.81	-21.59	61.47	86.56
BW	3.429	1.550	2.952	1.871	2.926	1.705
Obs.	161163	42428	89211	31659	30550	9953
Placebo zone approach						
	54.76***	123.9***	28.33**	-22.72	6.545	94.89
	(12.65)	(25.66)	(14.18)	(29.51)	(21.93)	(83.75)
BW	2.800	1.500	2.900	1.500	2.900	0.500
Obs.	131060	40896	87533	26020	30237	2885

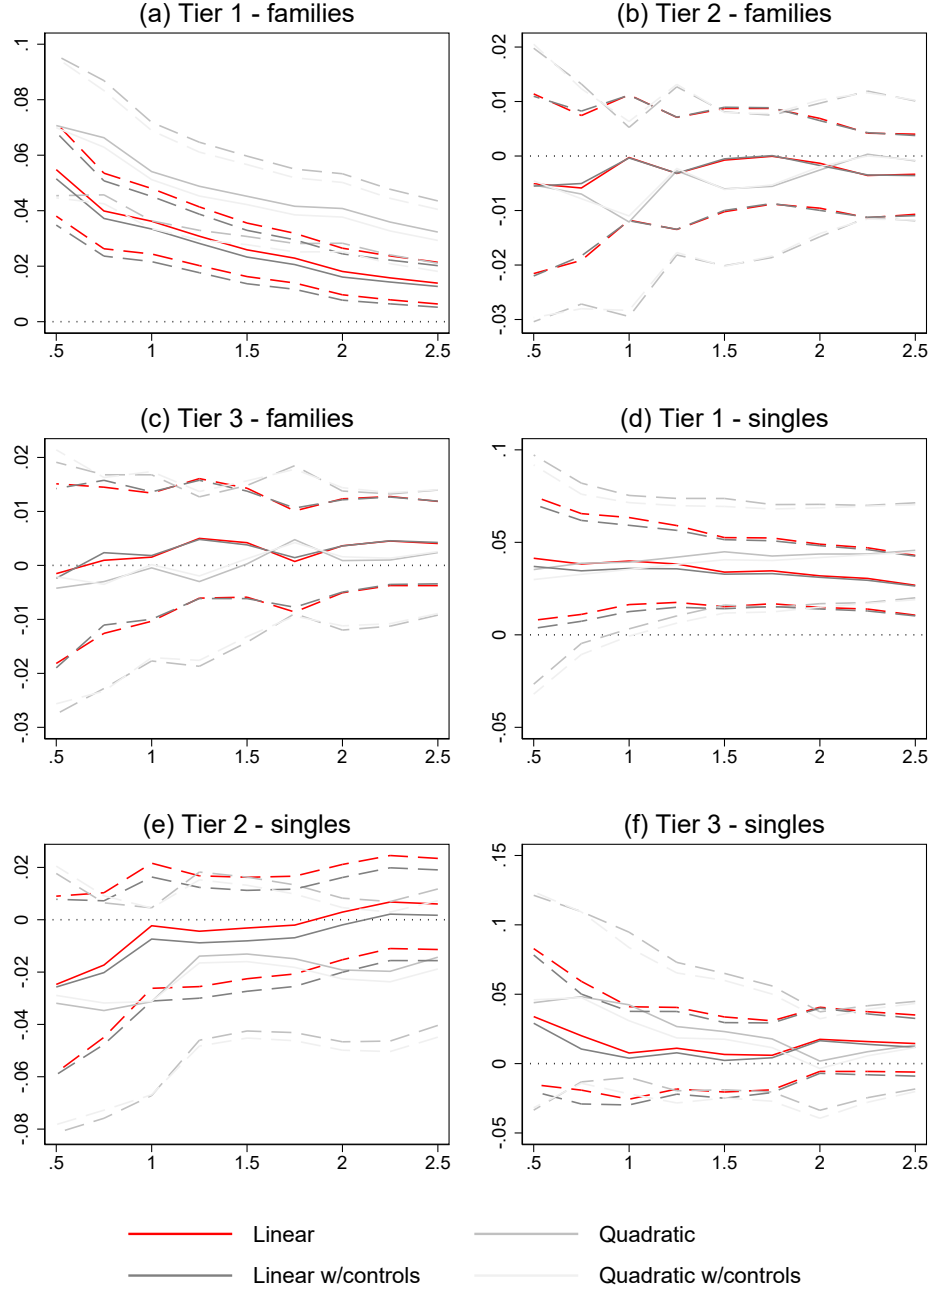
Notes: Dependent variable is annualized premium expenditure. Heteroskedasticity robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Bandwidths are expressed in \$0000s. (A) Local linear RDD estimates obtained using the *rdrobust* package for Stata (Calonico et al., 2017) with uniform kernel, with the robust confidence intervals reported in rows CI LL and CI UL (95% level). (B) local linear RDD estimates obtained using the *pzms* package for Stata with uniform kernel. 50 Placebo thresholds in the income range [\$50,000,\$350,000] (families) and [\$30,000,\$200,000] (singles) are used to select the preferred bandwidth (the one with the lowest RMSE across those thresholds, where the true treatment effect is assumed zero) and bandwidths between \$5,000-\$30,000 are considered in steps of \$1,000. For the singles tier 1 (2) threshold, the maximum bandwidth on the right (left) is set to \$15,000.

and C).

For singles, the jump in take-up at tier 1 for singles (Figure 3 Panel D) is quite robust to specification and bandwidth. The estimates for premium spending (Figure 4 Panel D) are somewhat sensitive however; the estimated jump is smaller when controls are included and only significantly different from zero when the bandwidth exceeds \$20,000 for the quadratic control function with controls specification.

Effect on insurance policy type. An interesting feature of the preceding results is that at the tier 1 thresholds, the relative jump in premium spending exceeds that for take-up, which can be seen in Appendix Figure A3 where we recast our estimates as the predicted percentage change in Y_i (i.e., $\frac{\hat{\beta}}{\alpha_1}$). One prominent concern with the MLS is that

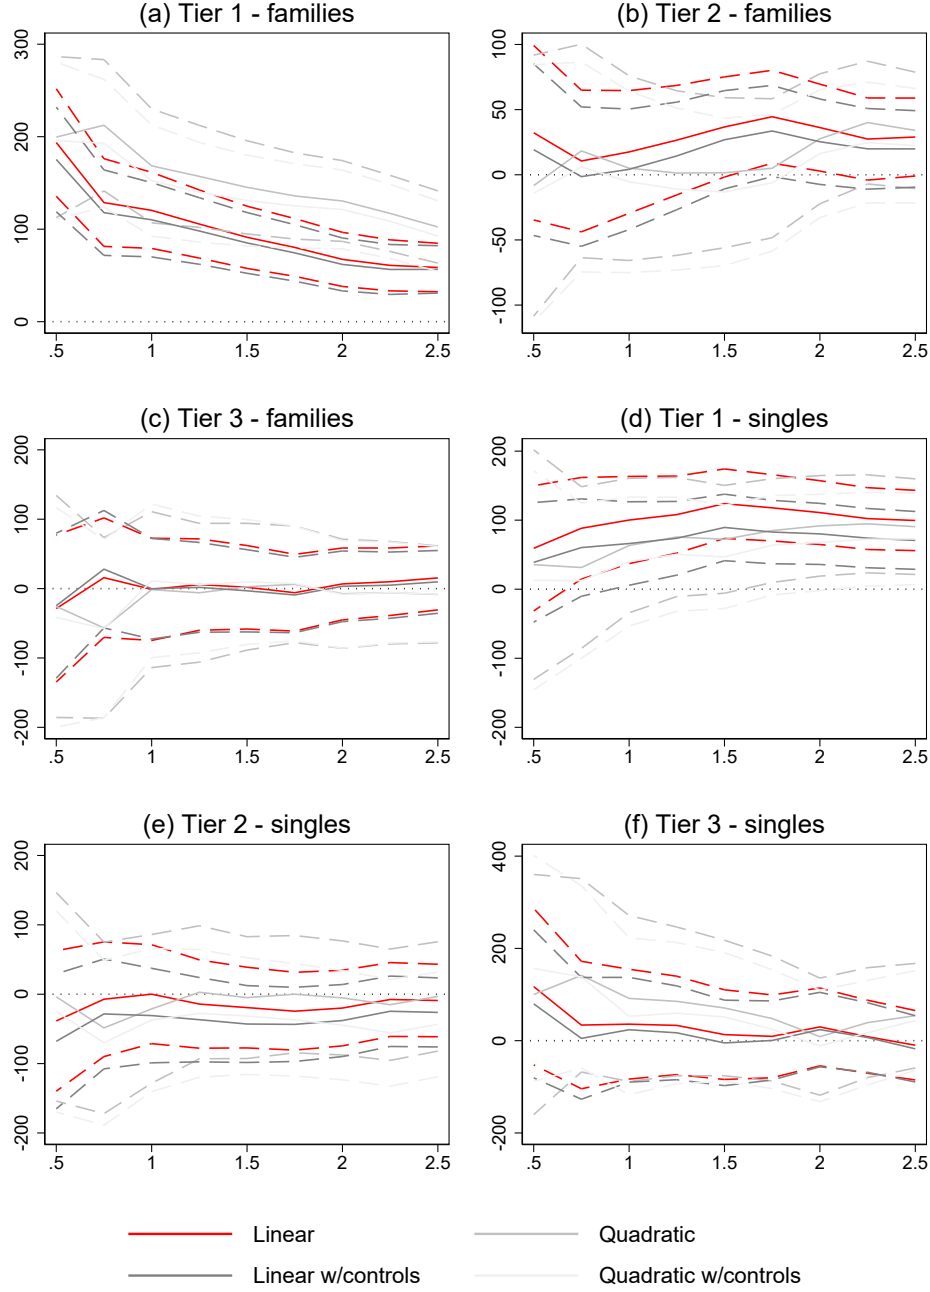
Figure 3: RDD estimates: Dependent variable = Has PHI



Notes: Each figure shows the estimated treatment effect and 95% confidence interval (based on the heteroskedasticity robust standard error) for different bandwidths between \$5,000 and \$25,000 (estimates are in \$2,500 increments). For Panel D (E) the maximum bandwidth on the right (left) is set to \$15,000. Linear/quadratic refers to the polynomial order of the control function.

it encourages people to purchase low-quality ('junk') PHI simply to avoid the penalty. If so, we might expect people purchasing PHI because of the MLS (i.e., just above the threshold)

Figure 4: RDD estimates: Dependent variable = Premiums

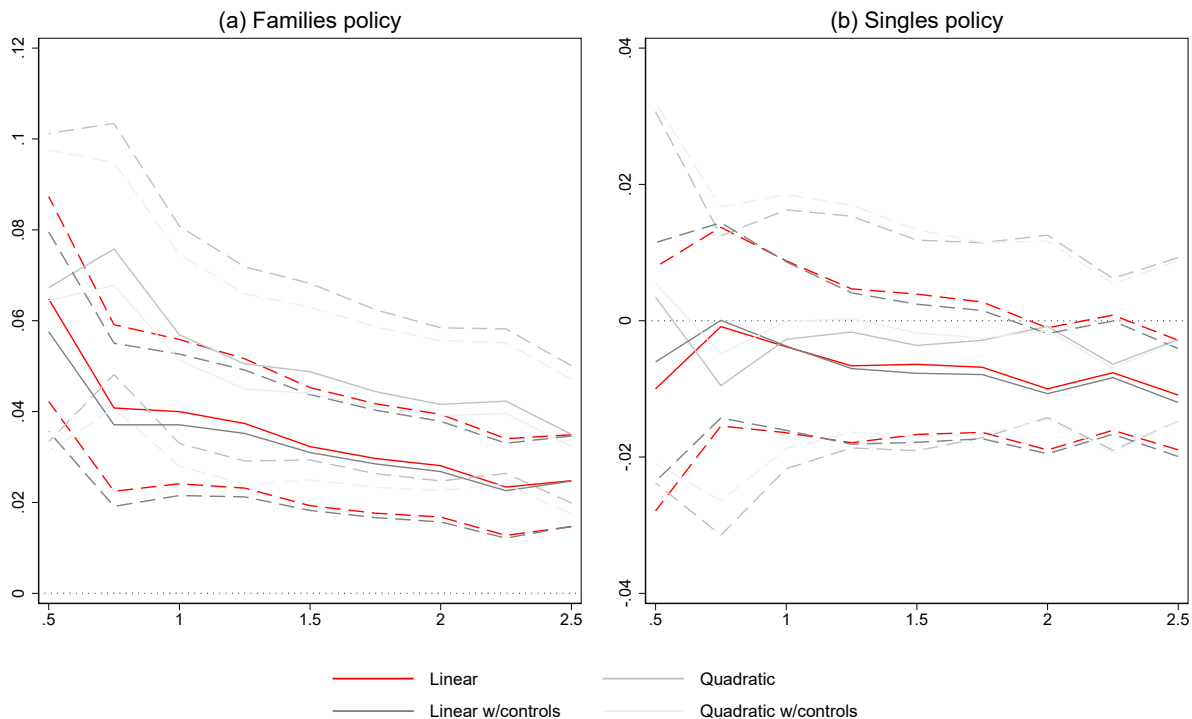


Notes: Each figure shows the estimated treatment effect and 95% confidence interval (based on the heteroskedasticity robust standard error) for different bandwidths between \$5,000 and \$25,000 (estimates are in \$2,500 increments). For Panel D (E) the maximum bandwidth on the right (left) is set to \$15,000. Linear/quadratic refers to the polynomial order of the control function.

to purchase cheaper insurance such that the jump in premiums is lower in percentage terms than the jump in take-up; the opposite of what we find.

At the family threshold, one potential explanation for the relatively larger jump in premiums spending is that PHI coverage must extend to one's spouse and dependent children to avoid the MLS. Consequently, there could be switching from singles to family policies at the tier 1 threshold. We test for this by estimating RDD models with dependent variables being (i) having a family PHI policy and (ii) having a singles policy. Our estimates show that there is indeed switching at the threshold; take-up of family policies increases by around 2-4 ppt (around 3-6%) while take-up of singles policies is small but negative (up to around 1 ppt) and marginally significant (for the linear RDD) at certain bandwidths. This pattern of results suggests marginal joiners are purchasing family plans, which are more expensive than single plans. This pattern is consistent with the relatively larger jump in premiums spending than take-up.

Figure 5: RDD estimates: Dependent variable = Has PHI type (families tier 1 threshold)



Notes: Each figure shows the estimated treatment effect and 95% confidence interval (based on the heteroskedasticity robust standard error) for different bandwidths between \$5,000 and \$25,000 (estimates are in \$2,500 increments). Linear/quadratic refers to the polynomial order of the control function.

Appendix Figure A3 also shows that the relative jump in premiums spending is larger than the jump in take-up for singles as well when we restrict attention to the RDD estimates with linear control function without controls. However, this is not robust to using a quadratic control function (Appendix Figure A3, Panel D) and later we will show evidence that a quadratic fit may be more appropriate for the singles tier 1 threshold.¹³

Kink effects. So far, our results suggest some effects on take-up and premiums at the initial threshold, but no effect at subsequent thresholds. However, we have only focused on the ‘jump’ at these thresholds; the MLS penalty is increasing in income such that we might also expect a ‘kink’ response – specifically, we might expect the slope to be steeper on the right of the cut-off on the extensive margin, as the penalty increases linearly with income. [Dong and Lewbel \(2015\)](#) also suggest examining whether the treatment effect derivative (TED) – which in our specifications is simply the coefficient on the interaction term between income and treatment – as a test for treatment effect stability. A small TED suggests that changes in the MLS thresholds would not greatly affect the magnitude of the treatment effect from the ‘jump’ in incentives at each tier.

We plot estimates for take-up of PHI in Appendix Figure A5.¹⁴ For both families and singles, there is no robust evidence that the TED changes at any threshold. For some bandwidths, the TED is statistically significant for the linear control function with larger bandwidths, but the point estimates are negative, implying, if anything, a positive price elasticity. Moreover, none of these estimates are robust to using a quadratic control function, which is arguably more reasonable for this analysis since the rate of take-up is capped at

¹³We also estimated whether there was switching from single to family policies at the singles tier 1 threshold (see Appendix Figure A4). Intuitively, there should be no switching as this sample is restricted to single people, who we do not expect to purchase family policies. In reality, we do see some single people with family policies (around 10% at the threshold), which could come about in several ways, such as joining a parent’s plan, staying on a joint plan after a relationship ends, or data entry mistakes. While there is some evidence for switching out of family plans, this is sensitive to the inclusion of controls and using local quadratic regression.

¹⁴We do not report estimates for bandwidths less than \$10,000 as the confidence intervals were extremely large. Since the rebate does not change linearly with income, there is no reason to expect a kink response. Nevertheless, we plot the estimates for premiums in Appendix Figure A6 for completeness. The results mirror those for take-up.

100% leading to a naturally concave relationship with income over a large enough range.

4.2 Robustness

Misreported marital status. In the earlier study by [Stavrunova and Yerokhin \(2014\)](#), the authors observed a discontinuous increase in the probability of being coupled for income just above the singles tier 1 threshold. This is explainable by the fact that, for people with income below that threshold, the reporting of relationship status is largely inconsequential (income tax is levied at the individual level in Australia). The authors dealt with this by estimating their model on all taxpayers and scaling the treatment effect by the increase in probability of being treated (similar to a fuzzy RDD). We also tested for misreporting in our data by replacing single status as the dependent variable and re-estimating our RDD specifications, but found no robust evidence of it (see Figure A7). This is likely due to the increasing integration of administrative data across both government and private systems, such as with the welfare system, which has improved reporting by ensuring that many personal details are pre-filled when completing a tax return.

Falsification tests. A common robustness check in RDD is to replace the dependent variable with covariates and test whether they change discontinuously at the policy threshold. We do this for each of the controls in our analysis and plot the estimates in Appendix Figures A8-A13. For families, there is little evidence for covariate imbalance. For some bandwidths and thresholds there are small statistically significant effects for the linear RDD specification (e.g., 0.4 higher age for a \$20,000 bandwidth at tier 2). However, these results are unsurprising given the number of hypotheses being tested, and are never robust to alternate bandwidths or using a quadratic control function. Moreover, we have already seen in Figures 3 and 4 that controlling for those differences makes little difference to our estimates.

For singles, the results using a linear control function are less robust, specifically at the tiers 1 and 2 thresholds. At the tier 2 threshold, there is evidence for a decrease in

salary/wage income and in the probability of being male. We have seen that controlling for those differences does not affect our point estimates (which suggest no treatment effect at that threshold). At the tier 1 threshold there are differences in salary/wage income, sex and age, and we saw earlier that controlling for those factors does impact the magnitude of the estimated treatment effects (although not by enough to reduce the treatment effects to zero for sufficiently large bandwidths). The estimates using a quadratic control function are more favorable; none of the covariates change discontinuously at the threshold. Importantly as well, Figures 3 and 4 demonstrate that our results for PHI coverage are robust to using a quadratic fit, particularly on the extensive margin.

Altogether, this exercise provides evidence supporting the robustness of our main findings.

4.3 Heterogeneity

We explore whether the policy effects differ by age and geographic location by estimating our regression models separately for different subgroups. For age, we split the sample into four groups: under 31, 31-49 years, 50-64 years, and 65 and older. Those aged under 31 are not affected by lifetime health cover loading, so might be more influenced by the MLS if the two policies crowd each other out. This is also the group with the lowest expected claims, so financial incentives may play a greater role in their demand than health. People aged 31-49 constitute an age group where the expected benefit of insurance (expected claims minus premiums) is negative for most people ([Donato and Onur, 2018](#)). For people aged 50-64 insurance is an increasingly attractive option. For people 65 years and older, insurance has positive expected value on average, and the premium rebate is higher. Our expectation is that the MLS incentive will be stronger for the younger age groups as health will explain less variation in demand for younger people relative to older people. We also split the sample by where they live – rural or urban – since the value of PHI is likely to differ along this dimension. Specifically, there is less access to private hospitals in rural areas, so we might expect the MLS to explain a significant fraction of demand, and therefore have a larger

impact for this subgroup.

As expected, we find the effect of the MLS is decreasing in age.¹⁵ For both families and singles, the effect is strongest both in absolute and relative terms for those aged under 31 years. It is generally significant but with a treatment effect approximately half as large for people aged 31-49 years. The estimates are generally insignificant for the older age groups.

This result implies that marginal joiners due to the MLS are likely to be people who put downward pressure on premiums. This in turn could generate additional demand for insurance, although most studies (and our own evidence) suggest the price elasticity of demand for PHI is low – with existing estimates in the range of 0.35-0.5 ([Duckett et al., 2019](#)).

Lastly, we obtain very similar estimates for people in urban and rural areas, although the point estimates are imprecise for rural due to limited observations.

5 Discussion

Penalty and subsidies are frequently used in health insurance markets to encourage people to insure. The impact of these policies rests on how responsive people are to prices and how the response varies by income. We have studied how people respond to a strong penalty (the MLS) and moderate withdrawal of subsidy in the context of a dual public/private healthcare system. Our results suggest that the high income earners subject to these policies are largely unresponsive to prices. There is a small increase in take-up and premium expenditure at the thresholds where the MLS first kicks in, but no response at other thresholds, where the penalty becomes larger and the subsidy becomes smaller.

To better appreciate the magnitude of our estimates, we note that approximately 90% of people in families earning \$AU180,000 (MLS tier 1 threshold) and around 70% of singles earning \$AU90,000 already have PHI. The MLS therefore explains only a small fraction of

¹⁵This result was also found in studies that evaluated the earlier version of the MLS policy ([Buchmueller et al., 2021](#); [Stavrunova and Yerokhin, 2014](#)).

demand for those liable for it. Because RDD estimates local effects only, it is difficult to extrapolate our estimates, but a crude approach is to assume a common treatment effect of 3 ppts (our estimate for singles) for everyone with income above the MLS tier 1 threshold and calculate the percentage change in take-up with and without that effect. In our tax data, this is equal to $0.84\% = 0.1808 * 0.03 / 0.6462$ where 0.1808 is the decimal fraction of people liable for the MLS and 0.6462 is the fraction of tax-filers with PHI. This exercise makes clear that given the small treatment effects we estimate and the relatively small fraction of people liable for the MLS, its overall impact is also likely small. One caveat to this is that, if income fluctuates above and below the MLS threshold, some people may choose to retain private cover due to state dependence ([Buchmueller et al., 2021](#)), which would imply larger effects. However, since the MLS induces only a small response from those who become liable, this second order effect is likely to be small as well. Indeed, [Buchmueller et al. \(2021\)](#) simulate that had the MLS been abolished completely in 2008, PHI take-up would have only decreased trivially from 61.3% to 60.2% in 2012, taking into account the effect of state dependence.

The financing of healthcare systems is becoming an increasingly pressing concern for governments around the world dealing with ageing populations, increasing health costs and tighter fiscal constraints. Australia’s experience with encouraging PHI to reduce pressure on the public health care system is an important opportunity to learn. Our results suggest that penalties targeting high income earners – a population group with already high PHI coverage – does increase take-up of PHI but only by a modest amount. The benefits of this need to be considered against its costs, such as money spent by consumers on insurance they do not necessarily want or need, and costs associated with tax avoidant behavior.

Our results also speak to the (lack of) efficacy of premium subsidies for high-income earners. The fact that we do not find evidence of a reduction in premium spending at any of the policy thresholds suggests that targeting benefits towards lower income earners is an efficient way to reduce government outlays.

References

- Abhayaratna, T., A. Carter, and S. Johnson (2022). The ATO Longitudinal Information Files (ALife): Individuals—A new dataset for public policy research. *Australian Economic Review* 55(4), 541–557.
- Australian Institute of Health and Welfare (2022). Health expenditure Australia 2020–21. Retrieved from <https://www.aihw.gov.au/reports/health-welfare-expenditure/health-expenditure-australia>.
- Australian Prudential Regulation Authority (2022). Quarterly private health insurance statistics.
- Barreca, A. I., J. M. Lindo, and G. R. Waddell (2016). Heaping-induced bias in regression-discontinuity designs. *Economic Inquiry* 54(1), 268–293.
- Bilgrami, A., H. Cutler, K. Sinha, and Z. Cheng (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.
- Buchmueller, T. C., T. C. Cheng, N. T. Pham, and K. E. Staub (2021). The effect of income-based mandates on the demand for private hospital insurance and its dynamics. *Journal of Health Economics* 75, 102403.
- Calonico, S., M. D. Cattaneo, M. H. Farrell, and R. Titiunik (2017). Rdrobust: Software for regression-discontinuity designs. *The Stata Journal* 17(2), 372–404.
- Calonico, S., M. D. Cattaneo, and R. Titiunik (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica* 82, 2295–2326.
- 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.
- Colombo, F. and N. Tapay (2004). Private health insurance in OECD countries: the benefits and costs for individuals and health systems.
- Department of Health and Aged Care (2023). Average annual price changes in private health insurance premiums. <https://www.health.gov.au/resources/publications/average-annual-price-changes-in-private-health-insurance-premiums?language=und>. Accessed 31 May 2023.
- Donato, R. and I. Onur (2018). Private health insurance in Australia: policy reform approaches towards greater competition and efficiency to improve health system performance. Melbourne, Vic: Australian Centre for Financial Studies.
- Dong, Y. and A. Lewbel (2015). Identifying the Effect of Changing the Policy Threshold in Regression Discontinuity Models. *The Review of Economics and Statistics* 97(5), 1081–1092.

- Duckett, S., M. Cowgill, and K. Nemet (2019). Saving private health 2: Making private health insurance viable. Grattan Institute.
- Ellis, R. P. and E. Savage (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.
- 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.
- Finkelstein, A., N. Hendren, and M. Shepard (2019). Subsidizing health insurance for low-income adults: Evidence from Massachusetts. *American Economic Review* 109(4), 1530–67.
- Frean, M., J. Gruber, and B. D. Sommers (2017). Premium subsidies, the mandate, and Medicaid expansion: Coverage effects of the Affordable Care Act. *Journal of Health Economics* 53, 72–86.
- Frech, H., S. Hopkins, and G. MacDonald (2003). The Australian private health insurance boom: was it subsidies or liberalised regulation? *Economic Papers-Economic Society of Australia* 22(1), 58–64.
- Gong, X. and J. Gao (2018). Nonparametric kernel estimation of the impact of tax policy on the demand for private health insurance in Australia. *Australian & New Zealand Journal of Statistics* 60(3), 374–393.
- Gruber, J. and E. Washington (2005). Subsidies to employee health insurance premiums and the health insurance market. *Journal of Health Economics* 24(2), 253–276.
- Hackmann, M. B., J. T. Kolstad, and A. E. Kowalski (2015). Adverse selection and an individual mandate: When theory meets practice. *American Economic Review* 105(3), 1030–1066.
- Hahn, J., P. Todd, and W. Van der Klaauw (2001). Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica* 69(1), 201–209.
- Hamilton, S. (2018). Optimal deductibility: Theory, and evidence from a bunching decomposition. The Tax and Transfer Policy Institute Working paper.
- Hunt, G. (2021). Delivering Australia’s lowest private health insurance premium change in 21 years. Media release: Minister for Health and Aged Care, 23 December 2021.
- Imbens, G. and K. Kalyanaraman (2012). Optimal bandwidth choice for the regression discontinuity estimator. *Review of Economic Studies* 79, 933–959.
- Kang, Y., X. Gong, J. Gao, and P. Qiu (2015). Jump detection in generalized error-in-variables regression with an application to Australian health tax policies. *The Annals of Applied Statistics* 9(2), 883–900.


- Kang, Y., X. Gong, J. Gao, and P. Qiu (2019). Errors-in-variables jump regression using local clustering. *Statistics in Medicine* 38(19), 3642–3655.
- Kettlewell, N. and P. Siminski (2022). Optimal model selection in rdd and related settings using placebo zones.
- Kettlewell, N., O. Stavrunova, and O. Yerokhin (2018). Premium subsidies and demand for private health insurance: results from a regression discontinuity design. *Applied Economics Letters* 25(2), 96–101.
- King, D. and E. Mossialos (2005). The determinants of private medical insurance prevalence in England, 1997–2000. *Health Services Research* 40(1), 195–212.
- Lee, D. S. and T. Lemieux (2010). Regression discontinuity designs in economics. *Journal of Economic Literature* 48(2), 281–355.
- Lorenzoni, L., A. Marino, D. Morgan, and C. James (2019). Health spending projections to 2030. OECD Health Working Papers No. 110.
- Nicolás, Á. L. and M. Vera-Hernández (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 J. Yong (2005). Effects of recent carrot-and-stick policy initiatives on private health insurance coverage in Australia. *Economic Record* 81(254), 262–272.
- Rodríguez, M. and A. Stoyanova (2008). Changes in the demand for private medical insurance following a shift in tax incentives. *Health Economics* 17(2), 185–202.
- Smits, M., J. Henriquez, A. Melia, and F. Paolucci (2022). Stability of private health insurance in Australia: Analysis and proposal of a legislative framework to boost participation. *Australian Business Law Review* 50(4), 291–300.
- Stavrunova, O. and O. Yerokhin (2014). Tax incentives and the demand for private health insurance. *Journal of Health Economics* 34, 121–130.
- WHO (2020). Global spending on health 2020: weathering the storm. World Health Organization.

Online Appendix


A Additional tables and figures

Figure A1: Density around MLS thresholds: Families


(a) Tier 1 threshold




(b) After donut




(c) Tier 2 threshold




(d) After donut



(e) Tier 3 threshold

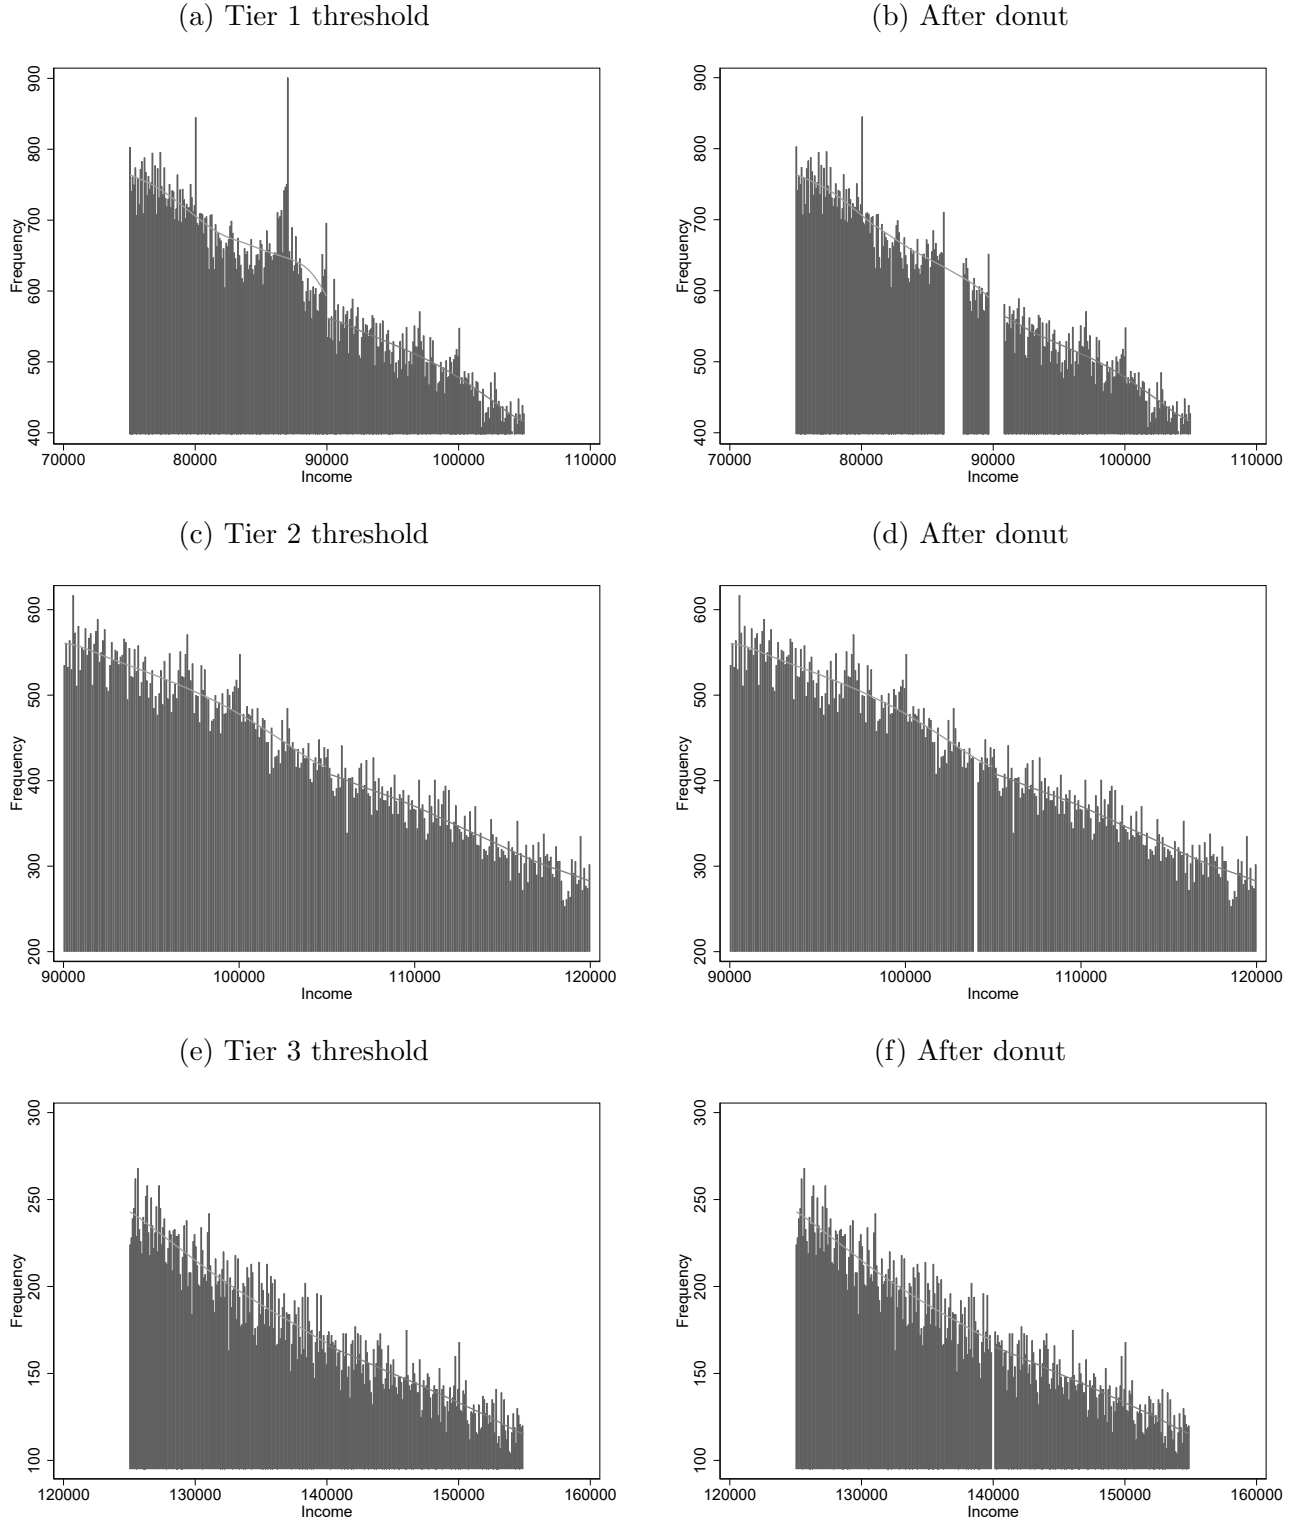


(f) After donut



Note: ALife 2018 release. Figures show counts of observations in \$100 income bins. Donuts are generated following the approach discussed in Section 3.


Figure A2: Density around MLS thresholds: Singles




Note: ALife 2018 release. Figures show counts of observations in \$100 income bins. Donuts are generated following the approach discussed in Section 3.

Figure A3: RDD relative treatment effect sizes


(a) Families – linear RDD




(b) Singles – linear RDD



(c) Families – quadratic RDD




(d) Singles – quadratic RDD



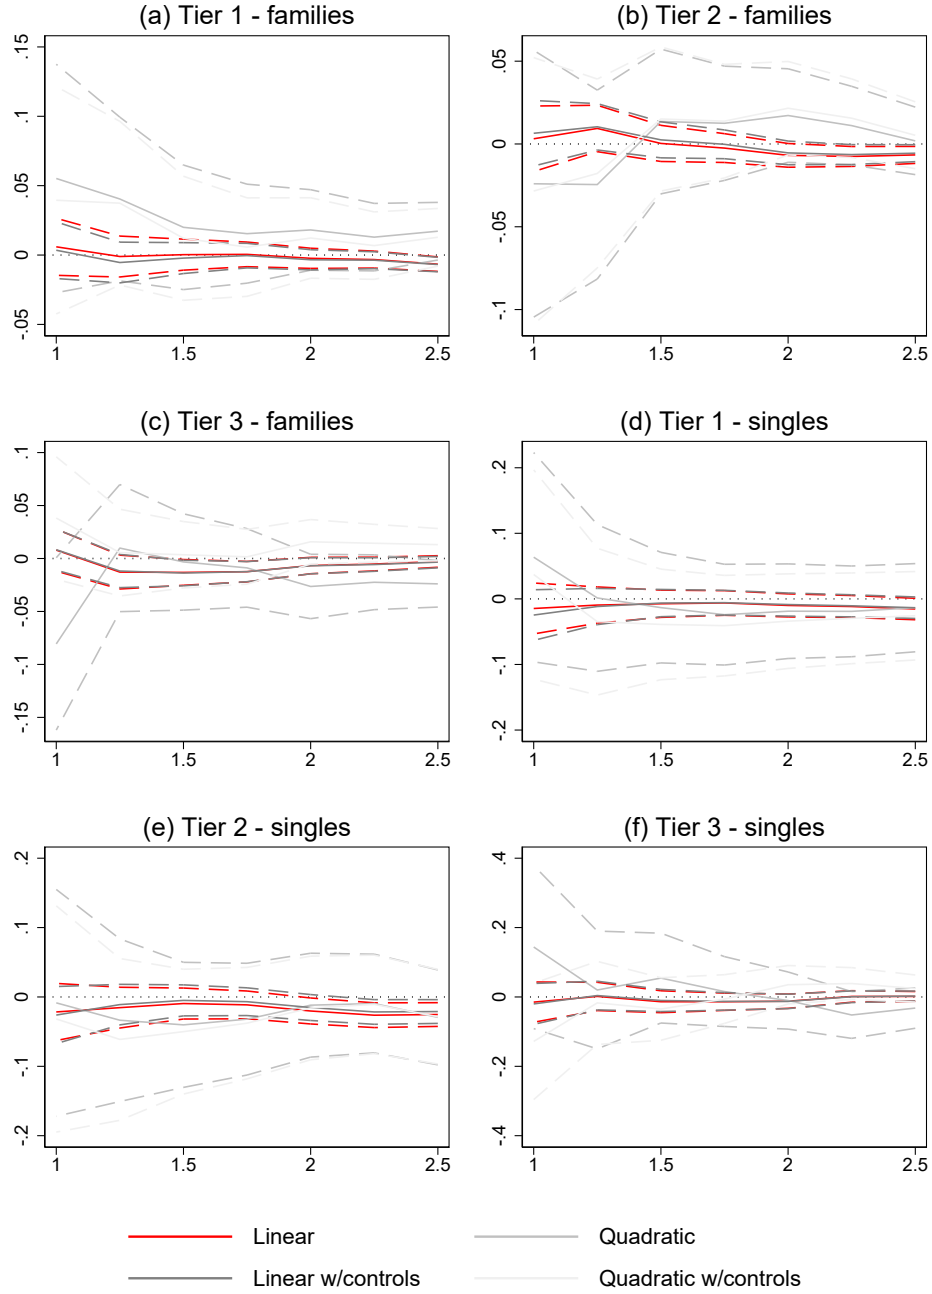
Notes: Each figure shows the estimated ratio between the treatment effect and the constant term (i.e., predicted value at the threshold) for the RDD estimates at the tier 1 thresholds, separately for take-up (Has PHI) and premiums. The estimates are reported for different bandwidths between \$5,000 and \$25,000 (estimates are in \$2,500 increments). In Panels B and D the maximum bandwidth on the right is set to \$15,000. Linear/quadratic refers to the polynomial order of the control function.

Figure A4: RDD estimates: Dependent variable = Has PHI type (singles tier 1 threshold)



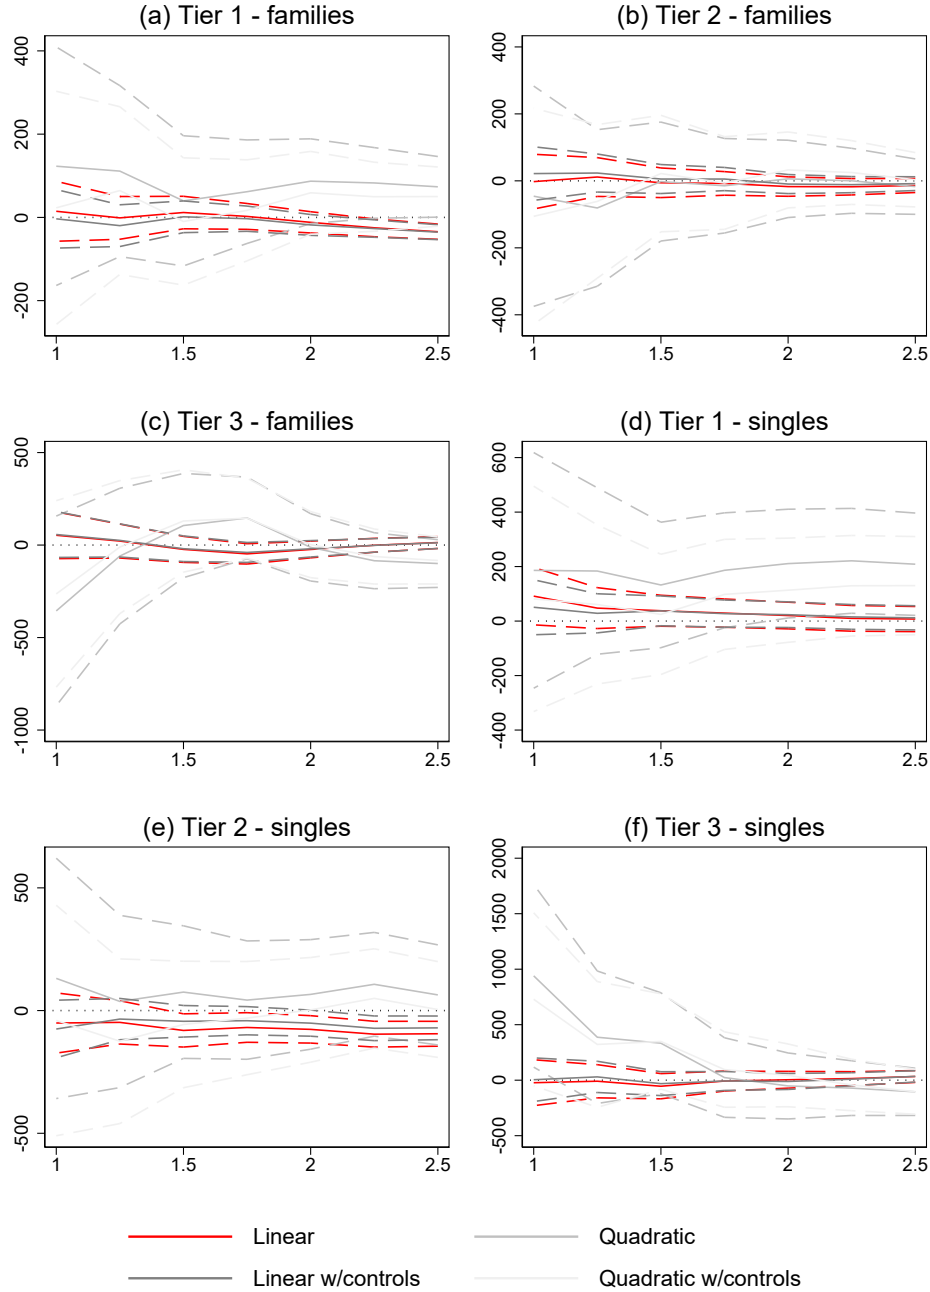
Notes: Each figure shows the estimated treatment effect and 95% confidence interval (based on the heteroskedasticity robust standard error) for different bandwidths between \$5,000 and \$25,000 (estimates are in \$2,500 increments). The maximum bandwidth on the right is set to \$15,000. Linear/quadratic refers to the polynomial order of the control function.

Figure A5: RDD kink estimates: Dependent variable = Has PHI




Notes: Each figure shows the estimated treatment effect derivative (coefficient on the linear interaction term between treatment and running variable (income) centered at zero where income equals the policy threshold) and 95% confidence interval (based on the heteroskedasticity robust standard error) for different bandwidths between \$5,000 and \$25,000 (estimates are in \$2,500 increments). For Panel D (E) the maximum bandwidth on the right (left) is set to \$15,000. Linear/quadratic refers to the polynomial order of the control function.

Figure A6: RDD kink estimates: Dependent variable = Premiums



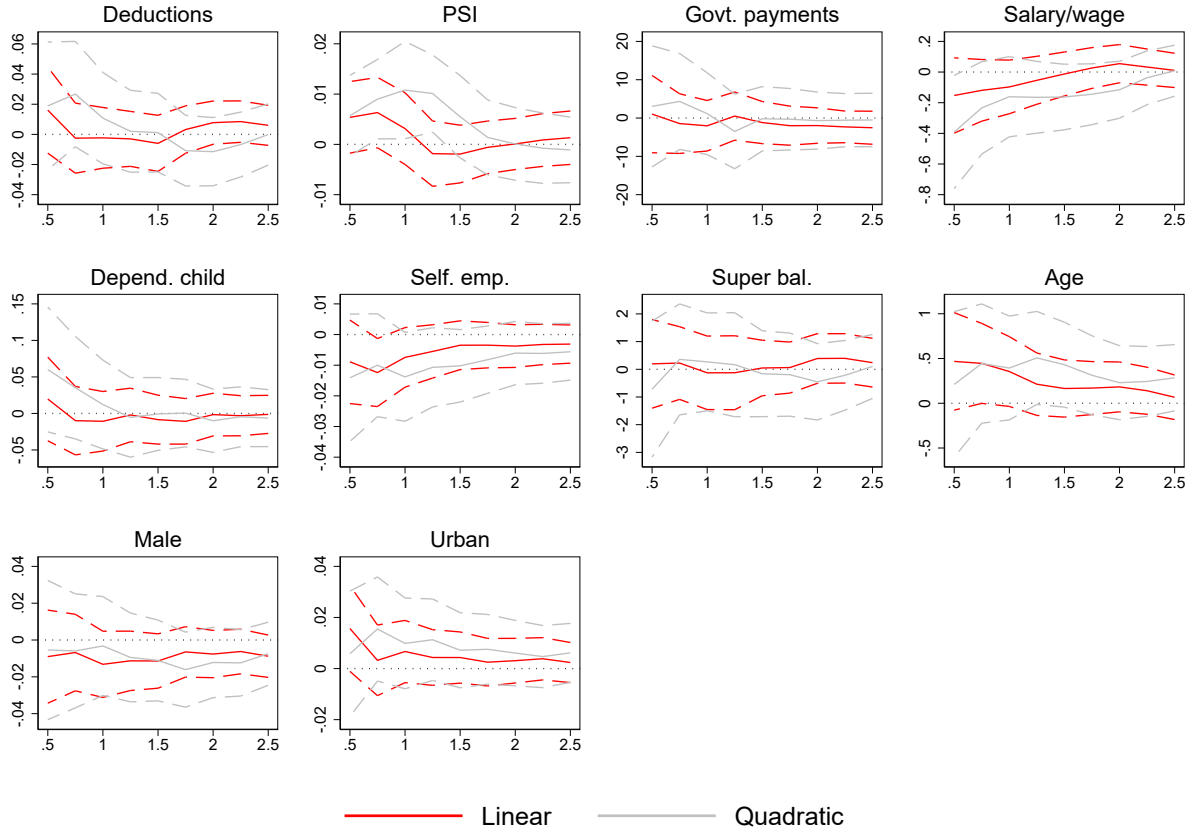
Notes: Each figure shows the estimated treatment effect derivative (coefficient on the linear interaction term between treatment and running variable (income) centered at zero where income equals the policy threshold) and 95% confidence interval (based on the heteroskedasticity robust standard error) for different bandwidths between \$5,000 and \$25,000 (estimates are in \$2,500 increments). For Panel D (E) the maximum bandwidth on the right (left) is set to \$15,000. Linear/quadratic refers to the polynomial order of the control function.

Figure A7: RDD estimates: Dependent variable = Single



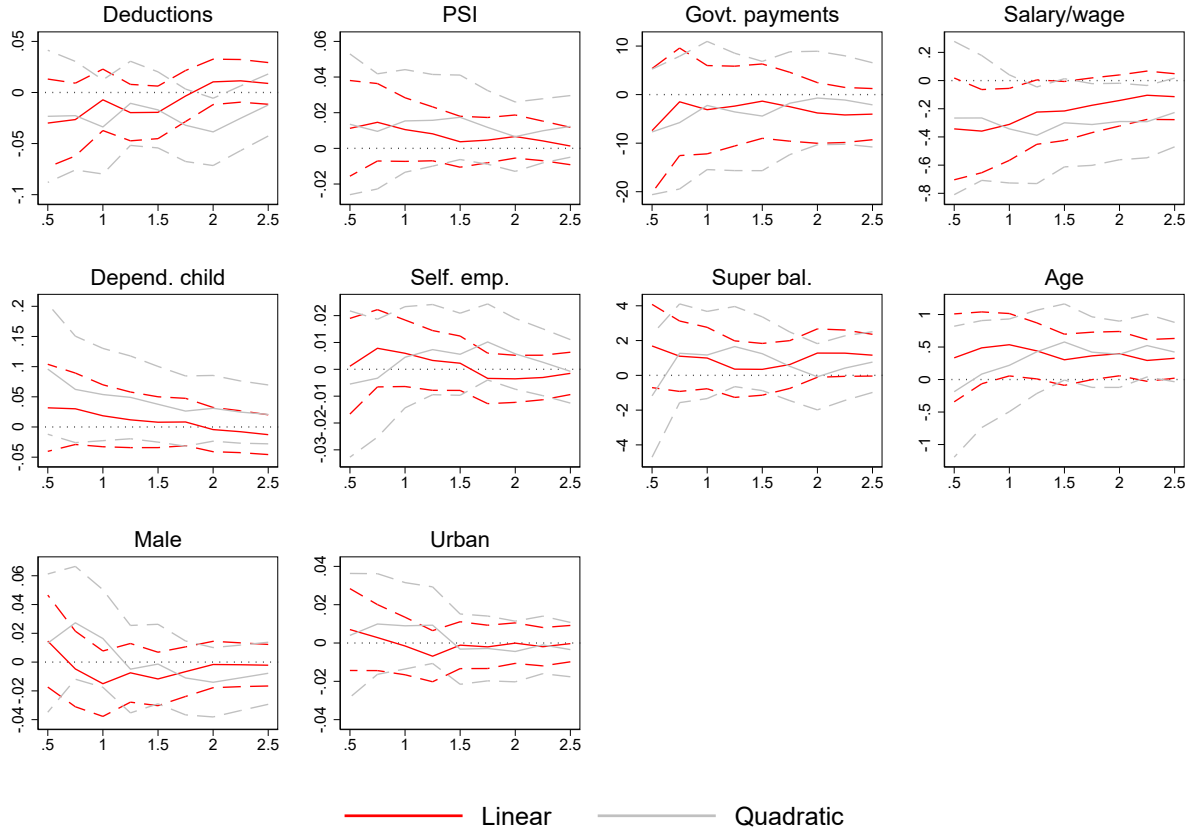
Notes: The figure shows the estimated treatment effect and 95% confidence interval (based on the heteroskedasticity robust standard error) for different bandwidths between \$5,000 and \$25,000 (estimates are in \$2,500 increments). The maximum bandwidth on the right is set to \$15,000. Linear/quadratic refers to the polynomial order of the control function.

Figure A8: RDD estimates: Covariates (families tier 1 threshold)



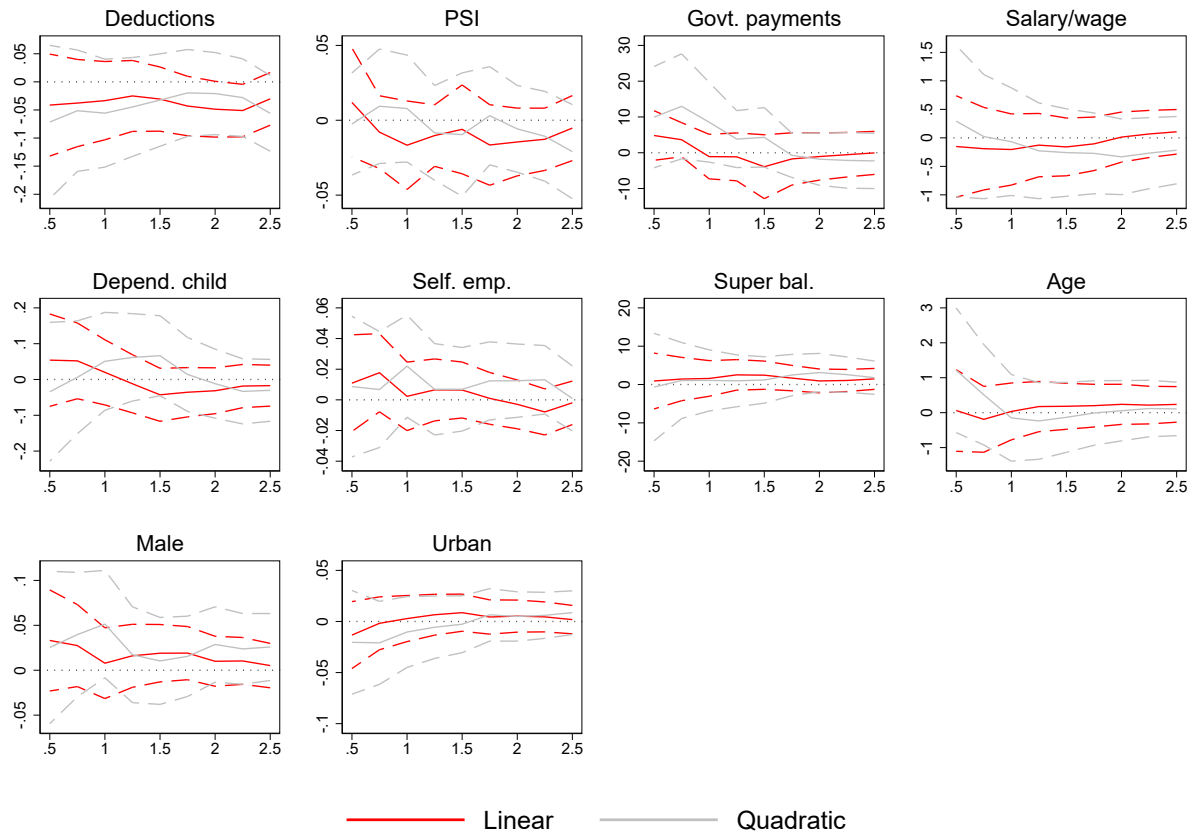
Notes: The figure shows the estimated treatment effect and 95% confidence interval (based on the heteroskedasticity robust standard error) for different bandwidths between \$5,000 and \$25,000 (estimates are in \$2,500 increments). Linear/quadratic refers to the polynomial order of the control function. Estimates for Deduction, Personal Services Income (PSI), Salary/wage income and Super balance are all expressed in \$0000s.

Figure A9: RDD estimates: Covariates (families tier 2 threshold)



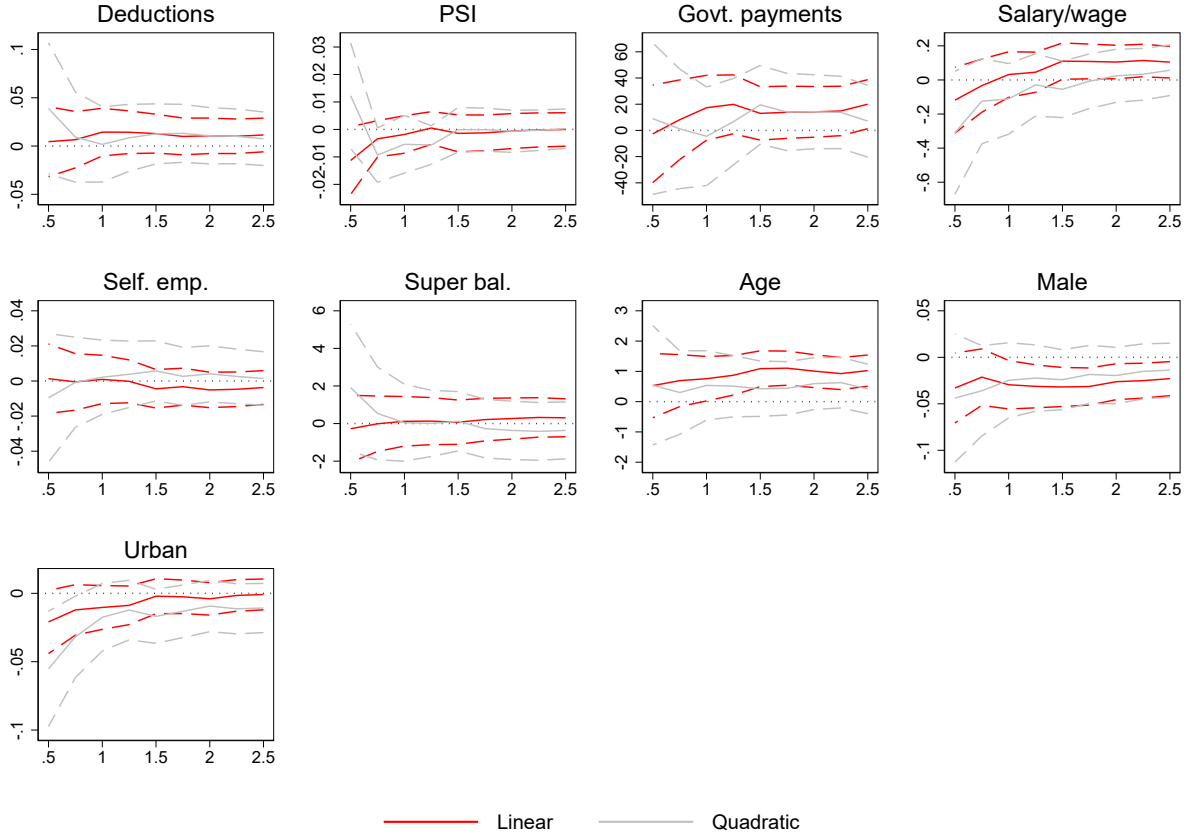
Notes: The figure shows the estimated treatment effect and 95% confidence interval (based on the heteroskedasticity robust standard error) for different bandwidths between \$5,000 and \$25,000 (estimates are in \$2,500 increments). Linear/quadratic refers to the polynomial order of the control function. Estimates for Deduction, Personal Services Income (PSI), Salary/wage income and Super balance are all expressed in \$0000s.

Figure A10: RDD estimates: Covariates (families tier 3 threshold)



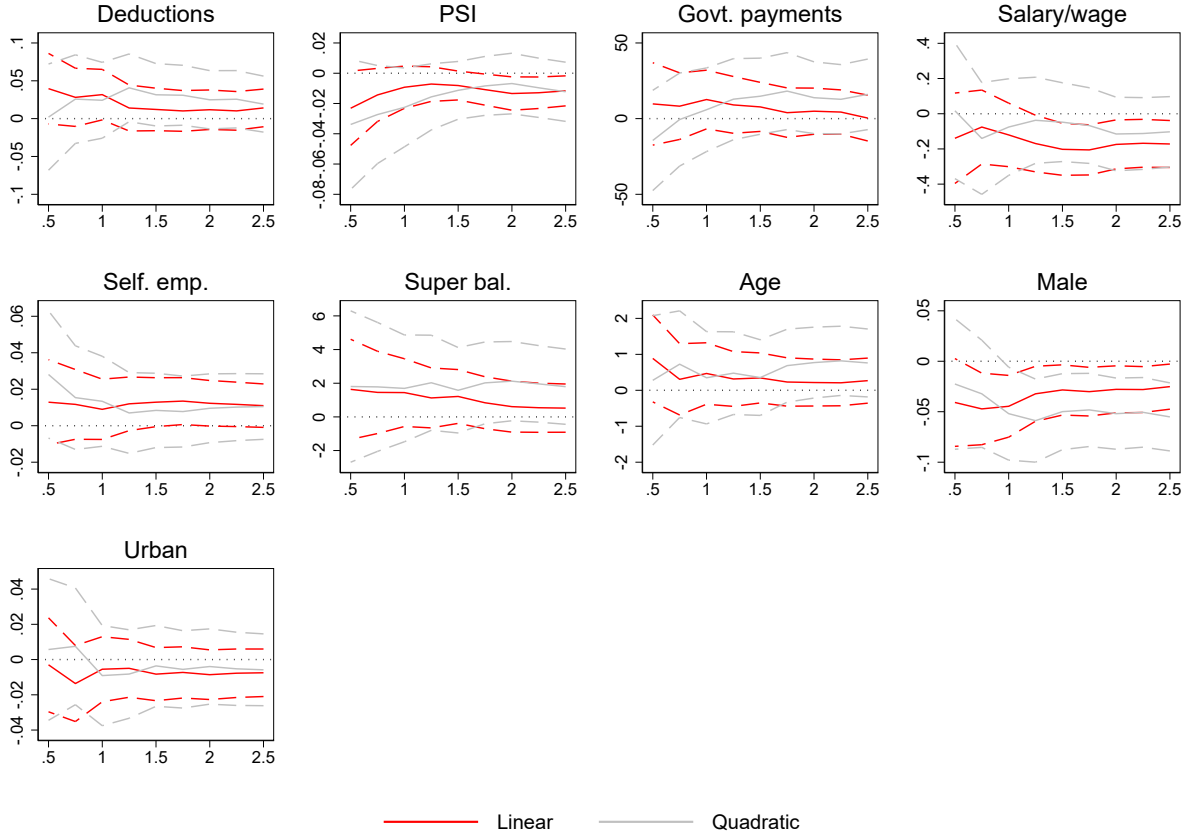
Notes: The figure shows the estimated treatment effect and 95% confidence interval (based on the heteroskedasticity robust standard error) for different bandwidths between \$5,000 and \$25,000 (estimates are in \$2,500 increments). Linear/quadratic refers to the polynomial order of the control function. Estimates for Deduction, Personal Services Income (PSI), Salary/wage income and Super balance are all expressed in \$0000s.

Figure A11: RDD estimates: Covariates (singles tier 1 threshold)



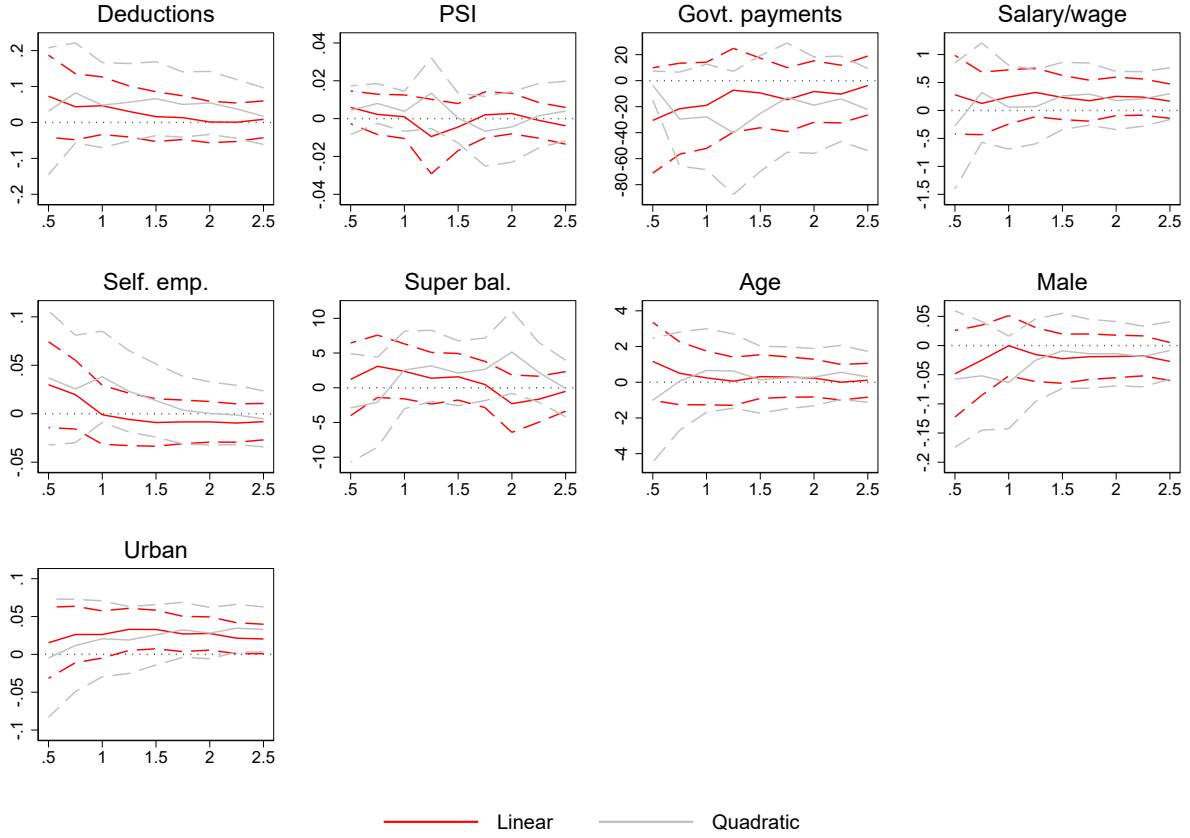
Notes: The figure shows the estimated treatment effect and 95% confidence interval (based on the heteroskedasticity robust standard error) for different bandwidths between \$5,000 and \$25,000 (estimates are in \$2,500 increments). The maximum bandwidth on the right is set to \$15,000. Linear/quadratic refers to the polynomial order of the control function. Estimates for Deduction, Personal Services Income (PSI), Salary/wage income and Super balance are all expressed in \$0000s.

Figure A12: RDD estimates: Covariates (singles tier 2 threshold)



Notes: The figure shows the estimated treatment effect and 95% confidence interval (based on the heteroskedasticity robust standard error) for different bandwidths between \$5,000 and \$25,000 (estimates are in \$2,500 increments). The maximum bandwidth on the left is set to \$15,000. Linear/quadratic refers to the polynomial order of the control function. Estimates for Deduction, Personal Services Income (PSI), Salary/wage income and Super balance are all expressed in \$0000s.

Figure A13: RDD estimates: Covariates (singles tier 3 threshold)



Notes: The figure shows the estimated treatment effect and 95% confidence interval (based on the heteroskedasticity robust standard error) for different bandwidths between \$5,000 and \$25,000 (estimates are in \$2,500 increments). Linear/quadratic refers to the polynomial order of the control function. Estimates for Deduction, Personal Services Income (PSI), Salary/wage income and Super balance are all expressed in \$0000s.