Euler Equalities Near Ruin

Zsolt Becsi *
Southern Illinois University Carbondale

August 2025

Abstract. The Euler equality is the most widely used condition in consumption economics, and it is valid on a subset of the state space. Near ruin, it stops being the right first-order condition. When ruin is irreversible and insurance is insufficient, the first-order conditions split into equalities far from ruin and inequalities near it. The gap is the value of survival. Far from ruin, the textbook holds, and a household satisfies orthogonality, adjusts buffers and consumption proportionally, with insurance remaining flat. Near ruin, the same household violates orthogonality, moves buffers before consumption, and shifts demand toward front-loaded payouts. Four decades of Euler puzzles spawned separate literatures, each with its own explanation. A common domain error runs through them: the equality was applied where only the inequality holds.

JEL Classification: C61, D15, D81, E21, G22
Keywords: Euler inequality; safe set; ruin risk; absorbing boundary; bounded insurance; dynamic programming
Replication Materials: `https://osf.io/67uyj/?view_only=94a3a5dedd9d459e88675178fc8a57eb`

1 Introduction

The textbook Euler equality is local: near ruin, it stops characterizing the optimum. Orthogonality can look plausible in safe states and then break sharply in stress states; risk-sharing benchmarks can look close in some regimes and fail in others. That selectivity

*Email: `becsi@siu.edu`. I thank my teachers Rao Aiyagari and John Rust for setting this work in motion, and Sunnar Hohlfeld and Leslie John for helpful discussions.

is the symptom. The underlying mistake is using a local condition globally. The profession has attributed this pattern to behavioral biases, heterogeneous types, or missing markets. This paper shows that many of these patterns do not require present bias, heterogeneous types, or missing markets. The Euler equality is exact only on part of the state space. In a dynamic program with an absorbing survival threshold and bounded insurance, the correct conditions switch from equalities to inequalities at a sharp frontier, and the gap is the value of continued existence. What looks like preference failure can instead reflect applying the equality outside its domain. The argument needs neither new preferences nor behavioral parameters; only a constraint the standard equality tools were not designed to see.

What is new here is not that constraints can turn Euler equalities into inequalities. That logic is standard. Standard constrained-control logic tells the modeler how to write the Kuhn–Tucker system when a constraint binds. The point here is different: absorption changes the geometry of the problem. What is new is the geometry created by an *absorbing* boundary combined with bounded state-contingent insurance. A reflecting borrowing limit is a spring: the process stays in the admissible set, the constraint can bind repeatedly, and the extra term values having a next period in which to adjust. An absorbing survival threshold is a wall: once wealth crosses the ruin threshold, the continuation problem ends, so the extra term values having a next period at all. First, it induces a sharp partition of the state space into a safe set where the survival term is slack and the Euler-equality characterization is valid, and a danger set where the survival term is active and the correct characterization is an inequality. Second, it changes aggregation: any equality-based rule that ignores the survival term generates one-way drift into the absorbing set, so small local hazard differences compound into large cross-sectional effects. Those two implications, safe-set validity and drift into the absorbing set, are the paper's contribution. Rules that ignore ruin are like flight plans that ignore the ground: they work until they don't. There is no second chance. Much of what we "know" from life-cycle Euler logic is safe-set knowledge: it describes behavior in periods when households are not close to the edge.

The absorbing boundary is not an abstraction. For an individual household, it is the threshold below which economic life as a going concern ends. The cleanest case is bankruptcy or exclusion: cross the line, lose the asset base, lose ordinary credit access, and the original self-insurance problem ends. This is not an optimal stopping problem. The agent does not choose when to exit; the agent is trying not to be ejected. The same structure appears whenever crossing a threshold ends the problem rather than merely constraining it: a chronic health shock whose costs permanently exceed the capacity to self-insure, exhaustion of unemployment benefits with no reemployment, a single bad

harvest that liquidates the subsistence farmer's productive asset. Nothing in the analysis requires aggregation, a representative agent, or equilibrium prices; the unit of analysis is one household solving one dynamic program, and the distinction from prior work turns on what happens at the boundary. Here "ruin" means the end of the original decision problem, not necessarily death.

Strotz (1955) saw that long plans unravel when the future arrives; Kydland and Prescott (1977) warned that such unraveling undermines rules. Both took survival for granted: they debated whether plans were credible, not whether credibility itself was feasible. Their inconsistency is strategic, requiring two players and vanishing with commitment. The inconsistency we identify requires one agent and no strategic interaction: a plan chosen in safety becomes suboptimal when the state moves toward an absorbing boundary, because the survival term enters the first-order condition and changes what the marginal dollar buys. Their inconsistency vanishes with commitment; ours vanishes when ruin is ruled out.

A standard clarification. In any single-agent dynamic program, the optimal Markov policy is the committed policy: re-optimization and commitment deliver the same decision rule. This is the principle of optimality, and we do not contradict it. What we show is that the Euler equality economists use to read that policy, and from which they calibrate rules, build tests, and diagnose anomalies across the consumption, insurance, and life-cycle literatures, is valid only on a strict subset of the state space.

Prior work has documented conditional Euler failures across three literatures without unifying them. In the consumption–savings literature, Zeldes (1989) showed that liquidity-constrained households exhibit the Euler inequality $u'(c_t) \geq \beta(1+r)\mathbb{E}_t[u'(c_{t+1})]$, with equality only for unconstrained households. Deaton (1991) formalized this in the buffer-stock model, proving that the no-borrowing constraint generates an inequality condition rather than a global equality. Carroll and Kimball (1996) showed how concavity and prudence shape the consumption rule under the same inequality logic. Attanasio and Weber (1995) documented the empirical counterpart: constrained groups systematically violate Euler orthogonality. In risk-sharing and contract models, similar gaps appear: Kocherlakota (1996) and Ligon et al. (2002) derived state-contingent inequalities from limited-commitment and participation constraints. These literatures taught us that the Euler equality is conditional, but they did not elevate that switch to the organizing principle of time consistency. In risk-sharing tests, the same selectivity appears: Mace (1991), Cochrane (1991), and Townsend (1994) rejected complete markets, yet Blundell et al. (2008) found systematic state dependence in partial-insurance coefficients. In our model these are not anomalies of preferences but signatures of the survival frontier; the

same object, the survival term, generates the selective pattern across all three literatures.

The incomplete-markets models of Bewley (1986), Aiyagari (1994), and Huggett (1993) also generate Euler inequalities, but from a reflecting constraint: the agent hits a borrowing limit, is pushed back, and visits the constraint repeatedly. A stationary ergodic distribution exists. Our constraint is absorbing: crossing the subsistence floor ends the process. The reflecting case produces an extra term that prices the cost of postponing consumption. The absorbing case produces an extra term that prices the probability of continuing to exist. This distinction matters for results. Reflecting constraints do not partition the state space into safe and danger regions, do not generate one-way drift into the absorbing set under fixed rules, and do not imply the strict dominance of state-contingent insurance over self-insurance near the boundary. The distinction is not technical; it determines whether the agent bounces back or ceases to exist. Our framework exploits the latter: crossing the boundary is irreversible.

The rest of the paper proceeds accordingly. Section 2 sets up the dynamic program with bounded insurance and an explicit survival floor. Section 3 proves the main theorems: the Euler-equality partition, characterization of the safe set, and a backward-induction lemma that replaces global Euler equalities with the proper inequality system outside $\mathcal{S}$. Section 4 develops corollaries for insurance reallocation, the collapse of smoothing, continuation-value divergence, and distributional dynamics. Section 5 translates the safe-set switch into the diagnostic objects used in the Euler and insurance literatures. Section 6 provides a transparent three-period quantitative illustration. We do not claim this is the only source of these patterns; we claim it generates a common regime-switching pattern across puzzles without behavioral parameters.

2 Model

The textbook consumption model has preferences, shocks, borrowing limits, and a Bellman equation. It has everything except the possibility that the agent stops existing. The model here is deliberately spare. One agent faces shocks, can buy imperfect insurance, and is ruined if wealth falls below a subsistence floor. The only real question is where the Euler-equality characterization of the optimum is valid and where it is not. By nailing down timing, feasibility, and what "insurance" actually does, we make that question precise. Preferences are standard throughout: no present bias, no hyperbolic discounting, no probability weighting, no heterogeneous types. Everything that follows comes from one

omitted state variable in the textbook characterization: survival.

Time is discrete, $t = 1, \ldots, T$. Wealth at the start of t is k_t, with $k_1 \geq \kappa$. Let $\kappa > 0$ denote the ruin threshold, the boundary beyond which there is no continuation value. The agent chooses consumption $c_t \geq 0$ and an insurance premium $m_t \geq 0$ at the start of t, based on the information available then. Choices must leave at least the subsistence floor in place within the period, so feasibility requires $k_t - c_t - m_t \geq \kappa$; equivalently, $m_t \leq k_t - c_t - \kappa$. This explicit cap matters later: near the boundary the agent simply cannot buy much coverage.

Preferences are standard. Instantaneous utility $u : \mathbb{R}_+ \to \mathbb{R}$ is strictly increasing, strictly concave, continuously differentiable, and $\lim_{c \downarrow 0} u'(c) = +\infty$; the discount factor is $\beta \in (0, 1)$. The gross return is normalized to $R = 1$; we retain R in the Euler expressions for familiarity. Ruin carries a finite penalty $U_{\text{ruin}} \in \mathbb{R}$ unless otherwise stated, so collapse can occur with positive probability while expected utility remains well defined. The catastrophic case $U_{\text{ruin}} = -\infty$ simply rules out any policy that allows ruin and collapses the choice set to survival paths.

Uncertainty is summarized by i.i.d. shocks $\{Y_t\}$ with law F on $(Y, \mathcal{G})$. A measurable map $g : Y \to \mathbb{R}$ delivers non-insured income $g(Y_{t+1})$. A designated "bad set" $B \subset Y$ has probability strictly between zero and one; these are the realizations that trigger insurance. The shock Y_{t+1} is realized at the *start* of period $t + 1$. We write $\mathbf{1}\{Y_{t+1} \in B\}$ for the indicator that the bad state materializes.

Insurance is a one-period state-contingent contract. The agent pays the premium m_t in period t. At the start of $t + 1$, the contract pays λm_t if and only if $Y_{t+1} \in B$, and zero otherwise. The parameter $\lambda \in (0, 1]$ is insurance efficiency: when $\lambda = 1$, a unit of premium buys a full unit of payout in the bad state (actuarially fair coverage); when $\lambda < 1$, administrative loadings or market frictions make coverage less than one-for-one (cf. Gruber, 1997; Chetty, 2008). This parameter is what separates the paper's world from the textbook. When $\lambda = 1$ and coverage is unlimited, ruin cannot happen and the paper has nothing to say. Every departure from that benchmark reintroduces the possibility that a bad draw overwhelms what the agent can finance. The entire paper lives in that gap. Lowering consumption can "self-insure" across time, but it cannot replicate the state-contingent payout when the bad state hits (see also Cochrane, 1991; Ligon et al., 2002).

With this timing, wealth carried into the next period is

$$k_{t+1} = k_t - c_t - m_t + g(Y_{t+1}) + \lambda m_t \, \mathbf{1}\{Y_{t+1} \in B\}. \tag{1}$$

Read from left to right: start-of-period wealth, minus choices, plus the realized non-insured income, plus the insurance payout if the bad state arrives. Because Y_{t+1} is drawn at the start of the next period, k_{t+1} is measurable with respect to that next period's information.

We use $h_t(c_t, m_t)$ to denote the one-step hazard of collapse,

$$h_t(c_t, m_t) := \Pr(k_{t+1} < \kappa \mid \mathcal{F}_t),$$

which will index near-frontier states and organize the empirical predictions in Section 4.

Lifetime utility under an admissible policy $\pi = \{(c_t, m_t)\}_{t=1}^{T}$ is

$$U^{\pi}(k_1) = \mathbb{E}\left[\sum_{t=1}^{T} \beta^{t-1} u(c_t) + U_{\text{ruin}} \cdot \mathbf{1}\{\exists\, s \le T :\, k_s < \kappa\}\right].$$

The second term is a literal switch: the indicator $\mathbf{1}\{\exists\, s \le T : k_s < \kappa\}$ equals one if wealth ever crosses the subsistence line during the horizon and zero otherwise, so the product contributes U_{ruin} exactly when collapse occurs. With $U_{\text{ruin}} \in \mathbb{R}$, ruin lowers welfare but does not send expected utility to $-\infty$; in the catastrophic (extended-real) case, any policy with positive ruin probability is excluded ex ante.

Dynamic programming proceeds in the usual way, following Rust (1987), who solved a structural Bellman program with an absorbing 'failure' state, a direct parallel to the ruin boundary in our model. Let $A(k) = \{(c, m) \in \mathbb{R}_+^2 :\, k - c - m \ge \kappa\}$ be the feasible set at wealth k. Set $V_{T+1}(k) = 0$ for $k \ge \kappa$ and $V_{T+1}(k) = U_{\text{ruin}}$ for $k < \kappa$. Then for $t = T, \ldots, 1$,

$$V_t(k) = \sup_{(c,m)\in A(k)} \left\{u(c) + \beta\, \mathbb{E}\big[V_{t+1}\big(k - c - m + g(Y_{t+1}) + \lambda m\, \mathbf{1}\{Y_{t+1} \in B\}\big)\big]\right\}.$$

For reference, Section 3 defines the safe set $\mathcal{S}$ as the wealth states from which one can choose $(c, m) \in A(k)$ that make the one-step hazard zero (avoid $k_{t+1} < \kappa$ almost surely) under the full Bellman optimum; outside $\mathcal{S}$ the one-step hazard is strictly positive.

Under the standing assumptions (concavity, integrability, compact feasibility), the Bellman problem is well defined and admits measurable selectors; technical details are standard (see Bertsekas and Shreve, 1978; Benveniste and Scheinkman, 1979; Stokey et al., 1989).

The Euler equality is valid only in safe states where survival is assured. The rest of the paper develops that claim; two polar cases make it concrete. First, if bad-state losses are bounded by $L < \infty$ and coverage is actuarially fair ($\lambda = 1$), then an agent far enough from

the boundary to afford $m_t \geq L$ within the feasibility cap can offset the worst loss in that period and keep $k_{t+1} \geq \kappa$. Maintaining such buffers period by period collapses uncertainty into a deterministic control problem in which the equality characterization is globally valid. That is the world most of the literature assumes. Second, any deviation, for example, inefficient insurance ($\lambda < 1$), binding premium caps near the boundary, or unbounded tails, reintroduces the possibility that a bad draw overwhelms what the agent can finance. That is the world most households actually inhabit. The mathematical setup above makes the boundary between them precise.

The finite horizon and i.i.d. assumptions are chosen for transparency, not necessity. The mechanism requires only an absorbing boundary and a feasibility constraint that tightens near it. An infinite-horizon version inherits the same safe set and the same equality-inside/inequality-outside structure, with the Bellman operator a contraction under standard conditions. Replacing i.i.d. shocks with a Markov chain makes the safe set state-dependent but leaves the within-period logic unchanged. A multi-dimensional state (multiple assets, multiple risk margins) turns the boundary into a surface, but the partition between safe and unsafe regions, and the switch from equalities to inequalities at that surface, survive intact. In each case, the engine is the same: an absorbing threshold from which recovery is impossible. What changes across specifications is the shape of the safe set, not the economics that the boundary organizes.

3 Results

Every Euler-equation test in the consumption literature implicitly assumes the answer to one question: is the agent in the safe set? This section proves that the answer determines everything. We first define the safe set, the states from which ruin is avoided with probability one under the full Bellman policy, and show that inside this set the equality characterization is correct, while outside it the full Bellman policy strictly dominates any equality-characterized rule. The object doing the work is feasibility, not preferences: once the lower bound on wealth can be hit with positive probability, reallocating toward survival raises continuation value more than any equality-characterized policy that spreads resources across time without pricing the state where survival is decided.

The sequence is deliberate. Theorem 1 draws the partition between states where the Euler equality is the correct first-order condition and states where it is not. Corollary 1 (Perfect insurance) gives the benchmark in which the partition disappears (survival is guaranteed

each period), collapsing the stochastic program to a deterministic one in which the equality characterization is globally valid. Theorem 2 characterizes the safe set itself, an upward-closed viability region, and explains why distance to its boundary organizes behavior. Read together, these results say something simple and powerful about standard dynamic programming: the Bellman recursion always solves the problem, but the Euler equalities we often use to characterize optimality hold only on the safe set. Outside it, the correct objects are inequalities that price survival. Sections 4–6 take this frontier to the margins we measure: insurance, consumption, continuation values, distributions. The answer is cleaner than one might expect. There is one set, one boundary, one test. Inside, the equality holds and every tool built from it works. Outside, it fails and every tool built from it is misspecified.

The logic is simple. An agent near the boundary faces a choice the Euler equality does not see: every dollar shifted from consumption to survival has a bounded utility cost (strict concavity) but a first-order survival gain (the boundary is absorbing). Any policy that ignores this tradeoff, that sets the survival term to zero, leaves value on the table exactly where the stakes are highest. The question is not whether this matters but where: at which states does the survival term switch on, and how large is the gap it creates?

The first theorem answers a precise question: when does the Euler-equality characterization deliver the Bellman optimum, and when does it not? We compare two Markov policies. The *full Bellman policy* solves the dynamic program (3) with the absorbing subsistence floor. The *equality-characterized policy* solves the same program but imposes the Euler equality globally, setting the survival term μ_t to zero in every state. Both are state-contingent and optimized. They differ in one thing: whether the first-order conditions price the possibility of ruin. That is the whole comparison, and the result is sharp: standard constrained-control logic tells us how to handle a binding constraint; it does not tell us that absorption shrinks the domain on which the Euler equality characterizes the optimum.

At the start of t the agent chooses (c_t, m_t) subject to $k_t - c_t - m_t \geq \kappa$; the shock Y_{t+1} arrives at the start of $t{+}1$ and wealth updates according to (1). Let V_t denote the full Bellman value and V_t^E the value induced by the equality-characterized policy. We write the *safe region* as the set of states $\mathcal{S} \subseteq [\kappa, \infty)$ from which ruin is avoided almost surely under the full Bellman optimum.

Theorem 1 (Euler Equalities Hold Only on the Safe Set)**.** *Under the assumptions in Section 2, there exists a nonempty safe region $\mathcal{S} \subseteq [\kappa, \infty)$ such that:*

(i) On $\mathcal{S}$, the two policies coincide: $\mu_t = 0$ along the full Bellman path, the Euler equality is the correct first-order condition, and the equality-characterized policy is the Bellman optimum.

(ii) Off $\mathcal{S}$, the full Bellman policy strictly dominates: $\mu_t > 0$ on a set of positive probability, the correct condition is $u'(c_t) \geq \beta R\, \mathbb{E}_t[u'(c_{t+1})]$, and the value gap is accounted for entirely by the survival terms the equality characterization omits.

Economic meaning. Both policies solve well-posed dynamic programs. They differ only in whether the first-order conditions carry a price for survival. Inside the safe set, that price is zero and the two programs are the same problem: the equality characterization is correct, and every rule, test, and calibration built from it performs as designed. Outside, the price is positive and the equality characterization becomes a misspecification: it omits the term that prices the one transition from which there is no recovery. The dominance in (ii) is not a failure of optimization. It is the cost of reading the Bellman solution through a lens that was ground for the interior. This theorem gives applied researchers a map. Before running an Euler-equation test, check whether the observation is in the safe set. If it is, the test is valid. If it is not, the test is measuring the survival term, not the preference parameter.

Read this theorem as a map rather than a verdict. It marks where the Euler-equality characterization is valid (inside the safe region) and where it is not (outside). In the danger set, the equality-characterized policy cannot restore the survival probability that the full Bellman policy preserves. That observation will power the later corollaries: the tilt toward insurance in exposed states, the breakdown of smoothing conditions when survival binds, and the gap in continuation values. The Euler equality, in short, is state-dependent: valid in safe states, misspecified where survival is at stake.

Proof Sketch. Let $\mathcal{S}$ be the largest forward-invariant subset of $[\kappa, \infty)$ under admissible policies. On $\mathcal{S}$, the constraint $k_t \geq \kappa$ never binds along the optimal path. The survival term is zero at every date, the Euler equality is the correct first-order condition, and by uniqueness of the Bellman maximizer the equality-characterized policy delivers the same value. The two policies coincide.

Off $\mathcal{S}$, the constraint binds with positive probability. The full Bellman policy prices this exposure: $\mu_t > 0$ values the marginal relaxation of survival risk. The equality-characterized policy, by construction, treats that risk as free. The resulting misallocation is not subtle. Near the boundary, the marginal value of wealth grows without bound (Lemma A.3), so a small reallocation toward survival yields a first-order gain; the consumption cost is

second-order by strict concavity. The dominance is driven entirely by the survival terms the equality characterization sets to zero: not by preferences, not by market frictions, and not by any failure of the Bellman recursion itself.[1] □

If the agent can purchase actuarially fair coverage each period and the within-period premium cap is high enough to cover the worst loss on the path, ruin is excluded. With collapse ruled out, uncertainty stops mattering for feasibility: the stochastic problem reduces to a deterministic control. In that benchmark, the equality-characterized policy and the full Bellman policy coincide globally.

Corollary 1 (Full Insurance Prevents Collapse)**.** *If insurance is actuarially fair ($\lambda = 1$) and full coverage of the worst loss is feasible subject to the within-period premium cap at every state visited on path, then ruin is excluded on all histories and the stochastic program reduces to a deterministic control problem. In this case, $V_t^E(k) = V_t(k)$ for all $k \geq \kappa$ and all t; the equality characterization is globally valid.*

Economic meaning. When insurance is complete and affordable, ruin is ruled out, and dynamic choice collapses to a deterministic problem. Any friction or cap restores fragility and reopens the gap between the equality characterization and the full Bellman solution.

The corollary presents the complete-markets benchmark in one period at a time. It clarifies what complete insurance does: it removes the states in which survival would bind. Any departure reintroduces fragility and reopens the gap between the equality characterization and the full Bellman policy. Inefficient insurance, tight premium caps near the boundary, or a loss distribution whose lower tail overwhelms what can be financed will each do the work.

Proof Sketch. With λ=1 and feasible m_t large enough to cover the worst loss each period, wealth never crosses κ. The post-decision state is deterministic conditional on (c_t, m_t) because the state-contingent payout neutralizes the bad set B. The Bellman equation reduces to a deterministic recursion with the same u and β, implying $V_t^E \equiv V_t$ by standard arguments. See Appendix A3. □

Corollary 1 identifies the world in which the partition disappears: when coverage is complete and affordable, ruin cannot happen, and the Euler equality holds everywhere. That is the special case the literature has been living in. The next result asks what happens when that benchmark fails, when insurance is bounded, loads are positive, or

[1] Full details are in Appendix A3, which establishes the necessary measurability, existence of selectors, and the stochastic viability arguments.

coverage caps bind. It turns out the answer has a clean geometry: the safe set is upward-closed, its boundary is sharp, and distance to that boundary organizes everything that follows. Theorem 2 is the map of the territory outside the textbook.

The subsequent theorem builds the safe region as the maximal forward-invariant set above the subsistence floor. It is upward-closed: more wealth never hurts feasibility. It also gives a practical membership test. If one-period losses are bounded by L and the within-period premium cap allows $m_t \geq L$ at a given state, then that state is safe in that period. The boundary is where small changes in exposure flip the feasibility of neutralizing the bad state. That is precisely where the full Bellman policy gains value, because reallocating toward survival changes whether the process continues.

Theorem 2 (Boundary of the Safe Region)**.** *There exists a maximal forward-invariant set $\mathcal{S} \subseteq [\kappa, \infty)$ with the following properties: (i) for any $k \in \mathcal{S}$ there exists an admissible policy that keeps $k_s \geq \kappa$ almost surely for all future dates; (ii) $\mathcal{S}$ is upward closed; and (iii) if one-period losses are bounded by $L < \infty$ and at state k the feasibility cap allows $m_t \geq L$, then $k \in \mathcal{S}$. Moreover, on the boundary $\partial\mathcal{S}$ any equality-characterized policy that does not reallocate toward survival is strictly dominated by the full Bellman policy.*

Economic meaning: The safe region is the set where survival is always feasible, and its boundary marks where small shocks can trigger collapse. Behavior changes predictably as agents move toward this boundary.

The boundary is the organizing object for the rest of the paper. It determines where insurance shares should tilt up, where familiar smoothing conditions should fail, and where continuation values diverge. Nothing else is needed to explain the switch in behavior.

Proof Sketch. Construct $\mathcal{S}$ as the largest set invariant under the wealth-transition correspondence induced by admissible policies. Upward closedness follows from monotonicity of the feasibility constraint. Condition (iii) holds because the ability to finance L neutralizes the worst bad-state loss in that period; iterating the invariance shows safety forward. Dominance on $\partial\mathcal{S}$ follows from a local perturbation that substitutes a marginal premium for consumption on histories where the constraint threatens to bind; the value gain from higher survival probability dominates by concavity and the finite U_{ruin}. Details are in Appendix A3. □

The distinction between the two types of boundary, reflecting and absorbing, is what makes the safe-set geometry nontrivial. The following proposition makes that distinction formal.

Proposition 1 (Absorbing $\neq$ reflecting)**.** *With a reflecting constraint (borrowing limit), binding changes the choice rule but the process continues and recovery is possible. With an absorbing ruin threshold κ, crossing κ ends the problem, so states split into a safe set where the Euler equality is valid and a danger set where it is not. Consequently, any equality-based characterization applied globally creates one-way drift into the absorbing set from the danger set, while the full Bellman policy reduces that drift by reallocating in danger states.*

Proof sketch. Reflecting constraints keep the state inside the feasible region after binding, so violations are local and do not create an exit channel. Absorption makes ruin a first-passage event: once κ is crossed, continuation value is terminated. The safe/danger partition follows from Theorems 1–2; one-way flow follows because ruin is absorbing. The Bellman policy lowers hazard in danger states (Corollaries 2–3) and sustains survival mass (Corollary 6). □

Remark (Comparative statics of the safe region). The safe set responds monotonically to primitives. If $\lambda' \geq \lambda$, the Bellman operator $T_{\lambda',F}$ dominates $T_{\lambda,F}$ pointwise because the transition k' is weakly higher in bad states under more efficient insurance. Monotonicity of the operator iterates gives $V_t^{\lambda'} \geq V_t^{\lambda}$, which enlarges the set of initial conditions from which survival is assured. A parallel argument holds when F' first-order stochastically dominates F on the lower tail: better shock distributions expand $\mathcal{S}$. Hence $\mathcal{S}(\lambda', F') \supseteq \mathcal{S}(\lambda, F)$. Inefficient insurance and fat left tails shrink the safe region; improving either primitive expands it.

The dynamic program is solved in the usual way by the Bellman recursion. What changes when ruin is admissible is not the recursion but the first-order characterization. Introduce a survival constraint that is slack inside the safe region and binding off it. Inside the safe region, the consumption first-order condition is the familiar equality that links today's marginal utility to its discounted expectation. When survival binds, a non-negative survival term enters the condition, and the equality becomes an inequality: the marginal utility of consuming today exceeds the discounted expected marginal utility tomorrow by exactly the amount needed to enforce survival.

Lemma 1 (Equalities Hold Only in Safe States)**.** *Consider a finite- or infinite-horizon dynamic program with state k, control $a = (c, m, \ldots)$, primitive feasibility $(k, a) \in F$, shocks z, and an absorbing ruin set $R = \{k < \kappa\}$. Let the safe set be $\mathcal{S} \subseteq \{k \geq \kappa\}$ as in Section 3. Suppose u is strictly concave and bounded above, $\beta \in (0, 1)$, and the transition $k_{t+1} = T(k_t, a_t, z_{t+1})$ is measurable. Then:*

(i) ***Bellman recursion.*** *The value function V exists and satisfies*

$$V(k) = \sup_{a \in A(k)} \{\, u(c) + \beta\, \mathbb{E}[V(k')]\},$$

with boundary condition $V(k) = -\infty$ for $k \in R$ (or any utility floor that tends to $-\infty$ along approximations). Optimal policies can be selected measurably. Backward induction / value iteration converges.

(ii) ***FOC structure.*** *Let $g(k, a) \geq 0$ encode survival (slack on $\mathcal{S}$, binding off $\mathcal{S}$). There exists a multiplier $\mu_t \geq 0$ such that along an optimal policy*

$$u'(c_t) - \beta R\, \mathbb{E}_t[u'(c_{t+1})] \;=\; \mu_t\, \frac{\partial g(k_t, a_t)}{\partial c_t},$$

with $\mu_t = 0$ whenever $k_t \in \operatorname{int} \mathcal{S}$. If $\partial g / \partial c_t < 0$ near danger, this implies the Euler inequality

$$u'(c_t) \;\geq\; \beta R\, \mathbb{E}_t[u'(c_{t+1})],$$

with equality if and only if the survival constraint is slack.

(iii) ***Propagation.*** *Backward induction propagates the Bellman recursion globally, but Euler equalities propagate only on histories that remain in $\mathcal{S}$. Where survival binds, the correct propagated objects are inequalities, not equalities.*

Economic meaning: Backward induction solves the problem everywhere, but classical Euler equalities hold only in the safe region. Where survival is threatened, inequalities replace equalities and the full Bellman policy departs from the equality characterization.

Proof Sketch. Approximate the absorbing set by a sequence of truncated problems that assign utility floor $-M$ on R, let $M \to \infty$, and apply monotone convergence of value iteration to obtain (i) and measurable selection. For (ii), write the Lagrangian with survival constraint $g(k_t, a_t) \geq 0$. Envelope + Euler give the stated condition with survival term $\mu_t \geq 0$; the sign of $\partial g / \partial c_t$ near danger yields the inequality. Part (iii) follows by noting $\mu_t = 0$ on $\operatorname{int} \mathcal{S}$ and $\mu_t > 0$ on the binding set. Full regularity and measurability details are in Appendix A3. □

4 Corollaries

The results here translate the feasibility frontier into the ordinary objects we use: insurance purchases, consumption smoothing, continuation values, and cross-sectional dynamics. Each corollary applies the same variation: shift resources from consumption toward survival near the boundary, show the survival gain is first-order while the consumption cost is second-order, and read the result through the margin it governs.

One object governs all five results: the survival term μ_t. When $\mu_t = 0$, every margin looks textbook. When $\mu_t > 0$, every margin departs. Because each corollary traces μ_t through a different observable, the premium share, the Euler residual, the pass-through coefficient, the cross-sectional survival rate, they provide independent readings of the same structural object. If the model is right, these readings should agree when sorted by the same hazard proxy. If they disagree, the proxy is wrong or an additional friction is present. Section 6 checks this agreement in a computed example.

In most intertemporal problems, "prudence" is read as cutting spending today to protect tomorrow. That logic moves dollars across dates, not across states. When the ruin threshold is close, the decisive question is *where* the next dollar lands. A dollar that pays in the bad realization raises survival; a dollar saved for the average realization does not. This is the simple asymmetry that drives the corollary. Throughout what follows, the behavioral pattern changes without any change in the agent, only in the agent's distance from the boundary. What changes when survival enters: the decisive margin shifts from when to consume to where the dollar lands.

The standard prescription for income risk is to save: cut spending today and build a buffer. That advice assumes the buffer survives long enough to help. Near the boundary, it may not, and the first result says when.

Corollary 2 (Self-Insurance Alone Cannot Prevent Ruin)**.** *Fix date t and a state $k \notin \mathcal{S}$. Along a set of histories of positive probability, the full Bellman policy that chooses (c_t, m_t) with $m_t > 0$ strictly dominates any equality-characterized rule that sets $m_t = 0$ and relies solely on intertemporal self-insurance via c_t. In particular, there exists $\varepsilon > 0$ such that replacing ε of consumption by premium raises the continuation value.*

Proof. See Appendix A4. □

Proof sketch (template used throughout). At $k \notin \mathcal{S}$, a small shift from consumption to

insurance has bounded utility cost (concavity) but a first-order survival gain (Lemma A.3); the full Bellman allocation exploits this. Full proof in Appendix A4.

In plain terms: near the boundary, saving alone is not enough. The agent must redirect resources toward coverage that pays in the bad state, because a dollar saved for the average future is worth less than a dollar that prevents ruin tomorrow.

Why do households near default spend more on insurance while safe households do not? The standard answer requires heterogeneous risk preferences. This corollary requires only proximity to the boundary.

Corollary 3 (Insurance Demand Rises Near Collapse)**.** *For any $k \notin \mathcal{S}$, the full Bellman policy assigns weakly more weight to insurance than any equality-characterized rule; the inequality is strict on a set of positive probability. Equivalently, the premium share in (c_t, m_t) rises under the full Bellman policy as the state approaches the boundary.*

Proof. See Appendix A4. □

What to look for in data: the share of resources allocated to insurance, $m_t/(c_t + m_t)$, should rise as the hazard proxy increases and flatten when it indicates safety.

Consumption smoothing tests pass in some subsamples and fail in others, often within the same dataset and the same households. The standard reading is that some agents smooth and some do not. This corollary says the test itself is state-dependent: it should pass when survival is slack and fail when survival binds, in the same agent, across different periods. The Euler equation rests on an implicit premise, that the survival constraint is slack, and the premise is not always true. When the constraint binds, the first-order condition carries an extra term that values movements in the survival probability. The equality turns into an inequality. What changes when survival enters: the Euler equality becomes a diagnostic of safe states, not a universal law.

Corollary 4 (Smoothing Fails When Survival Is at Risk)**.** *If $k \notin \mathcal{S}$, the Euler condition implied by any equality-characterized rule fails to describe optimal behavior; the full Bellman policy departs by reallocating resources toward survival. If $k \in \mathcal{S}$, the Euler condition holds and the two policies coincide. Any rule that enforces equality outside $\mathcal{S}$ is strictly dominated by the full Bellman policy, and equality resumes upon re-entry into* $\operatorname{int}\mathcal{S}$.

Proof. See Appendix A4. □

What to look for in data: the Euler residual $e_t := u'(c_t) - \beta R\,\mathbb{E}_t[u'(c_{t+1})]$ should be near zero when $h_t = 0$ and positive, increasing in h_t, when the hazard is active.

If a financial advisor tells a household near default to follow the Euler equality, how much does that advice cost? The previous corollaries describe how the agent's choices change near the boundary. The next result describes what those choices cost: not in current consumption but in the value of the entire continuation. The gap is not small, and Section 6 puts a number on it.

Corollary 5 (Full Bellman Policy Increases Survival Value)**.** *For any $k \notin \mathcal{S}$ and all t, $V_t(k) > V_t^E(k)$ with strict inequality on a set of positive probability, and the difference is accounted for by survival-probability terms in the law of motion rather than by curvature alone.*

Proof. See Appendix A4. □

In plain terms: the agent who follows the Euler equality near the boundary is not just making a slightly worse choice. The entire future is worth less, because more paths end in ruin, and each ruined path carries the full penalty.

The four results above are about one agent. The final corollary asks what happens when a population faces the same boundary: equality-characterized rules push mass into the absorbing set; the full Bellman policy sustains it.

Corollary 6 (Full Bellman Policy Sustains Survival Mass)**.** *Consider a continuum of agents with an initial distribution over $k \geq \kappa$ and absorbing ruin. Under the equality-characterized rule, if a positive-measure set begins outside $\mathcal{S}$, the cross-section monotonically loses mass to ruin over time. Under the full Bellman policy, with compact feasibility and continuous transition dynamics and with choices that return states to $\mathcal{S}$ with positive frequency, the cross-section converges to an invariant distribution with strictly positive survival mass.*

Proof. See Appendix A4. □

What to look for in data: the fraction of the cross-section above κ should decline steadily under equality-characterized rules and stabilize under the full Bellman policy. The rate at which mass drifts into the absorbing set is a direct measure of how large μ_t is in the aggregate.

5 Five Diagnostic Implications

The focus is not on adding mechanisms, but on restoring the missing state variable. Each margin below has been attributed, at various points, to myopia, present bias, probability distortion, or heterogeneous types. The safe-set partition ties all five to one device, survival risk near an absorbing threshold, and the discriminator is within-unit switching indexed by hazard: the same agent looks "behavioral" in danger states and textbook in safe states because the first-order condition changes, not the agent. The same agent, with the same preferences and information, moves between equality and inequality as the state moves. This converts what look like differences in "type" (e.g., hand-to-mouth versus optimizer) into differences in proximity to the boundary. It also tells readers what to expect from standard tests: inside the safe set, orthogonality and smoothing should pass; outside, they should fail for structural reasons. The question is whether that same switch appears across all five margins.

The empirical record on consumption and insurance is marked less by blanket rejection than by selective failure. Cross-section tests of risk sharing (Mace, 1991; Cochrane, 1991; Townsend, 1994) reject the complete-markets benchmark, yet in the same environments there are stretches where household consumption hardly moves with idiosyncratic shocks. Euler-equation tests show similar selectivity: predictable income sometimes has bite (Flavin, 1981; Hall and Mishkin, 1982; Campbell and Mankiw, 1989, 1990; Shea, 1995; Parker, 1999; Souleles, 1999, 2002), and sometimes it does not. This section makes that selectivity structural with one device already in the paper: a feasibility frontier that partitions the state space into a safe region, $\mathcal{S}$ (defined in Section 2), where survival is slack ($h_t = 0$), and a danger region, where survival binds ($h_t > 0$). Inside safety, the standard Euler equality and familiar risk-sharing logic apply. At the frontier, the survival term enters the optimality condition and turns equalities into inequalities; the agent reallocates toward survival, raises short-run MPC, and values liquidity for what it does to feasibility. What follows traces that switch through five familiar margins.

We now apply Lemma 1's safe-versus-danger switch to five diagnostic margins used in the Euler and insurance literatures; Appendix A5 provides the derivative-and-sign steps. In each case, the empirical discriminator is the same: the margin should move with hazard within the same unit.

Define the Euler residual $e_t := u'(c_t) - \beta\, R\, \mathbb{E}_t[u'(c_{t+1})]$. A practical measurement strategy follows from Lemma 1: sort periods by a one-step near-ruin proxy and expect $e_t \approx 0$ when the proxy indicates safety and $e_t > 0$ when it indicates danger. Distance to κ in liquid

resources or any short-horizon default probability will serve. The governing object is the survival term $\mu_t \left|\frac{\partial g}{\partial c}\right|$: when $\mu_t = 0$, the standard Euler object obtains; when $\mu_t > 0$, the survival term enters and the standard object changes. We trace it through five margins: excess sensitivity, partial insurance, borrowing constraints, buffer sequencing, and contract design.

Households sometimes spend tax refunds and bonus checks as if they had no savings account; other times the same households barely react to predictable payments. The profession calls this excess sensitivity and attributes it to rule-of-thumb behavior. When $h_t = 0$, predictable income does not move current consumption: the Euler orthogonality holds. When $h_t > 0$, the same predictable resources relax the survival constraint, and the survival term $\mu_t \left|\frac{\partial g}{\partial c}\right|$ tilts the adjustment toward c_t. "Excess sensitivity" (Flavin, 1981; Campbell and Mankiw, 1989) is therefore state-contingent: the model predicts that predictable payments raise spending strongly for exposed households but weakly for safe ones (Hall and Mishkin, 1982; Shea, 1995; Parker, 1999; Souleles, 1999, 2002). The standard interpretation, rule-of-thumb consumers who fail to optimize, requires a stable type; this model requires only proximity to the boundary. This differs from type-based explanations such as rule-of-thumb consumers, savers and spenders, or wealthy hand-to-mouth households, which sort behavior across households by balance-sheet structure; here the same household switches behavior as hazard rises toward the boundary. Within-household, the effect should load on months when the hazard proxy is high and vanish when it indicates safety.

Risk-sharing regressions sometimes reject complete markets and sometimes do not, in the same sample, across the same households, depending on the time period. Inside $\mathcal{S}$, transitory shocks are smoothed more than permanent shocks: the classic insurance ranking holds ($\theta_T \approx 1 > \theta_P$). Near the boundary, the survival term $\mu_t \left|\frac{\partial g}{\partial c}\right|$ raises pass-through for both shock types, so a larger fraction of each shows up in c_t. The model implies that the same person looks well insured in safe months and poorly insured near the frontier without any change in preferences. The split is not between stable "types" but between states (Blundell et al., 2008; Zeldes, 1989; Kaplan and Violante, 2014). What changes when the hazard rises: the same person crosses from textbook insurance to apparent under-insurance without any shift in preferences or information. The discriminator is within-person: the same household should show high pass-through in months with $h_t > 0$ and textbook insurance in months with $h_t = 0$.

The pass-through pattern in the previous margin looks like a binding borrowing constraint, but it originates in the agent's own first-order condition. Near the frontier, the shadow

value of liquid resources steepens, desired borrowing rises, and an exogenous borrowing cap binds more often. Terms and prices need not change. When h_t falls to zero, the corner fades and standard smoothing returns. The classic "constrained vs unconstrained" partition (Zeldes, 1989; Attanasio and Weber, 1995) is therefore a state-dependent corner, not a stable type. Rising corner frequency generates gap-like outcomes without any movement in rates. Feasibility is the first-pass explanation; price gaps are the residual once feasibility has been ruled in. The borrowing corner emerges from the agent's own first-order condition, not from the lender's terms.

If the borrowing corner is endogenous, how does the agent respond when resources arrive? Following a favorable resource shock near the boundary, liquid buffers rise on impact while nondurable consumption is muted (Deaton, 1991). The survival term explains the pattern: liquidity relaxes the constraint more effectively than nondurable spending, so the shadow premium on liquid balances is positive until μ_t falls to zero. What changes when survival binds: liquid balances buy safety when danger is high; nondurable spending does not. The same logic produces a kink in the marginal propensity to consume at the safety threshold, the geometric counterpart of the survival term entering the FOC, distinct from inertia or habits. The falsifier is sequencing: the liquid-share response should precede the nondurable response only when the hazard proxy is elevated; in safe periods, both adjust together.

If buffers matter because they relax the survival constraint, then insurance contracts that deliver liquidity at the right time should dominate. In safe states, actuarially equivalent contracts with different payout timing are near-substitutes: preferences and minor frictions govern. When $h_t > 0$, the immediate-liquid contract strictly dominates because it relaxes the survival constraint at the right time. The survival term is $\mu_t \frac{\partial g}{\partial \text{payout}}$, which is positive at binding states and zero in safety. Low take-up of delayed or illiquid coverage among high-hazard agents is therefore optimal, not anomalous. This state-contingent valuation explains patterns in insurance adoption without appealing to probability mistakes: demand shifts toward fast, liquid payouts only when exposure is high and is muted otherwise (Cole et al., 2013, 2014; Gallagher, 2014).

Three diagnostics make the switch testable. The Euler residual e_t is near zero when $h_t = 0$ and rises monotonically with h_t; the interaction of predictable income with h_t is positive and first-order in h_t. Borrowing-constraint frequency rises near the boundary and falls in safety, driven by the agent's FOC rather than lender terms. Insurance demand is sharply higher for front-loaded, liquid payouts when $h_t > 0$; indifference returns in safety. Each diagnostic is falsifiable on the same dimension: h_t sorts the anomalies; crossing into safety switches them off. No margin required a change in preferences; each required only the

recognition that the first-order condition changes at the frontier. The common prediction is not that every anomaly is always present, but that each should switch on in danger states and switch off in safe ones. The next section makes this partition visible in a computed instance.

In applications with a meaningful failure threshold (default, exit, or ruin), treat the Euler equality as an interior characterization: it should hold in safe states and fail in danger states. Empirically, sort observations by a short-horizon hazard proxy, whether distance to the threshold, distress or default risk, or any measure that predicts "no next period" states. The key discriminator is within-unit switching: the same household or firm should look textbook in safe periods and "anomalous" in danger periods because the first-order condition changes at the frontier, not because preferences or types change.

6 Quantitative Example

The theory makes a sharp prediction: sort any panel by a short-horizon hazard proxy and the anomalies should switch on in danger periods and switch off in safe ones. Section 6 shows that mechanism in the cleanest possible numerical setting. We solve the Bellman equation exactly by backward induction over a three-period, binary-shock environment and compare two policies: the equality-characterized rule, which sets the survival term to zero everywhere, and the full Bellman policy, which prices survival where it binds. The only difference between them is whether μ_t enters the first-order condition. Same preferences, same information, same shocks; the policies diverge because one prices survival and the other does not. The four figures then do four jobs: they show why fixed rules fail near the boundary, where the safe set ends, how the equality-characterized rule drifts into the absorbing set, and why the value gap is survival rather than smoothing. Binary shocks and log utility make every computation exact; the analytic results in Sections 3–5 hold for arbitrary horizons and distributions.

Parameters are chosen for transparency, not calibration. The agent begins with wealth $k_1 = 1.2$ and faces adverse shocks of size $\Delta = 0.6$, each with probability $q = 0.6$. The subsistence floor is $\kappa = 0.5$; ruin carries penalty $U_{\text{ruin}} = -10$. Insurance pays $\lambda = 0.8$ per premium dollar in the bad state. Expected values are computed exactly from the two-point distribution; no simulation is required.

Figure 1 plots wealth paths under three fixed rules, the bare buffer ($m = 0$), the wallet-burner ($m = \text{cap}$), and the prudent buffer (interior m), alongside the full Bellman

policy. No fixed rule guarantees survival; only the Bellman policy, which re-optimizes after shocks, avoids crossing κ (Theorem 1). The point is not calibration; it is demonstration. Near the boundary, no fixed allocation can jointly preserve a buffer and insure the downside. This is the feasibility cap of Theorem 2: near the boundary, no fixed allocation stays above κ. No fixed allocation can serve both goals at once: insuring against the bad state erodes the buffer that prevents ruin in the worse state.

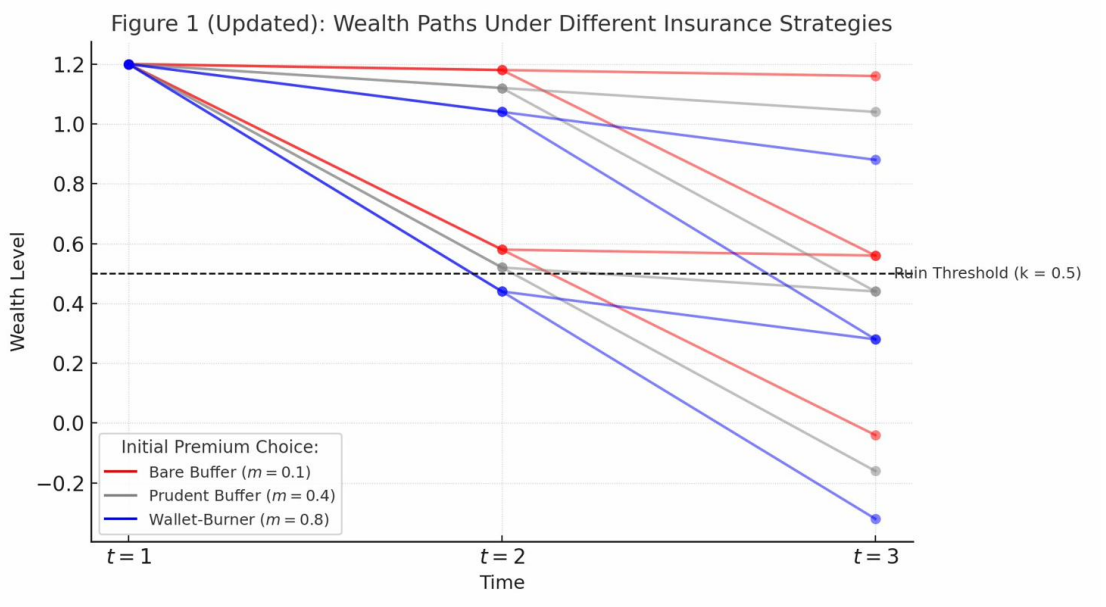


Figure 1: Wealth paths from $t = 1$ to $t = 3$ under several fixed insurance strategies. Parameters: $k_1 = 1.2$, $\Delta = 0.6$, $q = 0.6$, $\lambda = 0.8$, $\kappa = 0.5$, $U_{\text{ruin}} = -10$.

Figure 2 plots ruin probability and expected utility as functions of the fixed premium m. In this illustration with $\lambda = 0.8$, ruin probability is flat at 0.36 for $m < 0.25$, then steps to 0.84 when the premium erodes the buffer past the survival threshold. The step is the feasibility cap of Theorem 2 in action. This is the safe-set frontier translated into premium space. Above it, the agent has insured away the buffer that would have kept wealth above κ.

Figure 3 compares wealth paths under the equality-characterized rule and the full Bellman policy from identical initial conditions. The equality-characterized path crosses κ after consecutive bad shocks; the Bellman path survives. That is the dynamic consequence of omitting the survival term: once the equality is imposed outside its domain, the path drifts into an absorbing state from which it does not return. The divergence is the continuation-value gap of Corollary 5, driven entirely by the survival terms the equality characterization omits. The paths diverge not because the policies use different information but because one prices survival and the other does not.

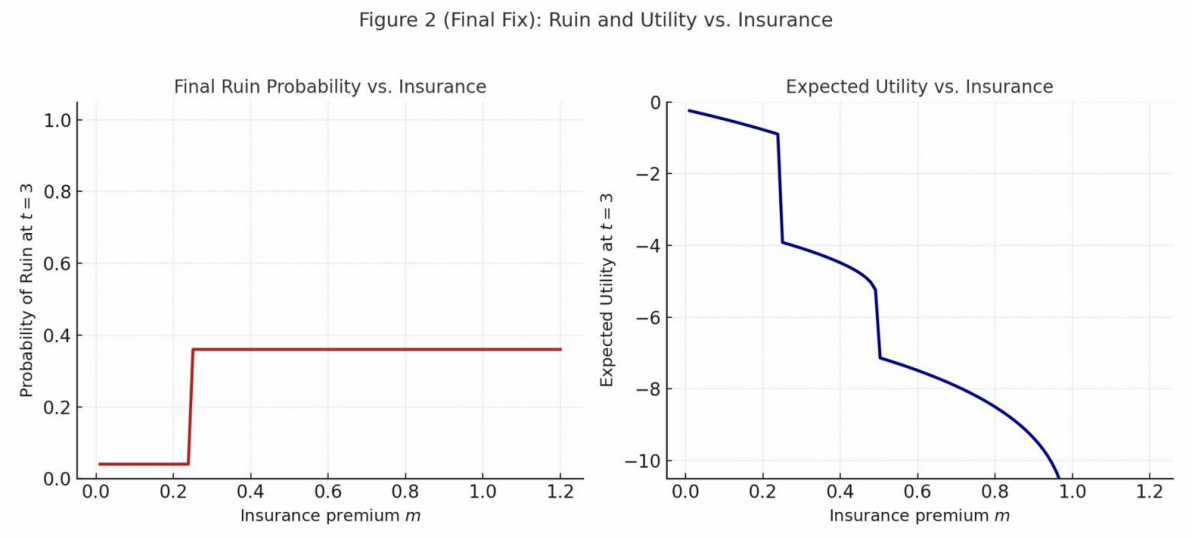


Figure 2: Ruin probability (left) and expected utility (right) at $t = 3$ as functions of fixed insurance premium m. Parameters: $k_1 = 1.2$, $\Delta = 0.6$, $q = 0.6$, $\lambda = 0.8$, $\kappa = 0.5$, $U_{\text{ruin}} = -10$.

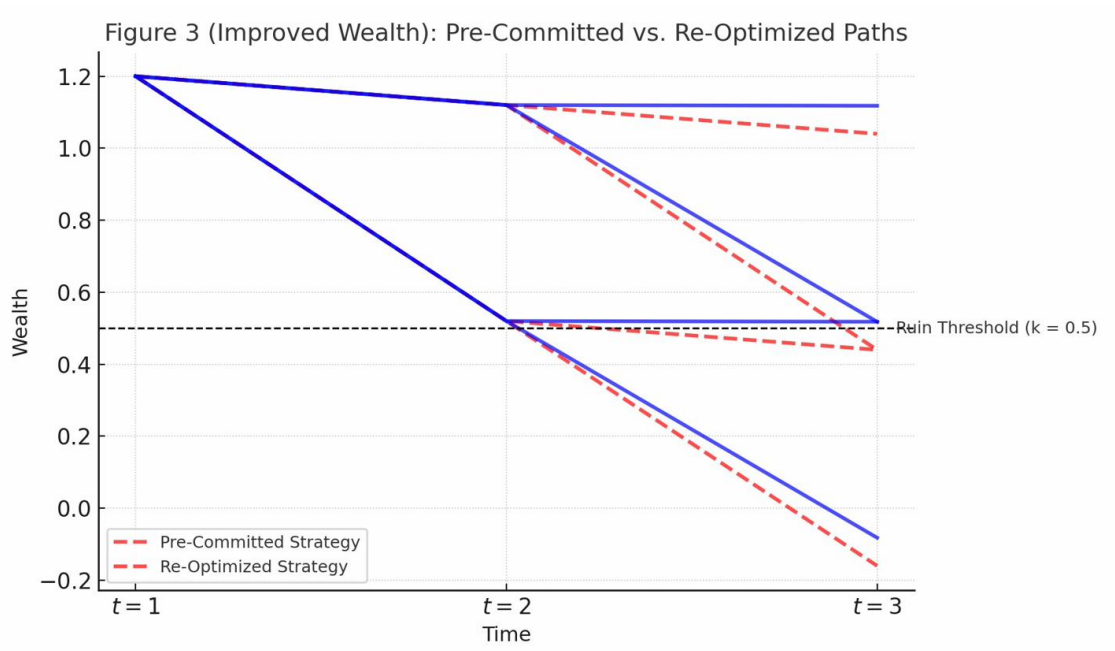


Figure 3: Wealth paths for equality-characterized (dashed) and full Bellman (solid) policies, starting from identical conditions. Parameters: $k_1 = 1.2$, $\Delta = 0.6$, $q = 0.6$, $\lambda = 0.8$, $\kappa = 0.5$, $U_{\text{ruin}} = -10$.

Figure 4 plots expected utility for both policies as U_{ruin} varies. The value gap grows in proportion to the survival-probability difference $P(\tau^E \leq T) - P(\tau^\star \leq T)$, reaching 15.1 utils at $U_{\text{ruin}} = -30$ for the $m_1 = 0.4$ strategy. The payoff of the section is here: the gap is a survival gap, not a consumption gap (Corollary 5).

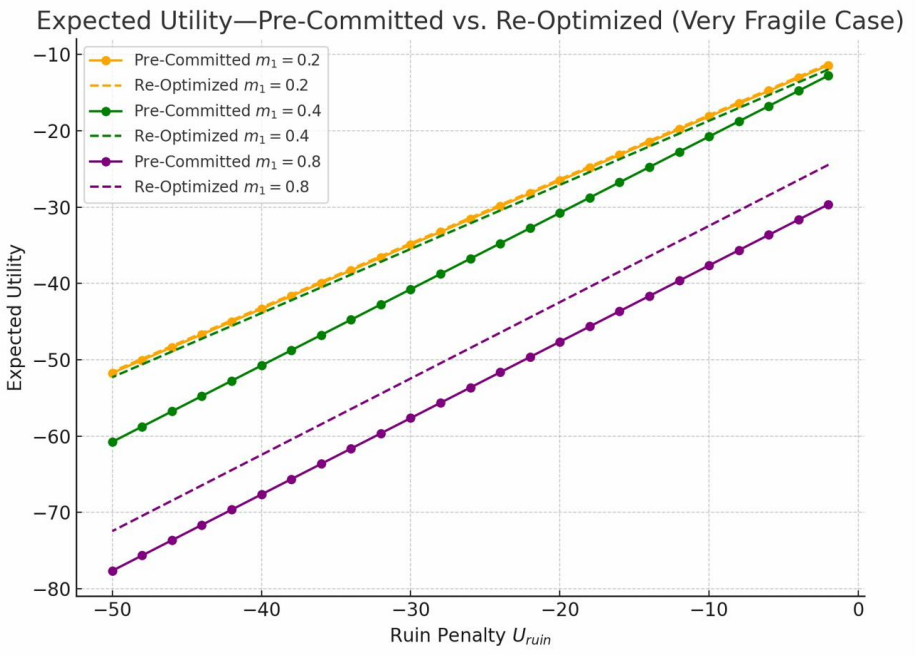


Figure 4: Expected utility for equality-characterized (dashed) and full Bellman (solid) policies across insurance levels ($m_1 = 0.2, 0.4, 0.8$) as ruin penalty U_{ruin} varies. Parameters: $k_1 = 1.2$, $\Delta = 0.6$, $q = 0.6$, $\lambda = 0.8$, $\kappa = 0.5$.

The computed example also lets us check whether the different margins from Section 4 point to the same object. After the first shock, the agent reaches one of two states: $k_2 = 1.12$ (good shock) or $k_2 = 0.52$ (bad shock). Under the full Bellman policy, the agent at $k_2 = 1.12$ sets $m_2 \approx 0$ and faces $h_2 = 0$: the state is safe, and every margin looks textbook. At $k_2 = 0.52$, wealth is only 0.02 above κ (less than 4 percent of initial wealth), the agent cannot finance meaningful coverage, and $h_2 = 0.60$: the state is dangerous. Both the insurance share and the Euler residual switch on at the same threshold, confirming that a single object, the survival term, governs both margins. Section 6 therefore does more than illustrate policy functions: it shows the safe-set split, the drift into the absorbing set under the wrong rule, and the common hazard index in one compact example.

Under the equality-characterized rule, the agent holds $m_2 = 0.4$ regardless of the state. At $k_2 = 1.12$, this premium is 36 percent of wealth and pushes the one-step hazard from zero to 0.60: the rule converts a safe state into a dangerous one. At $k_2 = 0.52$, the premium absorbs 77 percent of wealth and pushes the hazard to 1.0. The equality-characterized policy does not merely miss the survival term; by ignoring the feasibility cap, it actively creates the ruin it was not designed to see.

How large is the cost? At $U_{\text{ruin}} = -30$, the equality-characterized policy delivers ruin on 84 percent of histories. The full Bellman policy delivers ruin on 36 percent. The ruin rate

under the domain error is 2.3 times the rate under the correctly specified policy.

The expected-utility gap of 15.1 utils decomposes cleanly. Approximately 95 percent is accounted for by the survival-probability difference: more histories end in ruin, and each ruined history carries the full penalty. The remaining 5 percent reflects the consumption reallocation the Bellman policy makes in danger states. The gap is not a small correction to a marginal misspecification. It is the difference between losing most of the cross-section and retaining most of it.

The qualitative results survive wide parameter variation. Vary q over $[0.3, 0.9]$: the direction of every result is unchanged; only ruin levels shift. Vary λ over $[0.5, 1.0]$: the feasibility cap moves but the cap-driven step in Figure 2 persists at every $\lambda < 1$. Vary U_{ruin} over $[-50, 0]$: the value gap scales with the survival-probability difference, confirming Corollary 5. The only knife-edge is $\lambda = 1$ with feasible full coverage: ruin is eliminated, the equality characterization is globally valid, and the entire framework collapses to the textbook benchmark.

All experiments use a three-period dynamic program with log utility, a subsistence floor κ, and a binary shock with probability $q = 0.6$ and loss size $\Delta = 0.6$. Insurance pays in the bad state with efficiency $\lambda = 0.8$. Expected values are computed exactly from the two-point distribution; no Monte Carlo approximation is required. Together, the four figures make the partition between $\mathcal{S}$ and its complement visible in a setting where every number can be verified.

7 Life-Cycle Illustration

The three-period example makes the mechanism exact. This section puts it where it belongs: on the life-cycle profile that organizes textbook consumption theory. Figures 5 and 6 overlay three policies on that profile: the classic LCH benchmark (dashed), which enforces the Euler equality globally; the buffer-stock path under a reflecting borrowing constraint (gray), following Zeldes (1989) and Deaton (1991); and the adaptive path under absorbing ruin (solid), which prices survival where it binds. Each figure has three panels. The panels are identical except for one thing: the age at which danger arrives. Panel A places the danger zone in early working life, Panel B in midlife, Panel C in retirement. Reading across panels answers a question the three-period example cannot: how does the timing of danger shape the rest of the life cycle? Reading across lines within any panel answers a more fundamental question: how does what happens at the boundary determine

whether smoothing survives.

Start with the dashed line. It is the textbook: the agent who never approaches the absorbing boundary, whose survival term is zero at every date, whose Euler equality holds globally. This agent lives entirely in the safe set. The Euler equality is the correct first-order condition at every date, and consumption smoothing is the correct description of behavior at every date. Much of what the profession calls life-cycle theory is a description of these households.

Now add the gray line. The buffer-stock path departs from the benchmark during the danger episode but returns to it afterward. Look at Panel A, wealth: the gray line dips below the dashed line between ages 20 and 30, then climbs back and runs parallel to it by age 45. Look at Panel B, consumption: the gray line drops from 5.0 to about 3.2 at age 50, then recovers to 5.0 by age 60. The buffer-stock agent smooths most of the time. Smoothing fails only during the constrained episode, when the reflecting boundary binds and the agent cannot borrow against future income. The reflecting constraint is costly while it binds, but the cost is local: the agent is pushed back, income arrives, wealth rebuilds, and the original program resumes; the multiplier binds in the danger period and releases when income arrives, leaving no residual in the first-order condition. The gray line always returns to the dashed line. Smoothing is interrupted, not destroyed. That is the signature of a reflecting constraint: the crisis is a parenthesis, not a rewrite.

Now look at the solid black line, and notice that it does not come back. The adaptive path under absorbing ruin is the policy where the survival term is positive in the shaded danger region and zero outside it. The gap between the solid and dashed lines at any age is the accumulated cost of the survival terms the benchmark omits. The ruin-adaptive agent cannot smooth. Not during the danger episode, and not after it, because the resources consumed by survival are permanently subtracted from the wealth that smoothing requires. In Panel A (early-life danger), the agent enters working life near the absorbing boundary. Survival demands immediate reallocation: consumption drops sharply as the agent diverts resources to insurance and buffer preservation. The buffer-stock agent, facing the same shock, recovers by age 40. The ruin-adaptive agent is still below the benchmark at 60. They bought continued existence, not consumption. The post-danger decline is not a second crisis; it is the agent drawing down a wealth stock that was depleted by the cost of staying alive.

Panel B makes the difference in shape unmistakable. The buffer-stock consumption path drops to roughly 3.2 and recovers to 5.0 within ten years: a symmetric V that closes and

vanishes. The ruin-adaptive consumption path drops to 1.5 and recovers only to about 3.8, then continues declining through retirement. The buffer-stock path is a V. The ruin path is an L. One closes; the other does not. The buffer-stock agent's crisis is an episode. The ruin-adaptive agent's crisis is a level shift.

Panel C (retirement danger) isolates the mechanism in its starkest form. Before retirement, all three paths are identical: no danger has arrived, and the safe-set characterization holds for all agents. At retirement, income ends and danger begins. The classic benchmark dissaves smoothly. The buffer-stock agent's consumption declines gradually: wealth drifts toward the floor, but hitting the floor is not terminal, so the agent spends down without a survival tradeoff. The ruin-adaptive agent's consumption drops sharply, a cliff rather than a slope, because the agent is actively holding wealth above κ. The agent chooses poverty of consumption to prevent poverty of existence. Look at the wealth figure, Panel C: the solid line stays above κ while the gray line descends through it. That vertical gap is the survival premium, paid in forgone consumption every period.

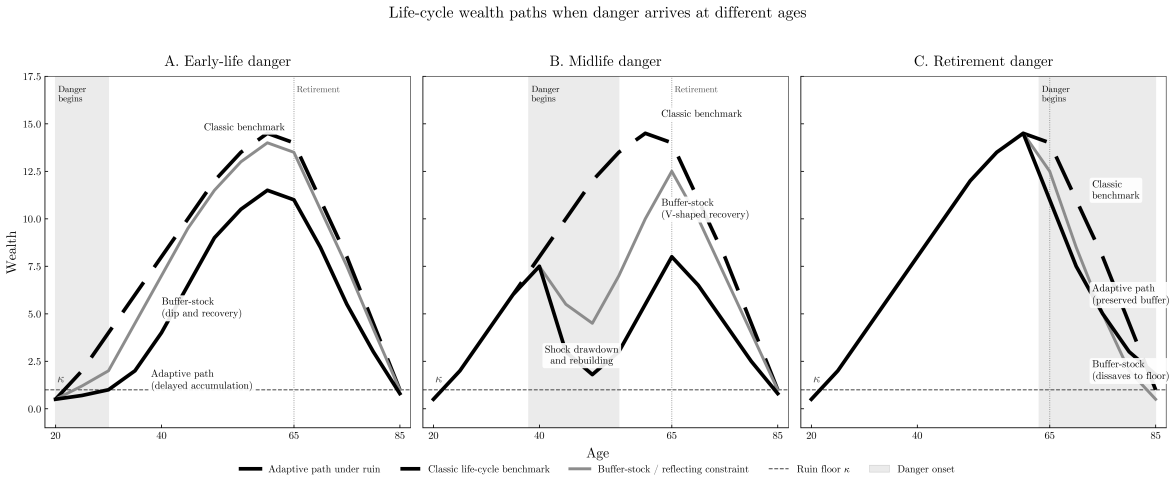


Figure 5: Life-cycle wealth paths when danger arrives at different ages. Solid: adaptive policy under absorbing ruin. Dashed: classic LCH benchmark. Gray: buffer-stock path under reflecting borrowing constraint. Dashed horizontal: ruin floor κ. Shaded region: danger episode. Three panels vary the timing of danger onset: early working life (A), midlife (B), retirement (C). The buffer-stock path (gray) dips during the danger episode and reconverges to the classic benchmark; the adaptive path (solid) departs permanently. Schematic; not drawn from a calibrated model.

The figures show only survivors. Every adaptive path depicted is a conditional path: the agent who followed the full Bellman policy and did not cross κ. The figures do not show the paths of agents who were less fortunate, whose bad draws overwhelmed even the survival-pricing reallocation, and who exited the problem permanently. Those agents are

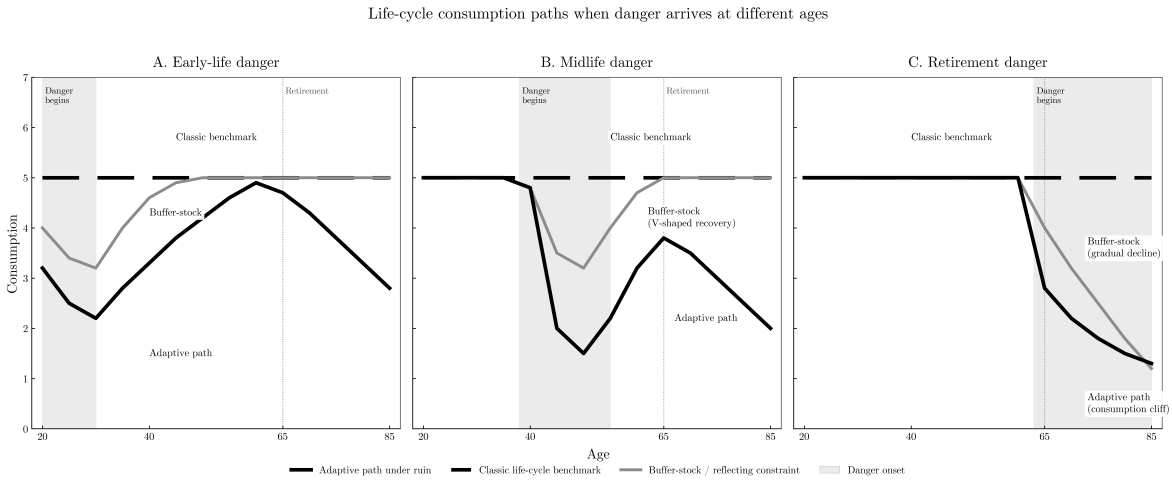


Figure 6: Life-cycle consumption paths when danger arrives at different ages. Solid: adaptive policy under absorbing ruin. Dashed: classic LCH benchmark (smooth consumption). Gray: buffer-stock path under reflecting borrowing constraint. Shaded region: danger episode. The classic benchmark smooths consumption across the entire lifetime. The buffer-stock agent smooths outside the danger episode; smoothing is interrupted but resumes when the constraint releases. The ruin-adaptive agent cannot smooth: consumption drops sharply during danger and never reconverges to the benchmark, because the resources spent on survival permanently reduce the wealth base. Schematic; not drawn from a calibrated model.

not on the graph because their paths ended. Corollary 6 quantifies their fate: under the equality-characterized rule, the exposed tail thins geometrically as mass drifts into the absorbing set with no mechanism to replenish it; under the full Bellman policy, reduced hazard and positive return rates arrest that drift and sustain an invariant distribution with positive survival mass. The visible paths are the success stories of the survival-pricing policy. The invisible paths, the agents who left the game, are the reason the survival term exists.

One mechanism produces three smoothing regimes. The three lines in each panel are not three models. They are three answers to one question: what does the boundary do to the agent who hits it? Nothing (dashed), a temporary interruption (gray), or a permanent scar (solid). Same preferences, same shocks, same income process. Only the boundary differs. What separates the solid line from the gray line is not degree but kind. Two signatures appear in every panel: asymmetry (the buffer-stock path dips and recovers symmetrically; the ruin-adaptive path drops fast and recovers slowly or not at all) and permanence (the buffer-stock path reconverges to the classic benchmark; the ruin-adaptive path does not). Read across any row of Figures 5 and 6: the gray line returns to the dashed line in every panel; the solid black line does not. That persistence is the absorbing boundary, and no reflecting constraint can replicate it. The constraint episode is not a parenthesis. It is a

scar.

8 Conclusion

The Euler equality has been the organizing condition of consumption economics for four decades. This paper shows it is valid on a subset of the state space. Outside that subset, the correct conditions are inequalities, and the profession's most persistent puzzles are the predictable consequence of ignoring them.

Relative to the literature, the contribution is precise. Backward induction survives ruin, but Euler equalities do not. They remain valid only on islands of safety. Once survival can fail, the inequalities that carry a survival term are the correct first-order conditions. This shows that one domain error can generate a common subset of the rejections in Euler tests and risk-sharing benchmarks (e.g., Hall and Mishkin, 1982; Shea, 1995; Souleles, 1999, 2002; Blundell et al., 2008): the puzzles emerge from applying the equality characterization outside its domain of validity. The paper does not resolve every empirical debate in these literatures; it isolates one domain error that generates a common pattern across them. Perfect insurance corresponds to the special case where ruin is ruled out and the deterministic benchmark applies. Everywhere else, bounded insurance makes the selective failures structural. Agents near the boundary are not irrational; they are responding to a constraint the equality-based tools do not see.

We do not claim this mechanism is the only source of Euler residuals or incomplete insurance; we claim that when collapse is feasible, the Euler-equality characterization is valid only on the safe set, so equality-based diagnostics become unreliable precisely in the states economists care most about.

The paper makes a falsifiable bet. If Euler residuals, insurance-share gradients, and pass-through coefficients do not switch on together when the hazard proxy rises, the survival term is not the common driver, and the pattern requires a different explanation. If they do, the profession has been testing preferences when it should have been testing feasibility. For applied work, the operational rule is simple: treat the Euler equality as an interior condition, sort observations by a short-horizon hazard proxy, and expect the equality to hold in safe periods and fail in danger periods.

The survival partition requires no equilibrium structure; embedding it in one is the natural next step. But the point of this paper is prior: most of what the profession calls life-cycle consumption theory is safe-set theory. It describes agents who were never close to the edge.

Appendix A. Technical Setup and Assumptions

A.1 Environment and probability space

Let $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t=0}^T, \mathbb{P})$ be a filtered probability space. Shocks $\{Y_t\}_{t=1}^T$ are i.i.d. with common law F on a measurable space $(\mathcal{Y}, \mathcal{G})$ and are $\mathcal{F}_t$–measurable. The agent starts with wealth $k_1 \geq \kappa$, where $\kappa > 0$ is the absorbing ruin threshold. At each $t = 1, \ldots, T$, the agent chooses consumption $c_t \geq 0$ and insurance $m_t \geq 0$, both $\mathcal{F}_t$–measurable. Wealth evolves by

$$k_{t+1} \;=\; k_t - c_t - m_t + g(Y_{t+1}) + \lambda\, m_t\, \mathbf{1}\{Y_{t+1} \in \mathcal{B}\}, \tag{2}$$

where $\mathcal{B} \subset \mathcal{Y}$ is the set of "bad" shocks and $\lambda \in (0, 1]$ is insurance efficiency; $\lambda < 1$ captures costly or actuarially unfair insurance, $\lambda = 1$ is fair insurance. If $k_{t+1} < \kappa$, the process enters the absorbing ruin state thereafter. Per–period utility is $u(c)$ with $u : \mathbb{R}_+ \to \mathbb{R}$ strictly increasing, strictly concave, continuously differentiable, and $\lim_{c \downarrow 0} u'(c) = +\infty$. Lifetime utility is

$$U^\pi(k_1) \;=\; \mathbb{E}\Big[\sum_{t=1}^{T} \beta^{t-1} u(c_t) \;+\; U_{\text{ruin}} \cdot \mathbf{1}\{\exists s \leq T : k_s < \kappa\}\Big],$$

where $\beta \in (0, 1)$ and $U_{\text{ruin}} > -\infty$ unless explicitly stated otherwise.

A.2 Admissible policies and feasibility

Let $A(k)$ denote the feasible set at wealth k:

$$A(k) \;=\; \{(c, m) \in \mathbb{R}_+^2 :\; k - c - m \geq \kappa\},$$

so that the agent never chooses to cross the boundary within the period. (If one wishes to allow $k - c - m < \kappa$, the same arguments go through with the understanding that k_{t+1} is then absorbing at ruin.) A policy $\pi = \{(c_t, m_t)\}_{t=1}^T$ is *admissible* if (c_t, m_t) is $\mathcal{F}_t$–measurable and $(c_t, m_t) \in A(k_t)$ a.s. for all t.

A.3 Value function, Bellman operator, and assumptions

Define $V_{T+1}(k) \equiv 0$ for $k \geq \kappa$ and $V_{T+1}(k) \equiv U_{\text{ruin}}$ for $k < \kappa$. For $t = T, \ldots, 1$ set

$$V_t(k) \;=\; \sup_{(c,m)\in A(k)} \Big\{ u(c) \;+\; \beta\, \mathbb{E}\big[V_{t+1}(k') \mid k, (c,m)\big] \Big\}, \tag{3}$$

with k' given by (2).

We impose standard conditions used in stochastic control (e.g., Stokey et al. 1989; Bertsekas and Shreve 1978; Benveniste and Scheinkman 1979):

(A1) u is strictly concave, continuous, $u'(0+) = +\infty$, and u has at most linear growth: $|u(c)| \leq a + bc$ for some $a, b < \infty$.

(A2) $g : \mathcal{Y} \to \mathbb{R}$ is measurable with $\mathbb{E}[|g(Y_1)|] < \infty$.

(A3) $A(k)$ is nonempty, convex, and compact for each $k \geq \kappa$ (compactness holds if one imposes an upper bound $\bar{c}$ or uses Fatou/monotone class arguments with linear growth; either approach is standard).

(A4) $\beta \in (0, 1)$.

(A5) The bad–shock event $\mathcal{B}$ has strictly positive probability: $p_b \equiv \mathbb{P}(Y_1 \in \mathcal{B}) \in (0, 1)$.

Under (A1)–(A5) the following are standard.

Lemma A.1 (Existence, Concavity, and DPP). *For each t, V_t exists, is finite on $[\kappa, \infty)$, is continuous and concave, and* (3) *holds (Dynamic Programming Principle). Moreover, there exists a measurable selector $\phi_t(k) \in A(k)$ attaining the supremum.*

Proof. Standard measurable selection and monotone operator arguments (Bertsekas & Shreve); concavity follows from concavity of u and linearity of the transition. □

Lemma A.2 (Envelope/Benveniste–Scheinkman). *For each t, V_t is differentiable a.e. on (κ, ∞). At differentiability points, the supergradient equals the Lagrange multiplier on the state constraint:*

$$V_t'(k) \;=\; \mu_t^\star(k) \quad \textit{for the optimal choice at } k,$$

and the FOCs read

$$u'(c_t^\star) \;=\; \mu_t^\star, \qquad \beta\,\mathbb{E}\left[V'_{t+1}(k_{t+1}^\star)\,(-1+\lambda\,\mathbf{1}\{Y_{t+1}\in\mathcal{B}\})\,\middle|\,\mathcal{F}_t\right] \;=\; \mu_t^\star-\theta_t^\star,$$

with $\theta_t^\star\geq 0$ *the multiplier on* $k_t\geq\kappa$ *and* $\theta_t^\star\,(k_t-\kappa)=0$.

Proof. Benveniste and Scheinkman (1979) envelope theorem applied to (3) with compact $A(k)$ and linear transition; see also Stokey et al. (1989, Ch. 4). □

Remark on multiplier notation. At interior states $k_t>\kappa$, the state constraint $k_t\geq\kappa$ is automatically slack; the binding constraint is within-period feasibility $k_t-c_t-m_t\geq\kappa$, whose multiplier is denoted μ_t in Appendix D. The proofs in A3–A4 write θ_t with complementary slackness on $k_t-c_t-m_t-\kappa$ to denote this within-period multiplier; it is θ_t that generates the Euler wedge in Sections 3–4.

Lemma A.3 (Boundary Behavior). *If* $u'(0+)=+\infty$, *then* $V_t'(k)\to+\infty$ *as* $k\downarrow\kappa$ *for each* t.

Proof. As $k\downarrow\kappa$, feasibility forces c to approach 0 in any optimal policy (else k' violates the constraint with positive probability). Since $u'(c)\to\infty$ and $u'(c)=\mu_t^\star=V_t'(k)$ at interior points, the result follows. □

Appendix A3. Proofs for Section 3 (Main Theorems)

Both proofs compare the equality-characterized policy π^E (which sets the survival multiplier to zero) against the full Bellman policy $\pi^\star$ (which prices survival where it binds). The value gap between them is entirely a survival-probability gap: the consumption-profile difference is bounded; the survival-probability difference is not.

Proof of Theorem 1 (Euler Equalities Hold Only on the Safe Set). The equality-characterized policy π^E solves the Bellman program as though the survival constraint were free: it satisfies the FOCs of Lemma A.2 with $\theta_t\equiv 0$. The full Bellman policy $\pi^\star$ prices that constraint: $\theta_t^\star\geq 0$ with complementary slackness. Both are state-contingent Markov strategies. They differ in whether ruin has a price.

(i) On $\mathcal{S}$, the constraint never binds along $\pi^\star$, so $\theta_t^\star=0$ a.s. The FOCs reduce to the Euler equality. By strict concavity and uniqueness (Lemma A.1), $\pi^\star=\pi^E$ and $V_t=V_t^E$.

(ii) Off $\mathcal{S}$, there exist t and $E_t \in \mathcal{F}_t$ with $\mathbb{P}(E_t) > 0$ such that $k_t^E \in (\kappa,\, \kappa + \varepsilon)$. At these states, π^E satisfies the Euler equality and allocates as though the boundary were inert. But the boundary is not inert: crossing it ends the program. Consider the variation that shifts δ from consumption to insurance on E_t:

$$D(\eta, \delta) \;=\; -\,\eta\, u'(c_t^E) \;+\; \beta\, \mathbb{E}\Big[V'_{t+1}(k_{t+1}^E) \cdot \big(-1 + \lambda\, \mathbf{1}\{Y_{t+1} \in \mathcal{B}\}\big) \,\Big|\, \mathcal{F}_t\Big].$$

The second term is the marginal value of surviving a bad shock. By Lemma A.3 this value grows without bound near κ, so for η near zero, $D > 0$: the gain from surviving one more period dominates the cost of consuming slightly less today. The full Bellman policy makes this reallocation; the equality-characterized policy cannot, because it has already set the price of survival to zero. Hence $V_t(k_1) > V_t^E(k_1)$, and the entire gap is the survival price that the equality characterization omits. ■

Proof of Theorem 2 (Boundary of the Safe Region). The safe region is the largest set from which the agent can guarantee survival; leaving it is costly because Theorem 1 applies at the exit point. Define the *safe region*

$$\mathcal{S}_{\text{safe}} \;=\; \Big\{k_1 \geq \kappa :\; \mathbb{P}\big(k_t^{\pi^\star} \geq \kappa \;\; \forall t\big) = 1\Big\},$$

where $\pi^\star$ denotes the full Bellman policy (a measurable selector of (3)).

(*Only if.*) Suppose the equality-characterized policy coincides with the full Bellman policy at k_1. If $\mathbb{P}(k_t < \kappa) > 0$ for some t under that plan, Theorem 1 yields a strict value gain for the full Bellman policy at the first date t where the state enters a danger neighborhood, contradicting coincidence. Hence $\mathbb{P}(k_t < \kappa) = 0$ for all t, so $k_1 \in \mathcal{S}_{\text{safe}}$.

(*If.*) If $k_1 \in \mathcal{S}_{\text{safe}}$, then the state constraint never binds along any history under $\pi^\star$. By Lemma A.2 the multipliers $\theta_t^\star = 0$ and V_t' remains bounded; the FOCs pin down a unique policy each period independent of the information that would be learned under the equality-characterized rule. Therefore the equality-characterized policy coincides with $\pi^\star$ and the Euler equality is globally valid. ■

Appendix A4. Proofs for Section 4 (Corollaries)

Each proof below uses the same variation: shift ε from consumption to insurance on the set where survival binds, then show the survival gain is first-order (unbounded marginal value

near κ, via Lemma A.3) while the consumption cost is second-order (bounded by concavity). The corollaries differ only in which margin the variation operates on: Corollary 2 on self-insurance vs state-contingent coverage, Corollary 3 on the premium share, Corollary 4 on the Euler wedge, Corollary 5 on the value decomposition, and Corollaries 6 on distributional mass.

Proof of Corollary 2 (Self-Insurance Alone Cannot Prevent Ruin). Fix t and $k \notin \mathcal{S}$. The equality-characterized policy π^E sets $\theta_t \equiv 0$ and $m_t = 0$: it saves but does not insure. Since $k \notin \mathcal{S}$, ruin has positive probability under π^E, and there exists $E_t \in \mathcal{F}_t$ with $\mathbb{P}(E_t) > 0$ on which the next-period state is near κ. Now consider a small reallocation: define π^ε as the policy that, on E_t, spends ε less on consumption and ε more on insurance; off E_t it follows π^E. The value difference on E_t has two parts:

The cost of consuming less: $u(c_t^E) - u(c_t^E - \varepsilon) \leq \varepsilon \cdot u'(c_t^E - \varepsilon)$, which is finite at any interior solution. This is a bounded sacrifice in current utility.

The gain from surviving a bad shock: on $\{Y_{t+1} \in \mathcal{B}\}$, the insurance payout $\lambda\varepsilon$ pushes next-period wealth away from κ. The marginal value of that upward shift is $\beta\, p_b\, \lambda\, V'_{t+1}(k^E_{t+1,\text{bad}}) \cdot \varepsilon$. By Lemma A.3, $V'_{t+1}(k) \to +\infty$ as $k \downarrow \kappa$: each dollar of survival is more valuable the closer the agent stands to ruin.

A finite cost against an unbounded gain: for ε small enough, the reallocation strictly improves value. The full Bellman policy is the global maximizer, so $V_t(k) \geq U^{\pi^\varepsilon} > V_t^E(k)$. Self-insurance alone leaves money on the table that state-contingent insurance picks up. ■

Proof of Corollary 3 (Insurance Demand Rises Near Collapse). Near the boundary, the survival premium lowers the effective price of insurance; a lower price means higher demand. The insurance FOC from Appendix D equates the marginal cost of a dollar of premium to its marginal benefit. Write this as $L(m) = R$, where $L(m)$ is the expected marginal value of the insurance position and R is the shadow price of wealth net of any survival premium. Under the full Bellman policy, $R = \mu_t^\star - \theta_t^\star$; under π^E, $R = \mu_t^E$ (since $\theta_t^E \equiv 0$). For $k \notin \mathcal{S}$, Theorem 1 gives $\theta_t^\star > 0$ on a positive-probability set: the full Bellman policy carries a survival premium that the equality-characterized policy does not. So the full Bellman R is strictly lower.

Why does a lower R mean more insurance? Because L is strictly decreasing in m. Each additional dollar of premium shifts wealth from the good state (where there is no payout) to the bad state (where the payout is λm). By strict concavity of V_{t+1} (Lemma A.1), the marginal value of wealth rises where wealth falls and falls where wealth rises. The weights

in L, which are negative in the good state and positive in the bad, ensure both shifts reduce L. So L is a downward-sloping demand curve for insurance, and a lower R moves the crossing point to the right.

The rest follows. The equation $L(m) = R$ has a unique solution because L is strictly decreasing and continuous. A lower R requires a higher $m_t^\star$ to restore equality: the Bellman policy buys more insurance. The consumption FOC $u'(c_t^\star) = \mu_t^\star$ and the tighter budget force $c_t^\star < c_t^E$, so the premium share $m_t^\star/(c_t^\star + m_t^\star)$ strictly exceeds the equality-characterized share wherever $\theta_t^\star > 0$. ■

Proof of Corollary 4 (Smoothing Fails When Survival Is at Risk). The Euler equality acquires an extra term when survival binds; that term is the wedge. From Lemma A.2, the consumption FOC at the optimum is $u'(c_t^\star) = V_t'(k_t)$. The envelope condition gives $V_t'(k_t) = \beta\,\mathbb{E}[V_{t+1}'(k_{t+1}^\star) \mid \mathcal{F}_t] + \theta_t^\star$, where $\theta_t^\star \geq 0$ is the multiplier on $k_t - c_t - m_t \geq \kappa$ with complementary slackness $\theta_t^\star(k_t - c_t^\star - m_t^\star - \kappa) = 0$. Combining:

$$u'(c_t^\star) \;=\; \beta R\,\mathbb{E}_t[u'(c_{t+1}^\star)] \;+\; \theta_t^\star.$$

On $\operatorname{int}\mathcal{S}$, the feasibility constraint is slack along all optimal histories: $\theta_t^\star = 0$ by complementary slackness. The FOC reduces to the Euler equality $u'(c_t) = \beta R\,\mathbb{E}_t[u'(c_{t+1})]$, and by the uniqueness in Lemma A.1 the two policies coincide.

Off $\mathcal{S}$, the constraint binds on E_t with $\mathbb{P}(E_t) > 0$: $\theta_t^\star > 0$ there. The FOC gives $u'(c_t^\star) > \beta R\,\mathbb{E}_t[u'(c_{t+1}^\star)]$, a strict Euler inequality. Any rule that enforces the equality off $\mathcal{S}$ sets $\theta_t = 0$ where it should be positive, consuming more than the Bellman optimum on the survival margin. By Corollary 2 this is strictly dominated.

When the state returns to $\operatorname{int}\mathcal{S}$, the constraint becomes slack, $\theta_t^\star$ returns to zero, and the Euler equality resumes. ■

Proof of Corollary 5 (Full Bellman Policy Increases Survival Value). Where does the value advantage of the Bellman policy come from: better consumption paths or better survival odds? This proof shows it is almost entirely survival.

Let $\tau^\star = \min\{s \geq t : k_s^\star < \kappa\}$ and $\tau^E = \min\{s \geq t : k_s^E < \kappa\}$, with $\min \varnothing = \infty$. Decompose

the value gap into two terms, each with a distinct economic source:

$$V_t(k) - V_t^E(k) \;=\; \underbrace{\sum_{s=t}^{T} \beta^{s-t}\, \mathbb{E}\big[u(c_s^\star)\mathbf{1}\{s < \tau^\star\} - u(c_s^E)\mathbf{1}\{s < \tau^E\}\big]}_{\text{consumption-profile difference}} \;+\; \underbrace{U_{\text{ruin}} \cdot \big[\mathbb{P}(\tau^E \le T) - \mathbb{P}(\tau^\star \le T)\big]}_{\text{survival-probability difference}}.$$

The first term captures any difference in the quality of consumption while alive. The second captures the difference in the probability of staying alive.

Survival term. For $k \notin \mathcal{S}$, Theorem 1 gives $\mathbb{P}(\tau^E \le T) > \mathbb{P}(\tau^\star \le T)$ on a positive-probability set: the equality-characterized policy lets agents fall into ruin that the Bellman policy prevents. Since $U_{\text{ruin}} < 0$, this term is strictly positive: it is the value of the ruin events that π^E causes and $\pi^\star$ avoids.

Consumption term. By (A1), $|u(c)| \le a + bc$ with $a, b < \infty$. Both policies draw from $A(k)$ with wealth bounded above by $\bar{K} = k_1 + T \cdot \sup|g|$, so the total consumption difference across T periods is bounded by $2T(a + b\bar{K})$, which is a finite number that does not depend on U_{ruin} or on survival probabilities.

One term is strictly positive on a positive-probability set; the other is finite. Hence $V_t(k) > V_t^E(k)$. The gap is a survival gap, not a consumption gap. ■

Proof of Corollary 6 (Full Bellman Policy Sustains Survival Mass). The two halves of this result are structurally opposite: the equality-characterized rule causes a population to drift into the absorbing set until it thins out; the full Bellman policy arrests that drift and builds a stationary population that equality-characterized rules cannot sustain.

Equality-characterized rule: monotone drift into the absorbing set. Under π^E with $\theta_t \equiv 0$, each agent at $k_t \notin \mathcal{S}$ faces one-step ruin probability $p(k_t) > 0$ (Theorem 1). Ruin is absorbing: mass that crosses κ never returns. So surviving mass falls every period:

$$m_{t+1} \;=\; m_t \;-\; \int p(x)\, d\mu_t(x) \;<\; m_t$$

whenever the exposed set carries positive measure. With $p(x) \ge \underline{p} > 0$ on a persisting subset, the loss compounds: $\mu_t(E) \le (1 - \underline{p})^t \cdot \mu_0(E)$. The exposed tail thins geometrically. No fixed rule can stop this, because the rule does not price survival.

Full Bellman policy: the drift is arrested. The full Bellman policy does price survival. Three facts combine to turn one-way drift into the absorbing set into a stationary distribution.

First, local hazard falls. By Corollary 2, the Bellman policy reallocates toward insurance wherever $k_t \notin \mathcal{S}$, strictly reducing ruin probability at every exposed state.

Second, agents return to safety. By (A2), $\mathbb{E}[|g(Y)|] < \infty$, and by (A5), $\mathbb{P}(Y \notin \mathcal{B}) = 1 - p_b > 0$. A good shock of size $g(Y) > \bar{k} - k_t$ for some $\bar{k} \in \mathcal{S}$ returns the agent to $\mathcal{S}$ in one step. This event has positive probability from any bounded state, providing a uniform return rate $\delta > 0$.

Third, return plus reduced hazard produce a stationary population. The return rate $\delta > 0$ to $\mathcal{S}$ and compactness of the feasible set $A(k)$ make the occupation measures $\nu_N = (1/N)\sum_{t=0}^{N-1} \mu_t$ tight on $[\kappa, \infty)$: the probability of remaining below $\kappa + \eta$ for $\lfloor 1/\delta \rfloor$ or more consecutive periods is at most $(1-\delta)^{\lfloor 1/\delta \rfloor}$. By Prokhorov's theorem, $\{\nu_N\}$ has a convergent subsequence whose limit ν is invariant for the transition kernel under $\pi^\star$. The return property ensures $\nu(\mathcal{S}) > 0$. Agents in $\mathcal{S}$ survive with probability one, so the invariant distribution carries strictly positive survival mass.

The contrast is the result. The equality-characterized rule leaks mass geometrically with no mechanism to replenish it. The full Bellman policy reduces hazard and recycles agents back into safety, sustaining a population that fixed rules cannot. ■

Appendix A5. Proofs for Section 5 (Five Diagnostic Implications)

This appendix derives the five implications in Section 5 from Lemma 1 and the KKT system in Appendix D. Each derivation starts from the same FOC and isolates the margin-specific derivative whose sign flips with μ_t.

FOC structure (from Lemma 1). With $\mu_t = 0$ the constraint is slack and the Kuhn–Tucker condition reduces to the Euler equality. With $\mu_t > 0$, the survival term enters, driving a wedge $\mu_t \left|\frac{\partial g}{\partial c}\right|$ between $u'(c_t)$ and $\beta\,(R)\,\mathbb{E}_t[u'(c_{t+1})]$.

Euler residual. Define $e_t := u'(c_t) - \beta\, R\, \mathbb{E}_t[u'(c_{t+1})]$. From the FOC in Lemma 1,

$$e_t \;=\; \mu_t \left|\frac{\partial g}{\partial c}\right| \;\geq\; 0,$$

with equality if and only if $\mu_t = 0$. The residual is zero in safety and strictly positive in danger; it rises monotonically with μ_t.

Derivation 1 (Excess sensitivity). Let Δy_{t+1} denote a predictable income change known at t. Differentiate the FOC with respect to Δy_{t+1}:

$$\frac{\partial c_t}{\partial \Delta y_{t+1}} \;=\; -\,\frac{\mu_t\,(\partial^2 g/\partial c\,\partial y)}{u''(c_t)}.$$

When $\mu_t = 0$, the numerator vanishes and $\partial c_t/\partial \Delta y_{t+1} = 0$: orthogonality holds. When $\mu_t > 0$, the cross-partial $\partial^2 g/\partial c\,\partial y < 0$ (predictable resources relax survival), so $\partial c_t/\partial \Delta y_{t+1} > 0$: predictable income moves current consumption. The sign flip is the formal basis for the state-contingent excess sensitivity in Section 5 (Hall and Mishkin, 1982; Souleles, 1999).

Derivation 2 (BPP pass-through). Let ε be a transitory shock and η a permanent shock. The pass-through coefficients are $\partial c_t/\partial \varepsilon$ and $\partial c_t/\partial \eta$. Inside $\mathcal{S}$ ($\mu_t = 0$), the standard ranking holds: transitory shocks are smoothed more than permanent shocks. When $\mu_t > 0$, both pass-through coefficients increase because $\mu_t\,|\partial g/\partial c|$ adds a common positive term to the FOC wedge; concavity of u ensures the increase is proportional to shock size, so both types show higher pass-through near the boundary. The ranking does not change; the region does (Blundell et al., 2008; Kaplan and Violante, 2014).

Derivation 3 (Buffers before consumption). After a favorable resource shock $dw > 0$, the agent allocates between liquid balances (dl) and nondurable consumption (dc). The FOC for the liquid margin carries the survival term $\mu_t\,\partial g/\partial l$, where $\partial g/\partial l > 0$ (liquidity relaxes survival). When $\mu_t > 0$, the shadow return on liquidity exceeds the shadow return on nondurables by $\mu_t\,|\partial g/\partial l|$, so $dl/dw > dc/dw$: buffers move first. When $\mu_t = 0$, the premium vanishes and standard smoothing allocates proportionally. Buffers lead not because of patience but because liquidity buys survival and consumption does not.

Derivation 4 (Timing of payouts). Compare two contracts: immediate payout π_1 at t and delayed payout π_2 at $t{+}1$, actuarially equivalent. The difference in marginal value is $\mu_t\,(\partial g/\partial \pi_1 - \beta\,\partial g/\partial \pi_2)$. Since $\partial g/\partial \pi_1 > 0$ at binding states (immediate payout relaxes today's constraint) and the delayed payout enters only tomorrow's constraint, the difference is strictly positive when $\mu_t > 0$: the immediate-liquid contract dominates. When $\mu_t = 0$, the ∂g terms vanish and the agent is near-indifferent. The immediate contract wins because it relaxes the binding constraint at the date it binds (Cole et al., 2013, 2014; Gallagher, 2014).

Derivation 5 (State-dependent corners). Let $\bar{b}$ denote an exogenous borrowing cap. The desired borrowing level $b^\star$ satisfies $u'(c_t) = \beta\,R\,\mathbb{E}_t[u'(c_{t+1})] + \mu_t\,|\partial g/\partial c|$. When $\mu_t > 0$,

the wedge raises $u'(c_t)$ relative to the discounted expectation, increasing desired current consumption and therefore desired borrowing. The cap binds when $b^\star > \bar{b}$. Since μ_t rises as $k_t \to \kappa$, corner frequency increases near the boundary even if $\bar{b}$ and prices are unchanged. When $\mu_t = 0$, the wedge vanishes, $b^\star$ falls, and the cap is slack. The corner is endogenous: the agent wants to borrow more because μ_t raises desired consumption, not because a lender tightened terms.

Appendix D. Lagrangian/Karush–Kuhn–Tucker (KKT) Representation

For completeness, we record the per–period KKT system at interior points. Given $k_t > \kappa$ and an optimal choice $(c_t^\star, m_t^\star)$ with multipliers $(\mu_t^\star, \theta_t^\star)$, the Lagrangian is

$$\mathcal{L}_t \;=\; u(c_t) + \beta\,\mathbb{E}[V_{t+1}(k_{t+1}) \mid \mathcal{F}_t] + \mu_t(k_t - c_t - m_t - \kappa) \;+\; \theta_t(k_t - \kappa),$$

yielding

$$u'(c_t^\star) = \mu_t^\star, \qquad \beta\,\mathbb{E}[V'_{t+1}(k_{t+1}^\star)(-1 + \lambda\,\mathbf{1}\{Y_{t+1} \in \mathcal{B}\}) \mid \mathcal{F}_t] = \mu_t^\star - \theta_t^\star,$$

with $\theta_t^\star \geq 0$, $\theta_t^\star(k_t - \kappa) = 0$. Lemma A.3 implies $\mu_t^\star = V_t'(k_t) \to \infty$ as $k_t \downarrow \kappa$. The multiplier $\mu_t^\star$ is the survival wedge: it equals zero inside $\mathcal{S}$ (constraint slack, Euler equality holds) and is strictly positive outside $\mathcal{S}$ (survival binds, Euler inequality holds). The complementary-slackness multiplier $\theta_t^\star$ prices the floor directly; $\mu_t^\star$ prices the marginal value of wealth that the floor makes scarce. This is the formal engine behind the insurance and smoothing corollaries in Section 4 and the five implications in Section 5.

References

Aiyagari, S. Rao (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *Quarterly Journal of Economics*, 109(3), 659–684. https://doi.org/10.2307/2118417

Attanasio, Orazio P., and Guglielmo Weber (1995). Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey. *Journal of Political Economy*, 103(6), 1121–1157. https://doi.org/10.1086/601447

Benveniste, Lawrence M., and José A. Scheinkman (1979). On the Differentiability of the Value Function in Dynamic Models of Economics. *Econometrica*, 47(3), 727–732. https://doi.org/10.2307/1914143

Bertsekas, Dimitri P., and Steven E. Shreve (1978). *Stochastic Optimal Control: The Discrete Time Case*. New York: Academic Press. https://www.worldcat.org/title/3449844

Bewley, Truman F. (1986). Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers. In: Hildenbrand, W. and Mas-Colell, A. (Eds.), *Contributions to Mathematical Economics in Honor of Gérard Debreu*, 79–102. Amsterdam: North-Holland.

Blundell, Richard, Luigi Pistaferri, and Ian Preston (2008). Consumption Inequality and Partial Insurance. *American Economic Review*, 98(5), 1887–1921. https://doi.org/10.1257/aer.98.5.1887

Campbell, John Y., and N. Gregory Mankiw (1989). Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence. *NBER Macroeconomics Annual*, 4, 185–216. https://www.nber.org/system/files/chapters/c10978/c10978.pdf

Campbell, John Y., and N. Gregory Mankiw (1990). Permanent Income, Current Income, and Consumption. *Journal of Business & Economic Statistics*, 8(3), 265–279. https://www.jstor.org/stable/1391906

Carroll, Christopher D., and Miles S. Kimball (1996). On the Concavity of the Consumption Function. *Econometrica*, 64(4), 981–992. https://doi.org/10.2307/2171853

Chetty, Raj (2008). Moral Hazard versus Liquidity and Optimal Unemployment Insurance. *American Economic Review*, 98(2), 139–145. https://doi.org/10.1257/aer.98.2.139

Cochrane, John H. (1991). A Simple Test of Consumption Insurance. *Journal of Political Economy*, 99(5), 957–976. https://doi.org/10.1086/261784

Cole, Shawn, Xavier Giné, Jeremy Tobacman, Petia Topalova, Robert Townsend, and James Vickery, 2013. Barriers to household risk management: Evidence from India.

American Economic Journal: Applied Economics, 5(1), 104–135. https://doi.org/10.1257/app.5.1.104

Cole, Shawn, Xavier Giné, Jeremy Tobacman, Petia Topalova, Robert Townsend, and James Vickery, 2014. Short-term effects of a weather insurance experiment in India. *American Economic Review*, 104(5), 284–290. https://doi.org/10.1257/aer.104.5.284

Deaton, Angus (1991). Saving and Liquidity Constraints. *Econometrica*, 59(5), 1221–1248. https://doi.org/10.2307/2938366

Flavin, Marjorie A. (1981). The Adjustment of Consumption to Changing Expectations about Future Income. *Journal of Political Economy*, 89(5), 974–1009. https://doi.org/10.1086/261014

Gallagher, Justin, 2014. Learning about an infrequent event: Evidence from flood insurance take-up in the United States. *American Economic Journal: Applied Economics*, 6(3), 206–233. https://doi.org/10.1257/app.6.3.206

Gruber, Jonathan (1997). The Consumption Smoothing Benefits of Unemployment Insurance. *American Economic Review*, 87(1), 192–205. https://www.jstor.org/stable/2950876

Hall, Robert E., and Frederic S. Mishkin (1982). The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households. *Econometrica*, 50(2), 461–481. https://doi.org/10.2307/1912638

Huggett, Mark (1993). The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies. *Journal of Economic Dynamics and Control*, 17(5–6), 953–969. https://doi.org/10.1016/0165-1889(93)90024-M

Kaplan, Greg, and Giovanni L. Violante (2014). A Model of the Consumption Response to Fiscal Stimulus Payments. *Econometrica*, 82(4), 1199–1239. https://doi.org/10.3982/ECTA10528

Kocherlakota, Narayana R. (1996). Implications of Efficient Risk Sharing without Commitment. *Review of Economic Studies*, 63(4), 595–609. https://doi.org/10.2307/2297795

Kydland, Finn E., and Edward C. Prescott, 1977. Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85(3), 473–491. https://www.journals.uchicago.edu/doi/10.1086/260580

Ligon, Ethan, Jonathan P. Thomas, and Tim Worrall (2002). Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies. *Review of Economic Studies*, 69(1), 209–244.

https://doi.org/10.1111/1467-937X.00203

Mace, Barbara J. (1991). Full Insurance in the Presence of Aggregate Uncertainty. *Journal of Political Economy*, 99(5), 928–956. https://doi.org/10.1086/261783

Parker, Jonathan A. (1999). The Reaction of Household Consumption to Predictable Changes in Social Security Taxes. *American Economic Review*, 89(4), 959–973. https://doi.org/10.1257/aer.89.4.959

Rust, John (1987). Optimal Replacement of GMC Bus Engines. *Econometrica*, 55(5), 999–1033. https://doi.org/10.2307/1911259

Rust, John (1994). Structural Estimation of Markov Decision Processes. In: Engle, R. F. and McFadden, D. L. (Eds.), *Handbook of Econometrics*, vol. 4, ch. 51, 3081–3143. Elsevier. https://doi.org/10.1016/S1573-4412(05)80018-8

Shea, John (1995). Union Contracts and the Life-Cycle/Permanent-Income Hypothesis. *American Economic Review*, 85(1), 186–200. https://www.jstor.org/stable/2118003

Souleles, Nicholas S. (1999). The Response of Household Consumption to Income Tax Refunds. *American Economic Review*, 89(4), 947–958. https://doi.org/10.1257/aer.89.4.947

Souleles, Nicholas S. (2002). Consumer Response to the Reagan Tax Cuts. *Journal of Public Economics*, 85(1), 99–120. https://doi.org/10.1016/S0047-2727(01)00098-1

Stokey, Nancy L., Robert E. Lucas, and Edward C. Prescott (1989). *Recursive Methods in Economic Dynamics*. Cambridge, MA: Harvard University Press. https://www.hup.harvard.edu/books/9780674750967

Strotz, Robert H. (1955). Myopia and Inconsistency in Dynamic Utility Maximization. *Review of Economic Studies*, 23(3), 165–180. https://doi.org/10.2307/2295722

Townsend, Robert M. (1994). Risk and Insurance in Village India. *Econometrica*, 62(3), 539–591. https://doi.org/10.2307/2951659

Zeldes, Stephen P. (1989). Consumption and Liquidity Constraints: An Empirical Investigation. *Journal of Political Economy*, 97(2), 305–346. https://doi.org/10.1086/261606