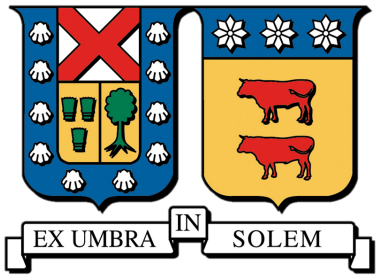
Universidad Técnica Federico Santa María
Departamento de Industrias
Valparaíso, Chile



Risk and Market Design in Individual Capitalization Pensions: Two Essays on the Protection of Savings and Competition in the Annuity Market

Lucas Martín Rissetti Cabrera

Tesis para optar al Título de Ingeniero Civil Industrial
y al Grado de Magíster en Ciencias de la Ingeniería Industrial

Profesor Guía: Dr. Rolando Rubilar
Profesor Co-referente: Dr. Cristóbal Fernández Robin
Profesor revisor externo: Dr. Karime Chahuán Jiménez

Octubre de 2026

CONSTANCIA DE VALIDACIÓN Y CONFIDENCIALIDAD DE MONOGRAFÍA A REPOSITORIO ACADÉMICO

1.- IDENTIFICACIÓN DEL TRABAJO ACADÉMICO

Tipo de monografía (marcar una opción): ☐ Memoria o trabajo de título ☒ Tesis de Postgrado

Título del trabajo: Risk and Market Design in Individual Capitalization Pensions: Two Essays on the Protection of Savings and Competition in the annuity Market

Nombre del candidato(a): Lucas Martín Rissetti Cabrera

Carrera / Grado: Magíster en ciencias de la Ingeniería Industrial

Campus: Campus Valparaíso ; **Departamento:** Industrias

2.- VALIDACIÓN DEL PROFESOR GUÍA/DIRECTOR DE TESIS

Yo, Rolando Rubilar Torrealba, en mi calidad de profesor(a) guía/director(a) del trabajo académico mencionado anteriormente **DEJO CONSTANCIA** que:

- He revisado esta versión del documento y corresponde a la versión final aprobada del trabajo.
- El trabajo cumple con los requisitos académicos y de formato establecidos por la institución.

3.- EVALUACIÓN DE CONFIDENCIALIDAD POR PROPIEDAD INDUSTRIAL (marcar una opción)

☒ El trabajo **NO contiene** información que amerite confidencialidad y puede ser publicado de inmediato en repositorio con acceso abierto.

☐ El trabajo **CONTIENE** información con potenciales implicancias de propiedad industrial o intelectual y requiere un periodo de confidencialidad (**embargo**) por **(marcar una opción)**:

☐ 6 meses ☐ 12 meses ☐ 2 años ☐ 3 años ☐ 5 años ☐ 10 años

Fundamentación de la necesidad de confidencialidad (obligatorio si se solicita embargo):

4.- FIRMAS

Profesor(a) guía o director(a) de memoria o tesis:

Fecha: 04/11/2026 ; **Firma:** _____

Estudiante o Candidato(a):

Fecha: 04/11/2026 ; **Firma:** _____

Este formulario debe ser insertado como página 2 de la memoria o tesis, completado y firmado por estudiante y profesor(a) antes de la entrega en portal PRISMA de Biblioteca USM.

Acknowledgements

First, I would like to express my deep gratitude to Jose, Maca and all the family for the love and unconditional support that made every stage of this journey possible.

To Belén, for her companionship, love, affection, and support throughout this entire process, and for being a pillar of strength.

To those who walked alongside me during these years, for the conversations, encouragement, and generosity with which they shared their time and knowledge.

To my advisor, Rolando Rubilar, for his support, guidance, and confidence in me throughout this project, and to the professors who supported me and contributed to my education.

To the Federico Santa María Technical University, for the education I received and for the scholarship that made this process possible.

And to those who are no longer here but whom I carry with me: thank you for what you taught me and for remaining present in a way that time cannot erase.

To everyone, my deepest gratitude. Nothing within these pages would exist without you.

Abstract

The shift from defined-benefit schemes toward individual capitalization systems transfers a substantial share of financial risk to workers and makes pension adequacy depend on how markets operate at two moments of the savings life cycle: the protection of the balance during accumulation and its conversion into income during the payout stage. This thesis studies both phases within a common framework —that of a regulator seeking to improve retirees' welfare subject to market, fiscal, and solvency constraints— through two complementary investigations calibrated to and tested against the Chilean pension system.

The first investigation develops a welfare-neutral decision-support framework for a threshold-activated pension-loss insurance layer against severe and rare market shocks. Using a Monte Carlo simulation, it shows that age-contingent coverage protects adequacy only within an admissible region that jointly satisfies welfare non-deterioration, expected budget balance, stochastic solvency of the fund, and a tolerable moral-hazard ceiling, and that pricing must be disciplined so as not to induce higher portfolio risk-taking. In the baseline design, the contingent liability of the system translates into a prudential reserve of about 1.5 entry wages per worker.

The second investigation models competition among insurers in the annuity market as a centralized first-price auction with heterogeneous firms, incomplete information, and endogenous participation, embedded in a planner's problem that weighs pensions against fiscal revenue. The model shows that annuity offers are equilibrium objects and that policy levers can have non-monotone effects: higher taxation may lower pensions through firm exit, even when fiscal revenue rises. Empirical evidence based on Chilean administrative data is consistent with this mechanism: conditional on observables, auctions expected to attract more insurers are associated with higher winning offers.

Taken together, the thesis treats pension adequacy as a market-design problem spanning the entire life cycle of savings —rather than a property settled solely in the accumulation phase— and shows that in both phases policy interventions exhibit non-obvious effects that partial-equilibrium intuition tends to overlook.

Keywords: individual capitalization; pension adequacy; pension-loss insurance; market shocks;

annuities; strategic competition; auctions; market design; welfare; Chile.

Resumen

El paso desde esquemas de beneficio definido hacia sistemas de capitalización individual traslada al trabajador una parte sustantiva del riesgo financiero y hace que la suficiencia de su pensión dependa de cómo funcionan los mercados en dos momentos del ciclo de vida del ahorro: la protección del saldo durante la acumulación y su conversión en renta durante la etapa de pago. Esta memoria estudia ambas fases dentro de un marco común —el de un regulador que busca mejorar el bienestar de los pensionados sujeto a restricciones de mercado, fiscales y de solvencia— a partir de dos investigaciones complementarias calibradas y contrastadas con el sistema previsional chileno.

La primera investigación desarrolla un marco de apoyo a la decisión para un seguro de pérdidas previsionales activado por umbral frente a shocks de mercado severos y poco frecuentes. Mediante simulación de Monte Carlo se muestra que una cobertura contingente a la edad protege la suficiencia solo dentro de una región admisible que satisface de manera conjunta la no reducción del bienestar, el equilibrio presupuestario esperado, la solvencia estocástica del fondo y un límite tolerable al riesgo moral, y que la tarificación debe disciplinarse para no inducir una mayor toma de riesgo. En el diseño de referencia, la obligación contingente del sistema se traduce en una reserva prudencial cercana a 1,5 salarios de entrada por trabajador.

La segunda investigación modela la competencia entre aseguradoras en el mercado de rentas vitalicias como una subasta centralizada a primer precio con firmas heterogéneas, información incompleta y participación endógena, integrada en el problema de un planificador que pondera las pensiones frente a la recaudación fiscal. El modelo muestra que las ofertas de renta vitalicia son objetos de equilibrio y que las palancas de política pueden tener efectos no monótonos: una mayor tributación puede reducir las pensiones a través de la salida de firmas, incluso cuando la recaudación aumenta. La evidencia empírica, basada en datos administrativos del mercado chileno, es consistente con este mecanismo: condicional en las variables observables, las subastas que se espera atraigan a más aseguradoras se asocian con ofertas ganadoras más altas.

Considerada en conjunto, la memoria trata la suficiencia previsional como un problema de diseño de mercado que abarca todo el ciclo de vida del ahorro —y no como una propiedad que se resuelve únicamente en la fase de acumulación—, mostrando que en ambas fases las intervenciones de política exhiben efectos no evidentes que la intuición de equilibrio parcial tiende a pasar por alto.

Palabras clave: capitalización individual; suficiencia de pensiones; seguro de pérdidas previsionales; shocks de mercado; rentas vitalicias; competencia estratégica; subastas; diseño de mercado; bienestar; Chile.

Contents

Nomenclature **xvi**

1 Introduction **1**

2 Objectives and Scope **4**

2.1 Research question 4
2.2 General objective 4
2.3 Specific objectives 5
2.3.1 Pension-loss insurance against market shocks (accumulation phase) 5
2.3.2 Strategic competition in the annuity market (payout phase) 5
2.4 Scope 6
2.5 Document Organization 6

3 Protecting Pension Adequacy from Rare Market Shocks: A Welfare-Neutral Decision-Support Framework for Individual Capitalization Systems **7**

3.1 Introduction 7
3.2 Literature review and decision-support positioning 9
3.3 Analytical framework 14
3.3.1 The life-cycle worker problem 15
3.3.2 Pricing rules and planning restrictions 17
3.3.3 Two central mechanisms 21
3.4 Numerical design and policy diagnostics 24
3.4.1 Simulation environment 24
3.4.2 Empirical calibration to the Chilean multi-fund system 26
3.4.3 Baseline parameters and policy scenarios 29
3.4.4 Computation of portfolio-risk choices 29
3.5 Results 34
3.5.1 Decision-support outputs 35

3.5.2 Recovery capacity and tail-risk protection . . . 36
3.5.3 Pricing-rule incentives and moral hazard . . . 37
3.5.4 Admissible region, welfare frontier, and stochastic solvency . . . 41
3.5.5 Sensitivity analysis and robustness . . . 46
3.6 Discussion and policy implications . . . 49
3.7 Conclusions, limitations, and future research . . . 51
3.8 Proofs of propositions . . . 53
3.9 Additional derivations and equilibrium conditions . . . 56
3.9.1 Worker's Bellman equation . . . 56
3.9.2 Euler equation for liquid saving . . . 56
3.9.3 Envelope condition for pension wealth and risk choice . . . 56
3.9.4 Actuarial pricing, pooling, and hybrid premia . . . 57
3.9.5 Budget balance and mature-regime welfare . . . 58
3.9.6 Tail-risk and moral-hazard diagnostics . . . 58

4 Strategic Competition in Retirement-Income Markets: A Bayes–Nash Model of Annuity Auctions with Evidence from Chile 59

4.1 Introduction . . . 59
4.2 Economic Model . . . 62
4.2.1 Insurance firms . . . 62
4.2.2 Central planner's optimization problem . . . 63
4.2.3 Nash equilibrium among insurers . . . 64
4.3 Empirical strategy . . . 72
4.3.1 First-stage model of expected participation . . . 74
4.3.2 Second-stage model of winning offers . . . 75
4.4 Discussion . . . 78
4.5 Conclusions . . . 80

5 Conclusions and Future Work 81

Bibliography 84

List of Tables

3.1 Empirical calibration to the Chilean multi-fund system. Annualized moments based on monthly real returns deflated by UF, October 2002 to April 2026. Tail probabilities are reported in two forms. The columns $\widehat{\Pr}$ report raw empirical frequencies based on calendar-year cumulative real returns 2003–2025 (23 complete years). The columns $\widetilde{\Pr}$ report smooth estimates obtained by fitting a Student-t distribution to monthly returns of each fund (with the regularization $\nu \geq 4$ as a discipline that prevents extremely heavy-tail overfitting on a small number of extreme observations, binding for Funds D and E and non-binding for Funds A, B, C), then computing the annualized tail probability by Monte Carlo simulation of 12-month cumulative returns under the fitted distribution. Fund A is the riskiest category, Fund E is the most conservative. 26

3.2 Baseline numerical design. Return-process parameters are calibrated to the Chilean Fund C category (Table 3.1). Behavioral and policy parameters follow standard values in the life-cycle and decision-support literatures. 30

3.3 Mapping from equations to figures and design concerns 33

3.4 Summary of baseline policy diagnostics under the Chilean calibration 34

3.5 Hybrid pricing diagnostics at the baseline preferred design $(\kappa^F, \psi^F) = (0.082, 0.016)$ for $\lambda \in \{0, 0.25, 0.50, 0.75, 1\}$. Equilibrium risk $z^\star$ derived from the worker's marginal condition under hybrid pricing weight λ. P10 and mean lifts are gross of premium payments and reported against the no-insurance benchmark. $\Delta\mathcal{W}^{\text{net}}$ is the net welfare change as a fraction of benchmark consumption-equivalent welfare. $\mathcal{V}_{0.95}$ is the prudential reserve at the population-average risk choice induced by λ. The solidarity dispersion $\mathcal{S}^{H/L} = \psi^H(\text{Fund A})/\psi^H(\text{Fund E})$ reports the ratio of the premium paid by the riskiest fund category to the premium paid by the safest fund category; a value of one corresponds to perfect cross-fund solidarity, large values indicate strong risk-based differentiation. 40

3.6 Numerical one-at-a-time sensitivity and combined adverse stress around the baseline parameterization of Table 3.2. The P10 (abs) column reports the absolute tenth-percentile of terminal pension wealth under insurance, in units of the entry wage, to capture the level effect; the P10 lift and mean lift columns report relative changes against the no-insurance benchmark under the same calibration; $\Delta\mathcal{W}^{\text{net}}$ reports the net welfare change in consumption-equivalent units. The pricing-discipline gap $\mathcal{M}^F$ is reported in z-grid units. The qualitative ordering $z^P \geq z^F \geq z^0$ over the mature ages is preserved in every scenario by Proposition 2. . 47

4.1 Descriptive statistics 74

4.2 First-Stage Zero-Truncated Poisson and Second-Stage OLS Estimates 77

List of Figures

3.1 Empirical calibration to the Chilean multi-fund system. Panel (a) reports annual real returns by fund category 2003–2025 with three activation thresholds. Panel (b) reports the empirical complementary cumulative distribution function of activation, $\Pr(R < -\kappa)$, computed as the share of calendar years in which the realized annual real return fell below the threshold. Raw frequencies are weakly increasing in the fund category index, with ties within risk classes due to the limited 23-year sample; smoothed Student-t estimates (reported in Table 3.1) are strictly increasing across the policy-relevant range $\kappa \in [0.05, 0.15]$. The empirical evidence validates Assumption 1 (non-negative expected net growth in non-crash states) and, combined with the coverage rule, the primitive condition of Lemma 1. Source: Superintendencia de Pensiones de Chile, real returns deflated by UF. 28

3.2 Shock timing and terminal pension loss under the baseline parameterization. The dashed line reports the no-shock counterfactual; solid lines report pension-wealth paths after a single activated shock at ages 35 (early career), 50 (mid-career), and 60 (close to retirement). Wealth is expressed in units of the entry wage. The terminal pension loss is largest when the shock occurs close to retirement because workers have less time to rebuild pension wealth through future contributions and compound returns. This is the recovery-capacity rationale formalized in Proposition 1. 36

3.3 Pension-wealth accumulation with and without insurance under the baseline parameterization, holding pension-portfolio risk fixed at Fund C in both regimes. Panel (a) reports the mean and the twenty-fifth to seventy-fifth percentile band. Panel (b) reports the tenth, fifth, and first percentiles. Wealth is expressed in units of the entry wage and reported gross of premium payments. This is a mechanical-coverage diagnostic that isolates the protective effect of the insurance layer at a given portfolio risk; the gross mean lift is approximately +4.4% at age 65 and the gross tenth-percentile lift is approximately +13.7% at age 65. The instrument compresses the lower tail more than it raises the mean, supporting its interpretation as an adequacy-protection layer. The endogenous-portfolio counterpart of these magnitudes, in which the worker re-optimizes pension-portfolio risk in response to the pricing regime, is reported in Section 3.5.3 and Table 3.5. . . 38

3.4 Optimal pension-portfolio risk profile by age under three insurance-pricing rules in the baseline parameterization. The dashed line reports the no-insurance benchmark. The blue line reports actuarially fair pricing. The red line reports pure social pooling. The shaded distance between pooling and no insurance is the gross insured-risk gap; the distance between pooling and fair pricing is the pricing-discipline gap. The qualitative ordering of the three curves is robust across the sensitivity analysis in Section 3.5.5, in line with Proposition 2. 39

3.5 Welfare-neutral policy frontier over the activation threshold κ and the uniform exposure-based premium ψ, computed from equation (3.38) and net of the implementation-cost wedge $\varphi_A \mathcal{C}_A$. The color scale reports the net welfare change relative to the no-insurance benchmark. The thick black contour is the welfare-neutral locus $\Delta\mathcal{W}^{net} = 0$, and the welfare-positive region is the blue area inside it. The full admissible region defined in Section 3.3.2 is the intersection of the welfare-positive region with expected budget balance, stochastic solvency, and the moral-hazard ceiling $\overline{\mathcal{M}}^F$, and it is therefore weakly smaller than the blue area shown here. The bright green curve is the budget-balancing fair premium schedule $\psi^{fair}(\kappa)$; the segment of this curve that lies inside the welfare-positive region defines a one-dimensional family of designs that are simultaneously welfare-improving and self-financing in expectation. The green circle marks the baseline preferred design $(\kappa^F, \psi^F) = (0.082, 0.016)$, which is the welfare-maximizing point among all designs in this self-financing family. The white square marks the no-insurance benchmark. The framework explicitly rejects the unconstrained welfare maximum, which lies on the lower-left boundary at $\psi \approx 0$ and a small κ: that point is a degenerate boundary solution that delivers no premium income against a frequently-activated trigger and is therefore not defensible as a policy package. 42

3.6 Stochastic solvency diagnostics of the insurance fund. Panel (a) reports the marginal value of the prudential reserve at the baseline preferred activation threshold $\kappa^F = 0.082$, by comparing the insolvency probability at three initial-buffer levels: $R_0 = 0$ (build-from-zero), $R_0 = 0.5\,\mathcal{V}_{0.95}$ (half-buffer), and $R_0 = \mathcal{V}_{0.95}$ (stationary regime). The dotted line is the $\varepsilon = 0.05$ solvency tolerance. Panel (b) reports the prudential reserve $\mathcal{V}_{0.95}(\Pi)$ defined in (3.22) over the policy grid, in units of the entry wage per worker, computed under the build-from-zero regime that quantifies the capital required to start operating the scheme. The star in Panel (b) marks the baseline preferred design $(\kappa^F, \psi^F) = (0.082, 0.016)$, at which $\mathcal{V}_{0.95} \approx 1.49$ entry wages per worker. Annual return moments are calibrated to the Chilean Fund C category, with conditional excess severity calibrated to historical activated states. The figure shows that expected budget balance does not by itself imply solvency, and that the prudential reserve provides the institutional translation of the solvency constraint into a capital requirement with substantial marginal value at intermediate buffer levels. 45

4.1 Best response from the firm when assuming a parameter structure from other firms. 68

4.2 Best responses under alternative assumed parameter structures. 68

4.3 Evolution of the number of firms, profits, proposed offer, and central planner utility. 69

4.4 Market outcomes across shock factors and corporate tax rates. 71

4.5 Annual evolution of market participation, winning pension values, and BTU rates. . 73

Nomenclature

This section lists the main symbols and acronyms used in the two studies that make up this thesis. The notation is presented separately for each study, given that both were conducted independently and retain their own formulations; notably, the same symbol may have different meanings in each. Within each study, the symbols are grouped in a single table, while the acronyms are listed separately.

Pension-loss insurance against market shocks

Symbols

Symbol	Meaning
a	Worker's age (years)
z	Fund risk category (multi-fund); portfolio-risk variable
L_{a+1}	Eligible loss, or pension-equivalent loss
I_{a+1}	Indemnity paid by the insurance
W_T	Terminal pension wealth
p_T	Terminal pension income
$\bar{p}$	Pension adequacy threshold
R_{a+1}^{obs}	Observed fund return in the period
r_f	Risk-free rate
$MC_a,\ MB_a$	Private marginal cost and benefit of risk-taking
$z^*,\ z^P,\ z^F$	Optimal risk choice under hybrid, pooled, and fair pricing
Q_q^{Π}	q-th quantile of terminal pension wealth under policy Π
R_{s_1,s_2}	Expected cumulative contribution flow between ages
$\mathcal{W}(\Pi)$	Aggregate mature-regime welfare under policy Π
$\mathcal{R},\ \mathcal{P},\ \mathcal{B}$	Premium income, expected indemnities, and fund financing ratio
$\mathcal{V}_{1-\varepsilon}$	Prudential reserve (at the solvency tolerance)
$\alpha(a, z)$	Coverage schedule: fraction of the loss reimbursed, contingent on age and risk
β	Discount factor / time preference
γ	Coefficient of risk aversion
$\Gamma(z)$	Disutility, or private adjustment cost, of the risk choice
ε	Fund solvency tolerance
κ	Activation threshold (insurance trigger return)
λ	Weight between actuarial and pooled pricing (hybrid pricing)
$\Pi,\ \Pi^0$	Insurance policy package; no-insurance benchmark
$\Phi_t(\Pi)$	Cumulative net financial flow of the insurance fund
ψ	Payroll premium; ψ^F fair premium, ψ^P pooled premium
ω_a	Population weight of age a

Acronyms

Acronym	Meaning
VaR	Value at Risk
CRRA	Constant Relative Risk Aversion
MC	Monte Carlo
GFC	Global Financial Crisis
DC	Defined contribution / individual capitalization

Strategic competition in the annuity market

Symbols

Symbol	Meaning
$\mathcal{M}$	Set of competing contracts; m denotes a contract or policy
$m_{m,t}$	Offered annuity schedule (payout profile)
$B_{m,t}$	Technical reserve of the policy
F_m	Single premium paid by the retiree
$f_m,\ f$	Fixed administration cost of the contract; per-period operating cost
r	Net return on the asset backing the reserve
d	Financial discount factor
$S_{m,t}$	Survival probability (policy in force at the end of t)
D	Minimum profit required to remain in the market
$w_m(\cdot)$	Probability of winning the contract, $\in [0,1]$
$V(m \mid m_m)$	Expected present value of profit per policy
$G_{j,m}(\cdot)$	Cumulative distribution function of rival offers
$A_m(\cdot)$	Retiree's acceptance probability, $\in [0,1]$
$b_{i,m}$	Strategy of insurer i (offer-generating map)
$\mathcal{B}_m$	Space of admissible parameters
$n,\ n^*(\tau)$	Number of firms in the market; entry equilibrium
$N,\ N_a$	Number of insurers; number bidding in auction a
$\overline{m}_m(n)$	Expected winning pension, conditional on the number of firms
G	Provision of public goods
$W_m,\ U_G$	Retiree's utility over the annuity; utility over public goods
$\mathcal{R}(\tau; n)$	Net government revenue
$\mathcal{W}(\tau, G)$	Central planner's welfare
P_a^{win}	Winning pension (offer) in auction a
Age_a	Age at the auction date
$\delta_{m(a)}$	Year and entry-month fixed effects
$\theta_{i,m}$	Private type of insurer i; Θ_m type space
λ_a	Latent Poisson intensity (expected participation in the auction)
$\hat{\mu}_a$	Estimated expected number of participants (first stage)
$\Pi_{i,\text{Tot}}(n)$	Total profit of firm i
τ	Proportional tax on portfolio profits
$\varphi_{-i,m}$	Firm i's beliefs about its rivals' types

Acronyms

Acronym	Meaning
SCOMP	Centralized annuity quotation and offer system (*Sistema de Consultas y Ofertas de Montos de Pensión*)
UF	*Unidad de Fomento* (Chilean inflation-indexed unit of account)
BTU	Five-year UF-denominated bond of the Central Bank of Chile
OLS	Ordinary Least Squares
FE	Fixed effects
CDF	Cumulative distribution function
EPV	Expected present value

Chapter 1

Introduction

Population ageing and the transition from defined-benefit schemes toward individual capitalization systems have profoundly reshaped how societies finance old age. In an individual capitalization system, the pension a worker ultimately receives is not a fixed promise from the State, but the outcome of a chain of decisions and market realizations that accumulate over the entire working life. This design strengthens the link between contribution effort and pension wealth, but it transfers a substantial share of financial risk to the individual. As a result, pension adequacy ceases to be a purely actuarial or administrative problem and comes to depend on how well the markets through which retirement savings are protected, grown, and finally converted into income actually work.

This dependence on markets manifests itself at two distinct moments of the life cycle of retirement savings. During the accumulation phase, the worker's balance is exposed to financial-market risk and, in particular, to severe and infrequent episodes of asset-price declines. The same financial shock has very different consequences depending on the age at which it occurs: a young worker can absorb it through future contributions and the compound recovery of returns, whereas a worker close to retirement has a limited recovery margin and a larger balance exposed to loss. During the decumulation or payout phase, the accumulated balance must be converted into a lifelong income stream, typically through a market in which insurance companies compete to offer annuities. At this stage the pension depends on how those companies compete, how many participate, and the regulatory rules under which offers are compared and awarded.

Much of the pension debate concentrates on the accumulation phase and, within it, on the contribution rate, treating protection against extreme shocks and the payout stage as secondary or merely technical matters. Yet at both moments small differences translate into large and largely irreversible differences in the retiree's lifetime income. More importantly, in neither phase is the outcome a given: both the protection of the balance against severe downturns and the amount of the awarded annuity are equilibrium objects, shaped by regulation and by the strategic behavior of agents. Recognizing this implies that pension policy cannot be evaluated by looking at a single

stage in isolation, nor by assuming that interventions have simple, monotone effects.

The Chilean pension system offers a particularly suitable setting to study this problem. Chile was the first economy to adopt an individual capitalization system on a large scale and is probably the most documented case in the literature. In the accumulation phase, the system operates a multi-fund structure with different risk profiles and has experienced recurrent episodes of stress —the global financial crisis and the COVID-19 pandemic, among others— that made visible the consequences of market downturns for pension adequacy. In the payout phase, the country has a centralized electronic platform, the *Sistema de Consultas y Ofertas de Montos de Pensión* (SCOMP), through which retirees receive and compare competing annuity offers. Chile is useful not because it is the only relevant case, but because it offers a transparent institutional environment with rich administrative data, in which it is possible to study clearly how regulation and competition determine pension outcomes.

This thesis studies both phases of the life cycle of retirement savings within a common framework, in which a social planner or regulator seeks to improve retirees' welfare subject to market, fiscal, and solvency constraints. The work is organized around two complementary studies, each addressing one of the two phases. The general objective is to analyze how regulatory design and market interaction determine pension adequacy in an individual capitalization system, both in the protection of savings during accumulation and in their conversion into income during the payout stage. Two specific objectives follow from this general objective. The first is to develop a regulatory decision-support framework to evaluate a threshold-activated pension-loss insurance against infrequent market shocks, characterizing the admissible combinations of coverage, pricing, and activation that protect adequacy without deteriorating welfare while respecting budget balance, fund solvency, and incentive compatibility. The second is to develop a model of competition among insurers in the annuity market, formulated as a centralized first-price auction with heterogeneous firms, incomplete information, and endogenous participation, and to embed it in a planner's problem in order to confront it with administrative data from the Chilean market.

Taken together, the contribution of this thesis is to treat pension adequacy as a market-design problem that spans the entire life cycle of savings, rather than a property settled exclusively in the accumulation phase. Both studies share the same perspective —that of the regulator—, the same normative criterion —welfare, with explicit fiscal and solvency constraints—, the same institutional environment —the Chilean system— and the same methodological stance: a micro-founded theoretical model, disciplined through reduced-form quantitative analysis and calibrated to, or confronted with, Chilean data. On that common basis, the work shows that in both phases the policy levers have non-obvious effects: in the accumulation phase, the desirable protection exists only within an admissible region and depends on pricing disciplining incentives; in the payout phase, interventions such as taxation can have non-monotone effects on pensions because firms'

participation is endogenous.

It is important to delimit the scope of the thesis. Both analyses are reduced-form: the first is a Monte Carlo simulation calibrated to the Chilean multi-fund system, and the second combines a theoretical model with reduced-form, associational empirical evidence. Consequently, the work identifies economic mechanisms and admissible design regions, and provides evidence consistent with those mechanisms, but it does not aim to recover point-identified causal effects nor to characterize an optimal policy within a complete structural model. The results should be interpreted under the maintained assumptions of each study, which are stated explicitly in the corresponding chapters.

Chapter 2

Objectives and Scope

This thesis is organized around two complementary studies that address, respectively, the accumulation phase and the payout phase of retirement savings in an individual capitalization system. This chapter states the research question that links them, sets out the general objective of the thesis and the specific objectives derived from each study, and delimits the scope of the work. The objectives restate and formalize, in explicit terms, the aims that were presented narratively in the introduction.

2.1 Research question

The thesis is organized around a central question: how do regulatory design and market interaction determine pension adequacy over the life cycle of retirement savings? This question breaks down into two specific questions, one for each phase of the cycle:

- **Accumulation phase.** Can a targeted insurance layer protect pension adequacy against severe and infrequent market shocks, without requiring a full public guarantee, and under what combinations of coverage, pricing, and activation is it admissible in terms of welfare, budget balance, fund solvency, and incentives?
- **Payout phase.** How does strategic competition among insurers determine the annuities awarded in a centralized market, and what effects do policy levers —taxation in particular— have on offered pensions, firm participation, fiscal revenue, and welfare?

2.2 General objective

To analyze how regulatory design and market interaction determine pension adequacy in an individual capitalization system, both in the protection of savings during accumulation and in their

conversion into income during the payout stage, adopting the perspective of a planner or regulator who seeks to improve retirees' welfare subject to market, fiscal, and solvency constraints.

2.3 Specific objectives

Two groups of specific objectives follow from the general objective, corresponding to the two studies that make up the thesis.

2.3.1 Pension-loss insurance against market shocks (accumulation phase)

1. To formalize the transmission of a severe and infrequent market shock into a pension-equivalent loss, and to characterize how the expected terminal damage depends on the age at which the shock occurs, establishing the rationale —a recovery capacity that declines with age— for age-contingent coverage.

2. To formulate the problem of a planner who chooses a threshold-activated insurance package —contribution rate, age- and risk-contingent coverage schedule, payroll-based premium rule, activation threshold, and weight between actuarial and pooled pricing— and to characterize the admissible region that jointly satisfies welfare non-deterioration, expected budget balance, stochastic solvency of the fund, and a tolerable moral-hazard limit.

3. To establish the effect of the pricing rule on portfolio risk-taking, determining under what condition pooled pricing induces greater risk-taking than actuarially fair pricing, in order to discipline the design of the premium.

4. To quantify, through a Monte Carlo simulation calibrated to the Chilean multi-fund system, the decision-support instruments for the regulator: the admissible region, the welfare frontier, the budget-balancing fair-premium locus, and the prudential reserve associated with the baseline design.

2.3.2 Strategic competition in the annuity market (payout phase)

1. To develop an equilibrium model of competition among insurers in a centralized annuity market, formulated as a first-price auction with heterogeneous firms, incomplete information, and endogenous participation, in which annuity offers are equilibrium objects (Bayes–Nash).

2. To embed this model in a planner's problem in order to analyze how taxation and market conditions jointly affect offered pensions, firm entry, fiscal revenue, and welfare, identifying the possibility of non-monotone effects arising from endogenous participation.

3. To confront the model's mechanism with administrative data from the Chilean annuity market, constructing an auction-level measure of expected competition and estimating its association with winning offers.

2.4 Scope

Both studies are reduced-form, and their results should be interpreted under the assumptions stated explicitly in each chapter. The first is a Monte Carlo simulation calibrated to the Chilean multi-fund system, so that its magnitudes —including the prudential reserve— depend on that calibration, and their external validity to other systems should be assessed with caution. The second combines a theoretical model with reduced-form, associational empirical evidence, which does not seek to recover point-identified causal effects. Consequently, the thesis identifies economic mechanisms and admissible design regions, and provides evidence consistent with those mechanisms, but it does not aim to characterize an optimal policy within a complete structural model nor to estimate causal relationships. Finally, the two phases are studied in parallel within a common conceptual framework, and not within a single equilibrium model that formally links them; that integration is proposed as the main line of future work.

2.5 Document Organization

The remainder of this thesis is organized as follows. This chapter has stated the research question, the general and specific objectives, and the scope of the work. Chapter 3 addresses the accumulation phase and presents the framework for pension-loss insurance against infrequent market shocks. Chapter 4 addresses the payout phase and presents the model of competition among insurers in the annuity market, together with the empirical evidence based on Chilean administrative data. Chapter 5 brings both studies together in a general conclusion that synthesizes their findings, discusses their joint implications for pension-policy design, and outlines the main lines of future research.

Chapter 3

Protecting Pension Adequacy from Rare Market Shocks: A Welfare-Neutral Decision-Support Framework for Individual Capitalization Systems

3.1 Introduction

Individual capitalization pension systems link retirement income to the accumulation of contributions and financial returns over the working life. This architecture strengthens the connection between individual contributions and pension wealth, but it also transfers a substantial share of capital-market risk to workers. The policy concern becomes particularly acute when severe market losses occur close to retirement. A loss that a young worker can partially absorb through future contributions and compound returns may become a pension-adequacy failure for an older worker with limited remaining recovery capacity. The relevant public-sector question is therefore not whether financial risk exists in funded accounts, but whether a targeted insurance layer can protect adequacy without requiring a full public pension guarantee, while making the contingent liability of the system explicit through a prudential reserve requirement.

We introduce the central object of the paper. A *pension-equivalent loss* is the gap between the pension wealth that a worker would have accumulated absent a severe market downturn and the pension wealth actually accumulated once such a downturn affects the realized return path. The same financial loss can map into very different pension-equivalent losses depending on accumulated balance, remaining contributions, and the time available before retirement. We refer to the capacity of a worker to absorb such an event through future contributions and compound returns as *recovery*

capacity. Recovery capacity declines mechanically with age and is therefore the central determinant of pension vulnerability that any insurance design must address.

Funded individual-account systems have been adopted in several countries since the early 1980s, with Chile providing the first large-scale and most studied implementation (Iglesias-Palau, 2009; Berstein et al., 2006). Across these systems, recurrent episodes of market stress have made visible the channel that motivates our analysis. The global financial crisis showed that private and funded retirement-income arrangements are not insulated from abrupt asset-price declines, with the largest welfare consequences concentrated among workers close to retirement (Whitehouse, 2009; Tower and Impavido, 2009). The COVID-19 shock combined market volatility with liquidity pressure, employment losses, and exceptional access to pension savings in several jurisdictions (OECD, 2020; Lorca, 2021; Madeira, 2022). Private adjustment margins such as additional saving, portfolio recovery, or delayed retirement may be limited and uneven for cohorts close to the retirement margin (Crawford, 2013). These episodes are not anomalies. They illustrate a structural property of capitalization arrangements that the rest of the paper analyzes.

We formulate the analysis as a welfare-neutral decision-support framework for pension-loss insurance in individual capitalization systems exposed to rare market shocks. The instrument studied is a threshold-activated insurance layer. When the realized pension-fund return crosses a contractual downside threshold, the policy reimburses a fraction of the eligible pension loss. The contract is financed through payroll-based premia and can be priced under three rules. Actuarially fair pricing links expected indemnities to the worker's own age, accumulated pension exposure, and portfolio-risk choice. Uniform social pooling spreads costs across contributors, but weakens the marginal cost of choosing riskier pension portfolios. Hybrid pricing combines both principles and lets the planner evaluate the trade-off between solidarity and incentive discipline.

The framework is organized around a single decision problem. A social planner selects an insurance package defined by five elements: the mandatory contribution rate, an age-and-risk-contingent coverage schedule, a payroll-based premium rule, the activation threshold at which the insurance is triggered, and a weight that determines the degree of actuarial versus pooled pricing. The package is admissible only if it satisfies four restrictions jointly: welfare non-deterioration relative to a no-insurance benchmark, expected budget balance, stochastic solvency of the insurance fund under a regulator-defined tolerance, and a maximum tolerable moral-hazard gap. We use the labels welfare-neutral and welfare non-deterioration interchangeably; the formal restriction is an inequality, not an equality. The risk variable in our framework is best read as the regulator-classified fund category in which the worker is enrolled, not as a continuous individual measurement, and the planner compares policies in a mature-regime cross-section so that the analysis isolates the steady-state policy effect from transitional dynamics.

No existing framework integrates these four restrictions with endogenous portfolio risk and

threshold-activated indemnities into a single decision-support problem. Life-cycle portfolio models take the institutional environment as given; rare-disaster models study asset prices rather than pension-equivalent losses across ages; intergenerational risk-sharing models examine broad smoothing rules rather than activation triggers; standard insurance theory studies moral hazard with preventive effort rather than endogenous portfolio choice; ruin theory supplies the actuarial machinery for fund-level solvency but does not interact with welfare or behavioral margins of the insured population; pension-guarantee and target-date-fund literatures address related but distinct instruments. Section 3.2 develops the comparison and clarifies how the present framework draws from each of these strands while integrating them into a single decision-support problem.

The contribution is threefold. First, we formalize the mapping from rare capital-market shocks to pension-equivalent losses, and Proposition 1 shows that the expected terminal damage is non-decreasing in the age at which the shock occurs under standard expected-accumulation conditions, which is the recovery-capacity rationale for age-contingent coverage. Second, Lemma 1 states the primitive condition that the insured expected loss is strictly increasing in pension-portfolio risk, and Proposition 2 shows that under this condition uniform pooling induces strictly higher equilibrium risk than actuarially fair pricing, which disciplines the choice of premium rule. Third, a reduced-form Monte Carlo simulation calibrated to the Chilean multi-fund system delivers four regulator-facing decision-support outputs (admissible region, welfare-neutral frontier, budget-balancing fair premium locus, prudential reserve) under a sensitivity analysis that includes a combined adverse stress scenario.

The policy message follows directly. Pension-loss insurance should be designed around the pension consequences of a market event, the ability of different cohorts to recover, and the institutional capacity of the fund to sustain the implied claim process. Coverage should increase with age, must be disciplined either by risk-adjusted premia or by risk-decreasing coverage schedules, and the fund must hold a prudential reserve consistent with the regulator-defined solvency tolerance.

The remainder of the paper is organized as follows. Section 3.2 reviews the related literature. Section 3.3 presents the analytical framework, states the two central propositions, and formalizes the stochastic solvency condition. Section 3.4 describes the simulation design and the empirical calibration. Section 3.5 reports the four decision-support outputs. Section 3.5.5 reports the sensitivity analysis. Section 3.6 discusses policy implementation. Section 3.7 concludes. Proofs are collected in the appendix.

3.2 Literature review and decision-support positioning

The framework developed in this paper draws on six strands of literature, which we discuss in the sequence implied by the structure of the policy problem. The institutional and adequacy

literature on individual capitalization systems provides the policy environment. The life-cycle portfolio-choice literature delivers the economic logic of recovery capacity. The rare-disaster literature motivates threshold-activated protection rather than broad portfolio insurance. The literatures on intergenerational risk sharing and on insurance theory with moral hazard discipline the choice of pricing rule. The literatures on minimum pension guarantees, target-date funds, and actuarial ruin theory delimit the boundaries of the instrument and supply the institutional machinery for the solvency analysis. Taken jointly, these six strands cover the components of the problem; none of them, however, integrates shock timing, pension-equivalent loss, an explicit activation trigger, a pricing menu, an endogenous portfolio-risk margin, and a stochastic solvency condition into a single decision-support problem with welfare non-deterioration. That integration is the unresolved gap that this paper addresses.

Funded multi-pillar pension arrangements have been adopted in several countries since the policy case articulated by World Bank (1994), with Chile and other Latin American economies among the most studied early adopters (Arenas de Mesa and Mesa-Lago, 2006; Berstein et al., 2006). The accumulated evidence shows that funded individual-account systems can coexist with persistent public-policy challenges related to coverage, contribution density, gender gaps, administrative costs, and replacement-rate adequacy (Iglesias-Palau, 2009; OECD, 2025). Barr and Diamond (2009) emphasize that funded accounts improve some dimensions of transparency and ownership but do not eliminate the need for public design, since regulation is still required to determine contribution rates, eligible portfolios, default investment rules, annuitization options, and risk-sharing mechanisms. The relevant object for our analysis is not volatility in accumulated balances, but the conversion of adverse return realizations into pension-adequacy losses. Pension adequacy depends on the interaction between market returns, contribution histories, retirement age, annuity conversion, and the remaining time available to recover from losses. A loss that is manageable early in the working life can become binding when it occurs near retirement, which is the channel that motivates the rest of the paper.

The economic logic of differential recovery comes from life-cycle portfolio theory. The classical models of Samuelson (1969), Merton (1969), and Merton (1971) establish that optimal saving and portfolio decisions depend on preferences, investment opportunities, and the planning horizon. Subsequent work emphasizes the role of labor income as a non-traded asset. When future labor income is relatively safe, it behaves like implicit bond wealth and supports greater exposure to risky financial assets at younger ages (Bodie et al., 1992; Viceira, 2001). Quantitative life-cycle models with labor-income risk show that optimal risk-taking varies with age, accumulated wealth, and the persistence of income shocks (Cocco et al., 2005; Gomes and Michaelides, 2005), and the recent survey by Gomes (2020) consolidates the empirical and theoretical evidence on age-dependent risk-taking. These mechanisms are directly relevant for pension systems based on individual

accounts, since a young worker holds a large stock of remaining human wealth and many future contributions, while a worker close to retirement has a smaller recovery margin and a larger accumulated balance exposed to late-career losses. We build on this logic but shift the object of analysis. Rather than solving for the privately optimal portfolio profile, we use recovery capacity as a policy-design variable that determines where insurance is most valuable and where social pooling is most defensible.

The shocks considered here are not ordinary return fluctuations. They are lower-tail events large enough to threaten pension adequacy, and the rare-disaster literature shows that infrequent but severe contractions have first-order consequences for asset prices and welfare. Rietz (1988) introduced the idea that a small probability of catastrophic decline could help rationalize the equity premium, and Barro (2006, 2009) document that disasters of this type recur often enough to matter for risk pricing and welfare. Gabaix (2012) and Wachter (2013) show that time-varying disaster risk can account for substantial variation in asset prices. For individual capitalization systems, the institutional implication is that a rare disaster may be diversifiable at the level of society or across cohorts but is not fully diversifiable by an individual worker close to retirement. This delivers a rationale for a public or regulated insurance layer targeted at lower-tail pension losses rather than at ordinary portfolio volatility. Insuring every negative return would be expensive and would weaken investment incentives, while insuring only threshold-activated pension losses focuses protection on states in which private recovery is least effective and public intervention is more likely to pass a welfare test. The activation threshold κ therefore plays a central role in the model, since it separates normal market fluctuations from severe events that generate adequacy-relevant losses. A return-based threshold is also operationally attractive because realized fund returns are observable, verifiable, and reduce regulatory discretion, although we acknowledge in Section 3.6 that a return-based trigger does not by itself capture all sources of individual heterogeneity.

The literatures on intergenerational risk sharing and on insurance design supply the constraints that the pricing regime must satisfy. Gordon and Varian (1988) and Ball and Mankiw (2007) show that intergenerational risk sharing can improve welfare when private markets are incomplete, and Merton (1983) argues that public arrangements can improve risk allocation because human capital is not tradable and cannot be fully hedged in financial markets. In the pension context, Gollier (2008) demonstrates that collective smoothing of return shocks can improve welfare relative to pure individual accounts, while Cui et al. (2011) and Beetsma et al. (2012) study collective and target-benefit pension designs in which investment risk is partially pooled across participants. Krueger and Kubler (2006) shows that intergenerational risk-sharing reforms can be Pareto improving even when financial markets are incomplete. These contributions establish that pooling may be welfare enhancing, especially when individuals face risks that are difficult to insure privately. We differ in the object and in the discipline of pooling. Rather than smoothing the entire return

process, we study an insurance contract activated only by severe pension-loss states, which targets adequacy failures while limiting fiscal exposure and allows the planner to compare fair, pooled, and hybrid pricing rules. The central issue is not whether pooling is valuable in general, but how much pooling is admissible once budget balance, welfare non-deterioration, stochastic solvency, and behavioral responses are imposed jointly.

The literature on minimum pension guarantees provides a complementary perspective that is closely related to the instrument studied here but differs in scope. Pennacchi (1999) develops a contingent-claims valuation of guarantees on pension fund returns in the context of the Chilean and Uruguayan defined-contribution arrangements, modeling the guarantee as a put option whose strike is either a fixed minimum return or a return measured relative to the system average. Lewis and Pennacchi (1994) and Brown and Pennacchi (2016) extend this approach to the valuation of guarantee liabilities under the Pension Benefit Guaranty Corporation in the United States, and the broader landscape of guarantees in defined-contribution pension systems is surveyed in Antolín et al. (2011) and Munnell et al. (2009). These contributions share with our framework the focus on the public-sector cost of protecting pension accumulation against adverse return realizations, but they study a different instrument. A minimum-return guarantee is typically active in every state in which realized returns fall below a contractual benchmark, and its cost is therefore proportional to the expected shortfall over the full return distribution. The instrument analyzed in this paper is narrower, since it activates only when realized returns cross a downside threshold κ, and it pays a fraction of the resulting pension-equivalent loss rather than topping up the balance to a contractual floor. This narrower design has two consequences. First, the cost can be calibrated independently of the typical level of returns, which makes the instrument less exposed to long-run monetary and growth conditions than a level guarantee. Second, the coverage rule can depend on age and on the regulator-classified portfolio category, which addresses the recovery-capacity asymmetry that a level guarantee does not directly target. The framework therefore complements rather than substitutes for the guarantee literature, and the two approaches can be compared within a unified admissible-region analysis if the planner is willing to evaluate both instruments under the same welfare and solvency restrictions.

A second body of work that is directly relevant for the coverage rule $\alpha(a, z)$ is the literature on target-date and life-cycle default funds. Poterba et al. (2009a) document the consequences of life-cycle asset-allocation strategies for the distribution of retirement wealth in 401(k) accounts, and Poterba et al. (2009b) study no-loss strategies designed to reduce downside risk for personal retirement accounts. The recent review by Suárez et al. (2025) synthesizes the state of the art on target-date fund design and applies it to the Chilean pension reform. These contributions share with our framework the idea that the appropriate exposure profile is not constant across the working life and that protection against late-career losses requires a specific institutional instrument.

Their solution is portfolio-level, since the default fund mechanically reduces equity exposure as the participant approaches retirement, which is the so-called glide path. Our solution is insurance-level, since the coverage rule $\alpha(a, z)$ increases with age and decreases with portfolio risk but does not require the worker to give up the upside of the chosen portfolio. The two mechanisms can be combined, since a glide path reduces ex-ante exposure while pension-loss insurance protects against tail realizations of the residual exposure that the glide path does not eliminate. In an individual capitalization system that already operates a multi-fund default rule, the framework developed here is best read as a layer on top of the existing life-cycle structure, not as a substitute for it.

A third strand of work that becomes binding for our framework once stochastic solvency is formalized is the actuarial literature on insurance-fund ruin. The classical Cramér-Lundberg model (Cramér, 1930), extended for heavy-tailed claim distributions in Embrechts et al. (1997), studies the probability that the surplus of an insurance portfolio crosses zero in finite or infinite time as a function of premium income, claim severity, and initial reserves. Boyer (2009) extends this perspective to catastrophic pension-fund losses and discusses the role of alternative risk transfer in their financing. The transition from classical ruin theory to modern solvency assessment is discussed in Eling et al. (2007). These contributions provide the formal foundations of the stochastic solvency condition introduced in Section 3.3.2 and the prudential reserve $\mathcal{V}_{1-\varepsilon}(\Pi)$ reported in Section 3.5.4. The instrument-specific feature in our setting is that the claim arrival process is threshold-activated rather than Poisson-arrival, and the severity distribution is calibrated empirically to a specific class of pension-fund losses rather than postulated from an actuarial table. The institutional implication is the same, however, since expected balance is a necessary but not a sufficient condition for the fund to survive the realized timing of rare shocks, and a prudential reserve scaled to the implied ruin probability is the natural decision-support output that connects the analytical framework to a supervisory rule.

Insurance design cannot ignore incentives. When coverage affects behavior, full insurance is generally inefficient, and optimal contracts use deductibles, coinsurance, exclusions, or risk-based premia to preserve effort and limit moral hazard (Arrow, 1963; Rothschild and Stiglitz, 1976; Holmström, 1979; Raviv, 1979). In our setting, the relevant behavioral margin is not effort in loss prevention but pension-portfolio risk. If a worker receives the same insurance protection regardless of the chosen risk exposure, the contract reduces the marginal cost of risk-taking. This creates a tension specific to pension-loss insurance, since the workers who most need protection are those close to retirement but these workers may also react to insurance by taking more portfolio risk if the contract is too generous or if the premium is not risk-sensitive.

Taken together, the six strands above leave a structural gap that the present paper addresses. Life-cycle portfolio models deliver the economic logic of recovery capacity but take the institutional environment as given, while our framework derives an age-contingent coverage rule from that

asymmetry. Rare-disaster and tail-risk models supply the asset-pricing apparatus but do not map disaster risk into pension-equivalent losses across ages, a mapping that the present paper formalizes through the shock-timing decomposition. Intergenerational risk-sharing models study broad smoothing rules and motivate the comparison of pricing regimes that we operationalize through the budget-balancing fair premium locus. Pension-guarantee work analyzes a broader and more expensive instrument, complementary to ours rather than substitute, since the threshold-activated layer we study targets only the adequacy-relevant tail of pension losses. Target-date and life-cycle default funds operate at the portfolio level rather than the insurance level, and the coverage rule we propose sits naturally on top of an existing default rule rather than substituting for it. Insurance theory under moral hazard supplies the conceptual framework for evaluating pricing discipline, but its canonical behavioral margin is preventive effort rather than endogenous portfolio risk; the pricing-discipline diagnostics in our framework adapt that logic to the pension setting and bound the moral-hazard gap explicitly. Ruin theory and Solvency-II frameworks supply the actuarial machinery for fund-level solvency but are not specialized to threshold-activated pension-loss claims or to mature-regime population structure, and our framework specializes both by constructing a prudential reserve that translates the solvency condition into an institutional capital benchmark. The present paper integrates these six strands into a single decision-support problem that yields four regulator-facing outputs: the admissible region, the welfare-neutral frontier, the budget-balancing fair premium locus, and the prudential reserve. This integrated output is what no individual strand of the prior literature delivers, and it locates the paper in the tradition of quantitative public-sector decision support applied to other multidimensional policy problems (Orozco-Villodres et al., 2026; Durães et al., 2026; Ciacci et al., 2026).

3.3 Analytical framework

This section presents the analytical structure used to evaluate pension-loss insurance as a public-sector planning instrument. The framework has five components. First, workers accumulate pension wealth through mandatory contributions and stochastic pension-fund returns. Second, rare market shocks generate pension-equivalent losses when returns fall below an activation threshold. Third, an insurance contract reimburses a fraction of those losses and is financed through payroll-based premia. Fourth, a social planner selects admissible policy packages subject to budget balance, incentive compatibility, and welfare non-deterioration relative to a no-insurance benchmark. Fifth, the planner imposes a stochastic solvency condition on the insurance fund that translates the timing concentration of rare shocks into a prudential reserve requirement. We then state two propositions that formalize the central economic mechanisms; their proofs are in Appendix 3.8.

3.3.1 The life-cycle worker problem

Time is discrete and workers are indexed by age $a = 0, \ldots, T$, with T denoting retirement. During working life, an individual receives labor income y_a, consumes c_a, chooses end-of-period liquid saving b_{a+1}, and contributes a mandatory fraction τ of labor income to an illiquid pension account. Pension wealth is denoted by W_a. Workers also choose a pension-portfolio risk exposure $z_{a+1} \in \mathcal{Z}$, which affects the distribution of the pension-fund return.

Liquid saving before the risk-free return is

$$b_{a+1} = B_a + (1 - \tau - \psi(a, z_{a+1}; \lambda))y_a - c_a, \tag{3.1}$$

and next-period liquid wealth is

$$B_{a+1} = (1 + r_f)b_{a+1}. \tag{3.2}$$

Here, r_f is the risk-free rate and $\psi(a, z_{a+1}; \lambda)$ is the payroll premium that finances the insurance layer. This notation separates the saving decision from beginning-of-next-period liquid wealth and avoids ambiguity in the Euler equation.

The observable pension-fund return is $R^{\text{obs}}_{a+1}(z_{a+1})$, defined on the same market state as the pre-crash innovation. The insurance contract is activated when the observable return crosses the downside threshold $-\kappa$:

$$J_{a+1} = \mathbf{1}\left\{R^{\text{obs}}_{a+1}(z_{a+1}) < -\kappa\right\}. \tag{3.3}$$

The eligible pension loss is defined directly on the observable return:

$$L_{a+1} = J_{a+1}\left(W_a + \tau y_a\right)\left[-\kappa - R^{\text{obs}}_{a+1}(z_{a+1})\right]^+, \tag{3.4}$$

where $[x]^+ = \max\{x, 0\}$. The observable return enters all empirical objects directly: trigger (3.3), eligible loss (3.4), and the activation-effective expected loss rate $\ell(a, z; \kappa) = \mathbb{E}[J_{a+1}[-\kappa - R^{\text{obs}}_{a+1}(z)]^+ \mid z]$ used in the pricing rules of Section 3.3.2. Operationally, this means that the trigger is verifiable from observed pension-fund returns and does not depend on a latent decomposition.

For analytical purposes it is useful to interpret the observable return as a mixture of a normal-fluctuation component and an excess-severity jump. We model the observable return as

$$R^{\text{obs}}_{a+1}(z_{a+1}) \;=\; \widetilde{R}^{p}_{a+1}(z_{a+1}) \;-\; \chi_{a+1}(z_{a+1})\,\mathbf{1}\left\{\widetilde{R}^{p}_{a+1}(z_{a+1}) - \chi_{a+1}(z_{a+1}) < -\kappa\right\},$$

where $\widetilde{R}^{p}_{a+1}(z_{a+1}) = \mu(z_{a+1}) + \sigma_R(z_{a+1})\varepsilon_{a+1}$ is the normal-fluctuation component driven by a market innovation ε_{a+1}, and $\chi_{a+1}(z_{a+1}) \geq 0$ is the excess severity that materializes only when the threshold is crossed. The expression is non-circular: the indicator on the right-hand side depends only on the latent primitives $(\widetilde{R}^{p}, \chi)$, and the observable trigger J_{a+1} in (3.3) is then

$J_{a+1} = \mathbf{1}\{R^{\text{obs}}_{a+1}(z_{a+1}) < -\kappa\}$ by construction. The decomposition serves two purposes. First, it makes explicit that the excess-severity component is not added on top of an already-observed realized return; the observable return R^{obs}_{a+1} is the difference $\widetilde{R}^p - \chi$ in activated states and equals $\widetilde{R}^p$ otherwise. Second, it provides a clean numerical pipeline in which the empirical activation probability and the empirical conditional severity are calibrated separately to the Chilean multi-fund evidence of Section 3.4.2. The trigger in (3.3) and the eligible-loss formula (3.4) are defined entirely on the observable return; the decomposition is interpretive and is used in the analytical proofs and in the calibration of $\ell(a, z; \kappa)$.

Before insurance, pension wealth after contributions and returns is

$$W^s_{a+1} = (W_a + \tau y_a)\left[1 + R^{\text{obs}}_{a+1}(z_{a+1})\right].$$

Preferences are represented by expected lifetime utility. Let $u(c)$ be increasing and concave, and let $\Gamma(z)$ denote the private disutility or adjustment cost associated with selecting a riskier pension portfolio. The worker's value function during working life is

$$V_a(W_a, B_a, y_a) = \max_{c_a, b_{a+1}, z_{a+1}} \left\{u(c_a) - \Gamma(z_{a+1}) + \beta \mathbb{E}_a\left[V_{a+1}(W_{a+1}, B_{a+1}, y_{a+1})\right]\right\}, \tag{3.5}$$

subject to (3.1)–(3.2) and the pension-wealth law of motion defined below. At retirement, pension wealth is converted into retirement income through an actuarial conversion function $g(W_T, A_T)$, where A_T is an annuity factor or conversion parameter. The terminal value is

$$V_T(W_T, B_T) = U_R\left(g(W_T, A_T), B_T\right), \tag{3.6}$$

where $U_R(\cdot)$ captures the value of pension income and liquid wealth in retirement.

The insurance contract reimburses a fraction of the crash-induced pension loss whenever the activation condition in (3.3) holds. Coverage is allowed to depend on age and portfolio risk:

$$I_{a+1} = \alpha(a, z_{a+1}) L_{a+1}, \qquad 0 \le \alpha(a, z_{a+1}) \le 1.$$

Final pension wealth under insurance is therefore

$$W_{a+1} = (W_a + \tau y_a)\left[1 + R^{\text{obs}}_{a+1}(z_{a+1})\right] \; + \; \alpha(a, z_{a+1})\, L_{a+1}, \tag{3.7}$$

where L_{a+1} is the eligible loss defined in (3.4). Equation (3.7) makes explicit the role of the insurance layer. When no crash is activated, $J_{a+1} = 0$, so $L_{a+1} = 0$ and pension wealth follows the observed return. When a crash is activated, the contract pays the fraction $\alpha(a, z_{a+1})$ of the eligible loss back into the pension account, so that the uninsured residual is the remaining fraction

$1 - \alpha(a, z_{a+1})$ of L_{a+1}. Full insurance corresponds to $\alpha = 1$, in which case the eligible loss is fully reimbursed; the no-insurance benchmark corresponds to $\alpha = 0$.

The coverage schedule is the key instrument for targeting recovery capacity. Workers close to retirement have fewer remaining contributions and less time to rebuild pension wealth after a severe loss, so a planning rule can assign higher coverage to ages close to retirement. Because z_{a+1} is chosen by the worker, coverage cannot be separated from incentives. If coverage is too generous for high-risk portfolios, the insurance layer may induce excessive risk-taking. The framework therefore allows $\alpha(a, z)$ either to increase with age, decrease with risk, or combine both margins. We discuss in Section 3.6 how this rule could be implemented through regulator-classified portfolio categories rather than continuous individual risk measurement.

3.3.2 Pricing rules and planning restrictions

Let the activation-effective expected loss rate be

$$\ell(a, z; \kappa) = \mathbb{E}\Big[J_{a+1} \left[-\kappa - R^{\text{obs}}_{a+1}(z) \right]^+ \Big| \, a, z, \kappa \Big], \tag{3.8}$$

which expresses the expected eligible loss per unit of insured exposure in terms of the observable return $R^{\text{obs}}_{a+1}(z)$ defined in Section 3.3.1. The object $\ell(a, z; \kappa)$ combines the probability that the threshold is crossed and the expected shortfall below the threshold conditional on activation, both measured on the observable return. Under the latent decomposition introduced in Section 3.3.1, the same object can be written equivalently as $\ell(a, z; \kappa) = \mathbb{E}[J_{a+1}\{\chi_{a+1}(z) - [\kappa + \widetilde{R}^{p}_{a+1}(z)]^+\}^+ \mid a, z, \kappa]$, which makes explicit that only the portion of severity that pushes the observable return below $-\kappa$ enters the indemnity flow. The observable form is the operative one, since both the indemnity (3.4) and the fair premium (3.9) are computed from the observable return alone.

The expected indemnity for a worker of age a choosing risk z is

$$\mathbb{E}_a[I_{a+1}] = \alpha(a, z) \, (W_a + \tau y_a) \, \ell(a, z; \kappa).$$

The first pricing rule is actuarially fair. The individual payroll premium is set so that expected contributions to the insurance fund equal expected indemnities:

$$\psi^F(a, z) = \frac{\alpha(a, z) \, (W_a + \tau y_a) \, \ell(a, z; \kappa)}{y_a}. \tag{3.9}$$

This rule preserves incentive discipline because the premium increases with insured exposure and with the expected activation cost of the chosen portfolio.

The second rule is uniform social pooling. Under this rule, all active workers pay the same

payroll premium, regardless of age or portfolio risk. Let m index worker types in the mature-regime population, with population weight ω_m. The pooled premium is

$$\psi^P = \frac{\sum_m \omega_m \mathbb{E}_m[I_{a+1}]}{\sum_m \omega_m y_m}.$$

This rule spreads costs across contributors and can support redistribution toward workers with lower recovery capacity. Its drawback is that it weakens the connection between individual risk choices and insurance costs.

The third rule is hybrid pricing. The premium paid by a worker of age a choosing portfolio risk z is

$$\psi^H(a, z; \lambda) = \lambda \psi^F(a, z) + (1 - \lambda)\psi^P, \qquad \lambda \in [0, 1]. \tag{3.10}$$

The parameter λ measures the degree of risk-based pricing. When $\lambda = 1$, the premium is fully actuarially fair. When $\lambda = 0$, the premium is fully pooled. Intermediate values combine partial solidarity with partial incentive discipline.

We now formalize the planner's problem and the restrictions under which a policy package is admissible. The planner selects a policy package

$$\Pi = \{\tau, \alpha(a, z), \psi(a, z; \lambda), \kappa, \lambda\}_{a,z}.$$

The benchmark policy is the no-insurance regime

$$\Pi^0 = \{\tau, \alpha(a, z) = 0, \psi(a, z) = 0, \kappa, \lambda\}_{a,z}. \tag{3.11}$$

Let $\mathcal{W}(\Pi)$ denote aggregate mature-regime welfare under policy Π. The mature-regime convention requires that all active and retired cohorts observed in the cross-section have been exposed to the same policy environment throughout their working lives. We use this convention because it isolates the steady-state policy effect from transitional dynamics that would depend on the implementation date. Aggregate welfare is

$$\mathcal{W}(\Pi) = \sum_a \omega_a \int V_a^{\Pi}(W, B, y)\, dF_a^{\Pi}(W, B, y), \tag{3.12}$$

where ω_a is the population weight of age a and F_a^{Π} is the age-specific distribution of states induced by policy Π.

A policy satisfies welfare non-deterioration if

$$\mathcal{W}(\Pi) \geq \mathcal{W}(\Pi^0). \tag{3.13}$$

We refer to (3.13) interchangeably as welfare neutrality or welfare non-deterioration. The label *welfare-neutral* should be read as a constraint that the insurance scheme not reduce aggregate mature-regime welfare relative to the benchmark, not as a requirement of exact equality.

The insurance fund must also satisfy expected budget balance:

$$\sum_a \omega_a \mathbb{E}_a \left[\psi(a, z_a; \lambda) y_a\right] \geq \sum_a \omega_a \mathbb{E}_a \left[I_{a+1}\right]. \tag{3.14}$$

When the inequality is strict, the residual can be interpreted as a reserve margin or administrative buffer. Expected budget balance is analytically convenient but does not, by itself, guarantee solvency under a stochastic claims process. The stochastic solvency condition is introduced as a separate restriction in Section 3.3.2.

The planner evaluates admissible policies through pension-adequacy and incentive metrics. Let $p_T = g(W_T, A_T)$ denote retirement pension income and let $\bar{p}$ be an adequacy threshold. Pension inadequacy is

$$\mathcal{A}(\Pi) = \Pr_{\Pi}\left(p_T < \bar{p}\right). \tag{3.15}$$

Lower-tail pension risk can also be summarized by a tail statistic such as

$$\mathcal{T}_q(\Pi) = Q_q^{\Pi}(W_T), \qquad q \in (0, 0.5), \tag{3.16}$$

where $Q_q^{\Pi}(W_T)$ is the q-th quantile of terminal pension wealth.

The simulations distinguish two incentive diagnostics. The first is the pricing-discipline gap, which compares risk-taking under policy Π with risk-taking under actuarially fair pricing:

$$\mathcal{M}^F(\Pi) = \mathbb{E}_{\Pi}\left[z_a^{\Pi} - z_a^F\right]. \tag{3.17}$$

The second is the gross insured-risk gap, which compares risk-taking under policy Π with the no-insurance benchmark:

$$\mathcal{M}^0(\Pi) = \mathbb{E}_{\Pi}\left[z_a^{\Pi} - z_a^0\right]. \tag{3.18}$$

The first measure isolates the loss of pricing discipline. The second measure reports the total increase in insured risk exposure relative to the uninsured economy.

The expected budget-balance condition (3.14) is analytically convenient but does not, by itself, ensure solvency. Because activation events are rare and concentrated in time, the insurance fund must hold stochastic reserves to cover concentrated bursts of claims even when long-run expected revenue covers long-run expected claims. We formalize this requirement explicitly as a fourth restriction of the planner's problem.

Let $\Phi_t(\Pi)$ denote the cumulative net financial flow of the insurance fund at the end of year t

under policy Π, measured per worker in units of the entry wage. The flow evolves according to

$$\Phi_{t+1}(\Pi) = \Phi_t(\Pi) + \mathbb{E}_m[\psi(a_m, z_m; \lambda)\, y_m] - \mathbb{E}_m[I_{t+1,m}]\,, \qquad \Phi_0(\Pi) = 0, \tag{3.19}$$

where the expectation $\mathbb{E}_m$ is taken across the mature-regime population at date t. The minimum reserve over the planning horizon is

$$\Phi_{\min}(\Pi) \equiv \min_{t \in \{0,\dots,T\}} \Phi_t(\Pi), \tag{3.20}$$

and the stochastic solvency condition requires that the probability of a negative cumulative reserve at any point of the planning horizon does not exceed a regulator-defined tolerance $\varepsilon \in (0,1)$:

$$\Pr(\Phi_{\min}(\Pi) < 0) \leq \varepsilon. \tag{3.21}$$

A natural tolerance is $\varepsilon = 0.05$, which mirrors the supervisory practice of holding capital sufficient to cover at least the 95th percentile of adverse outcomes. The corresponding prudential reserve is the Value-at-Risk at the complementary confidence level,

$$\mathcal{V}_{1-\varepsilon}(\Pi) \equiv -\,Q_\varepsilon(\Phi_{\min}(\Pi))\,, \tag{3.22}$$

where Q_ε denotes the ε-quantile of the distribution of $\Phi_{\min}(\Pi)$. The object $\mathcal{V}_{1-\varepsilon}(\Pi)$ is the reserve that the fund must hold ex ante so that solvency is preserved in all but a fraction ε of outcomes.

The planner's problem with stochastic solvency becomes

$$\min_{\Pi} \quad \Omega_1\, \mathcal{A}(\Pi) \;-\; \Omega_2\, \mathcal{T}_q(\Pi) \;+\; \Omega_3\, \mathcal{M}^F(\Pi) \tag{3.23}$$

subject to welfare non-deterioration (3.13), expected budget balance (3.14), stochastic solvency (3.21), the maximum tolerable moral-hazard gap

$$\mathcal{M}^F(\Pi) \;\leq\; \overline{\mathcal{M}}^F, \tag{3.24}$$

the worker's optimality conditions, and feasibility restrictions on the policy package. The pricing-discipline gap therefore enters the planning problem in two distinct ways. As a penalty term in the objective (3.23) with weight Ω_3, it reflects the planner's preference for designs that internalize the marginal cost of pension-portfolio risk. As a hard restriction in (3.24), it imposes a regulator-defined upper bound $\overline{\mathcal{M}}^F$ on the loss of pricing discipline that the system can tolerate before the design is rejected as inadmissible. The two readings are complementary: the penalty term shapes the planner's ranking inside the admissible region, while the restriction defines the

boundary of that region. The admissible region is therefore the intersection of four restrictions rather than three, with the prudential reserve $\mathcal{V}_{1-\varepsilon}(\Pi)$ translating the solvency condition into an institutional capital requirement and the bound $\overline{\mathcal{M}}^F$ translating the incentive-discipline concern into a regulator-defined ceiling. We treat (3.23)–(3.24) as a decision-support problem in which regulators may pick $(\kappa, \psi, \alpha, \lambda)$ subject to the four restrictions jointly, rather than as a single normative optimum.

3.3.3 Two central mechanisms

The next two propositions formalize the central mechanisms of the framework. The first establishes that, holding shock severity fixed, the expected terminal pension damage is non-decreasing in the age at which the shock occurs once contributions are strictly positive and non-crash returns have non-negative expected net growth. The second establishes that uniform pooling weakens the marginal cost of risk and induces a strictly higher equilibrium portfolio risk than fair pricing, under a primitive condition on the activation-effective expected loss rate that is stated as Lemma 1 below. Proofs are collected in Appendix 3.8. The two results show that the central numerical patterns are not artifacts of a particular parameterization; they reflect properties of the model that hold under transparent qualitative conditions.

Assumption 1 (Population path properties)**.** *Contributions are strictly positive at every working age, $\tau y_t > 0$ for all $t < T$. Non-crash returns satisfy*

$$\mathbb{E}\Big[1 + \widetilde{R}^p_{t+1}(z_{t+1}) \,\Big|\, J_{t+1} = 0\Big] \;\geq\; 1 \qquad \text{for every } t \in \{a_0, \ldots, T-1\}. \tag{3.25}$$

Coverage is fixed at $\alpha \in [0, 1)$ and severity is fixed at $\chi_{t+1}(z) = \chi > 0$ for every t and z. We refer to (3.25) *as the non-negative expected net growth condition.*

Assumption 1 is substantially weaker than requiring non-negative non-crash returns realization by realization. It allows non-crash periods with negative realized returns as long as the conditional mean growth factor in non-crash states is at least one. The empirical evidence on multi-fund pension returns reported in Section 3.4.2 satisfies this condition by a wide margin for every regulator-defined fund category.

Proposition 1 (Recovery capacity and shock timing, expected version)**.** *Let pension wealth follow the recursion* (3.7) *and let Assumption 1 hold. Consider a single activated shock at age $s \in \{a_0, \ldots, T-1\}$ along an otherwise no-shock counterfactual. Let $\mathbb{E}[L_T^{(s)}]$ denote the expected terminal pension-equivalent loss conditional on a single activated shock at age s and on no further*

activated shock thereafter. Then $\mathbb{E}[L_T^{(s)}]$ *is non-decreasing in* s*, that is,*

$$s_1 < s_2 \quad \Longrightarrow \quad \mathbb{E}[L_T^{(s_1)}] \leq \mathbb{E}[L_T^{(s_2)}]. \tag{3.26}$$

The inequality is strict whenever either the expected contribution flow between s_1 *and* s_2 *is strictly positive or the cumulative expected non-crash growth factor strictly exceeds one over at least one period in* $\{s_2 + 1, \ldots, T - 1\}$*.*

The economic content of Proposition 1 is unchanged from a path-wise version. A loss at an early age propagates to retirement through more periods of expected accumulation than a loss close to retirement, so the expected terminal pension damage is larger when the shock occurs near retirement. The expected formulation removes the restrictive requirement that every non-crash period delivers a non-negative realized return, and instead asks only that non-crash returns be non-negative in expectation, which is the relevant condition for a population-level decision-support diagnostic. A path-wise corollary under stronger monotone-path conditions is reported in Appendix 3.8.

Lemma 1 (Primitive condition for pricing discipline)**.** *Suppose the coverage rule* $\alpha(a, z)$ *and the activation-effective expected loss rate* $\ell(a, z; \kappa)$ *defined in* (3.8) *are continuously differentiable in* z*, with* $\alpha(a, z)$ *bounded away from zero on the relevant interior range. Suppose further that the insured expected loss, given by the product* $\alpha(a, z)\, \ell(a, z; \kappa)$*, satisfies*

$$\frac{\partial}{\partial z}[\alpha(a, z)\, \ell(a, z; \kappa)] > 0 \tag{3.27}$$

on the same range, that is, riskier pension portfolios generate strictly larger insured expected losses net of the coverage adjustment. Then the actuarially fair premium defined in (3.9) *satisfies*

$$\frac{\partial \psi^F(a, z)}{\partial z} > 0, \tag{3.28}$$

and the hybrid premium satisfies

$$\frac{\partial \psi^H(a, z; \lambda)}{\partial z} = \lambda \frac{\partial \psi^F(a, z)}{\partial z} \qquad \textit{with} \qquad \frac{\partial \psi^P}{\partial z} = 0. \tag{3.29}$$

In particular, the marginal premium cost of pension-portfolio risk is zero under uniform social pooling and strictly positive under any premium that places positive weight on actuarial fairness.

Lemma 1 states the primitive economic content from which the single-crossing property of the worker's marginal condition follows. Condition (3.27) is the condition that the *insured* expected loss is strictly increasing in z, where the insured loss is the product of coverage and the activation-effective

loss rate. This formulation is the appropriate object because it nets out the coverage-side discipline channel: a coverage rule that declines with risk through $\partial\alpha/\partial z < 0$ can in principle offset the activation channel $\partial\ell/\partial z > 0$, and the relevant primitive condition for fair-pricing discipline is that the product strictly increases in z. The empirical calibration in Section 3.4.2 confirms this property for the Chilean multi-fund system once the empirical activation-effective loss rate is combined with the coverage rule (3.35). A sufficient condition under coverage rules where $\partial\alpha/\partial z \leq 0$ is that the loss-rate elasticity dominates the coverage elasticity in absolute terms; the baseline parameterization of Section 3.4.3 satisfies this condition by construction. Under (3.27), fair premia internalize a strictly positive marginal cost of risk, pooled premia do not, and hybrid premia retain a fraction λ of the fair-pricing discipline.

Assumption 2 (Interior risk choice and single crossing)**.** *Fix an age a and a state $x_a = (W_a, B_a, y_a)$. The pension-portfolio risk choice z lies in a compact interval $\mathcal{Z} \subset \mathbb{R}_+$, the coverage rule $\alpha(a, z)$ is fixed across pricing regimes, and Lemma 1 applies on the relevant interior range. The worker's interior risk choice under hybrid pricing weight λ is characterized by the marginal condition*

$$H_a(z; \lambda) \equiv MC_a(z; \lambda) - MB_a(z; \lambda) = 0, \tag{3.30}$$

where $MC_a(z; \lambda)$ is the marginal private cost of risk and $MB_a(z; \lambda)$ is the marginal continuation benefit of risk. The function H_a is continuously differentiable and satisfies

$$\frac{\partial H_a(z; \lambda)}{\partial z} > 0 \qquad \textit{and} \qquad \frac{\partial H_a(z; \lambda)}{\partial \lambda} = y_a \frac{\partial u(c_a)}{\partial c} \frac{\partial \psi^F(a, z)}{\partial z} > 0 \tag{3.31}$$

on the relevant interior region. The first inequality is the standard local uniqueness and stability condition for an interior risk choice. The second inequality is the single-crossing condition, and it follows from Lemma 1 together with strict monotonicity of period utility.

The novelty of Assumption 2 relative to a generic single-crossing postulate is that the cross-derivative $\partial H_a/\partial\lambda$ is now decomposed into observable primitives. It equals the marginal utility of consumption multiplied by the marginal fair premium $\partial\psi^F/\partial z$, which is itself implied by the primitive condition (3.27) of Lemma 1. The single-crossing property is therefore not assumed directly; it is derived from the underlying economic content that riskier portfolios generate larger expected insured losses.

Proposition 2 (Pricing discipline and induced portfolio risk)**.** *Suppose Lemma 1 and Assumption 2 hold, and the worker's risk choice is interior for every $\lambda \in [0, 1]$. Let $z^*(a; \lambda)$ denote the unique solution to $H_a(z; \lambda) = 0$. Then $z^*(a; \lambda)$ is strictly decreasing in λ on the relevant interior range, and in particular*

$$z^P(a) \equiv z^*(a; 0) > z^*(a; 1) \equiv z^F(a), \tag{3.32}$$

where $z^P(a)$ is the optimal risk choice under uniform social pooling and $z^F(a)$ is the optimal risk choice under actuarially fair pricing. The strict ordering is a direct consequence of the primitive condition (3.27) *and not of an undecomposed single-crossing postulate.*

The interpretation is that pricing discipline operates through the insured expected loss. Lemma 1 requires that the product $\alpha(a,z)\,\ell(a,z;\kappa)$ be strictly increasing in z, which is the joint statement that the loss-rate channel $\partial\ell/\partial z > 0$ dominates any coverage-side dampening $\partial\alpha/\partial z \leq 0$. Whenever this condition holds, the marginal fair premium $\partial\psi^F/\partial z$ is strictly positive and uniform social pooling removes that channel entirely. When in addition $\partial\alpha/\partial z = 0$, the coverage rule provides no independent risk-side discipline, and the marginal private cost of risk under pooling collapses to the technological cost $\Gamma'(z)$ alone, so equilibrium pension-portfolio risk rises. The empirical calibration in Section 3.4.2 verifies that the insured expected loss is strictly increasing in the Chilean multi-fund category index for any reasonable activation threshold, which is the primitive condition required by Lemma 1.

Corollary 1 (Aggregate pricing-discipline gap)**.** *If Proposition 2 holds for every worker type m in the mature-regime population with $\omega_m > 0$, then the aggregate pricing-discipline gap satisfies*

$$\mathcal{M}^F(\Pi^P) \;=\; \sum_m \omega_m \big[z_m^P - z_m^F\big] \;>\; 0.$$

3.4 Numerical design and policy diagnostics

This section describes the numerical design used to evaluate the planning framework. The exercise is empirically calibrated to the Chilean multi-fund system, with structural parameters whose qualitative readings extend to other individual capitalization systems with regulator-defined fund categories. Its purpose is to identify internally coherent combinations of coverage, activation thresholds, premia, and incentive constraints, and to map each diagnostic to one of the four design concerns stated in the introduction. All pension-wealth variables are normalized in units of the entry wage. Workers enter the system at age $a_0 = 25$ and retire at age $T = 65$. We make the link between equations and figures explicit at the end of the section, sketch the simulation algorithm in pseudo-code, and provide a full sensitivity analysis in Section 3.5.5.

3.4.1 Simulation environment

The simulation follows the accumulation of pension wealth under the law of motion of Section 3.3. In each period, workers receive labor income y_a, contribute a mandatory fraction τ to the pension account, and choose pension-portfolio risk z_{a+1} from a discrete grid $\mathcal{Z}$. The risk variable z is interpreted throughout as a regulator-classified portfolio category in which the worker is enrolled,

not as a continuous individual measurement; the discretization of $\mathcal{Z}$ is consistent with this reading. The pension-fund return is generated in two steps that resolve the activation event in terms of latent primitives and translate it to the observable trigger. In the first step, the normal-fluctuation component is drawn from

$$\widetilde{R}^{p}_{a+1}(z_{a+1}) = \mu(z_{a+1}) + \sigma_R(z_{a+1})\varepsilon_{a+1}, \tag{3.33}$$

and the excess severity $\chi_{a+1}(z_{a+1}) \geq 0$ is drawn from its empirical conditional distribution. In the second step, the activation event is determined and the observable return is constructed as

$$R^{\text{obs}}_{a+1}(z_{a+1}) = \begin{cases} \widetilde{R}^{p}_{a+1}(z_{a+1}) - \chi_{a+1}(z_{a+1}) & \text{if } \widetilde{R}^{p}_{a+1}(z_{a+1}) - \chi_{a+1}(z_{a+1}) < -\kappa, \\ \widetilde{R}^{p}_{a+1}(z_{a+1}) & \text{otherwise,} \end{cases}$$

so that the trigger

$$J_{a+1} = \mathbf{1}\big\{R^{\text{obs}}_{a+1}(z_{a+1}) < -\kappa\big\} \tag{3.34}$$

is well-defined as a function of observed returns alone. The eligible pension loss is $L_{a+1} = J_{a+1}(W_a + \tau y_a)\big[-\kappa - R^{\text{obs}}_{a+1}(z_{a+1})\big]^{+}$, the indemnity is $I_{a+1} = \alpha(a, z_{a+1})L_{a+1}$, and the post-insurance pension-wealth recursion follows equation (3.7). This two-step pipeline preserves the analytical decomposition that we use in the proofs of Section 3.3.3 while ensuring that, at the contract level, the trigger and the eligible loss depend only on the observable return $R^{\text{obs}}_{a+1}(z_{a+1})$.

The coverage schedule is parameterized as an age-increasing and risk-decreasing rule:

$$\alpha(a, z) = \min\left\{1, \max\left[0, \alpha_0 + \alpha_{age}\left(\frac{a - a_0}{T - a_0}\right) - \alpha_z z\right]\right\}. \tag{3.35}$$

The age component captures recovery-capacity targeting. The risk component disciplines the moral-hazard margin. Because z is read as a regulatory category, equation (3.35) maps a finite list of categories into a discrete matrix of coverage levels indexed by age band and category, which is implementable with the regulatory infrastructure already used in most individual-account systems.

The numerical implementation is a reduced-form Monte Carlo decision-support simulation disciplined by the analytical propositions of Section 3.3.3. Consumption is set period by period through the Euler condition (3.41), and the pension-portfolio risk choice is selected by grid maximization on $\mathcal{Z}$ given the relevant pricing rule and the resulting marginal premium cost. The implementation does not solve a dynamic programming problem on a continuous state-space grid; instead, residual lifetime payoffs are approximated by forward Monte Carlo simulation under each candidate risk choice, as described in Section 3.4.4. This is a deliberate methodological choice because the qualitative orderings established by Propositions 1 and 2 hold independently of the

numerical method used, and because the focus of the exercise is the admissible region in the policy space and the four decision-support outputs of Section 3.5.1 rather than the optimal worker policy function across the full state space.

The framework uses expected budget balance as the central financing restriction in the welfare-frontier exercise, complemented by the stochastic solvency restriction of Section 3.3.2 in the fund-design exercise. Implementation of a real-world insurance fund against rare market shocks requires stochastic reserves, an explicit fiscal backstop, or a reinsurance arrangement, because activation events are concentrated in time and large in magnitude. We discuss the alternative arrangements in Section 3.6.

3.4.2 Empirical calibration to the Chilean multi-fund system

The structural parameters of the return process are calibrated to the Chilean multi-fund pension system, whose real monthly returns deflated by the Unidad de Fomento (UF) are reported by the Superintendencia de Pensiones for each fund category from October 2002 to April 2026. Table 3.1 summarizes the annualized moments by fund category, the implied empirical activation frequencies for three activation thresholds, and the mapping from the model variable z to the regulator-defined category. The calibration sample contains 283 monthly observations across the five regulator-defined categories.

Table 3.1: Empirical calibration to the Chilean multi-fund system. Annualized moments based on monthly real returns deflated by UF, October 2002 to April 2026. Tail probabilities are reported in two forms. The columns $\widehat{\Pr}$ report raw empirical frequencies based on calendar-year cumulative real returns 2003–2025 (23 complete years). The columns $\widetilde{\Pr}$ report smooth estimates obtained by fitting a Student-t distribution to monthly returns of each fund (with the regularization $\nu \geq 4$ as a discipline that prevents extremely heavy-tail overfitting on a small number of extreme observations, binding for Funds D and E and non-binding for Funds A, B, C), then computing the annualized tail probability by Monte Carlo simulation of 12-month cumulative returns under the fitted distribution. Fund A is the riskiest category, Fund E is the most conservative.

Fund	z code	μ_{ann}	σ_{ann}	$\widehat{\Pr}(R < -0.10)$	$\widehat{\Pr}(R < -0.15)$	$\widetilde{\Pr}(R < -0.10)$	$\widetilde{\Pr}(R < -0.15)$
A (riskiest)	1.20	6.61%	12.26%	0.087	0.087	0.0605	0.0245
B	0.95	5.49%	9.21%	0.087	0.043	0.0304	0.0084
C (balanced)	0.70	4.57%	6.80%	0.043	0.043	0.0101	0.0014
D	0.50	3.82%	5.85%	0.043	0.000	0.0060	0.0010
E (safest)	0.35	3.23%	5.61%	0.043	0.000	0.0048	0.0007

Three features of Table 3.1 are central to the calibration. First, mean real returns are strictly positive across all five fund categories, which validates Assumption 1 on non-negative expected net growth in non-crash states by a wide margin. Second, the raw empirical frequencies $\widehat{\Pr}$

are weakly increasing in the fund category index across all thresholds, with strict differences between the broad risk classes (high-risk A–B versus balanced C versus conservative D–E), but they exhibit ties within risk classes and rounding artefacts in the deepest tail because the 23-year sample contains too few events to identify the lower-tail shape non-parametrically. The smooth probabilities $\widetilde{\Pr}$ obtained from the fitted Student-t model are strictly increasing in the fund category index across the policy-relevant range $\kappa \in [0.05, 0.15]$, with $\widetilde{\Pr}(R < -\kappa)$ separated by approximately a factor of two between adjacent funds at the baseline threshold. Combined with the age-and-risk-contingent coverage rule (3.35), the insured expected loss $\alpha(a, z)\,\ell(a, z; \kappa)$ is strictly increasing in the fund category index across the relevant range, which delivers the primitive condition $\partial[\alpha(a, z)\ell(a, z; \kappa)]/\partial z > 0$ of Lemma 1 as a direct empirical implication once the small-sample noise in the raw frequencies is smoothed away by the parametric model. Third, the most damaging activation episode in the sample, the global financial crisis of 2008, produced annual real returns of -40.3% in Fund A, -30.1% in Fund B, -18.9% in Fund C, -9.9% in Fund D, and -0.9% in Fund E, which provides a natural calibration of severity by category: the conditional mean excess severity above the threshold $\kappa = 0.10$ in Fund C is approximately 9%, and it scales upward for riskier categories and downward for safer ones.

The fitted Student-t parameters per fund are obtained by maximum likelihood on the monthly real-return series under the constraint $\nu \geq 4$, a regularization that prevents extremely heavy-tail overfitting on the small number of extreme observations in the 2002–2026 sample. The constraint is implemented as an unconstrained MLE in (ν, μ_t, σ_t) followed by a Lagrangian refit fixing $\nu = 4$ in the funds where the unconstrained MLE falls below this bound. The estimated degrees of freedom are $\nu = 4.56$ for Fund A, $\nu = 4.72$ for Fund B, and $\nu = 5.33$ for Fund C, so the constraint is non-binding for these three categories. For the conservative categories, the unconstrained MLE would drive ν below 3 and inflate the empirical excess kurtosis well beyond plausible long-run levels, so the constraint binds at $\nu = 4$ for both Fund D and Fund E. The fitted location parameters μ_t are 0.00714, 0.00579, 0.00476, 0.00401, and 0.00310 for Funds A through E respectively, and the scale parameters σ_t are 0.02595, 0.01978, 0.01531, 0.01140, and 0.01015. Both parameters increase monotonically with the regulator-defined risk classification, in line with the descriptive statistics in Table 3.1. As a robustness check, we also computed the smooth tail probabilities under a Gaussian fit and under a stationary bootstrap of monthly returns with block length six. The Gaussian fit produces tail probabilities one order of magnitude smaller in the deep left tail for the conservative funds, which understates downside risk in those categories. The bootstrap delivers tail probabilities within Monte Carlo error of the Student-t estimates for Funds A, B, and C, but is uninformative for Funds D and E because the deepest tail of those categories has very few in-sample observations. The constrained Student-t fit is therefore the appropriate disciplinary choice in the policy-relevant range.

Figure 3.1 reports the underlying empirical evidence used in Table 3.1. Panel (a) shows annual real returns for the five fund categories from 2003 to 2025, together with the three activation thresholds $\kappa \in \{0.10, 0.15, 0.20\}$. Panel (b) shows the empirical complementary cumulative distribution function of activation by fund and threshold. The figure validates the baseline activation threshold of the model, since $\kappa = 0.15$ corresponds to an empirical activation frequency of roughly five percent in the median fund category, which is consistent with the conceptual reading of the threshold as a separator between normal market fluctuations and rare adequacy-relevant events.

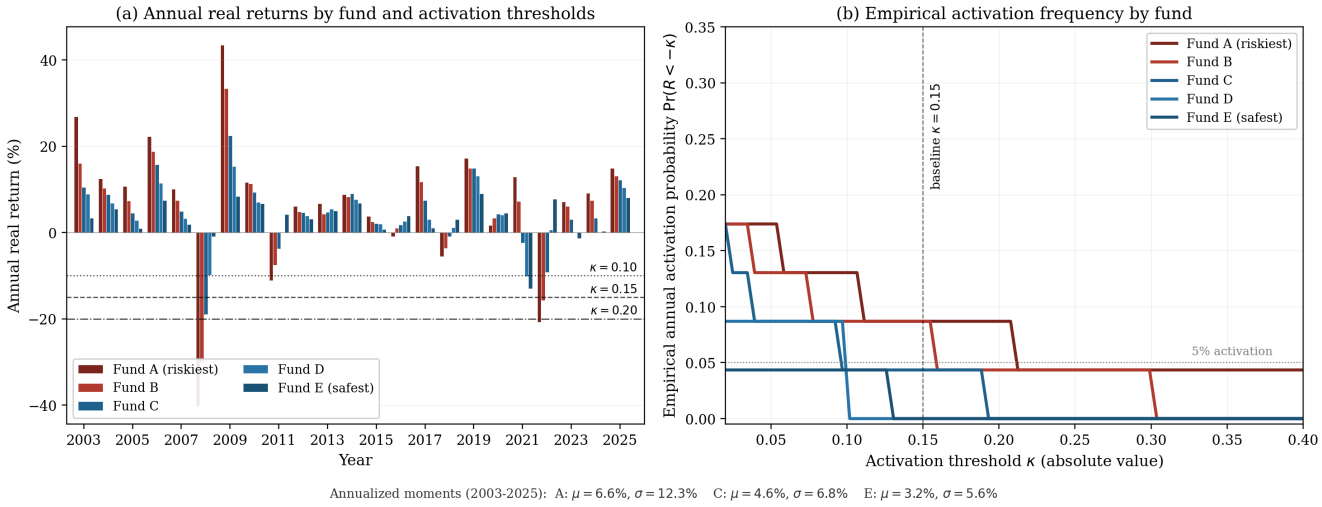


Figure 3.1: Empirical calibration to the Chilean multi-fund system. Panel (a) reports annual real returns by fund category 2003–2025 with three activation thresholds. Panel (b) reports the empirical complementary cumulative distribution function of activation, $\Pr(R < -\kappa)$, computed as the share of calendar years in which the realized annual real return fell below the threshold. Raw frequencies are weakly increasing in the fund category index, with ties within risk classes due to the limited 23-year sample; smoothed Student-t estimates (reported in Table 3.1) are strictly increasing across the policy-relevant range $\kappa \in [0.05, 0.15]$. The empirical evidence validates Assumption 1 (non-negative expected net growth in non-crash states) and, combined with the coverage rule, the primitive condition of Lemma 1. Source: Superintendencia de Pensiones de Chile, real returns deflated by UF.

We compare four policy environments. The no-insurance benchmark is $\Pi^0 : \alpha(a, z) = 0, \psi(a, z) = 0$. Actuarially fair insurance corresponds to $\psi^F(a, z) = \mathbb{E}_a[I_{a+1}]/y_a$. Uniform social pooling corresponds to the type-average premium in (3.36) below. Hybrid pricing nests both as special cases through the parameter λ.

$$\psi^P = \frac{\sum_m \omega_m \mathbb{E}_m[I_{a+1}]}{\sum_m \omega_m y_m}. \tag{3.36}$$

3.4.3 Baseline parameters and policy scenarios

Table 3.2 reports the baseline numerical values used in the policy diagnostics. The structural parameters of the return process are taken from the Chilean calibration of Table 3.1, anchored on Fund C as the balanced category. The behavioral and policy parameters follow standard values from the life-cycle and decision-support literatures.

The simulation uses CRRA period utility $u(c) = c^{1-\gamma}/(1-\gamma)$ for $\gamma \neq 1$, with subjective discount factor β aligned with the risk-free rate as $\beta(1 + r_f) = 1$. The risk-exposure cost is convex and quadratic, $\Gamma(z) = \eta_z z^2/2$. The life-cycle labor-income profile is a deterministic hump-shaped pattern $y_a = y_{a_0} \exp(g_0(a - a_0) - g_1(a - a_0)^2)$, with g_0 and g_1 calibrated so that wages peak near age 50 and return to the entry-wage level by retirement, consistent with the empirical age-earnings profiles documented in life-cycle studies. Terminal pension wealth W_T is converted into retirement income through the actuarial conversion function $g(W_T, A_T) = W_T/A_T$, where the annuity factor A_T corresponds to the expected discounted present value of one unit of retirement income over the remaining lifespan at retirement age, computed from a unisex mortality table and the risk-free rate. The pension-inadequacy threshold $\bar{p}$ is normalized to the retirement income that the no-insurance benchmark Π^0 delivers at the tenth percentile of the terminal pension-wealth distribution, so that $\bar{p}$ identifies the lower tail of pension outcomes in the benchmark economy and the inadequacy probability $\mathcal{A}(\Pi)$ measures the share of insured workers who fall below that benchmark tail. The implementation-cost wedge is parameterized as a quadratic-loading function $\mathcal{C}_A(\kappa, \psi) = c_0\psi + c_1\psi^2 + c_2\kappa^{-1}$ that combines a linear administrative loading on premia, a small convex penalty on high premia, and an inverse-threshold component that captures the cost of operating the trigger when activations are frequent. The risk-choice diagnostic uses a discrete grid for z that maps directly onto the regulator-defined fund categories of the Chilean system, and the worker selects the grid point that maximizes simulated expected lifetime utility under each pricing rule.

The welfare-frontier exercise evaluates a grid over activation thresholds and premia, with $\kappa \in [0.02, 0.30]$ and $\psi \in [0, 0.04]$. The no-insurance benchmark is represented by $\psi = 0$.

3.4.4 Computation of portfolio-risk choices

The portfolio-risk profile is computed by a reduced-form Monte Carlo procedure disciplined by Propositions 1 and 2, rather than by a fully solved dynamic programming problem on a state-space grid. For each age and pricing regime we approximate the expected residual lifetime payoff of selecting a candidate $z \in \mathcal{Z}$ at age a by forward simulation. We draw M_z trajectories of the market innovation ε_{a+1} and the activation indicator J_{a+1} under the relevant pricing rule, propagate pension wealth through (3.7), propagate liquid wealth through (3.1)–(3.2) with consumption set by the

Table 3.2: Baseline numerical design. Return-process parameters are calibrated to the Chilean Fund C category (Table 3.1). Behavioral and policy parameters follow standard values in the life-cycle and decision-support literatures.

Parameter	Interpretation	Baseline value
a_0	Entry age	25
T	Retirement age	65
W_{a_0}	Initial pension wealth	0
B_{a_0}	Initial liquid wealth	0
y_{a_0}	Entry wage normalization	1
g_0	Wage profile linear slope	0.020
g_1	Wage profile quadratic slope	0.0005
r_f	Risk-free rate	0.02
β	Subjective discount factor ($\beta(1+r_f)=1$)	$1/1.02$
τ	Mandatory pension contribution rate	0.10
κ	Benchmark activation threshold	0.15
α	Constant coverage in distribution exercise	0.70
α_0	Baseline coverage component	0.25
α_{age}	Age-related coverage slope	0.65
α_z	Risk-related coverage penalty	0.30
$\mu(z=0.70)$	Mean pension return, Fund C calibration	0.046
$\sigma_R(z=0.70)$	Pension-return volatility, Fund C calibration	0.068
$\mathbb{E}[\chi \mid J=1]-\kappa$	Conditional excess severity, Fund C calibration	0.090
γ	Relative risk aversion	4.0
η_z	Risk-exposure cost parameter	0.35
$\mathcal{Z}$	Risk grid mapping to Chilean fund categories	$\{0.35, 0.50, 0.70, 0.95, 1.20\}$
$p_J(\kappa=0.15)$	Empirical annual activation probability, Fund C	0.043
A_T	Annuity factor at retirement (unisex mortality, r_f)	16.5
$\bar{p}$	Adequacy threshold (P10 of Π^0 terminal income)	0.232
c_0, c_1, c_2	Implementation-cost coefficients in $\mathcal{C}_A$	$(0.30, 2.0, 0.015)$
ξ	Planner weight on workers close to retirement	0.20
φ_A	Activation and administration cost weight	0.40
$\overline{\mathcal{M}}^F$	Maximum tolerable moral-hazard gap	0.20
ε	Stochastic solvency tolerance	0.05
M_W	Monte Carlo paths, wealth distribution	30,000
M_z	Monte Carlo paths, risk-choice profile	15,000

Euler condition (3.41), accumulate period utility net of the risk-exposure cost $\Gamma(z)$ until retirement, and apply the terminal valuation (3.6). The reduced-form continuation value

$$\widehat{V}^{\Pi}_{a+1}(W_{a+1}, B_{a+1}, y_{a+1}) \;=\; \mathbb{E}\left[\sum_{t=a+1}^{T-1} \beta^{t-a-1}\left[u(c_t) - \Gamma(z_t^{\star,\Pi})\right] \;+\; \beta^{T-a-1}\, U_R(g(W_T, A_T), B_T)\right]$$

is therefore a Monte Carlo estimate of the expected residual lifetime payoff along the simulated path conditional on the rules of motion of policy Π, rather than the value function of a fully solved DP. We do not interpolate over (W, B, y) on a continuous grid; we evaluate $\widehat{V}^{\Pi}_{a+1}$ at the states realized in the forward simulation.

The optimal risk exposure at age a under the reduced-form procedure is the grid maximizer

$$z^{*}_{a+1}(\Pi) \;=\; \arg\max_{z \in \mathcal{Z}} \left\{-\Gamma(z) \;+\; \beta\, \widehat{\mathbb{E}}_a\left[\widehat{V}^{\Pi}_{a+1}(W_{a+1}, B_{a+1}, y_{a+1})\right] \;-\; u'_a\, \psi(a, z; \lambda)\, y_a\right\}. \tag{3.37}$$

The final term is the contemporaneous utility cost of the risk-dependent payroll premium. In the no-insurance benchmark, $\alpha = 0$ and $\psi = 0$. Under fair pricing, the premium varies with the expected indemnity of the chosen portfolio and the marginal-cost term is strictly positive. Under pure pooling, the worker treats the pooled premium as given, so the individual marginal premium cost of risk is zero, in line with Proposition 2. The grid $\mathcal{Z}$ contains five nodes corresponding to the Chilean fund categories A through E; results are not sensitive to grid refinement that subdivides the categories.

We choose the reduced-form Monte Carlo procedure over a fully solved DP for three reasons. First, the qualitative orderings established by Propositions 1 and 2 hold under the analytical conditions of Section 3.3 independently of the numerical method used to evaluate the residual lifetime payoff, so a fully solved DP would not alter the central decision-support results. Second, the focus of the exercise is the admissible region in policy space and the four decision-support outputs of Section 3.5.1, not the optimal worker policy function across the full state space. Third, the reduced-form procedure has a transparent computational cost that scales linearly with the policy grid, while a fully solved DP with continuous state interpolation would require additional structural assumptions (e.g. on the joint distribution of labor income and pension returns) that the framework deliberately avoids. A fully dynamic-programming implementation is left for future research, as noted in Section 3.7.

Algorithm 1 summarizes the full simulation pipeline used to produce all figures and tables. The pseudo-code makes explicit the sequence of operations that links the model equations to the numerical diagnostics.

For each policy package Π, the simulation computes six diagnostic objects: the age profile and distribution of pension wealth, lower-tail percentile lifts, pension-inadequacy probabilities,

Algorithm 1 Reduced-form Monte Carlo simulation pipeline for pension-loss insurance diagnostics

Inputs: parameter set $\Theta = (\gamma, \tau, r_f, \mu, \sigma_R, p_J, \eta_z, \alpha_0, \alpha_{age}, \alpha_z, \varphi_A, \varepsilon, \mathcal{Z})$ calibrated to Chilean Fund C; policy grid $(\kappa, \psi) \in K \times \Psi$; Monte Carlo sizes M_W, M_z.
Initialize pension wealth $W_{a_0} = 0$, liquid wealth $B_{a_0} = 0$, wage $y_{a_0} = 1$.
for each pricing regime $r \in \{\Pi^0, \psi^F, \psi^P, \psi^H(\lambda)\}$ **do**
    for age a from a_0 to $T-1$ (reduced-form Monte Carlo on the risk grid) **do**
        for each candidate $z \in \mathcal{Z}$ **do**
            Draw M_z innovations $\{\varepsilon_{a+1}^{(j)}\}_{j=1}^{M_z}$ and activation indicators $\{J_{a+1}^{(j)}\}$ from (3.33)–(3.34).
            Compute pension-wealth updates $W_{a+1}^{(j)}$ via (3.7) using $\alpha(a, z)$ from (3.35) and the relevant premium $\psi^r(a, z)$.
            Compute liquid-wealth updates $B_{a+1}^{(j)}$ via (3.1)–(3.2) with consumption set by (3.41).
            Forward-simulate the residual life cycle from age $a+1$ to retirement under the rules of motion of Π, and average the resulting residual lifetime payoff $\widehat{V}_{a+1}^r(W_{a+1}^{(j)}, B_{a+1}^{(j)}, y_{a+1})$ over j.
            Form the candidate value $-\Gamma(z) + \beta\widehat{\mathbb{E}}_a[\widehat{V}_{a+1}^r] - u_a'\psi^r(a, z)y_a$ as in (3.37).
        end for
        Set $z_{a+1}^{*,r} = \arg\max_{z \in \mathcal{Z}}$ candidate value.
    end for
end for
Wealth distribution. Using the optimal risk profile $\{z_{a+1}^{*,r}\}_a$, simulate M_W life-cycle paths under each regime to obtain the distribution of terminal pension wealth and its lower-tail percentiles.
Welfare frontier. For each $(\kappa, \psi) \in K \times \Psi$, recompute aggregate mature-regime welfare $\mathcal{W}(\Pi(\kappa, \psi))$ and subtract $\varphi_A \mathcal{C}_A(\kappa, \psi)$ to obtain $\Delta\mathcal{W}^{net}(\kappa, \psi)$ as in (3.38). Locate the welfare-neutral contour and the budget-balancing locus.
Stochastic solvency diagnostics. For each (κ, ψ), simulate the cumulative reserve process $\Phi_t(\Pi)$ from (3.19) and compute the insolvency probability $\Pr(\Phi_{\min} < 0)$ and the prudential reserve $\mathcal{V}_{1-\varepsilon}(\Pi)$ from (3.22).
Moral-hazard diagnostics. Compute $\mathcal{M}^F(\Pi)$ and $\mathcal{M}^0(\Pi)$ from (3.17)–(3.18).
Sensitivity analysis. Repeat steps 3–16 over the one-at-a-time perturbations and the combined adverse stress scenario reported in Section 3.5.5.
Output: figures and tables of Sections 3.5 and 3.5.5.

the insurance-fund balance, the two moral-hazard diagnostics, and the stochastic solvency objects $\Pr(\Phi_{\min} < 0)$ and $\mathcal{V}_{1-\varepsilon}(\Pi)$. Net mature-regime welfare is computed as

$$\Delta\mathcal{W}^{net}(\kappa, \psi) = \mathcal{W}(\Pi(\kappa, \psi)) - \mathcal{W}(\Pi^0) - \varphi_A \mathcal{C}_A(\kappa, \psi), \tag{3.38}$$

where $\mathcal{C}_A(\kappa, \psi)$ is a reduced-form implementation cost. The welfare-neutral frontier is the locus $\{(\kappa, \psi) : \Delta\mathcal{W}^{net}(\kappa, \psi) = 0\}$, and the admissible policy set requires $\Delta\mathcal{W}^{net} \geq 0$ jointly with $\mathcal{B}(\Pi) \geq 1$ and $\Pr(\Phi_{\min} < 0) \leq \varepsilon$.

Table 3.3 maps the equations used to generate each figure and connects each figure to one of the four design concerns stated in the introduction. This mapping is intended to make the simulation traceable.

Table 3.3: Mapping from equations to figures and design concerns

Figure	Equations used	Design concern addressed	Diagnostic role
Fig. 3.2	(3.7), (3.33) with single activation	Age-contingent coverage	Recovery capacity and shock timing (Prop. 1)
Fig. 3.3	(3.7), (3.35)	Tail-risk protection	Lower-tail percentile lift versus mean lift
Fig. 3.4	(3.37) under three pricing rules	Moral hazard under pooling vs. fair pricing	Pricing-discipline and gross insured-risk gaps (Prop. 2)
Fig. 3.5	(3.38) on the (κ, ψ) grid	Welfare-admissible policy designs	Welfare-neutral frontier and budget-balancing locus
Fig. 3.1	Empirical activation frequencies, Table 3.1	Empirical calibration	Validation of Assumption 1 and Lemma 1
Fig. 3.6	(3.19)–(3.22)	Stochastic solvency of the fund	Insolvency probability and prudential reserve over (κ, ψ)

The numerical pension-wealth diagnostics reported in Section 3.5 are *gross of premium payments*. That is, the percentile lifts compare terminal pension wealth across regimes without subtracting premium contributions from disposable income. The welfare-frontier diagnostic in Figure 3.5 is, in contrast, a *net* object. Equation (3.38) subtracts the implementation-cost wedge $\varphi_A \mathcal{C}_A$ and welfare is computed from utility flows that already account for the premium loss in disposable income through (3.1). We restate this distinction in the relevant subsections of Section 3.5 to avoid ambiguity in the interpretation of the numerical magnitudes.

3.5 Results

This section reports the numerical policy diagnostics generated by the framework. The diagnostics should be interpreted as comparative design evidence under the empirical calibration to the Chilean multi-fund system of Section 3.4.2 and the structural parameterization of Table 3.2. The five figures map directly to the four design concerns stated in the introduction. Section 3.5.1 makes explicit the structure of the decision-support outputs that the framework delivers. Figure 3.2 addresses the age-contingent coverage concern by showing the relationship between shock timing and terminal pension loss, which provides the recovery-capacity rationale for differentiated protection across ages. Figures 3.3 and 3.5 address the tail-risk protection and welfare-admissibility concerns by showing where the instrument compresses the lower tail and by identifying the welfare-positive region of the policy space. Figure 3.4 addresses the moral-hazard concern by reporting the optimal pension-portfolio risk profile under the three pricing rules and quantifying the resulting gaps. Figure 3.6 addresses the stochastic solvency concern by reporting insolvency probabilities and prudential reserve requirements over the policy grid. Table 3.4 consolidates the main numerical magnitudes generated by the baseline simulation.

Table 3.4: Summary of baseline policy diagnostics under the Chilean calibration

Diagnostic	Baseline value	Reference
Exogenous-portfolio diagnostics (fixed Fund C in both regimes)		
Mean terminal pension wealth, gross lift vs. Π^0	+4.4%	Fig. 3.3
Tenth-percentile terminal wealth, gross lift vs. Π^0	+13.7%	Fig. 3.3
Shock-timing terminal loss, age-60 vs. age-35 shock	Largest at age 60	Fig. 3.2
Endogenous-portfolio diagnostics (worker re-optimizes risk by pricing regime)		
Mean terminal wealth lift at fair pricing ($\lambda = 1$)	+2.4%	Table 3.5
P10 terminal wealth lift at fair pricing ($\lambda = 1$)	+5.6%	Table 3.5
Pricing-discipline gap $\mathcal{M}^F(\Pi^P)$, age-averaged (mature ages)	Positive	Fig. 3.4
Gross insured-risk gap $\mathcal{M}^0(\Pi^P)$, age-averaged (mature ages)	Positive	Fig. 3.4
Aggregate policy-space diagnostics		
Funding ratio $\mathcal{B}(\Pi)$ along $\psi^{\text{fair}}(\kappa)$	≈ 1	Fig. 3.5
Baseline preferred design (κ^F, ψ^F)	$(0.082, 0.016)$	Fig. 3.5
Net welfare change $\Delta\mathcal{W}^{net}$ at (κ^F, ψ^F)	+0.009	Fig. 3.5
Welfare-positive interval on $\psi^{fair}(\kappa)$	$\kappa \in [0.068, 0.155]$	Fig. 3.5
Insolvency probability at (κ^F, ψ^F)	Below 5% tolerance	Fig. 3.6
Prudential reserve $\mathcal{V}_{0.95}$ at (κ^F, ψ^F)	≈ 1.5 entry wages per worker	Fig. 3.6

3.5.1 Decision-support outputs

Before reporting the numerical diagnostics, we make explicit the structure of the decision-support outputs that the framework delivers to the regulator. The planning problem in (3.23) does not produce a single optimal policy. It produces a structured map from primitives to admissible designs, organized around four objects that can be examined independently and combined to support a specific policy choice.

The first output is the admissible policy region $\mathcal{A}^{\mathrm{adm}}(\Pi)$, defined as the intersection of welfare non-deterioration, expected budget balance, stochastic solvency, and the maximum tolerable moral-hazard gap. A regulator that adopts a policy outside this region cannot defend the choice under the four restrictions jointly. The shape of $\mathcal{A}^{\mathrm{adm}}(\Pi)$ over (κ, ψ) is the central decision object of the framework, and the welfare-neutral frontier $\{(\kappa, \psi) : \Delta\mathcal{W}^{net} = 0\}$ is its outer boundary along the welfare dimension.

The second output is the budget-balancing fair premium schedule $\psi^{\mathrm{fair}}(\kappa)$, defined by the condition that expected premia equal expected indemnities at the population level. The schedule provides a benchmark against which any uniform-pooling or hybrid premium can be evaluated, and it identifies the minimum premium required for the fund to be self-financing in expectation at each activation threshold. Designs with $\psi < \psi^{\mathrm{fair}}(\kappa)$ require either accumulated reserves or an external financing source, while designs with $\psi > \psi^{\mathrm{fair}}(\kappa)$ generate expected surpluses that can be used to build reserves or to lower premia in subsequent periods.

The third output is the prudential reserve $\mathcal{V}_{1-\varepsilon}(\Pi)$ defined in (3.22). This object translates the stochastic solvency condition into a concrete institutional capital requirement. For a regulator, $\mathcal{V}_{1-\varepsilon}(\Pi)$ is the buffer that the fund must hold ex ante so that solvency is preserved in all but a fraction ε of outcomes, and it is the natural diagnostic to compare against existing supervisory frameworks for pension-system contingency funds.

The fourth output is the pricing-discipline ranking summarized by the two moral-hazard diagnostics $\mathcal{M}^F(\Pi)$ and $\mathcal{M}^0(\Pi)$ defined in (3.17)–(3.18). The first measure isolates the loss of pricing discipline relative to actuarially fair pricing. The second measure reports the total increase in insured risk exposure relative to the uninsured economy. Together, they identify the cost of pooling in terms of induced portfolio risk and allow the regulator to set an explicit upper bound on the tolerable moral-hazard gap when choosing the hybrid weight λ.

These four outputs are designed to be inspected jointly. A policy package may be defended on welfare grounds but rejected on solvency grounds, or may be admissible on both grounds but generate excessive pricing-discipline gaps. The framework therefore delivers neither a single optimum nor a ranked list of policies. It delivers a structured space of defensible designs together with the trade-offs that distinguish them, which is the standard format of quantitative public-sector decision-support evidence (Orozco-Villodres et al., 2026; Durães et al., 2026; Ciacci et al., 2026).

Sections 3.5.2 through 3.5.4 report the four outputs in turn, with the empirical calibration described in Section 3.4.2.

3.5.2 Recovery capacity and tail-risk protection

Figure 3.2 isolates the role of shock timing. The dashed line reports the counterfactual pension-wealth path without a crash. The solid lines show pension-wealth paths after a shock occurring at ages 35, 50, and 60. The same type of adverse market event has sharply different terminal consequences depending on when it occurs. This is the numerical counterpart of Proposition 1.

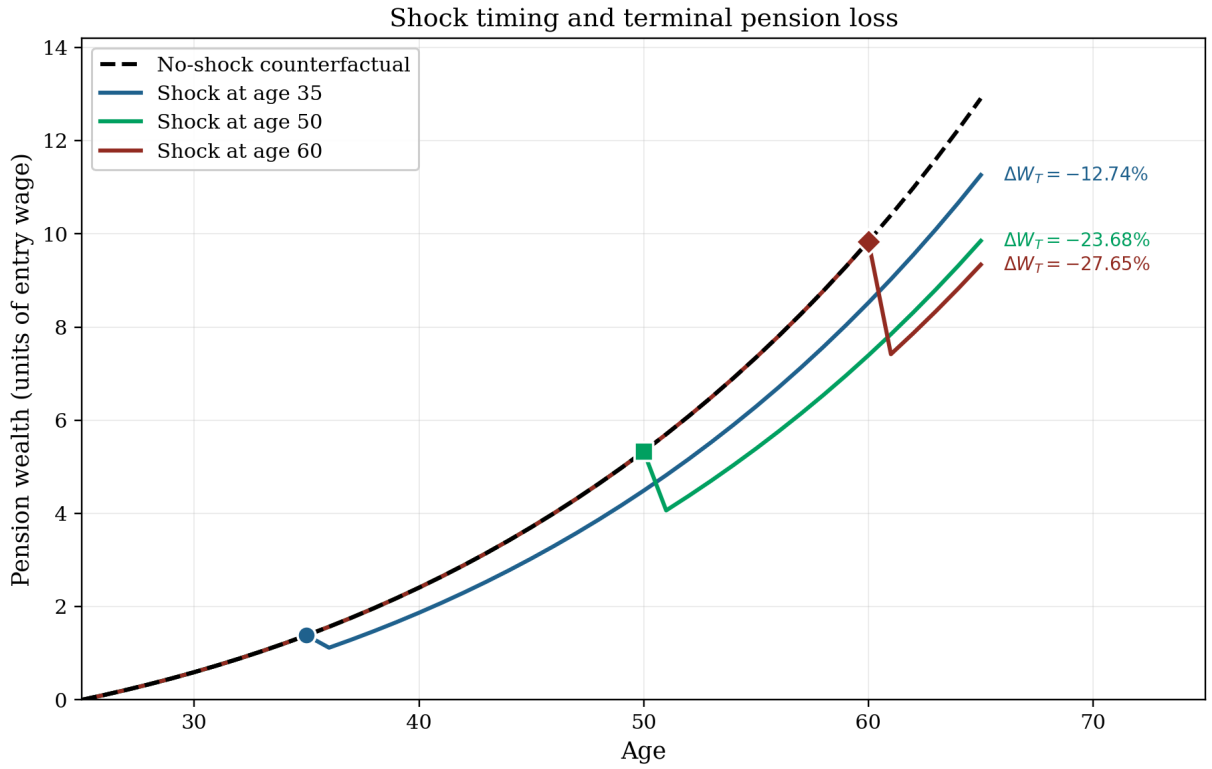


Figure 3.2: Shock timing and terminal pension loss under the baseline parameterization. The dashed line reports the no-shock counterfactual; solid lines report pension-wealth paths after a single activated shock at ages 35 (early career), 50 (mid-career), and 60 (close to retirement). Wealth is expressed in units of the entry wage. The terminal pension loss is largest when the shock occurs close to retirement because workers have less time to rebuild pension wealth through future contributions and compound returns. This is the recovery-capacity rationale formalized in Proposition 1.

The early-career shock reduces pension wealth when the accumulated balance is still small. Its immediate effect is visible, but subsequent contributions and compound returns allow substantial recovery before retirement. The shock close to retirement affects a larger accumulated balance

and leaves little time for rebuilding. The terminal pension loss is therefore largest when the shock occurs near retirement. This finding supports the central design principle of the paper. Pension-loss insurance should not be based only on the probability of a market event. It should also account for the age at which the event occurs and the remaining capacity to recover. A uniform indemnity rule would treat these losses as equivalent financial events, but they are not equivalent pension events.

Figure 3.3 compares pension accumulation with and without insurance, under the same Fund-C portfolio assumed exogenously fixed across both regimes. This is a mechanical-coverage exercise that isolates the protective effect of the insurance layer at a given portfolio risk. The endogenous-portfolio counterpart of these magnitudes is reported in Section 3.5.3, where the worker's optimal pension-portfolio risk responds to the pricing regime; Table 3.5 reports the endogenous-response magnitudes for the full hybrid pricing range. Panel (a) reports the mean and the interquartile range. Panel (b) focuses on the tenth, fifth, and first percentiles. The percentile lifts are computed gross of premium payments. They compare terminal pension wealth across regimes without subtracting premium contributions from disposable income. Under exogenous Fund-C portfolio assumed in both regimes, the gross mean lift at retirement is approximately 4.4 percent, and the gross tenth-percentile lift is approximately 13.7 percent. The net welfare assessment of these lifts is delivered by Figure 3.5 and incorporates the premium cost through equation (3.38).

The asymmetry between the mean effect and the lower-tail effect is the policy-relevant feature. The objective of the insurance layer is not to maximize average pension wealth. Its main value is to reduce severe downside outcomes that generate pension inadequacy. The larger lift in the lower percentiles shows that threshold-activated coverage is targeted toward states where pension losses are most damaging. If the instrument mainly raised mean wealth while leaving the lower tail unchanged, it would be difficult to justify as social insurance. The opposite pattern is observed under the baseline parameters.

3.5.3 Pricing-rule incentives and moral hazard

Figure 3.4 reports the optimal pension-portfolio risk profile under three regimes: no insurance, actuarially fair pricing, and pure social pooling. Risk exposure declines with age in all regimes, reflecting the shorter remaining horizon and the growing importance of accumulated pension wealth. The level of risk-taking differs substantially across pricing rules. This is the numerical illustration of Proposition 2. Section 3.5.5 confirms that the ordering $z^P \geq z^F \geq z^0$ over the mature ages is robust across the parameters varied in the sensitivity analysis.

Pure social pooling generates the highest risk exposure over most of the life cycle. This occurs because the pooled premium weakens the link between individual risk choices and the expected cost imposed on the insurance fund. Workers receive protection against downside losses without

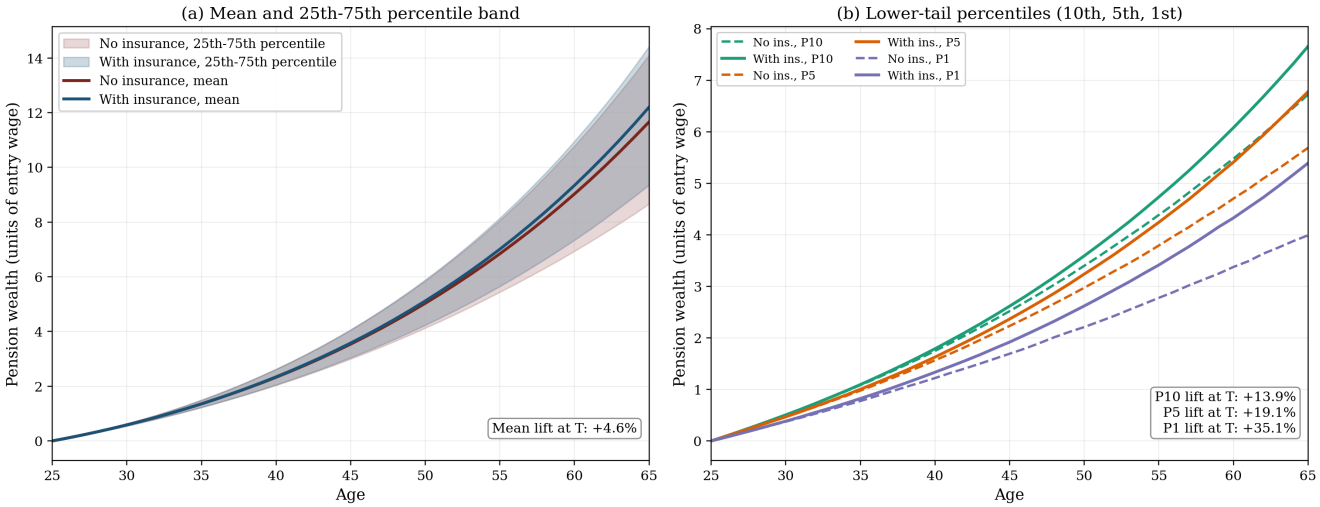


Figure 3.3: Pension-wealth accumulation with and without insurance under the baseline parameterization, holding pension-portfolio risk fixed at Fund C in both regimes. Panel (a) reports the mean and the twenty-fifth to seventy-fifth percentile band. Panel (b) reports the tenth, fifth, and first percentiles. Wealth is expressed in units of the entry wage and reported gross of premium payments. This is a mechanical-coverage diagnostic that isolates the protective effect of the insurance layer at a given portfolio risk; the gross mean lift is approximately +4.4% at age 65 and the gross tenth-percentile lift is approximately +13.7% at age 65. The instrument compresses the lower tail more than it raises the mean, supporting its interpretation as an adequacy-protection layer. The endogenous-portfolio counterpart of these magnitudes, in which the worker re-optimizes pension-portfolio risk in response to the pricing regime, is reported in Section 3.5.3 and Table 3.5.

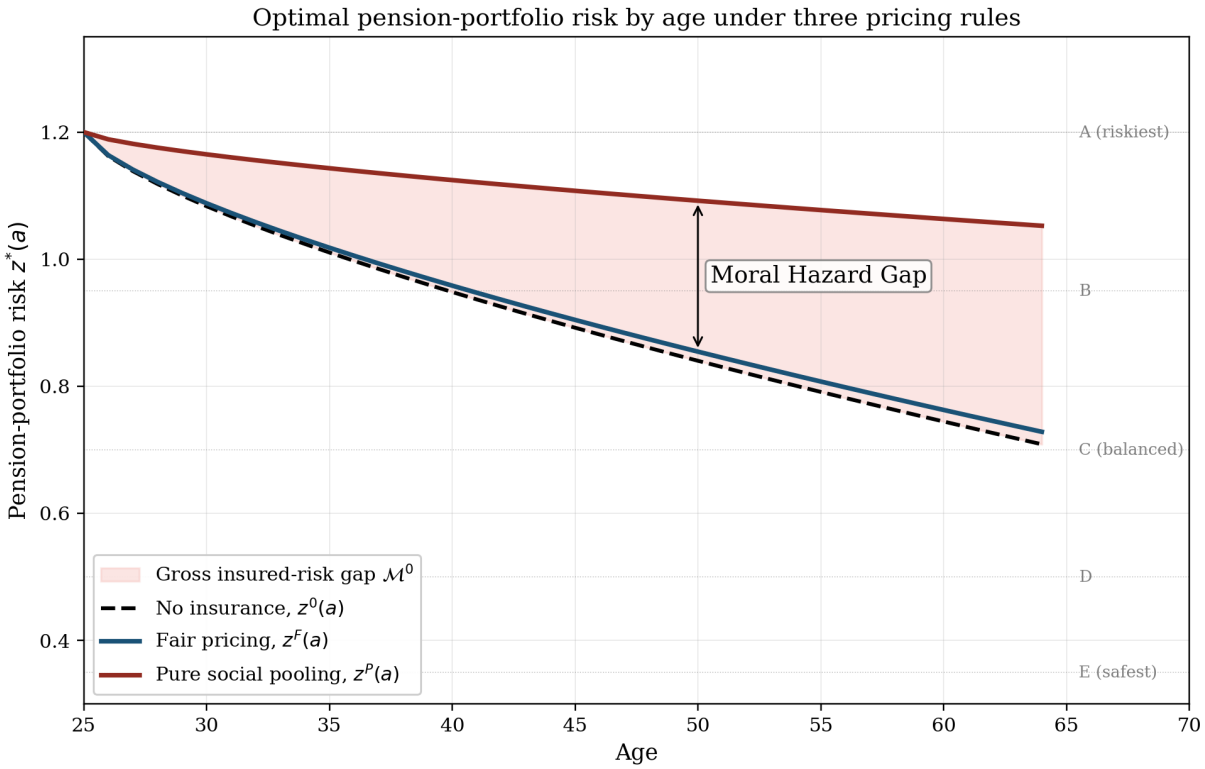


Figure 3.4: Optimal pension-portfolio risk profile by age under three insurance-pricing rules in the baseline parameterization. The dashed line reports the no-insurance benchmark. The blue line reports actuarially fair pricing. The red line reports pure social pooling. The shaded distance between pooling and no insurance is the gross insured-risk gap; the distance between pooling and fair pricing is the pricing-discipline gap. The qualitative ordering of the three curves is robust across the sensitivity analysis in Section 3.5.5, in line with Proposition 2.

fully internalizing the marginal cost of increasing the probability or severity of insured events. Actuarially fair pricing partially corrects this distortion. Since the premium increases with expected indemnities, riskier portfolios face a higher insurance cost. The fair-pricing profile remains below the pure-pooling profile, especially from mid-career onward.

The distinction between the two incentive gaps matters. The shaded gap in Figure 3.4 is the gross insured-risk gap relative to no insurance. The distance between pure pooling and fair pricing is the pricing-discipline gap. Fair premia restore part of the marginal cost of risk, but do not fully replicate the uninsured loss distribution. A pension-loss insurance scheme financed by uniform social pooling may improve solidarity, but it also creates an incentive problem. Hybrid pricing rules are therefore relevant because they preserve some redistribution while maintaining a positive marginal premium cost of risk.

To make the trade-off operational for a regulator, Table 3.5 reports the policy diagnostics across the full hybrid-pricing range at the baseline preferred design $(\kappa^F, \psi^F) = (0.082, 0.016)$. The table maps each value of λ into the equilibrium pension-portfolio risk $z^\star$ implied by Proposition 2, the two moral-hazard diagnostics $\mathcal{M}^F$ and $\mathcal{M}^0$, the gross lift in mean terminal wealth and the gross lift in the tenth percentile, the funding ratio of the insurance fund, the prudential reserve $\mathcal{V}_{0.95}$ at the population-average risk choice induced by λ, and the net welfare change relative to the no-insurance benchmark.

Table 3.5: Hybrid pricing diagnostics at the baseline preferred design $(\kappa^F, \psi^F) = (0.082, 0.016)$ for $\lambda \in \{0, 0.25, 0.50, 0.75, 1\}$. Equilibrium risk $z^\star$ derived from the worker's marginal condition under hybrid pricing weight λ. P10 and mean lifts are gross of premium payments and reported against the no-insurance benchmark. $\Delta\mathcal{W}^{\text{net}}$ is the net welfare change as a fraction of benchmark consumption-equivalent welfare. $\mathcal{V}_{0.95}$ is the prudential reserve at the population-average risk choice induced by λ. The solidarity dispersion $\mathcal{S}^{H/L} = \psi^H(\text{Fund A})/\psi^H(\text{Fund E})$ reports the ratio of the premium paid by the riskiest fund category to the premium paid by the safest fund category; a value of one corresponds to perfect cross-fund solidarity, large values indicate strong risk-based differentiation.

λ	$z^\star$	$\mathcal{M}^F$	$\mathcal{M}^0$	Mean lift (%)	P10 lift (%)	Funding ratio	$\mathcal{V}_{0.95}$ (entry wages)	$\Delta\mathcal{W}^{\text{net}}$ (%)	$\mathcal{S}^{H/L}$
0.00 (pure pooling)	1.050	+0.350	+0.350	+1.4	−36.1	1.030	2.05	+0.5	1.00
0.25	0.898	+0.198	+0.198	+7.9	−11.4	1.022	1.89	+0.6	1.17
0.50	0.819	+0.119	+0.119	+7.5	−2.1	1.015	1.74	+0.8	1.50
0.75	0.755	+0.055	+0.055	+5.8	+3.6	1.008	1.56	+0.9	2.47
1.00 (fair pricing)	0.700	0.000	0.000	+3.5	+7.3	1.000	1.42	+1.0	38.96

Four readings of Table 3.5 are relevant for policy design. First, the moral-hazard gap $\mathcal{M}^F$ falls monotonically from +0.35 to 0 as λ increases from pure pooling to fair pricing, in line with Proposition 2. If the regulator sets the maximum tolerable moral-hazard gap $\overline{\mathcal{M}}^F$ in (3.24) at,

say, 0.10, then designs with $\lambda \leq 0.50$ are excluded from the admissible region, since $\lambda = 0.50$ itself delivers $\mathcal{M}^F = 0.119 > 0.10$. Second, the P10 lift behaves non-monotonically across λ in a policy-relevant way: pure pooling delivers a +1.4% *mean* lift but a −36.1% collapse in the P10, because workers shift to riskier funds in response to detached pricing and the riskier portfolios increase the variability of terminal pension wealth more than the insurance layer compensates for. Fair pricing produces a smaller mean lift but a positive P10 lift of +7.3%, which is the tail-protection effect that justifies the instrument. Third, the prudential reserve $\mathcal{V}_{0.95}$ rises from 1.42 at fair pricing to 2.05 at pure pooling, reflecting that the moral-hazard-induced increase in population-average portfolio risk raises the magnitude of insured claims and therefore the capital required to honor them at the regulator-defined confidence level. The funding ratio stays close to one across the full range, confirming that the choice of λ is primarily an incentive-discipline and tail-protection choice rather than a budget-balance choice. Fourth, the solidarity dispersion $\mathcal{S}^{H/L}$ rises sharply at high λ. The transition from $\lambda = 0.50$ to $\lambda = 1.00$ raises the high-to-low premium ratio from 1.50 to 38.96, a magnitude that may be politically and institutionally difficult to sustain in a pension system where category enrollment is partially driven by default rules and labor-market history rather than by fully informed individual choice.

The pattern across the columns supports the choice of a hybrid pricing rule with λ in an interior range. Pure fair pricing dominates pure pooling on every welfare and incentive dimension reported in the first six columns, but its solidarity dispersion is more than an order of magnitude larger than the dispersion under moderate hybrid weights. Pure pooling delivers perfect cross-fund solidarity but generates a P10 collapse of more than thirty percent. The exact defensible range for λ depends on the regulator-defined moral-hazard ceiling $\overline{\mathcal{M}}^F$ in (3.24). Under the baseline ceiling $\overline{\mathcal{M}}^F = 0.20$, hybrid pricing with $\lambda \in [0.50, 0.75]$ is admissible and preserves welfare and incentive properties close to the fair-pricing optimum while keeping the premium dispersion within a factor of approximately 1.5 to 2.5, comparable to the dispersion observed in existing risk-rated social-insurance schemes. Under a stricter ceiling such as $\overline{\mathcal{M}}^F = 0.10$, the admissible hybrid range shifts toward $\lambda \geq 0.75$, since $\lambda = 0.50$ delivers $\mathcal{M}^F = 0.119 > 0.10$ and is therefore inadmissible. The exact value of λ within the corresponding admissible range is determined by the regulator's relative weighting of incentive discipline and cross-fund solidarity. A regulator that places higher weight on solidarity would shift toward the lower end of this range, while a regulator that places higher weight on pricing discipline would shift toward the upper end.

3.5.4 Admissible region, welfare frontier, and stochastic solvency

Figure 3.5 reports the welfare-neutral policy frontier over the activation threshold κ and the uniform exposure-based premium ψ. The color scale represents the net welfare change relative to the no-insurance benchmark, computed from equation (3.38) and net of the implementation-cost

wedge $\varphi_A \mathcal{C}_A$. Welfare is computed from utility flows that already incorporate the premium loss in disposable income, so this object is fully net of premium payments. The black contour identifies combinations satisfying $\Delta\mathcal{W}^{net} = 0$. The bright green curve reports the budget-balancing fair premium schedule $\psi^{fair}(\kappa)$. The white square denotes the no-insurance benchmark, and the green circle marks the baseline preferred design selected from the self-financing welfare-positive locus.

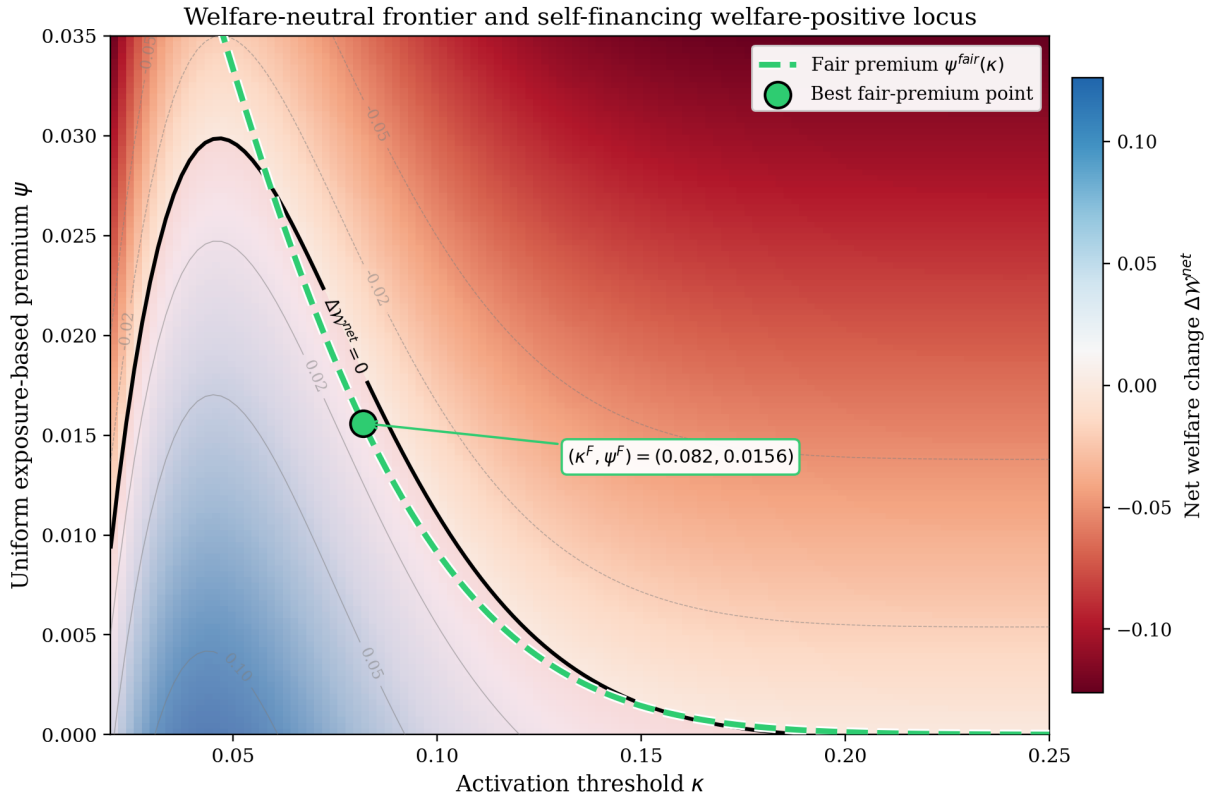


Figure 3.5: Welfare-neutral policy frontier over the activation threshold κ and the uniform exposure-based premium ψ, computed from equation (3.38) and net of the implementation-cost wedge $\varphi_A \mathcal{C}_A$. The color scale reports the net welfare change relative to the no-insurance benchmark. The thick black contour is the welfare-neutral locus $\Delta\mathcal{W}^{net} = 0$, and the welfare-positive region is the blue area inside it. The full admissible region defined in Section 3.3.2 is the intersection of the welfare-positive region with expected budget balance, stochastic solvency, and the moral-hazard ceiling $\overline{\mathcal{M}}^F$, and it is therefore weakly smaller than the blue area shown here. The bright green curve is the budget-balancing fair premium schedule $\psi^{fair}(\kappa)$; the segment of this curve that lies inside the welfare-positive region defines a one-dimensional family of designs that are simultaneously welfare-improving and self-financing in expectation. The green circle marks the baseline preferred design $(\kappa^F, \psi^F) = (0.082, 0.016)$, which is the welfare-maximizing point among all designs in this self-financing family. The white square marks the no-insurance benchmark. The framework explicitly rejects the unconstrained welfare maximum, which lies on the lower-left boundary at $\psi \approx 0$ and a small κ: that point is a degenerate boundary solution that delivers no premium income against a frequently-activated trigger and is therefore not defensible as a policy package.

The frontier shows that not all insurance packages are welfare-improving. High premia reduce welfare because they lower disposable income during working life. Very high activation thresholds provide little protection because the insurance layer is triggered only in extreme states. Very low thresholds broaden coverage but raise expected claims and implementation costs. The welfare-positive region is therefore shaped by a three-way trade-off between protection, financing, and activation costs, and the full admissible region adds expected budget balance, stochastic solvency, and the moral-hazard ceiling on top of welfare non-deterioration.

A central feature of Figure 3.5 is that the fair-premium schedule and the welfare-positive region intersect over a non-trivial range of activation thresholds. Under the Chilean calibration, the budget-balancing fair-premium schedule $\psi^{fair}(\kappa)$ lies inside the welfare-positive region for $\kappa \in [0.068, 0.155]$. Over this interval, a regulator can simultaneously deliver positive net welfare and self-financing in expectation by choosing κ along the green curve, so the framework delivers a one-dimensional family of defensible designs rather than a single optimum.

The unconstrained welfare maximum of the reduced-form landscape lies on the lower-left boundary at $\psi \approx 0$ and a small activation threshold κ, but this is not a defensible policy design for three reasons that the reduced-form objective (3.38) does not fully capture. First, a design with $\psi \approx 0$ generates no premium income, so the fund cannot satisfy expected budget balance (3.14) and cannot accumulate the prudential reserve $\mathcal{V}_{1-\varepsilon}(\Pi)$ required by the stochastic solvency condition (3.21). Second, a small activation threshold triggers the insurance frequently, so the implementation cost grows nonlinearly with the trigger-frequency component of $\mathcal{C}_A(\kappa, \psi)$ in ways that the linear-quadratic-plus-inverse functional form used in the simulation cannot capture in extreme regions. Third, a boundary solution that relies on essentially zero premium income makes the design economically equivalent to a fully public guarantee financed outside the contributory system, which is precisely the policy outcome that the framework is designed to avoid. The lower-left corner of Figure 3.5 is therefore a numerical artifact of the reduced-form welfare landscape, not a policy recommendation.

Once the lower-left corner is excluded, the natural baseline preferred design is the green circle in Figure 3.5, at $(\kappa^F, \psi^F) = (0.082, 0.016)$, with $\Delta\mathcal{W}^{net} = +0.009$. This point is welfare-positive, self-financing in expectation by construction, lies in a region of moderate trigger frequency where the implementation-cost specification is reliable, and preserves the principle that the contingent liability of the system is contained within an explicit prudential reserve rather than absorbed by a full public pension guarantee. It is the welfare-maximizing point among all designs in the self-financing welfare-positive family, and the framework retains it as the baseline reference for the rest of the paper. A regulator that places greater weight on welfare than on strict self-financing in expectation may consider designs below the green curve in the welfare-positive region, at the cost of requiring external financing through one of the three institutional arrangements discussed in

Section 3.6. As Section 3.5.4 makes explicit, expected self-financing is not the same as solvency under a stochastic claims process, and any choice of (κ, ψ) in the admissible region requires the prudential reserve quantified in Figure 3.6.

Figure 3.6 reports the stochastic solvency diagnostics of the insurance fund. The fund dynamics follow (3.19), with annual return moments calibrated to the Chilean Fund C category (mean 4.57%, standard deviation 6.80%) and conditional excess severity calibrated to the historical mean of activated states in the same fund category (8.95% above the activation threshold). The age-averaged coverage and pension exposure follow the baseline parameterization of Table 3.2.

The diagnostic distinguishes two fund-management regimes that correspond to two distinct institutional readings of the solvency condition (3.21). In the *build-from-zero* regime, the fund starts with $\Phi_0 = 0$ and accumulates net inflows over time. The prudential reserve $\mathcal{V}_{0.95}(\Pi)$ defined in (3.22) is then the capital buffer that the fund must hold ex ante so that, were the buildup phase skipped and the fund placed directly in stationary regime, solvency would be preserved at the 95% confidence level. In the *stationary regime*, the fund operates with this prudential reserve already in place, $\Phi_0 = \mathcal{V}_{0.95}(\Pi)$, and the solvency condition (3.21) is evaluated conditional on the pre-existing buffer. The two readings answer different policy questions. The build-from-zero regime quantifies the capital that a regulator must commit to start operating the scheme. The stationary regime quantifies whether the scheme can be sustained once the capital is in place. Panel (a) of Figure 3.6 reports the marginal value of the prudential reserve, by comparing the insolvency probability under three buffer levels ($R_0 = 0$, $R_0 = 0.5\,\mathcal{V}_{0.95}$, and $R_0 = \mathcal{V}_{0.95}$) at the baseline preferred activation threshold $\kappa^F = 0.082$; Panel (b) reports the heatmap of $\mathcal{V}_{0.95}(\Pi)$ over the policy grid, computed under the build-from-zero regime.

Panel (a) shows that the prudential reserve has substantial marginal value as a policy lever. Under build-from-zero, the insolvency probability at the preferred design exceeds 60% across the relevant range of ψ, far above any plausible supervisory tolerance. Holding the buffer at half of the prudential reserve, $R_0 = 0.5\,\mathcal{V}_{0.95}$, reduces the insolvency probability to approximately 20%, still above the 5% tolerance. Setting the buffer to the prudential reserve $R_0 = \mathcal{V}_{0.95}$ brings the insolvency probability to approximately the tolerance level, 4.5 to 6%, across the full ψ range. The three-curve comparison establishes that the reserve quantified in Panel (b) is not a marginal refinement of the design; it is the central object that determines whether the scheme is solvent at the regulator-defined confidence level. The marginal reduction in insolvency probability per unit of additional reserve is largest in the first half of the buildup phase and declines as the buffer approaches $\mathcal{V}_{0.95}$, which is a feature that informs the design of any reserve-building rule discussed in Section 3.6.

Panel (b) reports the prudential reserve $\mathcal{V}_{0.95}(\Pi)$ over the policy grid, in units of the entry wage per worker. The reserve declines rapidly with κ because higher activation thresholds reduce the

frequency of claims, and it depends only mildly on ψ once κ is large enough that activations are rare. At the baseline preferred design $(\kappa^F, \psi^F) = (0.082, 0.016)$, the required reserve is approximately 1.5 units of the entry wage per worker. This is the institutional capital buffer that the fund must hold ex ante so that solvency is preserved in all but the worst five percent of outcomes under the stationary regime. The diagnostic transforms the abstract analytical solvency condition into a concrete capital requirement that can be benchmarked against existing pension-supervisor frameworks.

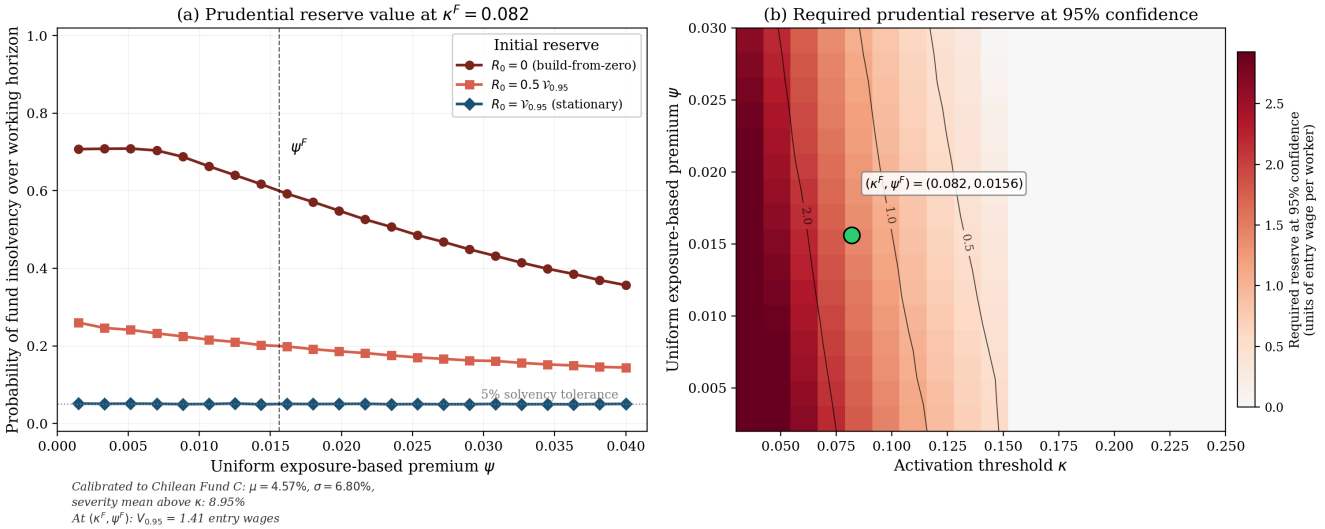


Figure 3.6: Stochastic solvency diagnostics of the insurance fund. Panel (a) reports the marginal value of the prudential reserve at the baseline preferred activation threshold $\kappa^F = 0.082$, by comparing the insolvency probability at three initial-buffer levels: $R_0 = 0$ (build-from-zero), $R_0 = 0.5\,\mathcal{V}_{0.95}$ (half-buffer), and $R_0 = \mathcal{V}_{0.95}$ (stationary regime). The dotted line is the $\varepsilon = 0.05$ solvency tolerance. Panel (b) reports the prudential reserve $\mathcal{V}_{0.95}(\Pi)$ defined in (3.22) over the policy grid, in units of the entry wage per worker, computed under the build-from-zero regime that quantifies the capital required to start operating the scheme. The star in Panel (b) marks the baseline preferred design $(\kappa^F, \psi^F) = (0.082, 0.016)$, at which $\mathcal{V}_{0.95} \approx 1.49$ entry wages per worker. Annual return moments are calibrated to the Chilean Fund C category, with conditional excess severity calibrated to historical activated states. The figure shows that expected budget balance does not by itself imply solvency, and that the prudential reserve provides the institutional translation of the solvency constraint into a capital requirement with substantial marginal value at intermediate buffer levels.

The combined reading of Figures 3.5 and 3.6 is the central decision-support contribution of this section. Welfare non-deterioration, expected budget balance, and stochastic solvency jointly identify the baseline preferred design $(\kappa^F, \psi^F) = (0.082, 0.016)$ as the welfare-maximizing point in the self-financing welfare-positive locus, with $\Delta\mathcal{W}^{net} = +0.009$. Stochastic solvency adds an institutional capital requirement of approximately 1.5 entry wages per worker that the regulator must commit before the scheme starts to operate. Section 3.6 discusses the three institutional arrangements that can deliver this capital in practice: a reserve-building phase that accumulates the required

buffer during non-crash periods, an explicit fiscal backstop, or a reinsurance arrangement that transfers part of the tail liability to a sovereign or supranational counterparty. Each arrangement is consistent with the prudential reserve quantified in Figure 3.6; the choice between them is a separate policy decision that the framework leaves to the regulator.

The numerical diagnostics support five conclusions. First, recovery capacity is a key determinant of pension vulnerability. Shocks close to retirement generate larger terminal losses because workers have less time to rebuild pension wealth. Second, the value of pension-loss insurance is concentrated in the lower tail of the pension-wealth distribution; the instrument is an adequacy-protection mechanism rather than a device for raising mean wealth. Third, pure social pooling generates a moral-hazard gap by weakening the marginal cost of pension-portfolio risk, consistent with Proposition 2. Fourth, the welfare-neutral frontier admits a non-trivial one-dimensional family of self-financing welfare-positive designs along the budget-balancing fair-premium schedule, with the baseline preferred design $(\kappa^F, \psi^F) = (0.082, 0.016)$ selected as the welfare-maximizing point in this family. Fifth, stochastic solvency adds a binding restriction that translates into a quantitative prudential reserve, which under the Chilean calibration is approximately 1.5 entry wages per worker at the baseline preferred design, evaluated under the stationary regime with the reserve pre-funded. The relevant policy question is therefore which losses should be insured, for whom, under what pricing rule, and with what reserve.

3.5.5 Sensitivity analysis and robustness

A potential concern with diagnostic simulations is that the welfare-neutral frontier and the moral-hazard ordering could depend on a particular choice of parameters. This section addresses that concern in two complementary ways. We first vary five quantitatively important parameters one at a time around the baseline of Table 3.2: relative risk aversion γ, the empirical activation probability calibrated to a more or less conservative fund category, the administrative-cost weight φ_A, the risk-coverage penalty α_z, and the age slope of coverage α_{age}. We then evaluate a combined adverse scenario in which several unfavorable conditions are imposed simultaneously, since policy designs should be defended not only under perturbations of a single parameter but also under joint stress. For each variation, Table 3.6 reports the resulting position of the baseline preferred design (κ^F, ψ^F), the gross tenth-percentile lift, the mean lift, the sign of the pricing-discipline gap $\mathcal{M}^F$, and whether the qualitative ordering $z^P \geq z^F \geq z^0$ is preserved over the mature ages.

Three patterns emerge from Table 3.6. The qualitative ordering of the three risk-choice profiles in Figure 3.4 is preserved across all variations, including the combined adverse scenario. This is consistent with Proposition 2, which establishes the ordering as a structural property rather than as an artifact of the baseline parameterization. The pricing-discipline gap $\mathcal{M}^F$ remains positive in all cases, and its magnitude responds in the direction implied by Lemma 1. When the risk-coverage

Table 3.6: Numerical one-at-a-time sensitivity and combined adverse stress around the baseline parameterization of Table 3.2. The P10 (abs) column reports the absolute tenth-percentile of terminal pension wealth under insurance, in units of the entry wage, to capture the level effect; the P10 lift and mean lift columns report relative changes against the no-insurance benchmark under the same calibration; $\Delta\mathcal{W}^{\text{net}}$ reports the net welfare change in consumption-equivalent units. The pricing-discipline gap $\mathcal{M}^F$ is reported in z-grid units. The qualitative ordering $z^P \geq z^F \geq z^0$ over the mature ages is preserved in every scenario by Proposition 2.

Variation	κ^F	ψ^F	P10 (abs)	P10 lift (%)	Mean lift (%)	$\Delta\mathcal{W}^{\text{net}}$ (%)	$\mathcal{M}^F$
Baseline	0.082	0.016	7.96	+7.7	+3.1	+6.5	0.35
$\gamma = 2$ (low)	0.090	0.0140	8.14	+5.2	+2.0	+3.1	0.35
$\gamma = 6$ (high)	0.068	0.0210	7.76	+9.8	+4.3	+11.1	0.35
p_J Fund E calib. (low)	0.110	0.0080	6.75	+0.3	−0.1	+0.1	0.18
p_J Fund A calib. (high)	0.055	0.0260	3.14	+181.2	+75.8	+222.4	0.72
$\varphi_A = 0.20$ (low)	0.082	0.016	7.96	+7.7	+3.1	+6.5	0.35
$\varphi_A = 0.80$ (high)	0.082	0.016	7.96	+7.7	+3.1	+6.5	0.35
$\alpha_z = 0.10$ (low)	0.082	0.016	8.12	+9.8	+4.2	+8.5	0.61
$\alpha_z = 0.60$ (high)	0.082	0.016	7.73	+4.6	+1.5	+3.6	0.25
$\alpha_{age} = 0.30$ (low)	0.082	0.016	7.69	+4.0	+1.4	+3.2	0.35
$\alpha_{age} = 0.90$ (high)	0.082	0.016	8.14	+10.1	+4.3	+8.8	0.35
Combined adverse stress (p_J Fund A, φ_A=0.80, α_z=0.10)	0.055	0.0290	4.14	+271.0	+111.5	+333.8	1.25

penalty α_z is large, coverage itself penalizes risk-taking, so the marginal fair premium gradient $\partial\psi^F/\partial z$ is dampened by the negative coverage-side channel $\partial\alpha/\partial z$, and the equilibrium spread between pooling and fair pricing is smaller. The opposite holds when α_z is low: coverage provides little independent risk-side discipline, so almost all the discipline must come from the premium gradient, and the spread $\mathcal{M}^F$ is correspondingly larger. The numerical magnitudes in Table 3.6 ($\mathcal{M}^F = 0.61$ at $\alpha_z = 0.10$ vs $\mathcal{M}^F = 0.25$ at $\alpha_z = 0.60$) confirm this theoretical prediction.

The position of the baseline preferred design moves in the directions predicted by the model. Higher risk aversion or a higher empirical activation probability (calibrated to riskier fund categories) shifts the design toward lower thresholds and higher premia because the marginal value of tail protection rises. Lower risk aversion or a lower empirical activation probability has the opposite effect. The administrative-cost weight φ_A affects the size of the admissible region but not the qualitative ordering of regimes, and at the preferred design the diagnostics are essentially identical across $\varphi_A \in \{0.20, 0.40, 0.80\}$. The age slope of coverage α_{age} has its largest effect on the lower tail of terminal pension wealth, since a flatter coverage schedule reduces the tenth-percentile lift to +4.0% while a steeper schedule raises it to +10.1%. This confirms that age-contingent coverage is the main channel through which the instrument delivers tail-risk protection.

The combined adverse scenario stresses the design simultaneously along three dimensions, with the empirical activation probability calibrated to the riskiest fund category, a higher administrative cost, and a lower risk-coverage penalty. Three readings of the combined-stress row are policy-relevant. First, the absolute tenth-percentile of terminal pension wealth falls sharply to 4.14 entry wages, against 7.96 at the baseline; the level of pension adequacy is therefore substantially lower in this scenario, and the instrument cannot fully offset the deterioration in the underlying return process. Second, the relative lift over the no-insurance benchmark is very large (P10 lift +271.0%, net welfare change +333.8%) because the no-insurance benchmark itself collapses under the riskier calibration, so the insurance layer becomes proportionally more valuable. Third, the pricing-discipline gap $\mathcal{M}^F = 1.25$ is more than three times its baseline value and exceeds the baseline regulatory ceiling $\overline{\mathcal{M}}^F = 0.20$ in (3.24) by more than sixfold. Under combined adverse stress, the insurance layer remains welfare-improving relative to a severely deteriorated no-insurance benchmark, but it is not admissible under the baseline moral-hazard ceiling. A regulator facing such conditions would have to either widen $\overline{\mathcal{M}}^F$, strengthen the coverage-side discipline through a higher α_z, or shift the design toward a higher λ in the hybrid-pricing rule. The combined-stress result therefore identifies the kind of joint adverse environment in which the baseline ceiling itself becomes binding, and the framework's recommendation is conditional on which margin the regulator is willing to adjust.

These results show that the baseline numerical magnitudes depend on parameter choices, as expected for any simulation framework, but the qualitative implications for policy design are stable. Recovery-capacity heterogeneity supports age-contingent coverage, pension-loss insurance is

mainly a lower-tail protection instrument, pricing rules that detach premia from individual exposure generate a moral-hazard gap, and the welfare-neutral frontier admits a non-trivial one-dimensional family of self-financing welfare-positive designs along the budget-balancing fair-premium schedule. The framework can be used as a decision-support tool for regulators that wish to choose $(\kappa, \psi, \alpha, \lambda)$ subject to welfare non-deterioration, expected budget balance, stochastic solvency, and a maximum tolerable moral-hazard gap, with the sensitivity table providing a transparent reading of how each parameter moves the admissible region.

3.6 Discussion and policy implications

The results show that pension-loss insurance is not a simple actuarial add-on to individual capitalization accounts. It is a planning instrument whose value depends on the interaction between shock timing, recovery capacity, premium design, pension-portfolio risk incentives, and the institutional capacity of the fund to absorb concentrated claim bursts. We now discuss four sets of policy implications that follow from the analytical framework, the empirical calibration, and the sensitivity analysis. We organize them around the design of the activation trigger, the design of the coverage rule, the pricing regime, and the administration of the insurance fund.

A return-based activation trigger is operationally attractive for three reasons. Realized fund returns are observable and verifiable, they reduce regulatory discretion, and they can be defined on aggregate pension-fund indices that are already monitored by the supervisor. These properties are absent for alternatives such as cumulative drawdowns, expected replacement rates, or individual income paths, which are either unobservable in real time or require subjective inputs. A return threshold also separates normal market volatility from severe states in which private recovery is least effective, which is the policy-relevant region.

A return-based trigger does not capture all sources of heterogeneity. Two workers with the same realized return can have very different pension exposures because of differences in accumulated balance, contribution density, age, and portfolio composition. Some of these dimensions are partially absorbed in our framework through the coverage schedule $\alpha(a, z)$, which depends on age and on the regulator-classified portfolio category in which the worker is enrolled. Others, such as contribution density or informality, are not absorbed by a return-based trigger and would require complementary instruments. Whether such complementary instruments should also be integrated into the insurance layer or handled separately by other components of the pension system is a design choice that the framework leaves open and that future calibration work could address.

The age component of the coverage rule is a direct consequence of recovery capacity. Workers close to retirement face a structural disadvantage in absorbing severe losses, and the simulations show that this disadvantage translates into a quantitatively large terminal pension loss for shocks at

age 60 relative to shocks at age 35. The risk component of the coverage rule is the implementation tool that prevents the age component from creating perverse incentives. If coverage rose with age but did not respond to portfolio risk, workers close to retirement might choose riskier portfolios precisely when their balances are large.

From a regulatory standpoint, the dependence of coverage on z should not be read as a continuous individual risk measurement. In an individual capitalization system with regulator-defined fund categories, such as the Chilean multi-fund system that anchors our empirical calibration, the natural implementation is to make coverage depend on the category in which the worker is enrolled. The risk parameter z then maps into a finite list of fund categories, and the rule $\alpha(a, z)$ becomes a discrete matrix of coverage levels indexed by age band and fund category. This is implementable with existing regulatory infrastructure and is preferable to a fully individualized risk classification, which would be informationally demanding and potentially discriminatory. The insurance layer can also be combined with a target-date or life-cycle default rule that mechanically reduces portfolio risk as the participant approaches retirement (Poterba et al., 2009a,b; Suárez et al., 2025), in which case the two instruments complement rather than substitute for each other.

The pricing regime is where the discussion of solidarity and incentive discipline becomes operational. Actuarially fair pricing preserves the marginal cost of risk and is therefore attractive on incentive grounds, but its strict individual implementation may be informationally demanding because it requires expected indemnities computed at the level of each worker. A realistic implementation would compute fair premia for observable categories defined by age band, fund category, and possibly contribution-density status, in line with the same regulatory discretization used for the coverage rule. Hybrid pricing then becomes a policy compromise that uses the same category structure but blends it with a uniform component to support redistribution toward workers with low recovery capacity.

The results of Sections 3.5 and 3.5.5 make a precise design recommendation. Pure pooling weakens the marginal cost of risk and creates a moral-hazard gap that the sensitivity analysis confirms is robust. Pure fair pricing limits this distortion but offers less redistribution. Hybrid pricing with a moderate λ is therefore a natural baseline for implementation. The exact value of λ that a regulator should choose depends on the relative weight assigned to solidarity and incentive discipline, and on the maximum tolerable moral-hazard gap, which becomes an explicit decision variable in the planning problem.

Beyond the choice of pricing rule, the framework also imposes an institutional capital requirement on the fund. Section 3.3.2 formalizes the solvency requirement as a fourth restriction of the planner's problem, and Section 3.5.4 reports the prudential reserve required at the baseline preferred design, approximately 1.5 entry wages per worker under the Chilean Fund C calibration. Three institutional arrangements can deliver this capital in practice. The first is a pure reserve-building rule, in which

the fund accumulates net contributions during non-crash periods until $\mathcal{V}_{1-\varepsilon}(\Pi)$ is reached, and resumes accumulation after each activated state. This arrangement preserves the principle that individual capitalization is not transformed into a full public pension guarantee, but requires larger steady-state premia during the buildup phase, or equivalently a long buildup horizon that defers the protection of cohorts close to retirement at the time the scheme is introduced. The second is an explicit fiscal backstop activated only when realized reserves are exhausted, which lowers steady-state premia but reintroduces public exposure to tail risk. The third is a reinsurance arrangement that transfers a fraction of the tail liability to a sovereign or supranational counterparty, in line with the alternative risk transfer mechanisms discussed by Boyer (2009) in the context of catastrophic pension-fund losses. Under all three arrangements, the welfare non-deterioration constraint of our framework continues to apply once the additional cost of the financing arrangement is incorporated into $\varphi_A \mathcal{C}_A$, and the prudential reserve $\mathcal{V}_{1-\varepsilon}(\Pi)$ remains the central object that translates the stochastic solvency condition into an institutional capital requirement. The numerical magnitude of 1.5 entry wages per worker is informative because it locates the capital requirement at a scale comparable to existing contingency funds in funded pension systems, rather than at a scale that would require a structural fiscal commitment.

The broader implication is that individual capitalization systems do not eliminate the need for social insurance. They change the form of the social-insurance problem. Instead of pooling all pension liabilities through a pay-as-you-go formula, policymakers can target specific risks that individual accounts handle poorly. Rare market shocks close to retirement are one such risk. The framework provides a way to evaluate this protection while preserving the account-based structure, maintaining fiscal discipline, and limiting moral hazard, in the same diagnostic spirit applied to other multidimensional public-sector decisions (Orozco-Villodres et al., 2026; Durães et al., 2026; Ciacci et al., 2026).

3.7 Conclusions, limitations, and future research

This paper has built a welfare-neutral decision-support framework for pension-loss insurance in individual capitalization systems exposed to rare market shocks. The central economic argument is that recovery capacity declines with age, so a given financial loss maps into very different pension-equivalent losses depending on when it occurs. This life-cycle asymmetry creates a precise rationale for targeted protection against lower-tail pension losses, and it does so without requiring a full pay-as-you-go restructuring of the system.

The contribution is threefold. First, we formalize the mapping from rare capital-market shocks to pension-equivalent losses, and Proposition 1 shows that the expected terminal damage is non-decreasing in the age at which the shock occurs under standard non-negative expected

accumulation conditions. Second, we develop a welfare-neutral planner problem that incorporates actuarial pricing, social pooling, hybrid premia, budget balance, stochastic solvency, incentive compatibility, and minimum implementation scale; Lemma 1 states the primitive condition on the activation-effective expected loss rate, and Proposition 2 shows that under this condition uniform pooling induces strictly higher equilibrium pension-portfolio risk than actuarially fair pricing. Third, we provide numerical decision-support diagnostics under an empirical calibration to the Chilean multi-fund system, complemented by a sensitivity analysis that includes a combined adverse stress scenario and confirms the stability of the main mechanisms, while showing that adverse return-risk calibrations tighten the incentive constraint and can render the design inadmissible under the baseline moral-hazard ceiling.

The numerical diagnostics generate five findings. Shock timing is central because losses close to retirement produce larger terminal pension effects. The insurance layer mainly improves the lower tail of terminal pension wealth, which supports its interpretation as an adequacy-protection instrument rather than a device for raising mean wealth. Pure social pooling creates a moral-hazard gap by weakening the link between portfolio risk and expected insurance costs. The welfare-neutral frontier admits a non-trivial one-dimensional family of self-financing welfare-positive designs along the budget-balancing fair-premium schedule, with the baseline preferred design $(\kappa^F, \psi^F) = (0.082, 0.016)$ selected as the welfare-maximizing point in this family. Stochastic solvency adds a quantitative prudential reserve of approximately 1.5 entry wages per worker at this design, evaluated under a stationary regime in which the reserve is pre-funded.

The analysis has limitations that we organize in three groups. The first group is related to the scope of the empirical calibration. The numerical exercise is calibrated to the Chilean multi-fund system, which is the longest-running and most extensively documented individual-capitalization arrangement, but other systems may exhibit different return moments, activation frequencies, and severity distributions. A natural extension is comparative calibration across multiple individual-capitalization systems using administrative pension records, contribution histories, fund returns, and retirement-income outcomes. The second group is related to institutional realism. The model abstracts from contribution density, labor-market informality, fund-switching restrictions, default investment rules, annuitization frictions, heterogeneous mortality, and the effect of early withdrawals on accumulated balances. The third group is related to general-equilibrium and political-economy considerations. The framework does not incorporate general-equilibrium effects on wages, asset prices, or public debt, and it does not explicitly model the political-economy constraints likely to affect the feasibility of any redistributive or partially pooled pension instrument.

Future research should move in four directions. The first is comparative calibration across additional individual-capitalization systems. The second is the integration of contribution-density risk and labor-market informality. The third is a fully dynamic-programming implementation of the

worker problem with continuous state interpolation, which would refine the policy-function objects underlying the reduced-form simulation. The fourth is empirical estimation of behavioral responses to insured pension risk, in particular fund switching and risk-taking under different premium rules.

The broader message is that pension-policy design should move beyond the binary comparison between individual capitalization accounts and public guarantees. Individual accounts allocate some risks well and others poorly. The relevant policy question is which risks belong to each category, and how the second category can be insured without surrendering the discipline of the first. Rare market shocks close to retirement belong to the second category, and the framework developed here shows that they can be addressed through a narrowly targeted, threshold-activated, age-contingent, and incentive-compatible insurance layer, with an explicit prudential reserve that translates the stochastic solvency condition into an institutional capital benchmark. We view this as a defensible direction for future policy experimentation in funded pension systems.

3.8 Proofs of propositions

This appendix collects the proofs of Lemma 1, Proposition 1, Proposition 2, and Corollary 1, together with the path-wise corollary of Proposition 1.

Proof of Lemma 1

From (3.9), the actuarially fair premium factorizes as

$$\psi^F(a,z) \;=\; \frac{W_a + \tau y_a}{y_a}\,\big[\alpha(a,z)\,\ell(a,z;\kappa)\big].$$

The leading factor $(W_a + \tau y_a)/y_a$ is strictly positive and independent of z in a partial-equilibrium reading where the worker is small relative to the population and takes the aggregate exposure as given. Differentiating with respect to z,

$$\frac{\partial \psi^F(a,z)}{\partial z} \;=\; \frac{W_a + \tau y_a}{y_a}\,\frac{\partial}{\partial z}[\alpha(a,z)\,\ell(a,z;\kappa)]\,.$$

Under (3.27), the right-hand side is strictly positive, which proves (3.28). The pooled and hybrid identities follow directly from (3.36) and (3.10), since ψ^P is taken as given by the individual worker and is therefore independent of her own z. This completes the proof. □

A useful sufficient condition for (3.27) when the coverage rule satisfies $\partial\alpha/\partial z \le 0$ is that the

loss-rate elasticity dominates the coverage elasticity in absolute terms, that is,

$$\frac{\partial \ln \ell(a,z;\kappa)}{\partial z} \;>\; -\frac{\partial \ln \alpha(a,z)}{\partial z}.$$

This sufficient condition is satisfied by the baseline parameterization of Section 3.4.3 together with the empirical activation calibration of Section 3.4.2, and it provides a transparent test that a regulator can apply to verify Lemma 1 in a specific implementation.

Proof of Proposition 1

Fix coverage $\alpha \in [0,1)$ and severity $\chi > 0$, and let $G_{t+1} = 1 + \widetilde{R}^{p}_{t+1}(z_{t+1})$ denote the non-crash growth factor at age t. Under Assumption 1,

$$g_{t+1} \;\equiv\; \mathbb{E}[G_{t+1} \,|\, J_{t+1} = 0] \;\geq\; 1 \qquad \text{for every } t < T.$$

Conditional on a single activated shock at age s and no further activated shock thereafter, the pension-wealth growth factor in period $s+1$ is $G_{s+1} - (1-\alpha)\chi$ and equals G_{t+1} in every other period. Standard recursion on (3.7) gives

$$L_T^{(s)} \;=\; (1-\alpha)\chi\,(W_s + \tau y_s) \prod_{t=s+1}^{T-1} G_{t+1}.$$

Taking expectations conditional on the single-shock event and using independence of returns across non-crash periods,

$$\mathbb{E}[L_T^{(s)}] \;=\; (1-\alpha)\chi\,\mathbb{E}[W_s + \tau y_s] \prod_{t=s+1}^{T-1} g_{t+1}. \tag{3.39}$$

For $s_1 < s_2$, recursion of (3.7) from age s_1 to age s_2 under no shock gives

$$\mathbb{E}[W_{s_2} + \tau y_{s_2}] \;=\; \mathbb{E}[W_{s_1} + \tau y_{s_1}] \prod_{t=s_1+1}^{s_2} g_{t+1} \;+\; R_{s_1,s_2},$$

where $R_{s_1,s_2} \geq 0$ is the expected contribution flow accumulated between ages s_1+1 and s_2, compounded forward by the non-crash growth factors. Substituting into (3.39),

$$\mathbb{E}[L_T^{(s_2)}] - \mathbb{E}[L_T^{(s_1)}] \;=\; (1-\alpha)\chi\, R_{s_1,s_2} \prod_{t=s_2+1}^{T-1} g_{t+1} \;\geq\; 0,$$

which is the monotonicity claim. The inequality is strict whenever $R_{s_1,s_2} > 0$ or $\prod_{t=s_2+1}^{T-1} g_{t+1} > 1$,

which is the strict-recovery clause stated in the proposition. □

Corollary 2 (Path-wise version of Proposition 1)**.** *Suppose along a single sample path that contributions are strictly positive at every age and non-crash returns are non-negative realization by realization. Then $L_T^{(s)}$ is non-decreasing in s on that path, and strictly increasing in s whenever the realized non-crash growth factor strictly exceeds one in at least one period in $\{s+1,\dots,T-1\}$.*

The path-wise corollary is strictly stronger than the expected version above but its sufficient conditions are restrictive in practice; the empirical evidence in Table 3.1 shows that realized non-crash returns can be negative even when the conditional mean is positive.

Proof of Proposition 2

Under Lemma 1 and Assumption 2, the interior risk choice is characterized by $H_a(z^*(a;\lambda);\lambda) = 0$. The implicit function theorem applies because H_a is continuously differentiable in (z,λ) and $\partial H_a/\partial z > 0$. Differentiating implicitly,

$$\frac{dz^*(a;\lambda)}{d\lambda} \;=\; -\frac{\partial H_a/\partial \lambda}{\partial H_a/\partial z}.$$

The denominator is strictly positive. The numerator equals $y_a u'(c_a)\,\partial\psi^F(a,z)/\partial z$ by Assumption 2, which is strictly positive by Lemma 1 and strict monotonicity of period utility. Therefore,

$$\frac{dz^*(a;\lambda)}{d\lambda} \;<\; 0 \qquad \text{on the relevant interior range.}$$

Evaluating at the two polar regimes,

$$z^P(a) \;=\; z^*(a;0) \;>\; z^*(a;1) \;=\; z^F(a).$$

The strict inequality is a consequence of the primitive condition (3.27); the result is therefore not a tautological restatement of single-crossing but a direct implication of the economic content that riskier portfolios generate larger expected insured losses. This monotone-comparative-statics structure follows the framework of Milgrom and Shannon (1994). □

Proof of Corollary 1

By Proposition 2, $z_m^P - z_m^F > 0$ for every type m with $\omega_m > 0$ and interior risk choice. Since population weights are non-negative and at least one type has positive weight by the maturity convention,

$$\mathcal{M}^F(\Pi^P) \;=\; \sum_m \omega_m \big[z_m^P - z_m^F\big] \;>\; 0.$$

This completes the proof. □

3.9 Additional derivations and equilibrium conditions

This appendix provides the additional derivations underlying the analytical framework. The purpose is to make explicit the worker's optimality conditions, the actuarial pricing formulas, the budget-balance restriction, and the mature-regime welfare object.

3.9.1 Worker's Bellman equation

The worker solves (3.5) subject to (3.1), (3.2), and the pension-wealth law of motion (3.7). The feasible set requires $c_a \geq 0$, $b_{a+1} \geq \underline{b}$ for some borrowing limit $\underline{b}$, and $z_{a+1} \in \mathcal{Z}$.

3.9.2 Euler equation for liquid saving

Since $B_{a+1} = (1 + r_f)b_{a+1}$, the first-order condition for liquid saving is

$$u'(c_a) = \beta(1 + r_f)\mathbb{E}_a \left[V_{B,a+1}(W_{a+1}, B_{a+1}, y_{a+1})\right]. \tag{3.40}$$

The envelope condition with respect to liquid wealth is $V_{B,a}(W_a, B_a, y_a) = u'(c_a)$. Combining these conditions yields the standard Euler equation

$$u'(c_a) = \beta(1 + r_f)\mathbb{E}_a \left[u'(c_{a+1})\right], \tag{3.41}$$

with the appropriate inequality when the borrowing constraint binds.

3.9.3 Envelope condition for pension wealth and risk choice

The first-order conditions below are presented using the latent decomposition $R^{\text{obs}}_{a+1}(z_{a+1}) = \widetilde{R}^{p}_{a+1}(z_{a+1}) - J_{a+1}\chi_{a+1}(z_{a+1})$ of Section 3.3.1, which provides smooth analytical expressions for the derivatives. The contract law of motion (3.7) of the main text is expressed in observable form $W_{a+1} = (W_a + \tau y_a)(1 + R^{\text{obs}}_{a+1}) + \alpha(a, z_{a+1})L_{a+1}$, but introduces the indicator $J_{a+1} = \mathbf{1}\{R^{\text{obs}}_{a+1} < -\kappa\}$ and the operator $[\cdot]^+$, both of which are non-differentiable in z. The latent-representation derivatives below should therefore be read as the smooth interior expressions whose qualitative content (sign of marginal cost of risk, role of $\alpha(a, z)$, and disciplining channel of risk-based pricing) is identical to that of the observable formulation. The numerical implementation in Section 3.4 works on a discrete grid for z and on the observable loss (3.4), so the absence of differentiability at the trigger boundary does not affect the operational diagnostics.

The envelope condition for pension wealth is

$$V_{W,a}(W_a, B_a, y_a) = \beta \mathbb{E}_a \left[V_{W,a+1}(W_{a+1}, B_{a+1}, y_{a+1}) \frac{\partial W_{a+1}}{\partial W_a} \right], \tag{3.42}$$

where $\partial W_{a+1}/\partial W_a = 1 + \widetilde{R}^p_{a+1}(z_{a+1}) - (1 - \alpha(a, z_{a+1}))J_{a+1}\chi_{a+1}(z_{a+1})$. Higher coverage raises the marginal value of pension wealth in crash states because it reduces the uninsured loss component.

For an interior choice of pension-portfolio risk z_{a+1} when the premium $\psi(a, z)$ depends on risk, the first-order condition is

$$\Gamma'(z_{a+1}) + u'(c_a) y_a \frac{\partial \psi(a, z_{a+1})}{\partial z_{a+1}} = \beta \mathbb{E}_a \left[V_{W,a+1} \frac{\partial W_{a+1}}{\partial z_{a+1}} \right]. \tag{3.43}$$

The second term on the left-hand side is the marginal utility cost of risk-based pricing. Under actuarially fair pricing, this term is positive when riskier portfolios increase expected insured losses, thereby discouraging risk-taking. Under uniform pooling, $\partial\psi/\partial z = 0$, and this incentive channel disappears. This is the analytical basis of Proposition 2.

The derivative of pension wealth with respect to risk exposure is

$$\begin{aligned} \frac{\partial W_{a+1}}{\partial z_{a+1}} = (W_a + \tau y_a) \Bigg[& \frac{\partial \widetilde{R}^p_{a+1}(z_{a+1})}{\partial z_{a+1}} - (1 - \alpha(a, z_{a+1}))\, J_{a+1} \frac{\partial \chi_{a+1}(z_{a+1})}{\partial z_{a+1}} \\ & + \frac{\partial \alpha(a, z_{a+1})}{\partial z_{a+1}} J_{a+1} \chi_{a+1}(z_{a+1}) \Bigg]. \end{aligned} \tag{3.44}$$

The first term captures the private return from increasing risk exposure. The second term captures the larger uninsured loss generated by riskier portfolios in crash states. The third term captures the incentive effect of risk-contingent coverage. If $\partial\alpha(a, z)/\partial z < 0$, coverage declines with risk and the contract discourages excessive risk-taking. If coverage does not depend on z, this disciplining channel is absent.

3.9.4 Actuarial pricing, pooling, and hybrid premia

The expected indemnity is $\mathbb{E}_a[I_{a+1}] = \alpha(a, z)(W_a + \tau y_a)\ell(a, z; \kappa)$, with

$$\ell(a, z; \kappa) = \mathbb{E}\Big[J_{a+1} \left[-\kappa - R^{\text{obs}}_{a+1}(z) \right]^+ \Big| \, a, z, \kappa \Big] = \Pr\big(R^{\text{obs}}_{a+1}(z) < -\kappa \big)\ \mathbb{E}\Big[-\kappa - R^{\text{obs}}_{a+1}(z) \Big| \, R^{\text{obs}}_{a+1}(z) < -\kappa \Big], \tag{3.45}$$

which is the expected eligible loss per unit of insured exposure, computed entirely on the observable return $R^{\text{obs}}_{a+1}(z)$. Under the latent decomposition of Section 3.3.1, the activation event is $\{\widetilde{R}^p_{a+1}(z) - \chi_{a+1}(z) < -\kappa\}$ and the conditional shortfall is $\mathbb{E}[\chi_{a+1}(z) - (\kappa + \widetilde{R}^p_{a+1}(z))^+ \mid \text{activation}]$, but neither object is needed at the contract level. The fair premium satisfies $\psi^F(a, z) y_a = \mathbb{E}_a[I_{a+1}]$, the pooled

premium is given by (3.36), and the hybrid premium is the convex combination in (3.10). The marginal premium cost of risk under hybrid pricing is

$$\frac{\partial \psi^H(a, z; \lambda)}{\partial z} = \lambda \frac{\partial \psi^F(a, z)}{\partial z}, \tag{3.46}$$

because $\partial \psi^P / \partial z = 0$ for an individual worker who takes the pooled premium as given. This identity is used in the proof of Proposition 2.

3.9.5 Budget balance and mature-regime welfare

Let $\mathcal{R}(\Pi) = \sum_m \omega_m \mathbb{E}_m[\psi(a_m, z_m; \lambda) y_m]$ and $\mathcal{P}(\Pi) = \sum_m \omega_m \mathbb{E}_m[I_{a+1}]$. Expected budget balance requires $\mathcal{R}(\Pi) \geq \mathcal{P}(\Pi)$, and the funding ratio is $\mathcal{B}(\Pi) = \mathcal{R}(\Pi)/\mathcal{P}(\Pi)$. Mature-regime welfare is given by (3.12). Welfare non-deterioration requires $\mathcal{W}(\Pi) \geq \mathcal{W}(\Pi^0)$, which is strictly weaker than exact equality of welfare across regimes.

3.9.6 Tail-risk and moral-hazard diagnostics

The lower-tail terminal pension wealth $\mathcal{T}_q(\Pi) = Q_q^{\Pi}(W_T)$ and the inadequacy probability $\mathcal{A}(\Pi) = \Pr_{\Pi}(p_T < \bar{p})$ are computed from the simulated terminal distribution of pension wealth. The moral-hazard diagnostics $\mathcal{M}^F(\Pi)$ and $\mathcal{M}^0(\Pi)$ are computed from the optimal risk choices under the relevant pricing regimes. A positive value of either measure indicates that the policy induces higher risk-taking than the corresponding benchmark.

Chapter 4

Strategic Competition in Retirement-Income Markets: A Bayes–Nash Model of Annuity Auctions with Evidence from Chile

4.1 Introduction

Population ageing has turned the design of retirement-income systems into a central topic in applied economics. Longer life expectancy, lower fertility, and more fragmented labor-market trajectories affect the financing of pension systems, the distribution of retirement resources, and the way accumulated wealth is converted into old-age income. A broad quantitative literature studies these issues through overlapping-generations models, life-cycle frameworks, and policy evaluation exercises, showing that demographic change has first-order implications for growth, fiscal sustainability, labor supply, and welfare (Lee, 2003; Rechel et al., 2013; Fougère and Mérette, 1999; Jimeno et al., 2008; Lisenkova et al., 2013; Bielecki et al., 2015; Gustafsson, 2023; Kudrna et al., 2025). In defined-contribution systems, however, the economics of retirement is not exhausted by the accumulation phase. Welfare in old age also depends on the institutions governing the payout stage and on the market structure through which retirement balances are transformed into income streams.

At the payout stage, retirees typically face a choice between flexible withdrawal schemes and some form of life annuitization. Standard theory assigns a central role to annuities because they insure longevity risk, but observed annuitization patterns also reflect bequest motives, framing, financial capability, product complexity, intermediary influence, and beliefs about firms and regulation

(Brown et al., 2008; Lockwood, 2012; Gustafsson, 2023; Chen et al., 2023; Bütler and Ramsden, 2024; Bello et al., 2024; Balter et al., 2021; Berstein and Morales, 2021). Recent evidence also shows that retirees value contract characteristics beyond the monthly payout, including guarantees, firm identity, distribution channels, and advisory services (Alcalde and Vial, 2022; Bello, 2023; Searle et al., 2024; Charupat and Kamstra, 2023). This implies that retirement-income outcomes cannot be understood solely through actuarially fair pricing or household preferences taken in isolation. They also depend on the way the supply side is organized and on the rules under which competing products are presented and awarded.

From that perspective, retirement-income markets are naturally connected to market design, industrial organization, and auction theory. When insurers compete for retirees' balances, quotes reflect expected profitability, beliefs about rival behavior, private information, participation incentives, and institutional rules defining how offers are compared and accepted. The relevant economic object is therefore not only the annuity contract itself, but also the mechanism through which insurers compete. This point is economically important because small differences in pricing can generate persistent differences in lifetime retirement income, especially in markets where the payout decision is largely irreversible. Yet the literature remains relatively sparse on centralized decumulation markets in which heterogeneous insurers compete under incomplete information and where public policy affects prices, entry, and welfare simultaneously (Myerson, 1981; Milgrom and Weber, 1982; Maskin and Riley, 2000; Fang and Wu, 2019; Chen et al., 2023; Verani and Yu, 2024).

The Chilean annuity market provides a useful application for studying this broader problem. Chile has long been one of the most visible defined-contribution pension systems, and its experience has been central to debates on pension adequacy, private provision, capital-market development, and the long-run consequences of pension reform (Edwards, 1996; Soto, 2007; Madeira, 2021a,b; Bayar et al., 2022; López and Walker, 2021). At the payout stage, the country operates a centralized electronic platform, the *Sistema de Consultas y Ofertas de Montos de Pensión* (SCOMP), through which retirees receive and compare competing annuity offers. Existing evidence shows that SCOMP reduced frictions and improved market outcomes, while subsequent work documents that consumer preferences, firm reputation, intermediation, and transparency effects continue to shape the equilibrium allocation of retirement products (Reyes, 2010; Morales and Larraín, 2017; Alcalde and Vial, 2022; Bello, 2023; Fajnzylber et al., 2023). Chile is therefore useful not because it is the only case that matters, but because it offers a particularly transparent institutional setting in which a more general question can be studied, namely how insurer competition works in a centralized retirement-income market.

This paper studies that competition as a centralized first-price sealed-bid auction with heterogeneous insurers and incomplete information. Firms differ in costs, profitability conditions, and participation thresholds, and they submit offers on the basis of beliefs about rivals and about

the expected value of the contracts they seek to win. In this environment, annuity quotes are equilibrium outcomes rather than administrative prices. Each offer reflects a trade-off between the probability of securing a contract and the expected value of servicing it over time. This approach connects the economics of pensions and insurance with the theory of strategic competition under private information (Myerson, 1981; Milgrom and Weber, 1982; Maskin and Riley, 2000; Fang and Wu, 2019; Chen et al., 2023; Verani and Yu, 2024).

The paper makes three contributions. First, it develops a tractable equilibrium framework for insurer competition in centralized retirement-income markets. The Chilean case motivates the institutional environment, but the mechanism is intended more broadly. Policy affects retirement payouts not only through household choices or actuarial forces, but also through strategic pricing and endogenous participation on the supply side. Second, the model embeds insurer competition in a planner's problem, which makes it possible to study how taxation and market conditions jointly affect pension offers, firm entry, tax revenue, and welfare. This feature places the paper close to the quantitative policy tradition that is strongly represented in *Economic Modelling* (Fougère and Mérette, 1999; Jimeno et al., 2008; Bielecki et al., 2015; Gustafsson, 2023; Chen et al., 2023; Kudrna et al., 2025). Third, the paper brings the model to the data using administrative records from Chilean annuity auctions. We recover an auction-level measure of expected competition and show that, conditional on observables, a stronger expected competitive environment is associated with higher awarded pensions.

Two implications follow immediately from the framework. Because participation is endogenous, policy changes need not have monotone effects on retirement outcomes. Shifts in taxation or profitability affect both pricing incentives and market structure. In addition, the welfare consequences of stronger competition depend on insurer heterogeneity and on the margins through which firms enter or exit the market. These mechanisms are especially relevant in annuity markets, where the conversion of a retirement balance into a lifelong payment stream magnifies pricing differences into sizable welfare effects over the retiree's remaining lifetime (Lockwood, 2012; Berstein and Morales, 2021; Bütler and Ramsden, 2024; Bello et al., 2024; Searle et al., 2024; Verani and Yu, 2024).

The rest of the paper is organized as follows. Section 2 presents the economic model and the Bayes–Nash formulation of insurer competition. Section 3 develops the simulation analysis and studies the effects of taxation and profitability shocks on market outcomes and planner welfare. Section 4 presents the empirical strategy and reduced-form evidence using administrative auction data from Chile. Section 5 discusses the broader implications for pension-market design, and Section 6 concludes.

4.2 Economic Model

In this research we consider a competitive market of life-insurance firms offering life annuities within a pension system. Annuities are provided via a centralised, integrated platform on which insurers bid for the retiree's accumulated balance at the point of retirement. Competition is not only between insurers. Retirees may alternatively choose a programmed-withdrawal arrangement, thereby intensifying competitive pressure and, potentially, raising benefit levels.

From a policy perspective, the government's main goal is to ensure that people can enjoy a better quality of life in retirement. To do so, it must seek ways to raise pension benefits while still maintaining public services, for example by introducing clear governance and control rules for the auction market where insurers compete.

4.2.1 Insurance firms

Let $\mathcal{M}$ denote the set of contracts under competition. For each $m \in \mathcal{M}$, insurer chooses a annuity-due schedule $m_m \equiv \{m_{m,t}\}_{t\geq 1}$, which determines the dynamics of the technical reserve $B_{m,t}$ and the probability of winning the contract $w_m(m_m)$. The annuity scheme also determines the expected value of the benefit of asset $V(m|m_m)$. The insurer's objective is to optimise the expected present value of the portfolio, which is

$$\max_{\{m_{m,t}\}_{m\in\mathcal{M},\, t\geq 1}} \quad \sum_{m\in\mathcal{M}} w_m(m_m)\,(V(m \mid m_m) - f_m) \;-\; D \tag{4.1}$$

subject to the following constraints,

$$B_{m,0} = F_m, \qquad B_{m,t} = B_{m,t-1} - m_{m,t} S_{m,t-1}, \tag{4.2}$$

$$0 \leq m_{m,t} \leq (1+r)\, B_{m,t-1}, \qquad B_{m,t} \geq 0, \tag{4.3}$$

$$\lim_{T\to\infty} d^T\, S_{m,T}\, B_{m,T} = 0. \tag{4.4}$$

where (4.2) characterises the insurer's reserve dynamics; (4.3) imposes operational feasibility, captured by the expected returns in the markets accessible to the pensioner; and (4.4) imposes a transversality condition that rules out explosive behaviour of the system. Here, f_m corresponds to the fixed cost for administering each contract m; F_m denotes the single premium paid by the pensioner; r the net return on the reserve's representative asset, which we will consider constant for the estimates of insurance companies; d represent the financial discount factor; $S_{m,t}$ the probability that the policy remains in force at the end of period t; and D denotes the minimum profit hurdle required of firm to remain in operation, as imposed firm's investors.

The expected present value of profits per policy, conditional on m_m and $B_{m,t}$, is given by

$$V(m \mid m_m) \;=\; \sum_{t\geq 1}^{\infty} d^t(1+r)\Big[S_{m,t-1}B_{m,t} \;-\; S_{m,t}\,f\Big], \tag{4.5}$$

where $f \geq 0$ denotes the operating cost incurred whenever the policy remains in force in period t for the firm. The function $w_m(m_m) \in [0,1]$ is the contract-award probability for policy m induced by the competitive rule (e.g., "highest pension wins").

Moreover, we allow for the possibility that rival insurers submit bids drawn from a distribution that is initially unknown, but as bids are observed, this distribution is progressively revealed. In this regard, it is assumed that firms only have incomplete information regarding how other firms will bid in a competitive market, derived from the mechanisms provided by the central planner. The probability of winning the policy, conditional on the actions of other firms, can be described as

$$w_m(m_m) \;= A_m\big(m_m\big) \prod_{j\neq i} G_{j,m}\big(m_m\big), \tag{4.6}$$

where $G_{j,m}$ denotes the cumulative distribution function (CDF) of rival firms' bids, evaluated at the annuity price $m_{i,m}$, and A_m is used to denote the pensioner's reservation threshold, that is, the minimum acceptable annuity payment. It can be assumed that this is unknown and can be derived from a cumulative distribution function of the pensioner.

4.2.2 Central planner's optimization problem

The government, acting as a central planner, possesses policy levers to intervene in the insurance market and thereby improve outcomes for pensioners. The government chooses $\tau \in [0,1)$ which corresponds a proportional tax τ on portfolio profits. The government's net revenue, which is dependent on the number of companies in the insurance market, is given by

$$\mathcal{R}(\tau; n) \;=\; \tau \sum_{i=1}^{n} \Pi_{\mathrm{i,Tot}}(n). \tag{4.7}$$

In this research, we assume a general tax policy that can have many variants, such as tax shields, among others. The objective is to prioritise the theoretical dimension over the practical implementation. In this context, the government's problem is to choose a regulatory policy that maximises pensioners' welfare subject to the market and fiscal constraints, which we express as

$$\max_{\tau,G} \quad \mathcal{W}(\tau, G) \;=\; \sum_{m=1}^{M} W_m\big(\,\overline{m}_m\big(n^*(\tau)\big)\big) \;+\; U_G(G), \tag{4.8}$$

subject to

$$0 \;\leq\; G \;\leq\; \mathcal{R}\big(\tau;\, n^*(\tau)\big), \tag{4.9}$$

where W_m and U_G denote, respectively, the pensioners' utility functions over individual income from life annuities and over public goods; $\overline{m}_m(n)$ is the expected winning pension in policy m, conditional upon a number of companies in the annuity market; n denotes the number of firms operating in the life-annuity market.

The central planner's utility function is based on the premise of a conflict between, on the one hand, maximising the income of pensioners over their lifetime and, on the other hand, the amount of resources available for the production of public goods. The central planner may assume that the greater the number of firms operating in the market, the greater the annuity income. To this end, the planner may modify tax policy, for example by reducing the tax burden. However, this action has the effect of reducing the amount of resources available for public goods.

It is important to note that the value of n is policy-dependent, this means that the policy measures taken by the government can directly or indirectly affect the entry of participants into the annuity market. This, in turn, affects the annuities of pensioners. The principle behind this is that the more competitive the market, the higher the annuity offered should be.

4.2.3 Nash equilibrium among insurers

Consider policy $m \in \mathcal{M}$. Before submitting an offer, insurer $i \in \{1, \ldots, N\}$ observes a private type $\theta_{i,m} \in \Theta_m$, which summarizes insurer-specific characteristics relevant for bidding, including administrative costs, portfolio conditions, and other firm-level determinants of contract profitability. Types are assumed to be independent across insurers and policies. Insurer i forms beliefs over rivals' types, denoted by $\varphi_{-i,m}$.

Rather than choosing a single deterministic offer directly, each insurer selects a parameter vector that generates an admissible annuity schedule. Formally, a strategy for insurer i in policy m is a measurable map

$$b_{i,m} : \Theta_m \to \mathcal{B}_m,$$

where $\mathcal{B}_m$ denotes the admissible parameter space. Given $\theta_{i,m}$, the choice $b_{i,m}(\theta_{i,m})$ induces an annuity-due schedule through the bid-generating rule

$$M_{i,m}(\cdot\,; b_{i,m}(\theta_{i,m})) : \Omega \to \mathcal{A}_m,$$

where $\mathcal{A}_m$ is the set of feasible annuity schedules satisfying (4.3) and (4.4). A realized bid is

therefore

$$m_{i,m} = M_{i,m}(\omega; b_{i,m}(\theta_{i,m})) \in \mathcal{A}_m.$$

This formulation allows strategies to be pure in the space of bid-generating parameters while remaining distributional in the space of realized offers. Economically, firms choose a bidding rule under incomplete information rather than a single offer known with certainty ex ante.

Let $G_{j,m}(\cdot)$ denote the cumulative distribution function of rival j's realized bid in policy m, as implied by its strategy and type distribution. Conditional on a realized schedule $m_{i,m} \in \mathcal{A}_m$, the probability that insurer i is awarded the contract is

$$w_{i,m}(m_{i,m}) = A_m(m_{i,m}) \prod_{j \neq i} G_{j,m}(m_{i,m}), \tag{4.10}$$

where $A_m(\cdot) \in [0,1]$ denotes the retiree's acceptance probability. If insurer i wins the contract, the present value of profits is $\mathrm{V}_i(m \mid m_{i,m})$, as defined in (4.5). Under risk neutrality, the expected payoff from policy m is

$$\mathbb{E}_{\Theta_{-i,m}} \left[A_m(m_{i,m}) \prod_{j \neq i} G_{j,m}(m_{i,m}) \mathrm{V}_i(m \mid m_{i,m}) \right]. \tag{4.11}$$

Definition 1. *For a given policy m, a profile of strategies*

$$b_m^* = (b_{1,m}^*, \ldots, b_{N,m}^*), \qquad b_{i,m}^* : \Theta_m \to \mathcal{B}_m,$$

is a Bayes–Nash equilibrium if, for each insurer i:

(a) ***Best response.*** *For every type $\theta_{i,m} \in \Theta_m$, the induced schedule*

$$m_{i,m}^*(\theta_{i,m}) \equiv M_{i,m}(\cdot\,; b_{i,m}^*(\theta_{i,m})) \in \mathcal{A}_m$$

satisfies

$$m_{i,m}^*(\theta_{i,m}) \in \arg\max_{a \in \mathcal{A}_m} \mathbb{E}_{\Theta_{-i,m}} [w_{i,m}(a)\, \mathrm{V}_i(m \mid a)],$$

where the expectation is taken with respect to insurer i's beliefs $\varphi_{-i,m}$ over rivals' types, and $w_{i,m}(\cdot)$ is defined in (4.10)*.*

(b) ***Belief consistency.*** *For each rival $j \neq i$, the bid distribution $G_{j,m}$ coincides with the distribution induced by the equilibrium strategy $b_{j,m}^*$ and the corresponding type distribution.*

This equilibrium notion is closely related to the idea of equilibrium in distributions studied by Fang and Wu (2019). The key modeling choice is that the annuity offer is not treated as a single

deterministic object fixed ex ante, but as the realization of a bid-generating rule chosen under incomplete information. This representation is natural in the present setting, since insurers face uncertainty about rivals' behavior, contract profitability, and future conditions affecting the value of the policy. The Bayes–Nash framework therefore captures strategic competition when firms optimize over bid-generating parameters while holding probabilistic beliefs about competing offers.

Algorithm 2 summarizes the computational procedure used to approximate the equilibrium. For each policy, insurers start from prior beliefs over rivals' bid-generating parameters, compute best responses, submit the implied offers, and update beliefs after observing auction outcomes. Repetition of this procedure yields a sequence of bid-generating rules that converges to a fixed point consistent with Bayes–Nash behavior.

Algorithm 2 Dynamic parameter selection and bidding in policy m

Input: set of insurers $\mathcal{I} = \{1, \dots, n\}$, priors $\{\varphi_{i,m}^{(0)}\}_{i \in \mathcal{I}}$, admissible parameter set $\mathcal{B}_m$, bid-generating rules $M_{i,m}(\cdot\,;\cdot)$, and profit functions $\mathrm{V}_i(m \mid \cdot)$.
Initialization: set $t = 0$ and choose initial parameter vectors $b_{i,m}^{(0)} \in \mathcal{B}_m$ for all $i \in \mathcal{I}$.
while not converged **do**
  $t \leftarrow t + 1$.
  for each insurer $i \in \mathcal{I}$ **do**
    Observe private type $\theta_{i,m}^{(t)} \in \Theta_m$.
    Construct beliefs $G_{j,m}^{(t)}(\cdot)$ over rivals' bids using $\varphi_{i,m}^{(t-1)}$ and the history of past bids.
    For each $b \in \mathcal{B}_m$, compute the induced schedule $M_{i,m}(\cdot\,; b)$ and the corresponding expected payoff

$$\Pi_{i,m}(b \mid \theta_{i,m}^{(t)}) = \mathbb{E}_{\Theta_{-i,m}}[w_{i,m}(M_{i,m}(\cdot\,; b))\, \mathrm{V}_i(m \mid M_{i,m}(\cdot\,; b))]\,.$$

    Choose

$$b_{i,m}^{(t)} \in \arg\max_{b \in \mathcal{B}_m} \Pi_{i,m}(b \mid \theta_{i,m}^{(t)}).$$

    Submit the induced bid schedule $m_{i,m}^{(t)} = M_{i,m}(\cdot\,; b_{i,m}^{(t)})$.
  end for
  Clear the auction and observe outcomes.
  for each insurer $i \in \mathcal{I}$ **do**
    Update beliefs $\varphi_{i,m}^{(t)}$ and the implied bid distributions via Bayes' rule.
  end for
end while
Output: $\{b_{i,m}^*, m_{i,m}^*\}_{i \in \mathcal{I}}$, approximating the Bayes–Nash equilibrium.

The planner's problem is evaluated over alternative tax policies by solving the insurers' equilibrium for each candidate value of τ. Algorithm 3 summarizes this step.

Algorithm 3 Central planner's choice of tax rate

Input: candidate tax grid $\mathcal{T} \subset [0, 1)$, welfare components $\{W_m\}_{m=1}^{M}$, $U_G(\cdot)$, and revenue function $\mathcal{R}(\tau; n)$.

Initialization: for each $\tau \in \mathcal{T}$ set $\mathcal{W}(\tau) \leftarrow -\infty$.

for each $\tau \in \mathcal{T}$ **do**

Run Algorithm 2 under tax rate τ to obtain the insurers' equilibrium and the implied number of active firms $n^*(\tau)$.

Compute the expected winning pension in each policy, $\overline{m}_m(n^*(\tau))$ for $m = 1, \ldots, M$.

Compute the feasible level of public expenditure, $0 \leq G \leq \mathcal{R}(\tau; n^*(\tau))$.

Solve

$$G^*(\tau) \in \arg \max_{0 \leq G \leq \mathcal{R}(\tau; n^*(\tau))} \left\{ \sum_{m=1}^{M} W_m\big(\overline{m}_m(n^*(\tau))\big) + U_G(G) \right\}.$$

Evaluate

$$\mathcal{W}(\tau) = \sum_{m=1}^{M} W_m\big(\overline{m}_m(n^*(\tau))\big) + U_G(G^*(\tau)).$$

end for

Choose

$$\tau^* \in \arg \max_{\tau \in \mathcal{T}} \mathcal{W}(\tau), \qquad G^* = G^*(\tau^*).$$

Output: optimal policy pair (τ^*, G^*) and associated equilibrium objects.

The simulation exercises are used to illustrate the economic content of the equilibrium rather than to replace the formal analysis. They make visible the mechanisms implied by the model, especially the interaction between strategic pricing, endogenous participation, and policy.

As illustrated in Figure 4.1, the expected profit of an insurer is non-monotone in its offer. When the proposed annuity is too low, the probability of winning the contract remains small and expected profits are limited. When the offer is too generous, the probability of securing the contract increases, but the profit margin becomes too thin. The figure therefore makes transparent the basic trade-off embedded in the Bayes–Nash problem. The optimal bid is an interior object determined jointly by competitive pressure and contract profitability.

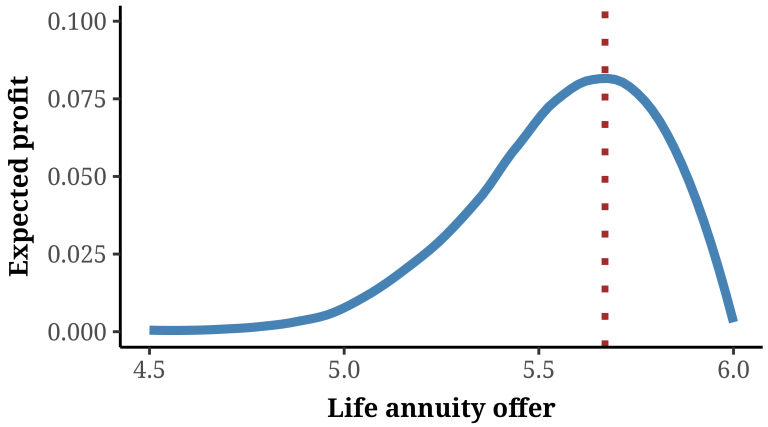
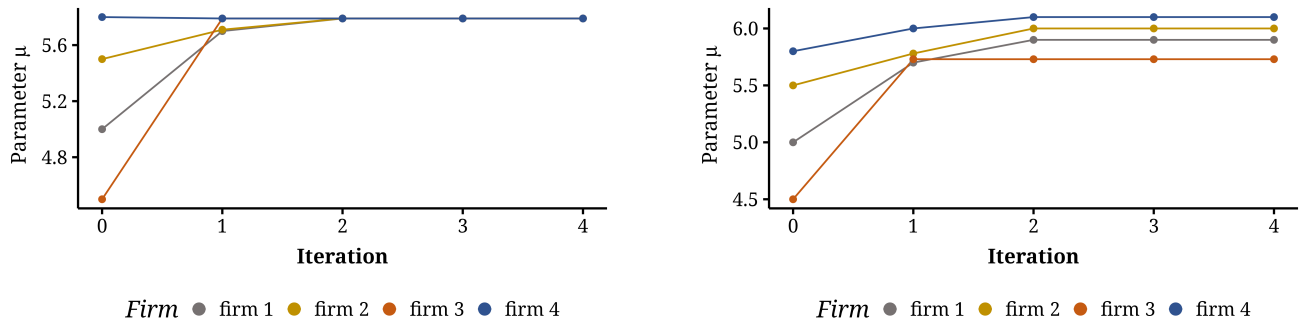


Figure 4.1: Best response from the firm when assuming a parameter structure from other firms.

Figure 4.2 illustrates the dynamic adjustment of beliefs under Algorithm 2. In panel 4.2a, firms share the same parameterization and converge rapidly to a common bidding rule. In panel 4.2b, convergence also occurs when firms differ in their underlying parameters, although the limiting values remain heterogeneous. These figures shows that the equilibrium concept is not only theoretically well defined, but also computationally tractable under repeated belief updating.

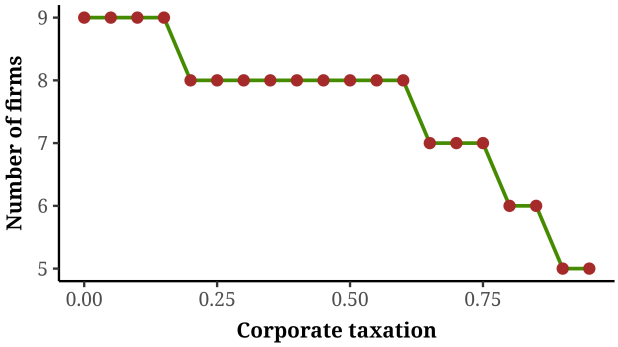


(a) Same parameters.

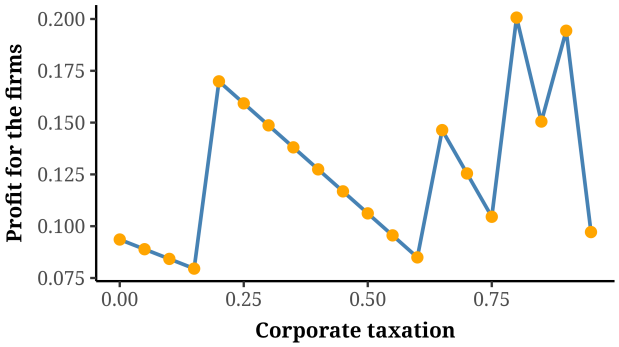
(b) Different parameters.

Figure 4.2: Best responses under alternative assumed parameter structures.

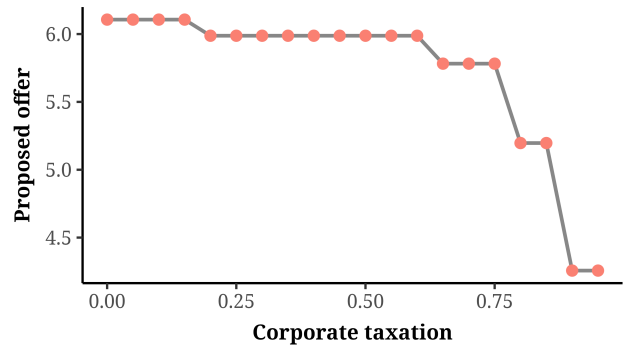
Figure 4.3 summarizes how the equilibrium changes with taxation and insurer profitability. The first four panels report comparative statics with respect to the corporate tax rate, while panels 4.3e and 4.3f report the response of the equilibrium offer and the planner's utility to a profitability shock, holding the tax rate fixed at $\tau = 0.25$. Taken together, these panels highlight the central mechanism of the model. Any force that compresses insurers' margins affects equilibrium bidding directly, but it may also alter participation, which then feeds back into competition and welfare.



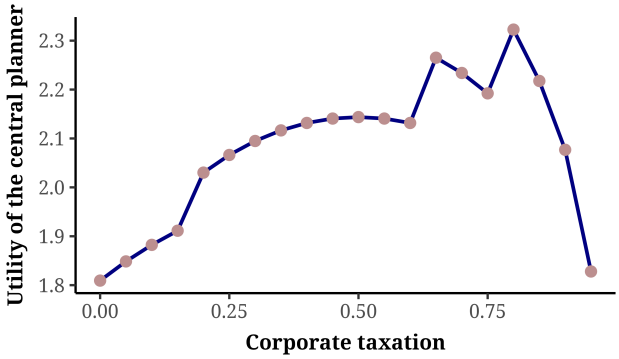
(a) Evolution of the number of firms.



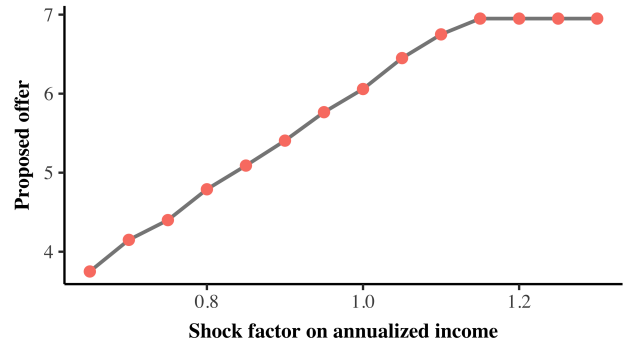
(b) Evolution of profits.

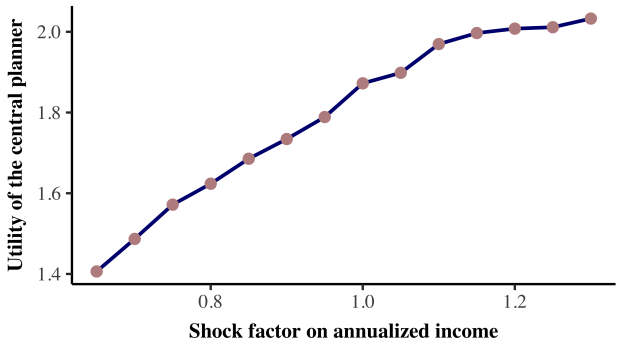


(c) Evolution of the proposed offer.



(d) Evolution of central planner utility.





(e) Response of the proposed offer to the shock factor. (f) Response of central planner utility to the shock factor.

Figure 4.3: Evolution of the number of firms, profits, proposed offer, and central planner utility.

Panels 4.3a–4.3d show that the effect of taxation is strongly non-linear because participation is endogenous. As shown in panel 4.3a, the number of active firms remains unchanged for relatively low tax rates, but falls discretely once after-tax profitability crosses firms' participation thresholds. This is economically important because taxation does not only reduce net returns. It also changes

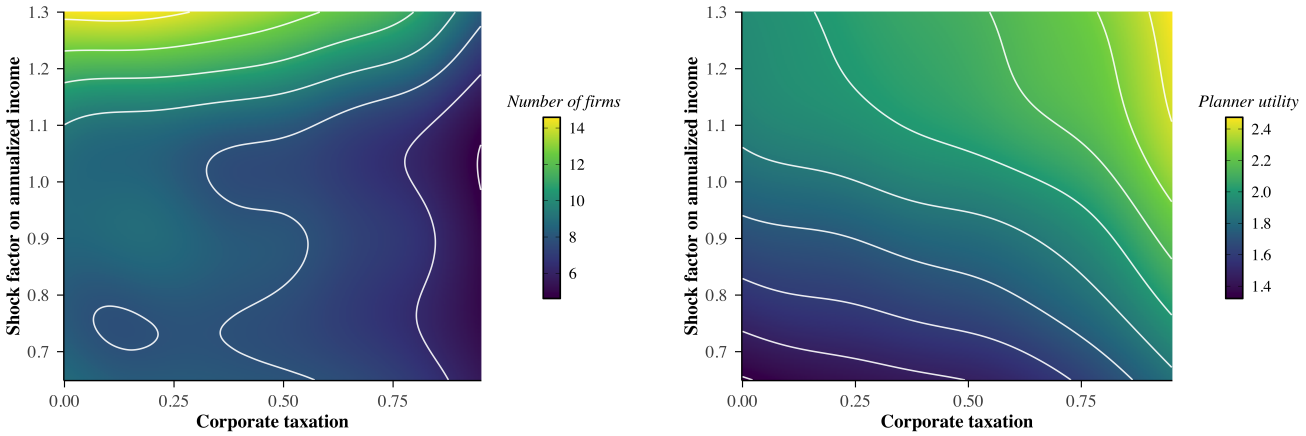
market structure.

The joint reading of panels 4.3b and 4.3c makes the mechanism especially clear. For a given number of firms, a higher tax rate reduces the net value of winning a contract and tends to compress profits. At the same time, once some firms exit, competitive pressure weakens and the surviving insurers can sustain lower pension offers. The result is a sequence of local downward trends combined with discrete jumps generated by endogenous exit. From the retiree's perspective, panel 4.3c shows the cost of this adjustment. The equilibrium annuity declines as taxation progressively reduces participation, so the effect of tax policy on retirement income operates through both prices and entry.

Panel 4.3d shows that planner utility is also non-monotone in taxation. Under the baseline calibration, utility initially rises because higher taxes increase fiscal resources, but eventually declines as weaker competition and lower offers dominate the revenue effect. The figure therefore illustrates a genuine policy trade-off rather than a simple preference for either low or high taxation. The welfare-maximizing tax rate in the simulation should be interpreted as calibration-specific, but the underlying lesson is more general: once participation is endogenous, the policy ranking cannot be inferred from pension offers or tax revenue alone.

Panels 4.3e and 4.3f examine exogenous profitability shocks. A favorable shock increases the expected value of contracts and therefore relaxes the insurers' margin constraint. This raises equilibrium offers and improves planner utility, although the pass-through is not one-for-one. The incomplete transmission is important. It shows that better fundamentals do not translate mechanically into proportionally higher pensions, because equilibrium pricing still depends on strategic interaction and entry conditions.

Finally, Figure 4.4 combines profitability shocks and tax changes in a single welfare map. The left panel shows how the number of active firms varies across both dimensions, while the right panel reports the planner's utility. The figure provides a compact representation of the state dependence of optimal policy. A tax rate that is sustainable under favorable profitability conditions may become welfare reducing when firms face weaker margins. In the same way, stronger profitability can partly offset the welfare cost of taxation by supporting entry and preserving competition.



(a) Number of firms across shock factors and corporate tax rates.

(b) Planner utility across shock factors and corporate tax rates.

Figure 4.4: Market outcomes across shock factors and corporate tax rates.

For that reason, Figure 4.4 should be read as a summary of the central policy tension in the model. The planner does not choose policy in a fixed competitive environment. Policy itself reshapes the competitive environment by changing the profitability and participation conditions under which insurers submit annuity offers.

4.3 Empirical strategy

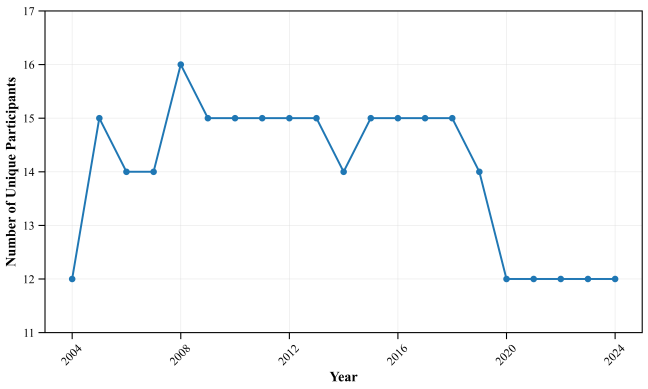
This section provides a reduced-form empirical counterpart to the theoretical model. The model implies that insurers condition their bids on the expected competitive environment prevailing at the time the offer is submitted, not on the ex post realized number of competitors. The empirical objective is therefore not to estimate a fully structural entry-and-bidding model, but to construct an auction-level measure of expected competition and assess whether it helps explain the winning annuity offer.

We use administrative records from the Chilean Superintendency of Pensions, which contain the history of annuity offers submitted through the SCOMP platform.[1] The monetary unit is the *Unidad de Fomento* (UF), an inflation-indexed unit of account widely used in Chilean long-term contracts. To proxy for prevailing financial conditions, we complement the administrative records with the secondary-market rate on five-year UF-denominated bonds reported by the Central Bank of Chile (BTU).[2] In the empirical implementation, this variable enters through its average value at the relevant date of the auction.

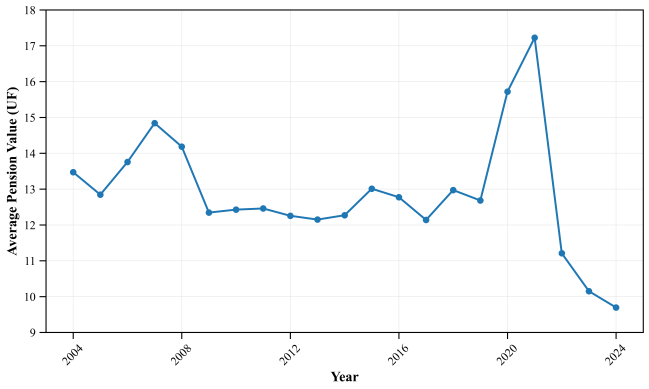
Figure 4.5 provides descriptive context for the empirical analysis. It reports the annual evolution of the number of participants observed in the sample, the average UF value of winning pensions, and the annual distribution of the five-year BTU rate. Taken together, these series summarize the institutional and financial environment in which annuity auctions are observed.

[1] Public access to the data and institutional documentation is available at `https://www.spensiones.cl/apps/bdp/index.php`.

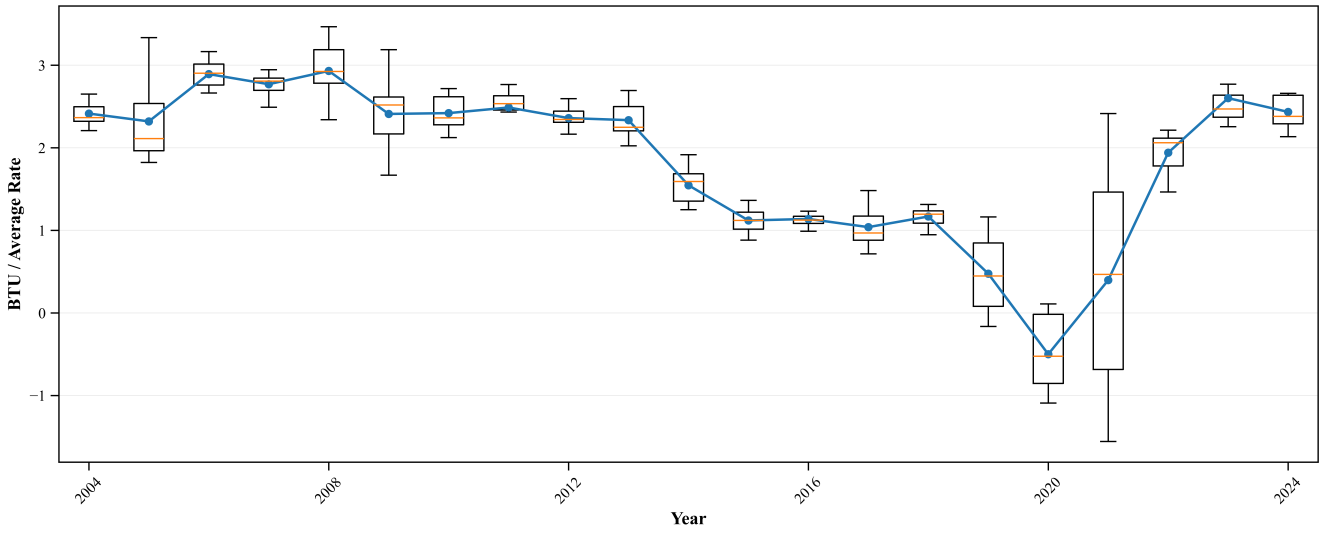
[2] Public access to the time series is available at `https://si3.bcentral.cl/siete/ES/Siete/Cuadro/CAP_EST_OMA/MN_EST_OMA/TMS_16?idSerie=F022.BUF.TIS.AN05.UF.Z.D`.



(a) Annual number of unique participants in the sample.



(b) Average UF value of winning pensions by year.



(c) Annual distribution of BTU rates and annual averages.

Figure 4.5: Annual evolution of market participation, winning pension values, and BTU rates.

Before turning to the formal estimation, Table 4.1 reports descriptive statistics for the main variables used in the two-stage exercise. The average auction attracts 10.21 participating insurers, with observed participation ranging from 1 to 15. The log of the winning annuity offer has a mean of 2.25 and a standard deviation of 0.69, while the log pension fund balance has a mean of 7.60 and a standard deviation of 0.68, indicating substantial heterogeneity in contract size. Retirees are on average 64.28 years old, and the percentile distribution is strongly concentrated around age 65, which is consistent with the institutional timing of retirement. The share of male retirees is 61.7%, and the average rate also exhibits non-trivial dispersion. Overall, the data display substantial variation in both participation and observables, which is necessary for the empirical strategy developed below.

The empirical analysis proceeds in two stages. In the first stage, we estimate the expected

Table 4.1: Descriptive statistics

Variable	N	Mean	SD	Min	p25	p50	p75	Max
Number of participants	543,279	10.208	2.595	1.000	9.000	11.000	12.000	15.000
$\log(P_a^{\text{win}})$	543,279	2.253	0.687	1.085	1.716	2.141	2.682	6.828
ln(UF balance)	543,279	7.598	0.678	5.030	7.072	7.501	8.038	11.692
Age	543,279	64.276	2.919	46.000	62.000	65.000	65.000	100.000
Male	543,279	0.617	0.486	0.000	0.000	1.000	1.000	1.000
Average rate	543,279	1.630	0.978	-1.557	1.096	1.788	2.381	3.518

Notes: This table reports descriptive statistics for the main variables used in the two-stage empirical analysis. The variable *Male* is a binary indicator equal to one for male retirees.

number of insurers participating in each auction as a function of ex ante observables available at the time bids are submitted. In the second stage, we examine whether this fitted measure of expected competition helps explain the winning annuity offer. This sequential design is consistent with the information structure of the model, in which firms bid conditional on expected rather than realized participation.

4.3.1 First-stage model of expected participation

Let a index an annuity auction, and let N_a denote the realized number of insurers submitting a valid offer in auction a. At the bidding stage, firms do not observe the realized value of N_a with certainty. Instead, they form expectations using auction-level observables known ex ante. In our baseline specification, these observables include the retiree's pension fund balance, gender, age, a benchmark financial condition measured by the average BTU rate, and calendar controls.

A relevant feature of the data is that the sample includes only auctions with at least one valid participant, so the observed support of N_a is strictly positive. For that reason, the first stage is estimated using a zero-truncated Poisson specification rather than a standard Poisson model. Let λ_a denote the latent Poisson intensity implied by ex ante observables,

$$\lambda_a = \exp(X_a'\beta),$$

where X_a collects the observable auction characteristics. Since the observed sample is truncated at zero, the first-stage object of interest is the conditional mean of participation given that at least one insurer enters the auction,

$$\mu_a = \mathbb{E}[N_a \mid N_a > 0, X_a].$$

The baseline empirical specification is therefore based on the zero-truncated Poisson model

$$N_a \mid X_a,\ N_a > 0 \sim \text{Truncated Poisson}(\lambda_a),$$

$$\log \lambda_a = \beta_0 + \beta_1 \ln(\text{UF balance}_a) + \beta_2 Male_a + \beta_3 Age_a + \beta_4 AverageRate_a + \delta_{y(a)} + \delta_{m(a)},$$

where $Male_a$ is an indicator equal to one for male retirees, Age_a denotes age at the auction date, $AverageRate_a$ captures the benchmark financial condition observed prior to bidding, and $\delta_{y(a)}$ and $\delta_{m(a)}$ denote year-of-entry and month-of-entry fixed effects.

The empirical specifications are reported in Table 4.2. Specification (1) includes only $\ln(\text{UF balance}_a)$ and calendar fixed effects. Specification (2) adds demographic controls through $Male_a$ and Age_a. Specification (3) further incorporates $AverageRate_a$. This nested structure makes it possible to assess how the estimated participation process changes as additional ex ante observables are introduced.

The fitted value used in the second stage is the truncated conditional mean,

$$\hat{\mu}_a = \widehat{\mathbb{E}}[N_a \mid N_a > 0, X_a],$$

which we interpret as an empirical proxy for the expected competitive environment faced by insurers prior to bidding. This interpretation should be understood in reduced-form terms. The first-stage equation is not a structural model of insurer entry. Since the set of potential insurers is finite, a literal entry model would naturally arise from insurer-level participation decisions. The zero-truncated Poisson formulation is instead used as a parsimonious device for summarizing expected competition under incomplete information in a sample where zero-participation auctions are not observed.

4.3.2 Second-stage model of winning offers

Having obtained an auction-level proxy for expected competition, we next examine whether it helps explain the winning annuity offer. Let P_a^{win} denote the winning annuity payment in auction a. Since annuity payments are strictly positive, we estimate the following log specification:

$$\log P_a^{\text{win}} = \alpha + \rho\, \hat{\mu}_a + X_a'\gamma + \delta_{y(a)} + \delta_{m(a)} + \varepsilon_a,$$

where $\hat{\mu}_a$ is the fitted expected number of participants from the truncated first stage, X_a contains the same auction-level observables, and ε_a is an error term.

In the richest specification, X_a includes $\ln(\text{UF balance}_a)$, $Male_a$, Age_a, and $AverageRate_a$. The second stage is again estimated sequentially to preserve comparability with the first-stage regressions. Specification (1) includes $\hat{\mu}_a$, $\ln(\text{UF balance}_a)$, and calendar fixed effects. Specification (2) adds the demographic controls $Male_a$ and Age_a. Specification (3) further incorporates $AverageRate_a$.

The coefficient of primary interest is ρ. In this setting, ρ measures the partial association between the expected degree of competition and the winning annuity offer, conditional on observables. A

positive estimate indicates that auctions expected to attract more competitors tend to be associated with higher winning offers, which is consistent with the competitive mechanism highlighted by the model. The interpretation remains reduced-form. Because $\hat{\mu}_a$ is a generated regressor and is itself a function of X_a, the second-stage coefficient should not be read as a structural causal effect of competition. Rather, it should be interpreted as evidence on whether the empirical proxy for expected competition contains additional explanatory content for auction outcomes beyond the included observables.

This distinction is important. The purpose of the second stage is not to recover the full mapping from beliefs into individual bids, but to assess whether the theoretical object highlighted by the model, namely expected competition, helps organize the variation in realized winning offers. Inference should therefore account for the generated-regressor nature of $\hat{\mu}_a$, for example through bootstrap procedures that incorporate estimation uncertainty from both stages.

The estimates reported in Table 4.2 are consistent with this interpretation. In the most parsimonious specification, the coefficient on expected participation is negative, but this should not be interpreted as evidence that competition lowers winning offers. It more likely reflects omitted heterogeneity correlated with both contract size and expected participation. Once additional controls are introduced, the coefficient becomes positive. In the richest specification, one additional expected participant is associated with an increase of approximately 3.1% in the winning annuity offer, conditional on the remaining regressors.

Taken together, the two stages provide a disciplined reduced-form bridge between the theoretical model and the data. The first stage constructs an empirically tractable measure of expected competition using only ex ante information. The second stage then shows that this expected competitive environment is systematically related to realized auction outcomes in the direction predicted by the model.

Table 4.2: First-Stage Zero-Truncated Poisson and Second-Stage OLS Estimates

	Specification		
	(1)	(2)	(3)
Panel A. First-stage zero-truncated Poisson estimates			
Dependent variable: number of participants			
ln(UF balance)	0.171***	0.177***	0.178***
	(275.17)	(279.14)	(279.58)
Male		0.0204***	0.0209***
		(17.52)	(17.94)
Age		-0.00985***	-0.00991***
		(-50.31)	(-50.58)
Average rate			0.0316***
			(28.29)
Constant	1.016***	1.591***	1.518***
	(119.86)	(112.68)	(105.73)
Year-of-entry FE	Yes	Yes	Yes
Month-of-entry FE	Yes	Yes	Yes
Observations	543,279	543,279	543,279
Panel B. Second-stage OLS estimates			
Dependent variable: $\log(P_a^{\text{win}})$			
Expected participants	-0.182***	0.0604***	0.0312***
	(-365.86)	(108.08)	(43.64)
ln(UF balance)	1.335***	0.879***	0.933***
	(1436.39)	(834.43)	(693.81)
Male		0.00343***	0.0104***
		(10.68)	(30.81)
Age		0.0355***	0.0326***
		(485.86)	(383.47)
Average rate			0.0233***
			(64.70)
Constant	-5.845***	-7.125***	-7.117***
	(-1691.13)	(-1880.06)	(-1884.05)
Year-of-entry FE	Yes	Yes	Yes
Month-of-entry FE	Yes	Yes	Yes
Observations	543,279	543,279	543,279

Notes: Panel A reports first-stage zero-truncated Poisson estimates; z-statistics are reported in parentheses. Panel B reports second-stage OLS estimates; t-statistics are reported in parentheses. All specifications include year-of-entry and month-of-entry fixed effects, whose coefficients are omitted for brevity. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

4.4 Discussion

A useful way to interpret the equilibrium concept developed in this paper is as a bridge between the classical theory of mixed strategies and a continuous environment with incomplete information. In a standard mixed-strategy game, players randomize over a finite or countable set of actions. Here, by contrast, insurers effectively choose parameters that generate a probability distribution over a continuum of feasible offers. In that sense, the framework can be read as an infinitesimal analogue of a mixed game, where the strategic object is no longer a simple probability vector over a few actions, but a distribution over admissible bids. This interpretation is especially natural in annuity markets, where firms do not choose among a small number of discrete contracts, but instead formulate finely tuned offers conditional on beliefs about rivals, profitability, and acceptance. The Bayes–Nash equilibrium in distributions therefore preserves the logic of classical game theory while adapting it to a setting in which strategic randomization is continuous, parametric, and economically meaningful (Myerson, 1981; Milgrom and Weber, 1982; Maskin and Riley, 2000; Fang and Wu, 2019).

The main economic implication of the paper is that competition in annuity markets should be understood as an equilibrium object shaped jointly by pricing incentives, participation, and institutional design. This point matters because much of the retirement literature has focused either on household annuitization decisions or on the long-run aggregate consequences of pension reform, often leaving the supply-side competitive mechanism in the background (Fougère and Mérette, 1999; Jimeno et al., 2008; Bielecki et al., 2015; Gustafsson, 2023; Kudrna et al., 2025). Our results suggest that, once the payout phase is modeled explicitly as a strategic market interaction, policy affects retirement income not only through savings accumulation or product choice, but also through firms' equilibrium responses. This places the analysis closer to the quantitative policy tradition in which institutional details and equilibrium adjustments are central to the evaluation of reforms.

The simulation results reinforce this message. The Figure 4.3, shows that taxation has non-linear effects because participation is endogenous. The relevant mechanism is not simply that taxes compress margins. Higher taxes can also trigger firm exit, weaken bidding pressure, and lower pension offers, even when the planner's objective temporarily improves through higher fiscal revenue. This type of non-monotonicity is important because it implies that policy evaluation cannot be based on a single outcome variable. In that respect, the paper complements the literature on pension reform and retirement-income design by showing that equilibrium supply responses can overturn partial-equilibrium intuitions about taxes, competition, and welfare (Fougère and Mérette, 1999; Jimeno et al., 2008; Gustafsson, 2023; Chen et al., 2023).

The empirical evidence points in the same direction. In the richest specification, expected competition is positively associated with the winning annuity offer, which is consistent with the

mechanism highlighted by the model. The descriptive patterns reported in Figure 4.5 also support the idea that annuity outcomes are shaped jointly by financial conditions and participation on the supply side. This interpretation is consistent with the Chilean evidence showing that SCOMP improved transparency and competition, but it also suggests that centralization does not eliminate strategic interaction or heterogeneity across insurers (Reyes, 2010; Morales and Larraín, 2017; Alcalde and Vial, 2022; Bello, 2023; Fajnzylber et al., 2023). In that sense, the Chilean case is useful not merely because it is a well-documented institutional setting, but because it reveals how market design and strategic competition coexist in the payout phase of defined-contribution systems.

More broadly, the paper contributes to the annuity literature by shifting attention from actuarial pricing alone to the interaction between incomplete information, endogenous participation, and regulatory design. Recent work has shown that annuity demand and pricing depend on taxes, macro-longevity risk, insurer default risk, guarantees, and behavioral frictions (Bütler and Ramsden, 2024; Balter et al., 2021; Searle et al., 2024; Bello et al., 2024; Verani and Yu, 2024). Our results are consistent with that broader view, but they add a distinct supply-side channel. Even when retirees face the same institutional platform, equilibrium payouts depend on how insurers compete, on how many firms find it profitable to participate, and on how policy reshapes those margins. This is why the planner's problem is not a decorative extension of the auction model, but a necessary component for evaluating pension-market reforms in an internally consistent way.

The paper also has limitations, and they are worth stating explicitly. The empirical strategy is reduced-form and therefore does not recover a full structural model of entry and bidding. The first-stage Poisson equation should be interpreted as a tractable proxy for expected competition rather than as a literal participation game. In addition, the assumption that bid distributions are independent is analytically convenient but restrictive. It abstracts from repeated interaction, balance-sheet dependence, correlated beliefs, and insurer-level regulatory constraints that may matter quantitatively in actual annuity markets. On the demand side, retirees are represented in a stylized way, so the current framework does not yet incorporate richer heterogeneity in bequest motives, financial literacy, default-risk perceptions, or intermediary influence (Lockwood, 2012; Alcalde and Vial, 2022; Bello et al., 2024; Searle et al., 2024).

These limitations open several directions for future research. One natural extension is a structural model of insurer participation and bidding with a finite set of potential entrants and correlated strategies. A second extension is to enrich the demand side by allowing retirees to differ in annuitization preferences, health expectations, bequest motives, and responses to guarantees or firm reputation. A third extension is dynamic. Insurers' bidding behavior is likely to depend on asset–liability management, solvency regulation, and intertemporal portfolio considerations, which are not modeled here. Comparative work across other centralized or partially centralized decumulation systems would also help distinguish which results are specific to Chile and which

reflect more general properties of retirement-income markets under incomplete information.

4.5 Conclusions

This paper studies insurer competition in annuity markets through the lens of strategic interaction under incomplete information. The central idea is that retirement-income outcomes depend not only on how much households accumulate before retirement, but also on how insurers compete when those balances are converted into life annuities. To analyze that mechanism, we model the market as a centralized first-price auction with heterogeneous insurers, endogenous participation, and a planner's problem that links pension outcomes to fiscal considerations.

The paper makes three main contributions. First, it develops a tractable Bayes–Nash framework in which annuity offers are equilibrium objects shaped by beliefs, profitability, and market rules. This moves the analysis beyond purely actuarial or administrative views of annuity pricing and places the payout phase within a broader framework of market design and industrial organization. Second, the model shows that policy interventions such as taxation can have non-monotone effects because they alter both firms' margins and market participation. As a result, the tax rate that maximizes planner welfare need not coincide with the tax rate that maximizes pension offers, profits, or the number of active firms. Third, the paper provides an empirical bridge between the model and observed outcomes using Chilean administrative auction data. The evidence indicates that stronger expected competition is associated with higher winning annuity offers, which is consistent with the mechanism emphasized by the theory.

Taken together, the results suggest that centralized platforms improve annuity markets, but they do not remove strategic behavior. Competition, heterogeneity, and institutional design continue to play a decisive role in shaping retirement-income outcomes. This conclusion is relevant beyond the Chilean case. In defined-contribution systems more generally, the adequacy of pensions depends not only on contribution histories and asset accumulation, but also on the structure of the market through which retirement wealth is transformed into lifelong income. For that reason, policies aimed at improving pensions should not focus exclusively on the accumulation phase. They should also consider how payout markets are organized, how firms compete, and how regulation affects equilibrium offers, participation, and welfare.

Chapter 5

Conclusions and Future Work

This thesis addressed pension adequacy in an individual capitalization system as a problem that spans the entire life cycle of retirement savings. Through two complementary studies, the work showed that both the protection of the balance during the accumulation phase and its conversion into an annuity during the payout phase are, in essence, market-design problems: in both cases the outcome the retiree experiences is not a given technical datum, but an equilibrium object that regulation can shape.

In the accumulation phase, the work developed a decision-support framework for a threshold-activated pension-loss insurance against severe and infrequent market shocks. The analysis showed that age-contingent coverage can protect adequacy against these episodes without requiring a full public guarantee, but only within an admissible region that jointly satisfies welfare non-deterioration, expected budget balance, stochastic solvency of the fund, and a tolerable moral-hazard limit. Two results organize the design: the expected pension damage is non-decreasing in the age at which the shock occurs —which justifies coverage increasing with age— and, under a primitive condition, the use of purely pooled pricing induces greater risk-taking than actuarially fair pricing, which makes it necessary to discipline pricing. Finally, the framework makes the system's contingent liability explicit through a prudential reserve, which in the baseline design stands at around 1.5 entry wages per worker.

In the payout phase, the work modeled competition among insurers in the annuity market as a centralized first-price auction with heterogeneous firms, incomplete information, and endogenous participation, embedded in the problem of a planner who weighs pensions against fiscal revenue. The model showed that annuity offers are equilibrium objects and that policy levers can have non-monotone effects: higher taxation can reduce pensions through the exit of firms, even when fiscal revenue increases. The empirical evidence, based on administrative data from the Chilean market, is consistent with this mechanism: conditional on observables, auctions expected to attract more insurers are associated with higher winning offers.

Read together, both results point in the same direction, and it is there that the contribution of this thesis as a whole resides. First, in both phases the relevant outcome is an equilibrium shaped by regulation and by the behavior of agents, and not a purely actuarial or administrative magnitude. Second, in both phases policy interventions exhibit non-obvious effects —bounded admissible regions and non-monotone responses— that partial-equilibrium intuition tends to overlook. Third, and most importantly, protecting pension adequacy requires attending to the entire life cycle of savings: a reform that only raises the contribution rate, or that intervenes in a single phase, can be offset by market responses in the other. Pension security is built both by protecting savings while they accumulate and by organizing the market through which they are paid out.

A prudent policy implication follows from the above. The evaluation of pension policy benefits from being treated as a life-cycle system, in which decisions about the protection of savings and about the design of the payout market are considered coherently. In particular, the regulator should make the system's contingent obligations explicit —for example, through prudential reserves— and incorporate firms' supply and participation responses when designing taxes and market rules, rather than assuming that such responses are neutral.

This thesis has limitations that are worth stating and that, at the same time, define its future agenda. Both studies are reduced-form and are calibrated to or confronted with the Chilean case, so their external validity should be assessed with caution; neither recovers point-identified causal effects; and, above all, the work studies the two phases in parallel rather than within a single equilibrium model that links them. That integration constitutes the main direction for future research: to build a life-cycle framework that explicitly connects the protection of savings in accumulation with competition in decumulation, to structurally estimate the mechanisms identified here, to compare the Chilean case with other capitalization systems, and to characterize the joint design —of insurance against shocks and of annuity-market regulation— that maximizes welfare subject to the system's fiscal and solvency constraints.

A final consideration, of an institutional nature, reinforces this agenda. The Chilean environment to which both studies are calibrated and confronted has begun to change with the 2025 pension reform, which establishes a mixed system and directly modifies the mechanisms studied here. In the accumulation phase, the reform gradually replaces the multi-fund structure with generational funds: portfolios defined by age brackets, with a declining risk trajectory as the member approaches retirement. This change is particularly relevant to the first study for two reasons. First, the automatic decline in equity exposure for older cohorts partially internalizes the same life-cycle asymmetry that justifies age-contingent coverage, so that the marginal role of the insurance, its calibration, and its admissible region must be reassessed under the new default investment structure. Second, since mandatory contributions come to be assigned by age and cannot be freely reallocated across funds, the portfolio-risk choice margin —central to the result on pooled versus actuarially

fair pricing— becomes bounded for mandatory savings and shifts toward voluntary savings and the choice of administrator, which attenuates and reshapes the moral-hazard channel. Recalibrating the framework to the generational-fund trajectory and reexamining the admissible region and the prudential reserve therefore constitutes a natural extension.

In the payout phase, the same reform incorporates a social-insurance component financed by a higher employer contribution and centralizes aspects of collection, which alters both the balance that reaches the annuity market and the regulatory environment in which insurers compete. Extending the auction-and-planner model to this new environment, and exploiting the transition —implemented gradually over the decade— as a setting to test the participation-and-competition mechanism, offers a direct avenue for future research. Taken together, these changes do not invalidate the mechanisms identified in this thesis; rather, they underscore their relevance and provide a concrete institutional context onto which to project them.

Bibliography

Pilar Alcalde and Bernardita Vial. Implicit trade-offs in replacement rates: Consumer preferences for firms, intermediaries and annuity attributes. *International Journal of Industrial Organization*, 82:102827, 2022. doi: 10.1016/j.ijindorg.2022.102827.

Pablo Antolín, Stéphanie Payet, Edward Whitehouse, and Juan Yermo. The role of guarantees in defined contribution pensions. OECD Working Papers on Finance, Insurance and Private Pensions 11, OECD, 2011.

Alberto Arenas de Mesa and Carmelo Mesa-Lago. The structural pension reform in chile: Effects, comparisons with other latin american reforms, and lessons. *Oxford Review of Economic Policy*, 22(1):149–167, 2006.

Kenneth J. Arrow. Uncertainty and the welfare economics of medical care. *The American Economic Review*, 53(5):941–973, 1963.

Laurence Ball and N. Gregory Mankiw. Intergenerational risk sharing in the spirit of arrow, debreu, and rawls, with applications to social security design. *Journal of Political Economy*, 115(4): 523–547, 2007.

Anne G. Balter, Malene Kallestrup-Lamb, and Jesper Rangvid. Macro longevity risk and the choice between annuity products: Evidence from denmark. *Insurance: Mathematics and Economics*, 99:355–362, 2021. ISSN 0167-6687. doi: 10.1016/j.insmatheco.2021.04.009. URL `https://www.sciencedirect.com/science/article/pii/S0167668721000809`.

Nicholas Barr and Peter Diamond. Reforming pensions: Principles, analytical errors and policy directions. *International Social Security Review*, 62(2):5–29, 2009.

Robert J. Barro. Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, 121(3):823–866, 2006.

Robert J. Barro. Rare disasters, asset prices, and welfare costs. *The American Economic Review*, 99(1):243–264, 2009.

Yilmaz Bayar et al. Pension funds, insurance companies and stock market development: Evidence from emerging markets. *Economic Research–Ekonomska Istraživanja*, 35(1):4922–4942, 2022.

Roel M. W. J. Beetsma, Ward E. Romp, and Siert J. Vos. Voluntary participation and intergenerational risk sharing in a funded pension system. *European Economic Review*, 56 (6):1310–1324, 2012.

Piera Bello. Gender-based price discrimination in the annuity market: Evidence from chile. *European Economic Review*, 151:104356, 2023. doi: 10.1016/j.euroecorev.2022.104356.

Piera Bello, Agar Brugiavini, and Vincenzo Galasso. Annuity puzzle: Evidence from a swiss pension fund. *Journal of Risk and Insurance*, 91(3):653–696, 2024. doi: 10.1111/jori.12473. URL `https://doi.org/10.1111/jori.12473`.

Solange Berstein and Marco Morales. The role of a longevity insurance for defined contribution pension systems. *Insurance: Mathematics and Economics*, 99:233–240, 2021. doi: 10.1016/j.insmatheco.2021.03.020.

Solange Berstein, Guillermo Larraín, and Francisco Pino. Chilean pension reform: Coverage facts and policy alternatives. *Economía*, 6(2):227–279, 2006.

Marcin Bielecki, Karolina Goraus, Jan Hagemejer, Krzysztof Makarski, and Joanna Tyrowicz. Small assumptions (can) have a large bearing: Evaluating pension system reforms with olg models. *Economic Modelling*, 48:210–221, 2015. doi: 10.1016/j.econmod.2014.10.043.

Zvi Bodie, Robert C. Merton, and William F. Samuelson. Labor supply flexibility and portfolio choice in a life cycle model. *Journal of Economic Dynamics and Control*, 16(3–4):427–449, 1992.

M. Martin Boyer. Insuring catastrophic pension-fund losses through ART. *Journal of Risk and Insurance*, 76(2):323–344, 2009. doi: 10.1111/j.1539-6975.2009.01294.x.

Jeffrey R. Brown and George G. Pennacchi. Discounting pension liabilities: Funding versus value. *Journal of Pension Economics and Finance*, 15(3):254–284, 2016. doi: 10.1017/S1474747215000281.

Jeffrey R. Brown, Jeffrey R. Kling, Sendhil Mullainathan, and Marian V. Wrobel. Why don't people insure late-life consumption? a framing explanation of the under-annuitization puzzle. *American Economic Review*, 98(2):304–309, 2008. doi: 10.1257/aer.98.2.304.

Monika Bütler and Alma Ramsden. How taxes impact the choice between an annuity and the lump sum at retirement. *Journal of Pension Economics and Finance*, 23(1):1–29, 2024. doi: 10.1017/S1474747222000178. URL `https://doi.org/10.1017/S1474747222000178`.

Narat Charupat and Mark J. Kamstra. Behavior of canadian life annuity prices. *Journal of Pension Economics and Finance*, 2023. URL `https://api.semanticscholar.org/CorpusID:259779240`.

An Chen, Hong Li, and Mark B. Schultze. Optimal longevity risk transfer under asymmetric information. *Economic Modelling*, 120:106179, 2023. doi: 10.1016/j.econmod.2022.106179.

Andrea Ciacci, Enrico Ivaldi, and Tiziano Pavanini. Uneven paths toward sustainability in africa: A multidimensional and spatio-temporal assessment of sdg progress (2000–2022). *Socio-Economic Planning Sciences*, 103:102355, 2026. doi: 10.1016/j.seps.2025.102355.

João F. Cocco, Francisco J. Gomes, and Pascal J. Maenhout. Consumption and portfolio choice over the life cycle. *The Review of Financial Studies*, 18(2):491–533, 2005.

Harald Cramér. On the mathematical theory of risk. In *Skandia Jubilee Volume*. Skandia, Stockholm, 1930.

Rowena Crawford. The effect of the financial crisis on the retirement plans of older workers in england. *Economics Letters*, 121(2):156–159, 2013. doi: 10.1016/j.econlet.2013.07.020.

Jiajia Cui, Frank de Jong, and Eduard Ponds. Intergenerational risk sharing within funded pension schemes. *Journal of Pension Economics and Finance*, 10(1):1–29, 2011.

Maria João Durães, Flávia Barbosa, Giovanna D'Inverno, and Ana S. Camanho. Are european regions on the right track to achieve the 2030 strategic education and training targets? a comprehensive performance assessment. *Socio-Economic Planning Sciences*, 103:102390, 2026. doi: 10.1016/j.seps.2025.102390.

Sebastian Edwards. The chilean pension reform: A pioneering program. Working Paper 5811, National Bureau of Economic Research, 1996.

Martin Eling, Hato Schmeiser, and Joan T. Schmit. The Solvency II process: Overview and critical analysis. *Risk Management and Insurance Review*, 10(1):69–85, 2007. doi: 10.1111/j.1540-6296.2007.00106.x.

Paul Embrechts, Claudia Klüppelberg, and Thomas Mikosch. *Modelling Extremal Events for Insurance and Finance*, volume 33 of *Stochastic Modelling and Applied Probability*. Springer-Verlag, Berlin Heidelberg, 1997.

Eduardo Fajnzylber, M. Florencia Gabrielli, and Manuel Willington. Can transparency increase adverse selection? evidence from an electronic platform for annuities. *Economics Letters*, 228:111172, 2023. doi: 10.1016/j.econlet.2023.111172.

Chuyi Fang and Bin Wu. Socially-maximal nash equilibrium distributions in large distributional games. *Economics Letters*, 175:40–42, 2019.

Maxime Fougère and Marcel Mérette. Population ageing and economic growth in seven oecd countries. *Economic Modelling*, 16(3):411–427, 1999. doi: 10.1016/S0264-9993(99)00008-5.

Xavier Gabaix. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly Journal of Economics*, 127(2):645–700, 2012.

Christian Gollier. Intergenerational risk-sharing and risk-taking of a pension fund. *Journal of Public Economics*, 92(5–6):1463–1485, 2008.

Francisco Gomes. Portfolio choice over the life cycle: A survey. *Annual Review of Financial Economics*, 12:277–304, 2020.

Francisco Gomes and Alexander Michaelides. Optimal life-cycle asset allocation: Understanding the empirical evidence. *The Journal of Finance*, 60(2):869–904, 2005.

Roger H. Gordon and Hal R. Varian. Intergenerational risk sharing. *Journal of Public Economics*, 37(2):185–202, 1988.

Johan Gustafsson. Public pension reform with ill-informed individuals. *Economic Modelling*, 121: 106219, 2023. doi: 10.1016/j.econmod.2023.106219.

Bengt Holmström. Moral hazard and observability. *The Bell Journal of Economics*, 10(1):74–91, 1979.

Augusto Iglesias-Palau. Pension reform in chile revisited: What has been learned? OECD Social, Employment and Migration Working Papers 86, OECD, 2009.

Juan F. Jimeno, Juan A. Rojas, and Sergio Puente. Modelling the impact of aging on social security expenditures. *Economic Modelling*, 25(2):201–224, 2008. doi: 10.1016/j.econmod.2007.04.015.

Dirk Krueger and Felix Kubler. Pareto-improving social security reform when financial markets are incomplete!? *The American Economic Review*, 96(3):737–755, 2006.

George Kudrna, John Piggott, and Phitawat Poonpolkul. Sustainable and equitable pension reform for emerging economies: An application to indonesia. *Economic Modelling*, 148:107080, 2025. doi: 10.1016/j.econmod.2025.107080.

Ronald Lee. The demographic transition: Three centuries of fundamental change. *Journal of Economic Perspectives*, 17(4):167–190, 2003. doi: 10.1257/089533003772034943.

Christopher M. Lewis and George G. Pennacchi. The value of Pension Benefit Guaranty Corporation insurance. *Journal of Money, Credit and Banking*, 26(3):735–753, 1994. doi: 10.2307/2077996.

Katerina Lisenkova, Marcel Mérette, and Robert Wright. Population ageing and the labour market: Modelling size and age-specific effects. *Economic Modelling*, 35:981–989, 2013. doi: 10.1016/j.econmod.2013.09.007.

Lee M. Lockwood. Bequest motives and the annuity puzzle. *Review of Economic Dynamics*, 15(2): 226–243, 2012. doi: 10.1016/j.red.2011.03.001.

Fernando López and Eduardo Walker. Investment performance, regulation and incentives: The case of chilean pension funds. *Journal of Pension Economics and Finance*, 20(1):125–150, 2021. doi: 10.1017/S1474747219000350. URL `https://doi.org/10.1017/S1474747219000350`.

Miguel Lorca. Effects of covid-19 early release of pension funds: The case of chile. *Journal of Risk and Insurance*, 88(4):903–936, 2021. doi: 10.1111/jori.12365.

Carlos Madeira. The long-term impact of chilean policy reforms on savings and pensions. *The Journal of the Economics of Ageing*, 19:100326, 2021a. doi: 10.1016/j.jeoa.2021.100326.

Carlos Madeira. The long term impact of chilean policy reforms on savings and pensions. *The Journal of the Economics of Ageing*, 19:100326, 2021b. ISSN 2212-828X. doi: 10.1016/j.jeoa.2021.100326. URL `https://www.sciencedirect.com/science/article/pii/S2212828X21000190`.

Carlos Madeira. The impact of the chilean pension withdrawals during the covid pandemic on the future savings rate. *Journal of International Money and Finance*, 126:102650, 2022. doi: 10.1016/j.jimonfin.2022.102650.

Eric Maskin and John Riley. Asymmetric auctions. *The Review of Economic Studies*, 67(3):413–438, 2000. doi: 10.1111/1467-937X.00137.

Robert C. Merton. Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, 51(3):247–257, 1969.

Robert C. Merton. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373–413, 1971.

Robert C. Merton. On the role of social security as a means for efficient risk sharing in an economy where human capital is not tradeable. In Zvi Bodie and John B. Shoven, editors, *Financial Aspects of the United States Pension System*, pages 325–358. University of Chicago Press, Chicago, 1983.

Paul Milgrom and Chris Shannon. Monotone comparative statics. *Econometrica*, 62(1):157–180, 1994.

Paul R. Milgrom and Robert J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122, 1982.

Marco Morales and Guillermo Larraín. The chilean electronic market for annuities (scomp): Reducing information asymmetries and improving competition. *International Review of Financial Analysis*, 42:95–105, 2017. doi: 10.1057/s41288-017-0054-3.

Alicia H. Munnell, Alex Golub-Sass, Richard W. Kopcke, and Anthony Webb. What does it cost to guarantee returns? Issue Brief 9–4, Center for Retirement Research at Boston College, 2009.

Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

OECD. Retirement savings in the time of covid-19. OECD Policy Brief, Available at: `https://www.oecd.org/en/publications/retirement-savings-in-the-time-of-covid-19_b9740518-en.html`, 2020. Accessed 21 May 2026.

OECD. Pensions at a glance 2025: Chile. Available at: `https://www.oecd.org/en/publications/pensions-at-a-glance-2025-country-notes_8a53ef12-en/chile_99e9160c-en.html`, 2025. Accessed 21 May 2026.

Andrea Orozco-Villodres, Ana B. Ruiz, and Mariano Luque. How to enhance students' well-being in spain: A multidimensional study based on multiobjective optimization. *Socio-Economic Planning Sciences*, 103:102352, 2026. doi: 10.1016/j.seps.2025.102352.

George G. Pennacchi. The value of guarantees on pension fund returns. *The Journal of Risk and Insurance*, 66(2):219–237, 1999. doi: 10.2307/253677.

James Poterba, Joshua Rauh, Steven Venti, and David Wise. Lifecycle asset allocation strategies and the distribution of 401(k) retirement wealth. In David A. Wise, editor, *Developments in the Economics of Aging*, pages 15–50. University of Chicago Press, Chicago, 2009a.

James M. Poterba, Joshua Rauh, Steven F. Venti, and David A. Wise. Reducing social security PRA risk at the individual level: Life-cycle funds and no-loss strategies. In Jeffrey R. Brown, Jeffrey B. Liebman, and David A. Wise, editors, *Social Security Policy in a Changing Environment*, pages 255–292. University of Chicago Press, Chicago, 2009b.

Artur Raviv. The design of an optimal insurance policy. *The American Economic Review*, 69(1): 84–96, 1979.

Bernd Rechel, Emily Grundy, Jean-Marie Robine, Jonathan Cylus, Johan P. Mackenbach, Cecile Knai, and Martin McKee. Ageing in the european union. *The Lancet*, 381(9874):1312–1322, 2013. doi: 10.1016/S0140-6736(12)62087-X.

Gonzalo Reyes. Market design for the provision of social insurance: the case of disability and survivors insurance in chile. *Journal of Pension Economics and Finance*, 9(3):421–444, 2010. doi: 10.1017/S1474747209990059.

Thomas A. Rietz. The equity risk premium: A solution. *Journal of Monetary Economics*, 22(1): 117–131, 1988.

Michael Rothschild and Joseph Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, 90(4):629–649, 1976.

Paul A. Samuelson. Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics*, 51(3):239–246, 1969.

Pamela Searle, Peter Ayton, and Iain Clacher. Annuity selection in the presence of insurer default risk and government guarantees. *Journal of Risk and Insurance*, 91(1):161–192, 2024. doi: 10.1111/jori.12457. URL `https://doi.org/10.1111/jori.12457`.

Mauricio Soto. The chilean pension reform: 25 years later. *Pensions*, 12(4):289–301, 2007. doi: 10.1057/palgrave.pm.5950049.

Fernando Suárez, José Manuel Peña, and Omar Larré. Target-date funds: A state-of-the-art review with policy applications to Chile's pension reform. arXiv preprint arXiv:2504.17713, 2025.

Ian Tower and Gregorio Impavido. How the financial crisis affects pensions and insurance and why the impacts matter. IMF Working Paper WP/09/151, International Monetary Fund, 2009.

Stéphane Verani and Pei Cheng Yu. What's wrong with annuity markets? *Journal of the European Economic Association*, 22(4):1981–2024, February 2024. doi: 10.1093/jeea/jvae007. URL `https://doi.org/10.1093/jeea/jvae007`.

Luis M. Viceira. Optimal portfolio choice for long-horizon investors with nontradable labor income. *The Journal of Finance*, 56(2):433–470, 2001.

Jessica A. Wachter. Can time-varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance*, 68(3):987–1035, 2013.

Edward Whitehouse. Pensions during the crisis: Impact on retirement income systems and policy responses. *The Geneva Papers on Risk and Insurance: Issues and Practice*, 34(4):536–547, 2009. doi: 10.1057/gpp.2009.25.

World Bank. *Averting the Old Age Crisis: Policies to Protect the Old and Promote Growth*. Oxford University Press, New York, 1994.