

Centre for Population Ageing Research (CEPAR)

Working Paper 2026/02

Household Home Care Financing: Long-Term Care Insurance, Home Equity Release and Means-Tested Public Support

Lingfeng Lyu, Michael Sherris

This paper can be downloaded without charge from the CEPAR Working Paper Series available at cepar.edu.au.

Household Home Care Financing: Long-Term Care Insurance, Home Equity Release and Means-Tested Public Support

Lingfeng Lyu[1,2], Michael Sherris[1,2]

[1]School of Risk and Actuarial Studies, UNSW Sydney, NSW 2052, Australia.
[2]Centre for Population Ageing Research (CEPAR), UNSW Sydney, NSW 2052, Australia.

Abstract

Australia has no private long-term care insurance (LTCI) market. The recent aged care reforms, which shift more of the cost of home care onto households, raise the question whether LTCI could complement public support. We develop a continuous-time lifecycle model of retired homeowners under a means-tested age pension, a means-tested home care subsidy with a lifetime cap on fees, and the government reverse mortgage scheme. These homeowners choose consumption and portfolio allocation continuously, and at retirement fix their LTCI coverage and the fraction of the permitted home equity release that they draw. The means tests that set the pension and the home care fee would count the LTCI contract and its indemnity. The model is solved by policy iteration and calibrated to Australian panel data. At an actuarially fair premium, almost no household buys LTCI when home equity release is also offered, and fewer than one in five buys it when release is not offered. Exempting the LTCI contract from the means tests raises the insured share of households to nearly half, with or without release. With the contract assessed and a bequest motive, release lowers LTCI demand, because the home is kept during home care and drawing can then provide money when LTCI pays. For the government, a premium rebate for LTCI raises the insured share of households. The government could recover the rebate through the means tests, but at a rebate of 30 per cent of the loaded price, the tests on average give insured households more than they take back.

Keywords: Means testing; Long-term care insurance; Home equity release; Multi-state model; Public–private risk sharing.

JEL classification: G22, I13, D15, H55, G51

1 Introduction

Across the OECD the share of the population aged 65 and over has doubled since 1960 (OECD, 2023), and the demand for long-term care (LTC) has risen with it. LTC is the sustained help with daily living that older people need when chronic illness or cognitive decline reduces their ability to look after themselves. Unlike acute care, whose aim is recovery, LTC is open-ended. A household entering it does not know how long it will last, how intensive it will become or what it will cost. Even with public subsidies, the out-of-pocket cost of severe need absorbs more than 70 per cent of older people's median income, on average across the OECD (OECD, 2024). For older homeowners the home is often the largest asset (Wood et al., 2022), but it is illiquid and, in Australia, exempt from the means tests that determine what the household pays for care. Two instruments address the mismatch between an open-ended care cost and wealth held in an illiquid home. One is an insurance contract whose benefit, the indemnity, is paid in the care states. The other is home equity release, a loan that converts home equity into cash while the household keeps the home. Both have been studied alongside a public safety net (Brown and Finkelstein, 2008; Nakajima and Telyukova, 2017). This paper studies them in a system in which one set of means tests links the age pension, the care subsidy and home equity release.

In Australia, Support at Home (SAH) has been in force since November 2025. Under SAH, the household pays a home care fee that depends on its age pension and is capped over a lifetime (DHDA, 2025b,c). Both the fee and the lifetime cap are higher than under the previous Home Care Package (Park and Sherris, 2026; DHDA, 2025b), so the household pays a larger share of the care costs. The government pays the rest as the care subsidy. Australia has no private LTC insurance (LTCI) market. This paper asks whether households would buy an LTCI contract if one were offered under the new rules, and what such a contract would cost the government.

The LTCI contract is therefore hypothetical. We specify its terms: a single premium at retirement, an 18 per cent expense loading (Brown and Finkelstein, 2007), which gives the loaded price, and a government premium rebate. At the baseline the rebate offsets the loading, which makes the premium actuarially fair, and it is about 15 per cent of the loaded price. We assume that the contract would be assessed under the rules for the closest existing product, a lifetime income stream. Under those rules, sixty per cent of the indemnity is assessable income, and the same share of the purchase price, before the rebate, is an assessable asset (Services Australia, 2026b). Both amounts enter the income and assets tests that set the age pension, and the pension links the three programmes. The fee rises as the pension falls, and the Home Equity Access Scheme (HEAS), a government reverse mortgage, has an annual drawing cap that is reduced by the pension the household receives. Unlike the indemnity, a drawing under the HEAS is exempt from the income test.

We consider a retired couple who own the home they live in, hold liquid wealth outside it, and face the care costs that follow a decline in either member's health.[1] The household chooses consumption and portfolio allocation continuously. At retirement it fixes how much LTCI cover to buy and what fraction of its permitted HEAS drawing to take each year. Cover is the share of the whole care cost that the contract pays, rather than a share of the fee. The state variables are liquid wealth, home equity and joint health, together with three path-dependent ones: the HEAS loan balance, which accumulates past drawings with interest, and each member's cumulative fees, to which the lifetime cap applies. We estimate the joint-life health generator on twenty waves of the Household, Income and Labour Dynamics in Australia (HILDA) survey, correct it for reporting error, and then solve the resulting Hamilton–Jacobi–Bellman equation by policy iteration. We evaluate the solved model for sixty-four household cells, thirty-two partnered and thirty-two single. For each type, we divide the HILDA homeowners into eight equal groups by liquid wealth at retirement and each group into four by home equity.

This paper makes three contributions.

[1] A single household is modelled in the same way, with one member deceased at retirement.

The first is the framework. Our continuous-time model includes the three programmes: the age pension's income and assets tests, the home care fee and its lifetime cap, and the HEAS's pension-linked drawing cap and no-negative-equity guarantee (Section 3). Under the guarantee, the estate never repays more of the HEAS loan than its home equity. Every rule depends on the household's choices and feeds back into its liquid wealth, so a change in any choice moves the pension, the fee and the drawing cap together. Continuous time makes the three path-dependent states easier to handle. The loan balance and each member's cumulative fees can only rise, so the solver treats their movement without interpolating the value function (Section 3.4.3).

Existing studies that jointly model home equity and LTCI assume that households sell their home when they enter residential care. Home equity therefore becomes liquid at the same time as LTCI benefits are paid (Davidoff, 2010; Shao et al., 2019; Xu et al., 2023). The most closely related Australian studies use discrete-time models that include the means-tested age pension (Andréasson et al., 2017). A later model also includes a commercial reverse mortgage, for which borrowing is capped by a loan-to-value ratio rather than by the amount of pension received (Andréasson and Shevchenko, 2024). We assume that the home is kept, and residential care is outside the model. To the best of our knowledge, this paper is the first to model the three linked means-tested programmes together in continuous time.

The second is a measure of the clawback, the net amount that the means tests take from a household, in age pension withdrawn and home care fees raised, because it holds a private LTCI contract. A negative clawback is an implicit subsidy. Brown and Finkelstein (2008) measure the corresponding implicit tax under Medicaid as one ratio, the share of a private contract's benefits that replace what the programme would otherwise have paid. Under the Australian rules, by contrast, the means tests both charge and subsidise the contract. Part of the indemnity counts as income, so the indemnity lowers the pension and raises the fee. Paying the premium instead reduces assessable wealth at retirement, and the pension is higher from then on.

We measure the clawback in two ways (Section 4.4). The first, the instantaneous clawback, applies the rules at a given liquid wealth and health state at a single point in time. The second, the realised clawback, adds up what the rules take and give along the household's own path, comparing the same household with and without LTCI cover on identical histories of health and asset returns. Because it follows the household's wealth, the realised clawback also counts what paying the premium gives back through lower assessable wealth, and what an unspent indemnity takes away through higher assessable wealth. Since most households buy no cover at the baseline, we measure the realised clawback with the same imposed coverage for every household. In the middle of the distribution of liquid wealth at retirement, the instantaneous clawback charges cover, whereas over a lifetime the realised clawback is an implicit subsidy.

The third is an evaluation of how LTCI cover, a premium rebate and release change the distribution of the government's net expenditure. The government is the residual insurer of home care, through the age pension, the hardship waiver of home care fees a household cannot afford, the care subsidy, the lifetime cap and the no-negative-equity guarantee. Net expenditure on a household is the discounted sum of the age pension, care subsidy, premium rebate and guarantee shortfall that the government pays for that household, less the spread it earns on the HEAS loan. The evaluation extends the method of Lyu et al. (2026), which tracks long-term government expenditure along simulated household paths, to a setting that includes private instruments. The far tail of net expenditure is what the government spends on a household in its costliest outcomes, which come from long lives and long spells of high need. In the model, health and survival risks are idiosyncratic, so they diversify across households, and the far tail is not a risk to the budget. The government still pays for these outcomes as residual insurer, and a private contract that pays in the long spells of high need takes part of that role. The change in the far tail with cover therefore shows whether the contract relieves the government in those outcomes. We judge the rebate by two criteria: whether the means tests recover the rebate paid on average, and how much the far tail rises per dollar of the rebate's net fiscal cost,

the rise in mean net expenditure (Section 7).

We find that demand for LTCI cover is low even at an actuarially fair premium, and lower where home equity release is offered. Where release is not offered, fewer than one household in five buys cover, and none of them buys more than 8 per cent of full cover. Most of these households are in the middle of the wealth distribution at retirement. Where release is also offered, only 2.4 per cent of households insure (Section 6.1). Those that still insure form one group of partnered households at the top of the wealth distribution, for whom release is worth little and who do not draw on it.

We also find that, with the contract assessed and a bequest motive, home equity release displaces LTCI cover. Offering release lowers the insured share by about sixteen percentage points and raises cover for no household, whereas offering cover leaves every household's drawing unchanged. Every household that stops insuring draws instead. Because the home is kept during care, drawing provides funds in every health state, including the care states in which the indemnity is paid. Section 6.4.1 shows, by allowing drawing only in high need, that these funds lower the value of cover. In the models of Davidoff (2010) and Shao et al. (2019), the home is sold on entering residential care, and release raises cover instead. The displacement depends on the bequest motive and on the assessment of the contract. Without the bequest motive, release raises the insured share by about one and a half percentage points (Section 6.2.2). With the contract exempt from the means tests, release lowers that share by under one percentage point, and about one household in five holds both instruments (Section 6.2.1).

Drawing in the model rests partly on risky investment. Almost all households hold all their liquid wealth in the risky asset at retirement, and the wealth that drawing frees is held there, earning more than the HEAS interest rate. With at most half of liquid wealth in the risky asset, about two in five households insure even where release is offered, and release still lowers cover for some households and raises it for none (Table 3 and Section 6.3). In practice only about twenty thousand participants had a HEAS loan at 30 June 2026 (Department of Social Services, 2026), so for most households the results where release is not offered are the closer guide (Section 6.3).

Apart from its displacement by release, three reasons hold demand for LTCI cover down. The first is the assessment of the contract under the means tests (Brown and Finkelstein, 2008). The means tests take back part of what paying the premium gains the household, through the income and asset assessments of the contract described above. With an imposed coverage and full drawing, the implicit subsidy in the middle of the wealth distribution would be about two to six times as large without these two assessments, depending on the household group (Section 6.4.2). Exempting the contract from both tests raises the insured share of households to nearly half where release is offered (Section 6.2.1). The second reason is that the care subsidy already pays most of the care cost, 77 per cent on average, including what it pays after a member's fees reach the lifetime cap. Because the contract pays a share of the whole care cost rather than of the fee, full cover would pay the fee and would also pay out in cash the part that the care subsidy already meets (Section 6.4.3). Cover therefore insures the household only against the fee. We do not study a contract written on the fee itself. The third is the bequest motive (Lockwood, 2018). Households prefer keeping wealth for their heirs to paying premiums for the remaining risk. Removing the motive raises the insured share where release is offered from 2.4 per cent to nearly three quarters (Section 6.2.2).

Finally, we find that a premium rebate raises the insured share. A rebate of 30 per cent of the loaded price, about twice the baseline rebate, raises it to a quarter of households where release is offered, and to about three in five where it is not. The means tests do not recover the rebate paid on average. Where release is offered, the households in the middle of the wealth distribution that would insure draw instead, so the insured are wealthier and hold more cover, and at a rebate of 30 per cent the net fiscal cost per insured household is higher. The rebate also raises the far tail of net expenditure, which we measure by the expected shortfall beyond the

ninety-ninth percentile, the average of the costliest one per cent of outcomes, among households of the same type and wealth at retirement. Where release is not offered, this rise is smaller than the rebate's net fiscal cost at every rebate level, so the far tail rises less than the mean. At a rebate of 30 per cent the rise is about a third of the net fiscal cost. Where release is offered, the rise is also smaller until the rebate is high enough that one group of partnered households holding cover stops drawing. The government then pays that group more age pension and care subsidy, and loses the net income from its HEAS loan. From that rebate on, the rise is about as large as the net fiscal cost or larger, mainly because of that group (Section 7).

Section 2 describes the three public programmes and the rules that link them. Section 3 sets out the household's problem, Section 4 the estimation on HILDA, the instantaneous clawback and the realised clawback, and Section 5 the calibration. Section 6 reports the results, and Section 7 considers what they mean for the government. Section 8 concludes. Sections, tables, figures and equations numbered with an S prefix are in the online supplement.

2 Institutional features

Three public programmes shape an Australian retiree's cash flows, and one set of means-testing rules links them. The central programme, the non-contributory age pension, is subject to an income test, under which financial assets are deemed to earn income at statutory rates, and to an assets test. The household receives the lower of the two resulting rates. Neither test counts the owner-occupied home. Each test withdraws the pension linearly in the quantity it assesses, between a lower threshold, below which the full pension is paid, and an upper threshold, at which no pension is paid. We call this withdrawal the test's taper, and the range of liquid wealth in which the household receives a part pension the taper band. Section S9 states both tests in closed form, with the statutory values.

The HEAS is a government reverse mortgage whose drawing cap falls as the pension received rises, so that pension plus drawing cannot exceed 1.5 times the maximum pension rate. Drawings are exempt from the income test, but any drawing not yet spent is assessable under the assets test. For independence and everyday-living services under SAH, the household pays a home care fee, which rises as its pension falls and is subject to a lifetime cap per member. The government pays the rest of the care cost, and we call this payment the care subsidy. Table 1 states the three rules and their dependencies.

A private LTCI contract, if one were offered, would affect this system through the premium and the indemnity. We assume that the statute assesses the contract as it assesses the closest existing product, a lifetime income stream. A share ν of the indemnity is then assessable income, and the same share of the purchase price is an assessable asset while the contract is held. We call ν the means-test intensity, and this part of the purchase price the deemed asset, which earns no deemed income. The assets test then counts assessable wealth, which is liquid wealth plus the deemed asset. The premium is paid out of liquid wealth at retirement, so paying it lowers assessable wealth, raises the pension, lowers the fee and narrows the HEAS drawing cap. While a member needs care, the indemnity is paid and is assessable income, so it has the opposite effects. The member in care also reaches the lifetime cap sooner, after which the whole of that member's care cost falls on the public budget. Joint health transitions set the timing and size of both sets of effects.

3 Model

A retired couple in full health faces disability risk and pays any care costs out of the wealth that also supports its precautionary saving and bequest. The household retires at age 67, and no member lives beyond age 117. A single household is modelled by placing the other member in the absorbing death state at retirement (Section 4.2.2), and Section 6 reports both household types. Throughout, a label ind or sp distinguishes the HILDA respondent from the spouse, and a star marks an optimised control.

Table 1. Statutory rules and joint dependency structure of the three public programmes.

Component	Notation	Means-testing rule	Constraint	Depends on
Age pension	$I^{\mathrm{AP}}(t)$	Income & assets tests (stricter binds). Assessable income includes deemed income on liquid wealth $X(t)$ and a share $\nu = 0.6$ of a lifetime income stream, here the LTCI indemnity.[a]	Maximum rate per person ($\bar{M}$) differs by marital status.	$X(t)$, $I^{\mathrm{LTCI}}(t)$, the deemed asset.
HEAS (reverse mortgage)	$I^{\mathrm{HEAS}}(t)$, $\kappa(t)$	Drawing cap: $1.5(1+\mathbb{1}_{\mathrm{couple}})\bar{M} - I^{\mathrm{AP}}(t)$; the Maximum Loan Amount (Section 3) rises with age and home equity $Y(t)$.	$\kappa(t)$ is the HEAS loan balance, $\mathrm{d}\kappa = (r^{\kappa}\kappa + I^{\mathrm{HEAS}})\,\mathrm{d}t$, and the household draws a constant fraction $\phi \in [0,1]$ of the smaller of the drawing cap and the headroom below the Maximum Loan Amount. Drawings are exempt from the income test.	$Y(t)$, $I^{\mathrm{AP}}(t)$, $\kappa(t)$, ϕ.
Home care under SAH[b]	$F^{\mathrm{HC}}_{\mathrm{ind}}(t)$, $F^{\mathrm{HC}}_{\mathrm{sp}}(t)$	Fee rate falls linearly in $I^{\mathrm{AP}}(t)$; clinical care is fully subsidised.	Lifetime cap $\bar{\zeta} =$ A$130,000 of DHDA (2025b) indexed to 1 November 2025, and is treated as fixed in real terms.	$I^{\mathrm{AP}}(t)$, $H^{\mathrm{ind}}(t)$, $H^{\mathrm{sp}}(t)$, $\zeta^{\mathrm{ind}}(t)$, $\zeta^{\mathrm{sp}}(t)$.

Notes: Statutory parameters in force on 1 November 2025, with the maximum pension rate $\bar{M}$ at the basic rate, without the pension and energy supplements; the hypothetical LTCI contract is assessed as a lifetime income stream. Symbols are defined in Section 3: $H^{\mathrm{ind}}(t)$ and $H^{\mathrm{sp}}(t)$ are the two members' health states, and each $I(t)$ and $F(t)$ is a flow that the named programme pays or charges.

[a] While the LTCI contract is held, assessable assets include ν times its loaded price, the deemed asset, which earns no deemed income.

[b] The lifetime cap $\bar{\zeta}$ is the A$130,000 of DHDA (2025b) indexed to 1 November 2025, and is treated as fixed in real terms.

3.1 Joint health

Each member's health lies in $\mathcal{H}^{\mathrm{ind}} = \mathcal{H}^{\mathrm{sp}} = \{1, \ldots, n\}$, with 1 full health, 2 to $n-1$ disability states of increasing care need, and n death. Section 4.2.2 estimates $n = 5$. The couple's joint state is numbered by

$$\iota : \mathcal{H}^{\mathrm{ind}} \times \mathcal{H}^{\mathrm{sp}} \longrightarrow \mathcal{H} := \left\{1, \ldots, n^2\right\}, \quad \iota(i,j) = (i-1)n + j, \tag{1}$$

which preserves lexicographic order. The joint states with at least one member in a disability state are the care states. The joint state $H(t) = \iota\big(H^{\mathrm{ind}}(t), H^{\mathrm{sp}}(t)\big)$ follows a continuous-time Markov chain with time-inhomogeneous generator

$$Q(t) = \big(q_{h,h'}(t)\big)_{h,h' \in \mathcal{H}}, \quad q_{h,h'}(t) \geq 0 \ (h \neq h'), \quad q_{h,h}(t) = -\sum_{h' \neq h} q_{h,h'}(t),$$

whose off-diagonal entries are the transition intensities and whose time dependence represents ageing. Dependence between spouses' lives has been modelled for joint-life annuities as a Markov process on the couple's survival states (Ji et al., 2011), and the care of both spouses has been priced into a life-care annuity (Ventura-Marco et al., 2023). Here one generator on the joint state governs both lives. Health changes in jumps, as in a cardiovascular event, rather than along a smooth path.

3.2 LTCI

Following Xu et al. (2023), the household buys the LTCI contract with a single premium at retirement. A couple's contract is joint-life, covering both members until the last death, and a single household's is single-life. As in Shao et al. (2019), the insurer pays an indemnity whenever the household occupies an insured state (Section 4.2.2).

3.2.1 Pricing

We assume that the insurer values the couple's care cost net of the informal care each spouse can provide the other (Costa-Font, 2010; Mommaerts, 2025). The maximum informal care a healthy spouse provides is IC, and that capacity declines at the rate $\varpi > 0$ per step of the carer's own deterioration. For joint state $h = \iota(h^{\text{ind}}, h^{\text{sp}})$ the care cost net of informal care is

$$\mathrm{LTC}_h = \left(\widetilde{\mathrm{LTC}}_{h^{\text{ind}}} - \mathrm{IC}(1-(h^{\text{sp}}-1)\varpi)^+\right)^+ + \left(\widetilde{\mathrm{LTC}}_{h^{\text{sp}}} - \mathrm{IC}(1-(h^{\text{ind}}-1)\varpi)^+\right)^+, \tag{2}$$

where $\widetilde{\mathrm{LTC}}$ is the gross care cost of a member's own state. At $\varpi = 1/3$ (Table 2), the inner positive part sets a dead spouse's capacity to zero, so a single household receives no informal care. The indemnity and the premium are both based on LTC_h, whoever pays that cost, and not on a member's gross cost or on the home care fee. The contract pays in the insured states $\mathcal{H}' \subset \mathcal{H}$ and can be bought in their complement $\complement_{\mathcal{H}}(\mathcal{H}')$, the eligible joint states. For a household that buys in eligible state h_0 at time t, the actuarially fair premium for full cover from t, valued at time 0, is

$$P(\complement_{\mathcal{H}}(\mathcal{H}');t) = \int_t^T e^{-rs} \sum_{h' \in \mathcal{H}'} \Pr\left(H(s) = h' \mid H(t) = h_0\right) \mathrm{LTC}_{h'} \mathrm{d}s, \tag{3}$$

where T is the maximum survival time and r the continuously compounded real risk-free rate. Below, P without arguments denotes $P(\complement_{\mathcal{H}}(\mathcal{H}');0)$. It does not depend on h_0 within a household type, because a couple can buy only with both members in full health, and a single household only with one member healthy and the other dead, two mirror-image states with pooled transitions (Section 4.2.2). The household chooses a coverage share $\psi \in [0,1]$, the share of its care cost that the contract pays. The loaded price is $(1+\theta_L)\psi P$, where $\theta_L \geq 0$ is the expense loading.

3.2.2 Means tests and the premium rebate

The means-test intensity $\nu \in [0,1]$ is the share of the indemnity that is assessable income (Section 2). The same share of the loaded price, $\nu(1+\theta_L)\psi P$, is the deemed asset, which only the assets test counts while the contract is held, although the premium has left liquid wealth. A premium rebate $\eta \in [0,1]$ reduces the premium paid to $(1+\theta_L)(1-\eta)\psi P$. Australian statute offers no LTCI rebate, so η is hypothetical, and its treatment borrows from the rules for lifetime income streams. Those rules assess a share of the purchase price, so we assume that the rebate, which meets part of the price without changing it, leaves the deemed asset unchanged. The part of the price that the rebate covers stays in liquid wealth, where it earns deemed income and is assessable like any other financial asset. Consequently η moves only the point at which the value function of Equation (9) is evaluated, so the model need not be solved again for any value of η.

3.3 Financial market and home equity

Let $W = (W_1, W_2)$ be a standard two-dimensional Brownian motion on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ satisfying the usual conditions, with $\mathbb{F}$ the augmentation of the natural filtration of W and H, which are independent. The household invests in the risk-free asset at rate r and in a traded risky asset

$$\mathrm{d}S(t) = S(t)\left(\mu_1 \mathrm{d}t + \sigma_1 \mathrm{d}W_1(t)\right), \quad S(0) > 0. \tag{4}$$

The household also owns non-traded home equity Y, the home's value net of any mortgage debt but before the HEAS loan balance κ. The return shocks of home equity have correlation ρ with

the traded asset's,

$$\mathrm{d}Y(t) = Y(t)\left(\mu_2 \mathrm{d}t + \sigma_2 \rho\, \mathrm{d}W_1(t) + \sigma_2\sqrt{1-\rho^2}\, \mathrm{d}W_2(t)\right), \quad Y(0) \geq 0, \tag{5}$$

with $\mu_1 \in \mathbb{R}$ and $\sigma_1 > 0$ the traded asset's expected total return and volatility, $\mu_2 \in \mathbb{R}$ and $\sigma_2 > 0$ the constant-quality drift and capital-gain volatility of home equity, and $\rho \in [-1, 1]$. Home equity's idiosyncratic shock W_2 cannot be hedged, as with the unhedgeable income of Munk (2000). We assume that the household does not sell, trade down or let the home, as older owners seldom move (Davidoff, 2010). Imputed rent is not assessable under either test, and an additively separable value of living in the home would depend only on Y, which no control affects, so that value is left out of the objective.[2]

3.4 Household optimisation

The household chooses consumption $c(t)$ and the risky share $\pi(t) \in [0, \bar{\pi}]$ continuously and fixes two choices at retirement, the coverage share ψ and the release intensity $\phi \in [0, 1]$. Each year the household draws the constant fraction ϕ of the smaller of the HEAS drawing cap and the headroom below the Maximum Loan Amount. The Maximum Loan Amount is the largest balance the scheme lends, a share of home equity that rises with age to 0.675 at age 90 (Section S1 gives the formula). Drawing stops when the HEAS loan balance reaches the Maximum Loan Amount. Consumption is bounded below by subsistence consumption $c^{\min}$, 0.85 of the household's maximum pension rate, and above by a liquidation constraint,

$$c^{\min} \leq c(t) \;\leq\; X(t)/\tau_c + I(t) - F(t),$$

with τ_c the shortest horizon over which liquid wealth X can be run down (Table 2). We write $c^{\max}$ for the upper bound, and the inflows $I(t)$ and home care fees $F(t)$ are given with the wealth process below.[3]

3.4.1 Wealth process and boundary conditions

Liquid wealth X satisfies, for $t \in [0, \tau)$ with τ the death of the last surviving member,

$$\mathrm{d}X(t) = \big(rX(t) + \pi(t)X(t)(\mu_1 - r) - c(t) + I(t) - F(t)\big)\,\mathrm{d}t + \pi(t)X(t)\sigma_1\,\mathrm{d}W_1(t), \tag{6}$$

where $I(t) = I^{\mathrm{AP}}(t) + I^{\mathrm{HEAS}}(t) + I^{\mathrm{LTCI}}(t)$ collects the household's inflows other than investment returns. The indemnity is $I^{\mathrm{LTCI}}(t) = \psi\,\mathrm{LTC}_{H(t)}$ in the insured states $\mathcal{H}'$ and zero elsewhere. Total home care fees are $F(t) = F^{\mathrm{HC}}_{\mathrm{ind}}(t) + F^{\mathrm{HC}}_{\mathrm{sp}}(t)$. Each member pays a share of that member's own part of LTC_h that rises as the age pension falls (Section 5), until the fees that member has paid reach the lifetime cap $\bar{\zeta}$. Where the fee exceeds what the household can pay, any shortfall against subsistence is met by the hardship waiver SAH provides, which also keeps the drift of X bounded. Wealth starts at $X(0) = X(0^-) - (1+\theta_L)(1-\eta)\psi P$ after the premium. At death $X(\tau)$ denotes the estate, which adds home equity net of the loan balance to liquid wealth,

$$X(\tau) = X(\tau^-) + \big(Y(\tau) - \kappa(\tau)\big)^+,$$

the positive part being the no-negative-equity guarantee, which caps estate recovery of the loan balance κ at home equity Y. The loan balance accrues at the HEAS interest rate r^{κ}, $\mathrm{d}\kappa = (r^{\kappa}\kappa + I^{\mathrm{HEAS}})\,\mathrm{d}t$, and the solver carries it as the loan-to-value ratio $\tilde{\kappa} = \kappa/Y$. The conditions stated here are the initial and terminal ones, and the spatial boundaries of Equation (11) are

[2] The omission leaves every decision and the sign of every welfare comparison unchanged, though not the welfare index of Section 6.1.1 (Section S10).

[3] The upper bound rules out consumption rates that are unbounded near the borrowing limit (Appendix A.3). It never falls below $c^{\min}$, because the hardship waiver stated with the wealth process meets any shortfall against subsistence.

imposed by the discretisation (Appendix A).

3.4.2 Preferences

Consumption utility is of constant relative risk aversion (CRRA) form and bequest utility is a shifted CRRA function,

$$U_1(c,h)=\frac{(\xi_1(h)c)^{1-\gamma}}{1-\gamma},\qquad U_2(x)=b\frac{(x+\bar{\xi}_2)^{1-\gamma}}{1-\gamma},\quad c,x>0, \tag{7}$$

with $\gamma\in(0,1)\cup(1,\infty)$ the relative risk aversion and $b>0$ the bequest intensity. The bequest shift $\bar{\xi}_2>0$ makes bequests rise more than in proportion to wealth, so that many households leave none (De Nardi et al., 2025). The joint-life weight ξ_1 makes consumption utility a sum over the living members. Each member's utility is weighted down as that member's health worsens, and while both members are alive, consumption is divided by an equivalence scale. For $h=\iota(i,j)$ the weight is

$$\xi_1(h)=\frac{1}{\bar{\xi}_1^{\mathbb{1}(i<n,j<n)}}\left[\frac{(n-\varepsilon(i-1))\mathbb{1}(i<n)+(n-\varepsilon(j-1))\mathbb{1}(j<n)}{n}\right]^{\frac{1}{1-\gamma}}, \tag{8}$$

where the indicator $\mathbb{1}(i<n)$ equals one while member i is alive, ε is the health dependence of consumption utility of Finkelstein et al. (2013) and $\bar{\xi}_1\in(1,2)$ the couple equivalence scale of De Nardi et al. (2025). Each member in state i contributes utility weighted by $1-\varepsilon(i-1)/n$.

3.4.3 Objective

The household's problem has an inner continuous-time part and an outer choice at retirement: (c,π) are adapted processes, while (ψ,ϕ) are $\mathcal{F}_0$-measurable constants. For each member, a fee accumulator, ζ^{ind} or ζ^{sp}, records the fees that member has paid and stops growing at the lifetime cap $\bar{\zeta}$ of Table 1. The pair (ψ,ϕ) enters the dynamics through I^{LTCI} and I^{HEAS} and, for ψ, through the deemed asset. For fixed (ψ,ϕ), with δ the subjective discount rate, the value function is

$$\begin{aligned}V^{\psi,\phi}(t,x,y,h,\kappa,\zeta^{\text{ind}},\zeta^{\text{sp}})=\sup_{(c(u),\pi(u))_{u\geq t}}\mathbb{E}\Big[&\int_t^\tau e^{-\delta(u-t)}U_1(c(u),H(u))\,\mathrm{d}u\\&+\;e^{-\delta(\tau-t)}U_2\big(X(\tau)\big)\Big].\end{aligned} \tag{9}$$

The household maximises over this family at $t=0$. With $\Pi(\psi)=(1+\theta_L)(1-\eta)\,\psi P$ the premium paid and $h_0=H(0)$,

$$\begin{aligned}(\psi^*,\phi^*)&\in\operatorname*{arg\,max}_{(\psi,\phi)\in\mathcal{A}}\;V^{\psi,\phi}\big(0,\;X(0^-)-\Pi(\psi),\;Y(0),\;h_0,\;0,\;0,\;0\big),\\ \mathcal{A}&=\Big\{(\psi,\phi)\in[0,1]^2:\;\Pi(\psi)\leq X(0^-)\Big\}.\end{aligned} \tag{10}$$

The affordability bound $\Pi(\psi)\leq X(0^-)$ keeps the premium paid within liquid wealth and binds at the bottom of the wealth distribution. The set $\mathcal{A}$ is compact and $V^{\psi,\phi}$ continuous in (ψ,ϕ), so the maximum is attained but need not be interior. In the computations the maximum is taken over ϕ in steps of 0.1 and over ψ on $\{0,0.02,\ldots,0.2,0.3,\ldots,1\}$, so $\psi^*=0$ means that no positive share in this set does better than none.

By dynamic programming, the value function $V^{\psi,\phi}$ satisfies the Hamilton–Jacobi–Bellman

(HJB) equation, superscripts suppressed,

$$
\begin{aligned}
0 = \sup_{c,\pi} \Bigg\{ & U_1(c,h) - \delta V + V_t + \underbrace{V_x(rx + \pi x(\mu_1 - r) - c + I(t) - F(t))}_{\text{drift of liquid wealth}} \quad + \quad \underbrace{V_y \mu_2 y}_{\text{drift of home equity}} \\
& + \underbrace{V_\kappa \Big(r^\kappa \kappa + I^{\text{HEAS}}(t) \Big)}_{\text{loan balance}} + \underbrace{V_{\zeta^{\text{ind}}} \mathbb{1}_{\{\zeta^{\text{ind}} < \bar{\zeta}\}} F^{\text{HC}}_{\text{ind}} \; + \; V_{\zeta^{\text{sp}}} \mathbb{1}_{\{\zeta^{\text{sp}} < \bar{\zeta}\}} F^{\text{HC}}_{\text{sp}}}_{\text{lifetime caps on fees}} \\
& + \frac{1}{2} V_{xx} (\pi \sigma_1 x)^2 + \frac{1}{2} V_{yy} (\sigma_2 y)^2 + V_{xy} \, \pi \rho \, \sigma_1 \sigma_2 x y \Bigg\} \\
& + \underbrace{\sum_{h' \neq h, \; h' \neq n^2} q_{h,h'}(t) \Big(V(t,x,y,h',\kappa,\zeta^{\text{ind}},\zeta^{\text{sp}}) - V(t,x,y,h,\kappa,\zeta^{\text{ind}},\zeta^{\text{sp}}) \Big)}_{\text{health transitions}} \\
& + \underbrace{q_{h,n^2}(t) \Big(U_2(x + (y-\kappa)^+) - V(t,x,y,h,\kappa,\zeta^{\text{ind}},\zeta^{\text{sp}}) \Big)}_{\text{death and bequest}} .
\end{aligned}
\tag{11}
$$

The means-tested budget set is non-convex (Hubbard et al., 1995), and the tapers act as a marginal tax on saving (Sefton et al., 2008), so the value function is locally non-concave. Spending down to the full age pension and saving through the taper band can yield comparable utility. We solve Equation (11) by backward-in-time policy iteration. The Hamiltonian is strictly concave in consumption, so its maximiser in consumption is the first-order solution projected onto $[c^{\min}, c^{\max}]$. The risky share is chosen by a direct search. A Craig–Sneyd alternating-direction implicit (ADI) scheme (Craig and Sneyd, 1988) treats the two diffusing states in logarithmic coordinates, and monotone upwind discretisations on a tensor-product grid treat the path-dependent states (Achdou et al., 2022). A vanishing-viscosity regularisation handles the kinks the means tests induce (Barles and Souganidis, 1991). Health transitions and mortality are treated implicitly, on one pre-factorised system per time level. On a means-tested problem these techniques need adjusting for two features, a value function nearly flat in $\ln x$ at the bottom of the solved grid and kinks inside the state space. Appendix A.3 records both and the four choices that followed: a liquidation horizon τ_c not tied to the time step, relaxed control updates, the hardship waiver, and a grid for $\tilde{\kappa}$ over the range a household can reach.

The Australian models closest to ours solve an annual-step problem by backward recursion (Andréasson et al., 2017) or by least-squares Monte Carlo (Andréasson and Shevchenko, 2024), and neither includes the path-dependent states that the lifetime cap and the loan balance force into the state vector here. Each of the three path-dependent states moves in one direction: the loan balance accrues interest, and each member's fees accumulate up to the lifetime cap. In continuous time each therefore enters Equation (11) through a transport term whose drift keeps one sign. A monotone upwind difference then treats each such term without an expectation or an interpolation (Appendix A). Equation (11) is written in the loan balance κ, which does not diffuse. The solver instead carries the loan-to-value ratio $\tilde{\kappa}$, whose drift at the calibrated rates keeps its sign too. Because $\tilde{\kappa}$ also moves with home equity, it adds second-order terms, which the solver evaluates explicitly. An annual recursion would instead interpolate the continuation value across all three states at every node of the solved grid.

4 Empirical analysis

This section estimates from HILDA and administrative data how much care a person with given health conditions needs, how health changes with age and how wealth is spread across homeowners at retirement. It then defines the clawback (Section 4.4).

4.1 Data and sample

We take the health and wealth data from the HILDA survey (Watson and Wooden, 2012), whose twenty annual waves (2004–2023) record demographics and long-term health conditions. The six quadrennial wealth modules add liquid financial assets net of non-housing debt and the equity in the owner-occupied home (Summerfield et al., 2024).[4] The public subsidies come from the SAH programme, which from November 2025 integrates the Commonwealth Home Support Programme (CHSP) and the four-level Home Care Packages (HCP) into one graded continuum of budget tiers, the lowest of which is Tier 1. Figure S2 sets out, with sources, how the two programmes map to SAH tiers, the age profiles of recipients, the means-tested fee rates and the assumed split of each tier's budget (DHDA, 2025a; Productivity Commission, 2025; AIHW, 2025; DHDA, 2025c; Nektarios et al., 2025). HILDA does not record an individual's care level, so the targets that link reported conditions to care levels are the administrative shares of care recipients in each age band across the CHSP level and the four HCP levels (Section 4.2).

We keep permanent sample members who either report specific long-term conditions or explicitly report none. For the joint-state sample, a cohabiting spouse aged 65 or over must meet the same criterion. Section S6.4 explains how missed waves are filled, how each trajectory ends, and how transitions are weighted and pooled.

Because the household of Section 3 ages in place, and home equity release presupposes that it keeps the home, we estimate the proportional odds model of Section 4.2 on community-dwelling respondents only. Residential care is means-tested under different rules and is outside the model. Only the proportional odds fit is restricted, and the classification of health states and the transition intensities use the full panel aged 65 and over, including respondents in residential care.[5]

4.2 Disability classification

4.2.1 Proportional odds model

We represent each respondent's latent care-need burden by a single index of the reported conditions, the care-need burden index. The probability of being at or above each care level rises with that index. We use a proportional odds model (McCullagh, 1980) with non-negative condition weights because it reflects three features of the aged-care needs assessment: care levels are ordered, the same burden index is compared with every cutpoint, and an additional condition can only raise the burden. Let $L \in \{1, \dots, K\}$ be the required care level, unobserved in HILDA, with $K = 5$ the CHSP level and the four HCP levels. Let $\mathbf{x}$ be the vector of the sixteen condition indicators, each binary report centred at its weighted prevalence and divided by the standard deviation that prevalence implies. Then

$$\Pr(L \geq k+1 \mid \mathbf{x}) = \left[1 + \exp\{-(\alpha + \beta^{\top}\mathbf{x} - c_k)\}\right]^{-1}, \qquad k = 1, \dots, K-1, \tag{12}$$

with intercept α and condition weights $\beta \geq 0$, whose entries sum to e^{σ_β} for a log burden scale σ_β estimated alongside them, and ordered cutpoints $c_1 < \cdots < c_{K-1}$. Only the differences $\alpha - c_k$ enter, so the intercept and the cutpoints are identified only up to a common shift.

We estimate the parameters by penalised grouped maximum likelihood on the community-dwelling person-years that report a condition. The fit matches the care-level shares the model predicts, weighted by HILDA cross-sectional weights, to the administrative shares in six five-year age bands (Section S6.2). Because those shares count care recipients only, and the model has no level for needing no care, the fit assumes that everyone who reports a condition needs at

[4] All financial variables are deflated to 2022 prices by realised consumer price index (CPI) inflation (World Bank, 2026), pooled over the six wealth waves on that base, and projected to 2026 prices. Section S5 lists the variables used.

[5] Institutionalised respondents are 2.0% of person-years with a reported condition, and removing them leaves both upper state boundaries of Section 4.2.2 unchanged (Section S6.1).

least CHSP-level care. Predicted and observed shares agree closely in every band and at every level (Section S6.3).

A few function-limiting conditions dominate the fit. Mobility and neurological impairment carry most of the weight, while "other long-term conditions", reported in 57.5% of the community-dwelling person-years that report a condition, carries half a per cent of it. The burden index sorts respondents into care levels sharply, because a single condition can raise it by more than the distance between the first and the last cutpoint (Tables S4 and S3). The care-level shares the fit predicts are stable under the bootstrap. The individual weights are not stable: the penalised likelihood is multi-modal, so the reported weight vector is one of several that fit the administrative shares almost equally well, and only the ranking of the largest weights is stable. The classification is made on the index that the reported vector produces: the state boundaries of Section 4.2.2 are placed on that index and, conditional on the reported vector, hold under the sample perturbations of Section S6.

4.2.2 Health states and transition dynamics

We model couples as two parallel lives starting at retirement, and the household's joint state is the ordered pair of the two individual states (Equation (1)). Starting both spouses at the same time removes gender and age-gap effects from the model, so that differences later attributed to health do not also reflect who is older. We pool transitions and exposure within six pooling groups: $[65, 70)$, $[70, 75)$, $[75, 80)$, $[80, 90)$, $[90, 100)$ and $[100, 117)$. Within each group, each joint state is pooled with its mirror image, the spouses swapped, because the health process should not depend on which spouse is the respondent. The grouped hazards are then smoothed with non-negative splines, corrected for reporting error with the measurement model of Jackson et al. (2003) (Section S7), and closed into a generator $Q(t)$ with valid intensities at every age. The economic results are computed on this one generator, so the uncertainty in the classification and in the reporting-error correction is not carried into them.

The burden index is the linear predictor $\alpha+\beta^{\top}\mathbf{x}$ of Equation (12) plus 20, which keeps every value strictly between 0 and 100. The values 0 and 100 are codes, for an interview that reports no condition and for death or censoring. In Figure 1 the annual transitions of the burden index fall into blocks, and a return to full health becomes less likely as severity and age rise, the age pattern that Fong et al. (2015) report. The blocks suggest three disability states on the burden index, $[5.8, 25.1)$, $[25.1, 51.4)$ and $[51.4, 93.1]$, with a region below 5.8 that behaves like health. Pale or white vertical and horizontal bands separate the blocks, and the dotted rules of the two upper boundaries lie near the bands' high-index side.

Each upper boundary is placed at the midpoint of the widest qualifying range of the burden index that no person-year occupies. For the boundary at 25.1 the range must contain the first care-level cutpoint, and for the boundary at 51.4 it must lie above the last. The ranges are measured on the panel with missed waves interpolated, the one on which the intensities are estimated, and each must meet further conditions. The lower boundary of health, 5.8, is set by the transitions. It lies just below 5.85, the lowest index of an interview that reports a condition, so below it lie only interviews that report none and a few person-years interpolated across missed waves. Section S8 gives these conditions, the margins of the rule, the lower boundary and the terminal condition on the fitted generator.

The boundary at 51.4 lies above the cutpoint for HCP Level 4 (33.5), so the top state isolates profound disability. The result is a five-state classification:

- *State 1:* No long-term condition reported (burden index below 5.8); healthy, without home care needs;
- *State 2:* Burden index in $[5.8, 25.1)$; the whole CHSP caseload, not the subset that SAH Tier 1 funds;
- *State 3:* Burden index in $[25.1, 51.4)$; the HCP caseload below the boundary at 51.4, including 58% of HCP Level-4 person-years;

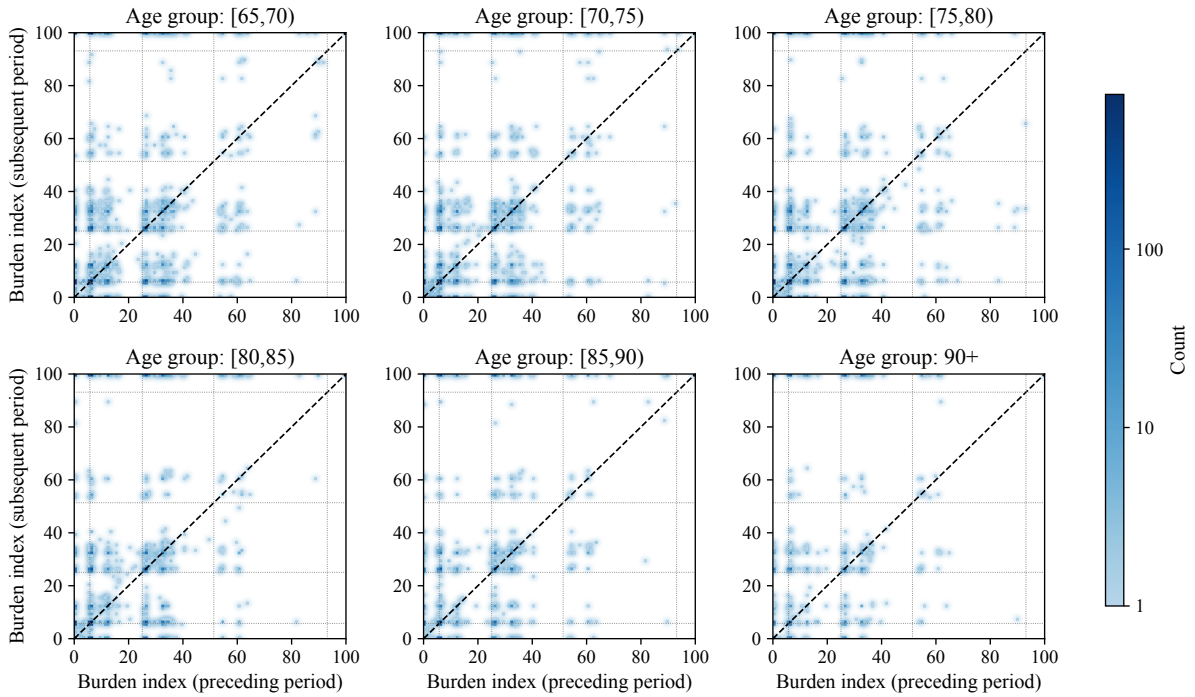


Figure 1. Annual transitions of the latent care-need burden index, HILDA person-years aged 65 and over, waves 4–23, by the five-year age bands of Section 4.2.1, not the pooling groups of Section 4.2.2. Each panel is a two-dimensional histogram of next year's index against this year's, with the number of person-years on a logarithmic colour scale; deteriorations lie above the diagonal and recoveries below it. The values 0 and 100 are codes, not index values, for an interview that reports no condition and for death or censoring. Dotted rules mark the state boundaries 5.8, 25.1 and 51.4 (Section 4.2.2) and the largest observed index, 93.1.

- *State 4:* Burden index in $[51.4, 93.1]$; profound disability, the top 42% of Level-4 person-years;
- *State 5:* Death.

States 3 and 4 are the high-need states, as is any joint state in which a living member is in one of them. Underwriting is confined to State 1, so the eligible joint states $\mathsf{C}_{\mathcal{H}}(\mathcal{H}')$ in Equation (3) are $\{1\}$ for couples and $\{5, 21\}$ for single households, and every other joint state is insured, although the indemnity is positive only in the care states. Each of the two single-household states has one member healthy and the other dead. Aggregate mortality is calibrated to the Australian Life Tables (AGA, 2024) from age 85, so that the life table supplies the old-age level of mortality. Below that age it is calibrated to the death intensity of the uncorrected generator, which the reporting-error correction leaves nearly unchanged (Section S7). Survival to each age up to 85 is therefore HILDA's own, and it lies above the life table's by up to 6.1 percentage points, reached at age 85. At every age the differences in mortality across health states are the panel's own (Section S8.2).

4.3 Household cells and the wealth grid

We evaluate the solved model in two ways: on the wealth grid defined below, which shows where a choice is made, and on household cells, which show how many households make it. The cells partition the joint distribution of liquid wealth X and home equity Y at retirement, estimated on HILDA homeowner households aged 60 to 69 and pooled over six wealth waves. Single homeowners hold less, and more dispersed, liquid wealth than partnered ones, and similar home equity (Section S5). The partition uses the wealth data alone. The health process estimated above enters through the solved model, which is read and simulated at each cell's median point, so a cell's result combines a position in the wealth distribution with the health risks that follow

it.

Each household type, partnered and single, is divided at its octiles of X and, within each octile, at the quartiles of Y conditional on that octile, giving 32 household cells per type and 64 in total. This partition is the Rosenblatt construction (Rosenblatt, 1952), so each cell nominally holds 1/32 of its household type by population weight. Of the households the model represents, 75.8 per cent are partnered and 24.2 per cent single, so a partnered cell nominally holds about 2.4 per cent of households and a single cell about 0.8 per cent.

We cut first on liquid wealth because both means tests assess it and the instantaneous clawback of Section 4.4 has all its kinks in liquid wealth. A home-equity quartile is therefore a rank within an octile, not a level: the top quartile begins at A$934,000 in the first partnered octile and at A$2,062,000 in the eighth. The value function is read at each cell's median point, the population-weighted median (X, Y) of the cell, marked in Figure 2. Because the partition uses nothing computed from the model, it does not change with the calibration.[6]

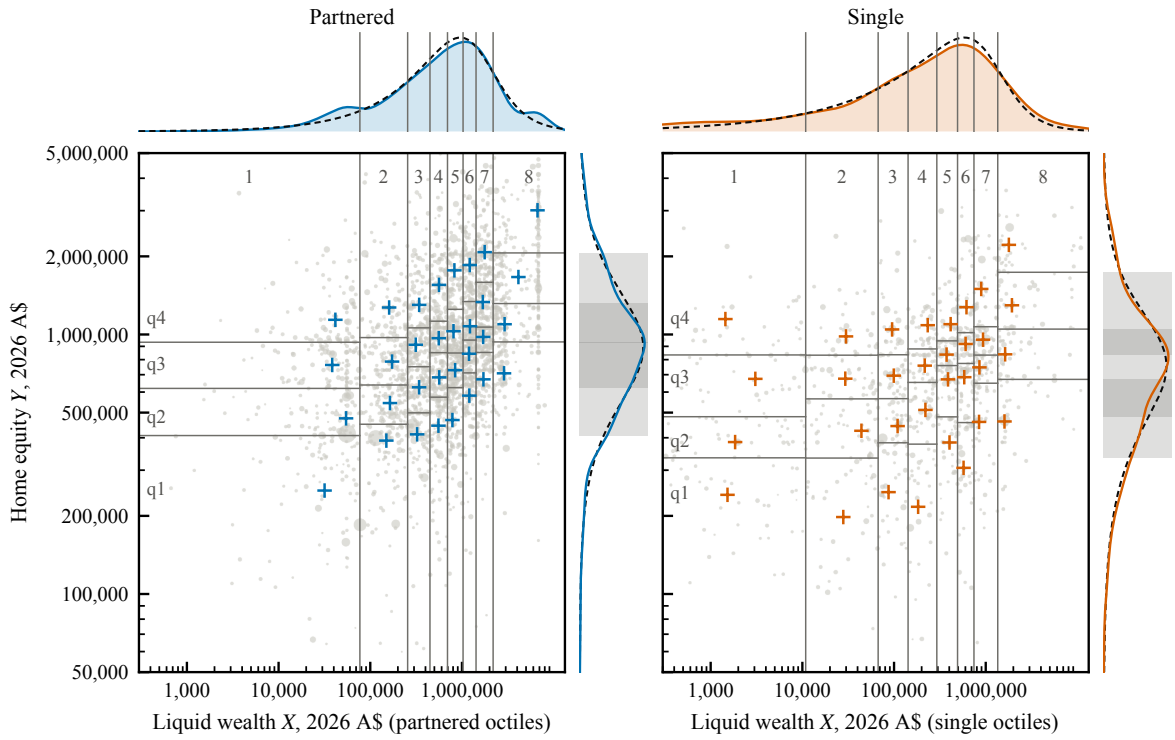


Figure 2. Joint distribution of liquid wealth X and home equity Y at retirement and the 64 household cells, for HILDA homeowner households aged 60 to 69, pooled over six wealth waves, in 2026 Australian dollars. Grey dots are households, sized by population weight; vertical rules cut X at its octiles, stepped rules cut Y at its quartiles within each octile, and crosses mark the population-weighted median (X, Y) of each cell. In the margins, solid curves are weighted kernel densities, dashed curves the fitted Dagum margins (Dagum, 1977), and each grey band the range over which one of the three home-equity quartile cuts moves across the eight octiles, darker where two bands overlap. The cuts also use the 1.2 per cent of modelled households that fall outside the plotted window.

The wealth grid is the lowest 60 liquid-wealth and 27 home-equity nodes of the solved grid, the grid on which the model is solved, spaced evenly in the logarithm of wealth. It runs from A$100 to A$13.3 million in liquid wealth and from A$10,000 to A$4.66 million in home equity, and every household cell's median point lies inside it (Appendix A). Sections 6.1 and 6.2

[6] Renters and homeowners with liquid wealth outside the span of the grid on which the model is solved are excluded (Section S5). All shares of households reported below refer to the homeowner population that the model can represent.

evaluate the optimal choices at each of its 1,620 nodes, with equal weight on each node. A share of the wealth grid therefore counts nodes, not households. Shares of households come from the household cells, on which both sections also report the optima. Each cell's optimum is read by interpolating the solved value functions between the nodes, so it can differ from the choices at the surrounding nodes (Section 6.1.1). Every table and figure from Section 6.3 onwards is evaluated on the 64 cells with the measured weights, and a result "by octile" aggregates that octile's four home-equity quartiles by weight. Present values are discounted to retirement at rate r.

4.4 The clawback of a private contract

A household that buys LTCI does not retain all of the indemnity. The indemnity is assessable income, so the age pension falls and the home care fee rises. The clawback is the net amount that the means tests take from a household, in age pension withdrawn and home care fees raised, because it holds a private LTCI contract. A negative clawback is an implicit subsidy. No statute defines the clawback. We measure it in two ways, which are not interchangeable. The first, the instantaneous clawback, is a point-in-time accounting measure that applies the statutory rules at one liquid wealth, one health state and one value of LTCI cover, with no optimised policy entering it. The second, the realised clawback, is the lifetime measure that Section 6.4.2 and Table 6 report, the expected discounted difference between an insured and an uninsured simulation of the household's own path.

4.4.1 The instantaneous clawback

Consider a household without home equity release, in joint health state h, with liquid wealth X and cover ψ, below the lifetime cap. The pension reduction and the fee increase combine into the instantaneous clawback

$$\mathcal{W}^{\varnothing}(X;h,\psi) = \underbrace{-\Delta I^{\mathrm{AP}}(X;h,\psi)}_{\substack{\text{age pension}\\ \text{reduction}}} + \underbrace{\Delta F(X;h,\psi)}_{\substack{\text{home care}\\ \text{fee increase}}}, \tag{13}$$

where each Δ compares cover $\psi > 0$ with no cover, and each is a flow in A$ a year. The signs are fixed by construction, $\Delta I^{\mathrm{AP}} \leq 0$ and $\Delta F \geq 0$, since the indemnity raises assessable income, the contract is a deemed asset while held, and both the pension and the fee move on those assessments. Hence $\mathcal{W}^{\varnothing} \geq 0$: the instantaneous clawback is never a subsidy. The realised clawback RC of Section 4.4.2 can be negative.

Equation (13) is piecewise linear in X, with at most seven kinks set by the legislated parameters of the means tests. The fee adds no kink of its own, because it is affine in the pension. The instantaneous clawback does not depend on home equity, because the home is exempt from the assets test.

Figure 3 draws the instantaneous clawback for three households whose living members are all in State 2, all in State 3 or all in State 4. The two dotted rules are an uninsured household's thresholds: the income test's lower threshold, expressed in liquid wealth, and the upper threshold at which its pension ends. Because the fee rises as the pension falls, they also mark where that household's fee begins and stops rising. The indemnity itself counts as income, so in States 3 and 4 the insured household's pension is reduced even at zero liquid wealth, and those curves start above zero. The State-4 curve is hump-shaped, peaking at or just above the first rule, where the uninsured household's pension begins to fall as well. In State 2 the assessed indemnity alone is below the income test's lower threshold, so that curve stays at zero until deemed income from liquid wealth makes up the difference, below the first rule. Section S9 places the kinks in closed form and gives each curve's value at the State-4 peak.

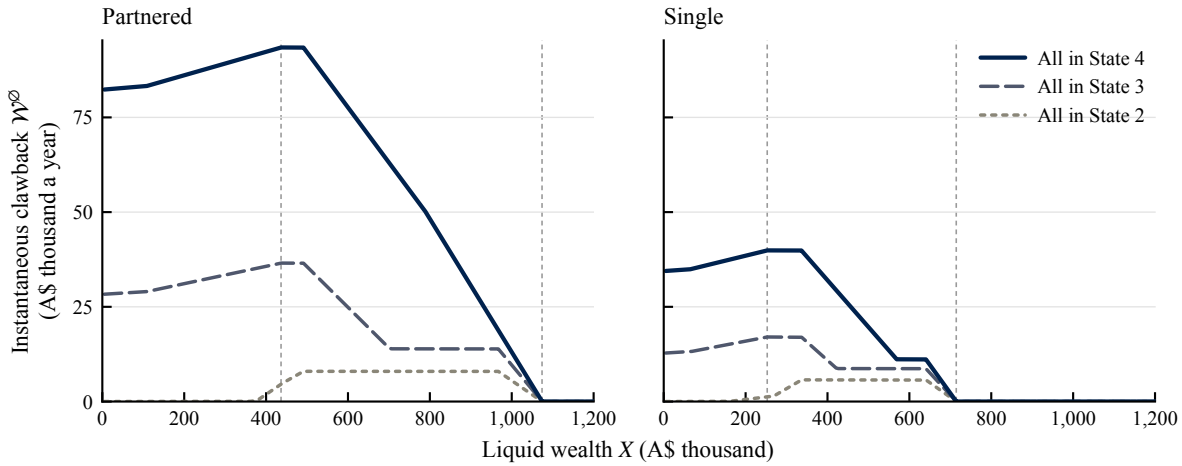


Figure 3. The instantaneous clawback $\mathcal{W}^{\varnothing}$ of Equation (13) against liquid wealth, at full cover and without home equity release. Each curve is a joint state in which every living member is in State 2, 3 or 4 (Section 4.2.2). The instantaneous clawback is an amount per year and liquid wealth an amount held, both in 2026 Australian dollars. The dotted vertical rules are an uninsured household's two thresholds, the income test's lower threshold expressed in liquid wealth and the upper threshold at which the pension ends: A$436,509 and A$1,074,000 for couples (left), A$252,800 and A$714,500 for singles (right).

4.4.2 The realised clawback

Let $(c^*, \pi^*, \psi^*, \phi^*)$ denote the optimal policy. Let $I_t^{\mathrm{AP}}(\psi, \phi)$ and $F_t(\psi, \phi)$ denote the pension received and the home care fee charged at date t, the fee measured before the lifetime cap, so that the realised clawback reflects the means tests alone (Section S14). Both are evaluated on the wealth path that the household's optimal consumption and risky share generate at the pair (ψ, ϕ). Let $\Delta_\phi Z_t = Z_t(\psi, \phi) - Z_t(0, \phi)$ be the difference between the insured and uninsured simulations under the same history of health and asset returns. Then the realised clawback without home equity release is

$$RC^{\varnothing} = \mathbb{E}\Big[\int_0^{\infty} e^{-rt}\left(\Delta_0 F_t - \Delta_0 I_t^{\mathrm{AP}}\right) \mathrm{d}t\Big], \tag{14}$$

and with home equity release,

$$RC = \mathbb{E}\Big[\int_0^{\infty} e^{-rt}\left(\Delta_\phi F_t - \Delta_\phi I_t^{\mathrm{AP}}\right) \mathrm{d}t\Big]. \tag{15}$$

Section 6.4.2 evaluates the realised clawback with cover set at a fifth of full cover for every household and ϕ held at 0 or at 1. Table 6 reports it at the household's own (ψ^*, ϕ^*) at a premium rebate of $\eta = 0.30$.

Because the insured and the uninsured household hold different liquid wealth from the premium onwards, the realised clawback includes an effect that the instantaneous clawback does not: paying the premium reduces assessable wealth, and an unspent indemnity increases it. Like $\mathcal{W}^{\varnothing}$, the realised clawback also contains the deemed asset and the assessed indemnity. Because the home care fee is measured before the lifetime cap, a change in when a member's fees reach the lifetime cap appears in the government's net expenditure instead of the realised clawback (Table 6 and Section S17.1). The expectation runs over the health process, survival included, over asset returns, and over the policy they induce. The realised clawback has no closed form, and we compute it by simulation.

5 Calibration

Table 2 lists the parameters. The model is stated in real terms: wealth is in 2026 Australian dollars, statutory rates are held fixed, and every rate is net of inflation. This section records the choices that required judgement. Section S10 gives the derivations behind them.

Table 2. Baseline model calibration (all rates real, monetary values in A$).

Symbol	Description	Value	Source / reference
Market parameters			
r	Risk-free rate	2.5%	Treasury Indexed Bond yield; see text
μ_1	Equity total return	7.8%	Jordà et al. (2019)
σ_1	Equity volatility	16.9%	Jordà et al. (2019)
μ_2	Constant-quality home-equity drift	1.9%	Jordà et al. (2019); see text
σ_2	Home-equity capital-gain volatility	11.9%	Jordà et al. (2019)
ρ	Equity–home-equity correlation	0.10	Jordà et al. (2019); see text
r^{κ}	HEAS interest rate	5.41%	Adverse spread over r; see text
Preference parameters			
γ	Relative risk aversion	3.701	De Nardi et al. (2025)
δ	Subjective discount rate	0.03	De Nardi et al. (2025)
$\bar{\xi}_1$	Couple equivalence scale	1.528	De Nardi et al. (2025)
b	Bequest intensity	8.651×10^7	De Nardi et al. (2025); see text
$\bar{\xi}_2$	Bequest shift	A$7,114,000	De Nardi et al. (2025); see text
ε	Health dependence of consumption utility	1	see text; Finkelstein et al. (2013)
LTCI parameters			
$\widetilde{\text{LTC}}_2$	Gross care cost before informal care, State 2	A$4,375	see text
$\widetilde{\text{LTC}}_3$	Gross care cost before informal care, State 3	A$37,784	see text
$\widetilde{\text{LTC}}_4$	Gross care cost before informal care, State 4	A$69,021	see text
IC	Maximum informal care	A$3,850	CHSP average (assumed); see text
ϖ	Decline rate of informal care	1/3	assumption; see text
θ_L	Expense loading on fair premium	0.18	Brown and Finkelstein (2007)
η	Premium rebate on loaded price	0.1525	$\theta_L/(1+\theta_L)$; see text
ν	Means-test intensity, lifetime income stream	0.6	Services Australia (2026b)
Other parameters			
$\bar{\pi}$	Upper bound on the risky share	1	Catherine (2022)
τ_c	Liquidation horizon of liquid wealth (years)	1.0	Regularisation; Appendix A.3

Notes: Rates and drifts are real. Care costs and informal care are annual amounts. For the gross care costs, see the text and Figure S2(a). Figures numbered with an S prefix are in the online supplement.

5.1 Market parameters

The equity and housing return parameters are the long-run Australian estimates of Jordà et al. (2019), so the model is calibrated to long-run averages rather than to the period since 2010. The correlation ρ is that of the real equity total return and the real housing capital gain in the same data, 0.095, rounded to 0.10. Equity enters as a total return. Home equity enters at a constant-quality drift through its capital-gain component alone, because an owner-occupier never receives the rental component. Jordà et al. (2019) publish a capital gain of 2.53%, the growth of median transacted prices in the house price series of Knoll et al. (2017). The quality-and-composition adjustment of -0.6 percentage points (Stapledon, 2007; Fox and Tulip, 2014) gives $\mu_2 = 1.9\%$. Maintenance that keeps the home at constant quality is not charged against liquid wealth, which may understate the household's need to draw cash from home equity release (Cocco and Lopes, 2020).

The risk-free rate is the current real yield on Treasury Indexed Bonds, $r = 2.5\%$, and the post-1950 Australian real bond return of Jordà et al. (2019), 2.45%, is close to it. The HEAS interest rate is set by the Minister in a determination rather than by a market (Minister for Social Services, 2026), and it enters the model as a spread over r, because a spread does not depend on a deflator. Since the statutory rate is nominal, the spread is measured against a nominal yield, the ten-year Commonwealth bond yield. The spread averaged 0.96 percentage points from January 2020 to June 2026 (Reserve Bank of Australia, 2026), but its means over

the three statutory periods since 1998 range from −0.12 to 3.30 percentage points. A household that fixes its release intensity at retirement cannot rule out a less favourable period. It therefore plans on an adverse spread, the average plus a margin of 1.95 percentage points, and r^{κ} is then 5.41% (Section S10). The margin guards against a widening of the spread, as in 2020 and 2021, and not against a rise in the statutory rate. It is set from the dispersion of the three period means, so that a likelihood-ratio test would confuse the two spreads with probability 0.10. The robustness literature would still class a margin of this size as hard to detect (Anderson et al., 2003; Maenhout, 2004). Only the HEAS interest rate carries a margin, because the Minister can reset it at discretion.

At $r^{\kappa} = 3.46\%$, the rate without the margin, the insured share of households where home equity release is offered is unchanged, and the mean release intensity rises from about three fifths to over nine tenths (Table 3). The current statutory rate is 3.95% in nominal terms (Services Australia, 2026a), about 1.4% in real terms at the midpoint of the Reserve Bank of Australia's inflation target band, below both r and the rate without the margin.

5.2 Preference parameters

Preferences follow De Nardi et al. (2025). Relative risk aversion γ and the equivalence scale $\bar{\xi}_1$ are dimensionless and transfer without adjustment, and $\delta = 0.03$ corresponds to their annual discount factor. The bequest intensity b and the bequest shift $\bar{\xi}_2$ are re-anchored through two quantities of Equation (7): the marginal propensity to bequeath, 0.993, and the activation threshold $\bar{c} = \bar{\xi}_2 b^{-1/\gamma}$, the consumption rate below which a household leaves no bequest. Above the threshold, the marginal propensity to bequeath is the share of an extra dollar that a household sure to die within the year would bequeath rather than consume. The single-household threshold is anchored at $\bar{c} =$ A\$51,000, from the replication files of De Nardi et al. (2025). Section S10 derives 0.993 and $\bar{\xi}_2$. Compared with Lockwood (2018) and Ameriks et al. (2020) on these two quantities, the motive is stronger once active (0.993 against 0.96 and 0.48) but starts at a higher consumption rate (A\$51,000 against A\$37,000 and A\$17,000).

We depart from De Nardi et al. (2025) in omitting side bequests at the first death. Australia levies no inheritance tax, the survivor's residence stays exempt, and superannuation passes tax-free to a spouse, so the first death forces no reallocation. The health dependence $\varepsilon = 1$ lowers a single person's marginal utility of consumption by 20% of its healthy level for each state by which health worsens, to 0.4 of that level in State 4. The fall per state is close to the 17.9% fall per additional chronic condition that Finkelstein et al. (2013) estimate. Their fall is also measured from the marginal utility of a healthy person.

5.3 LTCI parameters

Each state's gross care cost, before informal care, is the average service cost of the care levels that fall in it, weighted by their numbers of recipients, at the SAH budgets of November 2025. A care level's service cost is the total cost of its services. Assuming the government budget is fully used, we infer it as that budget plus the home care fees of full pensioners. No fee is charged on the 10% care-management share, which the government funds in full. State 2 is costed differently, because, with total funding unchanged, only about a third of the CHSP caseload it holds receives Tier 1 funding. We cost it at A\$4,375, the average CHSP spend converted to a total service cost by the same rule, and not at Tier 1's service cost (Section S10.3). The A\$4,375 therefore measures care received rather than care needed.

Informal care from a healthy spouse reduces the gross care cost by up to IC = A\$3,850, the average annual CHSP spend, less $\varpi = 1/3$ of that amount for each state by which the carer's own health has worsened. The level and the rate of decline are both assumptions rather than estimates; the level treats a healthy spouse as able to give the low-level help that CHSP pays for. A healthy spouse's informal care meets about seven eighths of the State-2 cost. A single household has none, so it faces roughly eight times the couple's State-2 care cost net of informal care (the role of family care in Mommaerts, 2025). Under SAH, the household's share of its care

cost net of informal care weights its fee rates for independence services and for everyday-living services by their shares of the budget. Both rates rise as the age pension falls. A full pensioner, who receives the maximum rate of the age pension, therefore pays 12.0%, 9.6% and 7.6% of the care cost net of informal care in States 2, 3 and 4. A self-funded household, which receives no age pension, pays 63.0%, 57.2% and 52.5%.

Under the fitted generator $Q(t)$, the actuarially fair premium for full LTCI cover at retirement is A$151,101 for the joint-life contract and A$104,861 for the single-life contract. With the 18% expense loading of Brown and Finkelstein (2007), the loaded prices are A$178,299 and A$123,736. The baseline premium rebate is $\eta = \theta_L/(1+\theta_L) = 0.1525$, at which the government meets the expense loading and the household pays the actuarially fair premium.

6 Results

We solve the model using the calibration described in Section 5 and report the results by marital status, wealth and the instruments available to the household. Single and partnered households face different pension rates, care costs and contracts. Wealth consists of liquid wealth X, which both means tests assess, and home equity Y, which both exempt. We define four markets by restricting the choice set $\mathcal{A}$ in Equation (10): neither instrument ($\psi = \phi = 0$), cover only ($\phi = 0$), release only ($\psi = 0$) and both instruments (the full set $\mathcal{A}$). The four markets use the same value functions $\{V^{\psi,\phi}\}$ from Equation (9) and the same histories. They differ only in the choices over which the household optimises. We summarise demand for LTCI cover by the insured share of households, the population-weighted share of households whose optimal cover is positive. Section 6.1 reports baseline demand, and Section 6.2 examines its sensitivity. Section 6.3 traces the paths to the estate, and Section 6.4 explains the low demand for cover. Section 7 considers what these results mean for the government.

6.1 Demand for the two instruments at the baseline

Figure 4 reports the baseline demand for LTCI cover and home equity release at every node of the wealth grid inside the window of Figure 2, under the three markets that offer at least one instrument.

6.1.1 The level of demand

We find that demand for the hypothetical LTCI contract is low. In the market with both instruments, one household cell, 2.4 per cent of households, holds cover at its optimum. It is a cell of partnered households at the top of the distribution of liquid wealth at retirement, with high home equity, and it holds 4 per cent of full cover without drawing on the HEAS. We compare a household's choices by the welfare index $\left((1-\gamma)V^{\psi,\phi}\right)^{1/(1-\gamma)}$, a monotone transform of the value at retirement $V^{\psi,\phi}$ in Equation (10).[7] The cell's cover raises its welfare index by A$0.10, under one hundred-thousandth of the welfare index, so the cell is close to indifferent between that cover and none. In the market with cover only, the insured share of households is 18.2 per cent, and each insured household buys between 2 and 8 per cent of full cover. Apart from the households of the cell above, the insured households lie in the middle of the wealth distribution, and most of them have home equity above the median for their liquid wealth (Section 6.3).[8]

[7] The welfare index is in 2026 Australian dollars and ranks choices under one set of preferences. Because utility depends on health and the value at retirement includes the utility of the bequest, differences in the welfare index are not amounts a household would pay.

[8] In the market with cover only, cover raises the welfare index of each household cell that insures by A$0.05 to A$5.23, under five hundredths of a per cent of the welfare index. Three of these cells, including the cell above, lie between nodes of the wealth grid that all hold zero cover, so Figure 4 shows their cover only by the crosses at their median points (Section 4.3). With the ceiling of the solved grid raised to A$218 million of liquid wealth or to A$526.5 million of home equity (Appendix A), the gain of the cell above moves by at most A$0.02, and its choice stays at 4 per cent of full cover.

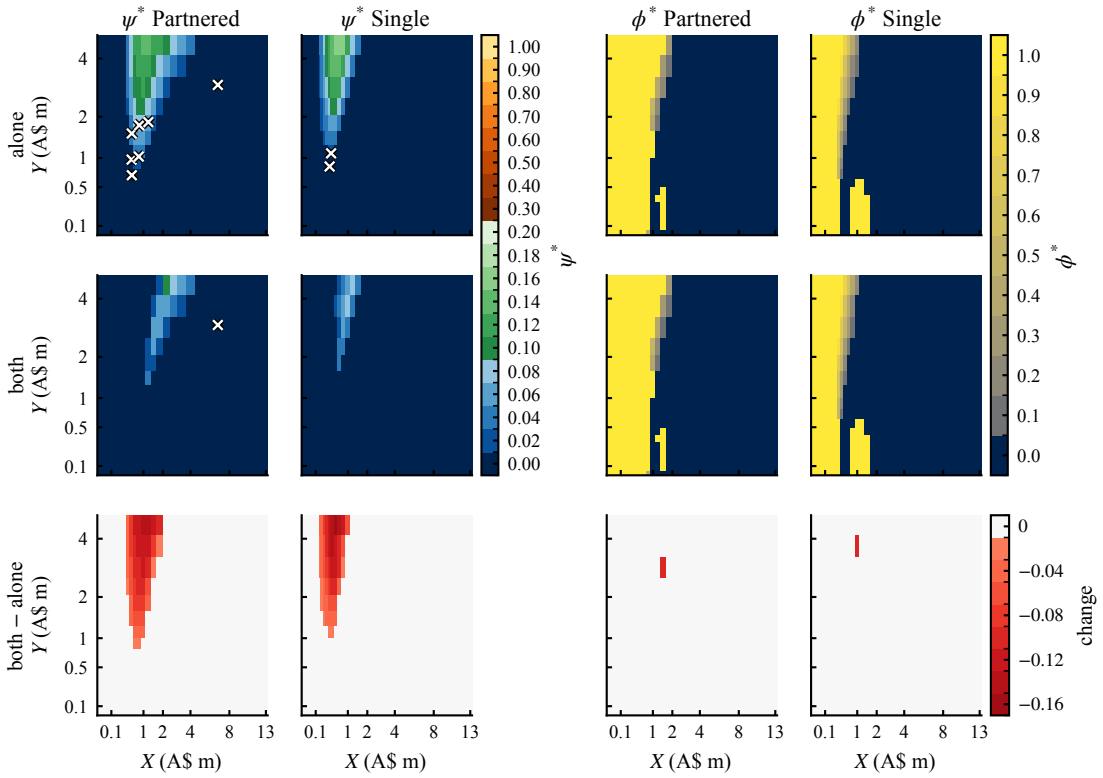


Figure 4. Optimal LTCI cover ψ^* and release intensity ϕ^* at the baseline on the unweighted nodes of the wealth grid (Section 4.3), 1,080 of its 1,620 for each household type, so the panels show where a choice is made, not how many households make it. Top row: each instrument offered alone, with $\phi = 0$ or $\psi = 0$ imposed; middle row: both offered, the optimum over $\mathcal{A}$ of Equation (10); bottom row: the middle row less the top row, red where the other instrument reduces a holding. In the cover panels of the top two rows, crosses mark the median points of the household cells that hold cover in that market. Liquid wealth X and home equity Y, in 2026 Australian dollars, are on square-root scales over the window of Figure 2, and each colour bar of the top two rows has one band for each value of ψ or ϕ over which the maximum is taken (Section 3.4.3).

The expense loading is not what keeps demand for LTCI cover low. At the baseline the premium rebate meets the loading (Section 5), so households pay the actuarially fair premium. Without the rebate they pay 1.18 times that premium, and demand falls further. No household insures in the market with both instruments, and in the market with cover only the insured share falls from 18.2 to 4.7 per cent of households, all of them partnered (Figure S6(a)). Brown and Finkelstein (2007) find that failures on the insurers' side do not by themselves explain why the private LTCI market is small. Under institutions that they did not study, the small amount of cover bought here at the fair premium matches that finding. The model contains most of the other explanations offered for a small market. They are the means tests' assessment of the contract, the affordability bound, under which the premium cannot exceed liquid wealth at retirement (Section 3), and the health dependence $\varepsilon = 1$, which lowers the marginal utility of consumption in the care states (Finkelstein et al., 2013). They also include a bequest motive (Lockwood, 2018) with a marginal propensity to bequeath at the strong end of the literature. Section 6.2 shows that two of them, the assessment and the bequest motive, each hold demand down.

6.1.2 The two choices over the wealth grid

In each panel of Figure 4, liquid wealth at retirement increases to the right and home equity upwards. The first two columns show LTCI cover and the last two release intensity, for partnered

and then for single households. In the market with cover only (top row), the nodes at which households buy cover form a blue and green region in the upper part of each cover panel, where home equity is high. The region is widest at the top and narrows further down until it ends. Cover is largest in the green upper part and falls towards the blue edges and lower tip. Elsewhere the panel is dark blue, and optimal cover is zero at every node.[9] Cover is worth little because the care subsidy already pays most of the care cost (Section 6.4.3).

Home equity release depends mainly on liquid wealth. Each release panel has a yellow region on the left, the drawing region, in which nearly all households draw as much as the HEAS allows. Nearly all households to the right of its edge draw nothing (dark blue). Drawing is worth most to households with little liquid wealth (Section 6.4.3). Each drawing adds to the HEAS loan, which is repaid, with interest, from the estate, and that repayment discourages drawing (Section 6.2.2). The edge of the drawing region leans to the right as home equity rises, which fits the rule that a larger home allows a larger loan (Table 1). Where home equity is larger, a grey band of partial drawing runs along the edge. In the lower part of each release panel, a small yellow patch of full drawing lies to the right of the edge, separated from the drawing region by dark blue. The patch begins just above the liquid wealth at which the assets test withdraws the age pension, so drawing there costs no pension even though drawings not yet spent count as assessable wealth. The loan can also be invested at an expected return above its interest rate.[10]

6.1.3 The displacement of cover at the baseline

At the baseline, release reduces cover and raises it for no household, whereas offering cover leaves every household's drawing unchanged. We measure the displacement of LTCI cover by home equity release as the fall in the insured share of households, in percentage points, when release is also offered. The insured share falls from 18.2 to 2.4 per cent. Every household that stops insuring draws on the HEAS instead, and most of them draw as much as it allows.

The wealth grid shows the same asymmetry. With both instruments (middle row of Figure 4), the region of cover shrinks to a narrow, mostly blue wedge along its right-hand edge. Over most of the region that holds cover in the top row, households draw on the HEAS and hold no cover, because drawing there lowers the value of cover below its price. Section 6.4.1 traces this to drawing in the care states by allowing drawing only while a member is in high need. In the wedge, households that do not draw keep their cover, and those that draw in part hold less than with cover only. In the bottom row, red in the cover panels marks the nodes at which offering release lowers cover, over most of the region of cover in the top row. Release therefore lowers cover on a small part of the wealth grid and raises it nowhere. The release panels of the bottom row are white except for a small red mark in each.[11]

At the baseline, home equity release lowers LTCI cover in our model, whereas it raises cover in Davidoff (2010) and Shao et al. (2019). The models differ in what happens to the home: in both studies the household sells it on entering residential care, so the sale proceeds and the LTCI indemnity provide funds in the same state. Borrowing against the home before care reduces the equity left to sell and so weakens the self-insurance that housing provides, which makes LTCI more valuable. Achou (2021) similarly finds that greater housing liquidity modestly raises cover when Medicaid counts the indemnity in its means test. In our model, drawing on the HEAS instead provides funds while the household keeps the home, including in the care states. In Xu et al. (2023), which does not include borrowing against the home, a larger stock of home equity lowers cover, in the same direction as release does here, although that result concerns housing

[9] With cover only, optimal cover is exactly zero on more than 96 per cent of the wealth grid.

[10] Full drawing beats none in the patch by at most A$18, under two tenths of a per cent of the welfare index, and with the risky share bounded at 0.5 the patch largely disappears.

[11] Offering release lowers cover on 2.7 per cent of the wealth grid for partnered households and 3.0 per cent for single ones, by 7 per cent of full cover on average and by at most 16 per cent. At the one node of each type where offering cover lowers release intensity, it falls from 0.2 to 0.1 for partnered households and from 0.1 to 0 for single ones.

wealth rather than access to release. Section 6.4.1 develops this explanation, and Sections 6.2.1 and 6.2.2 examine how the result depends on the assessment of the contract and on the bequest motive.

6.2 Sensitivity analysis

Four exercises each change one feature of the model and hold the rest of the calibration fixed, except that a change in risk aversion also re-anchors the bequest intensity b (Section 6.2.3). The first, the exemption exercise, exempts the LTCI contract from the means tests. The next two are the bequest exercise and the risk-aversion exercise, which change preferences: the bequest motive sets what it costs to draw on home equity, and relative risk aversion sets what the indemnity is worth as insurance. The fourth, the timing exercise, changes when drawing under the HEAS is allowed. Each exercise is compared with the baseline at the same node of the wealth grid. For a node (x, y) and a scenario s, let

$$\mathcal{V}_s(x,y) = \max_{(\psi,\phi)\in\mathcal{A}} V_s^{\psi,\phi}(0,\ x-\Pi(\psi),\ y,\ h_0,\ 0,\ 0,\ 0) \tag{16}$$

be the value at retirement under that scenario, with $\Pi(\psi)$ the premium of Equation (10) and $s = 0$ the baseline.

The two preference exercises are compared through the optimal choices (ψ^*, ϕ^*) alone. A scenario that changes the utility function describes a different household, so its value at retirement $\mathcal{V}_s$ cannot be compared with the baseline's. The exemption exercise leaves preferences unchanged but asks how demand responds to the assessment of the contract, so it too is compared through the optimal choices. The timing exercise holds preferences fixed, so its rules are compared by the change in the welfare index of Section 6.1.1, $((1-\gamma)\mathcal{V}_s)^{1/(1-\gamma)}$.[12] Table 3 reports the demand outcomes of the first three exercises and, as additional checks, of the bound on the risky share of Section 6.3 and of the HEAS interest rate without its robustness margin (Section 5).

Table 3. Demand for LTCI cover and home equity release by scenario, in the market with both instruments.

	Wealth grid		Household cells			
	Share with zero cover (%)					
Scenario	Partnered	Single	Cells with cover	Insured share (%)	Mean ϕ^*	Mean ψ^*
Baseline	98.8	99.0	1	2.4	0.618	0.001
Contract exempt from the means tests, $\nu = 0$	89.7	91.7	26	47.1	0.607	0.079
No bequest motive, $b = 0$	68.1	59.2	46	73.5	1.000	0.190
Relative risk aversion $\gamma = 3.33$	99.9	99.7	0	0.0	0.756	0.000
Relative risk aversion $\gamma = 4.07$	97.8	98.0	7	15.0	0.604	0.009
Risky share bounded at $\bar{\pi} = 0.5$	92.8	94.5	22	39.3	0.477	0.027
HEAS interest rate $r^{\kappa} = 3.46\%$	99.9	100.0	1	2.4	0.929	0.001

Notes: The two wealth-grid columns are shares of the 1,620 nodes of the wealth grid (Section 4.3). The four household-cell columns are over the 64 household cells, and the last three of them are population-weighted, with households that do not insure counted at zero in the mean of ψ^*. Each row after the first solves the model again with the change named and holds the rest of the calibration fixed, except that the two γ rows re-anchor the bequest intensity b (Section 6.2.3).

6.2.1 The assessment of the contract

With the contract assessed as we assume, the means tests over a lifetime charge LTCI cover very little anywhere and give it an implicit subsidy in the middle of the distribution of liquid wealth at retirement (Section 6.4.2). To test whether assessing the contract nevertheless lowers demand for cover, we exempt the contract from both means tests ($\nu = 0$). The tests then count neither the contract as an asset nor the indemnity as income, so the whole premium leaves

[12] The choices alone do not rank the rules, because the release intensity ϕ applies only in years in which a rule allows drawing (Section 3), so the same ϕ means less drawing under a rule that allows fewer years. Section S12 gives the loss in the welfare index under the delayed rules.

assessable wealth, and paying it raises the age pension and lowers the home care fee by more. In the market with both instruments, the exemption raises the insured share of households from 2.4 to 47.1 per cent (Table 3).

The exemption brings cover mainly to the middle of the wealth distribution. It lowers the release intensity ϕ^* for 4.8 per cent of households, all of whom now buy cover, and leaves it unchanged for the rest. The mean of ϕ^* therefore barely moves (Table 3), and almost all of the response is in cover.

The displacement of LTCI cover by home equity release at the baseline depends on the assessment of the contract. With the contract exempt from the means tests, release lowers the insured share by 0.8 percentage points, from 47.9 to 47.1 per cent, against 15.8 points at the baseline (Section 6.1.3). The whole fall comes from one group of single households, 0.8 per cent of all households, which stops insuring when release is also offered. Among the other households that insure, release lowers cover for some and raises it for others.[13] The households that hold both instruments make up 22.0 per cent of all households, nearly half of those that insure, and most of them draw as much as the HEAS allows. The two instruments are still not complements in value: in every household cell that holds both, the gain from holding both is smaller than the sum of the gains from each alone (Sections S15 and S16).

6.2.2 Bequest motive

We remove the bequest motive, setting $b = 0$ at $\gamma = 3.701$. The motive lowers the demand for both instruments. Wealth left unspent at death has value as a bequest (Lockwood, 2018), and every drawing accrues interest at r^{κ} against the estate. Without the motive, the insured share is 73.5 per cent, against 2.4 per cent at the baseline (Table 3). Every household then draws as much as the HEAS allows, against 60.8 per cent of households at the baseline (Section 6.3).

Without the bequest motive, release raises the insured share by 1.5 percentage points, from 72.0 to 73.5 per cent, against a fall of 15.8 points at the baseline. Release then raises cover for 32.7 per cent of households and lowers it without removing it for 5.6 per cent. The finding that release displaces cover therefore depends on both the bequest motive and the assessment of the contract (Section 6.2.1), and is not a property of the instruments themselves.[14] In Xu et al. (2023) the bequest motive weakens the displacement of cover by housing wealth and reverses it for low-wealth homeowners, because the home is sold on entering residential care and cover preserves the home as a bequest. Here the home is kept, and residential care is outside the model, so the motive acts on the drawing instead and strengthens the displacement of cover by release.

6.2.3 Relative risk aversion

We move relative risk aversion ten per cent either side of the baseline, to $\gamma = 3.33$ and 4.07, and re-anchor the bequest intensity b so that the marginal propensity to bequeath and the activation threshold of Section 5 keep their baseline values. In the market with both instruments, cover falls when risk aversion is lowered, and rises when it is raised (Table 3). The households that insure at $\gamma = 4.07$ lie in the upper part of the wealth distribution. In the market with cover only, 4.7 per cent of households, all of them partnered, still buy cover at $\gamma = 3.33$, against 18.2 per cent at the baseline. At the lower value, the release intensity ϕ^* changes for 14.9 per cent of households, and nearly all of them move to full drawing.

On the wealth grid, cover in the market with both instruments begins at a higher liquid wealth than at the baseline when $\gamma = 3.33$, and at a lower one when $\gamma = 4.07$ (Section S11). Under either change, households with little liquid wealth still draw as much as the HEAS allows, and release still lowers cover on part of the wealth grid and raises it nowhere (Sections 6.1.2

[13] The group of single households above buys 2 per cent of full cover with cover only, the smallest positive value of cover solved. Release lowers cover without removing it for 7.8 per cent of households and raises it for 7.1 per cent.

[14] The model is solved only at the two endpoints, b at its baseline and at zero, so the exercise gives the sign of the motive's effect and not how that effect varies with the strength of the motive.

and 6.1.3). The insured share runs from zero at $\gamma = 3.33$ to 15.0 per cent at $\gamma = 4.07$, whereas Shao et al. (2019) find cover stable across preference values.

6.2.4 The timing of HEAS access

We solve the model under three rules of HEAS access. Under the retirement rule, the baseline, the household may draw from retirement. Under the high-need rule it may draw whenever a living member is in State 3 or 4, and under the widowhood rule it may draw from widowhood, which for a single household is retirement itself. The retirement rule gives the highest welfare index on most of the wealth grid, and each delayed rule raises it on less than a twentieth of the grid (Table S7). The HEAS drawing cap limits a yearly flow, and ϕ is a fraction of the drawing cap, so drawings forgone in years without access cannot be made up by drawing faster afterwards. With the release intensity chosen afresh under each rule, a delayed rule does better mainly where the household does not draw under the retirement rule (Panel C of Table S7), and by at most 0.143 per cent of the welfare index. Drawing is assumed to begin at retirement, and the ranking of the three rules supports this assumption as long as the release intensity is constant over time.[15]

6.3 Optimal choices across household cells, and the estate

We now report the optimal choices at the baseline across the 64 household cells rather than the nodes of the wealth grid. In the market with LTCI cover only, 18.2 per cent of households buy it. The buyers are partnered households in the fourth, fifth, sixth and eighth octiles of liquid wealth at retirement and single households in the fifth. Most of them lie in the upper half of home equity for their octile. When home equity release is also offered, only the partnered households of the eighth octile with the most home equity (2.4 per cent of households) keep their cover, and they draw nothing (Section 6.1.1). The other buyers draw on the HEAS instead and hold no cover (Section 6.1.3).

Drawing on the HEAS is concentrated in the first five octiles, where every household draws. For the market with release only, Table S8 reports medians among surviving households at retirement and twenty years later, averaged over each octile's four home-equity quartiles at their population weights. Results in that market closely approximate those in the market with both instruments, because the optimal choices differ only for the households that keep their cover.[16] Twenty years after retirement, the ratio of the HEAS loan balance to home equity is between 0.58 and 0.70 in the first five partnered and first four single octiles, and lower in the rest.

This high level of drawing is closely related to the model's portfolio allocation. At retirement, the risky share of liquid wealth reaches its upper bound, $\bar{\pi} = 1$, in every octile except the eighth partnered one. A drawing adds to liquid wealth, and when the risky share is at this bound the drawing also increases the amount invested in the risky asset. Liquid wealth accounts for 0.4 to 73 per cent of net worth, and the remainder is home equity, which is itself risky but almost uncorrelated with the traded risky asset (Tables S8 and 2). Because households cannot rebalance their home equity or borrow outside the HEAS, the risky share reaches its upper bound when the desired holding of the traded risky asset exceeds liquid wealth.[17]

[15] Holding the release intensity ϕ constant over time is a separate assumption that this paper does not test. A household that could vary its drawing over time could still draw a constant fraction of the cap, so the value of home equity release reported here is a lower bound.

[16] For the eighth partnered octile, the entries in Tables S8 and 4 differ by less than 0.05 per cent between the two markets.

[17] The Merton share, the fraction of net worth that an investor with the households' relative risk aversion would hold in the risky asset if free to rebalance all its wealth, is $(\mu_1 - r)/(\gamma\sigma_1^2) = 0.501$ (Section S13). In the first five partnered and first six single octiles, even investing all liquid wealth in the risky asset leaves the holding below this share of net worth, which explains the bound there. In the higher octiles the comparison does not explain it, and in both eighth octiles the bequest motive gives a separate reason for keeping the risky share at or near its upper bound.

The same pattern of drawing produces a gradient in the home equity left in the estate (Table 4). In the first five octiles the loan sometimes grows beyond home equity, and under the no-negative-equity guarantee the government bears the difference, which we call the guarantee shortfall. In the household's final year of life, the loan exceeds home equity on 28 to 36 per cent of partnered paths and 7 to 15 per cent of single paths. Among these octiles, this share is lowest in the fifth single octile, because its households draw less.[18] Households in the sixth to eighth octiles draw little or nothing and leave their home equity largely intact, so the guarantee is triggered on few of their paths.

The drawing and the risky investment reported in this section describe a benchmark favourable to drawing on the HEAS rather than a forecast of take-up. A drawing can finance consumption, but one that instead keeps more liquid wealth invested is worthwhile only if that wealth earns more than the HEAS interest rate. In the model the risk-free asset does not, so this motive for drawing depends on investment in the risky asset. In practice, limited financial knowledge (van Rooij et al., 2011) may reduce risky investment, while minimum withdrawals from superannuation pensions reduce the balances that households can leave invested. The model has neither feature, and both would weaken the incentive to draw. Nakajima and Telyukova (2017) also count the risk of moving into a nursing home among the reasons why few households in the United States take out reverse mortgages. Our model leaves out residential care, so it omits this reason as well.

Lowering the upper bound on the risky share shows how much of the drawing rests on risky investment. At $\bar{\pi} = 0.5$, the share of households that draw as much as the HEAS allows falls from 60.8 to 40.5 per cent, and demand for cover rises (Table 3). At the baseline, a household with all its liquid wealth in the risky asset would borrow at the risk-free rate to invest more if it could. An LTCI premium priced at that rate therefore carries a higher opportunity cost than under the bound of one half. Under that bound, home equity release still lowers cover for some households and raises it for none.[19] At 30 June 2026, only 20,264 participants had a HEAS loan, with an average balance of A$38,212 (Department of Social Services, 2026). Section 7 therefore also reports results for the market with LTCI cover only.

6.4 Why the demand for cover is low

At the baseline, households buy little or no LTCI cover, and home equity release lowers optimal cover on part of the wealth grid and raises it nowhere (Section 6.1.3). This section sets out why, first for release (Section 6.4.1) and then, in the three subsections after it, for the LTCI contract itself. They show how the means tests treat cover over a lifetime (Section 6.4.2), why cover has little to insure (Section 6.4.3) and how the realised clawback compares with the implicit tax in the literature (Section 6.4.4).

6.4.1 Release and the value of cover

Where both instruments are affordable, release supplies funds in every health state, including the care states, where the LTCI indemnity is paid. The value of cover is low to begin with. The indemnity is paid mostly in the high-need States 3 and 4, which households enter late or not at all, and in those states the care subsidy already pays most of the care cost (Section 6.4.3). Drawing in the care states lowers it further in most household cells.

We test this by measuring how drawing changes the value of LTCI cover. Because optimal cover is zero for most households at the baseline, we impose the same cover on every household, a fifth of full cover, the largest that every household cell outside the poorest single octile can afford. Section 6.4.2 measures the realised clawback at the same cover. The value of cover is the gain in the welfare index from that cover over none. Of the 60 household cells that can

[18] A smaller drawing reduces the probability that the loan will exceed home equity, but the realised outcome also depends on survival and house prices. Measured at the start of the final year of life, as in the text, the share over all paths is 16.7 per cent. Table S14 measures it one simulated year later, at the last death, as 18.5 per cent.

[19] With $\bar{\pi} = 0.5$, 47.8 per cent of households buy cover in the market with cover only.

Table 4. Expected present value of the estate at death and probability that the guarantee is triggered, by liquid-wealth octile.

Liquid-wealth octile	Liquid component ($X_{\tau-}$)[a] (A$)	Home-equity component[b] $[Y_{\tau-}(1-\tilde{\kappa}_{\tau-})]^+$ (A$)	Total estate (A$)	Guarantee triggered (%)
Panel A: Partnered households				
1	47,800	237,042	284,842	30.3
2	125,137	275,415	400,552	28.2
3	261,506	290,487	551,993	28.1
4	486,923	296,900	783,823	31.5
5	833,223	283,048	1,116,271	36.1
6	852,275	900,399	1,752,674	0.0
7	1,704,429	1,081,320	2,785,749	0.0
8	5,904,997	1,370,281	7,275,279	0.0
Panel B: Single households				
1	29,649	327,843	357,491	14.2
2	49,564	295,831	345,395	14.5
3	105,663	318,202	423,865	13.9
4	220,200	328,158	548,358	14.5
5	309,703	507,202	816,906	6.9
6	377,181	707,203	1,084,384	0.0
7	865,839	758,813	1,624,653	6.3
8	2,461,991	1,075,144	3,537,135	0.0

Notes: Population-weighted expectations over 10,000 simulated paths per cell under release only, discounted to age 67 at the risk-free rate from the start of the last year a member is alive. Every simulated path ends with the last death before age 117 (Section S4).
[a] Liquid component $X_{\tau-}$ is liquid wealth in the last year alive.
[b] The home-equity component is zero where the loan balance exceeds home equity, which triggers the no-negative-equity guarantee.

afford it, full drawing lowers that gain in 30, which lie in the middle of the distribution of liquid wealth at retirement (Section S16). Under the high-need rule of Section 6.2.4, which lets the household draw only in states in which cover pays, full drawing lowers the gain in 55 cells. Drawing therefore lowers the value of cover more widely when its funds arrive only in the states in which cover pays. Without the bequest motive, full drawing under the retirement rule lowers the gain in only 10 cells, in line with the finding of Section 6.2.2 that the displacement of LTCI cover by home equity release depends on the motive.

6.4.2 Realised clawback rates

In the middle of the wealth distribution the means tests charge LTCI cover at a point in time but subsidise it over a lifetime. This subsection measures the lifetime amount, in age pension withdrawn and home care fees raised along a household's own path, and asks what changes when the contract is exempt from the tests. Figure 3 draws the instantaneous clawback (Section 4.4) at a single point in time and at full cover, and its largest values arise when every living member is in State 4.[20] We instead measure the realised clawback of Section 4.4.2 over the whole lifecycle, at an imposed coverage and release intensity.

Table 5 reports the realised clawback as a rate on the fair value of cover at the imposed coverage of Section 6.4.1, a fifth of full cover ($\psi = 0.2$). A rate on the fair value of optimal cover, $\psi^* P$, would be undefined wherever ψ^* is zero. The first column is computed without home equity release ($\phi = 0$) and the second with full drawing ($\phi = 1$), rather than at each cell's own optimal release intensity ϕ^*. Consumption and the risky share remain the household's own choices. The table therefore describes the means tests facing an insured household, not the demand for cover.

Table 5 places the subsidy in the middle of the wealth distribution and a small charge at the top and in some of the lowest octiles. The households whose liquid wealth at retirement lies in the taper band are in the fourth and fifth partnered octiles and the fifth and sixth single

[20] These choices isolate the shape of the instantaneous clawback. Full cover is well above what any household buys at the baseline.

Table 5. Realised clawback rates on private insurance contracts by liquid-wealth octile.

Liquid-wealth octile	Without home equity release $RC^{\varnothing}/\psi P$ (%)	With release $RC/\psi P$ (%)
Panel A: Partnered households		
1	2.5	0.4
2	−1.6	−13.8
3	−19.3	−31.4
4	−40.7	−42.3
5	−40.6	−23.4
6	−18.1	1.1
7	−1.1	4.6
8	1.7	1.2
Panel B: Single households[a]		
2	1.9	2.2
3	1.2	−3.7
4	−10.6	−21.2
5	−31.5	−29.7
6	−24.4	−7.8
7	−2.6	5.2
8	3.3	1.8

Notes: At the baseline, each row is the population-weighted mean of the rates of that octile's four household cells, at an imposed coverage $\psi = 0.2$, as a percentage of the fair value of cover, ψP (Equation (3)); a negative rate is an implicit subsidy to the household. The columns are $RC^{\varnothing}$ (Equation (14)) and RC at full drawing, $\phi = 1$ (Equation (15)). Both are expected lifetime paired differences on identical simulated paths; Section S15 gives the change in the HEAS loan balance at death.

[a] Row 1 of the single panel is omitted: the poorest single octile cannot afford the premium at $\psi = 0.2$.

ones. The deepest implicit subsidy is in the fourth or fifth octile of each household type. With full drawing, the realised clawback rate is −42.3 per cent in the fourth partnered octile, and in the top two octiles of either household type it is a charge of at most 5.2 per cent.

The realised clawback has four components. The first is what paying the premium restores by lowering assessable wealth: a higher age pension and, because the fee falls as the pension rises, a lower home care fee, most of all in the taper band. The other three take pension back and raise the home care fee. They are the deemed asset, the indemnity counted as income in the care states, and the indemnity counted as assessable wealth while it is unspent. The deemed asset is the share $\nu = 0.6$ of the loaded price that the assets test counts while the contract is held. Paying the premium can at most give back the pension that the means tests would otherwise withdraw and the fee that they would otherwise add. What the deemed asset takes back has no such limit, because it grows with the cover held. Which side prevails depends on where assessable wealth stands on the taper.

Home equity release changes the realised clawback rate because drawings not yet spent count as assessable wealth and lower the pension through the taper. How much the premium restores and how much the indemnity removes both depend on where assessable wealth stands against the taper's thresholds, so release changes both. Release therefore changes the level of the rate in Table 5 and, in four octiles, its sign.

A negative realised clawback means that what the premium restores, in age pension and home care fees, exceeds what the means tests take back through the contract and the indemnity. It does not mean that assessing the contract subsidises cover. Exempting the contract (Section 6.2.1) changes only how it is treated. Liquid wealth remains assessable, so a rise in it still lowers the pension, which in turn changes the home care fee. Setting $\nu = 0$ removes the deemed asset and the indemnity counted as income, the two components that assess the contract itself. Without the deemed asset, the whole premium leaves assessable wealth, so paying it restores more pension and lowers the fee further.

At the imposed coverage and full drawing of Table 5, exempting the contract makes the implicit subsidy about two to six times its baseline size in the taper-band octiles.[21] At the

[21] The two ends of the range are the fourth partnered octile, where the rate moves from −42.3 to −83.8 per cent, and the sixth single octile, from −7.8 to −45.9 per cent (Section S14).

baseline, most of the buyers in the market with cover only lie in these octiles, and none of the buyers there keeps cover when release is also offered (Section 6.3). With the contract exempt, cover is chosen in each of these octiles when both instruments are offered.[22] The assessment of the contract under the means tests therefore holds demand down through what it takes back, even where the realised clawback is an implicit subsidy.

6.4.3 What the contract insures

The only risk that LTCI cover removes for the household is its home care fee, because the care subsidy already pays most of the care cost. Cover is therefore worth little. In every octile in which home equity release has value, release raises the welfare index by more than LTCI cover does, and its gain is largest where liquid wealth is lowest (Table S9).

The care cost net of informal care, LTC_h of Equation (2), averages A$140,884 per household in present value at retirement, in the market with both instruments, at each household cell's optimum. The government pays 77 per cent of that cost as the care subsidy (Table S11), and the household pays the remaining 23 per cent as its home care fee. Before either instrument is held, the care subsidy pays 79 per cent of the care cost (Section S16). The care subsidy has two parts. The base subsidy is the care cost less the fee that the means tests would charge without the lifetime cap, and the post-cap subsidy is the fee that the government pays once a member's fees reach that cap. The LTCI contract indemnifies LTC_h, not the fee, so full cover would pay the fee and a further 77 per cent of the care cost as cash in the care states. The means tests count that whole indemnity as income, scaled by the means-test intensity ν, although the household bears only the fee. In the market with cover only, the households that buy the most cover receive an expected indemnity of about two fifths to a half of their expected fee.

6.4.4 The clawback in the literature

The realised clawback is much smaller than the implicit tax that Brown and Finkelstein (2008) measure under Medicaid, and it is made up differently. Their tax is the share of a private contract's benefits that replace what Medicaid would otherwise have paid, about 60 to 75 per cent at median wealth. It is large because Medicaid pays only what private insurance has not paid. No Australian rule cuts the care subsidy by what a contract pays, so full cover would pay out in cash the part of the care cost that the subsidy already meets (Section 6.4.3). The only part of the realised clawback that works like Medicaid's tax is the indemnity counted as income, and it is small. For the one household cell that insures at the baseline with release offered, it takes back 0.05 per cent of the fair value of cover (Section S14). At a premium rebate of $\eta = 0.30$, where a quarter of households insure, their realised clawback is an implicit subsidy in the sixth and seventh octiles and a small charge in the top partnered octile (Section S17.2).

The subsidy side of the realised clawback comes from the Australian rules for lifetime income streams, not from means-tested care in general. It needs an assets test with a taper, so that lower assessable wealth restores pension. It also needs a rule that counts only part of the purchase price, so that paying the premium takes the rest out of assessment. The same purchase brings no such gain under a safety net that pays only once the household has spent down its wealth (Brown and Finkelstein, 2008), or under England's current means test for care (Hu et al., 2026). Neither system has a counterpart to the rising limb of the instantaneous clawback in Figure 3, along which the charge grows with liquid wealth.

7 LTCI cover, the premium rebate and the government's net expenditure

For the government, LTCI cover at the baseline costs little, and a premium rebate that raises the insured share does not pay for itself through the means tests. Where home equity release is not offered, the rebate raises the far tail of net expenditure by less than it raises the mean. Where

[22] The exemption removes both assessed components at once, so the change in demand is their joint effect.

release is also offered, the rebate reaches fewer and wealthier households, and from $\eta = 0.24$ the net fiscal cost per insured household is higher and the far tail rises about as much as the mean or more.

7.1 Net expenditure and the cost of cover at the baseline

As in Lyu et al. (2026), we follow the government's expenditure along simulated household paths. We compare the same households with and without LTCI cover on identical histories of health and asset returns. Throughout this section the LTCI contract is assessed under the means tests, as at the baseline. The government's net expenditure on a household is the discounted sum of the age pension, care subsidy, premium rebate and guarantee shortfall that it pays for that household, less the spread it earns on the HEAS loan (Equation (S10)). The rebate's net fiscal cost at a given level is the rise in the government's mean net expenditure from the LTCI cover that households choose at that level, against the same households uninsured in the same market. We measure the far tail of net expenditure by the expected shortfall beyond the ninety-ninth percentile, ES_{99}, the average of the costliest one per cent of outcomes among households of the same type and wealth at retirement.[23]

Through the age pension, the care subsidy and the no-negative-equity guarantee, the government is the residual insurer of each household's costliest outcomes, which come from long lives and long spells of high need. A private LTCI contract that pays its indemnity in those spells takes part of that role. In the model, health and survival risks are idiosyncratic, so they diversify across households, and the far tail is not a risk to the budget. Its change with LTCI cover shows whether a private contract relieves the government, as the residual insurer, in a household's costliest outcomes.

At the baseline rebate, where release is not offered, LTCI cover adds A$889 per household to mean net expenditure, of which A$251 is the rebate paid. It adds less in the far tail than over the body of the distribution (Table S10). Home equity release, by comparison, lowers mean net expenditure by A$113,506 per household at the model's HEAS interest rate.[24]

7.2 The net fiscal cost of the premium rebate

A premium rebate raises the insured share of households but does not pay for itself. At $\eta = 0.30$, about twice the baseline rebate, a quarter of households insure where home equity release is offered and about three in five where it is not (Table 6).[25] On average, the means tests do not recover the rebate paid at any level up to $\eta = 0.30$, in either market, and at $\eta = 0.30$ they add to what the government pays beyond the rebate itself (Table 6 and Section 6.4.2).

Where release is also offered, fewer households insure, but at $\eta = 0.30$ the net fiscal cost per insured household is about twice as large, A$19,758 against A$9,320. Two things raise it: release changes who insures, and one group of insured households stops drawing on the HEAS. The households in the middle of the wealth distribution that insure only where release is not offered

[23] Net expenditure leaves out the hardship waiver, which the government also bears. Where home equity release is not offered, LTCI cover at the baseline rebate raises ES_{99} by A$361 per household, and with the waiver counted it lowers ES_{99} by A$241 (Section S17.1). At $\eta = 0.30$, counting the waiver lowers the rise in ES_{99} from A$1,971 to A$91 where release is not offered, or from 0.34 to 0.02 per dollar of net fiscal cost. Where release is offered, it lowers the rise from A$4,995 to A$4,829, or from 1.00 to 0.97. With all households pooled rather than grouped into household cells, the rebate $\eta = 0.30$ raises ES_{99} by A$3,468 in the market with both instruments and by A$1,082 in the market with cover only, against the rises in Table 6.

[24] About two thirds of that saving is the government's net income from the HEAS loan, the spread it earns at the model's HEAS interest rate of 5.41 per cent less the guarantee shortfall. In the model the loan accrues at that rate for the household and the government alike, and the government's accounts in this section and in Sections S17 to S19 are computed only at it. The model's rate includes the margin that households plan on (Section 5). At the rate without the margin the spread is 0.96 rather than 2.91 percentage points, so the net income from the loan, and with it the saving, would be smaller. The current statutory HEAS interest rate is below the risk-free rate in real terms, so at that rate the net income from the loan would be negative.

[25] The rebate $\eta = 0.30$ lies between the highest (base-tier) rates of Australia's private health insurance rebate for people aged 65 to 69 and for those aged 70 and over (Section S17.2).

draw on the HEAS instead. The insured then include a larger share of wealthy households, and two household cells in the eighth partnered octile take full LTCI cover. The rebate is a share of the premium, so full cover attracts more of it, and the means tests take back little from households this wealthy. Those two cells receive four fifths of the rebate paid. A third cell, the partnered households of the sixth octile with the most home equity, stops drawing on the HEAS when it insures. The government then loses the net income from its HEAS loan and pays it more age pension and care subsidy (Section S17.2).

Table 6. Insured share, and fiscal and tail effects against the same households uninsured, at $\eta = 0.30$ with and without home equity release.

Per household unless labelled (expected present value at age 67, A$)	Both instruments	Cover only
Panel A: Demand		
Insured share of households (%)	25.2	62.0
Mean cover ψ^*	0.058	0.083
Panel B: Fiscal effects		
LTCI premium rebate	3,116	4,379
Age pension	693	1,392
Care subsidy, base	56	−44
Care subsidy, post-cap	−7	48
Net income from the loan, forgone	1,124	0
Net fiscal cost	4,982	5,775
Net fiscal cost per insured household	19,758	9,320
Realised clawback, RC ($RC^{\varnothing}$ with cover only)	−121	−1,348
Panel C: Expected shortfall of the government's net expenditure		
Change in ES_{95} (beyond the ninety-fifth percentile)	4,389	3,162
Change in ES_{99}	4,995	1,971
Rise in ES_{99} per A$1 of net fiscal cost	1.00	0.34

Notes: Population-weighted means per household over the 64 cells, 200,000 paths per cell; mean cover ψ^* counts the uninsured at zero, spending is positive, RC is positive for a charge, and expected shortfall, ES, is computed within cells. Money rows other than RC compare each household insured and uninsured on identical histories, uninsured at its optimal release intensity (release only in the first column, neither instrument in the second). RC keeps the insured household's release intensity in both runs (Equations (15) and (14)), so with cover only it equals minus the sum of the age pension and base subsidy rows. In the first column one cell stops drawing when it insures, so those two rows, not RC, include its extra pension and base subsidy.

7.3 The premium rebate and the far tail

Where home equity release is not offered, the rebate raises the far tail of the government's net expenditure by less than its net fiscal cost at every rebate level up to $\eta = 0.30$ (Figure S6(b)). The far tail therefore rises less than the mean, and at $\eta = 0.30$ the rise is about a third of the net fiscal cost. In the long care spells that make up much of the far tail, the indemnity counted as income lowers the household's age pension and raises its home care fee. The higher fee lowers the care subsidy until a member's fees reach the lifetime cap, so the rise in what the government pays is smaller on those histories (Section S17.1).

Where release is also offered, the rise in the far tail is smaller than the net fiscal cost at lower rebates, but from $\eta = 0.24$ it is about as large or larger. The change comes mainly from the sixth-octile household cell of Section 7.2, which stops drawing from $\eta = 0.24$. The net income from its loan that the government then loses is largest in the longest lives, which make up much of the far tail. For most households, who in practice do not draw on the HEAS, the results where release is not offered are the closer guide (Section 6.3).

8 Conclusion

This paper asks whether a private long-term care insurance (LTCI) contract could complement Australia's public support for home care. Within the model and its calibration, it would do so only to a limited extent under the current means tests. Even at an actuarially fair premium, only 2.4 per cent of households would buy the LTCI contract where home equity release is also offered, and 18.2 per cent where it is not. The contract therefore has a role where households do not draw on the HEAS, as most do not in practice. The first reason for the low demand is that the means tests count part of the contract's purchase price as an asset and part of its indemnity as income. Exempting the contract from this assessment raises the insured share of households to nearly half where release is offered (Section 6.2.1). The other two reasons are the care subsidy, which already pays most of the care cost, and the bequest motive (Sections 6.4.3 and 6.2.2).

At the baseline, offering home equity release lowers the insured share of households by about sixteen percentage points, because the home is kept during care and drawing therefore provides funds in the same care states in which LTCI cover pays (Section 6.4.1). With the contract exempt from the means tests, release barely lowers that share, and without the bequest motive it raises that share (Sections 6.2.1 and 6.2.2). The size of the fall depends on how households invest, because drawing earns more than the HEAS loan costs only when the liquid wealth it frees is held in the risky asset. With at most half of liquid wealth there, release still lowers the insured share, by less than at the baseline (Section 6.3). Release also lowers the government's mean net expenditure on a household by more than LTCI cover changes it, at the model's HEAS interest rate (Section 7.1).

A premium rebate of 30 per cent of the loaded price would raise the insured share to a quarter of households where release is offered, and to about three in five where it is not. At that rebate the means tests would on average add to what the government pays beyond the rebate itself, whether or not release is offered (Section 7.2). The rebate would also raise the government's net expenditure on a household in the costliest outcomes, measured by expected shortfall. Where release is not offered, this rise is about a third of the rebate's net fiscal cost, the rise in mean net expenditure. Where release is offered, the two are about equal (Section 7.3).

The model leaves out the insurers' side of the LTCI market, such as private information and administrative costs (Finkelstein and McGarry, 2006; Braun et al., 2019). The home is neither sold nor traded down, residential care is outside the model, and households choose cover and drawing once, at retirement. The contract pays a share of the whole care cost, and a contract on the household's own means-tested home care fee is not studied. The signs and magnitudes found here should not be assumed to hold in another system. For example, the means tests subsidise LTCI cover over a lifetime in the middle of the wealth distribution, a result that rests on how the Australian assets test treats the purchase price (Section 6.4.4). The two measures behind these results apply to any means-tested system that assesses a private contract: the realised clawback along the household's own path, and the government's net expenditure read as a distribution rather than as a mean.

Appendix A Numerical solution method

This appendix describes how the HJB system of Equation (11) is solved. The discrete operators, the policy-iteration algorithm with its control updates, and the forward simulation are in Sections S1, S2 and S4.

A.1 Discretisation

Liquid wealth and home equity diffuse on the solved grid, which is uniform in logarithmic coordinates $(\mathscr{X}, \mathscr{Y}) = (\ln x, \ln y)$ and so suits the multiplicative shocks of the financial returns. One letter is local to this appendix: a under a state subscript is a drift, not an age. Discrete health states are placed alongside the grid in $(\mathscr{X}, \mathscr{Y})$. The path-dependent states, namely the

loan-to-value ratio $\tilde{\kappa}$ and the two fee accumulators, lie on a tensor-product Clenshaw–Curtis grid.[26] The reported solutions use 66 nodes in $\mathscr{X}$, spanning liquid wealth from A\$100 to A\$44.1 million, and 31 in $\mathscr{Y}$, spanning home equity from A\$10,000 to A\$12 million.[27] They use the twenty-four joint health states in which at least one member is alive, and 501 time levels over the fifty years from age 67 to 117, so $\Delta t = 0.1$. The path-dependent states take 125 nodes at Clenshaw–Curtis levels $(3, 3, 3)$, five per axis, and 25 nodes when the household never draws under the HEAS.

The solver carries the HEAS loan balance as the loan-to-value ratio $\tilde{\kappa} = \kappa/Y$, although Equation (11) is written in κ. By Itô's lemma,

$$\mathrm{d}\tilde{\kappa} = \Big[\tilde{\kappa}(r^{\kappa} - \mu_2 + \sigma_2^2) + \frac{I^{\mathrm{HEAS}}}{Y}\Big]\mathrm{d}t \;-\; \tilde{\kappa}\,\sigma_2\,\mathrm{d}W^{Y}, \tag{17}$$

where $W^{Y} = \rho W_1 + \sqrt{1-\rho^2}\,W_2$ is the Brownian motion driving home equity in Equation (5). The change of variable is exact, since $\tilde{\kappa}_t Y_t$ recovers the balance path by path. It puts the no-negative-equity guarantee at the fixed boundary $\tilde{\kappa} = 1$, so the range the household's problem needs is $[0, 1]$ at every level of home equity. Section S1 gives the numerical reasons for the choice.

At the top of the $\mathscr{X}$ grid the value function takes its far-field form. Far above every threshold the age pension is zero, fees and care costs are capped, and consumption, the risky position and the bequest scale with wealth, so $V \approx A\,(x + \Xi)^{1-\gamma}$. Here $\Xi = \bar{\xi}_2 + (y(1-\tilde{\kappa}))^{+}$ is the part of the estate that does not scale with x: the bequest shift and home equity net of the loan. The top row of the operator uses a ghost node on that profile, $V_N = \big((x_N + \Xi)/(x_{N-1} + \Xi)\big)^{1-\gamma} V_{N-1}$ over one grid step. The implicit matrix stays an M-matrix, and the controls at the top column are computed from the same ghost node.

A reflecting condition, $V_{\mathscr{X}} = 0$ through a mirror node, would set the marginal value of wealth to zero at the ceiling, so a premium paid there would cost nothing. We check the far-field condition against a reference solution with the liquid-wealth ceiling at A\$218 million, eight columns beyond the top of the solved grid. The reference solution is computed at thirteen values of LTCI cover in the market with cover only. The grid step is held fixed, so the grids share their nodes wherever they overlap. On the wealth grid the value function at retirement differs from the reference by at most 0.012 per cent up to A\$13.3 million, in every health state. The cover-only optimum over those values is the same at every node from A\$0.3 million to A\$13.3 million for both household types. A reflecting condition on a grid cut at A\$13.3 million instead differs by 2 to 7 per cent at A\$4 million, and its cover-only optimum for partnered households reaches full cover in its own top two columns. That region of full cover is the localisation error of a truncated state space (Barles et al., 1995).

The bottom of the liquid-wealth axis and the top of the home-equity axis assign both parts of the convection term, which the upwind discretisation below splits into two weights, to their interior neighbour, whichever way the drift points. The condition is a zero normal derivative, $V_{\mathscr{X}} = 0$ at the bottom of the liquid-wealth axis, where $V_{\mathscr{X}} = xV_x$ is small because transfers keep V_x bounded there while x is only A\$100, and $V_{\mathscr{Y}} = 0$ at the top of the home-equity axis. On the bottom home-equity row, home equity is absorbing and no home-equity term remains (Table S1). On the three axes of the path-dependent states, a node with no upwind neighbour takes a zero derivative, the same condition at the inflow boundary. At the horizon the value is the bequest payoff of the estate, $V(T, \cdot) = U_2\big(x + (y(1-\tilde{\kappa}))^{+}\big)$.

The wealth grid of Section 4.3 ends six columns and four rows below the top of the solved grid; those columns and rows are a numerical guard region. The four rows above it run to the

[26] We do not use a sparse Smolyak construction, because the upwind stencil on these states is well posed only where every node has an axis-aligned upwind neighbour, which a Smolyak grid does not guarantee.

[27] A choice that leaves less than A\$100 of liquid wealth after the premium is valued at that floor, and none of the reported optima is such a choice.

top of the home-equity axis, where the zero normal derivative is least accurate. A reference solution with the home-equity ceiling at A$526.5 million, sixteen rows further out on the same step, is solved with release intensity 0 or 1. Against it the value function differs by up to 58 per cent at A$12 million and 8 per cent at A$5.9 million, both in those four rows. On the wealth grid it differs by at most 4 per cent, in its top row, at A$4.66 million. With release intensity restricted as in the reference, the optimum differs at ten nodes of that row, eight of them by one step of 0.02 in cover, and at no household cell. Population-weighted, 1.4 per cent of simulated household-years are spent with home equity above A$4.66 million and 0.6 per cent above A$5.9 million. Within a single cell, the shares are at most 22 and 11 per cent, both in the only cell that insures where home equity release is also offered, the partnered households of the eighth octile with the most home equity (Section 6.1.1). Against both reference solutions its welfare index moves by nearly the same amount at every value of cover, so its gain from cover moves by at most A$0.02 and its choice stays at 4 per cent of full cover.

A.2 Four numerical techniques

We combine four techniques to keep the procedure tractable and stable.

- *Policy iteration.* Policy iteration converts the nonlinear HJB into a sequence of linear problems. We alternate between (a) improving the controls and (b) advancing the value function one implicit step under the current policy, and repeat the pair until the value function converges at that time level. Convergence is tested on the value function alone, not on the controls, against a time-dependent tolerance. We improve consumption by inverting its first-order condition and projecting the result onto the admissible interval, and the risky share by a direct search over a discrete approximation to the Hamiltonian (Section S2.1).
- *Craig–Sneyd ADI in log-space.* Under each policy we solve the two-dimensional diffusion with the Craig–Sneyd ADI scheme (Craig and Sneyd, 1988) in log coordinates $(\mathscr{X}, \mathscr{Y})$. The scheme keeps second-order accuracy for the pure diffusion terms and handles the mixed derivative through a Craig–Sneyd correction.
- *Upwind discretisations for stability.* First-order transport components need monotone discretisations or the scheme can lose stability. In $(\mathscr{X}, \mathscr{Y})$ we split the convection with a smoothed absolute value of the drift $\tilde{a}$ in those coordinates, $|\tilde{a}|_{\epsilon_{\text{up}}} = \sqrt{\tilde{a}^2 + \epsilon_{\text{up}}^2}$, where ϵ_{up} is a small constant. The split keeps both off-diagonals non-negative and lets the stencil vary continuously with the drift. The standard upwind split (Hanson, 2007, pp. 232–233), by contrast, switches with the drift's sign. The path-dependent states take monotone tensor-grid upwind operators, which do switch, since their drifts are evaluated one node at a time.
- *Vanishing-viscosity regularisation.* Means-tested transfers introduce kinks where a text-book policy-improvement step can stall or oscillate. We add $\varsigma(V_{\mathscr{X}\mathscr{X}} + V_{\mathscr{Y}\mathscr{Y}})$ as a vanishing-viscosity term (Yong and Zhou, 1999, p. 197), with $\varsigma = \mathcal{O}(\Delta\mathscr{X}^2 + \Delta\mathscr{Y}^2)$ in the coordinates the operators are assembled in, so the perturbation disappears as the solved grid refines and the scheme stays consistent (Barles and Souganidis, 1991).

A.3 Choices required by the means tests and the HEAS

Four choices go beyond the standard techniques above. The liquidation constraint of Section 3 forces the first, two features of the means tests call for the second and third, and the rules of the HEAS call for the fourth. One feature is that a transfer replacing private wealth leaves the value function nearly flat in the logarithm of liquid wealth at the bottom of the solved grid. The poorest households are located there, and the optimal choices are evaluated on the wealth grid, which starts there too. The other feature is that the means-test rules put kinks inside the state space, some of which move with age. The diagnostics that isolate the first two choices are in Section S3.

First, the liquidation horizon τ_c is a parameter of the model and is not tied to the time step. It is set at one year (Table 2). The consumption bound $c^{\max} = X/\tau_c + I - F$ is needed because without it the optimisation admits consumption rates that are unbounded near the borrowing limit, a singular arc whose discrete approximation depends on the time step rather than the model. Setting τ_c equal to Δt would make the bound itself grow like $1/\Delta t$. With $\tau_c = \Delta t$, one backward step moved the bottom of the wealth axis by more than a grid interval, a Courant number above one. The implicit scheme remained stable, but it understated the marginal value of a dollar at the bottom of the axis.

Second, relaxing the consumption update at 0.3 removes a two-cycle, and damping the risky share limits the size of its cycle without removing it. The two-cycle comes from the consumption first-order condition, which raises $V_{\mathscr{X}}$ to the power $-1/\gamma$, so the gradient feeds into the drift and the drift back into the gradient. The risky share is chosen by a search over a discrete set of candidates, so the cycle that remains is an alternation between two neighbouring candidates.

Third, the hardship waiver keeps the drift bounded. In log coordinates a fixed-dollar outflow gives the drift $a_{\mathscr{X}} = -\text{outflow}/X$, which grows without limit as liquid wealth falls. We model the waiver by cutting the home care fee until $c^{\max}$ reaches $c^{\min}$, so the consumption interval is never empty and the net outflow never exceeds X/τ_c. The drift then cannot fall below $r + \pi(\mu_1 - r) - 1/\tau_c$, and the Courant number this lower bound gives on the solved grid is 0.49.

Fourth, the loan-to-value ratio $\tilde{\kappa}$, a path-dependent state, is discretised over the range a household can reach. The Maximum Loan Amount caps the drawing, not the balance. The Maximum Loan Amount reaches 0.675 of home equity at about age 90. After that the balance keeps accruing at the drift of Equation (17), and at the last death the ratio is above 1 on about a fifth of paths (Table S14). Treating the Maximum Loan Amount as a bound on the balance would therefore cut off a third of the range $[0, 1]$. The $\tilde{\kappa}$ axis instead ends at $\tilde{\kappa} = 1$, where the no-negative-equity guarantee takes effect and above which the estate no longer depends on $\tilde{\kappa}$, so the zero derivative imposed there is an approximation.

Financial support

This work was supported by the Centre for Population Ageing Research (CEPAR) at UNSW Sydney through a grant from the Department of Health, Disability and Ageing (DHDA).

Competing interests

The authors declare none.

Acknowledgements

This paper uses unit record data from the Household, Income and Labour Dynamics in Australia (HILDA) Survey. The unit record data from the HILDA Survey was obtained from the Australian Data Archive, which is hosted by The Australian National University. The HILDA Survey was initiated and is funded by the Australian Government Department of Social Services (DSS) and is managed by the Melbourne Institute of Applied Economic and Social Research (Melbourne Institute). The findings and views based on the data, however, are those of the authors and should not be attributed to the Australian Government, DSS, the Melbourne Institute, the Australian Data Archive or The Australian National University and none of those entities bear any responsibility for the analysis or interpretation of the unit record data from the HILDA Survey provided by the authors. Computations were performed on the UNSW Compute Cluster Katana (DOI: 10.26190/669X-A286). The authors thank Dr Bei Lu for help with proofreading the manuscript.

References

Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2022). Income and wealth distribution in macroeconomics: A continuous-time approach. *Review of Economic Studies*, 89:45–86.

Achou, B. (2021). Housing liquidity and long-term care insurance demand: A quantitative evaluation. *Journal of Public Economics*, 194:104353.

AGA (2024). *Australian Life Tables 2020–22*. Australian Government, Canberra.

AIHW (2025). *GEN data: People Using Aged Care CURF*. Australian Government, Canberra.

Ameriks, J., Briggs, J., Caplin, A., Shapiro, M. D., and Tonetti, C. (2020). Long-Term-Care Utility and Late-in-Life Saving. *Journal of Political Economy*, 128(6):2375–2451.

Anderson, E. W., Hansen, L. P., and Sargent, T. J. (2003). A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association*, 1(1):68–123.

Andréasson, J. G. and Shevchenko, P. V. (2024). Optimal annuitisation, housing and reverse mortgage in retirement in the presence of a means-tested public pension. *European Actuarial Journal*, 14(3):871–904.

Andréasson, J. G., Shevchenko, P. V., and Novikov, A. (2017). Optimal consumption, investment and housing with means-tested public pension in retirement. *Insurance: Mathematics and Economics*, 75:32–47.

Barles, G., Daher, C., and Romano, M. (1995). Convergence of numerical schemes for parabolic equations arising in finance theory. *Mathematical Models and Methods in Applied Sciences*, 5(1):125–143.

Barles, G. and Souganidis, P. E. (1991). Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic Analysis*, 4(3):271–283.

Braun, R. A., Kopecky, K. A., and Koreshkova, T. (2019). Old, frail, and uninsured: Accounting for features of the US long-term care insurance market. *Econometrica*, 87(3):981–1019.

Brown, J. R. and Finkelstein, A. (2007). Why is the market for long-term care insurance so small? *Journal of Public Economics*, 91(10):1967–1991.

Brown, J. R. and Finkelstein, A. (2008). The interaction of public and private insurance: Medicaid and the long-term care insurance market. *American Economic Review*, 98(3):1083–1102.

Catherine, S. (2022). Countercyclical labor income risk and portfolio choices over the life cycle. *The Review of Financial Studies*, 35(9):4016–4054.

Cocco, J. F. and Lopes, P. (2020). Aging in place, housing maintenance, and reverse mortgages. *Review of Economic Studies*, 87(4):1799–1836.

Costa-Font, J. (2010). Family ties and the crowding out of long-term care insurance. *Oxford Review of Economic Policy*, 26(4):691–712.

Craig, I. J. and Sneyd, A. D. (1988). An alternating-direction implicit scheme for parabolic equations with mixed derivatives. *Computers & Mathematics with Applications*, 16(4):341–350.

Dagum, C. (1977). A new model of personal income distribution: specification and estimation. *Economie Appliquée*, 30(3):413–437.

Davidoff, T. (2010). Home equity commitment and long-term care insurance demand. *Journal of Public Economics*, 94(1-2):44–49.

De Nardi, M., French, E., Jones, J. B., and McGee, R. (2025). Why do couples and singles save during retirement? Household heterogeneity and its aggregate implications. *Journal of Political Economy*, 133(3):750–792.

Department of Social Services (2026). Home Equity Access Scheme – Quarterly Report. Dataset, data as at 30 June 2026. Available at: https://data.gov.au/data/dataset/home-equity-access-scheme-quarterly-report (accessed 24 September 2026).

DHDA (2025a). *Aged Care Subsidies and Supplements: Support at Home*. Australian Government, Canberra.

DHDA (2025b). *Support at Home Program Handbook*. Australian Government, Canberra.

DHDA (2025c). *Support at Home Program: Pooled Care Management Funding and Service Delivery Branches*. Australian Government, Canberra.

Finkelstein, A., Luttmer, E. F., and Notowidigdo, M. J. (2013). What good is wealth without health? The effect of health on the marginal utility of consumption. *Journal of the European Economic Association*, 11(s1):221–258.

Finkelstein, A. and McGarry, K. (2006). Multiple dimensions of private information: Evidence from the long-term care insurance market. *American Economic Review*, 96(4):938–958.

Fong, J. H., Shao, A. W., and Sherris, M. (2015). Multistate actuarial models of functional disability. *North American Actuarial Journal*, 19(1):41–59.

Fox, R. and Tulip, P. (2014). Is Housing Overvalued? Research Discussion Paper 2014-06, Reserve Bank of Australia.

Hanson, F. B. (2007). *Applied Stochastic Processes and Control for Jump-Diffusions: Modeling, Analysis and Computation*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA.

Hu, B., Hancock, R., Wittenberg, R., King, D., and Morciano, M. (2026). Reforming the funding of long-term care for older people: costs and distributional impacts of planned changes in England. *Health Economics, Policy and Law*, 21(3):285–305.

Hubbard, R. G., Skinner, J., and Zeldes, S. P. (1995). Precautionary saving and social insurance. *Journal of Political Economy*, 103(2):360–399.

Jackson, C. H., Sharples, L. D., Thompson, S. G., Duffy, S. W., and Couto, E. (2003). Multistate Markov models for disease progression with classification error. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 52(2):193–209.

Ji, M., Hardy, M. R., and Li, J. S.-H. (2011). Markovian approaches to joint-life mortality. *North American Actuarial Journal*, 15(3):357–376.

Jordà, Ò., Knoll, K., Kuvshinov, D., Schularick, M., and Taylor, A. M. (2019). The rate of return on everything, 1870–2015. *The Quarterly Journal of Economics*, 134(3):1225–1298.

Knoll, K., Schularick, M., and Steger, T. (2017). No price like home: Global house prices, 1870–2012. *American Economic Review*, 107(2):331–353.

Lockwood, L. M. (2018). Incidental bequests and the choice to self-insure late-life risks. *American Economic Review*, 108(9):2513–2550.

Lyu, L., Shen, Y., Sherris, M., and Ziveyi, J. (2026). Financing aged care with home equity allowing for government age pension and aged care support. *Insurance: Mathematics and Economics*, 126:103193.

Maenhout, P. J. (2004). Robust portfolio rules and asset pricing. *Review of Financial Studies*, 17(4):951–983.

McCullagh, P. (1980). Regression models for ordinal data. *Journal of the Royal Statistical Society: Series B (Methodological)*, 42(2):109–127.

Minister for Social Services (2026). Social Security (Home Equity Access Scheme – Rate of Compound Interest) Determination 2026. Available at: https://www.legislation.gov.au/F2026L00945 (accessed 12 October 2026).

Mommaerts, C. (2025). Long-term care insurance and the family. *Journal of Political Economy*, 133(1):1–52.

Munk, C. (2000). Optimal consumption/investment policies with undiversifiable income risk and liquidity constraints. *Journal of Economic Dynamics and Control*, 24(9):1315–1343.

Nakajima, M. and Telyukova, I. A. (2017). Reverse mortgage loans: A quantitative analysis. *The Journal of Finance*, 72(2):911–950.

Nektarios, M., Valvis, Z., and Georgiadis, T. (2025). A methodology for estimating the demand for long-term care in hours: The baseline scenario for Greece. *Journal of Long-Term Care*, pages 194–207.

OECD (2023). *Health at a Glance 2023: OECD Indicators*. OECD Publishing, Paris.

OECD (2024). *Is Care Affordable for Older People?* OECD Publishing, Paris.

Park, K. and Sherris, M. (2026). Design and pricing of private long-term care insurance: An Australian analysis. *Journal of Risk and Insurance*.

Productivity Commission (2025). *Report on Government Services 2025: Aged Care Services*. Australian Government, Canberra.

Reserve Bank of Australia (2026). Statistical Tables F2 and F2.1: Capital Market Yields – Government Bonds. Available at: https://www.rba.gov.au/statistics/tables/ (accessed 24 September 2026).

Rosenblatt, M. (1952). Remarks on a multivariate transformation. *The Annals of Mathematical Statistics*, 23(3):470–472.

Sefton, J., van de Ven, J., and Weale, M. (2008). Means testing retirement benefits: Fostering equity or discouraging savings? *The Economic Journal*, 118(528):556–590.

Services Australia (2026a). Home Equity Access Scheme. Available at: https://www.servicesaustralia.gov.au/home-equity-access-scheme (accessed 13 January 2026).

Services Australia (2026b). Income streams. Available at: https://www.servicesaustralia.gov.au/income-streams?context=22526 (accessed 23 September 2026).

Shao, A. W., Chen, H., and Sherris, M. (2019). To borrow or insure? Long term care costs and the impact of housing. *Insurance: Mathematics and Economics*, 85:15–34.

Stapledon, N. D. (2007). *Long term housing prices in Australia and some economic perspectives*. PhD thesis, University of New South Wales, Sydney.

Summerfield, M., Garrard, B., Nesa, M. K., Kamath, R., Macalalad, N., Watson, N., Wilkins, R., and Wooden, M. (2024). HILDA user manual – Release 23. Technical report, Melbourne Institute: Applied Economic and Social Research, University of Melbourne.

van Rooij, M. C. J., Lusardi, A., and Alessie, R. (2011). Financial literacy and stock market participation. *Journal of Financial Economics*, 101(2):449–472.

Ventura-Marco, M., Vidal-Meliá, C., and Pérez-Salamero González, J. M. (2023). Joint life care annuities to help retired couples to finance the cost of long-term care. *Insurance: Mathematics and Economics*, 113:122–139.

Watson, N. and Wooden, M. P. (2012). The HILDA survey: A case study in the design and development of a successful household panel survey. *Longitudinal and Life Course Studies*, 3(3):369–381.

Wood, G. A., Ong, R., and Haffner, M. E. (2022). Housing wealth and aged care: Asset-based welfare in practice in three OECD countries. *Housing Studies*, 37(4):511–536.

World Bank (2026). Inflation, consumer prices (annual %), indicator FP.CPI.TOTL.ZG. Available at: https://data.worldbank.org/indicator/FP.CPI.TOTL.ZG (accessed 12 January 2026).

Xu, M., Alonso-García, J., Sherris, M., and Shao, A. W. (2023). Insuring longevity risk and long-term care: Bequest, housing and liquidity. *Insurance: Mathematics and Economics*, 111:121–141.

Yong, J. and Zhou, X. Y. (1999). *Stochastic Controls: Hamiltonian Systems and HJB Equations*, volume 43 of *Applications of Mathematics*. Springer-Verlag, New York, NY.