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

Population ageing is driving up long-term care (LTC) costs and straining public budgets. Yet private LTC insurance (LTCI) markets remain small, and retirees' substantial home equity is largely untapped as a financing source. We develop a continuous-time life-cycle framework that writes the cross-dependencies among nested means-tested programmes (age pension, home care, and home equity release) directly into the household optimisation, with health states aligned ex ante to eligibility thresholds via a proportional-odds model. By generating both household decisions and government fiscal exposure as endogenous outputs, our framework allows us to jointly evaluate private LTCI and public support across the wealth distribution. Three findings emerge. First, the means-tested Home Equity Access Scheme (HEAS) functions as a liquidity-bridging channel for the LTCI premium; however, this channel narrows along the wealth distribution due to HEAS's compounding borrowing rate, with bequest motives reinforcing the closure at the top. Second, private LTCI exhibits a fiscal divide: it generates net public savings for lower-wealth households (1.7%–2.7%), yet results in no discernible change to government expenditure for wealthier cohorts. Third, nested means tests impose a hump-shaped clawback on private LTCI indemnities that peaks at the income-to-asset test transition and scales with health severity, hollowing out the middle of the wealth distribution. These findings imply that expanding private LTCI is fiscally sustainable, but requires coordinated reform of the means-testing architecture: flattening the clawback peak and differentially recalibrating HEAS access across the wealth distribution.

Keywords: Long-term care insurance; Home equity release; Means-tested home care; Middle-wealth trap; Fiscal incidence.

JEL classification: D14, D15, G22, H55, I13, J26, R31

1 Introduction

Population ageing is accelerating. Across OECD countries, the share of people aged 65 and above has doubled (OECD, 2023), and the demand for long-term care (LTC) is rising with it. LTC covers the sustained assistance that older people need with daily living when chronic illness or cognitive decline erodes functional capacity. Unlike acute care, where the goal is recovery, LTC is open-ended: a household entering it does not know how long it will last, how intensive it will become, or how much it will cost. Public budgets are straining under the resulting demand. Even with public subsidies, out-of-pocket LTC expenses for individuals with severe needs remain prohibitive, depleting over 70% of a household's average cash inflow (OECD, 2024). While retirees often possess substantial wealth, a huge portion is tied up in illiquid assets, predominantly home equity (Wood et al., 2022). The mismatch between rising care costs and locked-up wealth raises a question that remains inadequately addressed by existing public-private frameworks: how can households combine private insurance, public support, and home equity to finance LTC without sacrificing financial security or planned bequests?

Despite the high costs of late-life care, private long-term care insurance (LTCI) markets remain small. This "LTCI puzzle" has been explained through three forces: the luxury-bequest and precautionary-saving motives of wealthier households (Lockwood, 2018; Ameriks et al., 2020), the availability of informal family care (Costa-Font, 2010; Mommaerts, 2025), and the crowd-out effect of means-tested public programs (Brown and Finkelstein, 2008). The home equity release market exhibits a similar stagnation. Although home equity is the dominant late-life illiquid asset and a natural candidate for financing care (Wood et al., 2022), take-up of private reverse mortgages remains low: high upfront costs (Nakajima and Telyukova, 2017), strict maintenance rules (Cocco and Lopes, 2020), and precautionary hoarding for severe health shocks (Davidoff, 2010) keep the asset locked up. Existing work has examined these two markets in parallel (Shao et al., 2019) but has stopped short of evaluating their interaction within the public-policy architecture that governs both. Modelling either market in isolation, or modelling them jointly without their institutional environment, misses the channels through which the two markets either reinforce or undercut one another.

The recent Australian home care reforms shift more non-clinical care costs onto retirees, bringing renewed attention to financing channels such as LTCI and home equity release. However, evaluating their viability in Australia requires navigating a highly nested policy architecture. Three statutory programmes, namely the age pension, the Home Equity Access Scheme (HEAS), and home care subsidies, are tightly coupled (DHDA, 2025b). Private financial decisions trigger a feedback loop because a retiree's financial profile simultaneously determines their pension entitlement, HEAS borrowing capacity, and care self-contributions. Additionally, the regulatory framework introduces a structural asymmetry. While means testing includes LTCI indemnities, it explicitly exempts HEAS drawdowns. We exploit this one-way interaction to isolate LTCI as the primary driver of the feedback loop. Since HEAS drawdowns do not trigger further clawbacks, they function as a neutral liquidity buffer. In contrast, LTCI indemnities are recursive because they simultaneously provide funding while eroding pension entitlements and inflating net costs. This insulation of the HEAS enables us to filter out debt-side noise and quantify LTCI's structural impact on household welfare.

Consequently, purchasing LTCI subjects policyholders to a cascade of non-monotonic implicit taxes. By treating private LTCI indemnities as assessable lifetime income, our model captures how these payouts simultaneously erode pension entitlements and inflate out-of-pocket home care costs, imposing an implicit tax on middle-wealth households. This phenomenon unites several forms of policy-induced distortion documented internationally: it mirrors the implicit tax that US Medicaid imposes on private LTCI under its insurance-first rules (Brown and Finkelstein, 2008); its steep, hump-shaped profile is structurally analogous to the "tax torpedo" affecting US retirees' Social Security benefits (Reichenstein and Meyer, 2018); and the resulting hollowing-out of the middle class exemplifies the classic savings trap generated by UK means-tested pensions (Sefton et al., 2008). Although institutional drivers differ, the economic consequence is consistent: a severe, non-linear marginal penalty concentrated on middle-wealth households. This cross-jurisdictional consistency reinforces our central argument: evaluating late-life risks requires moving beyond isolated constraints to capture the joint dynamics of nested policy feedback loops.

The primary contribution of this paper is to extend the theoretical framework of analyses that treat policy distortions in isolation (Hubbard et al., 1995; Brown and Finkelstein, 2008; Lockwood, 2018; Reichenstein and Meyer, 2018) to encompass a multi-layered retirement welfare architecture. We achieve this by integrating multiple nested and cross-dependent policy rules directly into the household's optimisation problem within a unified continuous-time framework. By explicitly modelling these interdependent rules, we characterise the non-monotonic, hump-shaped trajectory of government benefit clawbacks on private LTCI indemnities. We show that what we term the "middle-wealth trap", the wealth trap confronting middle-wealth households, is a structural byproduct of this nested policy architecture.

Second, we provide a policy-aligned joint-life model that ties health dynamics directly to public-coverage rules. Traditional multi-state approaches (Fong et al., 2015) track clinically defined transitions, which diverge from the statutory thresholds that actually govern subsidy eligibility; health and policy are then spliced together only after the fact. We instead align the two ex ante: drawing on twenty waves (2004–2023) of the HILDA survey (Watson and Wooden, 2012), we estimate a proportional-odds model (McCullagh, 1980) that maps long-term health conditions onto the care levels recognised by the subsidy rules, and integrate this estimated mapping into a continuous-time joint-life Markov process with a risk-equivalent age (Fong et al., 2025). Two payoffs follow. Health-driven robustness exercises map directly onto eligibility transitions rather than onto clinical labels, so every health perturbation has an immediate fiscal interpretation. And because the construction is estimated rather than assumed, it transfers to any system whose subsidies are organised around comparable care-level thresholds.

Third, we link household optimisation to government long-term fiscal exposure within a single state-space. The household-side LTCI literature (Lockwood, 2018; Ameriks et al., 2020) characterises private demand without modelling the public balance sheet, while fiscal projections of long-term care take household behaviour as given rather than as an endogenous response to policy parameters. This separation masks the structural channels through which private insurance reshapes public liability. Our baseline framework (Lyu et al., 2026) establishes this household–government linkage exclusively for lower-wealth households, thereby precluding the analysis of the middle-wealth trap. Here, we extend the model across the full wealth distribution, ensuring that household decisions and the resulting public liabilities are recovered jointly for every wealth segment rather than for the bottom alone.

The framework yields three findings that, together, sharpen what is currently known about household LTC choice. First, HEAS functions not as a substitute for LTCI but as a liquidity-bridging channel for the premium: dynamic access to home equity through HEAS relaxes the liquidity constraint that otherwise prevents lower-wealth households from purchasing LTCI. Driven by strong bequest motives and deterred by HEAS's compounding borrowing rate, wealthy households bypass the scheme, keeping their home equity intact once LTCI absorbs the care-cost shock. Second, private LTCI exhibits a fiscal divide across the wealth distribution. Among lower-wealth households, LTCI adoption reduces public outlays by substituting government subsidies with self-financed care. Among wealthier cohorts, by contrast, the aggregate fiscal effect plateaus near neutrality: sustained premium outlays deplete assessable wealth and accelerate means-tested eligibility, while LTCI indemnities prematurely trigger the lifetime care cap; these two channels largely offset, leaving government expenditure broadly unchanged. Expanding LTCI thus remains fiscally net-positive on average. Third, overlapping means tests impose a non-monotonic implicit tax on private LTCI indemnities. The clawback is hump-shaped in liquid wealth, peaks at the transition from the income test to the asset test, and scales with health severity. The peak is steep enough that the optimal life-cycle response in the middle-wealth band is to leave the band rather than save through it, hollowing out the middle of the wealth distribution and pushing households towards the wealth extremes.

These findings point to a layered policy response. At the level of instrument choice, expanding private LTCI is more fiscally defensible than broadly relaxing home equity release, since LTCI substitutes private spending for public transfers whereas a more generous release scheme creates new tail exposure for the public budget. Systemically, expanding this coverage requires dismantling the middle-wealth trap by flattening the implicit-tax peak of overlapping means tests, introducing complementary financial instruments, and calibrating release-scheme access to prevent strategic behaviour among the wealthy.

The paper proceeds as follows: Section 2 outlines the institutional setting; Section 3 describes the model; Sections 4–5 detail the data and strategy; and Sections 6–8 present results, policy implications, and conclusions.

2 Institutional features

2.1 Three public programmes, one means-testing architecture

We formalise the three pillars of Australian retirees' cash flows: the age pension, HEAS, and home care subsidies. Their deeply nested means-testing architecture distinguishes the Australian system internationally.

- **The age pension as the central anchor.** Serving as the baseline for retiree cash flows, the pension is governed by non-linear means tests that, despite exempting home equity (assuming owner-occupancy), trigger a statutory clawback via steep taper rates, ultimately producing the middle-wealth trap detailed in Section 6.4.

- **Pension-inverse liquidity channel via HEAS.** The instantaneous HEAS draw cap is set inversely to the pension entitlement. As the pension shrinks, housing liquidity expands, capping combined cash flows at 1.5 times the maximum pension. HEAS withdrawals remain exempt from income tests.
- **Pension-linked home care self-contributions.** Out-of-pocket care obligations are tethered to pension status. A reduced pension shifts a larger share of care costs from the public budget onto the household, subject to a lifetime cap.

Table 1 summarises the notation, cash-flow roles, and dependency structures of these programmes, briefly contrasting them with US and UK systems. Table 7 extends this comparative analysis later in the paper, situating our contribution within the broader literature on means-testing and implicit taxation.

2.2 How LTCI propagates through the system

To integrate private LTCI into the public welfare system, we avoid arbitrary assumptions and strictly replicate how the Australian government treats analogous financial products. On the premium side, we mirror Australia's private health insurance framework by applying a flat tax rebate. On the claims side, we adopt the means-testing rules for lifetime income streams, treating a fraction of the LTCI indemnities as assessable income under the age pension income test. By anchoring our model in actual market products, we explicitly link households' private insurance decisions to the three public pillars.

Premium payments work through the assets test. Because premiums are partially excluded from assessable wealth, paying a premium lowers a household's assessable assets and raises the age pension it receives. The higher pension, in turn, reduces means-tested home care self-contribution, but it also tightens the instantaneous HEAS draw cap, since HEAS borrowing capacity is set inversely to pension. Therefore, LTCI premium payment reshapes all three public programmes at once.

LTCI indemnities work through the income test, and the cascade runs in reverse. A fraction of the LTCI indemnities enters assessable income and erodes the age pension. The lower pension raises out-of-pocket home care payments and expands HEAS borrowing capacity. Concurrently, it accelerates the household's path towards the lifetime self-contribution cap for home care. Once that cap binds, severe-state care costs shift from the household onto the public budget.

Table 1. Model setup: institutional rules and corresponding mathematical notation.

Panel A: Notation.

Private market cash inflows

$I^{\mathrm{LTCI}}(t)$ — Benefits received from long-term care insurance.

$I^{\mathrm{inv}}(t)$ — Investment income on assessable assets. ▷ Operationalised as deemed income computed from $X(t)$ via statutory deeming rates.

State variables and stocks

$X(t)$, $Y(t)$ — Non-home equity household assets and home equity.

$\kappa(t)$, $\dot{\kappa}(t)$ — Outstanding HEAS loan balance and its rate of change. ▷ Accruing with a fixed interest rate r^{κ}.

$H^{\mathrm{ind}}(t)$, $H^{\mathrm{sp}}(t)$ — Individual's and spouse's health states.

$\zeta^{\mathrm{ind}}(t)$, $\zeta^{\mathrm{sp}}(t)$ — Cumulative self-contribution to home care services for both the individual and the spouse. ▷ Lifetime self-contribution cap $\bar{\zeta}$ (AUD 130,000 p.p.).

Indicators and parameters

$\mathbb{1}^{\mathrm{HFS}}(t)$ — Housing-and-family-status indicator: housing tenure (homeowner vs non-homeowner) and marital status.

$\phi \in [0, 1]$ — Constant fraction of the instantaneous maximum admissible HEAS drawdown rate at which the household draws.

Panel B: Institutional features and dependency structure.

Component	Notation	Cash-flow role	Means-testing rules	Constraints	Inputs	International comparison
Age pension (AP)	$I^{\mathrm{AP}}(t)$	Inflow: non-contributory income support for retirees.	Subject to both income and assets tests; entitlement is determined by the stricter of the two.[a]	The maximum pension rate differs by $\mathbb{1}^{\mathrm{HFS}}(t)$	$X(t)$, $I^{\mathrm{LTCI}}(t)$, $I^{\mathrm{inv}}(t)$, $\mathbb{1}^{\mathrm{HFS}}(t)$.	The US and UK combine a primary contributory pension system with secondary means-tested support (Braun et al., 2017; Sefton et al., 2008).
Home Equity Access Scheme (HEAS)	$I^{\mathrm{HEAS}}(t)$	Liquidity inflow: government-provided reverse mortgage unlocking illiquid home equity.	The total borrowing limit relaxes as home equity and age rise; the instantaneous draw cap relaxes as $I^{\mathrm{AP}}(t)$ falls, subject to a binding cap.[b]	$\dot{\kappa}(t) = r^{\kappa}\kappa(t) + I^{\mathrm{HEAS}}(t)$; $I^{\mathrm{HEAS}}(t)$ is excluded from the AP income test, but any unspent balance accumulating in $X(t)$ enters the assets test.	$Y(t)$, $I^{\mathrm{AP}}(t)$, $\kappa(t)$, ϕ.	This contrasts with the privately issued but government-insured Home Equity Conversion Mortgage in the US (Davidoff, 2015) and the entirely private, industry-regulated equity release sector in the UK (Dowd et al., 2019).
Home care	$F^{\mathrm{HC}}_{\mathrm{ind}}(t)$, $F^{\mathrm{HC}}_{\mathrm{sp}}(t)$	Outflow: subsidised home care services for retirees ageing in place.	Means-tested self-contributions are proxied by $I^{\mathrm{AP}}(t)$, with a higher $I^{\mathrm{AP}}(t)$ corresponding to a lower self-contribution proportion.[c]	Lifetime self-contribution cap $\bar{\zeta}$ for each individual.	$I^{\mathrm{AP}}(t)$, $H^{\mathrm{ind}}(t)$, $H^{\mathrm{sp}}(t)$, $\zeta^{\mathrm{ind}}(t)$.	US Medicaid relies on an impoverishment-based safety net (Brown and Finkelstein, 2008), whereas the UK system imposes strict means-testing without a lifetime self-contribution cap (Hu et al., 2025).

Notes:

[a] Higher assessable income and assets, home-ownership, and cohabiting marital status all signal higher socio-economic status and therefore reduce $I^{\mathrm{AP}}(t)$.

[b] The binding cap requires that the combined total of the instantaneous draw cap and $I^{\mathrm{AP}}(t)$ not exceed 1.5 times the maximum pension rate.

[c] More severe health states $H^{\mathrm{ind}}(t)$ raise the total fee jointly funded by government and individual; a healthier spouse offsets formal-care use through informal-care capacity, with this substitution attenuating as $H^{\mathrm{sp}}(t)$ deteriorates. The specific relationship between $I^{\mathrm{AP}}(t)$ and the self-contribution proportion is illustrated in Figure 1(c).

The timing and magnitude of both cascades are disciplined by joint health transitions, as LTCI indemnities are strictly state-contingent, with each distinct health state triggering a different indemnity level. LTCI therefore operates not as a passive payout vehicle, but as a dynamic, multi-channel instrument that continuously reshapes the household's position across the entire means-tested architecture.

3 Model

Before detailing the formulas, we first conceptualise the household's core dilemma. A typical retired couple, currently not severely disabled but facing stochastic disability risks, must secure funding for potential long-term care shocks. Their goal is to finance these future liabilities, safeguarding their precautionary savings and bequest motives.

3.1 Continuous-time health state transition

3.1.1 Definition of states

To capture a couple's journey, we define their individual health state space $\mathscr{H}^{\text{ind}} = \{1, \ldots, n\}$, where 1 denotes full health, $n-1$ the most severe disability, and n death. We track the spouse by the state space $\mathscr{H}^{\text{sp}} = \{1, \ldots, n\}$. We then define the joint health state space of the couple $\mathscr{H}$. For notational convenience, we introduce the ordering map from the Cartesian product of individual health state spaces:

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

which preserves lexicographic order on the joint health state. A continuous-time Markov chain describes the couple's joint health:

$$H = \{H(t)\}_{t \geq 0}, \quad H(t) = \iota\left(H^{\text{ind}}(t), H^{\text{sp}}(t)\right), \tag{2}$$

where $\left(H^{\text{ind}}(t), H^{\text{sp}}(t)\right) \in (\mathscr{H}^{\text{ind}} \times \mathscr{H}^{\text{sp}})$. The time-inhomogeneous generator governs these transitions:

$$Q(t) = \left(q_{h,h'}(t)\right)_{h,h' \in \mathscr{H}}, \quad q_{h,h'}(t) \geq 0\,(h \neq h'), \quad q_{h,h}(t) = -\sum_{h' \neq h} q_{h,h'}(t).$$

Here, the quantity $q_{h,h'}(t)\mathrm{d}t + o(\mathrm{d}t)$ is the conditional probability that the process jumps from h to h' during the interval $[t, t+\mathrm{d}t)$. Time dependence in $Q(t)$ captures the impact of ageing and the transition probability is determined via Kolmogorov equations.

3.1.2 Stochastic differential representation

Rather than a smooth decline, we model their health as taking sudden jumps (e.g., a cardiovascular event). We represent this using a jump stochastic differential equation:

$$\mathrm{d}H(t) = \sum_{h' \neq H(t^-)} (h' - H(t^-))\, \mathrm{d}N^{H(t^-),h'}(t), \quad H(0) \in \mathcal{H}, \tag{3}$$

where $H(t^-) := \lim_{u \uparrow t} H(u)$ and $N^{h,h'}(t)$ is a non-homogeneous Poisson process of rate $q_{h,h'}(t)$. More details for this representation are provided in Appendix A.

3.2 LTCI

Following previous research (Shao et al., 2019; Xu et al., 2023), we assume the household purchases the policy with a single premium at retirement, carrying no expense loading. The LTCI insurer pays indemnities whenever the household enters an insured state, the definitions of which are detailed in Section 5.2.3.

3.2.1 Pricing model

To intuitively understand the pricing framework, consider how the insurer evaluates the couple at retirement. Rather than simply summing their independent expected care costs, the model explicitly accounts for mutual spousal support by netting out the informal care capacity that the healthier spouse can provide (Mommaerts, 2025), which directly reduces the household's net insurance liability. Mathematically, for a household in joint state $h = \iota(h^{\text{ind}}, h^{\text{sp}})$, we define the net long-term care cost LTC_h as:

$$\text{LTC}_h = \left(\widetilde{\text{LTC}}_{h^{\text{ind}}} - \tilde{v}(1 - (h^{\text{sp}} - 1)\varpi)^+\right)^+ + \left(\widetilde{\text{LTC}}_{h^{\text{sp}}} - \tilde{v}(1 - (h^{\text{ind}} - 1)\varpi)^+\right)^+, \tag{4}$$

where $\widetilde{\text{LTC}}_{h^{\text{ind}}}$ and $\widetilde{\text{LTC}}_{h^{\text{sp}}}$ denote the gross care costs for the individual and spouse, respectively. The term $\tilde{v}$ represents the maximum informal care a healthy spouse can provide, while $\varpi > 0$ governs the rate at which caregiving capacity declines with the caregiver's own health deterioration.

The actuarially fair premium P for full coverage of a household is given by

$$P(\mathsf{C}_{\mathcal{H}}(\mathcal{H}'); t) = \int_t^T e^{-rs} \sum_{h' \in \mathcal{H}'} \Pr\left(H(s) = h' \mid H(t) \in \mathsf{C}_{\mathcal{H}}(\mathcal{H}')\right) \text{LTC}_{h'} \mathrm{d}s, \tag{5}$$

where $\mathcal{H}' \subset \mathcal{H}$ is the set of indemnities-paying states, $\text{LTC}_{h'}$ is the pre-subsidy care cost in state h', T represents the maximum possible survival time, and r is the risk-free interest rate. The household chooses a coverage proportion $\psi \in [0, 1]$ and pays the corresponding premium $\psi P(\mathsf{C}_{\mathcal{H}}(\mathcal{H}'); 0)$.

3.2.2 Interaction with means tests and tax treatment

The interaction between private insurance and public support creates a financial tug-of-war for the household. We model an indemnity-type LTCI contract. As illustrated in Section 2, a portion $\nu \in [0,1]$ of the LTCI indemnities is treated as assessable income. Consequently, larger private LTCI indemnities trigger a reduction in public subsidies, creating a pronounced implicit tax.

To capture the policy reality, we incorporate the asymmetric tax treatment of LTCI directly into the household's budget constraints. On the front end, the government incentivises uptake via an upfront premium rebate of $\eta\psi P(\complement_{\mathcal{H}}(\mathcal{H}');0)$, where $\eta \in [0,1]$ is the rebate rate. This acts as a contemporaneous increase in initial liquid wealth, though the rebated capital remains subject to the prevailing assets test.

3.3 Underlying financial market

Turning to the household's balance sheet, retirees manage two primary stores of wealth to fund their late-life needs: liquid financial investments and illiquid home equity. Because the valuations of both asset classes are subject to unpredictable market shocks, we capture these dynamics through a continuous-time stochastic framework. Formally, 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, where $\mathbb{F} = \{\mathcal{F}_t : t \geq 0\}$ is the augmentation of the natural filtration of W and H. The household can invest in a traded risky asset with a price process $S(\cdot)$ given by

$$\mathrm{d}S(t) = S(t)\left(\mu_1 \mathrm{d}t + \sigma_1\, \mathrm{d}W_1(t)\right), \quad S(0) > 0, \tag{6}$$

and in a riskless asset with a constant continuously compounded rate of return r.

The household also owns a non-traded home equity $Y(\cdot)$, whose value follows a geometric Brownian motion with correlation ρ to the traded asset:

$$\mathrm{d}Y(t) = Y(t)\left(\mu_2 \mathrm{d}t + \sigma_2\rho\, \mathrm{d}W_1(t) + \sigma_2\sqrt{1-\rho^2}\, \mathrm{d}W_2(t)\right), \quad Y(0) \geq 0. \tag{7}$$

Here, $\mu_1, \mu_2 \in \mathbb{R}, \sigma_1, \sigma_2 > 0$, and $\rho \in [-1,1]$; then $\operatorname{corr}(\mathrm{d}\log S,\ \mathrm{d}\log Y) = \rho$. Returns load on the market shock W_1 with weight ρ and on an orthogonal idiosyncratic shock W_2, capturing systematic co-movement with the general risky asset while retaining home equity risk that cannot be hedged and maintaining log-normal tractability for portfolio choice, mathematically analogous to the unhedgeable income risk framework in Munk (2000).

3.4 Household optimisation

Navigating late-life financial risks requires the household to perform a balancing act: allocating liquid resources between instantaneous non-durable consumption $c(t) > 0$ and the risky portfolio share $\pi(t)$. We additionally impose the portfolio bound $0 < \pi(t) \leq \bar{\pi}$ for all t, with $\bar{\pi}$ specified by household segment in Section 5.

Beyond these continuous controls, we assume that the household determines its non-instantaneous strategies exactly at retirement. These fixed decisions include the LTCI coverage share ψ (and the resulting premium payment ψP) and the home equity release intensity ϕ, defined as the instantaneous maximum admissible HEAS drawdown rate.

3.4.1 Household wealth process

The planning horizon terminates at the stochastic time of death, τ, of the surviving spouse. Under the self-financing condition, the household wealth process $X(t)$ satisfies the following stochastic differential equation for $t \in [0, \tau)$:

$$\mathrm{d}X(t) = \big(rX(t) + \pi(t)X(t)(\mu_1 - r) - c(t) + I(t) - F(t)\big)\,\mathrm{d}t + \pi(t)X(t)\sigma_1\,\mathrm{d}W_1(t). \tag{8}$$

Here, $I(t)$ represents the income flows and $F(t)$ represents home care expenses, defined as:

$$I(t) = I^{\mathrm{AP}}(t) + I^{\mathrm{HEAS}}(t) + I^{\mathrm{LTCI}}(t) \quad \text{and} \quad F(t) = F^{\mathrm{HC}}_{\mathrm{ind}}(t) + F^{\mathrm{HC}}_{\mathrm{sp}}(t).$$

3.4.2 Boundary conditions

The wealth process starts at $t = 0$ with an initial wealth $X(0)$ after paying the LTCI premium $(1-\eta)\psi P(\complement_{\mathscr{H}}(\mathscr{H}');0)$. At the random time of death τ, the terminal wealth $X(\tau)$ is determined by the wealth just prior to death $X(\tau^-)$ plus the net bequest from home equity $Y(\tau)$ less the outstanding HEAS balance $\kappa(\tau)$:

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

3.4.3 Household preference

The preferences over consumption and bequests follow a standard constant relative risk aversion (CRRA) utility form:

$$U_1(c,h) = \frac{(\xi_1(h)c)^{1-\gamma}}{1-\gamma}, \qquad U_2(x) = b\frac{(x+\bar{\xi}_2)^{1-\gamma}}{1-\gamma}, \quad c, x > 0, \tag{9}$$

where $\gamma \in (0,1) \cup (1,\infty)$ is relative risk aversion, $b > 0$ scales the bequest motive. Here, $\xi_1(h)$ is tailored to the joint-life setting with $h = \iota(i,j)$ as defined in Equation (1):

$$\xi_1(h) = \frac{1}{\bar{\xi}_1^{\mathbb{1}(i<n,j<n)}}\left[\frac{(n-\varepsilon(i-1))\mathbb{1}(i<n) + (n-\varepsilon(j-1))\mathbb{1}(j<n)}{n}\right]^{\frac{1}{1-\gamma}}, \tag{10}$$

where ε follows the health-dependent consumption motivation in Finkelstein et al. (2013), while the equivalence-scale adjustment in the joint-life setting $\bar{\xi}_1 \in (1,2)$ is aligned with the standard specification in De Nardi et al. (2025). In addition, modelling $\bar{\xi}_2 > 0$ as an additive shift inside U_2 preserves the CRRA curvature, preventing overstatement of precautionary saving and bequest intensity (De Nardi et al., 2025).

3.4.4 Household objective

The objective is to choose the controls $(c(u), \pi(u))$ to maximise the household's expected lifetime utility. The value function V represents this maximum achievable utility, discounted at rate δ, given the current state:

$$V(t,x,y,h,\kappa,\zeta^{\text{ind}},\zeta^{\text{sp}}) = \sup_{(c(u),\pi(u))} \mathbb{E}\left[\int_t^\tau e^{-\delta(u-t)}U_1\big(c(u),h\big)\,du \;+\; e^{-\delta(\tau-t)}U_2\big(X(\tau)\big)\right]. \tag{11}$$

By the principle of dynamic programming, this value function satisfies the Hamilton-Jacobi-Bellman (HJB) equation, which characterises the instantaneous optimal tradeoff:

$$\begin{aligned}
0 = \sup_{c,\pi}\Bigg\{ & U_1(c,h) - \delta V + V_t + \underbrace{V_x(rx + \pi x(\mu_1 - r) - c + I(t) - F(t))}_{\text{drift of liquid wealth}} \quad + \quad \underbrace{V_y \mu_2 y}_{\text{drift of home equity}} \\
& + \underbrace{V_\kappa\Big(r^\kappa \kappa + I^{\text{HEAS}}(t)\Big)}_{\text{HEAS balance}} + \underbrace{V_{\zeta^{\text{ind}}}\,\mathbb{1}_{\{\zeta^{\text{ind}}<\bar{\zeta}\}}F^{\text{HC}}_{\text{ind}} \;+\; V_{\zeta^{\text{sp}}}\,\mathbb{1}_{\{\zeta^{\text{sp}}<\bar{\zeta}\}}F^{\text{HC}}_{\text{sp}}}_{\text{caps on self-contributions}} \\
& + \frac{1}{2}V_{xx}(\pi\sigma_1 x)^2 + \frac{1}{2}V_{yy}(\sigma_2 y)^2 + V_{xy}\,\pi\rho\,\sigma_1\sigma_2 xy \Bigg\} \\
& + \underbrace{\sum_{h'\neq h} q_{h,h'}(t)\Big(V(t,x,y,h',\kappa,\zeta^{\text{ind}},\zeta^{\text{sp}}) - V(t,x,y,h,\kappa,\zeta^{\text{ind}},\zeta^{\text{sp}})\Big)}_{\text{health transitions}}.
\end{aligned} \tag{12}$$

Here, local non-concavities ($V_{xx} > 0$) are structurally induced by the non-convexity of the budget set inherent in means-tested age pension and home care costs (Hubbard et al., 1995). The statutory tapering of benefits acts as a steep implicit marginal tax on savings: as private wealth increases, subsidies are withdrawn, creating a "savings trap" (Sefton et al., 2008). To realise a net increase in future consumption, households must accumulate a critical mass of wealth to break out of this withdrawal zone. This threshold effect creates areas where the marginal incentive to save diminishes or reverses, generating local non-concavities in the value function where divergent strategies, such as rapidly spending down to rely on the public safety net versus aggressively saving, can yield comparable utility.

We solve the high-dimensional, nonlinear HJB equation via backward-in-time policy iteration, alternating between updating controls and evaluating the value function. For continuous financial states, an accelerated Craig-Sneyd Alternating Direction Implicit scheme handles diffusion and correlation in logarithmic coordinates. Auxiliary and path-dependent states are managed on a sparse Smolyak grid, utilising monotone upwind discretisations to ensure numerical stability (Hanson, 2007). To address the local non-concavities and kinks induced by institutional rules, we introduce a vanishing-viscosity regularisation (Yong and Zhou, 1999; Barles and Souganidis, 1991), which stabilises the numerical optimisation and captures corner solutions (Fella, 2014). Finally, regime-switching dynamics from health transitions and mortality are resolved implicitly via a pre-factorised, coupled linear system at each time step. Implementation details are provided in Appendix B.

4 Data

4.1 Individual-level data

To capture the real-world trajectories of older Australians as they navigate ageing, health shocks, and financial decisions, we construct our primary sample using the Household, Income and Labour Dynamics in Australia (HILDA) survey (Watson and Wooden, 2012). Rather than relying on a static snapshot, we track individuals across twenty annual waves (2004–2023), monitoring shifts in their demographic profiles, long-term health conditions, and household finances.[1] When modelling retirement resources, we focus on the wealth retirees can actually deploy: following the HILDA wealth model (Summerfield et al., 2024), we calculate assessable household assets by taking total assets, deducting all outstanding debt, and excluding the value of the principal dwelling.[2] Appendix C provides a comprehensive breakdown of our HILDA variable selection.

4.2 Aggregate-level data

To model the financial reality of ageing, we must account for the evolving landscape of government subsidies that retirees navigate. Prior to 2025, older Australians faced a fragmented framework split between entry-level support (the Commonwealth Home Support Programme, CHSP) and tiered, higher-level care (Home Care Packages, HCP). The new Support at Home (SAH) programme fundamentally reshapes this environment by integrating these legacy streams into a single, seamless continuum of care (as mapped out in Figure 1(a)). By extracting budget allocations for both historical and upcoming care levels directly from DHDA (2025a), we capture this structural shift in the Australian home care system. Full details regarding data sources and parameter calibrations are provided in the figure caption.

But how exactly does a decline in an individual's health translate into actual care utilisation? To bridge this gap, we bring in administrative records detailing the aggregate allocation of home care services by age group. By linking these macro-level service statistics (Figure 1(b)) with the micro-level health indicators from our HILDA sample (Figure 2), we empirically recover the relationship between health shocks and care needs (we detail this identification strategy in Section 5).

5 Empirical analysis and calibration

5.1 Sample construction

To build our analysis sample, we track individuals through the HILDA survey to map their health states. First, we isolate respondents who provide a definitive health status: whether

[1]When analysing financial behaviours, we restrict our sample to the eleven waves between 2013 and 2023. This specific timeframe ensures our data captures the market environment shaped by the 2013 Living Longer Living Better (LLLB) reforms, which governed the home care sector until the introduction of the 2025 new Aged Care Act.

[2]We adjust all financial variables to reflect a standard trend inflation rate of 2% per annum, aligning with advanced economy inflation targets (Ascari and Sbordone, 2014).

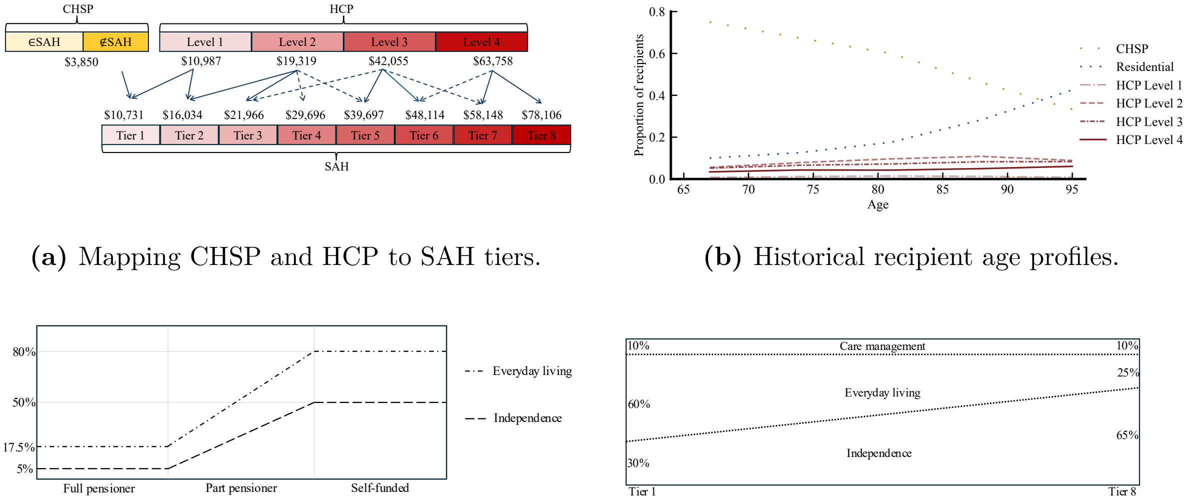

(a) Mapping CHSP and HCP to SAH tiers.

(b) Historical recipient age profiles.

(c) SAH contribution rates.

(d) Assumed SAH budget composition.

Figure 1. Mapping the transition from legacy systems (CHSP and HCP) to SAH tiers: age profiles and budget composition.

Tier mapping and age profiles: Figure 1(a) maps legacy programmes to SAH tiers based on assigned budgets (DHDA, 2025a), assuming SAH Tier 1 covers only a subset of former CHSP recipients due to budget disparities. Figure 1(b) calibrates SAH age profiles using historical recipient shares (Productivity Commision, 2025; AIHW, 2025).

Means-testing and self-contributions: Under SAH, retiree contributions are means-tested against the age pension (Figure 1(c)). 'Clinical care' is fully subsidised, whereas 'Independence' and 'Everyday living' services require self-contributions. Contribution rates are linearly interpolated for part-pensioners based on their pension income (I^{AP}) and are subject to an indexed lifetime cap.

Demand and budget allocation: We define a tier's estimated home care demand as the sum of government funding (assumed fully utilised) and full pensioners' contributions. After deducting care management fees (DHDA, 2025c), Figure 1(d) details the budget split between 'Independence' and 'Everyday living', using adapted empirical data from Nektarios et al. (2025).

they report specific long-term conditions or explicitly confirm they have none.[3] We then filter this cohort to include only those designated as permanent sample members. With our core group established, we construct a health trajectory for each person by linking their annual observations.[4] Finally, we classify each trajectory endpoint into one of three states. A trajectory is coded as ending in death whenever HILDA records a death in the final follow-up year or the subsequent year, with the reported year of death retained as the endpoint. When the final observed wave attributes nonresponse specifically to severe disability, we impute the maximum severity score and right-censor the series at that observation. All remaining trajectories are right-censored at the final observed wave, with the HILDA longitudinal weight for that observation halved to mitigate censoring-induced bias (Klein and Moeschberger, 2003, p. 152).

[3] We also ensure that the corresponding health status for each respondent's spouse is fully recorded.

[4] If a respondent misses a survey wave (interval censoring), we bridge the gap by linearly interpolating between their adjacent observed time points.

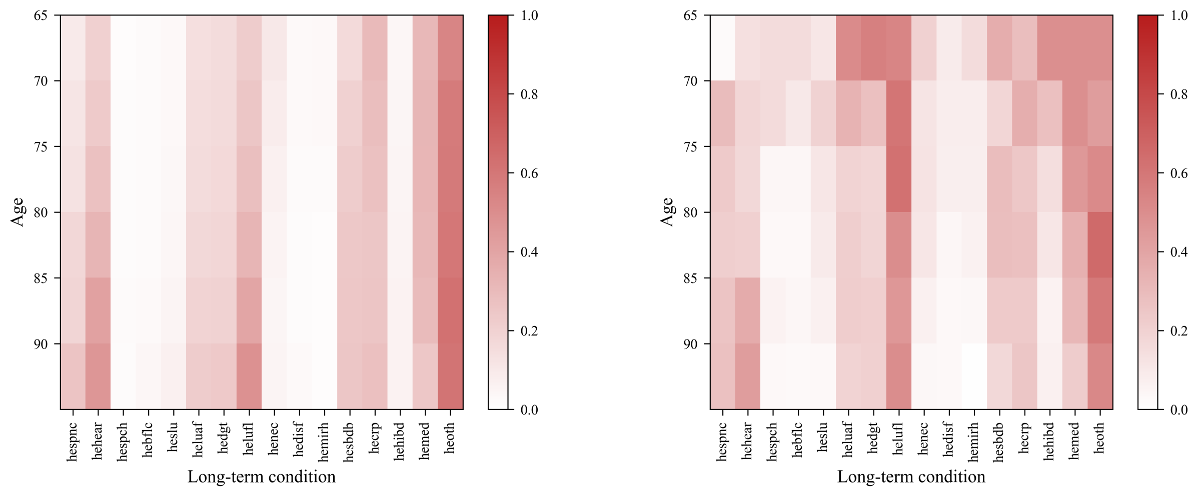

(a) Ageing in place. **(b)** Ageing in residential home care facilities.

Figure 2. Age-related prevalence of long-term health conditions by residential status. This figure displays the proportion of individuals reporting specific long-term health conditions, stratified by age. Proportions are calculated within each residential group. The residential care subpopulation is explicitly included here to establish a baseline for evaluating the model's empirical fit, the details of which are presented in Table 2.

Notes: Data are sourced from HILDA Survey, Waves 4–23. Conditions correspond to HILDA variables: Sight problems (`hespnc`); Hearing problems (`hehear`); Speech problems (`hespch`); Blackouts/fits (`hebflc`); Learning difficulties (`heslu`); Limited use of arms/fingers (`heluaf`); Difficulty gripping (`hedgt`); Limited use of feet/legs (`helufl`); Nervous/emotional condition (`henec`); Disfigurement (`hedisf`); Mental illness (`hemirh`); Shortness of breath (`hesbdb`); Chronic pain (`hecrp`); Effects of brain injury/stroke (`hehibd`); Restrictive condition despite medication (`hemed`); and Other long-term conditions (`heoth`). Detailed descriptions of these variables are provided in Appendix C.

5.2 Disability classification

We need to assign each HILDA respondent to a care level, but HILDA does not record who actually receives which service, or why. We work around this gap in two steps: first, we use historical home care allocation proportions (how many people in each age band sit at each care level) as a population-level anchor (Figure 1(b)); then we combine these proportions with each respondent's reported long-term conditions to map individuals to home care levels (Figure 2). These mapped levels become the health states that drive our transition model.

5.2.1 Proportional odds model

We can think of each respondent carrying an invisible health burden, which is a single number that grows as conditions accumulate. For example, a respondent with mild arthritis sits low on the scale; one with diabetes, heart disease, and a recent stroke sits much higher. Whenever this burden crosses a threshold, the respondent moves up to the next required care level. This is the intuition behind the proportional odds model (McCullagh, 1980), which we adopt here: long-term conditions shift individuals along a latent burden scale rather than triggering arbitrary jumps between discrete states.

Formally, let $L \in \{1, \ldots, K\}$ denote the required care level and $\mathbf{x}$ a vector of binary indicators for long-term conditions. The model gives the probability of needing at least level k as:

$$\Pr(L \geq k \mid \mathbf{x}) = \left[1 + \exp\{-(\alpha + \beta^\top \mathbf{x} - c_k)\}\right]^{-1}. \tag{13}$$

The latent burden itself is $\alpha + \beta^\top \mathbf{x}$: α sets the baseline, β (constrained non-negative) tells us how much each condition adds to the burden, and the ordered cutpoints $c_1 < \cdots < c_{K-1}$ mark where one care level gives way to the next.

5.2.2 Parameter estimation

Following the methodology of Berry et al. (1995), we use a grouped data maximum likelihood approach to estimate individual-level parameters, $\Theta = \left(\alpha, \beta^\top, c_1, \ldots, c_{K-1}\right)'$, from aggregate data. The estimator is defined by maximising the grouped multinomial log-likelihood function:

$$\widehat{\Theta} \in \arg\max_{\Theta} \sum_{g=1}^{G} \sum_{k=1}^{K} N_g \widehat{\Pi}_{gk} \log \left(\frac{1}{N_g} \sum_{i \in \mathcal{I}_g} \Pr\left(L = k \mid \mathbf{x}_i, \Theta\right) \right).$$

Here, the data are aggregated into G age bands, indexed by $g = 1, \ldots, G$. $\mathcal{I}_g$ denotes the set of N_g individuals within age group g. $\widehat{\Pi}_{gk}$ represents the empirical share of individuals in age group g selecting the home care service category k. Maximising the log-likelihood is equivalent to minimising the total weighted cross-entropy (Kullback and Leibler, 1951), between this empirical share, $\widehat{\Pi}_{gk}$, and the corresponding predicted share, Π_{gk}.

5.2.3 Health state classification and transition dynamics

How well does the proportional odds model recover what we observe? Predicted and observed shares track across all age bands and care levels (Table 2), giving us confidence to use the model as a basis for state classification.

Joint-life setup. Before detailing the specific five-state classification, we must first establish how individual trajectories combine. We model couples as two parallel lives, both starting at retirement. For each spouse we build a separate household health trajectory; the household's joint state at any time is the ordered pair of the two individual states (Equation (1)). Starting both spouses at the same calendar time is simplification: it strips gender and spousal-age-gap effects out of the design so that the differences we later attribute to health are not contaminated by who is older.[5]

From individual transitions to a continuous-age generator. We aggregate individual transitions and exposure times within six age groups ($[65, 70)$, $[70, 75)$, $[75, 80)$, $[80, 90)$, $[90, 100)$, $[100, 117)$), and read off the hazard for each ordered pair of states as total events divided by total exposure. We then smooth these grouped hazards with non-negative spline interpolants, fix the diagonals so each row sums to zero, and treat death as absorbing. The result is a continuous-age generator $Q(t)$ whose rows are valid probability intensities at every age: non-negative off-diagonals, finite row sums, no negative survival probabilities at long horizons.

[5] Take one household and align two event-time trajectories: one with $t = 0$ at the older spouse's retirement, the other with $t = 0$ at the younger spouse's. In the first alignment the partner is younger by the spousal age gap; in the second, older by the same amount. Averaging across both alignments centres the partner's age on the retirement benchmark.

Table 2. Comparison of observed and predicted population shares by age group and care level.

Age band	CHSP Obs (%)	CHSP Pred (%)	HCP Level 1 Obs (%)	HCP Level 1 Pred (%)	HCP Level 2 Obs (%)	HCP Level 2 Pred (%)	HCP Level 3 Obs (%)	HCP Level 3 Pred (%)	HCP Level 4 Obs (%)	HCP Level 4 Pred (%)
Panel A: Ageing in place										
65–69	80.446	78.553	0.486	0.623	6.634	7.670	7.827	8.333	4.606	4.820
70–74	74.243	75.426	0.705	0.673	8.989	8.674	10.127	9.624	5.936	5.604
75–79	72.221	73.110	0.815	0.776	10.530	10.434	10.398	9.929	6.037	5.750
80–84	67.378	67.364	0.975	0.854	13.002	12.440	12.305	12.813	6.340	6.529
85–89	59.897	60.749	0.992	0.911	15.844	14.509	15.276	15.171	7.990	8.659
90+	50.235	50.540	0.765	0.960	16.555	16.579	19.647	19.840	12.798	12.080
Panel B: Ageing in residential home care facilities										
65–69		39.181		0.826		5.577		31.474		22.943
70–74		39.376		1.148		11.127		28.245		20.104
75–79		46.302		1.646		14.892		22.554		14.605
80–84		51.530		1.680		15.408		22.262		9.120
85–89		52.140		1.490		15.800		19.387		11.183
90+		52.477		1.941		18.264		17.293		10.025

Notes: This table presents observed and predicted shares for five ordered levels of home care needs (CHSP and HCP Levels 1-4). The predictions are derived from a proportional-odds model (McCullagh, 1980), estimated by grouped multinomial likelihood (Berry et al., 1995). To ensure model validity, cutpoints are ordered via a monotone reparameterisation, and the non-negativity of burden weights (β) is enforced through a shape-and-scale representation. The data, from the same source as Figure 1(b), are calibrated for wait times using 2023 recipient and waitlist data (AIHW, 2024). Panel A compares observed group shares with in-sample model predictions. Uncertainty is quantified using a non-parametric bootstrap that resamples individuals within groups. Panel B reports out-of-sample predictions for the subpopulation in residential aged care facilities. As observed empirical shares are unavailable for this group, the reported values represent the cross-sectional average of individual-level predicted probabilities for each care state across ages. The projections capture a well-documented empirical regularity: the proportion of residents with severe care needs inversely correlates with advancing age (AIHW, 2025).

Health state classification The annual transitions of the latent burden index display a clear block structure with minimal crossover between clustered regions. As shown in Figure 3, these visible transition blocks suggest three distinct burden index groups: $[8.5, 17.3)$, $[17.3, 24.8)$, $[24.8, 90.0)$. Here, 17.3 falls below the cutpoint for HCP Level 1, while 24.8 represents the cutpoint for HCP Level 3 (Equation (13)). In addition to these thresholds, the CHSP-to-SAH transition is a critical consideration for state classification. Because the SAH Tier 1 budget falls between those of CHSP and HCP Level 1 (Figure 1(a)), we assume that some current CHSP recipients will not qualify for SAH Tier 1 under the new regime.

These observations support a five-state classification model (healthy, three disability states, and death), with disability thresholds defined by the cutpoints in Equation (13).

- *State 1:* Healthy without home care needs;
- *State 2:* CHSP components not absorbed into SAH Tier 1;
- *State 3:* Remaining CHSP components together with HCP Levels 1 and 2;
- *State 4:* HCP Levels 3 and 4;
- *State 5:* Death.

We keep State 2 outside the severely disabled category: retirees in States 1 and 2 are still healthy enough to buy LTCI. Consequently, the eligible joint health states $\mathsf{C}_{\mathscr{H}}(\mathscr{H}')$ in Equation (5) for purchasing LTCI are defined as $\{1, 2, 6, 7\}$ for coupled households, and $\{5, 10, 21, 22\}$ for single households.

The marginal distribution of health states among survivors at each age is shown in Figure 4. A comparison against the empirical survival curve shows a close match overall, with a marginal overestimation near age 90. Overall, this supports the validity of our health transition model.

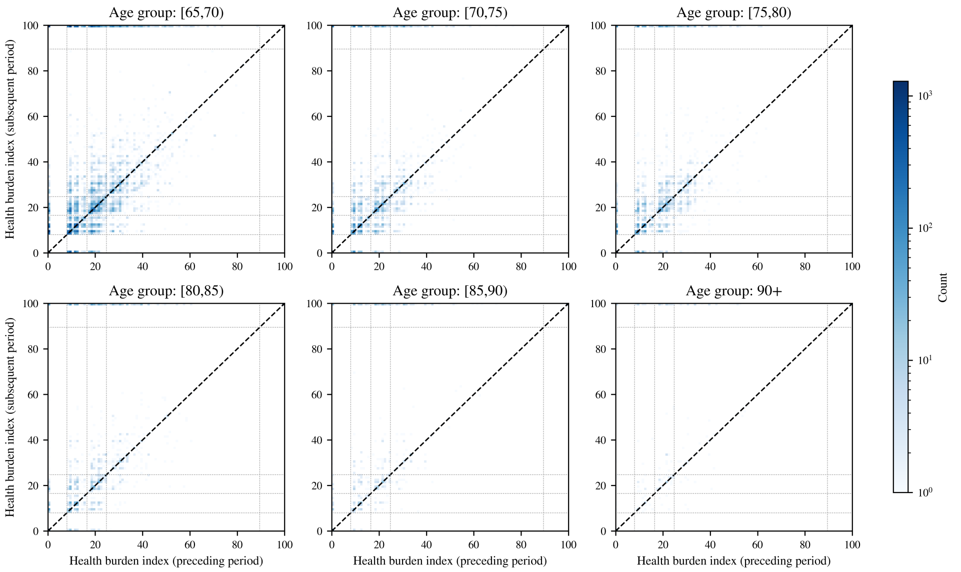

Figure 3. Individual-level annual transition of the latent health burden index stratified by age groups. The index, derived from a proportional odds model's linear predictor ($\alpha + \beta^\top \mathbf{x}$ in Equation (13)), is scaled from 0 (healthy) to 100 (dead or censored). Points above the diagonal indicate health deterioration; points below indicate recovery. The panels show that recovery becomes less likely as disability severity increases, with a stronger effect in older age groups. This pattern aligns with evidence from other countries; see, for example, Fong et al. (2015). Visible transition blocks motivate three disability states, $[8.5, 17.3)$, $[17.3, 24.8)$, $[24.8, 90.0)$, where 17.3 is smaller than the cutpoint for HCP Level 1, and 24.8 is the cutpoint for HCP Level 3. The sample minimum and maximum of the index among disabled individuals define the left and right bounds of the disabled states.

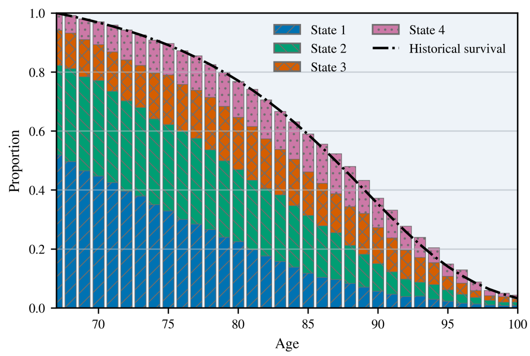

(a) Respondent

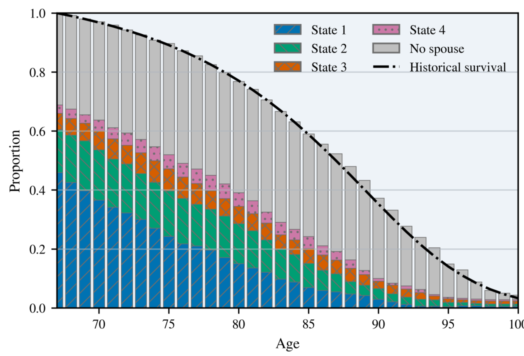

(b) Spouse

Figure 4. Predicted marginal distributions of health state among survivors (ages 67–100) from a piecewise-constant continuous-time Markov model. The stacked bars show the proportion of individuals in each health state at a given age, conditional on survival. The initial distribution is derived from a longitudinally weighted average over the $(65, 70]$ window. The model's generator, $Q(t)$, uses a stable polynomial ramp to approximate Gompertz-like mortality. Figures 4(a) and 4(b) show the health states for respondents and spouses, respectively, independent of the partner's state. In Figure 4(b), individuals without a spouse (i.e., single or widowed) are grouped into a single category. The historical survivor distribution (AGA, 2025) across the HILDA waves utilised supports our health transition model.

5.3 Segmentation of households

We slice the HILDA sample along three dimensions: wealth, marital status, and homeownership. Households first land in one of three wealth tertiles, defined by the lower and upper tercile thresholds of inflation-adjusted total net wealth. Within each tertile, we then split households by marital status (couple vs. single) and by homeownership (owner vs. renter).

Liquid wealth and home equity follow different trajectories across these segments, with homeowners in the top tertile pulling away from the rest after age 70 (Figure 5). One cell of the grid is, however, almost empty: very few middle- and top-tertile households rent rather than own, because wealthier Australians predominantly channel their savings into home equity. We therefore restrict our analysis of non-homeowners to the bottom wealth tertile. Table 3 reports the predicted asset levels we use for each segment.

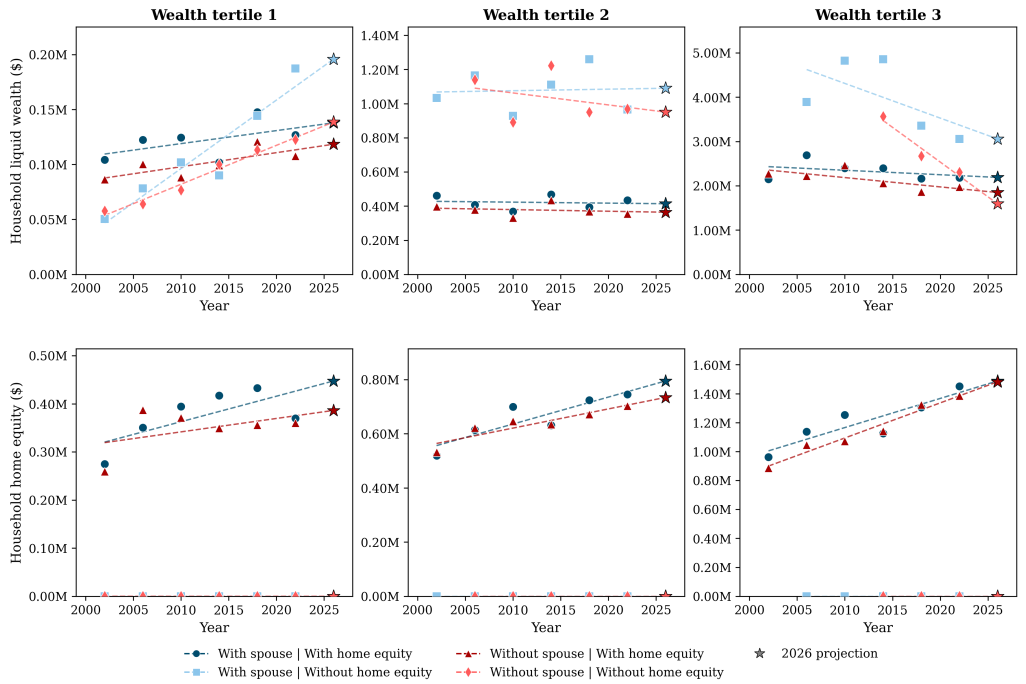

Figure 5. Projected liquid and illiquid wealth profiles by household segment. We stratify households from HILDA survey data into three wealth tertiles based on total net wealth, using the lower and upper tercile thresholds of the distribution. Estimates of wealth at retirement are derived from retirees aged 60 to 69, adjusted for HILDA sample weights and inflation (World Bank, 2026), and projected to 2026 via linear regression.

5.4 Calibration

Table 4 collects every parameter that enters the model. We use this subsection to flag the parameters that drive the results, the choices that required judgement, and the points where our setup departs from the source papers.

Market parameters. We calibrate to long-run, cross-country averages from Jordà et al. (2019) rather than recent data, anchoring the model on a full cycle of asset performance rather than the post-2010 anomaly. Equity offers an expected return of $\mu_1 = 7.8\%$ with volatility $\sigma_1 = 16.9\%$; housing yields a slightly lower return of $\mu_2 = 6.4\%$ with a lower volatility of

Table 3. Estimated median liquid wealth and home equity at retirement, stratified by household segment and wealth tertile.

Wealth tertile	Household segment		Median asset value ($)		Share
	Spouse present	Home equity > 0	Liquid wealth	Home equity	
1	Yes	Yes	156,914	460,098	32.4%
	Yes	No	234,452	0	13.8%
	No	Yes	135,690	396,924	25.2%
	No	No	153,793	0	28.6%
2	Yes	Yes	487,412	810,549	67.0%
	No	Yes	377,305	752,156	28.5%
		No	—	—	4.5%
3	Yes	Yes	2,298,295	1,495,618	76.0%
	No	Yes	1,864,200	1,514,919	21.0%
		No	—	—	3.0%

Notes: Data represent estimated median asset levels for retirees aged 60–69, derived from the Household, Income and Labour Dynamics in Australia (HILDA) survey. "Spouse present" indicates partnered households, whereas "Home equity > 0" designates households holding positive equity in owner-occupied housing. Due to sparse observations, non-homeowners in wealth tertiles 2 and 3 are excluded from subsequent analyses. The majority of middle-wealth households (Tertile 2) retain home equity, with their liquid wealth clustering near the threshold for full age pension eligibility.

$\sigma_2 = 11.9\%$. The liquid-illiquid asset correlation is set to $\rho = 0.2$ (Flavin and Yamashita, 2002). Additionally, the borrowing rate for the HEAS is assumed to be $r^{\kappa} = 4.0\%$, which closely aligns with the compound rate charged by the scheme (Services Australia, 2026a).

Preference parameters. For preferences we lean on De Nardi et al. (2025), who structurally estimate household preferences under explicit late-life risk. Their estimates give a relative risk aversion $\gamma = 3.701$ (moderate by the macro-finance standard, well within the $[2, 5]$ range commonly used), discount rate $\delta = 0.03$, and a couple equivalence scale $\bar{\xi}_1 = 1.528$. Following the unified terminal bequest specification in De Nardi et al. (2025), we set the bequest intensity and curvature parameters to generate a realistic effective wealth threshold (around AUD $42,000$ to AUD $70,000$) below which households exhibit no active bequest motive. The health sensitivity $\varepsilon = 1$ links health to preferences: each one-step health deterioration cuts the marginal utility of consumption by 20%, which is what generates precautionary saving aimed at health and long-term care rather than at discretionary spending (Finkelstein et al., 2013).

LTCI parameters. Long-term care costs rise with disability state: gross annual care costs are AUD 2,100 in State 2, AUD 15,000 in State 3, and AUD 55,000 in State 4 (Figure 1(a)). We assume informal care from a healthy partner offsets up to AUD 2,100 of these costs, declining at rate $\varpi = 1/3$ as the carer's own health deteriorates. This assumption captures the documented shift towards expensive formal care when intra-household support breaks down (Mommaerts, 2025). One modelling choice deserves emphasis here: we assume a healthy partner can fully cover a spouse's State 2 needs. This captures the low-intensity informal support historically delivered through the CHSP that the new SAH cannot fully replicate. The mechanical consequence is that single retirees enjoy less implicit insurance than coupled ones, so a single-life LTCI premium ends up higher than a naive 50% scaling of the joint-life premium would suggest. The joint-life contract prices at AUD 188,925 at retirement, the single-life contract at AUD 138,411, which is

Table 4. Summary of baseline parameter calibration.

Symbol	Description	Value	Remark
Market parameters			
r	Risk free interest rate	2.5%	Jordà et al. (2019)
μ_1	Expected return on risky asset	7.8%	
σ_1	Volatility of return on risky asset	16.9%	
μ_2	Expected return on home equity	6.4%	
σ_2	Volatility of return on home equity	11.9%	
ρ	Correlation between risks	0.2	Flavin and Yamashita (2002)
r^{κ}	Compound borrowing rate (HEAS)	4.0%	Services Australia (2026a)
Preference parameters			
γ	Relative risk aversion	3.701	De Nardi et al. (2025)
δ	Subjective discount rate	0.03	
$\bar{\xi}_1$	Couple equivalence scale	1.528	
$\bar{\xi}_2$	Bequest curvature (in 000s)	AUD 5,000	
b	Bequest intensity (in 000s)	6,826	
ε	Sensitivity of consumption utility to health	1	Finkelstein et al. (2013)
LTCI parameters			
$\widetilde{\mathrm{LTC}}_2$	Individual gross care costs	AUD 2,100	See Figure 1(a)
$\widetilde{\mathrm{LTC}}_3$		AUD 15,000	
$\widetilde{\mathrm{LTC}}_4$		AUD 55,000	
$\tilde{v}$	Maximum informal care	AUD 2,100	
ϖ	Decline rate of informal care	1/3	Mommaerts (2025)
$P\left(\mathtt{C}_{\mathcal{H}}\left(\mathcal{H}'\right);0\right)$	Joint LTCI price at retirement	AUD 188,925	See Equation (5)
	Single LTCI price at retirement	AUD 138,411	
η	Tax rebate rate	0.3	The Treasury (2025)
ν	Assessable income fraction	0.6	Services Australia (2026b)
Other parameters			
$\bar{\pi}$	Upper bound on risky-asset share: T1	0.4	Van Rooij et al. (2011)
	Upper bound on risky-asset share: T2 & T3	1	Catherine (2022)

Notes: Market parameters are annualised estimates based on historical data from Jordà et al. (2019). A one-step deterioration in health status reduces the marginal-utility weight by 20%, consistent with Finkelstein et al. (2013). We assume that a healthy partner can provide the necessary support for a spouse in State 2. This provision is designed to capture low-intensity informal care historically accessible under the CHSP that cannot be fully covered under new SAH arrangements. The proportional decline rate of informal care ($\varpi = 1/3$) is a structural parameter chosen to reflect the empirical evidence that a spouse's caregiving capacity is strictly bounded by their own physical health, capturing the endogenous breakdown of intra-household support modelled in Mommaerts (2025). We calibrate the tax rebate rate using Australia's private health insurance rebate rate and set the testable income fraction based on the age pension income test treatment of lifetime income streams (Services Australia, 2026b).

about 73% of the joint premium, not 50%(Equation (5)). We calibrate the tax rebate $\eta = 0.3$ to Australia's private health insurance rebate rate (The Treasury, 2025) and set the testable income fraction $\nu = 0.6$ to match the age pension income test treatment of lifetime income streams (Services Australia, 2026b).

Portfolio bounds. We cap the risky-asset share at $\bar{\pi} = 0.4$ for the bottom wealth tertile (Van Rooij et al., 2011) and at $\bar{\pi} = 1$ for the middle and top tertiles (Catherine, 2022), reflecting the more limited equity participation observed among lower-wealth households. We discuss the rationale behind this distinction in detail in Section 6.1.1.

6 Results

A single regularity runs through the analysis below: lower- and middle-wealth (T1 and T2) homeowners behave analogously, while top-wealth (T3) homeowners diverge from both. We label this convergence the middle-wealth trap and document it across four settings of increasing realism: first in a parsimonious baseline that replaces HEAS with frictionless asset diversification to isolate the LTCI–home-care interaction (Sections 6.1.1–6.1.2); then under the HEAS in its institutional form (Section 6.2); next under health and HEAS-timing perturbations (Section 6.3); and as a mechanism in Section 6.4.

6.1 The dual gains of LTCI

LTCI delivers benefits on two fronts: household welfare and public finance. We unpack these in turn.

6.1.1 Household welfare gains

Retirees facing stochastic long-term care shocks confront a tradeoff when they lack access to LTCI and asset diversification between the liquid and illiquid. They can either leave catastrophic health risk uninsured, devastating late-life consumption, or hoard liquid assets as self-insurance. Because most of their wealth is locked in illiquid housing, the latter strategy compresses consumption during the healthy, active years of retirement. Both extremes carry welfare losses, motivating the analysis of financial tools below.

To quantify how effectively households can escape this tradeoff, we measure the welfare recovered through two financial channels: LTCI and frictionless asset diversification. Following Ghili et al. (2024), we define the welfare recovery rate as:

$$\text{Welfare recovery rate} = \frac{V_{\text{Option}} - V_{\text{Autarky}}}{V_{\text{Benchmark}} - V_{\text{Autarky}}} \times 100\%, \tag{14}$$

which benchmarks the option under evaluation (V_{Option}) against two reference points: a no-tools world (V_{Autarky}) and an ideal in which both instruments are available ($V_{\text{Benchmark}}$).

Two features of this measure are:

1. The rate is normalised within each segment. Because the denominator $V_{\text{Benchmark}} - V_{\text{Autarky}}$ depends on the household's wealth, tenure, and demographic position, a high recovery rate indicates that the option closes a large share of the achievable welfare gap, not that the household ends up well off in absolute terms. Cross-segment comparisons therefore speak to how effectively an instrument mitigates risk for a given household type, rather than to how well that household fares overall.

2. A recovery rate near the ceiling admits two interpretations: the option may be uniquely valuable, or the household may simply lack alternative channels for recovering welfare. Which reading applies depends on the household's broader balance sheet, a distinction that becomes important for the renter segment analysed below.

To isolate LTCI's contribution from that of asset diversification, we score two single-tool counterfactuals against the same benchmark. The first is an *LTCI-only* world, in which households can purchase LTCI but cannot rebalance between liquid assets and home equity. The second is a *diversification-only* world, which reverses the restriction. The benchmark permits the simultaneous use of both instruments. All value functions are evaluated at the retirement states detailed in Table 3. We condition on initial health states eligible for LTCI purchase and integrate across all subsequent shocks, with the resulting recovery rates reported in Table 5

Table 5. Welfare recovery from asset diversification and LTCI uptake across household segments.

Household segment	Liquid wealth share (%)			LTCI holding (%)		Welfare recovery (%)	
	Initial	Div.-only	Benchmark	LTCI-only	Benchmark	Div.-only	LTCI-only
Partnered							
T1 Homeowner	25.4	71.9	65.4	60.7	37.8	77.3	36.4
T2 Homeowner	37.6	62.9	61.6	61.4	48.5	72.0	41.1
T3 Homeowner	60.6	60.1	44.7	86.9	97.4	3.6	78.1
T1 Renter	100.0	100.0	100.0	44.3	44.3	0.0	100.0
Single							
T1 Homeowner	25.5	75.4	67.0	55.2	32.1	76.4	33.1
T2 Homeowner	33.4	66.5	62.9	56.8	47.0	72.7	37.1
T3 Homeowner	55.2	57.9	46.8	82.6	94.8	3.8	79.5
T1 Renter	100.0	100.0	100.0	41.6	41.6	0.0	100.0

Notes: Segments are classified by marital status, housing tenure, and initial wealth tertile (T1–T3). "Initial" reflects the liquid wealth share upon entering retirement. "Div.-only" denotes the optimal liquid share under a diversification-only counterfactual, which permits frictionless asset diversification between liquid assets and home equity at retirement, but excludes LTCI. "LTCI-only" indicates the optimal insurance holding when LTCI is available but dynamic portfolio reallocation is restricted. "Benchmark" represents the unconstrained model featuring both dynamic asset reallocation and LTCI availability.

LTCI reshapes the T3 portfolio. T3 homeowners are the only segment whose liquid share declines once LTCI becomes available: T3 partnered households drop from 60.1% (div-only) to 44.7% (benchmark). LTCI takes over the care-risk hedging role previously served by liquid wealth, freeing T3 households to reallocate towards home equity.

Substitution at the bottom, complementarity at the top. The two instruments interact differently across the wealth ladder. For T1 and T2 homeowners, opening diversification cuts optimal LTCI holding by 13–23 percentage points, which indicates diversification substitutes for insurance. For T3 homeowners, opening diversification raises LTCI holding by 10–12 percentage points: the released liquidity becomes premium-paying resources. The two tools are substitutes at the bottom and complements at the top.

Couples gain more from LTCI than singles, and the gap widens at the bottom. Couples carry a risk that singles do not: one spouse may outlive the other and inherit a sharp drop in resources when their own care needs are escalating. LTCI insures against this survivor-state vulnerability on top of insuring against care-cost shocks themselves. The Table 5 numbers bear this out at the bottom of the wealth distribution: T1 partnered homeowners recover 36% of the welfare gap from LTCI alone versus 33% for their single counterparts. In contrast, the gap narrows for wealthier segments where the survivor-state risk is partly self-insured by financial wealth, supported by the empirical findings on wealth and self insurance in Van Houtven et al. (2015).

LTCI as the sole margin for T1 renters. Table 5 shows a 100% recovery rate for T1 renters under LTCI-only. Because these households lack home equity and face steep implicit taxes on liquid self-insurance, LTCI acts as their sole unconstrained margin for welfare recovery. We flag two caveats. First, this relies on a 40% portfolio cap on risky assets; without it, T1 renters would tilt aggressively into equities and rely on public backstops, weakening the case for LTCI. We maintain this cap to reflect their historically limited formal-market participation (Van Rooij et al., 2011).[6] Second, the result assumes frictionless access to public housing and home care, which are practically rationed through eligibility assessments.

In summary, from a demand-side perspective, the welfare value of LTCI diverges sharply across the household segments. For the poorest (T1 renters), who lack home equity and face strict means-testing penalties on liquid savings, LTCI serves as the sole instrument for hedging long-term care risk. Conversely, for the wealthiest (T3 homeowners), LTCI functions primarily as an asset optimisation tool; by absorbing tail-end care risks, it frees up liquidity previously hoarded for self-insurance, enabling more efficient wealth allocation. For the middle segments (T1 and T2 homeowners), however, asset diversification (the ability to flexibly draw upon home equity) acts as a strong substitute, suppressing their demand for LTCI.

6.1.2 Government financial expenditures

The household-side gains from LTCI carry a fiscal counterpart. As insured households self-finance more late-life care, they trigger fewer means-tested transfers, and the government's long-run liability shrinks. We trace this counterpart by projecting cohort-level paths of household wealth and home care expenditure under the means-testing rules across segments (Lyu et al., 2026). For analytical clarity, we consolidate the T1 and T2 segments into a single "non-T3" category, given that their simulated behaviours are nearly identical. This convergence is a direct consequence of the middle-wealth trap discussed previously; a comprehensive justification for this grouping is provided in Section 6.4.

T3 households enter retirement with substantially higher assets; consequently, the government pays them a lower age pension during early retirement compared to the non-T3 category (Figure 6, applicable to both spouse statuses). These households self-finance a larger proportion of their early care (Lockwood, 2018; Ameriks et al., 2020). However, this gap narrows over time, and the age pension and home care subsidy trajectories both intersect for T3 households. This occurs because LTCI premiums reduce assessable resources, enabling insured T3 households to become eligible for the pension sooner than their uninsured counterparts. Exacerbating the government's long-term liabilities, T3 households are also the most likely to reach the lifetime self-contribution cap, at which point late-life care costs revert to public funding. This dynamic

[6] Although their formal financial participation is low, higher gambling prevalence in this group suggests potential risk-seeking behaviour in non-financial domains (Wardle et al., 2024). This remains a suggestive parallel rather than direct evidence of financial risk tolerance.

explains why the gap between T3 and non-T3 households is narrower for home care subsidies than it is for age pension entitlements.

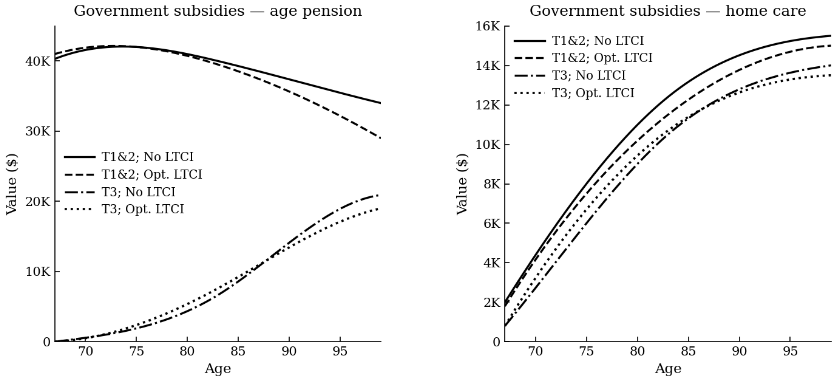

(a) With spouse.

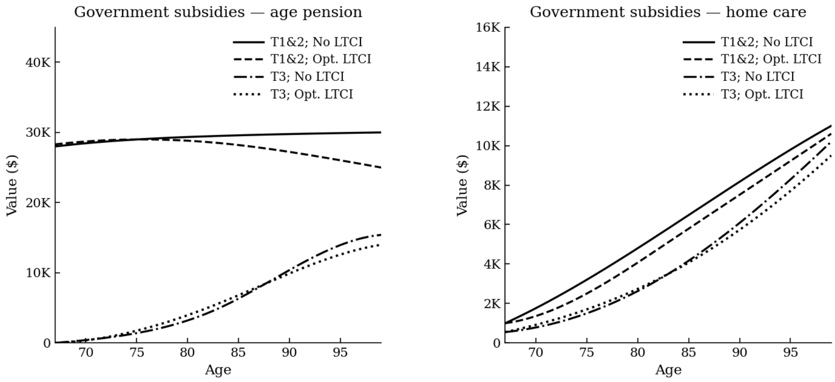

(b) Without spouse.

Figure 6. Cohort-level trajectories of public support under means tests. The figure plots projected age pension entitlements and home care subsidies over retirement for households with a spouse (top) and without a spouse (bottom). For clarity, the three non-T3 cases are pooled within each spouse status.

When these flows are aggregated, Table 6 reports how expected lifetime transfers change when households take up LTCI optimally, and the pattern divides sharply by wealth.

- For lower-wealth segments (T1 and T2 homeowners and T1 renters, both partnered and single), optimal LTCI uptake generates net public savings: total expected transfers fall by 1.7% to 2.7%, with the largest reductions concentrated on the home care subsidy line (approximately -4%).
- For wealthier T3 homeowners, by contrast, the aggregate fiscal effect is essentially neutral, with total transfers shifting by only -0.18% (partnered) and +0.05% (single). This plateau in public liability reflects an interplay between LTCI premium payments, the means tests governing the age pension, and the lifetime cap on home care contributions.
  - On the age pension line, premium outflows erode assessable wealth sufficiently to push some T3 households into the taper zone, raising entitlements by 1.2% (partnered) and 1.7% (single).
  - On the home care line, the T3 reduction (-1.43% for partnered, -2.77% for single) is primarily driven by the income test mechanism. As LTCI indemnities supple-

ment household resources, they reduce the government's required subsidies under the means-testing framework. Ultimately, the broad-based fiscal savings generated by this income test effect outweigh the increased public liabilities associated with late-life tail risks.

The two effects largely offset, leaving aggregate T3 transfers broadly unchanged.

Table 6. Expected lifetime government transfers under optimal LTCI uptake by household segment.

Household segment	Age pension ($)	Home care subsidy ($)	Total transfer ($)
Partnered			
T1 Homeowner	820,469.7 (−2.27%)	179,248.7 (−4.28%)	999,718.4 (−2.64%)
T2 Homeowner	814,462.8 (−1.21%)	177,763.0 (−4.06%)	992,225.9 (−1.73%)
T3 Homeowner	134,305.4 (+1.20%)	144,785.9 (−1.43%)	279,091.3 (−0.18%)
T1 Renter	824,108.4 (−2.31%)	180,057.0 (−4.31%)	1,004,165.4 (−2.67%)
Single			
T1 Homeowner	505,959.8 (−2.10%)	110,550.1 (−3.85%)	616,509.9 (−2.42%)
T2 Homeowner	505,217.7 (−2.18%)	110,291.9 (−3.99%)	615,509.6 (−2.51%)
T3 Homeowner	153,739.3 (+1.70%)	86,173.1 (−2.77%)	239,912.4 (+0.05%)
T1 Renter	506,142.5 (−2.19%)	110,578.5 (−3.10%)	616,720.9 (−2.35%)

Notes: Entries report expected lifetime government payments for a representative cohort within each household segment under optimal LTCI uptake. Values in parentheses denote proportional changes relative to the benchmark scenario without LTCI availability. "Total transfer" represents the sum of age pension entitlements and home care subsidies. These figures do not account for the 30% premium rebate.

6.2 The strategic interaction between home equity release and LTCI

While the previous subsection employed frictionless asset diversification as a theoretical proxy for home equity liquidity, we now relax this assumption by explicitly modelling the HEAS. This introduces two critical departures from the frictionless benchmark:

- Unlike two-way portfolio diversification, the HEAS operates strictly as a unidirectional equity drawdown mechanism;
- We capture the actual policy instrument, explicitly incorporating its compounding borrowing rate and its dynamic, yet constrained, drawdown structure.

The first half of this subsection isolates the HEAS as the exclusive home equity channel to examine its specific impact. The second half then evaluates the HEAS alongside the frictionless diversification benchmark to assess whether they play equivalent roles.

6.2.1 HEAS as the sole liquidity channel

Households face two ways to fund late-life care: insure against the shock with LTCI, or release liquidity from the home through HEAS. The two instruments are partial substitutes, but the substitution rate is far from uniform across the wealth distribution.

To map this, we evaluate all strategy pairs (ψ, ϕ) on a discrete grid, with $\psi \in [0, 1]$ representing the LTCI coverage proportion and $\phi \in [0, 1]$ denoting the HEAS drawdown intensity.

For each pair we solve the household's optimal policies and compute the welfare recovery rate of Section 6.1.1, now taking the no-option pair $(0, 0)$ as the autarky reference and the best-performing (ψ, ϕ) on the grid as the benchmark. A rate of 100% therefore marks the best combination available to that segment, and negative values mark combinations that leave the household worse off than doing nothing. Figure 7 plots the resulting surfaces, with panels (a) and (b) splitting the sample by spouse status.

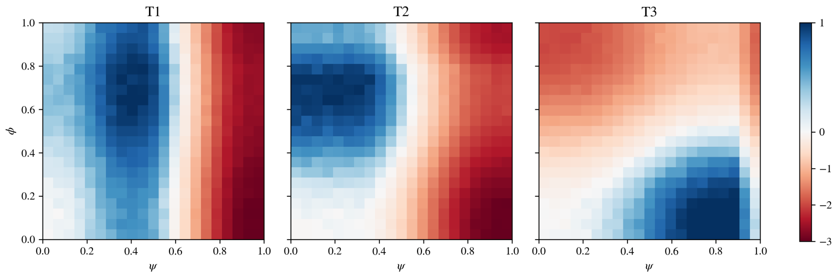

(a) With spouse.

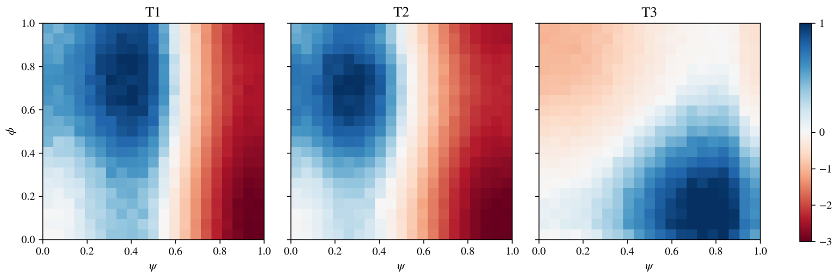

(b) Without spouse.

Figure 7. Normalised welfare gain from joint utilisation of LTCI and HEAS. Each grid point corresponds to a strategy pair (ψ, ϕ), where ψ is the LTCI purchase proportion and ϕ is the HEAS utilisation rate. For each (ψ, ϕ) we solve the household's optimal policy functions and compute welfare relative to the no-option baseline $(0, 0)$; gains are then normalised by the maximum welfare improvement over the grid. Negative values indicate welfare losses relative to $(0, 0)$. Results are shown by household segment (T1–T3) and spouse status.

T1 and T2: nearly identical, save for a middle-wealth signal at T2. Initially, the welfare surfaces for T1 and T2 homeowners appear broadly similar, with both segments concentrating their welfare gains along the LTCI axis and using HEAS to fill liquidity gaps. Yet, we observe a distinct shift in the welfare surface for T2 households with a spouse when ϕ reaches the upper-left range. This shift reveals that for wealthier middle-class cohorts, the demand for LTCI weakens unless complemented by HEAS. It suggests a negative wealth effect on LTCI demand within the middle class, which the availability of HEAS helps alleviate. This phenomenon is closely tied to the middle-wealth trap, which is explored in detail in Sections 6.4 and 7.2. By contrast, singles flatten this T1-T2 disparity: without a spouse to provide informal care, the marginal value of liquidity versus insurance remains structurally aligned across T1 and T2 single homeowners (Mommaerts, 2025).

T3: HEAS adds almost nothing once LTCI is available. The standout pattern sits in the T3 panel. For wealthy homeowners, the welfare surface is essentially flat along the ϕ axis once LTCI is in place. Raising HEAS intensity buys these households almost no extra welfare. The reason is that T3 households rely little on means-tested transfers to begin with (Table 6), so the substitute role HEAS plays for those transfers has very little to substitute. With LTCI absorbing the care-cost shock, T3 households prefer to keep the home intact and use it as a store of wealth rather than a liquidity tap.

Frictionless diversification versus HEAS. How does the institutional HEAS compare against the frictionless diversification benchmark of Section 6.1.1? The wealth-segment ordering carries over, with HEAS delivering measurable welfare for T1 and T2 homeowners but adding almost nothing for T3. Because affluent households strongly favour housing as a wealth repository, they are naturally deterred by the unidirectional equity drawdown of the HEAS. The scheme's compounding borrowing rate further amplifies this unpopularity, ensuring that their home equity remains anchored in place once LTCI absorbs the care-cost shock.

6.2.2 Counterfactual: when HEAS is no longer just a stand-in for retirement diversification

Up to this point we have used optimal HEAS as a reduced-form proxy for the asset diversification at retirement. However, when HEAS and diversification are both available, do they pull in the same direction, or do households actually use them differently?

To answer this, we run a counterfactual in which optimal HEAS is available and recompute welfare recovery from (i) frictionless asset diversification at retirement and (ii) optimal LTCI takeup. The result, plotted in Figure 8, is pronounced. As total wealth increases, the combined benefits of diversification and LTCI diminish because HEAS captures a larger share of the welfare gains. Consequently, the optimal level of retained home equity rises sharply.

The interpretation matters for policy: under this counterfactual scenario, wealthy households would not simply liquidate the home and substitute LTCI for home equity. Instead, their optimal strategy would pivot: they would retain a large home equity position, draw on HEAS dynamically for liquidity as care needs emerge, and still purchase LTCI. By sheltering their wealth in the means-test-exempt home equity, this strategy artificially suppresses their assessable assets, which would translate into an escalation in public transfer liabilities.

Nevertheless, because real-world asset diversification is not frictionless, and households derive great non-financial utility from not moving (Cocco and Lopes, 2020), the extreme theoretical substitution drawn from this counterfactual may be muted in reality, provided that there are no further policy perturbations, such as a relaxation of HEAS borrowing conditions.

6.3 Robustness: Initial health heterogeneity and HEAS timing

Couples do not arrive at retirement equally healthy. While the baseline model assumes an LTCI-eligible health profile at retirement (specifically, individual health states 1 and 2), actual health histories are heterogeneous, leading to different degrees of LTC risk. Moreover, HEAS uptake may be state-dependent, shifting in response to subsequent health shocks. We therefore perform two robustness exercises in tandem: relaxing the initial health assumption via four joint

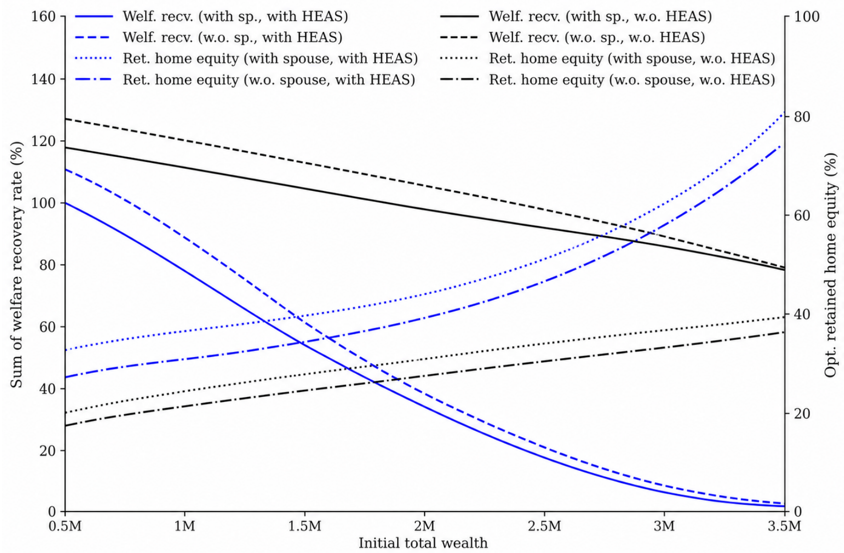

Figure 8. Wealth gradients in option complementarities under optimal HEAS. The figure plots (i) the sum of welfare recovery ratios from diversification at retirement and optimal LTCI take-up, and (ii) the corresponding optimal retained home equity, as functions of total wealth (cash plus home equity). Recovery ratios are evaluated relative to the no-option baseline with HEAS available at its optimal intensity. For comparison, we also report the corresponding no-HEAS cases, highlighting the extent to which optimal HEAS reshapes the recovery profile and the retained housing position. The sharp decline in the combined recovery ratio and the steep rise in retained home equity shows that wealthier households optimally retain a large housing position and use HEAS for liquidity, which would in turn raise potential public exposure.

health configurations and testing the sensitivity of HEAS access timing through three distinct trigger scenarios.

Initial joint health. We examine how an unequal health start affects household strategies through the spousal gap in remaining life expectancy it generates, an interpretation consistent with the risk-equivalent age framework of Fong et al. (2025). We let each spouse start in one of two health states, healthy (1) or mildly impaired (2), giving four initial household configurations: $(1, 1)$ for two healthy spouses, $(1, 2)$ and $(2, 1)$ for one healthy and one mildly impaired (the order identifies which spouse is which), and $(2, 2)$ for two mildly impaired spouses. Throughout this exercise, LTCI pricing remains uniform across all health combinations for couples.

HEAS access scenarios. We also compare three HEAS access scenarios that differ only in when this option becomes available:

- Scenario 1 mirrors the baseline setting: HEAS opens at retirement, so households can use home equity from day one to fund consumption or LTCI premia.
- Scenario 2 locks HEAS until a spouse enters a high-need state (State 3 or 4), turning home equity into a care-triggered liquidity instrument.
- Scenario 3 locks HEAS until widowhood, so the partnered care phase must be financed entirely from cash and LTCI.

Figure 9 plots the optimal (ψ, ϕ) combinations for each initial joint health state across the three HEAS access scenarios, with arrows of distinct styles tracing the transitions between scenarios.

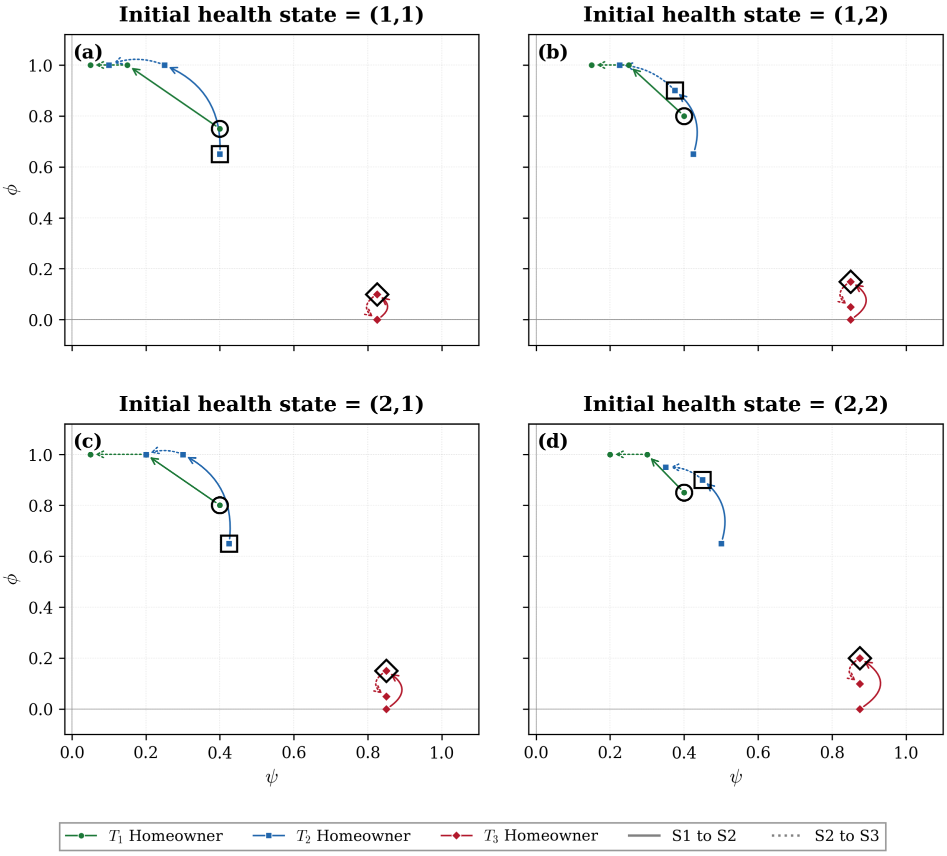

Figure 9. Optimal combinations of LTCI intensity ψ and HEAS take-up ϕ across initial joint health states and HEAS-trigger scenarios. Panels compare the welfare-maximising choices for partnered homeowners starting from $(1,1), (1,2), (2,1)$, and $(2,2)$ across wealth segments (T1–T3). Scenarios differ only in the HEAS access trigger: immediate access after retirement (Scenario 1), access when the spouse enters high-need states (Scenario 2), or access only upon widowhood (Scenario 3). Black markers indicate welfare-dominant scenario within each asset group.

6.3.1 Asymmetric initial health tilts the balance towards HEAS for the poor and LTCI for the rich

A natural question follows: does the wealth-segment ordering for LTCI and HEAS, established in previous sections, hold when households arrive at retirement in unequal health? If so, how? We hold HEAS timing fixed (at Scenario 1, 2, or 3) and let initial joint health vary across the four configurations. That is, we do not examine transitions between scenarios, but observe how the optimal coordinates within each fixed scenario shift across the four health panels. Figure 9 shows that health asymmetry does not blur the wealth-segment patterns; it sharpens them.

Wealthier households buy more LTCI when they start in worse health. LTCI take-up rises with initial illness. Under our pricing structure, a household starting from a worse joint health state has a higher near-term probability of claiming, which raises the actuarial value of coverage. For households whose liquidity position is comfortable enough to afford the premia, this raises optimal LTCI intensity. However, T1 households cannot follow this logic: worse initial health raises the actuarial value of LTCI for them too, but the liquidity-bridging constraint identified above keeps their LTCI position close to flat across health states. They want more coverage; they cannot afford it.

Lower-wealth households substitute towards HEAS when informal care is fragile. T1 households move in the opposite direction on the HEAS margin: poorer initial health raises HEAS take-up. The logic runs through informal care. Fragile spousal health signals an elevated near-term risk that informal care breaks down and formal care costs arrive sooner than expected. Those costs are liquidity-intensive and front-loaded, so the option value of converting home equity into cash through HEAS rises. The same liquidity tightness that blocks T1 from scaling up LTCI thus makes home equity the only margin they can scale up at all. When starting health worsens, T1 households do not buy less of LTCI and HEAS together; they shift mass from the former onto the latter.

6.3.2 Delaying HEAS collapses LTCI demand at the bottom

Holding each initial health configuration fixed, we now examine within Figure 9 how delaying HEAS access reshapes the optimal policy mix. We find the patterns here remain consistent with the wealth-segment ordering established in earlier sections.

A complementarity, not a substitution. Start with two healthy spouses. T1 and T2 homeowners shift their optimal mix across the three scenarios: when HEAS is locked early, optimal LTCI intensity collapses, and the household relies on cash plus delayed HEAS instead. The driver is the liquidity-bridging constraint. Without early HEAS access, liquidity-constrained households cannot fund a large LTCI position out of cash alone; the policy that would be valuable ex post is unaffordable ex ante, so they buy less of it. T3 households reveal the opposite pattern: their LTCI intensity stays high across all three scenarios and HEAS take-up stays low. T3 households finance premia from their own balance sheet and never depend on HEAS for liquidity, so the timing of HEAS access barely matters to them.

T3 reveals what HEAS is actually for at the top of the wealth distribution. The T3 panel shows a non-monotone pattern in ϕ^*: T3 households use HEAS most heavily in Scenario 2 but pull back in Scenario 3. This tells us what HEAS is doing for them. While both spouses are alive, HEAS works as a discretionary liquidity margin that supports higher consumption during the partnered phase. After widowhood, the surviving spouse's consumption needs fall and the relative weight on bequest rises, so balance-sheet preservation takes over from liquidity. HEAS is the right tool for the first phase but the wrong tool for the second.

6.4 The middle-wealth trap

A recurring feature of the numerical results is the close parallel between T1 and T2 homeowners. This subsection explains why.

6.4.1 Definition of the clawback

A household with joint health state h, liquid assets X, and a fixed home value Y that buys LTCI at coverage level ψ pays for that coverage three times over: once through reduced age pension income, once through higher means-tested home care fees, and (when HEAS is available) once again through expanded HEAS borrowing capacity that the lower pension creates. The total clawback is the sum of these three offsets:

$$C(X;h,\psi,Y) = |\underbrace{\Delta I^{\text{AP}}(X;h,\psi)}_{\substack{\text{age pension}\\ \text{reduction}}} - \underbrace{\Delta F(X;h,\psi)}_{\substack{\text{home care}\\ \text{fee increase}}} + \underbrace{\Delta I^{\text{HEAS}}(X;h,\psi,\phi,Y)}_{\substack{\text{HEAS constraint}\\ \text{easing}}}|, \tag{15}$$

where each Δ compares the outcome at coverage $\psi > 0$ against the no-coverage baseline $\psi = 0$. The signs are unambiguous by construction: $\Delta I^{\text{AP}} \leq 0$ (pension income falls when private indemnities raise assessable income), $\Delta F \geq 0$ (home care fees rise on the same assessable income), and $\Delta I^{\text{HEAS}} \geq 0$ (a lower pension expands HEAS borrowing capacity).

6.4.2 Patterns observed

Figure 10 traces C against liquid wealth X for the full grid of LTCI coverage rates $\psi \in \{0.1, 0.2, \ldots, 1.0\}$ and the nine symmetric joint health states with positive LTCI indemnities. Panel (a) sets $\phi = 0$ (no HEAS), so the third term in Equation (15) drops out. Panel (b) sets $\phi = 1$, letting the household borrow up to the instantaneous HEAS cap with no lifetime accumulation limit. For the purposes of this visualisation, we abstract from the maximum withdrawal cap tied to home equity. This omission is methodologically justified, as households with limited home equity reach an interior solution in their utilisation of the HEAS. Panels (a) and (b) therefore bracket reality, and within those brackets C reduces to $C(X;h,\psi)$ in X alone.

The clawback architecture exhibits three defining characteristics. First, it follows a hump-shaped trajectory relative to liquid wealth: it remains near zero for households below the income test free area, climbs steadily through the income test segment, peaks where the income test transitions to the assets test, and drops back to zero once the asset test threshold eliminates the pension entirely (Appendix D provides closed-form expressions for each kink). Second, the magnitude of this peak scales with health severity. Households anticipating severe joint care states face a peak clawback several times higher than those in mild states, meaning that the middle-wealth band loses a greater proportion of each LTCI indemnity as prospective care needs intensify. Finally, LTCI flows through cleanly at both extremes of the wealth distribution; below the income test free area, the means test does not bind, whilst above the asset test cutoff, government subsidies are already exhausted, leaving nothing further to claw back.

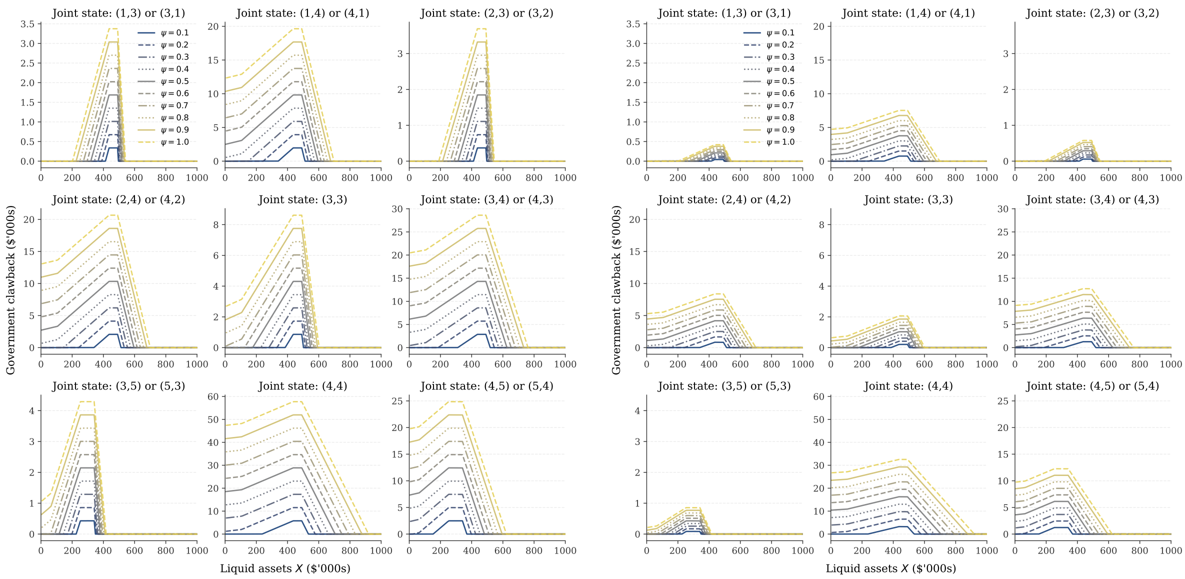

(a) Without HEAS **(b)** With HEAS

Figure 10. Total clawback of LTCI indemnities versus liquid assets X by joint health state and LTCI coverage level $\psi \in \{0.1, \ldots, 1.0\}$. Panel (a) presents the scenario without HEAS; Panel (b) incorporates the HEAS at a 1.5× maximum pension cap, excluding maximum loan amount constraints. Closed form expressions for each kink are given in Appendix D.

6.4.3 Why the middle-wealth band is dynamically unstable

Together these three patterns explain the middle-wealth trap. Households in the middle-wealth band sit on the steep flank of a non-monotone payoff: a small positive shock pushes them up the clawback peak, where the marginal value of saving collapses, and a small negative shock can push them across the peak entirely. The optimal life-cycle response is to leave the band rather than stay in it. The simulation captures this through the induced curvature in the value function, which empties the middle of the wealth distribution as households drift towards either the full-pensioner plateau on the left or the asset-test-exempt region on the right.

The trap is harsher for households below or near the peak, because buying LTCI exacerbates their financial penalty. The upfront premium drains liquid assets, moving the household further into the income test band where the next dollar of LTCI indemnity will be heavily taxed. The household pays a certain premium today for an uncertain, partially clawed-back indemnity tomorrow. The asymmetry is severe enough that lower- and middle-wealth households optimally choose minimal LTCI coverage, even when the actuarially fair premium-to-expected-benefit ratio would justify more. High-wealth households face the mirror image: their initial assets keep them safely on the right-most flat segment of the clawback curve, the premium does not drag them into the binding region, and indemnities flow through largely free of means-tested erosion.

6.4.4 A familiar shape from different policy domains

Table 7 situates our contribution within the international literature on means-testing, implicit taxation, and private insurance demand, complementing Table 1.

The clawback we derive shares the structural form of these implicit-tax penalties: a non-convex distortion concentrating a marginal penalty on a specific wealth band, with a hump-

Table 7. International literature on means-tested implicit taxation.

Study	Setting	Core mechanism	Limitation addressed
Hubbard et al. (1995)	US; transfers	Asset tested floor imposes a near-confiscatory implicit wealth tax at the threshold.	Single asset test; no role for private insurance or LTC risk.
Brown and Finkelstein (2008)	US; Medicaid	Insurance-first rule crowds out LTCI demand for the bottom two-thirds.	Aggregate focus; misses non-monotonic geometry of implicit tax.
Lockwood (2018)	US; Medicaid	Bequest motives reduce self-insurance opportunity costs for the wealthy.	Isolates single constraint; ignores nested policy interactions.
Reichenstein and Meyer (2018)	US; Soc. Sec.	"Tax torpedo": hump-shaped marginal tax profile for middle-income retirees.	Pure tax-planning focus; no LTC or private insurance engagement.
Sefton et al. (2008)	UK; Pension	High withdrawal rates suppress saving among middle earners (savings trap).	Single means test; no LTC or private insurance engagement.

Notes: Our paper focuses on Australia's age pension, home care fees, and HEAS. We demonstrate that nested means-testing generates a hump-shaped clawback on private LTCI indemnities, hollowing out the middle-wealth distribution. Unlike prior studies, our dynamic framework integrates multiple policy layers and identifies HEAS as a complementary liquidity channel.

shaped profile structurally analogous to the US tax torpedo and a middle-class hollowing-out parallel to the UK means-tested savings trap. The institutional details, however, differ. In the US, the implicit tax on private LTCI arises from Medicaid's explicit insurance-first rule, requiring private coverage to be exhausted before public assistance engages. In Australia, no equivalent rule exists; the same penalty emerges from treating LTCI indemnities as assessable lifetime income streams, which propagate through the income test for age pension and home care. The structural form of the defined clawback is thus robust to institutional form: it can be triggered by direct insurance-coordination rules or by the income test of insurance indemnities, with equivalent distributional consequences.

7 Policy implications

Our findings identify two structural distortions that demand policy response: an implicit-tax clawback that hollows out the middle of the wealth distribution, and a HEAS-LTCI complementarity that opens strategic loopholes at the top. The two subsections below address each in turn.

7.1 Flatten the clawback peak that hollows out the middle

Section 6.4 traced the middle-wealth trap to a specific feature of the means-tested architecture: the clawback on private LTCI indemnities peaks sharply where the income test gives way to the assets test, and households on that flank face value-function curvature steep enough that the optimal life-cycle response is to leave the band rather than save through it. Smoothing this clawback peak eliminates the trap, aligning household behaviour with classical models that abstract from the means testing architecture and treat public safety nets as a simple consumption floor (Shao et al., 2019; Xu et al., 2023). Two reforms address this directly:

- First, the government can flatten the peak by decoupling LTCI indemnities from the income test, assessing the maximum of deemed and LTCI income rather than their sum. This change removes the double-counting that drives the climb through the income test segment, without affecting households who sit safely below the free area or above the asset test cutoff.

- Second, revenue lost from the flatter peak can be recovered from the segment that currently escapes the clawback altogether. Wealthy households whose assets push them past the asset test cutoff receive their LTCI indemnities free of any means-tested erosion. A means-tested LTCI levy or a sliding-scale premium subsidy targeted at this segment would recover revenue from households for whom the private-public substitution is otherwise one-way.

7.2 Expanding LTCI, recalibrating HEAS

Panel (b) of Figure 10 demonstrates that HEAS structurally mitigates the means-testing clawback that typically erodes the value of private LTCI. This mitigating effect arises because the scheme is exempt from the income test and scales positively with household means. In isolation, however, the HEAS remains unattractive: it entails high costs and reduces bequests. Furthermore, without extra means-tested income sources such as LTCI indemnities, low assessable means limit the equity that lower-wealth households can draw.

From a fiscal perspective, our baseline results position LTCI expansion as a prudent policy lever. Provided HEAS borrowing constraints remain tight, our simulations indicate that broader LTCI coverage imposes only modest additional public costs. However, Section 6.2.2 identifies a strategic vulnerability: if HEAS constraints are relaxed, wealthier households can exploit this policy interaction to engineer an "asset-rich, cash-poor" profile. This deliberate suppression of assessable financial wealth bypasses means-testing penalties and accelerates access to public subsidies, shifting the fiscal burden directly back onto the government budget.

To insulate the public purse against this strategic behaviour, we recommend complementing any LTCI expansion with tightened HEAS borrowing caps for higher-wealth households, while explicitly preserving or easing access for middle- and lower-wealth cohorts. This targeted calibration upholds the redistributive integrity of the means test, ensuring HEAS continues to serve as a vital liquidity bridge for those who genuinely need it, rather than an optimisation tool for the affluent.

8 Conclusion

This paper writes the cross-dependencies of LTCI, HEAS, and means-tested public support directly into household optimisation, and recovers their joint consequences as endogenous objects. Three patterns characterise these consequences. First, HEAS acts as a liquidity channel that makes LTCI premia affordable for lower-wealth households, but the compounding borrowing rate and bequest motives progressively close the bridge as wealth rises; wealthy households consequently anchor home equity in place once LTCI absorbs the care-cost shock. Second, while LTCI uptake generates net public savings at the bottom of the wealth distribution, the aggregate fiscal impact at the top remains modest. Third, overlapping means tests impose a hump-shaped clawback on LTCI indemnities that peaks at the income-to-asset test transition and scales with health severity; the resulting non-convexity makes the middle band dynamically unstable, hollowing out the middle through endogenous portfolio response rather than preference heterogeneity.

Methodologically, three elements underpin these findings. First, nested means-testing rules with cross-dependent eligibility are integrated into a life-cycle system, with implicit-tax wedges recovered as endogenous objects rather than imposed reduced-form taxes. Second, the proportional-odds construction aligns health states with eligibility thresholds rather than clinical labels, making the joint-life risk model directly transferable to other systems with comparable rule structures. Third, treating the public balance sheet as an endogenous output of the same state-space that produces household policy exposes channels that neither single-agent household models nor reduced-form fiscal projections can see.

These results rest on a counterfactual contract: Australia has no private LTCI market, so the policy is not calibrated to an existing product but constructed to be institutionally disciplined, with a single premium at retirement, no loading, actuarially fair pricing, and a statutory treatment inferred by analogy with the private health insurance rebate and the age pension treatment of lifetime income streams. The findings should therefore be read as the behaviour and fiscal incidence a fully optimising retiree would generate under this stylised contract, not as predictions for an operating market; a market with loading, lapse, or periodic premiums would weaken the household gains and compress the fiscal divide, though the clawback geometry, being a property of the means-testing architecture rather than the contract, would survive. Comparable caveats attach to the newly legislated Support at Home calibration, the imputation of health states from HILDA conditions, and the assumption that both spouses retire simultaneously.

The framework focuses on home care. Natural extensions are: endogenising the transition to residential care to quantify the bequest-versus-quality tradeoff; comparative-statics analysis across alternative pension-test architectures to benchmark the middle-wealth trap cross-nationally; and the design of mortality-linked risk-pooling products that operate efficiently within the means-testing architecture, letting middle-wealth households save through the trap rather than around it. The message is conditional: private LTCI is fiscally defensible only if the surrounding means tests are reformed in tandem, and because the clawback is generated by nested eligibility rules rather than anything unique to Australia, the lesson generalises.

Appendix A Stochastic differential representation for a continuous Markov chain

The following proof is based on Skorokhod (2009, pp. 104–105). For every ordered pair (h, h') with $h \neq h'$ define an independent non-homogeneous Poisson process $N^{hh'}(t)$ of rate $q_{hh'}(t)$. At each calendar time t construct disjoint half-open intervals

$$\Gamma_{hh'}(t) = \left[\sum_{j<h'} q_{hj}(t), \sum_{j\le h'} q_{hj}(t) \right), \quad h \neq h', \tag{16}$$

where the pairs (i, j) are ordered lexicographically. Each interval has length

$$|\Gamma_{hh'}(t)| = q_{hh'}(t),$$

and the collection $\{\Gamma_{hh'}(t)\}_{h\neq h'}$ partitions $\mathbb{R}$.

Let $y \in \mathbb{R}$ be the mark (spatial coordinate) of a point of a unit-intensity Poisson random measure $\upsilon(dy, dt)$. Define

$$\hbar(h, y) = \sum_{h'\neq h} (h' - h)\, \mathbb{1}_{\Gamma_{hh'}(t)}(y), \quad (h, y) \in \mathscr{H} \times \mathbb{R}. \tag{17}$$

When a Poisson point (y, t) falls inside $\Gamma_{hh'}(t)$ and the individual is currently in state h, the indicator equals one and $\hbar(h, y) = h' - h$; otherwise $\hbar(h, y) = 0$. Thus the auxiliary variable y selects the ordered pair (h, h') corresponding to the jump that occurs at time t. The health process $H(t)$ evolves according to

$$dH(t) = \int_{\mathbb{R}} \hbar\left(H(t^-), y\right) \upsilon(dy, dt), \tag{18}$$

because υ places at most one point in $(t, t + dt]$, the integral reduces to the familiar sum of Poisson counts:

$$dH(t) = \sum_{h'\neq H(t^-)} (h' - H_{t^-})\, dN^{H(t^-)h'}(t), \quad N^{hh'}(t) = \upsilon\left((0, t] \times \Gamma_{hh'}(\cdot)\right). \tag{19}$$

Appendix B Numerical solution method: policy iteration on a sparse grid

This appendix describes how we solve the household's continuous-time HJB system numerically (Equation (12)). The system combines discrete health transitions with multiple continuous state variables, and three features make a textbook approach unworkable: financial states diffuse and correlate with each other, means-tested transfers and institutional rules introduce kinks in the payoff and continuation value, and a generator with death intensities switches between regimes. We describe the discretisation, then the four numerical ingredients we use to keep the procedure stable, and close with the policy-iteration algorithm itself.

B.1 Discretisation

The state vector splits into two parts. Liquid wealth and housing wealth diffuse on a uniform grid in logarithmic coordinates, where log-space discretisation suits the multiplicative shock structure of the financial returns. Discrete health states sit alongside, together with slowly moving or path-dependent variables (such as the reverse-mortgage balance and expenditure accumulators), which we handle on a sparse Smolyak grid to avoid the curse of dimensionality. The control vector reduces to consumption and the risky-asset share, both subject to feasibility constraints we list below.

B.2 Four numerical ingredients

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

- *Policy iteration.* Policy iteration converts the nonlinear HJB into a sequence of linear problems. We alternate between (a) improving controls via discrete Hamiltonian maximisation, initialised by the interior first-order conditions whenever they are informative, and (b) evaluating the value function under the current policy until convergence at each time step.

- *Craig-Sneyd ADI in log-space.* Conditional on a policy, the two-dimensional diffusion block in (x, y) dominates the computation. We solve it with an accelerated Craig-Sneyd ADI scheme in log-coordinates (ξ, η). 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. We use standard upwinding in (ξ, η) based on the sign of the effective log drifts (Hanson, 2007), and monotone sparse-grid upwind operators for the auxiliary states on the Smolyak grid.

- *Vanishing-viscosity regularisation.* Means-tested transfers introduce kinks where a textbook policy-improvement step can stall or oscillate. We add $\varsigma(V_{xx} + V_{yy})$ as a vanishing-viscosity term (Yong and Zhou, 1999), with $\varsigma = \mathcal{O}(\Delta x^2 + \Delta y^2)$ so the perturbation disappears as the grid refines (Barles and Souganidis, 1991). This keeps the iteration stable near kinks and steers convergence towards the viscosity solution.

B.3 Household strategies and control updates

Within each policy-improvement step, we update consumption and the risky share separately because they have different concavity structure.

Interior candidate controls. Under the standard regularity conditions (local smoothness of the value function and local concavity of the Hamiltonian), pointwise maximisation of the HJB delivers two interior candidate rules. The consumption first-order condition gives

$$c^{\mathrm{FOC}}(t) = \left(\frac{\xi_1 (h)^{1-\gamma}}{V_x(t)} \right)^{1/\gamma}, \tag{20}$$

and the risky-share first-order condition gives

$$\pi^{\mathrm{FOC}}(t) = -\frac{(\mu_1 - r)V_x(t)}{\sigma_1^2 x V_{xx}(t)} - \frac{\rho\sigma_2 y}{\sigma_1 x} \frac{V_{xy}(t)}{V_{xx}(t)}. \tag{21}$$

The first term in π^{FOC} is the familiar myopic demand. The second captures hedging demand against shocks to the non-traded home-equity component. Both candidates rely on local smoothness; we use them whenever the underlying value-function estimates are well-behaved, and fall back to direct boundary search when they are not.

Consumption update. The Hamiltonian is strictly concave in c, and the admissible set is a closed interval $\mathcal{C}$ pinned down by the subsistence floor and liquidity constraints. The global maximiser is therefore the projected interior candidate $\tilde{c} = \Pi_{\mathcal{C}}(c^{\text{FOC}})$. For numerical robustness we also evaluate the discrete Hamiltonian at the endpoints of $\mathcal{C}$ and keep whichever value is largest. This procedure delivers corner solutions automatically in kinked regions (Fella, 2014).

Portfolio update. The risky share π enters the Hamiltonian through diffusion terms and cross-derivatives, so the discrete objective in π is only locally quadratic. Its concavity depends on the sign and conditioning of the discretised curvature terms (V_{xx} and V_{xy}). Near kinks these quantities can be noisy, sign-indefinite, or close to zero, which makes the closed-form rule in Equation (21) numerically unstable: the denominator V_{xx} becomes ill-conditioned and the formula blows up. We therefore update π by direct maximisation of the one-dimensional discrete Hamiltonian over the bounded admissible set $[0, \bar{\pi}]$. We evaluate the boundary points and include the interior candidate only when the implied discrete quadratic is well-behaved and concave.[7]

B.4 Policy-iteration routine

Algorithm 1 summarises the full routine. For each time level we warm-start from the next time step, pre-factorise the health-transition linear system (reused across all spatial nodes), and iterate on the policy until the sup-norm change in the value function falls below a time-dependent tolerance. The closed-form interior candidates serve double duty: as initial guesses that accelerate convergence, and as diagnostics for when the discrete maximisation has departed from the local-smoothness regime.

Algorithm 1. Policy iteration on a sparse grid with accelerated Craig–Sneyd ADI and vanishing-viscosity regularisation.

Grids:
Uniform logarithmic grids (ξ_u, η_v) with $\xi_u = \xi_{\min} + u\Delta\xi$, $\eta_v = \eta_{\min} + v\Delta\eta$. ▷ Physical nodes are $x_i = e^{\xi_u}$ and $y_j = e^{\eta_v}$.
Discrete health states $h \in \mathcal{H}$.
Sparse–Smolyak grid $\mathcal{S}_L$ in $(\kappa, \zeta^{\text{ind}}, \zeta^{\text{sp}})$, nodes $(\kappa_p, \zeta_p^{\text{ind}}, \zeta_p^{\text{sp}})$, $p = 1, \ldots, N_{\text{SG}}$.
Store the value function as $V^m[i, j, h, p]$ at time level t_m over (x_i, y_j, h, p).
Control sets: $c \in [\underline{c}(x_i, \cdot), \bar{c}(x_i, \cdot)]$, $\pi \in \Pi = [\underline{\pi}, \bar{\pi}]$.
for $m = N_T - 1, \ldots, 0$ **do** ▷ Backward induction in time.
    $V^m \leftarrow V^{m+1}$ ▷ Warm start at t_m.

[7] For low-wealth non-homeowners we impose a tighter admissible set on the portfolio share, motivated by empirical evidence of limited risky-asset participation. This reduced-form constraint captures market-participation frictions faced by this subgroup, including lower financial literacy and higher relative participation and decision costs (Van Rooij et al., 2011).

Define alive-state generator $Q^S(t_m)$ and death intensities $q_{h,\dagger}(t_m)$.
Precompute LU of the health-update matrix

$$M_h(t_m) = (1 + \Delta t\,\delta)\mathrm{Id}_h - \Delta t\,(Q^S(t_m))^\top.$$

▷ Reused for all (p, i, j) at time t_m.

repeat ▷ Policy iteration at fixed t_m

$V^{m,\mathrm{old}} \leftarrow V^m$.

for all $p \in \{1, \ldots, N_{\mathrm{SG}}\}$ **do**

(A) CS-ADI update in (ξ, η) **for each health state** h

for all $h \in \mathscr{H}$ **do**

(1) Derivatives on (ξ, η) **from** $V^{m,\mathrm{old}}[h, p, :, :]$

Compute $V_\xi, V_\eta, V_{\xi\xi}, V_{\eta\eta}$ (central) and $V_{\xi\eta}$ (standard 9-point cross stencil).

Recover physical derivatives:

$$V_x = \tfrac{1}{x} V_\xi, \quad V_{xx} = \tfrac{1}{x^2}(V_{\xi\xi} - V_\xi), \quad V_y = \tfrac{1}{y} V_\eta, \quad V_{yy} = \tfrac{1}{y^2}(V_{\eta\eta} - V_\eta), \quad V_{xy} = \tfrac{1}{xy} V_{\xi\eta}.$$

(2) Candidate controls from FOCs (diagnostics / initial guesses)

Compute $(c^{\mathrm{FOC}}, \pi^{\mathrm{FOC}})$ from Equations (20)–(21) using local derivatives.

Project $(c^{\mathrm{FOC}}, \pi^{\mathrm{FOC}})$ onto admissible sets.

(3) Vanishing viscosity

Set $\varepsilon = C_\varepsilon(\Delta\xi^2 + \Delta\eta^2)$.

(4) Policy improvement by discrete Hamiltonian maximisation ▷ Maximisation uses derivatives evaluated from $V^{m,\mathrm{old}}$.

Choose $c^\star \in [\underline{c}, \bar{c}]$ that maximises the discrete Hamiltonian w.r.t. c, initialised at c^{FOC}.

Choose $\pi^\star \in [\underline{\pi}, \bar{\pi}]$ that maximises the discrete Hamiltonian w.r.t. π, initialised at π^{FOC}.

(5) Drifts and diffusions (log-space form)

Compute physical drifts:

$$a_x = r x_i + \pi^\star x_i(\mu_1 - r) - c^\star + I(\cdot) - F(\cdot), \qquad a_y = \mu_2 y_j.$$

Diffusions (with vanishing viscosity):

$$\nu_x = (\pi^\star \sigma_1)^2 + 2\varepsilon, \qquad \nu_y = \sigma_2^2 + 2\varepsilon, \qquad \nu_{xy} = \rho\sigma_1\sigma_2\,\pi^\star.$$

Effective log-space convection coefficients:

$$\tilde{a}_\xi = \frac{a_x}{x_i} - \frac{1}{2}\nu_x, \qquad \tilde{a}_\eta = \frac{a_y}{y_j} - \frac{1}{2}\nu_y.$$

(6) Operators in (ξ, η)

Build diffusion operators (implicit, central):

$$T_x = \tfrac{1}{2}\nu_x\, D_{\xi\xi}, \qquad T_y = \tfrac{1}{2}\nu_y\, D_{\eta\eta}, \qquad (CV) = \nu_{xy}\, V_{\xi\eta}.$$

▷ Mixed derivative uses a symmetric cross stencil; boundaries use one-sided variants.

Compute upwind first derivatives $V_\xi^{\mathrm{up}}, V_\eta^{\mathrm{up}}$ based on the signs of $\tilde{a}_\xi, \tilde{a}_\eta$.

(7) Advections in κ and ζ (sparse-grid upwind, monotone)

Compute sparse-grid upwind derivatives using monotone operators:

$$b_\kappa = r^\kappa \kappa_p + I^{\mathrm{HELOC}}(\cdot), \qquad V_\kappa \approx \mathscr{D}_{\mathrm{up}}^{\kappa,\mathrm{SG}} V^{m,\mathrm{old}},$$

$$b_{\zeta^{\mathrm{ind}}} = \mathbb{1}_{\{\zeta_p^{\mathrm{ind}} < \bar{\zeta}\}}\big(F_{\mathrm{HC}}^{\mathrm{ind}}(\cdot) + F_{\mathrm{RC}}^{\mathrm{ind}}(\cdot)\big), \qquad V_{\zeta^{\mathrm{ind}}} \approx \mathscr{D}_{\mathrm{up}}^{\zeta^{\mathrm{ind}},\mathrm{SG}} V^{m,\mathrm{old}},$$

$$b_{\zeta^{\mathrm{sp}}} = \mathbb{1}_{\{\zeta_p^{\mathrm{sp}} < \bar{\zeta}\}}\big(F_{\mathrm{HC}}^{\mathrm{sp}}(\cdot) + F_{\mathrm{RC}}^{\mathrm{sp}}(\cdot)\big), \qquad V_{\zeta^{\mathrm{sp}}} \approx \mathscr{D}_{\mathrm{up}}^{\zeta^{\mathrm{sp}},\mathrm{SG}} V^{m,\mathrm{old}}.$$

(8) Bequest at death

$$B(x_i, y_j, \kappa_p) = U_2\big(x_i + (y_j - \kappa_p)^+,\ h,\ t_m\big).$$

(9) Local source term (explicit first-order terms)

$$G = U_1(c^\star, h) + \tilde{a}_\xi\, V_\xi^{\mathrm{up}} + \tilde{a}_\eta\, V_\eta^{\mathrm{up}} + b_\kappa V_\kappa + b_{\zeta^{\mathrm{ind}}} V_{\zeta^{\mathrm{ind}}} + b_{\zeta^{\mathrm{sp}}} V_{\zeta^{\mathrm{sp}}}.$$

▷ Discounting and health transitions are handled implicitly in Step (B) below.

(10) Craig–Sneyd ADI in (ξ, η) with $\vartheta = \frac{1}{2}$

Mixed term uses previous policy iterate:

$$R^{(0)} = V^{m+1}[h, p, :, :] + \Delta t\big(G + CV^{m,\mathrm{old}}[h, p, :, :]\big).$$

Form (and factorise) the tridiagonal matrices for the current policy iterate:

$$A_x = \mathrm{Id} - \vartheta \Delta t\, T_x, \qquad A_y = \mathrm{Id} - \vartheta \Delta t\, T_y.$$

▷ LU factorisations are cached per (h, p) and recomputed when $(c^\star, \pi^\star)$ changes.

x-implicit step:

$$A_x\, \widetilde{V} = \big(\mathrm{Id} + (1-\vartheta)\Delta t\, T_x\big) V^{m+1}[h, p, :, :] + \Delta t\, R^{(0)}.$$

y-implicit step:

$$A_y\, \widehat{V} = \big(\mathrm{Id} + (1-\vartheta)\Delta t\, T_y\big) V^{m+1}[h, p, :, :] + \Delta t\big(\widetilde{V} - T_y V^{m+1}[h, p, :, :]\big).$$

Craig–Sneyd correction:

$$\widetilde{V}^{\,\mathrm{CS}}[h, p, :, :] = \widehat{V} + \big(\tfrac{1}{2} - \vartheta\big)\Delta t\, C\Big(\widehat{V} - V^{m+1}[h, p, :, :]\Big).$$

end for ▷ end loop over h

(B) Implicit update in health state h (coupled system)
For each (i,j), solve the $|\mathcal{H}| \times |\mathcal{H}|$ system:

$$M_h(t_m)\, V^m[:,p,i,j] = \widetilde{V}^{\,\mathrm{CS}}[:,p,i,j] + \Delta t\, q_{\cdot,\dagger}(t_m)\, B[:,p,i,j].$$

$\triangleright$ Same LU of $M_h(t_m)$ is reused for all (p,i,j).
end for $\triangleright$ end loop over p
until $\|V^m - V^{m,\mathrm{old}}\|_\infty < \epsilon_{\mathrm{pol}}^{(m)}$
end for $\triangleright$ end loop over m

Appendix C HILDA data

This appendix provides detailed information on the variables extracted from the Household, Income and Labour Dynamics in Australia (HILDA) survey used in our empirical analysis. Specifically, we utilise twenty data waves spanning from 2004 to 2023 to capture respondents' general information, long-term health conditions, and financial characteristics.

Appendix D Details of the clawback function

To state the kink locations in closed form, we introduce the following means-testing parameters, all expressed on a per-capita basis consistent with the age pension schedule. Let θ denote the deeming threshold, and $r_L < r_U$ the lower and upper deeming rates. The per-capita deeming function is $d(\bar{X}) = r_L \min(\bar{X},\theta) + r_U \max(\bar{X}-\theta,0)$, with per-capita assets $\bar{X} = X/\,(1+\mathbb{1}_{\mathrm{couple}})$. The per-capita assessable LTCI income is $e = \nu I^{\mathrm{LTCI}}/\,(1+\mathbb{1}_{\mathrm{couple}})$. For the income test, let $\mathcal{F}_I$ and $\mathcal{C}_I$ denote the lower and upper taper thresholds, respectively; for the asset test, let $\mathcal{F}_A$ and $\mathcal{C}_A$ denote the corresponding thresholds. No reduction applies when assessable income (or assets) falls below the lower threshold; the pension is fully extinguished at the upper threshold. The per-capita income test and asset test reductions are

$$\mathcal{R}_I(\bar{X},e) = \frac{\bar{M}}{\mathcal{C}_I - \mathcal{F}_I} \max\left(d(\bar{X}) + e - \mathcal{F}_I, 0\right), \quad \mathcal{R}_A(\bar{X}) = \frac{\bar{M}}{\mathcal{C}_A - \mathcal{F}_A} \max\left(\bar{X} - \mathcal{F}_A, 0\right),$$

where $\bar{M}$ is the maximum pension rate. The pension is determined by whichever test is more stringent: $I^{\mathrm{AP}} = (1+\mathbb{1}_{\mathrm{couple}})\max\left(\bar{M} - \max\left(\mathcal{R}_I,\mathcal{R}_A\right),0\right)$. When needed, the inverse deeming function is

$$d^{-1}(y) = \begin{cases} y/r_L, & y \le r_L\theta, \\ \theta + \left(y - r_L\theta\right)/r_U, & y > r_L\theta. \end{cases}$$

and the kinks of the clawback function occur at the following household-level liquid asset values:

K1. Income test activation under ψ. The LTCI indemnity pushes total assessable income past $\mathcal{F}_I$. When $e > \mathcal{F}_I$ this occurs at $X = 0$; otherwise

$$X = (1+\mathbb{1}_{\mathrm{couple}})\, d^{-1}\left(\mathcal{F}_I - e\right).$$

Table 8. HILDA variables used and their collection frequency.

Variable name	Description	Frequency
hgyob	Year of birth	Fixed
chkms	Check marital status	Annual
hstenr	Own, rent or live rent free	Annual
hsbedrm	Number of bedrooms	Annual
xpfoodi	Household nondurable groceries	Annual
xposmli	Household meals eaten outside	Annual
hifapti	Household public transfers	Annual
hifpiip	Household regular income (+)	Annual
hifpiin	Household regular income (−)	Annual
helth	Long term health condition	Annual
hespnc	Sight problems not corrected	Annual
hehear	Hearing problems	Annual
hespch	Speech problems	Annual
hebflc	Blackouts, fits or loss of consciousness	Annual
heslu	Difficulty learning or understanding things	Annual
heluaf	Limited use of arms or fingers	Annual
hedgt	Difficulty gripping things	Annual
helufl	Limited use of feet or legs	Annual
henec	A nervous or emotional condition	Annual
hedisf	Disfigurement or deformity	Annual
hemirh	Mental illness	Annual
hesbdb	Shortness of breath or difficulty breathing	Annual
hecrp	Chronic or recurring pain	Annual
hehibd	Result of a head injury, stroke or other brain damage	Annual
hemed	Restrictive long-term condition or ailment though being treated	Annual
heoth	Other long-term conditions	Annual
hwassei	Household total assets	Quadrennial
hwdebti	Household total debt	Quadrennial
hwhmvai	Home: Apportioned value	Quadrennial

Notes: This table summarises the HILDA variables used in the study dataset and shows how often each variable is collected. Descriptions have been edited for clarity, but the original variable codes are unchanged. 'Fixed' variables are recorded once and do not vary across waves. 'Annual' variables are updated in every survey wave, whereas 'Quadrennial' variables are collected once every four waves. We apply a cubic smoothing spline to each individual's quadrennial wealth observations to interpolate the intervening annual values.

K2. Deeming break-point. The deeming rate switches from r_L to r_U:

$$X = (1 + \mathbb{1}_{\text{couple}})\, \theta.$$

K3. Income test activation under baseline. The same threshold as **K1** evaluated at $e = 0$:

$$X = (1 + \mathbb{1}_{\text{couple}})\, d^{-1}(\mathcal{I}_I)\,.$$

K4. Income-versus-assets crossover under ψ. The asset level at which $\mathcal{R}_I(\bar{X}, e) = \mathcal{R}_A(\bar{X})$, i.e.

$$\frac{d(\bar{X}) + e - \mathcal{F}_I}{\mathcal{C}_I - \mathcal{F}_I} = \frac{\bar{X} - \mathcal{F}_A}{\mathcal{C}_A - \mathcal{F}_A},$$

solved piecewise over the two segments of $d(\cdot)$ and converted to household level via $X = (1 + \mathbb{1}_{\text{couple}})\,\bar{X}$. Above this threshold, the asset test binds and the LTCI indemnity becomes infra-marginal for the age pension.

K5. Income-versus-assets crossover under baseline. The same equation as **K4** evaluated at $e = 0$.

K6. Pension extinction. The pension is fully tapered to zero under whichever test binds first:

$$X = (1 + \mathbb{1}_{\text{couple}}) \min\Big(\mathcal{C}_A, d^{-1}(\mathcal{C}_I - e)\Big),$$

and analogously at $e = 0$ for the baseline. Beyond this point $\Delta I^{\text{AP}} = 0$, so no further pension-channel clawback is possible.

Not all six kinks produce a visible slope change in every configuration: when **K1** falls at $X = 0$, it is absent from the figure; when **K4** and **K5** lie close together, they may appear as a single kink at the resolution of the plot. Because the SaH fee contribution rate is an affine function of the pension-to-maximum-pension ratio $I^{\text{AP}} / \Big((1 + \mathbb{1}_{\text{couple}})\,\bar{M}\Big)$, and because the HEAS drawable amount equals $1.5\bar{M} - I^{\text{AP}} / (1 + \mathbb{1}_{\text{couple}})$, neither channel introduces additional kinks: both inherit exactly the same breakpoints, affecting only the slope of clawback on each segment. Collectively, these thresholds partition the wealth axis into segments on which the clawback is affine in X, fully characterising the piecewise linear structure displayed in Figure 10.

References

AGA (2025). *Australian Life Tables.* Australian Government, Canberra.

AIHW (2024). *Series: Home Care Packages Program Data Report.* 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.

Ascari, G. and Sbordone, A. M. (2014). The macroeconomics of trend inflation. *Journal of Economic Literature*, 52(3):679–739.

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

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.

Braun, R. A., Kopecky, K. A., and Koreshkova, T. (2017). Old, sick, alone, and poor: A welfare analysis of old-age social insurance programmes. *The Review of Economic Studies*, 84(2):580–612.

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.

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

Davidoff, T. (2015). Can "high costs" justify weak demand for the Home Equity Conversion Mortgage? *The Review of Financial Studies*, 28(8):2364–2398.

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.

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.

Dowd, K., Buckner, D., Blake, D., and Fry, J. (2019). The valuation of no-negative equity guarantees and equity release mortgages. *Economics Letters*, 184:108669.

Fella, G. (2014). A generalized endogenous grid method for non-smooth and non-concave problems. *Review of Economic Dynamics*, 17(2):329–344.

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 Suppl 1:221–258.

Flavin, M. and Yamashita, T. (2002). Owner-occupied housing and the composition of the household portfolio. *American Economic Review*, 92(1):345–362.

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

Fong, S., Koh, W.-P., and Gruber, J. (2025). Reframing biological age as risk-equivalent age. *Nature Aging*, pages 1–4.

Ghili, S., Handel, B., Hendel, I., and Whinston, M. D. (2024). Optimal long-term health insurance contracts: Characterization, computation, and welfare effects. *Review of Economic Studies*, 91(2):1085–1121.

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. (2025). Reforming the funding of long-term care for older people: costs and distributional impacts of planned changes in England. *Health Economics, Policy and Law*, pages 1–21.

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

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.

Klein, J. P. and Moeschberger, M. L. (2003). *Survival Analysis: Techniques for Censored and Truncated Data.* Springer, 2nd edition.

Kullback, S. and Leibler, R. A. (1951). On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86.

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*, page 103193.

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

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.*

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

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

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

Reichenstein, W. and Meyer, W. (2018). Understanding the tax torpedo and its implications for various retirees. *Journal of Financial Planning*, 31(7):38–45.

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: www.servicesaustralia.gov.au/home-equity-access-scheme (accessed 13 January 2026).

Services Australia (2026b). Income stream. Available at: https://www.servicesaustralia.gov.au/income-streams?context=22526 (accessed 2 Feb 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.

Skorokhod, A. V. (2009). *Asymptotic Methods in the Theory of Stochastic Differential Equations*, volume 78. American Mathematical Society.

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

The Treasury (2025). *Tax Expenditures and Insights Statement.* Australian Government, Canberra.

Van Houtven, C. H., Coe, N. B., and Konetzka, R. T. (2015). Family structure and long-term care insurance purchase. *Health Economics*, 24:58–73.

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

Wardle, H., Degenhardt, L., Marionneau, V., Reith, G., Livingstone, C., Sparrow, M., Tran, L. T., Biggar, B., Bunn, C., Farrell, M., et al. (2024). The lancet public health commission on gambling. *The Lancet Public Health*, 9(11):e950–e994.

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. 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.