

Centre for Population Ageing Research (CEPAR)

Working Paper 2026/05

The Implicit Tax On Long-Term Care Insurance: Optimal Coverage Under Means-Tested Public Support

Lingfeng Lyu, Michael Sherris, Mengyi Xu

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

The Implicit Tax On Long-Term Care Insurance: Optimal Coverage Under Means-Tested Public Support

Lingfeng Lyu[1,2*] Michael Sherris[1,2] Mengyi Xu[1,2]

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

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

November 2, 2026

Abstract

Means tests reserve public support, including care, for retirees of limited means. Income tests reduce it as income rises, and assets tests stop it above a wealth cut-off. Buyers of private long-term care insurance face an implicit tax, the share of each payout lost to reduced support. We find that the tax lowers coverage even for retirees who start far above the cut-off. Although no payout is taxed above it, consumption and risk can later push wealth below it. In the model a retiree buys coverage once while healthy, consumes and invests, may need care and may value bequests. Support stops at the cut-off and wealth cannot go negative, so Merton's solution does not apply. Without a bequest motive, we use convex duality to solve the care phase. In our calibration a retiree below the cut-off keeps under half of each payout, so the premium is over twice what they expect to keep. With unbounded leverage, coverage stays under a third of the cost of care up to five times the cut-off. With a bequest motive added, the tax almost removes coverage. Even then, weakening income tests enough raises coverage at all these wealths for little extra public cost. A leverage cap lowers coverage, but exempting payouts from income tests still raises it. Lowering the cut-off raises coverage a little, with a bequest motive only at higher wealth. Yet retirees above the present cut-off lose more than the government saves, since the support removed comes only when they are poorer.

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 care on average across its member countries. The same report puts the share above one in five in Germany, Israel, Lithuania and Switzerland. Over a lifetime the share is far higher. In the US, two in five men and more than half 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). LTCI is unavailable in many countries, Australia among them (Park & Sherris, 2026). 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 assets test reduces or stops it as their wealth rises (OECD, 2024, pp. 53, 60). The US, the UK, France and Australia means test both this support and the minimum income of retirees (the cash safety net). When a test counts a payout from LTCI as income or the policy as an asset, a retiree who holds LTCI loses part of the support they would otherwise receive. This lost support is the clawback. The share of each payout that it takes acts as an implicit tax on LTCI. Brown and Finkelstein (2008) use the term 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, in which 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).

We ask two questions about the spillover, namely 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 our model, the answer to the first is timing. A retiree chooses coverage while still healthy (which we simplify to a single purchase at retirement), 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. A retiree who starts retirement above the range where the test pays support has no public support. Yet by the time disability begins, consumption and these risks may already have run their wealth down into that range, where payouts reduce support. Hence 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. Its assets test sets the range of wealth on which payouts are clawed back and its income tests set the rate, so it determines what coverage is worth to every retiree who may qualify for support. Coverage rises in our calibration when the income tests are weakened, at little extra public cost, and also when the assets test stops support at a lower wealth. The lower cut-off costs retirees who start above the present range more than it saves in public cost. The gain from exempting payouts survives a cap on the risky share, which lowers coverage itself.

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. The model's assets test withdraws support once wealth passes the assets-test cut-off. Below the cut-off the income tests claw back part of each payout through two channels. One withdraws the cash safety net and the other withdraws 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 cut-off, a point inside the state space. Merton's solution fitted on each side of the cut-off does not join there. If we join the two sides smoothly by matching value and marginal value at the cut-off, we obtain a solution of the Hamilton–Jacobi–Bellman (HJB) equation, but not yet the value function. Because leverage is unbounded (no constraint caps the share of wealth held in the risky asset), the optimal risky position in that equation is finite only where the solution is strictly concave in wealth. Only a strictly concave solution can therefore be the value function. The step in income and costs at the cut-off can break that concavity, so a smooth join alone does not settle the problem. Second, the retiree cannot borrow against future income, so wealth cannot go negative. This limits wealth, not leverage. At zero wealth Merton's portfolio would require borrowing because it counts future income as wealth. Third, the income tests withdraw a fixed share of each payout until the support they pay is used up. The amount withdrawn is thus 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. The retiree's objective is therefore not concave in coverage, and the first-order condition no longer identifies the optimum.

Without a bequest motive we solve the disabled state by convex duality, in which a Legendre transform turns the HJB equation on each side of the cut-off into a linear Cauchy–Euler equation. In earlier uses of the dual method (Dybvig & Liu, 2010, 2011; El Karoui & Jeanblanc-Picqué, 1998; He & Pagès, 1993) the boundary between regions is free, set by the agent's own policy and located by smooth fit. Here the means test fixes the cut-off, and income and costs jump there, so we join the two sides across that jump. By policy iteration on a finite-difference grid we solve the healthy state, and the disabled state with a bequest motive. To find optimal coverage we compare the best coverage on each side of a kink. On each side an interior optimum solves a demand equation, which equates the premium's marginal cost with the marginal benefit of coverage in disability, averaged over the wealth the retiree may hold when care begins and while it lasts.

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. LTCI is not sold in either country, so the coverage we compute is what a retiree facing means tests like theirs would buy if such a market existed.

On the first question, within our calibrated model, we find that at five times the cut-off and without a bequest motive the government expects to pay a retiree almost nothing over their remaining life. Yet exempting payouts from the income tests would raise their coverage by more than half. Below the cut-off, on coverage up to the kink where the cash safety net runs out, the test claws back more than half of every dollar paid out. Each dollar the retiree keeps therefore

carries a loading of 139%, whereas the market charges 10%. Without a bequest motive, optimal coverage insures less than a third of the annual cost of care at every wealth up to five times the cut-off. Without the implicit tax a bequest motive lowers it by a fifth to a third, but with the tax it brings coverage close to zero.

On the second question, in the same model, we find that weaker income tests raise coverage without a bequest motive at every wealth we compute and with a bequest motive once the cut in their rates reaches a fifth. Exempting payouts does so at little extra public cost. Each retiree without a bequest motive values the exemption at two to twenty times its extra public cost. A cut-off 5% lower raises coverage a little, and with a bequest motive only from three times the cut-off. It lowers the public cost at every wealth, but at each wealth above the calibrated cut-off the retiree loses more than the government saves, and below it less.

This paper contributes to the literature on stochastic control models of insurance and retirement, and to that on LTCI demand under means testing. 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.

A further contribution, certified numerically for our calibration, is the value function of the disabled state when wealth can cross the cut-off in either direction but cannot fall below zero. Without a bequest motive the solution joined at the cut-off is in closed form apart from two numbers, the marginal value of wealth at the cut-off and at zero wealth. These two numbers are set by the conditions for joining the two sides at the cut-off and by the borrowing constraint at zero wealth. The joined solution is the value function if the two numbers exist and make it strictly concave. Sign conditions on the two numbers check this, and they hold for our parameters. The demand equation and the pricing of the implicit tax are not derived when the sign conditions fail. The construction is not specific to LTCI. Under the same conditions, it applies to problems with constant relative risk aversion, unbounded leverage and no bequest motive in which income or costs change by a step when wealth crosses an exogenous threshold.

We also 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. Brown and Finkelstein (2008) measure the implicit tax on a policy as a whole, whereas we price the tax on a marginal coverage dollar. Under the means test alone and without the borrowing constraint, a coverage dollar received below the assets-test cut-off would be worth the share of the payout the test leaves plus the option of rising above the cut-off (where payouts are kept in full). At zero wealth, where the borrowing constraint binds and cash today counts for more than a later payout, the constraint removes a fraction of that share. The fraction does not depend on the means test, so the two sources of loss can be separated.

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 and Section 9 reports the loadings and coverage. Section 10 varies the income tests and the cut-off, and Section 11 concludes. We give the proofs, the numerical reporting and the country schedules in the Online Supplement.

2 Related literature

Our 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 our retiree 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. The finite-horizon retirement problem of Kim et al. (2026) has two free boundaries, one from the constraint and one from the retirement decision. There, income stops for good once the agent retires. In our model, support stops above the assets-test cut-off 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 (Pauly, 1990), 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. Neither that model nor the two US papers below price the tax on a marginal coverage dollar. We do, and without a bequest motive we obtain the price in closed form.

In Hubbard et al. (1995) the implicit tax falls on saving. Their assets-tested consumption floor makes the budget set non-convex, so the optimum is found by comparing candidates rather than from a first-order condition alone, as with coverage in our problem (Theorem 6.4). In Brown and Finkelstein (2008) the tax 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. This order of payment plays the part of the clawback. With the insurance policy held fixed under each reform, 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). By contrast, our retiree chooses coverage under each design, and exempting payouts from the income tests but keeping the assets test (Section 9.3) raises coverage in our calibration.

The bequest motive, the other influence on coverage here, affects LTCI demand differently across models. 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 range of published values over which we vary the motive's strength in Section 9.3. The motive lowers demand in Lockwood (2018) as it does in ours, 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. The household can also buy life insurance, which pays at death. Here no life insurance restores the estate, so a premium paid by one who never needs care is lost to the estate (Section 9.3).

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. The insured share in their model falls with wealth, whereas coverage in our calibration rises with it (Section 9.3). In operational research two papers are closest to ours, 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 our

model (Section 4.1). They 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 the UK and Australia the clawback, the support a retiree loses by holding LTCI, extends to retirees whose wealth exceeds the value of their home and keeps rising with coverage. Section 3.4 carries it into the model as a schedule with two channels.

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, which puts the safety net at 0.625 in both countries.[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.

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. The safety-net channel is the reduction in the cash safety net when LTCI payouts are assessed as income. Their sum is the net clawback.

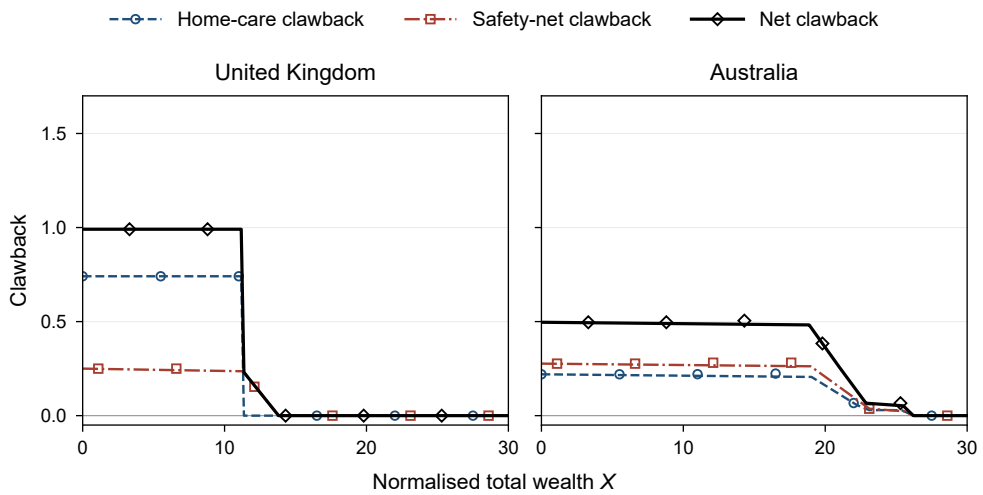


Figure 1. The clawback in the UK and Australia as a function of normalised total wealth.
Notes: The clawback, a yearly flow, is drawn at an assumed LTCI indemnity of 1. The retiree receives private income of 0.375 a year in the UK. In Australia the retiree receives 0.15, set just above the income-test 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.

[1] The value 0.625 is our choice, 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.

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]

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

Because the home is the same in both countries, Figure 1 compares institutions and not wealth composition.[4]

For low-wealth retirees who rely on public support, an assessed LTCI payout triggers a clawback. In the UK, for example, 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. It 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 assets 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 the net clawback rises gradually with the indemnity and 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 assets test's capital limit. That 0.25 is all the cash support the retiree in the UK receives, since the safety net tops a private income of 0.375 up to 0.625. The indemnity uses up this cash support first, so further coverage is clawed back by the home-care channel alone. In Australia both channels are still rising at $D = 1$.

3.3 The clawback and the size of private LTCI markets

We offer a hypothesis, not an explanation, of where private LTCI markets form. 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. The US and France, drawn in Section S.9.7 for a retiree with no private income, recover benefits from the estate after death. Net of that recovery counted at face value, their clawback is confined to wealth up to about the value of the home and stops growing at an indemnity of 0.625. In the UK and Australia the primary residence is exempt and the cash safety net is withdrawn gradually, so the clawback extends to total wealth above the value of the home (Figure 1) and keeps rising with coverage. Much of a payout then replaces public support. Private LTCI covered about 7.1 million people in France

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

[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 but does not move the retiree along the axis.

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).

3.4 From the cross-country clawback to the model

We model the two clawback channels of Figure 1 as two parts of one public schedule. Both parts count the LTCI payout as income, and both apply only below the assets-test cut-off, above which the assets test removes eligibility. This cut-off stands in for the gradual fall of the clawback with wealth in Figure 1. 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. 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 of which jump at the cut-off.

4 Model

In this section we set out the health states, the wealth and income dynamics and the LTCI premium. We then turn to the implicit tax created by the means test and to the retiree's problem, a single choice of coverage at retirement followed by consumption and portfolio choice.

4.1 Health state transitions

The agent is a single retiree. Time $t \geq 0$ is continuous and measured in years, and the agent's health state S_t is a continuous-time Markov chain on $\{0, 1, \dagger\}$, which captures the morbidity and mortality risks associated with long-term care (LTC). In State 0 (healthy) the agent needs no LTC services, and in State 1 (disabled) severe cognitive impairment or limitations in activities of daily living require formal LTC. State † (dead) is absorbing. Transitions occur at constant intensities, and we write ν for the morbidity rate, the intensity of transition from State 0 to State 1. Constant intensities keep the problem time-homogeneous, which the closed form of Section 5 requires. We rule out recovery from State 1 to State 0, consistent with the persistence of late-stage LTC needs. The mortality intensities in the healthy and disabled states, $\lambda^{(0)}$ and $\lambda^{(1)}$, satisfy

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

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

4.2 Wealth and income dynamics

We work on a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ 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.

The agent can invest in two assets, namely a riskless bond paying a constant interest rate $r > 0$ and a risky asset with expected return $\mu_R > r$ and volatility $\sigma_R > 0$. We denote by $x_t \in \mathbb{R}_+$ 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 $\pi_t \geq 0$ is the proportion of wealth held in the risky asset, and we require wealth to stay non-negative (Section 4.4).

Before we introduce any insurance product, wealth follows the stochastic differential equation

$$dx_t = \big[(r + \pi_t(\mu_R - r))x_t + y(x_t) - c_t\big]\, dt + \pi_t \sigma_R x_t\, dW_t. \tag{3}$$

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). Because the public flows are means tested, y has a two-level structure in wealth

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

where $x_U > 0$ is the assets-test cut-off, the total wealth (home included) above which the assets test removes eligibility for means-tested support.

In $[0, x_U]$ the agent passes the assets test and receives the public flows on top of any private ones. Thus y_{in} is the cash flow received in $[0, x_U]$ before any clawback of the indemnity, and we treat this flow as one constant (Section S.1.1). Above x_U the assets test withdraws the public flows and leaves 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 cut-off. These tests assess wealth and income, not disability, so (4) holds in health as in disability.

Assumption 4.1 (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.*[5]

We take a single premium so that coverage is one choice made in health, and no premium enters the wealth dynamics afterwards. The single premium $P(D) = a_D D$ is linear in coverage, with market loading $\eta \geq 0$ and the unit premium

$$a_D = (1 + \eta)\, \frac{\nu}{\big(r + \nu + \lambda^{(0)}\big)\big(r + \lambda^{(1)}\big)}, \tag{5}$$

the expected present value at inception of a unit indemnity flow over the care spell, times $1 + \eta$ (derived in 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 $O(x) \geq 0$, the care cost net of the subsidy received, which depends on wealth.

[5] 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.

The assets-test cut-off x_U of Section 4.2 separates two regions, and we call $\mathcal{Z} = [0, x_U]$ the implicit tax band, or simply the band (Assumption 4.2 below). In $\mathcal{Z}$ the agent qualifies for the subsidy subject to the income test. The baseline cost there is the constant O_{in}. Above x_U 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,

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

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

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

the support that a retiree with no coverage receives in $\mathcal{Z}$ when disabled, and so the most that the means test can withdraw.

We write Π for the clawback of Section 3, the withdrawal with coverage less the withdrawal without it. Assumption 4.2 below gives it the two-channel form of Section 3.4.

Assumption 4.2 (The implicit tax band: Two channels, two kinks)**.** (i) (Two channels.) *In $\mathcal{Z}$ each channel withholds a constant share of the indemnity ($m_{\text{SN}} > 0$ for the safety-net channel and $m_{\text{HC}} > 0$ 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 $\mathcal{Z}$ the assets test has removed the support and nothing is withheld:*

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

The cash safety net runs out at the coverage D_{SN} and the public care subsidy at D_{HC}, where

$$D_{\text{SN}} := \frac{\Delta_y}{m_{\text{SN}}}, \qquad D_{\text{HC}} := \frac{\Delta_O}{m_{\text{HC}}}. \tag{9}$$

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

$$D_{\text{SN}} \;<\; D_{\text{HC}}. \tag{10}$$

This ordering holds when the cash step Δ_y is the smaller support and is clawed back at the steeper rate, so it is used up first, as in the UK schedule of Section 3.2. We call D_{SN} and D_{HC} the kinks. Coverage on $\mathcal{Z}$ therefore falls into three layers. On the first layer $[0, D_{\text{SN}})$ both channels claw back payouts. On the intermediate layer $(D_{\text{SN}}, D_{\text{HC}})$ the home-care channel alone does so. On the top layer beyond D_{HC} nothing more is clawed back. We derive (8) from a general schedule in Section S.1.1. The two-channel form holds when the assets test and the income tests act separately, the withdrawals of the two channels add, and each channel's rate is an average over $\mathcal{Z}$. The condition $m \le 1$ is the non-confiscatory condition (at most one dollar of support withdrawn per dollar of income) applied to the averaged rates.

Under Assumption 4.2 the share of a marginal indemnity dollar that the agent keeps after

the test in $\mathcal{Z}$, the layer retention, is the two-step schedule

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

with Π' the derivative in coverage (one-sided at the kinks). The complement $\Pi'(D) = 1-\delta(D)$ of the retention is the implicit tax, the share of a marginal indemnity dollar that the test withholds. It is m on the first layer (the only layer where it equals $\Pi(D)/D$, the share of the whole payout withheld), m_{HC} on the intermediate layer and zero on the top layer. At the kink D_k retention rises by exactly m_k, the rate of the channel that saturates there.

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

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

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,

$$\partial_D\left[D - \Pi(x,D)\right] = \delta(x;D) := \delta(D)\,\mathbf{1}\{x \le x_U\} + \mathbf{1}\{x > x_U\}. \tag{13}$$

4.4 The optimisation problem

4.4.1 Objective and value functions

The agent maximises expected discounted lifetime utility from consumption and, through a bequest motive, from the wealth left at death. Let $\tau_\dagger$ be the time of death. The objective below is maximised over admissible strategies. A strategy $\Gamma = (c_t,\pi_t)_{t\ge 0}$ is admissible if it is progressively measurable with respect to $\{\mathcal{F}_t\}$ and satisfies $c_t > 0$, $\pi_t \ge 0$ 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 has a unique strong solution with $x_t \ge 0$ on $[0,\tau_\dagger]$.[6] The wealth equation is (3) in health and (12) in disability. We denote the set of admissible strategies by $\mathcal{A}$.

The portfolio enters the wealth dynamics only through the dollar exposure $\wp_t := \pi_t x_t$, which is therefore the effective control. Admissibility rules out short sales but leaves leverage unbounded. That is, the only restriction on the portfolio apart from $\pi_t \ge 0$ is the state constraint $x_t \ge 0$, which we call the borrowing constraint. It bars borrowing against future income, though not borrowing to invest. The closed form of Section 5 needs unbounded leverage, since the dual transform linearises the HJB equation only where no leverage cap binds (Vila & Zariphopoulou, 1997, Appendix C, eqs. (C.1)–(C.5)). A leverage cap would bind near $x = 0$, where the optimal share diverges (Section S.1.4). Under a cap of one, coverage in our calibration falls to zero at most wealths (Section 9.3). To obtain a closed form under a cap, Vila and Zariphopoulou (1997, Proposition 4.2) require the cap to be proportional to wealth, with no credit line.

[6] The integrability condition is the standard one (Karatzas et al., 1987, Definition 2.2).

The objective functional under initial state (x, S_0) is

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

where $\beta > 0$ is the subjective discount rate and $\Gamma \in \mathcal{A}$ is an admissible strategy. Coverage $D \geq 0$ is fixed by the purchase of Assumption 4.1.

For analytical tractability we take both utilities to be constant relative risk aversion (CRRA) functions, the second shifted as in Lockwood (2018),

$$U(c) = \frac{c^{1-\gamma}}{1-\gamma}, \qquad B(x) = \kappa\, \frac{(x + \underline{x})^{1-\gamma}}{1-\gamma}, \tag{15}$$

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

For a fixed coverage level $D \geq 0$, the state-dependent value functions are

$$V^{(j)}(x; D) = \sup_{\Gamma \in \mathcal{A}} J(x, j; \Gamma, D), \qquad j \in \{0, 1\}. \tag{16}$$

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 4.1), an optimal coverage solves the static problem

$$D^*(x_0) \in \arg\max_{0 \leq D \leq D_{\max}} V^{(0)}(x_0 - a_D D;\, D). \tag{17}$$

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 $\mathrm{SR} := (\mu_R - r)/\sigma_R$ for the Sharpe ratio. On functions a subscript x, z or ζ denotes the derivative in that variable. The partial derivative in coverage is ∂_D and the ordinary one a prime, as in $\Pi'(D)$. On an indexed scalar a dot also denotes the ordinary derivative in coverage, as in $\dot{\zeta}_U$ in Section 6.

In the healthy state (State 0) the agent pays no premium flow (Assumption 4.1), so the HJB equation has the baseline drift and two jump terms. One is for the onset of disability, which moves the agent to State 1. The other is for death, which pays the bequest. The equation is

$$\begin{aligned} A_0\, V^{(0)}(x; D) = \max_{(c,\pi) \in (0,\infty) \times [0,\infty)} \left\{ \frac{c^{1-\gamma}}{1-\gamma} + [\mu(x, \pi) - c] V_x^{(0)} + \tfrac{1}{2} \pi^2 \sigma_R^2 x^2 V_{xx}^{(0)} \right\} \\ + \nu V^{(1)}(x; D) + \lambda^{(0)} B(x). \end{aligned} \tag{18}$$

In the disabled state (State 1) the agent pays $O(x)$ for care, receives D and incurs the clawback $\Pi(x, D)$ (8). In this state death is the only jump (Section 4.1), and the HJB equation

is

$$A_1\, V^{(1)}(x;D) = \max_{(c,\pi)\in(0,\infty)\times[0,\infty)} \Big\{ \frac{c^{1-\gamma}}{1-\gamma} + [\mu(x,\pi) - c - O(x) + D - \Pi(x,D)]V_x^{(1)} + \tfrac{1}{2}\pi^2\sigma_R^2 x^2 V_{xx}^{(1)} \Big\} + \lambda^{(1)} B(x). \tag{19}$$

When $V^{(j)}$ is C^2 on $(0,\infty)$ off x_U, with $V_x^{(j)} > 0$ and $V_{xx}^{(j)} < 0$, the Hamiltonian in (18)–(19) attains its maximum over $(c,\pi) \in (0,\infty)\times[0,\infty)$ for $x > 0$ off x_U at

$$c^\star = (V_x^{(j)})^{-1/\gamma}, \qquad \wp^\star = \pi^\star x = -\frac{(\mu_R - r)\, V_x^{(j)}}{\sigma_R^2\, V_{xx}^{(j)}} \;>\; 0, \tag{20}$$

which are Borel functions of x (Section S.5). The condition $V_{xx}^{(j)} < 0$ is necessary as well as sufficient for a maximiser to exist. If $V_{xx}^{(j)} \geq 0$ at some $x > 0$, the supremum over π is $+\infty$.

Since the State 1 equation (19) does not involve $V^{(0)}$, we can solve the coupled system recursively, with State 1 first. We construct the disabled value function in Section 5.

5 The disabled value function

The demand for coverage depends on what a marginal indemnity dollar is worth to a disabled agent. Without a bequest motive ($\kappa = 0$) the disabled value function is Merton's power law in the sum of wealth and the capitalised net flow (its present value at the riskless rate), with two corrections. One accounts for the value of crossing the assets-test cut-off x_U. The other, on the band $\mathcal{Z}$, accounts for the cost of the borrowing constraint at zero wealth. Because the clawback depends on coverage alone on each side of the cut-off (Assumption 4.2), we can reduce State 1 to a Merton problem on each region (Lemma 5.1). The value is explicit apart from two numbers, the marginal values of wealth at the cut-off and at zero wealth, which we find by root-finding. A reader who wants the economics can take the construction and the standing assumptions of Section 5.3 as given and go on to Section 6.

5.1 The two-region Merton reduction

Lemma 5.1 (Two-region Merton reduction and saturation)**.** *Under Assumption 4.2, the following hold.*

(a) (The clawback is a capped piecewise-linear function of coverage.) *The clawback (8) is continuous, concave and piecewise linear in coverage. Under the ordering (10) of the kinks, its slope is m on $[0, D_{\mathrm{SN}})$, m_{HC} on $(D_{\mathrm{SN}}, D_{\mathrm{HC}})$ and zero beyond. Under the same ordering $\Pi(D_{\mathrm{HC}})$ equals $\theta_{\max}$, the support lost at the cut-off.*

(b) (A two-region Merton problem.) *For $D < D_{\mathrm{HC}}$, State 1 (19) is a Merton consumption–portfolio problem with the piecewise-constant retained drift*

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

capitalised into the shifts $h_{\text{in}} = (y_{\text{in}} - O_{\text{in}} + D - \Pi(D))/r$ *and* $h_{\text{out}} = (y_{\text{out}} - O_{\text{out}} + D)/r$.[7] *The shift* h_{in} *increases with coverage at the retention of the current layer (with slope* $(1-m)/r$ *on the first layer and* $(1-m_{\text{HC}})/r$ *on the intermediate layer) and has a kink at* D_{SN}. *The cut-off* x_U *lies inside the domain of* $V^{(1)}$, *so the two regions must be joined there, and the retained drift changes across it by the net resource step*

$$r\,(h_{\text{out}} - h_{\text{in}}) = \Pi(D) - \theta_{\max} \;<\; 0. \tag{21}$$

(c) (The two regions merge.) *For* $D \geq D_{\text{HC}}$ *both channels are exhausted and* $\Pi(D) = \theta_{\max}$. *The drift* $y_{\text{in}} - O_{\text{in}} + D - \theta_{\max}$ *in the band* $\mathcal{Z}$ *then 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 at* $\kappa = 0$ *equals the power law only when that constraint is dropped.*

Proof. Section S.1.1.

This step (21) is negative because on crossing the cut-off the agent escapes the clawback $\Pi(D)$ but loses the support $\theta_{\max}$, which is larger. We write $\Delta h := h_{\text{out}} - h_{\text{in}}$ for this step capitalised at r. The construction needs resources to stay strictly positive on $[0, x_U]$.

Assumption 5.2 (Well-posed band)**.** *For every coverage* $D \in [0, D_{\text{HC}}]$, $h_{\text{in}}(D) > 0$, *or 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. It holds in practice because means-tested schedules cap the out-of-pocket cost by capacity to pay.

5.2 The dual construction across the assets-test cut-off

On each region the State 1 HJB equation (19) with $\kappa = 0$ is Merton's equation. 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))

$$A_1\, V^{(1)} = r z\, V_z^{(1)} - \frac{\text{SR}^2}{2}\frac{(V_z^{(1)})^2}{V_{zz}^{(1)}} + \frac{\gamma}{1-\gamma}\,(V_z^{(1)})^{1-1/\gamma}, \qquad \text{SR}^2 = \frac{(\mu_R - r)^2}{\sigma_R^2}. \tag{22}$$

With a single region and no constraint, (22) is solved by the power law $\frac{\omega^{-\gamma}}{1-\gamma} z^{1-\gamma}$ with consumption–wealth ratio $\omega = \frac{1}{\gamma}\big[A_1 - (1-\gamma)(r + \frac{\text{SR}^2}{2\gamma})\big] > 0$.[8] 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 fitted on each side does not join there because the two sides capitalise different net flows and $\Delta h \neq 0$ (21). In our calibration $x_U + h_{\text{out}} \leq 0$ at low coverage, which leaves the power law from above undefined at the cut-off (Section S.1.4). Second, wealth may not fall below $x = 0$ and the power law is inadmissible at zero wealth. It

[7] Mortality already enters through the discount A_1 of (19), so each shift is capitalised at r. Capitalising a shift at $r + \lambda^{(1)}$ would count mortality twice.

[8] Positivity requires no restriction beyond $\gamma > 1$. Rearranging gives $\gamma\omega = A_1 + (\gamma - 1)\big(r + \frac{\text{SR}^2}{2\gamma}\big)$, a sum of strictly positive terms, so the analogue of the Merton condition holds for the CRRA preferences (15).

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 $\wp^\star(0) = 0$ (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 power law is a particular solution, so the value is that solution plus homogeneous corrections for the two failures.[9]

Three versions of the disabled problem appear below, differing in what they keep. Merton's problem drops both the assets-test cut-off and the borrowing constraint, and its value is the power law above. The unconstrained benchmark keeps the cut-off and drops the constraint, so wealth may fall below zero. The region below the cut-off is then the extended band $\tilde{\mathcal{Z}} = (-h_{\text{in}}, x_U]$, which ends where discretionary wealth runs out and depends on coverage through h_{in}. This benchmark isolates the means test. Its value is a closed form (Section S.1.3), in which we price the implicit tax in Section 7. The constrained problem keeps both. Figure 2 draws the three.

Proposition 5.3 (Disabled value across the assets-test cut-off)**.** *Let $\kappa = 0$ and $D < D_{\text{HC}}$, and let Assumption 5.2 hold. Write $\zeta := V_z^{(1)}$ for the dual variable and $\widehat{V}(\zeta) := V^{(1)} - z\zeta$ 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)). Set $s := 1 - 1/\gamma$ and $\Lambda := \gamma/((1-\gamma)\omega)$. Then the following hold.*

(a) *On each region, (22) is equivalent under the transform to the linear Cauchy–Euler equation*

$$\frac{\text{SR}^2}{2}\,\zeta^2\widehat{V}_{\zeta\zeta} + (A_1 - r)\,\zeta\widehat{V}_\zeta - A_1\,\widehat{V} + \frac{\gamma}{1-\gamma}\,\zeta^{\,s} = 0, \tag{23}$$

whose general solution (the particular solution $\Lambda\zeta^s$ plus two homogeneous modes) is

$$\widehat{V}(\zeta) = \Lambda\,\zeta^{\,s} + G_+\zeta^{\,p_+} + G_-\zeta^{\,p_-}, \tag{24}$$

where $p_\pm$ are the roots of the indicial quadratic

$$\text{Ind}(p) := \frac{\text{SR}^2}{2}\,p^2 + \Big(A_1 - r - \frac{\text{SR}^2}{2}\Big)p - A_1 = 0. \tag{25}$$

(b) *The roots are real and satisfy $p_- < 0$, $p_+ > 1$ and $p_- < s < p_+$. Moreover, $\text{Ind}(s) = -\omega$.*

(c) *Let $\widehat{V}^{\text{in}}$ on $[0, x_U]$ and $\widehat{V}^{\text{out}}$ on (x_U, ∞) be solutions of the form (24), with coefficients $(G_\pm^{\text{in}})$ and $(G_\pm^{\text{out}})$. Transversality as $x \to \infty$ forces $G_-^{\text{out}} = 0$, while the branch on $\mathcal{Z}$ 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 $\bar{\zeta} > \zeta_U$, satisfy the five conditions*

$$\begin{aligned}
&\text{(I1)}\quad -\widehat{V}_\zeta^{\text{in}}(\zeta_U) = x_U + h_{\text{in}}, \qquad \text{(I2)}\quad -\widehat{V}_\zeta^{\text{out}}(\zeta_U) = x_U + h_{\text{out}},\\
&\text{(I3)}\quad \widehat{V}^{\text{in}}(\zeta_U) - \widehat{V}^{\text{out}}(\zeta_U) = \zeta_U\,\Delta h,\\
&\text{(B1)}\quad -\widehat{V}_\zeta^{\text{in}}(\bar{\zeta}) = h_{\text{in}}, \qquad \text{(B2)}\quad \widehat{V}_{\zeta\zeta}^{\text{in}}(\bar{\zeta}) = 0.
\end{aligned} \tag{26}$$

[9] The construction of Proposition 5.3 follows Dybvig and Liu (2010, 2011). The dual 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).

Both branches reach x_U at the marginal value ζ_U and their values agree there (I1)–(I3). *The branch on $\mathcal{Z}$ reaches the constraint* (B1) *and no risky asset is held at zero wealth* (B2)*, 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.1.3).*

(d) *A branch is strictly concave in wealth if and only if $\widehat{V}_{\zeta\zeta} > 0$ on its interior. Strict concavity of the constructed function reduces to the endpoint conditions $\widehat{V}_{\zeta\zeta}(\zeta_U) > 0$ on both branches and $\widehat{V}_{\zeta\zeta\zeta}(\bar{\zeta}) < 0$. In the unconstrained benchmark they reduce to a window of admissible step sizes Δh, given in Section S.2, that contains $\Delta h = 0$ in its interior.*

(e) *If the conditions of* (c) *have a root at which those of* (d) *hold, the function defined by*

$$z = -\widehat{V}_\zeta(\zeta), \qquad V^{(1)} = \widehat{V}(\zeta) - \zeta\widehat{V}_\zeta(\zeta), \qquad x = z - h_\bullet, \tag{27}$$

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 $V^{(1)}_{xx}$ 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, with the concavity conditions of part (d) proved in Section S.2.

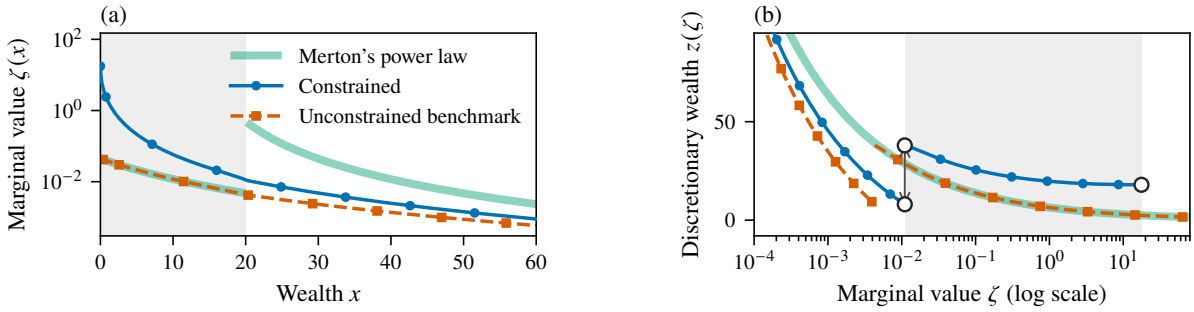


Figure 2. The disabled marginal value in wealth and in the dual variable.

Notes: Calibration of Table 1, coverage $D = 0.3$. The shaded strip is the band $\mathcal{Z}$, in Figure 2(b) the constrained problem's dual interval $[\zeta_U, \bar{\zeta}]$. Figure 2(a) plots the marginal value $\zeta(x)$ on a logarithmic scale and Merton's power law fitted on each side of the cut-off x_U. Figure 2(b) plots discretionary wealth z against ζ, where the power law is one curve. The upper circle at ζ_U marks (I1), on the band branch, and the lower (I2). The circle at $\bar{\zeta}$ marks (B1) with (B2), where the branch is tangent at $z = h_{\text{in}}$. The arrow is the step Δh at the cut-off. In the unconstrained benchmark the band branch runs on towards $z = 0$ and crosses the cut-off left of the strip.

Below the cut-off the borrowing constraint adds the ζ^{p_+} mode, whereas the unconstrained benchmark has $G^{\text{in}}_+ = 0$ and only the ζ^{p_-} mode, the value of crossing the cut-off. At zero wealth the borrowing constraint binds. The agent there holds no risky asset and consumes less than the net flow rh_{in}, so wealth is pushed back up. Zero wealth is therefore a reflecting boundary, not an absorbing one. We derive these properties in Remark 5.4.

Remark 5.4 (Borrowing constraint at zero wealth)**.** *Assumption 5.2 gives $h_{\text{in}} > 0$, so at $x = 0$ the disabled agent still receives the net flow $rh_{\text{in}} > 0$, the marginal value $\bar{\zeta}$ is finite and consumption $c^\star(0) = \bar{\zeta}^{-1/\gamma}$ is positive. The constraint at $x = 0$ is therefore only a borrowing constraint. The feedback controls* (20) *are themselves admissible at $x = 0$. Because $\wp^\star(0) = 0$, which is* (B2) *read in wealth, the drift at $x = 0$ is $rh_{\text{in}} - c^\star(0)$. At and above D_{HC} the regions*

merge and only the ζ^{p_+} mode remains. Conditions (B1) *and* (B2) *of Proposition 5.3(c) keep their form in the single region. As in the constant-income case of El Karoui and Jeanblanc-Picqué (1998, Theorem 4.3) with A_1 for their discount rate, eliminating the mode's coefficient between the two conditions gives*

$$c^\star(0) = \frac{\omega h_{\text{in}} \gamma (p_+ - 1)}{1 + \gamma (p_+ - 1)},$$

so

$$r\, h_{\text{in}} - c^\star(0) = \frac{\text{SR}^2}{2} \frac{(p_+ - 1)\,(p_+ - s)}{1 + \gamma\,(p_+ - 1)}\, h_{\text{in}} \;>\; 0, \tag{28}$$

because $p_+ > 1$ and $p_+ > s$ (Proposition 5.3(b)). Below D_{HC} the ζ^{p_-} mode adds a correction at $x = 0$, and the drift there stays positive under the concavity conditions of Proposition 5.3(d). In our calibration the correction's weight is negligible up to the coverage cap.[10]

5.3 Standing assumptions and scope

The construction of Proposition 5.3 rests on two properties that we check numerically rather than prove, namely that the conditions (26) have a root and that the candidate is concave. We state both as standing assumptions and assume regularity where no closed form exists.[11]

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, continuous in coverage and C^1 within each layer.*

In the unconstrained benchmark the root is explicit at every coverage. In the constrained problem the step Δh at the cut-off vanishes at D_{HC}, so the root there is the single-region root of the constant-income case of He and Pagès (1993, p. 684) and El Karoui and Jeanblanc-Picqué (1998, Theorem 4.2). From there we can continue it locally by the implicit function theorem, layer by layer. With the parameters of Table 1, D_{HC} lies beyond the coverage cap $\bar{D}$, so the continuation runs from outside the feasible range towards $D = 0$. We check the assumption numerically from zero coverage to beyond D_{HC}.[12] Under it and Assumption 5.6 the map $D \mapsto V^{(1)}(\cdot; D)$ is continuous, and C^1 off the kinks D_{SN} and D_{HC}, which Lemma 6.2 and the kink jump (36) use.

Assumption 5.6 (Concave regime)**.** *Wherever the construction of Proposition 5.3 is used, the parameters satisfy the concavity conditions of its part (d), which we call the concavity certificate.*

In the unconstrained benchmark the window of Proposition 5.3(d) is tightest at $D = 0$. In the constrained problem the operative check is the numerical one at the root, over the feasible coverage range.[13] This assumption does not cover the healthy value $V^{(0)}$, which must be concave because leverage is unbounded, as for State 1. For our parameters we verify numerically

[10] Section S.4 proves that the feedback controls are admissible at $x = 0$, and Section S.1.2 reads (B2) in wealth. Section S.3 shows that (B1) and (B2) keep their form in the single region, and Section S.8 reports the correction's weight.

[11] For each property, Section S.8 lists what is proved, what is checked numerically and what is assumed.

[12] Section S.1.3 gives the benchmark's root, Section S.3 the single-region root and Section S.8 the check.

[13] The root is that of the five conditions (26).

the strict concavity that the maximisers (20) need on each side of the cut-off, where the cash step makes the curvature of $V^{(0)}$ jump. Assumption 5.7(i) records this concavity.[14]

Where the certificate fails, the smooth-fit candidate is not the value function. With the other parameters as calibrated, the constrained certificate fails at zero coverage on the side of the cut-off in the band $\mathcal{Z}$ when the Sharpe ratio falls below a floor, which Section 8 reports. The condition $\widehat{V}^{\mathrm{in}}_{\zeta\zeta}(\zeta_U) > 0$ on that side says an agent just below the cut-off holds the risky asset, and it fails where that position falls to zero: The support lost above the cut-off then outweighs the reward for the risk of crossing it. Below the floor and at low coverage, under conditions that hold in every case we compute, an agent in $\mathcal{Z}$ remains there and the value stays concave but its marginal value drops at the cut-off (Section S.2). Because Sections 6 and 7 use a continuous marginal value at the cut-off, we exclude the case below the floor from them.

With a bequest motive the two-region structure survives, but the bequest term makes the dual equation nonlinear, so solutions no longer superpose and the closed form is lost. We compute the value numerically (Section S.1.7). The disabled value enters Section 6 only through its regularity, which the closed form supplies at $\kappa = 0$. With a bequest motive, and for the healthy value at every κ, we assume that regularity instead.

Assumption 5.7 (Regularity without a closed form)**.** *For every coverage D in the feasible range, the HJB system (18)–(19) has a solution with the following properties, for $j = 0$ at every $\kappa \geq 0$ and for $j = 1$ when $\kappa > 0$.*

(i) (Each state.) *For each such j, $V^{(j)} < 0$ and $V^{(j)}_x > 0$. The value $V^{(j)}$ is C^1 on $\mathbb{R}_+$ and C^2 on $(0, \infty)$ off x_U, with $V^{(j)}_{xx}$ $<$ 0 there and finite negative one-sided limits at x_U. At zero wealth $V^{(j)}_x(0) < \infty$, the drift of wealth under the optimal controls is positive and the risky position vanishes like $\sqrt{x}$, as (B2) of Proposition 5.3(c) and the sign condition of its part (d) make it for $j = 1$ at $\kappa = 0$. The optimal controls grow at most linearly in wealth.*

(ii) (In coverage.) *The map $D \mapsto (V^{(j)}, V^{(j)}_x, V^{(j)}_{xx})$ is continuous, and C^1 on each layer with one-sided derivatives at the kink D_{SN}, locally uniformly in x. The optimal controls are continuous in D uniformly on compact subsets of $\mathbb{R}_+$, zero wealth included. The coverage derivative $\partial_D V^{(j)}$ is bounded, C^1 on $\mathbb{R}_+$ and C^2 off x_U. For $j = 1$ we write it ψ.*

(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.*

With Assumption 5.7 for $j = 1$ in place of Assumptions 5.5 and 5.6, the results of Section 6 hold at $\kappa > 0$ with the same proofs, since the bequest terms of (18)–(19) do not depend on coverage and drop out of every derivative in D. Proposition 5.3, Remark 5.4 and the pricing of Section 7 rest on the closed form and we state them at $\kappa = 0$ only. We verify State 1 as the value function by the supermartingale argument of Dybvig and Liu (2011, Lemma 3.3), which covers the state constraint $x_t \geq 0$, and State 0 by the same argument given part (i). Pointwise attainment of the maximum in (20) justifies the envelope step of Section 6, and the Feynman–Kac representations there use the bounds on the coverage derivatives of part (ii). The closed-loop wealth equation is well posed, in State 0 given parts (i) and (iii), of which we check

[14] Section S.2 gives the window, Section S.8 the check at the root and Section S.1.5 the jump in curvature.

the positive drift and the Lipschitz continuity near $x = 0$ numerically at $\kappa = 0$. At $\kappa > 0$ we check the signs, the concavity and the drift at zero wealth at bequest strengths that span the published estimates. At the baseline strength we also check the rate $\sqrt{x}$, the regularity of ψ in D and its jump at D_{SN} (36).[15] We assume the rest, including the regularity of the healthy value in D and the growth condition, and at $\kappa > 0$ the Lipschitz condition. 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 $\partial_D V^{(0)}$ of coverage at inception and not directly on the disabled-state marginal benefit $\psi = \partial_D V^{(1)}$. 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 wealth rises.

For the results of this section we need the second kink to lie beyond every feasible coverage.

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

$$D_{\max} \;<\; D_{\mathrm{HC}} \;=\; \frac{\Delta_O}{m_{\mathrm{HC}}}, \tag{29}$$

for which $m_{\mathrm{HC}}\,\bar{D} < \Delta_O$ is sufficient.

That is, even the largest indemnity that insurers offer cannot exhaust the public care subsidy at the averaged care rate. We place no restriction on D_{SN}. Both $D_{\max} \le D_{\mathrm{SN}}$ and $D_{\mathrm{SN}} < D_{\max}$ are admitted, Theorem 6.4 covers both, and our calibration determines which holds. The assumption matches the UK and Australian side of the country contrast of Sections 3.2 and 3.3.[16]

We write $\zeta(x) = V_x^{(1)}(x; D)$ for the disabled marginal value of Proposition 5.3, and $\delta(x; D)$ is the after-test retention (13), built from the layer retention $\delta(D)$ of (11) and the assets-test cut-off x_U. A star marks an optimally controlled process, and $X^\star$ denotes disabled wealth under the controls (20). Its drift $b^\star(x) = \mu(x, \pi^\star) - c^\star - O(x) + D - \Pi(x, D)$ and diffusion $\sigma^\star(x) = \pi^\star \sigma_R x$ are both piecewise across x_U. Subscripts on $\mathbb{E}$ give the initial wealth of the displayed process.

Lemma 6.2 (State-1 marginal benefit)**.** *Let Assumption 6.1 hold, with Assumptions 5.2, 5.5 and 5.6 at $\kappa = 0$ or Assumption 5.7 at $\kappa > 0$. Let $D \le D_{\max}$ with $D \neq D_{\mathrm{SN}}$, 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 and is C^1 on $\mathbb{R}_+$ and C^2 on $(0, \infty)$ off x_U (at $\kappa = 0$ by the closed form and at $\kappa > 0$ by Assumption 5.7(ii)), and the following hold.*

(a) *The marginal benefit ψ solves the linearisation of (19) along the optimal control,*

$$A_1\,\psi = b^\star \psi_x + \tfrac{1}{2}(\sigma^\star)^2 \psi_{xx} + \delta(x; D)\,\zeta(x), \tag{30}$$

[15] Section S.1.7 gives the bequest case, Section S.1.6 the verification argument, Section S.4 the closed-loop wealth equation and Section S.8 the numerical checks.

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

with $b^\star, \sigma^\star$ as above.

(b) *The marginal benefit ψ admits the Feynman–Kac representation*

$$\psi(x) = \mathbb{E}_x\left[\int_0^\infty e^{-A_1 t}\,\delta(X_t^\star; D)\,\zeta(X_t^\star)\,dt\right] \;\geq\; 0. \tag{31}$$

In the constrained problem the integrand is bounded, by $\delta(D)\bar{\zeta}$ on the band $\mathcal{Z}$ and by $\bar{\zeta} = V_x^{(1)}(0) < \infty$ above it (Section S.5).

(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 (31) lets us read ψ as an expected discounted integral along the optimal wealth path. Each instant is weighted by the marginal value ζ and by the retention, $\delta(D)$ while wealth is at or below x_U and 1 while it is above x_U. Hence $\psi(x)$ is depressed only through the time 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 $D \geq 0$ with $D \neq D_{\mathrm{SN}}$, and let $X^{(0)\star}$ denote the optimally controlled healthy wealth of (18). Its drift $b^{(0)\star}(x) = \mu(x, \pi^\star) - c^\star$ and diffusion $\sigma^{(0)\star}(x) = \pi^\star \sigma_R x$ are both piecewise across x_U through $y(x)$ and the controls. Then $\psi^{(0)} := \partial_D V^{(0)}(\cdot; D)$ solves*

$$A_0\,\psi^{(0)} = b^{(0)\star}\psi_x^{(0)} + \tfrac{1}{2}\big(\sigma^{(0)\star}\big)^2\psi_{xx}^{(0)} + \nu\,\psi, \tag{32}$$

and admits the representation

$$\partial_D V^{(0)}(x; D) = \nu\int_0^\infty e^{-A_0 t}\,\mathbb{E}_x[\psi(X_t^{(0)\star})]\,dt = \mathbb{E}_x\left[e^{-\beta\tau_{\mathrm{morb}}}\mathbf{1}_{\{\tau_{\mathrm{morb}} < \tau_\dagger^{(0)}\}}\,\psi\big(X_{\tau_{\mathrm{morb}}}^{(0)\star}\big)\right] \geq 0, \tag{33}$$

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

Proof. Section S.5.

Proposition 6.3 averages the marginal benefit ψ over the wealth held when care begins. The demand equation below sets the marginal premium cost at inception equal to that average.

Theorem 6.4 (Optimal coverage and the demand equation)**.** *Let the utilities be the CRRA functions (15) and let Assumptions 4.1, 6.1 and 5.7 hold. When $\kappa = 0$, let Assumptions 5.2, 5.5 and 5.6 hold as well. Then the inception objective of (17)*

$$\Phi(D) := V^{(0)}(x_0 - a_D D; D)$$

is continuous on $[0, D_{\max}]$ and C^1 off the kink D_{SN}, with

$$\Phi'(D) = -\,a_D\,V_x^{(0)}(x_0 - a_D D; D) + \partial_D V^{(0)}(x_0 - a_D D; D) \tag{34}$$

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

$$\underbrace{a_D\, V_x^{(0)}(\tilde{x}_0; D^*)}_{\text{marginal premium cost at inception}} \quad = \quad \underbrace{\nu \int_0^\infty e^{-A_0 t}\, \mathbb{E}_{\tilde{x}_0}[\psi(X_t^{(0)\star})]\, dt}_{\text{marginal benefit averaged over wealth at onset}} \quad , \qquad \tilde{x}_0 := x_0 - a_D D^*, \tag{35}$$

with ψ given by Lemma 6.2 at D^*. When $D_{\mathrm{SN}} < D_{\max}$, the disabled-state marginal benefit ψ jumps upward across the kink at every x by the safety-net rate times the expected discounted marginal value of wealth accumulated while disabled wealth is in $\mathcal{Z}$,

$$\psi(x; D_{\mathrm{SN}}^+) - \psi(x; D_{\mathrm{SN}}^-) = m_{\mathrm{SN}}\, \mathbb{E}_x\left[\int_0^\infty e^{-A_1 t}\, \mathbf{1}\{X_t^\star \le x_U\}\, \zeta(X_t^\star)\, dt\right] \; > \; 0. \tag{36}$$

Here the law of $X^\star$ and the marginal value ζ are those of the two-region problem at D_{SN}, common to both sides. Hence Φ' jumps upward across D_{SN} by $\nu \int_0^\infty e^{-A_0 t}\, \mathbb{E}_{x_0 - a_D D_{\mathrm{SN}}}[\psi(X_t^{(0)\star}; D_{\mathrm{SN}}^+) - \psi(X_t^{(0)\star}; D_{\mathrm{SN}}^-)]\, dt$, no optimum sits at D_{SN}, and a global optimum is the better of the layer optima. On the first layer, $\Phi'(0) > 0$ rules out the corner $D^* = 0$. If instead Φ is quasi-concave there and $\Phi'(0) < 0$, then $D = 0$ is the best coverage on that layer.

Proof. Section S.5.

This jump of Φ' follows because the law of $X^{(0)\star}$ at D_{SN} is common to both sides, so (33) carries the jump of ψ through, and the kink is convex. Assumption 6.1 keeps the further kink at D_{HC} (convex at $\kappa = 0$) out of the feasible range. For the calibrated parameters we verify numerically that Φ is quasi-concave on the first layer. The low retention in $\mathcal{Z}$ does not by itself make zero coverage optimal, because care that begins with wealth above the cut-off keeps part of the benefit independent of that retention.[17]

Optimal coverage $D^*(x_0)$ varies with inception wealth in two ways. First, within a layer it is continuous in x_0 where the layer optimum is unique, and has no step at x_U. In our calibration and without a bequest motive, coverage under unbounded leverage increases over the whole range of x_0 in Figure 4. Second, where D_{SN} lies inside the feasible range, the global optimum can switch layer as wealth rises. The optimum $D^*(x_0)$ then jumps from low first-layer coverage to coverage past the kink.

7 The implicit tax priced

On the band $\mathcal{Z}$ below the assets-test cut-off, 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 (we read Theorem 7.1 below as a loading). In this section we compute 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.[18] We read the marginal benefit through the money's worth of a marginal indemnity dollar, its value as a share of ζ/r (the value of a dollar

[17] Section S.6 shows that the kink at D_{HC} is convex at $\kappa = 0$ and gives this argument, Section S.8.11 the check.
[18] We compute in the dual variable ζ of Proposition 5.3, where the equation for ψ has explicit coefficients.

kept in full in the unconstrained benchmark),

$$\mathrm{MW}(x) := \frac{r\,\psi(x)}{\zeta(x)}. \tag{37}$$

In the unconstrained benchmark of Section 5.2 the marginal benefit has a closed form. There the money's worth rises with wealth across the extended band $\tilde{\mathcal{Z}}$, from the retention $\delta(D)$ in the limit at $-h_{\mathrm{in}}$, and tends to one as wealth grows without bound (the benchmark proposition of Section S.5). Its excess over $\delta(D)$ on $\tilde{\mathcal{Z}}$ is the value of the option of escaping across the cut-off, which an agent deep in $\tilde{\mathcal{Z}}$ no longer has. Theorem 7.1 gives the marginal benefit when borrowing is ruled out, the problem behind the numbers we report in Section 9.

Theorem 7.1 (The implicit tax priced: Closed-form marginal benefit)**.** *Let $\kappa = 0$, and let Assumptions 5.2, 5.5 and 5.6 hold. Write a dot for d/dD, and let the layer pair $(\Pi'(D), \delta(D))$ equal $(m, 1-m)$ on $[0, D_{\mathrm{SN}})$ and $(m_{\mathrm{HC}}, 1-m_{\mathrm{HC}})$ on $(D_{\mathrm{SN}}, D_{\mathrm{HC}})$. In the constrained problem, for D off the kink D_{SN} and below the kink D_{HC} (the last sentence treats D above it), the disabled-state marginal benefit on $\mathcal{Z}$ is*

$$\psi(x) = \frac{\delta(D)}{r}\,\zeta(x) + \dot{G}^{\mathrm{in}}_{+}\,\zeta(x)^{p_+} + \dot{G}^{\mathrm{in}}_{-}\,\zeta(x)^{p_-}, \tag{38}$$

with $p_\pm$ the indicial roots of Proposition 5.3 and $\dot{G}^{\mathrm{in}}_{\pm}$ the D-derivatives of the coefficients $G^{\mathrm{in}}_{\pm}$ of its part (c)*. Conditions* (B1) *and* (B2) *of* (26) *fix $\dot{G}^{\mathrm{in}}_{+}$ by making ψ flat in ζ at $\bar{\zeta}$, the marginal value at zero wealth. The coefficient derivative $\dot{G}^{\mathrm{in}}_{-}$ is positive. The money's worth* (37) *on $\mathcal{Z}$ is*

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

For D above D_{HC}, where $\delta = 1$ and the two regions merge, the second term is absent and $\mathrm{MW}(0) = 1 - 1/p_+$.

Proof. Section S.5.

In Theorem 7.1 the money's worth is the retention $\delta(D)$ times a factor below one rising with wealth from $1 - 1/p_+$ at zero wealth, plus the escape option.

The premium (5) is priced on the gross indemnity, at retention one. Measured against $1/r$, the same benchmark value in dollars, the effective loading on a marginal dollar of coverage paid out below the cut-off is $(1+\eta)/\mathrm{MW}(x) - 1$ and varies with wealth. In the unconstrained benchmark its supremum over $\tilde{\mathcal{Z}}$ on the first layer is $(1+\eta)/(1-m) - 1$, approached as wealth falls towards $-h_{\mathrm{in}}$. 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 layer 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 $(1+\eta)/(1-m_{\mathrm{HC}}) - 1$, and it reaches the market loading η only at D_{HC}, which Assumption 6.1 places out of reach.

With Theorem 7.1 we can separate 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 Theorem 7.1 below $\delta(D)$ at every wealth, most at low wealth. Above D_{HC}, where retention is one and no escape option is left, the

constraint holds the money's worth itself below one. At zero wealth the constraint removes the fraction $1/p_+$ of the retention. That fraction depends on the interest rate, the Sharpe ratio and the discount rate in disability $A_1 = \beta + \lambda^{(1)}$ but not on the means test, on coverage or on risk aversion. On $\mathcal{Z}$ 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. In our calibration the escape option is the smaller at every wealth we compute in $\mathcal{Z}$ (Section 9.2), so there the money's worth lies below $\delta(D)$, the reverse of the benchmark. The static loadings above are therefore lower bounds on the constrained effective loading.

8 Calibration

We use two baselines that differ only in the bequest parameters. The no-bequest baseline sets $\kappa = 0$, for which Sections 5 and 7 give closed forms. The bequest baseline adds the pair $(\kappa, \underline{x}) = (1000, 7.42)$, and we solve it numerically. Table 1 gives the inputs with their sources in Panel A and the derived quantities in Panel B, among them the two kinks and the coverage cap. Section S.9 documents each source.

Risk aversion and the discount rate follow Brown and Finkelstein (2008). Expressed as a marginal propensity to bequeath $\kappa^{1/\gamma}/(1+\kappa^{1/\gamma})$, published estimates of the bequest strength are about 0.48 in Ameriks et al. (2020), 0.96 in Andréasson et al. (2017) and 0.98 in Ameriks et al. (2011). The baseline strength $\kappa^{1/\gamma} = 10$ gives 0.91, and we vary the strength across this range in Section 9.3. We take the threshold $\underline{x}/\kappa^{1/\gamma} = 0.74$, the consumption rate below which the motive is inactive, from Andréasson et al. (2017).

For the riskless rate we use 2.3%, the UK–Australia average real return on long-term government bonds in 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 lies between the equity-only and equity-and-housing ratios of the two countries in these data, from 0.32 to 0.53. Below a Sharpe ratio of 0.398 the concavity certificate of Assumption 5.6 fails at zero coverage (Section S.8.9 and Section 5.3). The UK's equity-only ratio of 0.32 is below this floor. With home equity in the risky asset, however, Proposition 5.3 applies to our calibration.

For a retiree aged 65 we set the three health transition intensities ν, $\lambda^{(0)}$ and α to reproduce 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 men and women, probabilities of 0.40 and 0.54 and mean years in care of 2.9 and 4.2. Our model treats those years in care as one spell. 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 inputs come from the UK and Australian schedules. The assets-test cut-off $x_U = 20$ and the care step $\Delta_O = 0.60$ are two-country means, and we derive the channel rates from the two schedules averaged at each wealth. The slope of the averaged clawback in coverage gives $m = 0.54$, and the safety-net channel's share of the first-layer clawback splits it into $m_{\mathrm{SN}} = 0.44$ and $m_{\mathrm{HC}} = 0.10$. The income in the band $y_{\mathrm{in}} = 0.625$ and the cash step $\Delta_y = 0.25$ are chosen rather than estimated and describe a stylised UK retiree, a Pension

Table 1. Baseline calibration and derived quantities.

Panel A: Input parameters	Value	Source
Preferences and market		
CRRA γ	3.00	Brown and Finkelstein (2008)
Discount rate β	0.030	Brown and Finkelstein (2008)
Bequest weight κ	1000	moderate, between 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, mid-range in Jordà et al. (2019)
Market loading η	0.10	model choice
Morbidity and mortality (set to three moments)		
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
Means-tested support and care costs (UK and Australian schedules, Section S.9)		
Channel rates m_{SN}, m_{HC}	0.44, 0.10	slope and split of the averaged schedule
Cash step Δ_y, care step Δ_O	0.25, 0.60	UK Pension Credit top-up, two-country mean
Out-of-pocket cost in the band O_{in}	0.35	model choice, above both countries' levels
Assets-test cut-off x_U	20.0	mean wealth where each country's clawback ends
Income in the band 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-layer rate
Support lost at the cut-off $\theta_{\max} = \Delta_y + \Delta_O$	0.85	(7)
Out-of-pocket cost above the band $O_{\text{out}} = O_{\text{in}} + \Delta_O$	0.95	(6)
Income above the band $y_{\text{out}} = y_{\text{in}} - \Delta_y$	0.375	(4)
Coverage cap $\bar{D} = O_{\text{out}}$	0.95	Assumption 4.1 at its bound
Kink $D_{\text{SN}} = \Delta_y / m_{\text{SN}}$	0.568	cash safety net used up, below $\bar{D}$ and so feasible
Kink $D_{\text{HC}} = \Delta_O / m_{\text{HC}}$	6.000	care subsidy used up, above $\bar{D}$ (Assumption 6.1)
Layer retention $1 - m$, $1 - m_{\text{HC}}$	0.46, 0.90	(11)
Static loading $(1 + \eta)/\delta(D) - 1$	139%, 22%, 10%	first, intermediate, top (Section 7)
Unit premium a_D	0.9483	(5)

Notes: Units are those of Section 3. The bequest rows enter the $\kappa > 0$ results only, and the no-bequest baseline sets $\kappa = 0$ and drops $\underline{x}$.

Credit claimant whose private income $y_{\text{out}} = 0.375$ is the median occupational pension of single pensioners who have one. The retiree has no State Pension, since a full one would lift them above the guarantee and empty the safety-net channel.

We set the out-of-pocket cost in the band $O_{\text{in}} = 0.35$ above both countries' levels, the market loading at $\eta = 0.10$ and the coverage cap at the bound of Assumption 4.1, $\bar{D} = O_{\text{out}} = 0.95$. Near the countries' own out-of-pocket levels the calibrated tests leave almost no coverage at any wealth we compute (Table S.11). That is consistent with the absence of a market in both countries, but it leaves no demand to explain, and our calibration is a stylised blend of the two schedules rather than either of them.

9 Numerical results

In our calibration (Table 1) the implicit tax more than doubles the premium per dollar of payout that a retiree below the cut-off keeps, and the borrowing constraint lowers what a kept dollar is worth, most at low wealth. Optimal coverage then stays under a third of the annual cost of care at every wealth we draw, and near zero with a bequest motive.

9.1 Numerical method

Our method has two nested levels. For a candidate coverage D, the inner level solves the two-region State 1 problem for the disabled value $V^{(1)}(\cdot; D)$ and its marginal value ζ (by Proposition 5.3 at $\kappa = 0$ and numerically with a bequest motive) and then computes the marginal benefit ψ from its linear equation (30). It also solves (18) for the healthy value $V^{(0)}(\cdot; D)$ at the same coverage by policy iteration, with forcing $\nu V^{(1)}$, on a finite-difference grid that has a grid point at the assets-test cut-off x_U. Then one solve of the linear equation (32) gives $\partial_D V^{(0)}$ at every wealth on the grid, so the average of ψ over the wealth held at onset in (33) needs no simulation. The outer level maximises the objective Φ over the scalar $D \in [0, D_{\max}]$ by one-dimensional root-finding on (34). It does so separately on each side of the kink D_{SN}, where the cash safety net runs out (Theorem 6.4). As checks we use the representation (31), finite differences in D and Monte Carlo simulation, and in the unconstrained benchmark the closed form.[19]

We take no finite-difference or root-finding step in D across a kink, or across the coverage $D_U = (x_0 - x_U)/a_D$ at which post-premium wealth lies exactly at the cut-off. A retiree just above the cut-off who buys more than D_U falls below it, where the withdrawn support is available again, so near the cut-off the premium partly pays for itself. The objective Φ nevertheless stays C^1 at D_U and only its curvature jumps, since $V_x^{(0)}$ and $\partial_D V^{(0)}$ are continuous across the cut-off (Assumption 5.7(i)–(ii)). The coverage D_U thus matters for the computation rather than for the optimum. It is feasible only for $x_0 \in (x_U,\, x_U + a_D \bar{D})$, which is (20, 20.9) in our calibration.[20]

9.2 The implicit tax, layer by layer

The means test applies on the band $\mathcal{Z}$ below the cut-off $x_U = 20$. On the first layer it claws back 54 cents of each marginal dollar that coverage pays in disability, up to the kink D_{SN} (Assumption 4.2). The retiree keeps 46 cents, so the premium per kept dollar is 2.39, a static loading of 139% against the market loading of 10%. On the intermediate layer the test claws back 10 cents. The kink D_{HC}, where the public care subsidy runs out, lies beyond the coverage cap.

What a kept dollar is worth to a disabled retiree depends on wealth, and the borrowing constraint lowers that worth most for the lowest-wealth retirees (Figure 3). The numbers that follow rest on the money's worth $\mathrm{MW}(x)$ of (37), the value of a marginal indemnity dollar as a share of $1/r$. The value $1/r$ is what a unit indemnity flow is worth to a retiree who can borrow

[19] Section S.1.7 gives the bequest case, Section S.5 the benchmark's closed form and Section S.8 the checks.

[20] Section S.6 gives the interval. Section S.8 gives the grids, tolerances and boundary treatment, and reports the convergence and run times.

against it at the riskless rate (the debt cancelled at death), as in the unconstrained benchmark.

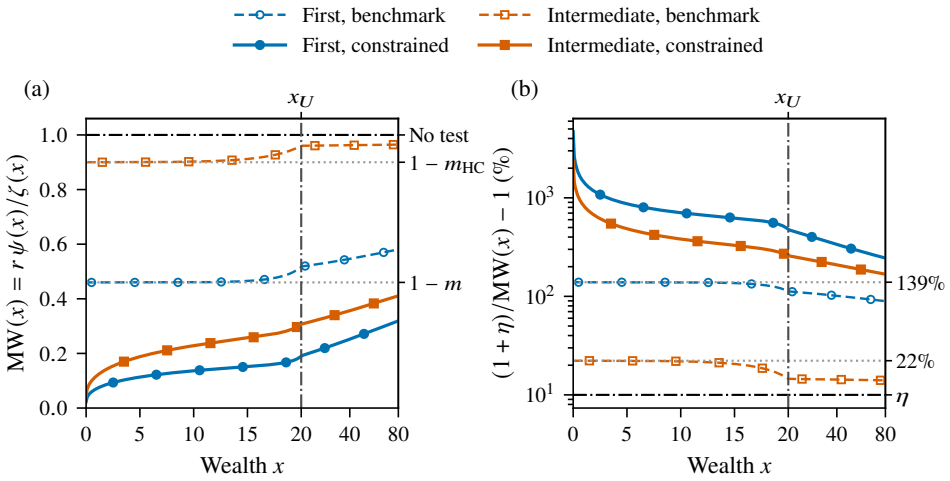


Figure 3. Money's worth of a marginal coverage dollar and the effective loading on it, by layer. *Notes:* Figure 3(a) plots the money's worth and Figure 3(b) the effective loading of Section 7 on a logarithmic axis, both against the wealth x of a disabled retiree. The first layer is drawn at coverage $D = 0.30$ and the intermediate layer at $D = 0.80$. Grey rules mark the two retentions and the dash-dot rule the no-test level 1, which is also the benchmark value on the top layer (not drawn). Figure 3(b) shows these levels as static loadings, the no-test level being the market loading η. The wealth axis is linear on the band $\mathcal{Z}$ and logarithmic above the cut-off.

In the unconstrained benchmark the money's worth is close to the layer retention (0.46 on the first layer and 0.90 on the intermediate one) over most of $\mathcal{Z}$ (Figure 3(a)). It rises well above the retention only near the cut-off, which the retiree may cross (the escape option of Section 7). By contrast, the borrowing constraint lowers it below the retention at every wealth we draw. On the first layer it is 0.191 at the cut-off and 0.023 at zero wealth, about two cents on the dollar and an effective loading of 4689% (Figure 3(b)).[21] Coverage is bought in health, before the wealth at the start of care is known, so the loss at zero wealth matters for coverage only to the extent that care may begin with little wealth (Theorem 6.4).

9.3 Optimal coverage, with and without a bequest motive

Even far above the cut-off, the implicit tax lowers optimal coverage by roughly a third to almost a half of its level with payouts exempted, and together with a bequest motive almost removes it (Figure 4).

Without a bequest motive, coverage rises with wealth and stays on the first layer at every wealth drawn, under a third of the annual cost of care and below the kink $D_{\mathrm{SN}} = 0.568$ (Figure 4(a)). Exempting payouts from the income tests removes the implicit tax and raises coverage at every wealth computed, from 0.31 to 0.49 at five times the cut-off (Figure 4(a) and (b)). Capping the risky share at one lowers coverage with payouts counted to zero at most wealths, but the exemption still raises it at every wealth.

With the bequest pair of Table 1 the retiree buys no coverage below $x_0 = 61$ and little above it, 0.04 at five times the cut-off (Figure 4(a)). Against the level with payouts exempted and no

[21] We measure against $1/r$ because on that base Theorem 7.1 separates what the means test leaves (the retention $\delta(D)$) from what the constraint removes. Against the flow's actuarial value $1/(r + \lambda^{(1)})$, the base on which the premium is priced, the loss at zero wealth is smaller and the means test removes more than the constraint (Section S.8.6).

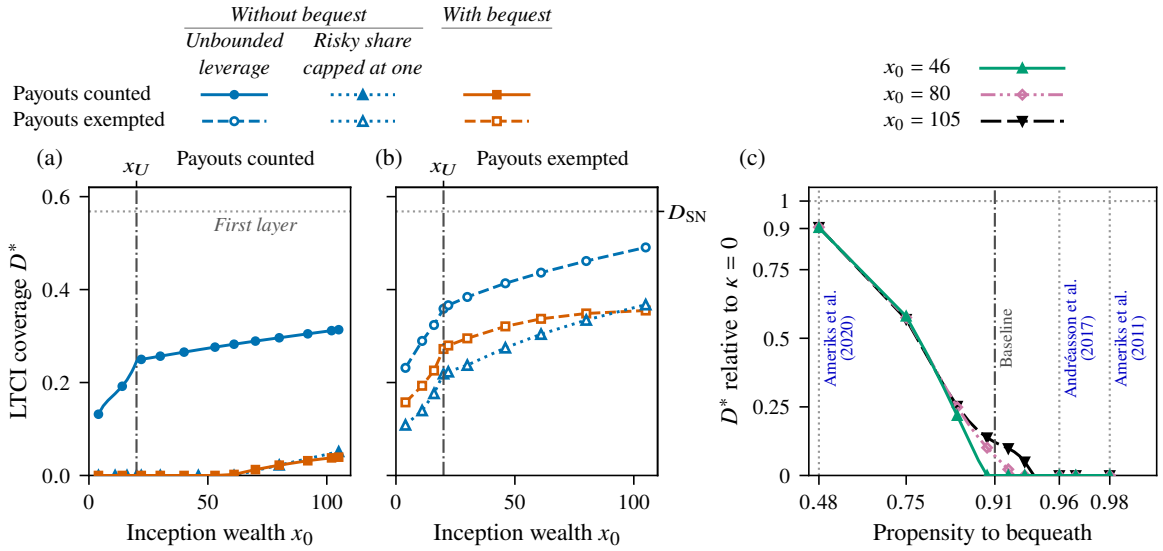


Figure 4. Optimal LTCI coverage by inception wealth with payouts counted or exempted, and by strength of the bequest motive.

Notes: Figure 4(a) and (b) each plot optimal coverage D^* against inception wealth x_0, without a bequest motive and with the bequest pair. The solid lines join the optima of Section S.7, and the dashed and dotted lines join those at the ten wealths of Section S.8.13. Figure 4(c) plots D^* with payouts counted relative to its level without a bequest motive. Its horizontal axis is the propensity to bequeath on a log-odds scale. The threshold $\underline{x}/\kappa^{1/\gamma}$ is held at 0.74.

bequest motive, the motive alone lowers coverage by less than the tax alone does, and the two together leave less than a tenth of it. Coverage enlarges the estate of a retiree who dies in care but takes the premium out of the estate of one who never needs care. In our calibration the second effect dominates: If the motive valued the estate only at a death in care, coverage would rise to the coverage cap from $x_0 = 61$ upward. If it valued the estate only at a death in health, coverage would fall to at most 0.013 (Section S.7). By Theorem 6.4, coverage is positive while the marginal benefit of the first coverage dollar exceeds its marginal premium cost. The motive narrows this margin as the tax does. Either alone leaves it positive, but the two together close it below $x_0 = 61$.

As we strengthen the motive across its published range, coverage falls (Figure 4(c)). At the propensity of 0.48 of Ameriks et al. (2020) it is about nine-tenths of its level without the motive. 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 we compute. Near the top of the published 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. If the motive valued the estate only at a death in health, the margin of the first coverage dollar would close at about the same strength at $x_0 = 80$ and at $x_0 = 105$. It is the effect at a death in care that lets coverage at $x_0 = 105$ last to a stronger motive than at $x_0 = 80$ (Section S.7).

9.4 Sensitivity

A ±5% change in any one of ten parameters of Table 1 moves coverage little at seven inception wealths from $x_0 = 15$ to 105, except where it carries coverage past the kink D_{SN}. There coverage jumps across the kink instead of moving a little (Figure 5). Past the kink the means test claws

back 10 cents of each further coverage dollar instead of 54, so the inception objective $\Phi(D)$ can rise again and have a second peak (Theorem 6.4). The optimum is the higher of the two peaks, the layer optima. At five times the cut-off the first is the higher in our calibration (Figure 5(a)). A rise of 5% in the out-of-pocket cost O_{in} in the band, the safety-net rate m_{SN} or mortality in health $\lambda^{(0)}$ puts the second above it, and so does a fall of 5% in the excess mortality in care α. Coverage jumps where the two peaks tie, after a move of 2% to 4% in any of the four (Figure 5(b)). The move needed grows as wealth falls, as Figure 5(c) shows for O_{in}. Within 5% only α and O_{in} carry coverage past the kink at $x_0 = 90$, and none does at lower wealths (Table S.7).

A jump means that the retiree faces two plans rather than one smooth trade-off. The first peak is a small cover that keeps the cash safety net, and the second a large cover that gives it up and faces only the lower tax on the dollars past the kink. Which plan is better depends on the parameters, so near a tie a small move switches the plan and coverage moves by about 0.4 while the retiree's value hardly moves.

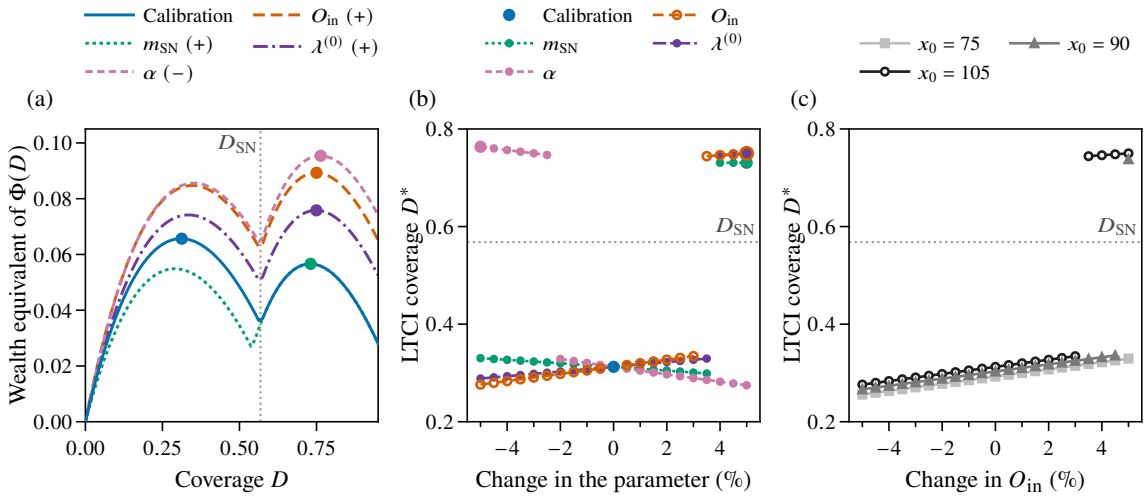


Figure 5. Inception objective and optimal coverage across the kink.

Notes: The whole figure sets no bequest motive and is solved as in Table S.7, and Figure 5(a) and (b) set $x_0 = 105$. Figure 5(a) plots the inception objective $\Phi(D)$ of Theorem 6.4 as a wealth equivalent. It is the wealth a retiree without coverage would need on top of x_0 to be as well off as one who buys D. In Figure 5(a) each parameter moves by 5% in the direction its legend entry gives, which carries coverage past the kink, and in Figure 5(b) from −5% to 5%. The larger dots mark the optima of Figure 5(a), and the same optima in Figure 5(b). Figure 5(b) and (c) join optima on the same side of the kink D_{SN}, where the cash safety net runs out. A higher m_{SN} moves the kink down, to 0.54 at 5% higher, and past D_{SN} the m_{SN} line lies on the calibration's.

Coverage with a bequest motive stays low, at most 0.09 in every case. The range ±5% lies inside the one over which the concavity certificate of Assumption 5.6 (on the construction of the disabled value) and the other conditions behind our results hold for each of the ten parameters. We also set four inputs to other values from their sources, one value at a time. The six values are two market loadings, an out-of-pocket cost in the band, a Sharpe ratio and a lower and a higher care need (probability and length of care). Coverage at these values stays under a third without a bequest motive and near zero with one, except at the higher care need.[22]

[22] Table S.7 gives the results at the seven wealths, Section S.8.9 measures the range of the conditions and Table S.11 gives the results at the six values.

10 Means-test design and optimal coverage

In our calibration the income tests move coverage at little public cost, whereas the cut-off x_U moves the public cost and changes coverage little except at the highest wealths (Figure 6).

10.1 Coverage, public cost and value as the tests change

Section 9.4 moves each of ten inputs by 5% alone to check the calibration, among them the rates m_{SN} and m_{HC} at which the income tests claw back a payout and the cut-off x_U of the assets test. Here we change the tests over a wider range, each alone and then the two together, with everything else as in Table 1. For the income tests we scale m_{SN} and m_{HC} together by up to 20% either way, and with them their total m. For the assets test we shift x_U by the same proportions, from 16 to 24. A lower x_U is a stricter assets test, since support ends at a lower wealth, but the band $\mathcal{Z} = [0, x_U]$ in which payouts are clawed back shrinks with it.

The public cost is the expected present value of the support paid along the optimal wealth path, net of the clawback and discounted at the riskless rate. The value of a change to the retiree is the most wealth they would give up at inception to have it, which is negative for a change that leaves them worse off.[23] We solve the moves of 5%, 10% and 20% at ten wealths from $x_0 = 4$ to 105, and each range quoted below runs over these wealths without a bequest motive. The case with a bequest motive gives the same reading of the two tests, so we give it here first. With a bequest motive coverage is lower, but a cut of 20% in m still raises it at every wealth, whereas a lower x_U raises it only at the higher wealths.[24]

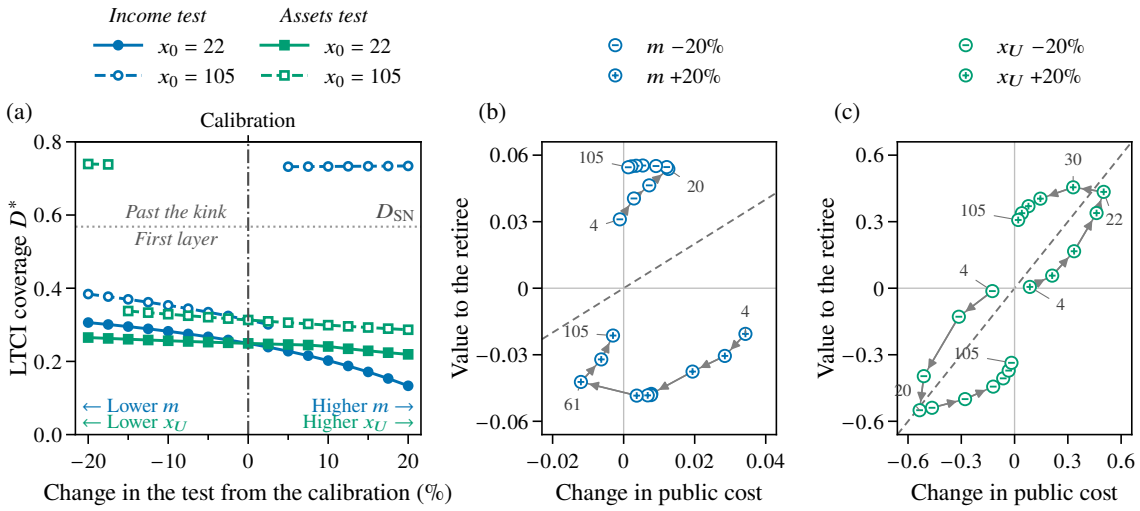


Figure 6. Optimal coverage as m or x_U moves, and value against public cost, without a bequest motive.

Notes: In Figure 6(a) m and x_U move in steps of 2.5%. Lines join optima on the same side of the kink D_{SN}, where the cash safety net runs out. The rule marks D_{SN} at the calibrated m, and a point above it lies past the kink of its own case. In Figure 6(b) and (c) the dashed line is value equal to cost, arrows join each case's points in order of rising inception wealth x_0, and the number beside a point is its x_0.

[23] Section S.8.11 gives the public cost's formula and compares further designs, and Table S.10 gives the values as compensating variations.

[24] Section S.8.14 gives the results with a bequest motive. Without a bequest motive each case passes the concavity certificate of Assumption 5.6 across its feasible coverage range.

Lowering m raises coverage at every wealth (Figure 6(a)). The income tests set the implicit tax, and with it the static loading on the first layer. That loading runs from 94% when m is 20% lower to 213% when it is 20% higher. It does not depend on x_U and stays at 139% wherever x_U is set. Cuts of 5% to 20% in m raise coverage by 0.02 to 0.07. A higher m lowers coverage on the first layer. However, at the highest wealths it carries coverage past the kink, as a rise in m_{SN} does in Figure 5. Section 9.4 explains such a crossing, and similar reasoning applies here. At $x_0 = 105$ a rise of 5% in m takes coverage from 0.31 to 0.73.

Moving x_U changes coverage little except at $x_0 = 105$ (Figure 6(a)). Lowering x_U by 20% raises coverage by at most 0.05 below $x_0 = 105$, less than lowering m by 20% does at each of those wealths. At $x_0 = 105$ the optimum passes the kink as x_U falls from 15% to 17.5% below its calibrated value. Raising x_U by 20% lowers coverage by up to 0.04.

Lowering m costs the government little (Figure 6(b)). The clawback is levied only in disability and only on the payouts, so a cut of 20% in m adds at most 0.013 to the public cost. A lower m is worth more to the retiree than its public cost at every wealth. That cost is largest at $x_0 = 20$, where the calibrated tests collect the most clawback, and the points of m −20% turn there in Figure 6(b). Above that wealth it falls, since a richer retiree is less often in the band $\mathcal{Z}$ when disabled. A higher m leaves the retiree worse off at every wealth. This is because each payout loses more to the clawback, so the retiree buys much less coverage. From $x_0 = 61$ a rise of 20% in m carries coverage past the kink instead. More is then clawed back from the far larger payout and the public cost is lower than the baseline's, which is where the points of m +20% turn in Figure 6(b).

Moves of 5% and 10% in m, which Figure 6(b) does not draw, change the retiree's value by 0.01 to 0.03 (a gain for a cut and a loss for a rise) and the public cost by about 0.01 at most. Figure 6(c) leaves out the moves of 5% and 10% in x_U too. They act like the moves of 20% over a smaller range, and we do not discuss them further.

Because x_U decides who gets support at all, a move of 20% in it changes the public cost by up to 0.53 (Figure 6(c)). A lower x_U saves public money at every wealth, but each retiree above the calibrated x_U loses more than the government saves, 1.2 to 19 times the saving when x_U is 20% lower. What is at stake is the support paid between the old and the new x_U. A risk-averse retiree who starts above both values it above its expected cost, since it is paid only where wealth has fallen. One who starts below the higher of the two values it below that cost, since it is paid where wealth has held up or risen. The same reasoning makes a rise of 20% in x_U, to 24, worth less than it costs up to $x_0 = 22$ and more from $x_0 = 30$.

When we move m and x_U together, they act almost independently at $x_0 = 22$ (just above the calibrated x_U) but not at $x_0 = 105$ (Figure 7). At $x_0 = 22$ coverage depends mainly on m. Raising m and x_U together by 20% lowers it from 0.25 to 0.10 (Figure 7(a)), close to the sum of the falls of 0.12 and 0.03 that each causes alone. At $x_0 = 105$ a boundary next to the calibration cell separates coverage on the first layer from coverage past the kink (Figure 7(b)). A higher m or a 20% lower x_U takes the retiree across it, but with a lower m the optimum stays on the first layer at every x_U in the figure, because a lower m raises the kink. On the boundary the two layer optima are equally good, so crossing it changes coverage by about 0.4 and the retiree's value hardly at all.

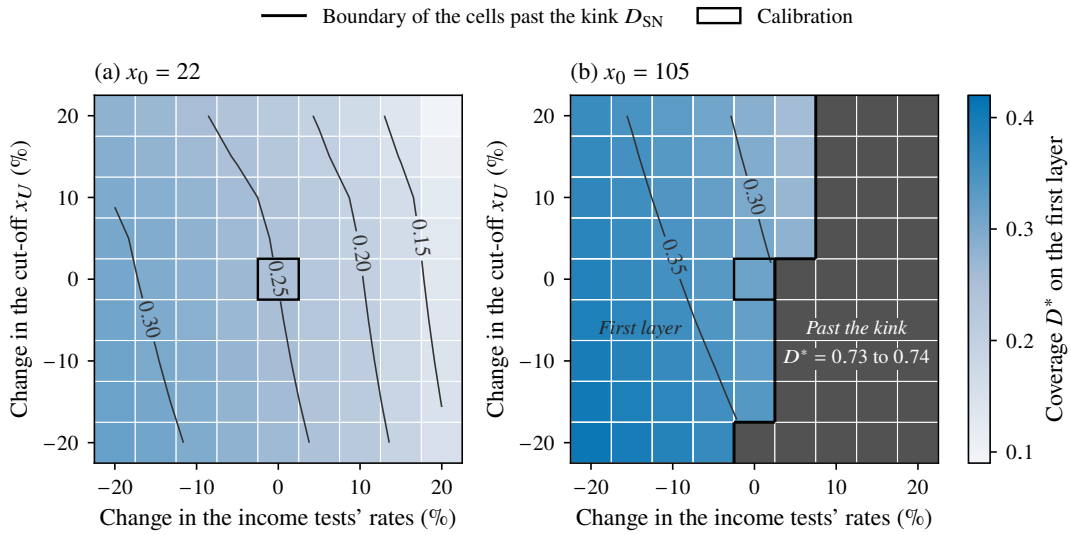


Figure 7. Optimal coverage as m and x_U move together.

Notes: Coverage without a bequest motive, with m and x_U moved in steps of 5%. Each cell is compared with its own kink D_{SN}, where the cash safety net runs out. Section S.8.14 gives the values.

10.2 Implications for the design of the means test

For the agency that sets the means test, our comparison points to the income tests. The income tests set the implicit tax. Lowering m costs the government little, is worth more than that cost to every retiree whose coverage it moves, and raises coverage even with a bequest motive if it goes far enough (Section S.8.14). The cut-off x_U instead decides who faces the tax. Lowering it does little for coverage except at the highest wealths, and with a bequest motive it raises coverage only at the higher wealths. It also takes more from the retirees above the calibrated x_U than the government saves, because they get the support only once their wealth has fallen. Even so, coverage does rise as the band narrows, which is the direction of the first condition of the cross-country hypothesis of Section 3.3 (a market can form where the clawback covers a narrow range of wealth).

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 assets-test cut-off reduces to two numbers, and we obtain a closed-form price for the implicit tax. We solve the remainder numerically.

Within our model calibrated to the UK and Australian schedules, the spillover is large: At five times the cut-off, exempting payouts would raise coverage by more than half its level. Timing is the reason. A payout arrives only in disability. By then a retiree unsupported above the cut-off may have spent or lost enough to fall below it, where part of each payout is clawed back. 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 we draw. These levels assume unbounded leverage. With the risky share capped at one, coverage under the calibrated tests is zero at most wealths, and without a bequest motive the gain from exempting payouts remains at every wealth. Weaker

income tests lift coverage at every wealth, for every cut without a bequest motive and for a cut of 20% with one. Without a bequest motive, exempting payouts is worth to the retiree 2 to 20 times its extra public cost. A slightly lower cut-off lifts coverage too, but the government saves less than a retiree above the calibrated cut-off loses.

Our 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 the assumptions behind them. The marginal benefit of coverage steps up where a channel's support runs out. 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. However, the loadings and coverage levels we report are specific to our calibration.

Two simplifications mark the scope of these results. First, we compare the designs for one stylised retiree by coverage, public cost and value to the retiree rather than by a social objective over a population. An objective weighing that value against the public cost would turn the comparison into a design problem over the schedule's rates and caps. The construction of the disabled value would also handle an assets test in several smaller steps, under which a saved payout costs support below the cut-off too. Second, we hold private income fixed, which leaves out an annuity bought with the wealth that also buys LTCI. A life annuity 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.

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

The code that solves the model, the stored optima behind the figures and tables, and the scripts that regenerate them are provided as a replication package with the submission.

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. https://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., & Koreshkova, 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

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

Pauly, M. V. (1990). The rational nonpurchase of long-term-care insurance. *Journal of Political Economy*, *98*(1), 153–168. https://doi.org/10.1086/261673

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