

The Kelly Criterion and Learning in Competitive Insurance Underwriting

Frank Schmid and Can Wang

frank.schmid@frankschmid.com ken.wang.can@gmail.com

This version: 26 October 2026

First version: 17 March 2026

Abstract. A rating plan gives an insurer a price for an account and no statement of what the price is for. This paper gives underwriting an objective, the long-run growth of the insurer's surplus, which the Kelly criterion states as the maximization of the expected mean geometric rate of growth, and asks what an insurer must know to pursue it and what of that it can learn from its own operations. The criterion names four things: the account's expected loss, the variability of that loss against the insurer's surplus, the value of keeping the account, and the chance that a quote wins it; the price it chooses from the four is the Kelly decision. A rating plan supplies a version of the first. The insurer can learn all four from what it already holds, its claim files, the loss runs of the accounts it is asked to quote, its own retention, and the broker's report of the quotes it lost at market, and it can put what it learns to work in two ways: through its rating plan, with the learned estimate in place of the plan's, or through the Kelly decision, as a deviation from the price the plan gives the learned estimate. In a stylized market for middle-market general liability beside three conventional insurers, both ways end with equivalent growth of surplus, the highest in the market, within a margin set after their difference was known. They differ in the book each builds, in the prices each quotes, and in the path by which each arrives. Learning is what grows the surplus beyond a conventional insurer's.

JEL classification: G22, D83, L11, C63.

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

Contents

Abbreviations and Acronyms	ii
Greek Letters	ii
Latin Symbols with More Than One Meaning	iii
1 Introduction	1
2 Growth as the Objective, and Learning as Its Source	4
3 A Stylized Market	6
4 What Learning Is Worth	8
5 Two Routes to Equivalent Growth	11
6 When the Routes Part	13
7 What Follows for an Insurer, and Conclusion	15
A The Market: Its Assumptions, Its Insurers, and Its Loss Process	17
B The Rating Plans	24
C The Capital Standard	25
D The Kelly Decision and the Estimators	26
E The Kelly Route's Computation, Step by Step	32
F The Registry of Predictions and Outcomes	34
G The Quote Data and the Further Markets	37
Bibliography	40

Abbreviations and Acronyms

ACORD	Association for Cooperative Operations Research and Development
GL	General liability
GLM	Generalized linear model
IRIS	Insurance Regulatory Information System
NAIC	National Association of Insurance Commissioners
RBC	Risk-based capital

Greek Letters

α	alpha	Permanent part of a trait (§A.6)
β	beta	Discount factor of the renewal credit: at the insurer's own growth of surplus, the return on funds included, or at the return on funds (§D.1.3; Section 4 names the two discounts)
Δ	Delta	Logarithm of the holder's quote over the lowest invited quote (§A.4)
ε	epsilon	Expenses as a share of premium (§A.3)
ζ	zeta	Latent log price of the hold curve, over the price its deviation is taken from (§D.2.2)
η	eta	Yearly shock to the drifting part of a trait (§A.6)
κ	kappa	The capital requirement's scale (Appendix C)
λ	lambda	An account's expected number of claims in a policy year (§A.1); $\hat{\lambda}$, the tracked claim rate (Equation D.5); λ_0 , the claims a year at the median size (§A.6)
μ_ζ	mu	Mean of the hold curve's ζ (§D.2.2)
ν	nu	Mean of the permanent traits' sum, in the schedule modification (Appendix B)
ξ	xi	Standard normal draw of a case reserve's error (§A.6)
π	pi	Provision for expenses and profit in a plan's premium (§A.3); the posterior density of the hold curve's parameters (Equation D.6)
ρ	rho	Persistence of the drifting part of a trait (§A.1); the capital requirement's reserve-charge ratio (Appendix C); with subscript f or s , the fitted persistence of an autocovariance (§D.3)
σ	sigma	Spread of the hold curve's ζ (§D.2.2); with subscript c , of a case reserve's error (§A.6)
τ	tau	Mean time from booking to payment, years (§D.1.2)
Σ	Sigma	Covariance of a file's observations (§D.3); the summation sign (Equation D.3, Appendix B)
ϕ	phi	Standard normal density (§D.3)
Φ	Phi	Standard normal distribution function (§D.2.2)
ψ	psi	Complement of credibility in the loss-ratio method's indication (§A.3)
ω	omega	Weight of a quote in the hold curve's likelihood (§D.3)

Latin Symbols with More Than One Meaning

The letters E, C, L, and T also name, in Table G.3, the four further markets of Section 6: E the market of smaller insurers, C the market of one tenth of the accounts, L the market with a second insurer that learns, and T the market with the severity trend (Appendix G).

- A* An account's capped losses on the file it is rated on (§A.3); a policy's capped reported losses, in the GLM insurer's fit (Appendix B); with subscript f or s , a fitted constant of an autocovariance (§D.3)
- a* Reported losses, in the loss-ratio method's indication (Equation A.2); the loss adjustment expense's share (Equation D.5)
- B* With subscript c , a class's expected capped loss of a year per unit of size (Appendix B); with subscript f or s , a fitted coefficient of an autocovariance (§D.3)
- c* An account's class (§A.3); as a subscript, a case reserve (σ_c); c_x , the return on the premium held (§D.1.2); c_f and c_s , covariances of the observations (§D.3)
- d* Base of the weights of the experience modification (§A.3); with subscript it , the drifting part of a trait (§A.6)
- E* A class's expected capped losses, E_{rep} their expected reported part (§A.3); expected reported losses, in the GLM insurer's fit (Appendix B); $\mathbb{E}$, expectation (§D.1.1)
- f* The indicated change of level (Equation A.2); the GLM insurer's change of level or a fitted factor (Appendix B)
- g* An insurer's growth (Equation A.3); $\hat{g}$, its mean yearly log growth of surplus over the credit's window, the return on funds included (§D.1.3)
- j* The age of a year, in the experience modification (§A.3); a quote, in the hold curve's likelihood (§D.3)
- k* The buyer's cost of shopping, k' its cost of switching (§A.4); years ahead (§D.1.3); a lag (§D.3)
- m* With subscript 1 or s , a margin per unit of net premium (§D.1.3); m_0 , the mean log claim on the bureau's curve (§D.3)
- n* Premium net of expenses (§D.1.2); reported claims, in the loss-ratio method's indication (Equation A.2) and in the estimator's observations, n_s^o the open claims among them (§D.3)
- s* The scale of the buyer's response (§A.4); as a subscript, later renewals (§D.1.3), the severity trait, or an account-year (§D.3); s_1^2 and s_2^2 , variances in the schedule modification (Appendix B)
- t* A year (§D.1.3, §A.6); the t statistic of a contrast (Section 3)
- u* The occurrence limit (Equation D.5); with subscript it , the frequency trait (§A.6)
- v* $\hat{v}$, the estimated variance of an account's loss, $\hat{v}_0$ before the rate level (§D.1.2, §D.2.1); v_0 , the variance of a log claim on the bureau's curve (§D.3)
- w* With subscript j , the weight of a year in the experience modification (Appendix B); a file's length in years (§D.3)
- x* $\bar{x}$, the plan's expected loss (Equation A.1); $\hat{x}$, the estimated expected loss, $\hat{x}_0$ before the rate level (§D.1.2, §D.2.1)
- Y* Claim size at the tracked scale (Equation D.5); $\bar{y}$, a mean log claim size (§D.3)
- Z* Classical credibility (Equation A.2, Appendix B); Z_{exp} , the experience modification's credibility (§A.3)

1 Introduction

An insurer that competes for an account decides two things at once: at what price to quote, and whether it wants the account at the price the market will bear. For neither does it know what it needs. The first requires the account's loss distribution, which no one observes; the second requires the chance of winning at each price, which depends on what competitors quote and on how the buyer chooses. The industry's answer is the rating plan: a class loss cost, an experience modification, a schedule modification, and a provision for expenses and profit, with the insurer's rate level moved once a year by what its book has earned against what it was priced to earn.¹ The plan states a price. It does not state what the price is for.

This paper gives underwriting an objective and asks what an insurer needs in order to pursue it. The objective is the long-run growth of surplus. The Kelly criterion states it: an insurer that maximizes the expected logarithm of its surplus at the end of a horizon maximizes the expected mean geometric rate at which the surplus grows, and in repeated independent bets on known odds no other rule grows wealth faster in the long run.² Applied to one account, the criterion gives a price, the *Kelly decision*, from four things: the account's expected loss; the variability of that loss, weighed against the insurer's surplus; the value of keeping the account in the years after this one; and the chance that the quote wins. A rating plan supplies a version of the first and none of the other three.

An insurer can learn all four from what it already holds. Its claim files record every claim of every account it has written; the loss runs of the accounts it is asked to quote record their claims elsewhere; its own retention of accounts records which stayed at which price; and, where an account went to market, the broker reports by how much a lost quote lost. From these the insurer tracks each account's claim frequency and claim size as the files lengthen and fits the chance that a quote wins the account at a given price. Learning in this sense is the source of what the objective needs. The insurer can put what it learns to work in two ways. It can price through its rating plan, with the learned estimate of each account's loss in place of the plan's: the *plan route* (in the registrations of Appendix F, learning at the plan's pricing). Or it can price by the Kelly decision, as a deviation from the plan route's price: the *Kelly route* (learning at the Kelly decision). The two routes price the same learned estimate and differ in the rule that turns it into a quote.

The two are measured in a stylized market for middle-market general liability (GL). Three conventional insurers compete in it, each named by how it sets its *rate level*, the multiplier it applies to the bureau's loss costs: a *Bureau insurer*, which prices on the bureau's plan and refiles its rate level when its experience departs from the plan; a *Loss-ratio insurer*, which moves its rate level each year by the loss-ratio method; and a *GLM insurer*, which fits a generalized linear model (GLM) to its own book and adds its relativities by class and size to the plan. The market has four places, which the paper calls *seats*; the fourth seat is filled in turn by the insurers under study, the three conventional insurers held as they are. The market has 4,000 accounts and runs for 28 record years. It is run 200 times, each run from one draw of its random streams, a *seed*, and every filling

¹Geoff Werner and Claudine Modlin, *Basic Ratemaking*, 5th ed. ([Arlington, VA]: Casualty Actuarial Society, 2016), chaps. 7–9, 12, and 15; Ginda Kaplan Fisher, Lawrence McTaggart, Jill Petker, and Rebecca Pettingell, *Individual Risk Rating* (Casualty Actuarial Society study note, 2019).

²John L. Kelly Jr., "A New Interpretation of Information Rate," *Bell System Technical Journal* 35, no. 4 (1956): 917–926; Leo Breiman, "Optimal Gambling Systems for Favorable Games," in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1, ed. Jerzy Neyman (Berkeley: University of California Press, 1961), 65–78.

of the fourth seat runs on the same 200 seeds, so that a comparison between two fillings is a comparison of the same market under two insurers; the body of these runs is *the record*. Growth is measured on the surplus at the close of the record years, the *terminal surplus*, against the surplus it started with, its *endowment*, grown at the return on funds, the yield that surplus and reserves earn, so that a surplus that earns that return and nothing else has not grown, and stated as a rate of growth a year or as a multiple of the endowment so grown; a difference between two fillings is stated as a percentage of terminal surplus, by how many percent the one's exceeds the other's; and the record *resolves* a difference, tells it from zero, where its *t* statistic, the mean over its standard error across seeds, exceeds 2 in magnitude. The market is stylized, and no observed market has an insurer that prices by the Kelly decision; the inputs follow industry practice where practice is documented, and each is listed with its grounding (Appendix A).

Three findings are on the record. First, learning is what grows surplus beyond a conventional insurer's. A fourth conventional insurer, a second of the Loss-ratio type, grows at 1.93 percent a year, a multiple of 1.71 over the record years. The plan route grows at 3.39 percent a year, a multiple of 2.55: learning each account's loss distribution and pricing it through the plan adds 49.2 percent to terminal surplus, and it adds to growth at four other settings of two dimensions of the market, a refiling threshold and the noise of an underwriter's read of an account (Section 4). Second, the two routes end the record years with equivalent growth. The Kelly route grows at 3.43 percent a year, a multiple of 2.57, and its difference from the plan route, 1.1 percent of terminal surplus, lies within an *equivalence margin*, set with the difference already on the record, of 2.8 percent of terminal surplus either side of zero, 0.10 percentage points of growth a year: two growths are read as equivalent where the 90 percent interval of their difference lies inside it (Section 5). Third, the two routes differ in their books and their paths. The Kelly route quotes renewals 12.6 percent above the plan route's price and new business 8.4 percent below it; it holds 49 percent of the market's accounts at the close of the record years where the plan route holds 67 percent; accounts change insurer around it at 26 percent a year against 16; and it trails the plan route through record year 23 by a difference the record resolves, and converges with it only over the final years, its growth steadier from year to year and its terminal surplus more dispersed across seeds. A management that reads growth at year ten or year twenty sees a route that trails.

Whether the two routes stay together is a property of the market. Where the *filed band* of schedule rating, inside which an insurer may move a price from the plan's, is 40 percent either way in place of 25, the Kelly route trails the plan route by 14.2 percent of terminal surplus. A buyer more responsive to price moves it behind the plan route, and a buyer less responsive moves it ahead; where the insurers are smaller it leads. The record resolves each of these three differences, and the interval of each crosses the equivalence margin. In a market with a tenth of the accounts and under a rising severity trend that no insurer in the market projects, it leads beyond the margin; beside a second insurer that learns and prices as the plan route does, the two routes stay within the margin of each other (Section 6 gives each difference). In the construction of record, which Sections 2 and 4 state, the value of keeping an account is anchored at the price the decision deviates from, the plan route's price. How that value enters the price decides where the Kelly route stands: the second finding is the Kelly route so constructed; with the value rising with the price quoted, it falls short of the plan route's growth (Section 4).

The limits are stated in full in §7.2; five of them are these. The market is one line of business, with no loss trend except in the one market run with one, no underwriting cycle, no reinsurance, and conventional insurers that respond to a learner only through their own rules; the two routes are run at the record's values of the filed band and the buyer's responsiveness and at one other value of the band and two of the responsiveness; the criterion's weight on risk is small at this market's

surplus, so the record tests the criterion's rule for pricing under uncertainty and not its weighting of risk; the horizon is known to the insurer, and growth is read at it; and every difference is measured within one model over 200 seeds, so that its t statistic states the precision of the simulation, not of an observed market.

The paper draws on three literatures. The first is the Kelly criterion of Kelly and Breiman, whose second-order form weighs an account's expected margin against the variance of its loss, with the weight fixed by the criterion and not by the insurer.³ The second is the theory of insurance pricing under asymmetric learning, in which the insurer that holds an account learns about it what competitors do not, and, where the insurer cannot bind itself to its later prices, the first price is low, the "lowballing" of that literature, recovered through the margins of later renewals on the accounts that stay.⁴ In the market of this paper the holder reads its own claim file where a competitor reads a broker's loss run, every contract is annual, and the insurers that value retention arrive at new-business prices below their renewal prices without a model of that equilibrium (§5.2). The empirical counterpart is the measured difference between new and renewal business: in each of the 25 books of business that Wu and Lin examine the loss ratio of new business is above that of renewals.⁵ The third is the study of pricing algorithms in competition, which asks whether algorithms that learn, or that respond to rivals' prices, raise prices.⁶ That question is about the level of prices among sellers whose algorithms meet one another; the question here is what an objective and the learning are worth to one insurer among insurers that do not learn as it does.

Section 2 states the approach: the objective, the four things the decision needs and how each enters the price, and what learning supplies. Section 3 describes the market. Section 4 reports what learning is worth and what the form in which the decision values keeping an account is worth, Section 5 the two routes and how they differ, and Section 6 the markets in which they part. Section 7 states what follows for an insurer, the limits, and the conclusion. The appendices carry the market's assumptions and its construction in full, the rating plans, the capital standard, the Kelly decision and the estimators with their equations, the Kelly route's computation step by step, the registry of predictions and outcomes, the quote data, and the further markets.

³Kelly, "Interpretation of Information Rate"; Breiman, "Optimal Gambling Systems"; Edward O. Thorp, "The Kelly Criterion in Blackjack Sports Betting, and the Stock Market," in *Handbook of Asset and Liability Management*, vol. 1, ed. Stavros A. Zenios and William T. Ziemba (Amsterdam: North-Holland, 2006), 385–428.

⁴Howard Kunreuther and Mark Pauly, "Market Equilibrium with Private Knowledge: An Insurance Example," *Journal of Public Economics* 26, no. 3 (1985): 269–288; Tore Nilssen, "Consumer Lock-in with Asymmetric Information," *International Journal of Industrial Organization* 18, no. 4 (2000): 641–666; Thomas de Garidel-Thoron, "Welfare-Improving Asymmetric Information in Dynamic Insurance Markets," *Journal of Political Economy* 113, no. 1 (2005): 121–150; Ruo Jia and Zenan Wu, "Insurer Commitment and Dynamic Pricing Pattern," *Geneva Risk and Insurance Review* 44, no. 1 (2019): 87–135; and, for evidence consistent with the theory, Stephen P. D'Arcy and Neil A. Doherty, "Adverse Selection, Private Information, and Lowballing in Insurance Markets," *Journal of Business* 63, no. 2 (1990): 145–164.

⁵Cheng-Sheng Peter Wu and Hua Lin, "Large Scale Analysis of Persistency and Renewal Discounts for Property and Casualty Insurance," *Casualty Actuarial Society E-Forum* (Winter 2009): 398; Stephen P. D'Arcy and Neil A. Doherty, "The Aging Phenomenon and Insurance Prices," *Proceedings of the Casualty Actuarial Society* 76 (1989): 24–44.

⁶Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello, "Artificial Intelligence, Algorithmic Pricing, and Collusion," *American Economic Review* 110, no. 10 (2020): 3267–3297; Stephanie Assad, Robert Clark, Daniel Ershov, and Lei Xu, "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market," *Journal of Political Economy* 132, no. 3 (2024): 723–771; Zach Y. Brown and Alexander MacKay, "Competition in Pricing Algorithms," *American Economic Journal: Microeconomics* 15, no. 2 (2023): 109–156.

2 Growth as the Objective, and Learning as Its Source

2.1 The Objective, and the Decision It Gives

The Kelly criterion takes the long-run growth of surplus as the objective: the insurer chooses its prices so as to maximize the expected logarithm of its surplus at the horizon, which is the same as maximizing the expected mean geometric rate at which the surplus grows from year to year. An insurer's surplus compounds, what it writes this year changing the surplus from which it writes next year, and the objective counts that. The criterion's property, in repeated independent bets on known odds, is that no other rule grows wealth faster in the long run; an insurer does not know the odds, and the record measures what the objective is worth to an insurer that has to estimate them. Appendix D states the approach of this section with its equations and the estimators in full.

For one account the objective reduces to a rule. Writing an account at a premium changes the surplus by the premium net of expenses less the account's loss. For an account small against the surplus, the expected change in the logarithm of surplus is, to a close approximation, the expected margin as a share of surplus, less half the variance of the loss divided by the square of the surplus: a mean–variance rule, with the weight on the variance set by the objective and not by the insurer. Two further terms enter. The premium is held until the claims are paid and earns the return on funds in the meantime. And in a market where the account can be kept, an account written today is likely to be held in the years to come, so the margins of those years count, discounted at the insurer's own rate of growth of surplus, the return on funds included, which the objective gives, or at the return on funds; the record runs both discounts (Section 4). With P the premium quoted, n the premium net of expenses, $\bar{n}$ the same at the undeviated price, the price the decision deviates from, $\hat{x}$ and $\hat{v}$ the account's estimated expected loss and its variance, S the surplus, c_x the return the premium earns before the loss is paid, netted from the loss, and V the value of keeping the account per unit of $\bar{n}$, the expected increment to the logarithm of surplus in the paper's construction is

$$G(P) = \frac{n - \hat{x}(1 - c_x) + \bar{n}V}{S} - \frac{\hat{v}}{2S^2}, \quad (1)$$

and the Kelly decision is the premium that maximizes the chance of winning the account at that premium times the increment it would earn. The decision answers both questions an insurer faces at a quote: the price, and, in the sign of the increment, whether it wants the account at that price. It needs four things of an account. The expected loss $\hat{x}$ is what a rating plan also estimates. The variance $\hat{v}$ enters through the criterion's weight on risk, the variance term of Equation 1: of two accounts with one expected loss, the one whose loss is more variable is charged no less, and an insurer with less surplus charges no less for the same variability. The value of keeping the account, V , is the margin of each later renewal times the chance of reaching it, discounted, summed over the years to the horizon; the paper calls it the *renewal credit*: a value in the insurer's objective, not a discount to the renewing buyer. And the chance of winning at each premium, a renewal kept or new business bound, is the insurer's own *hold curve*, fitted to the quotes it won and lost.

How the value of keeping an account enters the price is measured in Section 4, and the paper's construction anchors it. What a markup today buys is today's margin and the chance of winning the account at all; to first order it does not buy a larger margin next year, because next year's quote is next year's rule applied to next year's estimate, and the buyer who shops compares the holder's renewal quote with the competition's, not with the holder's own raise, an argument from the market's rules that the record does not measure (§D.1.5). The chance of winning multiplies the whole increment, so what a markup does to that chance is already counted there; the credit itself

therefore carries no markup and enters at the undeviated price, $\bar{n}$ in Equation 1, the plan route's price: the *anchored credit*. The other form, the *quote-anchored credit*, lets the credit rise with the premium quoted, n in place of $\bar{n}$, and so rewards a markup with a value the markup does not buy; the record measures what each form is worth (Section 4, Appendix F). The decision also sets a floor, the premium at which the increment is zero, below which the insurer does not write; the floor is written with the quote-anchored credit, and in the record's market it binds no quote of the Kelly route (Table G.1).

2.2 Learning from What an Insurer Holds

A rating plan reads an account's history through one number, the experience modification on its capped losses. The learning of this paper has two parts, each read from data the insurer holds: each account's loss distribution, estimated from the whole file, and the chance that a quote wins the account, fitted from the broker's figures. First, the account's claim frequency relative to its class and the scale of its claim sizes are tracked from the files the insurer reads, its own claim file on an account it holds and the broker's loss run on an account it is asked to quote, by credibility weighting, with the weight on each year of the file set by how much the year's claims carry and by how far back it lies, and with the persistence of an account's departure from its class estimated from the insurer's own files rather than assumed. The expected loss and the variance follow from the tracked frequency and scale, and the variance is the quantity a plan does not hold.

Second, the insurer fits its own hold curve, the chance that a quote at a given deviation from the price its rating plan gives the learned estimate wins the account, from the only information it has on competitors' prices: the broker's figure, by how much a quote lost where the account went to market, and which of its quotes held their accounts. The fit allows for the quotes that bound, whose competing prices the insurer never sees, and the decision averages the chance of winning over the uncertainty of the fit. Beside the two parts, the renewal credit is read from the insurer's own book over a window of years: the retention and the margin at a first renewal and at later renewals, and the insurer's own growth of surplus, the return on funds included. The four things of §2.1 come from these parts and that book: the expected loss and its variance from the first, the chance of winning from the second, and the value of keeping the account from the insurer's own book.

On the plan route the learned expected loss takes the place of the class loss cost with its experience modification, the underwriter's schedule modification multiplying it as it does the plan's estimate, and the plan's own rules turn it into a quote: the rate level, moved by the loss-ratio method with competitors' winning prices as the complement of credibility, the figure that receives the weight not given to the insurer's own experience, and the schedule credit on new business (Section 3). On the Kelly route the same estimate is priced by the Kelly decision as a deviation from the plan route's price, chosen on the hold curve, inside the filed band of 25 percent either way that also bounds the schedule modification and the schedule credit, so that the Kelly route contains the plan route as the choice of a zero deviation at every account. The plan route uses the first part of the learning alone, the learned expected loss; the Kelly route uses both, with the variance entering through the decision; both price at the plan's rate level (Appendix E states the Kelly route's computation step by step).

3 A Stylized Market

The learning and the decision are measured in a market of 4,000 accounts of middle-market general liability, written by four insurers without reinsurance; Figure 1 maps one year of it, and Appendix A states every assumption with its value, its grounding, and the variations of the market the paper runs (Section 6; Appendix F). Each account has a size and two traits, one of claim frequency and one of claim size, each with a permanent part and a part that drifts from year to year; claims occur, are reported with a lag, and are paid at settlement, with an open claim carried at a case reserve with error. The claim sizes follow the distribution increased-limits ratemaking uses, a mixed exponential, scaled so that the book's mean claim at the occurrence limit is the industry's published general liability severity, with limits per occurrence and in the aggregate and loss adjustment expense in proportion to the indemnity paid (Appendix A).

Three sources of information exist, and they are not equally shared. A bureau publishes, for each of five classes, the market's pooled loss costs and claim frequency of past policy years. An insurer asked to quote receives the broker's loss run, the account's claims of the past five policy years whoever wrote them, and its underwriter forms a read of the account of its own, which sets the schedule modification (Appendix B). The insurer that holds an account reads its own claim file for every year it has held it, current to the quarter and with a smaller reserve error than the broker's run carries. The holder therefore knows more about its accounts than a competitor invited to quote does, and the longer it has held an account the more; nothing else in the market favors the holder except the buyer's costs of shopping and switching.

Three conventional insurers compete in every market the paper runs but one, in which a second insurer that learns takes the GLM insurer's seat (Section 6), and they are the benchmark. Each prices by the bureau's plan, the class loss cost times the experience and schedule modifications times its own rate level, with a provision for expenses and profit taken from the industry's published expense ratios, and each sets its level in its own way: the Bureau insurer refiles its rate level when the indication on its own experience departs from the plan by a stated threshold; the Loss-ratio insurer moves its level every year by the loss-ratio method, by how its book's losses ran against what it was priced to earn, with the level of competitors' winning prices as the complement of credibility; and the GLM insurer moves its level by its book's losses against the losses its own relativities expect, credibility-weighted, and adds relativities by class and size band from a generalized linear model fitted to that book. None files a load on new business; each applies a schedule credit on new business in the larger size bands from the winning prices it has seen, inside the filed band of 25 percent either way, as Wu and Lin's evidence suggests insurers do, pricing new business below renewals through the plan's own flexible components.⁷ The three do not respond to a competitor that learns beyond these rules. Appendices A and B state the plans and the rules in full.

⁷Wu and Lin, "Large Scale Analysis," 399, 404–405.




Figure 1: One year of the market: each quarter the holder and challengers quote, the buyer chooses, an insurer that lost at market is told the broker's figure, and the policy is booked; at the year's close the rate levels, the GLM insurer's factors, the schedule credits, and the renewal credit are re-estimated and capital is tested (§A.5). Solid arrows are flows within a quarter, dashed arrows what returns to an insurer. The fourth seat holds each filling of Section 4 in turn. Source: authors' construction.

An account renews once a year; its holder quotes, and the broker invites one or more *challengers*, competitors asked to quote, more for a larger account. The chance that the buyer goes to market rises with the gap between the holder's renewal quote and the lowest competing quote, against a cost of shopping, and the chance that it then switches rises with the same gap against a cost of switching; with probability 0.10 an account is remarketed for reasons unrelated to price. The two costs were calibrated, on seeds outside the record, to a target of 14 percent of the accounts of a market of four conventional insurers changing insurer a year, half the share of middle-market companies that a 2007 survey of the segment found changing an insurer over two years, counting no company changing twice; the calibration gives a premium retention between two retentions of commercial business published in 2026 (Appendix A). After placement each insurer that quoted on an account that went to market and lost is told the difference between its quote and the bound premium, rounded, *the broker's figure*; the insurer that bound is told nothing. Every insurer starts the record years with the same endowment, set so that its opening book stands at the industry's aggregate capital ratio under the risk-based-capital standard, a ratio of written premium to surplus

of about a third against the industry's 0.773, and surplus and reserves earn the return on funds, the industry's published investment yield. Capital acts at three thresholds (Appendix C): a broker withholds new-business invitations from an insurer whose capital ratio falls into the band at which regulators watch it, *broker flight*; an insurer whose written premium exceeds three times its surplus stops quoting new business; and an insurer below the level at which regulators take control fails and is replaced. No insurer fails in any seed of any market the record runs.

After a build-up of the books and files (§A.5), the market runs 10 practice years before the 28 record years begin, so that every insurer opens the record with estimators that have run and a book its own pricing has sorted. An insurer's *growth* is the logarithm of its terminal surplus over its endowment, net of the return on funds, so that a surplus that earned that return and nothing else has a growth of zero; the paper states it as a rate of growth a year and as the multiple of the endowment grown at the return on funds. A difference between two fillings of the fourth seat is taken within each seed and averaged, with its t statistic, the mean over its standard error across seeds, stating how precisely the simulation resolves the difference; a difference whose t statistic exceeds 2 in magnitude is one the record resolves. The paper states a difference in growth as a percentage of terminal surplus, computed from the unrounded means; the tables carry each difference's t statistic, and a difference the text states without qualification is one the record resolves. Every difference is a total effect: what changes in one seat moves the other insurers' books, rate levels, and capital, and the difference carries their response.

4 What Learning Is Worth

Learning is the largest effect in the record's market. The fourth seat is filled in seven ways on the same 200 seeds, the three conventional insurers the same in each (Table 1). With the plan's estimate of each account's loss distribution and the plan's pricing as the Loss-ratio insurer applies it, the fourth seat is a second Loss-ratio insurer, and the market has no insurer that learns. With the learned estimate and the plan's pricing it is the plan route. The other fillings are the Kelly route, the learned estimate priced by the Kelly decision as a deviation from the plan route's price, in the construction defined below (the first of the Kelly route's rows in Table 1), and its variations, which vary the construction in three ways: the form of the credit; its discount, the *growth discount* at the insurer's own growth of surplus, the return on funds included, or the *funds discount* at the return on funds (§D.1.3); and a *pessimism variant* that shifts the fitted hold curve down by one standard deviation of its uncertainty, so that the decision quotes against a lower chance of winning. The *construction of record* is the anchored credit on the funds discount with the pessimism variant.

Table 1: The fillings of the fourth seat, means over 200 seeds. Growth is the logarithm of the fourth seat's terminal surplus over its endowment, net of the return on funds (Section 3); the rate a year and the multiple of the endowment grown at the return on funds restate the mean logarithm, as a yearly rate and as a factor. The share is the fourth seat's share of the accounts at the close, the turnover the share of the market's accounts changing insurer a year. The construction of record is the table's second row, the first of the Kelly route's. Source: authors' calculations.

	Growth, logarithm	Rate a year	Growth multiple	Share of accounts	Changing insurer a year
The fourth seat					
<i>The two routes, the Kelly route with the anchored credit</i>					
The plan route	+0.935	3.39	2.55	0.672	0.163
The Kelly route, with pessimism, funds discount	+0.945	3.43	2.57	0.494	0.263
The Kelly route, funds discount	+0.945	3.43	2.57	0.488	0.263
The Kelly route, with pessimism, growth discount	+0.941	3.42	2.56	0.461	0.270
The Kelly route, growth discount	+0.939	3.41	2.56	0.452	0.269
<i>The Kelly route with the quote-anchored credit</i>					
The Kelly route, funds discount	+0.836	3.03	2.31	0.308	0.258
<i>No learning</i>					
Second Loss-ratio insurer: plan's estimate, plan's pricing	+0.534	1.93	1.71	0.223	0.141

A second Loss-ratio insurer in the fourth seat grows at 1.93 percent a year, a multiple of 1.71 over the record years. The fourth seat so filled and the Loss-ratio insurer beside it, one rule in two seats, end 2.9 percent of terminal surplus apart ($t = +1.60$), a difference the record does not resolve. The plan route, the learned estimate priced as the Loss-ratio insurer prices, grows at 3.39 percent a year, a multiple of 2.55: learning alone adds 49.2 percent to terminal surplus. The three conventional insurers grow less beside the plan route than beside a second Loss-ratio insurer (Figure 2), and in that market the plan route ends with 71.8, 84.4, and 68.0 percent more terminal surplus than the Bureau insurer, the Loss-ratio insurer, and the GLM insurer beside it (Appendix F, Table F.1), and it holds 67 percent of the market's accounts at the close, a book at which the capital standard begins to act: brokers withhold 964 new-business invitations over the 28 record years and 200 seeds of a market of 4,000 accounts, where in the market with a second Loss-ratio insurer they withhold none. One source of the lead shows in the new business it writes. The new business a second Loss-ratio insurer wins has a true expected loss of 1.012 times its own estimate, because an insurer wins where its estimate is low;⁸ the new business the plan route wins has a true expected loss of 0.927 times its estimate, and the insurer is paid on its estimate. With the learned estimate the insurer does not win mainly where its estimate is low. Figure 2 gives every insurer's growth by seed in the two markets.

⁸Edward C. Capen, Robert V. Clapp, and William M. Campbell, "Competitive Bidding in High-Risk Situations," *Journal of Petroleum Technology* 23, no. 6 (1971): 641–653; Paul R. Milgrom and Robert J. Weber, "A Theory of Auctions and Competitive Bidding," *Econometrica* 50, no. 5 (1982): 1089–1122.




Figure 2: Growth by insurer and seed with a second Loss-ratio insurer in the fourth seat (upper panel) and with the plan route there (lower panel), over 200 seeds. Each insurer's row of points gives its growth, the logarithm of Equation A.3, on one horizontal scale, a vertical bar marking the mean, stated at right. In all figures the Bureau insurer is bright blue, the Loss-ratio insurer amber, the GLM insurer green, and the fourth seat deep blue. Source: authors' calculations.

The comparison was also run at four other settings of two dimensions of the market, the Bureau insurer's refiling threshold, 0.10 in the record's market and 0.05 and 0.15 at the other settings, and the noise of an underwriter's read of an account, the read noise, 0.7 there and 0.35 and 1.05 at the other settings. Learning adds to growth at every setting, by a difference the record resolves at each (Appendix F, Table F.1).

The form in which the value of keeping an account enters the decision moves terminal surplus by more than the decision itself moves it (+1.1 percent, the Kelly route less the plan route). With the anchored credit, the construction of record, the Kelly route reaches the plan route (Section 5); with the quote-anchored credit the same decision on the same estimate falls 9.4 percent of terminal surplus short of the plan route, and the anchoring is worth 11.6 percent of terminal surplus, each on the funds discount. The paper *registered* its predictions: each was written down and filed before the runs that tested it, or, for the one registration that analyzes runs already made, before any of its metrics was computed; Appendix F reports each registration's outcomes, every failed prediction among them named, in the order in which each was written and run. The anchoring was designed after the first registration's outcomes were on the record, and registered before any run of it on the record's seeds.

5 Two Routes to Equivalent Growth

The plan route and the Kelly route price the same learned estimate, and they end the record years with equivalent growth; they differ in the prices they quote, the book each builds, the turnover around them, and the path by which each arrives. Table 2 sets the two side by side, and Figure 3 draws the difference between them at the close and year by year.

5.1 Equivalent in Growth

The Kelly route grows at 3.43 percent a year, a multiple of 2.57, and the plan route at 3.39 percent, a multiple of 2.55, the highest growths of any filling of the fourth seat (Table 1). Their difference at the close of the record years, the Kelly route less the plan route paired by seed, is 1.1 percent of terminal surplus, with 113 of 200 seeds above zero and a t statistic of 1.97, below the 2 at which the record resolves a difference (Figure 3, left). A difference the record does not resolve is not a finding of equality, so the paper tests equivalence against a margin registered with the difference already on the record (Appendix F): 0.028 in the logarithm over the 28 record years, which is 2.8 percent of terminal surplus and 0.10 percentage points of the rate of growth a year, chosen in the objective's own unit for the reasons Appendix F states. Two growths are equivalent at that margin where the 90 percent interval of their difference lies between the margin below zero and the margin above it. The two routes' growth is equivalent at the margin on each discount of the credit and with or without the pessimism variant; in the construction of record the interval runs from 0.2 to 2.0 percent of terminal surplus (Appendix F, Table F.3). The two tests read one difference two ways. The rule that resolves a difference asks whether it can be told from zero, and at $t = 1.97$ it cannot; the equivalence test asks whether it can be told from the margin, and it can, the interval lying inside the margin. The interval also lies above zero, so at the equivalence test's own confidence the Kelly route's lead is positive and smaller than the margin. The second registration had predicted that lead at $t > 2$; at $t = 1.97$ the prediction failed (Appendix F). It does not take the Kelly decision to realize the gain from learning: the plan route computes no Kelly decision at quote time, and its growth lies within the equivalence margin of the Kelly route's.

The two routes share the estimate of each account and the filed band, inside which alone the Kelly route can move a price and which the plan route also observes. The record measures their equivalence in the record's market and does not measure what in the two constructions produces it (§7.2); Section 6 shows the markets in which their growth parts.

5.2 Where the Two Routes Differ

The routes are told apart in the quote data, in the market around them, and in time (Tables 1 and 2). The Kelly route prices renewals 12.6 percent above the plan route's price on average, with more than a quarter of its renewal quotes within 0.25 percent of the band's upper end, and it prices new business 8.4 percent below that price. The pattern, a low first price recovered on the accounts that stay, has the shape of the lowballing the theory of pricing under asymmetric learning predicts of an insurer that knows its own accounts better than its competitors do;⁹ the Kelly route arrives at it without a model of that equilibrium, from the value of keeping an account and the chance of winning it. Against the true expected loss of the business each writes, the Kelly route's premium is 1.42 times that loss on new business and 1.54 on renewals, and the plan route's 1.57 and 1.48

⁹D'Arcy and Doherty, "Adverse Selection"; Jia and Wu, "Insurer Commitment," 89.

(Table 2); each multiple is taken on the route's own book, where the deviations above compare the two routes' prices on the same account. Net of expenses of 0.30 of premium, the Kelly route's new business is written below its true expected loss, by a difference the record resolves, and above that loss discounted at the return on funds over the mean time to payment (§D.1.2); the plan route's is written above both. The Kelly route holds 49 percent of the market's accounts at the close of the record years against the plan route's 67 percent, and the accounts around it change insurer at 26 percent a year against 16: a smaller book at a higher renewal margin, in a market that turns over more.

Table 2: The plan route and the Kelly route side by side, means over 200 seeds; the Kelly route is the construction of record (Table 1, the second row). Table 1 gives the two routes' growth, shares, and turnover; the prices are, for the Kelly route, its mean deviation from the plan route's price, a factor less one, and, for each route, its premium over the true expected loss of the business it writes, the ratio of the two sums over the record years and seeds; the path's rows read the difference between the routes at the close of each record year, paired by seed, with the volatility the standard deviation over the record years of the yearly increment of growth and the dispersion the standard deviation across seeds of growth at the close, each with the t of the difference between the routes (for the dispersion, the Pitman–Morgan t of two paired standard deviations). Source: authors' calculations.

	The plan route	The Kelly route
<i>The prices</i>		
Renewals, percent above the plan route's price	—	12.6
New business, percent below the plan route's price	—	8.4
Premium written over the true expected loss, new business	1.57	1.42
Premium written over the true expected loss, renewals	1.48	1.54
<i>The path</i>		
Record years through which the Kelly route trails, by a difference the record resolves	—	23
Deepest shortfall of the Kelly route, percent of the surplus of that year, in year 10	—	−5.0
Within-path volatility of growth, logarithm	0.0204	0.0179
Its difference, the Kelly route less the plan route: t	—	−14.42
Dispersion of growth across seeds at the close, logarithm	0.0956	0.1218
The ratio of the two: Pitman–Morgan t	—	+5.40

In time the two routes are told apart through most of the record. Read at the close of each earlier record year, the Kelly route trails the plan route by a difference the record resolves in every year through year 23, deepest in year 10 (Table 2), and the 90 percent interval of the difference lies inside the equivalence margin from year 22 through the close (Figure 3, right); in years 22 and 23 both readings hold, the difference resolved against zero and inside the margin. Seed by seed, the longest run of consecutive years in which the Kelly route trails is 23.5 years at the median seed. Along any one path the Kelly route's growth is the steadier of the two, its year-to-year volatility the lower, the one a management sees along its own path; across seeds its terminal surplus is the more dispersed, the standard deviation of its growth at the close 1.274 times the plan route's,¹⁰ the spread of its possible outcomes at the close.

¹⁰The t of two paired standard deviations, Table 2: E. J. G. Pitman, "A Note on Normal Correlation," *Biometrika* 31, no. 1/2 (1939): 9–12; W. A. Morgan, "A Test for the Significance of the Difference between the Two Variances in a Sample from a Normal Bivariate Population," *Biometrika* 31, no. 1/2 (1939): 13–19.




Figure 3: At left, the Kelly route less the plan route, the fourth seat's growth at the close of the record years paired by seed, in percent of terminal surplus (horizontal axis), one jittered point a seed; the solid bar is the mean and the short-dashed line zero. At right, the same difference read at the close of each record year (horizontal axis), in percent of the surplus of that year, the mean over seeds as the line and its 90 percent interval as the shaded area; the dashed lines are the equivalence margin either side of zero. The Kelly route is the construction of record. Source: authors' calculations.

6 When the Routes Part

The equivalence of the two routes is a property of the market the record runs. Two dimensions of the market act on the Kelly decision: the filed band, which bounds how far the decision can move a price, as it bounds the conventional insurers' schedule credits; and the buyer's responsiveness to price (the scale of its response in Table 3, 0.05 at the record's, a smaller scale a more responsive buyer), which sets how much shopping and switching a given price gap produces and so the markup the decision chooses. The two routes were run at one other value of the band and two of the responsiveness, and in four further markets run beside the record: a market with a second insurer that learns and prices as the plan route does, in the GLM insurer's seat; smaller insurers, each endowment 0.467 of the record's, at which an insurer's opening book stands at the capital ratio at which broker flight begins; a market of one tenth of the accounts; and a market with a severity trend of 6.8 percent a year that no insurer projects, the yearly growth of the severity of other liability occurrence claims from 2015 through 2024,¹¹ applied to the mean claim limited per occurrence. Table 3 reports the Kelly route less the plan route in each, in percent of terminal surplus.

¹¹Jim Lynch and William Nibbelin, *Increasing Inflation on Liability Insurance – Impact as of Year-End 2024* (Insurance Information Institute and Casualty Actuarial Society, October 2025), 18; for other liability occurrence, net ultimate loss and defense and cost containment expense per claim, accident years 2015 through 2024.

Table 3: The Kelly route less the plan route at the close of the record years, paired by seed over 200 seeds, in percent of terminal surplus with its t statistic, in the record's market and in seven others that move one dimension of it; the Kelly route is the construction of record, the plan route in the same market. Appendix F names the registration each row belongs to, and Appendix G gives the quote data at each setting and learning's contrast in the four further markets. Source: authors' calculations.

The Kelly route less the plan route	Percent of terminal surplus	t
The record's market	+1.1	+1.97
A filed band of 40 percent either way, in place of 25	-14.2	-30.28
A buyer more responsive to price (scale 0.025 in place of 0.05)	-3.4	-5.56
A buyer less responsive to price (scale 0.075)	+3.1	+5.51
A second insurer that learns, in the GLM insurer's seat	+0.3	+0.44
Smaller insurers: each endowment at 0.467 of the record's	+2.3	+3.19
A market of one tenth of the accounts	+12.3	+8.50
A severity trend of 6.8 percent a year that no insurer projects	+7.2	+4.89

Where the band is wide the Kelly route trails. At a band of 40 percent either way the Kelly route ends 14.2 percent of terminal surplus below the plan route, beyond the equivalence margin; both routes grow less at the wider band than at the record's, the plan route by 11.5 and the Kelly route by 24.8 percent of terminal surplus (t of -24.91 and -46.99), and the record does not attribute the Kelly route's loss there between its own quotes and the conventional insurers' wider schedule credits. The buyer's responsiveness moves the routes apart in either direction: a buyer more responsive to price moves the Kelly route 3.4 percent behind the plan route, and a buyer less responsive moves it 3.1 percent ahead, each a difference the record resolves whose interval crosses the margin. The decision's quotes move with the buyer's responsiveness, the renewal deviation and the share of renewal quotes at the band's upper end lower where the buyer is more responsive and higher where it is less (Appendix G).

Where the insurers are smaller, the accounts fewer, or the severity trend rising with no insurer projecting it, the Kelly route leads, and the record does not separate what in each market gives the lead. Beside a second insurer that learns, the two routes stay within the equivalence margin of each other (+0.3 percent). Where the insurers are smaller the Kelly route leads the plan route by 2.3 percent of terminal surplus, a difference the record resolves whose interval crosses the margin; in the market of one tenth of the accounts it leads by 12.3 percent and under the severity trend by 7.2 percent, each beyond the margin. No prediction was registered for the Kelly route's differences in the four further markets (Appendix F). In each of the four further markets the plan route outgrows a second Loss-ratio insurer, as in the record's market: the gain from learning stands in every market in which the two were run, its size moving with the market, and the worth of the decision on top of it changes sign between markets.

7 What Follows for an Insurer, and Conclusion

7.1 The Two Routes

An insurer that learns each account's loss distribution from its own files and prices through its rating plan, the plan route, reaches the highest growth of any filling of the fourth seat in the record's market, shared with the Kelly route within the equivalence margin. Of its pricing architecture the rate level, the schedule credit on new business, and the provision for expenses and profit stand; the learned estimate takes the place of the class loss cost and the experience modification, with the schedule modification multiplying it as before. Its book is the largest of the fillings of Table 1, 67 percent of the market's accounts at the close of the record years.

An insurer that prices the same learned estimate by the Kelly decision, the Kelly route, arrives by a different book and a different path. It quotes renewals above the plan route's price and new business below it, it holds 49 percent of the market's accounts at the close, its renewal margin is the higher, and the market around it turns over more. Its surplus trails the plan route's through record year 23 by a difference the record resolves; the two meet only over the final years. Where the insurers are smaller, the accounts fewer, the severity trend rising with no insurer projecting it, or the buyer less responsive to price, this route leads; where the filed band is wide or the buyer more responsive to price, it trails. All of this holds with the value of keeping an account anchored at the plan route's price; with the value rising with the price quoted, the route falls short of the plan route's growth.

7.2 Limits

The market of the record is one line of business with independent traits, no loss trend (Section 6 reports a market with one), no shock common to all accounts, no reinsurance, no underwriting cycle.¹² The three conventional insurers are built to price as the sources describe a rating plan, a loss-ratio indication, and a generalized linear model; what they do not have is a response to a competitor that learns beyond the schedule credit and the complement of credibility set to competitors' rates, and an insurer in their place might acquire one. The market with a second insurer that learns gives the fourth seat a competitor that learns from the start, not one that responds to it. The buyer responds to price alone, through a rule of shopping and switching calibrated to one target, from a survey of 2007. The broker tells every insurer that lost at market its figure, rounded and without error; a figure that is partial, absent, or shaded is not run. Expenses are 0.30 of premium on new business and renewals alike, so a higher cost of acquiring an account is not measured. Surplus bears no income tax and pays no dividend, and terminal surplus carries each insurer's reserves at its own estimates, unrestated (§A.5). The schedule credits on new business and the Kelly route's deviations answer competitors' prices, which the rule the band's width is taken from, a Florida rule on schedule rating (Appendix A), excludes from what a schedule may rate; for personal lines, a task force of the National Association of Insurance Commissioners considers rate adjustments by a buyer's price elasticity of demand or propensity to shop inconsistent with the standard that rates not be unfairly discriminatory.¹³ The model follows the practice the sources report and does not assess compliance.

¹²On the explanations of underwriting cycles: Mary A. Weiss, "Underwriting Cycles: A Synthesis and Further Directions," *Journal of Insurance Issues* 30, no. 1 (2007): 31–45.

¹³National Association of Insurance Commissioners, Casualty Actuarial and Statistical (C) Task Force, *Price Optimization White Paper* (adopted by the task force, November 19, 2015), 15–16.

The two routes are run at the record's values of the filed band and the buyer's responsiveness and at one other value of the band and two of the responsiveness, and the equivalence of their growth is a statement within a margin set with the difference already on the record, at this record's precision, and at the close of the record years. The decision's variance term, the criterion's weight on risk, is small at this market's surplus; of the record's market and the four run beside it, it is largest in the market of one tenth of the accounts, at 0.6421 percent of the expected loss of an account at the bureau's claim-size curve's own scale, and its removal is run in the record's market alone, where it is shown not to move the Kelly route's growth by more than 0.010 in the logarithm (§D.4). The anchored credit is anchored at a price below the Kelly route's own renewal prices, so it understates the credit's level, and the deviations the decision chooses are higher than a credit at the realized scale would give; the record does not measure by how much (§D.1.5). The record therefore tests the criterion's rule for pricing under uncertainty about the odds, a margin against the chance of winning the account and the discounted value of keeping it, and not the property that separates the criterion from the maximization of expected profit when the stake is large against the surplus. The record runs the criterion in full and no fraction of it, the form the literature uses to give up growth for security, and the criterion's long-run growth does not make it the choice of an insurer that weighs a loss of surplus otherwise than the logarithm does.¹⁴ The horizon is known to the insurer, the value of keeping an account falls toward zero over the final record years, and growth is read at that horizon; the two routes converge in those years, and the record does not attribute the convergence between the decision and the approach of the horizon, nor the two routes' arriving together to anything in the two constructions, nor does it measure a market whose horizon ends earlier, or one under a going-concern horizon. The record does not separate a filling's growth into what the size of its book contributes and what its margin contributes. And every difference is a total effect within a market that responds through its competitors' capital and books, measured over 200 seeds of one model, so that its t statistic is a statement about the simulation's precision at that count and not about an observed market; the registered rule is applied to each prediction without adjustment for the number of predictions tested, and Appendix F reports the count and the reading that allows for it.

7.3 Conclusion

A rating plan tells an insurer what to charge. The Kelly criterion tells it what the charge is for: the growth of its surplus, which, in repeated independent bets on known odds, the criterion maximizes in the long run, and which requires, beyond the expected loss a plan also estimates, three things a plan does not hold, the variability of the loss against surplus, the value of keeping the account, and the chance that the quote wins. An insurer can learn all four from the records of its own operations, and this paper has measured what the learning is worth and the two ways it can be put to work, in one market, over 200 seeds.

The learning is the largest single effect in the record's market. An insurer that estimates each account's loss distribution from its own files and prices it through the rating plan it already has grows its surplus at 3.39 percent a year against 1.93 for a second Loss-ratio insurer in its place, which does not learn, and its growth is the highest of any filling of the fourth seat in the record's

¹⁴Paul A. Samuelson, "Why We Should Not Make Mean Log of Wealth Big Though Years to Act Are Long," *Journal of Banking and Finance* 3, no. 4 (1979): 305–307; L. C. MacLean, W. T. Ziemba, and G. Blazenko, "Growth versus Security in Dynamic Investment Analysis," *Management Science* 38, no. 11 (1992): 1562–1585; Leonard C. MacLean, Edward O. Thorp, and William T. Ziemba, eds., *The Kelly Capital Growth Investment Criterion: Theory and Practice* (Singapore: World Scientific, 2011).

market, shared within the equivalence margin with the Kelly route, which prices the same estimate by the Kelly decision. What the decision changes is the book and the path: a smaller book with a higher renewal margin, new business priced below the plan route's price, more turnover in the market, and a surplus that trails the plan route's through record year 23 before the two routes converge. In the markets where the Kelly route leads, smaller insurers, fewer accounts, a rising severity trend no insurer projects, a buyer less responsive to price, and in those where it trails, a wider filed band and a buyer more responsive to price, its construction is the same.

The learning of this paper is econometric: credibility weighting at estimated moments and a hold curve fitted to the quotes won and lost, each reading the data an insurer holds. The same data admit a deep learner. A neural network need not state the form of the relations among an account's characteristics, its claims, and the market's responses, as the estimators of this paper do: a network with one hidden layer and enough units can approximate any continuous relation on a bounded range of inputs to any stated accuracy.¹⁵ The two routes of this paper take such a learner's estimates as they take the paper's: whatever estimates each account's loss distribution and the chance of winning it can be priced through the plan or by the decision, and the record's measure of what the learning is worth is the measure that learner would be held to. What the criterion's weight on risk would add in a market whose accounts are large against its surplus, and how the two routes fare where all three conventional insurers learn, are questions this record leaves open.

A The Market: Its Assumptions, Its Insurers, and Its Loss Process

This appendix states the market of Section 3 in full: its assumptions (Table A.1), the accounts and their losses, what an insurer knows, the three conventional insurers and how each prices, the buyer and the broker, capital and the measure of growth, and the loss process.

A.1 Accounts and Losses

The market holds 4,000 accounts of one line, written by four insurers without reinsurance. An account has a size, lognormal with median 1 and a log standard deviation of 1, and two traits, one of claim frequency and one of claim size, independent of each other. Each trait, in logarithms, is the sum of a permanent part and a part that drifts as a first-order autoregression with persistence $\rho = 0.85$: the frequency trait's permanent part has a standard deviation of 0.40 and its yearly shock one of 0.30, and the severity trait's permanent part one of 0.632 (a variance of 0.40) and the same shock. In a policy year the account has a Poisson number of claims with mean λ , the product of its size, its frequency trait, and one expected claim a year at the median size. A claim's size is mixed exponential, the form used in increased-limits ratemaking, with means of \$9,277, \$92,766, \$927,662, and \$9,276,615 at weights of 0.700, 0.250, 0.045, and 0.005, scaled by the account's severity trait; the means are set, through one factor of 1.85532 on a stated base curve, so that the book's mean claim limited at the occurrence limit is Verisk's general liability indemnity severity of \$56,213 a claim.¹⁶ Indemnity is limited to \$1,000,000 an occurrence and \$2,000,000 in the aggregate of a

¹⁵Kurt Hornik, Maxwell Stinchcombe, and Halbert White, "Multilayer Feedforward Networks Are Universal Approximators," *Neural Networks* 2, no. 5 (1989): 359–366.

¹⁶The form: Clive L. Keatinge, "Modeling Losses with the Mixed Exponential Distribution," *Proceedings of the Casualty Actuarial Society* 86 (1999): 654–698; Joseph M. Palmer, *Increased Limits Ratemaking for Liability Insurance* (Casualty Actuarial Society study note, 2006). The severity: Verisk, *2024 Executive Insights: General Liability*, 1st ed., Verisk's Core Lines Services Report, Second-Quarter 2019 to First-Quarter 2024 (Verisk, 2024), 4.

policy year, and loss adjustment expense is 0.39156 of the indemnity paid, in addition to the limit, the ratio of Verisk's loss adjustment expense severity of \$22,011 to its indemnity severity. A claim is reported with a lag, 0.55 of a policy year's claims by the end of the year and all of them within four years, and is paid once, at settlement, up to 12 years after it is reported. Until it settles it stands on an insurer's books at a case reserve, its ultimate value times a lognormal error. §A.6 states the process in full.

A.2 What an Insurer Knows

Three sources of information exist, and they are not equally shared. The first is public. A classification score, the sum of an account's two permanent traits read with noise, places each account in one of five classes, and a bureau publishes for each class the market's pooled experience of five policy years ending two years before the rating year: a loss cost, a loss cost with each occurrence and its expense capped at \$100,000, and a claim frequency, each per unit of size. The second is the submission. An insurer that is asked to quote receives the broker's loss run, the account's claims of the past five policy years whoever wrote them, valued one quarter before the policy's inception,¹⁷ with an open claim at a case reserve whose error has a log standard deviation of 0.30. Each insurer's underwriter also forms a read of the account of its own, the permanent traits read with noise of the size of the score's, a standard deviation 0.7 times that of the traits. The third is private. The insurer that holds an account reads its own claim file for every year it has held it, current to the quarter and with half the case-reserve error, and keeps the loss run it received when it acquired the account. The holder therefore knows more about its accounts than a challenger does, and the longer it has held an account the more: this is the asymmetry of the literature on dynamic insurance contracts, and nothing else in the market favors a holder except the buyer's costs of shopping and switching.

A.3 Three Conventional Insurers

Each conventional insurer prices by the bureau's plan and differs from the others in how it sets its level. For an account of class c and size z the plan's expected loss is

$$\bar{x} = L_c z \times M_{\text{exp}} \times M_{\text{sch}} \times R, \quad (\text{A.1})$$

where L_c is the bureau's loss cost, M_{exp} the experience modification, M_{sch} the schedule modification, and R the insurer's rate level, and the premium is $\bar{x}/(1 - \pi)$ with a provision $\pi = 0.34$ for expenses of $\varepsilon = 0.30$ of premium and a profit of 0.04.¹⁸ The experience modification weights the account's capped losses against the class's by credibility: $M_{\text{exp}} = 1 + Z_{\text{exp}}(A - E_{\text{rep}})/E$, with A the account's losses on the file the insurer reads, each occurrence with its expense capped, E the class's expected capped losses for the period, E_{rep} their expected reported part, and $Z_{\text{exp}} = E/(E + K)$. The bureau's plan as filed uses three policy years, the latest year before the rating year left out, and K of \$66,000: the plan states K as \$100,000 of premium, and it enters on the expected-loss basis, times $1 - \pi$, so that a change of the provision moves K with it. An insurer applies the plan's formula over every year its file holds, the latest again left out, a year of age j weighted by d^j , and estimates

¹⁷Association for Cooperative Operations Research and Development, *Commercial Insurance Application*, ACORD 125 (2007/10), 3, which asks for the claims or losses of the prior five years.

¹⁸National Association of Insurance Commissioners, *Report on Profitability by Line by State in 2024* (Kansas City, Missouri: NAIC, 2026), 9: for other liability countrywide, a line wider than general liability, general expense, selling expense, and taxes, licenses, and fees of 8.2, 18.8, and 1.9 percent of net premiums earned, 28.9 percent in all.

d and K from its own book by least squares (Appendix B). The schedule modification is what the underwriter's own reading adds to the class, within a filed band of 25 percent either way.¹⁹

The *Bureau insurer* files a rate level, a multiplier on the bureau's loss costs, and refiles it when the indication on its own experience reaches the threshold below. At each year's close it computes the loss-ratio method's indication on its own experience of the past five policy years: with a its reported losses, e the expected reported losses it priced those years at, brought to its present level, and n its reported claims, the indicated change is

$$Zf + (1 - Z)\psi, \quad f = a/e, \quad Z = \min(1, \sqrt{n/1,082}), \quad (\text{A.2})$$

with the complement of credibility ψ equal to one.²⁰ The standard of 1,082 claims is the full-credibility standard of classical (limited-fluctuation) credibility for claim frequency, applied here to losses, as practice often applies it. Where the indicated change differs from one by 0.10 or more it refiles, and its level is multiplied by the change; otherwise its level stands.

The *Loss-ratio insurer* moves its level every year by Equation A.2 on a window of five policy years, the span of Werner and Modlin's worked indications, with the complement ψ set to competitors' rates: the level of the winning prices relative to its own plan prices, fitted as a censored normal to the broker's figures on the quotes it lost at market (§A.4) and to the quotes it bound (Appendix B). The complement takes that level once the window holds 100 lost quotes, and one before.

The *GLM insurer* adds relativities of its own to the plan. Each year it fits a multiplicative model with a logarithmic link to the capped reported losses of its own book over five policy years, with a factor for each class and for each of three size bands, thirds of the year's accounts by manual premium, the premium at the bureau's loss cost for the account's class and size before the experience and schedule modifications, and takes each fitted factor at its classical credibility; its level moves by the ratio of its book's reported losses to their expected amount at its present relativities, taken at the classical credibility of the book's claims (Appendix B).²¹

On new business, a quote on an account it does not hold, a conventional insurer files no load. In every one of Wu and Lin's books new business runs at a higher loss ratio than renewals. In their two books of general liability, new policies receive schedule debits about half as often as renewals and schedule credits as often: there, new business is priced in practice no higher than the plan states, and the loss-ratio difference is a cost the price does not carry. Wu and Lin also judge a surcharge on new business justifiable on their data, so a filed load of zero follows the practice their general liability books show and not their reading of what the data justify.²² What an underwriter may move without a filing is the schedule modification, and insurers do not always implement the rate their experience indicates.²³ In the middle and the largest band the insurer therefore applies a schedule credit or debit to its new-business quotes, equal to the level of the winning prices it has seen in the band relative to their level in the smallest band, from the broker's figures of the past

¹⁹Fisher et al., *Individual Risk Rating*; a limit of 25 percent either way is that of Florida Administrative Code, Rule 69O-170.004. The same rule, at (2)(b), excludes from the characteristics a schedule may rate "the degree of competition for the risk or the rates which may be offered by other insurers"; the schedule credit on new business described below and the Kelly route's deviation within the band (§D.1.4) answer competitors' prices, as underwriters use schedule rating to meet a competitor's quote in practice, and the model follows the practice: Lisa A. Hays, "Managing Commercial Lines Pricing Levels in a Loss Cost Environment," *Casualty Actuarial Society Forum* (Winter 2002): 11.

²⁰Laurence H. Longley-Cook, "An Introduction to Credibility Theory," *Proceedings of the Casualty Actuarial Society* 49 (1962): 194–221; Werner and Modlin, *Basic Ratemaking*, chaps. 8 and 12, and apps. A–D, the worked indications.

²¹Robert A. Bailey, "Insurance Rates with Minimum Bias," *Proceedings of the Casualty Actuarial Society* 50 (1963): 4–11; Werner and Modlin, *Basic Ratemaking*, chap. 10.

²²Wu and Lin, "Large Scale Analysis," 398, 402, and 405.

²³Werner and Modlin, *Basic Ratemaking*, chap. 13; Hays, "Commercial Lines Pricing Levels," 11.

three years, with the schedule credit and the schedule modification together kept inside the filed band. The smallest band is the reference and carries no schedule credit. The rule refers to account size only through the choice of the smallest band as the reference.

A.4 The Buyer and the Broker

An account renews once a year, in the quarter of its inception. Its holder quotes. The broker invites one challenger, drawn at random among the other insurers that are open to new business, and each further one with a probability equal to the account's rank by manual premium among the year's accounts, from zero for the smallest through one for the largest, so that the smallest account draws one competing quote and the largest draws all three. Let Δ be the logarithm of the holder's quote over the lowest invited quote. With probability 0.10 the account is remarketed for reasons unrelated to price and chooses among all the quotes by a logit on log price with the scale s of the response below. Otherwise it goes to market with probability $1/(1 + \exp(-(\Delta - k)/s))$ and, once there, stays with its holder with probability $1/(1 + \exp((\Delta - k')/s))$, taking the lowest invited quote if it does not; k is the cost of shopping, k' the cost of switching, and $s = 0.05$ the scale of the response, each in units of log price. The two costs are equal and calibrated rather than stated. On seeds outside the record, $k = k'$ was moved until 0.14 of the accounts of the market of four conventional insurers changed insurer a year, half the 28 percent of middle-market companies that a 2007 survey by Greenwich Associates found changing an insurer over two years; the halving counts no company changing twice, and independent yearly switching would give 0.151. The calibration gives $k = k' = 0.061184$. At that cost, in the calibration's own run, 0.1404 of the accounts change insurer a year (standard error over 20 seeds 0.0033), 0.2654 go to market, and 0.8313 of premium is retained, between two retentions of commercial business published in 2026.²⁴

After placement the broker tells an insurer that quoted on an account that went to market and did not bind the difference between its quote and the bound premium, as a fraction of the bound premium rounded to the nearest 0.05. The insurer that bound is told nothing, and nothing is said of an account that did not go to market. This is the only information about competitors' prices that any insurer receives.

A.5 Capital, Solvency, and the Measure

Within a seed every insurer starts the first record year with the same endowment, set so that a seat's expected written premium of the opening book and the steady-state reserves of a book of that premium stand at the industry's aggregate risk-based-capital (RBC) ratio of 6.42458 (Appendix C); over the seeds of the record it averages \$556.7 million, from \$519.3 through \$605.8 million.²⁵ A policy's premium is booked with expenses of 0.30 and a reserve at the writer's own expected loss;

²⁴Greenwich Associates, *Insurance Market Pulse: Looking at the Small Business and Middle Market Segments* (2007), the share of middle-market companies that changed a property-casualty or employee-benefits insurer over the two years before the survey, as reported in Douglas McLeod, "Changing Insurers Yields Surprise," *Business Insurance*, June 17, 2007. The retentions: The Hartford Insurance Group, Inc., *Annual Report on Form 10-K for the Fiscal Year Ended December 31, 2025* (filed with the U.S. Securities and Exchange Commission, February 20, 2026), a middle-market premium retention of 83 percent; The Travelers Companies, Inc., "Travelers Reports Excellent Second Quarter and Year-to-Date Results," press release, July 17, 2026, a retention of 86 percent in its Business Insurance segment, retention being the premium available for renewal that was retained, excluding rate and exposure changes, with the segment's National Accounts left out.

²⁵National Association of Insurance Commissioners, "Aggregated P&C RBC Data: 2025 Data as of June 1, 2026," *Summary: Aggregated P/C RBC Results by Year* (NAIC, 2026).

the reserve runs off as claims are reported and paid; and surplus and reserves alike earn the *return on funds*, $r = 0.0356$ a year, the industry's published investment yield.²⁶ At each year's close an RBC ratio is computed for each insurer, its surplus over a requirement of the form of the NAIC in its written premium and its reserves (Appendix C). Capital affects behavior at three thresholds. The broker reads the ratio: below 3 it withholds a new-business invitation with a probability that rises to one at 2 (broker flight), and a renewal is never affected. The insurer reads its own writings: in a quarter that opens with its written premium of the year above three times its surplus, the level at which the NAIC's Insurance Regulatory Information System (IRIS) ratio of premium to surplus is flagged, it quotes no new business, and it keeps quoting renewals. And below a ratio of 0.70, or with its surplus exhausted, the insurer fails and is replaced by an entrant, a Bureau insurer with the endowment, at most two times in a seed.²⁷

The books are built over 19 years in which accounts change holder at random at a rate of 0.15 a year, so that every insurer opens with a seasoned book and a file on each account. The market then runs 10 practice years, in which every surplus is reset to the endowment at the year's close while rate levels, estimators, and books carry forward, and 28 record years; at the close of the final practice year, before the reset, each insurer's reserves on the business it wrote before the record years are restated, the provision for unreported claims at its true expectation and each open claim's case reserve at what the claim will pay (Appendix C). An insurer's *growth* is

$$g = \ln \frac{S_T}{S_0} - T \ln(1 + r), \quad T = 28, \quad r = 0.0356, \quad (\text{A.3})$$

the terminal logarithm of its surplus over the endowment, net of the return on funds, so that g is zero for a surplus that earns that return and nothing else, and g/T , a logarithm, is a yearly rate of growth from writing business, the investment income on the reserves an insurer holds included; its exponential less one is the rate a year the tables state. Surplus bears no income tax and pays no dividend. The terminal surplus S_T is taken from the insurer's own books at the close of the final record year, its reserves at its own estimates: no restatement like that of the final practice year is made there, and the paper does not measure the difference one would make. Every *arm* of the record, one filling or variation of the market, runs on seeds 0 through 199, and every run is a function of its arm and its seed alone, in double precision on one thread. An insurer that fails leaves the market, and its growth is not defined; no insurer fails in any seed of any arm of the record, so every mean and contrast is over all 200 seeds, and the record does not test the decision's exposure to ruin. A contrast between two arms or two insurers is paired by seed, with its mean, its t statistic (the mean over its standard error, the sample standard deviation over seeds divided by the square root of the number of seeds), and the number of seeds in which it is above zero. A contrast between two arms is a total effect: what changes in one seat moves the other insurers' books, rate levels, and capital, and the contrast carries their response; it is not the effect of the change with the competitors held as they were.

²⁶National Association of Insurance Commissioners, *U.S. Property & Casualty and Title Insurance Industries—2025 Full Year Results* (NAIC Financial Regulatory Services, 2026).

²⁷The failure threshold is the level at which the NAIC's model act mandates control: National Association of Insurance Commissioners, *Risk-Based Capital (RBC) for Insurers Model Act*, Model 312 (2012).

A.6 The Loss Process

An account i has a size z_i , lognormal with median 1 and log standard deviation 1. Its frequency trait in year t is $u_{it} = \alpha_i + d_{it}$, with α_i normal with mean zero and standard deviation 0.40 and $d_{it} = \rho d_{i,t-1} + \eta_{it}$, $\rho = 0.85$, η_{it} normal with standard deviation 0.30, started from its stationary distribution. Its severity trait is of the same form, with a permanent standard deviation of 0.632 (a variance of 0.40) and the same persistence and shock, independent of the frequency trait. The stationary variance of a trait's drifting part is $0.30^2/(1 - \rho^2)$, so the frequency trait has a variance of 0.484 and the severity trait one of 0.724. Each trait enters as its exponential less half its variance, so that it has a mean of one across accounts in every year.

In policy year t the account's claim count is Poisson with mean

$$\lambda_{it} = \lambda_0 z_i \exp\left(u_{it} - \frac{1}{2}\text{Var}(u)\right), \quad \lambda_0 = 1. \quad (\text{A.4})$$

A claim's size before any limit is exponential with a mean drawn from the four means of §A.1 at their weights and multiplied by the account's severity trait; the curve's mean before any limit is \$117,813, and limited to \$1,000,000 an occurrence it is \$61,964. Within a policy year the aggregate limit of \$2,000,000 erodes in the order in which the claims are paid. Loss adjustment expense is 0.39156 of the indemnity paid and is outside the limits. An account's true expected loss in a year, the quantity the tables call by that name, is λ_{it} times the expected claim limited per occurrence at the account's severity trait, with loss adjustment expense, times the share of the year's expected occurrence-limited losses that the aggregate limit leaves.

A claim occurs in a quarter of the policy year drawn uniformly. It is reported within the policy year with probability 0.55, by the end of the second year with probability 0.85, of the third with 0.95, and of the fourth with certainty: in the policy year, in a quarter drawn uniformly from its occurrence quarter through the year's final quarter, and in a later year uniformly over that year's quarters. It settles between one quarter and 12 years after it is reported, 0.30 of claims within a year of report, 0.50 within two, 0.82 within five, and all within 12. An open claim is carried at a case reserve, its ultimate value times $\exp(\sigma_c \xi - \sigma_c^2/2)$ with ξ standard normal, one draw a claim, and σ_c of 0.30 on a broker's loss run and 0.15 on the holder's own file. A settled claim stands at what was paid.

Table A.1 collects the assumptions of this appendix.

Table A.1: The market’s assumptions, each with its value in the record’s market and its grounding. An assumption is grounded in a source, a calibration to a sourced target, a source’s form, the objective, or the simulation’s scale, or it is stated, where no source was found and the value is the model’s own. The settings of Table F.1 vary the Bureau insurer’s threshold and the read noise, the markets of Section 6 the filed band, the buyer’s scale, the accounts, the endowment, the severity trend, and the insurer in the GLM insurer’s seat, and the fillings of Table 1 the renewal credit’s discount. Amounts are per occurrence or per policy year and rates per year; the sources are in the text. Source: authors’ construction.

Assumption	Value	Grounding
Account size, log standard deviation	1	stated
Claims a year at the median size	1	stated
Frequency trait: spread, persistence, shock	0.40, 0.85, 0.30	stated
Severity trait: variance, persistence, shock	0.40, 0.85, 0.30	stated
Claim size, at Verisk’s severity	mixed exponential	Keatinge; Palmer; Verisk
Loss adjustment expense over indemnity	0.39156	Verisk
Limits: occurrence; aggregate (\$)	1,000,000; 2,000,000	Palmer
Share of claims reported by year	0.55, 0.85, 0.95, 1.00	stated
Case-reserve error: loss run; holder’s file	0.30; 0.15	stated
Classes; size bands	5; 3	scale
Read noise, over the traits’ spread	0.7	stated
Bureau’s statement: years; lag	5; 2	stated
Loss run: years; valuation lag, quarters	5; 1	ACORD; stated
Full credibility, claims	1,082	Longley-Cook, frequency
Filed experience plan: years; cap; K (\$)	3; 100,000; 66,000	bureau plan’s form
Filed band, either way	0.25	Florida rule
Provision π ; expenses ε ; profit	0.34; 0.30; 0.04	NAIC; stated
Filed load on new business	0.00	Wu and Lin, GL books
The load’s rule	schedule credit by band	stated
Lost quotes to fit: complement; schedule	100; 30	stated
Windows: loss-ratio method; GLM fit, years	5; 5	Werner and Modlin; stated
Renewal credit’s window, years	3	stated
Schedule fit’s window, years	3	stated
Bureau insurer: experience, years; threshold	5; 0.10	stated
Costs of shopping and switching, $k = k'$	0.061184	calibrated to Greenwich
Scale of the buyer’s response s	0.05	stated
Remarketing unrelated to price	0.10	stated
Invitations: power of the account’s rank	1	stated
Broker’s figure rounded to	0.05	stated
Capital requirement: ρ ; κ	1.72456; 0.0968742	NAIC 2025
Flight band, RBC ratio (NAIC trend test)	2 through 3	NAIC; slope stated
New-business stop, premium over surplus	3	NAIC IRIS; the stop stated
Failure, RBC ratio	0.70	NAIC Model Act
Return on funds, on surplus and reserves	0.0356	NAIC yield
Endowment, RBC ratio	6.42458	NAIC 2025
Entrants after failures, at most	2	stated
Hold curve: prior spreads; decay	0.25, 1.0; 0.98	stated
Renewal credit: discount; horizon	return on funds; final record year	the construction’s choice, both dis- counts run (Section 4); the objective
Accounts; years of build-up, practice, record	4,000; 19, 10, 28	scale
Build-up: holders changing a year	0.15	stated

B The Rating Plans

This appendix states the rating plans in their formulas: the classes and the bureau's statement, the experience and schedule modifications, the GLM insurer's fit, and the schedule credit on new business.

The classification score is α_i of the frequency trait plus the permanent part of the severity trait plus normal noise with a standard deviation 0.7 times that of the sum, and the classes are the quintiles of its distribution. For policies of year t the bureau's statement covers policy years $t - 6$ through $t - 2$, valued at the start of year t . For each class it gives the reported losses, the reported losses with each occurrence and its expense capped at \$100,000, and the reported claims, each divided by the public expected reported share of its year and summed over the market, per unit of size; a class's figure is weighted with the all-class figure at the classical credibility of the class's reported claims.

In the experience modification, for an account of class c and size z let $B_c z$ be the expected capped loss of a year. Over the years of the file the insurer reads, a year of age j (the latest year of the period at age zero, the period ending two years before the rating year) weighted by $w_j = d^j$,

$$M_{\text{exp}} = 1 + \frac{E}{E + K} \cdot \frac{A - E_{\text{rep}}}{E}, \quad A = \sum_j w_j A_j, \quad E = B_c z \sum_j w_j, \quad E_{\text{rep}} = B_c z \sum_j w_j F_j, \quad (\text{B.1})$$

with A_j the year's capped reported losses on the file and F_j the public expected reported share of the year at the file's valuation. The holder applies it to its own claim file and a challenger to the broker's loss run. Each insurer estimates (d, K) once a year by least squares over a grid, d from 0.5 through 1.0 and K from \$25,000 through \$500,000: for the accounts it held and renewed in the latest expired policy year it compares what each pair would have given as that year's modification, from the file as it stood when the year was quoted, with the year's capped losses as now reported, and it adds the weighted squared differences to those of the years before. Until its first estimate it uses the plan as filed, $d = 1$ over three years and K of \$66,000.

In the schedule modification, with (ν_1, s_1^2) the mean and variance of the permanent traits' sum in the account's class and (ν_2, s_2^2) its mean and variance given the class and the underwriter's reading, each under a normal approximation, the modification is $\exp(\nu_2 - \nu_1 + (s_2^2 - s_1^2)/2)$, held within 1 ± 0.25 .

On the policies the GLM insurer wrote in the five policy years ending at the year's close, it fits $A \approx E \times f_{\text{class}} \times f_{\text{band}} \times f_{\text{new}}$, with A a policy's capped reported losses, E their expected reported amount at its present relativities, and f_{new} a factor by band on new business, by iterative proportional fitting, the estimator of a Poisson likelihood with a logarithmic link. The band factors and then the class factors are brought to a mean of one, weighted by E . The change of level is the ratio of the book's reported losses over those years, uncapped and with expense, to their expected reported amount at the present relativities. The level, each class relativity, and each band relativity are then multiplied by $Zf + 1 - Z$, with f the change of level or the fitted factor and Z the classical credibility of the claims behind it. The factors on new business do not enter its price.

For the schedule credit on new business, the insurer takes for each band, from its quotes as a challenger of the past three years on accounts that went to market and were placed, the logarithm of the bound premium over its own plan price at its present level: observed where the broker gave a figure, and known only to be at or above its quote where the quote bound. A normal distribution is fitted by maximum likelihood to both kinds of observation, and the band's level of winning prices is the fitted mean plus half the fitted variance. The schedule credit in the middle and the largest band

is the exponential of the band's level less the smallest band's. A band's schedule credit changes only where the band and the smallest band each hold 30 lost quotes in the window; otherwise it keeps its previous one. The same fit over all bands, on its own window of five years, gives the Loss-ratio insurer's complement of credibility ψ once it has 100 lost quotes in the window.

C The Capital Standard

This appendix states the capital standard in its formulas: the risk-based-capital ratio, broker flight, the stop on new business, and the endowment the standard sets.

At a year's close an insurer with written premium W in the year, reserves Q , and surplus S has the risk-based-capital ratio

$$\text{RBC} = \frac{S}{\kappa \sqrt{W^2 + (\rho Q)^2}}, \quad \rho = 1.72456, \quad \kappa = 0.0968742. \quad (\text{C.1})$$

The square-root combination is the form of the NAIC's formula, in which the premium-risk and reserve-risk charges combine under a root; the model carries those two charges and absorbs the others in κ in the industry's proportion, the covariance term of the two underwriting charges alone standing at 0.470 of the industry's total risk-based capital after covariance, of which the Authorized Control Level is half. The two constants are calibrated on one year of the NAIC's aggregated data: ρ is the reserve charge per dollar of reserves over the premium charge per dollar of net written premium, and κ is set so that the industry's own point—net premium at 0.773 of surplus, reserves at 1.05664 times premium—stands at the industry's aggregate ratio of 6.42458.²⁸ A failed insurer's seat (§A.5) passes, while entries remain, to an entrant at a level of one with the endowment and no book. The model insurer's reserves run above the industry's all-lines ratio to premium, as a long-tail book's do, so at an equal ratio of premium to surplus its RBC ratio is lower than the industry's, as the formula's reserve charge makes it.

Broker flight is the broker's response to the ratio. At each year's close the broker reads each insurer's RBC ratio and, in the following year, extends a new-business invitation to it with probability $\min(\max(\text{RBC} - 2, 0), 1)$: one at a ratio of 3 or more, falling linearly to zero at 2, account by account by its own random stream. The band is the NAIC trend-test band and stands in the model as a proxy for the financial-strength thresholds brokers act on; the steepness of the withdrawal has no source. Flight is suspended in the practice years, and the first record year reads the final practice year's closing statement with the surplus reset to the endowment.

The stop on new business (§A.5) reads IRIS ratio 2, net premiums written over policyholders' surplus, at the level the NAIC flags, on the year's written premium to date; the test is made again each quarter against the surplus of the moment, and the year's closing ratio can end above the threshold on renewals alone.²⁹

The endowment (§A.5) is set so that the opening book's expected written premium W and the steady-state reserves of a book of that premium, $Q = (1 - \pi) \tau W$ with τ the mean time from booking to payment (§D.1.2), stand at the industry's aggregate RBC ratio, which gives an opening ratio of written premium to surplus of about a third in every seed. Surplus changes within a quarter by

²⁸National Association of Insurance Commissioners, "Aggregated P&C RBC Data"; National Association of Insurance Commissioners, 2025 *Full Year Results*.

²⁹National Association of Insurance Commissioners, "NAIC IRIS Ratios," data year 2025, property and casualty (NAIC, May 19, 2026).

the premium written less expenses and less the reserve set up at the writer's expected loss, by the release of that reserve along the public reporting pattern, by the case reserves set up on the claims reported in the quarter and released at their settlement, by the payments made, and by investment income on surplus and reserves (§A.5). In a practice year no insurer fails; the restatement of the reserves at the close of the final practice year (§A.5) takes the provision for unreported claims at its true expectation, not conditioned on the claims already reported.

D The Kelly Decision and the Estimators

This appendix states the approach of Section 2 with its equations: the objective and the decision (§D.1), what is learned, from what data, and by what estimator (§D.2), the estimators in detail (§D.3), and the decision's variance term (§D.4).

D.1 The Objective and the Decision

This section states the objective and the decision that follows from it: the criterion (§D.1.1), the increment one account adds to expected growth (§D.1.2), what keeping an account is worth (§D.1.3), the decision, the chance of winning at a price weighed against the increment, with the lowest premium at which the insurer writes (§D.1.4), and the form in which the value of keeping an account enters the decision, with its anchoring (§D.1.5). Each quantity is stated with the data it reads, and Appendix E gives the computation in the order in which it is carried out.

D.1.1 The Objective

Let S be the insurer's surplus. The Kelly criterion is the objective: the insurer chooses its actions to maximize the expected logarithm of its surplus at the horizon, $\mathbb{E}[\ln S_T]$. Since $\ln S_T - \ln S_0$ is T times the mean yearly growth of the logarithm of surplus, the mean geometric rate of growth in logarithms, this is the maximization of that expected rate; the rate a year the tables state is its exponential less one at the expected logarithm. Where an action leaves the insurer's later opportunities unchanged, the objective is the maximization of $\mathbb{E}[\ln S']$, the expected logarithm of next period's surplus, in each period; where it changes them, as when an account written today is held in later years (§D.1.3), the one-period objective carries what the action leaves for later. Two features of an insurer's position change the statement of the criterion's property in repeated independent bets on known odds (§2.1). The insurer does not choose a stake at given odds: it quotes a price, the buyer decides, and the account is written whole or not at all. And the insurer does not know the distribution of an account's loss; it has an estimate. The decision below approximates the criterion for the distribution the insurer is given: to second order in the account's share of surplus, one account and one year at a time, with the value of later years entered to first order (§D.1.3). The property of optimal growth is not claimed for this approximation in this market.

D.1.2 The Increment of One Account

Consider one account with loss X , estimated mean $\hat{x} = \mathbb{E}[X \mid \text{data}]$, and estimated variance $\hat{v}$, quoted at premium P . Written, it changes surplus by $n - X$, where $n = P(1 - \varepsilon)$ is the premium net of expenses. Expanding the logarithm to second order in the account's share of surplus, the

account taken alone against the current surplus, and dropping the square of the expected margin, which is small against the variance,

$$\mathbb{E} \left[\ln \frac{S + n - X}{S} \right] \approx \frac{n - \hat{x}}{S} - \frac{\hat{v}}{2S^2}, \quad (\text{D.1})$$

which is a mean–variance rule in which the weight on variance, one half over the surplus, is fixed by the objective and is not a parameter. Two further terms enter the increment. The premium is held against the policy until its claims are paid and earns the return on funds meanwhile: with $\tau = 4.083$ years the mean time from booking to payment under the market’s reporting and settlement patterns, the expected loss net of that return is $\hat{x}(1 - c_x)$ with $c_x = r\tau = 0.1454$, the return over the mean time to payment without compounding. And in a market where an account written today is likely to be held in later years, the increment carries what keeping the account is worth, V per unit of net premium, the renewal credit of §D.1.3. The increment of the decision, in the form first registered, the quote-anchored credit (§D.1.5; Equation 1 gives the anchored form of the construction of record), is

$$G(P) = \frac{n(1 + V) - \hat{x}(1 - c_x)}{S} - \frac{\hat{v}}{2S^2}. \quad (\text{D.2})$$

In the market of Section 3 the variance term is small. At the endowment’s mean over the seeds, \$556.7 million, it is 0.0641 percent of the expected loss for an account at the curve’s own scale of claim sizes, and since the expected loss and the variance both rise in proportion to the claim rate, it does not vary with the account’s size. §D.4 measures what removing it does.

D.1.3 What Keeping an Account Is Worth

Equation D.1 values an account by one year’s margin. In a market where the buyer bears a cost of switching and the holder knows the account better than a challenger does, an account written today is likely to be held next year, and the objective, which is growth over the whole horizon, counts the margins of the years to come. The objective also fixes how they are discounted. A margin m that an account adds k years ahead raises $\ln S_T$ by about m/S_{t+k} , and where the insurer’s surplus grows at its own yearly rate $\hat{g}$, gross of the return on funds, $S_{t+k} \approx S_t e^{\hat{g}k}$, the margin counts about $e^{-\hat{g}k}$ in units of today’s surplus: the discount factor is $\beta = e^{-\hat{g}}$. The step replaces the expectation of S_t/S_{t+k} by its value on the path of mean growth, which by Jensen’s inequality understates it; the funds discount of Section 4 takes another first-order route, which discounts at the return on funds. A margin that falls after the final record year adds nothing to $\ln S_T$, so the sum runs through the horizon H , the number of years from the year after the one quoted through the final record year. With p_1 and m_1 the retention and the expected margin per unit of net premium at an account’s first renewal, and p_s and m_s at later renewals,

$$V_{\text{ren}}(H) = m_s \sum_{k=1}^H (\beta p_s)^k \quad \text{and} \quad V_{\text{nb}}(H) = \beta p_1 [m_1 + V_{\text{ren}}(H - 1)], \quad (\text{D.3})$$

with $V_{\text{nb}}(0) = 0$. This is the *renewal credit*: for an account the insurer holds, the discounted margins of the renewals to come; for an account it does not, the first renewal’s margin and what follows it, reached with probability p_1 . The five quantities are estimated on the insurer’s own book over the credit’s window (three years): the retentions in accounts, the margins as Equation D.2’s numerator without the credit per unit of net premium, on first renewals kept and on later renewals

kept apart—the margin of a newly won account’s first renewal is the margin of a selected account, which is what the decision needs—and $\hat{g}$ as the window’s mean yearly log growth of surplus, the return on funds included, each year’s growth taken from the surplus at its opening to that at its close and the window reading the practice years as it reads the record years. The derivation carries no cap: a negative margin gives a negative credit, and a finite horizon, here the number of record years, needs no bound for any $\hat{g}$ (a going-concern sum exists only where $\beta p_s < 1$). Its discount is the insurer’s own growth on the growth discount and the return on funds on the funds discount. Until the window holds a first renewal, the credit on new business is zero.

D.1.4 The Chance of Winning, and the Decision

A quote is a bid, and the increment $G(P)$ is earned only if the quote holds the account. Let $h(\ell)$ be the probability that a quote at log deviation ℓ from the *undeviated price*, the plan route’s price for the account, holds. The expected increment of quoting is $h(\ell) G(P_\ell)$ with P_ℓ the undeviated price times e^ℓ , and the Kelly decision is the premium of the deviation that maximizes it:

$$\ell^* = \arg \max_{\ell} h(\ell) G(P_\ell), \quad P = \max \left\{ P_{\ell^*}, \frac{\hat{x}(1 - c_x) + \hat{v}/(2S)}{(1 - \varepsilon)(1 + V)} \right\}. \quad (\text{D.4})$$

A higher deviation earns more on an account that is held and holds fewer accounts; the two are weighed by the insurer’s own hold curve, taken as the posterior average that §D.2.2 states. Because the Kelly route prices through a schedule plan as the others do, its deviation, the schedule modification, and, on new business, the schedule credit together stay inside the filed band of 25 percent either way, so that ℓ^* ranges over the band and is at most $\ln 1.25 = 0.223$ on an account whose schedule modification and schedule credit are one; the zero deviation is always admissible. The second element of Equation D.4 is the *floor*: the premium at which Equation D.2 is zero, the lowest at which writing the account adds to expected growth, with the credit of the account’s phase, renewal or new business, inside it. The floor rises with the account’s variance, falls as surplus grows, and falls with the value of keeping the account, so that of two accounts with one expected loss the one with the wider distribution is quoted no lower, and an insurer with less surplus quotes an account of the same distribution no lower. Without the credit and the return on the premium held, the floor’s premium net of expenses is the expected loss plus a loading on the variance of one over twice the surplus, the premium of the variance principle at that loading; the fair premium of a discounted-cash-flow approach to rate regulation and the allocation of capital among lines of business are two other constructions of an insurer’s charge for risk.³⁰

Equation D.4, over the filed band around the undeviated price, is the decision. Its inputs other than the surplus S , the expense share ε , and the return on the premium held, c_x , are $\hat{x}$ and $\hat{v}$ for each account, the credits V_{nb} and V_{ren} , the hold curve h , the undeviated price, and the band. A rating plan supplies a version of $\hat{x}$, the undeviated price, and the band; it supplies none of the others. An insurer may hold a learned estimate and price it as a plan prices, or price it by the decision as a deviation from that price; the appendices name the first *learning at the plan’s pricing* and the second *learning at the Kelly decision*, the plan route and the Kelly route of Section 5.

³⁰Hans Bühlmann, *Mathematical Methods in Risk Theory* (Berlin: Springer-Verlag, 1970), chap. 4; Stewart C. Myers and Richard A. Cohn, “A Discounted Cash Flow Approach to Property-Liability Insurance Rate Regulation,” in *Fair Rate of Return in Property-Liability Insurance*, ed. J. David Cummins and Scott E. Harrington (Boston: Kluwer-Nijhoff, 1987), 55–78; J. David Cummins, “Allocation of Capital in the Insurance Industry,” *Risk Management and Insurance Review* 3, no. 1 (2000): 7–27.

D.1.5 The Form of the Credit, and Its Anchoring

Equation D.2 carries the credit as V per unit of the net premium n . In the decision of §D.1.4 the premium varies with the deviation under consideration, and the form in which the credit follows it is a choice of construction with a measured consequence. In the form first registered, the *quote-anchored credit*, the credit enters the increment as $Vn(\ell)$, with $n(\ell)$ the net premium of the quote at deviation ℓ , in proportion to the premium of the deviation under consideration, so that the gradient of the objective carries V times the premium's own derivative: a markup earns the value of keeping the account in proportion to its size. The continuation of a won account is invariant to today's deviation to first order: next year's quote is next year's rule applied to next year's estimate, and the buyer's decision to shop compares the holder's renewal quote with the lowest invited quote (§A.4), which excludes the holder's own raise. The invariance is argued from the market's rules; the record does not measure each won account's continuation by the deviation that won it. What today's deviation buys is today's margin and the chance of winning the account at all: the hold probability $h(\ell)$ of Equation D.4 multiplies the whole increment, so the value of keeping the account, $h(\ell) V n(0)$, is carried in full. The construction of record therefore anchors the credit at the undeviated price, the price of the zero deviation. This is the *anchored credit*, $Vn(0)$ in place of $Vn(\ell)$. It removes the gradient of the quote-anchored form, and at a deviation other than zero it also changes the credit's level. The quote-anchored form is the form registered first, in the first registration, on the construction the paper does not report, and the anchored form was registered in the second, before any run of it on the record's seeds (Appendix F); Section 4 measures what each form is worth.

The anchored form has a stated limit. The policy's own renewals are written above the zero deviation on average, so a credit anchored there understates the credit's level by the policy's mean renewal deviation factor, and since the hold probability multiplies the whole increment, an understated level biases the chosen deviations upward relative to a credit at the realized scale. An anchor at the realized scale would make the credit a function of the deviation it prices, the self-reference the anchor removes, so the understatement is a limit of the construction: the deviations of Section 5 are higher than a credit at the realized scale would give, and the record does not measure by how much.

D.2 Learning

This section states what the two routes learn and how: each account's loss distribution, tracked from the files the insurer reads (§D.2.1), which both routes price; and the hold curve, from the broker's figures (§D.2.2), which the Kelly route's decision reads with the renewal credit of §D.1.3, whose five quantities are estimated on the insurer's own book (Appendix E). Every estimator reads only data that an insurer in the market of Section 3 holds, and what the market treats as public: the bureau's curve of claim sizes and its moments at scale one, the reporting and settlement patterns (§A.1, §A.6), the spread of the case-reserve error on a broker's loss run and on a holder's own file (§A.2, §A.6), the share of loss adjustment expense (§A.1), and, as for every insurer, the distribution of the traits within a class that the schedule modification reads (Appendix B).

D.2.1 The Account's Loss Distribution

A rating plan reads an account's history through one number, the experience modification on its capped losses. The learning tracks two quantities for each account, its claim frequency relative to

its class and the scale of its claim sizes, and it does so by the best linear predictor of credibility theory, at moments it estimates, in a form that allows for a risk that changes over time.³¹ For each year of a file it forms two observations: the reported claim count over its expected number at the class's frequency, and the mean logarithm of the reported claim sizes less the curve's mean log size, an open claim's case reserve corrected for its error. Its estimate of the coming year's frequency relativity and log scale is the best linear predictor from these observations jointly, in which the weight on a year depends on how much the year's observation carries (the expected claims behind it) and on how far back it lies.

The predictor needs the second moments of the two traits: how much of an account's departure from its class stays and how much decays, and at what rate. The insurer estimates them. From every file and loss run it has read it keeps the autocovariance of each observation at lags of one through 20 years, fits to each a constant plus a geometrically decaying term, a part that stays and a part that decays with a persistence, and keeps one covariance between the two traits (§D.3). On an account it holds it reads its own file, every year it has held the account and the loss run from before; on an account it is asked to quote it reads the broker's loss run. The expected loss is then the class's claim frequency times the account's size, times the tracked relativity, times the expected limited claim at the tracked scale, with loss adjustment expense, and the variance is that of a compound Poisson sum at the same rate and scale, which leaves out the variance from the uncertainty of the rate and the part of the scale's uncertainty common to the account's claims:

$$\hat{x}_0 = \hat{\lambda} \mathbb{E}[\min(Y, u)] (1 + a), \quad \hat{v}_0 = \hat{\lambda} \mathbb{E}[\min(Y, u)^2] (1 + a)^2, \quad (\text{D.5})$$

with $\hat{\lambda}$ the tracked claim rate, Y the claim size at the tracked scale, u the occurrence limit, and a the loss adjustment expense's share. The underwriter's schedule modification multiplies $\hat{x}_0$ as it does a plan's estimate, and $\hat{v}_0$ with it; from here on $\hat{x}_0$ and $\hat{v}_0$ denote the two after it. On either route the price carries the insurer's rate level R , so that the expected loss in the price is $\hat{x} = R \hat{x}_0$, and on the Kelly route the variance is $\hat{v} = R^2 \hat{v}_0$. The schedule modification and the tracked estimate both read the account's permanent traits, the first through the underwriter's reading and the second through the account's claims; their product is not adjusted for what the two share, and the rate level, moved each year by the loss-ratio method (§A.3), is the one correction of its level on the business written. The variance is the quantity a plan does not hold; it enters the Kelly route's price through the increment of Equation D.2, which Equation D.4 multiplies by the hold probability and maximizes, and through the floor of Equation D.4.

D.2.2 The Hold Curve

The hold curve h of Equation D.4 is learned from the only data on competitors' prices an insurer has, the broker's figure. The figure is censored: it is given on the quotes that did not bind on an account that went to market and withheld on those that bound. Let ζ be a latent log price: the logarithm of the premium at which the account would bind with another insurer, over the undeviated price for it at the quote, in each *segment*, a size band in renewals or in new business. On a quote that was lost at market, ζ is observed, the quote's own deviation less the figure. Every quote that bound enters as the event $\zeta \geq \ell$, since its holder is told nothing, including a renewal kept against a lower competing quote and one kept without going to market, which the market may also have undercut;

³¹The predictor, for a risk that does not change over time: Hans Bühlmann, "Experience Rating and Credibility," *ASTIN Bulletin* 4, no. 3 (1967): 199–207. A risk that changes over time: Howard C. Mahler, "An Example of Credibility and Shifting Risk Parameters," *Proceedings of the Casualty Actuarial Society* 77 (1990): 225–308.

the buyer's costs of shopping and switching are thus absorbed in the fitted μ_ζ , and h is a fitted hold curve rather than the distribution of where the market stands. The insurer models ζ as normal, $\zeta \sim N(\mu_\zeta, \sigma^2)$, with a stated prior: $\mu_\zeta \sim N(0, 0.25^2)$, the market centered on the undeviated price with the spread of the filed band, and $\ln \sigma \sim N(\ln 0.10, 1.0^2)$. The posterior is taken at the mode of the decayed log likelihood plus the log prior, the lost quotes entering by their density and the bound quotes by the probability of the event, with weights that decline by 2 percent a quarter, and its covariance from the inverse of the Hessian of the negative log posterior at the mode, a Laplace approximation.

The decision then uses the hold probability averaged over that posterior,

$$\bar{h}(\ell) = \int \left[1 - \Phi\left(\frac{\ell - \mu_\zeta}{\sigma}\right) \right] \pi(\mu_\zeta, \sigma \mid \text{data}) d(\mu_\zeta, \sigma), \quad (\text{D.6})$$

exactly over μ_ζ given $\ln \sigma$ and by a 12-node Gauss–Hermite rule over $\ln \sigma$. For one period the average is what the objective implies under uncertainty about the curve's parameters: the increment G of Equation D.2 does not depend on them, so the expected increment of a deviation is the posterior-mean hold probability times G . The decision does not value the information a deviation adds about the curve in later quarters, and its deviations carry no rule that keeps them dispersed. In den Boer and Zwart's setting such a rule is an interval around the average of past prices, shrinking over time, that the price chosen on the current estimate may not enter, and under it the optimal price is eventually learned.³² Before any observation the prior predictive governs, so a quote can leave the undeviated price from the first quarter. The curve is fitted separately in each size band for renewals and for new business, six segments in all.

D.3 The Estimators

The estimator's observations are formed account-year by account-year. For an account-year s in a file let n_s be the reported claims with a positive value, q_s the expected reported claims at the class's frequency and the account's size, $\bar{y}_s$ the mean logarithm of the reported claim sizes, with $\sigma_c^2/2$ added for each open claim, and m_0 and v_0 the mean and variance of the logarithm of a claim on the bureau's curve at scale one before the occurrence limit, while the reported sizes are capped at it. The frequency observation is n_s/q_s , with noise variance $1/q_s$ around the year's relativity; the severity observation is $\bar{y}_s - m_0$ less the class's mean log scale, with noise variance $(n_s v_0 + n_s^o \sigma_c^2)/n_s^2$, n_s^o the open claims among them. A year without a reported claim carries no severity observation. The class's mean log scale is the average of $\bar{y}_s - m_0$, weighted by n_s , over the account-years of the class in every block the insurer has read, an account-year counting at each read.

Over every block of consecutive years the insurer has read, it accumulates the products of the frequency observations less one at each lag k from 1 through 20, weighted by the expected claims, and the products of the centered severity observations at each lag, a year's observation weighted by $n_s/(n_s + 5)$, and it accumulates the product of the two kinds within an account. To the frequency autocovariances it fits $A_f + B_f \rho_f^k$ by weighted least squares over a grid of ρ_f from 0.30 through 0.98 with A_f and B_f not below zero, and likewise $A_s + B_s \rho_s^k$ to the severity autocovariances. The cross moment is one constant, bounded by $\sqrt{A_f A_s}$ in absolute value. It refits after every quarter's reads.

³²Arnoud V. den Boer, "Dynamic Pricing and Learning: Historical Origins, Current Research, and New Directions," *Surveys in Operations Research and Management Science* 20, no. 1 (2015): 1–18; Arnoud V. den Boer and Bert Zwart, "Simultaneously Learning and Optimizing Using Controlled Variance Pricing," *Management Science* 60, no. 3 (2014): 770–783.

For the predictor, with an account's file of w years let o be the vector of its $2w$ observations, centered, Σ their covariance under the fitted moments (the fitted autocovariances between years, the cross moment between kinds, and each observation's noise variance on the diagonal), and c_f and c_s the covariances of the observations with the coming year's frequency relativity and log scale. The tracked relativity is $1 + c'_f \Sigma^{-1} o$, not below 0.05, and the tracked log scale is the class's mean plus $c'_s \Sigma^{-1} o$, plus the covariance of the two that remains given the observations, with variance $A_s + B_s - c'_s \Sigma^{-1} c_s$. The expected limited claim and its second moment in Equation D.5 are integrals of the curve's closed forms over a normal distribution of the log scale with that mean and variance, by Gauss–Hermite quadrature.

In the hold curve, for a segment, renewals or new business in a band, let quote j have log deviation ℓ_j from the undeviated price and weight ω_j . The log likelihood is

$$\sum_{j \text{ lost}} \omega_j \ln \phi\left(\frac{\zeta_j - \mu_\zeta}{\sigma}\right) - \sum_{j \text{ lost}} \omega_j \ln \sigma + \sum_{j \text{ bound}} \omega_j \ln \left[1 - \Phi\left(\frac{\ell_j - \mu_\zeta}{\sigma}\right)\right], \quad (\text{D.7})$$

with ζ_j the deviation ℓ_j less the broker's figure in logarithms and ϕ the standard normal density. With the log prior of §D.2.2 added, it is maximized in $(\mu_\zeta, \ln \sigma)$ by Newton's method, the covariance taken as §D.2.2 states. The pessimism variant of Section 4 lowers the conditional posterior mean of μ_ζ by one posterior standard deviation before the average of Equation D.6 is taken, which lowers the hold probability at every deviation.

D.4 The Variance Term

Without the variance term the Kelly route, in its construction of record, chooses another deviation on 0.028 of the quotes it computes over the record years, and the 90 percent interval of the mean of its growth less its growth with the term, paired by seed over 200 seeds, runs from -0.0059 to $+0.0051$ about a mean of 0.000 ($t = -0.13$), inside $(-0.010, +0.010)$. The fifth registration fixed a reading rule for this contrast before it was read: removing the term is shown not to move growth by more than 0.010 where the 90 percent interval lies inside that range; otherwise, where the interval excludes zero, the effect is reported at its sign; otherwise it is not shown. By the reading rule, removing the term is shown not to move the Kelly route's mean growth by more than 0.010. In the market of one tenth of the accounts the term is 0.6421 percent of the expected loss, 10.0 times its size in the record's market; the record does not run the Kelly route without it there.

E The Kelly Route's Computation, Step by Step

The steps are in the order in which they are carried out. Each names the data it reads; all of it is data the insurer holds or that the market treats as public (§D.2). Learning at the plan's pricing, the plan route, carries out the first three steps of a quote and quotes the price of the third, the first step after each quarter's placements, and the first and the third steps at the year's close.

At each quote, for each account:

1. Read the file. On an account the insurer holds: its own claim file for every year held, and the loss run received at acquisition for the five years before. On any other account: the broker's loss run of five years. Read the class and the bureau's claim frequency for it.

2. Track. Form the observations of §D.3 for each year of the file and compute the relativity, the log scale, and its variance from the fitted moments. Compute $\hat{x}_0$ and $\hat{v}_0$ by Equation D.5, and multiply both by the schedule modification.
3. Price through the plan. Multiply $\hat{x}_0$ by the insurer's rate level R , and $\hat{v}_0$ by R^2 , giving $\hat{x}$ and $\hat{v}$. The undeviated price is $R\hat{x}_0/(1 - \pi)$, on new business times the schedule credit of the account's size band, the schedule credit and the schedule modification together held within 1 ± 0.25 .
4. Deviate. Take the deviation ℓ , on a grid over the filed band and at zero, that maximizes $\bar{h}(\ell) G(P_\ell)$, with P_ℓ the undeviated price times e^ℓ , G of Equation 1 with the credit of the account's phase, renewal or new business, anchored at the undeviated price, and $\bar{h}$ the posterior-averaged hold probability of the account's segment (Equation D.6). Hold the deviation plus the logarithms of the schedule modification and the schedule credit within $\ln(1 \pm 0.25)$; the zero deviation is always admissible.
5. Apply the floor. Quote the larger of the deviated price and the floor of Equation D.4, with S the insurer's surplus at the quote.

After each quarter's placements:

1. Add every file and loss run read in the quarter to the sums of the second moments, and refit them.
2. For the segment's hold curve, record each quote that bound with its deviation, among them a renewal kept without going to market and one kept against a lower competing quote, and each quote lost on an account that went to market with its deviation and the broker's figure; decay the weights by 0.98 a quarter and refit the posterior of each segment.
3. For each renewal quoted, record whether it was a first renewal and whether it was kept; for each kept, the premium net of expenses and Equation D.2's margin, first renewals and later ones apart.

At each year's close:

1. Move the rate level by the loss-ratio method on the insurer's own experience, with competitors' winning prices as the complement of credibility (Equation A.2, §A.3).
2. Update the renewal credit. Over the past three years, take p_1 and p_s as the shares of first and of later renewals kept, and m_1 and m_s as the margins per unit of net premium on each; set β at $1/(1 + r)$ on the funds discount, or at $e^{-\hat{g}}$ on the growth discount, $\hat{g}$ the window's mean yearly log growth of surplus, the return on funds included; and set the horizon H of §D.1.3 and V_{ren} and V_{nb} by Equation D.3.
3. Update the schedule credit of each size band from the broker's figures on the insurer's own quotes on new business (Appendix B).

F The Registry of Predictions and Outcomes

Every prediction of the paper was written down before the runs that tested it, or, for the sixth registration, which analyzes runs already made, before any of its metrics was computed, under one rule: a predicted contrast holds where its paired mean over the 200 seeds has the predicted sign at $|t| > 2$, and a prediction that names another reading is read as it names. This appendix states the seven registrations in the order of their writing, each with its predictions on the fillings this paper reports, their outcomes under the rule, the order of writing against the order of running, and any reading of seeds outside the record that preceded it; the body points here for the outcomes it reports. The registrations also carried 101 predictions on an insurer that prices by the decision alone, with no rate level, on its own learned estimate and on the plan's, on the market around it, and on a program of variations of the market run on it, which this paper does not report: 91 of the first registration's, three of the second's, all four of the fourth's, and three of the fifth's. Of them 82 held, 18 failed, and one, the first registration's, was not tested.

The first registration made 97 predictions, registered after a reading of 20 seeds outside the record and before any arm ran on the record's seeds. The six reported here held: learning alone at the record's setting and at each of the four other settings of Section 4 (Table F.1), and, in the market with a second Loss-ratio insurer in the fourth seat, the buyer's calibration: on the record's seeds 0.141 of that market's accounts change insurer a year (standard error over seeds 0.0009), within two standard errors of the target of 0.14 (§A.4).

Table F.1: Learning's contrast at the five settings of Section 4 and the plan route's leads over the conventional insurers, paired by seed over 200 seeds: the mean on the logarithm (Equation A.3), its t statistic, and the seeds above zero. Learning alone is the plan route less a second Loss-ratio insurer, the fourth seat's growth in two arms run on the same seeds; the record's setting is the record's market, and each other setting moves one of two dimensions of it. Contrasts are taken on unrounded values and may differ by 0.001 from differences of the means Figure 2 states. Source: authors' calculations.

Contrast	Mean	t	Seeds above zero
<i>Learning alone, the fourth seat's growth in two arms, by setting</i>			
The record's setting	+0.400	+31.10	197
Refiling threshold 0.05	+0.394	+31.42	198
Refiling threshold 0.15	+0.391	+31.57	195
Read noise 0.35	+0.490	+36.08	199
Read noise 1.05	+0.384	+32.36	196
<i>Learning at the plan's pricing less a conventional insurer's</i>			
Bureau insurer	+0.541	+42.09	200
Loss-ratio insurer	+0.612	+48.62	198
GLM insurer	+0.519	+43.39	200

The second registration made 10 predictions, written after the first registration's outcomes were on the record. The seven reported here were written after a reading of seeds outside the record: one, that its arms carried over from the first registration, the plan route's among them, reproduce the first registration's runs, run for run; and six on four directions first read on those seeds, the anchoring on learning at the Kelly decision, its three leads over the conventional insurers, the pessimism variant on it, and learning at the Kelly decision with pessimism against the pricing it contains, the

primary prediction. The reading ran six constructions of learning at the Kelly decision: the construction as first built, on each discount of the credit, and, on the funds discount, that construction without the credit, with the anchored credit, with the pessimism variant, and with both. Of the seven, five held and two failed: the pessimism variant on learning at the Kelly decision and the primary prediction.

The third registration made two predictions for learning at the Kelly decision on the growth discount, written before any run of the arms it tests, and both failed: the discount's effect on learning at the Kelly decision and the primary prediction evaluated on that discount. The fourth registration's four predictions are among those the count above names.

The fifth registration ran the four markets beside the record and made six predictions for two of them, the market of smaller insurers and the market with a second insurer that learns, written after a reading of 20 seeds outside the record and before the arms ran on the record's seeds. The three reported here held: learning against a second Loss-ratio insurer in each of the two markets, and, in the market with a second insurer that learns, that insurer's leads where the fourth seat is a second Loss-ratio insurer (Appendix G). It made no prediction for the Kelly route, which it ran in the market with a second insurer that learns and under the severity trend, and for the Kelly route without the variance term it states a reading rule in place of predictions (§D.4).

The sixth registration is an analysis on the records in hand, with no run of the market: the equivalence margin of §5.1 and the path metrics of §5.2, with six predictions written before any metric was computed. The primary contrast's value was on the record and is stated in the registration as known; the margin is 0.028 in the logarithm over the 28 record years, 2.8 percent of terminal surplus and a tenth of a percentage point of the mean geometric rate of growth a year, chosen in the objective's own unit and set below the smallest effects on which registered findings of the record rest, the anchoring's worth on the insurer that prices by the decision alone, which the paper does not report: 0.054 in the logarithm on the growth discount and 0.033 on the funds discount; a margin of a tenth of learning's worth would have placed the second inside it. The first prediction is the equivalence reading itself at that margin; the second and third are read at $t > 2$, the paired within-path volatility and the Pitman–Morgan t of the dispersion across seeds; the fourth through sixth are read on medians and shares of the paired path, without a t . The equivalence was shown; of the five predictions that name a direction, four held and one failed, that on the within-path volatility, which the Kelly route has below the plan route's where above was predicted.

The seventh registration ran the two routes at other values of the filed band and the buyer's scale (Section 6), with six predictions written before its reading of 20 seeds outside the record, which left them unchanged, and before any arm ran on the record's seeds: the first, that its runs reproduce the record's arms at the record's values run for run, read on 60 runs outside the record and held there; the second, a loss for the Kelly route against the plan route at the wider band; the third and fourth, that the buyer's scale moves the Kelly route's mean renewal deviation and its band-edge share below their values at the record's scale at the lower scale and above them at the higher; and the fifth and sixth, that the routes' contrast less the same contrast at the record's scale is below zero at the lower scale and above it at the higher. All six held on the record; at the reading of 20 seeds outside the record one of them, the sixth, had failed. The registration also ran learning at the Kelly decision in the market of smaller insurers and the market of one tenth of the accounts, without a prediction.

Table F.2: The contrasts among the fillings of Table 1 that the second and third registrations read, the fourth seat's growth paired by seed over 200 seeds: the mean on the logarithm (Equation A.3), its t statistic, and the seeds above zero. The anchoring is an arm with the anchored credit less the same arm with the quote-anchored credit, the pessimism variant's contrast an arm with the variant less the same arm without it, and the growth discount's contrast an arm on the growth discount less the same arm on the funds discount. Above the subheading are the second registration's contrasts, on the funds discount: its predictions are on the anchoring, the pessimism variant, and learning at the Kelly decision with pessimism less learning at the plan's pricing, and the block's two other rows carry none. Below it are the third registration's, on the growth discount: its two predictions are on the growth discount on learning at the Kelly decision and on learning at the Kelly decision with pessimism less learning at the plan's pricing. Under the registered rule a predicted contrast holds where its paired mean has the predicted sign at $|t| > 2$; contrasts are taken on unrounded values. Source: authors' calculations.

Contrast, the fourth seat's growth	Mean	t	Seeds above zero
The anchoring, on learning at the Kelly decision	+0.110	+14.85	177
Pessimism, on learning at the Kelly decision	0.000	+0.03	100
<i>Less learning at the plan's pricing</i>			
Learning at the Kelly decision with pessimism	+0.011	+1.97	113
Learning at the Kelly decision	+0.011	+1.92	118
Learning at the Kelly decision, quote-anchored credit	-0.099	-12.01	34
<i>On the growth discount</i>			
The growth discount, on learning at the Kelly decision	-0.006	-1.36	86
The growth discount, on the same with pessimism	-0.005	-0.96	99
Pessimism, on learning at the Kelly decision	+0.002	+0.42	98
<i>Less learning at the plan's pricing</i>			
Learning at the Kelly decision with pessimism	+0.006	+1.03	119
Learning at the Kelly decision	+0.004	+0.71	113

The rule is applied to each prediction without adjustment for the number of predictions tested. A Bonferroni level divides the rule's two-sided 5 percent among a family of predictions read by a t , which is every prediction except those read otherwise: the first registration's calibration and its prediction not tested, the second's and the seventh's predictions that arms reproduce arms, and the sixth's equivalence reading and its three predictions on medians and shares. A prediction that rests on three leads is one prediction of its family, and it holds at an adjusted level where each of its leads clears it. Over the 123 predictions of the seven registrations read by a t , at $|t| > 3.60$, every prediction reported here that held still holds.

Table F.3: The two routes' difference at the close of the record years and the equivalence test of the sixth registration: learning at the Kelly decision less learning at the plan's pricing, the fourth seat's growth paired by seed over 200 seeds, on each discount of the credit with and without the pessimism variant. The mean is given on the logarithm (Equation A.3) and as the difference in terminal surplus in percent, the exponential of the mean less one; the standard error, the t statistic, and the 90 percent interval are on the logarithm, the interval the mean plus and minus 1.645 standard errors. Equivalence is shown at the registered margin m where the interval lies inside $(-m, +m)$; the margin was registered with the first row's value on the record (the sixth registration, above). Source: authors' calculations.

Less learning at the plan's pricing	Mean, logarithm	Percent of terminal surplus	s.e.	t	90 percent interval	At $m =$ 0.028
<i>Learning at the Kelly decision</i>						
With pessimism, the funds discount	+0.0108	+1.1	0.0055	+1.97	[+0.0018, +0.0198]	shown
The funds discount	+0.0107	+1.1	0.0055	+1.92	[+0.0015, +0.0198]	shown
With pessimism, the growth discount	+0.0062	+0.6	0.0061	+1.03	[-0.0037, +0.0162]	shown
The growth discount	+0.0043	+0.4	0.0061	+0.71	[-0.0057, +0.0144]	shown

G The Quote Data and the Further Markets

Table G.1 gives the quotes of learning at the Kelly decision by arm and phase (§5.2 reads its renewal share at the band's upper end); Table G.2 gives the quote data of learning at the Kelly decision with pessimism at each setting of the seventh registration (Section 6); and Table G.3 gives learning's contrast in the four further markets of the fifth registration (Section 6), which the paragraph before it reads.

Table G.1: The quotes of learning at the Kelly decision over the record years and the 200 seeds, by arm (above the subheading, on the funds discount, the arm with the quote-anchored credit, with the anchored credit, and with the anchored credit and the pessimism variant, the arms of Table 1) and phase: the mean deviation from the plan route's price (a logarithm); the share of quotes within one grid step (0.0025 in log price) of the filed band's upper end, the share at exactly the zero deviation, and the share raised to the floor; for the arms with the anchored credit, the posterior-mean hold probability at the chosen deviation (predicted hold); and the share of the renewal quotes that held the account (realized hold). On new business the predicted hold is over every account the arm priced and the record's share held is over the quotes it made, so no realized hold is set beside it. Below the subheading, the third registration's arms with the anchored credit, on the growth discount. Source: authors' calculations.

Arm	Phase	Mean deviation	At the band's upper end	At zero	Raised to the floor	Predicted hold	Realized hold
Quote-anchored	Renewals	+0.174	0.288	0.127	0.000	—	0.753
	New business	+0.021	0.105	0.169	0.000	—	—
Anchored	Renewals	+0.120	0.278	0.166	0.000	0.775	0.809
	New business	−0.081	0.039	0.136	0.000	0.708	—
Anchored, with pessimism	Renewals	+0.118	0.276	0.168	0.000	0.776	0.811
	New business	−0.088	0.037	0.133	0.000	0.714	—
<i>On the growth discount</i>							
Anchored	Renewals	+0.135	0.288	0.150	0.000	0.760	0.793
	New business	−0.057	0.039	0.146	0.000	0.680	—
Anchored, with pessimism	Renewals	+0.133	0.288	0.152	0.000	0.761	0.795
	New business	−0.063	0.037	0.145	0.000	0.686	—

Table G.2: The quote data of learning at the Kelly decision with pessimism at each setting of the seventh registration, each seed's mean over its record years and the mean of those over the 200 seeds (Table G.1 pools the quotes): the mean renewal deviation from the plan route's price (a logarithm), the share of renewal quotes within one grid step (0.0025 in log price) of the filed band's upper end, and the mean new-business deviation; and, for the mean renewal deviation and the band-edge share, the movement from the record's setting paired by seed, with its t statistic. Source: authors' calculations.

Setting	Mean renewal deviation	Share at the band's upper end	Mean new-business deviation	Movement from the record's setting (t)	
				Renewal deviation	Band-edge share
The record's setting	+0.119	0.275	−0.093	—	—
The band of 40 percent either way	+0.117	0.125	−0.023	−0.0013 (−3.48)	−0.1500 (−123.24)
The buyer's scale of 0.025	+0.113	0.254	−0.036	−0.0059 (−21.33)	−0.0209 (−15.74)
The buyer's scale of 0.075	+0.123	0.285	−0.113	+0.0041 (+15.56)	+0.0109 (+11.95)

In each of the four markets learning at the plan's pricing outgrows a second Loss-ratio insurer in the fourth seat at $|t| > 2$ (Table G.3); in the market of smaller insurers and beside a second insurer that learns this was a registered prediction. In the market with a second insurer that learns and a second Loss-ratio insurer in the fourth seat, the second insurer that learns leads the Bureau insurer, the Loss-ratio insurer, and the fourth seat, as predicted (+0.505, +0.624, and +0.612; t of +35.30, +44.34, and +48.19); in its arm whose fourth seat is learning at the plan's pricing, the two seats that learn differ in growth by a difference the record does not resolve (the fourth seat less the second, −0.019, $t = -0.79$). Table 3 gives the Kelly route against the plan route in the four markets,

without a prediction; no contrast of the four markets was predicted but those this paragraph names as predicted.

Table G.3: The four markets of the fifth registration: E, every insurer at an endowment at which a seat writing the opening book stands at the capital ratio at which broker flight begins; C, one tenth of the accounts; L, the GLM insurer's seat a second insurer that learns; T, a trend of 6.8 percent a year in the mean claim limited per occurrence. Learning at the plan's pricing less a second Loss-ratio insurer in the fourth seat, the fourth seat's growth (Equation A.3) paired by seed over 200 seeds, each mean on the logarithm above its t statistic in parentheses; Table 3 gives the Kelly route against the plan route in each. The variance term at each market's mean endowment is in percent of the expected loss of an account at the curve's own scale of claim sizes. Source: authors' calculations.

The fourth seat's growth	E	C	L	T
Learning at the plan's pricing less a second Loss-ratio insurer	+0.451 (+26.37)	+0.241 (+12.01)	+0.340 (+23.11)	+0.163 (+7.39)
Variance term, percent of expected loss	0.1372	0.6421	0.0641	0.0644

Bibliography

- Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu. "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market." *Journal of Political Economy* 132, no. 3 (2024): 723–771.
- Association for Cooperative Operations Research and Development. *Commercial Insurance Application*. ACORD 125, 2007/10.
- Bailey, Robert A. "Insurance Rates with Minimum Bias." *Proceedings of the Casualty Actuarial Society* 50 (1963): 4–11.
- Breiman, Leo. "Optimal Gambling Systems for Favorable Games." In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1, edited by Jerzy Neyman, 65–78. Berkeley: University of California Press, 1961.
- Brown, Zach Y., and Alexander MacKay. "Competition in Pricing Algorithms." *American Economic Journal: Microeconomics* 15, no. 2 (2023): 109–156.
- Bühlmann, Hans. "Experience Rating and Credibility." *ASTIN Bulletin* 4, no. 3 (1967): 199–207.
- Bühlmann, Hans. *Mathematical Methods in Risk Theory*. Berlin: Springer-Verlag, 1970.
- Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello. "Artificial Intelligence, Algorithmic Pricing, and Collusion." *American Economic Review* 110, no. 10 (2020): 3267–3297.
- Capen, Edward C., Robert V. Clapp, and William M. Campbell. "Competitive Bidding in High-Risk Situations." *Journal of Petroleum Technology* 23, no. 6 (1971): 641–653.
- Cummins, J. David. "Allocation of Capital in the Insurance Industry." *Risk Management and Insurance Review* 3, no. 1 (2000): 7–27.
- D'Arcy, Stephen P., and Neil A. Doherty. "Adverse Selection, Private Information, and Lowballing in Insurance Markets." *Journal of Business* 63, no. 2 (1990): 145–164.
- D'Arcy, Stephen P., and Neil A. Doherty. "The Aging Phenomenon and Insurance Prices." *Proceedings of the Casualty Actuarial Society* 76 (1989): 24–44.
- de Garidel-Thoron, Thomas. "Welfare-Improving Asymmetric Information in Dynamic Insurance Markets." *Journal of Political Economy* 113, no. 1 (2005): 121–150.
- den Boer, Arnoud V. "Dynamic Pricing and Learning: Historical Origins, Current Research, and New Directions." *Surveys in Operations Research and Management Science* 20, no. 1 (2015): 1–18.
- den Boer, Arnoud V., and Bert Zwart. "Simultaneously Learning and Optimizing Using Controlled Variance Pricing." *Management Science* 60, no. 3 (2014): 770–783.
- Fisher, Ginda Kaplan, Lawrence McTaggart, Jill Petker, and Rebecca Pettingell. *Individual Risk Rating*. Casualty Actuarial Society study note, 2019.
- Greenwich Associates. *Insurance Market Pulse: Looking at the Small Business and Middle Market Segments*. 2007. As reported in Douglas McLeod, "Changing Insurers Yields Surprise," *Business Insurance*, June 17, 2007.
- The Hartford Insurance Group, Inc. *Annual Report on Form 10-K for the Fiscal Year Ended December 31, 2025*. Filed with the U.S. Securities and Exchange Commission, February 20, 2026.
- Hays, Lisa A. "Managing Commercial Lines Pricing Levels in a Loss Cost Environment." *Casualty Actuarial Society Forum* (Winter 2002): 9–30.
- Hornik, Kurt, Maxwell Stinchcombe, and Halbert White. "Multilayer Feedforward Networks Are Universal Approximators." *Neural Networks* 2, no. 5 (1989): 359–366.
- Jia, Ruo, and Zenan Wu. "Insurer Commitment and Dynamic Pricing Pattern." *Geneva Risk and Insurance Review* 44, no. 1 (2019): 87–135.
- Keatinge, Clive L. "Modeling Losses with the Mixed Exponential Distribution." *Proceedings of the Casualty Actuarial Society* 86 (1999): 654–698.
- Kelly, John L., Jr. "A New Interpretation of Information Rate." *Bell System Technical Journal* 35, no. 4 (1956): 917–926.
- Kunreuther, Howard, and Mark Pauly. "Market Equilibrium with Private Knowledge: An Insurance Example." *Journal of Public Economics* 26, no. 3 (1985): 269–288.

- Longley-Cook, Laurence H. "An Introduction to Credibility Theory." *Proceedings of the Casualty Actuarial Society* 49 (1962): 194–221.
- Lynch, Jim, and William Nibbelin. *Increasing Inflation on Liability Insurance – Impact as of Year-End 2024*. Insurance Information Institute and Casualty Actuarial Society, October 2025.
- MacLean, Leonard C., Edward O. Thorp, and William T. Ziemba, eds. *The Kelly Capital Growth Investment Criterion: Theory and Practice*. Singapore: World Scientific, 2011.
- MacLean, L. C., W. T. Ziemba, and G. Blazenko. "Growth versus Security in Dynamic Investment Analysis." *Management Science* 38, no. 11 (1992): 1562–1585.
- Mahler, Howard C. "An Example of Credibility and Shifting Risk Parameters." *Proceedings of the Casualty Actuarial Society* 77 (1990): 225–308.
- Milgrom, Paul R., and Robert J. Weber. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50, no. 5 (1982): 1089–1122.
- Morgan, W. A. "A Test for the Significance of the Difference between the Two Variances in a Sample from a Normal Bivariate Population." *Biometrika* 31, no. 1/2 (1939): 13–19.
- Myers, Stewart C., and Richard A. Cohn. "A Discounted Cash Flow Approach to Property-Liability Insurance Rate Regulation." In *Fair Rate of Return in Property-Liability Insurance*, edited by J. David Cummins and Scott E. Harrington, 55–78. Boston: Kluwer-Nijhoff, 1987.
- National Association of Insurance Commissioners. "Aggregated P&C RBC Data: 2025 Data as of June 1, 2026." *Summary: Aggregate P/C RBC Results by Year*. NAIC, 2026.
- National Association of Insurance Commissioners. "NAIC IRIS Ratios," data year 2025, property and casualty. NAIC, May 19, 2026.
- National Association of Insurance Commissioners. *Report on Profitability by Line by State in 2024*. Kansas City, Missouri: NAIC, 2026.
- National Association of Insurance Commissioners. *Risk-Based Capital (RBC) for Insurers Model Act*. Model 312. 2012.
- National Association of Insurance Commissioners. *U.S. Property & Casualty and Title Insurance Industries—2025 Full Year Results*. NAIC Financial Regulatory Services, 2026.
- National Association of Insurance Commissioners, Casualty Actuarial and Statistical (C) Task Force. *Price Optimization White Paper*. Adopted by the task force, November 19, 2015.
- Nilssen, Tore. "Consumer Lock-in with Asymmetric Information." *International Journal of Industrial Organization* 18, no. 4 (2000): 641–666.
- Palmer, Joseph M. *Increased Limits Ratemaking for Liability Insurance*. Casualty Actuarial Society study note, 2006.
- Pitman, E. J. G. "A Note on Normal Correlation." *Biometrika* 31, no. 1/2 (1939): 9–12.
- Samuelson, Paul A. "Why We Should Not Make Mean Log of Wealth Big Though Years to Act Are Long." *Journal of Banking and Finance* 3, no. 4 (1979): 305–307.
- Thorp, Edward O. "The Kelly Criterion in Blackjack Sports Betting, and the Stock Market." In *Handbook of Asset and Liability Management*, vol. 1, edited by Stavros A. Zenios and William T. Ziemba, 385–428. Amsterdam: North-Holland, 2006.
- The Travelers Companies, Inc. "Travelers Reports Excellent Second Quarter and Year-to-Date Results." Press release, July 17, 2026.
- Verisk. *2024 Executive Insights: General Liability*. 1st ed. Verisk's Core Lines Services Report, Second-Quarter 2019 to First-Quarter 2024. Verisk, 2024.
- Weiss, Mary A. "Underwriting Cycles: A Synthesis and Further Directions." *Journal of Insurance Issues* 30, no. 1 (2007): 31–45.
- Werner, Geoff, and Claudine Modlin. *Basic Ratemaking*. 5th ed. [Arlington, VA]: Casualty Actuarial Society, 2016.
- Wu, Cheng-Sheng Peter, and Hua Lin. "Large Scale Analysis of Persistency and Renewal Discounts for Property and Casualty Insurance." *Casualty Actuarial Society E-Forum* (Winter 2009): 396–408.