

Algorithmic Learning in Competitive Insurance Underwriting

Frank Schmid and Can Wang

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

This version: 6 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 an insurer that prices by the Kelly criterion, and that learns what the criterion requires from its own claims and from the broker's word on the quotes it lost, in a stylized market for middle-market general liability beside three insurers that 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. The fourth seat of the market is filled in four ways, the plan's estimate or the learned estimate, priced as a plan prices or by the Kelly decision, so that the decision and the learning are measured apart and together. Learning alone raises the growth of surplus. The Kelly decision applied to the plan's estimate lowers it. Applied to the learned estimate it raises it, and the two together raise growth by more than the sum of their separate effects. Most of the Kelly underwriter's lead over the traditional insurers comes from one term of the decision, the value of keeping an account, which it estimates from its own retention and margin. It quotes new business at the lowest premium that covers its estimate of the first year's cost, below its own renewal price by more the larger the account. With such an insurer in the market, accounts change insurer more often, the traditional insurers' loads on new business fall with account size, and their rate levels rise. The lead holds in every arm of a program that moves the market's conditions and its fixed constants one at a time.

JEL classification: G22, D83, L11, C63.

This paper was developed with substantial 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	3
2.1 Accounts and Losses	3
2.2 What an Insurer Knows	4
2.3 Three Traditional Insurers	4
2.4 The Buyer and the Broker	6
2.5 Capital, Solvency, and the Measure	7
3 The Kelly Decision	8
3.1 The Objective	9
3.2 The Price of One Account	9
3.3 What Keeping an Account Is Worth	9
3.4 The Chance of Winning at a Price	10
4 Learning	11
4.1 The Account's Loss Distribution	11
4.2 The Business Won	11
4.3 The Demand Curve	12
4.4 What Was Learned	12
5 The Two Together	14
6 What Each Part Is Worth	16
7 The Market with a Kelly Underwriter	18
7.1 The Price of New Business	18
7.2 Turnover, Shares, and Rate Levels	20
8 Sensitivity	22
9 Discussion and Conclusion	25
9.1 Limits	25
9.2 Two Other Markets	25
9.3 Conclusion	26

A The Loss Process	27
B The Rating Plans	27
C The Capital Standard	28
D The Estimators	29
E The Kelly Underwriter's Computation, Step by Step	29
F The Arms of the Record	31
Bibliography	33

Abbreviations and Acronyms

GL	General liability
GLM	Generalized linear model
NAIC	National Association of Insurance Commissioners
RBC	Risk-based capital

Greek Letters

α	alpha	Permanent part of a trait (Appendix A)
δ	delta	Yearly discount factor of the renewal credit (§3.3)
ε	epsilon	Expenses as a share of premium (§2.3)
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)
λ	lambda	An account's expected number of claims in a policy year (§2.1)
μ	mu	Mean of where the market's price stands (§4.3)
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)
ρ	rho	Persistence of the drifting part of a trait (§2.1)
σ	sigma	Spread of where the market's price stands (§4.3); of a case reserve's error
Σ	Sigma	Covariance of a file's observations (Appendix D)
ϕ	phi	Standard normal density (Appendix D)
Φ	Phi	Standard normal distribution function (§4.3)
ω	omega	Weight of a quote in the demand model's likelihood (Appendix D)

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 among all strategies attains the highest long-run growth rate.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. None of the four is supplied by a rating plan, 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; the value of keeping an account from its own retention and margin; 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.

Four findings are on the record of 100 seeds, the same seeds in every arm. The measure is an insurer's growth, the logarithm of its surplus at the end of 28 years over its endowment, net of the return on surplus.

First, the decision and the learning are complements. In the fourth seat a second Loss-ratio insurer grows by +0.593. Learning alone, the learned estimate priced as that insurer prices, adds +0.495 ($t = +26.0$). The Kelly decision alone, on the plan's estimate, takes away: -0.190 ($t = -7.0$). The Kelly decision on the learned estimate adds +0.271 ($t = +11.9$) to learning alone, so that the Kelly underwriter grows by +1.359, and the interaction of the two is +0.461 ($t = +14.3$). The

¹Geoff Werner and Claudine Modlin, *Basic Ratemaking*, 5th ed. (Arlington, Virginia: Casualty Actuarial Society, 2016), chaps. 8 and 12; Ginda Kaplan Fisher, Lawrence McTaggart, Jill Petker, and Rebecca Pettingell, *Individual Risk Rating* (Casualty Actuarial Society study note, 2019).

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

Kelly underwriter leads the Bureau insurer, the Loss-ratio insurer, and the GLM insurer by +0.805, +0.905, and +0.853, and no insurer fails in any seed.

Second, one term of the decision accounts for most of the lead. Without the value of keeping an account, which we call the renewal credit, the Kelly underwriter's growth is lower by 0.533 and its leads over the Bureau insurer and the GLM insurer are not distinguishable from zero. Without its model of demand its growth is lower by 0.093; without its correction for the business won it is higher by 0.105. With all three removed, the learned estimate priced at the plan's provision with no rate level of its own, the insurer fails in 48 of 100 seeds.

Third, the market changes around such an insurer. The Kelly underwriter quotes new business below its own renewal price for the same estimate, by 15.6 percent in the middle third of accounts by size and by 26.0 percent in the largest third. The traditional insurers, which file a load of 10 percent on new business and may give a schedule credit against it, end up charging a load that falls with account size, by 6.8 points from the smallest third to the largest, against 0.3 points where the fourth seat is one of their own kind. Accounts change insurer at 0.209 a year against 0.062, the Kelly underwriter ends with 0.484 of the accounts, and the traditional insurers' rate levels end higher and their growth lower than in the market without it.

Fourth, the lead does not depend on the calibration. Each of the three leads is above zero at $t > 2$ in 20 arms that change one condition of the market, at two other levels of the noise in an underwriter's read of an account, and in 39 arms that move one of 20 fixed constants to about half and to about one and a half of its value; the lowest t among them is 6.4. The two by two keeps its four signs and the sign of its interaction at every one of the five settings at which it was run.

The paper draws on three literatures. The first is the Kelly criterion itself, whose growth-optimality Kelly and Breiman established for a gambler who knows the odds, and 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 first finding is a statement about that case: whether the decision adds growth or takes it away depends on the estimate it is given. 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 that pattern by valuing retention from its own book. 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

³Kelly, "A New 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; 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; 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 (2019): 87–135.

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); Stephen P. D'Arcy and Neil A. Doherty, "The Aging Phenomenon and Insurance Prices," *Proceedings of the Casualty Actuarial Society* 76 (1989).

learning algorithms raise prices.⁶ Here the learner prices against insurers that do not learn as it does, and what the record shows is a reallocation: lower prices on new business, higher prices on the renewals it keeps, and higher rate levels at the insurers that lose accounts to it.

The paper proceeds as follows. 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, and the two terms competition adds. Section 4 states what the Kelly underwriter learns, from what data, and by what estimator. Section 5 reports the central result, the fourth seat's two by two. Section 6 removes one part of the Kelly underwriter at a time. Section 7 describes the market with a Kelly underwriter against the market without one. Section 8 reports the sensitivity program, and Section 9 states the limits and concludes. The appendix carries the loss process, the rating plans, the capital standard, the estimators, the Kelly underwriter's computation step by step, and the arms of the record.

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 one 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 has a permanent 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 \$5,000, \$50,000, \$500,000, and \$5,000,000 at weights of 0.700, 0.250, 0.045, and 0.005, scaled by the account's severity trait.⁷ Indemnity is limited to \$1,000,000 an occurrence and \$2,000,000 in the aggregate of a policy year, and adjustment expense is 0.20 of the indemnity paid, in addition to the limit. 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, 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.

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.

⁷Clive 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).

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 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 last five policy years whoever wrote them, valued one quarter before the policy's inception, 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 cost of 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

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

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 $e/(1 - \pi)$ with a provision $\pi = 0.34$ for expenses of $\varepsilon = 0.30$ of premium and a profit of 0.04. The experience modification weights the account's capped losses against the class's by credibility: $M_{\text{exp}} = 1 + Z(A - E_{\text{rep}})/E$, with A the account's losses on the file the insurer reads, each occurrence capped, E the class's expected capped losses for the period, E_{rep} their expected reported part, and $Z = E/(E + K)$. The bureau's plan as filed uses three policy years, the latest year before the rating year left out, and K of \$100,000. An insurer applies it over every year its file holds, a year of age a weighted by d^a , 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.⁸

⁸Fisher and others, *Individual Risk Rating*; a limit of 25 percent either way is that of Florida Administrative Code, Rule 69O-170.004, among other state rules.




Figure 1: One year of the market. In each quarter the accounts that renew go to the holder and to the challengers the broker invites; each insurer quotes from the submission it reads; the buyer keeps its insurer or moves; an insurer whose quote did not bind is told by how much it missed; and the writer books premium, expenses, and a reserve. At the year's close each traditional insurer sets its rate level by its own rule, the Kelly underwriter re-estimates its cells and its renewal credit, and the capital standard decides failure and broker flight. Source: authors' construction.

The *Bureau insurer* files a multiplier on the bureau's loss costs and refiles it when its experience says so. At each year's close it computes the loss-ratio method's indication on its own experience of the last 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

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

with the complement c equal to one.⁹ 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 three policy years, with the complement of credibility set to competitors' rates: the level of the prices at which

⁹Laurence 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.

the accounts it quoted and lost were bound, relative to its own quotes, read from the broker's figure (§2.4) once it has 100 such figures in the window, 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, and takes each fitted factor at its classical credibility; the fitted overall level is its change of level.¹⁰

On new business, a quote on an account it does not hold, a traditional insurer charges a filed load of 10 percent. A load of this kind answers a documented fact, that new business runs at a higher loss ratio than renewals, and its size is set here within the range of the measured differences.¹¹ What an underwriter may move without a filing is the schedule modification, and insurers do not always implement the rate their experience indicates.¹² 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 last three years, with the credit and the schedule modification together kept inside the filed band. The smallest band carries the filed load and 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 (§7.1), and three other rules for the load are arms of the record (§8).

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 to one for the largest, so that the smallest account draws one competing quote and the largest draws all three. Let x 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 $1/(1 + \exp(-(x - k)/s))$ and, once there, stays with its holder with probability $1/(1 + \exp((x - k')/s))$, taking the lowest invited quote if it does not; $k = 0.10$ is the cost of shopping, $k' = 0.10$ the cost of switching, and $s = 0.05$ the scale of the response, each in units of log price.

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.

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

¹¹Wu and Lin, "Large Scale Analysis of Persistency and Renewal Discounts."

¹²Werner and Modlin, *Basic Ratemaking*, chap. 13; Lisa A. Hays, "Managing Commercial Lines Pricing Levels in a Loss Cost Environment," *Casualty Actuarial Society Forum* (Winter 2002).

2.5 Capital, Solvency, and the Measure

Every insurer starts the record with the same endowment, \$129.4 million, the industry's ratio of net premiums written to surplus of 0.773 applied to a quarter of the market's premium at \$100,000 an account.¹³ 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; surplus earns 0.03 a year and reserves 0.01. An insurer stops taking new business in a year in which its written premium reaches three times its surplus. At each year's close a risk-based capital (RBC) ratio is computed for each insurer, its surplus over a requirement that rises with its written premium and its reserves (Appendix C). Below a ratio of 2 the insurer receives no new business the following year; below 0.70, or with its surplus exhausted, it fails, the two thresholds being the company action level and the mandatory control level of the National Association of Insurance Commissioners (NAIC) standard.¹⁴ A failed insurer is replaced by an entrant, a Bureau insurer with the endowment, at most two times in a seed.

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 years of record. An insurer's *growth* is

$$g = \ln \frac{S_T}{S_0} - T \ln(1 + r), \quad T = 28, \quad r = 0.03, \quad (3)$$

the terminal logarithm of its surplus over the endowment, net of the return on surplus, so that an insurer that writes nothing has a growth of zero and g/T is a yearly rate of growth from underwriting. An insurer that fails has no growth in that seed and is left out of a mean over seeds; the tables state the seeds in which an insurer fails. Every arm of the record runs on seeds 0 through 99, 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.

¹³National Association of Insurance Commissioners, *U.S. Property & Casualty and Title Insurance Industries—2025 Full Year Results* (NAIC Financial Regulatory Services, 2026).

¹⁴National Association of Insurance Commissioners, *Risk-Based Capital (RBC) for Insurers Model Act*, Model 312 (2012).

Table 1: The assumptions of the market. Each is stated with its value in the base arm, its grounding, and the arms of the record that vary it; an arm moves one assumption and leaves the others. Dollar amounts are per occurrence or per policy year as stated; rates are per year. Sources are cited where the text introduces the assumption. 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: variance, persistence, shock	0.40, 0.85, 0.30	stated	six arms
Claim size	mixed exponential	Keatinge; Palmer	scale, two arms
Limits: occurrence; aggregate (\$)	1,000,000; 2,000,000	Palmer	no
Adjustment expense over indemnity	0.20	stated	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	stated	years, two arms
Full credibility, claims	1,082	Longley-Cook	no
Experience plan as filed: years; cap (\$); K (\$)	3; 100,000; 100,000	bureau plan's form	one arm
Schedule band, either way	0.25	state rules	no
Provision π ; expenses ε	0.34; 0.30	stated	no
Filed load on new business	0.10	Wu and Lin	three arms
The load's rule	schedule credit by band	stated	three arms
Loss-ratio method's window, years	3	stated	two arms
Bureau insurer: experience, years; threshold	5; 0.10	stated	four arms
Cost of shopping k ; of switching k'	0.10; 0.10	stated	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
Premium over surplus at the start	0.773	NAIC	no
Premium over surplus that stops new business	3	stated	no
Return on surplus; on reserves	0.03; 0.01	stated	four arms
RBC ratio: failure; broker flight	0.70; 2	NAIC	no
Entrants after failures, at most	2	stated	no
Renewal credit: discount rate; cap	0.10; 0.5	stated	four arms
Accounts; years of seeding, practice, record	4,000; 19, 10, 28	scale	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 price it gives for one account (§3.2), and the two terms that competition adds, what keeping an account is worth (§3.3) and the chance of winning at a price (§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. The Kelly criterion chooses the insurer's actions to maximize the expected logarithm of next period's surplus, $\mathbb{E}[\ln S']$. For a bettor who knows the distribution of outcomes, the strategy that does so 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.¹⁵ 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 is the criterion's for the distribution the insurer is given, and its property of optimal growth holds with respect to that distribution. How much growth it delivers on an estimate from a rating plan and on an estimate that is learned is the subject of Section 5.

3.2 The Price 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,

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

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. Writing the account adds to expected growth where $G(P) > 0$, that is, where

$$P \geq \frac{\hat{x} + \hat{v}/(2S)}{1 - \varepsilon}. \quad (5)$$

Equation 5 is the Kelly underwriter's *floor*: the lowest premium at which it writes an account. It reads three quantities, the account's expected loss, its variance, and the insurer's surplus, and it has two properties a plan's premium lacks. It rises with the account's variance, so that of two accounts with one expected loss the one with the wider distribution is quoted higher. And it falls as surplus grows, so that an insurer charges more for risk when its surplus is low. In the market of Section 2 the variance term is small. At the endowment of \$129.4 million it is 0.22 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 floor therefore lies close to the expected loss over $1 - \varepsilon$.

The Kelly underwriter's working price starts from the same expected loss at a plan's provision, the *manual premium* $\hat{x}/(1 - \pi)$, so that with nothing else learned it charges what a plan charges on its own estimate. The two terms below move it from there, and the floor bounds it from below.

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. Let a held account earn an expected margin m per unit of premium

¹⁵Breiman, "Optimal Gambling Systems"; Thorp, "The Kelly Criterion."

net of expenses in each year it is held, let an account be kept at its first renewal with probability p_1 and at each later renewal with probability p_s , and let δ be a yearly discount factor. To first order in the account's share of surplus, each future year adds its expected margin over surplus to expected growth, and the sum of the future years, per unit of this year's net premium, is

$$V_{\text{nb}} = \frac{m p_1 \delta}{1 - p_s \delta} \quad \text{and} \quad V_{\text{ren}} = \frac{m p_s \delta}{1 - p_s \delta} \quad (6)$$

for an account that is new to the insurer and for one it already holds. This is the *renewal credit*. With it the increment of Equation 4 becomes

$$G(P) = \frac{n - \hat{x} + r_f (n/2 + \hat{x}) + Vn}{S} - \frac{\hat{v}}{2S^2}, \quad (7)$$

where V is the credit of the account's phase and the term in $r_f = 0.01$ is an allowance for the return on the funds held against the policy, the return on reserves applied to half the net premium and to the expected loss. On new business the credit also enters the manual premium, which solves $P(1 - \pi) + V_{\text{nb}} P(1 - \varepsilon) = \hat{x}$: the insurer asks of a new account this year's provision less what the account is expected to return in later years. An account worth acquiring is thus quoted below the price the same estimate has as a renewal, the pattern the literature on dynamic contracts calls lowballing; here it is the consequence of an objective and of three quantities measured on the insurer's own book.

The growth objective contains no discount factor, and with $\delta = 1$ the credit would count every later year at its full expected margin. The Kelly underwriter uses a discount rate of 0.10, $\delta = 1/1.10$, and caps each credit at 0.5, so that a later year's margin, which depends on estimates of retention and margin that the market can move, counts for less than the current year's. Both constants are varied in Section 8.

3.4 The Chance of Winning at a Price

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

$$\ell^* = \arg \max_{\ell} h(\ell) G(P_\ell), \quad P = \max \left\{ P_{\ell^*}, \frac{\hat{x} + \hat{v}/(2S)}{1 - \varepsilon} \right\}. \quad (8)$$

A higher loading earns more on an account that is held and holds fewer accounts; the criterion weighs the two by the insurer's own hold curve. Because the Kelly underwriter rates on a schedule plan as the others do, its loading and its schedule modification together stay inside the filed band of 25 percent either way, so that ℓ^* is at most $\ln 1.25 = 0.223$ on an account whose schedule modification is one.

Equation 8 is the whole of the decision. Its inputs are $\hat{x}$ and $\hat{v}$ for each account, the credits V_{nb} and V_{ren} , and the hold curve h . A rating plan supplies a version of the first; 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 three quantities are sums over 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 none is given a constant of the loss process.

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 credibility that allows for a risk that changes over time.¹⁶ 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 adjustment expense, and the variance is that of a compound Poisson sum at the same rate and scale:

$$\hat{x}_0 = \hat{\lambda} \mathbb{E}[\min(Y, u)] (1 + a), \quad \hat{v} = \hat{\lambda} \mathbb{E}[\min(Y, u)^2] (1 + a)^2, \quad (9)$$

with $\hat{\lambda}$ the tracked claim rate, Y the claim size at the tracked scale, u the occurrence limit, and a the adjustment expense's share. The underwriter's schedule modification multiplies $\hat{x}_0$ as it does a plan's estimate. The variance is the quantity a plan does not hold, and Equation 5 is where it enters the price.

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

¹⁶Hans Bühlmann, "Experience Rating and Credibility," *ASTIN Bulletin* 4, no. 3 (1967): 199–207; Howard C. Mahler, "An Example of Credibility and Shifting Risk Parameters," *Proceedings of the Casualty Actuarial Society* 77 (1990): 225–308.

that won it.¹⁷ 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*, renewals and new business in each of the three size bands, 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 last 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

$$\kappa_c \leftarrow Z_c \frac{a_c}{e_c} + (1 - Z_c) \kappa_c, \quad (10)$$

and the expected loss in the price is $\hat{x} = \kappa_c \hat{x}_0$ for an account quoted in cell c . 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 Curve

The hold curve h of Equation 8 is learned from the only data on competitors' prices an insurer has, the broker's figure. The figure has a particular shape: it is given on the quotes that did not bind and withheld on those that did. Let z be the logarithm of the premium at which an account bound over the Kelly underwriter's manual premium for it. On a quote that was lost, z is observed, the quote's own loading less the figure. On a quote that bound, all that is known is that the market did not undercut it, $z \geq \ell$. The Kelly underwriter models where the market stands as normal, $z \sim N(\mu_z, \sigma^2)$, and estimates the two parameters by maximum likelihood from both kinds of observation, 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. The hold curve is then

$$h(\ell) = 1 - \Phi\left(\frac{\ell - \mu_z}{\sigma}\right), \quad (11)$$

with Φ the standard normal distribution function. It is fitted separately in each size band for renewals and for new business, six segments in all. Two safeguards bound the use of a curve that is known only near the loadings the insurer has tried: the loading maximizes Equation 8 with the hold probability taken one standard error toward the unfavorable side, and it moves by at most 0.02 a quarter from the segment's previous loading. Until a segment has 200 weighted observations, 30 of them lost quotes, the loading is zero and the quote is the manual premium.

4.4 What Was Learned

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

The cells' factors are above one in every cell: the losses of the business the Kelly underwriter writes run above the tracked estimates, by 43 percent on renewals and by 68, 43, and 24 percent on new business in the smallest, the middle, and the largest band. The factor on new business falls with account size. A loss run of five years carries more claims on a large account than on a small one, and the selection against the winner is weaker there: after the cells' correction the true expected loss of the new business the Kelly underwriter wins still exceeds its estimate, by 13.8 percent in the

¹⁷Edward 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.

smallest band and by 4.8 percent in the largest. In a single year a cell with few claims reaches a factor as high as 10.25.

On its retained renewals the Kelly underwriter earns a margin of 0.196 of premium net of expenses over its own expected loss; it keeps 0.596 of its accounts at their first renewal and 0.866 at each later one; and the credit averages 0.438 on new business and 0.495 on renewals against a cap of 0.5.

A credit of that size puts the price of new business at the floor. At the loading of the band's limit, the manual premium with the credit of Section 3.3 is below the floor of Equation 5 wherever V_{nb} exceeds 0.31, which it does in 0.89 of the years of record across the seeds. The premium the Kelly underwriter writes on new business is 1.43 times its own estimate of the expected loss, and 1.44, 1.44, and 1.43 in the three size bands, where the floor is $1/(1 - \varepsilon) = 1.43$ times the estimate before the variance term. The Kelly underwriter thus writes a new account at a premium that, net of expenses, covers its estimate of the first year's expected loss with the variance term and no more, and it earns its margin on the renewals it keeps. Above the threshold the size of the credit does not move the price of new business.

The loading stands at the limit of the filed band. In every band, on renewals and on new business, the demand model's loading averages 0.223 in logarithms, the band's limit of $\ln 1.25$. At the surplus of this market the variance term of Equation 7 is small against the margin, the expected increment rises in the loading up to the band's limit, and the hold curve learned from the broker's figures does not bring the optimum inside the band. What the demand model contributes in this market is therefore the finding that the manual premium is below the price that maximizes expected growth, and the move to the band's limit; it does not locate an interior price.

Table 2: What the Kelly underwriter learned in the base arm: means over the years of record and the 100 seeds. A cell's factor multiplies the tracked expected loss (Equation 10); the largest is over every year and seed. The renewal credit is per unit of premium net of expenses, on renewals and, in one figure for all bands, on new business (Equation 6). The loading is the demand model's, the logarithm of the loaded premium over the manual premium before the floor is applied, by band for renewals (smallest, middle, largest) and for new business; the filed band's limit is 0.223. Source: authors' calculations.

Quantity	Renewals	New business, by size band		
		Smallest	Middle	Largest
Cell's factor, mean	1.431	1.683	1.426	1.237
Cell's factor, largest in any year	2.77	10.25	3.53	3.12
Renewal credit, per unit of premium net of expenses	0.495		0.438	
Loading on the manual premium, logarithm	0.223, 0.223, 0.223	0.223	0.223	0.223

The tracker's fitted moments are of the kind the loss process has: at the end of the record it puts the persistence of the decaying part at 0.754 for frequency and 0.740 for the log scale, where the process has 0.85 for each trait in logarithms, and it finds a covariance between the two of -0.039 , 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 paper's central result: 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.

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 tracker's estimate of Section 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, the filed load on new business, and the schedule credit. With the plan's estimate and the Kelly decision it applies Equation 8, 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 100 seeds, and the fourth seat's share of the accounts at the end of the record. Growth is the terminal logarithm of surplus over the endowment, net of the return on surplus, over 28 years (Equation 3). No insurer fails in any seed of the four arms. 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	+0.765	+0.622	+0.815	+0.593	0.204
Learned estimate, plan's pricing	+0.453	+0.370	+0.534	+1.088	0.655
Plan's estimate, Kelly decision	+0.851	+0.711	+0.821	+0.403	0.062
Learned estimate, Kelly decision	+0.554	+0.454	+0.506	+1.359	0.484

Learning alone is worth +0.495 of growth ($t = +26.0$), above zero in 100 of 100 seeds. An insurer that tracks its accounts and otherwise prices as its competitors do ends the record with 0.655 of the accounts and binds 0.505 of its quotes as a challenger. The new business it wins has a true expected loss of 0.975 times its own estimate, where a second Loss-ratio insurer's has 1.132: 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 100 seeds: the mean, its t statistic, and the number of seeds in which the contrast is 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; both is the Kelly underwriter less a second Loss-ratio insurer; the interaction is both less the sum of learning alone and the Kelly decision on the plan's estimate. Source: authors' calculations.

Contrast	Mean	t	Seeds above zero
<i>The fourth seat's growth, one arm less another</i>			
Learning alone	+0.495	+26.0	100
The Kelly decision on the plan's estimate	-0.190	-7.0	24
The Kelly decision on the learned estimate	+0.271	+11.9	89
Both	+0.766	+24.3	97
Interaction	+0.461	+14.3	93
<i>The Kelly underwriter's growth less a traditional insurer's</i>			
Bureau insurer	+0.805	+18.4	94
Loss-ratio insurer	+0.905	+20.9	96
GLM insurer	+0.853	+21.9	97

The Kelly decision on the plan's estimate lowers growth, by 0.190 ($t = -7.0$). The decision requires more of its estimate than a plan's pricing does. It loads the price toward the limit of the filed band, and a price near the band's limit holds an account only where competitors price it higher still. The Kelly underwriter's estimate distinguishes those accounts and the plan's does not: the insurer that applies the decision to the plan's estimate keeps 0.580 of its renewal quotes where the Kelly underwriter keeps 0.814, and ends with 0.062 of the accounts. It does not fail in any seed, and its growth of +0.403 is positive; it is lower than the +0.593 of the insurer that prices the same estimate by the plan.

On the learned estimate the same decision raises growth, by +0.271 over learning alone ($t = +11.9$). The two together are worth +0.766, against +0.305 for the sum of the two taken alone, and the difference, the interaction, is +0.461 ($t = +14.3$), above zero in 93 seeds. The Kelly decision thus adds growth to a learned estimate and takes growth from a plan's. The growth-optimality of Section 3.1 is a property of the decision at the true distribution, and the record shows the sign of the decision's contribution differing between the two estimates.

Against the three traditional insurers in its own market the Kelly underwriter leads by +0.805, +0.905, and +0.853 (t of 18.4, 20.9, and 21.9), in 94, 96, and 97 of 100 seeds. The traditional insurers' own growth depends on what sits in the fourth seat (Table 3): it is highest beside the insurer that applies the Kelly decision to the plan's estimate, which cedes accounts to them, and lowest beside the learner that prices as they do or beside the Kelly underwriter.




Figure 2: Growth by insurer and seed in the four arms of the fourth seat. Each panel is one arm, the rows the pricing (the plan's above, the Kelly decision below) and the columns the estimate (the plan's on the left, the learned on the right). Within a panel each insurer has a row of 100 points, one a seed, with its mean shown by a vertical bar and stated at the right; the horizontal axis is growth, the terminal logarithm of surplus over the endowment net of the return on surplus over 28 years, on one scale in all panels, and the dashed vertical line is zero. Each insurer keeps one color in every figure: the Bureau insurer bright blue, the Loss-ratio insurer amber, the GLM insurer green, and the fourth seat deep blue. Source: authors' calculations.

6 What Each Part Is Worth

This section removes one part of the Kelly underwriter at a time, the renewal credit, the demand model, and the cells' correction, and then all three. Table 5 reports the fourth seat's growth in each arm, its difference from the Kelly underwriter's paired by seed, and its leads.

Table 5: The Kelly underwriter with one part removed, and with all three: the fourth seat's growth (mean over the seeds in which it survives), its growth less the Kelly underwriter's paired by seed, its lead over each traditional insurer paired by seed, each with its t statistic in parentheses, the seeds of 100 in which it fails, and its share of the accounts at the end of the record. Without all three the contrasts are read on the 52 seeds in which the insurer survives. Source: authors' calculations.

The fourth seat	Less the Kelly		Lead over (t)			Seeds failed	Share of accounts
	Growth	underwriter (t)	Bureau	Loss-ratio	GLM		
Kelly underwriter	+1.359	—	+0.805 (+18.4)	+0.905 (+20.9)	+0.853 (+21.9)	0	0.484
No renewal credit	+0.826	-0.533 (-22.4)	+0.028 (+0.7)	+0.181 (+4.3)	-0.004 (-0.1)	0	0.163
No demand model	+1.266	-0.093 (-3.8)	+0.995 (+38.0)	+1.063 (+49.3)	+1.007 (+37.7)	0	0.781
No cells' correction	+1.464	+0.105 (+4.0)	+1.391 (+72.4)	+1.417 (+69.2)	+1.443 (+69.0)	0	0.916
None of the three	-0.396	-1.830 (-16.5)	-0.506 (-4.2)	-0.480 (-3.8)	-0.450 (-3.7)	48	0.458

The renewal credit. Without the credit the Kelly underwriter's growth falls by 0.533 ($t = -22.4$), to +0.826. Its lead over the Loss-ratio insurer is +0.181 ($t = +4.3$); its leads over the Bureau insurer and the GLM insurer, +0.028 and -0.004, are not distinguishable from zero (t of 0.7 and -0.1). The credit is the part that makes the Kelly underwriter a competitor for accounts it does not hold: without it the insurer quotes new business above its own renewal price in the two smaller bands, by 25.6 and 24.0 percent, binds 0.283 of its quotes as a challenger where the Kelly underwriter binds 0.642, and ends with 0.163 of the accounts. It still learns each account's loss distribution and still charges the limit of the band on the accounts it holds: its renewals are written at 1.71 times their true expected loss, against 1.74 for the Kelly underwriter, on a smaller book.

The demand model. Without the demand model the loading is zero and the insurer quotes its manual premium. Its growth is lower by 0.093 ($t = -3.8$) and its three leads stay above zero with t of 37.7 or more. At the manual premium it holds more accounts, 0.781 of them at the end, at a lower margin: its renewals are written at 1.50 times their true expected loss against 1.74 for the Kelly underwriter.

The cells' correction. Without the cells the tracked estimate enters the price as it stands. The insurer's growth is higher by 0.105 ($t = +4.0$) than the Kelly underwriter's, and it ends with 0.916 of the accounts. On this record the correction for the business won lowers the growth of the insurer that applies it. The correction raises the Kelly underwriter's price, by the factors of Table 2; the insurer without it prices lower, holds most of the accounts, and at the loading of the band's limit still writes its renewals at 1.51 times their true expected loss. The three traditional insurers, left with 0.084 of the accounts between them, end with growth of +0.074, +0.048, and +0.021.

All three. With the credit, the demand model, and the cells all removed, what remains is the tracked estimate at a plan's provision with no rate level and no correction. That insurer fails in 48 of 100 seeds, and in the seeds in which it survives its growth is -0.396. The comparison with learning alone, which has the same estimate and grows by +1.088, isolates what a level tied to experience does: the Loss-ratio insurer's method on the tracked estimate, or the cells, or the loading, each keeps the insurer's price in line with the losses of the business it wins. The insurer with none of them writes its renewals at 1.32 times their true expected loss, a loss ratio of 0.748 beside expenses of 0.30 of premium.

7 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 price of new business by account size (§7.1), and turnover, shares, and the traditional insurers' rate levels and growth (§7.2).

7.1 The Price of New Business

Table 6 reports the load each insurer charges on new business by size band, and Figure 3 draws it. For a traditional insurer the load is one plus the filed 10 percent, times its schedule credit, less one. For the Kelly underwriter it is its quote on a new account over the price it would quote the same estimate as a renewal, less one.

The Kelly underwriter's load is negative and falls with size: -1.7 , -15.6 , and -26.0 percent in the smallest, the middle, and the largest band. In the smallest band it is not distinguishable from zero across seeds ($t = -0.5$); in the middle and the largest it is below zero in every reading (t of -10.7 and -46.1), and the smallest band's load exceeds the largest's by 24.3 points ($t = +7.4$). Two learned quantities produce the gradient. The renewal credit puts the price of new business at the floor in every band (Section 4.4), where the renewal price of the same estimate is at the limit of the filed band. The cells' factors, which multiply the estimate in the price, are higher on new business in the smaller bands than in the largest (Table 2), and that raises the price of new business on small accounts toward the renewal price. The load is a comparison with the insurer's own renewal price: against the true expected loss of the business written, the Kelly underwriter's price of new business is 1.26 in the smallest band and 1.37 in the largest, and its renewal price 1.65 and 1.77.

The traditional insurers' loads fall with size where they compete with the Kelly underwriter; where they compete with one another the fall is 0.3 points. With the Kelly underwriter in the market the three charge 3.4, 3.9, and 2.4 percent in the largest band against the filed 10; the load in the smallest band exceeds that in the largest, on the mean of the three, by 6.8 points ($t = +23.5$). Without it the same difference is 0.3 points ($t = +3.6$), and the largest band's load is higher by 6.5 points ($t = +21.8$). The schedule credit follows the winning prices an insurer sees, and winning prices on large accounts are low where an insurer prices large accounts on fuller information. A load on new business that falls with account size by more than that appears in this market where a competitor prices large accounts on fuller information; the number of competing quotes, which rises with size in both markets, produces a fall of 0.3 points.

Table 6: The load charged on new business, in percent, by insurer and size band (thirds of the year's accounts by manual premium), with the Kelly underwriter in the fourth seat and with a second Loss-ratio insurer there: the mean over the seeds of a seed's mean over its years of record. A traditional insurer's load is its filed load times its schedule credit, less one; the Kelly underwriter's is its quote on new business over its renewal price for the same estimate, less one. The smallest band carries the filed load by construction. Source: authors' calculations.

Insurer	With the Kelly underwriter			Without it		
	Smallest	Middle	Largest	Smallest	Middle	Largest
Bureau insurer	+10.0	+6.4	+3.4	+10.0	+9.6	+10.0
Loss-ratio insurer	+10.0	+6.8	+3.9	+10.0	+9.9	+10.2
GLM insurer	+10.0	+2.7	+2.4	+10.0	+7.4	+8.8
Fourth seat	-1.7	-15.6	-26.0	+10.0	+9.8	+10.2




Figure 3: The load charged on new business by size band. Each panel is one insurer; the horizontal axis is the size band (thirds of the year's accounts by manual premium, smallest to largest) and the vertical axis the load in percent, on one scale in all panels, with the filed load of 10 percent drawn as a dotted line and zero as a dashed line. The solid line in the insurer's color is the market with the Kelly underwriter and the dashed gray line the market with a second Loss-ratio insurer in the fourth seat; in the fourth seat's panel the dashed line is that insurer's load. Each point is the mean over 100 seeds of a seed's mean over 28 years of record. Source: authors' calculations.

Table 7 sets the two phases of a policy side by side. The price is the premium written over the true expected loss of the business written, which the simulation knows and no insurer does; the selection is the true expected loss over the writer's own estimate; the share bound is, on new business, the share of an insurer's quotes as a challenger that bind and, on renewals, the share of its renewal quotes that keep the account.

The Kelly underwriter writes new business at 1.36 times its true expected loss and renewals at 1.74, a loss ratio of 0.724 on new business and 0.570 on renewals. The traditional insurers, in the market without it, write both phases at about one price, 1.49 and 1.49 for the Bureau insurer, and at loss ratios within 2.6 points of one another. A new-business loss ratio above the renewal loss ratio, the pattern measured in insurers' books, is in this market the pattern of the insurer that values retention: its size, 15.4 points, is the distance between new business written at the floor and renewals written at the limit of the band. On every insurer's new business the selection is above one: the business an insurer wins runs above its own estimate, by 10.4 percent for the Bureau insurer in the market without a Kelly underwriter.

Table 7: New business and renewals by insurer, sums over the years of record and the 100 seeds, with the Kelly underwriter in the fourth seat and with a second Loss-ratio insurer there. Price: premium written over the true expected loss of the business written. Loss ratio: ultimate losses over premium. Selection: true expected loss over the writer's own estimate of it. Bound: on new business, quotes bound over quotes made as a challenger on accounts that went to market and were placed; on renewals, renewals kept over 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	1.50	0.651	1.034	0.191	1.44	0.671	1.050	0.775
Loss-ratio insurer	1.47	0.670	1.062	0.210	1.44	0.677	1.053	0.757
GLM insurer	1.44	0.681	1.032	0.238	1.43	0.680	1.063	0.773
Fourth seat	1.36	0.724	1.054	0.642	1.74	0.570	1.020	0.814
<i>Without it</i>								
Bureau insurer	1.49	0.660	1.104	0.147	1.49	0.658	1.020	0.942
Loss-ratio insurer	1.46	0.675	1.133	0.166	1.48	0.663	1.025	0.933
GLM insurer	1.43	0.684	1.103	0.216	1.47	0.665	1.028	0.943
Fourth seat	1.46	0.672	1.132	0.153	1.49	0.661	1.019	0.929

7.2 Turnover, Shares, and Rate Levels

Table 8 compares the two markets, and Figure 4 draws the shares and the rate levels over the years of record.

The market turns over more. Accounts change insurer at 0.209 per account and year with the Kelly underwriter and at 0.062 without it, and go to market at 0.313 against 0.173. The Kelly underwriter binds 0.642 of the quotes it makes as a challenger; the traditional insurers bind 0.191, 0.210, and 0.238, more than the 0.147, 0.166, and 0.216 they bind against one another. The model's shares of quotes bound are outcomes; no figure for middle-market general liability was found to set them against.

Table 8: The market with the Kelly underwriter in the fourth seat and with a second Loss-ratio insurer there, over 28 years of record and 100 seeds. Rates per account and year are sums over the seeds and years; a challenger's quote is a quote by an insurer other than the holder on an account that went to market. The share of quotes bound is over an insurer's quotes as a challenger on accounts that went to market and were placed. Shares of the accounts and rate levels are means at the end of the record. Growth is the mean over seeds (Equation 3). Source: authors' calculations.

Quantity	With the Kelly underwriter	Without it
<i>Per account and year of record</i>		
Accounts changing insurer	0.209	0.062
Accounts going to market	0.313	0.173
Challengers' quotes	0.696	0.366
<i>Share of its quotes as a challenger that an insurer binds</i>		
Bureau insurer	0.191	0.147
Loss-ratio insurer	0.210	0.166
GLM insurer	0.238	0.216
Fourth seat	0.642	0.153
<i>Share of the accounts at the end of the record</i>		
Bureau insurer	0.153	0.237
Loss-ratio insurer	0.174	0.235
GLM insurer	0.189	0.324
Fourth seat	0.484	0.204
<i>Rate level at the end of the record</i>		
Bureau insurer	1.387	1.219
Loss-ratio insurer	1.358	1.242
GLM insurer	1.314	1.146
<i>Growth</i>		
Bureau insurer	+0.554	+0.765
Loss-ratio insurer	+0.454	+0.622
GLM insurer	+0.506	+0.815
Fourth seat	+1.359	+0.593
Bureau insurer's refilings a seed, years of record	3.33	2.67

The Kelly underwriter ends the record with 0.484 of the accounts, and the three traditional insurers with 0.153, 0.174, and 0.189, against 0.237, 0.235, and 0.324 in the market without it. Their rate levels end higher, at 1.387, 1.358, and 1.314 against 1.219, 1.242, and 1.146. The book a traditional insurer keeps costs more against its plan, the true expected loss of the Bureau insurer's renewals standing at 1.050 times its own estimate with the Kelly underwriter in the market and at 1.020 without, and its method raises its level. The Bureau insurer refiles 3.33 times a seed over the years of record, against 2.67. The three insurers' growth is lower by 0.211, 0.168, and 0.308 than in the market without a Kelly underwriter (t of 9.5, 7.9, and 14.8), and it remains positive: each still earns more than the return on its surplus, and none fails. In two of the 100 seeds the Bureau insurer ends the record with under one percent of the accounts.

The two markets are apart from the first year of record (Figure 4). The practice years, in which estimates, rate levels, and books carry forward while surplus is reset, have sorted the accounts before the record opens: the Kelly underwriter holds 0.450 of them at the close of the first year of record and between 0.450 and 0.498 throughout. The traditional insurers' rate levels stand higher from the first year, at 1.310, 1.348, and 1.190 against 1.201, 1.229, and 1.128, and those of the Bureau

insurer and the GLM insurer go on rising through the record, while the Loss-ratio insurer's, which gives weight to competitors' rates, ends the record within 0.010 of where it began it.




Figure 4: Shares of the accounts and rate levels over the 28 years of record, means over 100 seeds. Upper row: each insurer's share of the accounts at the year's close. Lower row: each traditional insurer's rate level at the year's close, its multiplier on the bureau's loss costs, relative to one at the start of the practice years; the fourth cell of the lower row is the reading guide. In every panel the solid line in the insurer's color is the market with the Kelly underwriter and the dashed gray line the market with a second Loss-ratio insurer in the fourth seat. The panels of a row share one vertical scale. Source: authors' calculations.

8 Sensitivity

This section asks whether the results depend on the calibration. The record holds three kinds of arm beyond those above: arms that change one condition of the market or one rule, arms that move one fixed constant, and the two by two at other settings. Table 9 summarizes the first two; Tables F.1 and F.2 in the appendix list every arm; and Figure 5 shows the three leads in every arm.

Conditions and rules. The 20 arms of the first group change one condition or one rule: the buyer's costs of shopping and switching, 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, the filed load at zero, at half, and at one and a half of its value, the three other rules for the

load, and the Bureau insurer's threshold. In each, the three leads are above zero; the lowest, +0.213 ($t = +6.4$), is the lead over the Bureau insurer under the rule in which a traditional insurer sets its load from its own experience times the level of its competitors' winning prices, a rule under which the traditional insurers also charge a negative load in the largest band, of -16.1 , -15.5 , and -21.2 percent. The leads rise with the filed load: with no filed load the lowest lead is +0.291, and it rises with the load to +1.137 at a load of 0.15. At the two other levels of the noise in a read of the account the leads are above zero with a lowest t of 16.6.

Fixed constants. Of the assumptions that have neither a source nor an arm among the conditions, 20 were each moved to about half and to about one and a half of their values, one at a time; Table 1 marks those that no arm moves. The 20 are the six constants of the two traits, the two case-reserve errors, the scale of the buyer's response, and the renewal credit's discount rate and cap; the lengths in years of the bureau's statement and its lag, of the loss-ratio method's window, of the Bureau insurer's experience, and of the loss run; and the scale of the claim-size curve, the report pattern, and the two rates of return. A persistence, which cannot be one and a half times 0.85, was moved to half its value and to half its distance from one, and the bureau's lag, which cannot be shorter than two years in this market, was moved to three. In all 39 arms each lead is above zero at $t > 2$, 117 leads of 117, the lowest t is 13.5, on the lead over the GLM insurer where the frequency trait's yearly shock is 0.15, and no insurer fails in any seed. The renewal credit, which Section 6 shows to account for most of the lead, does not depend for that on its two constants, since over their range the price of new business stays at the floor or close to it (with the cap at 0.25 the premium on new business is 1.47 times the estimate): at a discount rate of 0.05 and of 0.15 and at a cap of 0.25 and of 0.75 the lowest of the twelve leads is +0.794. The Kelly underwriter's own growth moves most, among these arms, with the frequency trait: it is higher by 0.365 at a yearly shock of 0.45 and by 0.300 at a persistence of 0.925, and lower by 0.227 and 0.162 at a shock of 0.15 and a persistence of 0.425; at half the scale of claim sizes it is lower by 0.206 (Table F.2).

Table 9: The Kelly underwriter's lead over each traditional insurer across the sensitivity arms, 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 99 with one assumption changed; a lead is paired by seed. Source: authors' calculations.

Group of arms	Arms	Leads above zero		Lowest lead	
		at $t > 2$		Mean (t)	Over
Conditions, the load's rule, the threshold	20	60 of 60		+0.213 (+6.4)	Bureau
Noise of a read of the account	2	6 of 6		+0.827 (+16.6)	Bureau
Loss process, files, renewal credit, buyer's scale	22	66 of 66		+0.466 (+13.5)	GLM
Counts of years	9	27 of 27		+0.687 (+15.8)	Bureau
Claim sizes, report pattern, returns	8	24 of 24		+0.708 (+13.5)	Bureau

Two alternatives of design. With five years of history and the bureau's plan as filed in place of the estimated experience plan and the longer file, and with the second moments of the loss process given to the tracker in place of its estimates of them, the leads are above zero (Table F.1).

The two by two at other settings. The four arms of the fourth seat were run at the Bureau insurer's two other thresholds and at the two other levels of read noise. Table 10 reports the contrasts.

Learning alone is above zero, the Kelly decision on the plan's estimate below zero, the Kelly decision on the learned estimate above zero, and the interaction above zero at each of the five settings; the smallest t in absolute value among the 25 contrasts is 4.3.

Table 10: 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 100 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	+0.495 (+26.0)	-0.190 (-7.0)	+0.271 (+11.9)	+0.766 (+24.3)	+0.461 (+14.3)
Threshold 0.05	+0.503 (+29.2)	-0.160 (-5.2)	+0.315 (+13.3)	+0.818 (+26.8)	+0.475 (+13.8)
Threshold 0.15	+0.456 (+22.4)	-0.133 (-4.6)	+0.272 (+11.0)	+0.728 (+20.8)	+0.405 (+11.5)
Read noise 0.35	+0.533 (+23.7)	-0.182 (-6.0)	+0.199 (+7.4)	+0.732 (+21.4)	+0.381 (+9.5)
Read noise 1.05	+0.494 (+25.0)	-0.137 (-4.3)	+0.332 (+16.5)	+0.826 (+26.2)	+0.469 (+13.5)




Figure 5: The Kelly underwriter's lead in every sensitivity arm. Each panel is the lead over one traditional insurer, the Kelly underwriter's growth less that insurer's, paired by seed: one point an arm, the mean lead on the horizontal axis with a bar of two standard errors either way, the arms ordered by the size of the lead and the base arm drawn in deep blue. The 65 arms are those of Tables F.1 and F.2; the vertical line is zero, and the horizontal scale is common to the three panels. The fourth cell is the reading guide. Source: authors' calculations.

9 Discussion and Conclusion

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

9.1 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.

Three properties of the Kelly underwriter's result are properties of this market. Its loading is at the limit of the filed band throughout and its price of new business is at the floor, so that its two prices are the two bounds of Equation 8 and neither is an interior optimum; a wider or a narrower band would move the price it charges on the accounts it holds. The variance term of the decision is small at the surplus of this market, so that 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. And the smallest size band serves as the reference of the traditional insurers' schedule credit, a third of the accounts, where the eligibility rules of schedule rating plans exempt fewer.¹⁸

The Kelly decision on the plan's estimate lowers growth, and the cells' correction lowers the Kelly underwriter's growth relative to the insurer without it. Neither changes the central result, the interaction: at every setting tested the decision is worth more on the learned estimate than on the plan's. The second means that the Kelly underwriter as specified is not the fastest-growing insurer of its kind in this market: without the cells' correction it prices lower and grows by 0.105 more.

Every prediction of the study was written down before the run that tested it: the signs of the two by two and the leads in the base arm, and the leads in each sensitivity arm, 82 predictions in all, of which 82 held (Appendix F).

9.2 Two Other Markets

The Bureau insurer of Section 2.3 refiles its multiplier. Two other first seats are on the record. Where the Bureau insurer holds a level of one throughout, it fails in 24 of 100 seeds with the Kelly underwriter in the market, and the Kelly underwriter's lead over it in the seeds in which both survive is +0.967 ($t = +10.4$). Where the first seat is a second Loss-ratio insurer, so that no insurer holds a fixed level, the three leads are +0.949, +0.935, and +0.878. In each the leads are above zero.

¹⁸Florida Administrative Code, Rule 69O-170.004, sets eligibility at a manual premium of \$1,000.

9.3 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 it requires what a plan does not hold: the account's expected loss and the variance around it, what the account is worth if it is kept, and the chance that a quote wins. This paper has stated 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 a broker's rounded word on the quotes it lost. In a stylized market for general liability the insurer that both learns and decides in this way outgrows a bureau-based insurer, a loss-ratio insurer, and an insurer with a generalized linear model, in the base calibration and in every arm of a sensitivity program that moved the market's conditions and its unsourced constants.

The decision and the learning do not contribute alike, and the difference is the paper's main result. Learning an account's loss distribution raises growth under a plan's pricing. The Kelly decision lowers growth on a plan's estimate and raises it on a learned one, and the two together are worth more than the sum of the two apart. An insurer that adopts the criterion without the learning prices to the limit of what its filing allows on accounts it cannot tell apart, and loses them; an insurer that learns without the criterion wins the accounts a plan overprices and prices them as a plan would. Of the terms of the decision, the one that matters most in this market is the value of keeping an account. With it an insurer that knows its accounts competes for accounts it does not yet know: it quotes them at the lowest premium that covers its estimate of their first year's cost, below the price it will charge once it holds them, and the difference is larger on larger accounts, where its correction for the business won is smaller.

The market around such an insurer is a different market. Accounts move more often; the load on new business falls with account size at every insurer, including those whose filed load does not vary with size; and the insurers that price by a plan lose the accounts their plans overprice, raise their 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 it supplies is the decision, the estimators, and a market in which the two can be measured apart.

A The Loss Process

An account i has a size z_i , lognormal with median one 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 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 Section 2.1 at their weights and multiplied by the account's severity trait; the curve's mean before any limit is \$63,500, and limited to \$1,000,000 an occurrence it is \$39,987. Within a policy year the aggregate limit of \$2,000,000 erodes in the order in which the claims are paid. Adjustment expense is 0.20 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 adjustment expense.

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, uniformly over the quarters of its year of report. 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

Classes and the bureau's statement. 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.

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 a (the latest year of the period at age zero, the period ending two years before the rating year) weighted by $w_a = d^a$,

$$M_{\text{exp}} = 1 + \frac{E}{E + K} \cdot \frac{A - E_{\text{rep}}}{E}, \quad A = \sum_a w_a A_a, \quad E = B_c z \sum_a w_a, \quad E_{\text{rep}} = B_c z \sum_a w_a F_a, \quad (\text{B.1})$$

with A_a the year's capped reported losses on the file and F_a 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 \$100,000.

The schedule modification. With (m_1, s_1^2) the mean and variance of the permanent traits' sum in the account's class and (m_2, s_2^2) its mean and variance given the class and the underwriter's reading, the modification is $\exp(m_2 - m_1 + (s_2^2 - s_1^2)/2)$, held within 1 ± 0.25 .

The GLM insurer's fit. On the policies it wrote in the five policy years ending at the year's close, the insurer 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 , and what is taken out is the fitted change of level. The level, each class relativity, and each band relativity are then multiplied by $Zf + 1 - Z$, with f 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.

The schedule credit on new business. For each band, from its quotes as a challenger of the last three years on accounts that went to market and were placed, the insurer takes 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 with fewer than 30 lost quotes in the window keeps its last credit. The same fit over all bands, net of the load, gives the Loss-ratio insurer's complement of credibility 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 ratio

$$\text{RBC} = \frac{S}{(W + 3\rho_Q Q)/(0.773 \times 6.19 \times (1 + 0.86))}, \quad \rho_Q = \frac{0.86}{3 \times 3} = 0.0956. \quad (\text{C.1})$$

The requirement in the denominator is scaled so that an insurer at the industry's ratio of premium to surplus, 0.773, with reserves of three times its premium, has the industry's aggregate ratio of 6.19; 0.86 is the reference ratio of the charge on reserves to the charge on premium. An insurer whose ratio is below 2 at a year's close takes no new business in the following year. An insurer whose ratio is below 0.70, or whose surplus is not positive, fails: its seat passes to an entrant, a Bureau insurer at a level of one with the endowment and no book, while entries remain. Within a year an insurer takes no new business in a quarter that it enters with written premium for the year of three times its surplus or more. In a practice year no insurer fails and every surplus is reset to the endowment at the year's close.

Surplus changes within a quarter by the premium written less expenses of 0.30 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 at 0.03 a year on surplus and 0.01 a year on reserves.

D The Estimators

The tracker's observations. For an account-year s in a file let n_s be the reported claims, 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. The frequency observation is n_s/q_s , with noise variance $1/q_s$ around the year's relativity; the severity observation is $\bar{y}_s - m_0$ less the class's mean log scale, with noise variance $(n_s v_0 + n_s^o \sigma_c^2)/n_s^2$, n_s^o the open claims among them. A year without a reported claim carries no severity observation.

The second moments. 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.

The predictor. For an account with a 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 9 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.

The demand model. For a segment, renewals or new business in a band, let quote j have log loading ℓ_j over the manual premium and weight ω_j . The log likelihood is

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

with z_j the loading less the broker's figure in logarithms and ϕ the standard normal density, maximized in $(\mu_z, \ln \sigma)$ by Newton's method. The standard error of the logit of $h(\ell)$ is taken from the inverse Hessian by the delta method.

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.

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 9, and multiply $\hat{x}_0$ by the schedule modification.
3. Correct. Multiply by the cell's factor: $\hat{x} = \kappa_c \hat{x}_0$, the cell being renewals for an account it holds and new business in the account's size band otherwise.
4. Set the manual premium. $\hat{x}/(1 - \pi)$ on a renewal; on new business that amount divided by $1 + V_{nb}(1 - \varepsilon)/(1 - \pi)$.
5. Load. If the segment's demand model has its minimum of observations, take the loading on the segment's grid, within 0.02 of the segment's mean loading of the previous quarter in steps of 0.0025, that maximizes $h_-(\ell) G(P_\ell)$ where $G > 0$ and $h_+(\ell) G(P_\ell)$ where it is not, with G of Equation 7 and h_- and h_+ the hold probability one standard error of its logit below and above the fit. Otherwise the loading is zero. Hold the loading plus the logarithm of the schedule modification within $\ln(1 \pm 0.25)$.
6. Apply the floor. Quote the larger of the loaded manual premium and $(\hat{x} + \hat{v}/(2S))/(1 - \varepsilon)$, 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 loading, a renewal kept without going to market among them, and each quote lost on an account that went to market with its loading and the broker's figure. Decay the weights by the quarterly factor of 0.98 and refit 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 its margin over $\hat{x}$.

At each year's close.

1. Update the cells by Equation 10 on the policies of the last three policy years, each cell's reported losses against the expected reported losses on $\hat{x}_0$.
2. Update the renewal credit. Over the last three years, m is the summed margin over the summed net premium of the renewals kept, p_1 the first renewals kept over the first renewals quoted, and p_s the later renewals kept over the later renewals quoted. With $\delta = 1/1.10$ and $p_s \delta$ held at or below 0.98, set V_{ren} and V_{nb} by Equation 6, each held between zero and 0.5.

The fourth seat's two intermediate forms use parts of this list. Learning alone carries out the first two steps of a quote, multiplies $\hat{x}_0$ by a rate level that it moves as the Loss-ratio insurer moves its own, and prices at the plan's provision with the filed load and the schedule credit. The Kelly decision on the plan's estimate replaces the first two steps by the plan's expected loss before any rate level, $L_c z M_{exp} M_{sch}$, and a variance from the class's claim frequency and the class's distribution of the severity trait, and carries out the rest.

F The Arms of the Record

Table F.1: The arms that change one condition, one rule, or one alternative of design: the Kelly underwriter's growth less its growth in the base arm, +1.359, paired by seed, and its lead over each traditional insurer paired by seed, means with t statistics in parentheses, over 100 seeds. An arm's label gives the value it moves an assumption to (Table 1 gives the base value). Where the Bureau insurer holds a fixed level it fails in 24 seeds and the lead over it is read on the other 76. Source: authors' calculations.

Arm	Growth less the base arm's (t)	Lead over (t)		
		Bureau	Loss-ratio	GLM
The base arm	—	+0.805 (18.4)	+0.905 (20.9)	+0.853 (21.9)
Costs of shopping and switching 0.05	−0.072 (−3.5)	+0.722 (15.9)	+0.796 (17.1)	+0.806 (19.7)
Costs of shopping and switching 0.15	+0.047 (+2.5)	+0.863 (19.3)	+0.945 (20.5)	+0.871 (20.9)
Remarketing unrelated to price 0.05	−0.079 (−4.7)	+0.739 (17.1)	+0.844 (18.5)	+0.808 (20.5)
Remarketing unrelated to price 0.15	+0.090 (+5.4)	+0.922 (21.9)	+1.005 (24.1)	+0.929 (23.7)
Broker's figure rounded to 0.10	+0.035 (+2.3)	+0.868 (20.2)	+0.942 (23.8)	+0.890 (24.4)
Broker's figure exact	+0.033 (+2.5)	+0.878 (21.5)	+0.936 (20.5)	+0.894 (23.4)
Claims a year at the median size 0.50	−0.052 (−1.5)	+0.839 (16.6)	+0.872 (20.8)	+0.950 (24.4)
Claims a year at the median size 3	−0.170 (−6.6)	+0.491 (14.7)	+0.774 (25.2)	+0.631 (21.0)
Account size, log spread 0.50	+0.043 (+1.3)	+0.894 (22.5)	+0.947 (25.7)	+0.953 (25.5)
Account size, log spread 1.50	−0.288 (−9.4)	+0.350 (8.0)	+0.570 (12.9)	+0.540 (13.9)
Invitations, power of the rank 0.50	+0.260 (+12.2)	+1.163 (25.2)	+1.203 (27.1)	+1.130 (27.0)
Invitations, power of the rank 2	−0.212 (−9.5)	+0.584 (15.2)	+0.679 (16.8)	+0.671 (18.4)
Filed load on new business 0	−0.532 (−24.4)	+0.291 (7.0)	+0.405 (9.1)	+0.438 (10.5)
Filed load on new business 0.05	−0.270 (−11.9)	+0.528 (11.5)	+0.637 (12.3)	+0.628 (13.6)
Filed load on new business 0.15	+0.273 (+15.9)	+1.137 (27.2)	+1.208 (31.8)	+1.141 (30.1)
Load: the filed load, no credit	+0.332 (+15.7)	+1.197 (32.0)	+1.231 (37.4)	+1.252 (37.4)
Load: the insurer's own indication	−0.009 (−0.4)	+0.792 (18.3)	+0.921 (20.4)	+0.962 (22.4)
Load: indication and competitors' level	−0.727 (−32.1)	+0.213 (6.4)	+0.313 (8.4)	+0.391 (11.0)
Refiling threshold 0.05	+0.059 (+3.4)	+0.928 (21.6)	+0.952 (22.2)	+0.873 (21.4)
Refiling threshold 0.15	−0.051 (−2.9)	+0.744 (15.6)	+0.883 (19.7)	+0.844 (20.9)
Noise of a read of the account 0.35	−0.050 (−2.2)	+0.827 (16.6)	+0.909 (20.6)	+0.772 (17.4)
Noise of a read of the account 1.05	+0.047 (+2.4)	+0.873 (23.1)	+0.960 (25.0)	+0.858 (23.3)
Bureau insurer at a fixed level	−0.628 (−28.9)	+0.967 (10.4)	+0.538 (18.9)	+0.599 (19.4)
A Loss-ratio insurer in the first seat	+0.036 (+1.9)	+0.949 (24.5)	+0.935 (22.6)	+0.878 (22.4)
Second moments known to the tracker	−0.012 (−0.4)	+0.800 (25.4)	+0.903 (27.7)	+0.814 (29.5)
Five years of history, plan as filed	+0.290 (+10.3)	+1.309 (54.4)	+1.350 (59.7)	+1.314 (58.6)

The record is 88 arms of 100 seeds each, 8,800 runs of the market, in five panels: 35 arms that hold the two by two, the Kelly underwriter with parts removed, the two alternatives of design, the conditions and rules, the three rules without a Kelly underwriter, the Bureau insurer's two other thresholds, and the two other first seats; 14 that hold the two by two's three other arms at those two thresholds and its four arms at two other levels of read noise; and 22, 9, and 8 that move the fixed constants. 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 $|t| > 2$: 20 for the first panel, written after a reading of twenty other seeds, and 23, 22, 9, and 8 for the others, written with no reading of their arms. All 82 held. Two counts were left unvaried and are stated as such: the numbers of lost quotes an insurer needs before it uses the competitors' level as a complement and before it changes

a band's schedule credit, 100 and 30. The simulation code, the record of every run, and the scripts that produce every table and figure from them are available from the authors.

Table F.2: The arms that move one fixed constant: the Kelly underwriter's growth less its growth in the base arm, +1.359, paired by seed, and its lead over each traditional insurer paired by seed, means with t statistics in parentheses, over 100 seeds. An arm's label gives the value it moves the constant to (Table 1 gives the base value); the slower report pattern has 0.325, 0.775, and 0.925 of claims reported by the end of the first three years and the faster 0.775, 0.925, and 0.975. Source: authors' calculations.

Arm	Growth less the	Lead over (t)		
	base arm's (t)	Bureau	Loss-ratio	GLM
The base arm	—	+0.805 (18.4)	+0.905 (20.9)	+0.853 (21.9)
Frequency trait, spread 0.20	+0.032 (+1.0)	+0.895 (21.2)	+0.983 (26.6)	+0.858 (21.6)
Frequency trait, spread 0.60	+0.034 (+1.2)	+0.829 (15.7)	+0.959 (21.3)	+0.891 (20.1)
Frequency trait, persistence 0.425	−0.162 (−4.7)	+0.603 (18.7)	+0.698 (21.9)	+0.534 (18.0)
Frequency trait, persistence 0.925	+0.300 (+10.9)	+1.276 (32.2)	+1.355 (35.9)	+1.282 (38.1)
Frequency trait, yearly shock 0.15	−0.227 (−7.4)	+0.507 (14.0)	+0.602 (19.0)	+0.466 (13.5)
Frequency trait, yearly shock 0.45	+0.365 (+13.4)	+1.447 (47.4)	+1.510 (54.1)	+1.415 (47.1)
Severity trait, variance 0.20	+0.015 (+0.7)	+0.821 (17.4)	+0.943 (20.9)	+0.851 (20.2)
Severity trait, variance 0.60	−0.032 (−1.4)	+0.802 (19.5)	+0.862 (19.0)	+0.813 (21.1)
Severity trait, persistence 0.425	+0.124 (+5.4)	+0.961 (20.5)	+1.032 (22.9)	+0.972 (20.8)
Severity trait, persistence 0.925	+0.101 (+4.1)	+1.001 (25.9)	+1.105 (33.7)	+1.001 (27.1)
Severity trait, yearly shock 0.15	+0.061 (+3.0)	+0.860 (18.3)	+0.937 (22.0)	+0.880 (21.4)
Severity trait, yearly shock 0.45	+0.087 (+3.8)	+1.016 (25.2)	+1.076 (25.5)	+1.044 (30.9)
Case-reserve error, loss run 0.15	+0.020 (+1.1)	+0.838 (19.8)	+0.919 (20.0)	+0.878 (21.5)
Case-reserve error, loss run 0.45	+0.029 (+1.6)	+0.873 (21.8)	+0.930 (20.6)	+0.856 (21.5)
Case-reserve error, own file 0.075	+0.023 (+1.4)	+0.849 (21.9)	+0.930 (23.7)	+0.869 (24.7)
Case-reserve error, own file 0.225	−0.014 (−0.8)	+0.810 (18.3)	+0.863 (18.8)	+0.846 (20.1)
Renewal credit, discount rate 0.05	+0.024 (+1.4)	+0.852 (19.7)	+0.924 (20.1)	+0.893 (23.4)
Renewal credit, discount rate 0.15	+0.008 (+0.5)	+0.823 (19.5)	+0.923 (22.3)	+0.860 (23.3)
Renewal credit, cap 0.25	+0.001 (+0.1)	+0.794 (17.6)	+0.896 (19.9)	+0.834 (21.6)
Renewal credit, cap 0.75	+0.002 (+1.5)	+0.807 (18.3)	+0.908 (21.2)	+0.855 (21.9)
Scale of the buyer's response 0.025	−0.027 (−1.3)	+0.812 (16.7)	+0.874 (18.0)	+0.833 (18.7)
Scale of the buyer's response 0.075	+0.016 (+0.9)	+0.834 (21.8)	+0.903 (23.7)	+0.869 (24.9)
Bureau's statement, policy years 3	+0.024 (+1.4)	+0.861 (20.2)	+0.955 (22.9)	+0.895 (23.2)
Bureau's statement, policy years 7	+0.003 (+0.2)	+0.823 (19.0)	+0.886 (21.3)	+0.845 (22.1)
Bureau's statement, lag in years 3	+0.026 (+1.5)	+0.870 (21.2)	+0.928 (22.5)	+0.884 (23.1)
Loss-ratio method's window, years 2	−0.106 (−5.4)	+0.729 (16.5)	+0.756 (17.8)	+0.778 (18.9)
Loss-ratio method's window, years 4	+0.081 (+4.7)	+0.909 (21.1)	+0.987 (23.6)	+0.919 (21.8)
Bureau insurer's experience, years 3	+0.084 (+4.4)	+1.038 (20.4)	+0.964 (21.3)	+0.881 (20.4)
Bureau insurer's experience, years 7	−0.062 (−3.2)	+0.687 (15.8)	+0.882 (21.1)	+0.854 (21.3)
Loss run, years 3	−0.078 (−3.2)	+0.783 (18.6)	+0.836 (22.0)	+0.704 (17.6)
Loss run, years 7	+0.078 (+3.5)	+0.902 (23.1)	+1.002 (28.0)	+0.933 (27.3)
Claim sizes, times 0.5	−0.206 (−9.4)	+0.890 (26.8)	+0.917 (27.2)	+0.904 (28.5)
Claim sizes, times 1.5	+0.097 (+4.6)	+0.708 (13.5)	+0.835 (16.7)	+0.739 (17.0)
Report pattern, slower	−0.079 (−3.3)	+0.739 (16.6)	+0.771 (17.9)	+0.708 (18.5)
Report pattern, faster	+0.120 (+5.4)	+1.014 (22.5)	+1.069 (23.6)	+1.009 (23.9)
Return on surplus 0.015	+0.142 (+25.1)	+0.878 (18.1)	+0.975 (19.9)	+0.916 (20.8)
Return on surplus 0.045	−0.118 (−22.4)	+0.754 (18.7)	+0.851 (21.4)	+0.801 (22.4)
Return on reserves held 0.005	−0.081 (−12.9)	+0.840 (19.3)	+0.933 (22.1)	+0.907 (22.4)
Return on reserves held 0.015	+0.088 (+12.6)	+0.794 (17.4)	+0.908 (19.7)	+0.827 (21.4)

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).
- Breiman, Leo. "Optimal Gambling Systems for Favorable Games." *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability* 1 (1961): 65–78.
- 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).
- 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.
- Hays, Lisa A. "Managing Commercial Lines Pricing Levels in a Loss Cost Environment." *Casualty Actuarial Society Forum* (Winter 2002).
- Jia, Ruo, and Zenan Wu. "Insurer Commitment and Dynamic Pricing Pattern." *Geneva Risk and Insurance Review* 44 (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. "A New Interpretation of Information Rate." *Bell System Technical Journal* 35, no. 4 (1956): 917–926.
- Kunreuther, Howard, and Mark Pauly. "Market Equilibrium with Private Knowledge: An Insurance Example." *Journal of Public Economics* 26, no. 3 (1985): 269–288.
- Longley-Cook, Laurence H. "An Introduction to Credibility Theory." *Proceedings of the Casualty Actuarial Society* 49 (1962): 194–221.
- 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. *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.
- Werner, Geoff, and Claudine Modlin. *Basic Ratemaking*. 5th ed. Arlington, Virginia: 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).