Optimal Menu-of-Contracts Design in Reinsurance

Frank Schmid
frank.schmid@frankschmid.com

This version: 4 November 2026
First version: 15 March 2026

Abstract. A cedant knows more about the risk it cedes than its reinsurer, which can quote a single price for full coverage or offer a menu from which the cedant chooses. This paper asks which design prospers when the two compete for the same cedants, and answers through the growth of surplus: capacity is a multiple of surplus, so the design with the higher expected logarithmic growth rate, the Kelly criterion's measure, compounds into the larger market share. In a market of two cedant types with mean–variance valuations, a menu reinsurer offering one cream-skimming contract beside the rivals' pooled contract earns a positive expected margin on every low-risk line whenever the pooled price exceeds the low-risk type's expected loss and the high-risk type gains from the pooled contract; at a competitive pooled price, wherever one of the two reinsurers can write all the business that comes to it, the rivals' expected underwriting result is negative and their growth rate below the risk-free rate; and where the growth gap is bounded away from zero, the menu reinsurer's share of industry surplus converges to one almost surely. At the calibration, at every share of a grid of ten-thousandths, the gap is positive and the rivals' growth rate below the risk-free rate; a Monte Carlo drawing every loss agrees with the quadrature, and simulated paths show the convergence and its dispersion. If the rivals reprice to their loss experience, the outcome turns on their loading: at a zero loading their price rises toward the high-risk type's cost, their growth rate falls below the risk-free rate, and a small book's noise lifts the price above that type's willingness to pay on some paths, closing the market; at a small positive loading they survive on high-risk business while being diluted. With a common loss factor and demand tied to the economy, at a share of 0.25, the factor widens the gap up to the calibrated capitalization and narrows it where capital is abundant, and the Kelly criterion sizes the menu reinsurer's book below capacity where losses are almost entirely common and on some simulated paths. On a continuum of types a menu partitions them by thresholds, and one contract skimming the better risks earns most of the gap where the rivals' capacity does not bind in every period; with private information on the mean and the volatility, a second contract screening on the volatility adds less than the first; and heavy-tailed losses move the gap only through their moments after the floor and the cap, against rivals pricing on the Normal's.

JEL classification: D82, D86, G22, C73.

This paper was developed with assistance from generative AI tools across the design, analysis, coding, and drafting stages. The author supervised all stages, verified the results, and takes full responsibility for the content.

Contents

Abbreviations and Acronyms **ii**

1 Introduction **1**

2 Related Literature **2**

3 The Market **4**
3.1 Cedants and Contracts 4
3.2 Reinsurers and Contracts Offered 5
3.3 Periods, Capacity, and Placement 5
3.4 Returns and the Objective 6
3.5 The Menu Reinsurer's Choice and the Rivals' Price 6
3.6 Calibration 6

4 Cream Skimming **7**

5 The Rivals' Book **9**

6 Selection **11**

7 The Growth Gap at the Calibration **11**
7.1 The Quadrature 12
7.2 The Gap over the Share 12
7.3 The Quadrature against a Monte Carlo 14

8 Share Paths under Fixed and Experience-Rated Pricing **14**
8.1 The Fixed Pooled Price 15
8.2 Experience Rating at a Zero Loading: The Rivals' Exit 16
8.3 Experience Rating at a Positive Loading 17
8.4 The Monopolist's Menu 20

9 Capacity, Correlation, and the Cycle **21**
9.1 The Common Factor, Demand, and Exposure 21
9.2 The Gap and the Exposure over Correlation and Capitalization 22
9.3 Surplus Paths with Demand Tied to the Economy 23

10 The Continuum of Types **24**
10.1 Types, Contracts, and the Menu 24
10.2 The Computation 25
10.3 Results 26

11 Two Dimensions of Private Information **28**

12 Heavy-Tailed Losses **31**

13 Conclusion **33**

Bibliography **35**

Abbreviations and Acronyms

ASTIN	Actuarial Studies in Non-Life Insurance (the ASTIN Bulletin)
CESifo	Center for Economic Studies and the ifo Institute
HH	high mean, high volatility (cedant type, Section 11)
HL	high mean, low volatility (cedant type, Section 11)
JEL	Journal of Economic Literature (classification system)
LH	low mean, high volatility (cedant type, Section 11)
LL	low mean, low volatility (cedant type, Section 11)
NBER	National Bureau of Economic Research
U.S.	United States
WMS	Wilson–Miyazaki–Spence

Greek Letters and Pronunciation

α	alpha (AL-fuh): coverage, the share of the loss ceded
Γ	Gamma, capital (GAM-uh): gross return on surplus
δ	delta (DEL-tuh): loading on the rivals' price
Δ	Delta, capital (DEL-tuh): difference of the two types' willingness to pay
ε	epsilon (EP-suh-lon): ruin floor; idiosyncratic shock (Section 9)
η	eta (AY-tuh): common loss factor (Section 9)
θ	theta (THAY-tuh): cedant type; expected loss ratio before the cap (Section 10)
κ	kappa (KAP-uh): capacity multiple of surplus
λ	lambda (LAM-duh): low-risk population share
μ	mu (MYOO): expected loss ratio
ν	nu (NOO): degrees of freedom of Student's t (Section 12)
ξ	xi (KSEE): martingale difference (Section 6)
π	pi (PIE): expected margin per unit of subject premium (Section 4)
ρ	rho (ROE): risk aversion
σ	sigma (SIG-muh): loss ratio volatility
τ	tau (TAW): period index of a sum (Section 6)
ϕ	phi (FIE): exposure, the share of capacity used (Section 9)

1 Introduction

A cedant knows more about the risk it cedes than the reinsurer that takes it. The reinsurer's response to that asymmetry is a design choice: quote one price for full coverage and take what comes, or offer a menu of coverage and price from which the cedant chooses. The first is the practice of the single-price underwriter; the second is the screening mechanism of Rothschild and Stiglitz,[1] in which a contract with less coverage at a lower price is taken by the better risks and left by the worse. This paper asks which of the two designs prospers when the two compete for the same cedants, and it answers through the one channel by which a reinsurer's design decides its future: the growth of its surplus. A reinsurer's capacity to write business is a multiple of its surplus, so a design that earns a higher expected logarithmic growth rate compounds into a larger share of the market's capacity. The criterion that ranks designs by that growth rate is Kelly's,[2] and the argument that the market comes to be held by the design with the higher rate is the market-selection argument of Blume and Easley.[3]

The model has two types of cedant, low-risk and high-risk, each with a mean–variance valuation of the coverage it buys, and two reinsurers: a menu reinsurer and a set of rivals that quote one price for full coverage, the pooled price of the population. Each period the cedants arrive in random order and each takes the contract it values most, a reinsurer at capacity writes the line that fits and sends the rest to the cedant's next-best contract, and each reinsurer's surplus compounds with its underwriting result. The menu reinsurer offers one contract beside the rivals' pooled contract: a cream-skimming contract that the low-risk type values exactly as it values the pooled contract and that the high-risk type values less. The paper establishes four things about this market. First, the cream-skimming contract exists and earns a positive expected margin on every low-risk line whenever the pooled price exceeds the low-risk type's expected loss and the high-risk type gains from the pooled contract; the price exceeds that loss whenever the pool contains high-risk business (Theorem 1 and Section 4). Second, the rivals' book is left with the high-risk business and the low-risk business the menu reinsurer cannot absorb, and at the pooled price of a competitive market, wherever the rivals can write every submission that comes to them or the menu reinsurer every low-risk one, their expected underwriting result is negative and their expected growth rate lies below the risk-free rate (Propositions 1 and 2); at the calibration their growth rate lies below the risk-free rate at every share computed (Section 7). Third, where the menu reinsurer's expected growth rate exceeds the rivals' by a margin bounded away from zero at every division of the market's surplus, its share of that surplus converges to one almost surely (Theorem 2); at the calibration the gap is positive at every share on a grid of ten-thousandths, and simulated paths from a five percent share end with a mean share of 0.992 after 300 periods. Fourth, if the rivals reprice to the loss experience of the business they write, the outcome turns on their loading: at a zero loading their price rises toward the high-risk type's cost, their growth rate falls below the risk-free rate, and on some paths the noise of a small book lifts the price above that type's willingness to pay, after which no reinsurance is written; at a small positive loading they survive on the high-risk business while the menu reinsurer's share still grows. A cycle extension adds a common loss factor and demand that grows at the risk-free rate. At a share of 0.25 of industry surplus, correlated losses narrow the menu reinsurer's advantage at the higher capitalizations computed and widen it where the industry

[1] Michael 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.

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

[3] Lawrence Blume and David Easley, "Evolution and Market Behavior," *Journal of Economic Theory* 58, no. 1 (1992): 9–40.

is capitalized at or below the calibration's level, because the common factor's variance scales with the square of a book's exposure, and the rivals' book carries the higher volatility. At that share and a capitalization at which the industry's capacity equals the premium of writing every submission at the population mean loss ratio, the Kelly criterion keeps the menu reinsurer's exposure at its full capacity until the common share of the loss-ratio variance reaches 0.897, a threshold found by bisection; on simulated paths with a common factor the exposure leaves full capacity on some paths.

In Rothschild and Stiglitz the separating contracts are a competitive equilibrium's response to private information, and the question is whether such an equilibrium exists; here the separating contract is a technology one firm adopts against others that have not, and the question is who survives. The answer runs through capital. A single-price reinsurer whose book deteriorates does not merely earn less; its capacity to write the next period's business shrinks relative to the firm that screened, so the composition of the market moves against it at a rate set by the growth gap.

The route is as follows. Section 2 places the paper between the competitive-screening and the market-selection literatures. Section 3 sets out the market: the cedants, the contracts, the placement rule, the reinsurers' objective, and the two regimes of the rivals' price. Section 4 establishes the cream-skimming contract (Theorem 1) and Section 5 the rivals' expected result and growth rate (Propositions 1 and 2). Section 6 states the selection theorem (Theorem 2), whose hypothesis is a uniform growth gap, and Section 7 measures that gap at the calibration by a quadrature and compares the quadrature with a Monte Carlo that draws every loss. Section 8 reports simulated paths of the menu reinsurer's share under the fixed pooled price and under experience-rated pricing at a zero and at a positive loading, and the configuration the market reaches when the rivals price at the high-risk type's cost. Section 9 extends the market to correlated losses and to demand that grows at the risk-free rate, where capacity binds and the Kelly criterion sizes the book. Sections 10 through 12 carry the market onto a continuum of types, onto two dimensions of private information, and onto heavy-tailed losses, and Section 13 concludes.

2 Related Literature

Two literatures answer halves of the question. Competitive screening under adverse selection characterizes which contracts survive profitable deviation when every firm may offer any contract. In Rothschild and Stiglitz[4] the results are the separating pair in which the high-risk type buys full coverage at its own fair price and the low-risk type buys partial coverage at its own, and the nonexistence of equilibrium when the low-risk share is large enough for a pooling contract to attract both types; Pauly treats the competitive market in which an insurer cannot observe the coverage a buyer holds with other insurers and so quotes a price per unit of coverage, a single price.[5] Wilson's anticipatory equilibrium[6] and the Wilson–Miyazaki–Spence (WMS) allocation, the separating menu best for the low-risk type among those on which firms break even overall, the low-risk type's contract subsidizing the high-risk type's,[7] restore existence by restricting the deviations firms

[4] Rothschild and Stiglitz, "Equilibrium in Competitive Insurance Markets."

[5] Mark V. Pauly, "Overinsurance and Public Provision of Insurance: The Roles of Moral Hazard and Adverse Selection," *Quarterly Journal of Economics* 88, no. 1 (1974): 44–62.

[6] Charles Wilson, "A Model of Insurance Markets with Incomplete Information," *Journal of Economic Theory* 16, no. 2 (1977): 167–207.

[7] Hajime Miyazaki, "The Rat Race and Internal Labor Markets," *Bell Journal of Economics* 8, no. 2 (1977): 394–418; Michael Spence, "Product Differentiation and Performance in Insurance Markets," *Journal of Public Economics* 10, no. 3 (1978): 427–447.

anticipate, and Mimra and Wambach give the WMS allocation a game-theoretic foundation under entry.[8] Inderst and Wambach show that capacity constraints restore existence by another route: a firm that cannot serve every cedant who would accept a pooling deviation cannot profit from offering it.[9] The dynamic treatments of competitive screening hold firm size fixed: Ania, Tröger, and Wambach track contract offers that firms imitate and experiment with,[10] and Lund and Nilssen follow competitive screening over two periods, in which a firm's second-period offers can draw the better risks, draw the worse, or pool them.[11] In each, a firm's share of the market is a choice variable or a parameter, not the compounded result of what it earned.

Market selection supplies the other half. Blume and Easley show that among agents whose wealth compounds, those with the higher expected logarithmic growth rate come to hold the market,[12] and in complete markets the selection favors the most accurate beliefs, a conclusion they show can fail when markets are incomplete;[13] Sandroni extends the result to agents who differ in the accuracy of their predictions,[14] and Evstigneev, Hens, and Schenk-Hoppé survey the evolutionary finance that grew from it.[15] The objective that the selected agent maximizes is Kelly's expected logarithm of the return,[16] and the strong law that carries the selection argument is the strong law of large numbers for martingales, as Hall and Heyde state it.[17] This paper joins the two literatures at the point where a reinsurer's capacity is a multiple of its surplus: a screening technology is then selected through the growth rate it earns.

Three further strands bear on the model. The principal–agent treatments of reinsurance under adverse selection characterize a monopolistic reinsurer's optimal contract against a cedant's risk measure, for two types[18] and for a continuum, where the optimum excludes the better risks by a threshold rather than distorting their coverage.[19] Section 8 reports the same exclusion for the menu reinsurer once its rivals have left the market, and Section 10 sets the continuum against a posted

[8] Wanda 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).

[9] Roman Inderst and Achim Wambach, "Competitive Insurance Markets under Adverse Selection and Capacity Constraints," *European Economic Review* 45, no. 10 (2001): 1981–1992.

[10] Ana B. Ania, Thomas Tröger, and Achim Wambach, "An Evolutionary Analysis of Insurance Markets with Adverse Selection," *Games and Economic Behavior* 40, no. 2 (2002): 153–184.

[11] Diderik Lund and Tore Nilssen, "Cream Skimming, Dregs Skimming, and Pooling: On the Dynamics of Competitive Screening," *Geneva Papers on Risk and Insurance Theory* 29, no. 1 (2004): 23–41.

[12] Blume and Easley, "Evolution and Market Behavior."

[13] Lawrence Blume and David Easley, "If You're So Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets," *Econometrica* 74, no. 4 (2006): 929–966.

[14] Alvaro Sandroni, "Do Markets Favor Agents Able to Make Accurate Predictions?" *Econometrica* 68, no. 6 (2000): 1303–1341.

[15] Igor V. Evstigneev, Thorsten Hens, and Klaus Reiner Schenk-Hoppé, "Evolutionary Finance," in *Handbook of Financial Markets: Dynamics and Evolution*, ed. Thorsten Hens and Klaus Reiner Schenk-Hoppé (Amsterdam: North-Holland, 2009), 507–566.

[16] Kelly, "A New Interpretation of Information Rate"; Edward O. Thorp, "The Kelly Criterion in Blackjack, Sports Betting, and the Stock Market," in *Handbook of Asset and Liability Management*, vol. 1, *Theory and Methodology*, ed. Stavros A. Zenios and William T. Ziemba (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).

[17] P. Hall and C. C. Heyde, *Martingale Limit Theory and Its Application* (New York: Academic Press, 1980), Theorem 2.18.

[18] Ka 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; Ka Chun Cheung, Sheung Chi Phillip Yam, Fei Lung Yuen, and Yiying Zhang, "Concave Distortion Risk Minimizing Reinsurance Design under Adverse Selection," *Insurance: Mathematics and Economics* 91 (2020): 155–165.

[19] Ka Chun Cheung, Sheung Chi Phillip Yam, Fei Lung Yuen, and Yiying Zhang, "Optimal Design of Reinsurance Contracts under Adverse Selection with a Continuum of Types" (arXiv preprint arXiv:2504.17468v1, 24 April 2025).

price, where the threshold cuts the type space from the other end. Kong, Layton, and Shepard show adverse-selection price competition concentrating an insurance market;[20] here concentration runs through capital rather than through price. The capacity-constraint theories of the insurance cycle, Gron's and Winter's,[21] tie the supply of coverage to surplus and read the hard market as the state in which that constraint binds; Section 9 adopts the constraint and reports where it binds in the model. The experience rating of the rivals' price in Section 8 is a credibility formula of the updating type, in the form of Jones and Gerber, with a fixed weight;[22] Bühlmann derives the credibility weight from the variance components of the collective,[23] which the rivals here do not estimate.

3 The Market

This section sets out the cedants and their contracts, the two reinsurers, the placement of submissions under capacity, the reinsurers' objective, and the rivals' pricing rule.

3.1 Cedants and Contracts

Two types of cedant, $\theta \in \{L, H\}$, with population share $\lambda \in (0,1)$ of type L. A submission of type θ has loss ratio $X_\theta = \min\{\max\{0, Y_\theta\}, \bar{X}\}$, where $Y_\theta \sim N(\mu_\theta, \sigma_\theta^2)$ with $\sigma_\theta > 0$ and $\bar{X} > 0$ caps the loss ratio; write $m_\theta = \mathbb{E}X_\theta$ and $v_\theta = \operatorname{Var} X_\theta$, with $m_L < m_H$ and $v_L < v_H$. A contract is a pair (α, p): the cedant cedes the share $\alpha \in [0,1]$ of its loss and pays the premium p per unit of subject premium. The cedant is mean–variance with risk aversion $\rho > 0$, read as the cost of the capital it holds against retained risk, and its gain from (α, p) over no reinsurance is

$$u_\theta(\alpha, p) = \alpha m_\theta + \tfrac{\rho}{2}\big[1 - (1-\alpha)^2\big] v_\theta - p, \tag{1}$$

the expected loss transferred plus the value of the variance removed, less the premium. Write $W_\theta(\alpha) = u_\theta(\alpha, 0)$ for the type's willingness to pay for coverage α. Since $\partial u_\theta / \partial \alpha = m_\theta + \rho(1-\alpha)v_\theta$ and both $m_H > m_L$ and $v_H > v_L$, the high-risk type's marginal willingness to pay exceeds the low-risk type's at every coverage: the single-crossing property. Define $\Delta(\alpha) = W_H(\alpha) - W_L(\alpha)$, with $\Delta(0) = 0$ and $\Delta'(\alpha) = (m_H - m_L) + \rho(1-\alpha)(v_H - v_L) > 0$, so that Δ is strictly increasing and positive on $(0,1]$.

A cedant buys the contract that maximizes its gain among those that leave it a non-negative gain, and buys nothing otherwise. A cedant indifferent between two contracts takes the one with the lower premium, and a cedant indifferent between a contract and no reinsurance buys the contract; these are the limits of undercutting by an amount that tends to zero, and they resolve the ties the constructions below create.

[20] Edward Kong, Timothy Layton, and Mark Shepard, "Adverse Selection and (un)Natural Monopoly in Insurance Markets" (NBER Working Paper No. 33187, 2024).

[21] Anne Gron, "Capacity Constraints and Cycles in Property-Casualty Insurance Markets," *RAND Journal of Economics* 25, no. 1 (1994): 110–127; Ralph A. Winter, "The Dynamics of Competitive Insurance Markets," *Journal of Financial Intermediation* 3, no. 4 (1994): 379–415.

[22] Donald A. Jones and Hans U. Gerber, "Credibility Formulas of the Updating Type," *Transactions of the Society of Actuaries* 27 (1975): 31–46.

[23] Hans Bühlmann, "Experience Rating and Credibility," *ASTIN Bulletin* 4, no. 3 (1967): 199–207.

3.2 Reinsurers and Contracts Offered

Two reinsurers write the market, indexed 0 (the menu reinsurer) and R (the rivals: one reinsurer, or several that offer the same contract and share its demand in proportion to their capacity). The rivals offer full coverage at a pooled price, $C_P = (1, p_P)$. The menu reinsurer offers one contract $C_L = (\alpha_L, p_L)$, which with C_P forms the menu from which a cedant chooses (Sections 10 and 11 give the menu reinsurer more contracts), from the family

$$C_L(\alpha) = \big(\alpha, p(\alpha)\big), \qquad p(\alpha) = W_L(\alpha) - \max\{u_L(C_P), 0\}, \qquad \alpha \in [0,1], \tag{2}$$

the price at which the low-risk type is indifferent between $C_L(\alpha)$ and its best alternative, C_P where that contract leaves it a non-negative gain and no reinsurance otherwise. The admissible coverages at the rivals' price p_P are

$$\mathcal{A}(p_P) = \big\{\alpha \in (0,1) : p(\alpha) > 0,\ u_H(C_P) \geq 0,\ u_H(C_P) > u_H(C_L(\alpha))\big\}, \tag{3}$$

those at which the contract is bought by the low-risk type and left by the high-risk type for C_P. The menu reinsurer offers no contract to the high-risk type: at the pooled price of a competitive market no such contract earns more than it costs, since C_P subsidizes that type. The restriction is kept under the experience rating of Section 8, where the rivals' price can exceed the high-risk type's expected loss.

3.3 Periods, Capacity, and Placement

Each period, a year, $N \geq 1$ submissions arrive, their types independent draws with $\mathbb{P}(\theta = L) = \lambda$, their loss ratios independent given the types; write n_L and $n_H = N - n_L$ for the type counts. Industry surplus is normalized to one at the start of the period; the menu reinsurer holds the share $s \in (0,1)$ of it and the rivals $1 - s$. A submission carries B units of subject premium, B in units of industry surplus, so that demand scales with industry surplus (Assumption 1; Section 9 relaxes it). A reinsurer may write reinsurance premium up to κ times its surplus: on the lines $f_i \in (0,1]$ it writes, $\sum_i f_i B p_i \leq \kappa S$. Submissions arrive in a uniformly random order, independent of the types and the losses, and go to the reinsurer offering the cedant's chosen contract; a reinsurer at capacity writes the marginal submission in part, the line that fits, and the remainder of that submission and every later one go to the cedant's next-best contract where that contract leaves the cedant a non-negative gain, and are otherwise unplaced. Where the menu reinsurer offers $C_L(\alpha)$ with $\alpha \in \mathcal{A}(p_P)$, the low-risk type's chosen contract is C_L and its next-best C_P where $u_L(C_P) \geq 0$; the high-risk type's chosen contract is C_P and its next-best C_L, which leaves it the gain $u_H(C_L(\alpha)) = \Delta(\alpha) + \max\{u_L(C_P), 0\} > 0$. Write w_{0L} and w_{0H} for the low-risk and high-risk lines the menu reinsurer writes and w_{RH}, w_{RL} for the rivals'; in lines, the menu reinsurer's capacity is $K_0 = \kappa s/(B p_L)$ and the rivals' $K_R = \kappa(1-s)/(B p_P)$.

Assumption 1 (scale-free market)**.** Demand is proportional to industry surplus: the subject premium per submission is B units of industry surplus in every period.

3.4 Returns and the Objective

The gross return of a reinsurer with surplus S is

$$\Gamma = 1 + r_f + \frac{U}{S}, \qquad U = \sum_{i \in \text{book}} f_i B\,(p_i - \alpha_i X_i), \tag{4}$$

the risk-free return on the whole surplus plus the underwriting result U on the lines it wrote, and its surplus after the period is $S \max\{\varepsilon, \Gamma\}$ for a ruin floor $\varepsilon \in (0,1)$. A reinsurer's objective over the period is its expected logarithmic growth, $g = \mathbb{E} \ln \max\{\varepsilon, \Gamma\}$, the Kelly criterion. Because the premium written is at most κS and losses are non-negative, $\Gamma \le 1 + r_f + \kappa$; because the loss ratio is at most $\bar{X}$ and the subject premium written is at most $\kappa S / \min_i p_i$, $\Gamma \ge \Gamma_{\min} = 1 + r_f - \kappa(\bar{X} - \min_i p_i)/\min_i p_i$, which can be negative. Hence $\ln \max\{\varepsilon, \Gamma\}$ lies in $[\ln \varepsilon, \ln(1 + r_f + \kappa)]$.

3.5 The Menu Reinsurer's Choice and the Rivals' Price

At the start of each period the menu reinsurer, given its share s and the rivals' price p_P, chooses α_L from a finite set $\mathcal{G} \cap \mathcal{A}(p_P)$ ($\mathcal{G}$ the grid of 99 coverages from 0.01 to 0.99, refined by thirty golden-section steps between the grid's neighbors of the best, a refined coverage kept where it separates the types and earns at least the grid's best; the strict inequality of (3) is tested with a tolerance of 10^{-12}) to maximize its expected logarithmic growth over the period under the placement rule above, the spillover of rationed high-risk submissions included; where $\mathcal{G} \cap \mathcal{A}(p_P)$ is empty it writes nothing and earns $g_0 = \ln(1 + r_f)$. The choice is myopic and is made against the posted contract C_P. Write $g_0(s, p_P)$ and $g_R(s, p_P)$ for the two expected growth rates at the contract so chosen. The rivals' price follows one of two regimes. Under the fixed pooled price (R1) it is $p_P = (1 + \delta)\bar{m}$ with $\bar{m} = \lambda m_L + (1 - \lambda) m_H$ the population mean loss ratio and $\delta \ge 0$ a loading. Under experience rating (R2) it is $p_{P,t+1} = (1 + \delta)\hat{m}_t$, where $\hat{m}_t$ is a credibility-weighted mean of the realized loss ratio of the business the rivals wrote, $\hat{m}_t = c\, x_{R,t} + (1 - c)\hat{m}_{t-1}$ with $\hat{m}_{-1} = \bar{m}$, $c \in (0, 1]$, and $x_{R,t}$ the premium-weighted mean loss ratio of the lines the rivals wrote in period t; in a period in which the rivals wrote nothing, $\hat{m}_t = \hat{m}_{t-1}$.

3.6 Calibration

Table 1 states every primitive, its source, and whether the runs of Sections 7 through 9 vary it. The loss-ratio means and volatilities, the population share, the risk-free rate, and the number of submissions per period are the paper's baseline; the cap $\bar{X}$ and the ruin floor ε are stated bounds. The risk aversion ρ is set so that the high-risk type's willingness to pay for full coverage equals a loading of 15 percent on its expected loss, $\frac{\rho}{2} v_H = 0.15\, m_H$. The book size B is set so that full coverage of every submission at 1.12 times the population mean loss ratio uses the industry's capacity exactly, $NB \cdot 1.12\,\bar{m} = \kappa$. The low-risk mean μ_L is set so that, at a zero loading, the low-risk type's gain from the pooled contract is one half of its full-coverage risk premium, $u_L(C_P) = \frac{1}{2} \cdot \frac{\rho}{2} v_L$; the pooled contract then has low-risk buyers, which a pooled market under voluntary purchase requires. The calibration sets $\delta = 0$, the zero-expected-profit pooled price of a competitive market; Section 8 varies it.

Table 1: Primitives of the market. Each parameter's value at the calibration, its source, and whether the runs of Sections 7 through 9 vary it. The capped means and variances are computed from the Normal parameters; the risk aversion, the book size, and the low-risk mean follow the stated rules. Source: author's calculations.

Parameter	Value	Source	Varied
μ_L, μ_H	0.7654, 0.80	μ_L by the half-risk-premium rule; μ_H baseline	no
σ_L, σ_H	0.12, 0.25	baseline	no
m_L, m_H	0.7654, 0.8000	capped means	no
v_L, v_H	0.0144, 0.0624	capped variances	no
$\bar{X}$	3	cap on the loss ratio	no
λ	0.60	baseline	no
N	50	submissions per period, baseline	no
r_f	0.045	3-month U.S. Treasury bill, baseline	no
κ	1	capacity multiple	no
ρ	3.8451	willingness-to-pay rule	no
B	0.022915	capacity rule	Section 9 (demand grows at r_f)
ε	10^{-10}	ruin floor	no
δ	0.00	zero-expected-profit pooled price	Section 8
c	0.3	credibility weight, R2	no
s_0	0.05	initial share of the menu reinsurer	no

At the calibration the population mean loss ratio is $\bar{m} = 0.7793$, the fixed pooled price $p_P = 0.7793$, the low-risk type's gain from the pooled contract $u_L(C_P) = 0.01384$ against a full-coverage risk premium of $\frac{\rho}{2} v_L = 0.02769$, and the high-risk type's gain $u_H(C_P) = 0.1408$.

4 Cream Skimming

This section establishes the contract the menu reinsurer offers beside the rivals' pooled contract and the margin it earns.

Theorem 1 (cream skimming)**.** *Let the rivals offer $C_P = (1, p_P)$ with $u_L(C_P) \geq 0$ and $p_P > m_L$, so that $p(\alpha) = W_L(\alpha) - u_L(C_P)$ in* (2)*. Then:*

(i) for every $\alpha \in [0,1]$ the low-risk type is indifferent between $C_L(\alpha)$ and C_P, and weakly prefers either to no reinsurance;

(ii) for every $\alpha \in [0,1)$ the high-risk type strictly prefers C_P to $C_L(\alpha)$, and prefers C_P to no reinsurance;

(iii) the menu reinsurer's expected underwriting margin on $C_L(\alpha)$ per unit of subject premium written on a low-risk line,

$$\pi(\alpha) = p(\alpha) - \alpha m_L = \tfrac{\rho}{2}\big[1 - (1-\alpha)^2\big] v_L - u_L(C_P), \tag{5}$$

is continuous and strictly increasing on $[0,1]$, with $\pi(0) = -u_L(C_P) \leq 0$ and $\pi(1) = p_P - m_L > 0$; hence it has exactly one zero $\alpha^* \in [0,1)$, at which the low-risk type's value of the variance removed equals its gain from C_P, and $\pi(\alpha) > 0$ for every $\alpha \in (\alpha^*, 1]$. Every $\alpha \in (\alpha^*, 1)$ lies in $\mathcal{A}(p_P)$.

Proof. (i) By construction $u_L(\alpha, p(\alpha)) = W_L(\alpha) - p(\alpha) = u_L(C_P) \geq 0$.

(ii) $u_H(\alpha, p(\alpha)) = W_H(\alpha) - W_L(\alpha) + u_L(C_P) = \Delta(\alpha) + u_L(C_P)$, while $u_H(C_P) = W_H(1) - p_P = \Delta(1) + W_L(1) - p_P = \Delta(1) + u_L(C_P)$. Hence $u_H(C_P) - u_H(\alpha, p(\alpha)) = \Delta(1) - \Delta(\alpha) = \int_\alpha^1 \Delta'(a)\, da > 0$ for $\alpha < 1$, by single crossing. Also $u_H(C_P) = \Delta(1) + u_L(C_P) \geq \Delta(1) > 0$, since $\Delta(0) = 0$ and Δ is strictly increasing.

(iii) From (1) and (2), $p(\alpha) - \alpha m_L = \frac{\rho}{2}[1 - (1-\alpha)^2] v_L - u_L(C_P)$, a polynomial in α whose derivative $\rho(1-\alpha)v_L$ is positive on $[0, 1)$. At $\alpha = 1$, $\pi(1) = \frac{\rho}{2} v_L - u_L(C_P) = \frac{\rho}{2} v_L - (m_L + \frac{\rho}{2} v_L - p_P) = p_P - m_L > 0$. A continuous, strictly increasing function with $\pi(0) \leq 0 < \pi(1)$ has exactly one zero in $[0, 1)$ and is positive beyond it. For $\alpha \in (\alpha^*, 1)$, $p(\alpha) = \pi(\alpha) + \alpha m_L > 0$, and (i) and (ii) give the two remaining conditions of (3). □

The contract of Theorem 1 is a cream-skimming contract against a posted price: the low-risk type's participation constraint relative to C_P binds and the high-risk type's incentive constraint is slack. It differs from the separating contract of Rothschild and Stiglitz, in which the high-risk type's incentive constraint binds and the low-risk type pays an actuarially fair price; here the menu reinsurer keeps the low-risk type's surplus at $u_L(C_P)$. At the calibration $\alpha^* = 0.2929$: every coverage above it earns a positive margin, and the menu reinsurer's choice among them is the growth-optimal one of Section 7. Figure 1 draws the family's margin on a low-risk line and on a high-risk line, and the high-risk type's preference for the pooled contract, over the coverage.

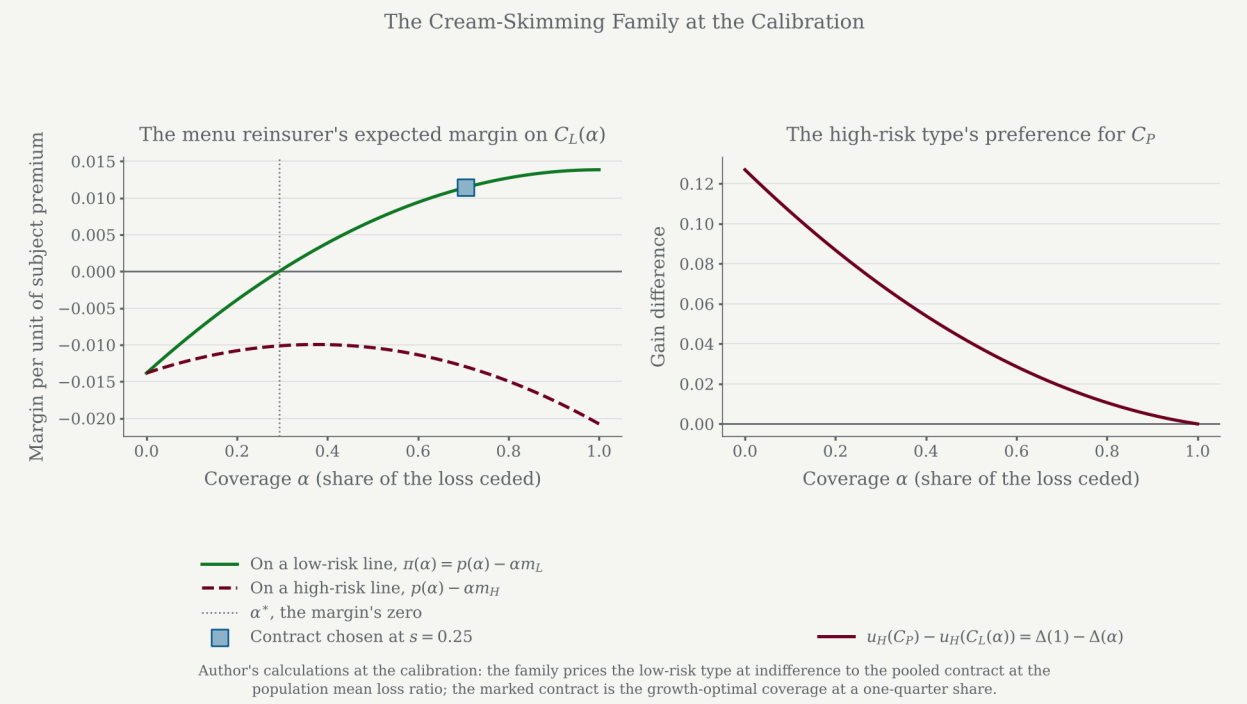


Figure 1: The cream-skimming family at the calibration. Left panel: the menu reinsurer's expected margin per unit of subject premium on $C_L(\alpha)$ when a low-risk cedant takes it, $\pi(\alpha)$ of (5) (green), and when a high-risk cedant would, $p(\alpha) - \alpha m_H$ (maroon, dashed); the dotted vertical marks $\alpha^* = 0.2929$, the margin's zero, and the square the contract the menu reinsurer chooses at a share of 0.25, $C_L(0.7045)$ at a premium of 0.5507. Right panel: the high-risk type's gain from the pooled contract less its gain from $C_L(\alpha)$, $\Delta(1) - \Delta(\alpha)$, positive at every coverage below one. Horizontal axes: coverage α. Source: author's calculations.

Where $u_L(C_P) < 0$, the low-risk type's alternative is no reinsurance, the price rule (2) gives $p(\alpha) = W_L(\alpha)$, the margin on a low-risk line is $\frac{\rho}{2}[1-(1-\alpha)^2]v_L > 0$ for every $\alpha > 0$, and the high-risk type prefers C_P to $C_L(\alpha)$ if and only if $\Delta(1) - \Delta(\alpha) > -u_L(C_P)$, that is, for α below the coverage $\hat{\alpha} \in (0,1)$ at which $\Delta(1) - \Delta(\hat{\alpha}) = -u_L(C_P)$, which exists if and only if $u_H(C_P) = \Delta(1) + u_L(C_P) > 0$. Where $u_H(C_P) < 0$ no cedant buys C_P, $\mathcal{A}(p_P)$ is empty, and the menu reinsurer writes nothing: a contract the low-risk type accepts is accepted by the high-risk type too, since $u_H(\alpha, p) - u_L(\alpha, p) = \Delta(\alpha) > 0$, so no contract of the family separates without C_P to draw the high-risk type away. Section 8 reaches both configurations.

5 The Rivals' Book

This section states what the cream-skimming contract leaves the rivals: their expected underwriting result on the book they write, and the bound it puts on their growth rate. In the vocabulary of underwriting, the rivals' losses are errors of commission—business written below its cost—and the menu reinsurer's retention of the low-risk type is the avoidance of errors of omission, profitable business declined by a price that does not distinguish it.

Corollary 1 (the rivals' margin at a zero loading)**.** *Under R1 with $\delta = 0$, $p_P = \bar{m}$, so the rivals' expected margin is $\bar{m} - m_H = -\lambda(m_H - m_L) < 0$ on a high-risk line and $\bar{m} - m_L = (1-\lambda)(m_H - m_L) > 0$ on a low-risk one. Their expected underwriting result on a book of w_{RH} high-risk and w_{RL} low-risk lines is $B(m_H - m_L)\big[(1-\lambda)w_{RL} - \lambda w_{RH}\big]$, non-positive if and only if the book's low-risk share $w_{RL}/(w_{RL} + w_{RH})$ is at most λ.*

Proof. Immediate from $p_P = \bar{m} = \lambda m_L + (1-\lambda)m_H$ and (4); the equivalence is $(1-\lambda)w_{RL} \le \lambda w_{RH} \iff w_{RL} \le \lambda(w_{RL} + w_{RH})$. □

Proposition 1 (the rivals' expected underwriting result)**.** *Under R1 with $\delta = 0$, let the menu reinsurer offer $C_L(\alpha)$ with $\alpha \in \mathcal{A}(p_P)$, so that every low-risk submission chooses C_L and every high-risk submission C_P. Given the type counts, the rivals' demand is n_H high-risk lines and the overflow $o_L = (n_L - K_0)^+$ of low-risk lines, and they write the fraction $q = \min\{1, K_R/(n_H + o_L)\}$ of it. Where at most one of the two reinsurers rations, $o_L = 0$ or $q = 1$, they accept each type in the proportion q, and*

$$\mathbb{E}[U_R \mid n_L] = B(m_H - m_L)\, q\, \big[(1-\lambda)o_L - \lambda n_H\big]. \tag{6}$$

Consequently:

- *(i) if the rivals' capacity covers every realization of their demand, $\kappa(1-s) \ge NBp_P$, then $q = 1$ and $\mathbb{E}U_R = -B(m_H - m_L)(1-\lambda)\,\mathbb{E}\min\{n_L, K_0\} < 0$;*
- *(ii) if the menu reinsurer's capacity covers every realization of the low-risk demand, $\kappa s \ge NBp_L$, then $o_L = 0$ and $\mathbb{E}U_R = -\lambda B(m_H - m_L)\,\mathbb{E}\min\{n_H, K_R\} < 0$.*

Between the two the sign of $\mathbb{E}U_R$ is not settled. Where the two reinsurers never ration together, (6) holds at every count, but the bracket is increasing and q non-decreasing in n_L, so the rationing weights the counts at which the rivals' demand is least high-risk. At a count at which both ration, the low-risk overflow reaches the rivals only from the arrival that fills the menu reinsurer's capacity, the lines the rivals accept are not an exchangeable sample of their demand, and (6) need not hold; the quadrature of Section 7 computes the expected counts exactly at such a count.

Proof. Where $q = 1$ the rivals write their whole demand. Where $o_L = 0$ their demand is n_H high-risk submissions, and the lines they accept, the first arrivals among them up to the premium $\kappa(1-s)$, the partial line included, number qn_H, all high-risk. In either case each type is accepted in the proportion q, and the expected underwriting result is $B[qn_H(p_P - m_H) + qo_L(p_P - m_L)]$, which Corollary 1 turns into (6). For (i), $q = 1$ for every realization, and $\mathbb{E}[(1-\lambda)o_L - \lambda n_H] = (1-\lambda)(\lambda N - \mathbb{E}\min\{n_L, K_0\}) - \lambda(1-\lambda)N = -(1-\lambda)\,\mathbb{E}\min\{n_L, K_0\}$, negative because $K_0 > 0$ and $n_L \geq 1$ with positive probability. For (ii), $o_L = 0$, $qn_H = \min\{n_H, K_R\}$, and the bracket is $-\lambda n_H$, negative with positive probability. For the monotonicity: o_L is non-decreasing in n_L and $n_H = N - n_L$ decreasing, so the bracket increases in n_L; the rivals' demand $n_H + o_L = N - \min\{n_L, K_0\}$ is non-increasing in n_L, so q is non-decreasing. That (6) need not hold where both ration: with $N = 3$, $n_L = 2$, and capacities of one line each, the rivals' demand is the high-risk submission and the second low-risk arrival, $q = 1/2$, and they accept the high-risk line where it arrives first or second among the three arrivals and not where it arrives third, in two of the three equally likely positions. □

At the calibration the hypothesis of (i) holds for every share $s \leq 0.107$, where the rivals can write all N submissions at the pooled price. At every share the two reinsurers' capacities in lines sum to at least $\kappa/(Bp_P) = 1.12\,N$, since $p_L < p_P$, so they never ration together and (6) holds at every count.

Lemma 1 (a tangent bound under the ruin floor)**.** *Let Γ be a random variable with $\Gamma \geq \Gamma_{\min}$ almost surely, and let $\varepsilon \in (0,1)$ satisfy*

$$\varepsilon < (1+r_f)\exp\bigl(-1 + \Gamma_{\min}/(1+r_f)\bigr). \tag{7}$$

Then $\ln\max\{\varepsilon, x\} \leq \ln(1+r_f) + \bigl(x - (1+r_f)\bigr)/(1+r_f)$ for every $x \geq \Gamma_{\min}$, with equality only at $x = 1 + r_f$, and therefore

$$\mathbb{E}\ln\max\{\varepsilon, \Gamma\} \leq \ln(1+r_f) + \frac{\mathbb{E}\Gamma - (1+r_f)}{1+r_f}, \tag{8}$$

with strict inequality unless $\Gamma = 1 + r_f$ almost surely.

Proof. For $x \geq \varepsilon$, $\ln\max\{\varepsilon, x\} = \ln x$, and the concavity of ln puts its graph below its tangent at $1 + r_f$, with equality only at the point of tangency. For $x \in [\Gamma_{\min}, \varepsilon)$, $\ln\max\{\varepsilon, x\} = \ln\varepsilon$, and (7) gives $\ln\varepsilon < \ln(1+r_f) - 1 + \Gamma_{\min}/(1+r_f) \leq \ln(1+r_f) + (x - (1+r_f))/(1+r_f)$. Taking expectations gives (8); the inequality is strict as soon as $\Gamma \neq 1 + r_f$ with positive probability. □

Proposition 2 (the rivals grow below the risk-free rate)**.** *Under R1 with $\delta = 0$, let the menu reinsurer offer $C_L(\alpha)$ with $\alpha \in \mathcal{A}(p_P)$, let s satisfy the hypothesis of (i) or of (ii) of Proposition 1, and let the ruin floor satisfy (7) with $\Gamma_{\min} = 1 + r_f - \kappa(\bar{X} - p_P)/p_P$. Then $g_R(s, p_P) < \ln(1+r_f)$.*

Proof. The rivals' book carries the single price p_P, so $U_R \geq -(\bar{X} - p_P)\sum_i f_i B \geq -\kappa(1-s)(\bar{X} - p_P)/p_P$ and $\Gamma_R \geq \Gamma_{\min}$. By Lemma 1, $g_R \leq \ln(1+r_f) + \mathbb{E}U_R/\bigl((1-s)(1+r_f)\bigr)$, strictly because Γ_R is non-degenerate (the rivals write at least one high-risk line with positive probability, and $v_H > 0$). Proposition 1 gives $\mathbb{E}U_R < 0$. □

At the calibration (7) reads $\varepsilon < 0.0684$ against the model's $\varepsilon = 10^{-10}$. Section 7 computes the rivals' growth rate, between the two hypotheses of Proposition 1 as under them, and finds it below $\ln(1+r_f) = 0.04402$ at every share of both its grids.

6 Selection

The argument is that of Blume and Easley for the selection of traders by the growth rates of their wealth, applied to the two reinsurers' surpluses. Let s_t be the menu reinsurer's share at the start of period t and $\ell_{0,t} = \ln\max\{\varepsilon, \Gamma_{0,t}\}$, $\ell_{R,t} = \ln\max\{\varepsilon, \Gamma_{R,t}\}$ the two logarithmic gross returns of period t. Since $s_{t+1} = s_t e^{\ell_{0,t}} / \big(s_t e^{\ell_{0,t}} + (1 - s_t) e^{\ell_{R,t}}\big)$, the log-odds $L_t = \ln\big(s_t/(1 - s_t)\big)$ obey

$$L_{t+1} = L_t + D_t, \qquad D_t = \ell_{0,t} - \ell_{R,t}, \tag{9}$$

exactly. Let $\mathcal{F}_t$ be the history up to the start of period t (the states and the draws of the periods before it), so that s_t and $p_{P,t}$ are $\mathcal{F}_t$-measurable and D_t is $\mathcal{F}_{t+1}$-measurable. Given $\mathcal{F}_t$, period t's draws are independent of the past and the menu reinsurer's contract is a function of the state $(s_t, p_{P,t})$, so the conditional mean of the increment is a function of the state alone,

$$d(s_t, p_{P,t}) = \mathbb{E}[D_t \mid \mathcal{F}_t] = g_0(s_t, p_{P,t}) - g_R(s_t, p_{P,t}), \tag{10}$$

the growth gap at the contract the menu reinsurer chooses at that state, where g_0 and g_R are the exact expected growth rates of the full placement rule. The increments are bounded: $|D_t| \le 2\bar{c}$ with $\bar{c} = \max\{|\ln\varepsilon|, \ln(1 + r_f + \kappa)\}$.

Theorem 2 (selection under the fixed pooled price)**.** *Under R1, suppose the growth gap is bounded away from zero on the state space: $d(s, p_P) \ge \underline{d} > 0$ for every $s \in (0, 1)$. Then, from every initial share $s_0 \in (0, 1)$, $\liminf_{t\to\infty} L_t/t \ge \underline{d}$ almost surely; in particular $L_t \to \infty$ and $s_t \to 1$ almost surely.*

Proof. Write $D_t = d(s_t, p_P) + \xi_t$ with $\xi_t = D_t - \mathbb{E}[D_t \mid \mathcal{F}_t]$. The partial sums $M_t = \sum_{\tau<t} \xi_\tau$ form a martingale with respect to $(\mathcal{F}_t)_{t\ge0}$ whose differences are bounded by $4\bar{c}$, so $\sum_{t\ge1} \mathbb{E}[\xi_t^2 \mid \mathcal{F}_t]/t^2 \le 16\bar{c}^2 \sum_{t\ge1} t^{-2} < \infty$, and the strong law of large numbers for martingales (Theorem 2.18 as Hall and Heyde state it, with the normalizing sequence t and $p = 2$) gives $M_t/t \to 0$ almost surely. By (9), $L_t/t = L_0/t + t^{-1}\sum_{\tau<t} d(s_\tau, p_P) + M_t/t \ge L_0/t + \underline{d} + M_t/t$, whose limit inferior is at least $\underline{d} > 0$ almost surely; the hypothesis applies along every path because the floor keeps every s_τ in $(0, 1)$. Hence $L_t \to \infty$ and $s_t = 1/(1 + e^{-L_t}) \to 1$ almost surely. □

The theorem's conclusion is about the limit. At any finite horizon the share's dispersion across paths is that of a random walk whose drift and conditional variance depend on the state; at small shares the menu reinsurer's underwriting result is magnified by B/s, up to κ/p_L at capacity, and the increment's standard deviation is a multiple of its drift, so that a path that falls to a small share recovers slowly against the noise. Section 8 shows that dispersion.

7 The Growth Gap at the Calibration

This section measures the hypothesis of Theorem 2 at the calibration: the two growth rates at the menu reinsurer's best contract, over the share, by a quadrature, and the quadrature against a Monte Carlo that draws every loss.

7.1 The Quadrature

Given the state and the contracts, the period's type counts are Binomial, and given the counts each reinsurer's book is determined up to the arrival order. The expected growth rate is computed as a finite sum over the counts under one approximation: each book's loss sum is replaced by the Normal with the same mean and variance, the variance of a count w of lines written in arrival order carrying the partial line as $\lfloor w \rfloor + (w - \lfloor w \rfloor)^2$. For a gross return Γ so replaced, $\mathbb{E} \ln \max\{\varepsilon, \Gamma\}$ is the probability that Γ lies below the floor times $\ln \varepsilon$ plus the integral of $\ln \Gamma$ above the floor; the integral is computed by a Gauss–Hermite rule with 96 nodes where the floor lies at least nine standard deviations below the mean of Γ, and otherwise by a tanh–sinh rule from the floor to twelve standard deviations above the mean. Where the rivals ration, their accepted counts enter at the fraction q of Proposition 1; under R1 at the calibration the two reinsurers never ration together (Section 5), and the fraction is then exact. Where the two capacities together fall short of the period's submissions, the expected counts are computed exactly over the arrival orders instead, and each book's Normal is taken at those expected counts, a second approximation, which leaves out the dispersion of the book's composition across arrival orders (Section 9 states the computation). The menu reinsurer's coverage is chosen on the grid of coverages refined by a golden section that Section 3 states.

7.2 The Gap over the Share

Figure 2 draws the two growth rates and their gap at the menu reinsurer's best contract over the share, and Table 2 reports them at selected shares. Under R1 at the calibration ($p_P = 0.7793$) the gap is positive at every share of the grid $s = 0.01, 0.02, \ldots, 0.99$: its smallest value is +0.00959 at $s = 0.01$ and its largest +0.10427 at $s = 0.99$. On the grid of ten-thousandths $s = 0.0001, 2 \cdot 0.0001, \ldots, 1 - 0.0001$ (9,999 points) the smallest gap is +0.00935 at $s = 0.0001$ and the largest +0.10442 at $s = 0.9822$, with 0 points at a non-positive gap; the best coverage lies between 0.5210 and 0.9648. The rivals' expected growth rate is below $\ln(1 + r_f) = 0.04402$ at every share of both grids, its largest value 0.04361 at $s = 0.0001$.

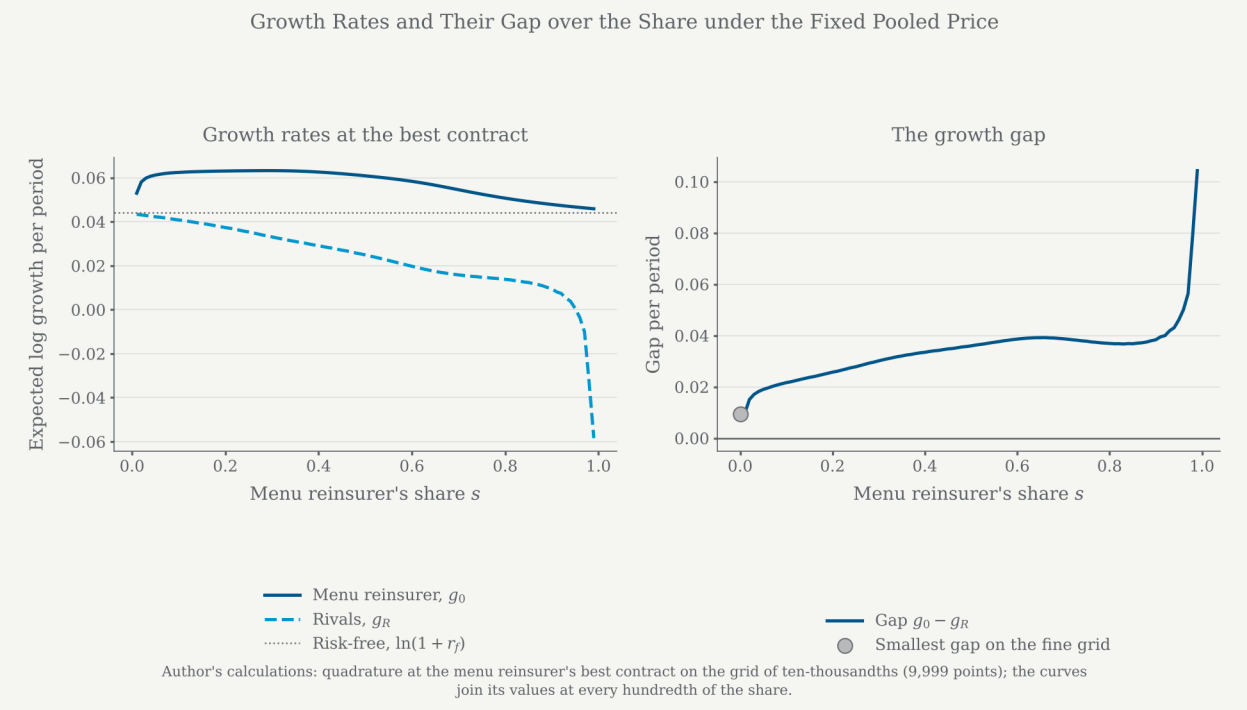


Figure 2: The two growth rates and their gap over the menu reinsurer's share under the fixed pooled price. Left panel: the expected logarithmic growth rate per period of the menu reinsurer (solid) and of the rivals (dashed) at the menu reinsurer's best contract, with the risk-free rate $\ln(1 + r_f)$ as the dotted reference; right panel: the gap $g_0 - g_R$. Horizontal axes: the menu reinsurer's share s of industry surplus; the curves join the values of the grid of ten-thousandths at every hundredth of the share. Source: author's calculations (quadrature).

Table 2: The quadrature at selected shares under the fixed pooled price $p_P = 0.7793$: the menu reinsurer's best coverage and price, the two expected growth rates per period, their gap, and the expected number of lines each reinsurer writes out of $N = 50$ submissions. Source: author's calculations.

s	α_L	p_L	g_0	g_R	$g_0 - g_R$	**Lines, menu**	**Lines, rivals**
0.01	0.7071	0.5527	0.05295	0.04336	+0.00959	0.79	49.21
0.05	0.6562	0.5128	0.06120	0.04222	+0.01898	4.25	45.75
0.10	0.6785	0.5303	0.06240	0.04077	+0.02163	8.23	41.77
0.25	0.7045	0.5507	0.06315	0.03535	+0.02780	19.81	30.19
0.50	0.8936	0.6975	0.06085	0.02490	+0.03596	29.18	20.82
0.75	0.8845	0.6905	0.05249	0.01472	+0.03778	36.05	13.95
0.90	0.7869	0.6149	0.04783	0.00949	+0.03834	44.40	5.60
0.95	0.7621	0.5956	0.04668	0.00055	+0.04613	47.20	2.80
0.99	0.7443	0.5818	0.04588	−0.05839	+0.10427	49.44	0.56

What the grids establish is the sign of the quadrature's gap on their points; the hypothesis of Theorem 2, a uniform bound over the continuum of shares, is not established by a grid, and between grid points the gap can move where the menu reinsurer's best contract jumps between two maximizers. As $s \to 0$ the menu reinsurer's book is a fraction of one low-risk submission and its growth is the most variable; the coarse grid's smallest gap stands there. As $s \to 1$ the rivals' capacity $\kappa(1 - s)$ vanishes, the high-risk submissions they turn away spill to $C_L(\alpha)$ at the low-risk

price, and the menu reinsurer's book becomes a pooled one: its margin per unit on a spilled line is $p(\alpha) - \alpha m_H = \pi(\alpha) - \alpha(m_H - m_L) \le (1 - \lambda - \alpha)(m_H - m_L)$, since $\pi \le \pi(1) = (1 - \lambda)(m_H - m_L)$, negative at every coverage above $1 - \lambda$, where the grids' best coverages lie. The gap is nonetheless largest near that end, at $s = 0.99$ on the coarse grid and 0.9822 on the grid of ten-thousandths, because the rivals' book is a fraction of one high-risk line priced at $\bar{m}$ on a surplus that vanishes with it: at $s = 0.99$ their expected gross return is 1.0184, whose logarithm lies 0.02583 below $\ln(1 + r_f)$, and the variance of that fraction of a line on their surplus, with the ruin floor, which values the return's lower tail at $\ln \varepsilon$, places their growth rate a further 0.07658 below the logarithm of the mean.

7.3 The Quadrature against a Monte Carlo

The quadrature was measured at four shares against a Monte Carlo of 200,000 periods each, in which every loss is drawn and every submission placed one by one in its arrival order. Table 3 reports the comparison: the largest deviation in absolute value is 1.5 standard errors, the two ends included, where one reinsurer's book is a fraction of one submission. At $s = 0.99$ the rivals are rationed to the same fraction of the first high-risk arrival in every period, so their book is a fixed fraction of one line, and the quadrature's approximation for that book is the replacement of the line's loss ratio, a Normal floored at zero and capped at $\bar{X}$, by the Normal with the same mean and variance; the high-risk Normal places 0.07 percent of its mass below zero (Section 12). The Monte Carlo gap at that share is +0.10254 against the quadrature's +0.10427.

Table 3: The quadrature against the Monte Carlo at four shares under the fixed pooled price: 200,000 periods at each share at the menu reinsurer's best contract; Monte Carlo standard errors in parentheses; the deviation, Monte Carlo less quadrature, in standard errors in brackets. Source: author's calculations.

s	α_L	p_L	g_0 quadrature	g_0 Monte Carlo	g_R quadrature	g_R Monte Carlo
0.01	0.7071	0.5527	0.05295	0.05293 (0.00033) [−0.1]	0.04336	0.04342 (0.00006) [+0.9]
0.25	0.7045	0.5507	0.06315	0.06307 (0.00007) [−1.1]	0.03535	0.03547 (0.00008) [+1.5]
0.90	0.7869	0.6149	0.04783	0.04787 (0.00005) [+0.8]	0.00949	0.00957 (0.00030) [+0.3]
0.99	0.7443	0.5818	0.04588	0.04592 (0.00005) [+0.8]	−0.05839	−0.05662 (0.00161) [+1.1]

8 Share Paths under Fixed and Experience-Rated Pricing

This section reports simulated paths of the market: the menu reinsurer's share of industry surplus under the fixed pooled price, and under experience-rated pricing at a zero and at a positive loading, with the configuration the market reaches when the rivals price at the high-risk type's cost. Each simulation draws every loss and places every submission in its arrival order under the rule of Section 3, on 1,000 paths of 300 periods from an initial share $s_0 = 0.05$; beside each, the expected-growth path replaces each period's realized growth rates by the quadrature's expectations and follows the same recursion without draws.

8.1 The Fixed Pooled Price

Figure 3 draws the share's median, its tenth-to-ninetieth percentile band, and its lowest value across paths, and Table 4 reports the cross-path means. Under R1 the cross-path mean of the share is 0.167 at period 59, 0.381 at period 100, 0.911 at period 200, and 0.992 at period 300; the median at the final period is 1.000 and 998 of the 1,000 paths stand above one half (Figure 6 draws the final shares). The cross-path mean of the log-odds increment is positive at every period, its smallest value +0.0142 at period 18. On the periods Figure 3 samples, one in every five, the expected-growth path lies below the cross-path median at all but one of the sampled periods before period 125 and above it at every sampled period from then on. The 2 paths at or below one half at the final period fell early to a small share, where the menu reinsurer's book is a fraction of one line and the increment's standard deviation is a multiple of its drift, and had not recovered within the horizon; their lowest share at the final period is 0.0028. At the calibration the cross-path mean share stands below one half at period 100 and above it at period 200.

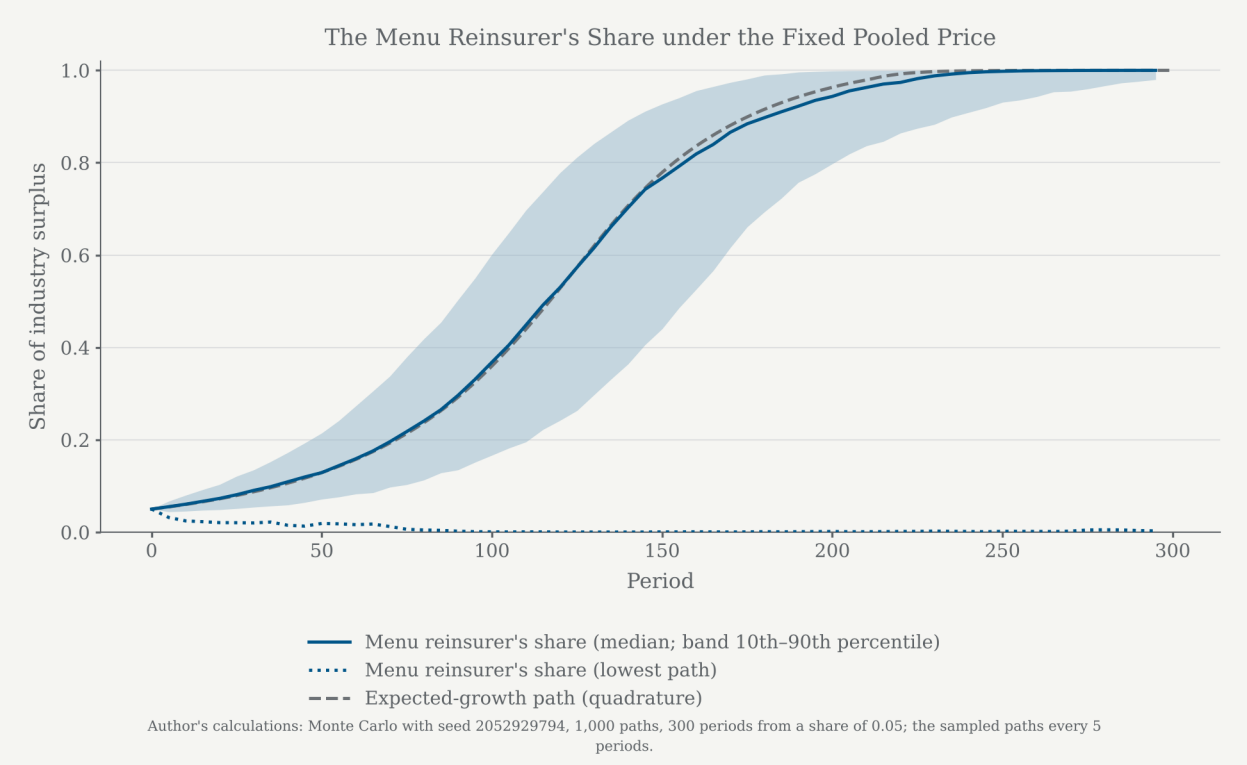


Figure 3: The menu reinsurer's share of industry surplus under the fixed pooled price. Horizontal axis: period; vertical axis: share. The solid line is the cross-path median and the shaded band spans the tenth to the ninetieth percentile; the dotted line is the lowest share across paths at each sampled period; the dashed line is the expected-growth path, in which each period's growth rates are the quadrature's expectations. Source: author's calculations (Monte Carlo, 1,000 paths, 300 periods, from a share of 0.05).

Table 4: The cross-path mean of the menu reinsurer's share at four periods under the three pricing regimes (period 59 is the period at which the text below compares experience rating with the fixed price), and at the final period the median, the number of paths at or below one half, and the number of paths on which the rivals' price stands above the high-risk type's willingness to pay for full coverage, $W_H(1)$ (none under the fixed price). Monte Carlo with 1,000 paths of 300 periods from a share of 0.05; the expected-growth path gives 0.154, 0.360, and 1.000 at periods 59, 100, and 300 under the fixed price, 0.180, 0.547, and 1.000 at $\delta = 0.00$, and 0.401, 0.652, and 0.998 at $\delta = 0.01$. Source: author's calculations.

Regime	$t = 59$	$t = 100$	$t = 200$	Final mean	Final median	$\leq \frac{1}{2}$	$p_P > W_H(1)$
Fixed price	0.167	0.381	0.911	0.992	1.000	2	0
Experience, $\delta = 0.00$	0.257	0.507	0.859	0.952	0.959	0	689
Experience, $\delta = 0.01$	0.310	0.540	0.804	0.899	0.903	0	352

8.2 Experience Rating at a Zero Loading: The Rivals' Exit

Under R2 with $\delta = 0.00$ the rivals' price is their experienced mean. As the menu reinsurer takes the low-risk business, the rivals' book tends to high-risk business alone, their experienced mean to m_H, and their price with it; on the paths that close (below) the noise of a small book lifts the price above the high-risk type's willingness to pay, and the cross-path mean of the price, which includes those paths, rises from 0.7793 to 0.8957 over the horizon, above $m_H = 0.8000$ (Figure 5, middle left). At the price m_H the rivals' expected margin on an unrationed high-risk book is zero, so by Lemma 1 their expected growth rate lies below the risk-free rate, and the low-risk type's gain from the pooled contract is $u_L(1, m_H) = m_L - m_H + \frac{\rho}{2} v_L = -0.00692 < 0$: its alternative is no reinsurance, the price rule gives $p(\alpha) = W_L(\alpha)$, and the admissible coverages are those below $\hat{\alpha} = 0.8556$. The share runs ahead of the fixed-price case early—0.257 against 0.167 at period 59—because the rising price removes the low-risk type's alternative and raises the menu reinsurer's margin.

Once the rivals' book is a fraction of one line, one year's realized loss ratio on that line moves their estimate by its credibility weight, and a realized loss ratio above 1.20 on that line, when the estimate stands at m_H, lifts the price above the high-risk type's willingness to pay for full coverage, $W_H(1) = 0.9201$. At that price no cedant buys the pooled contract; without the pooled contract the menu reinsurer cannot separate (Section 4) and writes nothing; the rivals write nothing; and an estimate that moves only on business written does not move again. No reinsurance is written on that path from that period on. In 38,191 path-periods no separating contract existed and the menu reinsurer wrote nothing; in 38,120 of them the rivals wrote nothing as well, and in the other 71 the rivals' price stood less than $\Delta(0.01) = 0.0022$ below $W_H(1)$, where every admissible coverage lies below 0.01, the grid's smallest. At the final period the rivals' price stood above $W_H(1)$ on 689 of the 1,000 paths (Figure 5, lower left); the cross-path mean share at the final period, 0.952, is the mean over paths that stopped where they stood when their market closed and paths that did not. The expected-growth path, which carries no draws, keeps the price at or below $W_H(1)$ and reaches a share of 1.000 at the final period (Table 4).

At the rivals' limiting price their expected margin is zero, and their growth rate is below the risk-free rate; the pricing rule, through the noise of a small book and a fixed credibility weight, decides which paths reach the closed-market state and when. In Akerlof's market for lemons[24] the better risks withdraw until no price clears the market; here the price that clears the pooled market is the high-risk type's cost, and at that price writing is worse than not writing.

[24] George A. Akerlof, "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics* 84, no. 3 (1970): 488–500.

Experience rating is the only learning in the model, and it is the rivals'. The menu reinsurer has nothing to learn from a cedant's losses: in a separating menu the cedant's choice reveals its type at once, and a menu that re-priced a cedant on its loss history would screen again what the choice has already screened. The rivals, who quote one price and observe the loss experience of the business they write, learn about their book and not about any cedant, and what they learn, the paths above show, is the composition the menu reinsurer leaves them. A market in which cedants persisted from period to period and a reinsurer observed a cedant's private loss history would add a second channel of information, held by whoever wrote the cedant last; the model's cedants are drawn afresh each period, and that channel is outside it.

8.3 Experience Rating at a Positive Loading

At a loading $\delta > 0$ on the experienced mean the rivals' limiting price is $(1+\delta)m_H$ and their margin on the high-risk book is δm_H. Two loadings bound the pooled market at the calibration. The pooled contract has low-risk buyers at the opening, where it is priced at $(1+\delta)\bar{m}$, if and only if $\delta \le \delta_L = W_L(1)/\bar{m} - 1 = 0.0178$; above that loading the market opens in the separated configuration, the rivals on the high-risk type alone and the menu reinsurer on the low-risk type, and the types are sorted between the two from the opening. The high-risk type buys the pooled contract priced on a high-risk book if and only if $\delta \le \delta_H = W_H(1)/m_H - 1 = 0.1500$ (at the opening, on the population mean, $\delta \le 0.1806$); δ_H is the loading of 15 percent on the high-risk type's expected loss by which the calibration sets ρ. Figure 4 draws the two reinsurers' growth rates over the loading in the separated configuration at three shares and at the opening. At the opening the rivals' growth rate falls where the loading crosses δ_L and the pooled contract loses its low-risk buyers, and rises with the loading above it, while the menu reinsurer's is constant on the grid above δ_L. In the separated configuration the menu reinsurer keeps the restriction of Section 3 and offers the high-risk type no contract, although the rivals' price $(1+\delta)m_H$ exceeds that type's expected loss there; the growth rates below are computed under that restriction. In the separated configuration at $s = 0.50$ the rivals' growth rate rises with the loading, from 0.04281 at $\delta = 0$ to 0.13871 at $\delta = 0.145$, the largest loading on the grid of 31 at which a separating contract exists (at δ_H the high-risk type is indifferent to the pooled contract and takes any contract of the family instead), while the menu reinsurer's falls from 0.07809 to 0.04535 as its best admissible coverage falls from 0.8460 to 0.0184; the gap is +0.03528 at $\delta = 0$ and -0.09337 at $\delta = 0.145$, and the smallest loading on the grid at which it is non-positive is $\delta^* = 0.045$. At $s = 0.50$, at δ^* and above it the rivals, pricing a high-risk book at a loading, outgrow the menu reinsurer in the separated configuration, and below it the gap is positive; δ^* is computed at that share alone.

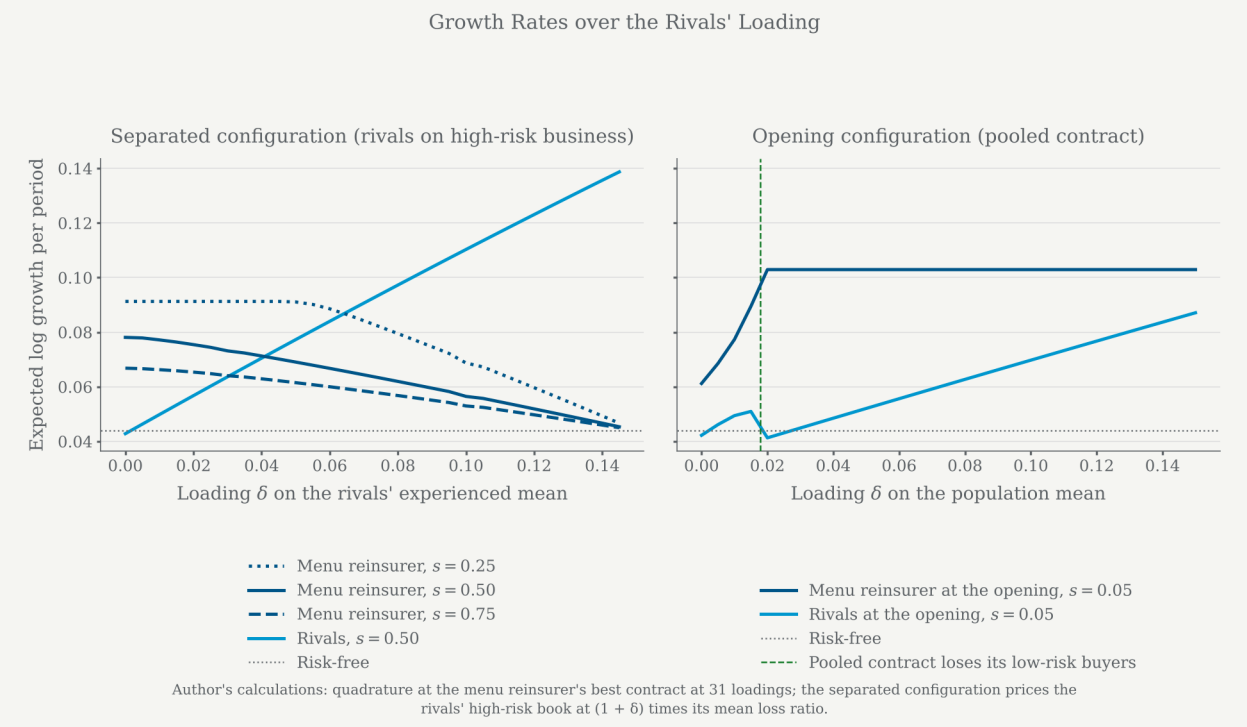


Figure 4: Expected logarithmic growth rates per period over the rivals' loading δ. Left panel: the separated configuration, the rivals' book high-risk business alone priced at $(1+\delta)m_H$, the menu reinsurer at its best admissible contract at shares of 0.25, 0.50, and 0.75 (dotted, solid, and dashed deep blue) and the rivals at $s = 0.50$ (bright blue); right panel: the opening configuration, the pooled contract priced at $(1+\delta)\bar{m}$ with the menu reinsurer at a share of 0.05; the dashed vertical marks δ_L, above which the pooled contract has no low-risk buyers; the dotted horizontal is the risk-free rate; the two panels share one vertical scale. Source: author's calculations (quadrature at 31 loadings).

The simulation at $\delta = 0.01$, inside the range in which the pooled contract has low-risk buyers at the opening, shows the regime with a positive margin for the rivals (Figure 5, right). Their price rises from 0.7871 to 0.8492 in cross-path mean; the cross-path mean of their realized logarithmic return over periods 50 through 150 is 0.04917, above $\ln(1+r_f) = 0.04402$: the rivals grow while they are being diluted. The share's cross-path mean is 0.540 at period 100 against 0.507 at a zero loading, and 0.899 at the final period. The closed-market state is reached here too, on 352 paths at the final period: the price exceeds $W_H(1)$ once the estimate exceeds $W_H(1)/(1+\delta)$, below the threshold $W_H(1)$ of a zero loading, while the rivals' book stays larger for longer. The two effects run against each other and the second dominates at this loading.



Figure 5: Experience-rated pricing at a zero loading (left column) and at a loading of one percent (right column). Upper panels: the menu reinsurer's share, as in Figure 3 (median, tenth-to-ninetieth percentile band, lowest share across paths, and the expected-growth path). Middle panels: the cross-path mean of the rivals' price per unit of subject premium (bright blue) and the high-risk type's willingness to pay for full coverage (dotted), on one vertical scale. Lower panels: the share of paths on which the rivals wrote no business in the period. Source: author's calculations (Monte Carlo, 1,000 paths, 300 periods, credibility 0.3).

8.4 The Monopolist's Menu

Once the rivals have left the market on a path, the menu reinsurer faces both types alone with no reinsurance as their outside option, and the cream-skimming contract has nothing to skim against. If it re-optimizes as a monopolist, offering the low-risk type a contract at its full willingness to pay and the high-risk type a contract whose price leaves that type the rent the low-risk contract would give it, its growth-optimal menu at the calibration, computed at a share of one on a grid of hundredths of each type's coverage, the two types sharing one queue up to its capacity, serves the high-risk type alone: coverage $\alpha_H = 0.99$, the grid's largest, at a price of 0.9120 per unit of subject premium, the low-risk type excluded ($\alpha_L = 0.00$), at an expected growth rate of 0.09500 per period; the best menu that serves both types earns 0.09444 at a low-risk coverage of 0.01. The exclusion is that of the monopoly screening problem of Stiglitz[25] and of the continuum result cited in Section 2: any coverage sold to the low-risk type hands the high-risk type a rent that costs more than the low-risk business earns. The reinsurer that won the market by serving the better risks serves, as a monopolist, the worse ones.

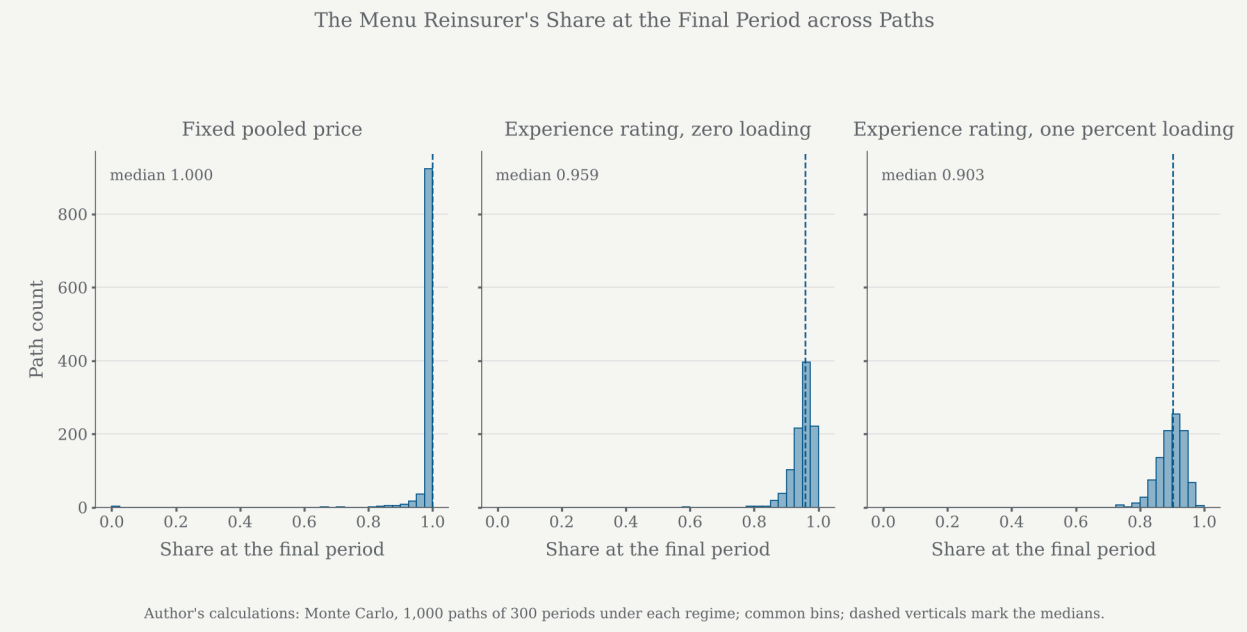


Figure 6: The menu reinsurer's share at the final period across 1,000 paths under the fixed pooled price (left), experience rating at a zero loading (center), and at a loading of one percent (right), on common bins; the dashed vertical in each panel marks the median. Under experience rating the mass below one holds the paths on which the market closed (Table 4), each at the share it had reached when it closed, beside paths on which it had not closed. Source: author's calculations (Monte Carlo, 300 periods).

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

9 Capacity, Correlation, and the Cycle

This section relaxes two of the market's assumptions: the independence of losses across submissions, and the scaling of demand with industry surplus. With a common loss factor the variance of a book's result no longer vanishes with diversification, and the growth-optimal choice then has a second dimension beside the contract: how much of its capacity the reinsurer uses. With demand tied to the economy rather than to surplus, a reinsurer's capacity relative to the business it is offered moves with its underwriting results, as in the capacity-constraint theories of the insurance cycle (Section 2).

9.1 The Common Factor, Demand, and Exposure

The loss ratio of submission i in a period is $X_i = \min\{\max\{0, Y_i\}, \bar{X}\}$ with

$$Y_i = \mu_\theta + \sigma_\theta\big(\sqrt{1-r}\,\varepsilon_i + \sqrt{r}\,\eta\big), \tag{11}$$

ε_i and η standard Normal, η common to every submission of the period and to both reinsurers' books, and $r \in [0, 1)$ the share of the loss-ratio variance that is common; the marginal distribution of each loss ratio is the baseline's at every r. Demand grows at the risk-free rate: each submission carries B units of subject premium in units of the initial industry surplus, and surpluses are carried in units of the demand's scale, so a reinsurer's capacity relative to the demand it faces moves with its underwriting results alone. The industry's capitalization is $C = \kappa(S_0 + S_R)/(NB\bar{m})$, capacity over the premium of writing every submission at the population mean loss ratio; at the calibration $C = 1.120$, and the scale-free market of Section 3 holds it there. The menu reinsurer writes premium up to $\phi\kappa S_0$ with the exposure $\phi \in [0.02, 1]$ chosen, with its coverage, to maximize its expected logarithmic growth over the period, by a coordinate search: the coverage as in Section 3 at full exposure, then the exposure by twenty-four golden-section steps at that coverage, full exposure kept where it earns at least the interior candidate, and the coverage once more where the exposure is below one; the rivals write what comes to them at the fixed pooled price, up to κS_R. The quadrature of Section 7 carries the common part of each book's variance as the square of the book's exposure,

$$\operatorname{Var}\frac{U}{S} = \Big(\frac{B}{S}\Big)^2\Big[(1-r)\big(\alpha_L^2\tilde{w}_L v_L + \alpha_H^2\tilde{w}_H v_H\big) + r\big(\alpha_L w_L\sqrt{v_L} + \alpha_H w_H\sqrt{v_H}\big)^2\Big], \tag{12}$$

with $\tilde{w}$ the arrival-order variance count of a count w of lines, so that the common part scales with the square of the book. It was measured at $r = 0.3$, a share of 0.25, and $C = 1$ against a Monte Carlo of 100,000 periods that draws every loss and the period's common factor: the menu reinsurer's expected growth rate is 0.06012 by the quadrature and 0.05994 by the Monte Carlo (-0.68 standard errors apart), the rivals' 0.02834 and 0.02825 (-0.23 standard errors apart). A reinsurer's capacity binds where the submissions that take its contract as their first choice, with the low-risk overflow in the rivals' case, exceed it. At that r and share the two capacities never bind together at $C = 1$, and they bind together with probability above 0.99999 at $C = 0.75$ and at $C = 0.5$, where the two capacities together fall short of the period's submissions. There the fraction q of Proposition 1 need not hold, and the quadrature computes each reinsurer's expected lines of each type exactly over the arrival orders and takes each book's Normal at those expected lines, without the dispersion of the book's composition across arrival orders: with the low-risk overflow taking C_P, as it does at the pooled price, after i low-risk and j high-risk arrivals the two reinsurers have used $\min\{K_0, i + (j - K_R)^+\}$ and $\min\{K_R, j + (i - K_0)^+\}$ lines, whatever the order of those arrivals, and the next arrival is low-risk with probability $(n_L - i)/(N - i - j)$. The same Monte Carlo at those

two capitalizations gives deviations of +1.55 and +0.68 standard errors for the menu reinsurer and +0.98 and +0.21 for the rivals.

9.2 The Gap and the Exposure over Correlation and Capitalization

Table 5 reports the growth gap and Figure 7 the gap and the rivals' growth rate at a share of 0.25 over the common share r and the capitalization C; the menu reinsurer's exposure at the same points is stated below. First, the exposure is the full capacity at every point (its smallest value 1.000), and the threshold at which the growth-optimal exposure leaves full capacity, by bisection in r on $[0, 0.98]$ in twelve steps, which presumes that it leaves full capacity once on that interval, is $r^* = 0.897$ at $C = 1$ and 0.877 at $C = 3$. At those two capitalizations the Kelly criterion sizes the menu reinsurer's book below its capacity only where losses are almost entirely common; among the simulated paths below, it does so on some at $r = 0.3$. Second, where the industry is capitalized at or below the calibration's level, the common factor widens the gap: at $C = 1$ the gap rises from +0.0279 at $r = 0$ to +0.0345 at $r = 0.5$, and at the calibration's capitalization from +0.0278 to +0.0312, because the common part of the variance scales with the square of a book's exposure and the rivals' book carries the higher volatility. Where the industry is capitalized at $C = 1.5$ and above, the common factor narrows the gap instead: at $C = 3$ it falls from +0.0179 to +0.0157, because the menu reinsurer, the smaller firm, then carries the larger book relative to its own surplus. Third, at every r the gap falls as capitalization rises from $C = 0.75$ through $C = 3$; at $C = 0.5$ it lies above its value at $C = 0.75$ for $r = 0$ through $r = 0.2$ and prints alike with it at $r = 0.3$ and $r = 0.5$. Up to the calibration's capitalization both reinsurers' growth rates rise with C, the rivals' by more; above it the menu reinsurer's falls while the rivals' goes on rising. At $C = 3$ the menu reinsurer's best coverage lies closer to one than at the calibration's capitalization at every r: the coverage distortion that separates the types is, in part, a response to its own capacity constraint. The rivals' growth rate (Figure 7, right) lies below $\ln(1 + r_f) = 0.04402$ at every point and rises with capitalization at every r.

Table 5: The growth gap $g_0 - g_R$ per period at a share of 0.25 under the fixed pooled price, over the common share r of the loss-ratio variance (rows) and the industry's capitalization C (columns; the column at 1.120 is the calibration's). At $C = 0.5$ and 0.75 both capacities bind, and the expected lines are computed exactly over the arrival orders (Section 9). The menu reinsurer's exposure is its full capacity at every cell (its smallest value 1.000). Source: author's calculations (quadrature).

r	$C = 0.5$	$C = 0.75$	$C = 1$	$C = 1.120$	$C = 1.5$	$C = 2$	$C = 3$
0.0	+0.0286	+0.0282	+0.0279	+0.0278	+0.0271	+0.0245	+0.0179
0.1	+0.0305	+0.0302	+0.0292	+0.0285	+0.0269	+0.0240	+0.0174
0.2	+0.0322	+0.0321	+0.0305	+0.0291	+0.0266	+0.0234	+0.0169
0.3	+0.0342	+0.0342	+0.0318	+0.0298	+0.0264	+0.0229	+0.0165
0.5	+0.0384	+0.0384	+0.0345	+0.0312	+0.0258	+0.0218	+0.0157

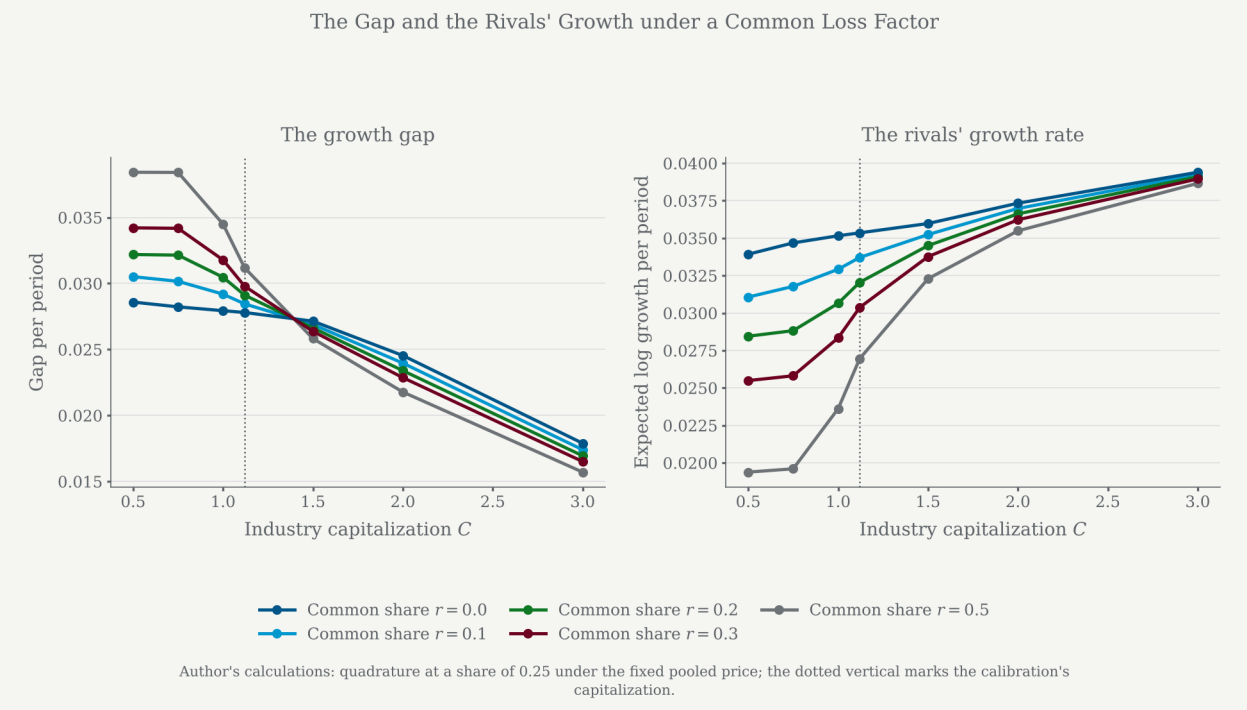


Figure 7: The growth gap (left) and the rivals' expected growth rate (right) over the industry's capitalization C, one line per common share r of the loss-ratio variance, at a share of 0.25 under the fixed pooled price; the dotted verticals mark the calibration's capitalization. Source: author's calculations (quadrature).

9.3 Surplus Paths with Demand Tied to the Economy

Figure 8 draws 500 surplus paths of 60 periods from a share of 0.05 at the calibration's capitalization, without a common factor and at $r = 0.3$, with demand growing at the risk-free rate; on the paths the coverage's golden section takes twenty steps. Without a common factor the industry's capitalization drifts down, from 1.120 to a cross-path mean of 1.088 at the final period: the rivals hold most of the surplus and grow below the risk-free rate at which demand grows (Proposition 2; their mean log gross return over the path-periods is 0.04080 against $\ln(1+r_f) = 0.04402$), and the menu reinsurer's faster growth does not offset that at its share; the rivals' capacity binds in 18.0 percent of the path-periods. At $r = 0.3$ a year in which the common factor is adverse reduces both surpluses together, the capitalization falls, and capacity binds: the rivals' capacity binds in 42.0 percent of the path-periods, in 970 spells of binding capacity whose mean length is 13.0 periods and whose longest is 59; the capitalization's cross-path mean is 1.101 at the final period. The menu reinsurer's share at the final period has a cross-path mean of 0.196 at $r = 0.3$ against 0.171 without a common factor. Its exposure leaves full capacity on some paths at $r = 0.3$, with a smallest value of 0.093, and on none without a common factor (smallest value 1.000); the paths do not identify the state of the market in the periods in which it does so. The hard market of the capacity-constraint theories appears here as a quantity: submissions the rivals cannot write after an adverse year, which spill to the menu reinsurer's contract or go unplaced. The price response that those theories add—a price that rises when capacity is short—is outside the two pricing regimes of Section 3, under which the rivals' price depends on the population mean or on their experience and not on their capacity.

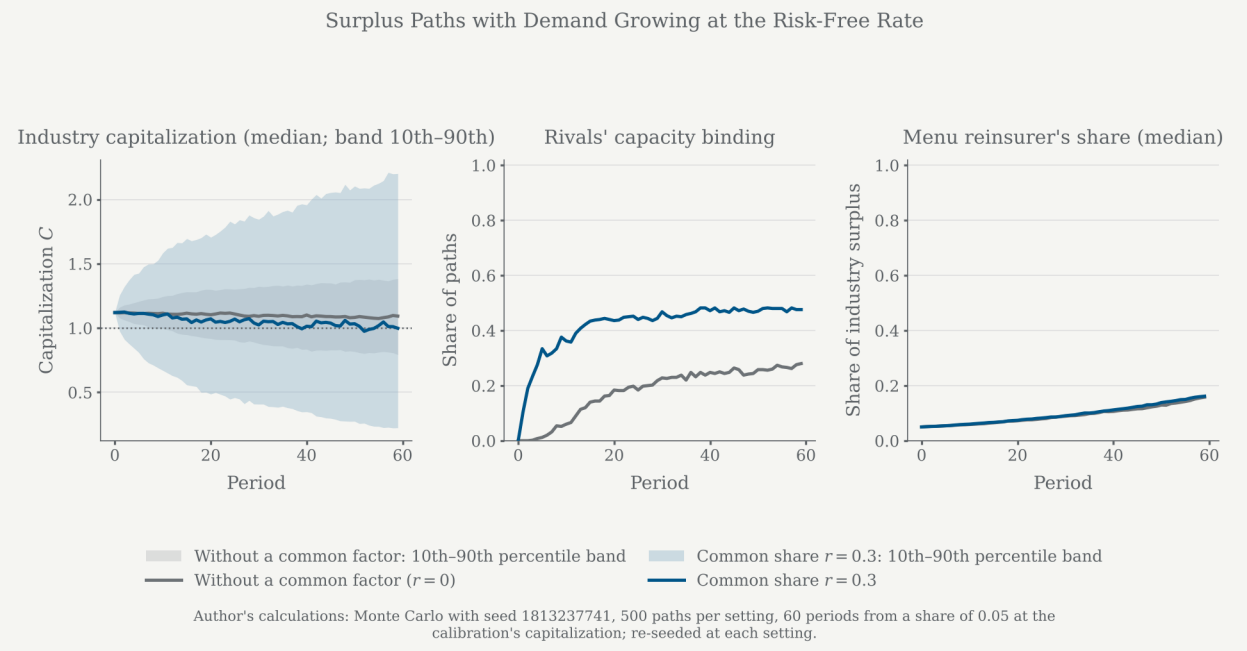


Figure 8: Surplus paths with demand growing at the risk-free rate, without a common factor (gray) and at a common share $r = 0.3$ of the loss-ratio variance (deep blue). Left panel: the industry's capitalization C (median; band tenth to ninetieth percentile, each setting's band in its own fill), the dotted line at $C = 1$; center panel: the share of paths on which the rivals' capacity binds in the period; right panel: the median of the menu reinsurer's share. The text reports cross-path means where the left and right panels draw medians. Source: author's calculations (Monte Carlo, 500 paths of 60 periods per setting from a share of 0.05).

10 The Continuum of Types

Sections 10 through 12 bring three features of the cedant population onto the market of Section 3, each under the calibration of Section 3, the rivals' pooled contract at the population mean of the types in hand, the capacity rule, the placement in arrival order, and the objective of Section 3: a continuum of types in place of two, two dimensions of private information in place of one, and heavy-tailed loss ratios in place of the Normal. Each section states what the menu reinsurer can do against the posted price, measures the growth gap, and compares the computation with a Monte Carlo that draws every loss.

10.1 Types, Contracts, and the Menu

The cedant's type is its expected loss ratio before the cap, θ, uniform on $[\underline{\theta}, \bar{\theta}] = [0.7499, 0.8086]$, the interval whose mean and variance are those of the two-type distribution of Section 3: the population mean $\lambda\mu_L + (1-\lambda)\mu_H$ and the between-type variance $\lambda(1-\lambda)(\mu_H - \mu_L)^2 = 0.29 \times 10^{-3}$. The volatility $\sigma(\theta)$ is affine through the two calibration points (μ_L, σ_L) and (μ_H, σ_H) and is held at σ_L below μ_L and at σ_H above μ_H: the calibration measures the volatility at two points, and the continuum does not extrapolate beyond them. The capped mean and variance $m(\theta)$ and $v(\theta)$ follow as in Section 3, and type θ's gain from a contract is (1) with $m(\theta)$ and $v(\theta)$ in place of m_θ and v_θ; write W_θ and u_θ for its willingness to pay and its gain. Since m is increasing and v non-decreasing

in θ, the marginal willingness to pay for coverage, $m(\theta) + \rho(1-\alpha)v(\theta)$, increases in θ: single crossing holds, and a type with a higher expected loss values any added coverage more. The rivals offer $C_P = (1, p_P)$ at the population mean of the capped means, $p_P = 0.7793$, which agrees at the printed precision with the price of the two-type market of Sections 3 through 8, the core below. The gain from C_P increases in θ and is zero at $\theta_c = 0.7516$: the 2.8 percent of types below θ_c would lose from the pooled contract, and their alternative to the menu is no reinsurance.

The menu reinsurer offers K partial-cover contracts (α_k, p_k) with $\alpha_1 < \cdots < \alpha_K < 1$. By single crossing a menu partitions the types by thresholds $\theta_1 < \cdots < \theta_K$: contract k is bought on $[\theta_{k-1}, \theta_k)$, with θ_0 the type indifferent between the first contract and no reinsurance (or $\underline{\theta}$ where every type gains from it), and the rivals' C_P is bought above θ_K. The prices follow from the thresholds by indifference, as the price of the family (2) follows from the low-risk type's: θ_K is indifferent between contract K and C_P, and θ_k between contracts k and $k+1$,

$$p_K = W_{\theta_K}(\alpha_K) - u_{\theta_K}(C_P), \qquad p_k = p_{k+1} - \big[W_{\theta_k}(\alpha_{k+1}) - W_{\theta_k}(\alpha_k)\big], \quad k < K. \tag{13}$$

Every type's gain from its own contract is non-negative along the chain: the lowest type of each class gains from its contract what it gains from the contract below, down to θ_0, where the gain is zero; the chain ends at $u_{\theta_K}(C_P) \geq 0$, so the construction requires $\theta_K \geq \theta_c$. A menu is admitted only if each of its contracts is bought by an interval of types of probability at least $1/N$, one submission a period in expectation; without that requirement a contract can be left to the rivals' rationed overflow alone, which is not a screening menu. The menu reinsurer chooses the $2K$ parameters (α_k, θ_k) to maximize its expected logarithmic growth over the period at its share s, as in Section 3, against the fixed pooled price.

10.2 The Computation

Given the menu, each submission falls into one of $K+2$ classes—uninsured below θ_0, the K contracts, the rivals' interval above θ_K—with the probabilities of the intervals, and the period's class counts are multinomial. Given the counts, the placement is that of Section 3 in expectation: each contract's class goes to the menu reinsurer up to its capacity and the overflow of types above θ_c to C_P; the rivals' class goes to C_P up to their capacity, and their rationed overflow to the menu's top contract, which leaves every type above θ_K a non-negative gain, up to the menu reinsurer's remaining capacity. Every contract of a menu is priced below p_P by (13), so the two capacities in lines sum to more than N, as in Section 5, and the two reinsurers never ration together. A class's book loss per submission is the Normal with the class's mean and variance—the mean of $m(\theta)$ over the interval, and the mean of $v(\theta)$ plus the variance of $m(\theta)$ over it, the between-type variance within the class—by a Gauss–Legendre quadrature of 32 nodes on the interval; the arrival-order variance and the expectation of the floored logarithm are those of Section 7. The expectation over the counts is a Monte Carlo over 20,000 count vectors, drawn once at each share and held fixed across every menu compared at that share, so that the search compares menus on common draws. The search adds one contract to the best menu of $K-1$ contracts on a grid of 40 coverages from 0.05 to 0.98 by 40 thresholds across the type interval and then refines every parameter by 3 rounds of golden sections of 24 iterations each.

10.3 Results

Table 6 reports the gap at each share for $K = 1, \dots, 5$ beside the two-type core's gap at the same share (Section 7), Table 7 the one- and two-contract menus, and Figure 9 draws the gap over K and three coverage schedules.

Table 6: The continuum of types: the growth gap $g_0 - g_R$ per period at each share for menus of up to 5 contracts, and the two-type core's gap at the same share. Source: author's calculations (count Monte Carlo over 20,000 count vectors at each share; the core's gap from the quadrature of Section 7).

s	$K = 1$	$K = 2$	$K = 3$	$K = 4$	$K = 5$	Two-type core
0.05	0.03827	0.03838	0.03844	0.03848	0.03852	0.01898
0.25	0.03193	0.03220	0.03222	0.03222	0.03223	0.02780
0.50	0.03132	0.03138	0.03140	0.03140	0.03140	0.03596
0.75	0.01048	0.01049	0.03395	0.03397	0.03400	0.03778

Table 7: The continuum of types: the one- and the two-contract menus at each share—coverages, prices, thresholds, and the share of types that buys from the menu. Source: author's calculations.

	One contract				Two contracts						
s	α_1	p_1	θ_1	Types, %	α_1	α_2	p_1	p_2	θ_1	θ_2	Types, %
0.05	0.555	0.4367	0.7574	12.8	0.546	0.589	0.4296	0.4637	0.7550	0.7574	12.8
0.25	0.911	0.7108	0.7681	30.9	0.832	0.953	0.6506	0.7431	0.7564	0.7681	31.0
0.50	0.947	0.7379	0.7796	50.5	0.943	0.962	0.7350	0.7495	0.7740	0.7795	50.5
0.75	0.479	0.3792	0.7531	5.4	0.478	0.483	0.3788	0.3826	0.7515	0.7531	5.4

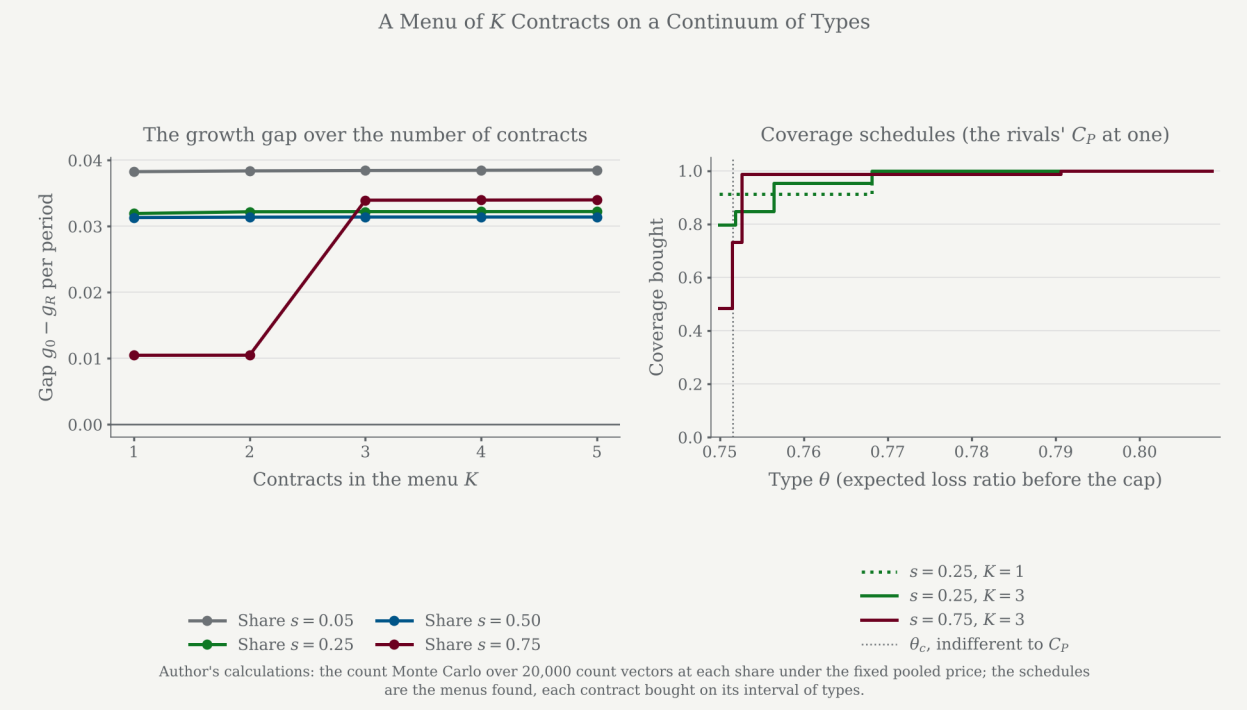


Figure 9: A menu of K contracts on a continuum of types. Left panel: the growth gap $g_0 - g_R$ per period over the number of contracts K, one line per share. Right panel: the coverage bought over the type θ under the menus found—at $s = 0.25$ with one contract (dotted) and three (solid green), and at $s = 0.75$ with three (maroon)—each contract on its interval of types, the rivals' C_P at a coverage of one above the top threshold, and the dotted vertical at θ_c, the type indifferent to C_P. Source: author's calculations (count Monte Carlo over 20,000 count vectors at each share under the fixed pooled price).

First, the gap is positive with one contract at every share, from 0.03827 at $s = 0.05$ to 0.01048 at $s = 0.75$, and it is non-decreasing in K at every share. Second, where the first contract is a cream-skimming contract, at $s = 0.05$, 0.25, and 0.50, the second contract adds 0.00011, 0.00027, and 0.00006 to the gap and the third 0.00006, 0.00002, and 0.00002, less than the second at each of the three shares. The menu reinsurer's own growth need not rise with a contract: at $s = 0.05$ the five-contract menu earns 0.07906 against the four-contract menu's 0.07910, because each contract's class must hold a submission in expectation, while the gap still rises, from 0.03848 to 0.03852, because the rivals' growth falls by more.

Third, with one contract the gap exceeds the two-type core's at $s = 0.05$ and 0.25 (0.03827 against 0.01898, and 0.03193 against 0.02780) and falls below it at $s = 0.50$ (0.03132 against 0.03596). The contract is priced at the marginal type's indifference, θ_1's, and every type below θ_1 pays that price on a lower expected loss, so the margin on the class rises with the distance between the marginal type and the class's mean and falls with the marginal type's gain from the pooled contract, which the price must leave it. At $s = 0.05$ the class is the 12.8 percent of types below $\theta_1 = 0.7574$, whose capped mean is 0.7537 against $m_L = 0.7654$ and whose marginal type gains $u_{\theta_1}(C_P) = 0.0059$ from the pooled contract against the core's low-risk type's $u_L(C_P) = 0.01384$: the menu reinsurer's expected margin per unit of subject premium on the class is 0.0184 against the core's 0.0106 at the same share. At $s = 0.50$ the class is the 50.5 percent of types below $\theta_1 = 0.7796$, whose capped mean is 0.7648 against $m_L = 0.7654$ and whose marginal type gains 0.0580 from the pooled contract; the margin is 0.0136 against the core's 0.0135, and the class's variance per

submission is 0.0180 against $v_L = 0.0144$, the class reaching types whose volatility lies above σ_L. The gap's shortfall at that share lies in the two books: the menu reinsurer writes 25.4 lines in expectation, fewer than the core's (Table 2), and its growth rate lies 0.00260 below the core's, while the rivals' lies 0.00203 above theirs. The between-type variance within a class is negligible beside the within-type variance: at $s = 0.05$, $(\theta_1 - \underline{\theta})^2/12 = 4.7 \times 10^{-6}$ against 0.0144.

Fourth, at $s = 0.75$ the one-contract menu is of another kind: coverage 0.479 at a price of 0.3792, bought by the 5.4 percent of types below $\theta_1 = 0.7531$ and, through the rivals' rationing, by the submissions the rivals cannot write. At that share the rivals' capacity is 14.0 lines of $N = 50$ and binds in every count draw, and the 36.0 lines the menu reinsurer writes are mostly rationed cedants of every type, who take its contract at a price of 0.7921 per unit of coverage, above the population mean. The menu reinsurer's growth is nearly the same under that contract and under a cream-skimming menu at high coverage: the three-contract menu, whose top contract at a coverage of 0.986 and a price of 0.7682 is bought by the 64.7 percent of types between θ_2 and $\theta_3 = 0.7906$ (the menu's three contracts by the 69.3 percent below θ_3), earns the menu reinsurer 0.05088 against the one contract's 0.05063, a difference of 0.00025, while the rivals' growth is 0.04015 under the first and 0.01693 under the second, and the gap jumps from 0.01049 to 0.03395 at the third contract. Where the rivals' capacity binds in every period, the menu reinsurer is nearly indifferent between skimming the better risks and pooling the rationed ones at a low coverage, and the rivals' growth, the gap's other half, is decided by that near-indifference. The three-contract gap at $s = 0.75$ lies below the core's 0.03778 in both books: the menu reinsurer's growth rate lies below the core's and the rivals' above theirs (Table 2). The menu's top class, 64.7 percent of the types, reaches types whose expected loss ratio is above the core's low-risk type's.

The count Monte Carlo was measured at $s = 0.50$ for $K = 1$ and $K = 2$ against a Monte Carlo of 100,000 periods that draws every type and every loss and places every submission in arrival order: the menu reinsurer's growth deviates by -0.34 and -0.47 standard errors and the rivals' by -2.24 and -2.14. The rivals ration in 14 percent of the count draws at that share; the cause of their deviation is not established.

In the monopoly of Cheung et al. with a continuum of types, cited in Section 2, the optimal menu excludes the better risks below a threshold rather than distorting their coverage. Here, against a posted price, the menu reinsurer serves the better risks below a threshold and leaves the worse ones to the rivals. The two thresholds cut the same type space from opposite ends: a monopolist excludes the types whose rent would cost more than their business earns, and a reinsurer that competes against a pooled price takes the types whose business earns more than the price leaves them.

11 Two Dimensions of Private Information

The cedant's private information in Section 3 is one number, its expected loss ratio, and its volatility moves with it. This section separates the two: the type is the pair $(\mu, \sigma) \in \{\mu_L, \mu_H\} \times \{\sigma_L, \sigma_H\}$, four types, the dimensions independent, with the share of μ_L the calibration's λ and the share of σ_L one half. Table 8 lists the four types' capped moments, shares, and gains from the pooled contract, which the rivals price at the population mean of the four capped means, $p_P = 0.7793$. The marginal willingness to pay for coverage, $m+\rho(1-\alpha)v$, ranks the types by the mean at coverages near one and by the volatility at lower coverages, so that a second contract can separate on the volatility where one contract separates on the mean. The high-volatility types value the pooled contract most: the low-mean, high-volatility type LH gains 0.1061 from it against the low-mean, low-volatility type

LL's 0.0138, because full coverage at a price near its mean removes a variance more than four times the other's.

Table 8: The four types: the Normal parameters, the capped mean and variance, the population share, and the gain from the pooled contract C_P at $p_P = 0.7793$. The first letter of a type is its mean (L for μ_L, H for μ_H) and the second its volatility (L for σ_L, H for σ_H). Source: author's calculations.

Type	μ	σ	m	v	**Share**	$u(C_P)$
LL	0.7654	0.12	0.7654	0.0144	0.30	0.0138
LH	0.7654	0.25	0.7655	0.0624	0.30	0.1061
HL	0.8000	0.12	0.8000	0.0144	0.20	0.0484
HH	0.8000	0.25	0.8000	0.0624	0.20	0.1408

The menu reinsurer offers one contract or two. For each coverage on the grid of hundredths from 0.05 to 0.99, the candidate prices are those at which one of the four types is indifferent between the contract and its best alternative—for the contract of the lower coverage the other contract, C_P, or no reinsurance, and for the contract of the higher coverage C_P or no reinsurance—every assignment of indifferent types enumerated, and each candidate menu is evaluated over every count: the period's counts of the four types are a multinomial sum enumerated in full, 23,426 count vectors; the placement is that of Section 3 in expectation, first choices and then each side's overflow to the other's best acceptable contract up to its remaining capacity; the books and the expectation are those of Section 7. Every contract of the menus found is priced below p_P (Table 9), so the two capacities in lines sum to more than N, and the two reinsurers never ration together. The two-contract menu is searched on a coarse grid of coverage pairs from 0.05 at a step of 0.04, each candidate admitted where each contract is some type's first choice, and then on the hundredths within 0.04 of the coarse best.

Table 9 reports the menus found and Figure 10 draws the gaps and the coverages. With one contract the menu is bought by LL at every share and never by HH: at $s = 0.05$ and 0.25 by LL alone, and at $s = 0.50$ and 0.75 by LL and LH together, at coverages of 0.96 and 0.97, near one, where a single contract screens on the mean alone and pools the two volatilities. With two contracts the contract with the higher coverage is bought by LH and the one with the lower coverage by LL at every share: the second contract separates on the volatility. The second contract adds 0.00103 to the gap at $s = 0.25$ and 0.00055 at $s = 0.50$, 4.6 and 1.6 percent of the one-contract gap; it adds nothing at $s = 0.75$, where the two coverages found differ by one hundredth and the gap by -0.00001; and at $s = 0.05$ any second contract that some type buys lowers the gap, from 0.01906 to 0.01079, because at that share the menu reinsurer's capacity is filled by the first contract's low-volatility business, and a contract the high-volatility type takes displaces that business in arrival order with lines of more than four times the variance. The one-contract gap at $s = 0.50$, 0.03378, lies below the two-type core's 0.03596: the pooled price now carries the low-mean, high-volatility type, which values the pooled contract more than the other low-mean type does, and a contract that draws that type must leave it its gain from C_P. The quadrature over the enumerated counts was measured at $s = 0.50$ against a Monte Carlo of 100,000 periods: the menu reinsurer's growth deviates by -0.31 and -0.29 standard errors under the one- and the two-contract menus and the rivals' by +0.00 and +0.05.

Table 9: Two dimensions of private information: the one- and the two-contract menus at each share—coverages, prices, the gap $g_0 - g_R$ per period, and the types whose first choice each contract is (for two contracts, the first contract's buyers before the semicolon and the second's after it). Source: author's calculations (quadrature over the 23,426 count vectors, enumerated).

	One contract				Two contracts					
s	α	p	Gap	Bought by	α_1	p_1	α_2	p_2	Gap	Bought by
0.05	0.66	0.5158	0.01906	LL	0.81	0.6324	0.93	0.7251	0.01079	LL; LH
0.25	0.87	0.6793	0.02242	LL	0.74	0.5783	0.96	0.7485	0.02344	LL; LH
0.50	0.96	0.7485	0.03378	LL, LH	0.87	0.6793	0.98	0.7639	0.03433	LL; LH
0.75	0.97	0.7562	0.03576	LL, LH	0.96	0.7485	0.97	0.7562	0.03575	LL; LH



Figure 10: Two dimensions of private information. Left panel: the growth gap $g_0 - g_R$ per period at each share with one contract (deep blue), with two (pastel), and in the two-type core of Section 7 (gray). Right panel: the coverages of the contracts found over the share—the one contract (deep blue), and the two-contract menu's lower coverage, bought by LL (green), and higher coverage, bought by LH (maroon). Source: author's calculations (quadrature over the 23,426 count vectors of the four types, enumerated, under the fixed pooled price).

The second dimension adds at most 4.6 percent to the one-contract gap. The type a second contract draws is the low-mean type for which the pooled contract is worth most. The low-mean, high-volatility type pays the pooled price, 0.7793, for a capped mean of 0.7655, and the variance the pooled contract removes is worth 0.1199 to it, a net gain of 0.1061; a contract that draws it from C_P must leave it that gain, and at $s = 0.25$ the contract it buys earns the menu reinsurer 0.0136 per unit of subject premium on its line against 0.0118 on the low-volatility type's.

12 Heavy-Tailed Losses

The loss ratio of Section 3 is a Normal floored at zero and capped at $\bar{X}$. This section replaces the Normal by two heavier-tailed families with the same mean and variance before the floor and the cap: the log-normal, and Student's t with $\nu = 4$ degrees of freedom scaled to the type's mean and variance (its variance is finite and its fourth moment is not), each then floored and capped as the Normal is. The contracts are held at those of Section 7, chosen on the Normal's capped moments, although each family's capped moments differ from the Normal's (Table 11): at $s \in \{0.05, 0.25, 0.50\}$ the menu reinsurer offers its best contract of Section 7 under the fixed pooled price, and a Monte Carlo of 100,000 periods at each share draws every loss ratio from each family in turn and places every submission in arrival order. The three families are drawn in every period from the same types, the same arrival order, and the same standard Normal variates—Student's t adds its chi-square draws—so that each family's difference from the Normal carries a standard error of its own, that of a paired difference.

Table 10 reports the growth rates and the gap by family and Figure 11 draws them. Under Student's t the gap is larger than under the Normal at every share, by 0.00080, 0.00074, and 0.00116 at $s = 0.05$, 0.25, and 0.50 (7.4, 11.0, and 14.9 standard errors of the paired difference); under the log-normal it is smaller, by 0.00006, 0.00007, and 0.00008 (2.0, 2.6, and 2.3 standard errors), a few tenths of a percent of the gap.

Table 10: Heavy-tailed losses: the two expected growth rates per period and their gap under each family at three shares, at the menu reinsurer's best contract under the fixed pooled price, from a Monte Carlo of 100,000 periods at each share; standard errors in parentheses; the final column is the family's gap less the Normal's on the same draws, with the standard error of that paired difference. Source: author's calculations.

s	Family	g_0	g_R	$g_0 - g_R$	Less the Normal's
0.05	Normal	0.06102 (0.00022)	0.04220 (0.00009)	0.01882 (0.00024)	
	Log-normal	0.06101 (0.00022)	0.04225 (0.00009)	0.01876 (0.00024)	−0.00006 (0.00003)
	Student's t	0.06106 (0.00022)	0.04144 (0.00009)	0.01963 (0.00023)	+0.00080 (0.00011)
0.25	Normal	0.06299 (0.00010)	0.03539 (0.00011)	0.02760 (0.00015)	
	Log-normal	0.06299 (0.00010)	0.03545 (0.00011)	0.02754 (0.00015)	−0.00007 (0.00003)
	Student's t	0.06284 (0.00010)	0.03450 (0.00011)	0.02834 (0.00015)	+0.00074 (0.00007)
0.50	Normal	0.06075 (0.00008)	0.02494 (0.00016)	0.03581 (0.00018)	
	Log-normal	0.06076 (0.00008)	0.02503 (0.00016)	0.03573 (0.00018)	−0.00008 (0.00004)
	Student's t	0.06062 (0.00008)	0.02365 (0.00016)	0.03697 (0.00017)	+0.00116 (0.00008)

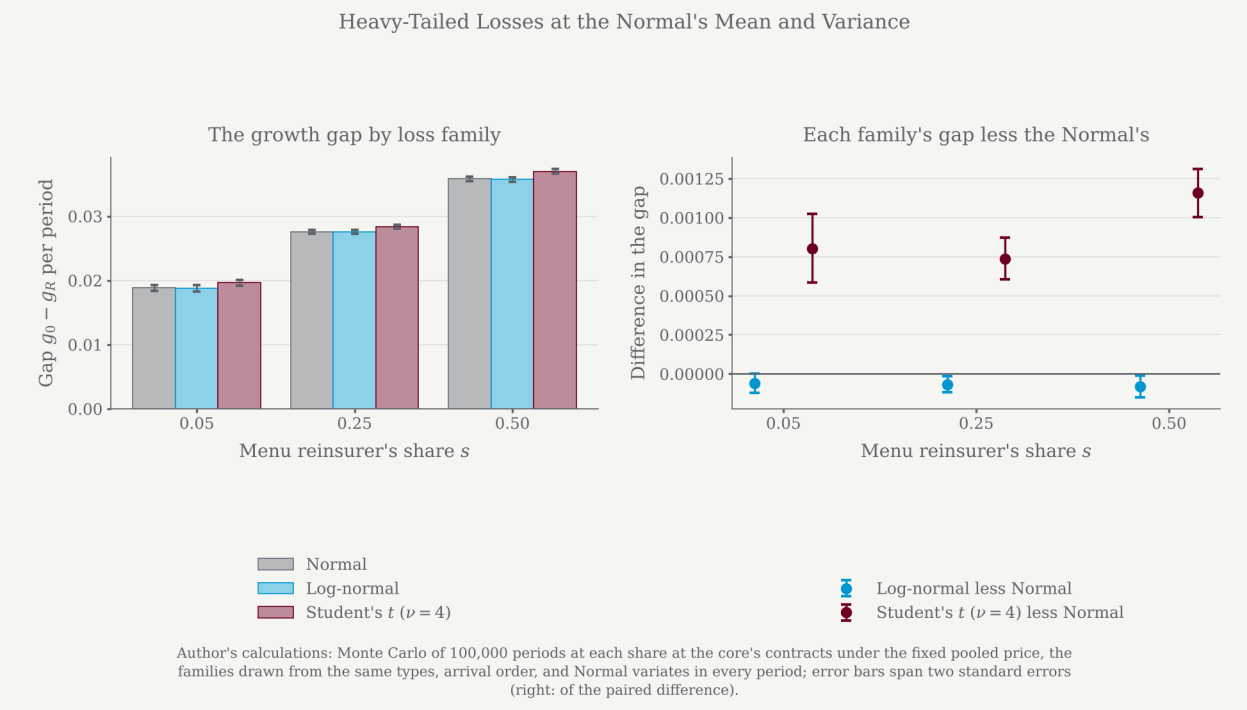


Figure 11: Heavy-tailed losses at the Normal's mean and variance. Left panel: the growth gap $g_0 - g_R$ per period at each share under the Normal (gray), the log-normal (bright blue), and Student's t with four degrees of freedom (maroon), with error bars of two standard errors. Right panel: each family's gap less the Normal's on the same draws, with error bars of two standard errors of the paired difference. Source: author's calculations (Monte Carlo of 100,000 periods at each share at the contracts of Section 7).

What the tail changes is the moments after the floor and the cap, and the growth rates follow those moments. Table 11 reports each family's capped mean and variance by type over the run's draws. Under Student's t the high-risk type's capped mean is higher than under the Normal by 0.0015 and its capped variance lower by 6.9 percent: with the same mean and variance before the floor, the t puts 0.53 percent of the high-risk loss ratios below zero against the Normal's 0.07 percent, the floor lifts them to zero, and the capped mean rises; the cap lies 8.8 standard deviations above the high-risk mean. The low-risk mean lies 6.4 standard deviations above zero, and the floor moves its capped mean by +0.0001. The rivals price at p_P whatever the family and carry the high-risk business (20.0 high-risk lines in expectation at $s = 0.50$, beside 0.8 low-risk), so the shift of the high-risk capped mean is a shift of their margin. The quadrature of Section 7, run with each family's capped moments in the books' returns and the contracts and choices standing on the Normal's, accounts for the measured differences: under Student's t it gives a change in the gap of +0.00066, +0.00077, and +0.00116 at the three shares against the measured +0.00080, +0.00074, and +0.00116, residuals of +1.4, −0.5, and −0.0 standard errors; under the log-normal, whose lower tail is lighter than the Normal's so that the floor removes less and the high-risk capped mean is lower by 0.00011 on the same draws, it gives −0.00005, −0.00006, and −0.00009 against the measured −0.00006, −0.00007, and −0.00008, residuals of −0.4, −0.3, and +0.3 standard errors. The shape of a book's loss distribution beyond its first two moments, through which a heavier tail would lower expected logarithmic growth at given moments, plays no measurable part: at the three shares the floor on the gross return lies at least 14.6 standard deviations below the mean of each book's gross return at every count.

Table 11: The capped mean and variance of each type's loss ratio under each family, over the draws of the Monte Carlo (100,000 periods of $N = 50$ submissions at each share, the same draws at every share); the Normal's exact values are $m_L = 0.7654$, $v_L = 0.0144$, $m_H = 0.8000$, and $v_H = 0.0624$. Source: author's calculations.

Family	m_L	v_L	m_H	v_H
Normal	0.7655	0.0144	0.8000	0.0623
Log-normal	0.7655	0.0144	0.7998	0.0624
Student's t	0.7657	0.0141	0.8015	0.0580

In this market a heavier tail moves the floored loss ratio's mean for the type that carries it, up under Student's t and down under the log-normal, and it is the single-price rivals, who write the high-volatility business at a price set on the Normal's moments, who carry that shift in their margin.

13 Conclusion

A reinsurer that screens with a menu and one that quotes a single price compete, in this paper, through the growth of their surplus, and the competition is decided by the growth gap between them. The cream-skimming contract that the menu reinsurer offers beside the rivals' pooled contract exists and earns a positive expected margin on every low-risk line whenever the pooled price exceeds the low-risk type's expected loss and the high-risk type gains from the pooled contract (Theorem 1 and Section 4); the rivals' book is left with the high-risk business and the low-risk overflow, and at the pooled price of a competitive market, wherever one of the two reinsurers can write all the business that comes to it, their expected underwriting result is negative and their growth rate lies below the risk-free rate (Propositions 1 and 2); and where the gap is bounded away from zero the menu reinsurer's share of industry surplus converges to one almost surely (Theorem 2). At the calibration the gap is positive and the rivals' growth rate below the risk-free rate at every share on a grid of ten-thousandths, and on simulated paths from a share of 0.05 the cross-path mean share reaches 0.992 after 300 periods, with the dispersion that a random walk from a small base carries. Under experience-rated pricing, with the menu reinsurer offering the high-risk type no contract, the outcome turns on the rivals' loading: at a zero loading their price rises toward the high-risk type's cost, their growth rate falls below the risk-free rate, and on 689 of 1,000 paths the noise of a small book lifts the price above that type's willingness to pay, a state in which no reinsurance is written; at a loading of 0.01 they grow while they are diluted; at a share of one half, at loadings of 0.045 and above, they outgrow the menu reinsurer in the separated configuration. The reinsurer that won the market by serving the better risks would, as a monopolist, serve the worse ones alone. At a share of 0.25 of industry surplus, correlated losses widen the gap where the industry is capitalized at or below the calibration's level and narrow it at $C = 1.5$ and above, and at $C = 1$ and $C = 3$ the Kelly criterion sizes the menu reinsurer's book below its capacity only where losses are almost entirely common; among simulated paths with a common factor, it does so on some.

The three further sections carry the mechanism onto richer cedant populations. On a continuum of types a menu partitions the types by thresholds, one contract priced at the marginal type's indifference earns a gap above the two-type core's at $s = 0.05$ and 0.25, and where the first contract skims the better risks the second adds less to the gap than the first and the third less than the second; where the rivals' capacity binds in every period, the menu reinsurer is nearly indifferent

between skimming the better risks and pooling the rationed ones, and the gap turns on that near-indifference. With two dimensions of private information, the mean and the volatility, one contract screens on the mean and a second separates on the volatility, and the second adds at most 4.6 percent to the one-contract gap; the type it draws is the low-mean type the pooled contract serves best. Heavy-tailed losses at the Normal's mean and variance change the gap only through the moments after the floor and the cap: against rivals who price on the Normal's moments, Student's t raises the high-risk type's capped mean and widens the gap, and the log-normal lowers that mean and narrows the gap by a few tenths of a percent.

Bibliography

Akerlof, George A. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics* 84, no. 3 (1970): 488–500.

Ania, Ana B., Thomas Tröger, and Achim Wambach. "An Evolutionary Analysis of Insurance Markets with Adverse Selection." *Games and Economic Behavior* 40, no. 2 (2002): 153–184.

Blume, Lawrence, and David Easley. "Evolution and Market Behavior." *Journal of Economic Theory* 58, no. 1 (1992): 9–40.

Blume, Lawrence, and David Easley. "If You're So Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets." *Econometrica* 74, no. 4 (2006): 929–966.

Bühlmann, Hans. "Experience Rating and Credibility." *ASTIN Bulletin* 4, no. 3 (1967): 199–207.

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 Yiying 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 Yiying Zhang. "Optimal Design of Reinsurance Contracts under Adverse Selection with a Continuum of Types." arXiv preprint arXiv:2504.17468v1, 24 April 2025.

Evstigneev, Igor V., Thorsten Hens, and Klaus Reiner Schenk-Hoppé. "Evolutionary Finance." In *Handbook of Financial Markets: Dynamics and Evolution*, edited by Thorsten Hens and Klaus Reiner Schenk-Hoppé, 507–566. Amsterdam: North-Holland, 2009.

Gron, Anne. "Capacity Constraints and Cycles in Property-Casualty Insurance Markets." *RAND Journal of Economics* 25, no. 1 (1994): 110–127.

Hall, P., and C. C. Heyde. *Martingale Limit Theory and Its Application*. New York: Academic Press, 1980.

Inderst, Roman, and Achim Wambach. "Competitive Insurance Markets under Adverse Selection and Capacity Constraints." *European Economic Review* 45, no. 10 (2001): 1981–1992.

Jones, Donald A., and Hans U. Gerber. "Credibility Formulas of the Updating Type." *Transactions of the Society of Actuaries* 27 (1975): 31–46.

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

Kong, Edward, Timothy Layton, and Mark Shepard. "Adverse Selection and (un)Natural Monopoly in Insurance Markets." NBER Working Paper No. 33187, 2024.

Lund, Diderik, and Tore Nilssen. "Cream Skimming, Dregs Skimming, and Pooling: On the Dynamics of Competitive Screening." *Geneva Papers on Risk and Insurance Theory* 29, no. 1 (2004): 23–41.

MacLean, Leonard C., Edward O. Thorp, and William T. Ziemba, eds. *The Kelly Capital Growth Investment Criterion: Theory and Practice*. Singapore: World Scientific, 2011.

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.

Miyazaki, Hajime. "The Rat Race and Internal Labor Markets." *Bell Journal of Economics* 8, no. 2 (1977): 394–418.

Pauly, Mark V. "Overinsurance and Public Provision of Insurance: The Roles of Moral Hazard and Adverse Selection." *Quarterly Journal of Economics* 88, no. 1 (1974): 44–62.

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.

Sandroni, Alvaro. "Do Markets Favor Agents Able to Make Accurate Predictions?" *Econometrica* 68, no. 6 (2000): 1303–1341.

Spence, Michael. "Product Differentiation and Performance in Insurance Markets." *Journal of Public Economics* 10, no. 3 (1978): 427–447.

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, *Theory and Methodology*, edited by Stavros A. Zenios and William T. Ziemba, 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.

Winter, Ralph A. "The Dynamics of Competitive Insurance Markets." *Journal of Financial Intermediation* 3, no. 4 (1994): 379–415.