

Algorithmic Learning in Competitive Insurance Underwriting

Frank Schmid and Can Wang

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

This version: 19 October 2026

First version: 17 March 2026

Abstract. An insurer that competes for each account does not know the account's loss distribution, and the price that maximizes the growth of its surplus depends on that distribution and on the chance of winning the account at the price. This paper places insurers that learn and price in different ways in a stylized market for middle-market general liability. Three price by a rating plan: one on the bureau's plan with a multiplier it refiles, one that moves its rate level by the loss-ratio method, and one that fits a generalized linear model to its own book. A fourth seat is filled in four ways, the plan's estimate or a learned estimate of each account's loss distribution, priced as a plan prices or by the Kelly decision, so that the learning and the decision are measured apart and together. Learning raises growth: the seat that tracks each account's distribution and otherwise prices as the insurer that uses the loss-ratio method does, the learned-estimate seat, outgrows every seat that does not. The Kelly decision lowers growth on the learned estimate, against the registered prediction, where the value of keeping an account enters it in proportion to the quoted premium; that form rewards a markup with continuation value the markup does not buy. With the value anchored at the undeviated premium, the Kelly underwriter's growth rises at either discount of the value, the insurer's own growth of surplus as registered and the return on funds. Applied as a deviation from the learned-estimate seat's own price, a deviation policy that contains that price as one of its choices, the corrected decision no longer trails that seat. Applied alone, by the Kelly underwriter, the seat that prices its learned estimate by the Kelly decision, it still trails the seat at either discount, by a gap the study's runs do not attribute among the underwriter's parts. The growth of a plan's pricing of a learned estimate is not distinguished, at the registered threshold, from that of the corrected Kelly decision applied to the same information. The two are told apart in the quote data, where the deviation policy prices renewals toward the upper end of the filed rate band and new business below the learned-estimate seat's price, and in the market around them, where the deviation policy holds fewer accounts and accounts change insurer more often. Around the Kelly underwriter the market changes: accounts move more often, every insurer's load on new business falls with account size, and the traditional insurers' rate levels end higher. The Kelly underwriter's leads over the traditional insurers hold, with one exception, across a program that moves the market's conditions and its fixed constants one at a time. The pattern of the fourth seat's four arms holds at the five settings at which they were run; the deviation policy was run at one calibration.

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	iii
1 Introduction	1
2 A Stylized Market	4
2.1 Accounts and Losses	4
2.2 What an Insurer Knows	5
2.3 Three Traditional Insurers	5
2.4 The Buyer and the Broker	7
2.5 Capital, Solvency, and the Measure	8
3 The Kelly Decision	11
3.1 The Objective	11
3.2 The Increment of One Account	11
3.3 What Keeping an Account Is Worth	12
3.4 The Chance of Winning, and the Decision	13
4 Learning	13
4.1 The Account's Loss Distribution	13
4.2 The Business Won	14
4.3 The Demand Model	15
4.4 What Was Learned	15
5 The Two Together	16
6 The Credit's Form, and a Second Construction	18
6.1 The Diagnosis and the Correction	19
6.2 The Deviation Policy, and Containment	19
6.3 Outcomes	21
6.4 Growth Not Distinguished, and Where the Two Differ	24
7 What Each Part Is Worth	25
8 The Market with a Kelly Underwriter	27
8.1 The Price of New Business	27
8.2 Turnover, Shares, and Rate Levels	29
9 Sensitivity	31

10 Discussion and Conclusion	34
10.1 Limits	34
10.2 Two Other Markets	35
10.3 Conclusion	35
A The Loss Process	37
B The Rating Plans	37
C The Capital Standard	38
D The Estimators	39
E The Kelly Underwriter’s Computation, Step by Step	40
F The Arms of the Record	42
G The Quote Data and the Loads by Band	45
Bibliography	47

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 (Appendix A)
β	beta	Discount factor of the renewal credit: at the insurer's own growth of surplus (§3.3) or at the return on funds (§6.2)
δ	delta	Deviation of the second construction's policy from the price it contains (§6.2)
Δ	Delta	Logarithm of the holder's quote over the lowest invited quote (§2.4)
ε	epsilon	Expenses as a share of premium (§2.3)
ζ	zeta	Latent log price of the demand model, over the price its deviation is taken from (§4.3)
η	eta	Yearly shock to the drifting part of a trait (Appendix A)
κ	kappa	A cell's factor on the tracked expected loss (§4.2); the capital requirement's scale (Appendix C)
λ	lambda	An account's expected number of claims in a policy year (§2.1); $\hat{\lambda}$, the tracked claim rate (Equation 8); λ_0 , the claims a year at the median size (Appendix A)
μ_ζ	mu	Mean of the demand model's ζ (§4.3)
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 (Appendix A)
π	pi	Provision for expenses and profit in a plan's premium (§2.3); the posterior density of the demand model's parameters (Equation 10)
ρ	rho	Persistence of the drifting part of a trait (§2.1); the capital requirement's reserve-charge ratio (Appendix C); with subscript f or s , the fitted persistence of an autocovariance (Appendix D)
σ	sigma	Spread of the demand model's ζ (§4.3); with subscript c , of a case reserve's error (Appendix A)
τ	tau	Mean time from booking to payment, years (§3.2)
Σ	Sigma	Covariance of a file's observations (Appendix D); the summation sign (Equation 6, Appendix B)
ϕ	phi	Standard normal density (Appendix D)
Φ	Phi	Standard normal distribution function (§4.3)
ψ	psi	Complement of credibility in the loss-ratio method's indication (§2.3)
ω	omega	Weight of a quote in the demand model's likelihood (Appendix D)

Latin Symbols

- A* An account's capped losses on the file it is rated on (§2.3); 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 (Appendix D)
- a* Reported losses, in the loss-ratio method's indication (Equation 2) and of a cell, a_c (Equation 9); the loss adjustment expense's share (Equation 8)
- 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 (Appendix D)
- c* An account's class (§2.3); as a subscript, a cell (§4.2) or a case reserve (σ_c); c_x , the return on the premium held (§3.2); c_f and c_s , covariances of the observations (Appendix D)
- d* Base of the weights of the experience modification (§2.3); with subscript *it*, the drifting part of a trait (Appendix A)
- E* A class's expected capped losses, E_{rep} their expected reported part (§2.3); expected reported losses, in the GLM insurer's fit (Appendix B); $\mathbb{E}$, expectation (§3.1)
- e* Expected reported losses, in the loss-ratio method's indication (Equation 2) and of a cell, e_c (Equation 9)
- F* With subscript *j*, the public expected reported share of a year (Appendix B)
- f* The indicated change of level (Equation 2); the GLM insurer's change of level or a fitted factor (Appendix B)
- G* The increment of the decision (Equation 5)
- g* An insurer's growth (Equation 3); $\hat{g}$, its mean yearly log growth of surplus over the credit's window (§3.3)
- H* The credit's horizon, years (§3.3)
- h* The hold curve, $\bar{h}$ its posterior average (§3.4, Equation 10)
- i* An account (Appendix A)
- j* The age of a year, in the experience modification (§2.3); a quote, in the demand model's likelihood (Appendix D)
- K* The credibility constant of the experience modification (§2.3)
- k* The buyer's cost of shopping, k' its cost of switching (§2.4); years ahead (§3.3); a lag (Appendix D)
- L* With subscript *c*, the bureau's loss cost (Equation 1)
- ℓ* Log deviation from the working premium (§3.4)
- M* M_{exp} and M_{sch} , the experience and the schedule modification (Equation 1)
- m* With subscript 1 or *s*, a margin per unit of net premium (§3.3); m_0 , the mean log claim on the bureau's curve (Appendix D)

N	The normal distribution (§4.3)
n	Premium net of expenses (§3.2); reported claims, in the loss-ratio method's indication (Equation 2) and in the tracker's observations, n_s^o the open claims among them (Appendix D)
o	A file's centered observations (Appendix D)
P	The premium quoted (§3.2)
p	With subscript 1 or s , the retention at the first and at later renewals (§3.3)
Q	Reserves (Appendix C)
q	Expected reported claims, in the tracker's observations (Appendix D)
R	The insurer's rate level (Equation 1)
r	The return on funds (§2.5)
S	Surplus (§3.1)
s	The scale of the buyer's response (§2.4); as a subscript, later renewals (§3.3), the severity trait, or an account-year (Appendix D); s_1^2 and s_2^2 , variances in the schedule modification (Appendix B)
T	The number of record years (Equation 3)
t	A year (§3.3, Appendix A); the t statistic of a contrast (§2.5)
u	The occurrence limit (Equation 8); with subscript it , the frequency trait (Appendix A)
V	The renewal credit, V_{nb} and V_{ren} (Equation 6)
v	$\hat{v}$, the estimated variance of an account's loss (§3.2); v_0 , the variance of a log claim on the bureau's curve (Appendix D)
W	Written premium in the year (Appendix C)
w	With subscript j , the weight of a year in the experience modification (Appendix B); a file's length in years (Appendix D)
X	An account's loss (§3.2)
x	$\bar{x}$, the plan's expected loss (Equation 1); $\hat{x}$, the estimated expected loss, $\hat{x}_0$ before a cell's factor (§3.2, §4.1)
Y	Claim size at the tracked scale (Equation 8); $\bar{y}$, a mean log claim size (Appendix D)
Z	Classical credibility (Equation 2, Appendix B), of a cell Z_c (Equation 9); Z_{exp} , the experience modification's credibility (§2.3)
z	An account's size (§2.3)

1 Introduction

An insurer that competes for an account has two things to decide and does not know what it needs for either. It decides at what price to quote, and thereby whether it wants the account at the price the market will bear. For the first it needs the account's loss distribution, which no one observes; for the second it needs the chance of winning at each price, which depends on what its competitors quote and on how the buyer chooses. The actuarial answer is a 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 an objective that the price serves.

This paper studies an insurer whose price follows from an objective, the growth of its surplus, and that learns from its own operations what the objective requires. The objective is the Kelly criterion, the maximization of the expected logarithm of wealth, which, in repeated independent bets on known odds, attains a long-run growth rate that no other strategy exceeds.² For one account the criterion gives a price from three quantities: the account's expected loss, the variance of its loss against the insurer's surplus, and the chance that the quote holds the account. Under competition it calls for a fourth, what keeping the account in later years is worth. A rating plan supplies a version of the first and none of the other three, and the insurer of this paper, the *Kelly underwriter*, estimates each: the expected loss and the variance by tracking every account's claim frequency and claim size from the files it reads, with a correction by cell for the business it wins (§4.2); the value of keeping an account from its own retention, margin, and growth of surplus; and the chance of winning from the figure a broker gives an insurer whose quote did not bind.

The two subjects are separable, and the paper separates them. An insurer can learn an account's loss distribution and price it as a plan prices; an insurer can price by the Kelly decision on the estimate a plan supplies. We place the Kelly underwriter in a stylized market for middle-market general liability (GL) beside three traditional insurers, each named by how it sets its price: a *Bureau insurer*, a *Loss-ratio insurer*, and a *GLM insurer*, which fits a generalized linear model (GLM) to its own book. The fourth seat of the market is then filled in four ways, the plan's estimate or the learned estimate, priced as the Loss-ratio insurer prices or by the Kelly decision, with the three traditional insurers held as they are. The market is stylized, and no observed market has a Kelly underwriter, so the outcomes are findings about the inputs and are not offered as a reproduction of a market; the inputs follow practice where practice is documented, and each is listed with its grounding.

Five findings are on the record of 200 seeds, the same seeds in every *arm*, an arm being one configuration of the market and a seed one draw of its random streams. The measure is an insurer's growth, the logarithm of its surplus at the end of the 28 *record years* over its endowment, net of the return on funds.

First, learning raises growth. In the fourth seat a second Loss-ratio insurer grows by +0.534. Learning alone, the learned estimate priced as that insurer prices (the *learned-estimate seat*), adds

¹Geoff Werner and Claudine Modlin, *Basic Ratemaking*, 5th ed. ([Arlington, VA]: Casualty Actuarial Society, 2016), chaps. 8, 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.

+0.400 ($t = +31.10$). The Kelly decision with the quote-anchored credit, the credit entering in proportion to the premium quoted (§6.1), does not: on the plan's estimate it moves growth by -0.003 ($t = -0.24$), and on the learned estimate by -0.244 ($t = -31.11$). The first registration, of the predictions written before the record of the fourth seat's four arms (the *two by two*) ran (Appendix F), predicted a loss at $t < -2$ from the decision on the plan's estimate, a gain at $t > 2$ on the learned estimate, and an interaction above zero at $t > 2$. All three failed.

Second, part of the decision's shortfall lies in the form of one term, and with the form corrected, the decision applied as a deviation from the learned-estimate seat's price no longer trails that seat. In the decision with the quote-anchored credit, the value of keeping an account enters the quote objective in proportion to the premium of the deviation under consideration, rewarding a markup with continuation value the markup does not buy, where by the market's own dynamics the continuation of a won account is invariant to today's deviation to first order. A second construction anchors the term at the undeviated premium and applies the corrected Kelly decision as a deviation from that seat's own price, a deviation policy that contains the seat's pricing as the zero deviation. On the Kelly underwriter as registered, which discounts the credit at its own growth, the correction is worth +0.054 ($t = +11.03$), and corrected, that insurer still trails the learned-estimate seat by 0.190 ($t = -30.88$). Of the second registration's eight arms, every one that prices by the Kelly objective, except the one that reproduces the registered Kelly underwriter, discounts the credit at the return on funds; on that discount the correction is worth +0.110 ($t = +14.85$) on the deviation policy and +0.033 ($t = +7.09$) on the Kelly underwriter. The pessimism variant lowers every hold probability, the probability that a quote holds the account (§6.2). The corrected policy with that variant is the arm of the second registration's primary prediction, and it ends +0.0108 above the seat it deviates from (standard error 0.0055). The registered threshold holds a contrast at a paired mean of the predicted sign at $|t| > 2$. At that threshold the two behaviors are not distinguished in growth, a statement at this record's precision and not a proof of zero. They are told apart in the quote data: the policy prices renewals toward the upper end of the filed band, 0.276 of its renewal quotes within one grid step of it, 0.0025 in log price, and new business below the learned-estimate seat's price. They are also told apart in the market around them: the policy holds 0.494 of the accounts at the end of the record years against the seat's 0.672.

Third, the Kelly underwriter outgrows the three traditional insurers. In the base arm, the market with the Kelly underwriter as registered in the fourth seat, its leads over the Bureau insurer, the Loss-ratio insurer, and the GLM insurer are +0.353, +0.430, and +0.323 (t of 22.73, 30.25, and 22.35), and no insurer fails in any seed. Its architecture nonetheless trails the learned-estimate seat, by 0.244 in the base arm (Table 4), and by 0.190 with the credit corrected (Table 6). The two run the same estimator of each account's distribution on different books and files, the record does not attribute the gap among the Kelly underwriter's parts (§6.3), and Section 7 removes the parts of the Kelly underwriter of record one at a time.

Fourth, the market changes around the Kelly underwriter. The Kelly underwriter quotes new business below its own renewal price for the same estimate, by more the larger the account; the traditional insurers, which file no load on new business, end up pricing new business below renewals through their schedule credits, by 13.5 percentage points more in the largest third of accounts than in the smallest, against 2.7 percentage points where the fourth seat is one of their own kind. Accounts change insurer at 0.394 a year against 0.141, and the traditional insurers' rate levels end higher and their growth lower than in the market without it.

Fifth, the Kelly underwriter's leads, with one exception, hold across the sensitivity program (Section 9); the second construction was run at one calibration. Each of the three leads is above zero

at $t > 2$ in 19 arms that change one condition, rule, or threshold of the market, at two other levels of the noise in an underwriter's read of an account, and in 54 arms that move one fixed constant or choice, most to about half and to about one and a half of its value; the lowest t among them is 5.80. The program's four other arms change the market in another way, two other insurers in the Bureau insurer's seat and the two changed inputs of Section 7, and among them is the program's one lead not above zero, the lead over a Bureau insurer that holds a fixed level, which the paper reports with the market it arises in (§10.2). The two by two keeps its pattern at every one of the five settings at which it was run.

The paper draws on three literatures. The first is the Kelly criterion of Kelly and Breiman, whose second-order form is a mean–variance rule with the weight on variance fixed by the criterion.³ An insurer does not know the odds, and the paper's findings are statements about that case: learning the odds raises growth, the corrected Kelly decision applied to a plan's pricing of what is learned adds an amount the record does not distinguish from zero at the registered threshold, and the Kelly underwriter's own architecture trails that pricing. 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 without commitment the equilibrium has a low first price recovered on the accounts that stay.⁴ In our market the holder reads its own claim file where a competitor reads a broker's loss run, every contract is annual, and the Kelly underwriter, which is given no model of an equilibrium, arrives at new-business prices below its renewal prices by valuing retention from its own book and by correcting new business cell by cell, while the deviations its demand model chooses for renewals and for new business move the two against each other in both directions by size band; the record does not separate the channels (§3.2, §4.2, §4.4). 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 Here the learner prices against insurers that do not learn as it does, and the outcome is a reallocation: lower prices on new business, higher prices on the renewals it keeps, and higher rate levels at the traditional insurers, whose prices over the true expected loss of the business they write nonetheless fall (Table 8).

Section 2 states the market and the three traditional insurers, which are the benchmark. Section 3 states the Kelly decision: the objective, the price it gives for one account, what keeping an

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.

account is worth, and the chance of winning it. Section 4 states what the Kelly underwriter learns, from what data, and by what estimator. Section 5 reports the fourth seat's two by two, with the registered predictions that failed. Section 6 traces part of the decision's shortfall to the form of the renewal credit, states the corrected form, and reports the second construction: the deviation policy, the containment argument, and the comparison in growth with its limits. Section 7 removes or replaces one part of the Kelly underwriter at a time. Section 8 describes the market with a Kelly underwriter against the market without one. Section 9 reports the sensitivity program, and Section 10 states the limits, sets the market beside two others on the record (§10.2), and concludes. The appendix carries the loss process, the rating plans, the capital standard, the estimators, the Kelly underwriter's computation step by step, the arms of the record, and the tables of the quote data and the loads that Figures 3 and 4 draw (Appendix G).

2 A Stylized Market

This section states the market in which the decision and the learning are tested: the accounts and their losses, what an insurer knows, the three traditional insurers, the buyer and the broker, and capital and the measure. Table 1 lists the assumptions with their values, their grounding, and whether an arm of the record varies them; Figure 1 maps one year.

2.1 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.⁷ Indemnity is limited to \$1,000,000 an occurrence and \$2,000,000 in the aggregate of a 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. Appendix A states the process in full.

⁷The 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.

2.2 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,⁸ 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 reading 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 an incumbent except the buyer's costs of shopping and switching.

2.3 Three Traditional Insurers

Each traditional 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

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

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 0.04.⁹ 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 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.¹⁰

⁸Association for Cooperative Operations Research and Development (ACORD), *Commercial Insurance Application*, ACORD 125 (2007/10), 3, which asks for the claims or losses of the prior five years.

⁹National Association of Insurance Commissioners (NAIC), *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.

¹⁰Fisher 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 underwriter's deviation within the band (§3.4) 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.

The *Bureau insurer* files 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

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

with the complement of credibility ψ equal to one.¹¹ The standard of 1,082 claims is the full-credibility standard of 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 2 on a window of five policy years, 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 (§2.4) 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).¹²

On new business, a quote on an account it does not hold, a traditional insurer files no load. The measured fact is that new business runs at a higher loss ratio than renewals in every one of Wu and Lin's books. 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.¹³ Filed loads of 0.05 and 0.1 are arms of the record (§9). What an underwriter may move without a filing is the schedule modification, and insurers do not always implement the rate their experience indicates.¹⁴ 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 three years, with the credit and the schedule modification together kept inside the filed band. The smallest band is the reference and carries no credit. The rule refers to account size only through the choice of the smallest band as the reference; the load that results in the larger bands is an outcome (§8.1), and three other rules for the load are arms of the record (§9).

¹¹Laurence 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.

¹²Robert A. Bailey, "Insurance Rates with Minimum Bias," *Proceedings of the Casualty Actuarial Society* 50 (1963): 4–11; Werner and Modlin, *Basic Ratemaking*, chap. 10.

¹³Wu and Lin, "Persistency and Renewal Discounts," 398, 402, and 405.

¹⁴Werner and Modlin, *Basic Ratemaking*, chap. 13; Hays, "Commercial Lines Pricing Levels," 11.




Figure 1: One year of the market: each quarter the holder and challengers quote, the buyer chooses, a losing insurer is told its miss, and the policy is booked; at the year's close levels, cells, and credit are re-estimated and capital is tested (§2.5). Solid arrows are flows within a quarter, dashed arrows what returns to an insurer. In all figures the Bureau insurer is bright blue, the Loss-ratio insurer amber, the GLM insurer green, the fourth seat deep blue. Source: authors' construction.

2.4 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. 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 plan-pricing insurers changed insurer a

year, half the 28 percent of middle-market companies that 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 published retentions of commercial business.¹⁵

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.

2.5 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.¹⁶ A policy's premium is booked with expenses of 0.30 and a reserve at the writer's own expected loss; 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.¹⁷ 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.¹⁸

The books are seeded 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

¹⁵Greenwich 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.

¹⁶National 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).

¹⁷National Association of Insurance Commissioners, *U.S. Property & Casualty and Title Insurance Industries—2025 Full Year Results* (NAIC Financial Regulatory Services, 2026).

¹⁸The 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).

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

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

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 is a yearly rate of growth from writing business, the investment income on the reserves an insurer holds included. An insurer that fails has no growth in that seed and would be left out of a mean over seeds; no insurer fails in any seed of any arm of the record, so every mean and contrast is over all 200 seeds. Every arm of the record runs on seeds 0 through 199, and a contrast between two arms or two insurers is paired by seed, with its mean, its t statistic, and the number of seeds in which it is above zero.

Table 1: The market’s assumptions, each with its value in the base arm, its grounding, and the arms that vary it. 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. An arm moves one assumption, except that the profit and expense arms move K with the provision, since K enters on the expected-loss basis (§2.3). Amounts are per occurrence or per policy year as stated and rates per year; the sources are in the text. Source: authors’ construction.

Assumption	Value	Grounding	Varied
Account size, log standard deviation	1	stated	two arms
Claims a year at the median size	1	stated	two arms
Frequency trait: spread, persistence, shock	0.40, 0.85, 0.30	stated	six arms
Severity trait: spread (variance), persistence, shock	0.632 (0.40), 0.85, 0.30	stated	six arms
Claim size, at Verisk’s severity	mixed exponential	Keatinge; Palmer; Verisk	scale, three arms
Loss adjustment expense over indemnity	0.39156	Verisk	two arms
Limits: occurrence; aggregate (\$)	1,000,000; 2,000,000	Palmer	no
Share of claims reported by year	0.55, 0.85, 0.95, 1.00	stated	two arms
Case-reserve error: loss run; holder’s file	0.30; 0.15	stated	four arms
Classes; size bands	5; 3	scale	no
Noise of a read over the traits’ spread	0.7	stated	two arms
Bureau’s statement: years; lag	5; 2	stated	three arms
Loss run: years; valuation lag, quarters	5; 1	ACORD; stated	years, two arms
Full credibility, claims	1,082	Longley-Cook, frequency	no
Filed experience plan: years; cap; K (\$)	3; 100,000; 66,000	bureau plan’s form	one arm
Schedule band, either way	0.25	Florida rule	one arm
Provision π ; expenses ε ; profit	0.34; 0.30; 0.04	NAIC; stated	three arms
Filed load on new business	0.00	Wu and Lin, GL books	two arms
The load’s rule	schedule credit by band	stated	three arms
Lost quotes to fit: complement; schedule	100; 30	stated	no
Windows: loss-ratio method; GLM fit, years	5; 5	Werner and Modlin; stated	two arms; no
Cells’ window; credit’s; schedule fit’s, years	3; 3; 3	stated	six arms
Bureau insurer: experience, years; threshold	5; 0.10	stated	four arms
Costs of shopping and switching, $k = k'$	0.061184	calibrated to Greenwich	two arms
Scale of the buyer’s response s	0.05	stated	two arms
Remarketing unrelated to price	0.10	stated	two arms
Invitations: power of the account’s rank	1	stated	two arms
Broker’s figure rounded to	0.05	stated	two arms
Capital requirement: ρ ; κ	1.72456; 0.0968742	NAIC 2025	no
Flight band, RBC ratio (NAIC trend test)	2 through 3	NAIC; slope stated	two arms
New-business stop, premium over surplus	3	NAIC IRIS; the stop stated	one arm
Failure, RBC ratio	0.70	NAIC Model Act	no
Return on funds, on surplus and reserves	0.0356	NAIC yield	two arms
Opening capital, RBC ratio	6.42458	NAIC 2025	no
Entrants after failures, at most	2	stated	no
Demand model: prior spreads; decay	0.25, 1.0; 0.98	stated	four arms
Renewal credit: discount, gross; horizon	own growth; final year	derived from the objective	the discount, \$6
Accounts; years of seeding, practice, record	4,000; 19, 10, 28	scale	no
Seeding: holders changing a year	0.15	stated	no

3 The Kelly Decision

This section states what an insurer would do if its objective were the growth of its surplus: the objective (§3.1), the increment one account adds to expected growth (§3.2), what keeping an account is worth (§3.3), and the decision, the chance of winning at a price weighed against the increment, with the lowest premium at which the insurer writes (§3.4). Each quantity is stated with the data it reads, and Appendix E gives the computation in the order in which it is carried out.

3.1 The Objective

Let S be the insurer's surplus. Under the Kelly criterion the insurer chooses its actions to maximize the expected logarithm of its surplus at the horizon, $\mathbb{E}[\ln S_T]$. Where an action leaves the insurer's later opportunities unchanged, this 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 (§3.3), the one-period objective carries what the action leaves for later. For a bettor who knows the distribution of outcomes and bets repeatedly on independent trials, the strategy that maximizes the expected logarithm attains, with probability one, a long-run growth rate that no other strategy exceeds, and it never stakes an amount whose loss would exhaust wealth.¹⁹ Two features of an insurer's position change the statement. 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's 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 (§3.3). The property of optimal growth is not claimed for this approximation in this market.

3.2 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,

$$\mathbb{E}\left[\ln \frac{S + n - X}{S}\right] \approx \frac{n - \hat{x}}{S} - \frac{\hat{v}}{2S^2}, \quad (4)$$

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 §3.3. The increment of the decision is

$$G(P) = \frac{n(1 + V) - \hat{x}(1 - c_x)}{S} - \frac{\hat{v}}{2S^2}. \quad (5)$$

¹⁹Breiman, "Optimal Gambling Systems"; Thorp, "The Kelly Criterion."

In the market of Section 2 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.

The Kelly underwriter's working price starts from the same expected loss at a plan's provision, the *working premium* $\hat{x}/(1-\pi)$, so that with nothing else learned it charges what a plan charges on its own estimate. On new business the credit enters the working premium, which solves $P(1-\pi) + V_{\text{nb}} P(1-\varepsilon) = \hat{x}$, with V_{nb} the credit on new business (Equation 6): the insurer asks of a new account this year's provision less what the account is expected to return in later years, and the divisor applies to a negative credit as it does to a positive one. An account worth acquiring is thus quoted below the price the same estimate has as a renewal: the divisor nets the credit from the plan's provision, a rule of this construction, and the decision of §3.4 chooses the deviation around the price so set. In the literature on dynamic insurance contracts “lowballing” names a first premium below the actuarially fair one, recovered on the accounts that stay.²⁰ The Kelly underwriter writes new business at 1.28 times its true expected loss (§8.1), which net of its expenses of 0.30 of premium is below that loss and above it discounted at the return on funds over the mean time to payment.

3.3 What Keeping an Account Is Worth

Equation 4 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 second construction takes another first-order route, which discounts at the return on funds (§6.2). 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,

$$V_{\text{ren}}(H) = m_s \sum_{k=1}^H (\beta p_s)^k \quad \text{and} \quad V_{\text{nb}}(H) = \beta p_1 [m_1 + V_{\text{ren}}(H-1)], \quad (6)$$

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 5'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. The derivation carries no cap and no stated discount rate: the discount is the insurer's own growth, a negative margin gives a negative credit, and a finite horizon, here the

²⁰D'Arcy and Doherty, “Adverse Selection”; Jia and Wu, “Insurer Commitment,” 89.

number of record years, needs no bound for any $\hat{g}$ (a going-concern sum exists only where $\beta p_s < 1$). Until the window holds a first renewal, the credit on new business is zero.

3.4 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 working premium holds. The expected increment of quoting is $h(\ell) G(P_\ell)$ with P_ℓ the working premium times e^ℓ , and the Kelly decision is the deviation that maximizes it:

$$\ell^* = \arg \max_{\ell} h(\ell) G(P_\ell), \quad P = \max \left\{ P_{\ell^*}, \frac{\hat{x}(1 - c_x) + \hat{v}/(2S)}{(1 - \varepsilon)(1 + V)} \right\}. \quad (7)$$

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 §4.3 states. Because the Kelly underwriter rates on a schedule plan as the others do, its deviation and its schedule modification 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 is one. The second element of Equation 7 is the *floor*: the premium at which Equation 5 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 charges more for risk when its surplus is low.

Equation 7, over the filed band around the working premium, 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 provision π , and the band. A rating plan supplies a version of $\hat{x}$, the provision, and the band; it supplies none of the others.

4 Learning

This section states what the Kelly underwriter learns and how: each account's loss distribution, tracked from the files it reads (§4.1); a correction for the business it wins (§4.2); and the hold curve, from the broker's figures (§4.3). The renewal credit's five quantities are estimated on its own book and are stated with the others in Appendix E. Section 4.4 reports what was learned in the base arm. Every estimator reads only data that an insurer in the market of Section 2 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 (§2.1, Appendix A), the spread of the case-reserve error on a broker's loss run and on a holder's own file (§2.2, Appendix A), the share of loss adjustment expense (§2.1), and, as for every insurer, the distribution of the traits within a class that the schedule modification reads (Appendix B).

4.1 The Account's Loss Distribution

A rating plan reads an account's history through one number, the experience modification on its capped losses. The Kelly underwriter 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.²¹ 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 Kelly underwriter 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 (Appendix D). 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:

$$\hat{x}_0 = \hat{\lambda} \mathbb{E}[\min(Y, u)] (1 + a), \quad \hat{v} = \hat{\lambda} \mathbb{E}[\min(Y, u)^2] (1 + a)^2, \quad (8)$$

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 $\hat{v}$ with it, as the cell's factor of §4.2 does. The variance is the quantity a plan does not hold; it enters the price through the increment of Equation 5, which Equation 7 multiplies by the hold probability and maximizes, and through the floor of Equation 7.

4.2 The Business Won

An insurer bears losses only on the accounts it wins, and winning is informative. Where several insurers quote from estimates with error and the buyer favors the lowest price, an insurer wins more often where its own estimate is low, so the losses of the business won run above the estimates that won it.²² The size of the effect depends on how many insurers quote and on what each knows, and it differs between an account the insurer holds and one it does not. The Kelly underwriter measures it. It sorts the policies it has written into four *cells*, one of renewals and three of new business, one in each size band, and at each year's close applies the loss-ratio method to each cell: with a_c the reported losses of the cell's policies of the past three policy years, e_c the expected reported losses on the tracked estimates $\hat{x}_0$, and Z_c the classical credibility of the cell's reported claims, the cell's factor becomes

$$\kappa_c \leftarrow Z_c \frac{a_c}{e_c} + (1 - Z_c) \kappa_c, \quad (9)$$

²¹The 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.

²²Edward 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.

and the expected loss in the price is $\hat{x} = \kappa_c \hat{x}_0$ for an account quoted in cell c , the variance taking the same factor. The cells take the place of a rate level: the Kelly underwriter has no multiplier on its whole book, and its estimate is brought to its experience cell by cell.

4.3 The Demand Model

The hold curve h of Equation 7 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 and withheld on those that did. Let ζ be a latent log price: the logarithm of the premium at which the account would bind with another insurer, over the Kelly underwriter's working premium for it, each observation brought on-level to the reference of its *segment*, a size band in renewals or in new business, before a fit (Appendix D), as the traditional insurers bring their quotes on-level. On a quote that was lost, 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; 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 Kelly underwriter 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 insurer's own working premium 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,

$$\bar{h}(\ell) = \int \left[1 - \Phi\left(\frac{\ell - \mu_\zeta}{\sigma}\right) \right] \pi(\mu_\zeta, \sigma \mid \text{data}) d(\mu_\zeta, \sigma), \quad (10)$$

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 5 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 what a deviation teaches about the curve in later quarters. Before any observation the prior predictive governs, so a quote can leave the working premium from the first quarter. The curve is fitted separately in each size band for renewals and for new business, six segments in all.

4.4 What Was Learned

Table 2 reports the learned quantities in the base arm, as means over the record years and the seeds.

The cells' factors are above one in every cell (Table 2): the losses of the business the Kelly underwriter writes run above the tracked estimates, and the factor on new business falls with account size. One reading, which no arm isolates, is that a loss run of five years carries more claims on a large account than on a small one. After the cells' correction the true expected loss of the new business the Kelly underwriter wins still exceeds its estimate, by 5.9 percent in the smallest band and by 0.8 percent in the largest. In a single year a cell with few claims reaches a factor as high as 4.64 (Table 2).

On its retained renewals the Kelly underwriter earns a margin of 0.272 of premium net of expenses at an account's first renewal and 0.287 at later ones; it keeps 0.385 of its accounts at their first renewal and 0.722 at each later one; and its surplus grows at 0.061 a year over the credit's window, the return on funds included, a discount factor β of 0.941. The credit (Table 2) falls toward zero over the final record years as the horizon closes.

The deviation chosen on the demand model stands near the band's limit of 0.223 on renewals in the two smaller bands and inside the band on the largest band's renewals; Table 2 gives it by phase and band. The floor of Equation 7 binds no quote of the base arm: with the credit in the floor, the deviated working premium stays above it, and the price of new business is set by the credit's divisor, the cells, and the deviation rather than by the floor. The premium the Kelly underwriter writes on new business is 1.30 times its own estimate of the expected loss, the same figure in each of the three size bands at the tables' precision; its renewals are written at 1.68 times the true expected loss of the renewals it keeps, and Table 8 sets the two phases against the true expected loss on one basis.

Table 2: What the Kelly underwriter learned in the base arm, means over record years and seeds. A cell's factor multiplies the tracked expected loss (Equation 9); its maximum is over every year and seed. The credit is per unit of net premium, one figure for all bands on new business (Equation 6). The deviation is the decision's grid choice of the logarithm of quoted over working premium, by band from smallest to largest, for renewals and new business, averaged over a year's fourth-quarter quotes before the floor and the filed band apply, the band bounding the deviation and schedule modification together to at most 0.223 in logs. Source: authors' calculations.

Quantity	Renewals, by size band			New business, by size band		
	Smallest	Middle	Largest	Smallest	Middle	Largest
Cell's factor, mean		1.430		1.385	1.255	1.111
Cell's factor, largest in any year		3.04		4.64	3.51	2.29
Renewal credit, per unit of premium net of expenses		0.523			0.280	
Deviation from the working premium, logarithm	0.221	0.205	0.134	0.200	0.212	0.216

The fitted moments of the tracker, the estimator of \$4.1, are of the kind the loss process has, and not all of its values: at the end of the record years it puts the persistence of the decaying part at 0.762 for frequency, where the process has 0.85, and at 0.356 for the log scale, below the process's value, and it finds a covariance between the two of -0.037 , the sign that a class defined on the sum of two independent traits induces within the class.

5 The Two Together

This section reports the fourth seat's two by two: the growth of the fourth seat when it holds the plan's estimate or the learned estimate and prices as a plan prices or by the Kelly decision, with the three traditional insurers the same in all four arms. Table 3 gives every insurer's growth in the four arms, Table 4 the contrasts, and Figure 2 the growth by seed. Three of the first registration's predictions for this section failed, and Section 6 traces part of the shortfall to one term of the decision.

With the plan's estimate and the plan's pricing the fourth seat is a second Loss-ratio insurer, and the market has no Kelly underwriter. With the learned estimate and the plan's pricing it holds the tracked estimate of \$4.1 in place of the plan's, and prices it as the Loss-ratio insurer prices:

a rate level moved by the loss-ratio method with the competitors' level as its complement, and the schedule credit on new business. With the plan's estimate and the Kelly decision it applies Equation 7, with its cells, its renewal credit, and its demand model, to the plan's expected loss and to a variance taken from the class. With both it is the Kelly underwriter.

Table 3: Growth by insurer in the four arms of the fourth seat, means over 200 seeds, and the fourth seat's share of the accounts at the end of the record years. Growth is the terminal logarithm of surplus over the endowment, net of the return on funds, over 28 years (Equation 3). Source: authors' calculations.

The fourth seat	Growth			Fourth seat	Fourth seat's share of accounts
	Bureau	Loss-ratio	GLM	Fourth seat	
Plan's estimate, plan's pricing	+0.612	+0.506	+0.584	+0.534	0.223
Learned estimate, plan's pricing	+0.393	+0.323	+0.416	+0.935	0.672
Plan's estimate, Kelly decision	+0.506	+0.372	+0.452	+0.532	0.142
Learned estimate, Kelly decision	+0.338	+0.261	+0.367	+0.691	0.250

Learning alone is worth +0.400 of growth ($t = +31.10$), and the learned-estimate seat has the highest growth of the four, at +0.935. It holds 0.672 of the accounts at the end of the record years and binds 0.618 of its quotes as a challenger. The new business it wins has a true expected loss of 0.927 times its own estimate, where a second Loss-ratio insurer's has 1.012: with the tracked estimate the insurer does not win mainly where its estimate is low.

Table 4: The fourth seat's contrasts and the Kelly underwriter's leads, paired by seed over 200 seeds: the mean, its t statistic, and the seeds above zero. Learning alone is the learned estimate less the plan's, both at the plan's pricing; the Kelly decision on an estimate is the Kelly decision less the plan's pricing on that estimate; the contrast Both is the Kelly underwriter less a second Loss-ratio insurer; and the interaction is Both less the sum of learning alone and the Kelly decision on the plan's estimate. Contrasts are taken on unrounded values and may differ by 0.001 from differences of Table 3's means. Source: authors' calculations.

Contrast	Mean	t	Seeds above zero
<i>The fourth seat's growth, one arm less another</i>			
Learning alone	+0.400	+31.10	197
The Kelly decision on the plan's estimate	-0.003	-0.24	108
The Kelly decision on the learned estimate	-0.244	-31.11	3
Both	+0.156	+11.25	153
Interaction	-0.241	-17.36	25
<i>The Kelly underwriter's growth less a traditional insurer's</i>			
Bureau insurer	+0.353	+22.73	190
Loss-ratio insurer	+0.430	+30.25	192
GLM insurer	+0.323	+22.35	186

The Kelly decision with the quote-anchored credit, the credit entering in proportion to the premium quoted (\$6.1), does not add what the first registration predicted of it. On the plan's estimate it moves growth by -0.003 ($t = -0.24$), where the registered prediction was a loss at $t < -2$; the prediction failed, the contrast being indistinguishable from zero. On the learned estimate it lowers

growth: -0.244 ($t = -31.11$), where the registered prediction was a gain at $t > 2$. The interaction, predicted above zero, is -0.241 ($t = -17.36$). The three predictions failed, and the pattern is not a property of one setting: at the Bureau insurer's two other refiling thresholds and at the two other levels of read noise, the decision with the quote-anchored credit lowers growth on the learned estimate at $t < -2$ every time (Section 9, Table 11).

The two together are still worth $+0.156$ over a second Loss-ratio insurer ($t = +11.25$), and against the three traditional insurers in its own market the Kelly underwriter leads by $+0.353$, $+0.430$, and $+0.323$ (Table 4). In the two by two the lead comes from the learning, and the decision with the quote-anchored credit lowers it.




Figure 2: Growth by insurer and seed, over 200 seeds, in the fourth seat's four arms, the pricing by row (plan's above, Kelly decision below) and the estimate by column (plan's left, learned right). Each insurer's row of points gives growth (Equation 3) on one horizontal scale, a vertical bar marking the mean, stated at right; zero is short-dashed. Source: authors' calculations.

6 The Credit's Form, and a Second Construction

This section traces part of the decision's shortfall of Section 5 to the form in which the renewal credit enters the decision and corrects the form (§6.1). It sets out the corrected decision as a deviation policy around the plan's pricing of the learned estimate, which contains that pricing as one of its choices (§6.2), and measures the correction on that policy and on the Kelly underwriter, with the

credit discounted at the insurer's own growth and at the return on funds (§6.3). It then shows where the two differ in the quote data and in the market around them (§6.4). The arms of this section were registered after the two by two's outcomes were on the record: the construction's eight arms in a second registration, and in a third and a fourth two arms each with the credit discounted at the insurer's own growth. Every prediction was written before any run of its arms on the record's seeds. The second registration's mostly followed a disclosed reading of seeds outside the record; the third's and the fourth's preceded any run of their arms (Appendix F). They run on the same market and the same seeds.

6.1 The Diagnosis and the Correction

In the decision with 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 continuation value 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 (§2.4), which excludes the holder's own raise. What today's deviation buys is today's margin and the chance of keeping the account at all— $\bar{h}(\ell)$ multiplies the whole increment, so the value of keeping the account, $\bar{h}(\ell) Vn(0)$, is carried in full. The correction therefore anchors the credit at the *undeviated premium*, the premium of the zero deviation: for the deviation policy of §6.2, the price of the seat it deviates from, and for the Kelly underwriter, the working premium. 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 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 this section are, if anything, higher than a credit at the realized scale would give, and the record does not measure by how much.

6.2 The Deviation Policy, and Containment

To measure what the corrected decision adds to the plan's pricing of the learned estimate, the second construction places in the fourth seat a policy that computes, account by account and by the same computation, the price the learned-estimate seat of Section 5 would quote. The policy quotes that price times e^δ , at the δ that maximizes the posterior-averaged hold probability times the increment of Equation 5 with the anchored credit, on a grid over the filed band in steps of 0.0025 in log price. The floor of Equation 7, in its form with the quote-anchored credit, limits δ from below, and the schedule band limits it from both sides. The increment is evaluated on the contained seat's own estimate, the tracked expected loss and variance at its rate level without the cells, and the hold probability is the policy's own demand model's, fitted on its deviations δ without bringing them on-level. Below, this policy, the Kelly decision applied as a deviation from the learned-estimate seat's price, is the *deviation policy*, and with the anchored credit the *corrected deviation policy*. The zero deviation at every account reproduces the learned-estimate seat's pricing, quote for quote, wherever

the floor does not bind, as it binds none on this record, so the policy contains that seat's pricing as one of its choices, and what it adds or subtracts is measured against a seat it could have reproduced. The containment holds quote by quote, not path by path: once a deviation is written the market's state diverges, and the policy maximizes an estimated one-period objective, so containment guarantees that the contained pricing can be reproduced, not that it will be improved on—the policy with the quote-anchored credit trails the seat it contains.

In every arm of the second registration's eight that prices by the Kelly objective, except the arm that reproduces the registered Kelly underwriter, the credit is discounted at the return on funds rather than at the insurer's own growth. At an interior optimum of the expected logarithm of surplus over a riskless position, the expected deflator of a sure dollar is $1/(1+r)$, with r the return on funds: the numeraire-portfolio property. The factor $e^{-\hat{g}}$ understates it by about $\hat{g} - r$ a year. The insurer of this market does not choose its riskless position freely, so neither derivation is exact here, and the choice between them is the construction's.²³ Below, the *funds discount* is $\beta = 1/(1+r)$ and the *growth discount* is the $\beta = e^{-\hat{g}}$ of §3.3. The choice was registered before the record's outcomes were read, its direction first read on seeds outside the record (Appendix F), with the prediction that on the *standalone* Kelly underwriter, the Kelly underwriter priced by its own decision rather than as a deviation policy, the funds discount raises growth.

Eight arms of the second registration carry the construction and its controls (Table 5, above its subheading): the learned-estimate seat and the Kelly underwriter exactly as Section 5 ran them, each reproducing its arm of the record run for run; the Kelly underwriter with the credit's growth discount replaced by a discount at the return on funds; the deviation policy with the quote-anchored credit and with the anchored credit; the corrected standalone Kelly underwriter, its own credit anchored at its undeviated premium; and, on the corrected deviation policy and the corrected standalone each, the *pessimism variant*, which lowers the conditional posterior mean of the demand model's μ_ζ by one posterior standard deviation before the average of Equation 10 is taken, and so lowers every hold probability. The third registration ran the corrected deviation policy, with and without the pessimism variant, on the growth discount, two arms over the same 200 seeds, with two predictions: that the growth discount raises growth on the corrected deviation policy, and that the corrected deviation policy with pessimism gains over the learned-estimate seat on that discount, the second registration's primary prediction (§6.3) evaluated on the growth discount. The fourth registration ran the corrected standalone Kelly underwriter, with and without the pessimism variant, on the growth discount, two arms over the same seeds, with four predictions: that the correction raises growth on the Kelly underwriter as registered, that corrected it still trails the learned-estimate seat, that the growth discount raises its growth over the funds discount, and that the pessimism variant raises it.

²³John B. Long Jr., "The Numeraire Portfolio," *Journal of Financial Economics* 26, no. 1 (1990): 29–69.

Table 5: The second construction's arms: the fourth seat's growth (mean over 200 seeds) and its renewal credit on new business and renewals, means over years and seeds; its share of the accounts at the end of the record years and, in the whole market, the accounts changing insurer a year, means over seeds. Above the subheading, the arms priced by the Kelly objective, the registered Kelly underwriter aside, discount the credit at the return on funds (\$6.2); below it are the third and fourth registrations' arms, which discount it at the insurer's own growth. A corrected arm carries the anchored credit. Table 6 gives the contrasts among the arms. A dash marks what an arm's record does not carry; the registered Kelly underwriter's credits are in Table 2. Source: authors' calculations.

The fourth seat	Fourth seat's	Renewal credit		Fourth seat's	Changing
	growth	V_{nb}	V_{ren}	share	insurer a year
The learned-estimate seat	+0.935	—	—	0.672	0.163
Kelly underwriter as registered	+0.691	—	—	0.250	0.394
Kelly underwriter, funds discount	+0.678	0.300	0.559	0.259	0.403
Deviation policy, quote-anchored credit	+0.836	—	—	0.308	0.258
Corrected deviation policy	+0.945	0.615	0.782	0.488	0.263
Corrected deviation policy with pessimism	+0.945	—	—	0.494	0.263
Corrected Kelly underwriter	+0.711	0.351	0.550	0.338	0.368
Corrected Kelly underwriter with pessimism	+0.723	—	—	0.353	0.365
<i>On the growth discount</i>					
Corrected deviation policy	+0.939	0.528	0.688	0.452	0.269
Corrected deviation policy with pessimism	+0.941	0.526	0.686	0.461	0.270
Corrected Kelly underwriter	+0.745	0.319	0.509	0.332	0.368
Corrected Kelly underwriter with pessimism	+0.748	0.323	0.511	0.333	0.366

6.3 Outcomes

Table 6 reports, paired by seed, the ten contrasts the second registration reads, six of them with a predicted direction; in the final row of its first block, one contrast outside the registration; and below its second subheading, the contrasts of the third and fourth registrations on the growth discount. The correction is worth +0.110 on the deviation policy ($t = +14.85$): the deviation policy with the quote-anchored credit trails the seat it contains by 0.099 ($t = -12.01$), and corrected it does not trail. The corrected deviation policy's leads over the three traditional insurers are +0.633, +0.705, and +0.611 (t of 43.45, 54.30, and 45.28), as predicted. The Kelly underwriter on the funds discount trails the learned-estimate seat by 0.257 ($t = -36.35$), and the correction recovers +0.033 of that gap ($t = +7.09$): corrected, that insurer still trails the seat by 0.224 ($t = -38.31$). On the growth discount, the discount at which the Kelly underwriter was registered, the fourth registration measures the correction on Section 5's arm itself: it is worth +0.054 ($t = +11.03$), and corrected, that insurer trails the seat by 0.190 ($t = -30.88$); both predictions held. The gap that remains at either discount shows that what separates the standalone architecture from the plan's pricing of the same estimate is not the credit's form alone. On the standalone's side the two differ in the working premium, the cells, the credit's levels, and the demand model's deviation; on the seat's side, in the rate level and the schedule credit on new business. The gap is not attributed among these parts on this record.

Besides the primary prediction (below), two predictions of the second registration failed: a discount of the credit at the return on funds, predicted from the numeraire-portfolio property at an

interior optimum (§6.2) to raise growth, lowers it (-0.013 , $t = -3.35$); and the pessimism variant, which the one-period posterior average does not call for, adds an amount the record does not distinguish from zero on the corrected deviation policy (0.000 , $t = +0.03$). On the corrected standalone it adds $+0.012$ ($t = +3.22$), and that prediction held; on the growth discount it adds $+0.003$ ($t = +0.80$), and the fourth registration's prediction of a gain there failed. The third registration's prediction on the discount failed: the corrected deviation policy on the growth discount less the same policy on the funds discount is -0.006 ($t = -1.36$), and with pessimism -0.005 ($t = -0.96$), both of the sign opposite to the prediction and neither distinguished from zero. On the corrected Kelly underwriter the growth discount raises growth by $+0.034$ ($t = +9.46$), and the fourth registration's prediction held; with the quote-anchored credit the funds discount lowers it (above), so on the standalone Kelly underwriter the growth discount gives the higher growth with either form of the credit.

The second registration's primary prediction was the corrected deviation policy with pessimism against the seat it contains: a gain at $t > 2$. The outcome is $+0.0108$ at a standard error of 0.0055 , a computed t of 1.9656 , below the rule's threshold, and the prediction failed. The rule's threshold is stricter than the one-sided 5 percent level of a test of the registered direction, whose critical value at this number of seeds is below the computed t (one-sided $p = 0.025$): at that level the gain would be above zero. Two standard errors above the mean, the addition to the plan's pricing of the learned estimate is 0.022 , under a tenth of what the learning itself is worth (Section 5); at the registered threshold the record does not distinguish the addition from zero. Evaluated on the growth discount, the corrected deviation policy with pessimism less the learned-estimate seat is $+0.006$ ($t = +1.03$; 119 of 200 seeds above zero), and the third registration's prediction of a gain there failed.

Table 6: The second construction's contrasts, the fourth seat's growth paired by seed over 200 seeds: the mean, its t statistic, and the number of seeds in which the contrast is above zero. The correction is an arm with the anchored credit less the same arm with the quote-anchored credit, the funds discount the Kelly underwriter on that discount less the registered one, and the pessimism variant's contrast an arm with the variant less the same arm without it, and the growth discount on an arm the arm on the growth discount less the same arm on the funds discount. The second registration's six predictions are on the two corrections, the funds discount, the two contrasts of the pessimism variant, and the corrected deviation policy with pessimism less the learned-estimate seat; the final row of the first block is not among its ten. Below the second subheading are the third and fourth registrations' contrasts, on the growth discount: the third's two predictions are on the growth discount on the corrected deviation policy and on the corrected deviation policy with pessimism less the learned-estimate seat, the fourth's four on the correction on the Kelly underwriter as registered, the growth discount on the corrected Kelly underwriter, the pessimism variant on it, and the corrected Kelly underwriter less the learned-estimate seat. The registered rule holds a predicted contrast at a paired mean of 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 correction, on the deviation policy	+0.110	+14.85	177
The correction, on the Kelly underwriter (funds discount)	+0.033	+7.09	134
The funds discount, on the Kelly underwriter	-0.013	-3.35	75
Pessimism, on the corrected deviation policy	0.000	+0.03	100
Pessimism, on the corrected Kelly underwriter	+0.012	+3.22	119
<i>Less the learned-estimate seat</i>			
The corrected deviation policy with pessimism	+0.011	+1.97	113
The corrected deviation policy	+0.011	+1.92	118
The quote-anchored deviation policy	-0.099	-12.01	34
The corrected Kelly underwriter	-0.224	-38.31	3
The corrected Kelly underwriter with pessimism	-0.212	-35.45	1
The Kelly underwriter, funds discount	-0.257	-36.35	1
<i>On the growth discount</i>			
The correction, on the Kelly underwriter as registered	+0.054	+11.03	164
The growth discount, on the corrected deviation policy	-0.006	-1.36	86
The growth discount, on the same with pessimism	-0.005	-0.96	99
The growth discount, on the corrected Kelly underwriter	+0.034	+9.46	153
The growth discount, on the same with pessimism	+0.025	+6.43	136
Pessimism, on the corrected deviation policy	+0.002	+0.42	98
Pessimism, on the corrected Kelly underwriter	+0.003	+0.80	110
<i>Less the learned-estimate seat</i>			
The corrected deviation policy with pessimism	+0.006	+1.03	119
The corrected deviation policy	+0.004	+0.71	113
The corrected Kelly underwriter	-0.190	-30.88	3
The corrected Kelly underwriter with pessimism	-0.187	-31.30	2

6.4 Growth Not Distinguished, and Where the Two Differ

Two readings follow, and the quote data rule out a third—that the two insurers are running the same pricing. First, the plan’s pricing of the learned estimate and the corrected deviation policy are not distinguished in growth on this record at the registered threshold: an observer of growth alone does not tell, on this record, an insurer that prices its learned estimate as a plan prices from one that applies the corrected Kelly decision to the same information. The statement holds at one calibration, at this record’s precision, and is not a proof of zero; it reads the same under either discount of the credit (§6.3), and the anchored credit stands below a credit at the realized scale, which biases the deviations upward (§6.1). Second, the statement is one of existence: there is a pricing that computes no Kelly decision at quote time and whose growth this record does not distinguish from that of the corrected Kelly decision on the same information, and other pricings may give the same growth by other means.

The two differ in the quote data (Figure 3 and Table G.1). The figures of this paragraph are those of the corrected deviation policy with pessimism, the primary prediction’s arm; the arm without pessimism is within 0.007 of each. The deviation policy prices renewals 0.118 in log price above the seat it contains on average, against the band’s limit of 0.223, with 0.276 of its renewal quotes within one grid step of the filed band’s upper end, and it prices new business 0.088 in log price below that seat. Under the growth discount (Table G.1) the pattern repeats: renewals 0.133 in log price above the seat on average, with 0.288 of renewal quotes within one grid step of the band’s upper end, and new business 0.063 in log price below that seat. The filed rate band, not the floor, bounds the policy’s renewal pricing, the band being the market’s regulatory constraint; the record runs no deviation policy at another band. The two also differ in the market around them (Table 5): the deviation policy with pessimism holds 0.494 of the accounts at the end of the record years against the learned-estimate seat’s 0.672, and accounts change insurer at 0.263 a year against 0.163. The behaviors are therefore told apart in quote microdata—the deviations, the band-edge mass, the sign split between renewals and new business—and in the shares and the turnover around them.




Figure 3: At left, the corrected deviation policy with pessimism less the learned-estimate seat, growth paired by seed (horizontal axis), one jittered point a seed; the solid bar is the mean and the short-dashed line zero. At right, for the deviation policy with the quote-anchored credit, the corrected deviation policy, and the corrected deviation policy with pessimism (rows), the share of renewal quotes (horizontal axis) within one grid step of the band's upper end (filled) and at the zero deviation (open), each pair joined in gray. Source: authors' calculations.

7 What Each Part Is Worth

The gap of \$6.3 is not attributed among the Kelly underwriter's parts; this section removes one part of the Kelly underwriter at a time, the renewal credit, the demand model, and the cells' correction, then all three, and then changes two inputs, the tracker's moments and the history every insurer reads. No direction was registered for these arms. Table 7 reports the fourth seat's growth in each arm, its difference from the Kelly underwriter's paired by seed, and its leads.

Without the renewal credit the Kelly underwriter's growth falls by 0.224 ($t = -29.91$), to +0.467, and its three leads turn negative: -0.179 , -0.067 , and -0.172 (t of -10.42 , -3.94 , and -10.31). The credit is the part that makes the Kelly underwriter a competitor for accounts it does not hold: without it the insurer binds 0.281 of its quotes as a challenger where the Kelly underwriter binds 0.679, and ends with 0.114 of the accounts. It is the one part whose absence turns the leads.

Table 7: The Kelly underwriter with one part removed, all three removed, or one of two inputs changed. The columns give the fourth seat's growth; its growth less the Kelly underwriter's and its leads, paired by seed (t in parentheses); and its final share of accounts. Settled-sample moments freeze the tracker's class mean log scales and second moments at its estimator's fit on a large settled sample; with five-year files every insurer reads five years of history, the traditional insurers on the plan as filed. Source: authors' calculations.

The fourth seat	Growth	Less the Kelly		Lead over (t)			Share of accounts
		underwriter (t)	Bureau	Loss-ratio	GLM		
Kelly underwriter	+0.691	—	+0.353 (+22.73)	+0.430 (+30.25)	+0.323 (+22.35)		0.250
No renewal credit	+0.467	-0.224 (-29.91)	-0.179 (-10.42)	-0.067 (-3.94)	-0.172 (-10.31)		0.114
No demand model	+0.645	-0.046 (-5.72)	+0.481 (+46.45)	+0.507 (+49.52)	+0.485 (+46.68)		0.678
No cells' correction	+0.648	-0.042 (-3.67)	+0.539 (+40.10)	+0.548 (+43.27)	+0.547 (+44.21)		0.802
None of the three	+0.768	+0.077 (+7.84)	+0.587 (+59.86)	+0.617 (+68.34)	+0.590 (+61.87)		0.928
Settled-sample moments	+0.830	+0.139 (+15.43)	+0.541 (+55.36)	+0.616 (+74.67)	+0.498 (+54.73)		0.327
Five-year files for all	+1.029	+0.339 (+38.58)	+0.849 (+97.42)	+0.896 (+142.63)	+0.791 (+106.41)		0.414

Without the demand model the deviation is zero, and the insurer quotes its working premium. Its growth is lower by 0.046 ($t = -5.72$), and its three leads stay above zero with t of 46.45 or more. At the working premium it holds more accounts, 0.678 of them at the end, at a lower margin: its renewals are written at 1.51 times their true expected loss against 1.68 for the Kelly underwriter. The market's capital standard binds on its competitors: 25,435 new-business invitations are withheld by broker flight over the record years against 0 in the base arm.

Without the cells' correction the tracked estimate enters the price as it stands. The insurer's growth is lower by 0.042 ($t = -3.67$), and it prices below its competitors on a book whose losses run above its estimates: it ends with 0.802 of the accounts, the three traditional insurers hold 0.198 between them with growth of +0.109, +0.100, and +0.101, and broker flight withholds 75,696 of their new-business invitations as their capital ratios fall toward the band. In no arm of this section does an insurer fail in any seed.

With the credit, the demand model, and the cells removed at once, what remains is the tracked estimate at a plan's provision with no rate level, no deviation, and no correction. That insurer grows by +0.768, which is +0.077 relative to the Kelly underwriter ($t = +7.84$). The removals are not additive: removing the three together raises growth where removing the credit alone lowers it, and in the construction the demand model's deviation and the cells' factors raise prices that the credit's divisor lowers. With the tracker alone in the fourth seat, accounts change insurer at 0.080 a year, the lowest rate of any arm of the record.

With the tracker's moments, each class's mean log scale and the second moments, frozen at its own estimator's fit to a sample of 40,000 accounts over 41 fully settled years, in place of the estimates it forms from the files it reads, the fourth seat grows by +0.139 more than the Kelly underwriter ($t = +15.43$): what estimating the moments on its own files costs the tracker in this market, against the same estimator on a large sample. The frozen moments are the estimator's own on that sample, in its form (Appendix D), and not the loss process's. With every insurer reading five years of an account's history and the traditional insurers rating by the experience plan as filed, in place of the files of every year held and the plan each estimates from its own book, the fourth seat grows by +0.339 more ($t = +38.58$), the largest movement of the section. The arm changes the reading of every insurer, not of the Kelly underwriter alone, so its leads are those of a different market, and the record carries the outcome and no mechanism for it.

8 The Market with a Kelly Underwriter

This section compares the market in which the fourth seat is the Kelly underwriter with the market in which it is a second Loss-ratio insurer: the load on new business by account size (§8.1), and turnover, shares, and the traditional insurers' rate levels (§8.2).

8.1 The Price of New Business

Figure 4 draws the load each insurer charges on new business by size band, and Table G.2 prints it with its construction. The smallest band is the reference of the schedule credit and carries none (§2.3).

The Kelly underwriter's load is negative in every band and falls with size: -26.2 , -31.5 , and -36.8 percent in the smallest, the middle, and the largest band (t of -26.67 , -87.18 , and -197.19), and the smallest band's load exceeds the largest's by 10.6 percentage points ($t = +10.58$). Three learned quantities enter the level and the gradient. The renewal credit's divisor prices every new account below the renewal price of the same estimate, one discount for all bands; the cells' factors, which multiply the estimate in the price, are below the renewal cell's factor in every band of new business and higher in the smaller bands than in the largest (Table 2), which lowers the price of new business against renewals and raises it on small accounts relative to large ones; and the deviations chosen on the demand model differ by phase, the largest band's renewals standing inside the band while its new business stands near the band's limit (Table 2), which offsets part of the cells' gradient. Against the true expected loss of the business written, the Kelly underwriter's price of new business is 1.23 in the smallest band and 1.29 in the largest, and its renewal price 1.63 and 1.70.

The traditional insurers' schedule credits also price new business below renewals, and by more where they compete with the Kelly underwriter: their load in the smallest band exceeds that in the largest by 13.5 percentage points on the mean of the three ($t = +97.98$), against 2.7 percentage points where the fourth seat is one of their own kind ($t = +29.91$), and the largest band's load is higher by 10.8 percentage points without the Kelly underwriter ($t = +65.80$). The schedule credit follows the winning prices an insurer sees (§2.3); the record does not separate, in either market, the causes of the gradient among the competitors' information, the number of competing quotes, which rises with size in both markets, and the experience modification's credibility, which also rises with size.




Figure 4: The load on new business in percent (vertical axis, one scale) by size band (horizontal axis: thirds of the year's accounts by manual premium), one panel an insurer; zero is short-dashed. Solid lines in the insurer's color are the market with the Kelly underwriter, dashed gray lines that with a second Loss-ratio insurer in the fourth seat, whose own load is the fourth panel's dashed line. Points are means over 200 seeds of seed means over the years. Source: authors' calculations.

Table 8 sets the two phases of a policy side by side, on the true expected loss of the business written, which the simulation records and no insurer observes.

The Kelly underwriter writes new business at 1.28 times its true expected loss and renews at 1.68, a loss ratio of 0.781 on new business and 0.595 on renewals. In the market without it, the traditional insurers' loss ratios, across insurers and phases, lie within 0.045 of premium of one another. A new-business loss ratio above the renewal loss ratio, the pattern measured in insurers' books, holds for every insurer in both markets, by between 0.020 and 0.040 of premium for those that price by a plan and by 0.185 for the insurer that values retention. The Kelly underwriter's gap comes through the three learned quantities stated above, channels the record does not separate. In the market without a Kelly underwriter, every insurer's new-business selection, the true expected loss of the business it writes over its own estimate, is above one: the business an insurer wins runs above its own estimate, by 1.5 percent for the Bureau insurer. With the Kelly underwriter in the market the traditional insurers' selections on new business fall below one, and on renewals they rise (1.064 for the Bureau insurer, against 1.007): the new business the plans win is business their estimates overstate, and the renewals they keep are renewals their estimates understate; the record does not separate the mechanism.

Table 8: New business and renewals by insurer, sums over record years and seeds, with the Kelly underwriter or a second Loss-ratio insurer in the fourth seat. Price is premium over its true expected loss, the loss ratio is ultimate losses with loss adjustment expense over premium, and selection is the true expected loss over the writer's estimate. The share bound is, on new business, the quotes bound over the challenger quotes on accounts that went to market and were placed, and on renewals, the renewals kept over the renewal quotes. Source: authors' calculations.

Insurer	New business				Renewals			
	Price	Loss ratio	Selection	Bound	Price	Loss ratio	Selection	Bound
<i>With the Kelly underwriter</i>								
Bureau insurer	1.39	0.721	0.976	0.241	1.42	0.701	1.064	0.581
Loss-ratio insurer	1.39	0.721	0.974	0.199	1.43	0.701	1.061	0.553
GLM insurer	1.34	0.745	0.937	0.406	1.40	0.714	1.081	0.662
Fourth seat	1.28	0.781	1.012	0.679	1.68	0.595	1.017	0.592
<i>Without it</i>								
Bureau insurer	1.45	0.691	1.015	0.252	1.51	0.665	1.007	0.870
Loss-ratio insurer	1.46	0.684	1.011	0.218	1.51	0.662	1.002	0.847
GLM insurer	1.42	0.707	1.007	0.274	1.50	0.667	1.013	0.862
Fourth seat	1.46	0.687	1.012	0.231	1.51	0.662	1.003	0.853

8.2 Turnover, Shares, and Rate Levels

Table 9 compares the two markets, and Figure 5 draws the shares and the rate levels over the record years.

In the market without a Kelly underwriter, accounts change insurer at the calibrated rate. The buyer's costs were calibrated, on seeds outside the record, so that 0.14 of the accounts of the market of four plan-pricing insurers change insurer a year; on the record's 200 seeds 0.1410 of that market's accounts change insurer a year (standard error over seeds 0.0009), within two standard errors of the target, and the first registration's prediction held. With the Kelly underwriter in the fourth seat 0.394 of the accounts change insurer a year, and 0.491 go to market against 0.266: the turnover of the market with the Kelly underwriter is an outcome, not a calibration. The Kelly underwriter binds 0.679 of the quotes it makes as a challenger, and Table 8 gives each insurer's share bound in the two markets.

The Kelly underwriter holds 0.250 of the accounts at the end of the record years (Table 9). The GLM insurer holds the largest share in both markets, 0.411 with the Kelly underwriter and 0.281 without it, and the Kelly underwriter, on a smaller book than the GLM insurer's, writes its renewals at a higher price over their true expected loss than any traditional insurer (Table 8). The traditional insurers' rate levels end higher than in the market without it (Table 9). The book a traditional insurer keeps costs more against its plan, as the Bureau insurer's renewal selection shows (§8.1), and its method raises its level: the Bureau insurer refills 2.18 times a seed over the record years, against 1.39. The three insurers' growth is lower by 0.275, 0.246, and 0.217 than in the market without a Kelly underwriter (t of 22.69, 21.54, and 17.79), and it remains positive: each still earns more than the return on its surplus. In 4 of the 200 seeds the Bureau insurer holds under 1 percent of the accounts at the end of the record years. The capital standard does not bind in either market of this comparison: broker flight withholds no invitation, the IRIS stop closes no quarter, and no insurer fails; it binds mainly where an arm concentrates the accounts on one book (§10.1), and

among the arms where it binds are the learned-estimate seat's, three of Section 7, one of the rules for the load (Table F.1), and the market with a Bureau insurer that holds a fixed level (§10.2).

Table 9: The market with the Kelly underwriter or a second Loss-ratio insurer in the fourth seat, over 28 record years and 200 seeds. Rates per account and year are totals over seeds and years divided by the account-years; a challenger's quote is one by an insurer other than the holder on an account that went to market. Shares and rate levels are means at the end of the record years, and refilings the Bureau insurer's mean count a seed over the record years. The shares bound are in Table 8, and growth in Table 3 (its first and fourth rows). Source: authors' calculations.

Quantity	With the Kelly underwriter	Without it
<i>Per account and record year</i>		
Accounts changing insurer	0.394	0.141
Accounts going to market	0.491	0.266
Challengers' quotes	1.082	0.580
<i>Share of the accounts at the end of the record years</i>		
Bureau insurer	0.177	0.274
Loss-ratio insurer	0.161	0.222
GLM insurer	0.411	0.281
Fourth seat	0.250	0.223
<i>Rate level at the end of the record years</i>		
Bureau insurer	1.385	1.236
Loss-ratio insurer	1.378	1.259
GLM insurer	1.377	1.335
Bureau insurer's refilings a seed, record years	2.18	1.39

The two markets differ from the first record year (Figure 5). The practice years, in which estimates, rate levels, and books carry forward while surplus is reset, have sorted the accounts before the record years begin: the Kelly underwriter holds 0.386 of them at the close of the first record year and between 0.250 and 0.422 throughout. Its share falls in each of the final nine record years, and each traditional insurer's rises in each of the final six; the record does not attribute the movement. The traditional insurers' rate levels stand higher from the first year, at 1.299, 1.360, and 1.346 against 1.225, 1.272, and 1.271, and after it they part: the Bureau insurer's reaches its highest in the final record year, the GLM insurer's in year 22, and the Loss-ratio insurer's in year 16. The Loss-ratio insurer, whose complement of credibility reads competitors' rates, ends the record years with its level 0.018 above that at the close of the first year.




Figure 5: Account shares (upper row) and traditional insurers' rate levels (lower row) at the close of each record year (horizontal axis), means over 200 seeds; a level multiplies the bureau's loss costs, one at the practice years' start. Solid lines in the insurer's color are the market with the Kelly underwriter, dashed gray lines that with a second Loss-ratio insurer as fourth seat. Each row has one vertical scale; the lower row's fourth cell holds the legend. Source: authors' calculations.

9 Sensitivity

This section reports the sensitivity program for the leads and the two by two. The program is the 79 arms of Tables F.1 through F.3 in the appendix. Table 10 groups 75 of them: arms that change one condition of the market, one rule, or one threshold; two that set the read noise at another level; and arms that move one fixed constant or choice of Table 1, most to about half and to about one and a half of its value and the rest to one or two other values. The remaining four change the market in another way: two other occupants of the *first seat*, the seat the Bureau insurer holds in the base arm (§10.2), and the two changed inputs of Section 7. Every lead of the program is above zero at $t > 2$ but one, the lead over a Bureau insurer that holds a fixed level (§10.2). The two by two was also run at other settings (Table 11). Figure 6 shows the three leads in every arm of the program.

Table 10: The Kelly underwriter's lead over each traditional insurer across the 75 arms that change one condition, rule, or threshold, the read noise, or one fixed constant or choice, by group: the number of arms, the leads (three an arm) that are above zero at $t > 2$, and the lowest lead of the group by its t statistic, with its mean and the insurer it is over. Each arm runs the Kelly underwriter's market on seeds 0 through 199 with one assumption changed (Table 1 states the one exception); a lead is paired by seed. Source: authors' calculations.

Group of arms	Arms	Leads above zero at $t > 2$	Lowest lead	
			Mean (t)	Over
Conditions, the loads, the rules, the threshold	19	57 of 57	+0.091 (+5.80)	Bureau
Noise of a read of the account	2	6 of 6	+0.319 (+21.82)	GLM
Loss process, files, buyer's scale	18	54 of 54	+0.131 (+10.49)	GLM
Counts of years	15	45 of 45	+0.271 (+17.59)	Bureau
Claim sizes, report pattern, the return, the expense	9	27 of 27	+0.267 (+16.89)	Bureau
Other choices of Table 1	12	36 of 36	+0.083 (+6.54)	GLM

The 19 arms of Table 10's first group, of conditions, the loads, the rules, and the threshold, change the buyer's costs (recalibrated to annual switching of 0.11, and set to zero, at which 0.216 of the traditional market's accounts change insurer a year, short of the target of 0.25 the first registration named), the rate of remarketing unrelated to price, the rounding of the broker's figure, the claim rate, the spread of account size, how steeply invitations rise with size, a filed load on new business at 0.05 and at 0.1, the three other rules for the load, and the Bureau insurer's refiling threshold. In each, the three leads are above zero at $t > 2$ (57 of 57); the lowest, +0.091 ($t = +5.80$), is the lead over the Bureau insurer where account sizes are spread at 1.5 log standard deviations. A filed load raises the leads: where the traditional insurers file 0.05 on new business the lowest lead is +0.441, and at 0.1 it is +0.604, against +0.323 with no load. Under the rule in which a traditional insurer sets its load from its own indication times the level of its competitors' winning prices, the three charge -22.6, -22.5, and -24.0 percent in the largest band. In the table's two noise arms, at the two other levels of the noise in a read of the account, the leads are above zero with a lowest t of 21.82, and one registered prediction of those arms failed: the Kelly underwriter's growth at a read noise of 1.05 less its growth at 0.35, predicted above zero, is -0.002 ($t = -0.27$), and the record does not distinguish the growth at the two levels.

The 54 arms that move one fixed constant or choice, the table's final four groups, carry one registered prediction each, that the arm's three leads are above zero at $t > 2$, and every one held: 162 of 162 leads above zero, the lowest t 6.54, on the lead over the GLM insurer where the schedule band is 0.4 either way. The Kelly underwriter's own growth moves most with the band's width, 0.217 lower at the wider band, over which its deviations range, and with the frequency trait: higher by 0.213 at a yearly shock of 0.45 and by 0.095 at a persistence of 0.925, lower by 0.142 and 0.128 at a shock of 0.15 and a persistence of 0.425: the larger the yearly change in an account's claim frequency and the longer the change persists, the more the Kelly underwriter grows; the severity trait's constants move its growth the other way, except at a yearly shock of 0.45, where the change, +0.005 ($t = +0.91$), has the frequency trait's sign and is not distinguished from zero (Table F.2). No other constant of the group moves the growth by more than 0.092.




Figure 6: The Kelly underwriter's lead over a traditional insurer, one panel an insurer, in every sensitivity arm: its growth less that insurer's, paired by seed, an arm's mean on one horizontal scale with a bar of two standard errors either way. The arms, ordered by the lead, are gray, the base arm deep blue; the 79 arms of Tables F.1–F.3 include the two other occupants of the Bureau insurer's seat and Section 7's changed inputs. Zero is short-dashed; the fourth cell holds the legend. Source: authors' calculations.

The four arms of the fourth seat were also run at the Bureau insurer's two other refiling thresholds and at the two other levels of read noise. Table 11 reports the contrasts. The pattern of the *registered setting*, the base arm's refiling threshold and read noise, holds at every one of the five settings: learning alone above zero at $t > 2$, the Kelly decision with the quote-anchored credit on the learned estimate below zero at $t < -2$, and the interaction below zero at $t < -2$. The Kelly decision with the quote-anchored credit on the plan's estimate is below zero at $t < -2$ at none of the settings, and at two it is above zero at $t > 2$: at the refiling threshold of 0.15, with a t of 2.01, and at the read noise of 0.35, with a t of 3.58. At each of the four other settings the first registration predicted the registered setting's signs. Its prediction for learning alone, above zero, held at every one. Its predictions for the decision on the plan's estimate, below zero, for the decision on the learned estimate, above zero, and for the interaction, above zero, failed at every one, 12 predictions in all.

Table 11: The fourth seat’s two by two at five settings: the registered one, the Bureau insurer’s refiling threshold at two other values, and the noise of a read of the account at two other values. Each cell is a contrast in the fourth seat’s growth paired by seed over 200 seeds, the mean, with its t statistic in parentheses; the contrasts are those of Table 4. Source: authors’ calculations.

Setting	Learning	The Kelly decision on		Both	Interaction
	alone	the plan’s estimate	the learned estimate		
Registered setting	+0.400 (+31.10)	−0.003 (−0.24)	−0.244 (−31.11)	+0.156 (+11.25)	−0.241 (−17.36)
Threshold 0.05	+0.394 (+31.42)	−0.008 (−0.75)	−0.241 (−32.74)	+0.153 (+10.94)	−0.233 (−17.39)
Threshold 0.15	+0.391 (+31.57)	+0.021 (+2.01)	−0.232 (−28.87)	+0.159 (+11.90)	−0.253 (−19.40)
Read noise 0.35	+0.490 (+36.08)	+0.044 (+3.58)	−0.278 (−31.83)	+0.212 (+13.67)	−0.322 (−22.16)
Read noise 1.05	+0.384 (+32.36)	−0.002 (−0.21)	−0.227 (−29.54)	+0.157 (+11.11)	−0.224 (−16.33)

10 Discussion and Conclusion

This section states what the paper does not show, sets the market beside two others on the record, and concludes.

10.1 Limits

The market is one line of business with independent traits, no reinsurance, no underwriting cycle, and at most two entrants. It has no counterpart in which a Kelly underwriter has been observed, and the paper does not claim that its outcomes reproduce a market. The three traditional 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 any response to a competitor that learns, beyond the schedule credit and the complement of competitors’ rates, and an insurer in their place might acquire one. An insurer retains all of its surplus, so its ratio of written premium to surplus falls over the record years, and the capital standard binds mainly where an arm concentrates the accounts on one book.

Six properties of the results belong to this market or to the paper’s design. The filed rate band bounds the deviation policy’s renewal pricing (§6.4), and the record runs no deviation policy at another band; it varies the band once for the Kelly underwriter: at a band of 0.4 either way its growth is lower by 0.217, and its leads narrow, the lead over the GLM insurer to 0.083, the smallest mean of any lead in the 75 arms of Table 10. The variance term of the decision is small at the surplus of this market, so the record does not test the criterion where it differs most from the maximization of expected profit, in an insurer whose surplus is small against the accounts it writes. The comparison in growth of Section 6 is a statement at one calibration and at this record’s precision, and not a proof of zero (§6.4). The anchored credit stands below a credit at the realized scale, which biases the deviation policy’s deviations upward by an amount the record does not measure (§6.1). The discount of the credit is the construction’s choice, neither derivation being exact in this market (§6.2): on the standalone Kelly underwriter the return on funds gives lower growth than the insurer’s own growth with either form of the credit, and on the corrected deviation policy the record does not distinguish the effect of the discount from zero (§6.3). And the credit’s horizon runs through the final record year, a date the objective $\ln S_T$ supplies but no insurer outside a model has: the credit falls toward zero over the final record years (§4.4), and the record does not measure the leads with those years set aside or under a going-concern horizon.

Every prediction of the study was written down before the record that tested it ran; Appendix F states the registry and its counts.

10.2 Two Other Markets

The first seat holds the Bureau insurer of \$2.3, which refles its multiplier; two other first seats are on the record. Where the Bureau insurer holds a level of one throughout, it holds 0.687 of the accounts at the end of the record years. The Loss-ratio insurer's and the GLM insurer's levels end at 1.406 and 1.752, against 1.378 and 1.377 in the base arm; broker flight withholds 39,555 new-business invitations over the record years; and no insurer fails. The Kelly underwriter's growth falls to +0.430, against +0.691 beside a refiling Bureau insurer (the base arm's growth less this arm's, paired by seed, +0.261 at $t = +45.26$; Table F.1 prints the reverse difference). Its lead over the fixed-level insurer, -0.009 ($t = -0.57$), is the one lead of the sensitivity program not above zero, while its leads over the other two stand (+0.312 and +0.307). Where the first seat is a second Loss-ratio insurer, so that no insurer holds a fixed level, the three leads are +0.426, +0.429, and +0.330, each above zero.

10.3 Conclusion

A rating plan gives an insurer a price for an account. The Kelly criterion gives it an objective, the growth of its surplus, from which a price follows, and the objective requires, beyond the expected loss a plan also estimates, what a plan does not hold: the variance around it, what the account is worth if it is kept, and the chance that a quote wins. This paper has stated an approximation to that decision for an insurer that competes account by account, with each quantity it needs and an estimator for each that reads only what an insurer has—its own claim files, the loss runs of the accounts it is asked to quote, its own retention and growth, a broker's rounded figure on the quotes it lost, and what the market treats as public (Section 4)—and it has measured the decision and the learning apart, on one market and one set of seeds, under predictions each registered before the record that tested it ran.

Learning an account's loss distribution is worth more than the whole of the Kelly underwriter's gain over a second Loss-ratio insurer: the learned-estimate seat, which learns and prices by a plan's own methods, has the highest growth of the two by two, and it outgrows every traditional insurer by more than the Kelly underwriter does (Table 3). The Kelly decision with the quote-anchored credit has lower growth than that seat, against the registered prediction; part of the shortfall traces to the form in which the value of keeping an account enters the objective, a part measured on the registered Kelly underwriter itself and with the credit discounted at the return on funds. With the form corrected, the decision applied as a deviation from the plan's pricing of the learned estimate no longer trails that seat, and the Kelly underwriter that prices by the corrected decision alone still does, at either discount. On this record the growth of surplus is the measure, and at the registered threshold, at one calibration and at this record's precision, the corrected decision built on the learned estimate adds no growth the record distinguishes over a plan's pricing of the learned estimate; the quote data—renewal quotes gathered at the filed band's upper end, new business priced below the learned-estimate seat's price—and the shares and turnover around the two insurers tell apart which decision an insurer is running.

The market around the Kelly underwriter is a different market. Accounts move more often; every insurer's load on new business falls with account size, the traditional insurers' through schedule

credits that follow the winning prices they see; and the insurers that price by a plan keep renewals their estimates understate and win new business their estimates overstate, end at higher rate levels on the books they keep, and grow more slowly without failing. Whether insurers of that kind would keep to their plans is a question the paper does not answer; what separates the standalone Kelly architecture from the plan's pricing of the same estimate, after the correction, is not attributed on this record. What the paper supplies is the decision, the estimators, the correction and its limit, and a market in which all of them can be measured apart.

A 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

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

A claim's size before any limit is exponential with a mean drawn from the four means of §2.1 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.

B The Rating Plans

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 under the base rule for the load.

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 credit in the middle and the largest band is the exponential of the band's level less the smallest band's. A band's credit changes only where the band and the smallest band each hold 30 lost quotes in the window; otherwise it keeps its last credit. 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

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 Authorized Control Level. 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 6.42458.²⁴ A failed insurer's seat (§2.5) 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, and the band shifted upward and flight removed are arms of the record. 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 (§2.5) 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.²⁵ The stop removed is an arm of the record.

The endowment (§2.5) 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 (§3.2), 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 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 (§2.5). In a practice year no insurer fails; the restatement of the reserves at the close of the final practice year (§2.5) takes the provision for unreported claims at its true expectation, not conditioned on the claims already reported.

D The Estimators

The tracker'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

²⁴National Association of Insurance Commissioners, "Aggregated P&C RBC Data"; National Association of Insurance Commissioners, 2025 *Full Year Results*.

²⁵National Association of Insurance Commissioners, "NAIC IRIS Ratios: Property and Casualty, Data Year 2025" (NAIC, May 19, 2026).

severity observation is $\bar{y}_s - m_0$ less the class's mean log scale, with noise variance $(n_s v_0 + n_s^2 \sigma_c^2)/n_s^2$, 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 tracker has read, an account-year counting at each read.

Over every block of consecutive years the Kelly underwriter 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.

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 8 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 demand model, a segment's reference is the logarithm of the working premium over the tracked estimate, whose part that changes between fits is the renewal cell's factor for renewals and, for new business, the band's cell factor over the credit's divisor. An observation recorded at one reference is moved, at each fit, by the difference between that reference and the present one, the on-level step of §4.3. For a segment, renewals or new business in a band, let quote j have log deviation ℓ_j from the working premium and weight ω_j . The log likelihood is

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

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 §4.3 added, it is maximized in $(\mu_\zeta, \ln \sigma)$ by Newton's method, the covariance taken as §4.3 states. The pessimism variant of Section 6 lowers the conditional posterior mean of μ_ζ by one posterior standard deviation before the average of Equation 10 is taken, which lowers the hold probability at every deviation.

E The Kelly Underwriter'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 (Section 4).

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 Appendix D 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 $\hat{v}$ by Equation 8, and multiply both by the schedule modification.
3. Correct. Multiply both by the cell's factor: $\hat{x} = \kappa_c \hat{x}_0$, and $\hat{v}$ likewise, the cell being renewals for an account it holds and new business in the account's size band otherwise.
4. Set the working premium. $\hat{x}/(1 - \pi)$ on a renewal; on new business that amount divided by $1 + V_{nb}(1 - \varepsilon)/(1 - \pi)$, whatever the credit's sign.
5. Deviate. Take the deviation ℓ , on a grid over the filed band around the working premium, that maximizes $\bar{h}(\ell) G(P_\ell)$, with G of Equation 5 and $\bar{h}$ the posterior-averaged hold probability of the account's segment (Equation 10). Hold the deviation plus the logarithm of the schedule modification within $\ln(1 \pm 0.25)$.
6. Apply the floor. Quote the larger of the deviated working premium and the floor of Equation 7, 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 demand model, 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, each brought on-level to the segment's reference; 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 5's margin, first renewals and later ones apart.

At each year's close:

1. Update the cells by Equation 9 on the policies of the past three policy years, each cell's reported losses against the expected reported losses on $\hat{x}_0$.
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, m_1 and m_s as the margins per unit of net premium on each, and $\hat{g}$ as the window's mean yearly log growth of surplus, the return on funds included; set $\beta = e^{-\hat{g}}$, the horizon H as the years to the final record year, and V_{ren} and V_{nb} by Equation 6.

The fourth seat's two intermediate forms use parts of this list. The learned estimate at the plan's pricing carries out the first two steps of a quote and then prices as the Loss-ratio insurer does (Section 5). The Kelly decision on the plan's estimate replaces the first two steps by the plan's expected loss before any rate level, $L_c \approx M_{exp} M_{sch}$, and a variance from the class's claim frequency and the class's distribution of the severity trait, and carries out the rest. The deviation policy of Section 6 computes the first form's price whole, by the same computation, and carries out step 5 around it with the anchored credit, and step 6 with the floor written with the quote-anchored credit (§6.2), on that form's expected loss and variance at its rate level, without the cells.

F The Arms of the Record

Tables F.1–F.3 give arm by arm the sensitivity program that Table 10 reports by group, and the paragraph after them states the sets of the record, its seeds, and its registered predictions.

Table F.1: The arms that change one condition, one rule, one threshold, or one alternative of design: the Kelly underwriter’s growth less its growth in the base arm, +0.691, paired by seed, and its lead over each traditional insurer paired by seed, means with t statistics in parentheses, over 200 seeds. An arm’s label gives the value it moves an assumption to (Table 1 gives the base value); Section 7 describes the arms of settled-sample moments and of five-year files. Source: authors’ calculations.

Arm	Growth less the		Lead over (t)		
	base arm’s (t)	Bureau	Loss-ratio	GLM	
The base arm	—	+0.353 (+22.73)	+0.430 (+30.25)	+0.323 (+22.35)	
Buyer’s costs at annual switching of 0.11	+0.033 (+6.86)	+0.405 (+25.19)	+0.466 (+30.57)	+0.367 (+24.88)	
Buyer’s costs at zero, annual switching of 0.216	−0.046 (−9.80)	+0.309 (+22.26)	+0.379 (+28.10)	+0.268 (+21.52)	
Remarketing unrelated to price 0.05	−0.020 (−4.79)	+0.349 (+22.31)	+0.415 (+27.63)	+0.316 (+22.19)	
Remarketing unrelated to price 0.15	+0.026 (+5.99)	+0.397 (+26.72)	+0.450 (+32.75)	+0.324 (+24.16)	
Broker’s figure rounded to 0.1	+0.007 (+1.95)	+0.362 (+22.62)	+0.440 (+30.76)	+0.334 (+22.98)	
Broker’s figure exact	+0.005 (+1.50)	+0.366 (+24.35)	+0.428 (+30.14)	+0.327 (+23.85)	
Claims a year at the median size 0.5	−0.028 (−3.06)	+0.378 (+24.02)	+0.372 (+24.50)	+0.339 (+21.46)	
Claims a year at the median size 1.5	+0.008 (+1.00)	+0.331 (+24.66)	+0.435 (+37.92)	+0.323 (+26.62)	
Account size, log spread 0.5	−0.038 (−3.87)	+0.309 (+23.53)	+0.351 (+31.32)	+0.301 (+26.81)	
Account size, log spread 1.5	−0.086 (−11.18)	+0.091 (+5.80)	+0.196 (+13.27)	+0.351 (+22.03)	
Invitations, power of the rank 0.5	+0.180 (+31.80)	+0.549 (+33.92)	+0.597 (+38.41)	+0.518 (+33.24)	
Invitations, power of the rank 2	−0.147 (−31.12)	+0.217 (+16.32)	+0.273 (+20.62)	+0.229 (+18.21)	
Filed load on new business 0.05	+0.130 (+24.50)	+0.505 (+34.08)	+0.563 (+40.56)	+0.441 (+34.11)	
Filed load on new business 0.1	+0.287 (+42.55)	+0.691 (+50.10)	+0.723 (+53.73)	+0.604 (+48.40)	
Load: the filed load, no credit	+0.402 (+65.00)	+0.726 (+51.02)	+0.796 (+63.18)	+0.824 (+63.49)	
Load: the insurer’s own indication	+0.159 (+18.54)	+0.544 (+39.84)	+0.582 (+41.48)	+0.642 (+49.69)	
Load: indication and competitors’ level	−0.185 (−28.70)	+0.241 (+21.24)	+0.254 (+24.16)	+0.286 (+26.08)	
Refiling threshold 0.05	+0.021 (+5.55)	+0.417 (+27.72)	+0.430 (+30.52)	+0.329 (+22.93)	
Refiling threshold 0.15	−0.036 (−8.25)	+0.256 (+15.49)	+0.426 (+30.12)	+0.333 (+23.33)	
Noise of a read of the account 0.35	+0.004 (+0.66)	+0.405 (+24.76)	+0.462 (+31.51)	+0.319 (+21.82)	
Noise of a read of the account 1.05	+0.002 (+0.42)	+0.343 (+23.44)	+0.411 (+28.00)	+0.328 (+23.60)	
Bureau insurer at a fixed level	−0.261 (−45.26)	−0.009 (−0.57)	+0.312 (+31.75)	+0.307 (+26.97)	
A Loss-ratio insurer in the first seat	+0.021 (+5.02)	+0.426 (+28.93)	+0.429 (+29.75)	+0.330 (+24.98)	
Settled-sample moments	+0.139 (+15.43)	+0.541 (+55.36)	+0.616 (+74.67)	+0.498 (+54.73)	
Five-year files for all	+0.339 (+38.58)	+0.849 (+97.42)	+0.896 (+142.63)	+0.791 (+106.41)	

Table F.2: The arms that move one fixed constant of the loss process, the files, the buyer, the counts of years, the claim sizes, the report pattern, the return, or loss adjustment expense: the Kelly underwriter's growth less its growth in the base arm, +0.691, and its lead over each traditional insurer, both paired by seed, means with t statistics in parentheses, over 200 seeds. An arm's label gives the value it moves its constant to (Table 1 gives the base value); the shares of claims reported by the ends of years one through three are 0.325, 0.775, 0.925 on the slower report pattern and 0.775, 0.925, 0.975 on the faster. The loss-run arms keep at five years an incumbent's run from acquisition, and the severity-trait arms the claim-size factor at its base value. Source: authors' calculations.

Arm	Growth less the base arm's (t)	Lead over (t)		
		Bureau	Loss-ratio	GLM
The base arm	—	+0.353 (+22.73)	+0.430 (+30.25)	+0.323 (+22.35)
Frequency trait, spread 0.2	0.000 (−0.04)	+0.385 (+32.41)	+0.432 (+39.34)	+0.317 (+35.11)
Frequency trait, spread 0.6	0.000 (−0.02)	+0.341 (+17.33)	+0.412 (+23.78)	+0.320 (+17.52)
Frequency trait, persistence 0.425	−0.128 (−14.54)	+0.177 (+13.32)	+0.253 (+22.25)	+0.186 (+15.01)
Frequency trait, persistence 0.925	+0.095 (+10.21)	+0.492 (+28.66)	+0.556 (+34.90)	+0.510 (+33.00)
Frequency trait, yearly shock 0.15	−0.142 (−16.00)	+0.164 (+12.28)	+0.227 (+20.70)	+0.131 (+10.49)
Frequency trait, yearly shock 0.45	+0.213 (+23.73)	+0.695 (+47.74)	+0.737 (+58.83)	+0.653 (+47.19)
Severity trait, variance 0.2	+0.042 (+7.46)	+0.438 (+26.99)	+0.487 (+33.96)	+0.361 (+27.56)
Severity trait, variance 0.6	−0.024 (−4.81)	+0.332 (+21.25)	+0.395 (+29.28)	+0.299 (+23.11)
Severity trait, persistence 0.425	+0.068 (+11.91)	+0.446 (+27.80)	+0.516 (+37.46)	+0.405 (+29.08)
Severity trait, persistence 0.925	−0.043 (−6.58)	+0.357 (+29.05)	+0.402 (+35.65)	+0.306 (+26.16)
Severity trait, yearly shock 0.15	+0.051 (+9.79)	+0.420 (+28.34)	+0.493 (+34.03)	+0.373 (+26.93)
Severity trait, yearly shock 0.45	+0.005 (+0.91)	+0.409 (+30.90)	+0.454 (+38.12)	+0.365 (+30.39)
Case-reserve error, loss run 0.15	+0.012 (+2.77)	+0.371 (+23.33)	+0.439 (+29.94)	+0.332 (+22.57)
Case-reserve error, loss run 0.45	+0.001 (+0.29)	+0.366 (+23.56)	+0.421 (+29.50)	+0.321 (+22.92)
Case-reserve error, own file 0.075	+0.007 (+1.74)	+0.368 (+22.81)	+0.428 (+30.23)	+0.331 (+24.28)
Case-reserve error, own file 0.225	+0.002 (+0.43)	+0.360 (+22.71)	+0.432 (+32.22)	+0.324 (+24.37)
Buyer's response scale 0.025	−0.021 (−5.10)	+0.348 (+22.06)	+0.422 (+28.90)	+0.291 (+20.69)
Buyer's response scale 0.075	+0.027 (+5.79)	+0.375 (+25.24)	+0.442 (+31.06)	+0.372 (+29.25)
Bureau's statement, policy years 3	−0.007 (−1.64)	+0.347 (+21.78)	+0.424 (+29.38)	+0.325 (+23.33)
Bureau's statement, policy years 7	+0.006 (+1.43)	+0.366 (+24.23)	+0.438 (+31.95)	+0.317 (+24.17)
Bureau's statement, lag in years 3	+0.006 (+1.46)	+0.369 (+22.74)	+0.439 (+31.30)	+0.328 (+23.23)
Loss-ratio method's window, years 3	−0.017 (−3.99)	+0.350 (+23.60)	+0.415 (+30.77)	+0.309 (+22.03)
Loss-ratio method's window, years 7	+0.008 (+2.12)	+0.387 (+24.98)	+0.395 (+26.24)	+0.336 (+24.28)
Cells' window, years 2	−0.028 (−6.76)	+0.320 (+20.37)	+0.398 (+26.80)	+0.286 (+19.04)
Cells' window, years 4	+0.018 (+4.98)	+0.381 (+24.96)	+0.450 (+32.13)	+0.349 (+25.17)
Credit's window, years 2	+0.001 (+0.39)	+0.369 (+24.49)	+0.428 (+29.78)	+0.330 (+23.08)
Credit's window, years 4	+0.004 (+1.21)	+0.360 (+23.34)	+0.432 (+30.29)	+0.322 (+22.06)
Schedule fit's window, years 2	+0.006 (+1.52)	+0.360 (+24.27)	+0.438 (+30.94)	+0.331 (+22.82)
Schedule fit's window, years 4	−0.002 (−0.75)	+0.356 (+23.27)	+0.429 (+31.56)	+0.316 (+23.12)
Bureau insurer's experience, years 3	+0.022 (+4.88)	+0.484 (+30.33)	+0.409 (+27.85)	+0.304 (+22.50)
Bureau insurer's experience, years 7	−0.021 (−5.24)	+0.271 (+17.59)	+0.444 (+31.67)	+0.344 (+24.02)
Loss run, years 3	−0.085 (−12.87)	+0.272 (+20.47)	+0.338 (+26.98)	+0.227 (+17.74)
Loss run, years 7	+0.054 (+9.30)	+0.429 (+32.47)	+0.482 (+38.28)	+0.386 (+29.81)
Claim sizes, times 0.5	+0.020 (+3.25)	+0.465 (+33.15)	+0.489 (+37.47)	+0.398 (+32.66)
Claim sizes, times 1.5	−0.016 (−3.34)	+0.287 (+17.34)	+0.386 (+24.82)	+0.285 (+19.29)
Claim-size curve unscaled	+0.018 (+3.21)	+0.455 (+32.87)	+0.475 (+36.94)	+0.395 (+32.17)
Report pattern, slower	−0.062 (−11.66)	+0.267 (+16.89)	+0.335 (+23.12)	+0.241 (+17.41)
Report pattern, faster	+0.062 (+10.75)	+0.446 (+31.03)	+0.532 (+42.39)	+0.416 (+34.60)
Return on funds 0.0258	−0.064 (−14.66)	+0.329 (+20.66)	+0.399 (+27.21)	+0.330 (+21.72)
Return on funds 0.0365	+0.011 (+3.32)	+0.368 (+23.48)	+0.432 (+29.51)	+0.334 (+24.72)
Loss adjustment expense over indemnity 0.337090	+0.007 (+1.80)	+0.372 (+24.64)	+0.431 (+30.58)	+0.331 (+23.67)
Loss adjustment expense over indemnity 0.590157	−0.007 (−1.65)	+0.337 (+21.53)	+0.415 (+30.11)	+0.308 (+21.54)

Table F.3: The arms that move one of the other choices of Table 1, among them three that move or remove a rule of the capital standard: the Kelly underwriter’s growth less its growth in the base arm, +0.691, paired by seed, and its lead over each traditional insurer paired by seed, means with t statistics in parentheses, over 200 seeds. An arm whose mechanism the base arm never reaches (broker flight removed, the IRIS stop removed) reproduces the base arm’s record seed for seed, and its contrast has no t . Source: authors’ calculations.

Arm	Growth less the base arm’s (t)	Lead over (t)		
		Bureau	Loss-ratio	GLM
The base arm	—	+0.353 (+22.73)	+0.430 (+30.25)	+0.323 (+22.35)
Schedule band, either way 0.4	−0.217 (−37.81)	+0.132 (+10.03)	+0.179 (+13.76)	+0.083 (+6.54)
Experience plan’s K (\$) 100,000	+0.003 (+0.66)	+0.368 (+24.43)	+0.427 (+28.19)	+0.326 (+24.47)
Profit provision 0.02	−0.081 (−20.42)	+0.312 (+21.02)	+0.380 (+27.67)	+0.311 (+22.55)
Profit provision 0.06	+0.092 (+23.65)	+0.426 (+26.71)	+0.483 (+34.23)	+0.351 (+27.01)
Expenses, share of premium 0.258	−0.008 (−2.00)	+0.357 (+22.73)	+0.420 (+29.61)	+0.315 (+23.08)
Broker flight, band shifted upward	0.000 (+0.10)	+0.353 (+22.79)	+0.431 (+30.20)	+0.321 (+22.01)
Broker flight removed	identical	+0.353 (+22.73)	+0.430 (+30.25)	+0.323 (+22.35)
The IRIS stop removed	identical	+0.353 (+22.73)	+0.430 (+30.25)	+0.323 (+22.35)
Demand model’s prior spreads times 0.5	+0.005 (+1.10)	+0.369 (+24.42)	+0.431 (+30.19)	+0.332 (+25.11)
Demand model’s prior spreads times 1.5	+0.001 (+0.16)	+0.356 (+23.12)	+0.429 (+29.97)	+0.320 (+20.80)
Demand model’s decay, half-life halved	+0.001 (+0.38)	+0.363 (+23.84)	+0.424 (+29.71)	+0.317 (+22.13)
Demand model’s decay, half-life doubled	+0.007 (+2.54)	+0.366 (+24.02)	+0.435 (+31.14)	+0.334 (+23.69)

The record is 114 arms of 200 seeds each, 22,800 runs of the market, in nine sets: 34 arms that hold the two by two, the Kelly underwriter’s parts and alternatives, the conditions, the loads, the rules, the thresholds, and the two other first seats; 14 that hold the two by two’s three other arms at the two other thresholds and all four of its arms at two other levels of read noise; 18, 15, and 9 that move the fixed constants of the loss process, the files, the buyer, the counts of years, the claim-size curve, the report pattern, the return, and the loss adjustment expense; 12 that move the other choices of Table 1; the second registration’s 8 of the second construction (Section 6); 2 that run its corrected deviation policy on the growth discount; and 2 that run its corrected standalone Kelly underwriter on that discount (§6.2). Every run is a function of its arm and its seed alone, in double precision on one thread. The predictions were written before the runs that test them, with the rule that a contrast holds where its paired mean has the predicted sign at $|t| > 2$; the rule is applied to each contrast without adjustment for the number of contrasts tested, and at a Bonferroni level over the second registration’s 10 contrasts of Table 6, $|t| > 2.84$, no contrast of the ten moves across the threshold. There are 97 for the two by two, the market around the Kelly underwriter, and the sensitivity program (Sections 5, 8, and 9), registered after a disclosed reading of 20 seeds outside the record. There are 10 for the second construction. Eight were written after a disclosed reading of seeds outside the record: one, that its two arms of record reproduce their arms of Section 5 run for run; six on four directions first read on those seeds, the correction on the deviation policy, the corrected deviation policy’s three leads, pessimism on the corrected deviation policy, and the corrected deviation policy with pessimism against the seat it contains; and one, the funds discount on the Kelly underwriter, argued from the numeraire-portfolio property and read on those seeds before it was written. The other two were written before the corrected Kelly underwriter had run anywhere: the correction and the pessimism variant on it. And there are two for the corrected deviation policy’s arms on the growth discount and four for the corrected Kelly underwriter’s, each written before any run of the arms it tests. Of the 97, 80 held, 16 failed—among them the Kelly decision on the plan’s estimate, the Kelly decision on the learned estimate, and the two by two’s

interaction, at the registered setting and at every other—and one, on the Bureau insurer’s failures under a fixed level, is not tested, no failure having occurred in either arm. Of the 10, seven held and three failed: the funds discount of the credit, the pessimism variant on the corrected deviation policy, and the primary prediction, the corrected deviation policy with pessimism against the seat it contains. Both predictions for the corrected deviation policy on the growth discount failed: the discount’s effect on it and the primary prediction evaluated on that discount. Of the four for the corrected Kelly underwriter on the growth discount, three held, those on the correction on the registered Kelly underwriter, on the gap to the learned-estimate seat, and on the discount’s effect, and one failed, that on the pessimism variant.

G The Quote Data and the Loads by Band

Table G.1 gives the deviation policy’s quotes by arm and phase; Figure 3 draws their renewal shares at the band’s upper end and at the zero deviation (§6.4). Table G.2 gives the load on new business by insurer and size band, which Figure 4 draws (§8.1).

Table G.1: The deviation policy’s quotes over the record years and the 200 seeds, by arm (above the subheading, on the funds discount, the deviation policy with the quote-anchored credit, the corrected deviation policy, and the corrected deviation policy with pessimism of Table 5) and phase: the mean deviation from the contained seat’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; and, for the corrected arms, the posterior-mean hold probability at the chosen deviation (predicted hold). Below the subheading, the corrected arms of the third registration, 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
Quote-anchored	Renewals	+0.174	0.288	0.127	0.000	—
	New business	+0.021	0.105	0.169	0.000	—
Corrected	Renewals	+0.120	0.278	0.166	0.000	0.775
	New business	−0.081	0.039	0.136	0.000	0.708
Corrected, with pessimism	Renewals	+0.118	0.276	0.168	0.000	0.776
	New business	−0.088	0.037	0.133	0.000	0.714
<i>On the growth discount</i>						
Corrected	Renewals	+0.135	0.288	0.150	0.000	0.760
	New business	−0.057	0.039	0.146	0.000	0.680
Corrected, with pessimism	Renewals	+0.133	0.288	0.152	0.000	0.761
	New business	−0.063	0.037	0.145	0.000	0.686

Table G.2: The load on new business in percent by insurer and size band (thirds of the year's accounts by manual premium), with the Kelly underwriter or a second Loss-ratio insurer in the fourth seat, each entry a mean over seeds of seed means over the record years. A traditional insurer's load is its new-business quote over its plan premium, less one: its schedule credit within the filed band, which bounds the credit and the schedule modification together, with the smallest band the credit's reference. The Kelly underwriter's is its new-business quote over its renewal price for the estimate, less one. Source: authors' calculations.

Insurer	With the Kelly underwriter			Without it		
	Smallest	Middle	Largest	Smallest	Middle	Largest
Bureau insurer	0.0	−7.5	−12.2	0.0	−1.2	−2.9
Loss-ratio insurer	0.0	−7.4	−12.1	0.0	−1.2	−2.8
GLM insurer	0.0	−11.0	−16.3	0.0	−2.7	−2.5
Fourth seat	−26.2	−31.5	−36.8	0.0	−1.3	−2.9

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.
- 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.
- 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.
- 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.
- 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.
- 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.

- Long, John B., Jr. "The Numeraire Portfolio." *Journal of Financial Economics* 26, no. 1 (1990): 29–69.
- Longley-Cook, Laurence H. "An Introduction to Credibility Theory." *Proceedings of the Casualty Actuarial Society* 49 (1962): 194–221.
- 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.
- 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: Property and Casualty, Data Year 2025." NAIC, May 19, 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.
- 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.
- 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.
- 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.