Insurance versus Moral Hazard in Income-Contingent Student Loan Repayment*

Tim de Silva[†]
Stanford University

March 29, 2025

Latest Version

Abstract

Student loans with income-contingent repayment insure borrowers against income risk but can reduce their incentives to earn more. Using a change in Australia's income-contingent repayment schedule, I show that borrowers reduce their labor supply to lower their repayments. These responses are larger among borrowers with more hourly flexibility, a lower probability of repayment, and tighter liquidity constraints. I use these responses to estimate a dynamic model of labor supply with frictions that generate imperfect adjustment. My estimates imply that the labor supply responses to income-contingent repayment limit the optimal amount of insurance in government-provided student loans. However, these responses are too small to justify fixed repayment contracts: restructuring outstanding student loans from fixed repayment to a constrained-optimal income-contingent loan—while keeping the tax and transfer system unchanged—increases borrower welfare by the equivalent of a 0.8% increase in lifetime consumption at no additional fiscal cost.

* This paper is a revised version of Chapter 1 of my Ph.D. dissertation at MIT. I'm grateful to my dissertation committee, Taha Choukhmane, Jonathan A. Parker, Lawrence D.W. Schmidt, Eric C. So, and David Thesmar for their continuous support. I thank David Autor, Anna Bedford, Nicholas Biddle, Michael Boutros (discussant), Bruce Chapman, Jan Eberly, Mehran Ebrahimian (discussant), Amy Finkelstein, Marina Gertsberg (discussant), Jonathan Gruber, Nathan Hendren, Greg Kaplan, Henrik Kleven, Debbie Lucas, Simone Melchionna, Simon Mongey, Andrew Norton, Jim Poterba, Andrei Shleifer, Stefanie Stantcheva, Toni Whited, Constantine Yannelis (discussant), four anonymous referees, and seminar participants at the FRA Conference, BlackRock Applied Research Award Panel, Olin Finance Conference, Inter-Finance Ph.D. Seminar, SFS Cavalcade, WFA, Policy Impacts Conference, SITE Market Failures, Brookings Conference on New Research and Tax Policy, NBER Public Economics, NBER Economics of Education, e61 Micro 4 Macro, AEA, MIT, Chicago Booth, Princeton, Yale SOM, Harvard, London School of Economics, London Business School, Wharton, UCLA Anderson, Stanford GSB, Northwestern Kellogg, Berkeley Haas, Columbia Business School, NYU Stern, Boston College Carroll, Swedish House of Finance, Bocconi, OSU Fisher, Claremont McKenna College, USC Marshall, and BYU for helpful comments. I also thank the ATO *ALife* and ABS DataLab teams for data assistance, in particular Andrew Carter, Justin Holland, and Son Nguyen. This research was supported by the Alfred P. Sloan Foundation through a Household Finance grant awarded to the National Bureau of Economic Research, Mark Kritzman and Elizabeth Gorman Research Fund, MIT Sloan Stone Fund, and Thomas Anthony Pappas Endowed Scholarship. The results of these studies are based, in part, on tax data supplied by the Australian Taxation Office to the Australian Bureau of Statistics under the Taxation Administration Act 1953. The final page of this paper contains a full required disclaimer.

† Stanford University, Graduate School of Business and Institute for Economic Policy Research, tdesilva@stanford.edu.

In many countries, students finance higher education through government-provided student loans. These loans are the second-largest household liability in the US at $1.6 trillion and account for 10% of household debt in the US and UK. Traditionally, government-provided student loans have resembled debt contracts, where borrowers make fixed payments after graduation to repay their loan balances. Because student loans are generally not dischargeable in bankruptcy, these contracts force borrowers to bear most of the risk associated with the returns to higher education. Unfortunately, the risk of low income after graduation materializes for many borrowers, with 25% of US borrowers defaulting within five years (Hanson 2022).

One potential policy to provide more insurance against income risk is to make student loans equity-like by linking repayments to borrowers' incomes. This idea has been discussed extensively (Friedman 1955; Shiller 2004; Palacios 2004; Chapman 2006), and governments in the US, UK, Canada, and Australia have recently implemented it by providing income-contingent loans. In contrast to nondischargeable debt contracts, income-contingent repayment provides insurance by reducing payments as a borrower's income declines. However, this insurance potentially comes at the cost of creating moral hazard: because repayments increase with income, borrowers have an incentive to reduce their labor supply to decrease repayments. Empirically, income-contingent repayment appears effective at providing insurance (Herbst 2023), but there is no consensus on the moral hazard effects that it creates (Yannelis and Tracey 2022).

The objective of this paper is to study two questions. First, how does income-contingent repayment affect borrowers' labor supply? Second, what is the optimal form of income-contingent repayment that balances this moral hazard, if it exists, with providing insurance? To identify labor supply responses empirically, I leverage administrative data and policy variation from the Australian Higher Education Loan Programme (HELP), the first program to provide income-contingent loans nationwide. I then use these responses to estimate a dynamic model of labor supply and study the implications of various repayment contracts. In my normative analyses, I consider a government that maximizes borrower welfare, taking education and borrowing choices as given. Therefore, my results are informative about the effects of a (mandatory) debt restructuring among existing borrowers (e.g., the $1.6 trillion of outstanding US student loans) whose ex-ante choices are fixed by definition. My analyses also treat the existing tax and transfer system, which is designed for the entire population and constrained by the political system, as given.[1]

[1]See Stantcheva (2017) for a joint analysis of the tax system and human capital financing policies.

My main empirical finding is that borrowers reduce their labor supply to lower repayments on income-contingent loans. These responses are larger among borrowers with more hourly flexibility, a lower probability of repayment, and who are more liquidity-constrained. However, my structural estimation shows that these responses are consistent with a moderate (Frisch) elasticity of labor supply and substantial frictions that limit labor supply adjustment. On the normative side, my estimates imply that moral hazard limits the optimal amount of insurance but that there are still significant welfare gains from income-contingent repayment. Specifically, restructuring from a fixed repayment contract to a constrained-optimal income-contingent loan increases borrower welfare by the equivalent of 0.8% of lifetime consumption at no additional fiscal cost. Adding forbearance to fixed repayment contracts is a poor substitute for income-contingent loans because it does not accelerate repayments from high-income borrowers. In sum, my results suggest that income-contingent repayment creates moral hazard that affects contract design but that it is too small to justify fixed repayment contracts.

There are several benefits to studying how income-contingent repayment affects labor supply in Australia. First, there is limited scope for selection because the only available contract is a government-provided income-contingent loan. This is useful for identifying moral hazard (Karlan and Zinman 2009) and contrasts with the US, where borrowers select into repayment contracts based on expected earnings (Karamcheva et al. 2020). Second, Australia was the first country to introduce income-contingent loans in 1989, meaning borrowers are familiar with the availability and design of these contracts, unlike in the US (Abraham et al. 2020; Mueller and Yannelis 2022). Finally, these loans can only be used to cover (government-controlled) tuition, implying that borrowers can only adjust their initial debt by changing their degree choices. This decision is likely less responsive than the other margins that borrowers in the US can adjust, such as room and board. Additionally, it suggests that the assumption in my normative analysis that ex-ante choices are fixed may be a reasonable approximation in this setting.

In the first part of this paper, I document evidence of moral hazard from income-contingent repayment: borrowers reduce their labor supply to lower repayments on income-contingent loans. I identify this behavioral response by leveraging a 2005 policy change that increased the income threshold above which all borrowers begin loan repayment. Figure 1 summarizes the effects of this policy change by showing that the income distribution of student debtholders exhibits significant bunching below the repayment threshold, both before and after the reform. I present two pieces of evidence that suggest this bunching reflects labor supply responses rather than solely income-shifting or tax evasion. First, the

Figure 1. Income Distribution for Debtholders around the Income-Contingent Repayment Threshold

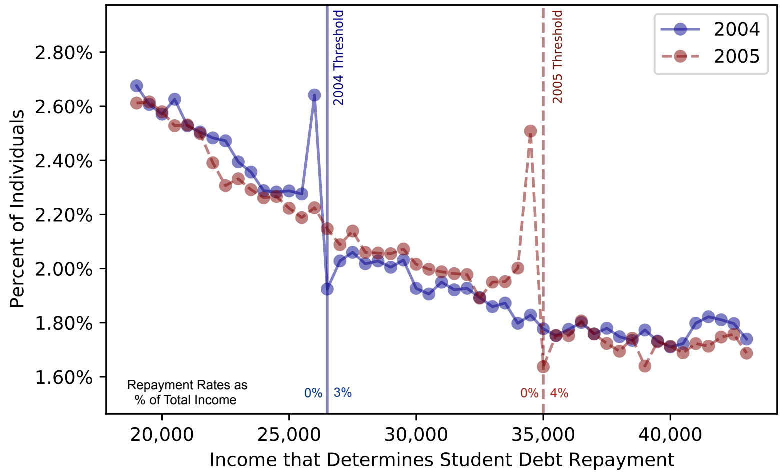

Notes: This figure shows the distribution of the income that determines repayments on income-contingent loans in 2004 and 2005, before and after the policy change. This income is called HELP income and equals taxable income (i.e., the sum of labor income, capital income, and deductions) plus investment losses, retirement contributions, foreign employment income, and fringe benefits. The vertical lines indicate the thresholds below which borrowers make no repayments and above which they repay 3% and 4% of their income. The sample is all debtholders subject to the criteria in Section 1.4. HELP income is deflated to 2005 AUD using the Consumer Price Index.

bunching is larger in occupations with high hourly flexibility (e.g., restaurant workers) and almost nonexistent in those with low flexibility (e.g., software engineers). Second, using data from Australia's Census, I find that borrowers below the repayment threshold work 2–3% fewer hours (i.e., 1–2 fewer weeks) per year than those above the threshold.

The second part of this paper develops a structural model of labor supply that quantitatively replicates the evidence in Figure 1. The purpose of the model is to translate this evidence into estimates of preference parameters and study the welfare implications of income-contingent repayment. In the model, borrowers choose consumption and labor supply over their life cycles. The two key ingredients are uninsurable income risk and endogenous labor supply, which create a trade-off between the insurance benefits and moral hazard costs of income-contingent repayment. Motivated by my empirical evidence, I do not explicitly model any forms of income-shifting or tax evasion. However, to the extent that the evidence in Figure 1 is driven by such "non-real" responses that create transfers to other agents, the model will overestimate the moral hazard costs of income-contingent repayment (Chetty 2009) and thus understate its benefits.

The evidence in Figure 1 is inconsistent with a frictionless formulation of this model. When borrowers' income crosses the repayment threshold, the fraction of *total* income that they repay increases from 0% to 3–4%. In a frictionless model, no borrowers would locate immediately above the threshold because locating below it delivers more leisure and cash on hand. Therefore, I introduce optimization frictions (Chetty 2012) to explain borrowers locating above the repayment threshold. In particular, adjusting labor supply requires paying a fixed cost, which could be monetary (e.g., a wage reduction) or psychological (e.g., hassle costs). Motivated by the variation in bunching across occupations, I allow this fixed cost to be stochastic and transition between two different values. A key benefit of modeling optimization frictions with a stochastic adjustment cost is that it nests a large set of alternative models, including the time-dependent Calvo (1983) model and the state-dependent (S, s) model (Caplin and Spulber 1987).

I estimate the model by simulating responses to the policy change in Figure 1 and find that they are consistent with a moderate labor supply elasticity and substantial optimization frictions. The key parameters that govern labor supply responses—the (Frisch) labor supply elasticity, two fixed costs, and their probabilities—are identified as follows. First, the labor supply elasticity is identified by the bunching below the repayment threshold: a larger elasticity implies more bunching. Second, the number of borrowers above the threshold jointly identifies the lower fixed cost and its probability because individuals with this lower cost are closer to their indifference condition for bunching. I then separately identify these two parameters by exploiting panel data: a higher probability of receiving the lower cost implies a larger fraction of individuals that were bunching below the policy change will also be bunching after the change. Finally, the larger adjustment cost is identified by the distribution of changes in hours worked from survey data. The estimation results show that the evidence in Figure 1 is quantitatively consistent with a labor supply elasticity of 0.15, fixed adjustment costs of 0.6% and 5% of average earnings, and a 15% probability of receiving the lower cost. Although I study labor supply responses to student loans rather than income taxes or wages, the estimated labor supply elasticity is close to the median of 0.14 from the meta-analyses in Keane (2011) and Chetty et al. (2012).

In the estimated model, two forces are quantitatively important for explaining the bunching in Figure 1: uncertainty about debt repayment and a demand for liquidity. Unlike income taxes, the incentives created by income-contingent repayment depend on the probability of debt repayment: for borrowers anticipating repayment, bunching below the repayment threshold transfers repayments over time rather than permanently reducing them. In a counterfactual where borrowers anticipate fully repaying their debt, the model predicts that the bunching decreases by 65%. Empirically, this is consistent with the fact that the amount of bunching is larger among borrowers with more debt and in occupations with lower lifetime incomes, both of whom have a lower probability of repayment. The remaining 35% of the bunching in the model is driven by a demand for liquidity: even when the present value of the change in repayments from locating below the repayment threshold is zero, borrowers may value the additional liquidity if they are liquidity-constrained. This importance of liquidity is supported empirically by the fact that borrowers below the repayment threshold have larger housing payments, which represent greater liquidity demands, and contribute less to tax-advantaged but illiquid retirement accounts.

In the final part of the paper, I use the estimated model to study the welfare impact of different repayment contracts. My analysis considers a social planner that maximizes borrower welfare by choosing one contract, holding fixed ex-ante choices and the tax system. This perspective speaks to how the loans of existing student debtholders (e.g., the \$1.6 trillion of US loans), whose ex-ante borrowing and education choices are fixed by definition, should be restructured. To the extent that non-pecuniary factors are the main drivers of education choices (as suggested by Patnaik et al. 2020), these results provide a good starting point for optimal contract design, more generally. As with any counterfactual analysis, a key concern is whether the estimated parameters are invariant to policy changes. However, the fact that the model is estimated using a policy change that did occur makes this more likely to be the case.

My main normative result is that income-contingent repayment generates meaningful welfare gains relative to fixed repayment. First, I show that the marginal value of public funds of moving from fixed repayment to several existing income-contingent loans ranges from 3.7 to 7.7. These values imply that the welfare gain of income-contingent repayment far exceeds its fiscal cost and are near the 75th percentile for over 100 other policies considered in Hendren and Sprung-Keyser (2020). Next, I solve for the income-contingent loan that maximizes the planner's objective subject to the constraint of raising the same revenue as fixed repayment. In my baseline analysis, I constrain the contract space to two dimensions, as in the US: an income threshold at which repayment begins and a repayment rate of income above this threshold. The resulting constrained-optimal income-contingent loan increases welfare relative to fixed repayment by the equivalent of a 0.8% increase in lifetime consumption at no additional fiscal cost. Although these gains imply that the moral hazard from income-contingent repayment is small relative to the benefits, it is still important for contract design. Absent labor supply responses, the constrained-optimal contract would provide more insurance with a repayment threshold that is over twice as high, doubling its welfare gain relative to fixed repayment.

Income-contingent loans perform well relative to three other methods of providing insurance: (anticipated) loan forgiveness, adding forbearance to fixed repayment contracts, and equity contracts. First, adding forgiveness after a fixed horizon, as in the US and

UK, decreases the welfare gains from income-contingent loans. Once income-contingent repayment has been implemented, forgiveness operates as a poorly targeted subsidy by transferring repayment burdens from older to younger, more liquidity-constrained borrowers. Second, a fixed repayment contract with forbearance, a form of income-contingency that pauses repayments for low-income borrowers, also underperforms relative to the constrained-optimal income-contingent loan. This is because income-contingent loans accelerate repayment from high-income borrowers, enabling them to provide more insurance at a given cost. Finally, equity contracts in which borrowers repay a share of their income for a fixed horizon can outperform income-contingent loans, but only if the horizon is longer than those implemented in practice. However, even with a longer horizon, equity contracts create significantly more redistribution than income-contingent loans because they decouple repayments from debt balances. This large redistribution suggests that equity contracts might cause ex-ante responses not captured by the model (e.g., additional borrowing) and, therefore, that income-contingent loans may be a more robust mechanism for implementing income-contingent repayment.

Related literature and contribution. This paper is most closely related to the literature on financing human capital, which spans household finance, public finance, and labor economics. Friedman (1955) popularized the idea that student loans should be equity-like and advocated using income-sharing agreements. Adverse selection prevents the private provision of these contracts (Herbst and Hendren 2021; Herbst et al. 2023), but a growing number of governments have attempted to correct this market failure by introducing income-contingent loans (Barr et al. 2019).[2] Theoretical work suggests that these loans provide a close approximation to Mirrlees (1974)–style optimal policies (Lochner and Monge-Naranjo 2016; Stantcheva 2017), which is supported by two empirical strands of literature on student loans (see Yannelis and Tracey 2022 for a review). The first documents debt overhang created by fixed repayment contracts, in which reductions in student debt decrease delinquencies and increase income and mobility (Di Maggio et al. 2021), increase homeownership (Mezza et al. 2020), and change education and occupation choices (Luo and Mongey 2024; Chakrabarti et al. 2020; Folch and Mazzone 2021; Hampole 2022; Murto 2022; Huang 2022).[3] The second studies income-contingent loans as a tool to mitigate these effects, finding reductions in delinquencies (Herbst 2023), mortgage defaults (Mueller

[2]Other government policies toward human capital include subsidies for educational expenses (Benabou 2002; Bovenberg and Jacobs 2005; Stantcheva 2017) and grants (Abbott et al. 2019; Ebrahimian 2020).

[3]A related literature emphasizes the importance of credit constraints for college attendance (Carneiro and Heckman 2002; Lochner and Monge-Naranjo 2012), which student loans help relax (Amromin and Eberly 2016; Black et al. 2022).

and Yannelis 2019), and the passthrough of income to consumption (Gervais et al. 2022).[4]

This paper makes three contributions to this literature. First, it empirically characterizes the moral hazard created by income-contingent repayment. Second, it provides a model of labor supply that replicates this evidence, finding an important role for optimization frictions, liquidity constraints, and dynamics. Finally, it quantifies the implications of this moral hazard for optimal contract design. Prior structural models of income-contingent repayment have emphasized its insurance benefits (Boutros et al. 2022) and its effects on other decisions, such as college enrollment (Matsuda and Mazur 2022), the wage-amenity trade-off (Luo and Mongey 2024), job search (Ji 2021), earnings profiles (Alon et al. 2024), and homeownership (Folch and Mazzone 2021). Relative to this literature, this paper provides the first model, to my knowledge, that replicates the moral hazard effects of income-contingent repayment. However, the fact that the model in the paper is rich enough to quantitatively match my empirical evidence requires it to abstract from these other mechanisms to maintain tractability. An important task for future work is to understand how these mechanisms interact in a single quantitatively realistic model.

The closest paper is Britton and Gruber (2020) (BG), who study bunching around the income-contingent repayment threshold in the UK. Unlike in Australia, the threshold in the UK changes the marginal repayment rate. BG find a small amount of bunching and show that it is consistent with an income elasticity in the static Saez (2010) model of essentially zero. Conversely, because HELP changes the average repayment rate, the responses in my setting allow me to estimate a model of labor supply that is dynamic. With this model, I show that the evidence in Figure 1 *and* BG are consistent with a labor supply elasticity of 0.15. In contrast, the evidence from BG alone does not say whether the lack of bunching is driven by a low structural elasticity, the dynamic incentives created by income-contingent repayment, or optimization frictions—each of which has different normative implications.

There are several other literatures to which this paper is related. First, it is part of the literature in public finance that studies the insurance–moral hazard trade-off in social insurance (Chetty and Finkelstein 2013), such as unemployment insurance (Gruber 1997), health insurance (Einav et al. 2015), disability insurance (Bound et al. 2004), and consumer bankruptcy (Indarte 2023). My finding that borrowers reduce their labor supply to locate below the repayment threshold, which increases liquidity more than wealth, is consistent with liquidity driving responses to other forms of social insurance (Chetty 2008; Indarte

[4] Alternative policies include making student debt dischargeable, which induces strategic default (Yannelis 2020); implementing universal forgiveness, which would be regressive (Catherine and Yannelis 2023); and offering targeted forgiveness, which borrowers appear to value but fail to take up (Jacob et al. 2023).

2023; Ganong and Noel 2023). Additionally, it complements the finding in Ganong and Noel (2020) that borrowers' decisions—in this case, labor supply instead of consumption—respond more to changes in short-term payments than long-term obligations.

Second, by studying state-contingent contracts, this paper is part of the literature in macro-finance on household security design. Motivated by evidence of imperfect risk-sharing (Cochrane 1991) and the household balance sheet channel (Mian and Sufi 2014), this literature studies contracts that make liabilities more state-contingent (Piskorski and Seru 2018), such as shared-appreciation mortgages (Caplin et al. 2007; Hartman-Glaser and Hébert 2020; Greenwald et al. 2021; Benetton et al. 2022) or adjustment-rate mortgages conditioned on aggregate shocks (Campbell et al. 2021). This paper contributes by studying one of the longest-running examples of such contracts and characterizing the welfare gains from alternative forms of state-contingent repayment. A distinguishing feature of my setting is limited strategic default, as student loans cannot be discharged in bankruptcy.

Finally, this paper contributes to the extensive literature on labor supply, which Appendix A reviews in detail. While my structural model could have been estimated using evidence from this literature, there are three benefits to using the evidence from the first part of this paper. First, it identifies responses in the relevant sample for the policy question of interest: college-graduates repaying their loans. Second, income-contingent loans are fundamentally different from taxes because of the intertemporal tradeoffs that they create. Both of these reasons imply the model's counterfactuals involve less extrapolation than there would be if I estimated it using other evidence. Third, a central challenge in the part of this literature that uses bunching at tax rate discontinuities to identify income elasticities (e.g., Saez 2010; Chetty et al. 2011) is that the patterns typically differ from the predictions of frictionless models, which has motivated various models with optimization frictions (Chetty 2012; Kleven and Waseem 2013; Anagol et al. 2022). In contrast, the empirical evidence in this paper that allows me to estimate the first, to my knowledge, dynamic model of labor supply that incorporates both time- and state-dependent adjustment.

1 Institutional Background and Data

1.1 Overview of Australia's Higher Education Loan Programme (HELP)

In Australia, higher education is primarily financed using government-provided student loans through the Higher Education Loan Programme (HELP), which was introduced in 1989. There are five HELP programs that provide income-contingent loans to Australian citizens for different purposes. This section provides an overview of the two largest programs called HECS-HELP and FEE-HELP, which historically have accounted for over 90% of HELP borrowing; Appendix C.1 presents additional details.[5] HELP loans provided through these two programs can be used to finance tuition for undergraduate and graduate degree programs. Tuition at public institutions is controlled by the government and varies by degree, while private universities generally charge higher tuition. Most degrees at public institutions are classified as Commonwealth Supported Places (CSPs), in which the government provides a subsidy in the form of a contribution to the tuition owed by the student. The tuition remaining after the government's contribution is deducted is paid by the student and is called the student contribution. As of 2023, student contributions ranged from $4,124 to $15,142 AUD per year ($2,700 to $10,100 USD), with undergraduate degrees typically lasting 3–4 years. The number of CSPs in Australia has generally been capped by the government, except from 2012–2017 (D'Souza 2018; Norton 2019).

Individuals who receive a CSP can either pay their student contribution upfront or borrow through the HECS-HELP program. Those who pursue degrees that are not CSPs are liable for full tuition and can either pay upfront or borrow through FEE-HELP. In both programs, most individuals choose to do the latter, with less than 10% of balances in 2022 being paid upfront (Department of Education and Training 2023). For borrowers who receive CSPs and access HECS-HELP, the largest program, their initial debt is equal to their student contribution. Given an average undergraduate student contribution of around $6,000 USD per year, tuition is comparable to that for US in-state public undergraduate degrees, which averages $9,000 USD per year (Hanson 2023). Figure A2 plots the time series of student contributions, aggregate HECS-HELP borrowing, and upfront payments.

HELP debt balances in subsequent years grow with the CPI net of repayments, implying

[5] Figure A1 plots aggregate borrowing and discusses the details of the different HELP programs.

that HELP debt has a zero real interest rate. Individual i's compulsory repayment in year t is

$$\text{HELP Repayment}_{it} = \min\{r_t(y_{it}) * y_{it}, D_{it}\},$$

where y_{it} denotes HELP income, $r_t(\cdot)$ is the income-dependent repayment rate, and D_{it} denotes the beginning-of-year debt balance. HELP income is the taxable income reported in a personal income tax return plus a few adjustments discussed in Section 1.5. Collection of HELP payments is integrated with the income tax system, which is crucial for HELP's success relative to other income-contingent loan programs (Barr et al. 2019). All individuals file tax returns in Australia, so y_{it} refers to *individual* rather than household HELP income. For most borrowers, HELP repayments are withheld by their employer during the year and deducted from their debt after they file their tax returns; see Appendix C.1 for additional details. Individuals also have the option to make voluntary repayments at any time.[6]

Repayment of HELP debt continues either until the remaining balance equals zero or until death. This means that HELP effectively forgives debt at the end of working life when borrowers stop generating sufficient income to make compulsory repayments, similar to the forgiveness embedded in US income-driven repayment plans. Partial repayment is common: as of 2004, approximately 25% of debt balances were forecast to be written off due to death (Martin 2004). As in the US, HELP debt cannot be discharged in bankruptcy, implying borrowers cannot default on their loans. Nevertheless, borrowers could choose to avoid repayment by not lodging tax returns or lying about whether they have HELP debt. In addition to being illegal, doing so implies large penalties, such as administrative penalties of 95% and additional interest charges.

1.2 2004–2005 Policy Change to HELP Repayment Rates

The policy change that I exploit is a 2004–2005 change in $r_t(\cdot)$. The left panel of Figure 2 plots repayment rates as a function of HELP income before the policy change in blue and after the change in red.[7] The most significant change was the movement of the repayment threshold, the point at which borrowers start making repayments, from approximately $26,000 AUD to $35,000 AUD. The median debtholder has HELP income between these two thresholds, so this policy change generated reductions in repayments for many borrowers.

[6] I do not have access to micro-data on voluntary repayments. To the extent that borrowers make such repayments in the data, despite the zero interest rate, there are two likely explanations: (i) individuals are debt-averse or (ii) paying down debts has other benefits, such as helping with mortgage qualification.

[7] Before 1998 and after 2018, there were other changes to the repayment schedule; see Appendix C.1.

Figure 2. HELP Repayment Rates as a Function of Income: Before and After the Policy Change

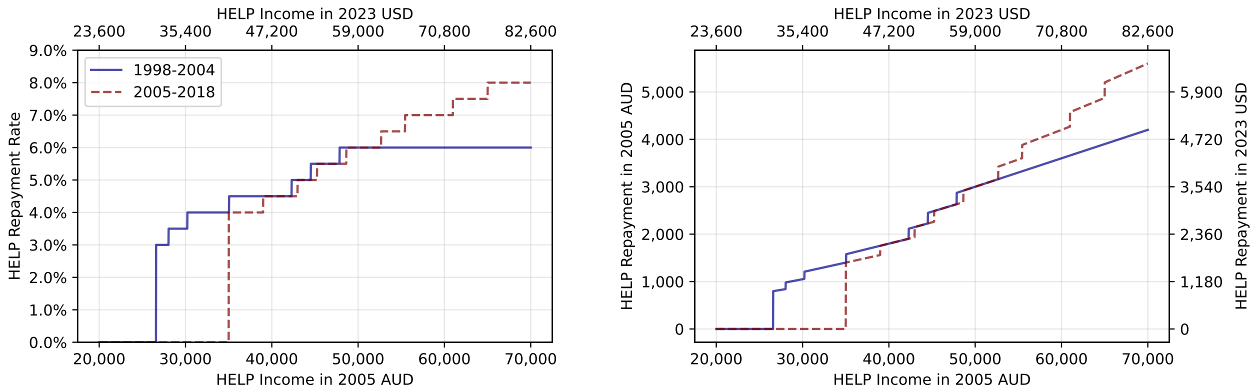

Notes: The left panel of this figure shows HELP repayment rates as a percentage of HELP income, which are average rather than marginal repayment rates. The right panel shows the required HELP payments implied by the repayment rates on the left in 2005 Australian dollars on the left axis and 2023 US dollars on the right axis. The blue and red lines correspond to the rates before and after the policy change, respectively. The bottom axis in both panels is HELP income measured in 2005 Australian dollars and the repayment schedule, which is constant in real terms. The top axis measures HELP income in 2023 US dollars calculated with the AUD/USD exchange rate from June 2005 and the US CPI inflation rate between June 2005 and January 2023.

Importantly, this policy change applied to all new and existing HELP debtholders.[8]

The right panel of Figure 2 plots required repayments in AUD, which illustrates that the repayment threshold creates a large incentive to reduce HELP income by generating a discontinuity in the *average* rather than marginal repayment rate. For example, consider a borrower with $35,000 of HELP income in 2005. For this borrower, earning an extra $1 of income results in a required HELP repayment of $35,001 × 4% ≈ $1,400 (i.e., the repayment threshold is a "notch" in the language of Kleven and Waseem 2013).

If borrowers chose their labor supply statically and treated repayments like an income tax, no borrowers would locate immediately above the repayment threshold because doing so would deliver less take-home pay and leisure. However, income-contingent repayment of debt differs from a tax in that it involves *dynamic*, in addition to static, trade-offs. For example, consider a borrower at $t = 0$ with a debt balance D_0 who is deciding between locating below versus above the 2005 repayment threshold. Locating below the threshold decreases her repayments at $t = 0$ by $1,400. However, under the assumption that this borrower's income at $t = 1$ will be high enough that the required payment is above D_0, this $1,400 repayment is simply transferred from $t = 0$ to $t = 1$. As a result, the present value of the reduction in repayments from locating below the repayment threshold is $\left(1 - \frac{p}{1+r}\right) \times \$1{,}400 = \frac{r+1-p}{1+r} \times \$1{,}400$, where r is the real interest rate and p is the probability

[8] This approach of identifying moral hazard by looking at the responses to changes in contract structure among individuals who have already taken up a contract has been applied in a variety of selection markets, such as consumer credit (Karlan and Zinman 2009) and mortgage markets (Gupta and Hansman 2022).

of repayment at $t = 1$.[9] In other words, locating below the threshold has a large impact on current payments, but a much smaller effect on the present value of payments as $p \to 1$. This is similar to the maturity extension program studied in Ganong and Noel (2020), which also increases borrowers' liquidity more than wealth.

There are several reasons to believe that the HELP repayment function and the changes to it are salient to debtholders. First, the repayment function is indexed to inflation, which means that it updates every year. When it is published at the beginning of each tax year, the government ensures that the change receives press coverage.[10] Second, the policy change received media coverage at the time of its implementation (Marshall 2003). Finally, the fact that HELP income determines repayment rates and features a repayment threshold has not changed since the program's introduction in 1989, meaning that debtholders are likely to understand the program's structure.

Government policy documents and media articles suggest that the primary reason for the policy change was to provide relief for lower-income borrowers, whose payments were burdensome and contributed little to the total HELP budget (Nelson 2003). In addition to changing the repayment function, other changes were implemented in 2004–2005, such as the introduction of HELP loans for non-CSPs through FEE-HELP and a 25% increase in student contributions (see Figure A2). These other changes, discussed in detail by Beer and Chapman (2004), were primarily aimed at those entering their degree programs rather than those repaying HELP debt. The simultaneous implementation of these other changes with the change to the repayment threshold is not ideal for my analysis. However, it likely has a minimal effect, given that I focus on identifying ex-post moral hazard.

1.3 Benefits of Studying Income-Contingent Repayment in Australia

In addition to the presence of high-quality administrative data and policy variation, there are several benefits to using HELP to identify labor supply responses to income-contingent repayment. First, there is limited selection on hidden information, such as unobservable types or expected moral hazard (Karlan and Zinman 2009), because HELP is the only government-provided student loan. The same is not true in the US (Karamcheva et al. 2020) or in countries with private providers of income-sharing agreements (Herbst et al. 2023). In principle, individuals in Australia could seek external financing from a bank or university. However, there is little economic incentive to do so because the interest rate

[9] Technically, r is the difference between the HELP interest rate and the borrower's private rate.
[10] For an example of an announcement, see here.

would exceed the zero real rate on HELP loans. The primary margin along which there is scope for selection is whether to pay upfront or borrow through HELP, but the zero interest rate on HELP loans again implies little incentive to pay upfront.[11]

A second benefit of this setting is the likely limited *ex-ante* moral hazard, in which borrowers increase their HELP debt in anticipation of a lower probability of future repayment. HELP can only be used to cover tuition at public undergraduate institutions, which make up over 94% of the domestic enrollment share and have government-controlled tuition. As a result, borrowers can only adjust their debt by changing their choice of degree or institution, which are likely less responsive than the other margins that borrowers in the US can adjust.

The third benefit of studying HELP is that it is the longest-running government-provided income-contingent repayment program. The fact that this program has been around since 1989 suggests that borrowers understand the repayment incentives. The same is not true in the US, where borrowers are unaware of the existence and structure of income-driven repayment options (Abraham et al. 2020; Mueller and Yannelis 2022; JPMorgan Chase 2022). Finally, there are likely limited responses on the supply-side due to government tuition control. If this were not the case, changes in government-provided contracts could pass through to tuition and thus debt balances (Kargar and Mann 2023).

The institutional differences between Australia and the US make the former advantageous for identifying labor supply responses to income-contingent repayment. However, Appendix C.2 presents a detailed discussion of how these and other differences would likely influence the effectiveness of income-contingent repayment in the US.

1.4 Data Sources

I use restricted-access de-identified administrative data from several sources. First, I use individual income tax returns from the Australian Taxation Office (ATO), which contain panel data on income components and basic demographic characteristics. Second, I use administrative data on HELP from the ATO that include debt balances, repayments, and a flag for whether individuals acquired new debt balances in a given year. Two limitations of these data are that they do not allow me to identify any information on the source of borrowing, such as degree choice, and they aggregate debt across all HELP programs. Third, I leverage administrative data on superannuation balances and contributions from the

[11] In earlier years of HELP, upfront payments were subject to a discount, which created a small incentive to pay upfront. See Appendix C.1 for additional discussion.

ATO. These three datasets are linked for the universe of Australian taxpayers between 1991 and 2019 in the ATO Longitudinal Information Files, known as *ALife*. Starting from the population dataset in *ALife*, I restrict attention to individual–year observations for which the individuals (i) are between ages 20 and 64, (ii) are residents in Australia for tax purposes, (iii) are not exempt from HELP repayment due to a Medicare exemption, and (iv) do not have any income from discretionary trusts.[12] I use this sample, which covers all 4 million unique debtholders between 1991 and 2019, for my main analysis.

To obtain data on hours worked and housing payments, I use a linkage of these ATO data with the 2016 Census of Population and Housing. This linkage cannot be performed with *ALife* directly, so I instead perform the merge through the Australian Bureau of Statistics Multi-Agency Data Integration Project (MADIP). The ATO data in MADIP have the same sample coverage as the population-level *ALife* data but a restricted set of variables. Due to data limitations, I use the first three filters from the *ALife* sample to construct a cross-sectional MADIP sample in 2016, the year in which the census was administered.

I supplement these administrative datasets with the Household, Income and Labour Dynamics in Australia Survey (HILDA), a household survey conducted by the Melbourne Institute between 2002 to 2021. HILDA is similar to the Survey of Consumer Finances in the US, except that it is a panel rather than a repeated cross-section.

1.5 Summary Statistics

Table 1 presents summary statistics on the *ALife* sample, the main sample in my analysis, for individuals with and without HELP debt. Relative to non-debtholders, debtholders tend to be younger, less likely to be wage-earners (defined as having any self-employment income from partnerships, sole-traders, or personal-services), and have lower taxable income. The most important variable is HELP income, which determines a borrower's HELP repayment rate. HELP income equals taxable income plus several other adjustments, such as adding back reportable superannuation contributions, investment losses, and fringe benefits. These adjustments are not relevant for most individuals: the difference between HELP and taxable income is less than \$100 for over 93% of the observations in 2004. I decompose HELP

[12]In Australia, there are unit trusts, in beneficiaries have no discretion over entitlements, and discretionary trusts, in which beneficiaries have full discretion. Discretionary trusts have been identified as potential sources of tax evasion (Australian Council of Social Service 2017), but *ALife* does not have information on the sources of trust income. I drop these observations to avoid attributing possible tax evasion to labor supply responses.

Table 1. Summary Statistics

	Non-Debtholders (1)	Debtholders (2)
Demographics		
Age	41.1	29.5
Female	0.46	0.60
Wage-Earner	0.85	0.91
Income Totals (in 2005 AUD)		
Taxable Income	37,695	27,796
HELP Income	38,756	28,586
Income Components (in 2005 AUD)		
Salary & Wages	32,415	26,068
Labor Income	35,480	27,136
Interest & Dividend Income	726	242
Capital Income	1,221	324
Net Deductions	-1,548	-1,099
HELP Variables		
HELP Debt (in 2005 AUD)	·	10,830
HELP Payment (in 2005 AUD)	·	991
HELP Debt at Age 26 (in 2005 AUD)	·	13,156
HELP Payment at Age 26 (in 2005 AUD)	·	1,305
HELP Income < 0% Threshold	0.50	0.65
HELP Income < 2004 0% Threshold	0.37	0.51
HELP Income < 2005 0% Threshold	0.52	0.67
Number of Unique Individuals	19,484,517	4,013,382
Number of Individual-Year Observations	247,118,713	27,316,037

Notes: This table presents summary statistics from the *ALife* sample from 1991 to 2019, subject to the sample selection criteria discussed in Section 1.4. Column (1) uses all individual–years with zero HELP debt; column (2) uses all individual–years with positive HELP debt. The values for all continuous variables represent means. All continuous variables are deflated to 2005 dollars based on the HELP threshold indexation rate. All continuous variables except HELP Debt and HELP Repayment are winsorized at 2%–98%. HELP Income < 0% Threshold corresponds to the mean of a dummy variable for whether HELP income in an individual–year was below the 0% HELP repayment threshold. HELP Income < 0% 2004 Threshold and HELP Income < 0% 2005 Threshold correspond to means between 1998–2004 and 2005—2018 for whether HELP income in an individual–year was below the HELP repayment threshold, respectively, after the thresholds are adjusted for inflation. Additional details on variable construction are presented in Appendix C.3.

Income into three terms:

$$\text{HELP Income} = \text{Labor Income} + \text{Capital Income} - \text{Net Deductions}. \tag{1}$$

Labor Income is defined as the sum of salary and wages, tips and allowances, and self-employment income. This represents the largest source of income for most individuals: 95% for debtholders and 91% for non-debtholders. Capital Income is defined as the sum of interest income, dividend income, capital gains, government superannuation and annuity income, rental income, and trust income. Net Deductions is defined as the residual in (1).

Table 1 shows that debtholders have lower income and claim fewer deductions than non-debtholders, which is not surprising given the age differences between the two groups. The average debt balance among debtholders is $10,800 in 2005 AUD ($12,800 in 2023 USD) and $13,200 in 2005 AUD ($15,600 in 2023 USD) among 26-year-old debtholders, which is the age at which most individuals have finished university in Australia. Notably, the 2004–2005 policy change had a large impact on the number of debtholders below the repayment threshold: the fraction below the threshold moved from 51% to 67% after the change. Among 26-year-old debtholders, this fraction moved from 35% to 55%, both of which are lower than among the full sample of borrowers because 26-year-olds are below the average age and hence have a lower average income.

Figure A3 shows how debt balances vary with age: most borrowers' debt balances peak in real terms between ages 24 and 26 and are paid down in their mid-30s. However, around 15% of borrowers who have debt at age 22 in 1991 still have debt at age 50 in 2019. Given the increase in real tuition over time, this number is forecasted to increase (Nelson 2003).

2 Empirical Evidence of Labor Supply Responses

This section uses the variation in HELP repayment rates from Figure 2 to characterize how labor supply responds to income-contingent repayment.

2.1 Bunching of HELP Income Below Repayment Threshold

The first result is the presence of bunching below the repayment threshold. Figure 3 plots the distribution of HELP income for borrowers with HELP debt in the three years before and after the policy change. HELP income is deflated to 2005 Australian dollars using the HELP threshold indexation rate. The vertical line in each plot corresponds to the HELP repayment threshold in that year, which is constant in real terms across the years in which there is no policy change. These plots focus on borrowers with HELP income within $8,000 of the two repayment thresholds—around 40% of the entire population of debtholders.

These results show that there is significant bunching below the repayment threshold from 2002 to 2007, but minimal bunching below the smaller thresholds (which will be useful in Section 3). For the three years before the policy change, shown in the left panels of Figure 3, the amount of bunching and shape of the income distribution remain relatively

Figure 3. Income Distribution of HELP Debtholders around the Repayment Threshold

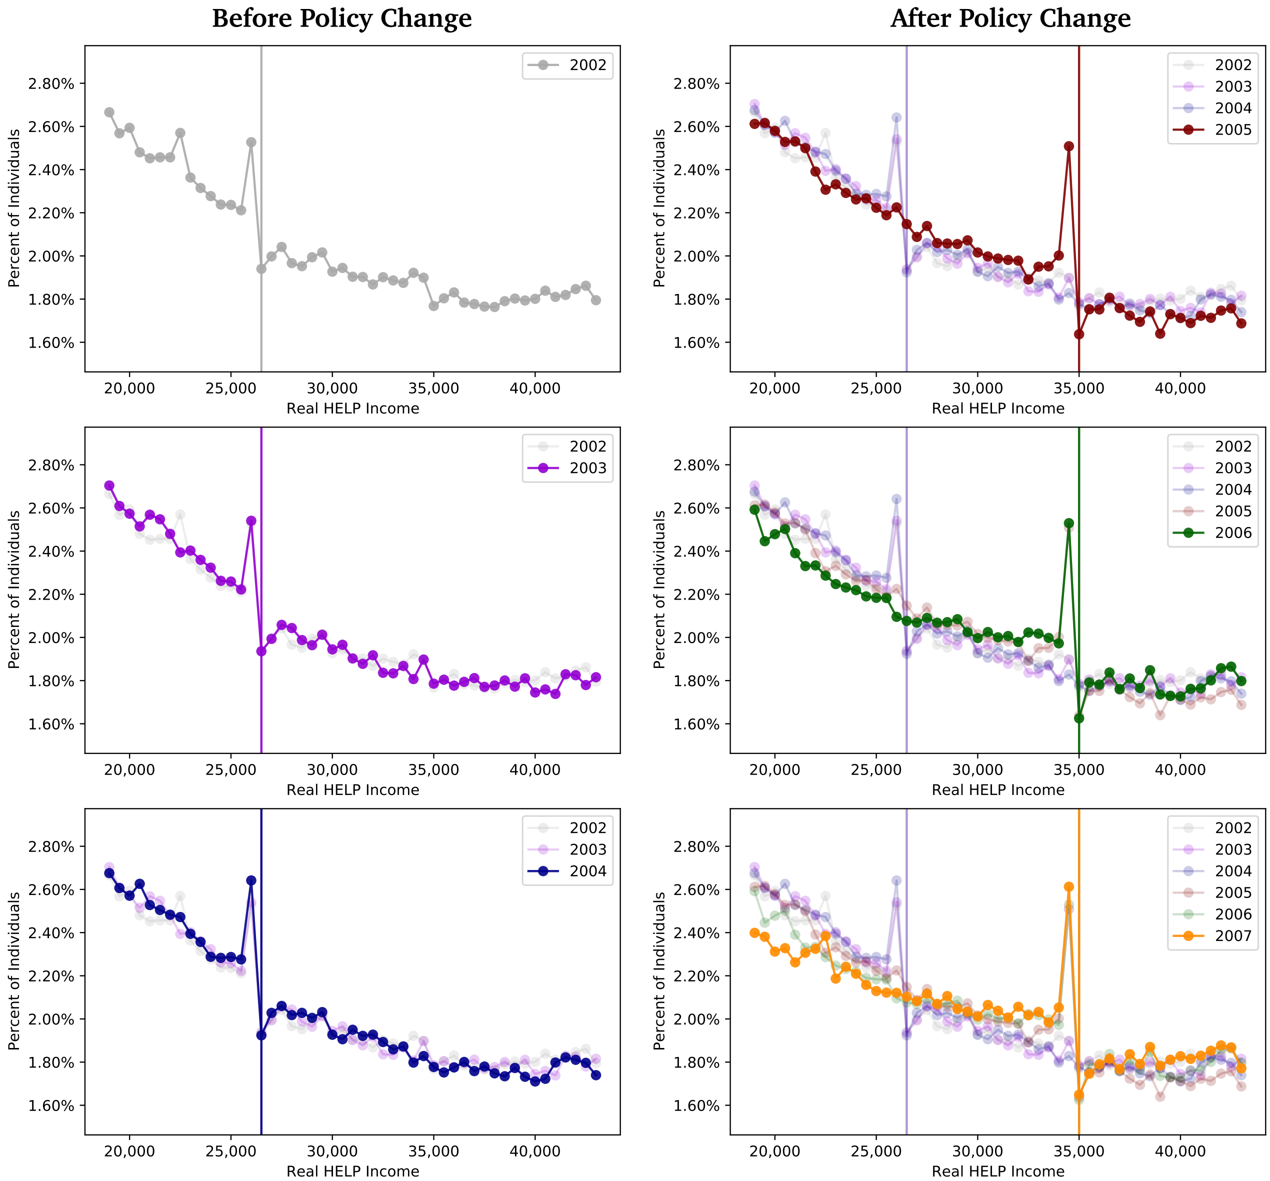

Notes: This figure shows the distribution of real HELP income in Australian dollars, which determines a borrower's repayment rate on her income-contingent loan, in the three years before and after the policy change to the repayment schedule between 2004 and 2005 that is illustrated in Figure 2. The vertical lines in the left (right) panel indicate the threshold above which borrowers begin making debt payments of 3% (4%) of their income before (after) the policy change. Each bin represents $500, and the plot focuses on borrowers within $8,000 of the two repayment thresholds. The bins are chosen so that they are centered around the 2005 repayment threshold. HELP income is deflated to 2005 Australian dollars using the HELP threshold indexation rate, which is based on the annual CPI. The sample is the *ALife* sample defined in Section 1.4, restricted to individuals with positive HELP debt balances in each year.

constant. However, the right panels show two changes to the income distribution after the policy change. First, the bunching at the 2004 repayment threshold disappears completely. Second, bunching reappears immediately below the new repayment threshold, providing clear evidence that borrowers adjust their income to reduce income-contingent repayments.

The fact that the bunching in Figure 3 responds quickly to the policy change shows that it is not driven by mechanical features of Australia's tax system, such as the tendency to report incomes at round numbers. However, a possible threat to identification is the presence of other changes between 2002 and 2007 that affected individuals' incentives to report incomes of certain values. Although it is unlikely that this could explain the evidence in Figure 3, given that the bunching is sharp around the repayment threshold, I assess this possibility by examining the income distribution of non-debtholders in Figure A4. In contrast to the income distribution of debtholders, this distribution shows no changes around the repayment threshold either before or after the policy change.[13]

The bunching in Figure 3 is also present in the distribution of labor income, one of the three components of HELP income in (1). Figure A5 follows Chetty et al. (2011) and examines a sample of borrowers whose primary source of income is labor income and who thus require similar values of labor income to generate HELP income at the threshold. I then compute a measure of the bunching from Chetty et al. (2011) and find that it is 83% as large for labor income as for HELP income.

2.2 Bunching Increases with Hourly Flexibility

Next, I show that the bunching in Figure 3 is greater in occupations with more hourly flexibility. Using HILDA, I measure the amount of hourly flexibility in each 2-digit ANZSCO occupation, the finest level at which *ALife* reports occupation codes, as the standard deviation of annual changes in log hours worked. This measure is highest for workers in occupations where it is relatively easy to adjust hours, such as hospitality workers (e.g., bartenders) and food preparation assistants (e.g., fast-food workers), and lowest for those where it is more difficult, such as ICT professionals (e.g., software engineers). Table A2 shows the value of this measure for each occupation in my sample.

Figure 4 plots the amount of bunching between 2005 and 2018 among wage-earners below the new repayment threshold relative to hourly flexibility. I focus on the period after the policy change because this is when *ALife* offers comprehensive coverage of occupation codes. Each point represents an occupation, and I measure the amount of bunching as the ratio of the number of borrowers in that occupation within $2,500 below to the number above the threshold so that a ratio of one indicates no bunching (similar to Chetty et al. 2013). The results show that bunching is more common in occupations with greater hourly

[13]There are small changes in the income distribution of non-debtholders at lower values of income, which reflect changes in real terms of the second income tax bracket.

Figure 4. Variation in Bunching across Occupations Based on Hourly Flexibility

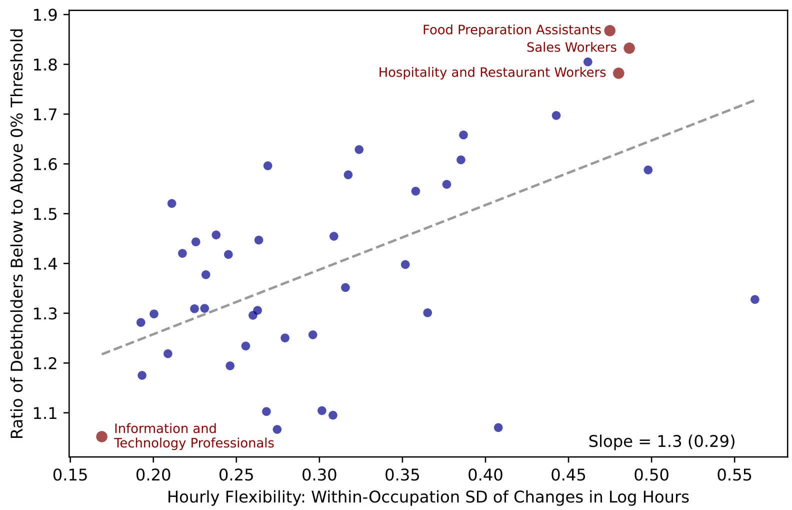

Notes: This figure plots the relationship between the amount of bunching below the repayment threshold and hourly flexibility by occupation, where each point represents a 2-digit ANZSCO occupation. Bunching is measured as the ratio of the number of borrowers in that occupation within \$2,500 below the repayment threshold to the number within \$2,500 above the threshold over 2005 to 2018. Hourly flexibility is measured as the standard deviation of annual changes in log hours worked from HILDA; see Figure A6 for an alternative measure. The highlighted points correspond to occupations described in the text. The gray dashed line is the regression line, with the estimated slope and standard error reported at bottom right. The sample is the *ALife* sample defined in Section 1.4, restricted to the subset of individual–years for which the individuals are wage-earners and have positive HELP debt balances.

flexibility. For example, ICT Professionals have the lowest hourly flexibility with a standard deviation of annual changes in log hours of 0.17. In this occupation, there is only 5% more borrowers below than above the threshold. In contrast, hospitality workers have almost three times more hourly flexibility and exhibit significantly more bunching, with 80% more borrowers below than above the threshold. Quantitatively, Table A3 shows that hourly flexibility explains 34% of the variation in bunching across occupations.

One concern with the evidence in Figure 4 is that hourly flexibility might be correlated with tax evasion or income-shifting across occupations. To assess the importance of evasion, I calculate the share of workers in each occupation that receives labor income from allowances, tips, director's fees, consulting fees, or bonuses. This variable is a proxy for tax evasion because it is easier to misreport these other sources of income relative to salary and wages (Paetzold and Winner 2016; Slemrod 2019). Figure A7 shows that this measure, unlike hourly flexibility, exhibits little correlation with the amount of bunching. Nevertheless, there are other differences across occupations that could explain these findings. Therefore, Section 2.3 provides complementary evidence of a labor supply response using a different source of variation.

Figure 5. Self-Reported Hours Worked around the Repayment Threshold

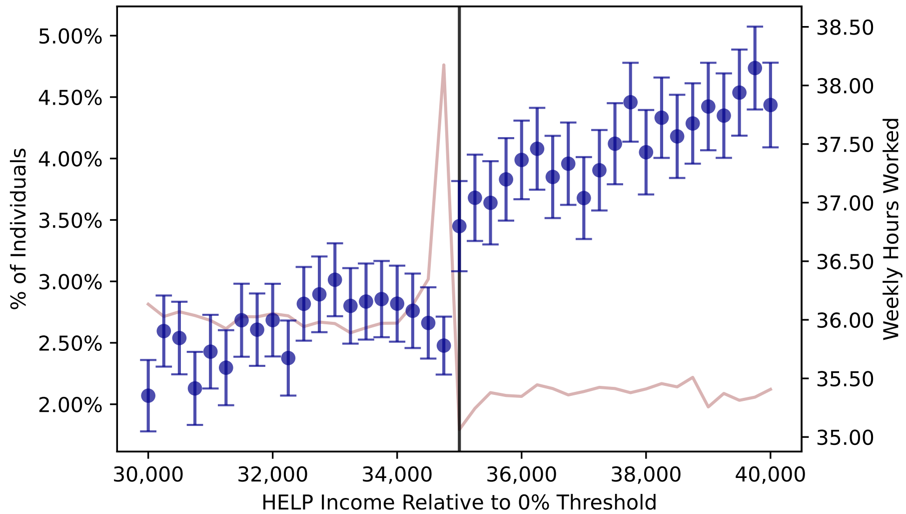

Notes: This figure plots the 2016 HELP income distribution in red, and measured on the left axis. HELP income is deflated to 2005 with the HELP threshold indexation rate, which is based on the annual CPI. Each bin represents $250, and the bins are chosen so that they are centered around the 2005 repayment threshold. The blue points present the average value of individuals' reported hours worked from the 2016 Census of Population and Housing within each bin, along with 95% confidence intervals. The sample is the cross-sectional MADIP sample described in Section 1.4, restricted to individuals with positive HELP debt balances.

2.3 Borrowers Below the Repayment Threshold Work Fewer Hours

A second piece of evidence that suggests that the bunching in Figure 3 reflects, at least in part, labor supply responses is that borrowers below the repayment threshold work fewer hours. I measure hours worked using a question in the 2016 Census of Population and Housing in which individuals report the number of hours worked during the week before the census night. Figure 5 plots the average hours worked in $250 bins of HELP income around the repayment threshold in the census year 2016, in addition to the distribution of HELP income in red. The results show that borrowers locating immediately below the threshold work on average 1 hour less per week than those immediately above it, which is 2.6% of the standard 38 hour workweek in Australia.[14] This adjustment occurs within a borrower's current occupation: Figure A8 finds little evidence that those below the repayment threshold are more likely to have switched occupations.

The results in Figure 5 are subject to two caveats. First, as discussed in Section 1.4, the MADIP and *ALife* samples differ slightly. To mitigate concerns about sample selection, Figure A11 shows that the distribution of HELP income in 2016 across the two samples is quantitatively similar. Second, these data on hours worked are self-reported by employees,

[14]These results are not driven by a group of borrowers outside the labor force: Figure A10 shows that the patterns are nearly identical in the sample of borrowers earning positive labor income.

which introduces concerns about reporting issues. For this reason, I do not target this evidence directly when estimating the structural model.

2.4 Bunching Decreases with Probability of Repayment

Next, I show that the amount of bunching below the repayment threshold increases with debt balances. To measure the amount of bunching, I construct a bunching statistic following the literature that uses discontinuities in tax rates to estimate income elasticities. First, I fit a five-piece spline to each distribution, leaving out the region $\mathcal{R} = [\$32,500, \$35,000 + X]$. The choice of $\$32,500$ represents a conservative estimate of where the bunching begins, and X is a constant intended to reach the upper bound at which the income distribution is affected by the threshold. This spline corresponds to an estimate of the counterfactual distribution absent the threshold. Next, I iterate on X so that this counterfactual density integrates to 1. Finally, I compute the bunching statistic, b, as:

$$b = \frac{\text{observed density in } \mathcal{R}}{\text{counterfactual density in } \mathcal{R}} - 1. \tag{2}$$

This bunching statistic is an estimate of the excess number of borrowers below the threshold relative to a counterfactual distribution in which it did not exist.[15]

Figure 6 shows the value of the estimated b across groups of borrowers with different ages and debt balances. I split ages into five-year bins, which gives a similar number of observations within each bin, and then split debt balances at their median value within each age and year. The results show two patterns. First, for all age groups, the estimated value of b is higher among borrowers with above-median debt balances. This finding suggests that the probability of eventual repayment is an important determinant of labor supply responses. The second pattern is that the amount of bunching decreases moderately with age: the estimated b is $22 - 33\%$ lower among borrowers above 40 than those below 25. This finding suggests that liquidity constraints, which are tightest among young borrowers, are important, which I test more directly in the next section.

[15] This statistic is a standard measure in the literature on bunching (e.g., Chetty et al. 2011; Kleven and Waseem 2013). Relative to simpler measures that count the number of individuals below and above a threshold, it has the advantage of being less sensitive to differences across groups of individuals in the shape of the income distribution away from the threshold. For this reason, I would prefer to use this measure for all analyses but cannot do so for two reasons. First, I cannot use it in the model estimation because it is too computationally intensive. Second, there are an insufficient number of observations within each occupation to estimate it at the occupation-level. Nevertheless, Figure A15 shows that the qualitative patterns in Figure 6 hold using a simpler measure of bunching.

Figure 6. Variation in Bunching by Debt Balances and Age

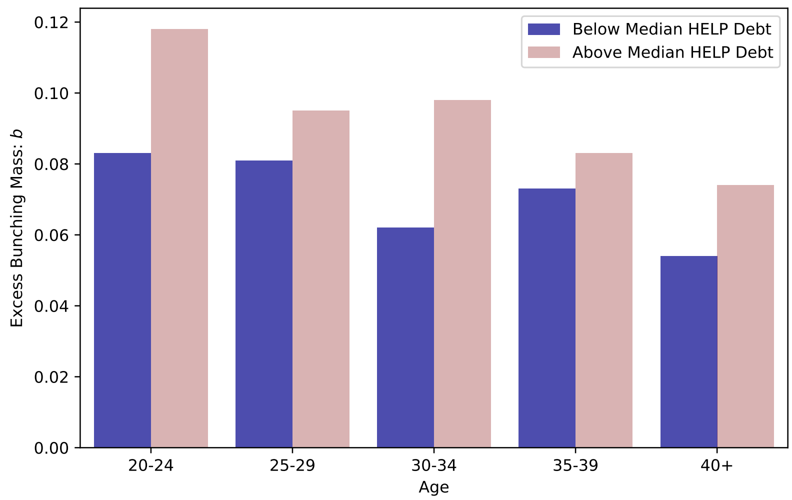

Notes: This figure plots the bunching statistic defined in (2) computed for different samples of debtholders based on age and debt balances. The age groups are listed on the horizontal axis. Within each age group, the blue (red) bars plot the estimated statistic for borrowers with below-median (above-median) debt balances, where the median is calculated separately for each year and age group. The calculation of b is detailed in Appendix C.4. Standard errors are omitted from this plot because the corresponding 95% confidence intervals overlap visually in the units of this plot. The sample is the *ALife* sample defined in Section 1.4 for the period between 2005 and 2018 after the policy change, restricted to individuals with positive HELP debt balances.

The amount of bunching below the repayment threshold also varies based on the properties of occupation-specific wage profiles. These wage profiles are plotted in Figure A9, which shows that there are some occupations in which the average individual will almost certainly earn enough income to pay her debt while there are others in which the average individual spends her entire life below the repayment threshold. Table A3 shows that the amount of bunching is larger in occupations with flatter income profiles and lower maximum incomes, both of which support the idea that a lower probability of eventual repayment increases borrowers' willingness to reduce their labor supply.

2.5 Bunching Decreases with Proxies for Liquidity

This section presents evidence that the responses in Figure 3 vary cross-sectionally with proxies for liquidity constraints. As discussed in Section 1.2, locating below the repayment threshold increases liquidity but has a smaller effect on wealth. Therefore, the evidence that borrowers reduce their labor supply to locate below the repayment threshold echoes the conclusion of Ganong and Noel (2020) that current budget constraints are important for understanding the behavior of indebted households. Absent direct measures of liquidity, I use several complementary measures to more directly assess its importance.

First, I find that the amount of bunching is larger among borrowers who reveal a preference for liquidity by holding less retirement wealth. The largest form of retirement savings in Australia is called superannuation ("super"), which is the second-largest source of household wealth (Australian Council of Social Service 2018). Contributions into super accounts primarily come from mandatory employer and voluntary employee contributions. Employee contributions, up to a limit, have generally been taxed at a rate lower than the personal income tax rate to incentivize saving. Therefore, super balances are a natural proxy for liquidity based on revealed preference: borrowers who are unwilling to contribute to a tax-advantaged but illiquid account are implicitly revealing a high valuation of liquidity (similar to Coyne et al. 2022). The left panel of Figure 7 plots the bunching statistic, b, based on quartiles of super balances from *ALife* that are defined within each year. The amount of bunching is highest for borrowers in the bottom quartile, approximately twice as large as the top quartile. This is also true within borrowers under age 30 (Figure A12).

Second, borrowers below the repayment threshold have larger housing payments. For most individuals, housing payments represent one of the largest sources of liquidity demand. Therefore, if liquidity influences labor supply responses, borrowers below the repayment threshold should have larger housing payments, or equivalently, borrowers with larger housing payments should be more likely to bunch below the repayment threshold. The right panel of Figure 7 shows that this prediction holds in the data: borrowers below the repayment threshold have larger housing payment–to–income ratios by approximately 2 percentage points, where housing payments are measured with combined mortgage and rent payments from the 2016 Census.

2.6 Additional Results and Discussion

Evasion. An obvious explanation for the bunching in Figure 3 is evasion, in which borrowers misreport their incomes. Although this is illegal and difficult to identify empirically, several facts (in addition to the direct evidence of a labor supply response in Figure 5 and the lack of evidence for evasion in Figure A7) suggest that it cannot explain all of the responses. First, Figure A13 shows that the distribution of salary and wages exhibits substantial bunching around the repayment threshold, which is generally interpreted as evidence of hours-worked responses (e.g., Chetty et al. 2013). This is because the literature on random audits finds that the majority of individual tax evasion comes from self-employment income, with an estimated noncompliance rate of less than 1% for items with withholding and substantial reporting information, such as salary and wages (Slemrod 2019). Second,

Figure 7. Bunching and Proxies for Liquidity Constraints

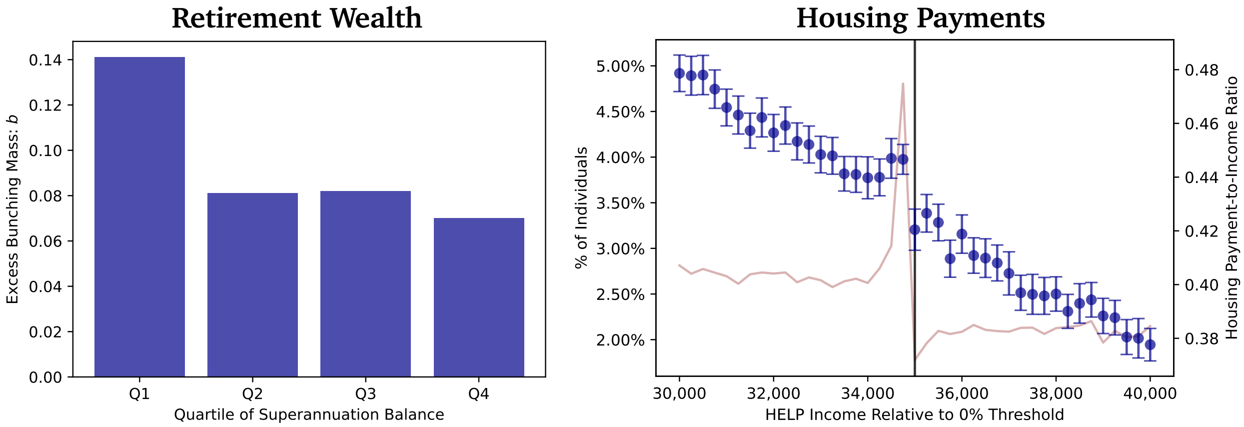

Notes: The left panel of this figure plots the bunching statistic defined in (2) computed for different samples of debtholders based on quartiles of superannuation balances computed within each year. The calculation of b is detailed in Appendix C.4. Standard errors are omitted because the corresponding 95% confidence intervals overlap visually in the units of this plot. The sample is the *ALife* sample defined in Section 1.4 between 2005 and 2018 after the policy change, restricted to individuals with positive HELP debt balances. The right panel replicates Figure 5 but plots the average housing payment–to–income ratios instead of hours worked within each bin. Error bars represent 95% confidence intervals.

Table A4 shows that the amount of bunching declines by only 4% when I restrict to the sample of wage-earners, who have substantially less flexibility in reporting their income, and is almost identical between borrowers who file their tax returns electronically and nonelectronically. When taxes are filed electronically, pure evasion is more difficult because the sources of labor income are often prefilled by the employer and, if they are not, the ATO compares what the individual reports with the employer's payment summary. Finally, the sample of borrowers near the repayment threshold is around median income, unlike the evidence from prior literature that evasion is largest among high-income individuals, who have more avoidance opportunities (Slemrod and Yitzhaki 2002; Saez et al. 2012).

Nevertheless, it is likely that some of the responses in Figure 3 reflect evasion rather than solely labor supply. In this case, the model I develop in Section 3 will overestimate labor supply responses to income-contingent repayment. There are two ways in which this could affect my normative results. First, if the costs of evasion are entirely real resource costs, then whether the responses in HELP income reflect labor supply or evasion is irrelevant as long as the model can replicate them (Feldstein 1999). However, in the more likely case that some of the costs of evasion are transfers to other agents or the government (e.g., fines), my model will overstate the welfare costs of the moral hazard created by income-contingent repayment (Chetty 2009; Gorodnichenko et al. 2009), reinforcing the qualitative conclusions from my normative analysis.

Other demographic heterogeneity. Table A4 examines heterogeneity in bunching based on the remaining demographic characteristics in the data. The results show little differences by gender, 5% less bunching among borrowers with a spouse, and 12% less bunching among borrowers with dependents. Although the first result contrasts with existing evidence that female labor supply is more elastic, an important caveat is that the responses that I estimate are local to the repayment threshold and thus do not capture extensive margin responses, which often drive the larger responses among women (Saez et al. 2012).

2.7 Taking Stock of Empirical Results

Summary of results. This section presents a series of empirical results that can be summarized as follows. First, borrowers reduce their income in response to income-contingent repayment. Second, these responses reflect, at least in part, labor supply responses rather than tax evasion or income-shifting, as borrowers below the repayment threshold work fewer hours and tend to be in occupations with more flexibility. Third, the size of these responses varies cross-sectionally based on two forces. The first is dynamics: borrowers with more debt and in occupations with lower income growth and maximum incomes, for whom the repayment reduction is more likely a permanent reduction rather than simply a transfer over time, exhibit greater responses. The second force is liquidity: borrowers who are likely liquidity-constrained, for whom the value of the repayment reduction is most valuable, are more willing to reduce their labor supply.

Implications for structural model. In Section 3, I develop a structural model that is designed to explain this evidence. Consistent with the bunching below the repayment threshold and the importance of dynamics and liquidity, borrowers choose their labor supply dynamically by trading off the disutility of work with the benefits of higher income, and they choose their consumption subject to borrowing constraints. However, the evidence in Figure 3 provides a rejection of a model in which labor supply is determined *solely* by the disutility of work and the benefits of higher income. Since utility increases in consumption and leisure, such a model cannot generate any borrowers immediately above the threshold because locating below it gives more consumption and leisure.[16] Motivated by the fact that bunching increases with hourly flexibility, the model in Section 3 introduces optimization frictions that prevent perfect adjustment of labor supply.

[16] One reason borrowers may locate above the repayment threshold is that, unlike a tax, income-contingent loans have an additional effect: increasing labor supply reduces the stock of debt. In Appendix B, I show that this is unlikely to explain the presence of borrowers above the threshold.

3 Life Cycle Model

The empirical analysis in Section 2 characterizes how labor supply responds to income-contingent repayment, but leaves open two important questions. First, how large are these responses quantitatively? Second, are these responses large enough to imply that the moral hazard created by income-contingent repayment outweighs the insurance benefits? The section presents and estimates a structural model designed to answer these two questions. The key ingredients in the model are endogenous labor supply, which creates moral hazard, uninsurable income risk, which creates a demand for insurance, and optimization frictions, which explain the presence of borrowers above the repayment threshold.

3.1 Model Description

3.1.1 Demographics

Time is discrete, and each period, t, corresponds to one calendar year. At $t = h \in \{\underline{h}, \underline{h}+1, \ldots \overline{h}\}$, a cohort h of individuals indexed by i are born at age a_0. The number of individuals is discrete and denoted by N, with a fraction μ_h born in cohort h. The initial age, a_0, should be interpreted as the age at which individuals exit college and enter the labor force. The age of an individual i in cohort h at time t is $a_{ht} = a_0 + t - h$. Before age a_T, individuals face age-dependent mortality risk, with the survival probability at age $a+1$ conditional on survival at age a denoted by m_a. Between ages a_0 and $a_R - 1$, individuals are in their working life and can supply labor to earn income. At age a_R, individuals transition to retirement and cannot supply labor; after age a_T, individuals die with probability one.

3.1.2 Preferences

During working life, individuals choose consumption, c, and labor supply, ℓ. An individual i at age a has preferences over consumption and labor supply that are time-separable with discount factor β and are expected utility with flow utility equal to:

$$\mathcal{U}_a(c_{ia}, \ell_{ia}) = \frac{n_a}{1-\gamma}\left(\frac{c_{ia}}{n_a} - \kappa\frac{\ell_{ia}^{1+\phi^{-1}}}{1+\phi^{-1}}\right)^{1-\gamma}. \tag{3}$$

In (3), γ is the coefficient of relative risk aversion (or, equivalently, the inverse of the EIS)[17], ϕ is the Frisch labor supply elasticity, κ is a scaling parameter, and n_a is an equivalence scale. The non-separability within-period follows Greenwood et al. (1988) (GHH) and eliminates wealth effects on labor supply, meaning the marginal rate of substitution between c and l is independent of c. This is consistent with empirical evidence that finds small labor supply responses to changes in wealth (Keane 2011; Cesarini et al. 2017; Gyöngyösi et al. 2022). The equivalence scale captures the evolution of household size over the life cycle, as in Lusardi et al. (2017). This generates a hump shape in consumption over the life cycle because the marginal utility of consumption increases with n_a and the calibrated values of n_a are hump shaped.

3.1.3 Labor Income Process

During working life, the labor income of individual i at age a, y_{ia}, is equal to the product of the individuals' wage rate, w_{ia}, and labor supply, ℓ_{ia}, where the latter is chosen endogenously. An individuals' wage rate is modeled in partial equilibrium and consists of three components:

$$\log w_{ia} = g_{ia} + \theta_{ia} + \epsilon_{ia}. \tag{4}$$

The first component, g_{ia}, is a deterministic life cycle component whose specific form is discussed later. The other two components, θ_a and ϵ_a, capture the stochastic components of the wage process, which take the following forms:

$$\begin{gathered} \theta_{ia} = \rho\theta_{ia-1} + \nu_{ia}, \quad \theta_{ia_0} = \delta_i, \\ \delta_i \sim \mathcal{N}(0, \sigma_i^2), \quad \nu_{ia} \sim \mathcal{N}(0, \sigma_\nu^2), \quad \epsilon_{ia} \sim \mathcal{N}(0, \sigma_\epsilon^2). \end{gathered} \tag{5}$$

The wage process in (5) incorporates idiosyncratic permanent and transitory shocks. The transitory component, ϵ_{ia}, is i.i.d. within and across individuals. The permanent component, θ_{ia}, depends on permanent shocks, ν_{ia}, which have persistence ρ, and an individual fixed effect, δ_i, which captures ex-ante heterogeneity across individuals. This specification for the wage process is similar to the permanent-transitory income processes used in canonical life cycle models (Gourinchas and Parker 2002). The key difference here is that income is endogenous, so these parameters cannot be estimated separately in a first-stage.

[17] In additional normative results, I use the rescursive generalization of these preferences from Guvenen (2009) to assess the role of risk and time preferences independently.

3.1.4 Education Levels

In addition to δ_i, individuals differ ex-ante based on their education levels. There are two education levels denoted by $\mathcal{E}_i \in \{0,1\}$, where

$$\mathcal{E}_i \sim \text{Bernoulli}(p_E). \tag{6}$$

Individuals' education level determines the deterministic component of their income process, g_{ia}, which takes the following form:

$$g_{ia} = \delta_0 + \delta_1 a + \delta_2 a^2 + \mathcal{E}_i \left(\delta_0^E + \delta_1^E a\right). \tag{7}$$

This specification captures that the returns to experience are quadratic (in logs), as in Mincer (1974), and that borrowers may have different wage levels and profiles.

Although education levels and borrowing (described below) are exogenous, heterogeneity in education levels is included in the model for two reasons. First, when I compare changing the structure of debt repayment contracts to changing the tax and transfer system, I need to account for the fact that the former only affects the college-educated, while the latter affects everyone. Second, the *ALife* panel is not long enough to separately identify the income process of the college-educated from the rest of the population.

3.1.5 Stochastic Fixed Cost of Labor Supply Adjustment

Individuals choose their labor supply at the same time that they choose consumption, which occurs at the end of each period after all shocks are realized. As discussed in Section 2.7, some type of adjustment friction is needed to generate borrowers above the repayment threshold. I introduce a fixed cost, f_{ia}, of choosing labor supply in the current period that is different from that in the past period, $\ell_{ia} \neq \ell_{ia-1}$. As in the "CalvoPlus" model of Nakamura and Steinsson (2010), this fixed cost is *stochastic* and evolves according to the following process:

$$f_{ia} = \left[\omega_{ia} f_L + (1-\omega_{ia}) f_H\right] \mathbf{1}_{a>a_0}, \quad \omega_{ia} \sim \text{Bernoulli}(\lambda), \quad f_L < f_H. \tag{8}$$

(8) allows the adjustment cost, f_{ia}, to vary over time between two values, f_L and f_H, with probabilities λ and $1-\lambda$, respectively. As a result, this specification nests two canonical models of imperfect adjustment. When $f_L = 0$ and $f_H = \infty$, it collapses to a Calvo (1983) model, which has been used in household finance to model mortgage refinancing (Andersen

et al. 2020). In contrast, when $\lambda = 1$, it corresponds to an (S, s) model, which has been used in many settings, such as portfolio choice (Abel et al. 2013), saving decisions (Choukhmane 2021), price-setting (Caplin and Spulber 1987), capital investment (Caballero and Engel 1999), and health insurance (Handel 2013). Finally, I model the fixed cost as a utility cost, as axiomatized by Masatlioglu and Ok (2005).

Modeling labor supply adjustment frictions using a stochastic fixed cost is reduced-form and warrants additional discussion. This choice is motivated by the evidence in Figure 4, which shows variation across occupations in labor supply responses. (8) is designed to capture this by allowing individuals to be in one of two "occupations" with different adjustment costs. However, the analogy between the different values of f_{ia} and occupations is incomplete in that f_{ia} is not associated with different wage processes. This restriction is made for tractability: allowing heterogeneity in wage profiles would make estimation infeasible because the wage process has to be jointly, as described in Section 3.2.

Ideally, the data would allow me to identify a more micro-founded model of adjustment frictions. Since this is not possible, my approach is to instead consider a reduced-form specification that allows for the two canonical types of imperfect adjustment, similar to Andersen et al. (2020).[18] State-dependent adjustment comes from the fixed costs that generate (S, s)–type behavior, in which individuals only adjust their labor supply only when the benefits of adjustment are sufficiently high. Economically, these costs could capture real costs associated with changing labor supply, such as wage reductions, or psychological costs, such as the hassle costs of adjusting a work schedule or search costs associated with changing jobs when hours are constrained by firms. However, adjustment in this model is also time-dependent in the sense that it depends on the realization of ω_{ia}. Economically, this can capture frictions on the demand-side of the labor market that result in the slow arrival of opportunities to adjust labor supply, as in models of job search à la Diamond-Mortensen-Pissarides or job transitions à la Kleven et al. (2023).[19]

The key concern with this reduced-form approach to modeling adjustment frictions is

[18]An alternative friction is optimization errors, which could take two forms, both inconsistent with the empirical evidence. The first is anticipated errors, in which individuals know that they cannot control labor supply perfectly. This, however, yields the prediction that there will be excess mass further to the left of the threshold as individuals reduce their labor supply even more to ensure that they do not end up above it, which is not the case in Figure 3. The second is unanticipated errors, where labor supply equals individuals' choice plus an error. This leads to the prediction that the bunching will be diffuse around the repayment threshold while the bunching in Figure 3 is sharp.

[19]Alternatively, $1 - \lambda$ could capture the fraction of inattentive individuals. However, this imperfectly captures inattention because agents are sophisticated about their inattention. Naive inattention introduces complications with violating budget constraints that are beyond the scope of this paper.

that the values of f_L, f_H, and λ that I estimate might not be policy-invariant when I study counterfactual repayment contracts. To address this concern, I explore how much these parameters would have to change in order to overturn the qualitative results from my counterfactuals in Section 4.3. However, the use of within-individual variation on a policy change that did occur in the data to estimate the model in Section 3.2 makes it more likely that the estimated parameters would be stable in the counterfactuals that I consider.

3.1.6 Liquid Assets

At age a_0, individuals are endowed with a stock of liquid assets, A_{ia_0}, where

$$A_{ia_0} \sim \begin{cases} 0, & \text{with probability } p_A(\mathcal{E}_i), \\ \text{Log-normal}\left(\mu_A(\mathcal{E}_i), \sigma_A(\mathcal{E}_i)^2\right), & \text{with probability } 1 - p_A(\mathcal{E}_i). \end{cases} \tag{9}$$

The dependence of this distribution on $\mathcal{E}_i$ allows for the possibility that initial liquidity varies with education levels. In subsequent periods, liquid asset balances at the end of the period at age $a-1$ are denoted by A_{ia}. Positive balances in the liquid asset pay a gross return of R. Individuals can also borrow using unsecured credit up to an age-dependent borrowing limit, $\underline{A}_a$.[20] The interest rate on borrowing is $R + \tau_b$, where τ_b captures the borrowing rate wedge. Asset income, i_{ia}, is received prior to consumption at age a and is equal to:

$$i_{ia} = r(A_{ia}) \times A_{ia}, \quad r(A_{ia}) = R - 1 + \tau_b \times \mathbf{1}_{A_{ia}<0}. \tag{10}$$

Both interest rates are taken as exogenous for tractability. This is unlikely to quantitatively affect the results because individuals with large debt balances, who are most affected by the policy changes that I consider, are young and hold a small share of aggregate wealth.

3.1.7 Student Debt

At age a_0, individuals are also endowed with debt balances, D_{ia_0}, where

$$D_{ia_0} \sim \begin{cases} 0, & \text{if } \mathcal{E}_i = 0, \\ \text{Log-normal}\left(\mu_d, \sigma_d^2\right), & \text{if } \mathcal{E}_i = 1. \end{cases} \tag{11}$$

[20] I do not allow agents to discharge this debt. Manso et al. (2024) show that doing so makes fixed repayment more attractive by limiting the reduction in labor supply created by the anticipation of the discharge.

These initial debt balances are exogenous because I focus on the trade-off between insurance and ex-post moral hazard. In subsequent periods, debt balances evolve according to:

$$D_{ia+1} = (1 + r_d)D_{ia} - d_{ia}, \quad d_{ia} = d(y_{ia}, i_{ia}, D_{ia}, a, t), \tag{12}$$

where r_d is the (net) interest rate on student debt and d_{ia} is the required debt payment determined by the repayment function, $d(\cdot)$. This function depends on borrowers' income and debt balance; any outstanding debt is discharged at $a = a_R$ or upon death.

3.1.8 Government

A government earns revenue from progressive taxes on labor and asset income, $\tau_{ia} = \tau(y_{ia}, i_{ia}, t)$, and student debt repayments. Expenditures include new student debt, D_{ia_0}, means-tested unemployment benefits, $ui_{ia} = ui(y_{ia}, i_{ia}, A_{ia})$, and a means-tested retirement pension, $\overline{y}_R(A_{ia})$. The government also pays a net consumption floor, $\underline{c}_{ia}$, to ensure that consumption exceeds zero. There is no deduction for interest paid on unsecured borrowing.

3.1.9 Recursive Formulation

Individuals solve a stochastic dynamic programming problem, which can be formulated recursively. There are five continuous states: A_{ia} = beginning-of-period liquid assets, ℓ_{ia-1} = past labor supply, D_{ia} = student debt, θ_{ia} = persistence component of wages, and ϵ_{ia} = transitory component of wages. There are four discrete states: t = current year, a = age, $\mathcal{E}_i$ = level of education, and f_{ia} = fixed cost. Denote $\mathbf{s}_{ia}$ as the vector of these state variables for individual i at age a and $\mathbf{E}_a(\cdot) = E(\cdot \mid \mathbf{s}_{ia+1})$ as the conditional expectation over the three shocks, ω_{ia+1}, ν_{ia+1}, and ϵ_{ia+1}. There are two controls: end-of-period liquid assets, A_{ia+1}, and labor supply, ℓ_{ia}. Consumption, c_{ia}, is pinned down by the budget constraint.

Suppressing i subscripts, individuals at age $a < a_R$ solve the following problem:

$$V_a(\mathbf{s}_a) = \max_{A_{a+1}, \ell_a} \left\{ \mathcal{U}_a \left(c_a - f \times \mathbf{1}_{\ell_a \neq \ell_{a-1}}, \ell_a \right) + \beta m_a \mathbf{E}_a V_{a+1}(\mathbf{s}_{a+1}) \right\}$$

subject to: (4), (5), (7), (8), (10), (12), and

$$c_a + A_{a+1} = y_a + A_a + i_a - d_a - \tau_a + ui_a + \underline{c}_a$$

constraints: $A_{a+1} \geq \underline{A}_{a+1}$ and $\ell_a \geq 0$

boundary conditions: (5), (6), (9), (11), and $\ell_{a_0-1} = \ell_{a_0}$

Retired individuals at age $a \geq a_R$ solve the following problem:

$$
\begin{aligned}
V_a(\mathbf{s}_a) = \max_{A_{a+1}} & \left\{ \mathcal{U}_a(c_a, 0) + \beta m_a \mathbf{E}_a V_{a+1}(\mathbf{s}_{a+1}) \right\} \\
& \text{subject to: (10), (12), and } c_a + A_{a+1} = \overline{y}_R(A_{ia}) + A_a + i_a - \tau(0, i_a, t) \\
& \text{constraint: } A_{a+1} \geq \underline{A}_{a+1} \\
& \text{boundary condition: } V_{a_T+1}(\mathbf{s}) = 0 \ \forall \mathbf{s}
\end{aligned}
$$

The model is solved using numerical dynamic programming; see Appendix D.1 for details.

3.2 Estimation Procedure

3.2.1 Calibrated Parameters

Table 2 shows the values of parameters that I calibrate directly using observed data, formulas from the Australian tax and transfer system, or prior literature. I provide a brief description of this calibration; see Appendix D.2 for additional details.

Demographics. Individuals are born at age 22 (the typical age at which students graduate from university in Australia), retire at age 65 (the age at which the Australian retirement pension began to be paid in 2004), and die with certainty after age 89. Prior to age 89, mortality risk is calibrated using Australia's life tables. Cohort-specific birth rates are calibrated to match the fraction of 22-year-olds in each year in *ALife*. I use data on household sizes from HILDA to compute equivalence scales as in Lusardi et al. (2017).

Interest rates and borrowing. There is no inflation in the model, and the numeraire is equal to \$1 AUD in 2005. When compared with the model, all empirical values are deflated to 2005 AUD using the HELP threshold indexation rate. The real interest rate is set to 1.84%, the (geometric) average deposit rate between 1991 and 2019 in Australia. The unsecured borrowing rate is set based on average credit card borrowing rates and age-specific borrowing limits are set based on credit card limits in HILDA. The real interest rate on student debt is set to zero, as in HELP.

Initial conditions. The distribution of initial assets is calibrated to match the liquid wealth distribution of individuals between ages 18 and 22. The fraction of borrowers, p_E, is equal to the fraction of 22-year-old individuals in *ALife* with positive debt. The distribution of initial debt is set based on the distribution among borrowers younger than age 26 in

ALife, the age by which most individuals have finished their undergraduate studies and debt balances peak in real terms.

Government taxes and transfers. Income and capital taxes are set to match the individual income tax schedules provided by the ATO in 2004 and 2005. Unemployment benefits are means-tested and calculated based on the Newstart Allowance, the primary form of government-provided income support in Australia for individuals above 22. The retirement pension is calculated following the Age Pension formula, the primary government-provided form of income-support for retirees in Australia. The age pension is available at age 65 and is means-tested based on assets and income.

Preference parameters. The preference parameter that I do not estimate due of a lack of identifying variation is the coefficient of relative risk aversion (RRA). I choose to set $\gamma = 2.23$ based on Choukhmane and de Silva (2023). In Section 4.3, I consider the effects of changing γ and the EIS independently using recursive Epstein–Zin preferences, which introduces a preference for timing of the resolution of uncertainty.

3.2.2 Simulated Minimum Distance Estimation

I estimate the remaining 15 parameters that cannot be calibrated directly, which I denote by Θ, using simulated minimum distance (SMD):

$$\Theta = \left(\underbrace{\phi \quad f_L \quad f_H \quad \lambda \quad \kappa \quad \beta}_{\text{preference parameters}} \quad \underbrace{\delta_0 \quad \delta_1 \quad \delta_2 \quad \delta_0^E \quad \delta_1^E}_{\text{wage profile parameters}} \quad \underbrace{\rho \quad \sigma_\nu \quad \sigma_\epsilon \quad \sigma_i}_{\text{wage risk parameters}} \right).$$

These parameters can be divided into three groups: preference parameters; parameters governing the age profile of wages, g_{ia}; parameters governing shocks to the wage process. In contrast to the standard approach for estimating life cycle models (e.g., Gourinchas and Parker 2002), I cannot estimate the latter two sets of parameters separately in a first stage because the income process is endogenous. I thus proceed by combining a standard set of estimation targets used to identify the latter two sets of parameters in models with exogenous income with the quasi-experimental variation from the HELP policy change.

Simulated policy change. I replicate the policy change in Figure 2 within the model by solving the model for two specifications of the student debt repayment function, $d(\cdot)$: the HELP 2004 schedule and the HELP 2005 schedule. Starting at $t = \underline{h} = 1963$, I simulate cohorts of individuals making choices under the 2004 schedule. At $t = T^* = 2005$, I then

Table 2. Values of Calibrated Model Parameters

Description	Parameter(s)	Values/Targets
Demographics		
Ages	$\{a_0, a_R, a_T\}$	$\{22, 65, 89\}$
Mortality rates	$\{m_a\}$	APA Life Tables
First and last cohorts	$\underline{h}$, $\overline{h}$	1963, 2019
Cohort birth probabilities	$\{\mu_h\}$	*ALife*
Equivalence scale	$\{n_a\}$	HILDA Household Size
Number of distinct individuals	N	1,600,000
Year of simulated policy change	T^*	2005
Assets		
Real interest rate	$R - 1$	1.84%
Unsecured borrowing wedge and limit	τ_b, $\{\underline{A}_a\}$	14.6%, HILDA Credit Card Limit
Probabilities of zero initial assets	$p_A(1)$, $p_A(0)$	0.197, 0.350
Distribution for $\log A_{ia_0}$	$\mu_A(1)$, $\mu_A(0)$,	7.42, 6.79
	$\sigma_A(1)$, $\sigma_A(0)$	1.72, 2.64
Student Debt		
Fraction of borrowers	p_E	0.308
Real interest rate on debt balances	r_d	0%
Distribution for $\log D_{ia_0}$	μ_d, σ_d	9.40, 0.86
Debt repayment function	$d(\cdot)$	HELP 2004 at $t < T^*$, HELP 2005 at $t \geq T^*$
Government		
Income and capital taxes	$\tau(\cdot)$	ATO Income Tax Formulas
Unemployment benefits	$ui(\cdot)$	ATO Newstart Allowance
Retirement pension	$\overline{y}_R(\cdot)$	ATO Age Pension
Net consumption floor	$\underline{c}$	\$40
Preference Parameters		
Relative risk aversion	γ	2.23

Notes: This table shows the parameters that are calibrated in a first-stage. See Appendix D.2 for additional details.

conduct a one-time unanticipated policy change in which all existing debtholders born at $t < T^*$ and subsequent debtholders start repaying under the 2005 schedule.

Estimator. I estimate Θ using SMD, which consists of choosing a set of estimation targets and a weighting matrix. Denote the empirical values of the estimation targets as $\hat{m}$, the vector of the estimation targets estimated in the model via simulation as $m(\Theta)$, and the weighting matrix as $W(\Theta)$. The estimate of Θ is then defined as Θ^*, where

$$\Theta^* = \arg\min_{\Theta} (\hat{m} - m(\Theta))' W(\Theta) (\hat{m} - m(\Theta)) .$$

I choose $W(\Theta)$ so the objective function is the sum of squared arc-sin deviations between $\hat{m}$ and $m(\Theta)$. The 44 estimation targets are listed in Appendix D.3 and discussed below.

3.2.3 Choice of Estimation Targets and Parameter Identification

This section discusses the identification of parameters in the SMD estimation. All parameters are jointly identified, but I choose the set of estimation targets so that each one is most sensitive to a subset of parameters. The discussion in this section is qualitative; Table A5 provides the elasticities of estimation targets with respect to parameters that support this discussion.

Labor supply elasticity, ϕ. The labor supply elasticity is primarily identified by bunching in the HELP income distribution below the repayment thresholds both before and after the policy change: a larger elasticity implies greater mass below these thresholds. To characterize this bunching, I use the distributions of HELP income among debtholders three years before and three years after the change. I pool distributions to minimize simulation error; Section 3.4 examines the model's fit in the two years surrounding the change. In estimation, I target the distribution within \$3,000 of the repayment thresholds so that these targets are primarily affected by the labor supply elasticity rather than wage profile parameters and use bins of \$500.

Lower adjustment cost, f_L. The lower value of the adjustment cost is primarily identified by the mass of the income distribution immediately *above* the repayment threshold. Since $f_L < f_H$, individuals that are marginal with respect to bunching below the threshold are more likely to have $f_{ia} = f_L$. Therefore, a higher value of f_L implies more marginal borrowers will choose not to bunch, increasing the mass of the income distribution above the threshold.

Adjustment cost probability, λ. Based on the income distribution alone, the probability of individuals receiving the lower adjustment cost, λ, is not separately identified from the cost itself, f_L: decreasing λ or increasing f_L would both generate more individuals above the repayment threshold. To separate these two parameters, I exploit an additional estimation target based on panel data: the probability of bunching in 2005 below the new repayment threshold conditional on bunching in 2004 below the old repayment threshold.[21] Because bunching in two subsequent periods is mostly concentrated among borrowers that have $f_{ia} = f_L$ in both periods, λ controls the likelihood of this happening. A higher value of lambda, therefore, implies greater persistence in bunching below the repayment threshold.

Upper adjustment cost, f_H. Identifying the upper adjustment cost, f_H, is more chal-

[21] I'm grateful to an anonymous referee for suggesting this identification strategy.

lenging because individuals with $f_{ia} = f_H$ are further from their indifference condition for bunching below the repayment threshold. This implies that changes in labor supply in response to the threshold are not very sensitive to f_H, once f_H is sufficiently larger than f_L. Therefore, to identify this parameter, I need to exploit information on labor supply adjustments made for reasons other reasons. In the model, these other adjustments are primarily due to wage fluctuations. To capture the distribution of these adjustments, I target the kurtosis of changes in (log) hours worked in the data and $\log \ell_{ia}$ in the model. Kurtosis is increasing in f_H because a higher value of f_H implies longer gaps between adjustments and hence a larger difference between the current and desired ℓ_{ia}.

In principle, I could target any (higher-order) moment of the distribution of hours worked, but I choose kurtosis for two reasons. First, kurtosis has the benefit of being a scale-free statistic that describes the peakedness of a distribution. Second and more importantly, Alvarez et al. (2016) show that kurtosis is the key moment of the distribution of price changes that determines the real effects of monetary policy shocks in models with sticky prices. This result comes from the fact that kurtosis embodies the amount of selection into adjustment à la Golosov and Lucas (2007), meaning the difference between the individuals (or firms in their case) who adjust and the full population. Since this same type of selection is important for the policy counterfactuals I consider, targeting kurtosis is a natural way to make sure the model accurately captures it.

To measure the kurtosis of changes in hours worked, I use annual data from HILDA. These data have the downside of being a survey, but they are the only available source of panel data on hours worked. Therefore, I follow Heathcote et al. (2014) and allow for measurement error in hours worked, which I take to be multiplicative in levels and normally distributed with mean one and variance ι^2.[22] To identify the measurement error parameter, I add as an estimation target the probability of zero adjustment in hours worked.

Time discount factor, β. To identify the time discount factor, I leverage the dynamic incentives created by an income-contingent loan. As discussed in Section 1.2, an income-contingent loan differs from a tax because reducing payments today leads to higher future payments when the debt is repaid. These dynamic incentives are larger for individuals with less debt, for whom the probability of repayment is higher. The extent to which individuals respond to these future incentives depends on their time preferences, which are controlled (partially) by β.[23] Therefore, I identify β by targeting heterogeneity in bunching with debt

[22] Allowing for such measurement error is also standard in the literature on durable consumption, where a similar combination of survey data and models with imperfect adjustment are used (Berger and Vavra 2015).

[23] To see this intuition more formally, consider a two-period model where borrowers have linear utility,

balances, where bunching is measured using the ratio of borrowers below to above the 2005 threshold after the policy change. I measure heterogeneity by taking the ratio of this measure for individuals in the top and bottom quartiles of debt balances within each year.

Scaling parameter, κ. This parameter is identified by the average value of ℓ_{ia}. A higher value increases the disutility of labor supply and thus lowers average values of ℓ_{ia}.

Wage profile parameters, δ_0, δ_1, δ_2, δ_0^E, and δ_1^E. These parameters are primarily identified by the regressions of log income onto polynomials in age and an education-level indicator. If labor supply were exogenous, they could be estimated separately with these estimation targets alone. However, with endogenous labor supply, these parameters control the *wage* rather than the income process and must be estimated jointly because the former is not observable.

Wage risk parameters, ρ, σ_ν, σ_ϵ, and σ_i. These parameters are identified by how the cross-sectional variance of log income varies with age and the percentiles of income growth at one-year and five-year horizons. This set of moments is standard in the literature used to estimate exogenous income processes (e.g., Guvenen et al. 2022), and the identification is similar here even though the income process is endogenous. The cross-sectional variance at age 22 identifies σ_i, the variance of the initial permanent income. The extent to which the cross-sectional variance increases with age identifies the persistence of income shocks, ρ: more persistent shocks generate a greater increase in variance over the life cycle (Deaton and Paxson 1994). The sum of the variances of permanent and transitory income shocks, σ_ν and σ_ϵ, are identified by the level of this cross-sectional variance at later ages. These two variances are then separated using the percentiles of income growth: a larger variance of permanent shocks, σ_ν, delivers fatter tails in 5-year than in 1-year income growth.

3.3 Estimation Results and Model Fit

Table 3 shows the results from estimating five different models, where the final column corresponds to the baseline model. Column (1) starts by estimating a model without labor supply adjustment frictions. This model does not generate any individuals locating above the repayment threshold, and delivers an unrealistically low estimate of the (Frisch) labor supply elasticity of $\phi = 0.003$. Column (2) estimates a model with a constant fixed adjustment

---

discount at rate $\beta = \frac{1}{1+r}$, and have a probability of repayment p in the second period. The value of locating below the 2005 repayment threshold is then $\$1400 \times \beta \times \left(\beta^{-1} - p\right)$. Differentiating with respect to p shows that this value depends on p if and only if $\beta \neq 0$, and increases in sensitivity to p as β increases.

cost, corresponding to a standard (S, s) model. Adding this fixed cost helps generate a more reasonable estimate of ϕ, and the estimated value of the fixed cost is approximately 2% of average earnings. Column (3) estimates a model in which the adjustment cost is either zero or infinity (i.e., a Calvo (1983)–style model). Compared to column (1), the estimate of ϕ is more reasonable, and the estimated value of λ implies individuals receive the opportunity to adjust their labor supply every 8-9 periods.

Column (4) shows the results from estimating a model that combines those in columns (2) and (3). This estimation attributes the lack of adjustment by individuals above the repayment threshold mostly to the adjustment cost shock, given the estimated value of λ is similar to that in column (3). This finding implies that labor supply adjustment appears to be more time- rather than state-dependent, similar to the findings of Andersen et al. (2020) in the context of mortgage refinancing. However, as illustrated by the estimate of ϕ, allowing $f_L > 0$ is important for the getting the right estimate of ϕ.

The results for the baseline model are reported in column (5). The estimate of the labor supply elasticity is 0.15. Appendix A shows this estimate is close to the median value of 0.14 for Frisch and Hicksian intensive margin elasticities reported in Keane (2011) and Chetty et al. (2012), and further discusses how my results relate to existing literature.[24] I estimate a fixed costs of $f_L = \$378$ and $f_H = \$3191$, approximately 0.6% and 5% of average income. Given the estimate of $\lambda = 0.15$, individuals receive the lower adjustment cost that allows them to adjust their labor supply more cheaply every 6-7 years. Comparing columns (4) and (5) shows that allowing $f_H < \infty$ does not have a major effect on other estimated parameters. This is consistent with the discussion in Section 3.2.3, which argues that f_H is primarily identified based on the distribution of changes in labor supply that is not targeted in the first four columns.

The baseline model provides a close fit to the bunching used to identify the key labor supply parameters. Figure 8 shows the model fits the distribution of HELP income before and after the policy change, including the mass of borrowers immediately below and above the repayment threshold. There are some differences in the shape of the distributions because the estimation is balancing improving this fit with matching the age profile of income. Table 4 shows that the model is also able to replicate the probability that individuals who bunch below the old repayment threshold in 2004 also bunch below the new repayment threshold in 2005, which is how the adjustment probability, λ, is identified. The model's

[24] Because I identify ϕ using bunching in HELP income, it should be interpreted as a reported income elasticity that aggregates both hours and non-hours responses (Feldstein 1999). Therefore, Appendix A compares my estimate of ϕ to existing estimates of both hours and taxable income elasticities.

Table 3. Simulated Minimum Distance Estimation Results

Parameter		Estimation (1)	(2)	(3)	(4)	(5)
Labor supply elasticity	ϕ	0.003	0.167	0.084	0.146	0.149
		(.000)	(.001)	(.001)	(.001)	(.001)
Lower adjustment cost	f_L	\$0	\$1377	\$0	\$454	\$378
		·	(\$6)	·	(\$9)	(\$16)
Adjustment cost probability	λ	1	1	0.124	0.161	0.153
		·	·	(.002)	(.002)	(.004)
Upper adjustment cost	f_H	∞	∞	∞	∞	\$3191
		·	·	·	·	(\$105)
Time discount factor	β	0.998	0.914	0.934	0.958	0.937
		(.000)	(.001)	(.003)	(.001)	(.001)
Scaling parameter	κ	0.179	1.233	0.236	0.697	2.667
		(.000)	(.007)	(.001)	(.006)	(.032)
Wage profile parameters	δ_0	10.170	9.360	9.089	9.243	9.667
		(.002)	(.004)	(.004)	(.004)	(.003)
	δ_1	0.067	0.074	0.073	0.078	0.064
		(.000)	(.000)	(.000)	(.000)	(.000)
	δ_2	−0.001	−0.001	−0.001	−0.001	−0.001
		(.000)	(.000)	(.000)	(.000)	(.000)
	δ_0^E	−0.442	−0.440	−0.480	−0.496	−0.473
		(.000)	(.001)	(.001)	(.001)	(.001)
	δ_1^E	0.025	0.019	0.022	0.021	0.019
		(.000)	(.000)	(.000)	(.000)	(.000)
Persistence of permanent shock	ρ	0.824	0.927	0.922	0.934	0.929
		(.000)	(.000)	(.000)	(.000)	(.000)
Std. deviation of permanent shock	σ_ν	0.057	0.223	0.252	0.222	0.224
		(.000)	(.000)	(.001)	(.001)	(.001)
Std. deviation of transitory shock	σ_ϵ	0.431	0.133	0.113	0.164	0.150
		(.000)	(.001)	(.001)	(.001)	(.001)
Std. deviation of individual FE	σ_i	0.575	0.569	0.541	0.591	0.569
		(.001)	(.001)	(.002)	(.002)	(.002)
Std. deviation of measurement error	ι	0	0	0	0	0.034
		·	·	·	·	(.000)

Notes: This table shows the results from simulated minimum distance (SMD) estimations. Each column corresponds to a separate estimation. Entries in the table are parameter estimates with standard errors below in parentheses. Parameters that are fixed and not estimated are indicated with "·" in place of a standard error. All estimations use the same set of estimation targets described in Appendix D.3, except for column (5) which uses the two additional targets described in Section 3.2.3 to identify f_H and ι.

ability to match the heterogeneity in bunching with debt balances is driven by the annual discount factor, β, which is estimated at 0.94. This estimate is between higher estimates from targeting consumption data (e.g., Gourinchas and Parker 2002) and lower estimates from targeting wealth accumulation and portfolio choices (e.g., Catherine 2022). The fact that this estimate is less than R^{-1} implies that individuals face a trade-off between wanting to consume at young ages due to impatience and accumulating precautionary savings, generating buffer-stock behavior (Carroll and Kimball 1996).

Figure 8. Model Fit: HELP Income Distribution around the Policy Change

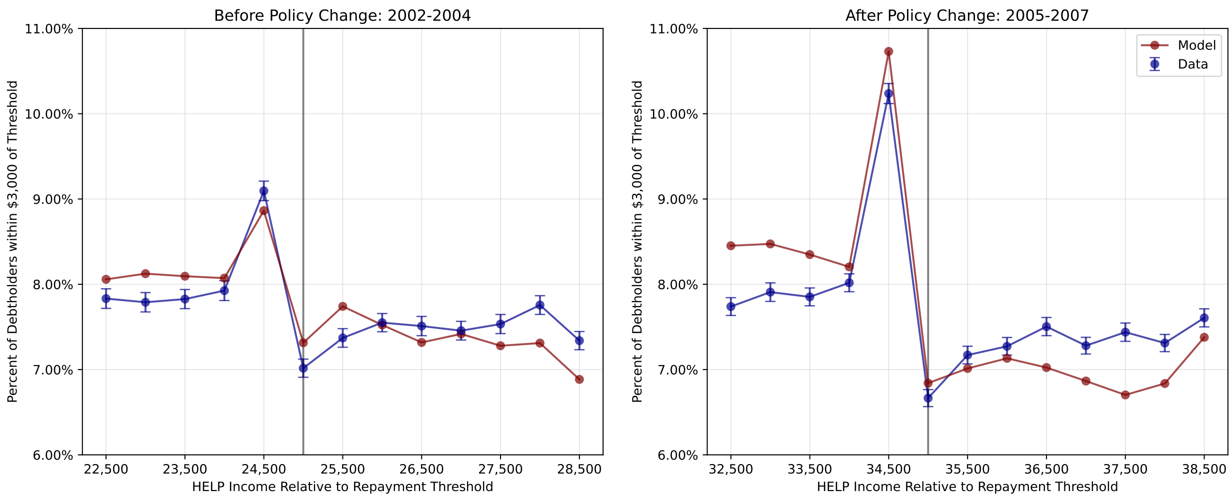

Notes: The left panel of this figure plots the HELP income distribution within $3,000 of the repayment threshold in bins of $500 for the period before the policy change from 2002 to 2004 in the data in blue. Bars represent 95% confidence intervals based on bootstrapped standard errors with 1000 iterations. The red line plots the same quantities from the model with parameters set at the estimated values in column (1) of Table 3. The right panel replicates the left panel for the period after the policy change between 2005 and 2007. The vertical gray line in each plot indicates the repayment threshold, which is the point at which repayment begins.

Table 4 shows the model provides a reasonably good fit to the remaining estimation targets. Through adjusting the upper adjustment cost, f_H, the model matches the distribution of changes in hours worked. The model can also replicate the age profile of labor income, which are most affected by the wage profile parameters. The fit is not perfect because income in the model is endogenous: if the age profile of labor supply varies over the life cycle for reasons outside the model, it will be unable to match these income profiles. The cross-sectional variance of income increases over the life cycle, and the model can replicate this pattern due to the high persistence of permanent shocks, $\rho = 0.93$.

3.4 Model Validation on Nontargeted Evidence

Before using the estimated model to perform counterfactual analyses, I show that it provides a reasonable fit to several pieces of evidence that were not targeted in estimation. The first set is the heterogeneity in bunching by debt balances and age around the repayment threshold shown in Section 2. The left panel of Figure 9 shows a scatterplot of the bunching for different groups based on age and debt balances in the data versus the model. The relationship between these points is positive with a slope coefficient of 1.35, indicating the model does a good job at qualitatively capturing this heterogeneity.

The right panel of Figure 9 shows that the model can also replicate responses to income

Table 4. Model Fit: Other Estimation Targets

	Data	Model
Cross-Sectional Variance of Log Labor Income at Age 22	0.453	0.448
Cross-Sectional Variance of Log Labor Income at Age 32	0.555	0.470
Cross-Sectional Variance of Log Labor Income at Age 42	0.577	0.503
Cross-Sectional Variance of Log Labor Income at Age 52	0.539	0.568
Cross-Sectional Variance of Log Labor Income at Age 62	0.608	0.665
Linear Age Profile Term	0.077	0.071
Quadratic Age Profile Term	−0.001	−0.001
Education Income Premium Constant	−0.574	−0.559
Education Income Premium Slope	0.023	0.022
10th Percentile of 1-Year Labor Income Growth	−0.387	−0.407
10th Percentile of 5-Year Labor Income Growth	−0.667	−0.702
90th Percentile of 1-Year Labor Income Growth	0.415	0.407
90th Percentile of 5-Year Labor Income Growth	0.698	0.706
Average Labor Supply	1.000	0.813
Probability that Labor Supply Not Adjusted	0.422	0.375
Kurtosis of Changes in Log Hours	5.637	5.721
Bunching Ratio: Q4 Debt to Q1 Debt	1.173	1.222
Bunching Probability in 2005 Conditional on Bunching in 2004	0.020	0.020

Notes: This table shows the value of the remaining estimation targets not shown in Figure 8 in the data and the model with parameters set at the estimated values in column (5) of Table 3.

taxes. This panel plots the bunching around the two discontinuities in marginal tax rates closest to the HELP repayment thresholds. The bunching around these tax "kinks" is smaller than around the repayment thresholds because these kinks induce a change in marginal rather than average rates. The model replicates this relatively small amount of bunching at these thresholds reasonably well.

Next, I assess whether the model can match the speed of the response to the policy change. As shown in Figure 1, the bunching in the data around the repayment thresholds responds rapidly to the change. A priori, it is not clear whether the model can generate an immediate response because the estimation targets income distributions that are pooled across the three years before and after the change. Nevertheless, Figure 10 shows that the immediate response in the data is present in the model. This may appear surprising given that some individuals who were bunching in 2004 receive $f_{ia} = f_H$ in 2005, and thus are less likely to adjust ℓ_{ia}. However, the bunching still disappears because $y_{ia} = w_{ia}\ell_{ia}$, which implies that y_{ia} can change even if ℓ_{ia} does not due to fluctuations in w_{ia}.

In addition to matching cross-sectional variation in income around the policy change, Figure 11 shows that the model does a reasonable job of capturing within-individual variation. The left panel shows that the model matches the probability that individuals

Figure 9. Fit of Model on Nontargeted Bunching Statistics

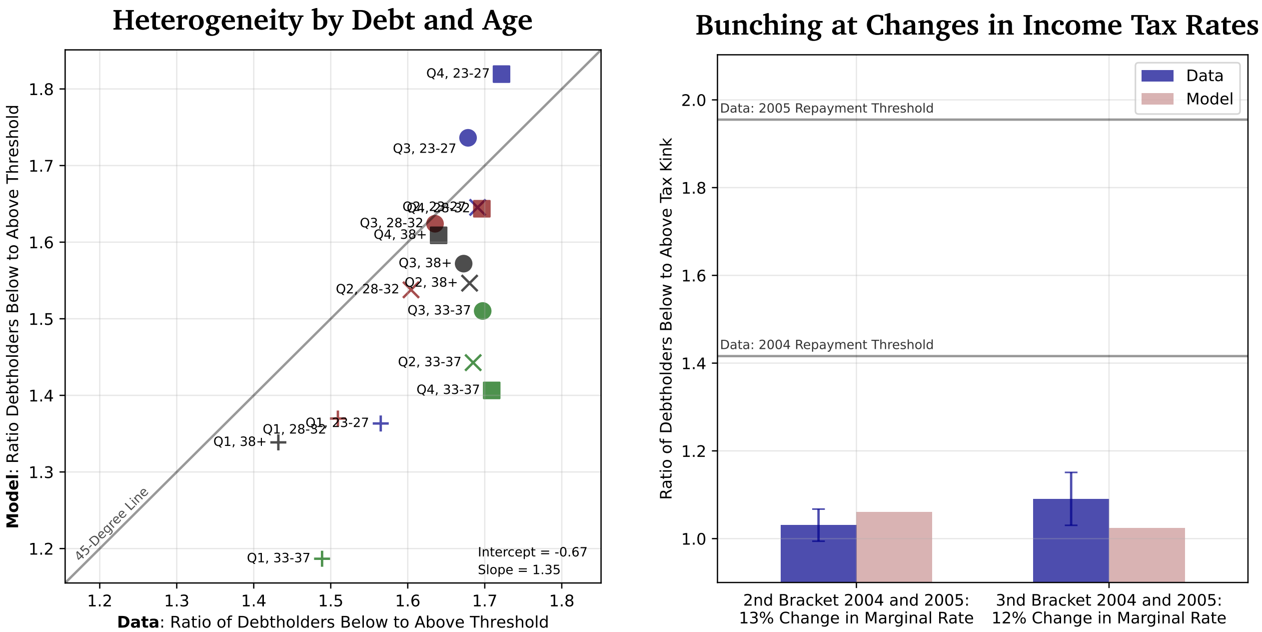

Notes: The left panel of this figure shows a scatterplot of bunching below the 2005 repayment threshold for different samples in the data versus the model. Each point corresponds to a different sample based on quartiles of debt and age labeled in the plot. The quartiles of debt are calculated in the data after taking out year fixed effects and adjusting for inflation. These same quartiles are used in the model. Each age group is plotted in a different color, and each quartile of debt has a differently shaped marker on the plot. For each sample, bunching is measured as the ratio of the number of debtholders with \$500 below to \$500 above different thresholds. The right panel shows the ratio of the number of debtholders with \$250 below to \$250 above different thresholds computed around two points with changes in marginal income tax rates in 2004 and 2005 using taxable income instead of HELP income in the data (there is no difference in the model). This panel also contains two horizonal lines on the values of the same bunching statistics computed around the two repayment thresholds for reference. Tax brackets are fixed in nominal terms, so when pooling 2004 and 2005, I adjust the thresholds and income using the HELP threshold indexation rate. Data values are presented in blue with 95% confidence intervals based on bootstrapped standard errors with 1000 iterations. Model values are presented in red. The sample is the *ALife* sample defined in Section 1.4 between 2005 and 2018, restricted to debtholders between 23 and 64. I impose the same sample filters in the model.

Figure 10. Fit of Model in Years Surrounding Policy Change

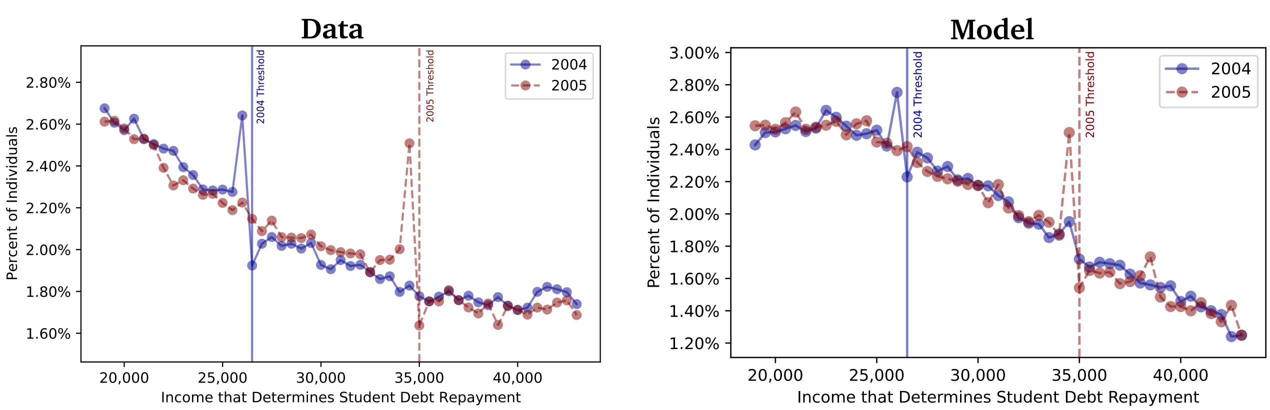

Notes: The left panel of this figure reproduces Figure 1. The right panel makes the analogous plot based on simulations from the baseline model with parameters set to the values in column (5) of Table 3.

who are bunching prior to the policy change are bunching after the policy change, which was targeted in the estimation, but also the probability that these individuals remain

Figure 11. Fit of Model on Within-Individual Moments around Policy Change

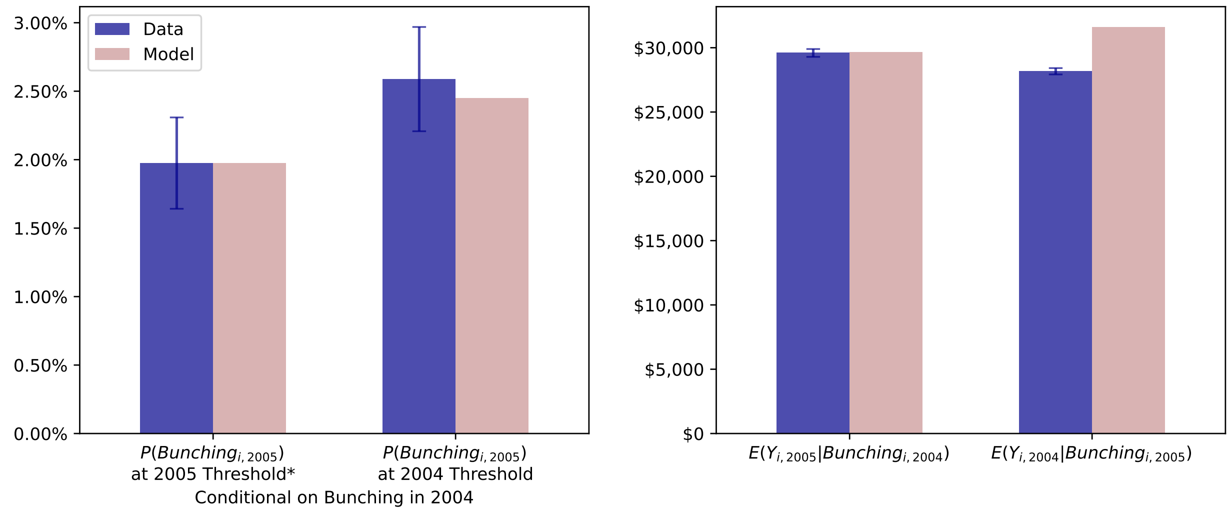

Notes: This figure shows how the model compares to the data on some panel-based moments in the years surrounding the policy change. In both panels, bunching is defined as individuals who are with $500 of the relevant threshold hold. The left panel restricts to individuals who are bunching below the 2004 repayment threshold in 2004, and plots two statistics: (i) the probability that they are bunching below the new 2005 repayment threshold in 2005, after the policy change; (ii) the probability that they remaining bunching below the old 2004 repayment threshold in 2005, after the policy change. The right panel plots two statistics: (i) the average income in 2005, after the policy change, of individuals that were bunching below the 2004 repayment threshold before the policy change; (ii) the average income in 2004, before the policy change, of individuals that were bunching below the 2005 repayment threshold after the policy change. Data values are presented in blue with 95% confidence intervals; model values are presented in red. The * in the left panel indicates that this moment was targeted in estimation; all other moments were not targeted.

bunching below the old threshold, which was not targeted. The right panel shows the model matches the average income after the change of individuals who are bunching before it. However, the model misses on the average income before the change of individuals who are bunching after it: in the model, these individuals tend to come from further up in the income disribution than in the data.

The final set of nontargeted evidence that I use to validate the model comes from Britton and Gruber (2020) (BG), who study taxable income responses to income-contingent loans in the UK. In the UK, there is a single government-provided income-contingent loan with one repayment rate that determines the marginal rate of income repaid above a repayment threshold. The left panel of Figure 12 reproduces Figure 5 from BG, which shows the income distribution for a 10% sample of debtholders between 2006 and 2012 when the threshold was £15,000 and the repayment rate was 9%. The conclusion from BG is that there is evidence of bunching below this threshold, but that the implied elasticity of taxable income (ETI) (estimated from the static model in Saez 2010) is very low: Figure 12 reproduces the 95% confidence interval of their estimate, which includes zero.

The right panel of Figure 12 shows that the baseline model replicates this conclusion from BG. To generate this figure, I change the debt repayment parameters in the model to those of the UK income-contingent loan. I then simulate from the model holding all structural parameters fixed at their estimated values in column (5) of Table 3. Like in BG, the model generates bunching immediately below this threshold, but the ETI implied by this bunching is essentially zero.[25] One difference is that the model cannot replicate the bunching at round numbers in BG. I have chosen not to add an ingredient to the model to match this fact because it is not present in my data. In BG, when the threshold is a round number, there is bunching in the income distribution of *non-borrowers*. In contrast, the income distribution of non-borrowers shows zero bunching around the HELP threshold in 2005, when it was a round number: see Figure A4.

It is important to emphasize that the lack of significant bunching in Figure 12 does not imply that labor supply responses are unimportant. As shown in Section 4, the labor supply responses to an income-contingent loan with a marginal repayment rate, like the one in the UK, that occur *away* from the threshold account for the majority of the fiscal cost of moving to income-contingent repayment and significantly change optimal contract design. This highlights the value of the empirical evidence in Section 2 relative to BG: because the incentives created by HELP are large enough to generate responses, it allows me to estimate a dynamic model of labor supply that captures the effects of income-contingent repayment both at and further from the threshold, which can then be used for counterfactuals. In contrast, the evidence from BG alone does not say whether the lack of bunching is driven by a low structural elasticity, the dynamic incentives created by income-contingent repayment, or optimization frictions. As with any structural model, the key assumption required is that the parameters are stable in the counterfactuals of interest. However, the model's ability to replicate the evidence in BG makes this a more plausible assumption.

3.5 Model-Based Decomposition of Bunching

As discussed in Section 1.2, the present value of bunching below the repayment threshold can be much different than the change in current repayments. To illustrate, assume that borrowers value repayments in two periods and discount cash flows in the second period with (net) interest rate r. Letting p denote the probability of repayment in the second

[25] There are some differences in the shape of the distribution that likely come from the fact that the wage process was estimated on Australian data.

Figure 12. Bunching around UK Income-Contingent Repayment Threshold

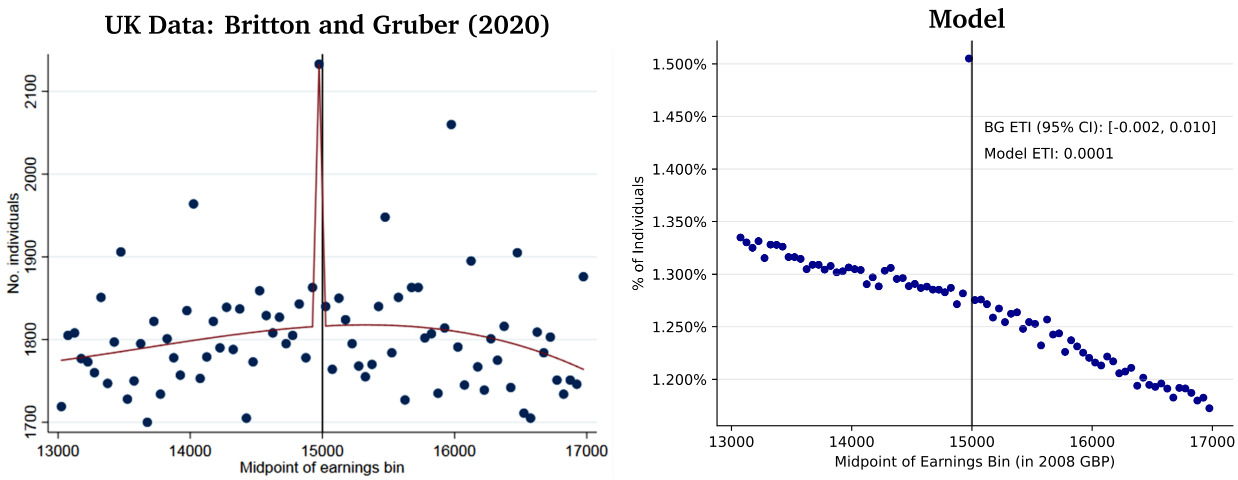

Notes: The left panel of this figure reproduces Figure 5 from Britton and Gruber (2020). This figure shows the income distribution in bins of £50 of student debtholders in years 2006-2012 around the £15,000 repayment threshold, at which the marginal repayment rate changes from 0% to 9% of taxable income. The sample is a 10% random sample of all students; see Britton and Gruber (2020) for additional details. The right panel shows the income distribution for debtholders generated by the model at the parameter values in column (5) of Table 3. To generate this plot, I change the debt repayment function in the model to be an income-contingent loan with a 9% marginal rate above $30,421 AUD, which corresponds to converting £15,000 from 2008 GBP to 2005 GBP using the CPI, and then adjusting to 2005 AUD using the exchange rate of 2.2 AUD/GBP, and an interest rate of $r_d = 1\%$, as is the case in the UK during this time period. The marginal tax rate at the repayment threshold in the model is 30% compared to 33% and 31% in the UK over this time period. The elasticity of taxable income (ETI) shown in the right panel for "BG" corresponds to the 95% confidence interval from Table 6 in Britton and Gruber (2020). The estimate for "Model" corresponds to apply the exact same approach on the model-generated data, adjusting for the differences in marginal income tax rate.

period, the net present value (NPV) of locating below the 2005 repayment threshold is

$$\underbrace{\$1400 \times \frac{r + (1-p)}{1+r}}_{\text{NPV gain from bunching}} \leq \underbrace{\$1400,}_{\text{liquidity gain from bunching}} \tag{13}$$

which is (weakly) smaller than the increase in liquidity.[26]

Motivated by (13), Figure 13 uses the estimated model to decompose the bunching below the 2005 repayment threshold into three distinct effects. The first effect is the bunching that arises from the difference between borrowers' discount rate and the debt interest rate (i.e., $r \neq 0$), which increases the NPV of bunching below the repayment threshold. Figure 13 shows that this has a negligible effect on the bunching below the repayment threshold. The second effect is that, even when $r = 0$, bunching below the repayment threshold has a positive NPV if borrowers do not anticipate repaying their debt (i.e., $p < 1$). The results in Figure 13 show that this channel accounts for the majority of the bunching: in a counterfactual where $p \approx 1$, bunching below the repayment threshold decreases by

[26]Deriving this expression requires two additional assumptions: (i) p is independent of bunching in the first period; (ii) the interest rate on outstanding debt is zero.

Figure 13. Decomposition of Bunching Below the Repayment Threshold

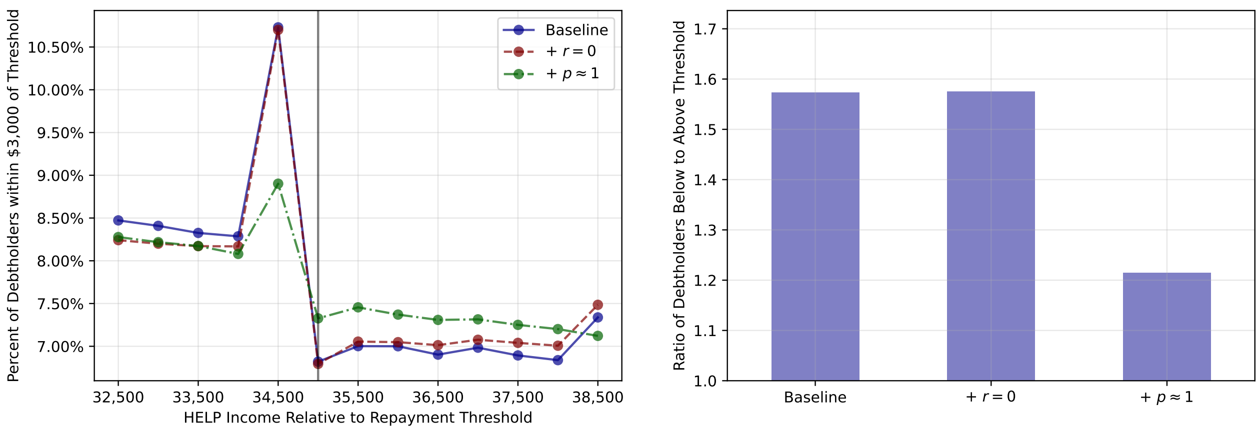

Notes: The left panel of this figure plots the income distribution in bins of $500 around the 2005 repayment threshold between 2005 and 2018 in three different models, and the right panel plots the ratio of the number of debtholders with $500 below to $500 above different thresholds. The first model, Baseline, corresponds to the baseline model estimated in column (5) of Table 3 after the calibrated value of R in Table 2 is replaced with β^{-1}. The second model, $+r = 0$, corresponds to additionally setting $r_d = \beta^{-1} - 1$ in the first model. The third model, $+p \approx 1$, corresponds to taking the second model and setting $D_0 = 4\% * \$35000 = \$1400$ for all borrowers. Then, for each year in which borrowers have debt in the second model, borrowers' debt balances in the third model are unanticipatedly reset to $1400, regardless of whether they paid it off in the prior period. The purpose of this third model is to (approximately) make borrowers anticipate repayment with probability one, while ensuring the set of borrowers who have positive debt balances in each year are the same as in the second model.

about 65%.[27] This model-based inference is consistent with the empirical evidence in Section 2 that the amount of bunching is larger among borrowers with a lower probability of repayment.

The remaining 35% of the bunching that remains even when $r = 0$ and $p \approx 1$ in Figure 13 reveals that a third effect is a quantitatively important driver of bunching below the repayment threshold: a demand for liquidity. When $r = 0$ and $p = 1$, the NPV of locating below the repayment threshold is zero. Nevertheless, locating below the repayment threshold still increases borrowers' current liquidity, which they may value if they are liquidity-constrained. This importance of liquidity is empirically supported by evidence in Section 2 that the amount of bunching increases with proxies for liquidity constraints, complementing evidence that a demand for liquidity created by incomplete markets amplifies the moral hazard created by other forms of social insurance (Chetty 2008; Ganong and Noel 2023; Indarte 2023). Additionally, it illustrates an important way in which the incentives created by income-contingent repayment differ from those of an income tax. Because most borrowers anticipate repaying their debt with some probability, the labor supply response created by an income-contingent loan is larger than that of a tax for a given repayment function.

[27] Formally, this counterfactual does not correspond precisely to setting $p = 1$ because p is an endogenous object in this dynamic model. Therefore, these results are a lower bound on the effect of p.

4 Normative Analysis of Income-Contingent Loans

This section uses the estimated model to study the welfare and fiscal impacts of moving from fixed repayment contracts to income-contingent loans. My analyses proceed in two steps. First, I study the effects of moving from fixed repayment to different forms of existing income-contingent loans. Because these contracts have different fiscal costs, I assess policies based on their marginal value of public funds (MVPF), as in Hendren and Sprung-Keyser (2020) and Finkelstein and Hendren (2020). Next, I solve a Ramsey (1927)–style problem to construct income-contingent contracts that maximize borrower welfare while raising the same revenue as a fixed repayment contract. By construction, the second step compares contracts that have the same fiscal cost. Throughout these analyses, borrowing, education choices, and prices (i.e., wages and interest rates) are held fixed. Therefore, these results are informative about the effects of a mandatory debt restructuring among existing borrowers whose ex-ante choices are fixed by definition.[28]

Details of policy environment. Before turning to the analyses, I describe a few additional details of the policy environment; in Section 4.3, I study how changing these features affects my results. The comparison of different repayment contracts is contingent on the tax system, $\tau_{ia} = \tau(y_{ia})$, which also provides insurance to borrowers and redistributes. I adopt the parametric income tax specification from Heathcote et al. (2017) calibrated to Australia's tax schedule; see Appendix D.2 for additional details. I then define the government budget, $\mathcal{G}$, as the expected discounted value of debt repayments and taxes net of transfers and debt issuance,

$$\mathcal{G} \equiv \mathbf{E}_0 \left(\sum_{a=a_0}^{a_T} \underbrace{\frac{\tau_{ia} - ui_{ia} - \underline{c}_{ia}}{R^{a-a_0}}}_{\text{taxes and transfers}} + \underbrace{\frac{d_{ia}}{R^{a-a_0}} - D_{ia_0}}_{\text{debt repayments}} \right), \tag{14}$$

where $\mathbf{E}_0(\cdot)$ denotes an expectation taken over all states, including the initial state.[29] I

[28]If these choices respond to the type of repayment contract independent of the government subsidy, the optimal contract will differ from in my restructuring exercise, but the direction of this bias is unclear. On the one hand, income-contingent repayment may encourage borrowers to pursue riskier degrees, which may be welfare-improving if fixed repayment distorts these choices. Alternatively, income-contingent repayment may incentivize excess borrowing from individuals who anticipate never earning enough to repay their debt. Importantly, these choices would need to respond to changes in the utility value of repayments associated with different majors or degrees, which existing literature suggests are likely small (Patnaik et al. 2020).

[29]I define the government budget in present-value terms rather than at the model's stationary distribution because the interpretation of the former is more intuitive, corresponding to the valuation implied by the first-order condition of a hypothetical lender with risk-free discounting. Additionally, this definition is preferable when I consider budget-neutral repayment policies in subsequent analyses because it ensures a reasonable

choose to discount at the risk-free rate since there is no aggregate risk in the model.

My analyses focus on subsidized contracts with a zero interest rate, like those available in Australia.[30] In principle, this interest rate is an additional policy parameter that could be optimized. However, a positive interest rate creates incentives for early repayment, which my model was not designed to capture. Finally, the benchmark contract for these analyses is a 25-year fixed repayment contract without forbearance (i.e., payment pauses for low-income borrowers), where borrowers make constant repayments for 25 years after graduation to repay their loan principal. I denote the government budget under this contract as $\overline{\mathcal{G}}$. This contract is a natural benchmark because it is available in the US and has a similar duration to existing income-contingent contracts while being a debt contract. I implement it without forbearance to create a realistic fiscal cost, but consider the effects of forbearance in Section 4.2.

Welfare metrics. To measure the welfare effects of moving from the benchmark contract to alternative (income-contingent) contracts, I use two welfare metrics. The first is the equivalent variation of contract p, defined as the size of the transfer that makes a borrower, prior to knowing her initial states, indifferent between repaying under contract p versus repaying under the benchmark contract with the additional transfer at $a = a_0$. The second is the consumption-equivalent welfare gain of contract p, as in Benabou (2002), defined as the value of g that makes a borrower, prior to knowing her initial states, indifferent between repaying under contract p versus repaying under the benchmark contract and having her consumption increased by $g\%$ in every state. I denote these two metrics by π_p and g_p, respectively. I compute these welfare metrics solely among college-educated borrowers with $\mathcal{E}_i = 1$; see Appendix D.5 for additional details on the computation of these metrics.

4.1 Effects of Existing Income-Contingent Contracts: Non-Budget-Neutral Comparison

I begin by studying the welfare and fiscal impacts of moving from the benchmark fixed repayment contract to various income-contingent contracts used in Australia and the US.

path for budget deficits in the transition between two policies without the difficulties associated with fully characterizing transition dynamics. In particular, this definition ensures that, if the government were to immediately start giving loans to people graduating from college under two policies with equal values of $\mathcal{G}$, there would be no change in expected costs for this group of individuals.

[30]Under the new "SAVE" income-driven repayment plan in the US, interest does not accrue for borrowers who make the required payments. Therefore, the interest rate is zero for many borrowers.

Figure 14. Repayment Functions for Existing and Constrained-Optimal Contracts

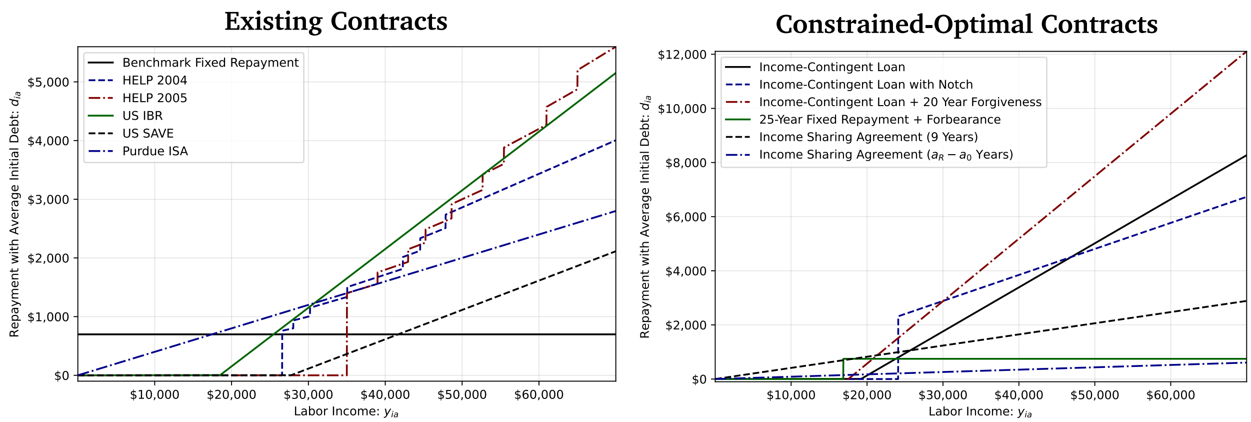

Notes: The left panel of this figure plots the required debt repayments as a function of income for the different existing income-contingent contracts that I consider. The horizonal black line corresponds to the benchmark fixed repayment contract. See Appendix D.4 for the exact implementation of each contract. The right panel of this figure plots the constrained-optimal repayment contracts that solve the constrained-planner's problem in (15) for the different contract spaces described in Section 4.2.

The repayment formulas for the contracts that I consider are shown in the left panel of Figure 14. The first two income-contingent contracts are the 2004 and 2005 HELP contracts. The next two are the income-based repayment (IBR) formula currently used in the US and the new IBR contract formula proposed by the Biden administration (known as SAVE), both of which set repayments equal to a fixed fraction of income above a threshold. Relative to the HELP contracts, these latter two contracts induce a change in marginal rather than average repayment rates. The final formula is a income-sharing agreement or equity contract offered by Purdue University, in which borrowers repay a share of their income for 9 years (Mumford 2022).

The first three rows of Table 5 show that there are significant welfare gains to moving from the benchmark fixed repayment contract to HELP and US IBR. These welfare gains are equivalent to cash transfers of \$$4000$-\$$6000$, or 23-35% of the average initial debt, and a 1.1-1.7% increase in lifetime consumption. These gains, however, come at a fiscal cost: in present value terms, the government collects \$$500$-\$$1600$ less in revenue. Given the differences in fiscal cost, a better way to compare these contracts to each other (and to other policies) is to compute their MVPFs, defined as $-\frac{\pi_p}{\Delta \mathcal{G}_p}$. Table 5 shows that these MVPFs range from 3.7 to 7.7, with US IBR having the highest value. The fact that these MVPFs are above one means that they generate a larger improvement in welfare to borrowers than they cost to the government. As an additional benchmark, the median MVPF for expenditure policies in the Policy Impacts Library is 1.5 and the 75th percentile is 6.7 (as of December 2024). Therefore, the MVPFs from income-contingent loans are sizeable, especially since the highest values typically come from policies that target children rather

Table 5. Effects of Moving from 25-Year Fixed Repayment to Alternative Contracts

Policy: p	π_p	Δ Repayments	Δ Taxes & Transfers	$\Delta\mathcal{G}_p$	MVPF	g_p
HELP 2004	$4,879	$217	-$1,200	-$983	4.96	1.39%
HELP 2005	$6,004	-$185	-$1,450	-$1,635	3.67	1.68%
US IBR	$3,989	$647	-$1,163	-$516	7.73	1.14%
US SAVE	$7,578	-$2,033	-$1,078	-$3,111	2.44	2.08%
US IBR + Fixed Cap	$5,518	-$1,006	-$447	-$1,453	3.80	1.55%
US SAVE + Fixed Cap	$8,171	-$3,018	-$683	-$3,702	2.21	2.23%
US IBR + Forgiveness	$4,265	-$573	-$1,171	-$1,745	2.44	1.22%
US SAVE + Forgiveness	$7,926	-$4,445	-$925	-$5,370	1.48	2.17%
Purdue ISA	$1,984	$1,069	-$1,808	-$738	2.69	0.59%
Debt Cancellation	$10,912	-$14,026	$391	-$13,634	0.80	2.89%

Notes: This table shows the effects of moving from the benchmark 25-year fixed repayment contract to alternative repayment contracts from the left panel of Figure 14 indicated in the first column. The second column shows the equivalent variation, π_p. The third and fourth columns show the change in the government budget defined in (14) that comes from changes in debt repayments and taxes and transfers individually; the fifth column sums these two columns to generate the total fiscal impact. The sixth column computes the marginal value of public funds (MVPF) computed by dividing the second and (negative one times) fifth columns. The final column shows the consumption-equivalent welfare gain, g_p. See Appendix D.4 for the exact implementation of each contract. The rows with "+ Fixed Cap" correspond to contracts where individuals cannot repay more than they would under the benchmark fixed repayment contract. The rows with "+ Forgiveness" correspond to contracts where all debt is forgiven after 20 years. The final row, Debt Cancellation, corresponds to borrowers not repaying any of their debt. The model used for these analyses is the model estimated in column (5) of Table 3.

than adults (Hendren and Sprung-Keyser 2020).

The fourth row of Table 5 shows that the new SAVE loan introduced by the Biden administration generates almost twice the welfare gains of the first three contracts, but is around three times as costly. This is because the SAVE repayment formula, shown in the left panel of Figure 14, has a very low repayment rate. Because the fiscal cost of this contract is so high, its MVPF is around 70% lower than that of the current US IBR program. This suggests that the SAVE's more generous repayment formula is not a well-targeted subsidy.

Next, I consider the effects of adding a cap to US IBR and US SAVE that prevents borrowers from repaying more than under the benchmark fixed repayment contract. This is a feature of most income-contingent loans in the US, but not in Australia. In the model, adding this cap lowers the MVPF of US IBR by over 50% to 3.8, suggesting that a desirable aspect of income-contingent loans is the faster repayment from higher-income borrowers. For US SAVE, adding this cap only decreases the MVPF by around 10% because the lower repayment rate implies the cap is binding for fewer borrowers.

The seventh and eighth rows of Table 5 show that adding forgiveness to US IBR and US SAVE significantly reduces the MVPF of both policies. I consider contracts where the forgiveness happens after 20 years, which applies to all income-contingent loans in the US

Table 6. Fiscal Cost Decomposition of Moving from 25-Year Fixed Repayment to Alternative Contracts

Policy: p	$\Delta\mathcal{G}_p$	$\Delta\mathcal{G}_p$ with ℓ Fixed	$\Delta\mathcal{G}_p$ from ℓ Response
HELP 2004	-\$983	\$261	-\$1,244
HELP 2005	-\$1,635	-\$105	-\$1,529
US IBR	-\$516	\$679	-\$1,195
US SAVE	-\$3,111	-\$2,000	-\$1,111
US IBR + Fixed Cap	-\$1,453	-\$990	-\$463
US SAVE + Fixed Cap	-\$3,702	-\$2,992	-\$710
US IBR + Forgiveness	-\$1,745	-\$516	-\$1,228
US SAVE + Forgiveness	-\$5,370	-\$4,403	-\$967
Purdue ISA	-\$738	\$1,195	-\$1,933

Notes: This table decomposes the total fiscal cost associated with moving from the benchmark 25-year fixed repayment contract to alternative repayment contracts from the left panel of Figure 14 indicated in the first column. The second column repeats the same values in column (5) of Table 5. The third column computes the change in fiscal cost assuming that ℓ_{ia} remains fixed it its value under the benchmark contract for all i and a. The final column reports the difference between the prior to columns, which represents the fiscal cost that comes from adjustments in labor supply.

but not in Australia.[31] For US IBR, this causes the MVPF to decline by around 70% to 2.4. This is because forgiveness is a poorly-targeted subsidy once income-contingent repayment has been implemented: lower-income borrowers who value it most already will repay less, so the subsidy primarily benefits higher-income borrowers who value it less.

The final two policies that I consider are the income-sharing agreement offered by Purdue University and full debt cancellation, in which all debt is forgiven. Table 5 shows that the income-sharing agreement has a lower MVPF than most of the income-contingent loans. This is because a pure income-sharing agreement requires repayments from *all* borrowers, while existing income-contingent loans allow zero payments from borrowers with sufficiently low incomes. The final row of Table 5 shows that debt cancellation increases borrower welfare by more than any of the other contracts considered, but it is sufficiently costly that it has an MVPF below one. This is because full debt forgiveness is a very inefficient policy: income-contingent loans can provide around 40% of the welfare gain of moving from fixed repayment to full forgiveness at around only 4% of the fiscal cost.

As illustrated in Table 5, the total fiscal cost associated with income-contingent loans comes from both a change in repayments *and* the fact that these contracts collect less in taxes net of transfers. Because the tax and transfer system is held fixed, changes in the latter indicate the presence of behavioral responses. To isolate the effect of the moral hazard created by income-contingent repayment, Table 6 decomposes the total fiscal cost associated with moving from the benchmark contract to repayment contract p, $\Delta\mathcal{G}_p$, into two components: (i) the change in fiscal cost assuming labor supply remains fixed; (ii)

[31] Some income-contingent loans in the US have forgiveness after 25 rather than 20 years.

the change in fiscal cost that comes from adjustments in labor supply. The results show that for the two HELP contracts and US IBR, which have the highest MVPFs, as well as the income-sharing agreement, more than 100% of fiscal cost is driven by adjustments in labor supply. For the remaining contracts which are significantly more subsidized, around 25-50% of the fiscal costs comes from moral hazard. These results highlight the importance of having a quantitative model of labor supply in order to correctly estimate the effects of income-contingent repayment.

4.2 Constrained-Optimal Income-Contingent Contracts: Budget-Neutral Comparison

Section 4.1 shows that moving from the benchmark fixed repayment contract to existing income-contingent loans increases welfare by a sizeable amount relative to the fiscal costs. In this section, I study optimal policy by solving a Ramsey (1927)–style constrained-planner's problem. This analysis has two goals. First, it quantifies the welfare gains of transitioning from fixed to income-contingent repayment without relying on other policy instruments to balance the government budget. Second, it provides insight on the shape and features of constrained-optimal income-contingent loans.

4.2.1 Definition of Constrained-Planner's Problem

I consider a social planner that maximizes borrower welfare by choosing a mandatory repayment contract that applies to all borrowers. The planner is constrained to choosing a contract p from the following contract spaces that have two parameters, ψ_p and K_p:

1. Income-Contingent Loan: $d_{ia}(\psi_p, K_p) = \min\left\{\psi_p * \max\left\{y_{ia} - K_p, 0\right\}, D_{ia}\right\} * \mathbf{1}_{a \leq a_R}$
2. Income-Contingent Loan with Notch: $d_{ia}(\psi_p, K_p) = \min\left\{\psi_p * y_{ia} * \mathbf{1}_{y_{ia} \geq K_p}, D_{ia}\right\} * \mathbf{1}_{a \leq a_R}$
3. Income-Sharing Agreement (T Years): $d_{ia}(\psi_p, K_p) = \psi_p * y_{ia} * \mathbf{1}_{a - a_0 \leq T}$

Aside from tractability, the restriction of the contract space is motivated by practical constraints that make implementing more complicated policies difficult (Piketty and Saez 2013). The first contract space corresponds to the class of income-contingent loans in the US and UK where individuals repay a fixed fraction of their marginal income, ψ_p, above a threshold, K_p. The second space is the same as the first, except that the income threshold at K_p changes the average rather than marginal repayment rate, as in Australia. The final contract space is income-sharing agreements, where individuals pay ψ_p share of their income for T years regardless of their debt balances.

Given a contract space, the constrained-planner's problem that determines ψ_p and K_p is:

$$\max_{\psi_p \in [0,1],\, K_p \geq 0} \mathbf{E}_0 \left(V_{ia_0} \mid \mathcal{E}_i = 1 \right), \tag{15}$$

subject to:

$$\mathbf{E}_0 \left(\sum_{a=a_0}^{a_T} \frac{\tau_{ia} - u i_{ia} - \underline{c}_{ia} + d_{ia}(\psi_p, K_p)}{R^{a-a_0}} - D_{ia_0} \right) \geq \overline{\mathcal{G}}.$$

The planner's objective function is the expected indirect utility of a hypothetical borrower who is "behind the veil of ignorance" with respect to her initial states and views the realization of these states as risk. This objective implicitly depends on the two policy parameters through the debt repayment function. By choosing this objective, redistribution across initial conditions and ex-post realizations of shocks are both viewed by the planner as providing insurance. In Section 4.2.7, I decompose welfare gains into the components that come from each of these two channels separately. The constraint that the planner faces is that the government budget under the chosen policy parameters be at least as large as under the benchmark contract. Solving (15) is numerically challenging; I leverage a combination of barrier methods and a global optimizer detailed in Appendix D.6.

4.2.2 Income-Contingent Loans Increase Welfare without Additional Fiscal Costs

The first row of Table 7 shows the results from solving for the constrained-optimal income-contingent loan, and the right panel of Figure 14 plots repayments as a function of income at the optimal contract parameters. The optimal ψ_p and K_p are around 16% and \$$19000$, implying that this contract collect zero repayments from borrowers in the bottom 11th percentile in the income distribution. Relative to the existing income-contingent loans in the left panel of Figure 14, this contract has a repayment threshold that is lower than both HELP contracts and US SAVE, but close to that of US IBR.[32] However, unlike US IBR, this contract has a higher repayment rate of 16% relative to 10% in order balance the government budget.

Table 7 also shows the welfare gain from this income-contingent loan is equivalent to a

[32] In US IBR, K is 1.5 times the US federal poverty line, which was \$14,580 USD for a single household in 2023. Deflating this to 2005 USD with the CPI and then converting to 2005 AUD with the USD/AUD exchange rate as of June 2005 delivers \$12,320, similar to the poverty line reported by the Melbourne Institute in 2005 of \$11,511. This implies the value of K for US IBR of 1.5∗\$12,320 = \$18,480 AUD.

Table 7. Parameters and Welfare Effects of Constrained-Optimal Contracts

Contract Space: p	ψ_p	K_p	π_p	g_p	$\psi_p^{\ell \text{ fixed}}$	$K_p^{\ell \text{ fixed}}$
Income-Contingent Loan	16%	\$19,188	\$2,778	0.79%	38%	\$39,702
Income-Contingent Loan with Notch	9.6%	\$24,093	\$1,508	0.46%	15%	\$47,001
Income-Contingent Loan + 20 Year Forgiveness	23%	\$17,533	\$1,128	0.36%	32%	\$29,516
25-Year Fixed Repayment + Forbearance	0.54%	.	\$267	0.10%	0.12%	.
Income Sharing Agreement (9 Years)	4.1%	.	\$1,730	0.52%	3.6%	.
Income Sharing Agreement ($a_R - a_0$ Years)	0.87%	.	\$6,549	1.82%	0.78%	.

Notes: This table shows the effects of moving from the benchmark 25-year fixed repayment contract to constrained-optimal repayment contracts that solve (15) within the different contract spaces indicated in the first column. The optimal contract parameters are shown in the second and third columns, and plotted in the right panel of Figure 14. For 25-Year Fixed Repayment + Forbearance, which is descrbed in Section 4.2.5, ψ_p corresponds to the (net) debt interest rate, r_d. The fourth and fifth columns show the two welfare metrics, π_p and g_p. The final two columns show the optimal contract parameters from solving (15) assuming that ℓ_{ia} remains fixed it its value under the benchmark contract for all i and a. The objective function in the latter column does not include the disutility of labor supply, given that labor supply is held fixed. The model used for these analyses is the model estimated in column (5) of Table 3.

cash transfer of $\$2800$ or 0.8% increase in lifetime consumption. These welfare gains are sizeable: they correspond to around 16% of the average initial debt and over a quarter of the gain from forgiving debt balances entirely. They are lower than the welfare gains in Table 5, but because this constrained-optimal income-contingent loan is budget-balanced it has an MVPF of infinity.

The fact that the welfare gain from the constrained-optimal income-contingent loan is positive implies that insurance benefits of income-contingent repayment outweigh the costs of the moral hazard and the distortions in consumption-saving decisions it creates. Nevertheless, labor supply responses still have a significant effect on contract design. The final two columns of Table 7 show the results from solving (15) under the assumption that labor supply remains fixed at its value under the baseline contract. This results in values of ψ_p and K_p that are over twice as large, generating a contract that provides substantially more insurance to borrowers, which increases welfare by an additional 0.9 pp of lifetime consumption (see Figure A16). However, this contract cannot raise a sufficient amount of revenue when labor supply is endogenous. Hence, the constrained-optimal contract lowers the threshold to collect repayments from more borrowers and the repayment rate to induce a smaller behavioral response.

4.2.3 Income-Contingent Loans with a Notch Still Increase Welfare

A distinguishing feature of HELP is that the repayment threshold induces a change in the average rather than marginal repayment rate. To assess the welfare effects of having a repayment function with notch rather than a kink (in the language of Kleven and Waseem 2013), the second row of Table 7 shows the results for the constrained-optimal income-contingent loan where ψ_p controls the average repayment rate when y_{ia} is above K_p. The resulting contract has a lower repayment rate and higher repayment threshold than the first income-contingent loan. Although having a notch causes a larger behavioral response, it also collects more revenue from individuals at the repayment threshold who do not adjust due to optimization frictions. As a result, the planner can collect substantially more revenue from these borrowers, allowing a higher value of K_p. However, even with the higher repayment threshold, this contract has a welfare gain equivalent to a 0.46% increase in lifetime consumption, around 40% lower than in the first income-contingent loan. This suggests that formulating an income-contingent loan with a notch is suboptimal, but it is still preferable to using a fixed repayment contract.

Importantly, the fact that a notch is suboptimal depends on the fact that the estimated value of f_H is finite. Table A6 shows that in the model estimated in column (4) of Table 3, which imposes $f_H = \infty$, having a notch does not decrease welfare. As the repayment threshold increases and hence the size of the notch increases, there are always borrowers who will not respond if $f_H = \infty$. In contrast, with $f_H < \infty$, the notch can only be so large before all borrowers start responding.

4.2.4 Adding Forgiveness to Income-Contingent Loans Reduces Welfare Gains

Next, I consider the effects of adding (anticipated) forgiveness after a fixed horizon to income-contingent loans, a feature of those available the US and UK. The third row of Table 7 shows the results from solving (15) using the same contract space in the first row, but adding forgiveness after twenty years At $a_0 + 20$, as in US IBR. Adding forgiveness reduces the welfare gain from the constrained-optimal income-contingent loan by over 50% to a 0.36% equivalent increase in lifetime consumption. This decline in the welfare gain from adding forgiveness reflects the fact that the repayment rate must be increased and the threshold decreased in order to balance the government budget. Given that older borrowers are those who receive the forgiveness, this effectively results in a transfer of repayments from older to younger borrowers. However, this transfer reduces welfare because younger borrowers have a higher marginal value of wealth from tighter borrowing constraints and a stronger precautionary saving motive (Gourinchas and Parker 2002; Boutros et al. 2022).

It is important to emphasize that the negative effects of anticipated forgiveness could be overturned in a behavioral model. For example, if individuals are present-biased and fail to adequately save for retirement, then forgiveness may be a useful way to target middle-aged individuals with inadequate retirement savings that are still paying off their debt. However, if this is the goal that forgiveness aims to achieve, other policy tools may be more desirable, such as reforming the retirement pension system, which could benefit all present-biased individuals independent of whether they have debt.

4.2.5 Fixed Repayment + Forbearance Underperforms Income-Contingent Loans

An alternative to income-contingent loans for providing insurance is to allow for default on fixed repayment contracts. In practice, this is how the fixed repayment contracts operate: repayment can be delayed (but not discharged) for low-income borrowers who enter deferment, forbearance, or default. As of 2019, 30% of student debt was in one of these non-repayment states (US Department of Education). Although formally modeling strategic default is beyond the scope of this paper (see Ji 2021), I evaluate its importance by adding forbearance to the benchmark fixed repayment contract that is available for borrowers receiving unemployment insurance and adjusting the (net) debt interest rate to satisfy the second constraint in (15). This resulting contract, however, likely overstates the actual benefits of forbearance in the US because it allows forbearance to be accessed frictionlessly for an unlimited number of times, while in the US it can only be used a fixed number of times.

The fourth row of Table 7 shows that adding forbearance to fixed repayment increases welfare by the equivalent of a 0.1% increase in lifetime consumption, only 13% of the gain from income-contingent loans. These smaller gains echo the results in Section 4.1 and reflect the benefits of the call option–like structure of an income-contingent loan, which collects repayments more quickly from high-income borrowers. Although these borrowers are likely to pay off their debt, the acceleration of these repayments forward in time increases their discounted value, which allows the planner to provide more insurance.

4.2.6 Long-Horizon ISAs Can Outperform Income-Contingent Loans

The final contract space I consider is income-sharing agreements (ISAs), which were originally proposed by Friedman (1955) and motivated the subsequent development of income-contingent loans. Although the private provision of these contracts has been limited by adverse selection (Herbst and Hendren 2021), my model serves as a natural laboratory to examine their effectiveness as a mandated government-provided contract that would be less subject to adverse selection. The fifth row of Table 7 shows the results for an ISA with a nine year duration, which is the duration of the ISAs offered by Purdue University (Mumford 2022) and is similar to the ISAs being used in developing countries (Herbst et al.

2023). Relative to the baseline fixed repayment contract, this 9-year ISA improves welfare by the equivalent of a 0.52% increase in lifetime consumption. However, this is still 35% lower than the constrained-optimal income-contingent loan.

One reason that the 9-year ISA underperforms the income-contingent loan is that it concentrates repayments in the first few years of borrowers' lives when they are more liquidity-constrained. The final row of Table 7 shows that using an ISA that has a repayment horizon of borrower's entire working life. This contract has a welfare gain that is over twice as large as that of the income-contingent loan. However, for reasons described in Section 4.2.7, these gains primarily reflect redistribution across initial states, especially based on debt balances since repayments do not depend on the amount borrowed. This suggests these gains are more likely to generate responses outside of my model, such as additional borrowing and selection, that would undermine the effectiveness of ISAs and make income-contingent loans a more robust implementation of income-contingent repayment.[33]

4.2.7 Decomposition of Welfare Gains into Insurance and Redistribution

The planner's objective function in (15) combines two distinct objectives: (i) providing borrowers with insurance against the realization of ex-post shocks; (ii) redistributing across borrowers with different initial conditions. The analysis thus far takes the perspective of a hypothetical borrower who does not know her initial states and views the realization of these states as risk. This perspective is natural because the initial states in the model are not primitive individual characteristics, but rather the outcomes of ex-ante borrowing and education decisions that are taken as given.[34] However, in reality, some of the variation in these states likely does not reflect risk and is probably driven by ex-ante heterogeneity in borrower types. In this case, the way the planner values redistribution *across* borrowers types depends on society's social preferences and need not be the same as how she values redistribution *within* borrowers, which depends on borrowers' preferences.

To decompose welfare gains into insurance and redistribution, I take an approach analogous to Berger et al. (2025) and introduce lump-sum transfers that differ based on initial

[33]Another reason these gains from ISAs are likely an upper bound is because ϕ is identified from median-income borrowers. However, the borrowers making the bulk of the repayments under ISAs have higher incomes, whom prior literature suggests have higher taxable income elasticities (e.g., Gruber and Saez 2002).

[34]For example, upon entering college, individuals do not know what their income will be upon graduation. Therefore, to the extent that they are risk-averse, they will value insurance against realizations of it, even though it corresponds to an initial state in my model.

Table 8. Welfare Effects Before and After Redistribution-Neutralizing Transfers

Contract Space: p	π_p^{Before}	π_p^{After}	g_p^{Before}	g_p^{After}
Income-Contingent Loan	\$4,012	\$1,616	1.03%	0.50%
Income Sharing Agreement ($a_R - a_0$ Years)	\$6,182	.	1.75%	.

Notes: This table shows the effects of moving from the benchmark 25-year fixed repayment contract to constrained-optimal repayment contracts that solve (15) within the different contract spaces indicated in the first column. The second and third columns show the two welfare metrics, π_p and g_p. These values differ from Table 7 because the set of initial conditions is different and has been discretized into $\mathcal{T}$ values. The final two columns show the same welfare metrics that come from solving (15) with lump-sum transfers made at a_0 to borrowers in each $\mathcal{T}$ possible initial state to ensure the government budget remains unchanged at each of these states. The values of these transfers are shown in Figure A17. See Appendix D.7 for additional details on this analysis.

conditions.[35] I begin by discretizing the set of initial states into a grid with $\mathcal{T}$ possible values. Next, I solve the constrained-planner's problem in (15) with two modifications. First, I introduce an additional $\mathcal{T}$ policy instruments that correspond to lump-sum transfers made at a_0 to borrowers based on their $\mathcal{T}$ possible initial states. Second, I introduce an additional $\mathcal{T}$ constraints, requiring that the government budget defined in (14) remains unchanged at each of the possible $\mathcal{T}$ initial states. As a result, the solution to this constrained-planner's problem with transfers does not involve any redistribution across initial states. Appendix D.7 provides additional computational details on this approach.

Table 8 shows that around half of the welfare gain from the constrained-optimal income-contingent loan comes from redistribution across initial states, while the other half comes from insurance. The second column and fourth columns show that the welfare gain from solving (15) without the additional transfers is equivalent to a cash transfer of \$$4000$ or a 1% increase in lifetime consumption. This differs slightly from the welfare gain in Table 7 because the distribution of initial states in this model is different than in the baseline model due to the discretization. In contrast, the third and fifth columns show the welfare gain from solving the (15) with the additional transfers is equivalent to a cash transfer of \$$1600$ or a 0.5% increase in lifetime consumption. This implies that 40-48% of the gains to moving from fixed repayment to the constrained-optimal income-contingent loan reflects insurance, while the remaining 52-60% reflects redistribution across initial states. To the extent that the initial states reflect the realization of some risk, these estimates provide an upper bound on the fraction of gains that come from redistribution.

Most of the redistribution that occurs when moving from the benchmark contract to the

[35] An alternative method to quantify the importance of redistribution separately from insurance would be to perform a decomposition in the spirit of Benabou (2002), Heathcote et al. (2017), or Abbott et al. (2019). However, these approaches are not feasible in my setting because they require computing certainty-equivalent consumption. Unlike Benabou (2002) and Heathcote et al. (2017), this cannot be done in closed form. Unlike Abbott et al. (2019), it is not computationally feasible to compute these certainty-equivalents numerically.

constrained-optimal income-contingent loan occurs across initial wage levels among borrowers with large initial debt balances. Figure A17 shows the transfers made to each of the $\mathcal{T}$ initial states. Eliminating redistribution requires transfers of around $\$3000$ from borrowers with lower wages and higher debt, who gain the most from the income-contingent loan, to those with higher wages and higher debt. In contrast, there is much less redistribution across different initial debt or asset levels.

The final row of Table 8 attempts to perform the same decomposition for the ISA in the final row of Table 7, which has the largest welfare gain. However, I find that it is not possible to find transfers that balance the government budget at every initial state. Figure A17 shows that, in the initial states where the budget can be balanced, the transfers required are much larger than under the income-contingent loan. The direction of redistribution is also quite different: while an income-contingent loan primarily redistributes based on initial wages, the ISA primarily redistributes based on initial debt. This is because ISAs decouple repayments from debt balances, resulting in a large transfer from low– to high–debt borrowers that needs to be undone with transfers. The large redistribution created by ISAs suggests they are more likely to generate responses outside of the model, such as additional borrowing and selection, that might make income-contingent loans a more robust implementation of income-contingent repayment.

4.3 Additional Results and Discussion

4.3.1 Sensitivity to Key Labor Supply Parameters

As with any counterfactual analysis, the key concern is whether the structural parameters are invariant to policy changes. The fact that the model is estimated using a policy change that did occur makes this more likely to be the case. Nevertheless, to assess this concern, I vary the four key parameters that govern labor supply responses–ϕ, λ, f_L, f_H–and resolve (15) for this range of possible values, holding all other parameters at their estimated values from Table 3. Figure 15 plots the resulting consumption-equivalent welfare gains, g_p, of the constrained-optimal income-contingent loan.

The top left panel shows that the welfare gain from the constrained-optimal income-contingent loan is decreasing in the labor supply elasticity, ϕ. This is natural: a higher ϕ increases the moral hazard created by income-contingent repayment, which reduces the amount of insurance that can be provided for a given fiscal cost. For the income-

Figure 15. Welfare Gains from Income-Contingent Loan as a Function of Model Parameters

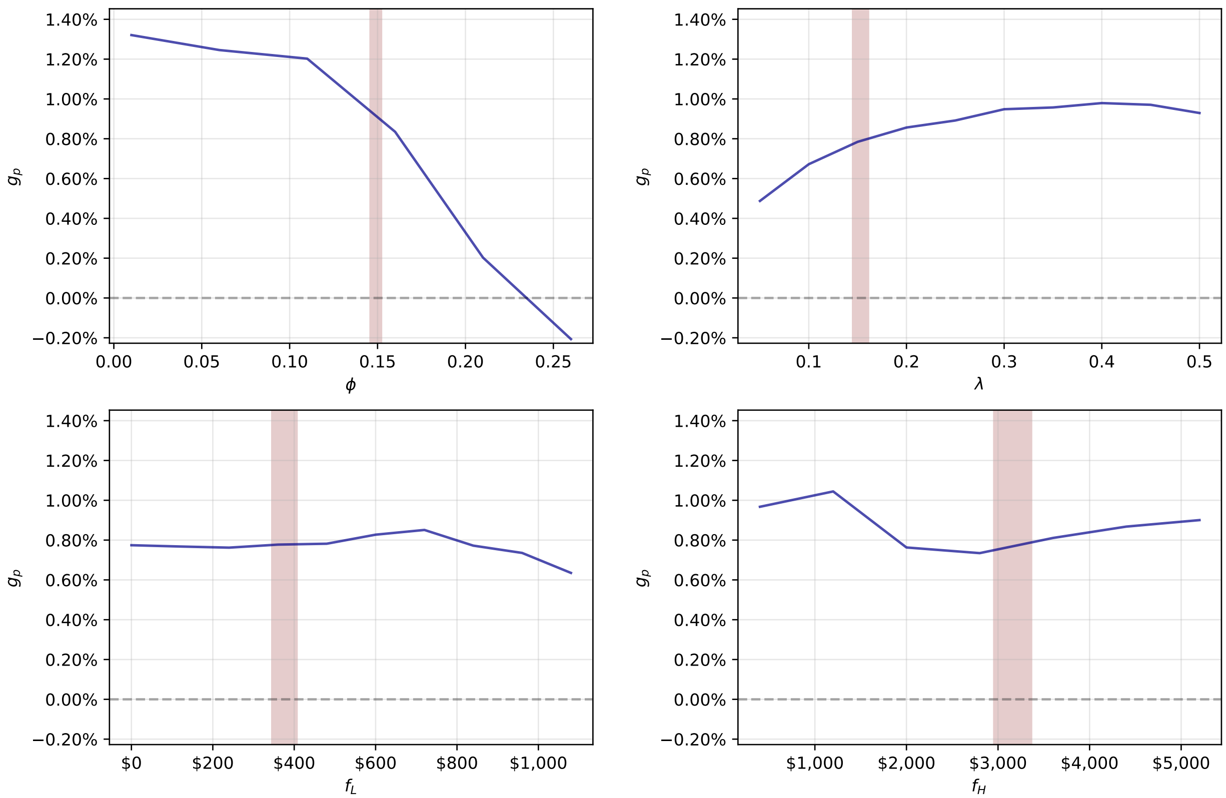

Notes: This figure shows the consumption-equivalent welfare gain from solving for the constrained-optimal income-contingent loan in the baseline model as a function of different model parameters. In each panel, the parameter on the horizontal axis is the only parameter that varies, while all other parameters are held fixed at their values in column (5) of Table 3. (15) is then resolved for each value of the parameter on the horizontal axis. The shaded red regions correspond to the 95% confidence intervals from column (5) of Table 3.

contingent loan to deliver a welfare loss relative to the benchmark fixed repayment contract, I estimate ϕ would need to be above 0.24, which is well outside the confidence interval for its estimated value. Nevertheless, Appendix D.9 shows that using a richer contract space of income-contingent loans restores welfare gains even when $\phi = 0.24$. This is consistent with Shavell (1979): the gains from insurance are first-order while the losses from moral hazard are second-order.

The remaining three panels of Figure 15 show that the welfare gains from income-contingent loans are substantially less sensitive to the adjustment cost probability, λ, and the adjustment costs, f_L and f_H. For all values of these parameters that I consider, which are well outside the range of their estimated values, the welfare gain is positive. This suggests that, although optimization frictions are important for individual-level labor supply responses, their precise values matter less for aggregate responses, which is what the planner cares about. This is not surprising because the responses that have the largest impact on the government budget are not year-to-year small adjustments, which are primarily controlled by frictions, but rather long-run steady-state responses, at which point the role of these frictions diminishes because the probability of labor supply adjustment approaches one. Nevertheless, distinguishing between different models of adjustment frictions is still quantitatively important: the top half of Table 9 shows how the welfare gains vary in the other models of adjustment frictions estimated in Table 3.

4.3.2 Interaction Between Income-Contingent Loans and the Tax System

The analysis thus far has taken the tax and transfer system as given, which is an alternative way to redistribute and provide insurance. I view this as a reasonable starting point because the tax system is designed for the entire population and constrained by the political system. As a result, government agencies, such as the Congressional Budget Office, typically evaluate policies in isolation. Nevertheless, it is clearly desirable to study the joint determination of taxes, transfers, and education financing, as in Stantcheva (2017). Although this is outside the scope of this paper because it requires a model that (at a minimum) also has endogenous college entry, this section performs additional analyses to shed light on how income-contingent loans interact with and differ from changes in the tax system.

The first row in the second panel of Table 9 shows how the solution to (15) changes when the parametric income tax schedule that was calibrated to Australia is changed to the match the US calibration in Heathcote et al. (2017). The shape and welfare gains of the constrained-optimal income-contingent loan are very similar. In the second row, I study how the results change when this tax schedule is optimized to maximize the expected utility of an individual that views all of her initial states as risk, including $\mathcal{E}_i$. With this optimized tax system, there is no gain to moving from income-contingent repayment.

Although the welfare gains from income-contingent repayment can be achieved through the tax system, the targeting of these two policy instruments is quite different. Figure A18 compares the distribution of the welfare gains across two policy experiments: (i) restructuring debt repayment from the baseline contract to the constrained-optimal income-contingent loan, holding taxes and transfers fixed; (ii) moving from the current to optimal tax system, holding the debt repayment fixed. These distributions differ in three ways. First, changing the tax system affects individuals of all education levels, while restructuring debt repayment only affects individuals that attend college. Second, higher-income individuals lose substantially from the change in the tax system because they have to make larger repayments throughout their life. In contrast, repayments by these individuals are capped

Table 9. Welfare Gains from Constrained-Optimal Income-Contingent Loans in Alternative Models

Estimated Models	ψ_p	K_p	π_p	g_p
Baseline Model	16%	\$19,188	\$2,778	0.79%
$f_L = f_H$ Model	16%	\$31,786	\$3,456	1.35%
$f_L = 0, f_H = \infty$ Model	37%	\$38,390	\$4,997	1.61%
$f_H = \infty$ Model	14%	\$31,055	\$4,821	1.18%
Deviation from Baseline Model	ψ_p	K_p	π_p	g_p
US Tax System	15%	\$18,539	\$2,599	0.65%
Optimized Tax System	6%	\$2,104	\$24	0.01%
Lower RRA = 1.5	14%	\$18,565	\$1,429	0.44%
Higher RRA = 4	22%	\$20,856	\$5,551	1.74%
Lower EIS = 0.25	18%	\$18,524	\$2,404	0.84%
Higher EIS = 1.5	11%	\$17,151	\$2,238	0.52%
Wealth Effects on ℓ	33%	\$34,083	\$3,129	0.76%
Less Persistence: $\rho = 0.8$	33%	\$37,518	\$2,963	0.83%
More Persistence: $\rho = 0.99$	8%	\$2,782	\$1,700	0.49%
US Initial Debt Levels	27%	\$16,994	\$9,838	3.03%
Higher Debt Interest Rate: $R_d = 2\%$	28%	\$43,863	\$6,776	1.88%
Government Discount Rate = $R + 2\%$	33%	\$33,095	\$5,044	1.43%

Notes: This table shows results from repeating the analysis in the first row of Table 7. The top panel of the table shows the results in the baseline model, as well as the three additional models with optimization frictions estimated in Table 3. The bottom panel shows the results in the baseline model with the deviations stated in the first column and described in further detail in Appendix D.8.

under income-contingent repayment by their debt balanaces. Finally, restructuring debt repayment more effectively targets individuals with high debt balances, who lose on average from the tax change. Given the policy discussion around student debt is particularly focused on highly-indebted individuals, this more precise targeting may be desirable.

4.3.3 Robustness to Model Misspecification

To assess the sensitivity of my results to model misspecification, Table 9 shows how the solution to (15) varies across several alternative models. A discussion of each model and the corresponding results are presented in Appendix D.8.

5 Conclusion

This paper studies the trade-off between insurance and moral hazard in student loans with income-contingent repayment. Empirically, I show that borrowers reduce their labor supply to lower income-contingent repayments and that these responses are consistent with a moderate elasticity of labor supply and substantial optimization frictions. Through the lens of a structural model, these estimates imply that income-contingent repayment provides significant welfare gains and that income-contingent loans are an effective and robust way of doing so. Relative to fixed repayment contracts with forbearance, income-contingent loans provide more insurance by accelerating payments from high-income borrowers. Relative to equity contracts, the welfare gains involve less redistribution, making them less likely to generate ex-ante responses (e.g., additional borrowing) and be adversely-selected.

The results in this paper speak to what has been labeled a “student debt crisis” in the US (Mitchell 2019). One possible solution is to use income-contingent repayment contracts, such as those introduced by the Biden administration. This paper provides empirical evidence and a structural model that can be used to calibrate the effects of student debt restructuring. Overall, the results suggest that a (mandatory) restructuring of the $1.6 trillion of US student debt in the US from fixed to income-contingent repayment would be beneficial. However, this analysis leaves open several questions, most importantly, how education, occupation, and borrowing choices respond. Income-contingent repayment may affect these choices on the intensive margin by encouraging borrowers to pursue degrees and occupations with riskier returns (Hampole 2022; Murto 2022; Abourezk-Pinkstone 2023) or on the extensive margin by encouraging more borrowers to pursue higher education. Quantifying these responses and their implications for optimal contract design is an important task for future research.

More broadly, the trade-off between insurance and moral hazard studied in this paper applies to the design of other state-contingent financing contracts. Two notable examples are shared-appreciation mortgages, which several public and private lenders have recently begun providing, and revenue-based loans (e.g., Russel et al. 2023), a growing source of financing for start-ups. As with student loans, a key question in designing these contracts is how to balance their insurance benefits with the behavioral distortions that they create. By carefully analyzing the insurance–moral hazard trade-off for student loans, this paper provides a template for studying these issues in other contexts.

References

Abbott, Brant, Giovanni Gallipoli, Costas Meghir, and Giovanni L. Violante (2019), "Education policy and intergenerational transfers in equilibrium." *Journal of Political Economy*, 127, 2569–2624.

Abel, Andrew B., Janice Eberly, and Stavros Panageas (2013), "Optimal Inattention to the Stock Market With Information Costs and Transactions Costs." *Econometrica*, 81, 1455–1481.

Abourezk-Pinkstone, Hayley (2023), "Student Loan Debt and Risk Preferences on the Job Market." *Working Paper*.

Abraham, Katharine G., Emel Filiz-Ozbay, Erkut Y. Ozbay, and Lesley J. Turner (2020), "Framing effects, earnings expectations, and the design of student loan repayment schemes." *Journal of Public Economics*, 183, 1040–1067.

Alon, Titan, Natalie Cox, and Arlene Wong (2024), "Debt, Human Capital Accumulation, and the Allocation of Talent." *Working Paper*.

Alvarez, Fernando, Hervé Le Bihan, and Francesco Lippi (2016), "The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach." *American Economic Review*, 106, 2817–2851.

Amromin, Gene and Janice Eberly (2016), "Education Financing and Student Lending." *Annual Review of Financial Economics*, 8, 289–315.

Anagol, Santosh, Benjamin B Lockwood, Allan Davids, and Tarun Ramadorai (2022), "Diffuse Bunching with Lumpy Incomes: Theory and Estimation." *Working Paper*.

Andersen, Steffen, John Y. Campbell, Kasper Meisner Nielsen, and Tarun Ramadorai (2020), "Sources of inaction in household finance: Evidence from the danish mortgage market." *American Economic Review*, 110, 3184–3230.

Australian Council of Social Service (2017), "Ending tax avoidance, evasion and money laundering through private trusts." ACOSS Policy Briefing.

Australian Council of Social Service (2018), "Components of Australia's wealth."

Barr, Nicholas, Bruce Chapman, Lorraine Dearden, and Susan Dynarski (2019), "The US college loans system: Lessons from Australia and England." *Economics of Education Review*, 71, 32–48.

Beer, Gillian and Bruce Chapman (2004), "HECS System Changes: Impact on Students." *Agenda*, 11, 157–174.

Benabou, Roland (2002), "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?" *Econometrica*, 70, 481–517.

Benetton, Matteo, Philippe Bracke, João F. Cocco, and Nicola Garbarino (2022), "Housing Consumption and Investment: Evidence from Shared Equity Mortgages." *Review of Financial Studies*, 35, 3525–3573.

Berger, David and Joseph Vavra (2015), "Consumption Dynamics During Recessions." *Econometrica*, 83, 101–154.

Berger, David W, Kyle F Herkenhoff, and Simon Mongey (2025), "Minimum Wages, Efficiency and Welfare." *Econometrica*, 93, 265–301.

Black, Sandra E., Jeffrey T. Denning, Lisa J. Dettling, Sarena Goodman, and Lesley J. Turner (2022), "Taking It to the Limit: Effects of Increased Student Loan Availability on Attainment, Earnings, and Financial Well-Being." *Working Paper*.

Bound, John, Julie Berry Cullen, Austin Nichols, and Lucie Schmidt (2004), "The welfare implications of increasing disability insurance benefit generosity." *Journal of Public Economics*, 88, 2487–2514.

Boutros, Michael, Nuno Clara, and Francisco Gomes (2022), "Borrow Now, Pay Even Later: A Quantitative Analysis of Student Debt Payment Plans." *Working Paper*.

Bovenberg, A. Lans and Bas Jacobs (2005), "Redistribution and education subsidies are Siamese twins." *Journal of Public Economics*, 89, 2005–2035.

Britton, Jack and Jonathan Gruber (2020), "Do income contingent student loans reduce labor supply?" *Economics of Education Review*, 79, 1020–1061.

Caballero, Ricardo J. and Eduardo M. R. A. Engel (1999), "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S, s) Approach." *Econometrica*, 67, 783–826.

Calvo, Guillermo A. (1983), "Staggered prices in a utility-maximizing framework." *Journal of Monetary Economics*, 12, 383–398.

Campbell, John Y., Nuno Clara, and João F. Cocco (2021), "Structuring Mortgages for Macroeconomic Stability." *The Journal of Finance*, 76, 2525–2576.

Caplin, Andrew, James H. Carr, Fredrick Pollock, and Zhong Yi Tong (2007), "Shared-Equity Mortgages, Housing Affordability, and Homeownership." *Working Paper*.

Caplin, Andrew and Daniel F. Spulber (1987), "Menu Costs and the Neutrality of Money." *Quarterly Journal of Economics*, 102, 703–726.

Carneiro, Pedro and James J. Heckman (2002), "The Evidence on Credit Constraints in Post-Secondary Schooling." *The Economic Journal*, 112, 705–734.

Carroll, Christopher D. and Miles S. Kimball (1996), "On the Concavity of the Consumption Function." *Econometrica*, 64, 981–992.

Catherine, Sylvain (2022), "Countercyclical Income Risk and Portfolio Choices over the Life-Cycle." *Review of Financial Studies*, 35, 4016–4054.

Catherine, Sylvain and Constantine Yannelis (2023), "The Distributional Effects of Student Loan Forgiveness." *Journal of Financial Economics*, 147, 297–316.

Cesarini, David, Erik Lindqvist, Matthew J. Notowidigdo, and Robert Ostling (2017), "The Effect of Wealth on Individual and Household Labor Supply: Evidence from Swedish Lotteries." *American Economic Review*, 107, 3917–3946.

Chakrabarti, Rajashri, Vyacheslav Fos, Andres Liberman, and Constantine Yannelis (2020), "Tuition, Debt, and Human Capital." *Working Paper*.

Chapman, Bruce (2006), "Chapter 25 Income Contingent Loans for Higher Education: International Reforms." In *Handbook of the Economics of Education*, volume 2, 1435–1503, Elsevier.

Chetty, R., J. N. Friedman, T. Olsen, and L. Pistaferri (2011), "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records." *The Quarterly Journal of Economics*, 126, 749–804.

Chetty, Raj (2008), "Moral hazard versus liquidity and optimal unemployment insurance." *Journal of Political Economy*, 116, 173–234.

Chetty, Raj (2009), "Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance." *American Economic Journal: Economic Policy*, 1, 31–52.

Chetty, Raj (2012), “Bounds on Elasticities With Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply.” *Econometrica*, 80, 969–1018.

Chetty, Raj and Amy Finkelstein (2013), “Social Insurance: Connecting Theory to Data.” In *Handbook of Public Economics*, volume 5, Elsevier B.V.

Chetty, Raj, John N Friedman, and Emmanuel Saez (2013), “Using Differences in Knowledge Across Neighborhoods to Uncover the Impacts of the EITC on Earnings.” *American Economic Review*, 103, 2683–2721.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber (2012), “Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities.” *NBER Macroeconomics Annual*, 27.

Choukhmane, Taha (2021), “Default Options and Retirement Saving Dynamics.” *Working Paper*, 1–78.

Choukhmane, Taha and Tim de Silva (2023), “What Drives Investors’ Portfolio Choices? Separating Risk Preferences from Frictions.” *Working Paper*.

Cochrane, John H. (1991), “A simple test of consumption insurance.” *Journal of Political Economy*, 93, 957–978.

Coyne, David, Itzik Fadlon, and Tommaso Porzio (2022), “Measuring Valuation of Liquidity with Penalized Withdrawals.” *Working Paper*.

Deaton, Angus and Christina H Paxson (1994), “Intertemporal Choice and Inequality.” *Journal of Political Economy*, 102, 437–468.

Department of Education and Training (2023), “HECS-HELP Determinations.” Technical report, Australian Government.

Di Maggio, Marco, Ankit Kalda, and Vincent W. Yao (2021), “Second Chance: Life without Student Debt.” *Journal of Finance*.

D’Souza, Gabriela (2018), “A higher education bubble?” *Working Paper*.

Ebrahimian, Mehran (2020), “Student Loans and Social Mobility.” *Working Paper*.

Einav, Liran, Amy Finkelstein, and Paul Schrimpf (2015), “The Response of Drug Expenditure to Nonlinear Contract Design: Evidence from Medicare Part D.” *The Quarterly Journal of Economics*, 130, 841–899.

Feldstein, Martin (1999), “Tax Avoidance and the Deadweight Loss of the Income Tax.” *Review of Economics and Statistics*, 81, 674–680.

Finkelstein, Amy and Nathaniel Hendren (2020), “Welfare analysis meets causal inference.” *Journal of Economic Perspectives*, 34, 146–167.

Folch, Marc and Luca Mazzone (2021), “Go Big or Buy a Home: Student Debt, Career Choices and Wealth Accumulation.” *Working Paper*.

Friedman, Milton (1955), “The Role of Government in Education.” In *Economics and the Public Interest* (Robert A. Solow, ed.), Rutgers University Press, New Brunswick, New Jersey.

Ganong, Peter and Pascal Noel (2020), “Liquidity versus wealth in household debt obligations: Evidence from housing policy in the great recession.” *American Economic Review*, 110, 3100–3138.

Ganong, Peter and Pascal Noel (2023), “Why Do Borrowers Default on Mortgages?” *The Quarterly Journal of Economics*, 138, 1001–1065.

Gervais, Martin, Qian Liu, and Lance Lochner (2022), "The Insurance Implications of Government Student Loan Repayment Schemes." *Working Paper*.

Golosov, Mikhail and Robert E. Lucas (2007), "Menu Costs and Phillips Curves." *Journal of Political Economy*, 115, 171–200.

Gorodnichenko, Yuriy, Jorge Martinez-Vazquez, and Klara Sabirianova Peter (2009), "Myth and Reality of Flat Tax Reform: Micro Estimates of Tax Evasion Response and Welfare Effects in Russia." *Journal of Political Economy*, 117, 504–555.

Gourinchas, Pierre-Olivier and Jonathan A. Parker (2002), "Consumption over the life cycle." *Econometrica*, 70, 47–89.

Greenwald, Daniel L., Tim Landvoigt, and Stijn Van Nieuwerburgh (2021), "Financial Fragility with SAM?" *Journal of Finance*, 76, 651–706.

Greenwood, Jeremy, Zvi Hercowitz, and Gregory W Huffman (1988), "Investment, Capacity Utilization, and the Real Business Cycle." *American Economic Review*, 78, 402–417.

Gruber, Jon and Emmanuel Saez (2002), "The elasticity of taxable income: Evidence and implications." *Journal of Public Economics*, 84, 1–32.

Gruber, Jonathan (1997), "The Consumption Smoothing Benefits of Unemployment Insurance." *American Economic Review*, 87, 192–205.

Gupta, Arpit and Christopher Hansman (2022), "Selection, Leverage, and Default in the Mortgage Market." *The Review of Financial Studies*, 35, 720–770.

Guvenen, Fatih (2009), "A Parsimonious Macroeconomic Model for Asset Pricing." *Econometrica*, 77, 1711–1750.

Guvenen, Fatih, Alisdair McKay, and Conor Ryan (2022), "A Tractable Income Process for Business Cycle Analysis." *Working Paper*.

Gyöngyösi, Gyozo, Judit Rariga, and Emil Verner (2022), "How do Borrowers Adjust in a Household Foreign Currency Debt Crisis?" *Working Paper*.

Hampole, Menaka V (2022), "Financial Frictions and Human Capital Investments." *Working Paper*.

Handel, Benjamin R. (2013), "Adverse selection and inertia in health insurance markets: When nudging hurts." *American Economic Review*, 103, 2643–2682.

Hanson, Melanie (2022), "Student Loan Default Rate." Technical report, Education Data Initiative.

Hanson, Melanie (2023), "Average Cost of College & Tuition." Technical report, Education Data Initiative.

Hartman-Glaser, Barney and Benjamin Hébert (2020), "The Insurance Is the Lemon: Failing to Index Contracts." *Journal of Finance*, 75, 463–506.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2014), "Consumption and Labor Supply with Partial Insurance: An Analytical Framework." *American Economic Review*, 104, 2075–2126.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2017), "Optimal Tax Progressivity: An Analytical Framework." *The Quarterly Journal of Economics*, 132, 1693–1754.

Hendren, Nathaniel and Ben Sprung-Keyser (2020), "A Unified Welfare Analysis of Government Policies." *The Quarterly Journal of Economics*, 135, 1209–1318.

Herbst, Daniel (2023), "Liquidity and Insurance in Student-Loan Contracts: The Effects of Income-Driven Repayment on Borrower Outcomes." *American Economic Journal: Applied Economics*, 15, 1–25.

Herbst, Daniel and Nathaniel Hendren (2021), "Opportunity Unraveled: Private Information and the Missing Markets for Financing Human Capital." *Working Paper*.

Herbst, Daniel, Miguel Palacios, and Constantine Yannelis (2023), "Equity and Incentives in Household Finance." *Working Paper*.

Huang, Yueling (2022), "Rethinking College Financing: Wealth, College Majors, and Macroeconomic Consequences." *Working Paper*.

Indarte, Sasha (2023), "Moral Hazard versus Liquidity in Household Bankruptcy." *Journal of Finance*, 78, 2421–2464.

Jacob, Brian A, Damon Jones, and Benjamin J Keys (2023), "The Value of Student Debt Relief and the Role of Administrative Barriers: Evidence from the Teacher Loan Forgiveness Program." *Working Paper*.

Ji, Yan (2021), "Job Search under Debt: Aggregate Implications of Student Loans." *Journal of Monetary Economics*, 117, 741–759.

JPMorgan Chase (2022), "Income Driven Repayment: Who needs student loan payment relief?" Technical report.

Karamcheva, Nadia, Jeffrey Perry, and Constantine Yannelis (2020), "Income-Driven Repayment Plans for Student Loans." *Working Paper*.

Kargar, Mahyar and William Mann (2023), "The Incidence of Student Loan Subsidies: Evidence from the PLUS Program." *The Review of Financial Studies*, 36, 1621–1666.

Karlan, Dean and Jonathan Zinman (2009), "Observing Unobservables: Identifying Information Asymmetries With a Consumer Credit Field Experiment." *Econometrica*, 77, 1993–2008.

Keane, Michael P (2011), "Labor Supply and Taxes: A Survey." *Journal of Economic Literature*, 49, 961–1075.

Kleven, Henrik, Claus Kreiner, Kristian Larsen, and Jakob Søgaard (2023), "Micro vs Macro Labor Supply Elasticities: The Role of Dynamic Returns to Effort." *Working Paper*.

Kleven, Henrik J. and Mazhar Waseem (2013), "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan*." *The Quarterly Journal of Economics*, 128, 669–723.

Lochner, Lance and A. Monge-Naranjo (2016), *Student Loans and Repayment: Theory, Evidence, and Policy*, volume 5. Elsevier B.V.

Lochner, Lance and Alexander Monge-Naranjo (2012), "Credit constraints in education." *Annual Review of Economics*, 4, 225–256.

Luo, Mi and Simon Mongey (2024), "Assets and Job Choice: Student Debt, Wages, and Amenities." *Review of Economic Studies*.

Lusardi, Annamaria, Pierre Carl Michaud, and Olivia S. Mitchell (2017), "Optimal financial knowledge and wealth inequality." *Journal of Political Economy*, 125, 431–477.

Manso, Gustavo, Alejandro Rivera, Hui Wang, and Han Xia (2024), "Household Debt Overhang and Human Capital Investment." *Working Paper*.

Marshall, Kate (2003), “Ease HECS burden on students, say universities.” *Australian Financial Review*.

Martin, Chelsey (2004), “For one in four, HECS now a lifelong debt.” *Australian Financial Review*.

Masatlioglu, Yusufcan and Efe A. Ok (2005), “Rational choice with status quo bias.” *Journal of Economic Theory*, 121, 1–29.

Matsuda, Kazushige and Karol Mazur (2022), “College education and income contingent loans in equilibrium.” *Journal of Monetary Economics*, 132, 100–117.

Mezza, Alvaro, Daniel Ringo, Shane Sherlund, and Kamila Sommer (2020), “Student loans and homeownership.” *Journal of Labor Economics*, 38, 215–260.

Mian, Atif and Amir Sufi (2014), *House of Debt: How They (and You) Caused the Great Recession, and How We Can Prevent It from Happening Again*. University of Chicago Press, Chicago, IL.

Mincer, Jacob (1974), *Schooling, Experience, and Earnings*. National Bureau of Economic Research.

Mirrlees, J. (1974), “Notes on Welfare Economics, Information and Uncertainty.” In *Essays on Economic Behavior under Uncertainty*, Amsterdam: North Holland.

Mitchell, Josh (2019), “The Long Road to the Student Debt Crisis.” *Wall Street Journal*.

Mueller, Holger M. and Constantine Yannelis (2019), “The rise in student loan defaults.” *Journal of Financial Economics*, 131, 1–19.

Mueller, Holger M. and Constantine Yannelis (2022), “Increasing Enrollment in Income-Driven Student Loan Repayment Plans: Evidence from the Navient Field Experiment.” *Journal of Finance*, 77, 367–402.

Mumford, Kevin J. (2022), “Student Selection into an Income Share Agreement.” *Working Paper*.

Murto, Michael J (2022), “Student Loans and Human Capital Investments.” *Working Paper*.

Nakamura, Emi and Jon Steinsson (2010), “Monetary Non-Neutrality in a Multi-Sector Menu Cost Model.” *Quarterly Journal of Economics*, 125, 961–1013.

Nelson, Brendan (2003), “Our universities: Backing Australia’s future.” Technical report, Commonwealth of Australia, Canberra, ACT.

Norton, Andrew (2019), “Demand-driven funding for universities is frozen. What does this mean and should the policy be restored?” *The Conversation*.

Paetzold, Jörg and Hannes Winner (2016), “Taking the high road? Compliance with commuter tax allowances and the role of evasion spillovers.” *Journal of Public Economics*, 143, 1–14.

Palacios, Miguel (2004), *Investing in Human Capital, A Capital Markets Approach to Student Funding*. Cambridge University Press.

Patnaik, Arpita, Matthew Wiswall, and Basit Zafar (2020), “College Majors.” *Routledge Handbook of Economics of Education*.

Piketty, Thomas and Emmanuel Saez (2013), “Optimal Labor Income Taxation.” In *Handbook of Public Economics*, volume 5, 391–474, Elsevier.

Piskorski, Tomasz and Amit Seru (2018), “Mortgage Market Design: Lessons from the Great Recession.” *Brookings Papers on Economic Activity*, 2018, 429–513.

Ramsey, F. P. (1927), “A Contribution to the Theory of Taxation.” *The Economic Journal*, 145, 47–61.

Russel, Dominic, Claire Shi, and Rowan P. Clarke (2023), "Revenue-Based Financing." *Working Paper*.

Saez, Emmanuel (2010), "Do Taxpayers Bunch at Kink Points?" *American Economic Journal: Economic Policy*, 2, 180–212.

Saez, Emmanuel, Joel Slemrod, and Seth H Giertz (2012), "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review." *Journal of Economic Literature*, 50, 3–50.

Shavell, Steven (1979), "On Moral Hazard and Insurance." *The Quarterly Journal of Economics*, 93, 541–562.

Shiller, Robert J. (2004), *The New Financial Order: Risks in the 21st Century*. Princeton University Press, Princeton, NJ.

Slemrod, Joel (2019), "Tax Compliance and Enforcement." *Journal of Economic Literature*, 57, 904–954.

Slemrod, Joel and Shlomo Yitzhaki (2002), "Tax Avoidance, Evasion, and Administration." *Handbook of Public Economics*, 3, 1425–1465.

Stantcheva, Stefanie (2017), "Optimal taxation and human capital policies over the life cycle." *Journal of Political Economy*, 125, 1931–1990.

Yannelis, Constantine (2020), "Strategic Default on Student Loans." *Working Paper*.

Yannelis, Constantine and Greg Tracey (2022), "Student Loans and Borrower Outcomes." *Annual Review of Financial Economics*, 14, 167–186.

Required Disclaimer for Use of MADIP Data

The results of these studies are based, in part, on Australian Business Registrar (ABR) data supplied by the Registrar to the ABS under A New Tax System (Australian Business Number) Act 1999 and tax data supplied by the ATO to the ABS under the Taxation Administration Act 1953. These require that such data is only used for the purpose of carrying out functions of the ABS. No individual information collected under the Census and Statistics Act 1905 is provided back to the Registrar or ATO for administrative or regulatory purposes. Any discussion of data limitations or weaknesses is in the context of using the data for statistical purposes, and is not related to the ability of the data to support the ABR or ATO's core operational requirements. Legislative requirements to ensure privacy and secrecy of these data have been followed. Source data are de-identified and so data about specific individuals or firms has not been viewed in conducting this analysis. In accordance with the Census and Statistics Act 1905, results have been treated where necessary to ensure that they are not likely to enable identification of a particular person or organisation.

Internet Appendix for "Insurance versus Moral Hazard in Income-Contingent Student Loan Repayment"

Tim de Silva[1]

For Online Publication Only

Contents

[1] Stanford Institute for Economic Policy Research and Graduate School of Business, tdesilva@stanford.edu.

Appendix A. Comparison with Existing Literature

A.1 Literature on Labor Supply

The literature on labor supply is extremely vast (see Blundell and MaCurdy 1999 and Keane 2011 for reviews) and can be divided into four strands (Chetty et al. 2012): the first uses data on hours worked to measure labor supply; the second uses income reported on tax returns to measure labor supply; the third also uses tax data, but focuses on top earners; and the fourth studies differences in hours worked in response to cross-sectional variation, such as variation in tax rates across countries. Because I identify ϕ using bunching in HELP income, it can also be interpreted as a reported income elasticity that aggregates both hours and non-hours responses (Feldstein 1999). Therefore, Figure A20 shows the distribution of labor supply elasticities estimated among studies in these first two strands of literature, which have the closest structural interpretation to ϕ. My baseline estimate of 0.15 is similar to the median of these estimates, 0.14. However, none of these studies explicitly account for optimization frictions, although some examine longer-run responses that might be less affected by such frictions. Assuming that these estimates do not account for frictions, the closer analog in my setting to these estimates would be my frictionless estimate of 0.003, which is smaller than most estimates.

There are several reasons why optimization frictions might be larger in my setting, making the frictionless elasticity smaller. First, my sample of individuals differs from the samples in most prior studies: they are college graduates early in their life cycles. These individuals are more likely to work in salaried jobs with less hourly flexibility and a less direct mapping between labor supply and income. Second, the variation that I exploit is the discontinuity in repayment rates at the threshold. As a result, the estimated elasticity applies to individuals with incomes near this threshold, which is around the median income. This suggests that my estimated elasticity should be smaller, given that I do not study high-income individuals, who typically have higher estimated elasticities (Gruber and Saez 2002). Finally, I cannot identify extensive margin responses, which are large in some populations such as married women (Saez et al. 2012). However, the individuals in my sample are likely to be less willing to make extensive margin adjustments, given that doing so would presumably have costs that would exceed the benefits of delayed debt repayment.

This paper builds on this extensive literature on labor supply in two ways. First, it empirically characterizes how labor supply responds to income-contingent repayment, which creates dynamic trade-offs that taxes do not. My finding that borrowers reduce their labor supply to locate below the repayment threshold, which, unlike a tax, increases liquidity more than wealth, connects this literature with evidence that consumption of indebted households responds to liquidity more than wealth (Ganong and Noel 2020).[2] Second, I estimate the first (to my knowledge) model of labor supply with both time- and state-dependent adjustment. In this model, the choice of labor supply is *dynamic* for two reasons: these optimization frictions and income-contingent debt repayment. In this sense, my contribution is analogous to that of Einav et al. (2015), who show that a dynamic model of drug expenditure is necessary for replicating bunching at the Medicare Part D "donut hole."

A.2 Literature on Labor Income Risk

A growing literature uses administrative data to estimate parametric models of labor income risk (see e.g., Guvenen et al. 2021; Catherine 2022). These income processes generally contain a richer set of stochastic shocks than those that individuals face in my model, which I omit because of computational constraints that arise with an endogenous income process. Nevertheless, it is instructive to compare my parameter estimates with those in the baseline specification from Guvenen et al. (2022), who estimate a similar model with exogenous income using US data.

My estimate of the standard deviation of the individual fixed effect is 0.57, lower than the 0.77 in Guvenen et al. (2022). This primarily reflects the cross-sectional standard deviation of income at age 22 being approximately 20% lower in Australia than in the US. Additionally, I estimate a standard deviation of transitory shocks that is approximately 30% smaller, which reflects the combination of two forces. First, the cross-sectional variance of income is lower, and the 10th/90th percentiles of income growth are less dispersed in Australia. Second, the fact that labor supply is endogenous implies that some transitory variation in income arises endogenously from labor supply adjustments rather than from transitory wage shocks.[3] Last, my estimate of the standard deviation of permanent shocks is approximately three times as large. In addition to differences in data, this primarily reflects that I estimate $\rho = 0.93$ rather than imposing $\rho = 1$. This lower ρ partly reflects the heterogeneity in income profiles across education groups, which requires a larger variance of permanent shocks to match the percentiles of 5-year income growth.

[2] See Moreno and Slavov (2024) for complementary recent evidence that labor supply responses to liquidity more than wealth.

[3] The fact that labor supply endogenously creates more volatility in income reflects the fact that GHH preferences have no wealth effects.

Appendix B. Theoretical Appendix

B.1 Debt and Tax Effects of Income-Contingent Loans

Consider an individual with HELP debt, D, who chooses consumption, c, and labor supply, ℓ, to maximize the discounted sum of utility subject to a standard budget constraint and the HELP repayment contract. This problem can be formulated recursively as follows:

$$V(A,D) = \max_{c,\ell} u(c,\ell) + \beta \int V(A',D')dF_{w'|w}$$

subject to:

$$c + A' = AR + y - d(y,D), \quad y = w\ell,$$
$$D' = (1+r_d)D - d(y,D), \quad w' = g(w,\omega), \quad \omega \sim F_\omega,$$

where $d(y,D)$ denotes the required debt payment that depends on income and debt. I assume throughout that utility is increasing in consumption, $u_c > 0$, decreasing in labor supply, $u_\ell < 0$, d is differentiable in both arguments, and the initial debt, D, is sufficiently high such that $D' > 0$. The first order condition for labor supply is:

$$-\frac{u_\ell}{u_c w} = \underbrace{(1-d_y)}_{\text{tax effect}} - \underbrace{\beta d_y \frac{\mathbf{E}V_{D'}}{u_c}}_{\text{debt effect}}.$$

This equation shows that income-contingent debt has two effects on labor supply. The first term captures that income-contingent repayments discourage labor supply by reducing the return on the marginal unit of labor supply, just like a tax. The second effect is specific to debt: increasing labor supply reduces the stock of future debt. If the value function decreases in debt, $V_{D'} < 0$, the debt effect implies that individuals may choose to locate above the threshold if the marginal value of repaying their debt is sufficiently high.

The first order condition for labor supply can be rewritten as:

$$-\frac{u_\ell}{w} = u_c + d_y\left(-\beta\mathbf{E}V_{D'} - u_c\right).$$

The previous expression shows that for the debt effect to dominate and make individuals locate above the repayment threshold, the (discounted) marginal value of reducing debt must be greater than the marginal utility of consumption. This is unlikely to be the case because HELP debt has a zero real rate, which means it is the lowest-cost source of borrowing that individuals can access. More formally, this can be shown as follows. Assume that debt repayment, d, is only a function of D

when debt is repaid:

$$d(y, D) = \tilde{d}(y) * \mathbf{1}_{\tilde{d}(y)<(1+r_d)D} + D * \mathbf{1}_{\tilde{d}(y)\geq(1+r_d)D}.$$

This is the case for all income-contingent loans, and it implies that

$$d_D = \mathbf{1}_{\tilde{d}(y)\geq(1+r_d)D}.$$

Given that the envelope theorem implies that

$$V_D = -d_D u_c + \beta(1 + r_d)\mathbf{E}V_{D'},$$

combining the last two lines gives the following result:

$$\beta(1 + r_d) < 1 \implies -V_D \leq u_c.$$

In other words, if borrowers' private discount rate is below the (gross) interest rate on debt, consumption is more valuable than debt repayment, and individuals will not locate above the repayment threshold. The fact that individuals can make voluntary repayments but many do not supports this claim: if the marginal value of reducing debt was higher than consumption, more individuals should make voluntary payments.

Appendix C. Empirical Appendix

C.1 Additional Institutional Details

Timing and collection of HELP repayments. Individuals can make compulsory HELP repayments, which are the repayments calculated according to the HELP repayment schedule when the individual's tax returns are filed, or voluntary HELP repayments, which are additional repayments made at any time. If individuals are working, they are required to advise their employer if they have HELP debt. The employer will then withhold the corresponding compulsory repayment amounts from an individual's pay throughout the year based on the individual's wage or salary. Based on discussions with the ATO, most employers use an ATO-approved payroll software that calculates withholding amounts using the Tax Withheld Calculator, which effectively computes withholding amounts by converting the wage (or salary) paid from whatever frequency it is paid at to an annual frequency and applying the HELP repayment schedule. These withheld amounts are used to cover any compulsory repayments due when the tax return is filed. The tax year in Australia runs from July 1st to June 30th (e.g., the 2023 income tax year runs from July 1st, 2022 to June 30th, 2023), and tax returns must be filed by October 31st. After tax returns are filed, the difference between the total amount withheld and the actual amount due results in an amount that is paid or refunded. Additional payments are due by November 21st; most refunds are issued within 50 days of the tax lodgement. This withholding procedure is identical to the procedure used for income tax withholding.

On June 1st, HELP debts are subject to indexation, which refers to increasing the outstanding debts based on the indexation rate. The indexation rate is the nominal interest rate on HELP debt, which is based on the year-on-year quarterly CPI calculated with the March quarter CPI. It is calculated by dividing the sum of the CPI for the four quarters ending in March of the current year by the sum of the index numbers for the four quarters ending in March for the preceding years.[4] For most individuals, indexation occurs prior to the deduction of compulsory repayments because these repayments are deducted at the time of tax filing, which generally occurs between July 1st and October 31st. This is true even if an employer withholds repayments, as these repayments are not applied until the individual's tax return is filed.

Other changes to HELP repayment schedule. Since HELP was introduced in 1989, there have been several changes to the repayment schedule detailed in Ey (2021). In the early years of the program, changes were more common: the schedule changed in 1991, 1994, 1996, and 1998. However, after 1998, there have been only two changes: the 2005 policy change that I study and a 2019 policy change that was phased in over two years. The fact that there have been several changes to the HELP repayment threshold is not ideal because it implies that the model will underestimate long-run labor supply responses: in the model, the policy change is unexpected and permanent, while empirically, individuals may expect other changes in the future that attenuate their responses. However, the size of this bias is likely small because news articles written at the time of the policy change suggest that the policy change was expected to last for several years (e.g., Marshall 2003). In contrast, empirically, I find that there is persistence of bunching below the repayment threshold for only around three years, likely shorter than when individuals expected a subsequent policy change. The same logic applies if the policy change was anticipated: because there is not a lot of persistence in individuals' responses, it is unlikely that they would not respond even if they expected a policy change in a few years.[5]

Discount for upfront and voluntary payments. In prior years, HELP provided discounts to individuals who paid their debt balances upfront and discounts for voluntary repayments. The upfront payment discount took the following values: 15% from 1989-1992, 25% from 1993-2004, 20% from 2005 to 2011, 10% from 2012 to 2016, and 0% after 2016. Unfortunately, *ALife* does not allow me to identify upfront payments, so I do not include this margin in the model. The fact

[4] See here for additional details.

[5] Figure 3 shows little evidence of anticipation in the years leading up to the policy change.

that most upfront payments came from high-income individuals with family support (Norton 2018) suggests this is likely to bias my results in one of two ways. On the one hand, existing literature finds taxable income elasticities increase with income (e.g., Gruber and Saez 2002), which would suggest the model understates the moral hazard created by income-contingent repayment. On the other hand, the probability of repayment is higher for high-income individuals. Given labor supply responses decrease with the probability of repayment, this suggests the model overestimates the moral hazard income-contingent repayment creates, reinforcing my qualitative conclusions. Nevertheless, the fact that aggregate upfront payments have been low and stable despite the variation in discounts (Figure A2) suggests any bias from omitting this margin is likely to be small.

The discount for voluntary repayments took the following values: 0% from 1989-1994, 15% from 1995-2004, 10% from 2005-2011, 5% from 2012-2015, and 0% after 2015. Voluntary repayments cannot be precisely estimated in *ALife*. The fact that I do not model voluntary repayments likely leads to an upward bias in the estimate of the labor supply elasticity: the benefit of locating below the repayment threshold is even higher in a model with an option for voluntary repayments because doing so allows any payments individuals make to be classified as voluntary and thus subject to a discount. Nevertheless, this bias is likely small because voluntary repayments are uncommon for most borrowers (Norton and Cherastidtham 2016). In fact, personal finance websites suggest that young HELP debtors should avoid making voluntary repayments if they have credit card or personal debts and that if a debtor earns below the threshold, voluntarily paying off HELP debt is probably not the best use of money (MoneySmart 2016).

Wage-setting in Australia. There are three wage-setting methods in Australia. The first method is through award-based wages, in which centralized bodies set the minimum terms and conditions for employment, including a minimum wage. The primary body responsible for setting these conditions is the Fair Work Commission, which operates at national level. The second method is through enterprise agreements, which set a rate of pay and conditions for a group of employees through negotiation. This method of wage setting is analogous to that used by labor unions in the US. Finally, individual arrangements set wages and conditions for employees on an individual basis. Individual arrangements and enterprise agreements are the dominant forms of wage-setting, accounting for approximately 40% each of total wage-setting arrangements, while award-based wages make up approximately 20%.[6]

C.2 Comparison of Institutional Environments in Australia and US

This section describes similarities and differences between Australia and the US, summarized in Table A1. Although these countries are similar in many ways, some institutional differences

[6]See, for example, here.

are important when considering whether welfare gains from income-contingent repayment would generalize in the US.

The first notable difference is the cost of higher education: the student contribution at a public undergraduate institution for a Commonwealth Supported Place in Australia is around $6,400 USD after subtracting the government subsidy. This is comparable to the average undergraduate tuition at a 4-year in-state public institution in the US but much smaller than tuition for a 4-year (non-profit) private degree. Unlike in the US, where many students receive scholarships and grants that reduce tuition below the "sticker price", this is extremely rare in Australia. In addition to differences in tuition, the cost of room and board and books and supplies are slightly higher in the US. These higher costs contribute to the second difference between Australia and the US: the amount individuals borrow from government-provided student loans. In Australia, this is around $20,000 on average, while in the US, it's around $50,000 (Catherine and Yannelis 2023). The fact that debt balances are higher in the US means that the scope for welfare gains from optimizing contract design is even larger, as shown in Table 9. However, the higher loan balances also reflect that undergraduate degrees last a year longer in the US and, more importantly, that student loans in Australia can only be used to cover tuition.[7] Although the latter is useful for identification, as discussed in Section 1.3, it implies that borrowers in the US have more flexibility to adjust their borrowing using discretionary expenses, such as room and board. This introduces scope for ex-ante moral hazard, in which individuals who anticipate low incomes borrow more in anticipation of low repayment. Quantifying the strength of this force is an important task for future research because it could undermine the effectiveness of income-contingent repayment in the US. It is also especially relevant for the equity contracts studied in Section 4.2.6, which create large incentives to adjust initial debt balances.

Like in Australia, the US government is the only provider of income-contingent loans and these loans are not dischargeable in bankruptcy. However, in the US, the government offers non-income-contingent contracts, and an active private market provides financing to high-income borrowers at lower rates (Bachas 2019). Both of these features are useful for my empirical analysis, as discussed in Section 1.3, and the former is not an issue for my normative analysis since I focus on the design of a single government-provided financing contract. In contrast, the presence of a private market implies that the degree of insurance that can be provided by income-contingent repayment in the US is limited: trying to collect repayments quickly from high-income borrowers to finance reduced payments from low-income borrowers may lead private lenders to cream-skim high-income borrowers with more favorable financing terms.

An additional difference between Australia and the US is that HELP loans are significantly more

[7] To finance nontuition expenses, students on income support can use a Student Start-Up Loan, but these loans only supported fewer than 100,000 borrowers in 2020–21. All other students must self-finance these expenses, which they generally do by using credit cards or taking jobs.

Table A1. Comparison of Australia and US

Feature of Environment	Australia	US
Cost of Higher Education		
Public Undergraduate Tuition Cost	$2,700–$10,100 USD per year for CSPs	$9,500 USD per year for 4-Year In-State $39,000 USD per year for 4-Year Private Nonprofit
Prevalence of Scholarships	Rare	Common
Cost of Books and Supplies	$850 USD per year	$1,200 USD per year
Cost of Room and Board	$9,000 USD per year	$12,000 USD per year
Total Cost of Attendance	$15,850 USD per year	$22,700 USD per year
Bachelors Degree Length	3 Years	4 Years
Financing of Higher Education		
Initial Student Debt Borrowed	$8,100–$30,300 USD	$51,800 USD (Average)
Uses of Student Debt	Tuition only	Tuition, textbooks, fees, room and board
Provider of Income-Contingent Loans	Government	Government
Eligibility for Income-Contingent Loans	Australian and NZ citizens, permanent humanitarian resident	US citizens, permanent residents, eligible non-citizens
Interest Rate on Debt	CPI	~2% above T-Bill rate
Student Debt Dischargeable	No	No
Other Contracts Available	No	Yes
Private Financing Available	No	Yes
Government-Regulated Tuition	Yes	No
Enrollment Caps	Yes (for CSPs)	No
Student Population		
% of Population with Undergraduate Degree	38%	32%
% of Undergraduates at Private Universities	6%	26%
% of Undergraduates from Abroad	16%	5%
% of Current Students Employed	50%	40%
% Dropout within First Year	20%	33%
Income Distribution and Taxes/Transfers		
Median Personal Income	$33,500 USD	$40,500 USD
Poverty Line for Single Individual	$16,200 USD	$14,580 USD
Gini Coefficient for Income	0.32	0.38
Marginal Tax Rate at Average Income	41%	41%
Heathcote et al. (2017) Tax Progressivity	0.133	0.184
1-Month Individual UI Replacement Rate	23%	35%
Union Membership Rate	13.7%	10.3%

Notes: The sources for various statistics are shown as hyperlinks. All statistics are computed in the most recent year available.

subsidized than student loans in the US because of the zero real interest rate. A less subsidized contract, such as those in the US, would only draw in individuals who place higher values on education. If the structural parameters governing labor supply are correlated with individuals' valuation of education, such a contract could generate different labor supply responses—this would be selection on moral hazard or an anticipated effort effect in the language of Karlan and Zinman (2009). Ex-ante, the sign of this correlation is unclear: individuals who place a higher value on education may be more motivated by non-pecuniary factors, which would lead to a negative correlation. Alternatively, these individuals may value education more because they have a higher labor supply elasticity and, thus, are more willing to work hard in response to higher wages, generating a positive correlation. Because of this concern, my counterfactual analysis focuses on repayment contracts with a similar fiscal cost to HELP. However, the caveat of this approach is that it limits the applicability of this analysis to the US, which provides a smaller subsidy.

The final important difference between the structure of higher education in Australia and the US is that the Australian government places caps on tuition at public universities[8] and has enrollment caps for Commonwealth Supported Places (the students who receive a government contribution to their tuition).[9] Because tuition is not government-regulated in the US, universities respond to changes in government subsidies by changing tuition, which is known as the "Bennett hypothesis" (Kargar and Mann 2023). In principle, universities could respond similarly to the adoption of government-provided income-contingent contracts, but, as my normative analysis shows, such contracts can be implemented even with the same subsidy level (i.e., fiscal cost) as fixed repayment contracts. Nevertheless, universities could still respond by changing tuition to select students with differential subsidies between the two types of repayment contracts. With no enrollment caps, universities could admit many borrowers with large subsidies, increasing the fiscal cost of income-contingent repayment to the government.

The bottom of Table A1 presents summary statistics on the income distribution and the social insurance system in Australia and the US. Median income and income inequality are lower in Australia: Australia has a Gini coefficient around halfway between France and the US. The personal income tax schedules are similar in terms of average level and progressivity, but Australia has a lower unemployment benefit replacement rate than the US, one of the lowest among OECD countries. Overall, Australia and the US are broadly similar in these aggregate statistics, suggesting differences in the institutional structure of higher education are more important when considering the applicability of my results to the US.

[8] Private institutions play a relatively small role in Australia, comprising only 3 out of the country's 42 universities and 6% of the domestic enrollment share as of 2021. These institutions are slightly more popular among international students, with 11.7% of the enrollment share. Private institutions are much more expensive than public ones, especially for domestic students, and primarily compete by offering more niche products.

[9] An exception is that during 2012–2017, these caps were not in place and the system was "demand-driven" (D'Souza 2018; Norton 2019).

C.3 Data and Variable Construction

C.3.1 *ALife*

ALife provides access to a 10% random sample for approved projects. My code and analysis were tested on this sample and then were executed on the population sample by research professionals at *ALife*. The remainder of this section provides additional details on variable definitions based on the underlying variables that I construct. For description of these underlying variables, see the following link: https://alife-research.app/research/search/list. Variable definitions are presented in Python 3.9, where df refers to the underlying *ALife* dataset as a Pandas DataFrame. When variables are missing from *ALife* in a given year, they are replaced with zero unless otherwise mentioned in the text.

Demographic variables. Age is defined as c_age_30_june. Gender is defined based on c_gender. Additional demographic variables for whether an individual files a tax return electronically, has a child, or has a spouse are defined as follows:

```
df['electronic'] = df['c_lodgement_type'].isin(['MYTAX','ETAX']).astype(
                                int)
df['has_child'] = (df['c_depend_child'].fillna(0) > 0).astype(int)
df['has_spouse'] = (df['sp_status_reported'] != '0_no_information').astype
                                (int)
```

Salary & Wages. Defined as i_salary_wage. This item is technically reported by taxpayers, but it is third-party reported in the sense that the ATO receives pay-as-you-go payment summary data from employers that includes this item. This item is pre-filled if the taxpayer files electronically and the ATO cross-checks discrepancies between taxpayer- and employer-reported values.

Taxable Income. Defined as ic_taxable_income_loss.

HELP Income. The definition of HELP income has changed since the introduction of HECS in 1989. For the 1989 to 1996 Australian tax years, HELP income was equal to taxable income. Between 1996 and 1999, net rental losses were added back. Between 2000 and 2005, net rental losses and total reportable fringe benefits amounts were added back. Between 2006 and 2009, net rental losses, total reportable fringe benefits amounts, and exempt foreign employment income were added back. After 2010, net rental losses, total reportable fringe benefits amounts, exempt foreign employment income, net investment losses, and reportable superannuation contributions were added back. In *ALife*, I construct this variable as follows:

```
df['help_income'] = np.maximum(df['ic_taxable_income_loss'], 0)
adds = ['help_income']
```

```
if yr >= 2000:
  adds += ['it_rept_fringe_benefit']
if yr >= 2006:
  adds += ['isn_fsi_exempt_empl']
if yr >= 2010:
  adds += ['it_property_loss', 'it_invest_loss',
      'it_rept_empl_super_cont']
df[adds] = df[adds].fillna(0)
if yr >= 2000:
  df['it_rept_fringe_benefit'] *= ((df['it_rept_fringe_benefit'] >=
                    fringeb_tsh[yr]).astype(int))
df['help_income'] = df[adds].sum(axis = 1)
```

In this variable definition, `fringeb_tsh` refers to the reporting threshold for fringe benefits, which varies by year. This variable definition is not a perfect replication of HELP income due to a lack of data availability on certain items from the ATO. However, discussions with *ALife* suggest that any error in measurement is likely to be relatively small. Additionally, I find quantitatively similar results across years in which there is a change in the HELP repayment definition, suggesting that changes in the components added back to taxable income are not driving my main results.

Labor Income and Wage-Earner.

```
df['psi_b9'] = df['i_attributed_psi'].fillna(0)
df['psi_b14'] = df['is_psi_net'].fillna(0)
df['pship_b13'] = df[['pt_is_pship_dist_pp', 'pt_is_pship_dist_npp']].
                    fillna(0).sum(axis = 1)
df['solet_b15'] = df[['is_bus_pp', 'is_bus_npp']].fillna(0).sum(axis = 1)
df['wage_earner'] = (np.abs(df[['psi_b9', 'pship_b13', 'solet_b15']]).max(
                    axis = 1) == 0).astype(int)
laborvars = ['i_salary_wage', 'i_allowances', 'psi_b9', 'psi_b14',
        'pship_b13', 'solet_b15']
df['labor_income'] = df[laborvars].fillna(0).sum(axis = 1)
```

Interest & Dividend Income.

```
df['interest_dividend'] = df[['i_interest', 'i_div_frank', 'i_div_unfrank'
                    ]].sum(axis = 1)
```

Capital Income.

```
capitalvars = ['i_annuities_txd', 'i_annuities_untxd',
        'i_annuities_lsum_txd', 'i_annuities_lsum_untxd',
        'i_super_lsum_txd', 'i_super_lsum_untxd',
```

```
        'i_interest', 'i_div_frank', 'i_div_unfrank',
        'pt_is_trust_dist_npp', 'pt_is_frank_dist_trust_npp',
        'is_cg_net', 'is_net_rent']
df['capital_income'] = df[capitalvars].fillna(0).sum(axis = 1)
```

Net Deductions.

```
df['net_deduc'] = -(df['help_income'] - df[['labor_income', '
                                    capital_income']].sum(axis = 1))
```

HELP Debt and Repayment. HELP Debt and HELP Repayment correspond to the variables `help_debt_bal` and `hc_repayment`, respectively.

Superannuation balances. Defined as `sb_mem_bal`.

Occupation-level measure of evasion. The sample of individuals used to calculate this measure of evasion is the *ALife* 10% random sample of individuals in the population *ALife* dataset who satisfy the sample selection criteria in Section 1, are wage-earners, and have annual salary and wages greater than one-half the legal minimum wage times 13 full-time weeks (Guvenen et al. 2014). The evasion measure is then computed as the share of all workers in each occupation, `c_occupation`, who receive income from working in the form of allowances, tips, director's fees, consulting fees, or bonuses, which are reported jointly in `i_allowances`. This item is subject to the same reporting requirements as Salary & Wages.

Indicator variable for switching occupations. Equals one if the value of `c_occupation` changes from one year to the next for a given individual.

C.3.2 MADIP

MADIP provides access to population-level data on health, education, government payments, income and taxation, employment, and population demographics (including the census) over time for approved projects. I obtained access to the datasets from the ATO and the 2016 Census of Population and Housing, which I merge using a unique identifier known as the MADIP Spine. Based on the 2016 Census of Population and Housing, I construct the following variables.

HELP Income. Computed using same definition as in *ALife*.

Hours Worked. I measure hours worked using `HRSP`, which corresponds to individuals' reported hours worked in all jobs during the week before the census night.

Housing Payment–to–Income Ratio. This is calculated by annualizing monthly mortgage payments from the census files, `MRED`, and weekly rent payments, `RNTD`, by multiplying by 12 and 52, respectively. I adjust for inflation, converting these to 2005 AUD, using the HELP threshold indexation rate. I define total housing payments as the sum of the two. For the majority of individuals, only one is positive. I then divide by HELP Income to obtain the payment-to-income ratio.

C.3.3 HILDA

I construct the following variables from HILDA, which is publicly available.

Hourly Flexibility: panel measure. Hourly flexibility is measured as the standard deviation of annual changes in log hours worked per week across all jobs, `jbhruc`. Before computing this measure at the occupation-level, I restrict the sample to individuals in the 2002–2019 HILDA survey waves who satisfy the following conditions: (i) report being employed; (ii) earn a positive weekly wage; (iii) do not switch occupations between two subsequent years; and (iv) are between ages 23 and 64. Prior to computing the standard deviation, I winsorize annual changes in log hours at 1%–99%. The standard deviation within each occupation is computed with longitudinal survey weights.

Hourly Flexibility: cross-sectional measure. I construct an alternative measure of hourly flexibility as the cross-sectional standard deviation of log hours worked per week across all jobs, `jbhruc`. I impose the same sample filters as when I compute the panel-based measure. Prior to computing the standard deviation, I winsorize log hours at 1%–99%. The standard deviation within each occupation is computed with cross-sectional survey weights.

C.4 Computation of Excess Bunching Mass Statistic, b

The bunching statistic that I compute follows Chetty et al. (2011) and Kleven and Waseem (2013). First, I fit a five-piece spline to each distribution, leaving out the region $\mathcal{R} = [\$32,500, \$35,000 + X]$. When fitting this spline, I calculate the distribution in bins of $\$250$ and center the bins so that one bin is $(\$34,750, \$35,000]$. The choice of $\$32,500$ as a lower point of the bunching region represents a conservative estimate of where the bunching begins, and X is a constant intended to reach the upper bound at which the income distribution is affected by the threshold. This spline corresponds to an estimate of the counterfactual distribution absent the threshold. Formally, this counterfactual distribution is estimated by regressing the distribution onto the spline features along with separate indicator variables for each $\$250$ bin in $\mathcal{R}$.

Next, for each possible $X > 0$, I sum all the estimated coefficients on all the indicator variables and normalize by the sum of the estimated coefficients on the indicator variables below the threshold.

Taking the absolute value of this delivers an estimate of the error in the estimate of the counterfactual density, since the sum of these coefficients should be zero under a proper counterfactual density. I then choose the value of X that minimizes this absolute error. Finally, I compute the bunching statistic, b, as:

$$b = \frac{\text{observed density in } \mathcal{R}}{\text{counterfactual density in } \mathcal{R}} - 1.$$

This bunching statistic is an estimate of the excess number of borrowers below the repayment threshold relative to a counterfactual distribution in which the threshold did not exist.

Computing this bunching statistic requires specifying the area of the income distribution that is being approximated with the counterfactual density. In all figures that present the bunching statistic along with an income distribution, I approximate the counterfactual density on the same range as the plot. In all other figures, I approximate between $[\$30,000, \$40,000]$. This smaller window is chosen because in these other plots, in which I split the sample to explore heterogeneity, the income distribution is more noisy. Including points further away from this threshold causes the estimate of the counterfactual density to be poorly behaved.

Appendix D. Structural Model Appendix

D.1 Model Solution and Simulation

Discretization of state variables. I have five continuous state variables that I discretize. During retirement, liquid wealth, $A_a R$, is placed onto a grid with 101 points that varies with age. The lower point of the grid linearly decreases from the minimum allowed value based on the borrowing constraint $a = a_R$ to 0 at $a = a_T$. During working life, the grid has 31 points, and the lower point on the grid is set to the lowest value allowed by the borrowing constraint. At all ages, the upper point of the liquid wealth grid is 100 times the numeraire, which is \$40,000 AUD in 2005, and the points are on a power grid with curvature parameter 0.2.[10] Debt, D_a, is placed onto a power grid that varies with age with 11 grid points, curvature parameter 0.35, a lower value of 0, and an upper value that starts at 3.67 at $a = a_0$ and is multiplied by $1 + r_d$ in each subsequent period. Past labor supply, ℓ_a, is placed on a grid with 25 grid points. The grid is centered at 1 and ranges from 0 to 2. The upper and lower halves of the grid are split into 2 and are power grids with curvature parameter 0.5. The grid for θ_i depends on the parameter values and has 21 points. The grid is centered at zero with upper and lower bounds equal to $\pm 4\sqrt{\sigma_i^2 + \sigma_\nu^2}$. Each half of the grid is a power-spaced grid with curvature

[10]A power grid for an array of values x is a grid that is evenly spaced on the unit interval for the function $x^{k^{-1}}$, where k is the curvature parameter. The grid is the adjusted from the unit interval based on the specified lower and upper grid points.

parameter 0.7. The grid for ϵ_a is computed as the nodes from a Gauss–Hermite quadrature with 7 nodes. The remaining states are age, which is discretized on a grid that is evenly spaced from a_0 to a_T with increments of one; time, which takes two values $t \in \{2004, 2005\}$ to index before and after the policy change; the adjustment cost shock, which takes a value of zero or one; and $\mathcal{E}_i \in \{0, 1\}$.

Solution algorithm. The model has a finite horizon and a terminal condition and hence can be solved by means of backward induction in age starting with the terminal condition in the final year of life. There are two notable aspects of the solution algorithm that are crucial for getting the SMD objective function to be smooth in the set of parameters. First, no choice variables are discretized, meaning I use continuous optimization routines rather than grid searches to find optimal policies. Second, I use Gauss–Hermite quadratures to integrate all continuous shocks, which means that continuous shocks are drawn from continuous rather than discretized distributions when I simulate from the model. Additionally, when solving the model, I work with the Epstein–Zin recursive generalization of (3) in Guvenen (2009). With slight abuse of notation, I refer to this value function using the same notation as the value function in the main text.

For the period during retirement, I keep track of one value function that is a function of two states: wealth and age. The terminal condition for the model is that $E_{a_T-1} V_{a_T}^{1-\gamma} = 0$, which embeds the assumption that $u_d^{1-\gamma} = 0$, where u_d is the utility upon death. This assumption is standard in life cycle models with recursive preferences.[11] Starting with this condition, I then solve the model in prior periods by finding the optimal consumption-saving choices using a golden-section search with boundaries set based on the borrowing constraint and positive consumption. I continue this backward induction until $a = a_R - 1$.

During working life, I keep track of two value functions that are solved separately for each $\mathcal{E}_i \in \{0, 1\}$. I describe how I solve for one of these, since the approach is the same, with the only difference that a different value of $\mathcal{E}_i$ changes the state transition equations. This backward induction during working life begins with the value function at retirement, $a = a_R$, as the terminal condition. At each age, for each of the grid points in the seven-dimensional state space that excludes the adjustment cost shock, I solve for optimal choices of savings and labor supply. I do this twice: once where I solve for savings using a golden-section search and labor supply is held fixed, and once where I solve for savings *and* labor supply using a Nelder–Mead algorithm. The bounds for the Nelder–Mead algorithm are set based on the budget constraint for assets and between 0 and 10 for labor supply. The starting point is set equal to β times cash-on-hand for assets and 1 for labor supply. I perform the Nelder–Mead up to three times, varying the starting point for labor supply, until the result passes a convergence check. The value function is then computed as the maximum from these two maximization problems, taking into account the fact that f_a is only paid when ℓ_a

[11] With $\gamma > 1$, it implies that $u_d = \infty$. Bommier et al. (2020) point out some undesirable implications of this assumption in models where mortality is endogenous, which is not the case in my model.

is adjusted. When solving for these optimal policy functions at a, I have to integrate V_{a+1} over θ_{a+1}, which depends on the stochastic shock, ν_{a+1}, and have to interpolate the value function in the continuous states. I perform the integration using a Gauss–Hermite quadrature with 9 nodes and use linear interpolation (and extrapolation, if necessary).[12] Linear interpolation is extremely accurate, which allows me to use few grid points as long as choice variables are not discretized, because the Epstein–Zin value function is approximately linear in wealth. Having solved for optimal choices and hence the value function in the seven-dimensional state space at each age, I then integrate out ω_a and ϵ_a to obtain a value function that depends on five states for each age: past labor supply, debt, permanent income, liquid savings, and t.[13] I continue this backward induction until $a = a_0$ and perform it twice for each $\mathcal{E}_i \in \{0, 1\}$.

Simulation procedure. I simulate N individuals, where q_e have debt at age 22 and $q_e = 0.9 > p_e$ so that I oversample individuals with $\mathcal{E}_i = 1$ to obtain a smaller approximation error among most of the estimation targets, which are computed among this group. To ensure comparability with the data, I then compute only the estimation targets that have observations on both individuals with $\mathcal{E}_i = 0$ and those with $\mathcal{E}_i = 1$ using all $(1 - q_e)N$ model observations for individuals with $\mathcal{E}_i = 0$ but only x observations for individuals with $\mathcal{E}_i = 1$, where x is given by:

$$\frac{x}{N(1 - q_e) + x} = p_e \Rightarrow x = N(1 - q_e)\frac{p_e}{1 - p_e}.$$

Software and hardware. The code to solve and estimate the model was compiled with the `mpiifort` compiler from the January 2023 version of Intel oneAPI. Each solution and simulation was parallelized across 768 CPUs using MPI and then double-threaded across the two threads on each CPU using OpenMP, using a total of 1536 threads on the MIT SuperCloud (Reuther et al. 2018). For a given set of parameters, each iteration of solving the model, simulating from it, and calculating the SMD objective function took approximately 30 seconds in total when parallelized across all these threads. The number of simulations, N, was chosen to be as large as possible while still being able to fit the necessary outputs in double precision in RAM of each CPU, which is 4GB.

D.2 First-Stage Calibration

This section provides a detailed description on the calibration of the parameters discussed in Section 3.2.1. Whenever possible, I calibrate parameters to match their observed values during the *ALife* sample period.

[12]When solving the model with learning-by-doing, I add a constant of 0.001 to l_{ia-1} in (5) when integrating over θ_{a+1} to prevent numerical instability.

[13]At all places where I integrate, I compute certainty-equivalents rather than expectations since I am using Epstein–Zin preferences.

Demographics. Individuals are born at age 22 (the typical age at which students graduate university in Australia), retire at age 65 (the age at which the Australian retirement pension began to be paid in 2004), and die with certainty after age 89. Survival probabilities prior to age 89 are taken from the APA life tables.[14] I calculate the cohort-specific birth rates, $\{\mu_h\}$, by constructing a dataset of individuals from *ALife* at $a = a_0$ and then calculating the fraction of individuals who are age a_0 in each year between $\underline{h}$ and $\overline{h}$. I set the number of distinct individuals to 1.6 million, which is the largest value that allows me to store simulated results from the model in double precision and stay within memory constraints.

To compute equivalence scales, I use data from the HILDA Household-Level File on the number of the adults in each household, `hhadult`, the number of children, defined as the sum of `hh0_4`, `hh5_9`, and `hh10_14`, and the age of the head of the household, `hgage1`. Following Lusardi et al. (2017), I compute the average number of adults and children for each age of the head of the household, denoted by adults_a and children_a. I then compute the equivalence scale at each age using the formula in Lusardi et al. (2017):

$$\tilde{n}_a = (\text{adults}_a + 0.7 * \text{children}_a)^{0.75}.$$

Finally, I normalize equivalence scales such that the average value is one, so that a household in the model corresponds to the size of the average household in the data:

$$n_a = \frac{\tilde{n}_a}{\sum_a \tilde{n}_a} * a_T.$$

Numeraire. The numeraire in the model is equal to \$1 AUD in 2005. There is no inflation in the model, so all empirical estimation targets, when they are compared with model values, are deflated to 2005 AUD with the indexation rates for HELP thresholds.

Interest rates. To calculate the real interest rate, I compute the average (gross) deposit interest rate in Australia in each year between 1991 and 2019, which is the time period of my *ALife* sample. I then divide these deposit rates in each year in each year by the (gross) inflation rate based on the CPI.[15] I take the geometric average of the resulting time series of real deposit rates between 1991 and 2019, which delivers $R = 1.0184$. To calculate the borrowing rate, I use the average standard credit card rate reported by the Reserve Bank of Australia between 2000 and 2019.[16] After deflating by the same CPI series and computing the geometric average, I obtain an average real credit card rate of 15.4%. Over 2000–2019, the geometric average of the real deposit rate was 0.8%, so I set

[14]See `https://aga.gov.au/publications/life-tables/australian-life-tables-2005-07`.

[15]See `https://data.worldbank.org/indicator/FR.INR.DPST?locations=AU` and `https://data.worldbank.org/indicator/FP.CPI.TOTL.ZG?locations=AU` for these two data series.

[16]See `https://www.finder.com.au/credit-cards/credit-card-statistics#interest-rates`.

$\tau_b = 15.4\% - 0.8\% = 14.6\%$.

Borrowing limit. I calculate the age-specific borrowing limit, $\{\underline{A}_a\}_{a=a_0}^{a_T}$, using data on credit card borrowing limits from HILDA. I start from the combined household level files from the 2002, 2006, 2010, 2014, and 2018 waves, which have Wealth modules that contain the total credit limit on all credit cards in the responding person's name, `crymbl`. Filtering the sample to individuals between 22 and 90, I deflate this variable to 2005 AUD and winsorize at 1%–99%. I then estimate a linear regression of this variable onto a constant and a fourth-order polynomial in age using weighted least squares, where the weights are the cross-sectional survey weights normalized to weight each year equally. Finally, I use the predicted value from this regression for each age as $\underline{A}_a$. The resulting values are:

$$\underline{A}_a = 1.402 \times 10^4 - 1401.63 * a + 33.14 * a^2 - 0.3682 * a^3 + 0.0017 * a^4.$$

Initial assets. I calculate the parameters that govern the initial asset distribution using data on asset holdings from HILDA. I start from the combined household-level files from the 2002, 2006, 2010, 2014, and 2018 waves, which have Wealth modules that contain household-level information on asset holdings. Among individuals who are lone persons (`hhtype` = 24) between ages 18 and 22, I compute liquid assets as the sum of bank account balances (`hwtbani`), cash, money market and debt investments (`hwcaini`), and equity investments (`hweqini`) minus credit card debt (`hwccdti`) and other personal debt (`hwothdi`), deflate the resulting estimates to 2005 AUD, and winsorize at 1%–99%. I split the sample into individuals with HELP debt, who correspond to $\mathcal{E}_i = 1$ in the model, and those wihout HELP debt, who correspond to $\mathcal{E}_i = 0$. I then estimate the fraction of individuals with nonpositive asset balances, $p_A(\mathcal{E}_i)$. Among the individuals in each group with posiitve asset balances, I estimate $\mu_A(\mathcal{E}_i)$ and $\sigma_A(\mathcal{E}_i)$ by fitting a normal distribution to the distribution of positive asset balances among individuals in each group, adjusting for the cross-sectional survey weights that are normalized to weight each year equally. The resulting estimates are shown in Table 2. When simulating from this distribution, I impose an upper bound equal to the largest value that I observe empirically. Additionally, because A_{ia} represents end-of-period savings, I scale A_{ia_0} by R^{-1} so that the liquid assets at $a = a_0$ in the model match the data.

Preference parameters. I set $\gamma = 2.23$ based on the results in Choukhmane and de Silva (2023).

Interest rate on student debt. I set the (net) interest rate on student debt, r_d, equal to zero, which is the case for HELP debt. In all counterfactuals that I consider, I leave this interest rate set to zero. This is done because the model does not include endogenous early repayment of debt balances. With a zero interest rate, this abstraction is without loss of generality, since borrowers have no incentive to pay their debt early.

Distribution of education levels. I set the fraction of individuals with that are borrowers, p_E, equal to the fraction of 22-year-old individuals in *ALife* who have positive debt balances (22 is the age by which most individuals have started their undergraduate degrees in Australia).

Initial student debt balances. I calculate the parameters that govern the initial debt distribution using data on HELP debt balances from *ALife*. First, I deflate the debt balances for all individual–years to 2005 AUD and then calculate the year in which each individual had her maximum real debt balance. From these debt balances, I drop observations in which (i) individuals are not classified by *ALife* as having acquired new debt balances, (ii) the maximum occurs in the year 2019, which is the final year of data, and (iii) individuals are older than 26 years old, which is the age by which most individuals have finished undergraduate studies in Australia and debt balances reach their maximum in real terms. Finally, I estimate μ_d and σ_d by fitting a normal distribution to the logarithm of these debt balances. When simulating from this distribution, I impose an upper bound equal to the largest value that I observe empirically.

Student debt repayment function. When estimating the model, I use the HELP 2004 repayment function at $t < T^*$ and the HELP 2005 repayment function at $t \geq T^*$.[17] Formally, I set $d(y,i,D,a,t) = \mathbf{1}_{a<a_R} * \min\{HELP_t(y+\max\{i,0\}) * (y+\max\{i,0\}), (1+r_d)D\}$, where

$$HELP_t(x) = \mathbf{1}_{t<T^*} HELP_{04}(x/\pi_{05}) + \mathbf{1}_{t\geq T^*} HELP_{05}(x),$$

$$HELP_{04}(x) = \begin{cases} 0 & \text{if } x \leq 25347, \\ 0.03 & \text{else if } x \leq 26371, \\ 0.035 & \text{else if } x \leq 28805, \\ 0.04 & \text{else if } x \leq 33414, \\ 0.045 & \text{else if } x \leq 40328, \\ 0.05 & \text{else if } x \leq 42447, \\ 0.055 & \text{else if } x \leq 45628, \\ 0.06 & \text{else,} \end{cases} \qquad HELP_{05}(x) = \begin{cases} 0 & \text{if } x \leq 35000, \\ 0.04 & \text{else if } x \leq 38987, \\ 0.045 & \text{else if } x \leq 42972, \\ 0.05 & \text{else if } x \leq 45232, \\ 0.055 & \text{else if } x \leq 48621, \\ 0.06 & \text{else if } x \leq 52657, \\ 0.065 & \text{else if } x \leq 55429, \\ 0.07 & \text{else if } x \leq 60971, \\ 0.075 & \text{else if } x \leq 64999, \\ 0.08 & \text{else,} \end{cases}$$

where π_{05} is the inflation rate used for the HELP indexation thresholds between 2004 and 2005. In counterfactuals, I consider alternative repayment contracts. In these counterfactuals, I consider repayments that are contingent only on wage income, y_{ia}, and not capital income, i_{ia}.

Income and capital taxation. In Australia, income taxes are paid on taxable income, which

[17] See https://atotaxrates.info/individual-tax-rates-resident/hecs-repayment/.

aggregates both wage income and capital income. The marginal tax rate that individuals pay increases in their income according to a schedule provided by the ATO.[18] When I estimate the model, I set $\tau(y,i,t) = T_t(y + \max\{i,0\})$, where T_t is equal to the ATO 2003/04 Income Tax Formula at $t < T^*$ and the ATO 2004/05 Formula at $t \geq T^*$:

$$T_t(x) = \mathbf{1}_{t<T^*} T_{04}(x/\pi_{05}) + \mathbf{1}_{t\geq T^*} T_{05}(x),$$

$$T_{04}(x) = \begin{cases} 0 & \text{if } x \leq 6000, \\ 0.17 * (x - 6000) & \text{else if } x \leq 21600, \\ 2652 + 0.3 * (x - 21600) & \text{else if } x \leq 52000, \\ 11952 + 0.42 * (x - 52000) & \text{else if } x \leq 62500, \\ 16362 + 0.47 * (x - 62500) & \text{else,} \end{cases}$$

$$T_{05}(x) = \begin{cases} 0 & \text{if } x \leq 6000, \\ 0.17 * (x - 6000) & \text{else if } x \leq 21600, \\ 2652 + 0.3 * (x - 21600) & \text{else if } x \leq 58000, \\ 13752 + 0.42 * (x - 58000) & \text{else if } x \leq 70000, \\ 18792 + 0.47 * (x - 70000) & \text{else,} \end{cases}$$

where π_{05} is the inflation rate used for the HELP indexation thresholds between 2004 and 2005. For individuals in retirement with $a \geq a_R$, I do not change the income tax schedule to avoid keeping track of an additional state variable. When comparing across student debt repayment policies, I eliminate taxes on capital income and adopt the following parametric income tax schedule, which Heathcote and Tsujiyama (2021) show provides a close approximation to unconstrained Mirrlees solutions, which is unlikely to be the case for the actual ATO schedule:

$$\tau(y,i,t) = y - ay^b.$$

I estimate a and b using the methodology from Heathcote et al. (2017) applied on the 2005 ATO tax schedule, which delivers $a = 1.1296$ and $b = 0.8678$.

Unemployment benefits and net consumption floor. Unemployment benefits are set equal to the payments provided by the Newstart allowance, which is the primary form of government-provided income support for individuals above 22 with low income due to unemployment. These benefits are means-tested based on income and assets. I use the formula for payments in 2005 for a single

[18] See https://www.ato.gov.au/Rates/Individual-income-tax-for-prior-years/.

individual with no children.[19] This formula is:

$$\frac{ui(y,i,A)}{26} = \begin{cases} 0 & \text{if } A \geq 153000 \text{ or } (y+\max\{i,0\})/26 > 648.57, \\ 394.6 & \text{else if } (y+\max\{i,0\})/26 \leq 62, \\ 394.6 - 0.5*(y+\max\{i,0\}-62) & \text{else if } (y+\max\{i,0\})/26 \leq 142, \\ 354.6 - 0.7*(y+\max\{i,0\}-142) & \text{else.} \end{cases}$$

When comparing across student debt repayment policies, I adopt the following smoothed specification of this formula and eliminate dependence on capital income and assets to remove the impact of changes in student debt repayments on the government budget constraint through changes in asset accumulation:

$$ui(y,i,A) = 26 * \max\left\{394.60 - y * \frac{394.60}{16863}, 0\right\}.$$

In addition to unemployment benefits, individuals receive a net consumption floor payment. This floor is needed to ensure that individuals' consumption net of labor supply disutility, $c_{ia} - \kappa\frac{\ell_{ia}^{1+\phi^{-1}}}{1+\phi^{-1}}$, remains positive in the event that they do not adjust their labor supply. The consumption floor is set equal to:

$$\underline{c}_a = \max\left\{\underline{c} - \Big(y_a + A_a + i_a - d_a - \tau(y_a, i_a, t) + ui(y_a, i_a, A_a)\Big), 0\right\},$$

where $\underline{c}$ is the minimum value of net consumption. I set $\underline{c} = \$40$ but have experimented with higher values up to $\$400$ and have found that the results remain unchanged.

Retirement pension. Individuals in retirement receive a retirement pension from the government that is based on the age pension, which is the primary form of government-provided income support for retirees in Australia. The age pension is available to individuals at age 65 and is means-tested based on assets and income. I use the formula for payments in 2005 for a single individual who is a homeowner based on assets, but I exclude means-testing on income since individuals earn no labor income in retirement. This formula is:

$$\overline{y}(A) = \begin{cases} 12402 & \text{if } A \leq 153000, \\ 12402 - 3*26*\left\lfloor\frac{A-153000}{1000}\right\rfloor & \text{else if } A \leq 312000, \\ 0 & \text{else.} \end{cases}$$

When comparing across student debt repayment policies, I remove means-testing and give everyone the full pension of $\$12402$ to remove the impact of changes in student debt payments on the government budget constraint through changes in asset accumulation.

[19] See https://melbourneinstitute.unimelb.edu.au/__data/assets/pdf_file/0006/2378733/co029_0501en.pdf.

D.3 Second-Stage Simulated Minimum Distance Estimation

Construction of estimation targets. The set of estimation targets that I use is:

1. OLS estimates of β_1 and β_2 from estimating the following equation among employed individuals between ages 22 and 64:
$$\log y_{ia} = \beta_0 + \beta_1 a + \beta_2 a^2$$
2. OLS estimates β_0^E and β_1^E from estimating the following equation among individuals that reach age 22 at $t \geq 1991$:[20]
$$\log y_{ia} = \beta_0 + \beta_1 a + \beta_2 a^2 + \beta_0^E \mathcal{E}_i + \beta_1^E \mathcal{E}_i a$$
3. Within-cohort cross-sectional variance of $\log y_{ia}$ at age 22, 32, 42, 52, and 62
4. 10th and 90th percentiles of $y_{ia+1} - y_{ia}$ and $y_{ia+5} - y_{ia}$
5. Average ℓ_{ia} among employed individuals between ages 23 and 64, which is normalized to 1 in the data
6. Real distribution of HELP income among debtholders aged 23 to 64 in 2002–2004 within \$3000 of the 2004 repayment threshold in bins of \$500
7. Real distribution of HELP income among debtholders aged 23 to 64 in 2005–2007 within \$3000 of the 2005 repayment threshold in bins of \$500
8. The following statistic, where quartiles of debt balances are calculated within each year and the number of debtholders is pooled from 2005–2018:
$$\frac{\dfrac{\text{\# of debtholders in top quartile of debt within \$500 below 2005 threshold}}{\text{\# of debtholders in top quartile of debt within \$500 above 2005 threshold}}}{\dfrac{\text{\# of debtholders in bottom quartile of debt within \$500 below 2005 threshold}}{\text{\# of debtholders in bottom quartile of debt within \$500 above 2005 threshold}}}.$$
9. Probability of bunching below the 2005 repayment threshold in 2005 conditional on bunching below the 2004 repayment threshold in 2004
10. Fraction of individuals that do not adjust their annual hours worked from HILDA
11. Kurtosis of annual changes in log hours worked from HILDA

[20] I do not allow for the possibility that the quadratic component of g_{ia} differs with $\mathcal{E}_i$. This is because *ALife* covers only 1991–2019 and does not have direct measures of education. Since I instead infer education level based on the presence of HELP debt, the oldest individual whom I observe in the sample with $\mathcal{E}_i = 1$ is around age 50–55. Without the final 5–10 years of working life, it is difficult to identify this additional parameter.

In these definitions, y_{ia} refers to the value of Salary and Wages in *ALife*, and i_{ia} refers to Capital Income defined in Appendix C.3. Because of data access restrictions, I construct the first six set of estimation targets using a 10% random sample of *ALife* data. This likely has little effect on the results because these estimation targets are very precisely estimated and are not the primary targets responsible for identifying the structural parameters of interest. For these estimation targets, I restrict to wage-earners between 22, the first age in the model, and 64, the age at which individuals retire in the model, and winsorize both y_{ia} and i_{ia} from above at 99.999% following Guvenen et al. (2014). When computing the estimation targets based on y_{ia}, I restrict to individuals who have annual salary and wages greater than one-half the legal minimum wage times 13 full-time weeks following Guvenen et al. (2014). When calculating all estimation targets in the data, I also restrict to individuals who were age 22 between 1963 and 2019 to match the cohorts simulated in the model. For the final two moments from HILDA, I restrict the sample to data in two subsequent years among individuals that are employed, earning a positive weekly wage, non-business owners, and between age 22 and 64. I also adjust for longitudinal survey weights, and compute hours worked using total reported hours worked across all jobs.

Weighting matrix. I choose the weighting matrix, $W(\Theta)$, such that the SMD objective function corresponds to the sum of squared arc-sin deviations between $\hat{m}$ and $m(\Theta)$. Specifically, I set $W(\Theta) = \text{diag}\left(w(\Theta)\right)$, where

$$w(\Theta) = \left(0.5 \times \max\left\{\underline{w}, |\hat{m}| + |m(\Theta)|\right\}\right)^{-2}.$$

This choice follows Guvenen et al. (2021) and is made because I have many estimation targets that differ greatly in scale.[21] I do not use the optimal weighting matrix because some of these targets are estimated from population-level data and thus have very small asymptotic variances that make the objective function unstable. I also follow Guvenen et al. (2021) and adjust $w(\Theta)$ so that the following blocks of estimation targets receive equal weight.

1. Block #1: Heterogeneity in bunching with debt balances, persistence of bunching below the repayment threshold, fraction of individuals immediately below and above repayment threshold prior to policy change, fraction of individuals immediately below and above repayment threshold after to policy change.

2. Block #2: All remaining estimation targets.

This is done to ensure that the first block of moments, which are most important for the structural

[21]The choice of constant $\underline{w}$ is made to ensure that the objective function remains well-behaved even as the targets become small and possibly differ in sign between the model and data. I set $\underline{w} = 0.01$ based on experimentation, but at the global optimum, this lower bound does not bind and thus does not meaningfully affect the results.

parameters of interest, get equal weight to the remaining moments, even though there are a fewer number of them.

Global optimization algorithm. I compute the value of Θ that minimizes the SMD objective function using a variant of the TikTak algorithm from Arnoud et al. (2019). I start by evaluating the objective function at 8000 pseudorandom Halton points that cover the parameter space. I then take the top 10 candidate points and perform a Nelder–Mead optimization at each of these 10 points. Finally, I use the Nelder–Mead solutions at each of these 10 points to perform a second round of 10 additional Nelder–Mead optimizations. Specifically, I rank the 10 solutions from the first set of optimizations and start the first optimization of the second round at the best point. Then, to start each of the remaining $i = 2, \ldots, 10$ optimizations, I use as a starting point the weighted average of the current candidate optimum and the ith ranked point, with the weighting function and parameters chosen exactly as in Arnoud et al. (2019). In each of these Nelder-Meads, the convergence criteria are a relative objective tolerance of 0.01 or 400 iterations. In a final polishing phase, I perform a Nelder-Mead with a tolerance of 0.001 and a maximum of 1000 iterations.

Calculation of standard errors. To apply standard asymptotic theory to calculate standard errors, I rewrite the SMD objective function as

$$\Theta^* = \arg\min_{\Theta} g(\Theta)' g(\Theta),$$

where

$$g(\Theta) = \text{diag}\left(\sqrt{w(\Theta)}\right)(m(\Theta) - \hat{m}).$$

Denote the true value of the parameters, Θ, as Θ_0. Under standard regularity conditions (e.g., McFadden 1989; Duffie and Singleton 1993),

$$\sqrt{N}\left(\Theta^* - \Theta_0\right) \xrightarrow{d} N(0, V),$$

where $\xrightarrow{d}$ denotes convergence in distribution as the number of sample observations, N, tends to infinity for a ratio of the number of model simulations to data observations, S. The asymptotic variance, V, is given by

$$V = \left(1 + \frac{1}{S}\right)\left[GG'\right]^{-1} G\Omega G' \left[GG'\right]^{-1},$$

where $G = \frac{\partial}{\partial\Theta} g(\Theta)$,

$$\Omega = \Omega_0 \Lambda, \quad \sqrt{N}\hat{m} \xrightarrow{d} N(m_0, \Omega_0),$$

$$\Lambda = \text{diag}\left(4 * c_0 * \left[\mathbf{1}_{\underline{w} \le |\hat{m}| + |m(\Theta)|} * \frac{|m(\Theta)||\hat{m}| + m(\Theta)\hat{m}}{|\hat{m}|\left(|m(\Theta)| + |\hat{m}|\right)^2} + \mathbf{1}_{\underline{w} > |\hat{m}| + |m(\Theta)|} * \underline{w}^{-1}\right]^2\right),$$

all multiplication and division in the definition of Λ is performed element-wise, all quantities are evaluated at Θ_0, and c_0 is a vector that accounts for the reweighting of the different blocks of estimation targets discussed above. The previous two equations define the asymptotic variance of $g(\Theta)$, denoted by Ω, which is derived by means of the delta method and the asymptotic distribution of $\hat{m}$.

By the continuous mapping theorem, each component of V can be estimated by replacing population quantities with sample analogs evaulated at the SMD estimate of Θ. I estimate Ω_0 via bootstrap assuming that all off-diagonal elements are zero[22] and compute G using two-sided finite differentiation.[23] The standard errors for Θ^* are then $\sqrt{N^{-1}\text{diag}(\hat{V})}$.

D.4 Description of Repayment Contracts

25-year Fixed Repayment. For a borrower i at age a, the required payment on a fixed repayment contract is:

$$d_{Fixed}(a, D_{ia}) = \begin{cases} 0, & \text{if } a < a_S \\ D_{ia} * \frac{r_d}{1-(1+r_d)^{-(a_E-(a-a_0+1)+1)}}, & \text{else,} \end{cases}$$

where a_S is the first age at which payments start and a_E is the age at which payments enda. In the event that borrowers' cash-on-hand prior to making debt payments falls below $d_{Fixed}(\cdot)$, I make borrowers pay only their cash-on-hand. In this case, borrowers will also receive the consumption floor payment since they have no resources for consumption. A 25-year fixed repayment contract corresponds to $a_S = a_0$, $a_E = a_0 + 25$, and $r_d = 0\%$.

US Income-Contingent Loans. For a borrower i at age a, the required payment on the US–style income-contingent that I consider are:

$$d_{ICL}(D_{ia}, y_{ia}) = \min\{\psi * \max\{y_a - K, 0\}, (1 + r_d)D_{ia}\} * \mathbf{1}_{a \le \overline{T}}.$$

The following specifies the parameters for the different IBR contracts that I implement in the text:

- US IBR: $\psi = 10\%$, $K = 1.5 * pov$, $\overline{T} = a_R$, $r_d = 0\%$
- US SAVE: $\psi = 5\%$, $K = 2.25 * pov$, $\overline{T} = a_R$, $r_d = 0\%$
- US IBR + Forgiveness: $\psi = 10\%$, $K = 1.5 * pov$, $\overline{T} = a_0 + 20$, $r_d = 0\%$
- US SAVE + Forgiveness: $\psi = 5\%$, $K = 2.25 * pov$, $\overline{T} = a_0 + 20$, $r_d = 0\%$

[22] I cannot compute off-diagonal elements because the estimation targets are calculated from different samples, which do not all fit in the RAM of the virtual machine used to access the data.

[23] I compute the standard error of average labor supply using the hours worked reported in HILDA, after normalizing it to have a mean of one.

where pov is the 2023 US poverty line of \$14,580 USD converted into AUD by adjusting for US CPI inflation from June 2005 to January 2023 by the exchange rate in June 2005.[24] The final two versions of these contracts that I consider are the US IBR + Fixed Cap and US SAVE + Fixed Cap. These correspond to the US IBR and US SAVE contracts defined above with the modification that $d_{ICL}(D_{ia}, y_{ia})$ cannot exceed $d_{Fixed}(a, D_{ia})$ for the 25-year fixed repayment contract.

Purdue Income-Sharing Agreement. For a borrower i at age a, the required payment is:

$$d_{ISA}(a, y_{ia}) = \begin{cases} 0, & \text{if } a > T, \\ \psi * y_{ia}, & \text{else.} \end{cases}$$

In this expression, T_{ISA} is the term of the ISA contract and ψ is the income-share rate. For the Purdue ISA contract, I set $T = 9$ and $\psi = 4\%$, which closely matches that of the ISAs provided by Purdue University in 2016–2017 (Mumford 2022).

D.5 Computation of Welfare Metrics

Equivalent variation. Let s_0 be the vector of four stochastic initial conditions in the model: education-level $\mathcal{E}_i$, permanent income δ_i, assets, A_{ia_0}, and debt balances D_{ia_0}. Let $s_0(\pi)$ be the same vector with initial assets $A_{ia_0} + \pi$ instead of A_{ia_0}. Denote the value function at $a = a_0$ and initial states s_0 with education level $\mathcal{E}_i = E$ under repayment policy p as $V_p(s_0 \mid \mathcal{E}_i = E)$, and denote the joint conditional distribution of the four stochastic initial conditions as $F(\mathbf{s_0} \mid \mathcal{E}_i = E)$.

The *equivalent variation* of policy p, π_p, relative to the 25-year fixed repayment contract is computed as the fixed point of the following equation in π:

$$\int V_p(s_0 \mid \mathcal{E}_i = 1) dF(s_0 \mid \mathcal{E}_i = 1) = \int V_{\text{25-Year Fixed}}(s_0(\pi) \mid \mathcal{E}_i = 1) dF(s_0 \mid \mathcal{E}_i = 1).$$

This left-hand side of this equation corresponds to the expected utility of random consumption and labor supply streams under repayment policy p to an agent with education level $\mathcal{E}_i = 1$ who is "behind the veil of ignorance" with respect to s_0 and views the realization of these states as risk. The right-hand side corresponds to the same quantity calculated under the 25-year fixed repayment contract when borrowers receive a deterministic cash transfer of π at $a = a_0$. I compute this fixed point using a standard bisection root-finding algorithm.

Consumption-equivalent welfare gain. Let $V_p(s_0 \mid \mathcal{E}_i = E)$ and $F(\mathbf{s_0} \mid \mathcal{E}_i = E)$ denote the same quantities as above. Let $V_p^g(s_0 \mid \mathcal{E}_i = E)$ denote $V_p(s_0 \mid \mathcal{E}_i = E)$ evaluated in a model in which, for

[24] This equals \$12,320, which is almost identical to the \$11,511 poverty line reported by the Melbourne Institute.

all ages a, borrowers i get to consume $(1+g)c_{ia}$. The *consumption-equivalent gain* of policy p, g_p, relative to the 25-year fixed repayment contract is computed as the fixed point to the following equation in g:

$$\int V_p(\mathbf{s}_0 \mid \mathcal{E}_i = 1) dF(\mathbf{s}_0 \mid \mathcal{E}_i = 1) = \int V^g_{\text{25-Year Fixed}}(\mathbf{s}_0 \mid \mathcal{E}_i = 1) dF(\mathbf{s}_0 \mid \mathcal{E}_i = 1).$$

This metric corresponds to the value of g that would make borrowers with $\mathcal{E}_i = 1$ indifferent between having to (i) pay their debt under repayment policy p and (ii) pay their debt under 25-year fixed repayment *and* having their consumption increased by $g\%$ in every state during their lifetime. I compute this fixed point using a standard bisection root-finding algorithm. This metric takes the perspective of a hypothetical borrower who does not know her initial values of D_{ia_0}, A_{ia_0}, and δ_i, and views the realization of these states as risk.

D.6 Computation of Constrained-Optimal Repayment Contracts

Solving (15) is numerically challenging, especially when higher-dimensional contracts are considered, because it is a nonlinear constrained optimization problem in which the objective and constraints do not have closed forms. I use a combination of a standard barrier method in numerical optimization (Nocedal and Wright 2006) and a global optimizer. Specifically, I set the objective function in (15) to an extremely large value in the event that the first constraint, which corresponds to the government budget constraint, is violated by more than a tolerance of \$1. I then perform the minimization of this objective function using the TikTak optimizer from Arnoud et al. (2019). Due to memory and computational constraints, I set N = 50,000 when solving for constrained-optimal policies and only simulate individuals with $\mathcal{E}_i = 1$ (individuals with $\mathcal{E}_i = 0$ do not affect the planner's problem).

D.7 Insurance versus Redistribution Decomposition

This section describes the details underlying the analysis in Section 4.2.7. I begin by discretizing the following initial states, which are continuous in the baseline model: $D_{ia_0}, A_{i_a 0}$ and δ_i. For each initial state X, I discretize it to take on one of two values, X_- and X_+, so that the total number of initial states is $\mathcal{T} = 2^3 = 8$. I set these values as follows:

- $A_- = E(A_{ia_0} | A_{ia_0} \leq \text{median}(A_{ia_0})) = \$246.13$,
- $A_+ = E(A_{ia_0} | A_{ia_0} > \text{median}(A_{ia_0})) = \$8772.68$,
- $D_- = E(D_{ia_0} | D_{ia_0} \leq \text{median}(D_{ia_0})) = \$6859.07$,
- $D_+ = E(D_{ia_0} | D_{ia_0} > \text{median}(D_{ia_0})) = \$28026.08$,

- δ_-, δ_+ = gridpoints on a two-dimensional Gauss-Hermite quadrature grid that approximates δ_i

This discretization of the initial states provides a parsimonious representation of the distribution of initial states while having the fewest possible number of values. In principle, the discretization could be finer, but each additional dimension makes the constrained-optimization problem that I solve significantly more complex by introducing another choice variable and constraint.

Given this discretization of the initial states, in the second step I resolve the constrained-planner's problem in (15) with two modifications. First, I introduce an additional $\mathcal{T}$ policy instruments that correspond to lump-sum transfers made at a_0 to borrowers based on their $\mathcal{T}$ possible initial states. Second, I introduce an additional $\mathcal{T}$ constraints, requiring that the government budget defined in (14) remains unchanged at each of the possible $\mathcal{T}$ initial states between a given repayment contract p and the benchmark 25-year fixed repayment contract. To evaluate the government budget at a given initial state, I simply replace the unconditional expectation in (14) with the expectation taken over all individuals with the given initial state.

D.8 Sensitivity of Welfare Gains to Model Misspecification

This section describes the details of the various models shown in the bottom half of Table 9, which represent deviations from the baseline model shown in the first row and estiamted in column (5) of Table 3.

US tax system. I change the parameters of the 2-parameter tax function defined in Appendix D.2 to the parameters defined in Section IIA of Heathcote et al. (2017).

Optimized tax system. I compute the optimized tax function by solving the constrained-planner's problem in (15) with two modifications. First, I change the objective function to be $\mathbf{E}_0(V_{ia_0})$, where the expectation is taken across all initial states. Second, I optimize over the parameters of the 2-parameter tax function from Heathcote et al. (2017) defined in Appendix D.2. The results with the optimized tax system then correspond to solve (15) with the parameters of the 2-parameter tax and transfer function already optimized.

Alternative risk and time preferences. To assess the effects of moving the RRA and EIS independently, I use the the recursive generalization of (3) in Guvenen (2009). I then change these two parameters independently, holding all others fixed.

Wealth effects on labor supply. Existing literature disagrees on the size of wealth effects on labor supply: Cesarini et al. (2017) find small wealth effects from lottery winnings in Sweden, while Golosov et al. (2023) find larger effects from lottery winnings in the US. To assess the importance of wealth effects, I adjust the flow utility in (3) to be

$$\frac{1}{\eta}\left(\frac{c_{ia}}{n_a}\right)^{\eta} - \kappa\frac{\ell_{ia}^{1+\phi^{-1}}}{1+\phi^{-1}}.$$

I set $\eta = 0.5$ following the calibration in Keane (2011).

Persistence of income risk. Because individuals can self-insure against transitory but not permanent shocks in incomplete markets, correctly estimating the persistence of income shocks is crucial for assessing the welfare impact of income-contingent repayment. Because estimates of this persistence vary between 0.8 and close to 1, depending on the degree of heterogeneity in income profiles (Guvenen 2009a), I consider alternative values of ρ, holding all other parameters fixed.

US initial debt levels. An important difference between the US and Australia is the level of initial debt that borrowers take on. In the 2019 Survey of Consumer Finances, the average initial debt among borrowers was $51,800 USD (Catherine and Yannelis 2023), while in the model, it is $17,400 in 2005 AUD ($20,500 in 2023 USD). I consider the effect of multiplying all initial debt balances by 2.51, the ratio of the previous two values.

Higher interest rate on debt. In my analysis, I set the real interest rate on debt to zero, as in HELP. However, in the US, debt balances have historically been subject to interest accumulation (although the new SAVE plan changes this). Alternatively, I consider an interest rate of 2% above the real interest rate, similar to the markup on student loans above Treasury bill rates in the US (Ji 2021) and above the Bank of England base rate in the UK (Britton and Gruber 2020).

Government discount rate. The model does not have aggregate risk, so the correct discount rate for debt repayments is the risk-free rate. I consider the effects of a higher discount rate, the risk-free rate plus a 2% risk premium.

D.9 Welfare Gains from using a Richer Contract Space

In this section, I consider the effects of solving (15) using the following three richer contract spaces:

1. Quadratic Income-Contingent Loan: $d_{ia}(\boldsymbol{\theta}) = \min\left\{\max\left\{\theta_1 + \theta_2 y_{ia} + \theta_3 y_{ia}^2, 0\right\}, D_{ia}\right\}$
2. Quadratic Income-Contingent Loan + Age: $d_{ia}(\boldsymbol{\theta}) = \min\left\{\max\left\{\theta_1 + \theta_2 y_{ia} + \theta_3 y_{ia}^2 + \theta_4 a, 0\right\}, D_{ia}\right\}$
3. Quadratic Income-Contingent Loan + Debt: $d_{ia}(\boldsymbol{\theta}) = \min\left\{\max\left\{\theta_1 + \theta_2 y_{ia} + \theta_3 y_{ia}^2 + \theta_4 D_{ia}, 0\right\}, D_{ia}\right\}$

The first contract corresponds to a smoothed version of the income-contingent loans considered in Section 4.2, in which repayments are a quadratic function of income. The latter two contracts make payments conditional on age and debt, respectively. For each of these alternative contracts, I solve the planner's problem in (15), optimizing over $\boldsymbol{\theta}$ instead of ψ and K. Figure A19 shows the results. In the baseline model, using a quadratic repayment function has no effect on the welfare gain. While making payments debt-contingent also has no effect, making payments age-contingent increases the welfare gain to a 0.94% equivalent increase in lifetime consumption. In the baseline model with a higher value of $\phi = 0.24$, where the baseline income-contingent loan leads to a welfare loss, the quadratic repayment function helps restore part of the welfare gain of income-contingent repayment. This is consistent with Shavell (1979), who shows that the unconstrained solution to (15) features some insurance because the gains from insurance are first-order while the losses from moral hazard are second-order. However, making payments age-contingent helps even further, since this allows the planner to condition payments on a variable that is correlated with the marginal value of wealth but that cannot be manipulated.

Additional Figures and Tables

Figure A1. Student Contributions and Aggregate HELP Borrowing over Time

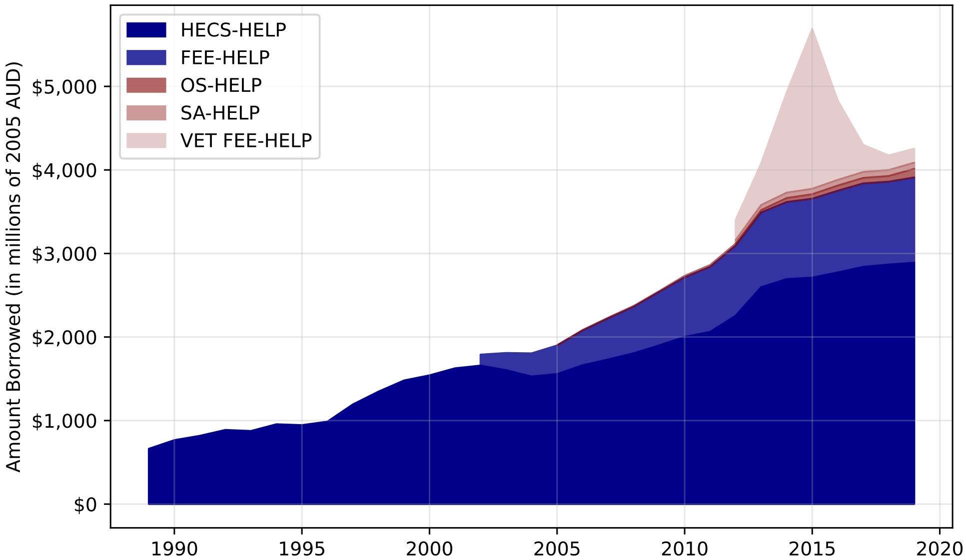

Notes: This figure plots the time series of the total amount borrowed each year among the five different HELP programs in millions of 2005 AUD. HECS-HELP refers to the primary HELP program that provides loans to cover student contribution amounts for Commonwealth Supported Places (CSPs), which cover mostly undergraduate and postgraduate degrees at public institutions. FEE-HELP loans are used to cover the fees associated with non-CSP degrees, such as undergraduate degrees at private institutions, which must be covered in full. FEE-HELP was introduced in 2005 and between 2002 and 2004 was formally called PELS. OS-HELP loans are used to cover expenses for students enrolled in a CSP degree who want to study overseas. SA-HELP loans are used to pay student services and amenities fees. VET FEE-HELP covers tuition fees for vocational education and training courses. VET FEE-HELP was closed on December 31st, 2016, and formally replaced by a different program called VET Student Loans on January 1st, 2017. The rapid increase in debt balances and subsequent closing of VET FEE-HELP was driven by fraud and corrupt behavior among vocational education providers (Australian National Audit Office 2016). A significant fraction of this debt has been written off in recent years (HELP Receivable Report 2021, DESE Annual Report 2022). Along with FEE-HELP and OS-HELP, borrowing through VET FEE-HELP has historically required incurring a loan fee that is around 20% of the amount borrowed. These data were obtained from Andrew Norton Higher Education Commentary.

Figure A2. Student Contributions and Aggregate HELP Borrowing over Time

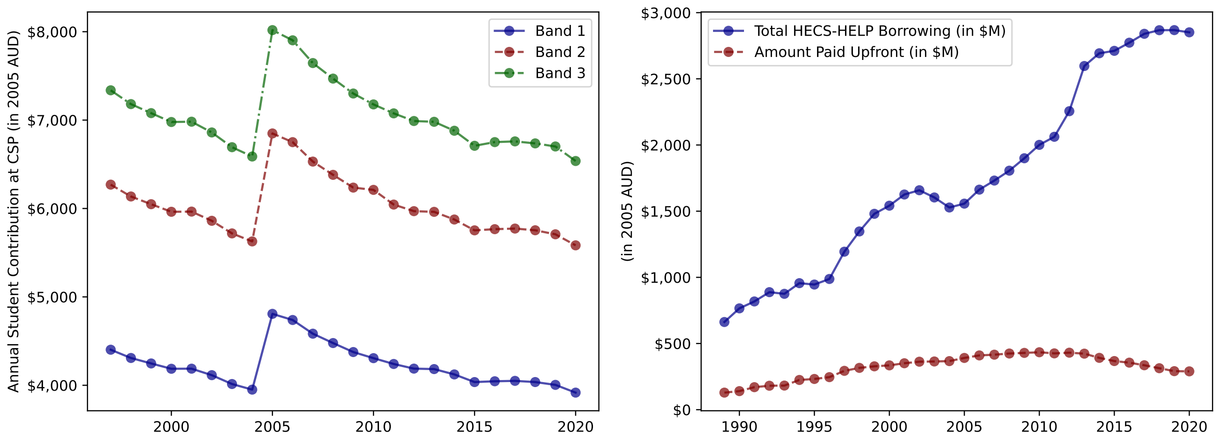

Notes: The left plot shows the time series of student contributions in 2005 AUD for Commonwealth Supported Places (CSPs) based on the three separate bands of study classified by the Australian government. These rates correspond to the cost of one year of coursework that must be covered with a HELP loan or by paying upfront. Prior to 2005, these rates were set by the government. After 2005, the rates were set by universities up to the maximum specified in this table, with most universities electing to charge the maximum. These three bands were introduced in 1997 and phased out in 2021 with the introduction of the Job Ready Graduates Package. Band 1 covers humanities, behavioral science, social studies, education, clinical psychology, foreign languages, visual and performing arts, education, and nursing. Band 2 covers computing, built environment, other health, allied health, engineering, surveying, agriculture, science, and maths. Band 3 covers law, dentistry, medicine, veterinary science, accounting, administration, economics, and commerce. Business and economics were Band 2 prior to 2008. Between 2005 and 2009, the government also had separate tuition for nursing and education and, from 2009 to 2012, for mathematics, statistics, and science, which were labeled national priorities. The right plot shows the time series of the aggregate amount of HECS-HELP borrowing and upfront payments in 2005 AUD. These data were obtained from Andrew Norton Higher Education Commentary.

Figure A3. Average Debt Balances by Age

Panel A: All Individuals

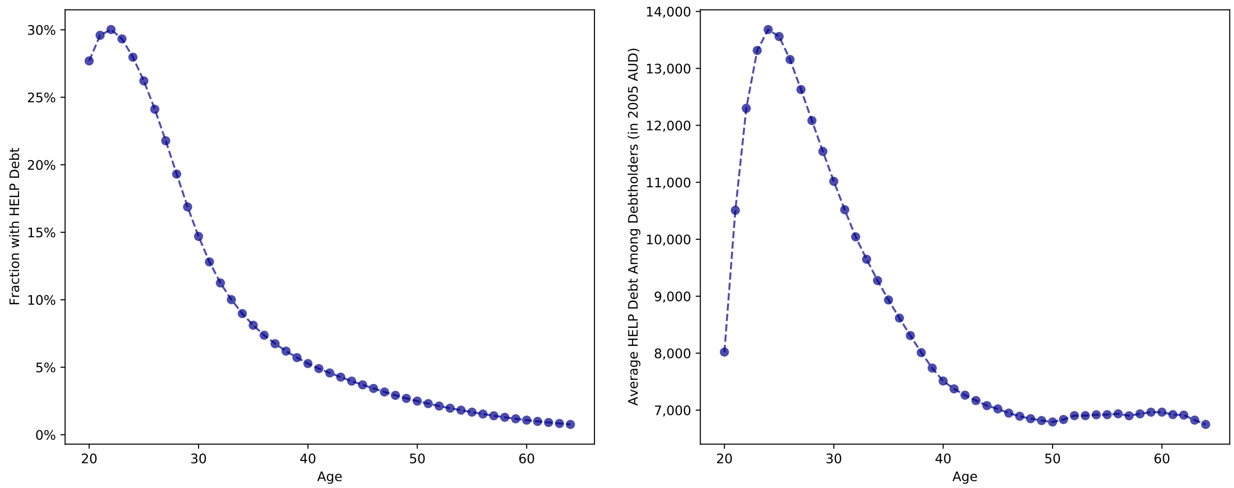

Panel B: Individuals with Positive Debt Balances at Age 22

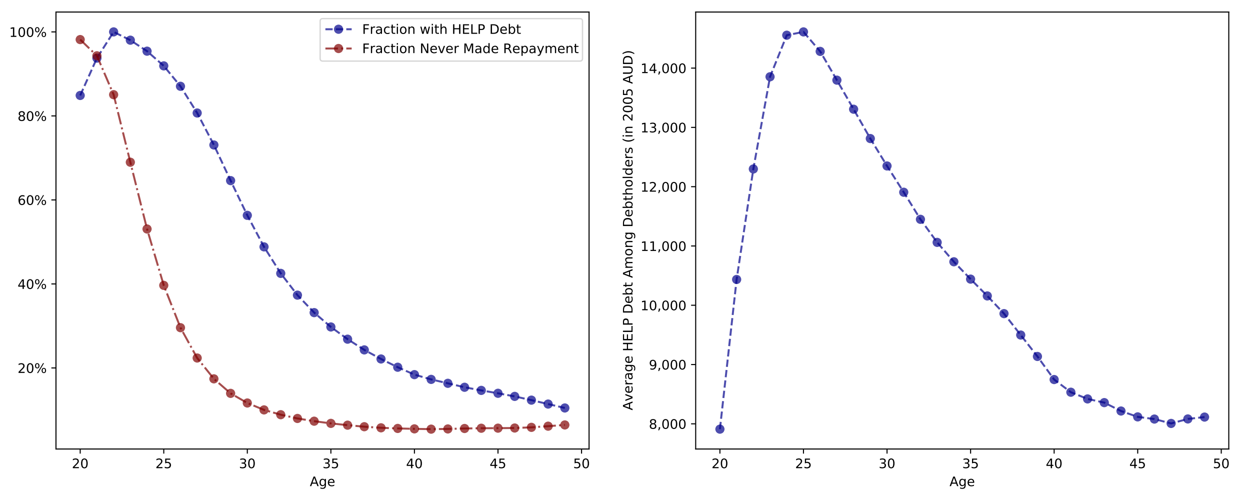

Notes: Panel A of this figure plots the fraction of individuals with HELP debt at each age in the left panel and the average HELP debt balances in 2005 AUD by age on the right. Panel B plots, in blue, the same quantity in Panel A among the subset of individuals who have positive debt balances at age 22 at some point during 1991–2019. The fraction of borrowers who have never made a HELP payment is also shown in the left panel in red. Debt balances are winsorized at 2% and 98%. The sample is the *ALife* sample defined in Section 1.4 from 1991 to 2019.

Figure A4. Comparison of HELP Income Distribution for Debtholders and Non-Debtholders

Panel A: Full Sample

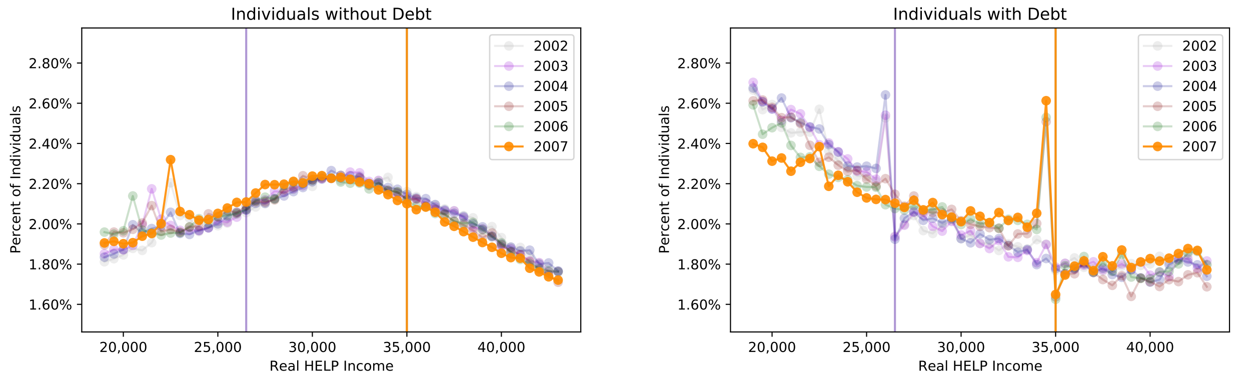

Panel B: Sample of Borrowers Held Fixed from 2002

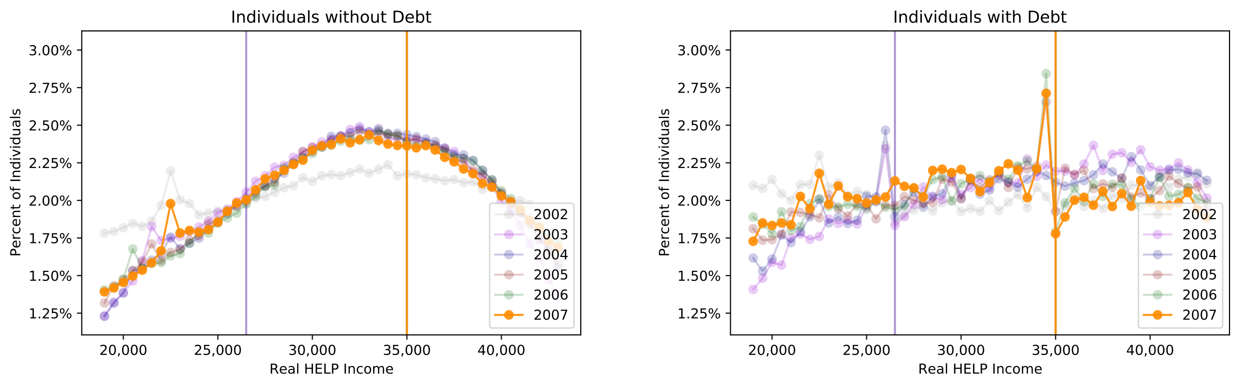

Notes: The right panel in Panel A of this figure replicates the bottom-right figure in Figure 3. The left panel in Panel A replicates the same analysis among individuals who do not have debt in each year. Panel B replicates the analysis in Panel A holding the sample of borrowers fixed to those who were present in the sample with HELP income (in 2005 AUD) between $20,000 and $50,000 in 2002.

Figure A5. Distributions of HELP Income and Labor Income

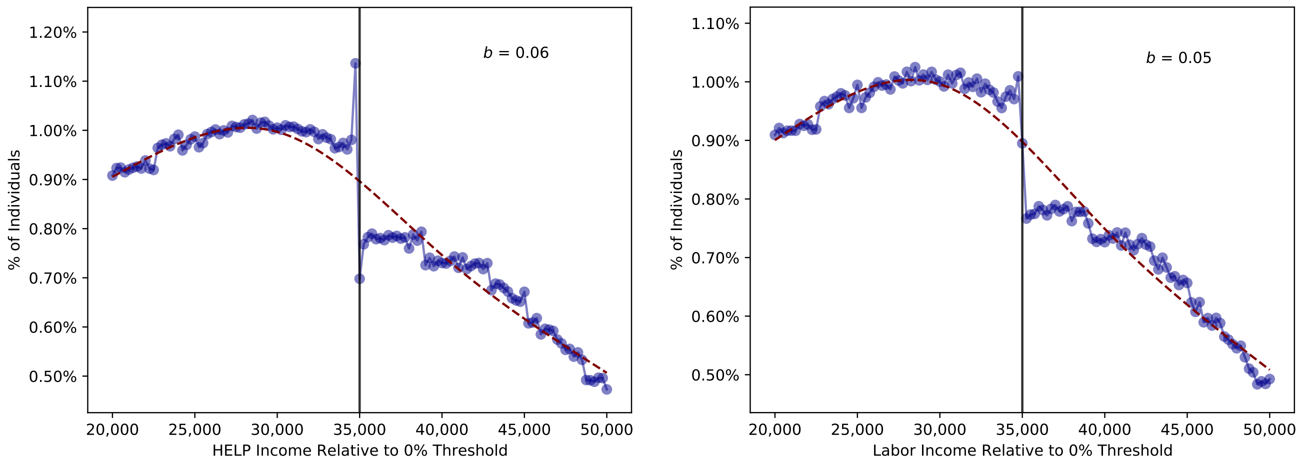

Notes: This figure plots the distributions of HELP and labor income (in 2005 AUD) relative to the repayment threshold after the policy change. This figure also plots the bunching statistic defined in (2) computed for the different distributions. Each bin corresponds to $250 AUD, and bins are chosen so that they center on the 2005 repayment threshold. The calculation of b is detailed in Appendix C.4, and the counterfactual density estimated in this procedure is plotted in the dashed red line. The sample is the *ALife* sample defined in Section 1.4 for the period between 2005 and 2018 after the policy change, restricted to individuals with positive HELP debt balances and less than 1% of HELP income from sources other than labor income.

Table A2. Hourly Flexibility Measures by 2-Digit ANZSCO Occupation

Occupation Title	SD Change in Log Hours	SD Log Hours
ICT Professionals	0.169	0.197
Electrotechnology and Telecommunications Trades Workers	0.192	0.209
Specialist Managers	0.193	0.265
Chief Executives, General Managers and Legislators	0.2	0.298
Engineering, ICT and Science Technicians	0.209	0.33
Factory Process Workers	0.211	0.309
Sales Representatives and Agents	0.218	0.316
Automotive and Engineering Trades Workers	0.225	0.226
Hospitality, Retail and Service Managers	0.226	0.347
Other Clerical and Administrative Workers	0.231	0.36
Machine and Stationary Plant Operators	0.232	0.269
Construction Trades Workers	0.238	0.213
Mobile Plant Operators	0.245	0.24
Health and Welfare Support Workers	0.246	0.408
Business, Human Resource and Marketing Professionals	0.256	0.33
Personal Assistants and Secretaries	0.26	0.503
Office Managers and Program Administrators	0.263	0.381
Road and Rail Drivers	0.263	0.394
Design, Engineering, Science and Transport Professionals	0.268	0.334
Inquiry Clerks and Receptionists	0.269	0.477
Protective Service Workers	0.275	0.274
Clerical and Office Support Workers	0.279	0.399
Numerical Clerks	0.296	0.483
Legal, Social and Welfare Professionals	0.302	0.378
Health Professionals	0.308	0.417
Construction and Mining Labourers	0.309	0.332
Other Technicians and Trades Workers	0.316	0.403
Skilled Animal and Horticultural Workers	0.317	0.517
Storepersons	0.324	0.356
General Clerical Workers	0.352	0.498
Food Trades Workers	0.358	0.42
Farmers and Farm Managers	0.365	0.441
Other Labourers	0.377	0.619
Carers and Aides	0.385	0.484
Farm, Forestry and Garden Workers	0.387	0.507
Education Professionals	0.408	0.529
Sales Support Workers	0.443	0.664
Cleaners and Laundry Workers	0.462	0.588
Food Preparation Assistants	0.475	0.637
Hospitality Workers	0.48	0.614
Sales Assistants and Salespersons	0.487	0.631
Sports and Personal Service Workers	0.498	0.687
Arts and Media Professionals	0.562	0.55

Notes: This table shows the measures of hourly flexibility at the 2-digit ANZSCO occupation-level used in Figure 4 and Figure A6. Hourly flexibility is measured as the standard deviation of annual changes, or the cross-sectional standard deviation, in log hours worked per week from HILDA.

Figure A6. Variation in Bunching across Occupations Based on Hourly Flexibility: Alternative Measure

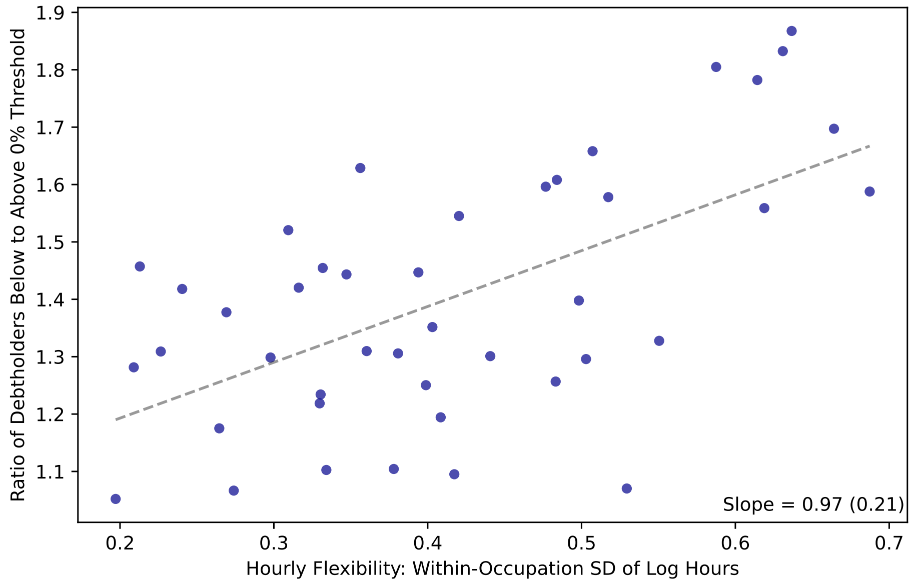

Notes: This figure plots the relationship between the amount of bunching below the repayment threshold and an alternative measure of hourly flexibility by occupation. Each point represents a 2-digit ANZSCO occupation code reported in *ALife*. The amount of bunching is measured as the ratio of the number of borrowers in that occupation within $2,500 below the repayment threshold to the number within $2,500 above the threshold for the period over 2005 to 2018. Hourly flexibility is measured as the cross-sectional standard deviation of log hours worked per week. The gray dashed line is the regression line with the estimated slope coefficient and standard error reported at bottom right. The sample is the *ALife* sample defined in Section 1.4, restricted to the subset of individual–years for which the borrowers are wage-earners.

Figure A7. Variation in Bunching across Occupations Based on Scope for Evasion

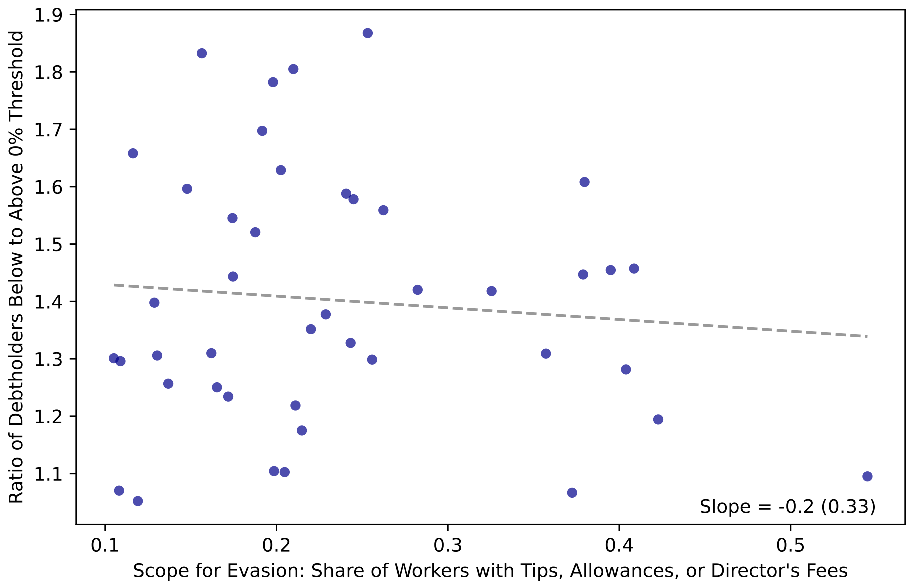

Notes: This figure replicates Figure 4 with a measure of evasion at the occupation level, instead of hourly flexibility on the horizontal axis. The measure of evasion is the fraction of individuals within each occupation who receive income from tips, allowances, or director's fees; see Appendix C.3 for additional details. This evasion measure is computed for the sample of individuals described in Figure A9.

Figure A8. Probability of Switching Occupations around the Repayment Threshold in 2005–2018

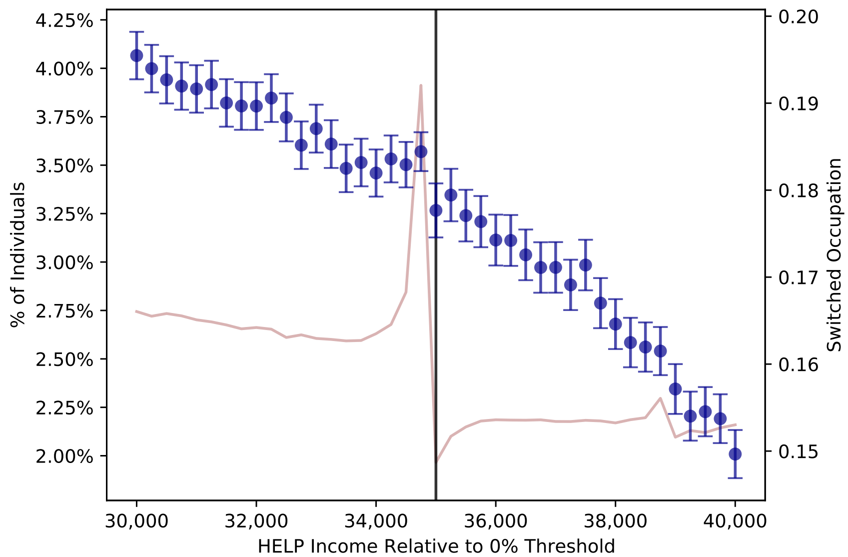

Notes: This figure plots the real HELP income distribution between 2005 and 2018 in red and measured on the left axis. HELP income is deflated to 2005 with the HELP threshold indexation rate, which is based on the annual CPI. Each bin represents $250, and the plot focuses on borrowers within $5,000 of the repayment threshold. The bins are chosen so that they are centered on the 2005 repayment threshold. The blue points present the fraction of individual–years in each bin in which borrowers' 2-digit ANZSCO occupation code differs from that of the previous year, along with 95% confidence intervals. The sample is the *ALife* sample defined in Section 1.4, restricted to the subset of individual–years with positive HELP debt balances between 2005 and 2018.

Figure A9. Age Profiles of Wage Income across Occupations

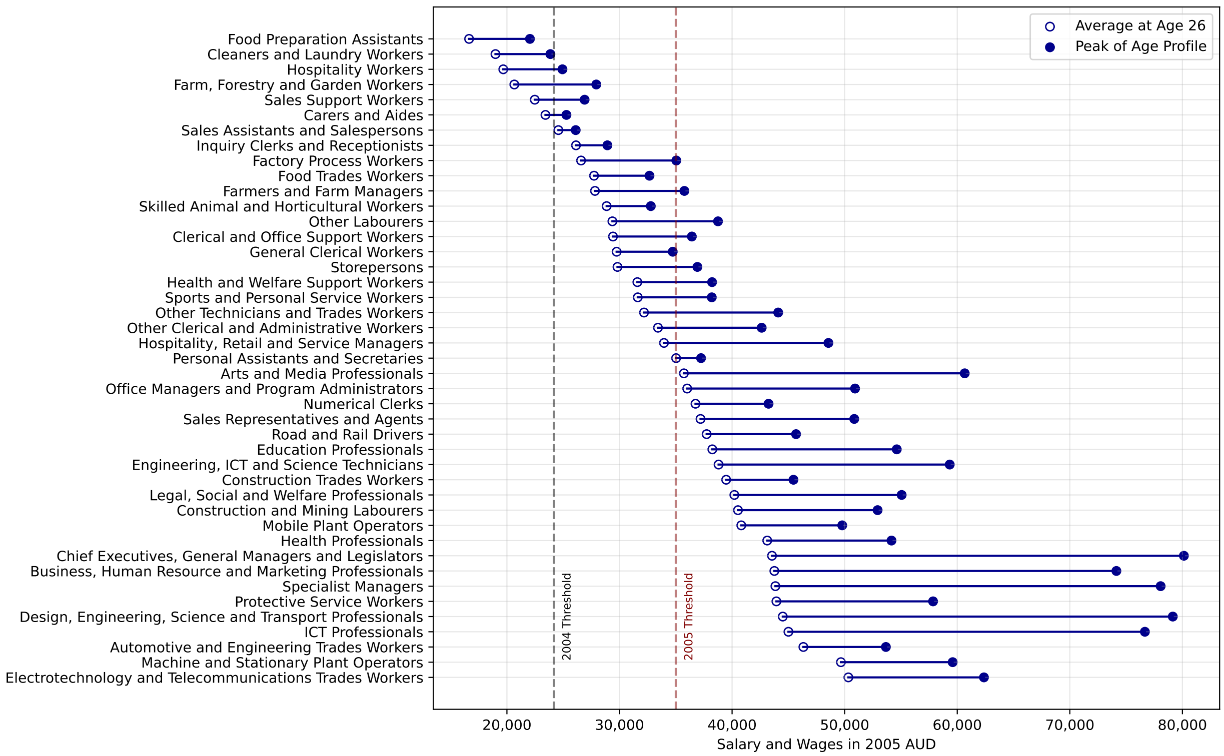

Notes: This figure plots characteristics of the age profile of salary and wages across 2-digit ANZSCO occupations. Occupation-specific age profiles are calculated by taking the average value of salary and wages across individuals in each occupation at a given age, after adjusting for inflation and removing year fixed effects. The figure then plots the value of each occupation profile at age 26 in white and the maximum value in the occupation profile in blue, with a blue line connecting the two. The sample of individuals used to calculate these age profiles is the *ALife* 10% random sample of individuals in the population *ALife* dataset who satisfy the sample selection criteria in Section 1, are wage-earners, and have annual salary and wages greater than one-half the legal minimum wage times 13 full-time weeks (Guvenen et al. 2014).

Table A3. Correlates of Bunching across Occupations

	Ratio of Debtholders Below to Above Threshold						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Hourly Flexibility: SD of Changes in Log Hours	1.30	·	·	·	1.30	1.05	0.50
	(0.35)	·	·	·	(0.35)	(0.28)	(0.23)
Evasion: Share with Non-Wage Income	·	−0.20	·	·	−0.02	−0.17	0.05
	·	(0.30)	·	·	(0.30)	(0.30)	(0.25)
Income Slope: Mean Wage at 45 / Mean Wage at 26	·	·	−0.53	·	·	−0.40	·
	·	·	(0.10)	·	·	(0.12)	·
Income Peak: Maximum Wage in Occupation Profile	·	·	·	−0.48	·	·	−0.40
	·	·	·	(0.06)	·	·	(0.07)
R^2	0.34	0.01	0.23	0.58	0.34	0.46	0.62
Number of Occupations	43	43	43	43	43	43	43

Notes: Each column of this table reports the results from an OLS regression run at the 2-digit ANZSCO occupation-level, with standard errors presented in parentheses below the coefficient estimates. The dependent variable in each column is the ratio of the number of debtholders within $2,500 below the repayment threshold to the number within $2,500 above the repayment threshold, as shown in Figure 4. Hourly Flexibility corresponds to the same measure used in Figure 4. Evasion corresponds to the share of all workers in each occupation who receive income from working in the form of allowances, tips, director's fees, consulting fees, or bonuses. Wage Slope corresponds to the occupation-specific average salary and wages at age 45, the age at which the pooled average of salary and wages reaches its maximum, divided by the average at 26 minus 1. Wage Peak corresponds to the maximum income in an occupation-specific age profile, normalized by the average value across all occupations. Salary and wages are adjusted for inflation, and year fixed effects are removed before computation of the occupation-specific age profiles used in the prior two measures. The Evasion, Wage Slope, and Wage Peak variables are calculated on the same sample of individuals used in Figure A9. Standard errors are computed with a heteroskedasticity-robust estimator.

Figure A10. Self-Reported Hours Worked around the Repayment Threshold: Borrowers with Positive Labor Income

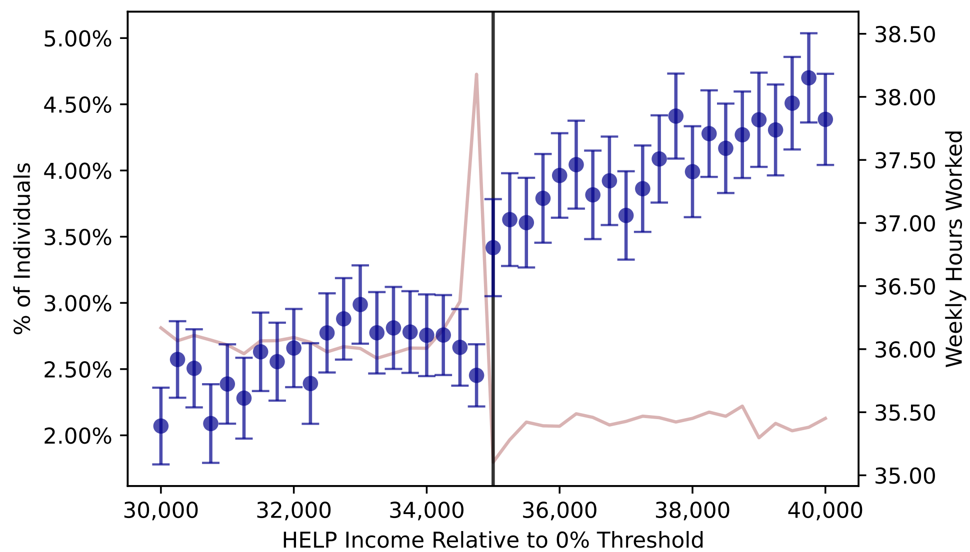

Notes: This figure replicates Figure 5 for the sample of borrowers with positive labor income.

Figure A11. Distribution of HELP Income in *ALife* versus MADIP Sample

Panel A: ALife Sample in 2016

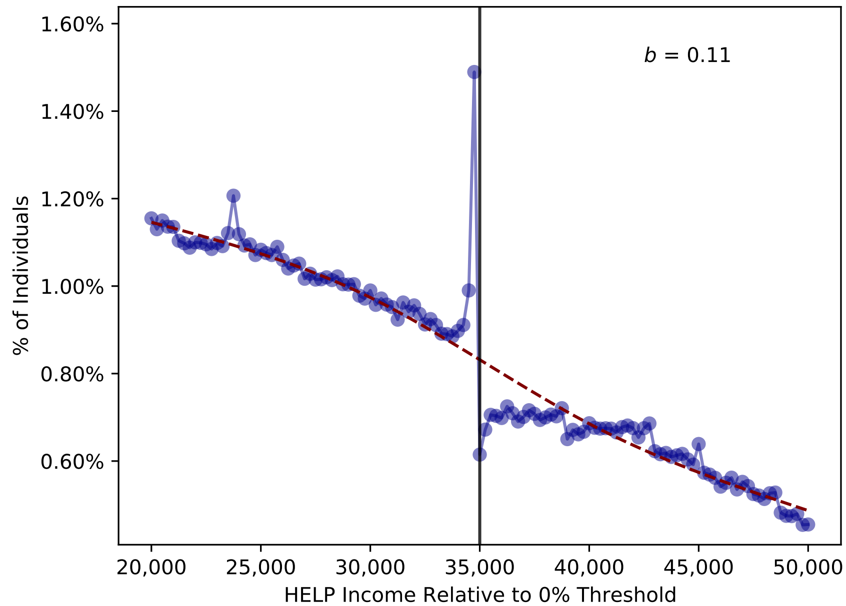

Panel B: MADIP Sample

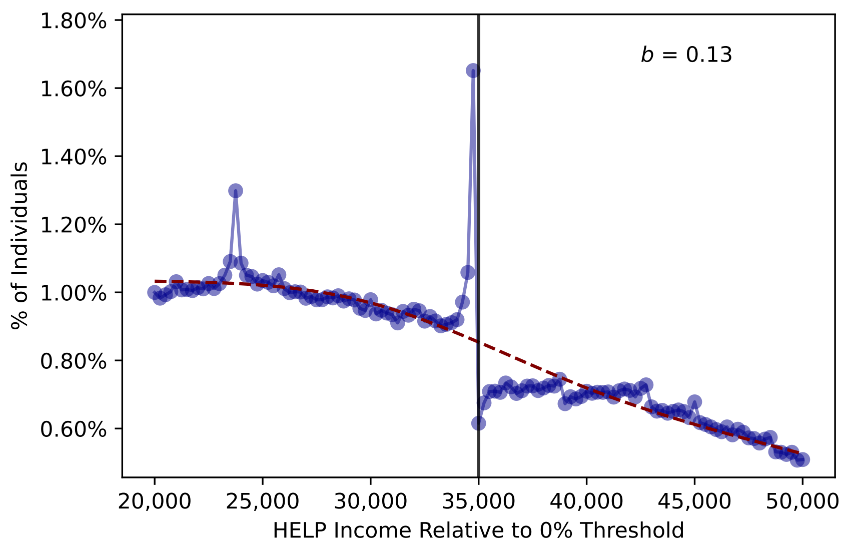

Notes: Panel A of this figure plots the distribution of HELP income (in 2005 AUD) in 2016 relative to the repayment threshold and the bunching statistic defined in (2). Each bin corresponds to $250 AUD, and bins are chosen so that they are centered around the 2005 repayment threshold. The calculation of b is detailed in Appendix C.4, and the counterfactual density estimated in this procedure is plotted in the dashed red line. The sample in this panel is the *ALife* sample defined in Section 1.4 in 2016, restricted to individuals with positive HELP debt balances. Panel B performs the same analysis in the cross-sectional MADIP sample, restricting to individuals with positive HELP debt balances.

Figure A12. Bunching Heterogeneity by Superannuation Balances: Ages 20–29

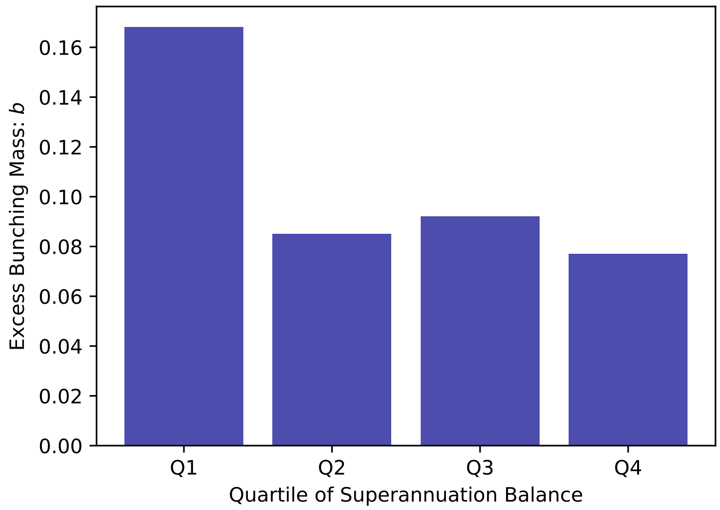

Notes: This figure replicates the analysis in the left panel of Figure 7 among borrowers who are ages 20–29.

Table A4. Additional Sources of Heterogeneity in Bunching

Sample	Estimated Bunching Statistic: b
Non-Electronic Filers	0.086
Electronic Filers	0.082
Wage-Earners	0.081
Entrepreneurs (Not Wage-Earners)	0.117
Females	0.081
Males	0.083
No Dependent Children	0.086
Has Dependent Children	0.077
No Spouse	0.085
Has Spouse	0.081
Full Sample	**0.084**

Notes: This table shows the bunching statistic defined in (2) computed for different samples of debtholders. The calculation of b is detailed in Appendix C.4. The sample in each row is the *ALife* sample defined in Section 1.4 for the period between 2005 and 2018 after the policy change, restricted to borrowers with positive HELP debt balances for whom the sample restrictions specified in each row are satisfied. The first two rows split borrowers based on whether they file their tax returns electronically; the third and fourth split the sample into wage-earners and non–wage-earners; the fifth and sixth split the sample based on gender; the seventh and eighth split the sample based on whether a borrower reports having a dependent child; and the ninth and tenth split the sample based on whether a borrower reports having a spouse.

Figure A13. Distributions of HELP Income and Salary and Wages

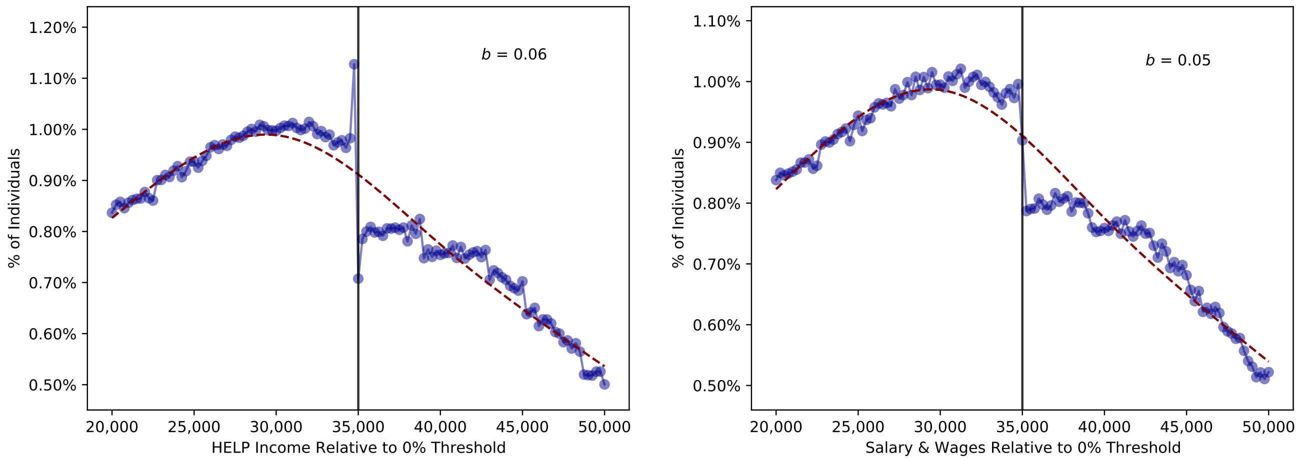

Notes: This figure replicates the analysis in Figure A5, replacing the right plot with salary and wages instead of labor income.

Figure A14. Distribution of HELP Income at Repayment Threshold versus Lowest 0.5% Threshold

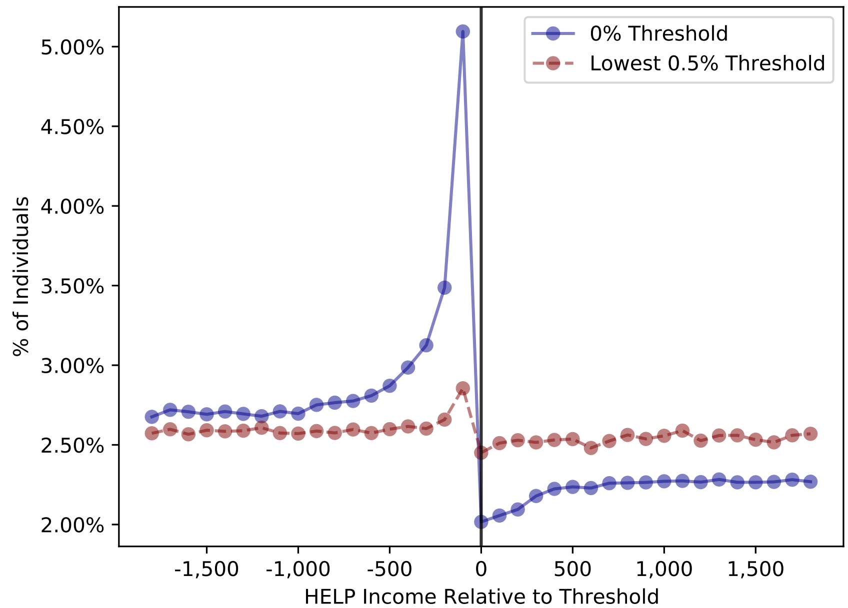

Notes: This figure plots the distribution of HELP income (in 2005 AUD) relative to the repayment threshold in solid blue and the lowest 0.5% threshold at $38,987 in dashed red. Each bin corresponds to $100 AUD, and bins are chosen so that they are centered around each threshold. The sample in this panel is the *ALife* sample defined in Section 1.4, restricted to individuals with positive HELP debt balances.

Table A5. Elasticity of Estimation Targets with Respect to Parameters

Panel A: Income Distribution Before the Policy Change

	y=22500	y=23000	y=23500	y=24000	y=24500	y=25000	y=25500	y=26000	y=26500	y=27000	y=27500	y=28000	y=28500
ϕ	0.04	0.00	0.04	0.05	0.31	-0.05	-0.06	-0.06	-0.11	-0.07	-0.06	-0.03	-0.10
λ	0.00	-0.00	-0.01	0.02	0.14	-0.10	-0.03	-0.01	-0.02	-0.01	-0.01	0.02	-0.02
f_L	0.01	0.00	0.00	0.01	-0.06	0.05	0.01	-0.01	0.01	-0.00	-0.00	-0.02	0.02
f_H	-0.00	-0.00	-0.00	-0.00	-0.00	0.00	0.01	0.00	-0.00	0.00	0.02	-0.02	-0.01
β	0.07	-0.03	0.19	-0.12	-0.87	0.68	0.05	0.25	-0.03	-0.00	0.17	0.09	-0.32
δ_0	-2.89	-2.14	-1.21	-1.36	-0.09	0.19	-0.77	0.56	1.09	0.91	1.91	2.57	2.25
δ_1	-0.57	-0.53	-0.39	-0.29	-0.03	-0.21	-0.04	-0.05	0.04	0.82	0.66	0.38	0.46
δ_2	-0.22	-0.14	-0.12	-0.11	-0.06	-0.09	0.01	0.00	0.06	0.00	0.16	0.42	0.18
δ_0^E	-0.01	-0.28	-0.12	-0.25	-0.04	-0.02	0.25	0.25	0.11	-0.08	0.03	0.15	0.08
δ_1^E	-0.12	-0.26	-0.16	-0.31	-0.08	0.08	0.15	0.31	0.05	0.07	0.00	0.07	0.32
ρ	1.15	0.65	0.84	0.07	0.33	0.01	-0.26	-0.40	-0.41	-0.43	-0.02	-1.15	-0.78
σ_ν	0.14	0.15	0.08	0.03	0.02	-0.04	0.01	-0.01	0.03	-0.09	-0.10	-0.17	-0.10
σ_ϵ	0.01	-0.00	-0.03	0.02	0.04	-0.01	-0.04	-0.03	0.04	0.01	0.02	-0.02	-0.00
σ_i	0.10	0.06	0.01	0.14	0.06	-0.07	-0.01	-0.17	0.08	-0.13	-0.01	-0.04	-0.06
κ	0.04	0.02	0.02	0.01	-0.03	0.02	0.02	0.00	-0.00	-0.01	-0.03	-0.04	-0.02
ι	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Panel B: Income Distribution After the Policy Change

	y=32500	y=33000	y=33500	y=34000	y=34500	y=35000	y=35500	y=36000	y=36500	y=37000	y=37500	y=38000	y=38500
ϕ	0.06	0.01	-0.01	-0.04	0.25	-0.06	-0.08	-0.08	-0.05	0.02	-0.05	-0.07	-0.01
λ	-0.02	0.01	-0.00	0.01	0.26	-0.13	-0.11	-0.08	-0.07	-0.04	-0.03	-0.01	0.06
f_L	0.01	0.01	0.01	0.01	-0.10	0.03	0.03	0.03	0.03	0.01	0.01	-0.01	-0.03
f_H	-0.01	-0.00	-0.00	0.00	-0.02	0.01	0.00	-0.00	0.00	0.01	0.00	0.01	0.01
β	-0.10	-0.07	0.05	-0.29	-0.90	0.18	0.72	0.19	0.22	0.40	0.16	0.02	0.00
δ_0	-1.74	-1.35	-1.51	-0.79	0.25	0.03	-0.42	-0.48	1.07	0.50	2.16	1.45	1.82
δ_1	-0.67	-0.21	-0.45	-0.33	0.41	-0.23	-0.26	0.34	0.27	-0.04	0.24	0.40	0.67
δ_2	-0.31	-0.27	-0.11	-0.04	0.17	-0.06	-0.15	-0.03	0.10	0.17	0.16	0.18	0.27
δ_0^E	-0.09	0.09	-0.06	0.05	0.04	-0.23	-0.30	-0.04	0.35	0.12	-0.22	0.06	0.19
δ_1^E	-0.05	-0.11	-0.04	-0.02	-0.10	-0.25	0.19	-0.03	-0.13	0.10	0.31	0.17	0.08
ρ	0.94	0.19	0.11	0.56	-0.84	-0.09	0.45	0.11	-0.03	-0.06	-0.69	0.09	-0.64
σ_ν	0.03	0.00	0.07	0.04	-0.05	0.02	0.03	0.02	-0.01	-0.04	-0.03	-0.13	0.03
σ_ϵ	0.00	-0.03	0.03	0.03	-0.02	0.05	-0.05	-0.01	0.01	0.02	0.01	-0.02	-0.04
σ_i	0.03	-0.04	-0.16	0.12	0.04	0.08	-0.03	-0.01	0.03	-0.10	-0.00	0.05	-0.01
κ	0.02	0.02	0.01	0.01	0.04	0.01	-0.01	-0.02	-0.02	-0.01	-0.03	-0.03	-0.01
ι	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table A5. Elasticity of Estimation Targets with Respect to Parameters (continued)

Panel C: Income Process Moments

	SD at 22	SD at 32	SD at 42	SD at 52	SD at 62	β_1	β_2	P10 1-Yr	P10 5-Yr	P90 1-Yr	P90 5-Yr	β_0^E	β_1^E
ϕ	0.27	0.24	0.25	0.28	0.31	0.01	-0.02	-0.03	-0.07	0.03	0.07	-0.14	0.14
λ	0.04	0.03	0.04	0.05	0.05	0.01	-0.01	-0.01	-0.02	0.01	0.02	-0.04	0.03
f_L	0.00	-0.00	-0.00	-0.00	-0.01	0.00	-0.00	0.00	0.00	-0.00	-0.00	0.00	-0.00
f_H	-0.04	-0.03	-0.03	-0.03	-0.05	-0.01	0.02	0.01	0.02	-0.01	-0.02	0.02	-0.02
β	-0.27	-0.21	-0.14	-0.15	-0.14	0.15	-0.12	0.03	0.06	-0.02	-0.06	0.10	-0.10
δ_0	-1.06	-0.64	-0.82	-1.13	-0.50	-0.36	0.43	0.02	0.06	0.02	-0.04	0.27	-0.35
δ_1	-0.13	-0.15	-0.24	-0.41	-0.22	0.87	0.21	0.02	0.04	0.01	-0.01	0.18	-0.17
δ_2	-0.10	-0.08	-0.15	-0.24	-0.04	-0.11	1.27	0.02	0.04	0.00	-0.01	0.10	-0.10
δ_0^E	-0.05	0.07	0.17	0.24	0.28	-0.00	0.00	-0.00	-0.00	0.00	-0.00	1.00	-0.02
δ_1^E	-0.05	0.08	0.28	0.50	0.71	0.08	0.02	-0.00	-0.01	0.00	0.01	0.06	0.95
ρ	0.71	9.63	11.64	11.20	8.87	-0.35	0.37	0.05	-0.69	-0.04	0.68	-0.32	0.28
σ_ν	0.04	1.49	1.76	1.63	1.28	-0.03	0.04	-0.55	-0.83	0.55	0.83	-0.08	0.07
σ_ϵ	0.10	0.09	0.09	0.07	0.04	-0.01	0.01	-0.44	-0.15	0.44	0.15	0.00	-0.00
σ_i	1.86	0.45	0.10	0.02	0.00	-0.00	0.00	-0.01	-0.04	0.01	0.04	-0.00	0.00
κ	0.01	0.01	0.01	0.01	0.00	0.00	-0.00	-0.00	-0.00	0.00	0.00	-0.00	0.00
ι	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Panel D: Remaining Estimation Targets

	Ratio: Q4 to Q1 Debt	Mean ℓ	Fraction No Adjustment	Kurtosis $\Delta \log \ell$	Persistence 2004-2005
ϕ	0.22	-0.20	-0.05	0.29	0.06
λ	0.19	0.01	-0.02	-0.03	0.36
f_L	0.04	-0.00	0.00	1.88	-0.13
f_H	0.02	-0.01	0.02	1.56	0.01
β	0.81	0.07	0.01	-6.30	3.43
δ_0	1.77	1.54	-0.11	9.16	6.03
δ_1	0.71	0.38	-0.03	0.75	-0.91
δ_2	0.15	0.22	-0.02	-1.06	0.44
δ_0^E	0.16	0.07	-0.01	1.03	0.27
δ_1^E	0.32	0.11	-0.01	-5.83	-0.00
ρ	-1.08	-0.12	-0.26	22.39	-1.69
σ_ν	-0.18	-0.02	-0.09	1.77	-0.87
σ_ϵ	-0.02	0.00	-0.01	1.52	-0.64
σ_i	-0.01	-0.00	-0.01	-1.94	1.01
κ	-0.04	-0.10	0.00	-1.16	-0.08
ι	0.00	0.00	-0.87	-3.33	0.00

Notes: This table reports the elasticity of the simulated estimation targets with respect to the estimated structural parameters. The four panels present the results for different sets of estimation targets. In each panel, the entry in row i and column j is an estimate of the derivative of the log of the estimation target in column j with respect to the log of the structural parameter in row i. I approximate this derivative locally around the estimated set of structural parameters in column (5) of Table 3 by central differencing. Since some estimation targets and parameters are negative, I take the absolute value before taking logarithms and then multiply the result by -1 if the parameter or estimation target is negative. The width between the lower and upper points in central differencing is set equal to the step size used in the Nelder–Mead optimization routine in estimating the model.

Figure A15. Variation in Bunching by Debt Balances and Age: Ratio Measure

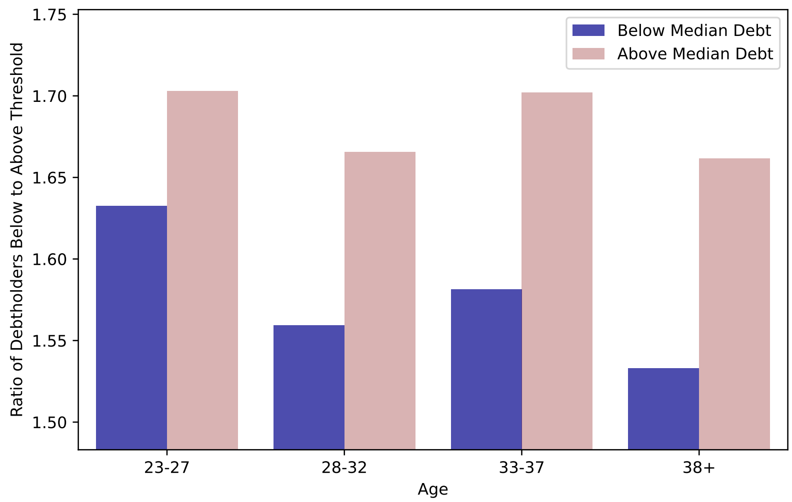

Notes: This figure shows the analogous plot to Figure 6 using the bunching measure used in Figure 9.

Figure A16. Effect of Moral Hazard on Welfare Gains

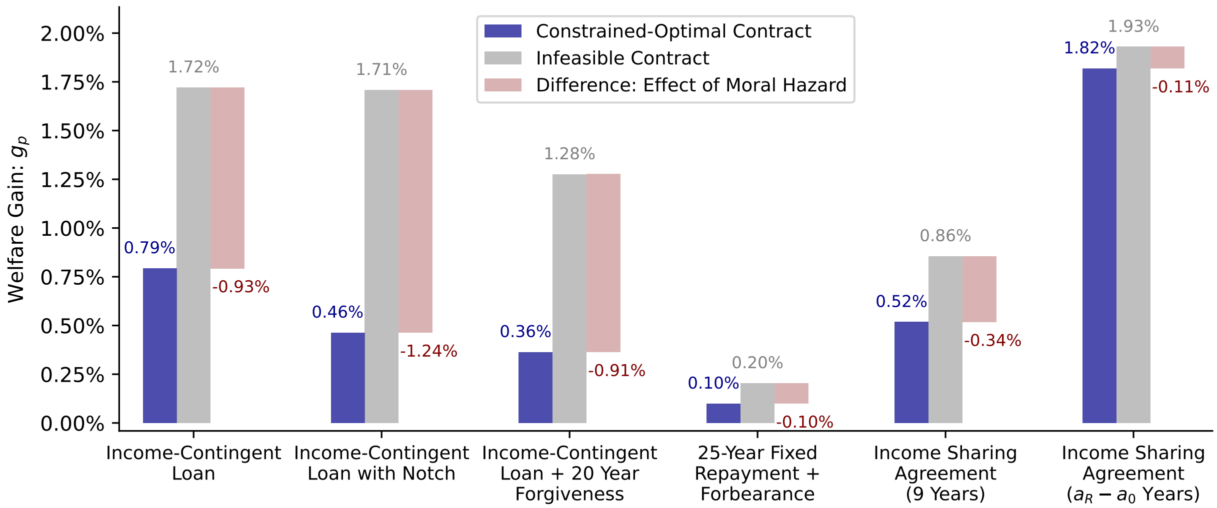

Notes: This figure decomposes the welfare gains, g_p, from Table 7, repeated in the left blue bar for each contract, into two components. The middle gray bar corresponds to the welfare gain that would exist if the contract shown in the final two columns of Table 7 was implemented in the baseline model with endogenous labor supply. This contract is not feasible because it was the solution to (15) assuming that ℓ_{ia} remains fixed at is value under the benchmark contract for all i and a. The right red bar plots the difference between the two bars, which corresponds to the lostt from moral hazard.

Table A6. Parameters and Welfare Effects of Constrained-Optimal Contracts: Model with $f_H = \infty$

Contract Space: p	ψ_p	K_p	π_p	g_p	$\psi_p^{\ell \text{ fixed}}$	$K_p^{\ell \text{ fixed}}$
Income-Contingent Loan	14%	\$31,055	\$4,821	1.18%	59%	\$62,022
Income-Contingent Loan with Notch	5.5%	\$37,704	\$4,978	1.21%	14%	\$67,315
Income-Contingent Loan + 20 Year Forgiveness	26%	\$27,877	\$3,047	0.76%	40%	\$42,285
25-Year Fixed Repayment + Forbearance	0.17%	.	\$1,558	0.40%	0.07%	.
Income Sharing Agreement (9 Years)	3.4%	.	\$2,494	0.63%	3.0%	.
Income Sharing Agreement ($a_R - a_0$ Years)	0.52%	.	\$7,374	1.75%	0.48%	.

Notes: This table reproduces Table 7 in the model estimated in column (4) of Table 3.

Figure A17. Redistribution-Neutralizing Transfers for Constrained-Optimal Contracts

Panel A: Income-Contingent Loan

Panel B: Income-Sharing Agreement ($a_R - a_0$ Years)

Notes: This figure shows the transfers in each of the $\mathcal{T} = 8$ initial states made to eliminate the redistributive effects of different constrained-optimal contracts described in Table 8. The missing value in Panel B corresponds to a case in which no transfer could be found to balance the government budget in that state. See Appendix D.7 for additional details and the discretized values of the three initial conditions.

Figure A18. Comparison of Debt Restructuring with Changing Taxes & Transfers

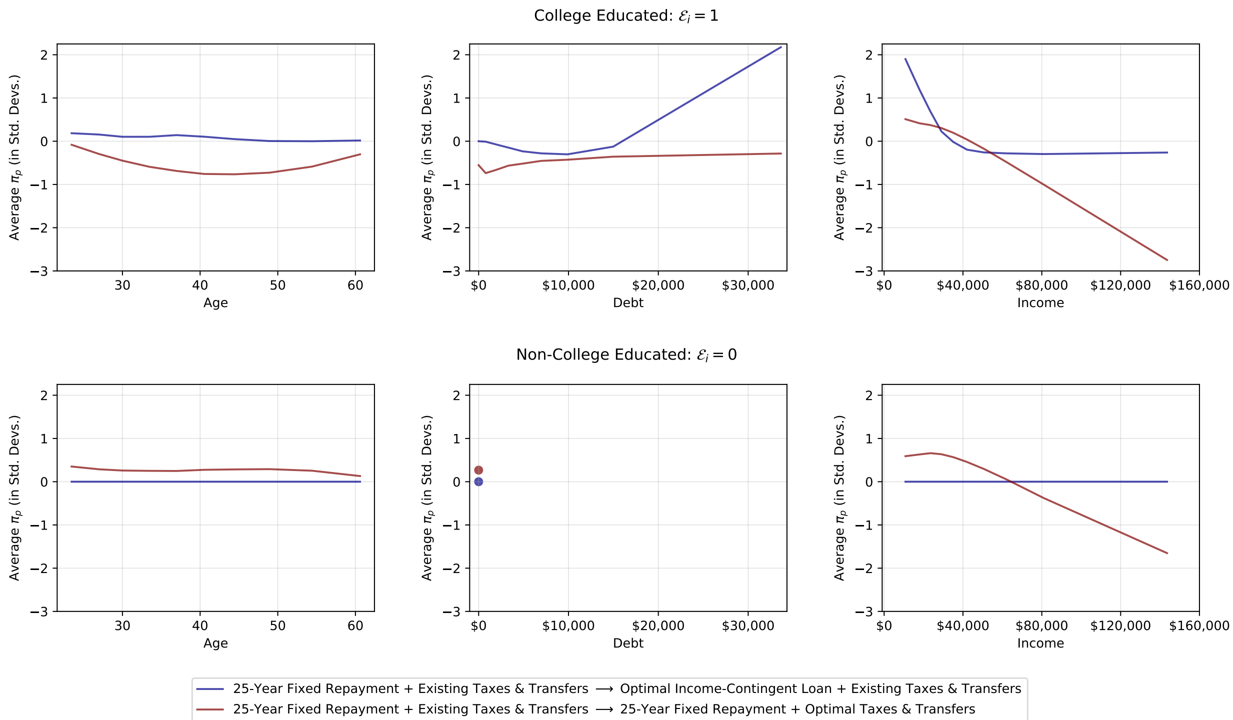

Notes: This figure compares the results from two experiments. The first experiment holds $\tau(\cdot)$ fixed and changes the debt repayment function from the benchmark fixed repayment contract to the constrained-optimal income-contingent loan in Table 7. The second holds the debt repayment function fixed and changes $\tau(\cdot)$. In the first experiment, $\tau(\cdot)$ is equal to the Heathcote et al. (2017) functional form calibrated to the Australian tax schedule, as described in Appendix D.2. In the second experiment, $\tau(\cdot)$ is equal to the Heathcote et al. (2017) functional form, where the two parameters of this tax function have been chosen to maximize the expected utility at $a = a_0$ of an individual that does not know any of her initial states and views their realizations as risk. In both experiments, the simulation procedure follows the same procedure used to estimate the model, where the policy change occurs at $t = T^*$. This figure then plots the average of π_p in each experiment across all individuals that have the value of the state shown on the horizonal axis. The top axis focuses on individuals with $\mathcal{E}_i = 1$, while the bottom focuses on those with $\mathcal{E}_i = 0$. In all panels, the welfare gains shown are normalized by the standard deviation of π_p across all states within each experiment so that the distribution of gains from the two experiments have similar magnitudes.

Figure A19. Welfare Gains from Smooth Repayment Contracts

Panel A: Repayment Functions for Constrained-Optimal Contracts

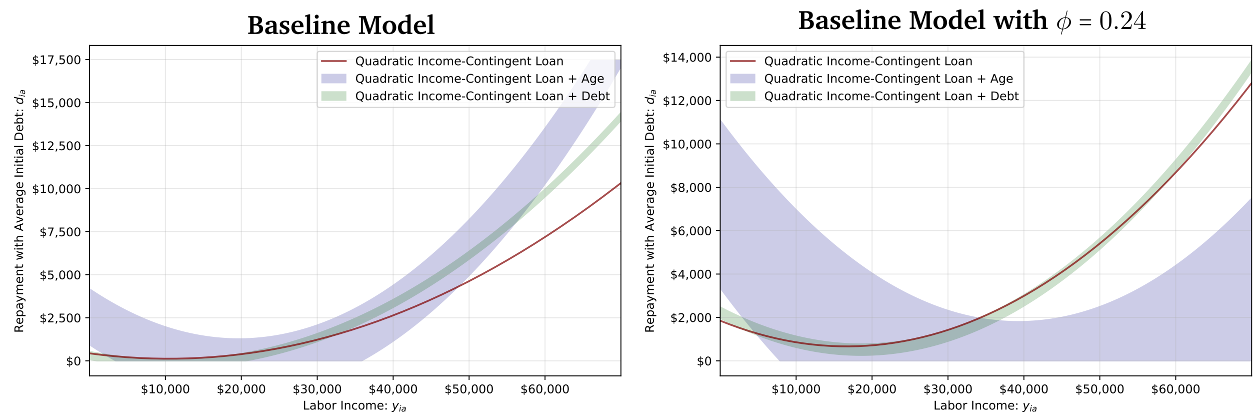

Panel B: Parameters and Welfare Effects of Constrained-Optimal Contracts: Baseline Model

Contract Space: p	θ_1	θ_2	$\theta_3 \times 1000$	θ_4	π_p	g_p
Quadratic Income-Contingent Loan	$435	-0.06	0.0029	.	$2,524	0.73%
Quadratic Income-Contingent Loan + Age	$365	-0.30	0.0075	0.005	$3,399	0.94%
Quadratic Income-Contingent Loan + Debt	$25	-0.09	0.0041	0.014	$2,657	0.77%

Panel C: Parameters and Welfare Effects of Constrained-Optimal Contracts: Baseline Model with $\phi = 0.24$

Contract Space: p	θ_1	θ_2	$\theta_3 \times 1000$	θ_4	π_p	g_p
Quadratic Income-Contingent Loan	$1,850	-0.14	0.0043	.	$614	0.19%
Quadratic Income-Contingent Loan + Age	$1,993	-0.47	0.0060	0.011	$2,011	0.60%
Quadratic Income-Contingent Loan + Debt	$1,898	-0.18	0.0049	0.016	$805	0.25%

Notes: Panel A of this figure shows repayments as a function of income for the richer contract spaces described in Appendix D.9. The shaded regions in this plot correspond to the repayments for individuals with initial debt balances between the 10th and 90th percentiles and with ages between the 10th and 90th percentiles of when the final debt repayment is made in the baseline model. Panels B and C show the corresponding welfare gains in the two different models used in Panel A: the baseline model, and the baseline model with $\phi = 0.24$ while all other parameters are held fixed.

Figure A20. Distribution of Estimated Labor Supply Elasticities from Prior Studies

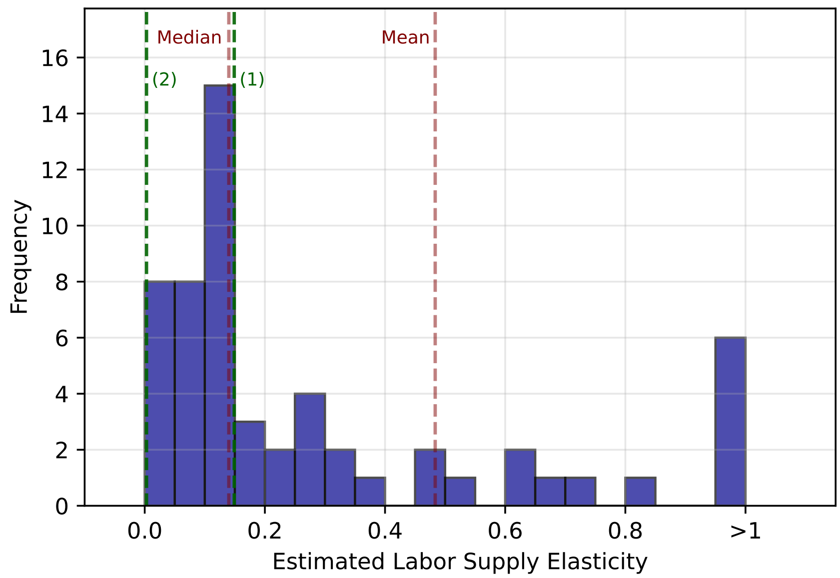

Notes: This figure plots a histogram of the intensive margin labor supply elasticities estimated in prior literature. I combine the estimates reported in Tables 6 and 7 of Keane (2011) and Table 1 of Chetty et al. (2012). These estimates include intensive margin Frisch (i.e., marginal utility–constant) and Hicksian (i.e., wealth-constant) elasticities estimated among studies that measure labor supply using hours worked or taxable income, which have the closest structural interpretation to my estimates. This graph pools all studies, some using full populations, others using just men or women. See Keane (2011) and Chetty et al. (2012) for a detailed discussion of the underlying studies. In the histogram, all studies that estimate a value above one are placed into the last bar, but the mean and median, shown in dashed red lines, are calculated before these observations are trimmed. The two dashed green lines plot the estimates from columns (1) and (5) of Table 3, respectively.

References

Arnoud, Antoine, Fatih Guvenen, and Tatjana Kleineberg (2019), "Benchmarking Global Optimizers." *Working Paper*.

Bachas, Natalie (2019), "The Impact of Risk-Based Pricing in the Student Loan Market: Evidence from Borrower Repayment Decisions." *Working Paper*.

Blundell, Richard and Thomas MaCurdy (1999), "Labor Supply: A Review of Alternative Approaches." In *Handbook of Labor Economics*, volume 3, 1559–1695, Elsevier.

Bommier, Antoine, Daniel Harenberg, François Le Grand, and Cormac O'Dea (2020), "Recursive Preferences, the Value of Life, and Household Finance." *Working Paper*.

Britton, Jack and Jonathan Gruber (2020), "Do income contingent student loans reduce labor supply?" *Economics of Education Review*, 79, 1020–1061.

Catherine, Sylvain (2022), "Countercyclical Income Risk and Portfolio Choices over the Life-Cycle." *Review of Financial Studies*, 35, 4016–4054.

Catherine, Sylvain and Constantine Yannelis (2023), "The Distributional Effects of Student Loan Forgiveness." *Journal of Financial Economics*, 147, 297–316.

Cesarini, David, Erik Lindqvist, Matthew J. Notowidigdo, and Robert Ostling (2017), "The Effect of Wealth on Individual and Household Labor Supply: Evidence from Swedish Lotteries." *American Economic Review*, 107, 3917–3946.

Chetty, R., J. N. Friedman, T. Olsen, and L. Pistaferri (2011), "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records." *The Quarterly Journal of Economics*, 126, 749–804.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber (2012), "Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities." *NBER Macroeconomics Annual*, 27.

Choukhmane, Taha and Tim de Silva (2023), "What Drives Investors' Portfolio Choices? Separating Risk Preferences from Frictions." *Working Paper*.

Duffie, Darrell and Kenneth J. Singleton (1993), "Simulated Moments Estimation of Markov Models of Asset Prices." *Econometrica*, 61, 929–952.

D'Souza, Gabriela (2018), "A higher education bubble?" *Working Paper*.

Einav, Liran, Amy Finkelstein, and Paul Schrimpf (2015), "The Response of Drug Expenditure to Nonlinear Contract Design: Evidence from Medicare Part D." *The Quarterly Journal of Economics*, 130, 841–899.

Ey, Carol (2021), "The Higher Education Loan Program (HELP) and related loans: A chronology." Technical report, Parliament of Australia.

Feldstein, Martin (1999), "Tax Avoidance and the Deadweight Loss of the Income Tax." *Review of Economics and Statistics*, 81, 674–680.

Ganong, Peter and Pascal Noel (2020), "Liquidity versus wealth in household debt obligations: Evidence from housing policy in the great recession." *American Economic Review*, 110, 3100–3138.

Golosov, Mikhail, Michael Graber, Magne Mogstad, and David Novgorodsky (2023), "How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income." *Working Paper*.

Gruber, Jon and Emmanuel Saez (2002), "The elasticity of taxable income: Evidence and implications." *Journal of Public Economics*, 84, 1–32.

Guvenen, Fatih (2009a), "An empirical investigation of labor income processes." *Review of Economic Dynamics*, 12, 58–79.

Guvenen, Fatih (2009b), "A Parsimonious Macroeconomic Model for Asset Pricing." *Econometrica*, 77, 1711–1750.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song (2021), "What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?" *Econometrica*, 89, 2303–2339.

Guvenen, Fatih, Serdar Ozkan, and Jae Song (2014), "The nature of countercyclical income risk." *Journal of Political Economy*, 122, 621–660.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2017), "Optimal Tax Progressivity: An Analytical Framework." *The Quarterly Journal of Economics*, 132, 1693–1754.

Heathcote, Jonathan and Hitoshi Tsujiyama (2021), "Optimal Income Taxation: Mirrlees Meets Ramsey." *Journal of Political Economy*, 129, 3141–3184.

Ji, Yan (2021), "Job Search under Debt: Aggregate Implications of Student Loans." *Journal of Monetary Economics*, 117, 741–759.

Kargar, Mahyar and William Mann (2023), "The Incidence of Student Loan Subsidies: Evidence from the PLUS Program." *The Review of Financial Studies*, 36, 1621–1666.

Karlan, Dean and Jonathan Zinman (2009), "Observing Unobservables: Identifying Information Asymmetries With a Consumer Credit Field Experiment." *Econometrica*, 77, 1993–2008.

Keane, Michael P (2011), "Labor Supply and Taxes: A Survey." *Journal of Economic Literature*, 49, 961–1075.

Kleven, Henrik J. and Mazhar Waseem (2013), "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan*." *The Quarterly Journal of Economics*, 128, 669–723.

Lusardi, Annamaria, Pierre Carl Michaud, and Olivia S. Mitchell (2017), "Optimal financial knowledge and wealth inequality." *Journal of Political Economy*, 125, 431–477.

Marshall, Kate (2003), "Ease HECS burden on students, say universities." *Australian Financial Review*.

McFadden, Daniel (1989), "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration." *Econometrica*, 57, 995–1026.

MoneySmart (2016), "Paying off your uni debt." Technical report.

Moreno, Carla and Sita Slavov (2024), "The Impact of Lump-Sum Retirement Withdrawals on Labor Supply: Evidence from Peru." *Working Paper*.

Mumford, Kevin J. (2022), "Student Selection into an Income Share Agreement." *Working Paper*.

Nocedal, Jorge and Stephen J. Wright (2006), *Numerical Optimization*, 2nd ed edition. Springer Series in Operations Research, Springer, New York.

Norton, Andrew (2018), “Has abolishing the discount for upfront payment of student contributions made a difference to upfront payment rates?” Technical report.

Norton, Andrew (2019), “Demand-driven funding for universities is frozen. What does this mean and should the policy be restored?” *The Conversation*.

Norton, Andrew and Ittima Cherastidtham (2016), *Shared Interest: A Universal Loan for HELP*. Grattan Institute, Carlton, Victoria.

Reuther, Albert, Jeremy Kepner, Chansup Byun, Siddharth Samsi, William Arcand, David Bestor, Bill Bergeron, Vijay Gadepally, Michael Houle, Matthew Hubbell, Michael Jones, Anna Klein, Lauren Milechin, Julia Mullen, Andrew Prout, Antonio Rosa, Charles Yee, and Peter Michaeleas (2018), “Interactive supercomputing on 40,000 cores for machine learning and data analysis.” In *2018 IEEE High Performance Extreme Computing Conference (HPEC)*, 1–6.