



Centre for Population Ageing Research (CEPAR)

Working Paper 2026/05

**Optimal Long-Term Care Insurance under Means-Tested
Public Support: A Stochastic Control Model**

Lingfeng Lyu, Michael Sherris and Mengyi Xu

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

Optimal long-term care insurance under means-tested public support: a stochastic control model

Lingfeng Lyu^{1,2*} Michael Sherris^{1,2} Mengyi Xu^{1,2}

¹School of Risk and Actuarial Studies, UNSW Sydney, NSW 2052, Australia.

²Centre for Population Ageing Research (CEPAR), UNSW Sydney, NSW 2052, Australia.

October 12, 2026

Abstract

Means tests reserve public support, care included, for retirees of limited means: income tests reduce it as income rises; an asset test stops it above a wealth threshold. Yet even wealthier retirees, seemingly beyond these tests, rarely buy long-term care insurance. Such insurance is penalised by an implicit tax: the support lost because payouts count as income. We show that the tax lowers coverage even for retirees far above the threshold: with no public support above it, consumption and risk can run wealth below it, where payouts reduce support. We derive this in a model where a healthy retiree buys coverage once, consumes and invests, may need care and may value bequests. Merton's solution does not apply: support stops abruptly at the threshold, and wealth cannot go negative. Without a bequest motive, convex duality solves the care phase explicitly up to two scalars, and the rest is numerical. In our calibration, a retiree below the threshold keeps under half of each payout, so the premium is over twice what they expect to keep. For wealth up to five times the threshold, optimal coverage is below a third of the cost of care, and a bequest motive can bring it near zero. Exempting payouts from income tests raises coverage at all these wealth levels, even with a bequest motive, at little extra public cost. Without a bequest motive, abolishing the asset test lowers coverage for all but the poorest, mainly because everyone then qualifies for support and payouts reduce it.

Keywords: Finance; Long-term care insurance; Means testing; Stochastic control; Convex duality

*Corresponding author. Email: lingfeng.lyu@unsw.edu.au

1 Introduction

The need for long-term care is common: OECD (2024, p. 11) reports that in 2021 about one person in nine aged 65 and over received it on average across its member countries, and more than one in five in Germany, Israel, Lithuania and Switzerland. Over a lifetime the chance is far higher: in the US, 40% of men and 54% of women aged 65 will use long-term care at some point (Brown & Finkelstein, 2008). Yet private cover is rare, even among the wealthy: in the US, where long-term care insurance (LTCI) is sold, only a fifth of those aged 62 and over in the top wealth quintile hold a policy (Braun et al., 2019). In many countries, Australia among them, LTCI is unavailable (Park & Sherris, 2026), and in the UK it is no longer sold (OECD, 2020, p. 18).

Part of the reason for low take-up lies in how public support is targeted. In most high-income countries, public support for older people’s long-term care is means tested, so that it goes to people of limited means: income tests lower it as a person’s income rises and, in many countries, an asset test reduces or stops it as their wealth rises (OECD, 2024, pp. 53, 60). Many countries, among them the US, the UK, France and Australia, also means test the minimum income of retirees. When a test counts a payout from LTCI as income, or the policy as an asset, a retiree who holds LTCI gives up part of the support they would otherwise receive. This lost support is the clawback. It acts as an implicit tax on LTCI, the term Brown and Finkelstein (2008) use for the share of a private policy’s benefits that replace what Medicaid, the US means-tested programme that pays for care, would otherwise have paid. Numerical studies find a spillover: a means test aimed at people of limited means also deters wealthier retirees from insuring. Given Medicaid, the bottom two-thirds of the US wealth distribution would not buy even fairly priced insurance (Brown & Finkelstein, 2008), and Medicaid, together with administrative costs, keeps take-up low among middle-income retirees (Braun et al., 2019).

This paper asks two questions about the spillover: why a means test lowers the coverage of retirees far above the wealth at which it pays support, and how the design of the test changes that. In the model, the answer to the first is timing. A retiree chooses coverage once, while healthy, without knowing whether or when disability will begin, how long they will live, or what their investments will return. Coverage pays out only in disability, possibly years later. Above the range where the test pays support a retiree has no public support, so by then consumption and these risks can have run wealth down into that range, where payouts reduce support. The value of a payout is an average over the wealth held while care lasts, weighted most where wealth is lowest, which is where the test claws payouts back (Lemma 6.2(b); Proposition 6.3); so the implicit tax also reaches retirees who are above that range when they buy, as Braun et al. (2019) observe of Medicaid. The second question concerns the government that sets the test and decides how it treats a private payout, and so what coverage is worth to every retiree who may qualify for support. In our calibration the lever is the income tests: exempting payouts from them raises coverage at little extra public cost, and abolishing the asset test does not.

We formulate the retiree’s problem as a continuous-time stochastic control problem on a three-state Markov chain (healthy, disabled and dead), with a riskless and a risky asset and a borrowing constraint at zero wealth (Section 4). In the model the asset test withdraws support once wealth passes a threshold, the asset-test edge. Below the edge the income tests claw back

part of each payout through two channels, one withdrawing the cash safety net, which pays the minimum income, and the other a care subsidy. We assume a competitive market in which coverage is bought for a single premium, set at the expected present value of the payouts plus a market loading.

Three features of the problem rule out the solution of Merton (1971) and the usual first-order condition for coverage. First, income and the out-of-pocket cost of care jump at the edge, a point inside the state space, and Merton’s solution, fitted on each side of the edge, does not join there. Joining solutions on the two sides smoothly at the edge does not settle the problem: the borrowing constraint bounds wealth, not leverage, so a solution of the Hamilton–Jacobi–Bellman (HJB) equation is the value function only if it is strictly concave, which keeps the risky position finite, and the step in income and costs at the edge can break concavity. Second, the retiree cannot borrow against future income, so wealth cannot go negative, and at zero wealth Merton’s portfolio, counting future income as wealth, would require borrowing. Third, the income tests withdraw a fixed share of each payout until the support they pay is used up, so the amount withdrawn is a concave, piecewise-linear function of coverage, with a kink where each channel’s support runs out. At each kink the marginal benefit of coverage jumps upward, so the retiree’s objective is not concave in coverage and the first-order condition no longer identifies the optimum.

We solve the disabled state, without a bequest motive, by convex duality: under a Legendre transform the HJB equation on each side of the edge becomes a linear Cauchy–Euler equation (Section 5). We solve the healthy state, and the disabled state with a bequest motive, by policy iteration on a finite-difference grid. Because the objective is not concave in coverage, we obtain optimal coverage by comparing the best coverage on each side of a kink. On each side, an interior optimum solves a demand equation: the premium’s marginal cost equals the marginal benefit of coverage in disability, averaged over the wealth the retiree may hold when care begins and while it lasts (Section 6). We calibrate the model jointly to the UK and Australian means tests, to long-run returns (Jordà et al., 2019), and to published preference, health and bequest parameters (Section 8). As LTCI is not sold in either country, the coverage we compute is what a retiree facing means tests like theirs would buy if such a market existed.

On the first question, we find, within the model and at the calibration, that at five times the edge and without a bequest motive the government expects to pay a retiree almost nothing over their remaining life, 0.05 of the annual cost of care, yet exempting payouts from the income tests would raise their coverage from 0.31 to 0.49 of the annual cost of care (Section 10). Below the edge, on coverage up to the kink where the cash safety net runs out, the test claws back 54 cents in every dollar paid out, so each dollar the retiree keeps costs 2.39 in premium: a loading of 139%, where the market charges 10% (Section 9.1). Without a bequest motive optimal coverage insures less than a third of the cost of care at every wealth up to five times the edge, and a bequest motive can bring it close to zero (Section 9.2).

On the second question, we find that, at the calibration, the exemption raises coverage at every wealth computed, with or without a bequest motive, at an extra public cost of at most 0.11 of the annual cost of care (Section 10); without a bequest motive, each retiree values the change at 2 to 20 times its extra public cost. Without a bequest motive, abolishing the asset

test lowers coverage for all but the poorest; with one, it removes what little coverage there is. The main reason is that, with no asset test, every retiree in effect passes it: all get support, and the income tests then claw back part of every payout. Abolishing the asset test also raises the public cost severalfold for retirees with wealth above the edge.

This paper contributes to the literature on stochastic control models of insurance and retirement, and to that on LTCI demand under means testing, in four ways. First, we formulate a means test as a step in the drift at an exogenous wealth threshold, and a withdrawal that is concave and piecewise linear in coverage (Section 4.3). Koo et al. (2022) solve means-tested models numerically; here, for a given coverage and without a bequest motive, the disabled state's HJB equation is Merton's on each side of the edge (Section 5.2), so the dual method applies on each side, and the two sides must then be joined at the edge.

Second, we construct the value function of the disabled state, in which wealth can fall below the edge and rise above it, under the borrowing constraint at zero wealth (Proposition 5.3). In earlier uses of the dual method (Dybvig & Liu, 2010, 2011; El Karoui & Jeanblanc-Picqué, 1998; He & Pagès, 1993), the boundary is a free boundary, set by the optimal policy and located by smooth fit; here the means test sets the edge, and value and marginal value are matched there across the step in the drift. Without a bequest motive the matched solution is explicit up to two scalars, and it is the value function if the two equations for those scalars have a root at which sign conditions for concavity hold (Section 5.3), as they do at the calibration (Section 8). Where a sign condition fails just below the edge, a retiree below the edge optimally stays below it in every case computed, and the demand equation and the pricing of the implicit tax are not carried to that case. The construction is not specific to care insurance: under the same conditions, it applies to problems with constant relative risk aversion, unbounded leverage and no bequest motive in which crossing an exogenous wealth threshold moves the drift by a step.

Third, we derive the demand equation for coverage and, without a bequest motive, price the implicit tax in closed form. Optimal coverage never sits at a kink but can jump past one as wealth rises (Theorem 6.4). Brown and Finkelstein (2008) measure the implicit tax on a policy as a whole; here the tax is priced on a marginal coverage dollar. Under the means test alone, without the borrowing constraint, a coverage dollar received below the asset-test edge would be worth the share of the payout the test leaves, plus the option of rising above the edge, where payouts are kept in full (Theorem 7.1). At zero wealth, where cash today counts for more than a later payout, the borrowing constraint removes a fraction of that share; the fraction does not depend on the means test, so the two losses can be told apart (Theorem 7.1(d)).

Fourth, we compare the baseline means test with seven alternatives, with and without a bequest motive, by the coverage each induces, by its expected public cost and by its value to the retiree (Section 10). Brown and Finkelstein (2008) hold the insurance policy fixed under each reform; here the retiree chooses coverage under each design.

Section 2 reviews the literature on consumption and portfolio choice under a borrowing constraint, on LTCI demand under means testing and on the bequest motive. Section 3 measures the clawback in the UK and Australia and reduces it to the two channels of the model of Section 4. Sections 5 to 7 construct the disabled value, derive the demand equation and price the implicit tax. Section 8 sets the calibration, Section 9 reports the loadings and coverage,

Section 10 compares means-test designs, and Section 11 concludes; proofs, numerical reporting and the country schedules are in the Online Supplement.

2 Related literature

The construction of the disabled value belongs to the literature that solves consumption and portfolio problems by duality (Karatzas et al., 1987). The papers closest to it treat an agent who cannot borrow against future income, as the retiree here cannot. He and Pagès (1993) and El Karoui and Jeanblanc-Picqué (1998) solve that problem in closed form over an infinite horizon, and Dybvig and Liu (2010, 2011) add a retirement decision and verify the solution. Kim et al. (2026) take the retirement problem to a finite horizon, where the constraint and the retirement decision give two free boundaries. In the retirement problem, income stops for good once the agent retires; here support stops above the asset-test edge and resumes below it.

The application is the demand for LTCI under means-tested public support. Small LTCI markets have been traced in part to loads (Brown & Finkelstein, 2007), to bequest motives, taken up below, and to means testing, which acts as an implicit tax. Lyu and Sherris (2026) write the Australian means tests into a life-cycle model of LTCI demand and solve it numerically. In Hubbard et al. (1995) the implicit tax falls on saving, and their asset-tested consumption floor makes the budget set non-convex, so that, as with coverage here (Theorem 6.4), the optimum is found by comparing candidates rather than from a first-order condition alone. In Brown and Finkelstein (2008) it falls on private LTCI: Medicaid pays only what a private policy does not, and a policy that protects assets makes its holder less likely to qualify, so part of its benefits replaces what Medicaid would have paid. This order of payment is the counterpart of the clawback; removing either the order or the means test alone has little effect on that tax, and hence on the willingness to pay for insurance (Brown & Finkelstein, 2008). In this paper, by contrast, Design 2 of Section 10, which exempts payouts from the income tests but keeps the asset test, raises coverage at the calibration.

The bequest motive, the other influence on coverage here, affects LTCI demand differently across models. It lowers demand in Lockwood (2018), as it does in this paper, but not in the next two models. In Xu et al. (2023), where home equity cannot be borrowed against, a bequest motive raises homeowners' demand for LTCI, to preserve their bequests. In the complete market of Koijen et al. (2016) a stronger bequest motive makes a household want more wealth in poor health relative to good health, and the household can also buy life insurance, which pays at death. Without life insurance to restore the estate, a premium paid here by one who never needs care is lost to it (Section 9.2). Ameriks et al. (2011, 2020) separate the bequest motive from an aversion to publicly funded care and from a higher marginal utility of spending in care; their two estimates bound the published range over which Section 9.2 varies the motive's strength.

Leporatti et al. (2025) solve consumption, portfolio and LTCI coverage in closed form, but in a complete market, with coverage adjusted continuously and no means test or borrowing constraint; their insured share falls with wealth, whereas coverage here rises with it at the calibration (Section 9.2). In operational research two papers are closest, and neither treats LTCI demand under a means test. Barucci et al. (2023) use convex duality for health insurance

and portfolio choice when a health shock raises mortality by a constant, as in Assumption 4.3, and find a sizeable demand for fairly priced health insurance in a model with no means test. Koo et al. (2022) solve retirement problems under the means-tested Australian Age Pension by dynamic programming; their analytical solutions fix the pension at its full rate, so the means test enters only their numerical solutions.

3 The clawback in the UK and Australia

In several OECD countries, means-tested public programmes count a private LTCI contract as an assessable asset, its payouts as assessable income, or both. A retiree who holds LTCI can therefore lose part of their public support. We call this lost support the clawback. In the UK and Australia it extends to retirees whose wealth exceeds the value of their home, and it keeps rising with coverage.

We measure all monetary quantities in units of the annual cost of care, so an LTCI policy that covers this cost in full pays 1 for each year of care. In each country we set this cost at the maximum annual cash safety net available to a single homeowner divided by 0.625, so the safety net is 0.625 in each country.^[1] The cash safety net is the means-tested retirement income: Pension Credit in the UK and the Age Pension in Australia. Each country’s clawback is our simplified reading of its published rules, at the parameter values listed in Section S.9 of the Online Supplement, to which every number with the prefix S refers. The calibration uses the UK and Australia, shown below.

3.1 Decomposition of the clawback across wealth levels

We restrict the analysis to home care.^[2] Figure 1 decomposes the clawback in the UK and Australia into two channels. The home-care channel is the public care subsidy forgone when LTCI payouts displace government funding, and the safety-net channel is the reduction in the cash safety net when LTCI payouts are assessed as income. Each channel is the support the retiree receives without LTCI less the support received with it, and the net clawback is the sum of the two.

The horizontal axis is normalised total wealth X , liquid assets plus home equity in the units above. Means tests exempt the family home, so we need a rule that splits X into the home and liquid wealth. For the means test we hold the home fixed at ten normalised units, and liquid wealth is the residual:^[3]

$$\text{Home equity} = \min(X, 10), \quad \text{Liquid wealth} = X - \min(X, 10). \quad (1)$$

Because the home is the same in both countries, the two panels of Figure 1 compare institutions

^[1]The value 0.625 is a figure we chose, not an estimate from care prices. Measuring each country in its own unit makes the two clawbacks comparable across currencies.

^[2]Residential-care means tests draw home equity, accommodation charges and deposits into the assessment, and they differ across regions within a country.

^[3]We assume retirees can turn home equity into liquid wealth through a reverse mortgage at no cost, repaid only on death, sale or a permanent move, and that home equity never exceeds the value of the home itself.




Figure 1. The clawback in the UK and Australia as a function of normalised total wealth.

Notes: Total wealth, liquid assets plus home equity, is in multiples of the annual cost of care, and the clawback, a yearly flow in the same unit, is drawn at an assumed LTCI indemnity of 1. The retiree receives private income of 0.375 a year in the UK, the mean private pension of single pensioners, and 0.15 in Australia, set just above the income-test free area, the income the test ignores. Australia’s income test therefore counts the first dollar of the payout. The income tests count this private income with and without LTCI (Section S.9). The sum of the home-care and safety-net clawbacks is the net clawback.

and not wealth composition.^[4]

For low-wealth retirees who rely on public support, an assessed LTCI payout triggers a clawback; in the UK a retiree with little besides the home loses almost a dollar of support for each dollar the policy pays. The clawback is flat while total wealth is below the value of the home, since the retiree then has no liquid assets to count, and stays flat while liquid assets are below the amounts the tests exempt. Once liquid assets pass the exempt amounts, the clawback shrinks, for two reasons. First, the income test deems liquid assets to earn income, so a wealthier retiree is already closer to the income limit of the cash safety net before the indemnity is counted. Second, the asset test cuts support directly once liquid assets pass its limit. Eventually the retiree is wealthy enough that no public support is left for the indemnity to displace.

3.2 The indemnity level and the intensity of the clawback

In both countries a larger indemnity raises the net clawback gradually, and the clawback is still rising at full coverage, an indemnity $D = 1$ (Figure S.2). In the UK the safety-net channel is flat at 0.25 in Figure 1 until liquid wealth reaches the asset test’s capital limit. That 0.25 is all the cash support the retiree in the UK panel receives, since the safety net tops a private income of 0.375 up to 0.625. The indemnity uses it up first, and further coverage is clawed back through the home-care channel alone. In Australia both channels are still rising at $D = 1$.

^[4]The LTCI contract itself is not part of total wealth X . Where a purchase creates assessable assets or income, as under the lifetime income stream rules used for Australia, the assessment changes the level of the clawback at a given X ; it does not move the retiree along the axis.

3.3 The clawback and the size of private LTCI markets

Private LTCI covered about 7.1 million people in France at the end of 2017 (Bonne & Meunier, 2019, p. 71) and about 7.5 million Americans at the start of 2020 (Colello, 2023, p. 1), while none is sold in the UK or Australia (Section 1).

In the UK and Australia the primary residence is exempt and the cash safety net is withdrawn gradually, so homeowners with savings above the exempt amounts still qualify for part of the support. This is why the clawback in these two countries extends past the value of the home (Figure 1) and keeps rising with coverage (Figure S.2). The US and France, shown in Section S.9.7 for a retiree with no private income, both recover benefits from the estate after death, which neither the UK nor Australia does. Net of that recovery, which we set against the clawback at face value, their clawback is confined to wealth up to about the value of the home. It also saturates at partial coverage: an indemnity of 0.625 exhausts the cash safety net, so further coverage adds nothing to the clawback.

We offer this contrast as a hypothesis, not as an explanation. Where the clawback covers a narrow range of wealth and stops growing at modest coverage, most of an LTCI payout escapes it, and a market can form. Where the clawback covers a wide range of wealth and keeps growing with coverage, the payout mostly replaces public support, and the policy is worth little.

3.4 From the cross-country clawback to the model

In the rest of the paper the two clawback channels of Figure 1 enter the model as two parts of one public schedule. Both parts count the LTCI payout as income, and both apply only below the same wealth threshold, the asset-test edge, above which the asset test removes eligibility. This threshold stands in for the gradual fall of the clawback with wealth in Figure 1. The asset test therefore sets the range of the clawback, and the income tests set its intensity. The safety-net channel claws back an LTCI payout at a higher rate and stops once the cash support is used up; the home-care channel claws it back at a lower rate and stops once the care subsidy is used up. We therefore give each channel its own rate and its own cap.^[5] The cash support enters the income schedule $y(x)$ of Section 4.2 and the care subsidy the cost schedule $O(x)$ of Section 4.3; both jump at the asset-test edge.

Section 4 formalises this two-part schedule. In the model we call this clawback the implicit tax, the name Brown and Finkelstein (2008) give to the same effect of Medicaid on private LTCI in the US. The wealth band on which Figure 1 shows a high clawback becomes the implicit tax zone, the interval on which the implicit tax applies.

4 Model

This section sets out the model: the health states, the wealth and income dynamics, the LTCI premium, the implicit tax created by the means test, and the retiree's optimisation problem.^[6]

^[5]Section S.9 constructs the clawback intensity of each channel and, from it, the rate used in the calibration.

^[6]Each result is followed by a reference to its proof. The proofs are in the Online Supplement, most of them in Section S.5.

4.1 Health state transitions

The agent is a single retiree. Time $t \geq 0$, in years, is continuous, and the agent's health state S_t is a continuous-time Markov chain.

Definition 4.1 (Health-state Markov chain). *The state space of S_t is $\{0, 1, \dagger\}$, which captures the morbidity and mortality risks associated with long-term care (LTC). In State 0 (healthy) the agent needs no LTC services; in State 1 (disabled) severe cognitive impairment or limitations in activities of daily living require formal LTC; State $\dagger$ (dead) is absorbing. Transitions occur at constant intensities, and ν denotes the morbidity rate, the intensity of transition from State 0 to State 1.*

Assumption 4.2 (Irreversibility of disability). *Recovery from State 1 to State 0 is ruled out, consistent with the persistence of late-stage LTC needs.*

Assumption 4.3 (Excess mortality penalty). *The mortality intensities in the healthy and disabled states, $\lambda^{(0)}$ and $\lambda^{(1)}$, satisfy*

$$\lambda^{(1)} = \lambda^{(0)} + \alpha, \quad \lambda^{(0)} \text{ and } \alpha > 0, \quad (2)$$

where α is the additional mortality intensity incurred on entering the disabled state.

4.2 Wealth and income dynamics

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space supporting a standard Brownian motion W_t independent of S_t , with the filtration $\{\mathcal{F}_t\}_{t \geq 0}$ generated by (W_t, S_t) and augmented to satisfy the usual conditions.

Definition 4.4 (Financial market and controls). *The agent can invest in two assets: a riskless bond paying a constant interest rate $r > 0$, and a risky asset with expected return $\mu_R > r$ and volatility $\sigma_R > 0$. The risky asset stands for the agent's risky wealth: equity and home equity together (Section 8). The means test treats the home differently: it counts the first ten units of total wealth as the exempt home (1). Let $x_t \in \mathbb{R}_+$ denote the agent's total wealth at time t , the X of Section 3. At each time the agent chooses a control $\Gamma_t = (c_t, \pi_t)$, where $c_t > 0$ is the consumption rate and π_t is the proportion of wealth held in the risky asset; Definition 4.8 states which controls are admissible.*

Before any insurance product is introduced, wealth follows the stochastic differential equation

$$dx_t = [(r + \pi_t(\mu_R - r))x_t + y(x_t) - c_t] dt + \pi_t \sigma_R x_t dW_t. \quad (3)$$

For brevity, we write the expected financial drift as $\mu(x, \pi) \equiv (r + \pi(\mu_R - r))x + y(x)$. The baseline income $y(x)$ collects the flows the agent receives whatever the portfolio: the cash safety net, together with any private flows. Investment income is carried separately by the term $(r + \pi(\mu_R - r))x$. Because the public flows are means tested, y has a two-level structure in wealth:

$$y(x) = y_{\text{in}} \mathbf{1}\{x \leq x_U\} + y_{\text{out}} \mathbf{1}\{x > x_U\}, \quad \Delta_y := y_{\text{in}} - y_{\text{out}} \geq 0, \quad (4)$$

where $x_U > 0$ is the asset-test edge, the wealth above which the asset test removes eligibility for means-tested support.

On $[0, x_U]$ the agent passes the asset test and receives the public flows on top of any private ones. Thus y_{in} is the cash flow received on $[0, x_U]$ before any clawback of the indemnity, treated as one constant there (Section S.1.1). Above x_U the asset test withdraws the public flows, leaving the private flows $y_{\text{out}} \geq 0$, mainly private pensions, which the income tests count but never withdraw. The cash step Δ_y is the public cash support lost at the edge. It equals the full safety net of Section 3 only for a retiree with no private income, and it is smaller for retirees with more private provision. The tests assess wealth and income, not disability, so (4) holds in health as in disability (Section S.1.5).

Assumption 4.5 (Single up-front LTCI premium). *The agent can buy LTCI in a competitive private market, and buys coverage $D \geq 0$ once, at inception in the healthy state, for a single premium $P(D)$ deducted from wealth at $t = 0$. From the onset of disability until death the policy pays a continuous indemnity flow at rate D . The premium cannot exceed inception wealth x_0 . Insurers offer indemnities up to the coverage cap $\bar{D}$, which is at most the unsupported care cost O_{out} of Section 4.3.^[7]*

Proposition 4.6 (Single premium at the market loading). *Under Assumption 4.5 with market loading $\eta \geq 0$, the single premium is linear in coverage, $P(D) = a_D D$, with the unit premium*

$$a_D = (1 + \eta) \frac{\nu}{(r + \nu + \lambda^{(0)})(r + \lambda^{(1)})}. \quad (5)$$

Proof. Section S.5.

4.3 The implicit tax under means testing

In the disabled state the government subsidises care, and the means test withdraws the subsidy as assessable income or assets rise. Even without LTCI the agent therefore bears a baseline out-of-pocket cost $O(x) > 0$, the care cost net of the subsidy received, which depends on wealth.

The asset-test edge x_U of Section 4.2 separates two regions, and we call $[0, x_U]$ the implicit tax zone, or simply the zone (Assumption 4.7 below). On the zone the agent qualifies for the subsidy, subject to the income test, and the baseline cost is the constant O_{in} , the out-of-pocket cost at zero LTCI coverage. Above x_U the asset test removes eligibility for the means-tested subsidy: the baseline cost rises to the unsupported care cost O_{out} , and income falls by the step Δ_y of (4). The cost schedule has the same form as the income schedule, with a single jump at x_U :

$$O(x) = O_{\text{in}} + \Delta_O \mathbf{1}\{x > x_U\}, \quad \Delta_O := O_{\text{out}} - O_{\text{in}} > 0. \quad (6)$$

The net resource flow $y(x) - O(x)$ therefore falls at x_U by

$$\theta_{\text{max}} = (y_{\text{in}} - y_{\text{out}}) + (O_{\text{out}} - O_{\text{in}}) = \Delta_y + \Delta_O, \quad (7)$$

^[7]An indemnity policy pays for care, so its cap is bounded by the most that care can cost the retiree, the cost O_{out} paid without means-tested support. Section 8 sets the cap at this bound, and Section S.9 explains why.

and we call this drop the support cliff. It is the support that a retiree with no coverage receives on the zone when disabled, and so the most that the means test can withdraw.^[8]

The implicit tax Π is the clawback of Section 3: the support that the test withdraws because of the indemnity, the withdrawal with coverage less the withdrawal without it. We take it in the two-channel form of Section 3.4.

Assumption 4.7 (The implicit tax zone: two channels, two kinks). (i) (Two channels.) *On the zone each channel withholds a constant share of the indemnity, $m_{\text{SN}} > 0$ for the safety-net channel or $m_{\text{HC}} > 0$ for the home-care channel, until it has taken the support it pays, the cash step Δ_y or the care step Δ_O ; above the zone the asset test has removed the support and nothing is withheld:*

$$\Pi(x, D) = \Pi(D) := \min\{m_{\text{SN}}D, \Delta_y\} + \min\{m_{\text{HC}}D, \Delta_O\} \text{ on } [0, x_U], \quad \Pi \equiv 0 \text{ for } x > x_U. \quad (8)$$

The channels have taken their whole support at the coverages

$$D_{\text{SN}} := \frac{\Delta_y}{m_{\text{SN}}}, \quad D_{\text{HC}} := \frac{\Delta_O}{m_{\text{HC}}}. \quad (9)$$

(ii) (Two kinks.) *The total rate $m := m_{\text{SN}} + m_{\text{HC}}$ satisfies $m \leq 1$, and the channels saturate in the order*

$$D_{\text{SN}} < D_{\text{HC}}. \quad (10)$$

The ordering holds when the cash step Δ_y is the smaller support and is clawed back at the steeper rate, so that it is used up first, as in the UK schedule of Section 3.2. We call D_{SN} and D_{HC} the kinks.^[9] Coverage on the zone therefore falls into three tranches: on the first tranche, $[0, D_{\text{SN}})$, both channels claw back the indemnity; on the intermediate tranche, $(D_{\text{SN}}, D_{\text{HC}})$, the home-care channel alone does; on the saturated tranche, beyond D_{HC} , nothing more is clawed back.

Section S.1.1 derives (8) from a general withdrawal schedule. The two-channel form holds when the asset test and the income tests act separately, the withdrawals of the two channels add, and each channel's rate is a zone average. The condition $m \leq 1$ is the non-confiscatory condition, at most one dollar of support withdrawn for each dollar of income, applied to the averaged rates.

Under Assumption 4.7, the share of a marginal indemnity dollar that the agent keeps after the test on the zone, the tranche retention, is the two-step schedule

$$1 - \Pi'(D) = \delta(D) := 1 - \sum_{k \in \{\text{SN}, \text{HC}\}} m_k \mathbf{1}\{D < D_k\} = \begin{cases} 1 - m, & 0 \leq D < D_{\text{SN}}, \\ 1 - m_{\text{HC}}, & D_{\text{SN}} < D < D_{\text{HC}}, \\ 1, & D > D_{\text{HC}}, \end{cases} \quad (11)$$

^[8]The zone-side pair $(y_{\text{in}}, O_{\text{in}})$ is the input and the values above the edge follow from it and the two steps, $y_{\text{out}} = y_{\text{in}} - \Delta_y$ and $O_{\text{out}} = O_{\text{in}} + \Delta_O$; Section S.1.1 records this decomposition.

^[9]Table 1 gives them at the calibration.

with one-sided derivatives at the kinks. At the kink D_k retention rises by exactly m_k , the rate of the channel that saturates there; this step is the source of the jumps in the marginal benefit of coverage in Section 6.

With the indemnity D and the means-tested schedule, the State 1 wealth dynamics are

$$dx_t^{(1)} = [\mu(x_t, \pi_t) - c_t - O(x_t) + D - \Pi(x_t, D)] dt + \pi_t \sigma_R x_t dW_t, \quad (12)$$

where $O(x_t)$ is the out-of-pocket care cost (6) and $D - \Pi(x_t, D)$ is the payout retained after the test. Its derivative in D is the after-test retention of an indemnity dollar,

$$\frac{\partial}{\partial D} [D - \Pi(x, D)] = \delta(D) \text{ on } [0, x_U], \quad 1 \text{ above.} \quad (13)$$

4.4 The optimisation problem

4.4.1 Objective and value functions

The agent maximises expected discounted lifetime utility, made up of utility from consumption and, through a bequest motive, utility from the wealth left at death. Both are constant relative risk aversion (CRRA) functions, the second shifted as in Lockwood (2018). Let $\tau_\dagger$ denote the random time of death.

Definition 4.8 (Admissible strategies). *A strategy $\Gamma = (c_t, \pi_t)_{t \geq 0}$ is admissible if it is progressively measurable with respect to $\{\mathcal{F}_t\}$, with $c_t > 0$, $\pi_t \geq 0$ and $\int_0^T (c_t + \pi_t^2 x_t^2) dt < \infty$ almost surely for every $T > 0$, and if the wealth equation, (3) in health and (12) in disability, has a unique strong solution with $x_t \geq 0$ on $[0, \tau_\dagger]$. Let $\mathcal{A}$ denote the set of admissible strategies.*

The portfolio enters the wealth dynamics only through the dollar exposure $\wp_t := \pi_t x_t$, which is therefore the effective control.^[10] The definition rules out short sales but leaves leverage unbounded: apart from $\pi_t \geq 0$, the only restriction on the portfolio is the state constraint $x_t \geq 0$, which we call the borrowing constraint. It bars borrowing against future income. A leverage cap would bind near $x = 0$, where the optimal share diverges (Section S.1.4), and would remove the closed form of Section 5, since the dual transform linearises the HJB equation only when leverage is unbounded (Vila & Zariphopoulou, 1997, Appendix C, eqs. (C.1)–(C.5)). Vila and Zariphopoulou (1997, Proposition 4.2) obtain a closed form under such a cap only when the cap is proportional to wealth, with no credit line.

Definition 4.9 (Lifetime utility objective). *The objective functional under initial state (x, S_0) is*

$$J(x, S_0; \Gamma, D) = \mathbb{E} \left[\int_0^{\tau_\dagger} e^{-\beta t} U(c_t) dt + e^{-\beta \tau_\dagger} B(x_{\tau_\dagger}) \right], \quad (14)$$

where $\beta > 0$ is the subjective discount rate and $\Gamma \in \mathcal{A}$ is an admissible strategy (Definition 4.8). The LTCI coverage $D \geq 0$ is fixed by the purchase of Assumption 4.5, against the single premium $P(D) = a_D D$ of Proposition 4.6.

^[10]The integrability condition of Definition 4.8 is the standard one (Karatzas et al., 1987, Definition 2.2).

Assumption 4.10 (CRRA preferences and bequest motive). *Preferences are CRRA, a choice made for analytical tractability:*

$$U(c) = \frac{c^{1-\gamma}}{1-\gamma}, \quad B(x) = \kappa \frac{(x + \underline{x})^{1-\gamma}}{1-\gamma},$$

where $\gamma > 1$ is the coefficient of relative risk aversion, $\kappa \geq 0$ is the weight of the bequest motive, and the shift $\underline{x} > 0$, a fixed constant distinct from the inception wealth x_0 , keeps the bequest utility bounded near zero wealth.

Definition 4.11 (Value functions and the coverage problem). *For a fixed coverage level $D \geq 0$, the state-dependent value functions are*

$$V^{(j)}(x; D) = \sup_{\Gamma \in \mathcal{A}} J(x, j; \Gamma, D), \quad j \in \{0, 1\},$$

where J is the objective functional (14) and the supremum is over admissible strategies $\Gamma \in \mathcal{A}$ only. The single premium is paid out of inception wealth, so a retiree with inception wealth x_0 who buys coverage D is left with post-premium wealth $x_0 - a_D D$. With $D_{\max} := \min\{x_0/a_D, \bar{D}\}$, the largest coverage the retiree can afford and the market offers (Assumption 4.5), optimal coverage is the solution of the static problem

$$D^*(x_0) = \arg \max_{0 \leq D \leq D_{\max}} V^{(0)}(x_0 - a_D D; D). \quad (15)$$

4.4.2 The HJB system and decoupling

For each fixed D , dynamic programming yields the coupled system of HJB equations. Throughout we write $A_0 := \beta + \nu + \lambda^{(0)}$ and $A_1 := \beta + \lambda^{(1)}$ for the discount rates in health and in disability, each the subjective rate plus the intensities of leaving that state, and $\text{SR} := (\mu_R - r)/\sigma_R$ for the Sharpe ratio.

In the healthy state (State 0) the agent pays no premium flow (Assumption 4.5), so the HJB equation has the baseline drift, plus two jump terms: the onset of disability, which moves the agent to State 1, and death, which pays the bequest:

$$\begin{aligned} A_0 V^{(0)}(x; D) = & \max_{(c, \pi) \in (0, \infty) \times [0, \infty)} \left\{ \frac{c^{1-\gamma}}{1-\gamma} + [\mu(x, \pi) - c] V_x^{(0)} + \frac{1}{2} \pi^2 \sigma_R^2 x^2 V_{xx}^{(0)} \right\} \\ & + \nu V^{(1)}(x; D) + \lambda^{(0)} B(x). \end{aligned} \quad (16)$$

In the disabled state (State 1) the agent pays $O(x)$ for care, receives D and incurs the implicit tax $\Pi(x, D)$ (8); death is the only jump (Assumption 4.2), and the HJB equation is

$$\begin{aligned} A_1 V^{(1)}(x; D) = & \max_{(c, \pi) \in (0, \infty) \times [0, \infty)} \left\{ \frac{c^{1-\gamma}}{1-\gamma} + [\mu(x, \pi) - c - O(x) + D - \Pi(x, D)] V_x^{(1)} \right. \\ & \left. + \frac{1}{2} \pi^2 \sigma_R^2 x^2 V_{xx}^{(1)} \right\} + \lambda^{(1)} B(x). \end{aligned} \quad (17)$$

The means test acts on the indemnity through the term $-\Pi(x, D) V_x^{(1)}$: part of the indemnity received on the zone is clawed back, while the premium is priced on the gross indemnity.

Lemma 4.12 (Pointwise maximisers). *Let $V^{(j)}$ be C^2 on $(0, \infty)$ off x_U , with $V_x^{(j)} > 0$ and $V_{xx}^{(j)} < 0$, $j \in \{0, 1\}$. For $x > 0$ off x_U the Hamiltonian in (16)–(17) attains its maximum over $(c, \pi) \in (0, \infty) \times [0, \infty)$ at*

$$c^* = (V_x^{(j)})^{-1/\gamma}, \quad \wp^* = \pi^* x = -\frac{(\mu_R - r) V_x^{(j)}}{\sigma_R^2 V_{xx}^{(j)}} > 0,$$

which are Borel functions of x . Strict concavity is necessary as well as sufficient: if $V_{xx}^{(j)} \geq 0$ at some $x > 0$, the supremum over π is $+\infty$ and no maximiser exists.

Proof. Section S.5.

Remark 4.13 (Verification). *Standard verification theorems do not cover the state constraint $x_t \geq 0$; Section S.1.6 verifies State 1 by the supermartingale argument of Dybvig and Liu (2011, Lemma 3.3), and State 0 by the same argument given a strictly concave solution of (16) with the other properties of Assumption 5.7 below. State 0 is solved numerically, and its solution is checked for strict concavity at the calibration (Assumption 5.6). Pointwise attainment of the maximum justifies the envelope step and fixes the coefficients of the linear equations used below; the Feynman–Kac representations of their solutions also use bounds on the coverage derivatives. Section S.4 shows that the closed-loop wealth equation is well posed, in State 0 given parts (i) and (iii) of Assumption 5.7 below; of these, the positive drift and the Lipschitz continuity near $x = 0$ are checked numerically at $\kappa = 0$.*

Since the State 1 equation (17) does not involve $V^{(0)}$, the coupled system can be solved recursively, State 1 first (Section S.5), and Section 5 constructs the disabled value function.

5 The disabled value function

The demand for coverage depends on what a marginal indemnity dollar is worth to a disabled agent; on the zone, the agent keeps only a fraction of that dollar. Without a bequest motive, $\kappa = 0$, the disabled value function is Merton’s power law, a power of the sum of wealth and capitalised income, with two corrections: the value of crossing the asset-test edge x_U and, on the zone, the cost of the borrowing constraint at zero wealth. Because the implicit tax depends on coverage alone on each side of the edge (Assumption 4.7), Lemma 5.1 reduces State 1 to a Merton problem on each region. Proposition 5.3 finds the two corrections through a dual transform; the value is explicit up to two scalars, the marginal values of wealth at the edge and at zero wealth, found by root-finding.

5.1 The two-region Merton reduction

Lemma 5.1 (Two-region Merton reduction and saturation). *Let Assumption 4.7 hold. Then:*

- (a) (The implicit tax is a capped piecewise-linear function of coverage.) *The implicit tax (8) is continuous, concave and piecewise linear in coverage, with slope m on $[0, D_{\text{SN}})$, m_{HC} on $(D_{\text{SN}}, D_{\text{HC}})$ and zero beyond, and $\Pi(D_{\text{HC}}) = \theta_{\text{max}}$, the last using the ordering (10).*

- (b) (A two-region Merton problem.) For $D < D_{\text{HC}}$, State 1 (17) is a Merton consumption–portfolio problem with the piecewise-constant retained drift

$$b_{\text{ret}}(x) = \begin{cases} y_{\text{in}} - O_{\text{in}} + D - \Pi(D), & x \in [0, x_U], \\ y_{\text{out}} - O_{\text{out}} + D, & x > x_U, \end{cases}$$

capitalised into the shifts $h_{\text{in}} = (y_{\text{in}} - O_{\text{in}} + D - \Pi(D))/r$ and $h_{\text{out}} = (y_{\text{out}} - O_{\text{out}} + D)/r$.^[11] The shift h_{in} increases with coverage at the retention of the current tranche, with slope $(1 - m)/r$ on the first tranche and $(1 - m_{\text{HC}})/r$ on the intermediate tranche, and has a kink at D_{SN} . The edge x_U lies inside the domain of $V^{(1)}$. The retained drift changes across it by the net resource step

$$r(h_{\text{out}} - h_{\text{in}}) = \Pi(D) - \theta_{\text{max}} < 0, \quad (18)$$

because on crossing the edge the agent escapes the clawback $\Pi(D)$ but loses the support θ_{max} , which is larger. We write $\Delta h := h_{\text{out}} - h_{\text{in}}$ for this step capitalised at r .

- (c) (The two regions merge.) For $D \geq D_{\text{HC}}$ both channels are exhausted, $\Pi(D) = \theta_{\text{max}}$, and the zone drift $y_{\text{in}} - O_{\text{in}} + D - \theta_{\text{max}}$ coincides with the drift above, $y_{\text{out}} - O_{\text{out}} + D$; hence $h_{\text{in}} = h_{\text{out}}$ and State 1 is a single-region Merton problem in $z = x + h_{\text{out}}$. The borrowing constraint at $x = 0$ remains (Remark 5.4), so even then the value equals the power law only when that constraint is dropped.

Proof. Section S.1.1.

Assumption 5.2 (Well-posed zone). For every coverage $D \in [0, D_{\text{HC}}]$, $h_{\text{in}}(D) > 0$; equivalently, $y_{\text{in}} + D - \Pi(D) > O_{\text{in}}$. Since h_{in} is non-decreasing in D , the single check $y_{\text{in}} > O_{\text{in}}$ covers the whole range, and it holds in practice because means-tested schedules cap the out-of-pocket cost by capacity to pay. The condition keeps resources strictly positive on $[0, x_U]$; Remark 5.4 gives its consequence at $x = 0$. No counterpart is imposed above the zone, where the net flow may be negative, as it is at the calibration (Section S.1.4).

5.2 The dual construction across the asset-test edge

On each region the State 1 HJB equation (17) with $\kappa = 0$ is Merton’s equation, but Merton’s solution is not the disabled value. In discretionary wealth $z = x + h_{\bullet}$, with $h_{\bullet} = h_{\text{in}}$ on $[0, x_U]$ and $h_{\bullet} = h_{\text{out}}$ above, the HJB equation reduces to the time-homogeneous equation (Merton, 1971, Section 7 and eq. (82))

$$A_1 V^{(1)} = rz V_z^{(1)} - \frac{\text{SR}^2 (V_z^{(1)})^2}{2 V_{zz}^{(1)}} + \frac{\gamma}{1 - \gamma} (V_z^{(1)})^{1-1/\gamma}, \quad \text{SR}^2 = \frac{(\mu_R - r)^2}{\sigma_R^2}. \quad (19)$$

^[11]Each shift is capitalised at r , not $r + \lambda^{(1)}$: mortality is already borne by the discount A_1 in (17), and discounting the constant flow a second time would double-count it.

With a single region and no constraint, (19) is solved by the power law $\frac{\omega^{-\gamma}}{1-\gamma}z^{1-\gamma}$, with consumption–wealth ratio $\omega = \frac{1}{\gamma}[A_1 - (1-\gamma)(r + \frac{SR^2}{2\gamma})] > 0$.^[12] The disabled problem differs from Merton’s in two respects, and the power law fails in each. First, wealth crosses x_U in both directions, and the power law does not join there: it equals $\frac{\omega^{-\gamma}}{1-\gamma}(x_U + h_{\text{in}})^{1-\gamma}$ from the zone and $\frac{\omega^{-\gamma}}{1-\gamma}(x_U + h_{\text{out}})^{1-\gamma}$ from above when $x_U + h_{\text{out}} > 0$, and the two differ whenever $\Delta h \neq 0$ (18). At the calibration $x_U + h_{\text{out}} \leq 0$ at low coverage, and the power law from above is then undefined at the edge (Section S.1.4). Second, wealth may not fall below $x = 0$, and the power law is inadmissible there: it prescribes a positive dollar exposure $(\mu_R - r)h_{\text{in}}/(\gamma\sigma_R^2)$, whose diffusion takes wealth below zero, whereas the borrowing constraint requires $\wp^*(0) = 0$ (Remark 5.4).

The dual transform takes the marginal value of wealth as the variable in place of discretionary wealth, and makes Merton’s equation linear. Under the transform, the value is the power law, a particular solution, plus homogeneous corrections for the two failures. The two regions are solved jointly, with value matching and smooth fit at x_U , the borrowing constraint at $x = 0$ and transversality at infinity.^[13]

Three versions of the disabled problem appear below. The Merton benchmark drops both the asset-test edge and the borrowing constraint, and its value is the power law above. The unconstrained benchmark keeps the edge and drops the constraint, so wealth may fall below zero and the zone extends down to $-h_{\text{in}}$, where discretionary wealth runs out. It is solved in closed form (Section S.1.3), and Theorem 7.1 prices the implicit tax in it. The constrained problem keeps both; it is the problem behind the numbers of Section 9, where Figure 2 also draws the unconstrained benchmark for comparison. Readers interested in the economics can take the construction below as given and go on to Section 6.

Proposition 5.3 (Disabled value across the asset-test edge). *Let $\kappa = 0$ and $D < D_{\text{HC}}$. Write $s := 1 - 1/\gamma$, $z = x + h_{\bullet}$ for discretionary wealth, $\zeta := V_z^{(1)}$ for the dual variable, $\widehat{V}(\zeta) := V^{(1)} - z\zeta$ for the Legendre transform of the disabled value (He & Pagès, 1993, Theorem 3 and eq. (11); Dybvig & Liu, 2010, Theorem 3 and eqs. (42)–(43)), and $\Lambda := \gamma/((1-\gamma)\omega)$. Then:*

(a) *On each region, (19) is equivalent under the transform to the linear Cauchy–Euler equation*

$$\frac{SR^2}{2} \zeta^2 \widehat{V}_{\zeta\zeta} + (A_1 - r) \zeta \widehat{V}_{\zeta} - A_1 \widehat{V} + \frac{\gamma}{1-\gamma} \zeta^s = 0, \quad (20)$$

whose general solution is

$$\widehat{V}(\zeta) = \Lambda \zeta^s + G_+ \zeta^{p+} + G_- \zeta^{p-}, \quad (21)$$

^[12]Positivity requires no restriction beyond $\gamma > 1$: rearranging gives $\gamma\omega = A_1 + (\gamma - 1)(r + \frac{SR^2}{2\gamma})$, a sum of strictly positive terms, so the analogue of the Merton condition holds under Assumption 4.10.

^[13]The construction of Proposition 5.3 follows Dybvig and Liu (2010, 2011), with the exogenous asset-test edge, at which means-tested support is withdrawn, in place of their endogenous stopping boundary; Section S.2 compares the two constructions. The same transform, with conditions at zero wealth of the same form as (B1) and (B2) of Proposition 5.3(c), is used by Milevsky and Young (2007).

where $\Lambda\zeta^s$ is the particular solution and $p_{\pm}$ are the roots of the indicial quadratic

$$\text{Ind}(p) := \frac{\text{SR}^2}{2} p^2 + \left(A_1 - r - \frac{\text{SR}^2}{2} \right) p - A_1 = 0. \quad (22)$$

- (b) The roots are real and satisfy $p_- < 0$, $p_+ > 1$ and $p_- < s < p_+$; moreover $\text{Ind}(s) = -\omega$.
- (c) Let $\hat{V}^{\text{in}}$ on $[0, x_U]$ and $\hat{V}^{\text{out}}$ on (x_U, ∞) be solutions of the form (21), with coefficients $(G_{\pm}^{\text{in}})$ and $(G_{\pm}^{\text{out}})$. Transversality as $x \rightarrow \infty$ forces $G_-^{\text{out}} = 0$, while the zone keeps both modes (Remark 5.4). The remaining coefficients $(G_+^{\text{in}}, G_-^{\text{in}}, G_+^{\text{out}})$ and the marginal values $\zeta_U := V_x^{(1)}(x_U)$ and $\bar{\zeta} := V_x^{(1)}(0) < \infty$, with $\bar{\zeta} > \zeta_U$, are determined by five conditions. Three interface conditions (I1)–(I3) require the two branches to reach x_U at the same marginal value ζ_U and with the same value, and two boundary conditions hold at $x = 0$: (B1) $-\hat{V}_{\zeta}^{\text{in}}(\bar{\zeta}) = h_{\text{in}}$, the zone branch reaches the constraint, and (B2) $\hat{V}_{\zeta\zeta}^{\text{in}}(\bar{\zeta}) = 0$, the borrowing constraint of Remark 5.4. Given $(\zeta_U, \bar{\zeta})$, the three coefficients solve a linear system with nonvanishing determinant, and the pair $(\zeta_U, \bar{\zeta})$ then solves two scalar equations; Section S.1.2 writes out the five conditions and Section S.1.3 solves them.
- (d) A branch is strictly concave in wealth if and only if $\hat{V}_{\zeta\zeta} > 0$ on its interior. Strict concavity of the constructed function reduces to the endpoint conditions $\hat{V}_{\zeta\zeta}(\zeta_U) > 0$ on both branches and $\hat{V}_{\zeta\zeta}(\bar{\zeta}) < 0$ (Section S.2); in the unconstrained benchmark they reduce to a window of step sizes Δh that contains $\Delta h = 0$ in its interior.
- (e) Under the conditions of (d), the function defined parametrically by

$$z = -\hat{V}_{\zeta}(\zeta), \quad V^{(1)} = \hat{V}(\zeta) - \zeta \hat{V}_{\zeta}(\zeta), \quad x = z - h_{\bullet}, \quad (23)$$

is the disabled value function on $\mathbb{R}_+$. It is strictly concave, C^1 on $\mathbb{R}_+$ and C^2 on $(0, \infty)$ off x_U , where $V_{xx}^{(1)}$ jumps by an amount that the HJB equation fixes (Section S.1.5). The map $x \mapsto \zeta$ is a strictly decreasing bijection of $\mathbb{R}_+$ onto $(0, \bar{\zeta}]$, so $V^{(1)}(x)$ is recovered by a one-dimensional inversion.

Proof. Section S.1.2; the concavity certificate of part (d) is proved in Section S.2.

On the zone the borrowing constraint adds the ζ^{p+} mode: the unconstrained benchmark has $G_+^{\text{in}} = 0$ and only the ζ^{p-} mode, the value of crossing the edge (Section S.1.3). At zero wealth, where the borrowing constraint binds, the agent holds no risky asset and consumes less than the net flow rh_{in} , so wealth is pushed back up: zero wealth is a reflecting boundary, not an absorbing one. Remark 5.4 and Section S.4 derive these properties.

Remark 5.4 (Borrowing constraint at zero wealth). Under Assumption 5.2, $h_{\text{in}} > 0$, so at $x = 0$ the disabled agent still receives the net flow $rh_{\text{in}} > 0$, the marginal value $\bar{\zeta} = V_x^{(1)}(0)$ is finite and consumption $c^*(0) = \bar{\zeta}^{-1/\gamma}$ is positive. The constraint at $x = 0$ is therefore only a borrowing constraint: the agent cannot borrow against the capitalised flow h_{in} , since admissibility forbids $x_t < 0$ (Definition 4.8); the feedback controls of Lemma 4.12 are themselves admissible at $x = 0$ (Section S.4). Because $\wp^*(0) = 0$, the drift at $x = 0$ is $rh_{\text{in}} - c^*(0)$. At and above D_{HC} , where the regions merge and only the ζ^{p+} mode remains, eliminating its coefficient between (B1)

and (B2) gives $c^*(0) = \omega h_{\text{in}} \gamma(p_+ - 1)/(1 + \gamma(p_+ - 1))$, the constant-income case of El Karoui and Jeanblanc-Picqué (1998, Theorem 4.3) with A_1 in place of their discount rate, and hence

$$r h_{\text{in}} - c^*(0) = \frac{\text{SR}^2}{2} \frac{(p_+ - 1)(p_+ - s)}{1 + \gamma(p_+ - 1)} h_{\text{in}} > 0, \quad (24)$$

because $p_+ > 1$ and $p_+ > s$ (Proposition 5.3(b)). Below D_{HC} the ζ^{p-} mode adds a correction at $x = 0$; at the calibration its weight is negligible up to the coverage cap (Section S.8).

5.3 Standing assumptions and scope

The construction of Proposition 5.3 rests on two properties that we verify numerically rather than prove: the five conditions of its part (c) have a root, and the candidate value function built from that root is concave. We state both as standing assumptions and describe the case in which the second fails.

Assumption 5.5 (Existence and regularity of the root). *At every coverage up to D_{HC} , the five conditions of Proposition 5.3(c) have a root $(\zeta_U, \bar{\zeta})$ with $0 < \zeta_U < \bar{\zeta}$. The root is unique in the concave regime of Assumption 5.6; it is continuous in coverage, and C^1 within each tranche. At D_{HC} its one-sided limits equal the explicit single-region root of Section S.3. That section computes this root, the constant-income case of the closed-form solutions of He and Pagès (1993, p. 684) and El Karoui and Jeanblanc-Picqué (1998, Theorem 4.2), and continues it from there by the implicit function theorem, tranche by tranche. At the calibration D_{HC} lies beyond the coverage cap $\bar{D}$, so the continuation runs from outside the feasible range towards $D = 0$. The assumption is verified numerically from zero coverage to beyond D_{HC} (Section S.8). Under it $D \mapsto V^{(1)}(\cdot; D)$ is continuous, and C^1 off the kinks D_{SN} and D_{HC} ; Lemma 6.2 uses the smoothness, and the kink jump (31) uses the continuity.*

Assumption 5.6 (Concave regime). *Wherever the construction of Section S.1.2 is used, the parameters satisfy the concavity conditions of Proposition 5.3(d), the concavity certificate. In the unconstrained benchmark these conditions are a window of step sizes, which binds at its lower bound and is tightest at $D = 0$, where $\Delta h = -\theta_{\text{max}}/r$ (Section S.2). In the constrained problem the window is indicative only, because the boundary mode G_+^{in} , which the borrowing constraint adds, also enters the interface curvature; the operative check is the numerical one at the root $(\zeta_U, \bar{\zeta})$, verified across the feasible coverage range (Section S.8). For the State-0 value, concavity holds because leverage is unbounded, as for State 1: a convex stretch would make the supremum over π infinite (Lemma 4.12; Section S.2). The strict concavity that Lemma 4.12 needs is verified numerically at the calibration, including near x_U . There the cash step in health, Δ_y , makes the curvature of $V^{(0)}$ jump, so concavity has to be checked on each side of the edge (Section S.1.5).*

Where the certificate of Proposition 5.3(d) fails, the smooth-fit candidate is not the value function. With the other parameters at the calibration, the constrained certificate fails at zero coverage, on the zone side of the edge, when the Sharpe ratio falls below a floor, which Section 8 reports. The zone-side condition, $\hat{V}_{\zeta\bar{\zeta}}^{\text{in}}(\zeta_U) > 0$, says an agent just below the edge holds the risky asset, and it fails where that position falls to zero: the support lost above the

edge then outweighs the reward for the risk of crossing it. Below the floor, at low coverage, an agent on the zone stays there, and the value stays concave but its marginal value drops at the edge (Section S.2). Sections 6 and 7 use a continuous marginal value at the edge and are not carried to that case (Assumption 5.6).

With $\kappa > 0$ the reduction of Lemma 5.1 and the two-region structure survive, but superposition and the closed form do not, and the value is computed numerically (Section S.1.7). Section 6 uses the disabled value only through its regularity, which the closed form supplies at $\kappa = 0$; with a bequest motive we assume this regularity. The healthy value is computed numerically at every κ , and we assume its regularity in the same way.

Assumption 5.7 (Regularity without a closed form). *For every coverage D in the feasible range, the HJB system (16)–(17) has a solution with the following properties, for $j = 0$ at every $\kappa \geq 0$ and for $j = 1$ when $\kappa > 0$. The computed value approximates that solution. At $\kappa = 0$ the closed form and Assumptions 5.5 and 5.6 give them for $j = 1$.*

- (i) (Each state.) *For each such j , $V^{(j)} < 0$ and $V_x^{(j)} > 0$; $V^{(j)}$ is C^1 on $\mathbb{R}_+$ and C^2 on $(0, \infty)$ off x_U , with $V_{xx}^{(j)} < 0$ there and finite negative one-sided limits at x_U ; $V_x^{(j)}(0) < \infty$; under the optimal controls the drift of wealth at $x = 0$ is positive; at $x = 0$ the risky position vanishes like $\sqrt{x}$, as condition (B2) of Proposition 5.3(c) makes it for $j = 1$ at $\kappa = 0$; and the optimal controls grow at most linearly in wealth.*
- (ii) (In coverage.) *The map $D \mapsto (V^{(j)}, V_x^{(j)}, V_{xx}^{(j)})$ is continuous, and C^1 on each tranche with one-sided derivatives at the kink D_{SN} , locally uniformly in x , and the optimal controls are continuous in D uniformly on compact subsets of $\mathbb{R}_+$, zero wealth included; the coverage derivative $\partial_D V^{(j)}$, which is ψ for $j = 1$, is bounded, C^1 on $\mathbb{R}_+$ and C^2 off x_U .*
- (iii) (Near zero wealth.) *The coefficients of the reflected equation for the square root of healthy closed-loop wealth (Section S.4) are Lipschitz in that square root near zero.*

The bequest terms $\lambda^{(j)}B(x)$ of (16)–(17) do not depend on coverage, so they drop out of every derivative in D . With Assumption 5.7 for $j = 1$ in place of Assumptions 5.5 and 5.6, Lemma 6.2, Proposition 6.3 and Theorem 6.4 therefore hold at $\kappa > 0$ with the same proofs (Section S.1.7). Proposition 5.3, Remark 5.4 and Theorem 7.1 rest on the closed form and are stated at $\kappa = 0$ only. At $\kappa > 0$, Section S.8 checks part of the assumption on the computed values: the signs, the concavity and the drift at zero wealth, at bequest strengths that span the published estimates of Section 8; and, at the baseline strength, the rate $\sqrt{x}$ at which the risky position vanishes at zero wealth, the regularity of ψ in D and its jump at D_{SN} (31). The rest, including the regularity of the healthy value in D , the growth condition and, at $\kappa > 0$, the Lipschitz condition, is assumed.

Section 6 carries the disabled value $V^{(1)}$ and its coverage derivative to the inception problem.

6 Optimal coverage at inception

Coverage is chosen in the healthy state at $t = 0$, so the choice depends on the marginal benefit of coverage at inception, $\partial_D V^{(0)}$, and not directly on the disabled-state marginal benefit $\psi =$

$\partial_D V^{(1)}$. This section links the two. An envelope argument gives the disabled-state marginal benefit (Lemma 6.2); averaging it over the distribution of wealth at disability onset, the onset-wealth law, carries it back to inception (Proposition 6.3); and the demand equation follows (Theorem 6.4). The marginal benefit of coverage jumps upward at the kink D_{SN} , where the cash safety net runs out, so optimal coverage never sits at it but can jump past it as inception wealth rises.

6.1 From the disabled-state marginal benefit to the demand equation

Assumption 6.1 (D_{HC} out of reach). *The kink D_{HC} lies beyond the largest feasible coverage,*

$$D_{\text{max}} < D_{\text{HC}} = \frac{\Delta_O}{m_{\text{HC}}}, \quad (25)$$

for which $m_{\text{HC}} \bar{D} < \Delta_O$ is sufficient: at the averaged care rate, even the largest indemnity insurers offer cannot exhaust the public care subsidy. No restriction is placed on D_{SN} : both $D_{\text{max}} \leq D_{\text{SN}}$ and $D_{\text{SN}} < D_{\text{max}}$ are admitted, Theorem 6.4 covers both, and the calibration determines which holds.

Assumption 6.1 matches the UK and Australian side of the country contrast of Sections 3.2 and 3.3: there the clawback grows gradually and is still rising at the largest indemnity (Figure S.2), so every feasible coverage lies in the two-region regime of Lemma 5.1(b) and faces the implicit tax at the margin (Section 7).^[14]

Throughout this section we impose Assumptions 5.2, 6.1 and 5.7, together with Assumption 5.6 when $\kappa = 0$. Recall the disabled marginal value $\zeta(x) = V_x^{(1)}(x; D)$ of Proposition 5.3, and write

$$\delta(x; D) := \delta(D) \mathbf{1}\{x \leq x_U\} + \mathbf{1}\{x > x_U\}$$

for the after-test retention (13) of a marginal indemnity dollar, with the tranche retention $\delta(D) \in \{1 - m, 1 - m_{\text{HC}}\}$ of (11). A star marks an optimally controlled process: X^* denotes disabled wealth under the controls of Lemma 4.12, with drift $b^*(x) = \mu(x, \pi^*) - c^* - O(x) + D - \Pi(x, D)$ and diffusion $\sigma^*(x) = \pi^* \sigma_R x$, both piecewise across x_U ; its healthy counterpart $X^{(0)*}$ enters at Proposition 6.3. Subscripts on $\mathbb{E}$ give the initial wealth of the displayed process.

Lemma 6.2 (State-1 marginal benefit). *Let Assumption 5.5 hold, or Assumption 5.7 when $\kappa > 0$, and let $D \neq D_{\text{SN}}$, so that $D \mapsto V^{(1)}(\cdot; D)$ is C^1 at D ; in the unconstrained benchmark this smoothness is explicit. Then $\psi := \partial_D V^{(1)}(\cdot; D)$ exists, is C^1 on $\mathbb{R}_+$ and C^2 on $(0, \infty)$ off x_U , and the following hold.*

(a) ψ solves the linearisation of (17) along the optimal control,

$$A_1 \psi = b^* \psi_x + \frac{1}{2} (\sigma^*)^2 \psi_{xx} + \delta(x; D) \zeta(x), \quad (26)$$

with b^*, σ^* as above.

^[14]Where the clawback saturates quickly, as in the US and France (Section S.9.7), D_{HC} lies inside the feasible range; Section S.6 treats that case. At fixed channel rates and a fixed support cliff θ_{max} , Assumption 6.1 becomes tighter as the cash step Δ_y rises, since the care step $\Delta_O = \theta_{\text{max}} - \Delta_y$ then falls, and D_{HC} with it (Section S.1.1).

(b) ψ admits the Feynman–Kac representation

$$\psi(x) = \mathbb{E}_x \left[\int_0^\infty e^{-A_1 t} \delta(X_t^*; D) \zeta(X_t^*) dt \right] \geq 0. \quad (27)$$

In the constrained problem the integrand is bounded by $\delta(D)\bar{\zeta}$ on the zone and by $\bar{\zeta}$ above it, with $\delta(D)$ from (11) and $\bar{\zeta} = V_x^{(1)}(0) < \infty$, by Remark 5.4 and Proposition 5.3(c) when $\kappa = 0$ and by Assumption 5.7(i) when $\kappa > 0$.

(c) At x_U , ψ and ψ_x are continuous, with $\psi_x(x_U) = \dot{\zeta}_U := d\zeta_U/dD$.

Proof. Section S.5.

Representation (27) expresses ψ as an occupancy average: an expected discounted integral along the optimal wealth path, in which each instant is weighted by the marginal value ζ and by the retention, $\delta(D)$ while wealth is on the zone and 1 while it is above x_U . The marginal benefit $\psi(x)$ is therefore depressed only through the time that the optimal wealth path from x spends below x_U .

Proposition 6.3 (Marginal benefit of coverage at inception). *Let the assumptions of Lemma 6.2 hold, and Assumption 5.7 for $j = 0$. Fix $D \geq 0$, $D \neq D_{\text{SN}}$, and let $X^{(0)*}$ denote the optimally controlled healthy wealth of (16), with drift $b^{(0)*}(x) = \mu(x, \pi^*) - c^*$ and diffusion $\sigma^{(0)*}(x) = \pi^* \sigma_R x$, both piecewise across x_U through $y(x)$ and the controls (Section S.1.5). Then $\psi^{(0)} := \partial_D V^{(0)}(\cdot; D)$ solves $A_0 \psi^{(0)} = b^{(0)*} \psi_x^{(0)} + \frac{1}{2} (\sigma^{(0)*})^2 \psi_{xx}^{(0)} + \nu \psi$, and admits the representation*

$$\partial_D V^{(0)}(x; D) = \nu \int_0^\infty e^{-A_0 t} \mathbb{E}_x[\psi(X_t^{(0)*})] dt = \mathbb{E}_x \left[e^{-\beta \tau_{\text{morb}}} \mathbf{1}_{\{\tau_{\text{morb}} < \tau_{\dagger}^{(0)}\}} \psi(X_{\tau_{\text{morb}}}^{(0)*}) \right] \geq 0, \quad (28)$$

where $\tau_{\text{morb}} \sim \text{Exp}(\nu)$ is the onset time and $\tau_{\dagger}^{(0)} \sim \text{Exp}(\lambda^{(0)})$ the death time in health, both independent of the diffusion.

Proof. Section S.5.

Theorem 6.4 (Optimal coverage and the demand equation). *Under Assumptions 4.5, 4.10 and 5.7, and Assumption 5.5 when $\kappa = 0$, the inception objective $\Phi(D) := V^{(0)}(x_0 - a_D D; D)$ of (15) is continuous on $[0, D_{\text{max}}]$ and C^1 off the kink D_{SN} , with*

$$\Phi'(D) = -a_D V_x^{(0)}(x_0 - a_D D; D) + \partial_D V^{(0)}(x_0 - a_D D; D) \quad (29)$$

on each tranche. Any optimum interior to a tranche satisfies the demand equation

$$\underbrace{a_D V_x^{(0)}(\tilde{x}_0; D^*)}_{\text{marginal premium cost at inception}} = \underbrace{\nu \int_0^\infty e^{-A_0 t} \mathbb{E}_{\tilde{x}_0}[\psi(X_t^{(0)*})] dt}_{\text{onset-wealth-averaged marginal benefit}}, \quad \tilde{x}_0 := x_0 - a_D D^*, \quad (30)$$

with ψ given by Lemma 6.2 at D^* . When $D_{\text{SN}} < D_{\text{max}}$, the disabled-state marginal benefit ψ jumps upward across the kink, at every x , by the occupancy value of the safety-net channel,

$$\psi(x; D_{\text{SN}}^+) - \psi(x; D_{\text{SN}}^-) = m_{\text{SN}} \mathbb{E}_x \left[\int_0^\infty e^{-A_1 t} \mathbf{1}_{\{X_t^* \leq x_U\}} \zeta(X_t^*) dt \right] > 0, \quad (31)$$

with the law of X^* and the marginal value ζ those of the two-region problem at D_{SN} , common to both sides. The law of $X^{(0)*}$ at D_{SN} is common to both sides, so (28) carries this jump through and Φ' jumps upward across D_{SN} by $\nu \int_0^\infty e^{-A_0 t} \mathbb{E}_{x_0 - a_D D_{\text{SN}}}[\cdot] dt$ of (31). The kink is convex, so no optimum sits at D_{SN} and the global optimum is the better of the tranche optima. On the first tranche, $\Phi'(0) > 0$ rules out the corner $D^* = 0$; conversely, if Φ is quasi-concave there, $\Phi'(0) < 0$ makes $D = 0$ the best coverage on that tranche. Assumption 6.1 keeps the further kink at D_{HC} , convex at $\kappa = 0$ (Section S.6), out of the feasible range.

Proof. Section S.5.

The demand equation sets the marginal premium cost at inception equal to the disabled-state marginal benefit, averaged over the wealth the agent will hold when care begins. At the calibration, Φ is verified numerically to be quasi-concave on the first tranche. The low retention on the zone cannot by itself force zero coverage, because care that begins with wealth above the edge keeps a part of the benefit that does not depend on that retention (Section S.6).

Optimal coverage $D^*(x_0)$ varies with inception wealth in three ways. First, within a tranche it is continuous in x_0 where the tranche optimum is unique, and has no step at x_U . Second, coverage is lower where the onset-wealth law places more weight on the zone, and higher where that law places more weight above x_U . At the calibration and without a bequest motive, coverage increases over the whole range of x_0 drawn in Figure 3(a). Third, where D_{SN} lies inside the feasible range, the global optimum can switch tranche as wealth rises, and $D^*(x_0)$ then jumps from low first-tranche coverage to coverage past the kink.^[15]

6.2 Numerical scheme

The scheme has two nested levels. For a trial coverage D , the inner level solves the two-region State 1 problem for $V^{(1)}(\cdot; D)$ and ζ , by Proposition 5.3 at $\kappa = 0$ and numerically with a bequest motive (Section S.1.7), and the marginal benefit ψ from its linear equation (26). Solving (16) by policy iteration on a finite-difference grid that has a grid point at x_U , with forcing $\nu V^{(1)}$, gives $V^{(0)}(\cdot; D)$, and the linear equation of Proposition 6.3, with source $\nu\psi$, gives $\partial_D V^{(0)}$, so the solver never simulates the onset-wealth law. The outer level maximises the objective Φ over the scalar $D \in [0, D_{\text{max}}]$ by one-dimensional root-finding on (29), separately on each side of D_{SN} (Theorem 6.4). The representation (27), finite differences in D , Monte Carlo simulation and, in the unconstrained benchmark, the closed form of Theorem 7.1 below serve as checks (Section S.8).

No finite difference or root-finding step in D is taken across a kink or across $D_U = (x_0 - x_U)/a_D$, the coverage at which post-premium wealth lies exactly at the edge; Section S.6 lists these points. If a retiree who starts just above the edge buys more coverage than D_U , the premium takes their wealth below the edge, where the withdrawn support is available again; near the edge the premium therefore partly pays for itself. The objective Φ nevertheless stays C^1 at D_U , because $V_x^{(0)}$ and $\partial_D V^{(0)}$ are continuous across the edge (Assumption 5.7(i)–(ii); Section S.1.5), and only the curvature of Φ jumps. The coverage D_U therefore matters for the computation rather than for the optimum, and it is feasible only for $x_0 \in (x_U, x_U + a_D \bar{D})$,

^[15]An empirical study could test for this jump in the demand profile.

which is (20, 20.9) at the calibration (Section S.6). Grids, tolerances, boundary treatment, convergence and run times are reported in Section S.8.

Section 7 evaluates the disabled-state marginal benefit ψ in closed form, which turns the implicit tax on the zone into an effective loading on a marginal coverage dollar.

7 The implicit tax priced

On the zone, below the asset-test edge, the implicit tax raises the loading on coverage above the market loading η , by an amount that steps down at the kink D_{SN} , where the cash safety net runs out (Theorem 7.1). This section computes the disabled-state marginal benefit $\psi = \partial_D V^{(1)}$ of Lemma 6.2 in closed form, without a bequest motive, in the setting and notation of Section 6. The marginal benefit is the value of a marginal indemnity dollar to the disabled agent, and it is the input that the demand equation (30) averages over the onset-wealth law.^[16] Parts (a) to (c) of Theorem 7.1 below hold in the unconstrained benchmark of Section 5.2, in which the retiree may borrow against future income. Part (d) gives the marginal benefit when borrowing is ruled out, the problem behind the numbers of Section 9.

Theorem 7.1 (The implicit tax priced: closed-form marginal benefit). *In the unconstrained benchmark, off the kink D_{SN} , with $\dot{\cdot} := d/dD$ and the tranche pair $(\Pi'(D), \delta(D))$ equal to $(m, 1 - m)$ on $[0, D_{\text{SN}})$ and to $(m_{\text{HC}}, 1 - m_{\text{HC}})$ on $(D_{\text{SN}}, D_{\text{HC}})$,*

$$\psi(x) = \frac{\delta(D)}{r} \zeta(x) + \dot{G}_{-}^{\text{in}} \zeta(x)^{p_{-}} \quad \text{on } (-h_{\text{in}}, x_U], \quad \psi(x) = \frac{1}{r} \zeta(x) + \dot{G}_{+}^{\text{out}} \zeta(x)^{p_{+}} \quad \text{above,} \quad (32)$$

with $p_{\pm}$ the indicial roots of Proposition 5.3, and with coefficient derivatives $\dot{G}_{-}^{\text{in}}$ and $\dot{G}_{+}^{\text{out}}$ obtained in closed form by differentiating the benchmark coefficients of Section S.1.3 in D , as written out in Section S.5. Moreover, in the benchmark:

- (a) (Uniqueness.) (32) is the unique solution of (26) with ψ and ψ_x continuous at x_U (Lemma 6.2), ψ/ζ bounded as $\zeta \rightarrow \infty$ and $\psi \rightarrow 0$ as $x \rightarrow \infty$.
- (b) (Signs.) $\dot{G}_{-}^{\text{in}} > 0$: on the zone the marginal benefit strictly exceeds $\delta(D)\zeta/r$, the retained flow $\delta(D)$ capitalised at r and multiplied by the marginal value ζ ; the excess is the value of the option of escaping the zone, and deep in the zone the option vanishes. Above the zone, $\psi - \zeta/r = \dot{G}_{+}^{\text{out}} \zeta^{p_{+}}$ is the net of two effects: the loss from re-entering the zone, and the gain from the upward push that ζ receives whenever wealth touches x_U , because the curvature of $V^{(1)}$ jumps there (Proposition 5.3(e); the local-time term in the dual dynamics of Section S.5). Its sign is not determined in general, and the gap vanishes as $x \rightarrow \infty$.
- (c) (Money's worth.) The money's worth, the value of a marginal indemnity dollar relative to ζ/r , the value of the same dollar kept in full,

$$\text{MW}(x) := \frac{r \psi(x)}{\zeta(x)} = \begin{cases} \delta(D) + r \dot{G}_{-}^{\text{in}} \zeta(x)^{p_{-}-1}, & x \leq x_U, \\ 1 + r \dot{G}_{+}^{\text{out}} \zeta(x)^{p_{+}-1}, & x > x_U, \end{cases} \quad (33)$$

^[16]The computation works in the dual variable ζ of Proposition 5.3, in which the equation for ψ has explicit coefficients (Section S.5).

is continuous, strictly increasing in wealth on the zone, equal to $\delta(D)$ in the deep-zone limit and to 1 at infinity; above the zone it is monotone with the sign of $-\dot{G}_+^{\text{out}}$. Across D_{SN} every displayed object steps with the tranche pair, and ψ jumps by (31).

- (d) (The borrowing constraint.) In the constrained problem instead, under Assumption 5.5, wealth cannot fall below zero, so ζ cannot rise above $\bar{\zeta}$, the marginal value at zero wealth (Proposition 5.3(c)). On the zone (32) then holds with the constrained coefficients, whose D -derivatives we still write $\dot{G}_{\pm}^{\text{in}}$, and with the further term $\dot{G}_+^{\text{in}}\zeta(x)^{p_+}$; conditions (B1) and (B2) fix it by making ψ , read as a function of ζ , flat at $\bar{\zeta}$. The sign $\dot{G}_-^{\text{in}} > 0$ still holds, and on the zone

$$\text{MW}(x) = \delta(D) \left[1 - \frac{1}{p_+} \left(\frac{\zeta(x)}{\bar{\zeta}} \right)^{p_+-1} \right] + r \dot{G}_-^{\text{in}} \zeta(x)^{p_- - 1} \left[1 - \frac{p_-}{p_+} \left(\frac{\zeta(x)}{\bar{\zeta}} \right)^{p_+ - p_-} \right].$$

The first term is the retention $\delta(D)$ times a factor below one that rises with wealth from $1 - 1/p_+$ at zero wealth; the second is the escape option, as in part (b). Above D_{HC} , where $\delta = 1$ and the two regions merge, the second term is absent and $\text{MW}(0) = 1 - 1/p_+$.

Proof. Section S.5.

Proposition 4.6 prices the premium on the gross indemnity, at retention one, whereas by Theorem 7.1 a marginal indemnity dollar received on the zone is worth the fraction $\text{MW}(x)$ of its full value, the money's worth (33). The effective loading on a marginal dollar of coverage paid out on the zone is therefore $(1 + \eta)/\text{MW}(x) - 1$, which varies with wealth. In the unconstrained benchmark its supremum on the first tranche is $(1 + \eta)/(1 - m) - 1$, approached deep in the zone. This supremum is the loading a static calculation gives, a premium of $1 + \eta$ for a dollar of which the test leaves $1 - m$; on any tranche we call $(1 + \eta)/\delta(D) - 1$ the static loading. Once the cash step is used up at D_{SN} the static loading falls to $(1 + \eta)/(1 - m_{\text{HC}}) - 1$, and it reaches the market loading η only at D_{HC} , which Assumption 6.1 places out of reach. What the model adds is how far the effective loading departs from the static loading: in the unconstrained benchmark the option to leave the zone lowers the effective loading near the edge (Theorem 7.1(b)), and the borrowing constraint raises it (part (d)).

Part (d) separates what the borrowing constraint does from what the means test does. Because the constraint binds at zero wealth, a dollar now is worth more there than a dollar that coverage pays later, as Bick et al. (2013) observe of future income. Hence the constraint holds the first term of the money's worth in part (d) below $\delta(D)$ at every wealth, most at low wealth, and at retention one, with no escape option, it lowers the money's worth itself. The fraction of the retention it removes at zero wealth, $1/p_+$, depends on the interest rate, the Sharpe ratio and $A_1 = \beta + \lambda^{(1)}$, the disabled agent's discount rate with mortality added, and not on the means test, on coverage or on risk aversion. On the zone the money's worth therefore lies below $\delta(D)$ exactly where the escape option is worth less than $\delta(D)(\zeta/\bar{\zeta})^{p_+-1}/p_+$, the part of the retention that the constraint removes. At the calibration the escape option is the smaller at every wealth computed on the zone (Section 9.1; Section S.8), so there the money's worth lies below $\delta(D)$, the reverse of part (b), and the static loadings above are lower bounds on the constrained effective loading. The upward jump of ψ at D_{SN} survives, because Theorem 6.4 establishes it in the constrained problem.

8 Calibration

Sections 4 to 7 establish signs, orderings and closed forms that hold for every parameter value satisfying their assumptions. This section sets the baseline calibration under which the numerical results of Sections 9 and 10 are computed.

We use two baseline specifications that differ only in the bequest parameters. The no-bequest baseline sets $\kappa = 0$, where Sections 5 and 7 give closed forms, and covers the no-bequest results. The bequest baseline adds the pair $(\kappa, \underline{x}) = (1000, 7.42)$, requires a numerical solution (Section S.1.7), and covers the bequest results. Table 1 reports the input parameters in Panel A along with their sources. Panel B lists the derived quantities, among them the two kinks and the coverage cap, which lies between them.

Table 1. Baseline calibration and derived quantities.

Panel A: input parameters	Value	Source
<i>Preferences and market</i>		
CRRA γ	3.00	Brown and Finkelstein (2008)
Discount rate β	0.030	Brown and Finkelstein (2008)
Bequest weight κ	1000	moderate, bracketed by Ameriks et al. (2020)
in years of consumption, $\kappa^{1/\gamma}$	10	and Ameriks et al. (2011)
Bequest shift $\underline{x}$	7.42	from the threshold of Andréasson et al. (2017)
Riskless rate r (real)	0.023	Jordà et al. (2019), UK–Australia
Risky drift μ_R , volatility σ_R	0.109, 0.200	Sharpe ratio 0.43, bracketed by Jordà et al. (2019)
Market loading η	0.10	below Brown and Finkelstein (2007)
<i>Morbidity and mortality (set to three moments)</i>		
Morbidity ν	0.0223	lifetime care probability 0.45
Mortality $\lambda^{(0)}$, excess α	0.0273, 0.306	life expectancy 21.5 years, care spell 3.0 years
<i>Means-tested support and care costs (UK and Australian schedules, Section S.9)</i>		
Channel rates $m_{\text{SN}}, m_{\text{HC}}$	0.44, 0.10	averaged schedule: slope and split
Cash step Δ_y , care step Δ_O	0.25, 0.60	UK Pension Credit top-up; two-country mean
Out-of-pocket cost on the zone O_{in}	0.35	model choice, above both countries' levels
Asset-test edge x_U	20.0	mean wealth where each country's clawback ends
Income on the zone y_{in}	0.625	Pension Credit guarantee, by the choice of unit
Panel B: derived quantities	Value	Role
Sharpe ratio SR (squared, SR^2)	0.43 (0.1849)	$(\mu_R - r)/\sigma_R$, Section 4.4
Merton consumption–wealth ratio ω	0.157	Section 5.2
Indicial roots p_-, p_+	−3.733, 1.053	modes of the dual equation, Proposition 5.3
Total channel rate $m = m_{\text{SN}} + m_{\text{HC}}$	0.54	first-tranche rate
Support cliff $\theta_{\text{max}} = \Delta_y + \Delta_O$	0.85	(7)
Out-of-pocket cost above the zone	0.95	(6)
$O_{\text{out}} = O_{\text{in}} + \Delta_O$		
Income above the zone $y_{\text{out}} = y_{\text{in}} - \Delta_y$	0.375	(4)
Coverage cap $\bar{D} = O_{\text{out}}$	0.95	Assumption 4.5 at its bound
Kink $D_{\text{SN}} = \Delta_y/m_{\text{SN}}$	0.568	cash safety net used up; below $\bar{D}$, so feasible
Kink $D_{\text{HC}} = \Delta_O/m_{\text{HC}}$	6.000	care subsidy used up; above $\bar{D}$: Assumption 6.1
Tranche retention $1 - m, 1 - m_{\text{HC}}$	0.46, 0.90	(11)
Static loading $(1 + \eta)/\delta(D) - 1$	139%, 22%, 10%	first, intermediate, saturated; Section 7
Unit premium a_D	0.9483	Proposition 4.6

Notes: Units are those of Section 3: flows per year and wealth in multiples of the annual cost of care. The health intensities reproduce the three moments in the Source column. The bequest rows enter the $\kappa > 0$ results only (Figure 3, Table 2 and Panel B of Table S.7); the no-bequest baseline sets $\kappa = 0$ and drops $\underline{x}$. Panel B is computed from the rows common to both baselines and checked against the values used in the computations.

Relative risk aversion and the discount rate follow Brown and Finkelstein (2008), who set $\gamma = 3$ alongside a 3% real interest rate and discount rate for a single retiree facing care risk;

Ameriks et al. (2011) take the same γ . Published estimates of the bequest strength, as a marginal propensity to bequeath $\kappa^{1/\gamma}/(1 + \kappa^{1/\gamma})$, are about 0.48 in Ameriks et al. (2020), 0.96 in Andréasson et al. (2017) and 0.98 in Ameriks et al. (2011). We take a moderate value as the baseline, $\kappa^{1/\gamma} = 10$: when a dollar bequeathed is worth as much as a dollar consumed, the estate plus the shift $\underline{x}$ equals ten years of consumption. The propensity is then 0.91, and Section 9.2 reports how the coverage results change across this range. Only the bequest threshold is taken from Andréasson et al. (2017): the motive is inactive below a consumption rate of $\underline{x}/\kappa^{1/\gamma} = 0.74$, where a dollar consumed is worth more than the first dollar bequeathed (Section S.9.6).

The riskless rate, 2.3%, is the UK–Australia average real return on long-term government bonds over the full samples of Jordà et al. (2019). The risky asset is the retiree’s total risky wealth, equity and home equity, with home equity released without friction (Section 3); its Sharpe ratio of 0.43, an expected excess return of 8.6% on a volatility of 20%, lies between the equity-only and the equity-and-housing ratios of the two countries in the same data (Section S.9.6). Below a Sharpe ratio of 0.398 the concavity certificate of Assumption 5.6 fails at zero coverage (Section S.8.9; Section 5.3). The calibrated ratio is 0.03 above this floor; Section 9.3 gives the margins in the risky return and its volatility. On equity alone the ratio is 0.32 for the UK, below the floor, and 0.40 for Australia; with home equity in the risky asset, Proposition 5.3 applies at the calibration.^[17]

We set the three health transition intensities ν , $\lambda^{(0)}$ and α to reproduce, for a retiree aged 65 at inception, a lifetime probability of needing care of 0.45, a mean care spell of 3.0 years and a life expectancy of 21.5 years. The first two lie between the values Brown and Finkelstein (2008) report for 65-year-old men and women: a probability of ever using LTC of 0.40 and 0.54, and mean years in care among users of 2.9 and 4.2. The third lies between the life expectancies at 65 of Australian men and women, 20.3 and 22.9 years (Australian Government Actuary, 2024, pp. 27, 29).

The means-test parameters are taken from the UK and Australian schedules of Section S.9. The asset-test edge, $x_U = 20$, is the mean of the wealths at which the two countries’ clawbacks end: a retiree at the edge holds the ten-unit home and ten units of liquid wealth. The zone income y_{in} is the income the UK safety net guarantees, which is 0.625 by the choice of unit (Section 3). The cash step $\Delta_y = 0.25$ is the UK retiree’s cash support (Section 3.2), and the care step $\Delta_O = 0.60$ the two-country mean; Section S.9.4 explains why. The channel rates come from the averaged schedule, the two countries’ schedules averaged at each wealth. The slope of the averaged net clawback in coverage gives $m = 0.54$, and the safety-net channel’s share of the first-tranche clawback splits m into $m_{\text{SN}} = 0.44$ and $m_{\text{HC}} = 0.10$ (Section S.9). Only the first tranche identifies this split, because the model’s two channels claw back from the first dollar while the UK’s act one after the other. The out-of-pocket cost on the zone is not taken from the schedules: we set $O_{\text{in}} = 0.35$, above both countries’ levels of 0.21 in Australia and 0.18 in the UK, so that disability on the zone is costly. Without a bequest motive coverage at low wealth responds to this input more than to any other, and rises with it (Table S.7).

The market loading $\eta = 0.10$ is below the loads Brown and Finkelstein (2007) estimate for

^[17]At Australia’s equity-only ratio, 0.40, optimal coverage at the four inception wealths of Table S.11 ranges from 0.15 to 0.31 of the annual cost of care without a bequest motive, against 0.13 to 0.31 at the calibration, and is the same to three decimals with one.

a policy bought at 65 and held until death, 0.18 in their base case and 0.11 to 0.27 across their assumptions; Section 9.3 reports coverage at $\eta = 0.18$ and 0.27. The coverage cap is not a free parameter: Assumption 4.5 bounds it by the unsupported care cost, and we impose that bound with equality, $\bar{D} = O_{\text{out}} = 0.95$. The cap therefore falls short of the annual cost of care, which is 1 in the units of Section 3. The calibrated $O_{\text{out}} = O_{\text{in}} + \Delta_O = 0.95$ lies between the two countries' unsupported care costs: 1.00 in the UK, and 0.61 in Australia, which funds part of the care budget at every level of means (Section S.9).

9 Numerical results

This section reports the numerical results at the calibration of Table 1. We first restate the setting. All amounts are measured in the normalised units of Section 3, in which the total annual cost of care is 1. The retiree buys coverage once, while healthy. The means test applies on the zone, the range of total wealth, home included, below the asset-test edge $x_U = 20$. There, in disability, it claws back 54 cents of each marginal dollar of coverage on the first tranche, up to the kink $D_{\text{SN}} = 0.568$, where the cash safety net runs out (Assumption 4.7). Above that kink, on the intermediate tranche, it claws back 10 cents. The clawback would stop only at the kink $D_{\text{HC}} = 6$, where the public care subsidy runs out, far above the coverage cap $\bar{D} = 0.95$, the largest coverage the market offers.

The coverage results are reported without a bequest motive and with one, the two baseline specifications of Section 8. Section 9.1 measures what the means test does to the value of a marginal coverage dollar, Section 9.2 reports what the retiree then buys, and Section 9.3 checks the coverage results against a $\pm 5\%$ change in each of ten calibrated parameters.

9.1 The implicit tax, tranche by tranche

On the first tranche of the zone the means test more than doubles the premium per dollar of payout the retiree keeps, to 2.39, a static loading of 139% against the market loading of 10%. What a kept dollar is worth to a disabled retiree depends on wealth, and the borrowing constraint lowers it, most for the lowest-wealth retirees (Figure 2).

The numbers that follow rest on one measure, the money's worth $\text{MW}(x)$ of Theorem 7.1(c). It is the value to a disabled retiree with wealth x of one more unit of indemnity flow, as a share of $1/r$, which is what the flow is worth to a retiree who can borrow against it at the riskless rate, the debt dying with them, as in the unconstrained benchmark. Without a means test or a borrowing constraint it equals one; a money's worth of 0.46 means that the retiree keeps 46 cents of the value of each marginal coverage dollar. The effective loading $(1 + \eta)/\text{MW}(x) - 1$ is the static loading with the retention replaced by the money's worth, so it is measured against the benchmark value $1/r$; it shows how far the borrowing constraint and the edge move a kept dollar's value from that benchmark. Figure 2 plots the money's worth in panel (a) and the effective loading in panel (b), on the first and the intermediate tranche, with the borrowing constraint and in the unconstrained benchmark, which removes it.

In the unconstrained benchmark the money's worth is close to the tranche retention over most of the zone, 0.46 on the first tranche and 0.90 on the intermediate one (panel (a) of




Figure 2. Money's worth of a marginal coverage dollar and the effective loading on it, by tranche.

Notes: Panel (a) plots the money's worth $MW(x) = r\psi(x)/\zeta(x)$ of Theorem 7.1(c) and panel (b) the effective loading $(1 + \eta)/MW(x) - 1$, on a logarithmic axis, against the wealth x of a disabled retiree. In panel (a), ψ is the disabled-state marginal benefit of coverage and ζ the marginal value of wealth when disabled; in panel (b), η is the market loading. The first tranche is drawn at $D = 0.30$ and the intermediate tranche at $D = 0.80$. Each is drawn both in the constrained solution, where the retiree cannot borrow, and in the unconstrained benchmark; other parameters as in Table 1. Grey rules mark the retentions $1 - m = 0.46$ and $1 - m_{HC} = 0.90$ and the dash-dot rule the no-test level 1, also the benchmark value on the saturated tranche, which is not drawn; panel (b) shows the three levels as static loadings. The wealth axis is linear on the zone, up to the asset-test edge x_U , and logarithmic above it.

Figure 2), and rises above it only near the edge, where the retiree may leave the zone (Theorem 7.1(b); Section S.8). The borrowing constraint lowers it below the retention at every wealth: on the first tranche it is 0.191 at the edge and 0.023 at zero wealth, an effective loading of 4689% (panel (b)).^[18]

Against the unconstrained benchmark, most of the effective loading at zero wealth comes from the borrowing constraint, not from the means test. Above D_{HC} , which coverage cannot reach at the calibration, the means test claws back nothing, yet the money's worth at zero wealth is only 0.0499, which is $1 - 1/p_+$ at the indicial root $p_+ = 1.053$ of Table 1. The constraint therefore removes about 95% of the value of a coverage dollar. On every tranche it removes this same share of what the means test leaves, because the escape option, the value of the option to leave the zone, is negligible at zero wealth (Theorem 7.1(d); Section S.8). On the first tranche the means test then multiplies what is left by the retention, $0.46 \times 0.0499 \approx 0.023$, which is the value above.

The premium is priced on a different base, the flow's actuarial value $1/(r + \lambda^{(1)})$ (Proposition 4.6), which is smaller than $1/r$ by the factor $(r + \lambda^{(1)})/r = 15.5$, because it also discounts at the disabled retiree's mortality intensity. Against this value the loss is smaller: at zero wealth the dollar is still worth about 77% of it, so on this base the means test, which leaves 46%, removes more than the constraint. Most of the 95% is therefore the benchmark's borrowing at the riskless rate against a flow that stops at death; the rest is the constraint binding at zero wealth, where a dollar now is worth more than a dollar paid later. On the first tranche the

^[18]The value at zero wealth, $MW(0) = r\psi(0)/\bar{\zeta}$, is finite because the marginal value of wealth there, $\bar{\zeta} = V_x^{(1)}(0)$, is finite (Remark 5.4).

value at zero wealth is about 36% of the actuarial value: the premium is roughly three times what a marginal coverage dollar is worth to the retiree there. Above a wealth of about 1 on the first tranche, and about 0.1 on the intermediate one, a disabled retiree values a marginal coverage dollar at more than its premium. Coverage is chosen in health, however, where the demand equation weighs that value, averaged over the wealth held when care begins, against the premium at the healthy retiree's marginal value of wealth (Theorem 6.4).

9.2 Optimal coverage, with and without a bequest motive

Optimal coverage is low at every wealth drawn, and a bequest motive lowers it further (Figure 3).




Figure 3. Optimal LTCI coverage by inception wealth and by strength of the bequest motive. *Notes:* Panel (a) plots optimal coverage D^* against inception wealth x_0 , without a bequest motive and with the bequest pair of Table 1, the other parameters as there. Optima are computed at thirty wealths from $x_0 = 4$ to $x_0 = 130$ and drawn up to $x_0 = 118$, the highest at which the grid ranks the optima on the two sides of the kink. The top axis marks the asset-test edge x_U , and the right axis the kink D_{SN} , where the cash safety net runs out. Panel (b) plots D^* at three inception wealths, $x_0 = 46, 80$ and 105 , relative to its level without a bequest motive ($\kappa = 0$, panel (a)), against the propensity to bequeath $\kappa^{1/\gamma}/(1 + \kappa^{1/\gamma})$ on a log-odds scale, with the threshold $\underline{x}/\kappa^{1/\gamma}$ held at 0.74; dotted rules mark the three published estimates of Section 8, and the dash-dot rule the baseline.

First, without a bequest motive coverage rises with wealth and stays on the first tranche (panel (a)). It is 0.13 at $x_0 = 4$, 0.25 at the asset-test edge and 0.31 at $x_0 = 105$: under a third of the annual cost of care, and below the kink $D_{SN} = 0.568$ at every wealth drawn. Just above the drawn range, at $x_0 = 130$, the best coverage past D_{SN} , 0.75, is worth almost exactly as much as the best coverage below it (Section S.7), so demand can jump there, as Theorem 6.4 allows.

Second, a moderate bequest motive almost removes coverage (panel (a)). With the bequest pair of Table 1 the retiree buys no coverage below $x_0 = 61$, and above it buys little: 0.022 at $x_0 = 80$ and 0.039 at $x_0 = 105$. Coverage enlarges the estate of a retiree who dies in care but shrinks, by the premium, the estate of one who never needs care. At the calibration the second effect wins: if the motive valued the estate only at a death in care, coverage would rise to the cap from $x_0 = 61$; if only at a death in health, it would fall to at most 0.013 (Section S.7).

Third, coverage falls as the motive strengthens across its published range (panel (b)): at the propensity of 0.48 of Ameriks et al. (2020) it is about nine-tenths of its level without the

motive, and at the propensities of 0.96 and 0.98 of Andréasson et al. (2017) and Ameriks et al. (2011) none is bought at any wealth computed (Section S.7). Near the top of that range the fall depends on wealth: coverage reaches zero at a propensity of 0.90 at $x_0 = 46$, 0.93 at $x_0 = 80$ and 0.94 at $x_0 = 105$, so a richer retiree keeps some coverage under a stronger motive. At $x_0 = 105$ the fall slows before coverage reaches zero, because both effects of coverage on the estate grow with the motive, and the marginal benefit of the first coverage dollar almost keeps pace with its marginal premium cost (Section S.7).

9.3 Sensitivity

A $\pm 5\%$ change in any one of ten parameters of Table 1 moves coverage little, except where the change carries coverage past the kink D_{SN} (Table S.7, at seven inception wealths from $x_0 = 15$ to $x_0 = 105$). With a bequest motive coverage stays low, at most 0.09 in every cell. Without one it stays on the first tranche at the five lower wealths, but at the two highest a 5% move in one of four parameters lifts it past the kink, the jump that Theorem 6.4 allows. A narrower zone and a lower channel rate each raise coverage, except without a bequest motive at $x_0 = 105$, where a 5% higher m_{SN} lets coverage cross D_{SN} ; the pattern fits the cross-country hypothesis of Section 3.3.

These variations stay inside the region where the conditions behind the results hold (Section S.8.9): the tightest margins are the risky return, at -6% , and its volatility, at $+8\%$, both bounded by concavity. The channel split, the input the country schedules determine least well (Section S.9), has more room, with m_{SN} admitting -40% and m_{HC} admitting $+459\%$. Three of the alternative estimates of the channel rates reported there, those that fit the averaged schedule over the whole coverage range, put m_{SN} below the range it admits. The kink D_{SN} then lies above the coverage cap, so the jump past it that Theorem 6.4 allows could not occur. Of these three estimates, the two that keep the total rate leave the first-tranche optima unchanged (Section S.9).

Table S.11 sets one input at a time to a value from its source, except that the care probability and care spell move together. At five of the six values coverage stays under a third without a bequest motive and near zero with one, and exempting payouts raises it. Among those five, at an out-of-pocket cost on the zone of 0.20, between the countries' levels, the means test leaves almost no coverage. The sixth is the higher care probability and care spell of Section 8, 0.54 and 4.2 years. There, coverage under the means test jumps past the kink D_{SN} and exceeds that bought when payouts are exempted, above the edge without a bequest motive and at the highest wealth with one. The payout retained after the test is nevertheless larger when payouts are exempted, in every cell of Table S.11. Both follow from the retentions: a coverage dollar keeps 90 cents past the kink and 46 on the first tranche, so a retiree buying past the kink needs more coverage for a given retained payout, and buys a smaller one because each retained dollar now costs more premium.

10 Means-test design and optimal coverage

At the calibration, exempting LTCI payouts from the income tests raises optimal coverage at every wealth computed, with a bequest motive and without one. Abolishing the asset test while keeping the income tests lowers coverage without a bequest motive except at the lowest wealths, and with one no coverage is bought. Changing the income tests costs the government about as much as the baseline, and without a bequest motive the retiree values the exemption of payouts at several times its extra cost; extending the support to every retiree costs the most.

10.1 Coverage, public cost and value under eight means tests

We compare the baseline means test of Table 1 with seven alternatives that change only the public schedule of Sections 4.2 and 4.3. The retiree, the market, the health process and the unit premium a_D are as in Table 1, and the coverage cap is set, as in the calibration, at the unsupported care cost, $\bar{D} = O_{\text{out}}$, which is 0.95 in the baseline. Designs 2 to 5 keep the asset test and change the income tests: Design 2 exempts payouts in both channels, Design 3 halves both channel rates, and Design 4 exempts payouts in the home-care channel only. Design 5 moves the whole first-tranche rate m into the safety-net channel: the first coverage dollar is clawed back at the baseline's rate, but the clawback stops at the cash step Δ_y , so the kink D_{SN} , where the cash safety net runs out, moves to $\Delta_y/m = 0.463$, against 0.568 in the baseline. Designs 6 to 8 change the asset test or the support itself: Design 6 removes the edge, so that every retiree faces the support and the clawback of the zone, the wealth range below the baseline's edge. Design 7 gives every retiree that support and withdraws nothing, which lowers the cap to the out-of-pocket cost on the zone, $\bar{D} = O_{\text{in}} = 0.35$. Design 8 makes care free on the zone, with no cash safety net and nothing withdrawn.

Panel A of Table 2 reports optimal coverage under each design at four inception wealths, without a bequest motive and with the bequest pair of Table 1; Panel B gives the expected public cost of each design at that coverage.^[19] The ranges quoted below are taken over the ten wealths of Section S.8.11, before rounding.

Exempting payouts (Design 2) raises coverage at every wealth (Panel A), by 0.10 to 0.18 without a bequest motive and, with one, from the near-zero coverage that the clawback and the bequest motive together leave in the baseline; with only the clawback gone, the design isolates the implicit tax. Capping the clawback at the cash step (Design 5) leaves coverage unchanged up to the edge, and at every wealth with a bequest motive; above the edge without one it moves coverage past the kink D_{SN} , the jump of Theorem 6.4. Abolishing the asset test (Design 6) lowers coverage without a bequest motive at every wealth from $x_0 = 16$ upward, and free care on the zone (Design 8) removes demand at and below the edge. Designs 5 and 6 bear on the cross-country hypothesis of Section 3.3, that a market can form where the clawback covers a narrow range of wealth, unlike Design 6, and stops growing at modest coverage, as in Design 5.

^[19]Every design satisfies Assumption 5.2; without a bequest motive each passes the concavity certificate of Assumption 5.6 across its feasible coverage range, and with one the designs rest on Assumption 5.7 in both states, which Section S.8 checks at the baseline schedule only. A direct evaluation of the inception objective Φ of Theorem 6.4 at 39 coverages and a finer grid confirm every optimum but do not rank the two tranche optima of Design 4 at $x_0 = 105$.

Table 2. Optimal coverage and expected public cost under eight means tests.

Design	Loading	Without a bequest motive				With a bequest motive			
		$x_0 = 4$	22	46	105	4	22	46	105
<i>Panel A: optimal coverage D^*</i>									
1. Baseline	139%	0.13	0.25	0.27	0.31	0.00	0.00	0.00	0.04
2. Payouts exempted	10%	0.23	0.37	0.41	0.49	0.16	0.28	0.32	0.35
3. Both rates halved	51%	0.21	0.35	0.38	0.44	0.11	0.17	0.20	0.24
4. Payouts exempted, home-care only	96%	0.18	0.30	0.33	0.72 ^a	0.04	0.06	0.08	0.13
5. Clawback capped at the cash step	139%	0.13	0.61	0.66	0.72	0.00	0.00	0.00	0.04
6. No asset test	139%	0.13	0.17	0.15	0.11	0.00	0.00	0.00	0.00
7. Support for all, nothing withdrawn	10%	0.23	0.32	0.35	0.35	0.16	0.20	0.20	0.19
8. Free care on the zone	10%	0.00	0.00	0.06	0.16	0.00	0.00	0.00	0.00
<i>Panel B: expected public cost at D^*</i>									
1. Baseline		3.69	1.66	0.39	0.05	3.53	1.39	0.26	0.02
2. Payouts exempted		3.76	1.77	0.44	0.06	3.54	1.41	0.26	0.02
3. Both rates halved		3.71	1.71	0.41	0.06	3.51	1.38	0.25	0.02
4. Payouts exempted, home-care only		3.69	1.67	0.40	0.05	3.52	1.38	0.25	0.02
5. Clawback capped at the cash step		3.69	1.66	0.39	0.05	3.53	1.39	0.26	0.02
6. No asset test		4.12	4.10	4.11	4.13	4.18	4.18	4.18	4.18
7. Support for all, nothing withdrawn		4.18	4.18	4.18	4.18	4.18	4.18	4.18	4.18
8. Free care on the zone		0.77	0.48	0.17	0.03	0.74	0.34	0.07	0.01

Notes: Optimal coverage D^* , as a share of the annual cost of care, at four inception wealths x_0 , without a bequest motive and with the bequest pair of Table 1; the asset-test edge is $x_U = 20$, and Design 6 has none. The loading is the static loading $(1 + \eta)/\delta(D) - 1$ on the first coverage dollar, with η the market loading and $\delta(D)$ the tranche retention (11). The grid has $N = 33601$ points; Section S.8.11 gives each design's parameters and ten wealths; the baseline row differs from Figure 3(a), computed on a finer grid, by less than 0.002.

Panel B: the expected public cost of the design at D^* , the present value of its support net of the clawback, in multiples of the annual cost of care (Section S.8.11); in Design 7 it does not depend on wealth, and in Design 6 it varies with wealth only through the coverage chosen.

^a The optima below and past the kink D_{SN} , 0.38 and 0.72, are equally good to within the precision of the grid (Section S.8.11); the cell and Panel B give the one past the kink.

Both move coverage the way the hypothesis predicts. Section S.8.11 reads the remaining rows.

The public cost in Panel B is the expected present value, at the riskless rate, of the support the design pays, net of the clawback and measured against the income y_{out} and the care cost O_{out} of a retiree above the edge in the baseline. It is counted along the optimal wealth path, in multiples of the annual cost of care (Section S.8.11).

Most of the baseline's cost at every wealth, 59% to 83% without a bequest motive and 70% to 82% with one, is the cash safety net paid while the retiree is healthy and on the zone. Designs 2 to 5, which change the income tests, therefore cost within 0.11 of the baseline at every wealth in both cases, because the clawback they change is levied only in disability and only on the payouts. Designs 6 and 7, which extend the support to every retiree, cost the most at every wealth, and free care on the zone (Design 8) the least (Table 2, Panel B).

What a design is worth to the retiree is the wealth they would give up at inception to be under it rather than the baseline, the compensating variation of Table S.10 (Section S.8.11). To the retiree without a bequest motive, exempting payouts (Design 2) is worth 2 to 20 times its extra public cost. Support for all (Design 7) is worth the most at every wealth, but Designs 6 and 7 are worth less than their extra public cost up to $x_0 = 61$, and about as much at $x_0 = 80$; at $x_0 = 105$ they are worth 1.1 to 1.2 times it. Free care on the zone (Design 8) costs the retiree more than it saves the government at every wealth; at high wealth the zone is reached only on rare paths, and a risk-averse retiree values the cash safety net on them far above its expected

cost.

10.2 Implications for the design of the means test

For the government agency that sets the means test, the comparison has the following implications, at the calibration of Table 1, with the premium unchanged and the retiree's choice as the only response. Coverage is raised through the income tests, not the asset test: keeping the asset test and exempting payouts (Design 2) gives the rise stated at the head of Section 10. A payout counts as income, so the income tests decide how much of it is clawed back; the asset test looks at wealth and only decides who gets support (Section S.8.11). Halving both channel rates (Design 3) gives more than half of that rise for half the cut to the zero rates of Design 2, although only the two ends of that cut and its midpoint are computed. A lower rate on the first tranche raises coverage at every wealth (Designs 2 to 4); the cap at the cash step does not (Design 5). Exempting payouts costs the agency little (Panel B of Table 2, Design 2), and the cap at the cash step costs within 0.01 of the baseline (Design 5).

Without the asset test, coverage falls more through the clawback that comes with the wider support (Design 6 against Design 7) than through the support itself (Design 7 against Design 2), although without a bequest motive Design 7 is at its cap from $x_0 = 46$. Abolishing the asset test raises the cost at every wealth, severalfold above the edge (Design 6). Ranked by value to the retiree per unit of extra public cost, the designs that change only the income tests (Designs 2 to 5) come first and the two that extend the support to everyone (Designs 6 and 7) last (Table S.10 of Section S.8.11).

11 Conclusions

We have cast the LTCI choice under a means test as stochastic control under a borrowing constraint and derived its demand equation. Without a bequest motive the disabled value across the asset-test edge reduces to two scalars, and the means test has a closed-form price; the rest is solved numerically.

Within the model, calibrated to the UK and Australian schedules, the spillover is large: at five times the edge, exempting payouts would raise coverage by more than half its level. The reason is timing: a payout arrives only in disability, and a retiree, unsupported above the edge, may by then have spent or lost enough to fall below it, where the first coverage dollar carries a loading of 139% against the market's 10%. Optimal coverage stays below a third of the annual cost of care without a bequest motive and close to zero with one, at every wealth reported. Of the seven alternative designs, two show which test is the lever. Exempting payouts lifts coverage at each wealth computed; without a bequest motive it is worth to the retiree 2 to 20 times its extra public cost. Abolishing the asset test alone does the opposite: it reduces coverage at all but the lowest wealths without that motive, and leaves none with it, mainly because every retiree then gets support, and with it the clawback.

The formulation applies to any public programme that withdraws a transfer at constant rates up to caps and stops at a wealth threshold. Two results hold for any means-test parameters that satisfy their assumptions. Where a channel's support runs out, the marginal benefit of coverage

steps up (Theorem 6.4). Without a bequest motive, the fraction of a coverage dollar’s retained value lost to the borrowing constraint at zero wealth is set by the interest rate, the Sharpe ratio and the disabled retiree’s mortality-adjusted discount rate alone (Theorem 7.1(d)). For CRRA utility without a bequest motive and with unbounded leverage, the construction of Section 5 applies wherever crossing a single exogenous threshold moves income or costs by a step, under the conditions of Section 5.3. The loadings and coverage levels of Sections 9 and 10 are specific to this calibration.

Two simplifications mark the scope of these results. First, private income is held fixed, which leaves out an annuity bought with the same wealth. A life annuity bought at inception would enter the same schedule, since the income tests assess its income as they assess a payout; how an annuity and LTCI are best combined under a means test is the subject of ongoing work. Second, the designs of Section 10 are compared by the coverage they induce, their public cost and their value to the retiree, not by a social objective. Such an objective, weighing that value against the public cost, would turn the comparison into a design problem over the schedule’s rates and caps. A further extension splits the step at the edge into smaller steps, so that a saved payout loses support below the edge too; the construction of Section 5 would handle each.

Acknowledgements

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

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRedit authorship contribution statement

Lingfeng Lyu: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. Michael Sherris: Funding acquisition, Methodology, Supervision, Validation, Writing – review & editing. Mengyi Xu: Methodology, Supervision, Validation, Writing – review & editing.

Data availability

Data will be made available on request.

References

Ameriks, J., Briggs, J., Caplin, A., Shapiro, M. D., & Tonetti, C. (2020). Long-term-care utility and late-in-life saving. *Journal of Political Economy*, 128(6), 2375–2451. <https://doi.org/10.1086/706686>

- Ameriks, J., Caplin, A., Laufer, S., & Van Nieuwerburgh, S. (2011). The joy of giving or assisted living? Using strategic surveys to separate public care aversion from bequest motives. *Journal of Finance*, 66(2), 519–561. <https://doi.org/10.1111/j.1540-6261.2010.01641.x>
- Andréasson, J. G., Shevchenko, P. V., & Novikov, A. (2017). Optimal consumption, investment and housing with means-tested public pension in retirement. *Insurance: Mathematics and Economics*, 75, 32–47. <https://doi.org/10.1016/j.insmatheco.2017.04.003>
- Australian Government Actuary. (2024). *Australian life tables 2020–22*. Retrieved October 6, 2026, from <https://aga.gov.au/publications/life-tables/australian-life-tables-2020-22>
- Barucci, E., Biffis, E., & Marazzina, D. (2023). Health insurance, portfolio choice, and retirement incentives. *European Journal of Operational Research*, 307(2), 910–921. <https://doi.org/10.1016/j.ejor.2022.09.016>
- Bick, B., Kraft, H., & Munk, C. (2013). Solving constrained consumption–investment problems by simulation of artificial market strategies. *Management Science*, 59(2), 485–503. <http://doi.org/10.1287/mnsc.1120.1623>
- Bonne, B., & Meunier, M. (2019). *Diminuer le reste à charge des personnes âgées dépendantes : c’est possible !* (Rapport d’information No. 428) [Lowering the out-of-pocket costs of dependent older people: It can be done]. Sénat. Retrieved October 5, 2026, from <https://www.senat.fr/rap/r18-428/r18-428.html>
- Braun, R. A., Kopecky, K. A., & Koreschkova, T. (2019). Old, frail, and uninsured: Accounting for features of the U.S. long-term care insurance market. *Econometrica*, 87(3), 981–1019. <https://doi.org/10.3982/ECTA15295>
- Brown, J. R., & Finkelstein, A. (2007). Why is the market for long-term care insurance so small? *Journal of Public Economics*, 91(10), 1967–1991. <https://doi.org/10.1016/j.jpubeco.2007.02.010>
- Brown, J. R., & 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. <https://doi.org/10.1257/aer.98.3.1083>
- Colello, K. J. (2023). *Long-term care insurance: Overview* (CRS In Focus No. IF11614). Congressional Research Service. <https://crsreports.congress.gov/product/pdf/IF/IF11614>
- Dybvig, P. H., & Liu, H. (2010). Lifetime consumption and investment: Retirement and constrained borrowing. *Journal of Economic Theory*, 145(3), 885–907. <https://doi.org/10.1016/j.jet.2009.08.003>
- Dybvig, P. H., & Liu, H. (2011). Verification theorems for models of optimal consumption and investment with retirement and constrained borrowing. *Mathematics of Operations Research*, 36(4), 620–635. <https://doi.org/10.1287/moor.1110.0507>
- El Karoui, N., & Jeanblanc-Picqué, M. (1998). Optimization of consumption with labor income. *Finance and Stochastics*, 2(4), 409–440. <https://doi.org/10.1007/s007800050048>
- He, H., & Pagès, H. F. (1993). Labor income, borrowing constraints, and equilibrium asset prices. *Economic Theory*, 3(4), 663–696. <https://doi.org/10.1007/BF01210265>
- Hubbard, R. G., Skinner, J., & Zeldes, S. P. (1995). Precautionary saving and social insurance. *Journal of Political Economy*, 103(2), 360–399. <https://doi.org/10.1086/261987>

- Jordà, Ò., Knoll, K., Kuvshinov, D., Schularick, M., & Taylor, A. M. (2019). The rate of return on everything, 1870–2015. *The Quarterly Journal of Economics*, 134(3), 1225–1298. <https://doi.org/10.1093/qje/qjz012>
- Karatzas, I., Lehoczky, J. P., & Shreve, S. E. (1987). Optimal portfolio and consumption decisions for a “small investor” on a finite horizon. *SIAM Journal on Control and Optimization*, 25(6), 1557–1586. <https://doi.org/10.1137/0325086>
- Kim, T., Yang, Z., & Jeon, J. (2026). The finite-horizon retirement problem with borrowing constraint: A zero-sum stopper versus singular-controller game. *Mathematics of Operations Research*, Advance online publication. <https://doi.org/10.1287/moor.2025.0873>
- Koijen, R. S. J., Van Nieuwerburgh, S., & Yogo, M. (2016). Health and mortality delta: Assessing the welfare cost of household insurance choice. *The Journal of Finance*, 71(2), 957–1010. <https://doi.org/10.1111/jofi.12273>
- Koo, B., Pantelous, A. A., & Wang, Y. (2022). Novel utility-based life cycle models to optimise income in retirement. *European Journal of Operational Research*, 299(1), 346–361. <https://doi.org/10.1016/j.ejor.2021.08.048>
- Leporatti, L., Levaggi, R., Menoncin, F., & Miniaci, R. (2025). Optimal consumption, portfolio, and long-term-care health insurance in a dynamic framework. *Risk Management and Insurance Review*, 28(4), 588–608. <https://doi.org/10.1111/rmir.70021>
- Lockwood, L. M. (2018). Incidental bequests and the choice to self-insure late-life risks. *American Economic Review*, 108(9), 2513–2550. <https://doi.org/10.1257/aer.20141651>
- Lyu, L., & Sherris, M. (2026). *Household home care financing: Long-term care insurance, home equity release, and means-tested public support* (Working Paper). Centre for Population Ageing Research (CEPAR).
- Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4), 373–413. [https://doi.org/10.1016/0022-0531\(71\)90038-X](https://doi.org/10.1016/0022-0531(71)90038-X)
- Milevsky, M. A., & Young, V. R. (2007). Annuitization and asset allocation. *Journal of Economic Dynamics and Control*, 31(9), 3138–3177. <https://doi.org/10.1016/j.jedc.2006.11.003>
- OECD. (2020). *Long-term care and health care insurance in OECD and other countries*. <https://doi.org/10.1787/3eabc286-en>
- OECD. (2024). *Is care affordable for older people?* <https://doi.org/10.1787/450ea778-en>
- Park, K., & Sherris, M. (2026). Design and pricing of private long-term care insurance: An Australian analysis. *Journal of Risk and Insurance*, 93(3), 628–668. <https://doi.org/10.1111/jori.70057>
- Vila, J.-L., & Zariphopoulou, T. (1997). Optimal consumption and portfolio choice with borrowing constraints. *Journal of Economic Theory*, 77(2), 402–431. <https://doi.org/10.1006/jeth.1997.2285>
- Xu, M., Alonso-García, J., Sherris, M., & Shao, A. W. (2023). Insuring longevity risk and long-term care: Bequest, housing and liquidity. *Insurance: Mathematics and Economics*, 111, 121–141. <https://doi.org/10.1016/j.insmatheco.2023.03.004>