

CORNELL FINANCIAL ENGINEERING MANHATTAN

The Price of Having Been Right

Capital Lock-In and Recovery in Costly Markets

Jan Vecer

Department of Probability and Mathematical Statistics
Charles University, Prague

Cornell CFEM seminar, October 15, 2026

The Price of Having Been Right

I study the behavior of agents in prediction markets with transaction costs. Agents use fractional-Kelly strategies based on competing probabilistic forecasts. In the frictionless benchmark, prices are wealth-weighted averages of forecasts, and wealth updates depend on likelihood ratios. Successful forecasts therefore come to dominate the price, making the market accurate but slow to adapt: after a change, repricing can take time proportional to the preceding regime's length.

Shrinking betting fractions can preserve learning while removing this linear lock-in when trading is free. With a per-contract fee, the power-law schedules studied instead produce linear recovery times, and faster shrinkage makes recovery slower. I identify a sharp lower bound on the delay when betting fractions never increase.

I then show how a policy combining shrinking baseline fractions with reversible increases in response to evidence of change restores sublinear recovery without new capital, while controlling false alarms.

I close with questions about how execution, fee incidence, and inventory constraints affect the speed at which prediction-market prices absorb new information.

Cornell seminar announcement: *The Price of Having Been Right: Capital Lock-In in Prediction Markets.*

A belief becomes a bet, and a bet becomes a likelihood ratio

Consider a complete betting market with no fees. Cash is the numeraire; ω denotes the outcome. All probabilities are positive.

Market probabilities $P_Y(\omega)$. A claim paying one dollar only in outcome ω costs $P_Y(\omega)$. A winning dollar staked returns gross odds $1/P_Y(\omega)$.

Investor beliefs $P_M(\omega)$. Allocate fraction $P_M(\omega)$ of initial wealth $M(0)$ to outcome ω . The fractions sum to one.

Only the bet on the realized outcome pays. Therefore

$$M(1, \omega) = \underbrace{M(0)P_M(\omega)}_{\text{stake}} \underbrace{\frac{1}{P_Y(\omega)}}_{\text{gross odds}},$$

$$\boxed{\frac{M(1, \omega)}{M(0)} = \frac{P_M(\omega)}{P_Y(\omega)}}.$$

This is the **likelihood ratio**. We will prove that it is log-optimal.

Soccer: three outcomes, one portfolio

Start with $M(0) = \$1,000$. Stake your own probabilities on win, draw and loss.

Outcome	Market P_Y	Belief P_M	Gross odds	Stake	Final wealth
Win	0.50	0.60	2	\$600	\$1,200
Draw	0.30	0.30	10/3	\$300	\$1,000
Loss	0.20	0.10	5	\$100	\$500

For a win: $\$600 \times 2 = \$1,200$, so $M(1)/M(0) = 1.2 = 0.60/0.50$.

The three wealth multiples are (1.2, 1, 0.5): exactly P_M/P_Y .

Even the loss outcome gives positive final wealth. Zero wealth on a possible outcome would give log wealth $-\infty$.

Vecer (2026), *Analytical solution for Kelly's criterion for multiple outcomes*, *Quantitative Finance* 26(5), 671–684.
[doi:10.1080/14697688.2026.2629952](https://doi.org/10.1080/14697688.2026.2629952).

Expected logarithmic performance is a measure of information

Take the expectation under the investor's beliefs P_M . Since wealth is a likelihood ratio,

$$\mathbb{E}_M \left[\log \frac{M(1)}{M(0)} \right] = \sum_{\omega} P_M(\omega) \log \frac{P_M(\omega)}{P_Y(\omega)} = D_{\text{KL}}(P_M \| P_Y).$$

D_{KL} is **Kullback–Leibler divergence**, also called **relative entropy**. It measures the expected log evidence for P_M against P_Y . It is nonnegative and is zero exactly when the distributions agree.

For the soccer portfolio,

$$0.60 \log(1.2) + 0.30 \log(1) + 0.10 \log(0.5) \approx \boxed{0.04008 \text{ nats}}.$$

Natural logarithms give nats per match. This is actual expected growth if P_M is the true law; otherwise it is the investor's assessment.

The inverse: a payoff also defines a probability distribution

Choose any positive payoff $V(1, \omega)$. Its initial price and associated distribution are

$$V(0) = \sum_{\omega} P_Y(\omega) V(1, \omega), \quad \boxed{P_V(\omega) = P_Y(\omega) \frac{V(1, \omega)}{V(0)}}.$$

The entries of P_V sum to one. Rearranging gives $V(1, \omega)/V(0) = P_V(\omega)/P_Y(\omega)$.

Soccer example: raise the loss-state payoff from \$500 to \$800.

$V(1) = (1200, 1000, 800)$ costs $V(0) = 0.5(1200) + 0.3(1000) + 0.2(800) = 1060$. Hence

$$P_V = \frac{(600, 300, 160)}{1060} \approx (0.5660, 0.2830, 0.1509).$$

P_V is implied by the payoff and market prices. It need not equal the investor's beliefs P_M or the true outcome probabilities.

Why the likelihood ratio is log-optimal

Compare any positive payoff V at its own cost $V(0)$. Its wealth multiple is P_V/P_Y . Insert P_M into that ratio:

$$\begin{aligned}\mathbb{E}_M \log \frac{V(1)}{V(0)} &= \sum_{\omega} P_M(\omega) \log \left(\frac{P_M(\omega)}{P_Y(\omega)} \frac{P_V(\omega)}{P_M(\omega)} \right) \\ &= D_{\text{KL}}(P_M \| P_Y) - D_{\text{KL}}(P_M \| P_V) \\ &\leq D_{\text{KL}}(P_M \| P_Y).\end{aligned}$$

The subtracted divergence is nonnegative. Equality holds **exactly when** $P_V = P_M$, giving $V(1)/V(0) = P_M/P_Y$.

Soccer: the \$800-loss payoff gives $0.0288 = 0.0401 - 0.0113$ nats. Its growth loss is exactly $D_{\text{KL}}(P_M \| P_V) \approx 0.0113$.

Vecer (2026), *Quantitative Finance*. doi:10.1080/14697688.2026.2629952.

Two opinions determine an equilibrium pricing measure

Two investors have beliefs P_1, P_2 and initial wealth $M_1(0), M_2(0)$. Both use full Kelly. At common no-fee prices P_E , their optimal final wealth is

$$M_1(1, \omega) = \frac{P_1(\omega)}{P_E(\omega)} M_1(0), \quad M_2(1, \omega) = \frac{P_2(\omega)}{P_E(\omega)} M_2(0).$$

Their trades must match. With no fees, total wealth is conserved in **every outcome**, so

$$M_1(1, \omega) + M_2(1, \omega) = \frac{M_1(0)P_1(\omega) + M_2(0)P_2(\omega)}{P_E(\omega)} = M_1(0) + M_2(0).$$

$$\boxed{P_E(\omega) = uP_1(\omega) + (1 - u)P_2(\omega)}, \quad u = \frac{M_1(0)}{M_1(0) + M_2(0)}.$$

The pricing measure is a **wealth-weighted average of opinions**. It sums to one because both opinions do.

Being right does not restore lost capital

For a binary event, investor 1 forecasts success probability 0.20 and investor 2 forecasts 0.80. The equilibrium success price is

$$P_E(\text{success}) = u(0.20) + (1 - u)(0.80).$$

Wealth split	Investor 1	Investor 2	Success price
Equal wealth	\$500	\$500	0.50
Investor 1 dominates	\$990	\$10	0.206

If the true success rate now changes to 0.80, beliefs and balances do not change instantly. **The price still starts at 0.206.** How quickly can the newly correct investor recover influence?

Notation from here on: $P_1 \rightarrow p$, $P_2 \rightarrow q$, $P_E \rightarrow \bar{\pi}$; wealth $M_1 \rightarrow I$ (incumbent), $M_2 \rightarrow E$ (challenger).

Four steps from optimal betting to fast recovery

1. **Free trading, full Kelly.** Optimal payoffs are likelihood ratios. Wealth follows Bayes, and a long history creates a long recovery delay.
2. **Free trading, smaller fractions.** Smaller invested fractions preserve more capital and can remove delay proportional to market age.
3. **Trading with fees.** Old losses become first order in the invested fraction. Merely reducing that fraction no longer removes linear delay.
4. **Evidence and a boost.** Keep a small baseline, detect the change, then raise the invested fraction. The same account recovers without new capital.

Forecasts and the change time

- Outcomes Y_t belong to a finite set $\{1, \dots, k\}$. The incumbent and challenger keep fixed probability vectors p and q .
- The entries p_y, q_y are their respective probabilities for outcome y . Each vector has positive entries and sums to one.
- Observations are independent: their distribution is p through round τ , then q . The trading rules do not know τ .
- $I_t, E_t > 0$ are the two accounts' wealth after round t . Cash is the numeraire. Every round settles before the next trade.

Running example: binary outcomes, with success probabilities $p = 0.2$ and $q = 0.8$. Here the scalar notation denotes Bernoulli probabilities.

Complete claims allow a payoff for each outcome.

Relative wealth is a Bayesian posterior

Use M_1, M_2 as model labels, with likelihoods $P(\omega \mid M_i) = P_i(\omega)$. The **prior** is the initial wealth share; the **joint probability** is prior times likelihood:

$$P(M_i) = \frac{M_i(0)}{M_1(0) + M_2(0)}, \quad P(M_i, \omega) = P(M_i)P(\omega \mid M_i).$$

The mixture marginal is $P(\omega) = \sum_j P(M_j, \omega) = P_E(\omega)$. Kelly and clearing give

$$\frac{M_i(1, \omega)}{M_1(1, \omega) + M_2(1, \omega)} = \frac{P(M_i, \omega)}{P(\omega)} = \boxed{P(M_i \mid \omega)}.$$

Both outcomes, from the same prior: $P(M_1) = 0.99$, $P(M_2) = 0.01$.

Outcome	$P(\omega \mid M_1)$	$P(\omega \mid M_2)$	$P(\omega)$	$P(M_1 \mid \omega)$	$P(M_2 \mid \omega)$
Success	0.20	0.80	0.206	96.12%	3.88%
Failure	0.80	0.20	0.794	99.75%	0.25%

Full Kelly learns the old law and becomes slow to change

Define the **log-wealth deficit** $Z_t = \log(I_t/E_t)$. Let τ count old-regime rounds and H count recovery rounds after the change.

True law	Expected change in Z	Effect on challenger
Old law p	$+D_{\text{KL}}(p q)$ per round	Builds a deficit
New law q	$-D_{\text{KL}}(q p)$ per round	Repairs that deficit

$$Z_\tau \approx \tau D_{\text{KL}}(p||q), \quad H \approx \tau \frac{D_{\text{KL}}(p||q)}{D_{\text{KL}}(q||p)}.$$

For $p = 0.2$, $q = 0.8$, both divergences equal $0.6 \log 4 \approx 0.832$. Thus **10,000 old rounds require about 10,000 repair rounds** at leading order.

Four steps from optimal betting to fast recovery

1. **Free trading, full Kelly.** Optimal payoffs are likelihood ratios. Wealth follows Bayes, and a long history creates a long recovery delay.
2. **Free trading, smaller fractions.** Smaller invested fractions preserve more capital and can remove delay proportional to market age.
3. **Trading with fees.** Old losses become first order in the invested fraction. Merely reducing that fraction no longer removes linear delay.
4. **Evidence and a boost.** Keep a small baseline, detect the change, then raise the invested fraction. The same account recovers without new capital.

Choosing an invested fraction adds a control beyond Bayes

Before round t , the incumbent invests fraction μ_t of wealth I_{t-1} ; the challenger invests fraction λ_t of wealth E_{t-1} . The rest stays in cash. These choices may depend on past outcomes.

$$w_t = \frac{\mu_t I_{t-1}}{\mu_t I_{t-1} + \lambda_t E_{t-1}}, \quad \bar{\pi}_t = w_t p + (1 - w_t) q.$$

Here w_t is the incumbent's share of **invested capital**, which determines the clearing price. With equal fractions, it equals the pre-round wealth share.

$$\frac{E_t}{E_{t-1}} = \underbrace{1 - \lambda_t}_{\text{cash kept back}} + \underbrace{\lambda_t \frac{q Y_t}{\bar{\pi}_t(Y_t)}}_{\text{invested portfolio}}, \quad \text{and similarly for } I_t.$$


The payoff now mixes cash with a likelihood ratio. **Ordinary Bayes is the special case** $\mu_t = \lambda_t = 1$. Changing a fraction changes exposure and can move the price before settlement changes wealth.

FREE TRADING: WHY USE A SMALLER INVESTED FRACTION?

Smaller fractions can remove linear delay when trading is free

Near incumbent dominance, the challenger's payoff is $f = q/p$. For a small common invested fraction λ :

$$\text{old loss} \approx \frac{1}{2} \lambda^2 \chi^2(q||p), \quad \text{new gain} \approx \lambda \chi^2(q||p).$$



$$\chi^2(q||p) = \sum_y (q_y - p_y)^2 / p_y = 2.25 \text{ here.}$$

Halving a small fraction quarters the old loss, but only halves repair.

For common $\lambda_t \propto t^{-\alpha}$, $0 < \alpha < 1$, the curve shows the power of τ . At $\alpha = 1/2$, recovery has order $\sqrt{\tau} \log \tau$.

For $\alpha > 1/2$, the deficit grows only logarithmically, but late invested fractions are too small to repair it quickly.

Four steps from optimal betting to fast recovery

1. **Free trading, full Kelly.** Optimal payoffs are likelihood ratios. Wealth follows Bayes, and a long history creates a long recovery delay.
2. **Free trading, smaller fractions.** Smaller invested fractions preserve more capital and can remove delay proportional to market age.
3. **Trading with fees.** Old losses become first order in the invested fraction. Merely reducing that fraction no longer removes linear delay.
4. **Evidence and a boost.** Keep a small baseline, detect the change, then raise the invested fraction. The same account recovers without new capital.

With fees, buyer and seller face different prices

After the change, success has true probability $q = 0.80$. The incumbent still forecasts $p = 0.20$ and supplies almost all invested capital: $w \rightarrow 1$. A **star** denotes this dominance limit.

Each side pays five cents per one-dollar claim, giving total spread $s = 0.10$. Write π for the midpoint with fees and $\bar{\pi}$ for the frictionless price.

One-dollar claim	Seller receives	Buyer pays	Midpoint π
Success	0.20	0.30	0.25
Failure	0.70	0.80	0.75

The dominant incumbent sells success claims at its valuation $p = 0.20$. Thus the buyer pays $p + s = 0.30$ and the midpoint is $p + s/2 = 0.25$.

Separate buying prices sum to $0.30 + 0.80 = 1.10$. We now build the challenger's portfolio using **cash and success claims bought at 0.30**.

The coefficients 0.70 and 0.30 bound every affordable payoff

Let $f_0, f_1 > 0$ be final dollars on failure and success. Write $C(f)$ for their purchase cost. Keep the smaller payoff in cash and buy claims for the difference.

$$f_1 \geq f_0 : \quad C(f) = f_0 + 0.30(f_1 - f_0) = 0.70f_0 + 0.30f_1,$$

$$f_1 < f_0 : \quad C(f) = f_1 + 0.80(f_0 - f_1) = 0.80f_0 + 0.20f_1.$$

In the second case, the extra cost is $0.10(f_0 - f_1) > 0$. Therefore **both cases** satisfy

$$C(f) \geq 0.70f_0 + 0.30f_1 = a^* \cdot f, \quad a^* = (0.70, 0.30).$$

Every affordable payoff has $C(f) \leq 1$, hence $a^* \cdot f \leq 1$. The next slide extends the prequel's proof to this budget inequality.

The same log-optimality proof works with fees

For any affordable positive payoff f , set $0 < c = a^* \cdot f \leq 1$ and $P_V(y) = a_y^* f_y / c$. The vector P_V sums to one, so

$$\mathbb{E}_q \log f = D_{\text{KL}}(q \| a^*) - D_{\text{KL}}(q \| P_V) + \log c \leq D_{\text{KL}}(q \| a^*).$$

Equality: $c = 1$ and $P_V = q$, hence $f^* = q/a^*$. Its actual cost is exactly one dollar.

Outcome y	New law q_y	Cost coefficient a_y^*	Payoff $f^*(y)$
Failure, $y = 0$	0.20	0.70	$2/7$
Success, $y = 1$	0.80	0.30	$8/3$

$$\mathbb{E}_q \log f^*(Y) = 0.20 \log \frac{2}{7} + 0.80 \log \frac{8}{3} \approx \boxed{0.534}.$$

$\mathbb{E}_q$ averages under the new law: expected log growth is 0.534 nats per round at full investment near incumbent dominance.

One dollar buys cash and success claims

Start with **one dollar assigned to the challenger's portfolio**. Its optimal allocation at these prices is:

Use of the dollar	Amount	What it buys
Keep cash	$2/7 \approx 0.286$	0.286 dollars in either outcome
Buy success claims	$5/7 \approx 0.714$	$(5/7)/0.30 = 50/21 \approx 2.381$ claims

Let $f^*(y)$ mean **final dollars per dollar assigned**, when outcome y occurs. Here $y = 0$ is failure and $y = 1$ is success.

$$\text{Failure: } f^*(0) = \frac{2}{7} \approx 0.286,$$

$$\text{Success: } f^*(1) = \frac{2}{7} + \frac{50}{21} = \frac{8}{3} \approx 2.667.$$

Investing ten percent scales the whole portfolio

λ is the fraction of current wealth **assigned to the portfolio just constructed**. Keep the remaining fraction $1 - \lambda$ as additional cash.

If that portfolio pays $f^*(y)$ per dollar, the entire account's wealth multiplier is

$$f_\lambda(y) = 1 - \lambda + \lambda f^*(y).$$

For an account starting with \$1	Failure	Success
Full investment, $\lambda = 1$	$2/7 \approx 0.286$	$8/3 \approx 2.667$
10% investment, $\lambda = 0.10$	$0.90 + 0.10(2/7) \approx 0.929$	$0.90 + 0.10(8/3) \approx 1.167$

Define $G_q(\lambda) = \mathbb{E}_q \log f_\lambda(Y)$, the expected log growth after the change. Here $G_q(0.10) \approx 0.1085$, compared with $G_q(1) \approx 0.5341$.

The full portfolio already contains some cash. A smaller invested fraction holds **extra cash outside it**.

Fees reduce total wealth while influence can recover

Let I_t and E_t be the accounts' wealth after round t . Let N_t be the signed number of success claims they trade that round.

Their risky payoffs cancel at settlement. The fee recipient collects $s|N_t|$, so

$$I_t + E_t = I_{t-1} + E_{t-1} - s|N_t|.$$

Example: combined wealth is \$100, the accounts trade one claim, and $s = 0.10$. Their combined wealth after settlement is \$99.90, whichever outcome occurs.

Fees scale with the number of claims traded. Smaller positions produce smaller dollar fees. The model keeps each account's wealth positive at every finite round.

Recovery concerns relative wealth and price influence. It does not restore the dollars already paid in fees.

The loss-to-repair ratio is ten cents divided by fifty cents

Near incumbent dominance, a success claim costs the challenger 0.30.

	True success rate	Expected payoff	Expected net gain
Before the change	0.20	0.20	−0.10 per claim
After the change	0.80	0.80	+0.50 per claim

The optimized one-dollar portfolio contains 50/21 claims and 2/7 cash. Therefore its expected net changes are $-A$ and $+B$, where

$$A = 0.10 \frac{50}{21} = 0.238, \quad B = 0.50 \frac{50}{21} = 1.190.$$

$$r = \frac{A}{B} = \frac{10 \text{ cents lost}}{50 \text{ cents gained}} = 0.20.$$

For small λ , expected log growth is approximately $-A\lambda$ before the change and $+B\lambda$ after it.

Betting less slows both losses and repair

Near incumbent dominance, a small common fraction λ gives these approximate changes in the expected log wealth of the challenger relative to the incumbent:

Before the change: the challenger is wrong	$-A\lambda$ per round
After the change: the challenger is right	$+B\lambda$ per round

If the **same small fraction** continues for τ old rounds and H repair rounds,

$$\underbrace{A\lambda\tau}_{\text{inherited deficit}} \approx \underbrace{B\lambda H}_{\text{repair}} \implies H \approx \frac{A}{B}\tau = r\tau.$$


The fraction cancels. Halving the invested fraction approximately halves both the damage and the speed of repair.

Increasing the invested fraction after detection can accelerate the comeback.

Compare recovery at the same invested fraction

Both accounts use the same constant fraction λ . For each fee choice, let $G_p(\lambda)$ be old-law expected log growth and $G_q(\lambda)$ the new-law growth, near incumbent dominance. Then

$$\theta_\lambda = \frac{-G_p(\lambda)}{G_q(\lambda)} = \frac{\text{old loss per round}}{\text{new gain per round}}, \quad \frac{H_\tau}{\tau} \xrightarrow{\mathbb{P}} \theta_\lambda.$$



λ	No fee	10-cent fee
$\downarrow 0$	0	0.200
0.1	0.051	0.262
0.2	0.105	0.330
1	1.000	1.509

At every fixed fraction, the fee raises the delay ratio. Its lower limit is zero without fees and $r = 0.20$ with fees.

What we measure during recovery

The law changes after round τ . Recovery is the **first next midpoint quote** $\pi \geq 0.50$. With equal invested fractions this means $E \geq I$.

$$Z_t = \log(I_t/E_t)$$

Log-wealth deficit. Wealths 99 and 1 give $Z = 4.60$.

$$H_\tau$$

Completed post-change rounds before the first qualifying next quote.

$$\Lambda_t = \sum_{j \leq t} \lambda_j$$

Cumulative exposure. Ten rounds at 10% add one unit.

$$R_\tau = \Lambda_{\tau+H_\tau} - \Lambda_\tau$$

Exposure used during recovery.

Change after round 100, first qualifying quote for round 121: $H_{100} = 20$. If all 20 fractions equal 0.1, then $R_{100} = 2$.

$\mathbb{P}_\tau$ denotes the law p through τ , then q . We count first passage; the quote may later fall back. With unequal fractions, invested capital determines the quote.

Continuing to reduce investment keeps recovery slow

Both accounts use the **same fraction of their respective wealth**. After a fixed starting date, that fraction never increases, including after the change. Its cumulative sum grows without bound.

For every $\varepsilon > 0$,

$$\mathbb{P}_\tau\{H_\tau > (r - \varepsilon)\tau\} \longrightarrow 1, \quad r = \frac{s}{q - p - s}.$$

After a late change, recovery takes at least approximately r times the market's age.

Example: $p = 0.2$, $q = 0.8$, $s = 0.1$ gives $r = 0.2$. A change after 10,000 rounds therefore has an asymptotic lower bound of about 2,000 recovery rounds.

The bound is **sharp**: the single schedule $\lambda_t = 1/\log(t + e)$ achieves $H_\tau/\tau \rightarrow r$ in probability.

A rule that raises investment after favorable evidence leaves this class.

Faster declines in investment mean longer delays

For common fractions $\lambda_t \sim \beta t^{-\alpha}$ with $0 < \alpha < 1$,

$$\Lambda_t \sim \frac{\beta}{1-\alpha} t^{1-\alpha}, \quad \Lambda_{\tau+H_\tau} \approx (1+r)\Lambda_\tau.$$

Recovery needs additional exposure of about r times the exposure accumulated before the change. Therefore

$$\left(1 + \frac{H_\tau}{\tau}\right)^{1-\alpha} \approx 1 + r, \quad \frac{H_\tau}{\tau} \xrightarrow{\mathbb{P}} (1+r)^{1/(1-\alpha)} - 1.$$

Schedule	H_τ/τ
$1/\log(t+e)$	0.200
Quarter-power	0.275
Square-root	0.440

For $r = 0.2$, the limits are shown at left. A steeper decline makes future rounds carry less exposure relative to the accumulated past exposure. It takes longer to supply the required repair.

The scale β cancels from the limiting delay ratio.

Four steps from optimal betting to fast recovery

1. **Free trading, full Kelly.** Optimal payoffs are likelihood ratios. Wealth follows Bayes, and a long history creates a long recovery delay.
2. **Free trading, smaller fractions.** Smaller invested fractions preserve more capital and can remove delay proportional to market age.
3. **Trading with fees.** Old losses become first order in the invested fraction. Merely reducing that fraction no longer removes linear delay.
4. **Evidence and a boost.** Keep a small baseline, detect the change, then raise the invested fraction. The same account recovers without new capital.

A harmonic baseline keeps earlier losses small

The default fraction of the challenger's current wealth invested in round t is $\ell_t = 1/t$.

Round t	10	100	1,000
Fraction invested	10%	1%	0.1%

Its cumulative exposure is the harmonic sum


$$L_\tau = \sum_{t=1}^{\tau} \frac{1}{t} \approx \log \tau.$$

The inherited log-wealth deficit therefore grows at logarithmic order.

When evidence favors the challenger, raise its fraction above the baseline.

A small baseline keeps the inherited deficit small but also makes repair slow. **Harmonic investment alone can be very slow; the boost is essential.**

Full investment gives the fastest capital repair



Keep the incumbent fully invested, $\mu_t = 1$. Under the new law,

$$D = G_q(1) = D_{\text{KL}}(q \| a^*) \\ = 0.534 \text{ nats per round.}$$

For a large deficit z , full investment $\lambda = 1$ is fastest within fractional-Kelly rules:

$$H \approx z/D = 1.87z.$$

Smaller fractions use less exposure but take longer. The limiting minimum is $R/z \rightarrow 1/B = 0.84$.

The baseline limits the old deficit. **A boost to $\lambda = 1$ gives the fastest capital repair.**

A rule the challenger can run alone

Before round t , the challenger averages the preceding

$$n_t = \lceil 12 \log(t + e) \rceil \quad \text{observations,} \quad \bar{Y}_t = \frac{1}{n_t} \sum_{j=t-n_t}^{t-1} Y_j.$$

$$\lambda_t = \begin{cases} 1, & \bar{Y}_t \geq 0.5, \\ \ell_t, & \bar{Y}_t < 0.5. \end{cases}$$

Below the threshold, use the cautious baseline. Above it, invest fully.

The constant 12 makes false-alarm probabilities summable: Hoeffding gives a bound of order $t^{-2.16}$. Choose $\ell_t \rightarrow 0$ with $L_t = \sum_{j \leq t} \ell_j \rightarrow \infty$. Use the baseline until a full window is available. Recheck each round.

The incumbent can keep $\mu_t = 1$. The challenger changes its invested fraction while retaining forecast $q = 0.8$.

The challenger can recover in less than market age

For any fixed predictable incumbent policy with $\mu_t \geq \underline{\mu} > 0$, the challenger rule gives

$$H_\tau = O_{\mathbb{P}}(L_\tau + \log \tau), \quad \frac{H_\tau}{L_\tau} \xrightarrow{\mathbb{P}} \frac{A}{D} \quad \text{if } \log \tau = o(L_\tau).$$

- Under unchanged law p , only finitely many false boosts occur almost surely. Longer windows can make the probability of any false boost as small as desired.
- After a late change, detection takes order $\log \tau$ with high probability. Full investment then repairs the inherited deficit at rate D .
- Since $\ell_t \rightarrow 0$, we have $L_\tau = o(\tau)$. Thus recovery takes a vanishing fraction of the market age.

$O_{\mathbb{P}}(a_\tau)$ means that H_τ/a_τ is bounded in probability. The limits concern late changes, $\tau \rightarrow \infty$.

The baseline controls how much capital must be repaired

Baseline	Cumulative exposure	Quote recovery after a late change
$\ell_t \sim \beta t^{-\alpha}, 0 < \alpha < 1$	$L_\tau \sim \frac{\beta}{1-\alpha} \tau^{1-\alpha}$	$H_\tau \sim \frac{A}{D} L_\tau$ in probability
Square-root	$L_\tau \sim 2\beta\sqrt{\tau}$	Order $\sqrt{\tau}$, with the sharp leading coefficient above
Harmonic, $\ell_t \sim \beta/t$	$L_\tau \sim \beta \log \tau$	$H_\tau = O_{\mathbb{P}}(\log \tau)$

The harmonic case has no sharp logarithmic constant in this theorem. Detection and capital repair both occur on the logarithmic scale.

Feedback helps by keeping the old deficit small and increasing exposure when evidence favors recovery. Once that deficit is fixed, the time–exposure frontier still applies.

Comparing feedback with the same incumbent

Binary outcomes use $p = 0.2$, $q = 0.8$, $s = 0.1$ and $I_0 = E_0 = 0.5$. The incumbent invests fully, $\mu_t = 1$, in **every comparison**.

Challenger baselines: $\ell_t = \frac{2}{3}(t+1)^{-1/2}$ or $\ell_t = (t+1)^{-1}$.


Without feedback: keep $\lambda_t = \ell_t$.

With feedback: raise λ_t to one when the recent success rate reaches 0.5.

- Target: the next midpoint quote reaches $\pi \geq 0.50$.
- At each change age, use 256 shared outcome paths. Recompute clearing and fees each round.
- Observe recovery for 2τ rounds. If fewer than half the paths recover, report a lower bound on the sample median.

The scale $2/3 = 1/\sqrt{\chi^2(q||p)}$ is the frictionless benchmark choice. It need not be optimal with fees.

Feedback shortens recovery with the incumbent unchanged



At age 32,000: square-root recovery falls from **14,639** to **173** rounds. Harmonic recovery falls from **more than 64,000** to **67** rounds.

Detection takes roughly half a window in this example

The old success rate is 0.2, the new rate is 0.8, and the threshold is 0.5. If h of the last n observations come from the new law, the window's expected average is

$$0.2 \frac{n-h}{n} + 0.8 \frac{h}{n} = 0.2 + 0.6 \frac{h}{n}.$$

It reaches 0.5 when $h \approx n/2$. At change age τ , use $n \approx n_\tau = \lceil 12 \log(\tau + e) \rceil$.

Age τ	Window n_τ	Simple approximation $n_\tau/2$	Observed median
2,000	92	46	45
8,000	108	54	54
32,000	125	62.5	61

This mean-window calculation explains the observed scale. Random crossing and the changing window affect the alarm time.

The theorem controls the order of delay and false boosts; it does not assert these exact medians.

After the alarm, remaining capital predicts the repair scale

At age 32,000, use the remaining log-wealth deficit Z at the first alarm and the limiting full-investment growth rate

$$D = D_{\text{KL}}(0.8||0.3) = 0.53411 \text{ nats per round,} \quad \text{repair estimate} \approx \frac{Z}{D}.$$

Baseline	Median Z at alarm	Approximation Z/D	Observed repair
Square-root	58.51	109.5 rounds	109 rounds
Harmonic	3.04	5.7 rounds	4 rounds

The square-root disadvantage is large, making the limiting drift informative. The harmonic account has little to rebuild, so detection accounts for most of its delay.

The price and growth rate change during repair. The harmonic disadvantage is small, so Z/D is a rough approximation. Rows use separate sample medians.

CONTEXT

Relation to market selection and adaptive learning

Market selection and price influence

Blume–Easley and Sandroni relate survival to beliefs, investment behavior and market structure. Kogan–Ross–Wang–Westerfield show that price impact and wealth survival can separate.

Adaptation after a change

The frictionless common-fraction update is Soft-Bayes. Easley–Kolumbus–Tardos show that concentrated Bayesian weights can incur linear regret after one shift, and study regularization that restores adaptation.

The contribution here

An attainable time–exposure frontier within fractional-Kelly rules for **separately owned capital**, with endogenous quotes and compulsory fees. A challenger can implement recovery using its own existing account.

The published *Methodology and Computing in Applied Probability* paper uses prefunded monitoring accounts and restarts. Here the same accounts persist and adjust their invested fractions.

Scope of the recovery results

Covered by the model

- Two fixed, full-support forecasts and complete contingent claims.
- Proportional fees on risky positions, with enough disagreement to make trading worthwhile.
- Fractional-Kelly portfolios and predictable invested fractions, with a relative lower bound on incumbent investment.
- One persistent distribution shift and first passage of a fixed quote target.

Capital lock-in here concerns wealth accumulated across settled rounds. Gebele–Matthes study capital lock-up caused by waiting for settlement, a different mechanism.

Questions for further work

- Many forecasts, where a third account can change execution prices.
- Recurrent regimes and sustained tracking.
- Risk budgets derived from intertemporal preferences or constraints.
- Fees that also charge a riskless bundle, requiring different normalization.

CONCLUSIONS

Four steps explain learning, lock-in and recovery

1. **Free trading, full Kelly.** The likelihood-ratio payoff is log-optimal. Wealth follows Bayes, but a long history creates a recovery delay proportional to market age.
2. **Free trading, smaller fractions.** Old losses become quadratic in the invested fraction while repair stays first order. Shrinking fractions can give sublinear recovery.
3. **Trading with fees.** Old losses become first order too. Common nonincreasing fractions face the positive delay floor $r\tau$, with $r = 0.20$ here.
4. **Evidence and a boost.** A harmonic baseline keeps the deficit small. Detection and an increase to full investment restore logarithmic recovery, using the same account and no new capital.

Why the likelihood ratio is log-optimal

Let $a = (a_y)$ be positive cost coefficients summing to one. Every affordable positive payoff $f = (f_y)$ satisfies $\sum_y a_y f_y \leq 1$. Assume $q_y > 0$.

$$\begin{aligned}\mathbb{E}_q \log f - D_{\text{KL}}(q \| a) &= \sum_y q_y \log \frac{a_y f_y}{q_y} \\ &\leq \log \left(\sum_y a_y f_y \right) \leq 0.\end{aligned}$$

The first inequality uses concavity of $\log$. The second uses the one-dollar budget. Equality requires $a_y f_y = q_y$ for every outcome.

If affordable, $f_y^{\text{opt}} = q_y / a_y$ is the unique optimum, with growth $D_{\text{KL}}(q \| a)$. Without fees, $a = \bar{\pi}$. With fees, the supporting coefficients will be a^* .

A trade specifies a payoff in each outcome

A binary position $v = (v_0, v_1)$ specifies its payoff on **failure** (v_0 , when $Y = 0$) and on **success** (v_1 , when $Y = 1$).

Position	v_0	v_1	Success claims held
One success claim	0	1	1
The same claim plus two dollars of cash	2	3	1
Two success claims	0	2	2

The realized payoff separates certain cash from the number of success claims:

$$v_Y = v_0 + (v_1 - v_0)Y.$$

Each side pays $c(v) = \frac{s}{2}|v_1 - v_0|$, where $s = 0.10$ is the total fee spread per claim.

One claim costs each side five cents. Two claims cost each side ten cents. Adding certain cash leaves the fee unchanged.

Portfolio choice and invested fractions

For general capital share w , write $m(w)$ and $a(w)$ for the incumbent's and challenger's supporting cost coefficients. Their optimal payoff multiples are

$$f_y^p(w) = \frac{p_y}{m_y(w)}, \quad f_y^q(w) = \frac{q_y}{a_y(w)}.$$

Investing fractions $\mu_t, \lambda_t \in (0, 1]$ gives

$$I_t = I_{t-1} [1 + \mu_t (f_{Y_t}^p(w_t) - 1)],$$

$$E_t = E_{t-1} [1 + \lambda_t (f_{Y_t}^q(w_t) - 1)].$$

These **fractional-Kelly rules** allocate committed capital log-optimally. We compare invested fractions within this family and assume enough disagreement to justify trade after fees.

Fractions depend only on past observations. We require $\mu_t \geq \kappa \lambda_t$, with fixed $0 < \kappa \leq 1$, so the incumbent cannot surrender price influence simply by withdrawing its capital much faster.

Each vector sums to one. The midpoint quote is $\pi(w) = [m(w) + a(w)]/2$.

The recovery clock counts completed settlements

Let τ be the last round under the old law. H_τ counts completed post-change rounds before the first qualifying next quote.

Example: the change occurs after round 100. The first quote reaching 50% is the quote for round 121. Then $H_{100} = 20$.

If each of those 20 rounds uses $\lambda = 0.10$, recovery exposure is $R_{100} = 20 \times 0.10 = 2$.

For a target region $\mathcal{V}$, the formal definitions are

$$H_\tau = \inf\{h \geq 0 : \pi_{\tau+h+1} \in \mathcal{V}\}, \quad R_\tau = \Lambda_{\tau+H_\tau} - \Lambda_\tau.$$

$\mathcal{V}$ contains the challenger-dominance quote and is separated from the incumbent-dominance quote. The $+1$ selects the *next* quote after the counted settlements.

$\mathbb{P}_\tau$ means the sequence law with probability p through round τ and q afterward. If the first post-change quote already qualifies, $H_\tau = R_\tau = 0$.

Three rates determine recovery

All three rates describe the challenger near incumbent dominance. The invested fraction is λ and the log-wealth deficit is $Z = \log(I/E)$.

Rate	Formula	Example
Old-law loss	$A = 1 - \mathbb{E}_p f^*$	0.238
New-law repair, small fractions	$B = \mathbb{E}_q(f^* - 1) = \chi^2(q a^*)$	1.190
New-law repair, full investment	$D = \mathbb{E}_q \log f^* = D_{\text{KL}}(q a^*)$	0.534

Expected change in Z : approximately $+A\lambda$ before the change, $-B\lambda$ after it for small fractions, and $-D$ after it at full investment.

The chi-squared divergence is $\chi^2(q||a^*) = \sum_y (q_y - a_y^*)^2 / a_y^*$. A positive loss coefficient A is the effect of fees.

A fixed common fraction produces age-proportional delay

Suppose both accounts use the same fixed fraction $\lambda \in (0, 1]$. For $a_1^* = p + s$ and $\delta = q - a_1^* > 0$,

$$G_\rho(\lambda) = \rho \log \frac{a_1^* + \lambda\delta}{a_1^*} + (1 - \rho) \log \frac{1 - a_1^* - \lambda\delta}{1 - a_1^*}, \quad \rho = p, q.$$

$$\frac{H_\tau}{\tau} \xrightarrow{\mathbb{P}} \theta_\lambda := \frac{-G_p(\lambda)}{G_q(\lambda)} > r := \frac{A}{B} = \frac{s}{q - p - s}.$$

The old regime creates a deficit proportional to its duration. After the change, a fixed fraction repairs it at a constant rate per round.

As $\lambda \downarrow 0$, $\theta_\lambda \rightarrow r$. For $p = 0.2, q = 0.8, s = 0.1$, the limiting ratio is $r = 0.2$: about one recovery round per five old-regime rounds.

Small invested fractions imply a fixed exposure requirement

For any deterministic schedule with $\lambda_t \rightarrow 0$ and $\Lambda_t \rightarrow \infty$,

$$\frac{\Lambda_{\tau+H_\tau} - \Lambda_\tau}{\Lambda_\tau} \xrightarrow{\mathbb{P}} \frac{A}{B}.$$

- This is a statement about **cumulative exposure**. The schedule translates that exposure into calendar time.
- The ratio is loss before the change divided by repair afterward, each per unit of small exposure.
- No monotonicity or regular variation is needed. The incumbent may choose a different predictable fraction, subject to $\mu_t \geq \kappa \lambda_t$.

The best tradeoff between recovery time and exposure

Under $\mathbb{P}_q$, all outcomes follow q . Start with $Z_0 = z$ large. Within the **fractional-Kelly policy class**, let H be recovery rounds and R cumulative exposure, with $\mu_t \geq \kappa \lambda_t$.

$$\mathbb{P}_q\{H < \infty, HG_q(R/H) < (1 - \varepsilon)z\} \longrightarrow 0, \quad z \rightarrow \infty.$$

For any fixed $\varepsilon > 0$, successful recovery cannot save that proportion of the required log-growth. A constant challenger fraction $u \in (0, 1]$ attains

$$\left(\frac{H}{z}, \frac{R}{z}\right) \xrightarrow{\mathbb{P}} \left(\frac{1}{G_q(u)}, \frac{u}{G_q(u)}\right).$$

Fastest recovery: $1/D$ rounds per nat, attained at full investment.

Least limiting exposure: $1/B$ units per nat, approached as $u \rightarrow 0$, while time per nat grows without bound.

These leading costs also hold for a fixed market-quote target. The exposure bound depends on the specified portfolio family.

Fractional commitment preserves the likelihood structure

Holding cash beside the full-investment payoff gives

$$f_\lambda = 1 - \lambda + \lambda \frac{q}{a^*} = \frac{q_\lambda}{a^*}, \quad q_\lambda = (1 - \lambda)a^* + \lambda q.$$

$$G_q(\lambda) := \mathbb{E}_q \log f_\lambda = D_{\text{KL}}(q \| a^*) - D_{\text{KL}}(q \| q_\lambda).$$

The forecast remains q . The mixture q_λ represents the **payoff produced by cash retention**. Reducing exposure lowers both the risk taken and the rate at which favorable evidence becomes log-capital growth.

At $\lambda = 1$ the full information-growth rate is available. At small exposure,

$$G_q(\lambda) = \lambda \chi^2(q \| a^*) + O(\lambda^2), \quad \chi^2(q \| a^*) = \sum_y \frac{(q_y - a_y^*)^2}{a_y^*}.$$

Why capital recovery implies quote recovery

With unequal fractions, the price responds to **committed capital**. Define

$$\widehat{Z}_t = \log \frac{\mu_t I_{t-1}}{\lambda_t E_{t-1}} = Z_{t-1} + \log \frac{\mu_t}{\lambda_t}.$$

Under the new law, the following wealth-ratio inequality holds

$$\mathbb{E}_q[\Delta(I/E) \mid \mathcal{F}_{t-1}] \leq -2\lambda_t c(f^q(w_t)).$$

After reaching a fixed capital-ratio level, the expected remaining exposure to the quote target is bounded independently of the original deficit. This transfers the leading frontier constants to actual quotes.

The restriction $\mu_t \geq \kappa \lambda_t$ prevents the incumbent from producing an apparent recovery merely by withdrawing its capital arbitrarily fast.

A boost changes commitment before wealth changes

At fixed account balances and incumbent exposure,

$$\hat{Z}^{\text{baseline}} = Z + \log(\mu/\ell), \quad \hat{Z}^{\text{boost}} = Z + \log \mu.$$

Raising the challenger's fraction from ℓ to 1 removes exactly $\log(1/\ell)$ from the log ratio of committed capital, immediately moving the quote.

For $\ell_t \sim \beta/t$ and $\mu_t \geq \underline{\mu} > 0$, compare commitments at the **same old-regime wealth state**:

$$Z \sim A_c \beta \log t, \quad \hat{Z}^{\text{baseline}} \sim (A_c \beta + 1) \log t, \quad \hat{Z}^{\text{boost}} \sim A_c \beta \log t.$$

The remaining log-wealth deficit still diverges. Subsequent settlement is needed to reach a fixed interior quote target. These frozen-wealth formulas do not describe wealth changes during the actual detection interval.

Zero cost is a singular limit

When $c = 0$, the boundary execution law becomes $a^* = p$ and $A_c = 0$. For $f^* = q/p$ and small exposure u ,

$$\begin{aligned} -\mathbb{E}_p \log[1 + u(f^* - 1)] &= \frac{1}{2}u^2\chi^2(q\|p) + O(u^3), \\ G_q(u) &= u\chi^2(q\|p) + O(u^2). \end{aligned}$$

Old-law attrition is now **second order** in exposure, while favorable-law repair remains first order. The positive-cost exposure clock therefore degenerates at zero cost.

Endogenous clearing also contributes a nonnegative forcing term to relative-wealth dynamics. The resulting frictionless phase diagram is a separate mechanism, summarized in the appendix.

The costly-market theorem cannot be specialized to the frictionless regime simply by setting $A_c = 0$.

How fees determine the limiting state prices

The dual cost set consists of zero-sum vectors:

$$\mathcal{C} = \{\xi : \xi \cdot v \leq c(v) \text{ for all } v\}, \quad \xi \cdot \mathbf{1} = 0.$$


At incumbent dominance,

$$a^* = \arg \min_{a \in (p+2\mathcal{C}) \cap \mathcal{P}_k} D_{\text{KL}}(q||a).$$

$\mathcal{P}_k$ is the probability simplex. The intersection enforces nonnegative coordinates. The factor two reflects charges on both sides.

Standing condition: $q - p \notin 2\mathcal{C}$. This excludes a no-trade equilibrium. In the binary case, a_1^* minimizes $D_{\text{KL}}(q||a_1)$ over $[p - s, p + s] \cap [0, 1]$. Here $a_1^* = 0.30$.

The static reverse-information projection is established duality. Here it determines the coefficients in a dynamic recovery problem. The appendix slide “Equilibrium payoffs and a supporting quote” explains the normalization.



An exposure rule the challenger can run alone

Choose a baseline $\ell_t \downarrow 0$ with $L_t = \sum_{j \leq t} \ell_j \rightarrow \infty$. Select an outcome y_0 for which $\delta_0 = q_{y_0} - p_{y_0} > 0$ and set $S_j = \mathbf{1}_{\{Y_j = y_0\}}$.

$$n_t = \lceil k_{\text{det}} \log(t + e) + c_0 \rceil, \quad k_{\text{det}} > 2/\delta_0^2, \quad \bar{S}_t = \frac{1}{n_t} \sum_{j=t-n_t}^{t-1} S_j.$$

$$\lambda_t = \begin{cases} 1, & \bar{S}_t \geq (p_{y_0} + q_{y_0})/2, \\ \ell_t, & \text{otherwise.} \end{cases}$$

Use the baseline while a full window is unavailable. Recheck every round, so an isolated false alarm need not cause a permanent boost.

Only **past public outcomes** enter the rule. The same account funds every position and fee. The incumbent may remain at $\mu_t = 1$ throughout.

In the binary example, the score is simply Y_j and the threshold is 0.5. The detector does not observe τ .

Why the frontier holds

Before the challenger repairs a large deficit, incumbent dominance makes the conditional drift approximately $-G_q(\lambda_t)$.

$$\mathbb{E}_q[\Delta Z_t \mid \mathcal{F}_{t-1}] \approx -G_q(\lambda_t), \quad \sum_{t=1}^H G_q(\lambda_t) \leq H G_q\left(\frac{R}{H}\right).$$

1. Strict concavity of G_q gives the second inequality by Jensen. It bounds total expected repair for given time and exposure.
2. Over exposure of order z , martingale fluctuations have scale $\sqrt{z}$. Beating the leading frontier would require an excursion of order z .
3. A constant fraction attains the boundary asymptotically as $z \rightarrow \infty$. The final quote passage has bounded expected remaining exposure after a fixed capital-ratio threshold.

This is an explanation of the proof mechanism. The paper supplies stopping, localization and remainder estimates for predictable policies.

Equilibrium payoffs and a supporting quote

Normalize committed capital to one, with shares $w, 1 - w$. Each account maximizes expected log payoff at given prices and fees. A zero-inventory venue matches their risky trades.

$$\max_{U, V > 0} \{w\mathbb{E}_p \log U + (1 - w)\mathbb{E}_q \log V\} \quad \text{subject to } U_y + V_y + 2c(V) \leq 1 \quad \forall y.$$

The solution has $U + V = (1 - 2c(V))1$. Its opposite risky positions imply $c(U) = c(V)$, so $2c(V)$ is the combined fee. Define

$$m_y = \frac{wp_y}{U_y}, \quad a_y = \frac{(1 - w)q_y}{V_y}.$$

The optimality and budget equations give

$$\sum_y m_y = \sum_y a_y = 1, \quad a - m \in 2\partial c(V), \quad \pi = \frac{1}{2}(m + a).$$

The allocations and effective state prices are unique. We **select this symmetric supporting quote**; nonsmooth fees can admit other supporting quotes.

Price-recovery claims use this selection. The supporting quote does not affect wealth dynamics.

A three-outcome execution calculation

$$p = (0.60, 0.25, 0.15), \quad q = (0.20, 0.30, 0.50), \quad c(v) = 0.04(\max_y v_y - \min_y v_y).$$

Here $\mathcal{C} = \{\xi : \xi \cdot \mathbf{1} = 0, \|\xi\|_1 \leq 0.08\}$. The projection shifts probability mass from outcome 1 to outcome 3:

$$a^* = (0.52, 0.25, 0.23), \quad f^* = (0.384615, 1.200000, 2.173913).$$

$$A_c = 0.143144, \quad B_c = 0.523880, \quad D_c = 0.251859.$$

The first-order old-law loss equals $2c(f^*)$, the small-exposure repair rate equals $\chi^2(q\|a^*)$, and full investment gives $D_{\text{KL}}(q\|a^*)$.

The general execution law is a probability vector because of the transaction technology. It is not an arbitrary collection of outcome-specific bid and ask quotes.

Comparable medians without incumbent cooperation

Both rules use the same challenger detector. Under common feedback, both accounts change exposure. Under autonomous feedback, only the challenger does and the incumbent stays at $\mu = 1$.

Age τ	Square-root baseline		Harmonic baseline	
	Common	Autonomous	Common	Autonomous
2,000	73	72.5	50	49
8,000	108	107	60	59
32,000	173.5	173	68	67

Median differences are at most one round in these experiments. This is a sample comparison, not a theorem ordering the two policies.

During boost rounds, both versions use $\mu = \lambda = 1$. Differences arise in baseline rounds and in the balances inherited when boosting begins.

256 paired paths per age. Target: midpoint $\pi \geq 0.50$. Common-feedback medians: earlier simulation supplied with the project.

Predetermined bursts help at a thin set of change times

A deterministic schedule can concentrate exposure just after selected dates. It can then recover unusually fast if the change occurs at a favorable date.

For a set of change times $S \subset \mathbb{N}$, its upper logarithmic density is


$$\bar{d}_{\log}(S) = \limsup_{T \rightarrow \infty} \frac{1}{\log T} \sum_{\substack{\tau \leq T \\ \tau \in S}} \frac{1}{\tau}.$$

In the binary model, for common deterministic exposure with divergent sum, the set of times with $\mathbb{P}_{\tau}(H_{\tau} \leq \delta\tau) \geq \eta$, $\delta > 0$ and $\eta \in (0, 1)$, has upper logarithmic density at most

$$\min \left\{ 1, \frac{\log(1 + \delta)}{\log(1 + r)} \right\}.$$

The nontrivial coefficient is sharp for $0 < \delta < r$. Change times permitting sublinear recovery with success probability bounded away from zero have logarithmic density zero.

Frictionless recovery has a different phase diagram



Common $\lambda_t \sim \beta t^{-\alpha}$

Assume no fee and $0 < \alpha < 1$.

$0 < \alpha < \frac{1}{2}$: recovery of order $\tau^{1-\alpha}$.

$\alpha = \frac{1}{2}$: order $\sqrt{\tau} \log \tau$, with best amplitude $\beta = \chi^2(q||p)^{-1/2}$.

$\frac{1}{2} < \alpha < 1$: order $\tau^\alpha \log \tau$.

Clearing adds a nonnegative forcing term to a random affine recursion. The critical half-normal correction comes from Brownian reflection. These results are outside the focused costly-market paper.

Notation for reading the results

p, q	One-round probability vectors; binary success probabilities when scalar.
I_t, E_t	Account wealth in cash units after settlement.
μ_t, λ_t	Incumbent and challenger invested fractions before Y_t .
$\bar{\pi}, \pi$	Frictionless price vector and midpoint quote with fees.
Z_t	Log-wealth deficit $\log(I_t/E_t)$ of the challenger.
Λ_t, L_t	Cumulative actual challenger exposure and cumulative baseline exposure.
a^*, f^*	Limiting challenger state-price vector and full-investment payoff q/a^* .
G_p, G_q	Expected log growth when the incumbent dominates under the old and new laws.
H, R, z	Recovery time in rounds, recovery exposure and initial log-wealth deficit.
$\xrightarrow{\mathbb{P}}, O_{\mathbb{P}}$	Convergence in probability and boundedness in probability after rescaling.

Main-slide A, B, D are the paper's A_c, B_c, D_c . The divergence is written $D_{\text{KL}}(\cdot \parallel \cdot)$.

Selected references on markets and learning

Blume and Easley (1992). Evolution and Market Behavior. *Journal of Economic Theory* 58, 9–40.

Blume and Easley (2006). If You're So Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets. *Econometrica* 74, 929–966.

Sandroni (2000). Do Markets Favor Agents Able to Make Accurate Predictions? *Econometrica* 68, 1303–1341.

Kogan, Ross, Wang and Westerfield (2006). The Price Impact and Survival of Irrational Traders. *Journal of Finance* 61, 195–229.

Easley, Kolumbus and Tardos (2026). Markets with Heterogeneous Agents: Dynamics and Survival of Bayesian vs. No-Regret Learners. *arXiv:2502.08597*.

Orseau, Lattimore and Legg (2017). Soft-Bayes: Prod for Mixtures of Experts with Log-Loss. *Algorithmic Learning Theory*.

Likelihood, execution and related prediction-market work

Vecer (2026). Analytical solution for Kelly's criterion for multiple outcomes. *Quantitative Finance* 26(5), 671–684. [doi:10.1080/14697688.2026.2629952](https://doi.org/10.1080/14697688.2026.2629952).

Kelly (1956). A New Interpretation of Information Rate. *Bell System Technical Journal* 35, 917–926.

Cover and Thomas (2006). *Elements of Information Theory*, 2nd ed., Chapter 6. Horse-race betting, fair odds and track take.


Larsson, Ramdas and Ruf (2025). The Numeraire E-Variable and Reverse Information Projection. *Annals of Statistics* 53, 1015–1043.

Wolfers and Zitzewitz (2006). Interpreting Prediction Market Prices as Probabilities. NBER Working Paper 12200.

Vecer (2026). Likelihood-Ratio Pricing, Capital Lock-In, and Prefunded Restarts in Parimutuel Prediction Markets. *Methodology and Computing in Applied Probability*.
[doi:10.1007/s11009-026-10329-z](https://doi.org/10.1007/s11009-026-10329-z).

Gebele and Matthes (2026). When Certainty Is Not Worth It: Capital Lock-Up and Settlement Discounting in Prediction Markets. *arXiv:2605.31431*.

Further reading and staying in touch



Principles of Portfolio Choice

An Information-Theoretic,
Likelihood-Based Perspective

Jan Vecer

janvecer.com



Website and research



Connect on LinkedIn