

# Computational Bayesian Portfolio Selection via Relative Entropy

Projection, Bayesian clearing, and adaptation

Jan Vecer  
Charles University

ICCF 2026, Oxford – 3 September 2026

---

# Abstract

This talk presents a computational framework for portfolio selection and hedging based on the interpretation of prices as likelihood ratios. In this formulation, optimal investment can be viewed as a projection of an investor's subjective model onto the set of tradable payoffs, measured by relative entropy.

The same perspective links Kelly growth, Bayesian model averaging, and incomplete-market hedging within a unified probabilistic framework. The approach provides tractable discrete-time implementations and clarifies continuous-time limits, including connections to minimum-variance hedging.

Its relevance to computational finance lies in turning portfolio construction and hedge design into explicit statistical optimization problems that can be implemented numerically under heterogeneous beliefs and incomplete models.

# A normalized payoff is a probability model

## Payoff–density correspondence

Fix a strictly positive normalized pricing mass  $m$ .  
A unit-value payoff  $R \geq 0$  satisfies

$$\mathbb{E}_m[R] = 1.$$

It induces a probability mass, and every mass induces a payoff:

$$p_R(\omega) = m(\omega)R(\omega), \quad R(\omega) = \frac{p_R(\omega)}{m(\omega)}.$$

## Two states

$$m = (0.5, 0.5), \quad p = (0.6, 0.4).$$

The associated gross payoff is

$$R = \frac{p}{m} = (1.2, 0.8).$$

Thus an investment of 100 pays

$$(120, 80).$$

The discounted gross return is a likelihood ratio.

# Expected log growth is information minus approximation error

## The identity

If the investor's subjective law is  $P$  and the chosen payoff induces  $p_R$ , then

$$\mathbb{E}_P[\log R] = D_{\text{KL}}(P \| m) - D_{\text{KL}}(P \| p_R).$$

## Complete market

If every positive state payoff is tradable, choose

$$p_R^* = P, \quad R^* = \frac{P}{m}.$$

The maximum expected log growth is  $D_{\text{KL}}(P \| m)$ .

## Incomplete market

Tradability restricts the attainable densities. The loss relative to the ideal payoff is exactly

$$\min_{p_R \in \mathcal{P}_{\text{trad}}} D_{\text{KL}}(P \| p_R).$$

Portfolio selection is an information projection onto tradable payoffs.

# The incomplete-market projection is directly computable

Let  $R_1, \dots, R_K$  be positive unit-value traded returns on scenarios  $\omega_1, \dots, \omega_N$ . For long-only weights  $\alpha \in \Delta_K$ ,

$$R_\alpha(\omega) = \sum_{j=1}^K \alpha_j R_j(\omega), \quad q_\alpha = m R_\alpha = \sum_{j=1}^K \alpha_j q_j.$$

Write  $f(\alpha) = \sum_n P_n \log(\sum_j \alpha_j R_{nj})$ .

## Scenario program

$$\alpha^* = \arg \max_{\alpha \in \Delta_K} f(\alpha) = \arg \min_{\alpha \in \Delta_K} D_{\text{KL}}(P \| q_\alpha).$$

$$\frac{\partial f}{\partial \alpha_j} = \mathbb{E}_P \left[ \frac{R_j}{R_\alpha} \right]$$

$$\frac{\partial^2 f}{\partial \alpha_j \partial \alpha_k} = -\mathbb{E}_P \left[ \frac{R_j R_k}{R_\alpha^2} \right].$$

The objective is concave: projected Newton, mirror descent, or standard convex solvers apply directly to the scenario table.

# Locally, KL hedging becomes a variance projection

For a target claim  $H > 0$ , put  $H_0 = \mathbb{E}_m[H]$ ,  $h = H/H_0$ , and  $p^H = mh$ . A positive attainable hedge  $V_\alpha$  has  $V_{\alpha,0} = \mathbb{E}_m[V_\alpha]$ ,  $v_\alpha = V_\alpha/V_{\alpha,0}$ , and  $p^{V_\alpha} = mv_\alpha$ .

$$\mathbb{E}_{p^H} \left[ \log \frac{v_\alpha}{h} \right] = -D_{\text{KL}}(p^H \| p^{V_\alpha}).$$

The hedge is the attainable density closest to the target density.

For normalized diffusive one-step innovations, write

$$h = 1 + \sqrt{\Delta t} \eta_H + o(\sqrt{\Delta t}), \quad v_\alpha = 1 + \sqrt{\Delta t} \eta_\alpha + o(\sqrt{\Delta t}), \quad \mathbb{E}_m \eta_H = \mathbb{E}_m \eta_\alpha = 0.$$

Then the local Taylor expansion gives

$$D_{\text{KL}}(p^H \| p^{V_\alpha}) = \frac{\Delta t}{2} \mathbb{E}_m [(\eta_H - \eta_\alpha)^2] + o(\Delta t).$$

Locally this is an  $m$ -weighted least-squares projection of normalized payoff innovations.

# Heterogeneous models can be traded as separate funds

## Model funds

Model  $i$  supplies a predictive distribution  $p_{i,t}$  and retains its own bankroll  $B_{i,t}$ . Given a common quote  $m_t$ , full commitment gives

$$\frac{B_{i,t+1}}{B_{i,t}} = \frac{p_{i,t}(X_{t+1})}{m_t(X_{t+1})}.$$

Gross return is a likelihood ratio; log return is a relative log score.

## Portfolio of models

With value weights  $\alpha \in \Delta_K$ ,

$$q_{\alpha,t} = \sum_{i=1}^K \alpha_i p_{i,t}.$$

The long-only portfolio of model funds is exactly the arithmetic mixture of their predictive densities.

Bayesian model averaging becomes a self-financing portfolio rule.

# Statewise budget balance fixes the parimutuel quote

Account  $i$  commits a fraction  $\lambda_{i,t}$  of its wealth and distributes that stake according to  $p_{i,t}$ . Put

$$\pi_{i,t} = \frac{B_{i,t}}{\sum_j B_{j,t}}, \quad \bar{\lambda}_t = \sum_i \pi_{i,t} \lambda_{i,t}.$$

Assume  $\bar{\lambda}_t > 0$ . Within a common positive-quote Arrow-ticket mechanism:

## Unique statewise self-financing quote

$$m_t^\lambda(x) = \sum_i \frac{\pi_{i,t} \lambda_{i,t}}{\bar{\lambda}_t} p_{i,t}(x), \quad B_{i,t+1} = B_{i,t} \left( 1 - \lambda_{i,t} + \lambda_{i,t} \frac{p_{i,t}(x)}{m_t^\lambda(x)} \right).$$



# With full commitment, wealth is Bayesian evidence

For  $\lambda_{i,t} = 1$ , the clearing quote and the settlement rule give

$$\text{predictive price: } m_t(x) = \sum_i \pi_{i,t} p_{i,t}(x),$$

$$\text{fund return: } \frac{B_{i,t+1}}{B_{i,t}} = \frac{p_{i,t}(X_{t+1})}{m_t(X_{t+1})},$$

$$\text{posterior share: } \pi_{i,t+1} = \frac{\pi_{i,t} p_{i,t}(X_{t+1})}{\sum_j \pi_{j,t} p_{j,t}(X_{t+1})}.$$

If full commitment has held through date  $t$ :


$$\frac{B_{i,t}}{B_{i,0}} = \frac{\mathcal{L}_{i,t}}{\mathcal{L}_{m,t}}$$

Fund-to-mixture wealth is an evidence ratio.

$$\sum_i \pi_{i,t} D_{\text{KL}}(p_{i,t} \| m_t) = I(\Theta; X_{t+1} \mid \mathcal{F}_t)$$

If  $\Theta \mid \mathcal{F}_t \sim \pi_t$  and  $X_{t+1} \mid \Theta = i, \mathcal{F}_t \sim p_{i,t}$ ,  
this is predictive mutual information.

# The online implementation is four accounting steps



## Inputs

- ▶ timestamped probability forecasts;
- ▶ separate bankrolls;
- ▶ declared commitment fractions.

## Outputs

- ▶ one bundled predictive price;
- ▶ self-financing model scores;
- ▶ automatically updated model weights.

The layer can run internally before live capital is deployed.

# Fractional commitment limits exposure – but does not forget

With a common commitment fraction  $\lambda_t \in (0, 1]$ , the quote remains the wealth-weighted mixture, while each account earns

$$\frac{B_{i,t+1}}{B_{i,t}} = (1 - \lambda_t) + \lambda_t \frac{p_{i,t}(X_{t+1})}{m_t(X_{t+1})} \geq 1 - \lambda_t.$$

## Soft-Bayes update – Orseau, Lattimore & Legg (2017)

$$\pi_{i,t+1} = (1 - \lambda_t)\pi_{i,t} + \lambda_t \pi_{i,t+1}^B,$$

where  $\pi_{i,t+1}^B$  is the full Bayesian posterior weight.

### What $\lambda$ does

It limits the loss in one outcome and attenuates each new likelihood increment.

### What $\lambda$ does not do

It never discounts, transfers, or deletes likelihood evidence already stored in wealth.

Risk control is not the same as adaptive forgetting.

# Partial commitment still leaves order- $\tau$ memory

Here  $p$  is the old-regime law and  $q$  the new-regime law. On a finite state space, observations are independent – i.i.d. from fixed, strictly positive  $p$  for  $\tau$  periods, then from  $q$ . Both accounts start with positive wealth and use one fixed  $\lambda \in (0, 1]$ . Let

$$r_\lambda = (1 - \lambda)p + \lambda q,$$

and let  $H_\tau$  be the post-change time until the  $q$ -account catches the  $p$ -account.

## Permanent-switch lock-in

$$\frac{H_\tau}{\tau} \xrightarrow{P} \kappa_\lambda(p, q) := \frac{D_{\text{KL}}(p \| r_\lambda)}{D_{\text{KL}}(q \| p) - D_{\text{KL}}(q \| r_\lambda)}.$$


Asymptotically, the clearing price remains close to obsolete  $p$  until approximately  $\kappa_\lambda \tau$ , then reprices toward  $q$ .

$$\kappa_1(p, q) = \frac{D_{\text{KL}}(p \| q)}{D_{\text{KL}}(q \| p)}$$

$$\kappa_\lambda(p, q) = \frac{\lambda}{2} + O(\lambda^2) \quad (\lambda \downarrow 0).$$

**Fixed  $\lambda > 0$ : lock-in remains  $\Theta(\tau)$ ; lowering  $\lambda$  makes deficit recovery take order  $1/\lambda$ .**

# Accumulated evidence becomes a stale market price



Source: Vecer (2026), MCAP manuscript, Figure 2.

## One fixed simulated path

$$\begin{aligned} p &= (0.15, 0.70, 0.15), \\ q &= (0.2615, 0.4500, 0.2885), \\ \pi_{p,0} &= \pi_{q,0} = \frac{1}{2}, \quad \lambda = 1. \end{aligned}$$

Independent draws:

$$X_t \sim \begin{cases} p, & 1 \leq t \leq 500, \\ q, & 501 \leq t \leq 1200, \end{cases}$$

Blue is determined by orange

$$\begin{aligned} m_t(2) &= 0.70\pi_{p,t} + 0.45(1 - \pi_{p,t}) \\ &= 0.45 + 0.25\pi_{p,t}. \end{aligned}$$

The truth jumps at 500; accumulated wealth, and hence the quote, does not.

$$\kappa_1 = 0.945 \implies 500\kappa_1 \approx 473.$$

**Benchmark:**  $t \approx 973$ ; this path crosses earlier.

# One restart account measures evidence on one recent window

Let  $p_{0,t}$  be the incumbent benchmark and  $p_{1,t}, \dots, p_{K,t}$  candidate forecasts. A pooling type  $\alpha \in \Delta_K$  defines  $q_{\alpha,t} = \sum_{i=1}^K \alpha_i p_{i,t}$ .

Account starting at  $s$  and active for  $d$  observations

$$L_t^{s,d,\alpha} = \prod_{u=s}^{\min\{t,s+d-1\}} \frac{q_{\alpha,u-1}(X_u)}{p_{0,u-1}(X_u)}$$

The empty product is one when  $t < s$ ; each active factor is the pool's one-step return relative to the benchmark.

## Before start

$L_t^{s,d,\alpha} = 1$ : its capital remains in the benchmark.

## During the window

It accumulates likelihood evidence for  $q_\alpha$  against  $p_0$ .

## After expiry

The relative wealth is frozen; no old evidence is added.

**The account asks whether  $q_\alpha$  predicted observations  $s, \dots, s + d - 1$  better than  $p_0$ .**

# Prefunding makes the whole search anytime-valid

First average over pooling types within each window; then combine all windows:

$$U_t^{s,d} = \int_{\Delta_K} L_t^{s,d,\alpha} \nu_{s,d}(d\alpha), \quad E_t = c + \sum_{s,d} \gamma_{s,d} U_t^{s,d}, \quad c = 1 - \sum_{s,d} \gamma_{s,d}.$$

## Notation in words

$\nu_{s,d}$  prior over pooling weights  $\alpha$  within window  $(s, d)$ ;

$\gamma_{s,d}$  fraction of initial capital assigned *in advance* to that window;

$c$  residual capital kept in the benchmark.

## Anytime validity

If  $p_{0,t}$  is the true conditional law, then  $E_0 = 1$ ,  $(E_t)$  is a nonnegative martingale, so

$$\mathbb{P}_0 \left( \sup_t E_t \geq \frac{1}{\delta} \right) \leq \delta.$$

Every start, duration, and pooling type is funded before its outcome is seen.


**continuum of  $(\alpha, s, d) \implies$  only  $K + 1$  traded funds**

Posterior-mean consolidation uses the  $K$  candidates plus the benchmark.

**Recovery scale under the stated switch model:**  $O(\tau) \longrightarrow O_{\mathbb{P}}(\log \tau)$ .

Related restart/e-detector methodology: Shin, Ramdas & Rinaldo (2024).

# Convex pooling helps between models – but it is not free



**Interior,  $d = 80$ :** pool 84.3% vs vertices 40.3%  
40.0%

**Listed,  $d = 10$ :** vertices 57.2% vs pool

Vecer (2026), MCAP manuscript, Table 3 and Figure 3; 5,000 paths per design point.

# One likelihood language, three computational tasks

**Portfolio selection:**  $p^* = \arg \min_{p_R \in \mathcal{P}_{\text{trad}}} D_{\text{KL}}(P \| p_R),$

**Model aggregation:**  $m_t = \sum_i \pi_{i,t} p_{i,t}, \quad \pi_{i,t+1} \propto \pi_{i,t} p_{i,t}(X_{t+1}),$

**Adaptation:** prefunded restarts, conditional on no prior alarm:  
 $O_{\mathbb{P}}(\log \tau)$  rather than  $O(\tau)$  under the switch model.

Prices as likelihood ratios turn portfolio and hedge design into statistical computation.

Market-selection antecedents: Blume–Easley (1992), Sandroni (2000); Bayesian price probabilities: Massari (2021). **New here:** unique budget-balancing clearing, the exact  $\kappa_\lambda$ , and a  $K + 1$ -fund restart implementation.  
*Principles of Portfolio Choice: An Information-Theoretic,* janvecer.com

*Likelihood-Based Perspective*

CRC Press, forthcoming 2026

*Likelihood-finance page coming soon*

vecer@karlin.mff.cuni.cz

# Time permitting: What exact Ville validity requires


The prefunded process is an exact test martingale under the law  $\mathbb{P}_0$  generated by the declared benchmark kernels  $p_{0,t}$ . This requires  $p_{0,t}$  to be the true conditional law under the null.

Under a fixed i.i.d. law  $P$ , benchmark misspecification instead gives the best monitored relative growth rate

$$g^* = \sup_{\alpha \in \Delta_K} \mathbb{E}_P \log \frac{q_\alpha(X)}{p_0(X)} = D_{\text{KL}}(P \| p_0) - \inf_{\alpha \in \Delta_K} D_{\text{KL}}(P \| q_\alpha).$$

If  $g^* > 0$ , properly prefunded unbounded-duration accounts with local prior support cross every fixed finite threshold. An alarm says that the monitored convex hull outpredicts the incumbent benchmark; by itself, it does not identify an abrupt change.

# Time permitting: Monte Carlo confirms the lock-in coefficient



Dashed lines are theoretical  $\kappa_\lambda$  values; solid lines are simulated medians of  $H_\tau/\tau$ .

## Time permitting: The measure-theoretic payoff identity

Let  $\xi > 0$  be a state-price density under a reference law  $\mathbb{P}^\circ$ . For a nonzero nonnegative payoff  $F$  and a strictly positive benchmark  $Y$ , assume the following expectations are finite and define

$$F_0 = \mathbb{E}_{\mathbb{P}^\circ}[\xi F], \quad Y_0 = \mathbb{E}_{\mathbb{P}^\circ}[\xi Y],$$

$$\frac{d\mathbb{P}^F}{d\mathbb{P}^\circ} = \frac{\xi F}{F_0}, \quad \frac{d\mathbb{P}^Y}{d\mathbb{P}^\circ} = \frac{\xi Y}{Y_0}.$$

Then

$$\frac{F/F_0}{Y/Y_0} = \frac{d\mathbb{P}^F}{d\mathbb{P}^Y}.$$

The finite-outcome formulas are the discrete version of this Radon–Nikodym identity.

# Time permitting: Ex-post selection creates an option

For model returns  $R_i = p_i/m$ , ex-post maximum-likelihood selection pays

$$R_{\max}(\omega) = \max_i R_i(\omega).$$

This is not a static portfolio weight: it is a nonlinear best-of contract with price

$$C_{\max} = \mathbb{E}_m[R_{\max}] = \sum_{\omega} \max_i p_i(\omega) \geq 1.$$

## Two models

For  $R_1$  and  $R_2$ ,

$$\max(R_1, R_2) = R_2 + (R_1 - R_2)^+,$$

and

$$\mathbb{E}_m[(R_1 - R_2)^+] = D_{\text{TV}}(p_1, p_2).$$

Hence  $C_{\max} = 1 + D_{\text{TV}}(p_1, p_2)$ .

# Time permitting: Restart experiment details

## Search budget

- ▶ 15 candidate start times;
- ▶ 5 durations: 10, 20, 40, 60, 80;
- ▶ 75 accounts, each prefunded with  $1/75$ ;
- ▶ crossing threshold 20, so Ville gives an anytime false-alarm bound of 0.05.

## Continuous pooling prior

- ▶ Dirichlet(1, 1, 1) on  $\Delta_3$ ;
- ▶  $41 \times 41 = 1681$  positive quadrature nodes;
- ▶ Gauss–Jacobi  $\times$  Gauss–Legendre;
- ▶ exact for the likelihood polynomials through window length 80 in exact arithmetic.

The numerical integration expands the statistical type space, not the number of traded funds.