

Columbia University

Portfolio Optimization through Information Theory

Continuous Models (Stock Markets)

Jan Vecer

Charles University

Frankfurt School of Finance & Management

16 October 2026

Price is information; portfolio choice has two parts

A positive asset represents a probability law. Its normalized price is a likelihood ratio, and its expected log return under that law is relative entropy.

1. What should the investor think?

- In Merton's model, the investor supplies a physical drift μ .
- Later, we construct a view by selecting existing market information channels.

2. How should that view be implemented?

Match the selected density using a bond, stock-delivery claims, options, or dynamic trading.

A choice of probability law becomes a portfolio problem.

From two outcomes to a continuous density

Recall the binary example

	Up	Down
Market $\mathbb{P}^Y$	1/2	1/2
Target $\mathbb{P}^M$	2/3	1/3
Payoff L_M	4/3	2/3

The price is one:

$$\frac{1}{2} \frac{4}{3} + \frac{1}{2} \frac{2}{3} = 1.$$

The same construction with densities

$$L_M(z) = \frac{p_M(z)}{p_Y(z)}.$$

$$\mathbb{E}^Y[L_M] = \int \frac{p_M(z)}{p_Y(z)} p_Y(z) dz = 1.$$

Conversely, a unit-cost nonnegative payoff V defines

$$p_V(z) = V(z)p_Y(z).$$

A payoff and its density describe the same investment.

V and L_M are unit-cost payoffs in delivery-bond Y units. $\mathbb{P}^M$ denotes a selected investment law.

Relative entropy is expected logarithmic growth

The Kullback–Leibler divergence, or **relative entropy**, is

$$D_{\text{KL}}(\mathbb{P}^M \parallel \mathbb{P}^Y) = \int p_M(z) \log \frac{p_M(z)}{p_Y(z)} dz \geq 0.$$

For the investment associated with $\mathbb{P}^M$,

$$\frac{M(T)}{M(0)} = L_M(Z(T)) = \frac{p_M(Z(T))}{p_Y(Z(T))},$$

so

$$\mathbb{E}^M \left[\log \frac{M(T)}{M(0)} \right] = D_{\text{KL}}(\mathbb{P}^M \parallel \mathbb{P}^Y).$$

A realized log return is evidence; its expectation is information.

M is benchmark-denominated; otherwise use $(M(T)/Y(T))/(M(0)/Y(0))$.

A mismatch of densities is a loss of log growth

A unit-cost positive terminal payoff V implements $p_V = Vp_Y$. Hence

$$\begin{aligned}\mathbb{E}^M[\log V] &= \mathbb{E}^M\left[\log \frac{p_V}{p_Y}\right] \\ &= D_{\text{KL}}(\mathbb{P}^M \parallel \mathbb{P}^Y) - D_{\text{KL}}(\mathbb{P}^M \parallel \mathbb{P}^V).\end{aligned}$$

- If $L_M = p_M/p_Y$ is attainable, it maximizes expected log wealth under $\mathbb{P}^M$.
- With a restricted menu, choose the attainable p_V closest to p_M in this KL direction.

After selecting a view, implementation becomes density matching.

The stock law is the bond law tilted by the price

Start with $Z(t) \sim N(0, \sigma^2 t)$ under the bond law $\mathbb{P}^Y$. The normalized price is

$$R(t) = e^{Z(t) - \frac{1}{2}\sigma^2 t}, \quad \mathbb{E}^Y[R(t)] = 1.$$

Apply $p_V = V p_Y$ to the stock payoff $V = R(t)$:

$$p_X(z; t) = e^{z - \frac{1}{2}\sigma^2 t} p_Y(z; t) \implies \mathbb{P}^X : Z(t) \sim N(\sigma^2 t, \sigma^2 t).$$

Thus geometric Brownian prices are normal likelihood ratios, with

$$D_{\text{KL}}(\mathbb{P}^X(t) \parallel \mathbb{P}^Y(t)) = \frac{1}{2}\sigma^2 t.$$

The upward mean shift follows from the price tilt.

R is a relative price with $R(0) = 1$. The next slide makes its bond normalization and discounting explicit.

Where are interest rates and discounting?

Fix the delivery date T . Let $X(t)$ and $Y(t)$ be the currency prices of contracts delivering **one share** and **one currency unit** at T . For a constant interest rate r_f ,

$$Y(T) = 1, \quad Y(0) = e^{-r_f T} Y(T) = e^{-r_f T}.$$

The formulas use the **normalized relative price**

$$R(t) = \frac{X(t)/Y(t)}{X(0)/Y(0)}.$$

At delivery, this is the discounted return:

$$R(T) = e^{-r_f T} \frac{X(T)}{X(0)}, \quad \frac{X(T)}{X(0)} = e^{r_f T} R(T).$$

Discounting is in the bond benchmark; interest is not assumed zero.

$Y(0)$ is today's zero-coupon bond price; $Y(T)$ is the currency delivered. The GBM term $-\frac{1}{2}\sigma^2 t$ normalizes the likelihood, not the interest rate.

Merton's problem begins with a third normal law

Here $\mathbb{P}^M$ is the specified **physical law**:

$$\mathbb{P}^Y : Z(T) \sim \mathcal{N}(0, \sigma^2 T),$$


$$\mathbb{P}^X : Z(T) \sim \mathcal{N}(\sigma^2 T, \sigma^2 T),$$

$$\mathbb{P}^M : Z(T) \sim \mathcal{N}(\mu T, \sigma^2 T).$$

Write $q = p_Y$, $m = p_M$ and $f_\star = \mu/\sigma^2$.

The intended unit-cost payoff is

$$h(z) = \frac{m(z)}{q(z)} = e^{f_\star z - \frac{1}{2} f_\star^2 \sigma^2 T}.$$



Can the bond and stock reproduce the investor's chosen density?

Example: $\mu = 0.06$, $\sigma = 0.30$, $T = 5$ years. Under $\mathbb{P}^M$, $dR(t)/R(t) = \mu dt + \sigma dW^M(t)$.

Three ways to match the physical law

For a **static stock weight** w , write $R = R(T)$:

$$v_w = V_w q = (1 + w(R - 1))q = (1 - w)q + wp_X.$$

Choose $v_w \approx m$. All three problems use $0 \leq w \leq 1$ and $\mathbb{E}^Y[V_w] = 1$.

K : exact KL $J_K(w) = \int m \log \frac{m}{v_w} dz$

F : Fisher local metric $J_F(w) = \frac{1}{2} \int \frac{(v_w - m)^2}{m} dz = \frac{1}{2} \mathbb{E}^M \left[\left(\frac{v_w}{m} - 1 \right)^2 \right]$

Q : bond-law quadratic $J_Q(w) = \frac{1}{2} \int \frac{(v_w - m)^2}{q} dz = \frac{1}{2} \mathbb{E}^Y [(V_w - h)^2]$

The Fisher local metric uses $\mathbb{P}^M$ as its benchmark.

w_K, w_F, w_Q minimize the corresponding J . F is the Fisher information metric; Q is the bond-law payoff hedge.

Quadratic hedging is a regression line

Fit $V_w = (1 - w) + wR$ to $h = m/q$. Both have bond-law mean one. Ordinary least squares under the **bond law** gives slope and intercept

$$\beta_Q = \frac{\text{Cov}^Y(h, R)}{\text{Var}^Y(R)}, \quad \beta_0 = \mathbb{E}^Y[h] - \beta_Q \mathbb{E}^Y[R] = 1 - \beta_Q.$$

The residual is orthogonal to the stock direction:

$$\mathbb{E}^Y[(h - V_{\beta_Q})(R - 1)] = 0, \quad \mathbb{E}^Y[V_{\beta_Q}] = 1.$$



For our Gaussian example,

$$\beta_Q = \frac{\mathbb{E}^M[R] - 1}{\text{Var}^Y(R)} = \frac{e^{\mu T} - 1}{e^{\sigma^2 T} - 1} \simeq 0.615610.$$

The best straight-line hedge is the bond-law regression of the target payoff.

$\mu = 0.06$, $\sigma = 0.30$, $T = 5$. Positivity requires $w_Q = \text{clip}(\beta_Q, 0, 1)$; this example is interior.

The density fit and the payoff hedge are the same choice




$T = 5$ years: $w_K = 0.681$, $w_F = 0.684$, $w_Q = 0.616$.

The same weight selects a density mixture and a straight-line payoff.

$\mu = 0.06$, $\sigma = 0.30$. The target is $m/q = e^{f_\star(1-f_\star)\sigma^2 T/2} R^{f_\star}$; exact and local KL fits nearly overlap.

Short-horizon hedges recover the Merton fraction



$$w_K(T), w_F(T), w_Q(T) \longrightarrow f_* = \frac{\mu}{\sigma^2} = \frac{2}{3} \quad \text{as } T \downarrow 0.$$

$\mu = 0.06$, $\sigma = 0.30$. The endpoint at $T = 0$ is a limit. Closed-form quadratic weights are in the backup.

Continuous rebalancing implements the target exactly

With wealth and prices in **bond units**, hold stock fraction f :

$$\frac{dV_f(t)}{V_f(t)} = f\sigma dW^Y(t) \xrightarrow{\text{It\^o}} d\log V_f(t) = f\sigma dW^Y(t) - \frac{1}{2}f^2\sigma^2 dt.$$

Thus $V_f(0) = 1$ gives

$$V_f(T) = e^{fZ(T) - \frac{1}{2}f^2\sigma^2 T}.$$

Under $\mathbb{P}^M$, **maximize** the expected log payoff:

$$\mathbb{E}^M[\log V_f(T)] = \left(f\mu - \frac{1}{2}f^2\sigma^2\right) T, \quad (\mu - f\sigma^2)T = 0.$$

$$f_\star = \frac{\mu}{\sigma^2}, \quad V_{f_\star}(T) = \frac{m(Z(T))}{q(Z(T))} = h(Z(T)).$$

Continuous rebalancing attains the target at every horizon.

The negative quadratic term is the It\^o correction. Bond normalization does not assume zero interest.


Two Gaussian tail probabilities give the option price

Since $R = p_X/q$, for $A = \{R > k\}$,

$$\mathbb{E}^Y[R\mathbf{1}_A] = \int_A p_X(z) dz = \mathbb{P}^X(A).$$

The normalized call price is

$$\begin{aligned} c &= \mathbb{E}^Y[(R - k)^+] \\ &= \mathbb{P}^X(R > k) - k\mathbb{P}^Y(R > k) \\ &= \boxed{\Phi(d_1) - k\Phi(d_2)}, \end{aligned}$$



$$d_{1,2} = \frac{-\log k \pm \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}.$$


The stock-law probability prices the stock delivered on exercise.

$k = KY(0)/X(0)$, $c = C_{\text{currency}}(0)/X(0)$. Here $\sigma = 0.30$, $T = 1$, $k = 1$; $c = 0.1192$.

An option becomes a density only after dividing by its price

$$p_C(z) = \frac{(R(z) - k)^+ q(z)}{c},$$

$$p_P(z) = \frac{(k - R(z))^+ q(z)}{c_{\text{put}}}.$$





The price makes each density integrate to one.

$c_{\text{put}} = k\Phi(-d_2) - \Phi(-d_1)$ is the normalized put price. Same driver and parameters as before.

A static stock–bond hedge leaves a finite-horizon gap

For the call target $m = p_C$, its normalized payoff is $h = m/q = (R - k)^+/c$.

$$v_w = (1 - w)q + wp_X \iff V_w = \frac{v_w}{q} = 1 + w(R - 1).$$



Under p_C , $R > k = 1$ almost surely: increasing w raises V_w in every target state.

On $[0, 1]$, exact KL and bond-law L^2 choose $w = 1$. Unconstrained L^2 : $w = 5.19$, negative wealth below $R = 0.807$. The local $\mathbb{P}^M$ quadratic is infinite because $p_C = 0$ below the strike.

Dynamic delta hedging closes the gap in the normal model

For $\tau = T - t > 0$ and $r = R(t)$,


$$C(t, r) = r\Phi(d_1) - k\Phi(d_2),$$

$$d_1 = \frac{\log(r/k) + \frac{1}{2}\sigma^2\tau}{\sigma\sqrt{\tau}},$$

$$d_2 = d_1 - \sigma\sqrt{\tau}.$$

Delta is the exercise probability:

$$\begin{aligned}\Delta(t, r) &= \frac{\partial C}{\partial r} = \Phi(d_1) \\ &= \mathbb{P}^X(R(T) > k \mid \mathcal{F}_t).\end{aligned}$$



Plot: unit-cost value $C(t, r)/c$, with $c = C(0, 1)$.

The finite-horizon gap is a restriction of static trading.

$C_{\text{currency}}(t) = X(0) \frac{Y(t)}{Y(0)} C(t, r)$. Trade the share-delivery claim X and maturity- T bond Y .

Utility makes the preference explicit

For $\nu > 0$, constant relative risk aversion uses

$$U_{\gamma}(\nu) = \begin{cases} \frac{\nu^{1-\gamma} - 1}{1-\gamma}, & \gamma > 0, \gamma \neq 1, \\ \log \nu, & \gamma = 1. \end{cases}$$

Choose a payoff to solve

$$\max_{V>0} \mathbb{E}^M[U_{\gamma}(V)] \quad \text{subject to} \quad \mathbb{E}^Y[V] = 1.$$

With $h = m/q$, the pointwise first-order condition is

$$U'_{\gamma}(V) = \lambda \frac{q}{m} \implies V^{-\gamma} = \frac{\lambda}{h} \implies V \propto h^{1/\gamma}.$$

$$V_{\gamma} = \frac{h^{1/\gamma}}{\mathbb{E}^Y[h^{1/\gamma}]}, \quad p_U = V_{\gamma} q.$$

With an attainable optimum, power-utility choice is log-optimal under a selected law.

$\gamma = 1$ gives $\mathbb{P}^U = \mathbb{P}^M$. Physical law: $\mathbb{P}^M$. Selected investment law: $\mathbb{P}^U$.

Merton with CRRA: risk aversion scales the mean

The physical and selected laws are

$$\mathbb{P}^M : \mathcal{N}(\mu T, \sigma^2 T),$$

$$\mathbb{P}^U : \mathcal{N}\left(\frac{\mu T}{\gamma}, \sigma^2 T\right).$$


The constant stock fraction is

$$f_\gamma = \frac{f_\star}{\gamma} = \frac{\mu}{\gamma \sigma^2}.$$

In the example, $f_1 = 2/3$, $f_2 = 1/3$ and $f_4 = 1/6$.

For $\gamma > 1$, the mean moves toward zero; the variance stays fixed.

$\mu = 0.06$, $\sigma = 0.30$, $T = 5$ years. For $0 < \gamma < 1$, the selected mean moves beyond the physical mean.



Linear information flow: one construction, many drivers

Let $Z(t)$ have independent, stationary increments and a natural exponential family:

$$p_A(z; t) = \rho_t(z) e^{\theta_A z - t\kappa(\theta_A)}, \quad A \in \{X, Y\}.$$

The common carrier ρ_t cancels:

$$\log R(t) = \log \frac{p_X(Z(t); t)}{p_Y(Z(t); t)} = (\theta_X - \theta_Y)Z(t) - t[\kappa(\theta_X) - \kappa(\theta_Y)].$$

If $\mathbb{E}^X[Z(t)] = \nu_X t$, then

$$D_{\text{KL}}(\mathbb{P}^X(t) \parallel \mathbb{P}^Y(t)) = [(\theta_X - \theta_Y)\nu_X - \kappa(\theta_X) + \kappa(\theta_Y)] t.$$

Log relative price is affine in the evidence; expected information is linear in time.

Normal: $\theta_Y = 0$, $\theta_X = 1$, $\kappa(\theta) = \sigma^2 \theta^2 / 2$; slope 1, time term $-\sigma^2 t / 2$.

Four analytical families, organized by their variance

At fixed $t > 0$, let $\bar{z} = \mathbb{E}^A[Z(t)]$. The variance power classifies how variance changes with the mean.

Family	Driver law	Variance	Power ξ
Normal	$\mathcal{N}(\mu_A t, \sigma^2 t)$	$\sigma^2 t$	0
Poisson	$\text{Pois}(\lambda_A t)$	$\bar{z}$	1
Gamma	$\text{Gamma}(\alpha t, \beta_A)$	$\bar{z}^2 / (\alpha t)$	2
Inverse Gaussian	$\text{IG}(\delta t, \eta_A)$	$\bar{z}^3 / (\delta^2 t^2)$	3

These Tweedie cases give explicit densities, likelihoods, and information rates.

Variance = $c(t)\bar{z}^\xi$; common activity/dispersion parameters are fixed within each pair. Other linear-flow families exist.

Poisson: arrivals and waiting both carry information

Under $\mathbb{P}^A$, $N(t) \sim \text{Poisson}(\lambda_A t)$, with mass

$$p_A(n; t) = e^{-\lambda_A t} \frac{(\lambda_A t)^n}{n!}, \quad n = 0, 1, \dots$$

The log relative price is the log likelihood ratio:

$$\log R(t) = \log \left(\frac{\lambda_X}{\lambda_Y} \right) N(t) - (\lambda_X - \lambda_Y)t.$$

$$D_{\text{KL}}(\mathbb{P}^X(t) \parallel \mathbb{P}^Y(t)) = t \left[\lambda_X \log \frac{\lambda_X}{\lambda_Y} - \lambda_X + \lambda_Y \right].$$

An arrival multiplies the price; waiting without an arrival changes it too.

NEF: $\theta_A = \log \lambda_A$, $\kappa(\theta) = e^\theta$. $\lambda_X < \lambda_Y$: downward jumps, upward compensation.

Gamma: a continuum of positive jump sizes

Use common activity $\alpha > 0$ and law-specific rates $\beta_A > 0$:

$$p_A(z; t) = \frac{\beta_A^{\alpha t}}{\Gamma(\alpha t)} z^{\alpha t - 1} e^{-\beta_A z}, \quad z > 0.$$

For the gamma driver $\Gamma(t)$, the common factors cancel:

$$\log R(t) = (\beta_Y - \beta_X)\Gamma(t) + \alpha t \log \frac{\beta_X}{\beta_Y}.$$

$$D_{\text{KL}}(\mathbb{P}^X(t) \parallel \mathbb{P}^Y(t)) = \alpha t \left[\log \frac{\beta_X}{\beta_Y} + \frac{\beta_Y}{\beta_X} - 1 \right].$$

$\beta_X > \beta_Y$: positive driver jumps produce downward price jumps.

NEF: $\theta_A = -\beta_A$, $\kappa(\theta) = -\alpha \log(-\theta)$. Infinite activity; finite variation.

Inverse Gaussian: another additive jump driver

With common scale $\delta > 0$ and law-specific $\eta_A > 0$, use

$$p_A(z; t) = \frac{\delta t}{\sqrt{2\pi} z^{3/2}} \exp\left\{\delta\eta_A t - \frac{\delta^2 t^2}{2z} - \frac{\eta_A^2 z}{2}\right\}, \quad z > 0.$$

The common $1/z$ term cancels in the likelihood ratio:


$$\log R(t) = \frac{1}{2}(\eta_Y^2 - \eta_X^2)I(t) + \delta(\eta_X - \eta_Y)t.$$

$$D_{\text{KL}}(\mathbb{P}^X(t) \parallel \mathbb{P}^Y(t)) = \frac{\delta t}{2\eta_X}(\eta_X - \eta_Y)^2.$$

A different jump-size law gives the same affine structure for log relative price.


NEF: $\theta_A = -\eta_A^2/2$, $\kappa(\theta) = -\delta\sqrt{-2\theta}$. $\mathbb{E}^A[I(t)] = \delta t/\eta_A$.

Four price evolutions: the original book figure



Book Figure 5.1, p. 216. Original $\mathbb{P}^Y$ paths: $\sigma = 0.10$, with KL rate approximately 0.005 per time unit.

The same book paths, viewed as OHLC candles



Book Figure 5.2, p. 217. Original panel scales. Visual similarity does not imply statistical indistinguishability.

Two ways to choose the investment law

Informative choice: Merton

$$\mathbb{P}^M : \quad Z(T) \sim N(\mu T, \sigma^2 T).$$

The investor supplies μ , a view about the physical drift.

Its payoff is p_M/p_Y ; preferences can temper this likelihood.

All stock laws give stock; all bond laws give the bond. Selection chooses between these endpoints.

No historical-return fit is added. Choosing channels still requires a model and investor judgment.

Uninformative channel selection

A **channel** is an independent driver with its own bond law and stock law.

Mute a channel: keep its bond law.

Retain it: use its stock law.

Independent information channels combine into one price

Let $B_0(t) = \sigma_0 W^Y(t)$, and let Γ_-, I_-, Γ_+ be independent positive jump drivers under $\mathbb{P}^Y$.

$$\ell(t) := \log R(t) = -\Gamma_-(t) - I_-(t) + \Gamma_+(t) + B_0(t) - \kappa t.$$

For each signed driver $D_j(t)$, set

$$\kappa_j = \log \mathbb{E}^Y[e^{D_j(1)}], \quad L_j(t) = e^{D_j(t) - \kappa_j t}, \quad \kappa = \sum_j \kappa_j.$$

$$R(t) = \prod_j L_j(t), \quad \frac{d\mathbb{P}_j^X}{d\mathbb{P}_j^Y} = L_j(t).$$

Likelihoods multiply; independent channel information adds.

Downward channels describe index skew; a small upward Gamma describes the right tail; the normal channel describes diffusion. These are modeled sources, not uniquely observed factors.

Options determine prices; the model describes channels

With $r = R(T) = e^{\ell(T)}$ and normalized call price $c(k)$,

$$c(k) = \int (r - k)^+ q(r) dr, \quad \boxed{c''(k) = q(k)}.$$

- A smooth continuum of calls determines the bond density q .
- Actual quotes constrain finitely many integrals; a convolution model fills the gaps and tails.
- The channel split comes from this model, rather than being separately observed in option prices.

No historical-return estimate is needed to construct the selected view.

The worked example fits one option cross-section at one maturity. Model identification and physical forecasting are separate questions.

Four channels in the S&P 500

18 January 2019; maturity 18 April 2019 (90 days).


Channel	Shape / scale	Bond	Stock	KL (log bps)
Gamma down	$\alpha_- T = 0.572$	$\beta_Y = 9.88$	$\beta_X = 10.9$	25.8
IG down	$\delta T = 0.0135$	$\eta_Y = 1.77$	$\eta_X = 2.26$	7.36
Gamma up	$\alpha_+ T = 0.0203$	$\beta_Y = 10.9$	$\beta_X = 9.88$	0.977
Normal	$\sigma_0^2 T = 0.00128$	mean 0	mean $\sigma_0^2 T$	6.42

$$B_0(T) \sim \begin{cases} N(0, \sigma_0^2 T), & \mathbb{P}^Y, \\ N(\sigma_0^2 T, \sigma_0^2 T), & \mathbb{P}^X, \end{cases} \quad \sigma_0^2 T = 0.00128.$$

Total information: 40.5 log basis points; 82% comes from downward channels.


One log bp is 10^{-4} of expected log growth over these 90 days. Gamma uses shape/rate; IG uses scale δT and parameter η .

The individual drivers on the same absolute scale



Raw signed drivers: negative for downward jumps, positive for upward jumps. Identical linear axes. Gamma peaks diverge at zero; peaks above 50 are clipped.

The channels combine into the market density



18 January–18 April 2019 | Gamma down + IG down + Gamma up + normal

$$p_X(r) = r q(r), \quad \int r q(r) dr = 1.$$

One choice: keep normal, mute all three jump channels

Law	Down G	Down IG	Up G	Normal
Bond	Y	Y	Y	Y
Selected view	Y	Y	Y	X
Stock	X	X	X	X

The jump drivers keep their bond laws. Only the normal driver changes from

$$N(0, \sigma_0^2 T) \quad \text{to} \quad N(\sigma_0^2 T, \sigma_0^2 T).$$

Jump randomness remains in the price; its stock-law tilt is muted.

Example mandate: accept the market-implied jump laws as fair and seek only the diffusive tilt. This does not remove jump risk from the portfolio.

The selected law gives a terminal-price payoff

Write q^ℓ for the density of $\ell(T) = \log r$. Only the normal mean changes, by $v = \sigma_0^2 T$, so


$$m^\ell(x) = q^\ell(x - v), \quad g(r) = \frac{m(r)}{q(r)} = \frac{q^\ell(\log r - v)}{q^\ell(\log r)}.$$

$$\mathbb{E}^Y[g] = 1, \quad D_{\text{KL}}(\mathbb{P}^M \| \mathbb{P}^Y) = \mathbb{E}^M[\log g] = 2.44 \log \text{bps}.$$


The full log-return distribution shifts; its variance stays unchanged.

m is the selected terminal-price density. B_0 remains the earlier unshifted normal driver under the bond law. The fixed compensator is already included in ℓ .

One selected density, one intended payoff



Retain the normal channel's stock law; use the bond law for all three jump channels.



In this fit, a large r increasingly points to an upward jump rather than normal variation; the conditional normal likelihood falls beyond the shoulder.

Only some channel information is visible in the total price

If we observed the normal driver separately, its selected likelihood would be

$$\Lambda_0 = e^{B_0(T) - \frac{1}{2}\sigma_0^2 T}.$$

An option observes only $r = e^{\ell(T)}$. The implementable target is

$$g(r) = \mathbb{E}^Y[\Lambda_0 \mid r].$$

The information chain rule gives

$$\underbrace{6.42}_{\text{normal channel}} = \underbrace{3.98}_{\text{unobserved component}} + \underbrace{2.44}_{\text{terminal-price view}} \quad (\log \text{ bps}).$$

The total price does not reveal every driver separately.

The unobserved component is information unavailable from terminal price, not market inefficiency.

A portfolio is a mixture of option-block densities

Use nonnegative piecewise-linear blocks B_i : butterflies and tail wedges. Set $V = \sum_i u_i B_i$, $c_i = \mathbb{E}^Y[B_i] > 0$ and $p^i = qB_i/c_i$.

$$\underbrace{\sum_i c_i u_i = 1}_{\text{unit budget}} \iff w_i = c_i u_i \geq 0, \quad \sum_i w_i = 1.$$


$$p_V = qV = \sum_i w_i p^i.$$

$$\min_w D_{\text{KL}}(\mathbb{P}^M \parallel \mathbb{P}_w^V) \quad \text{or locally} \quad \frac{1}{2} \mathbb{E}^M \left[\left(\frac{p_{V_w}}{p_M} - 1 \right)^2 \right].$$

The same matching problem as Merton, now with more instruments.

Block coefficients u_i are payoff units; $w_i = c_i u_i$ are capital weights. Bond and stock belong to the same attainable family.

Replicating the normal-only payoff with nine strikes



Dots mark observed strike levels; option prices are generated by the fitted bond law.

All three criteria give 0.0152 log bp loss and 2.42 log bps attained growth.

Exact KL fit with nine observed strikes, valued under the same Gamma-IG market law. Dots identify strike locations.

Which strikes and how many option units?

Asset / call strike K	k	Position
Bond	—	1
Stock	—	-0.0216
1870	0.7	0.00989
2060	0.771	-0.00517
2250	0.842	-0.00162
2445	0.915	-0.0678
2635	0.987	0.551
2825	1.06	0.371
3000	1.12	-0.783
3200	1.2	-0.842
3400	1.27	0.824

Normal-only, exact KL. K is the SPX strike; $k = K Y(0)/X(0)$, with $X(0)/Y(0) = 2671.05$. Rounded holdings multiply $(r - k)^+$; these are payoff units, not contract counts or capital weights.

From normal-channel information to an option portfolio




Main waterfall uses one linear scale; only the right-hand detail has a magnified, truncated vertical axis.

$$\underbrace{6.42}_{\text{selected source}} = \underbrace{3.98}_{\text{unobserved component}} + \underbrace{0.0152}_{\text{implementation loss}} + \underbrace{2.42}_{\text{attained growth}}$$

Choose a law, then implement its observable part with options.

Log basis points over 90 days under the normal-only view; model quantities, not realized performance.

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

The local information quadratic uses the target as benchmark

OPTIONAL DETAIL

Let $\varepsilon = p_V/p_M - 1$, with $\int p_V = \int p_M = 1$ and $\mathbb{P}^V \ll \mathbb{P}^M$.

$$\mathbb{E}^M[\varepsilon] = 0, \quad D_{\text{KL}}(\mathbb{P}^M \parallel \mathbb{P}^V) = \mathbb{E}^M[-\log(1 + \varepsilon)].$$

Expanding about $\varepsilon = 0$ gives

$$D_{\text{KL}}(\mathbb{P}^M \parallel \mathbb{P}^V) \simeq \frac{1}{2} \mathbb{E}^M[\varepsilon^2] = \frac{1}{2} \int \frac{(p_V - p_M)^2}{p_M} dz.$$

Only $\mathbb{P}^M$ and $\mathbb{P}^V$ enter this local metric.

The quadratic is exactly $\frac{1}{2} \chi^2(\mathbb{P}^V \parallel \mathbb{P}^M)$. The approximation requires small relative density error and suitable integrability.

The precise mean–variance connection

OPTIONAL DETAIL

For $m = qg$ and unit-cost V , expanding the squared error gives

$$\min_V \mathbb{E}^Y[(V - g)^2] \iff \max_V \left\{ \mathbb{E}^M[V] - \frac{1}{2} \text{Var}^Y(V) \right\}.$$

This is a **pricing-law quadratic hedge**: mean under $\mathbb{P}^M$, variance under $\mathbb{P}^Y$.

For $\varepsilon = V - 1$, expected log wealth near the bond instead gives

$$\mathbb{E}^M[\log(1 + \varepsilon)] \simeq \mathbb{E}^M[\varepsilon] - \frac{1}{2} \mathbb{E}^M[\varepsilon^2].$$

The latter second moment is under $\mathbb{P}^M$. Replacing it by a $\mathbb{P}^Y$ moment needs an additional near-bond/near-market approximation.

The Fisher hedge is a regression under the target law

OPTIONAL DETAIL

Define centered density-ratio coordinates

$$x_F = \frac{p_X - q}{m} = \frac{R - 1}{h}, \quad y_F = \frac{m - q}{m} = \frac{h - 1}{h}, \quad \mathbb{E}^M[x_F] = \mathbb{E}^M[y_F] = 0.$$

Then the same unit-cost portfolio satisfies

$$\frac{v_w}{m} - 1 = wx_F - y_F, \quad J_F(w) = \frac{1}{2} \mathbb{E}^M[(wx_F - y_F)^2].$$

Its unconstrained regression slope is

$$\beta_F = \frac{\text{Cov}^M(y_F, x_F)}{\text{Var}^M(x_F)}, \quad w_F = \text{clip}(\beta_F, 0, 1).$$

The benchmark stays $\mathbb{P}^M$; relative errors determine the regression weight.

In payoff coordinates: weights h^{-2} under $\mathbb{P}^M$, intercept $1 - w$ fixed. Require full target support and finite second moments.

Gaussian hedging: the two quadratic solutions

OPTIONAL DETAIL

Let $s = \sigma^2 T$, $f_\star = \mu/\sigma^2$, $q = \mathcal{N}(0, s)$, $p_X = \mathcal{N}(s, s)$ and $m = \mathcal{N}(f_\star s, s)$.

$$A = \int \frac{q^2}{m} = e^{f_\star^2 s}, \quad B = \int \frac{qp_X}{m} = e^{f_\star(f_\star - 1)s}, \quad C = \int \frac{p_X^2}{m} = e^{(f_\star - 1)^2 s}.$$

Consequently,

$$J_F(w) = \frac{1}{2}\{A(1-w)^2 + 2Bw(1-w) + Cw^2 - 1\}, \quad w_F = \frac{A - B}{A - 2B + C}.$$

For the bond-law payoff hedge,

$$J_Q(w) = \frac{1}{2}\{(e^s - 1)w^2 - 2(e^{f_\star s} - 1)w + e^{f_\star^2 s} - 1\},$$

$$w_Q = \frac{e^{f_\star s} - 1}{e^s - 1} = \frac{\mathbb{E}^M[R] - 1}{\text{Var}^Y(R)}.$$

Project the displayed unconstrained quadratic solutions onto $[0, 1]$ for static nonnegative wealth.

Exact KL and the short-horizon limit

OPTIONAL DETAIL

For $V_w = 1 + w(R - 1)$,

$$J'_K(w) = -\mathbb{E}^M \left[\frac{R - 1}{V_w} \right], \quad J''_K(w) = \mathbb{E}^M \left[\frac{(R - 1)^2}{V_w^2} \right] > 0.$$

For $0 < f_\star < 1$ the optimizer is the unique interior root. As $T \downarrow 0$,

$$J_K(w), J_F(w), J_Q(w) = \frac{1}{2} \sigma^2 T (w - f_\star)^2 + o(T).$$

For $s = \sigma^2 T$, the first corrections are


$$\begin{aligned} w_K &= w_F + O(s^2) = f_\star + \frac{1}{2} f_\star (1 - f_\star) (2f_\star - 1) s + O(s^2), \\ w_Q &= f_\star - \frac{1}{2} f_\star (1 - f_\star) s + O(s^2). \end{aligned}$$

The local information metric matches the first correction to exact KL.

With general $f_\star$, all three constrained limits are $\text{clip}(f_\star, 0, 1)$. At $T = 0$, every weight has zero loss.

The channel likelihoods reveal the price tilt

OPTIONAL DETAIL



Bond law supplies the benchmark. The density of each driver determines how often each likelihood occurs.

Here z is the raw signed driver. The fixed compensator normalizes the likelihood. The individual bond and stock densities appear in the main lecture.

The three criteria give almost the same normal-only fit

OPTIONAL DETAIL

Nine observed strikes; positive wealth; model prices.

Criterion	Exact KL loss	Attained growth
Exact KL	0.0152	2.42
Local information quadratic	0.0152	2.42
Bond-law quadratic	0.0152	2.42


$$\mathbb{E}^M[\log V] = \underbrace{D_{\text{KL}}(\mathbb{P}^M \parallel \mathbb{P}^Y)}_{2.44 \text{ log bps}} - \underbrace{D_{\text{KL}}(\mathbb{P}^M \parallel \mathbb{P}^V)}_{\text{implementation loss}}.$$

Compare all methods using exact KL loss under the selected view.

Expected-log basis points over 90 days. Rounded to three significant figures; small displayed differences are not economic evidence.

The fitted curve is an explicit option portfolio

OPTIONAL DETAIL




Call purchases build the shoulder; higher-strike sales reduce upside exposure.

Bond, stock and signed calls implement positive total wealth. Call holdings are changes in slope; the main-talk table identifies the actual strikes and holdings.

Positive blocks allow nonconvex payoffs and linear tails

OPTIONAL DETAIL



$$V(r) = \sum_{i=0}^n u_i B_i(r) + s_{\infty} (r - k_n)^+, \quad u_i \geq 0, \quad s_{\infty} \geq 0.$$

The hats interpolate nodes at $0, k_1, \dots, k_n$; the last hat has a flat right shoulder. The additional call sets the right-tail slope. Butterflies have signed vanilla legs.

From nodal values to bond, stock and signed calls

OPTIONAL DETAIL

Let $k_0 = 0 < k_1 < \dots < k_n$, with node values $u_i = V(k_i) \geq 0$. Define slopes $s_i = (u_{i+1} - u_i)/(k_{i+1} - k_i)$ for $i = 0, \dots, n-1$, and a right-tail slope $s_n = s_\infty \geq 0$.

$$V(r) = u_0 + s_0 r + \sum_{i=1}^n (s_i - s_{i-1})(r - k_i)^+.$$

- Nonnegative nodes keep every interpolation interval nonnegative.
- The right tail is $u_n + s_\infty(r - k_n)$; bond and stock are exact.
- Model cost: $u_0 + s_0 + \sum_i (s_i - s_{i-1})c(k_i) = 1$.

The code checks this representation against the positive blocks and prices it independently.

A continuum of options replicates a smooth payoff

OPTIONAL DETAIL

For a sufficiently regular h and an anchor $k_* > 0$,

$$\begin{aligned} h(r) &= h(k_*) + h'(k_*)(r - k_*) \\ &\quad + \int_0^{k_*} h''(k)(k - r)^+ dk \\ &\quad + \int_{k_*}^{\infty} h''(k)(r - k)^+ dk. \end{aligned}$$

Curvature determines the signed option holdings.

Breeden–Litzenberger; Carr–Madan. For exchanging price and integral, require

$$\int_0^{k_*} |h''(k)|c_{\text{put}}(k) dk + \int_{k_*}^{\infty} |h''(k)|c(k) dk < \infty.$$

What was checked in the numerical construction?

OPTIONAL DETAIL

- One January 2019 Gamma–IG hybrid calibration for every target and instrument.
- Density masses, the stock tilt $p_X(r) = r q(r)$, and the normal-only shift identity.
- The same observed strikes, positive blocks, linear tails and budget in all three optimizations.
- Independent refined quadrature, convex optimality residuals and signed-call reconstruction costs.
- Bond and stock are represented exactly; positivity holds for all positive terminal prices.

These are model-consistent calculations, not an executable bid/ask backtest.

Reproducible code, fitted parameters, nodal payoffs, option positions and numerical diagnostics accompany the source package.

Selected references

OPTIONAL DETAIL

- Kelly (1956), A new interpretation of information rate.
- Cover and Thomas (2006), *Elements of Information Theory*.
- Geman, El Karoui and Rochet (1995), Changes of numeraire, changes of probability measure and option pricing.
- Breeden and Litzenberger (1978), Prices of state-contingent claims implicit in option prices.
- Carr and Madan (2001), Optimal positioning in derivative securities.
- Vecer, *Principles of Portfolio Choice: An Information-Theoretic, Likelihood-Based Perspective*.
- Merton (1971), Optimum consumption and portfolio rules in a continuous-time model.