Revisiting Meucci's Entropy Pooling

Constructing Posterior Return Distributions from Economic Regime Views

Research Paper

July 29, 2026

Redouane Raki

redouane.raki@saa-forum.com

Key Words: entropy pooling, prior and posterior distribution, regime based entropy pooling, market regimes, probability of regimes, crash, bearish, bullish, boom, S&P 500.

Table of Contents

Abstract

Entropy Pooling has become one of the standard methodologies for incorporating subjective investor views into empirical return distributions while preserving maximum consistency with historical data. Existing applications typically formulate views directly on return moments, such as expected returns, volatilities, value at risk or correlations.

This paper proposes a regime-based application of Meucci's Entropy Pooling framework. Instead of expressing views on return statistics, investors formulate views directly on economically meaningful market regimes, such as Crash, Bearish, Bullish and Boom. The posterior regime probabilities are subsequently translated into a complete posterior return distribution by proportionally redistributing probability mass across historical observations belonging to each regime.

The proposed framework preserves the mathematical foundations of Entropy Pooling while offering a more intuitive and economically interpretable specification of investor beliefs. Since investment decisions are frequently expressed in terms of market conditions rather than precise return forecasts, regime-based views provide a natural interface between qualitative market expectations and quantitative portfolio construction.

An empirical application to monthly S&P 500 returns illustrates that the proposed approach generates coherent posterior return distributions, preserves broad scenario diversification and produces economically plausible changes in portfolio risk and return characteristics.

1 Meucci Entropy Pooling Model

Attilio Meucci's (2008) entropy pooling framework makes it possible to combine an investor's subjective views with a prior distribution without relying on parametric normal distribution assumptions. The method operates directly on a discrete scenario pool and is thus considerably more flexible than the classic Black-Litterman model. The central idea is as follows: One searches for a new probability distribution over the available return scenarios that (a) is as close as possible to the prior distribution and (b) satisfies all views. "Close" is measured here by the Kullback–Leibler divergence.

1.1 Model Specification

1.1.a The Prior:

We assume T historical or simulated scenarios. Each scenario t is assigned a prior probability: p = (p_1, p_2, …, p_T). All p_t are positive and sum to 1. In [1]this framework we work with an equally weighted uniform distribution is used (p_t = 1/T), i.e., each historical scenario contributes equally to the prior distribution. By working directly with scenarios, arbitrary distributions—including non-normal distributions, fat tails, and nonlinear dependencies—can be accurately represented. No parametric model is necessary.

1.1.b The Views:

The investor formulates K views. Unlike the Black-Litterman model, Entropy Pooling allows views to be defined directly on any measure of the distribution: expected values, variances, covariances, quantiles, or event probabilities. In this paper, we focus on views based on expected values. These are formulated as follows: Aq = b, where:

q ist the vector of the Posterior probabilities we are searching for

A ist the matrix of the available (empirical or simulated) returns

b ist the vector of the subjective Views

1.1.c The problem to solve:

q ist he solution oft he following "Kullback–Leibler"optimization problem:

$$q = argmin\{\sum_{k=1}^{T} q_k . \ln\left(\frac{q_k}{p_k}\right)\}$$

such that:

$\left[\right.$**Aq = b** (Views on returns)$\left.\right]$ and $\left[\right.$**1'.q = 1** (All probabilities sum to one)$\left.\right]$

[1] *Other models use for example the exponential weighted moving average method*

1.2 The Entropy solution

Step 1: Lagrange transformation: Inclusion of the constraints in the Utility function using Lagrange penalty parameter λ and μ

$$L(q, \lambda, \mu) = \sum_{k=1}^{T} q_k . \ln\left(\frac{q_k}{p_k}\right) + \lambda' . (Aq - \mathrm{b}) + \mu . (1'\mathrm{q} - 1)$$

Step 2: Optimality condition: The Gradient function relative to the k^{th}-q element is equal to zero

$$\frac{\partial L(q,\lambda,\mu)}{\partial q_k} = 0$$

$$==> \ln\left(\frac{q_k}{p_k}\right) + 1 + \lambda' . A_k + \mu = 0$$

$$==> q_k = p_k \exp(-1-\mu) \exp(-\lambda' A_k)$$

Since q_k sum to one, we necesserly have:

$$\sum_{k=1}^{T} p_k \exp(-1-\mu) \exp(-\lambda' A_k) = 1$$

$$==> \exp(-1-\mu) = \frac{1}{\sum_{k=1}^{T} p_k . \exp(-\lambda' A_k)}$$

$$==> q_k = \frac{p_k \exp(-\lambda' A_k)}{\sum_{k=1}^{T} p_k . \exp(-\lambda' A_k)}$$

This means that we have transformed our optimization problem from one with q as a variable into one with lambda as a variable, which represents a significant reduction in algebraic dimension.

Step 3: Views Dual Problem

We define: $Z(\lambda) = \sum_{k=1}^{T} p_k . \exp(-\lambda' A_k)$

$$==> q_k = \frac{p_k \exp(-\lambda' A_k)}{Z(\lambda)}$$

$$==> \frac{q_k}{p_k} = \frac{\exp(-\lambda' A_k)}{Z(\lambda)}$$

$$==> \ln\left(\frac{q_k}{p_k}\right) = -\lambda' A_k - \ln(Z(\lambda))$$

$$==> \sum_{k=1}^{T} q_k . \ln\left(\frac{q_k}{p_k}\right) = -\lambda' \sum_{k=1}^{T} q_k . A_k - \ln(Z(\lambda)) \sum_{k=1}^{T} q_k$$

$$==> \sum_{k=1}^{T} q_k . \ln\left(\frac{q_k}{p_k}\right) = -\lambda' \mathrm{A.q} - \ln(Z(\lambda))$$

and now, let´s go back to the Lagrange Function:

$$L(q, \lambda, \mu) = \sum_{k=1}^{T} q_k . \ln\left(\frac{q_k}{p_k}\right) + \lambda' . (Aq - \mathrm{b}) + \mu . (1'\mathrm{q} - 1)$$

$$L(q, \lambda, \mu) = -\lambda' \mathrm{A.q} - \ln\big(Z(\lambda)\big) + \lambda' . (Aq - \mathrm{b}) + \mu . (1'\mathrm{q} - 1)$$

$$L(q, \lambda, \mu) = -\ln\big(Z(\lambda)\big) - \lambda'.\,b$$

$$L(q, \lambda, \mu) = G(\lambda) = -\ln\big(Z(\lambda)\big) - \lambda'.\,b$$

Since we work with the dual problem, we consider F(l) = -G(l) as the new utility function to optimize:

$$\phi(\lambda) = \ln\big(Z(\lambda)\big) + \lambda'.\,b$$

$$\phi(\lambda) = \ln\left(\sum_{k=1}^{T} p_k.\exp(-\lambda' A_k)\right) + \lambda'.\,b$$

Step 4: Adding uncertainty to the views: Soft Views Pooling

The uncertain (soft) views are expressed through adding a stochastic term to the vector b:the equation: Aq=b+ε , ε~N(0,Ω) and the new Lagrange function takes the following form:

$$\phi(\lambda) = \ln\big(\textstyle\sum_{k=1}^{T} p_k.\exp(-\lambda' A_k)\big) + \lambda'.\,b + \frac{1}{2}\,\lambda'\Omega\lambda$$

Step 5: Entropy formula

$\phi(\lambda)$ is strictly convex and therefore has a unique minimum l*, which the Newton algorithm, for example, can easily find. The last step is simply substituting the vecktor l* in the formula of q_k:

$$q_k = \frac{p_k.\exp\,(-\lambda^{*\prime} A_k)}{\sum_{k=1}^{T} p_k.\exp\,(-\lambda^{*\prime} A_k)}$$

All distribution moments (means, volatilities, value at risk or correlations) will the be adjusted based on the new probabilities.

1.3 Entropy Quality

The *Entropy Diagnostics* section provides a quantitative assessment of how strongly the Entropy Pooling procedure has modified the original probability distribution. These diagnostics help evaluate whether the imposed views are consistent with the historical data and whether the resulting posterior distribution remains sufficiently diversified.

Effective Number of Scenarios: ENS = exp(-Σ q_i ln(q_i)). Measures the effective diversification of the posterior probability distribution. In uniform distribution with T observations, we have q_i =1/T. $==> ENS = T$

Effective Scenario Ratio: ESR = ENS / T. Measures the proportion of scenarios that remain effectively relevant after incorporating the views. High ratio means broad scenario diversification.

Kullback-Leibler Divergence: KLD(q , p) = Σ q_i ln(q_i/p_i). Measures the informational distance between prior and posterior distributions. Low KLD Values mean posterior remains close to prior.

Maximum Posterior Weight: max(q_i). Low maximum weight means no single scenario dominates.

A healthy Entropy Pooling solution typically exhibits a high Effective Scenario Ratio, low to moderate KL Divergence and moderate Maximum Posterior Weight.

1.4 Illustration and example

We consider the **monthly** returns of the us S&P500 equity index between January 1999 and March 2026 and a **bearish** view_return with a **high confidence**.

Historical_Return	Historical_Volatility
0,61%	4,35%

View_Return	View_Uncertainty
-1,00%	0,50%

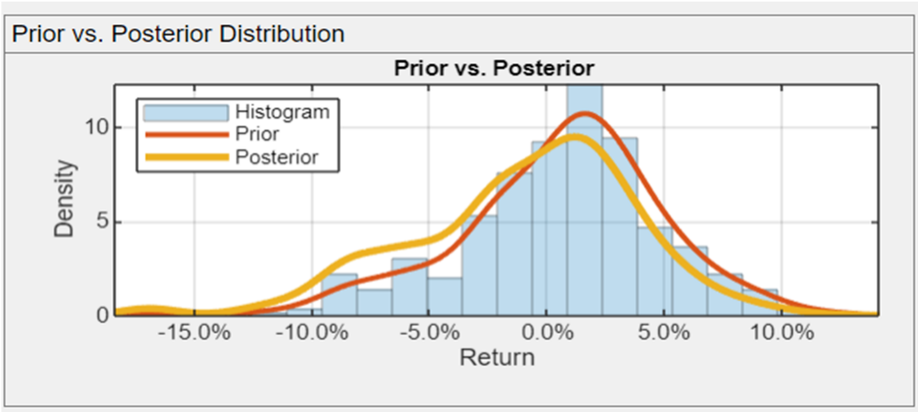

Entropy Impact on the distribution	Metric	Prior	Posterior	Delta
	Return	0,61%	-0,98%	-1,59%
	Volatility	4,34%	4,81%	0,47%
	VaR95	7,90%	9,18%	1,28%
	CVaR95	9,61%	12,01%	2,41%

- The posterior expected return moved from 0.61% to -0.98%. The view fulfillment ratio reached 98.80%, indicating that the posterior distribution strongly incorporated the view.
- the bearish view shifted probability mass toward adverse scenarios, increasing both tail risk measures (VaR and CvaR)

Entropy Diagnostics	Metric	Value
	Number of scenarios	327
	ENS: Effective number of scenarios	307
	ESR: Effective scenario ratio	93,9%
	Maximum posterior weight	1,2%
	KL divergence D(q||p)	6,3%

- The high ESR indicates that almost the entire scenario set remains active after conditioning.
- The maximum posterior weight of only about 1.24% shows that no single scenario dominates the posterior distribution.
- The KL divergence is very low, indicating that the posterior distribution remains close to the empirical prior.

2 Regime Entropy Pooling Model

Investors rarely formulate expectations in terms of expected monthly returns. Instead they think in terms of market states: recession, recovery , crisis or expansion. This motivates replacing moment-based views by regime-based views

2.1 Views on Regimes instead of returns

Let us start with the following notations:

P_q(Crash) = b_1, P_q (Bearish) = b_2, P_q (Bullish) = b_3. ==>P_q (Boom) $b_4 = 1- b_1- b_2- b_3$

The core philosophy of Entropy Pooling remains exactly the same. Posterior probabilities are obtained by minimizing the Kullback-Leibler divergence subject to regime constraints

Historical observations are classified into four regimes: Crash, Bearish, Bullish and Boom. Using a regime indicator matrix A, the constraints become: $Aq \approx b$ with uncertainty matrix Ω.

The optimization reallocates probability mass across regimes while remaining as close as possible to the empirical prior distribution.

2.2 Deriving posterior retruns from posterior regimes

After posterior regime probabilities have been determined, they are distributed across all historical observations belonging to each regime. Let g(i) denote the regime of observation i. The posterior probability of observation i can be approximated by:

$$q_i = p_i \frac{P_{post}(g(i))}{P_{prior}(g(i))}$$

where:

p_i = prior probability of observation I,

$P_{Post}(g(i))$ = posterior probability of its regime,

$P_{Prior}(g(i))$ = prior probability of its regime

All observations within a regime therefore receive the same multiplicative adjustment factor. Consquently, the regime posterior probability is distributed proportionally across all historical returns belonging to that regime. and once posterior probabilities q have been obtained, the complete posterior distribution is now available and the framework therefore allows a complete comparison of prior and posterior risk-return characteristics.

3 Case study

In this section we consider a "Post Crisis Scenario" regime, which is a transitional market environment following a major economic or financial downturn, characterized by recovering economic activity, improving investor sentiment, supportive policy conditions, and gradually declining market risk.

Input Views: Post Crisis Views with high Confidence

Regime	Prior_Probability	View_Probability	Uncertainty
Crash	19,81%	10%	5,00%
Bearish	30,35%	25%	5,00%
Bullish	30,99%	50%	5,00%

- The prior probabilities come directly from the empirical clustering of the considered regimes.
- The probability of a boom is derived from these tree probabilities as the residual value to one.

Regime Diagnostics:

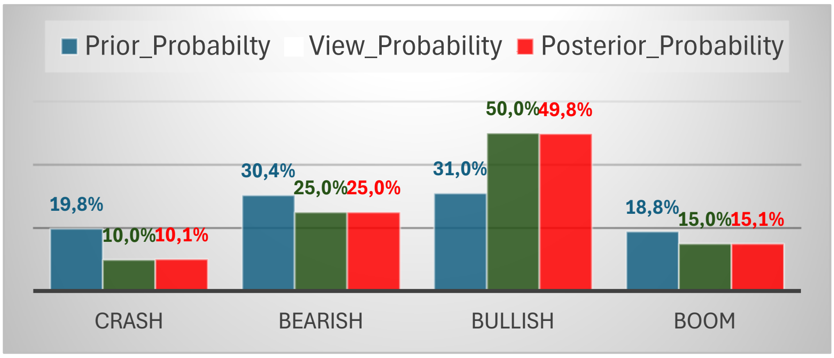

- The posterior regime probabilities suggest a constructive but balanced market outlook. The probability mass is redistributed toward favorable regimes while adverse regimes receive lower weights. Importantly, the resulting distribution remains diversified across multiple market states, preserving the economic realism of the analysis

Entropy Diagnostics:

Metric	Value
Number of scenarios	313
Effective number of scenarios	287
Effective scenario ratio	91,7%
Maximum posterior weight	0,5%
Minimum posterior weight	0,2%
KL divergence D(q||p)	8,6%

- The entropy diagnostics indicate a healthy posterior structure. A high Effective Scenario Ratio implies that most historical scenarios continue to contribute to the posterior distribution. A low maximum posterior weight confirms that no individual scenario dominates the results, while a moderate KL divergence suggests a meaningful but controlled departure from the empirical prior distribution.

Return Distributions:

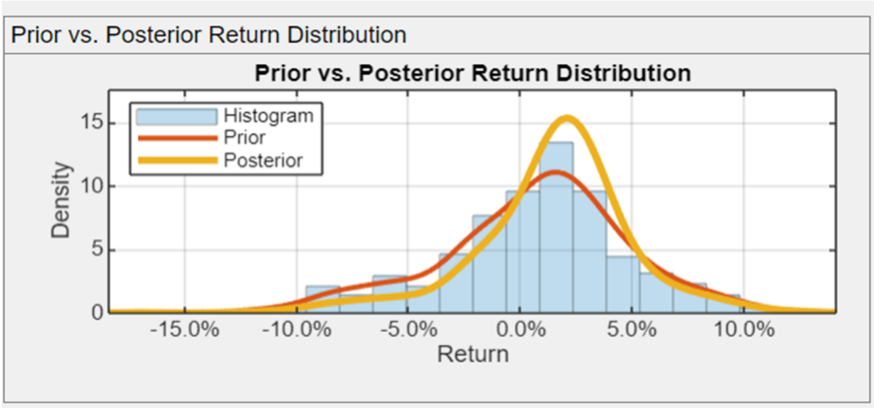

	Prior	Posterior
Return	0,6%	1,4%
Volatility	4,3%	3,5%
VaR_95	7,9%	5,7%
CVaR_95	9,6%	8,1%

- The posterior distribution exhibits an improved risk-return profile. Expected returns increase while downside risk measures improve. At the same time, volatility remains within a realistic range, indicating that the framework shifts probability mass toward more favorable market environments without generating an excessively concentrated posterior distribution.

Global Assesment:

The analysis demonstrates that Regime-Based Entropy Pooling provides an economically intuitive mechanism for incorporating forward-looking market beliefs. By expressing views directly through market regimes rather than through expected returns, the framework aligns more closely with the way investors typically formulate their expectations. The resulting posterior distribution remains diversified statistically coherent, and fully suitable for different scenario analysis.

4 Conclusion and Practical Applications

This paper revisits Meucci's Entropy Pooling framework from a regime-based perspective and could be particularly attractive because it translates qualitative economic expectations into complete posterior return distributions while preserving the historical scenario structure. Since investment professionals typically reason in terms of market environments rather than precise numerical return forecasts, regime-based views provide a more intuitive and operational interface between economic expectations and quantitative portfolio construction.

Field of application: Tactical Asset Allocation

Tactical allocation decisions are generally driven by macroeconomic expectations, business cycle assessments, valuation signals, or market sentiment. Investment committees rarely formulate their outlook in terms of expected monthly asset returns. Instead, they discuss questions such as:

- Is the probability of a recession increasing?
- Are financial markets entering a recovery phase?
- Is a bullish expansion becoming more likely?
- How likely is another market crash?

Regime-Based Entropy Pooling allows these qualitative assessments to be translated directly into posterior regime probabilities. The resulting posterior return distribution automatically reflects the new market outlook and immediately provides updated estimates of expected returns, volatilities, downside risk measures, and portfolio characteristics.

Field of application: Portfolio Management

Portfolio managers continuously adjust portfolios according to changing market conditions. Traditionally, these adjustments require explicit assumptions regarding future expected returns or covariance matrices, which are often difficult to estimate reliably.

The proposed framework offers an alternative approach. Portfolio managers only need to express their confidence regarding future market regimes. The Entropy Pooling procedure subsequently transforms these regime views into a complete posterior probability distribution over historical scenarios.

This posterior distribution naturally updates all relevant portfolio statistics, including:

- Expected returns
- Volatility
- Correlations
- Value-at-Risk (VaR)
- Conditional Value-at-Risk (CVaR)

without requiring direct forecasts of these variables

.

Field of application: Risk Management

Risk managers frequently assess changes in downside risk by discussing the likelihood of adverse market environments rather than specifying numerical return expectations. For example, increasing the probability of a Crash regime immediately reallocates probability mass toward historically adverse observations. As a consequence, downside risk measures such as Value-at-Risk and Expected Shortfall are updated automatically while remaining fully consistent with the empirical scenario distribution. This provides an intuitive mechanism for incorporating forward-looking stress assessments into scenario-based risk analysis.

Regime Entropy Pooling and inference of the posterior return distribution

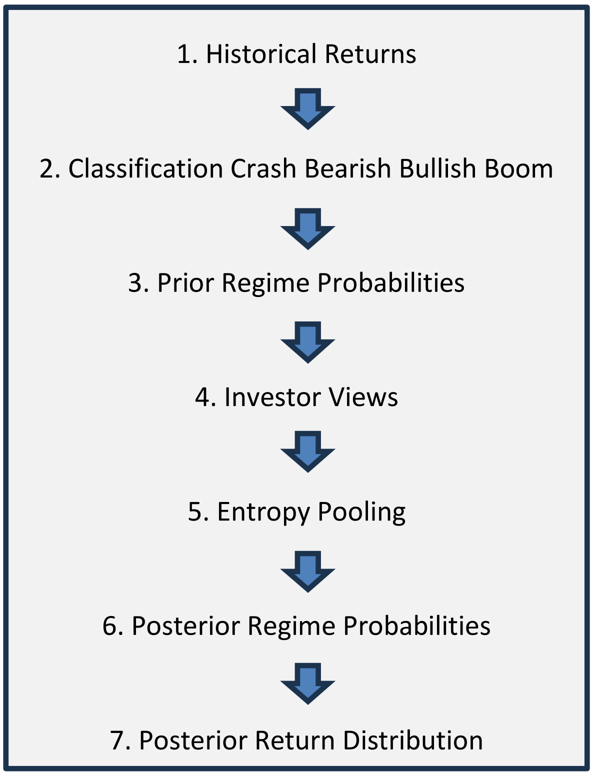

Main References

. Fully Flexible Views: Theory and Practice (2008) by Attilio Meucci .

. Mixing Probabilites, Priors and Kernels via Entropy Pooling by Attilio Meucci GARP Risk Professional, pp. 32-36, December 2011