

Optimal Menu-of-Contracts Design in Reinsurance

*A Principal–Agent Framework with
Kelly-Criterion Growth Maximization*

Frank Schmid

frank.schmid@frankschmid.com

Working Paper

This version: July 12, 2026

First version: March 15, 2026

Abstract

We develop a formal economic model of optimal reinsurance contract design in which a reinsurer (principal) faces a primary insurer (agent) possessing private information about the quality of the ceded risk. Drawing on the Rothschild–Stiglitz and Stiglitz screening frameworks, and the recent principal–agent reinsurance models of Cheung, Yam, and Yuen, we construct a menu-of-contracts mechanism that simultaneously mitigates adverse selection—reducing errors of commission (writing unprofitable business)—and avoids errors of omission (declining profitable business). The reinsurer’s objective is to maximize the long-run geometric growth rate of surplus, formalized via the Kelly criterion. We characterize the optimal separating menu for both proportional and non-proportional reinsurance structures, spanning property & casualty and life & health lines, and encompassing both treaty and facultative placements. We prove that a menu-of-contracts reinsurer achieves strictly higher long-run growth than competitors offering take-it-or-leave-it single contracts. A Monte Carlo simulation with 10,000 paths and 50 submissions per period validates the theoretical predictions. We further extend the model to the reinsurance market cycle, showing that the menu approach dominates across the entire cycle—from deep soft markets to hard markets—while exhibiting counter-cyclical stability.

JEL Classification: D82, D86, G22, C73

Keywords: menu of contracts, reinsurance, adverse selection, Kelly criterion, mechanism design, screening, principal–agent model, Monte Carlo simulation, market cycle

Contents

Abbreviations and Acronyms	iii
Glossary of Abbreviations	iv
1 Introduction	1
2 Related Literature	2
2.1 Adverse Selection in Insurance and Reinsurance	2
2.2 The Kelly Criterion and Growth-Optimal Strategies	2
2.3 Mechanism Design and Screening	3
3 Model Setup	4
3.1 Agents and Types	4
3.2 Contract Space	4
3.3 Reinsurer's Objective: Kelly-Criterion Growth Maximization	4
4 Optimal Menu of Contracts under Adverse Selection	5
4.1 The Screening Problem	5
4.2 Kelly-Optimal Separating Menu	5
4.3 Application Across Reinsurance Structures	5
5 Mitigating Errors of Commission and Omission	6
6 Competitive Dynamics: Menu vs. Take-It-or-Leave-It Rivals	7
7 Equilibrium Existence under Capacity Constraints	8
8 Computational Model and Monte Carlo Validation	9
8.1 Assumptions	9
8.2 Three Competing Strategies	9
8.3 Results	10
9 Market Cycle Analysis: Hardening and Softening	15
9.1 Market Hardness Parameterization	15
9.2 Theoretical Growth Rates across the Cycle	15
9.3 Menu Adjustment Mechanics	17
9.4 Simulated Surplus Trajectories across Regimes	18
9.5 Counter-Cyclical Properties	18
10 Extension: Continuum of Types	19
10.1 Optimal Schedule under IC	19
10.2 Optimal Number of Risk Groups	20
11 Extension: Dynamic Menus with Bayesian Experience Rating	22
11.1 Bayesian Updating Mechanism	22
11.2 Belief Convergence	22
11.3 The Ratchet Effect	23
11.4 Growth Advantage of Dynamic Menus	24

11.5 Learning Speed and Signal Quality	25
12 Extension: Multi-Dimensional Screening	27
12.1 Four-Type Model	27
12.2 Results	27
13 Extension: Heavy-Tailed Losses and Correlated Portfolios	29
13.1 Heavy-Tailed Loss Distributions	29
13.2 Correlated Losses	30
13.3 Joint Implications	31
14 Conclusion	32
Bibliography	33

Abbreviations and Acronyms

CESifo	Center for Economic Studies and the ifo Institute
IC	incentive compatibility
i.i.d.	independent and identically distributed
IR	individual rationality
JEL	Journal of Economic Literature (classification system)
L&H	life and health
MRS	marginal rate of substitution
NBER	National Bureau of Economic Research
P&C	property and casualty
VaR	Value-at-Risk
WMS	Wilson–Miyazaki–Spence

Greek Letters and Pronunciation

α	alpha (AL-fuh): coverage rate; cession fraction
δ	delta (DEL-tuh): premium loading schedule
Δ	Delta, capital (DEL-tuh): difference operator
ε	epsilon (EP-suh-lon): idiosyncratic shock
θ	theta (THAY-tuh): cedant type; expected loss ratio
Θ	Theta, capital (THAY-tuh): type space
κ	kappa (KAP-uh): distortion intensity
λ	lambda (LAM-duh): low-risk population share; posterior belief
μ	mu (MYOO): expected loss ratio
π	pi (PIE): reinsurance premium rate
ρ	rho (ROE): loss correlation
σ	sigma (SIG-muh): loss ratio volatility

Glossary of Abbreviations

Abbreviation	Definition
IC	Incentive Compatibility
IR	Individual Rationality
L&H	Life & Health
MRS	Marginal Rate of Substitution
P&C	Property & Casualty
VaR	Value-at-Risk
WMS	Wilson–Miyazaki–Spence

1. Introduction

The reinsurance market is characterized by pronounced information asymmetries. When a primary insurer submits a risk to a reinsurer—whether a property & casualty (P&C) exposure or a life & health (L&H) book, whether structured as a treaty or a facultative placement, whether proportional (quota share, surplus share) or non-proportional (excess of loss, stop-loss)—the cedant possesses superior information about the underlying risk quality. This is the classic adverse selection problem first formalized for insurance markets by Rothschild and Stiglitz¹ and extended to the monopolistic screening case by Stiglitz.²

The reinsurer confronts two symmetric decision errors. An **error of commission** occurs when the reinsurer writes a contract that turns out to be unprofitable—it has been adversely selected against by accepting a risk whose true loss characteristics are worse than priced. An **error of omission** occurs when the reinsurer declines or fails to compete for a submission that would have been profitable—the reinsurer’s pricing or terms are too conservative, and a profitable cedant is lost. Traditional single-price, take-it-or-leave-it underwriting conflates these two errors: tightening terms to avoid commission errors mechanically increases omission errors, and vice versa.

The menu-of-contracts approach, rooted in the mechanism design literature,³ resolves this tension by offering the cedant a *choice* among contracts that vary systematically in coverage, retention, and premium. By exploiting the single-crossing property—the fact that high-quality (low-risk) cedants and low-quality (high-risk) cedants have different marginal rates of substitution (MRS) between coverage and premium—the reinsurer can design a menu that induces truthful self-selection.

Our contribution is fivefold. First, we embed this screening mechanism in a reinsurance-specific model that accommodates the full taxonomy of reinsurance structures. Second, we adopt the Kelly criterion⁴ as the reinsurer’s objective function. Third, we analyze competitive dynamics when the menu-of-contracts reinsurer faces rivals offering single take-it-or-leave-it contracts. Fourth, we validate the theoretical predictions through a Monte Carlo simulation. Fifth, we extend the model to the reinsurance market cycle.

¹Michael Rothschild and Joseph E. Stiglitz, “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information,” *Quarterly Journal of Economics* 90, no. 4 (1976): 629–649.

²Joseph E. Stiglitz, “Monopoly, Non-Linear Pricing and Imperfect Information: The Insurance Market,” *Review of Economic Studies* 44, no. 3 (1977): 407–430.

³Roger B. Myerson, “Optimal Auction Design,” *Mathematics of Operations Research* 6, no. 1 (1981): 58–73; Eric Maskin and John Riley, “Monopoly with Incomplete Information,” *RAND Journal of Economics* 15, no. 2 (1984): 171–196; Jean-Jacques Laffont and David Martimort, *The Theory of Incentives: The Principal-Agent Model* (Princeton: Princeton University Press, 2001).

⁴John L. Kelly, Jr., “A New Interpretation of Information Rate,” *Bell System Technical Journal* 35, no. 4 (1956): 917–926.

2. Related Literature

2.1 Adverse Selection in Insurance and Reinsurance

The foundational model of competitive insurance markets under adverse selection is due to Rothschild and Stiglitz,⁵ who showed that a separating equilibrium, if it exists, involves high-risk types receiving full coverage and low-risk types receiving only partial coverage at actuarially fair rates. Stiglitz analyzed the monopoly case.⁶ Wilson proposed the anticipatory equilibrium concept,⁷ later given game-theoretic foundations by Mimra and Wambach,⁸ in which the Wilson–Miyazaki–Spence (WMS) allocation emerges as the second-best efficient outcome.

A recent literature applies these frameworks to reinsurance. Cheung, Yam, and Yuen employed a principal–agent model to study a monopolistic reinsurance market with adverse selection, modeling each insurer’s risk preference via Value-at-Risk (VaR) consistent with Solvency II.⁹ A subsequent paper extended this to concave distortion risk measures,¹⁰ and the continuum-of-types case has been solved.¹¹

2.2 The Kelly Criterion and Growth-Optimal Strategies

Kelly showed that a bettor¹² facing a sequence of favorable wagers maximizes the long-run geometric growth rate of wealth by maximizing the expected logarithm of the return. The criterion has been extensively applied to portfolio optimization.¹³ The Kelly criterion penalizes both over-exposure (errors of commission leading to surplus destruction) and under-exposure (errors of omission forgoing compounding opportunities).

⁵Rothschild and Stiglitz, “Equilibrium in Competitive Insurance Markets.”

⁶Stiglitz, “Monopoly, Non-Linear Pricing and Imperfect Information.”

⁷Charles Wilson, “A Model of Insurance Markets with Incomplete Information,” *Journal of Economic Theory* 16, no. 2 (1977): 167–207.

⁸Wanda Mimra and Achim Wambach, “A Game-Theoretic Foundation for the Wilson Equilibrium in Competitive Insurance Markets with Adverse Selection” (CESifo Working Paper No. 3412, 2011).

⁹Ka Chun Cheung, Sheung Chi Phillip Yam, and Fei Lung Yuen, “Reinsurance Contract Design with Adverse Selection,” *Scandinavian Actuarial Journal* 2019, no. 9 (2019): 784–798.

¹⁰Ka Chun Cheung, Sheung Chi Phillip Yam, Fei Lung Yuen, and Yiyang Zhang, “Concave Distortion Risk Minimizing Reinsurance Design under Adverse Selection,” *Insurance: Mathematics and Economics* 91 (2020): 155–165.

¹¹Ka Chun Cheung, Sheung Chi Phillip Yam, Fei Lung Yuen, and Yiyang Zhang, “Optimal Design of Reinsurance Contracts under Adverse Selection with a Continuum of Types” (arXiv preprint arXiv:2504.17468, 2025).

¹²Kelly, “A New Interpretation of Information Rate.”

¹³Edward O. Thorp, “The Kelly Criterion in Blackjack, Sports Betting, and the Stock Market,” in *Handbook of Asset and Liability Management*, vol. 1 (Amsterdam: North-Holland, 2006), 385–428; Leonard C. MacLean, Edward O. Thorp, and William T. Ziemba, eds., *The Kelly Capital Growth Investment Criterion: Theory and Practice* (Singapore: World Scientific, 2011).

2.3 Mechanism Design and Screening

The revelation principle¹⁴ establishes that the reinsurer can, without loss of generality, design a menu of contracts indexed by the cedant's type such that each type self-selects into its designated contract. The relevant constraints are incentive compatibility (IC)—no type prefers a contract designed for a different type—and individual rationality (IR)—each type weakly prefers its assigned contract to the outside option.

¹⁴Myerson, "Optimal Auction Design."

3. Model Setup

3.1 Agents and Types

A reinsurer (principal, R) receives an underwriting submission from a primary insurer (agent, I). The insurer possesses private information about its type $\theta \in \Theta = \{\theta_L, \theta_H\}$, where θ_L denotes a low-risk cedant and θ_H denotes a high-risk cedant. The reinsurer holds a prior belief that the proportion of low-risk types is $\lambda \in (0, 1)$.

Each type θ presents a loss distribution characterized by expected loss ratio $\mu(\theta)$ and volatility $\sigma(\theta)$, with $\mu(\theta_L) < \mu(\theta_H)$ and $\sigma(\theta_L) \leq \sigma(\theta_H)$.

3.2 Contract Space

A reinsurance contract is characterized by a triple $C = (\alpha, d, \pi)$, where $\alpha \in [0, 1]$ is the cession rate, $d \geq 0$ is the retention (attachment point), and $\pi > 0$ is the reinsurance premium. A quota share is $C = (\alpha, 0, \pi)$. An excess-of-loss contract is $C = (1, d, \pi)$.

3.3 Reinsurer's Objective: Kelly-Criterion Growth Maximization

Let S_t denote the reinsurer's surplus at time t , and let r_f denote the risk-free rate (e.g., the 3-month U.S. Treasury bill yield). Upon accepting a reinsurance contract C for a type- θ cedant, the reinsurer's absolute underwriting profit is $r(C, \theta) = \pi - L(C, \theta)$, where $L(C, \theta)$ is the realized indemnity payment. The fractional allocation is $f = \pi/S_t$. The fractional excess return is $R(C, \theta) = [\pi - L(C, \theta)]/\pi - r_f$.

The Kelly criterion directs the reinsurer to maximize:

$$G = \mathbb{E}[\ln(1 + f \cdot R(C, \theta))] \quad (1)$$

For small exposures, the second-order Taylor expansion yields:

$$G \approx f \cdot m - \frac{f^2}{2} \cdot v \quad (2)$$

where $m = \mathbb{E}[R]$ is the expected excess profit margin and $v = \mathbb{E}[R^2] = \text{Var}[R] + m^2$. The Kelly-optimal fraction is $f^* = m/v$, and $G^* = m^2/(2v)$.

4. Optimal Menu of Contracts under Adverse Selection

4.1 The Screening Problem

The reinsurer designs a menu $\{C_L, C_H\}$, one contract per type, subject to IC and IR constraints:

$$U(\theta_L, C_L) \geq U(\theta_L, C_H) \quad (\text{IC}_L)$$

$$U(\theta_H, C_H) \geq U(\theta_H, C_L) \quad (\text{IC}_H)$$

$$U(\theta_L, C_L) \geq U_0(\theta_L) \quad (\text{IR}_L)$$

$$U(\theta_H, C_H) \geq U_0(\theta_H) \quad (\text{IR}_H)$$

4.2 Kelly-Optimal Separating Menu

The reinsurer maximizes:

$$\max \lambda \cdot G(C_L, \theta_L) + (1 - \lambda) \cdot G(C_H, \theta_H) \quad \text{subject to } \text{IC}_L, \text{IC}_H, \text{IR}_L, \text{IR}_H \quad (3)$$

where $G(C, \theta) = f(C) \cdot m(C, \theta) - \frac{1}{2}f(C)^2 \cdot v(C, \theta)$.

Proposition 1 (Kelly-Optimal Separating Menu). *Under the single-crossing property:*
 (a) *The high-risk contract C_H^* offers the Kelly-optimal coverage level at the first-best level (efficiency at the top).* (b) *The low-risk contract C_L^* offers distorted coverage $\alpha_L^* < \alpha_L^{fb}$.*
 (c) *The reinsurer's expected growth rate under the menu strictly exceeds that under any single take-it-or-leave-it contract, provided $\lambda \in (0, 1)$.*

4.3 Application Across Reinsurance Structures

Proportional reinsurance. The menu varies the cession rate α and premium rate. The low-risk cedant receives a lower cession rate at a higher per-unit rate, with reduced coverage serving as the incentive-compatible distortion.

Non-proportional reinsurance. The menu varies the attachment point d and premium π . The high-risk cedant receives a lower attachment at a higher premium. The low-risk cedant receives a higher attachment at a lower premium.

P&C versus L&H. The single-crossing property holds in both settings.

5. Mitigating Errors of Commission and Omission

Error of commission. The menu mitigates commission errors through self-selection: C_H^* is priced for the high-risk distribution, and IC_H ensures high-risk cedants choose it.

Error of omission. The menu offers C_L^* specifically designed to attract low-risk cedants. Although coverage is distorted, the low-risk cedant still obtains strictly more surplus than under the outside option.

Information rent. The high-risk type earns a surplus beyond its outside option. The Kelly growth gap $\Delta G = G^{fb} - G^{sb}$ is strictly less than the growth destruction from either pooling or exclusion.

6. Competitive Dynamics: Menu vs. Take-It-or-Leave-It Rivals

Theorem 1 (Kelly-Growth Superiority of the Menu). *Let G^M denote the expected Kelly growth rate under the optimal separating menu, and G^T the rate under any single take-it-or-leave-it contract. Under the standard two-type adverse selection model with single-crossing and the Kelly objective: $G^M > G^T$ for all $\lambda \in (0, 1)$. The advantage ΔG is maximized at an interior $\hat{\lambda}$ and vanishes only in the degenerate cases $\lambda \rightarrow 0$ or $\lambda \rightarrow 1$.*

Proof. For any single pooling contract with premium π_P , the portfolio return variance decomposes as $v_{\text{pool}} = v_{\text{within}} + v_{\text{between}}$, where $v_{\text{within}} = [\lambda\sigma_L^2 + (1-\lambda)\sigma_H^2]/\pi_P^2$ is the within-type component and $v_{\text{between}} = \lambda(1-\lambda)(\mu_L - \mu_H)^2/\pi_P^2$ is the between-type component. The pooling Kelly growth rate is $G^T = \tilde{f} \cdot m_P - \frac{1}{2}\tilde{f}^2 \cdot (v_{\text{within}} + v_{\text{between}})$. Two effects follow, both at the level of approximation (3). First, for any common fraction $\tilde{f} > 0$, removing $v_{\text{between}} > 0$ (strict for $\lambda \in (0, 1)$ and $\mu_L \neq \mu_H$) raises the growth objective. Second, the menu prices and sizes each type separately: since the maximum of a sum is at most the sum of the maxima, $\max_f [\lambda G_L(f) + (1-\lambda)G_H(f)] \leq \lambda \max_f G_L(f) + (1-\lambda) \max_f G_H(f)$, and the type-specific premiums weakly enlarge each per-type margin relative to a single blended premium. Both effects favor the menu, the first strictly, so $G^M > G^T$. $\square$

7. Equilibrium Existence under Capacity Constraints

A classical concern with separating menus in competitive settings is non-existence: when the low-risk share is high, a pooling deviation can attract both types and break the separating allocation.¹⁵ Capacity constraints—natural in solvency-regulated reinsurance, where premium writings are bounded by surplus—limit the scale of any such deviation and stabilize the separating menu.

Proposition 2 (Equilibrium Existence). *Under the Kelly-criterion objective with endogenous capacity constraints (aggregate portfolio Kelly fraction bounded above), the separating menu $\{C_L^*, C_H^*\}$ constitutes a Nash equilibrium for all $\lambda \in (0, 1)$.*

¹⁵Rothschild and Stiglitz, “Equilibrium in Competitive Insurance Markets.”

8. Computational Model and Monte Carlo Validation

This section describes a Monte Carlo simulation that validates the theoretical predictions. All assumptions are stated explicitly.

8.1 Assumptions

Assumption 1 (Loss Distribution). $L(\theta) \sim \max\{0, \mathcal{N}(\mu(\theta), \sigma^2(\theta))\}$. We set $\mu(\theta_L) = 0.55$, $\mu(\theta_H) = 0.80$, $\sigma(\theta_L) = 0.12$, $\sigma(\theta_H) = 0.25$.

Assumption 2 (Population). $\lambda = 0.60$; $N = 50$ submissions per period.

Assumption 3 (Risk-Free Rate). $r_f = 4.5\%$ (3-month U.S. Treasury bill).

Assumption 4 (Loadings). 15% (high-risk), 10% (low-risk), 12% (pooling).

Assumption 5 (Coverage Distortion). $\alpha_L = 0.80$.

Assumption 6 (Kelly Sizing). $f^* = m/v$; per-submission allocation $= f^*/N$.

Assumption 7 (Independence). Losses are i.i.d. across submissions and periods.

Assumption 8 (Horizon). 10,000 paths, $T = 100$ periods, $S_0 = 100$, seed = 42.

8.2 Three Competing Strategies

Strategy 1: Kelly-Optimal Separating Menu. Low-risk: $C_L = (0.80, 0.605)$, $f_L = 0.933$, $m_L = 0.046$. High-risk: $C_H = (1.0, 0.920)$, $f_H = 1.157$, $m_H = 0.085$. $G^M = 0.0326$.

Strategy 2: Pooling. $\pi_P = 0.728$, $f_P = 0.677$. The between-type loss-ratio variance $\lambda(1-\lambda)(\mu_L - \mu_H)^2 = 0.015$ inflates the return variance to $v_{\text{pool}} = 0.092$ (versus $v_L = 0.039$ and $v_H = 0.074$ for the separated types). $G^T = 0.0210$.

Strategy 3: Conservative. $\pi = 0.920$. Only 40% of submissions written. $G^C = 0.0198$.

Table 1: Contract parameters by strategy at the baseline parameterization. Columns report the premium rate π , the excess profit margin m (net of $r_f = 4.5\%$, per unit premium), the return second moment v , the Kelly fraction $f^* = m/v$, the coverage rate α , and the per-type Kelly growth contribution G (per year). Portfolio $G = \lambda G_L + (1 - \lambda)G_H = 0.0326$ for the menu; the conservative strategy's portfolio value is $(1 - \lambda) \times G_H = 0.0198$. Source: closed-form calculation (workbook tab Contract Parameters).

Contract	π	m	v	f^*	α	G
Menu C_L	0.605	0.046	0.039	0.933	0.800	0.0214
Menu C_H	0.920	0.085	0.074	1.157	1.000	0.0494
Pooling	0.728	0.062	0.092	0.677	—	0.0210
Conservative	0.920	0.085	0.074	1.157	—	0.0198

8.3 Results

Table 1 exhibits the ordering $G^M = 0.0326 > G^T = 0.0210 > G^C = 0.0198$, consistent with Theorem 1.




Figure 1: Surplus growth trajectories by strategy. Horizontal axis: underwriting period (years); vertical axis: surplus in units of initial surplus $S_0 = 100$, logarithmic scale. Solid, dashed, and dotted lines mark the path-wise median for the menu (deep blue), pooling (bright blue), and conservative (amber) strategies; shaded bands span the 10th–90th percentile range in each strategy’s pastel tint. Source: `reinsurance_sim_v2.py` Monte Carlo, seed 42, 10,000 paths, 50 submissions per period, 100 periods; summary statistics in Table 2 and the workbook tab Monte Carlo Results.

The menu strategy reaches a median terminal surplus of 52,518 (in units of $S_0 = 100$). Pooling reaches 23,708—less than half the menu’s terminal surplus. The conservative strategy reaches 46,455.




Figure 2: Distribution of realized annualized growth rates across 10,000 Monte Carlo paths, one panel per strategy with common bins and axes. Horizontal axis: annualized geometric growth rate (dimensionless, per year); vertical axis: path count. Bars are filled in each strategy’s pastel tint with saturated edges; the dashed vertical line marks the median, annotated in each panel. Source: `reinsurance_sim_v2.py`, seed 42.

The menu achieves a median annualized growth rate of 6.26%, the conservative 6.14%, and pooling 5.47%.

Table 2: Monte Carlo summary statistics by strategy: 10,000 paths, $N = 50$ submissions per period, $T = 100$ periods, initial surplus $S_0 = 100$, seed 42. Growth rates are annualized geometric rates (dimensionless, per year); terminal-surplus rows are in units of S_0 . Source: `reinsurance_sim_v2.py`; mirrored in the workbook tab Monte Carlo Results.

Metric	Menu	Pooling	Conserv.
Median annualized growth	0.0626	0.0547	0.0614
Mean annualized growth	0.0627	0.0547	0.0614
Std dev of growth	0.0033	0.0028	0.0027
Median terminal surplus	52,518	23,708	46,455
10th pctl terminal surplus	34,609	16,542	32,899
90th pctl terminal surplus	81,122	33,617	65,652




Figure 3: Kelly growth as a function of the low-risk population share λ . Left panel: closed-form growth rate G (per year) for the menu (deep blue, solid), pooling (bright blue, dashed), and conservative (amber, dotted) strategies over 60 values of λ . Right panel: the menu's growth advantage ΔG over pooling (bright blue, solid) and over the conservative strategy (amber, dashed); the gray reference line marks zero. Source: reinsurance_sim_v2.py closed-form evaluation; baseline parameters otherwise.




Figure 4: The separating menu in coverage–premium space. Horizontal axis: coverage rate α (cession fraction); vertical axis: premium rate π . Zero-profit lines $\pi = \alpha\mu$ are drawn for the low-risk type (green, solid), high-risk type (maroon, dashed), and pooled population (bright blue, dotted). Diamonds mark first-best contracts at full coverage with a 12 percent loading; circles mark the second-best menu contracts; the dashed gray double arrow marks the information-rent distortion that moves C_L^* away from full coverage. Labels sit adjacent to their markers. Source: reinsurance_sim_v2.py contract design at baseline parameters (workbook tab Contract Parameters).




Figure 5: Optimal Kelly fractions by contract type. Vertical axis: Kelly fraction $f^* = m/v$ as a share of surplus, bars from a zero baseline with value annotations; the menu's low-risk bar is shown after the $\alpha_L = 0.80$ coverage distortion. Fills use each contract's pastel tint with saturated edges: low-risk (green), high-risk (maroon), pooling (bright blue), conservative (amber). The pooling fraction is depressed by the between-type variance. Source: reinsurance_sim_v2.py at baseline parameters.

The between-type loss-ratio variance $\lambda(1 - \lambda)(\mu_L - \mu_H)^2 = 0.015$, which enters the pooling return variance as $v_{\text{between}} = 0.015/\pi_P^2 = 0.028$, is the quantitative mechanism underlying Theorem 1.

9. Market Cycle Analysis: Hardening and Softening

The reinsurance market is cyclical. This section extends the baseline model to analyze how the Kelly-optimal separating menu adjusts across the cycle.

9.1 Market Hardness Parameterization

We introduce $h \in [-1, 1]$, where $h > 0$ denotes hardening and $h < 0$ denotes softening. The parameter shifts four primitives simultaneously:

(i) **Loading factors:** $\delta_H(h) = \max\{2\%, 15\% + 10h\}$, $\delta_L(h) = \max\{2\%, 10\% + 8h\}$; the pooling contract carries $\delta_P(h) = \max\{2\%, 12\% + 9h\}$.

(ii) **Population composition:** $\lambda(h) = 0.60 + 0.15h$.

(iii) **Submission volume:** $N(h) = 50 - 20h$.

(iv) **Coverage distortion:** $\alpha_L(h) = \min\{0.95, \max\{0.60, 0.80 + 0.15h\}\}$.

9.2 Theoretical Growth Rates across the Cycle

Table 3: Menu performance across the market cycle at five hardness levels h . Columns report the market phase, the low-risk share $\lambda(h)$, submissions per period $N(h)$, the coverage distortion $\alpha_L(h)$, closed-form Kelly growth rates (per year) for the three strategies, and the menu's advantage over pooling ΔG . ΔG is computed from unrounded growth rates and may differ in the last digit from subtraction of the displayed columns. Source: `market_cycle_sim.py` (workbook tab Market Cycle).

h	Mkt	λ	N	α_L	G_{menu}	G_{pool}	G_{cons}	ΔG
-0.6	Soft	0.51	62	0.710	0.0043	0.0014	0.0042	+0.0029
-0.3	Soft	0.55	56	0.755	0.0144	0.0082	0.0110	+0.0062
0.0	Mid	0.60	50	0.800	0.0326	0.0210	0.0198	+0.0116
+0.3	Hard	0.65	44	0.845	0.0601	0.0405	0.0293	+0.0196
+0.6	Hard	0.69	38	0.890	0.0983	0.0677	0.0384	+0.0305




Figure 6: Kelly growth across the market cycle. Horizontal axes: market hardness h (negative = soft, positive = hard), dotted vertical at mid-cycle $h = 0$. Left panel: closed-form growth rate G (per year) for the menu (deep blue, solid), pooling (bright blue, dashed), and conservative (amber, dotted) strategies on an 80-point hardness grid. Right panel: the menu's advantage ΔG over pooling and over the conservative strategy. Source: market_cycle_sim.py; parameter schedules per Section 9.1 (workbook tab Market Cycle).

9.3 Menu Adjustment Mechanics




Figure 7: Menu parameter adjustments across the market cycle, six panels on a common hardness axis h with dotted verticals at $h = 0$. Panels show premium loadings (percent) for the low-risk (green, solid) and high-risk (maroon, dashed) contracts; premium rates π against dotted reference lines at the expected loss ratios; the low-risk coverage distortion α_L (deep blue); Kelly fractions f^* for the two menu contracts and the pooling contract (bright blue, dotted); the low-risk population share λ (amber, percent); and submission volume N (gray). Source: market_cycle_sim.py parameter schedules per Section 9.1.

Hardening: raise loadings, relax coverage distortion, increase Kelly fractions, benefit from improved risk pool.

Softening: compress loadings, tighten coverage distortion, reduce Kelly fractions, accept deteriorated risk pool.

9.4 Simulated Surplus Trajectories across Regimes




Figure 8: Surplus trajectories by market regime: soft ($h = -0.6$), mid-cycle ($h = 0$), and hard ($h = +0.6$), on a common logarithmic surplus axis in units of $S_0 = 100$. Lines mark medians for the menu (deep blue, solid), pooling (bright blue, dashed), and conservative (amber, dotted) strategies; bands span the 10th–90th percentiles. Source: market_cycle_sim.py Monte Carlo, seed 42, 3,000 paths and 80 periods per regime.

9.5 Counter-Cyclical Properties

Proposition 3 (Counter-Cyclical Stability). (a) In softening markets, the menu’s growth rate declines less steeply than pooling’s in proportional terms: at $h = -0.6$ the menu retains 13 percent of its mid-cycle growth rate while pooling retains 7 percent. (b) In hardening markets, the menu’s growth rate increases by more than pooling’s in absolute terms. (c) The absolute advantage $\Delta G = G_{\text{menu}} - G_{\text{pool}}$ is monotonically increasing in h , while the ratio $G_{\text{menu}}/G_{\text{pool}}$ is monotonically decreasing in h (from $3.1\times$ at $h = -0.6$ to $1.5\times$ at $h = +0.6$): the menu’s relative advantage is greatest precisely where growth is scarcest, in soft markets.

10. Extension: Continuum of Types

The two-type model of Sections 3–8 admits a natural generalization to a continuum of types $\theta \sim F(\theta)$ on $[\underline{\theta}, \bar{\theta}]$, where θ represents the cedant’s expected loss ratio and higher θ denotes worse risk. The reinsurer designs a schedule of contracts $\{C(\theta)\}_{\theta \in \Theta}$, one for each announced type, specifying coverage $\alpha(\theta)$ and premium $\pi(\theta)$. This follows Cheung et al., who solved the continuum case for a monopolistic reinsurer.¹⁶

10.1 Optimal Schedule under IC

Incentive compatibility together with single crossing implies a monotonicity condition: coverage $\alpha(\theta)$ must be non-decreasing in θ (worse types receive more coverage). The envelope theorem then pins down the premium schedule, up to a constant, as the integral of the coverage schedule. The Kelly-optimal schedule preserves the qualitative features of the two-type model:

- **Efficiency at the top:** The highest-risk type $\bar{\theta}$ receives its first-best Kelly-optimal coverage: $\alpha(\bar{\theta}) = 1$.
- **Downward distortion:** For all $\theta < \bar{\theta}$, coverage is distorted below the first-best level: $\alpha(\theta) < \alpha^{fb}(\theta) = 1$. The distortion increases monotonically as θ moves away from $\bar{\theta}$ —the best risks receive the most distorted coverage.
- **Information rent:** The growth gap between first-best and second-best is concentrated at the low- θ end, where the distortion is largest.

Figure 9 displays the first-best and second-best schedules. The coverage schedule (left panel) shows the linear distortion: $\alpha(\theta) = 1 - \kappa \cdot (\bar{\theta} - \theta)/(\bar{\theta} - \underline{\theta})$ where $\kappa = 0.40$, calibrated so that the midpoint of the type range receives $\alpha = 0.80$, matching the two-type baseline. The market loading schedule (center panel) interpolates from 10 percent at $\underline{\theta}$ to 15 percent at $\bar{\theta}$, matching the two-type baseline’s loadings; the loading is set by market conditions and individual rationality rather than optimized, since the Kelly growth rate is monotonically increasing in the premium and an unconstrained loading choice would be unbounded. The growth panel (right) shows the information rent cost as the shaded area between the two curves.

¹⁶Cheung et al., “Optimal Design of Reinsurance Contracts under Adverse Selection with a Continuum of Types.”




Figure 9: Continuum of types: first-best vs. second-best schedules over θ uniform on $[0.40, 0.90]$. Left panel: coverage $\alpha(\theta)$, first-best (deep blue, solid) at full coverage and second-best (bright blue, dashed) with the linear IC distortion $\alpha(\theta) = 1 - 0.40(\bar{\theta} - \theta)/(\bar{\theta} - \underline{\theta})$. Center panel: the market loading schedule $\delta(\theta)$ from 10 to 15 percent. Right panel: per-type Kelly growth $G(\theta)$ (per year), with the information-rent cost shaded in the second-best pastel tint. Source: extensions_sim.corrected_continuum_analysis(); workbook tab Extensions.

10.2 Optimal Number of Risk Groups

A question raised by Cheung et al., who find a two-group classification optimal in their setting,¹⁷ is how much growth a finite menu of K risk groups captures relative to the continuous second-best schedule. Within each group, all types receive the same contract (pooled). Marginal gains diminish with K , while operational complexity increases; the rate of diminishment determines how coarse a menu the reinsurer can afford.

Figure 10 reports the portfolio Kelly growth rate under optimally partitioned K -group menus. Each group's contract is priced at the group-average cost with the market loading, carries the IC coverage distortion evaluated at the group midpoint, and faces an individual-rationality constraint: a type declines the contract when the pooled premium exceeds its willingness to pay (its standalone fair premium plus a 5 percent switching tolerance). The group return variance includes both the average within-type variance and the across-type dispersion of expected losses within the group—the between-type variance of Theorem 1, now operating within each group.

¹⁷Cheung et al., “Optimal Design of Reinsurance Contracts under Adverse Selection with a Continuum of Types.”

Three results emerge. First, growth is monotonically increasing in K : G rises from 0.00548 at $K = 1$ (pooling) to 0.02069 at $K = 15$, converging to the continuous second-best ($G = 0.02073$) from below and remaining below the first-best ($G = 0.02410$) throughout. Second, the first split delivers the largest single incremental gain: the optimal two-group menu roughly doubles the pooling growth rate ($G = 0.01052$, a 92 percent increase), with diminishing marginal gains thereafter. Third, a material gap to finer menus remains: $K = 2$ captures approximately one-third of the pooling-to-continuous-second-best gain, and convergence is nearly complete only by $K \approx 8$. The two-type baseline of Sections 3–8 is exact when the population is genuinely bimodal (two discrete types); under a uniform continuum, coarse menus forgo growth through two channels—residual within-group between-type variance and participation losses when the pooled price exceeds low types’ willingness to pay.




Figure 10: Portfolio Kelly growth (per year) under optimally partitioned K -group menus of the type continuum. Markers on the deep-blue line show $K \in \{1, 2, 3, 5, 8, 15\}$ with grid-searched boundaries for $K \leq 5$ and equal-width groups for $K \geq 8$; each group’s contract carries the IC coverage distortion at the group midpoint, prices at group-average cost with the market loading, includes the within-group between-type variance, and enforces IR participation with a 5 percent switching tolerance. The dashed green and dotted maroon reference lines mark the continuous first-best and second-best. The annotation reports the $K = 2$ result. Source: extensions_sim.corrected_continuum_analysis(); workbook tab Extensions.

11. Extension: Dynamic Menus with Bayesian Experience Rating

The static model of Sections 3–8 assumes the reinsurer’s belief λ is fixed. In practice, the reinsurer observes the cedant’s loss experience over time and updates its assessment. This section develops the dynamic extension in which the menu is re-optimized each period based on a Bayesian posterior.

11.1 Bayesian Updating Mechanism

At period $t = 0$, the reinsurer holds a prior belief $\lambda_0 = P(\theta = \theta_L)$. After observing the cedant’s loss realization L_t in period t , the reinsurer updates via Bayes’ rule:

$$\lambda_{t+1} = \frac{\lambda_t \cdot f(L_t | \theta_L)}{\lambda_t \cdot f(L_t | \theta_L) + (1 - \lambda_t) \cdot f(L_t | \theta_H)} \quad (4)$$

where $f(L | \theta)$ is the Normal density with mean $\mu(\theta)$ and standard deviation $\sigma(\theta)$. The menu at period $t+1$ is then re-optimized using the posterior λ_{t+1} . Specifically, the coverage reduction on Contract A adjusts with the posterior: $\alpha_L(\lambda) = \min\{0.95, \max\{0.60, 0.50 + 0.50\lambda\}\}$. As λ rises (the reinsurer becomes more confident the cedant is low-risk), it offers more generous terms; as λ falls, terms tighten.

11.2 Belief Convergence

Figure 11 displays the posterior belief evolution for two scenarios: a cedant whose true type is low-risk (left panel) and a cedant whose true type is high-risk (right panel). Starting from a diffuse prior of $\lambda_0 = 0.50$, the posterior converges to the true type within approximately 5–10 periods. Convergence is faster for high-risk types because a typical high-risk loss lies far in the thin upper tail of the low-risk distribution, so the likelihood ratio against θ_L is extreme; a typical low-risk loss, by contrast, retains non-negligible likelihood under the wider high-risk distribution, so evidence for θ_L accumulates more gradually.




Figure 11: Bayesian belief evolution by true type over 30 periods from the prior $\lambda_0 = 0.50$ (dotted gray reference). Left panel: cedants whose true type is low-risk (green); right panel: high-risk (maroon); solid lines mark median posteriors, bands the 10th–90th percentiles, and the dashed line the true-type limit ($\lambda \rightarrow 1$ left, $\lambda \rightarrow 0$ right). Common vertical axis: posterior $\lambda = P(\theta = \theta_L)$. Source: `extensions_sim.DynamicBayesianModel`, seed 42, 3,000 paths per panel.

11.3 The Ratchet Effect

The dynamic menu creates a ratchet: cedants with favorable loss experience gain access to more generous terms (higher α_L , approaching full coverage), while those with adverse experience face tighter terms (lower α_L , approaching the minimum). Figure 12 displays this relationship. The left panel shows how Contract A’s coverage rate adjusts with the posterior belief. The right panel shows the corresponding Kelly growth rate.




Figure 12: The ratchet effect under Bayesian learning. Left panel: Contract A's coverage rate $\alpha_L(\lambda) = \text{clip}(0.50 + 0.50\lambda, 0.60, 0.95)$ as a function of the posterior belief λ ; favorable experience raises coverage toward 0.95, adverse experience lowers it toward 0.60. Right panel: the portfolio Kelly growth rate (per year) at each belief; growth declines in λ because a book assessed as predominantly low-risk carries lower margins. Dotted verticals mark the prior $\lambda_0 = 0.50$. Source: extensions_sim.DynamicBayesianModel.

11.4 Growth Advantage of Dynamic Menus

The dynamic menu generates a persistent growth advantage over the static menu. Figure 13 compares surplus paths under the dynamic menu (Bayesian updating) and the static menu (fixed $\lambda_0 = 0.50$). The dynamic menu achieves a median annualized growth rate of 3.99% versus 3.95% for the static menu—an advantage of +0.04 percentage points per year. While modest in annual terms, this advantage compounds: over 30 years, it translates to a 1.3% higher terminal surplus.

The advantage arises because the dynamic menu converges toward the first-best allocation as belief converges to the truth: for low-risk cedants, α_L rises toward 0.95 (near full coverage), and for high-risk cedants, α_L falls toward 0.60 (tight screening). The static menu, locked at $\lambda_0 = 0.50$, distorts terms for both types relative to what they would receive under full information.




Figure 13: Dynamic vs. static menu: surplus over 30 periods in units of $S_0 = 100$. The dynamic menu (deep blue, solid) re-optimizes α_L each period from the Bayesian posterior; the static menu (bright blue, dashed) holds terms at the prior $\lambda_0 = 0.50$. Lines mark medians, bands the 10th–90th percentiles; the annotation reports median annualized growth for each and the difference. True population share of low-risk types: $\lambda = 0.60$. Source: extensions_sim.DynamicBayesianModel Monte Carlo, seed 42, 5,000 paths.

11.5 Learning Speed and Signal Quality

The speed of Bayesian convergence depends on the informativeness of the loss signal—the ratio of the type-separation ($\mu_H - \mu_L$) to the within-type noise (σ). Figure 14 displays median belief paths for a true low-risk cedant under three noise levels. At low noise ($\sigma/2$), the reinsurer identifies the type within 3 periods. At high noise (2σ), convergence takes 15–20 periods. This has practical implications: for lines with high inherent volatility (catastrophe-exposed property), Bayesian updating is slow and the static menu may be near-optimal; for lines with low volatility (workers’ compensation, commercial auto), updating is fast and the dynamic menu provides a material advantage.




Figure 14: Learning speed by signal noise for a true low-risk cedant. Vertical axis: median posterior $\lambda = P(\theta_L)$ across 2,000 paths from the prior $\lambda_0 = 0.50$ (dotted gray reference); horizontal axis: period. Loss volatilities are scaled to one-half (deep blue, solid), one (bright blue, dashed), and two times (amber, dotted) the baseline σ . Lower noise yields faster type identification. Source: extensions_sim.DynamicBayesianModel, seed 42.

12. Extension: Multi-Dimensional Screening

In the baseline model, the cedant's private information is one-dimensional: the expected loss ratio θ . In practice, a P&C cedant possesses private knowledge along at least two dimensions: claim *frequency* (how often losses occur) and claim *severity* (how large each loss is). A cedant with low frequency but high severity (e.g., a catastrophe-exposed property book) presents a different risk profile from one with high frequency but low severity (e.g., a commercial auto book), even if both have the same expected loss ratio. A one-dimensional screen conflates these types; a two-dimensional screen separates them.

12.1 Four-Type Model

The type space is $(f, s) \in \{\text{low}, \text{high}\} \times \{\text{low}, \text{high}\}$, yielding four types: LL (low frequency, low severity; $\mu = 0.30$), LH (low frequency, high severity; $\mu = 0.60$), HL (high frequency, low severity; $\mu = 0.75$), and HH (high frequency, high severity; $\mu = 1.50$). The reinsurer designs a menu of four contracts, one per type. Two screening instruments are available for separation: the attachment point d screens on severity (cedants with heavy tails prefer lower attachments), and the cession rate α screens on frequency (cedants with high frequency prefer higher cession). The computational model assigns each type its designated contract and verifies the growth accounting; an explicit incentive-compatibility construction over the two instruments is left for future work.

12.2 Results

Figure 15 displays the type space (Panel A) and the growth comparison (Panel B). The four-contract (2D) menu achieves $G = 0.0693$, compared to $G = 0.0405$ for the two-contract (1D) menu and $G = 0.0114$ for pooling. The 2D screening advantage over 1D is +288 basis points—a gain of 71%. This gain arises because the 1D menu, sorting on the aggregate loss ratio alone, pools LL with LH ($\mu = 0.30$ with 0.60) and HL with HH ($\mu = 0.75$ with 1.50), leaving large within-group dispersion of expected losses; the 2D menu separates all four types.




Figure 15: Multi-dimensional screening on frequency and severity. Left panel: the four types in frequency–severity space with $\mu = f \times s$; markers are filled in each type’s pastel tint (LL green, LH amber, HL bright blue, HH maroon) with labels adjacent to markers reporting μ and population share. Right panel: portfolio Kelly growth (per year, bars from a zero baseline with value annotations) for the four-contract two-dimensional menu (deep blue), the two-contract one-dimensional menu (bright blue), and pooling (amber). Source: extensions_cd_sim.MultiDimModel; workbook tab Extensions.

The practical implication: reinsurers with access to granular submission data (frequency and severity separately, rather than a blended loss ratio) can design richer menus that extract more information. The additional screening dimension is operationalized through the contract structure: varying both the attachment point (severity screen) and the cession rate (frequency screen) simultaneously.

13. Extension: Heavy-Tailed Losses and Correlated Portfolios

The baseline model assumes Normally distributed losses (Assumption 1) and independence across submissions (Assumption 7). Insurance losses are heavy-tailed, and within a portfolio, losses may be correlated through common factors (economic conditions, legal environment, catastrophe exposure). This section relaxes both assumptions.

13.1 Heavy-Tailed Loss Distributions

We compare three loss distributions, all calibrated to the same mean loss ratios ($\mu_L = 0.55$, $\mu_H = 0.80$) but with different tail behavior:

- **Normal** (baseline): symmetric, thin-tailed. Variance: $\sigma_L = 0.12$, $\sigma_H = 0.25$.
- **Log-Normal**: right-skewed, moderately heavy-tailed. Calibrated to match the Normal mean and variance.
- **Pareto-mixed**: 95% Normal body + 5% Pareto tail ($\alpha = 2.5$). Generates realistic extreme losses. Effective standard deviation is $2\text{--}3\times$ the Normal level.

The Kelly fractions are recalibrated to each distribution's empirical variance. Figure 16 reports the results. Under Normal and Log-Normal losses, the menu's advantage over pooling is stable at approximately 80 basis points per year. Under Pareto-mixed losses, the advantage narrows to 17 basis points. The narrowing reflects two effects: (i) the heavy tails increase within-type variance, reducing Kelly fractions for both strategies; (ii) the between-type variance (which the menu eliminates) becomes a smaller fraction of total portfolio variance. The menu's advantage persists but is diluted by tail risk.




Figure 16: Heavy-tailed losses and the menu advantage. Left panel: the menu's growth advantage over pooling in basis points per year (bars from a zero baseline) under Normal (deep blue), Log-Normal (bright blue), and Pareto-mixed (amber) losses, with Kelly fractions recalibrated to each distribution's empirical variance and means held equal. Right panel: median surplus (log scale) under Pareto-mixed losses for the menu (deep blue, solid) and pooling (bright blue, dashed), with the Normal baseline as dotted pastel reference lines. Source: extensions_cd_sim.HeavyTailModel, seed 42, 3,000 paths, 60 periods.

13.2 Correlated Losses

We introduce a common-factor correlation structure: each submission's loss is $L_{ij} = \mu + \sigma(\sqrt{\rho} \cdot Z_t + \sqrt{1 - \rho} \cdot \varepsilon_{ij})$, where Z_t is a period-level common factor and ε_{ij} is idiosyncratic. The parameter $\rho \in [0, 1]$ controls the correlation between submissions within a period.

Figure 17 reports the results for $\rho \in \{0, 0.10, 0.20, 0.30, 0.50\}$. Both the menu and pooling growth rates decline with ρ (Panel A), reflecting the loss of diversification. The menu's advantage narrows from 80 basis points at $\rho = 0$ to 9 basis points at $\rho = 0.50$ (Panel B). The economic interpretation: when a large common factor drives losses, the between-type variance (which the menu eliminates) becomes a smaller share of total portfolio risk. At high ρ , systematic risk dominates, and the informational advantage of type separation is diluted. The menu's advantage persists at all correlation levels tested, but its magnitude is inversely related to the strength of the common factor.




Figure 17: Loss correlation and the menu advantage under the common-factor model $L_{ij} = \mu + \sigma(\sqrt{\rho} Z_t + \sqrt{1 - \rho} \varepsilon_{ij})$, which preserves the marginal loss variance at every ρ . Left panel: median annualized growth (percent per year) for the menu (deep blue, circles, solid) and pooling (bright blue, squares, dashed) at $\rho \in \{0, 0.10, 0.20, 0.30, 0.50\}$. Right panel: the menu's advantage in basis points per year, narrowing as the common factor strengthens. Source: extensions_cd_sim.HeavyTailModel, seed 42, 3,000 paths, 60 periods per ρ .

13.3 Joint Implications

Heavy tails and correlation both reduce the menu's growth advantage relative to the baseline, but neither eliminates it. The menu's advantage is largest when adverse selection is the dominant source of portfolio risk—that is, when within-type loss volatility is moderate, tails are thin, and correlation is low. As other risk sources (tail events, systematic factors) grow, the between-type variance that the menu eliminates becomes proportionally less consequential. The practical implication: the menu-of-contracts approach is most valuable for books where adverse selection is the primary underwriting concern (long-tail casualty, professional liability, specialty lines) and less differentiating for books dominated by catastrophe or systematic risk (property, financial guarantee).

14. Conclusion

We have developed a formal economic model of optimal reinsurance contract design in which a reinsurer uses a menu-of-contracts approach to screen cedant types under asymmetric information, with the objective of maximizing the long-run geometric growth rate of surplus (the Kelly criterion). The model delivers nine key results: (1) the optimal separating menu simultaneously mitigates errors of commission and omission; (2) the Kelly criterion balances growth and safety; (3) the menu reinsurer achieves strictly higher long-run growth than take-it-or-leave-it rivals (Theorem 1); (4) Monte Carlo simulation validates these predictions with $2.2\times$ the median terminal surplus of pooling; (5) the menu approach exhibits counter-cyclical stability across the entire market cycle (Proposition 3); (6) the continuum-of-types extension shows that the first split ($K = 2$) roughly doubles the pooling growth rate and delivers the largest single incremental gain, with diminishing returns to finer menus and near-complete convergence to the continuous second-best by $K \approx 8$; (7) the dynamic Bayesian extension shows that experience-rated menu updating generates a persistent growth advantage, with type identification converging within 5–10 periods; (8) two-dimensional screening along frequency and severity adds +288 basis points over one-dimensional screening; (9) the menu’s advantage persists under heavy-tailed losses and correlated portfolios, though its magnitude is inversely related to the prevalence of non-informational risk sources.

Bibliography

- Akerlof, George A. “The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism.” *Quarterly Journal of Economics* 84, no. 3 (1970): 488–500.
- Arrow, Kenneth J. “Uncertainty and the Welfare Economics of Medical Care.” *American Economic Review* 53, no. 5 (1963): 941–973.
- Asheim, Geir B., Thomas Tröger, and Achim Wambach. “Competitive Insurance Markets under Adverse Selection and Capacity Constraints.” *European Economic Review* 46, no. 10 (2001): 1735–1753.
- Borch, Karl. “An Attempt to Determine the Optimum Amount of Stop-Loss Reinsurance.” *Transactions of the 16th International Congress of Actuaries* 1 (1960): 597–610.
- Chade, Hector, and Edward Schlee. “Optimal Insurance with Adverse Selection.” *Theoretical Economics* 7, no. 3 (2012): 571–607.
- Cheung, Ka Chun, Sheung Chi Phillip Yam, and Fei Lung Yuen. “Reinsurance Contract Design with Adverse Selection.” *Scandinavian Actuarial Journal* 2019, no. 9 (2019): 784–798.
- Cheung, Ka Chun, Sheung Chi Phillip Yam, Fei Lung Yuen, and Yiyang Zhang. “Concave Distortion Risk Minimizing Reinsurance Design under Adverse Selection.” *Insurance: Mathematics and Economics* 91 (2020): 155–165.
- Cheung, Ka Chun, Sheung Chi Phillip Yam, Fei Lung Yuen, and Yiyang Zhang. “Optimal Design of Reinsurance Contracts under Adverse Selection with a Continuum of Types.” arXiv preprint arXiv:2504.17468, 2025.
- Crocker, Keith J., and Arthur Snow. “Multidimensional Screening in Insurance Markets with Adverse Selection.” *Journal of Risk and Insurance* 78, no. 2 (2011): 287–307.
- Doherty, Neil, and Kent Smetters. “Moral Hazard in Reinsurance Markets.” *Journal of Risk and Insurance* 72, no. 3 (2005): 375–391.
- Handel, Benjamin R., Igal Hendel, and Michael D. Whinston. “Multidimensional Screening and Menu Design in Health Insurance.” NBER Working Paper No. 30542, 2022.
- Kelly, John L., Jr. “A New Interpretation of Information Rate.” *Bell System Technical Journal* 35, no. 4 (1956): 917–926.
- Laffont, Jean-Jacques, and David Martimort. *The Theory of Incentives: The Principal-Agent Model*. Princeton: Princeton University Press, 2001.
- MacLean, Leonard C., Edward O. Thorp, and William T. Ziemba, eds. *The Kelly Capital Growth Investment Criterion: Theory and Practice*. Singapore: World Scientific, 2011.
- Maskin, Eric, and John Riley. “Monopoly with Incomplete Information.” *RAND Journal of Economics* 15, no. 2 (1984): 171–196.

- Mimra, Wanda, and Achim Wambach. “A Game-Theoretic Foundation for the Wilson Equilibrium in Competitive Insurance Markets with Adverse Selection.” CESifo Working Paper No. 3412, 2011.
- Myerson, Roger B. “Optimal Auction Design.” *Mathematics of Operations Research* 6, no. 1 (1981): 58–73.
- Rothschild, Michael, and Joseph E. Stiglitz. “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information.” *Quarterly Journal of Economics* 90, no. 4 (1976): 629–649.
- Stiglitz, Joseph E. “Monopoly, Non-Linear Pricing and Imperfect Information: The Insurance Market.” *Review of Economic Studies* 44, no. 3 (1977): 407–430.
- Thorp, Edward O. “The Kelly Criterion in Blackjack, Sports Betting, and the Stock Market.” In *Handbook of Asset and Liability Management*, vol. 1, 385–428. Amsterdam: North-Holland, 2006.
- Wilson, Charles. “A Model of Insurance Markets with Incomplete Information.” *Journal of Economic Theory* 16, no. 2 (1977): 167–207.