

Dynamic Entropy Pooling: Portfolio Management with Views at Multiple Horizons

Attilio Meucci and Marco Nicolosi

Learn it: <http://www.symmys.com/node/831> (paper and code)

Live it: Advanced Risk and Portfolio Management[®] Bootcamp
www.symmys.com/arpm-bootcamp

Agenda

Background

The profit-and-loss (P&L)

The market model

Portfolio construction

Case studies

Background

| | Discretionary | Multiperiod | Mkt impact |
|--------------------------------|---------------|-------------|------------|
| Grinold ('89) | × | × | × |
| Black-Litterman ('90) | ✓ | × | × |
| Entropy Pooling ('08) | ✓ | × | × |
| Davis-Lleo ('13) | × | ✓ | × |
| Garleanu-Pedersen ('13) | × | ✓ | ✓ |
| Dynamic Entropy Pooling | ✓ | ✓ | ✓ |

- The standard approach to discretionary portfolio management (Black-Litterman, Entropy Pooling) processes subjective views that refer to the distribution of the market at a specific single investment horizon.
- The standard approach to multi-period portfolio management with market impact (Garleanu-Pedersen) processes non-discretionary (systematic) signals
- Dynamic Entropy Pooling is a quantitative approach to perform dynamic portfolio management with discretionary, multi-horizon views

Agenda

Background

The profit-and-loss (P&L)

The market model

Portfolio construction

Case studies

The profit-and-loss (P&L)

- We assume the single-period P&L is a set of exposures multiplied by the increments of the risk drivers over the rebalancing period:

$$\Pi_{t+1} = \mathbf{b}'_t \Delta \mathbf{X}_{t+1}$$

- The set of risk drivers can be extended to include also external factors that do not affect directly the P&L of the instruments. On such additional factors we can express views that influence the P&L through correlation. The corresponding entries in the exposures vector will be set to zero.

- We assume the single-period P&L is a set of exposures multiplied by the increments of the risk drivers over the rebalancing period:

$$\Pi_{t+1} = \mathbf{b}'_t \Delta \mathbf{X}_{t+1}$$

- The set of risk drivers can be extended to include also external factors that do not affect directly the P&L of the instruments. On such additional factors we can express views that influence the P&L through correlation. The corresponding entries in the exposures vector will be set to zero.

- Consider an equity share or an index. Then the risk driver is its log-value:

$$X_t = \ln V_t$$

- The P&L of a portfolio with $h_{n,t}$ shares in the n -th asset is:

$$\Pi_{t+1} = \sum_n \underbrace{h_{n,t} V_{n,t}}_{b_{n,t}} \times \underbrace{\left(\frac{V_{n,t+1}}{V_{n,t}} - 1 \right)}_{\Delta X_{n,t}} \approx \sum_n b_{n,t} \Delta X_{n,t+1}$$

- More in general, in terms of a style/risk linear factor model:

$$\Pi_{t+1} = \sum_k b_{k,t}^{style} \Delta X_{k,t+1}^{style}$$

- We assume the single-period P&L is a set of exposures multiplied by the increments of the risk drivers over the rebalancing period:

$$\Pi_{t+1} = \mathbf{b}'_t \Delta \mathbf{X}_{t+1}$$

- The set of risk drivers can be extended to include also external factors that do not affect directly the P&L of the instruments. On such additional factors we can express views that influence the P&L through correlation. The corresponding entries in the exposures vector will be set to zero.

- Suppose that the n -th asset is a fixed income instrument. Its value at the first order satisfies

$$\Pi_{n,t+1} \approx -\sum_k dv01_{n,k,t} \Delta Y_{k,t+1}$$

where $Y_{k,t}$ is the k -th key-rate on the yield curve; $dv01_{n,k,t}$ is the dollar-sensitivity of the n -th instrument to $Y_{k,t}$

- Then the P&L due to a set of fixed income instruments is:

$$\Pi_{t+1} \approx \sum_k \underbrace{\left(-\sum_n h_{n,t} dv01_{n,k,t} \right)}_{b_{k,t}} \Delta X_{k,t+1}$$

- We assume the single-period P&L is a set of exposures multiplied by the increments of the risk drivers over the rebalancing period:

$$\Pi_{t+1} = \mathbf{b}'_t \Delta \mathbf{X}_{t+1}$$

- The set of risk drivers can be extended to include also external factors that do not affect directly the P&L of the instruments. On such additional factors we can express views that influence the P&L through correlation. The corresponding entries in the exposures vector will be set to zero.

- For a stock option, the risk drivers are the log-value of the underlying and the implied volatility $X_t = \ln V_t$ and Σ_t^{impl}

- Then for a portfolio of stock options, the P&L is:

$$\Pi_{t+1} \approx \sum_n \underbrace{h_{n,t} \delta_{n,t} V_{n,t}}_{b_{n,t}^\delta} \Delta X_{n,t+1} + \sum_n \underbrace{h_{n,t} v_{n,t}}_{b_{n,t}^\sigma} \Delta \Sigma_{n,t+1}^{impl}$$

where $\delta_{n,t}$ and $v_{n,t}$ are the delta and vega of the n -th option.

Agenda

Background

The profit-and-loss (P&L)

The market model

Portfolio construction

Case studies

- Consider a book of assets driven by a set of \bar{n} risk drivers \mathbf{X}_t (interest rates, implied volatility surfaces, log-prices, etc.)
- We assume that the drivers follow a MVOU process:

$$d\mathbf{X}_t = (-\boldsymbol{\theta}\mathbf{X}_t + \boldsymbol{\mu}) dt + \boldsymbol{\sigma}d\mathbf{W}_t$$

- Choose a set of discrete monitoring dates $t, t+1, \dots, \bar{t}$
- Stack the process at the monitoring times as follows:

$$\mathbf{X}_{t \rightsquigarrow \bar{t}} \equiv \begin{pmatrix} \mathbf{X}_t \\ \mathbf{X}_{t+1} \\ \vdots \\ \mathbf{X}_{\bar{t}} \end{pmatrix}$$

- Then the process is jointly multivariate normal at all times

$$\mathbf{X}_{t \rightsquigarrow \bar{t}} | \mathbf{i}_t \sim N(\boldsymbol{\mu}_{t \rightsquigarrow \bar{t}}, \boldsymbol{\sigma}_{t \rightsquigarrow \bar{t}}^2)$$

- The expectation vector of the is

$$\boldsymbol{\mu}_{t \rightsquigarrow \bar{t}} \equiv \begin{pmatrix} e^{-0\boldsymbol{\theta}} \mathbf{x}_t + (\mathbb{I}_{\bar{n}} - e^{-0\boldsymbol{\theta}}) \boldsymbol{\theta}^{-1} \boldsymbol{\mu} \\ e^{-1\boldsymbol{\theta}} \mathbf{x}_t + (\mathbb{I}_{\bar{n}} - e^{-1\boldsymbol{\theta}}) \boldsymbol{\theta}^{-1} \boldsymbol{\mu} \\ \cdot \\ e^{-(\bar{t}-t)\boldsymbol{\theta}} \mathbf{x}_t + (\mathbb{I}_{\bar{n}} - e^{-(\bar{t}-t)\boldsymbol{\theta}}) \boldsymbol{\theta}^{-1} \boldsymbol{\mu} \end{pmatrix}$$

- The covariance matrix is

$$\boldsymbol{\sigma}_{t \rightsquigarrow \bar{t}}^2 \equiv \begin{pmatrix} \boldsymbol{\sigma}_0^2 & \boldsymbol{\sigma}_0^2 e^{-\boldsymbol{\theta}'} & \boldsymbol{\sigma}_0^2 e^{-2\boldsymbol{\theta}'} & \cdot & \boldsymbol{\sigma}_0^2 e^{-(\bar{t}-t)\boldsymbol{\theta}'} \\ e^{-\boldsymbol{\theta}} \boldsymbol{\sigma}_0^2 & \boldsymbol{\sigma}_1^2 & \boldsymbol{\sigma}_1^2 e^{-\boldsymbol{\theta}'} & \cdot & \boldsymbol{\sigma}_1^2 e^{-(\bar{t}-t-1)\boldsymbol{\theta}'} \\ e^{-2\boldsymbol{\theta}} \boldsymbol{\sigma}_0^2 & e^{-\boldsymbol{\theta}} \boldsymbol{\sigma}_1^2 & \boldsymbol{\sigma}_2^2 & \cdot & \\ \cdot & \cdot & \cdot & \cdot & \\ e^{-(\bar{t}-t)\boldsymbol{\theta}} \boldsymbol{\sigma}_0^2 & & & & \boldsymbol{\sigma}_{\bar{t}-t}^2 \end{pmatrix}$$

where

$$vec(\underline{\boldsymbol{\sigma}}_{\tau}^2) \equiv (\boldsymbol{\theta} \oplus \boldsymbol{\theta})^{-1} (\mathbb{I}_{\bar{n}^2} - e^{-(\boldsymbol{\theta} \oplus \boldsymbol{\theta})\tau}) vec(\boldsymbol{\sigma}^2)$$

We extend the Entropy Pooling approach in Meucci (2010) to the case of multiple horizons

- **The prior:** assume a model for the joint distribution of the process at the monitoring times:

$$\mathbf{X}_{t \rightsquigarrow \bar{t}} | \mathbf{i}_t \sim f$$

- **The views:** are statements (constraints) on the yet-to-be defined distribution of the process:

$$g \in \mathcal{V}_t$$

- **The posterior:** is the closest distribution to the prior that satisfies the views:

$$\bar{f} \equiv \operatorname{argmin}_{g \in \mathcal{V}_t} \{ \mathcal{E}(g, f) \}$$

where the “distance” is the relative entropy

$$\mathcal{E}(g, f) \equiv \int g(\mathbf{x}_t, \dots, \mathbf{x}_{\bar{t}}) \ln \frac{g(\mathbf{x}_t, \dots, \mathbf{x}_{\bar{t}})}{f(\mathbf{x}_t, \dots, \mathbf{x}_{\bar{t}})} d\mathbf{x}_t \cdots \mathbf{x}_{\bar{t}},$$

We extend the Entropy Pooling approach in Meucci (2010) to the case of multiple horizons

- **The prior:** assume a MVOU model for the joint distribution of the process at the monitoring times

$$\mathbf{X}_{t \rightsquigarrow \bar{t}} | \mathbf{i}_t \sim N(\boldsymbol{\mu}_{t \rightsquigarrow \bar{t}}, \boldsymbol{\sigma}_{t \rightsquigarrow \bar{t}}^2)$$

- **The views:** are statements (constraints) on the yet-to-be defined distribution of the process :

$$\mathcal{V}_t : \begin{cases} \mathbb{E}_t^g \{ \mathbf{v}_{\mu,t} \mathbf{X}_{t \rightsquigarrow \bar{t}} \} \equiv \boldsymbol{\mu}_{view;t} \\ \mathbb{C}v_t^g \{ \mathbf{v}_{\sigma,t} \mathbf{X}_{t \rightsquigarrow \bar{t}} \} \equiv \boldsymbol{\sigma}_{view;t}^2 \end{cases}$$

where $\mathbf{v}_{\mu,t}$ and $\mathbf{v}_{\sigma,t}$ are matrices that defines arbitrary linear combinations of the process at the times for the views.

- **The posterior:** is the closest distribution to the prior that satisfies the views:

$$\bar{f} \equiv \operatorname{argmin}_{g \in \mathcal{V}_t} \{ \mathcal{E}(g, f) \} \Rightarrow \mathbf{X}_{t \rightsquigarrow \bar{t}} | \mathbf{i}_t \sim N(\bar{\boldsymbol{\mu}}_{t \rightsquigarrow \bar{t}}, \bar{\boldsymbol{\sigma}}_{t \rightsquigarrow \bar{t}}^2)$$

$$\mathbf{X}_{t \rightsquigarrow \bar{t}} | \mathbf{i}_t \sim N(\bar{\boldsymbol{\mu}}_{t \rightsquigarrow \bar{t}}, \bar{\boldsymbol{\sigma}}_{t \rightsquigarrow \bar{t}}^2)$$

- For the expectation, we introduce the pseudo inverse matrix of $\mathbf{v}_{\mu,t}$

$$\mathbf{v}_{\mu,t}^+ \equiv \boldsymbol{\sigma}_{t \rightsquigarrow \bar{t}}^2 \mathbf{v}'_{\mu,t} (\mathbf{v}_{\mu,t} \boldsymbol{\sigma}_{t \rightsquigarrow \bar{t}}^2 \mathbf{v}'_{\mu,t})^{-1}$$

we define the two complementary projectors:

$$\mathbb{P}_{\mu,t} \equiv (\mathbb{I}_{\bar{n}(\bar{t}-t+1)} - \mathbf{v}_{\mu,t}^+ \mathbf{v}_{\mu,t}) \quad \mathbb{P}_{\mu,t}^\perp \equiv \mathbf{v}_{\mu,t}^+ \mathbf{v}_{\mu,t}$$

Then

$$\bar{\boldsymbol{\mu}}_{t \rightsquigarrow \bar{t}} \equiv \mathbb{P}_{\mu,t} \boldsymbol{\mu}_{t \rightsquigarrow \bar{t}} + \mathbb{P}_{\mu,t}^\perp (\mathbf{v}_{\mu,t}^+ \boldsymbol{\mu}_{view;t})$$

- Similar, for the covariance we introduce the pseudo inverse of $\mathbf{v}_{\sigma,t}$

$$\mathbf{v}_{\sigma,t}^+ \equiv \boldsymbol{\sigma}_{t \rightsquigarrow \bar{t}}^2 \mathbf{v}'_{\sigma,t} (\mathbf{v}_{\sigma,t} \boldsymbol{\sigma}_{t \rightsquigarrow \bar{t}}^2 \mathbf{v}'_{\sigma,t})^{-1}$$

and the two complementary projectors:

$$\mathbb{P}_{\sigma,t} \equiv \mathbb{I}_{\bar{n}(\bar{t}-t+1)} - \mathbf{v}_{\sigma,t}^+ \mathbf{v}_{\sigma,t} \quad \mathbb{P}_{\sigma,t}^\perp \equiv \mathbf{v}_{\sigma,t}^+ \mathbf{v}_{\sigma,t}$$

Then

$$\bar{\boldsymbol{\sigma}}_{t \rightsquigarrow \bar{t}}^2 \equiv \mathbb{P}_{\sigma,t} \boldsymbol{\sigma}_{t \rightsquigarrow \bar{t}}^2 \mathbb{P}'_{\sigma,t} + \mathbb{P}_{\sigma,t}^\perp (\mathbf{v}_{\sigma,t}^+ \boldsymbol{\sigma}_{view;t} (\mathbf{v}_{\sigma,t}^+)') (\mathbb{P}_{\sigma,t}^\perp)'$$

Agenda

Background

The profit-and-loss (P&L)

The market model

Portfolio construction

Case studies

- As in Garleanu and Pedersen (2013), the satisfaction functional is an infinite sum of discounted trade-offs:

$$\bar{S}_t^{(\gamma, \eta)} \equiv \sum_{s=t}^{\infty} e^{-\lambda(s-t)} [\bar{\mathbb{E}}\{\Pi_{(s,s+1)} | \mathbf{i}_t\} - \frac{\gamma}{2} \bar{\mathbb{V}}\{\Pi_{(s,s+1)} | \mathbf{i}_t\} - \frac{\eta}{2} \bar{\mathbb{E}}\{MI_s | \mathbf{i}_t\}]$$

where the market impact is a quadratic function of the exposure rebalancing

$$MI_t = a^2 + \Delta \mathbf{b}'_t \mathbf{c}^2 \Delta \mathbf{b}_t$$

with \mathbf{c}^2 a suitable positive definite matrix. Note the term a^2 , which represents the average cost of maintaining constant exposures

- Given that the P&L is linear in the exposures $\Pi_{t+1} = \mathbf{b}'_t \Delta \mathbf{X}_{t+1}$, we need to solve for the optimal policy of exposures as functions of information

$$\{\mathbf{b}_s^* = p_s^*(\mathbf{i}_s)\}_{s \geq t}$$

where

$$\begin{aligned} \{p_s^*\}_{s \geq t} = \operatorname{argmax}_{\{p_s\}_{s \geq t} \in \mathcal{C}} & \bar{\mathbb{E}}_t \left\{ \sum_{s=t}^{\infty} e^{-\lambda(s-t)} [p_s(\mathbf{I}_s)' \boldsymbol{\omega} \bar{\mathbb{E}}_s \{\Delta \mathbf{X}_{s+1}\} \right. \\ & \left. - \frac{\gamma}{2} p_s(\mathbf{I}_s)' \boldsymbol{\omega} \bar{\mathbb{C}} v_s \{\Delta \mathbf{X}_{s+1}\} \boldsymbol{\omega}' p_s(\mathbf{I}_s) - \frac{\eta}{2} \Delta p_s(\mathbf{I}_s)' \mathbf{c}^2 \Delta p_s(\mathbf{I}_s) \right\} \end{aligned}$$

- As in Garleanu and Pedersen (2013), the satisfaction functional is an infinite sum of discounted trade-offs:

$$\bar{\mathbb{S}}_t^{(\gamma, \eta)} \equiv \sum_{s=t}^{\infty} e^{-\lambda(s-t)} [\bar{\mathbb{E}}\{\Pi_{(s,s+1)} | \mathbf{i}_t\} - \frac{\gamma}{2} \bar{\mathbb{V}}\{\Pi_{(s,s+1)} | \mathbf{i}_t\} - \frac{\eta}{2} \bar{\mathbb{E}}\{MI_s | \mathbf{i}_t\}]$$

where the market impact is a quadratic function of the exposure rebalancing

$$MI_t = a^2 + \Delta \mathbf{b}'_t \mathbf{c}^2 \Delta \mathbf{b}_t$$

with \mathbf{c}^2 a suitable positive definite matrix. Note the term a^2 , which represents the average cost of maintaining constant exposures

- Given that the P&L is linear in the exposures $\Pi_{t+1} = \mathbf{b}'_t \Delta \mathbf{X}_{t+1}$, we need to solve for the optimal policy of exposures as functions of information

$$\{\mathbf{b}_s^* = p_s^*(\mathbf{i}_s)\}_{s \geq t}$$

where

$$\begin{aligned} \{p_s^*\}_{s \geq t} = \operatorname{argmax}_{\{p_s\}_{s \geq t} \in \mathcal{C}} & \mathbb{E}_t \left\{ \sum_{s=t}^{\infty} e^{-\lambda(s-t)} [p_s(\mathbf{I}_s)' \boldsymbol{\omega} \mathbb{E}_s \{\Delta \mathbf{X}_{s+1}\} \right. \\ & \left. - \frac{\gamma}{2} p_s(\mathbf{I}_s)' \boldsymbol{\omega} \bar{\mathbf{C}} v_s \{\Delta \mathbf{X}_{s+1}\} \boldsymbol{\omega}' p_s(\mathbf{I}_s) - \frac{\eta}{2} \Delta p_s(\mathbf{I}_s)' \mathbf{c}^2 \Delta p_s(\mathbf{I}_s) \right\} \end{aligned}$$

- Dynamic programming with a quadratic value function yields a recursive problem with time-dependent coefficients

$$\begin{aligned} v_{s+1}(\mathbf{b}_s, \mathbf{x}_{s+1}) &= -\frac{1}{2} \mathbf{b}'_s \psi_{bb,s} \mathbf{b}_s + \mathbf{b}'_s \psi_{bx,s} \mathbf{x}_{s+1} + \frac{1}{2} \mathbf{x}'_{s+1} \psi_{xx,s} \mathbf{x}_{s+1} + \psi'_{b,s} \mathbf{b}_s + \psi'_{x,s} \mathbf{x}_{s+1} + \psi_{0,s} \\ &\iff \psi_{s-1} = g_s(\psi_s) \end{aligned}$$

- The optimal policy of exposures then reads

$$\begin{aligned} \mathbf{b}_s^* &= (\gamma \boldsymbol{\omega} \bar{\boldsymbol{\sigma}}_s^2 \boldsymbol{\omega}' + \eta \mathbf{c}^2 + e^{-\lambda} \psi_{bb,s})^{-1} \left[\underbrace{\eta \mathbf{c}^2 \mathbf{b}_{s-1}}_{\text{legacy exposures}} \right. \\ &\quad \left. + \underbrace{(\boldsymbol{\omega} \boldsymbol{\beta}_s + e^{-\lambda} \psi_{bx,s} (\boldsymbol{\beta}_s + \mathbb{I}_{\bar{n}})) \mathbf{x}_s}_{\text{current risk drivers}} + \underbrace{(\boldsymbol{\omega} + e^{-\lambda} \psi_{bx,s}) \boldsymbol{\alpha}_s + e^{-\lambda} \psi_{b,s}}_{(\star) \text{ future views}} \right] \end{aligned}$$

- Given that the P&L is linear in the exposures $\Pi_{t+1} = \mathbf{b}'_t \Delta \mathbf{X}_{t+1}$, we need to solve for the optimal policy of exposures as functions of information

$$\{\mathbf{b}_s^* = p_s^*(\mathbf{i}_s)\}_{s \geq t}$$

where

$$\begin{aligned} \{p_s^*\}_{s \geq t} = \operatorname{argmax}_{\{p_s\}_{s \geq t} \in \mathcal{C}} \mathbb{E}_t \{ & \sum_{s=t}^{\infty} e^{-\lambda(s-t)} [p_s(\mathbf{I}_s)' \boldsymbol{\omega} \mathbb{E}_s \{\Delta \mathbf{X}_{s+1}\} \\ & - \frac{\gamma}{2} p_s(\mathbf{I}_s)' \boldsymbol{\omega} \overline{\mathbf{C}} v_s \{\Delta \mathbf{X}_{s+1}\} \boldsymbol{\omega}' p_s(\mathbf{I}_s) - \frac{\eta}{2} \Delta p_s(\mathbf{I}_s)' \mathbf{c}^2 \Delta p_s(\mathbf{I}_s)] \} \end{aligned}$$

- With no market impact, we obtain a series of myopic one-period problems
- The optimal policy is a sequence of mean-variance optimizations based on the posterior distribution of the risk drivers process

$$\begin{aligned} \mathbf{b}_s^* = & \underbrace{\frac{1}{\gamma} (\boldsymbol{\omega} \bar{\boldsymbol{\sigma}}_s^2 \boldsymbol{\omega}')^{-1} \boldsymbol{\omega} (\mathbb{P}_{\mu,s})_{s+1, \cdot} \Delta \boldsymbol{\mu}_{s \rightsquigarrow \bar{t}}^{LongTerm}}_{\mathbf{b}_s^{LongTerm}} \longleftarrow \begin{pmatrix} \mathbb{I}_{\bar{n}} - e^{-0\boldsymbol{\theta}} \\ \mathbb{I}_{\bar{n}} - e^{-1\boldsymbol{\theta}} \\ \vdots \\ \mathbb{I}_{\bar{n}} - e^{-(\bar{t}-s)\boldsymbol{\theta}} \end{pmatrix} (\boldsymbol{\theta}^{-1} \boldsymbol{\mu} - \mathbf{x}_s) \\ & + \underbrace{\frac{1}{\gamma} (\boldsymbol{\omega} \bar{\boldsymbol{\sigma}}_s^2 \boldsymbol{\omega}')^{-1} \boldsymbol{\omega} (\mathbb{P}_{\mu,s}^{\perp})_{s+1, \cdot} \Delta \boldsymbol{\mu}_{s \rightsquigarrow \bar{t}}^{ViewMean}}_{\mathbf{b}_s^{ViewMean}} \longleftarrow \mathbf{v}_{\mu,s}^+ \boldsymbol{\mu}_{view;s} - \mathbf{x}_s \end{aligned}$$

Agenda

Background

The profit-and-loss (P&L)

The market model

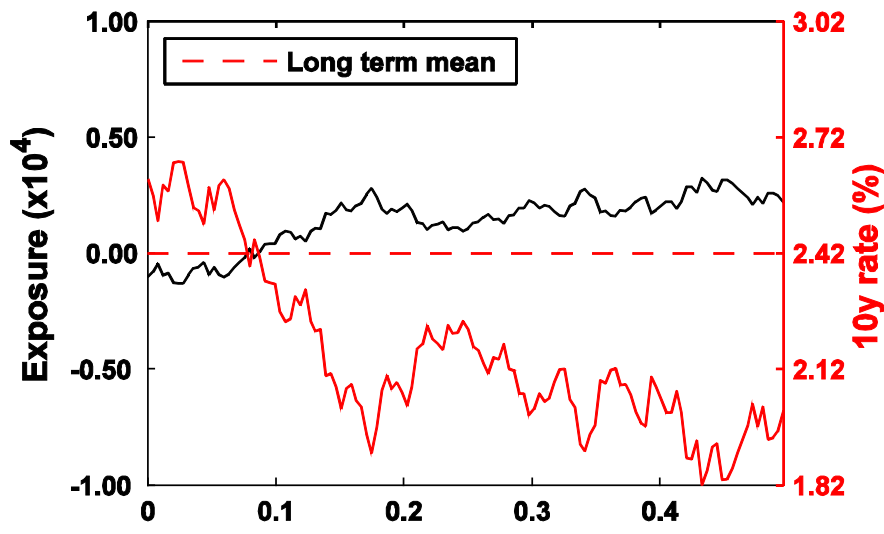
Portfolio construction

Case studies


Case studies

One risk driver, one view


b_t^{Prior}



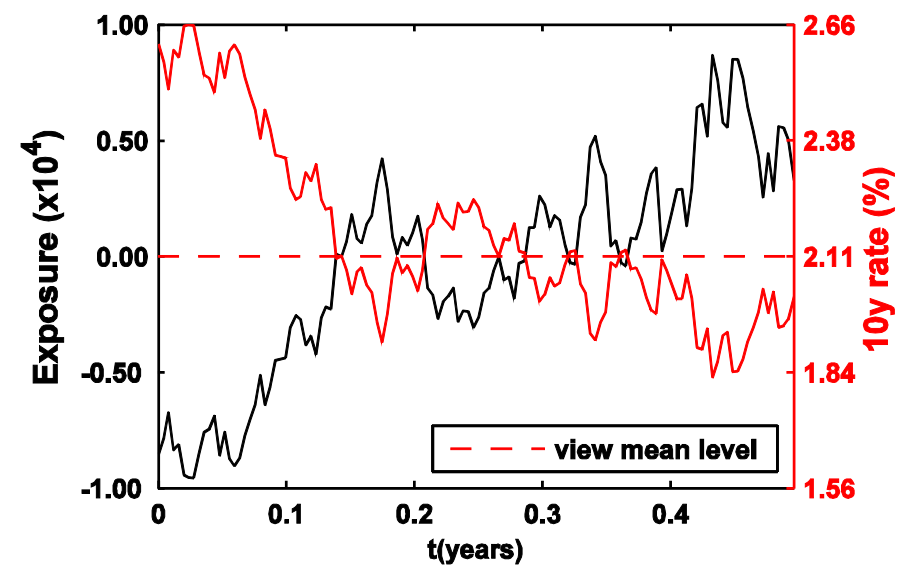
b_t^*



$b_t^{LongTerm}$



$b_t^{viewMean}$



Case studies

Two risk drivers (one investable), two views

