Risk-Sensitive Investment Management: A Guide for Quants

Sébastien Lleo, CQF

Explaina-what-again?

- Harry Markowitz's pioneering work on portfolio selection (Markowitz , 1952) opened the door to the development of an abundance of models.
- Fund managers now have at their disposal models in all shapes and sizes:
  - simple **static models** such as the mean-variance criterion;
  - supremely flexible **stochastic programming models** that embrace the dynamics of financial markets; and
  - **algorithmic models** showcasing the latest developments in **machine learning.**
- The more, the merrier? Right?
- Well… probably not.
  - The trouble starts when you have to pick a model (or an ensemble of models).
  - Which one should you choose? And why?
  - The irritation keeps growing when you are asked to explain why your prized black-box model produced a particular asset allocation.
  - **Explainability just went missing in action…**

Photo courtesy of Pixabay

- This short tutorial demonstrates how to bring the **explainability of dynamic investment models** back to life using **risk-sensitive investment management (RSIM)**.

  - RSIM applies **risk-sensitive control**, a branch of stochastic control, to solve portfolio selection problems dynamically.

  - It enables investors to **optimize their portfolios, modeled as a dynamical system, subject to random market noise**.

  - The control variable, i.e., the decision variable, for this optimization problem is the proportion of wealth the investors allocate to each security.

  - The objective function connects the dynamical system and the control variable to the investors' broader goals.

Risk-Sensitive Control vs. Standard Stochastic Control

(Standard) Stochastic Control

- In stochastic control, the standard formulation for the objective function is as an expected reward of the form $\mathbb{E}\left[r\right]$, where
  - $\mathbb{E}\left[\,\cdot\,\right]$ denotes the expectation and
  - r is a stochastic reward.
- However, this formulation does not account for the investors' risk preference.
- Robert Merton (1969) then defined the reward as the utility U of the investors' wealth w, that is, $\mathbb{E}\left[U(w)\right]$.
  - This approach is known as the `Merton model.'
  - The **upside** is obvious: inserting a utility function in the objective function ensures consistency with economic models of risk preferences.
  - The **downside** is that we now have a nonlinear function between the evolution of our dynamical system, the investors' wealth, and the expectation we seek to maximize.

Risk-Sensitive Control

- Risk-sensitive control proposes a more efficient formulation for the objective function:

$$J := -\frac{1}{\theta} \ln \mathbb{E}\left[e^{-\theta r}\right]$$

where $\theta \in (-1,0) \cup (0,\infty)$ parametrizes the investors' aversion toward risk.

  - Thus, no extra utility function is needed.

- Another important property for finance is that the risk-sensitive criterion naturally transposes **mean-variance optimisation** to a dynamical setting,

  - A simple **Taylor expansion** around $\theta = 0$:

$$J \approx \mathbb{E}[r] - \frac{\theta}{2}\text{Var}[r].$$

Three Use Cases

1. RSIM in the Black-Scholes-Merton World;
2. RSIM with Factors and Benchmark
3. RSIM with Unobservable Factors and Expert Opinions

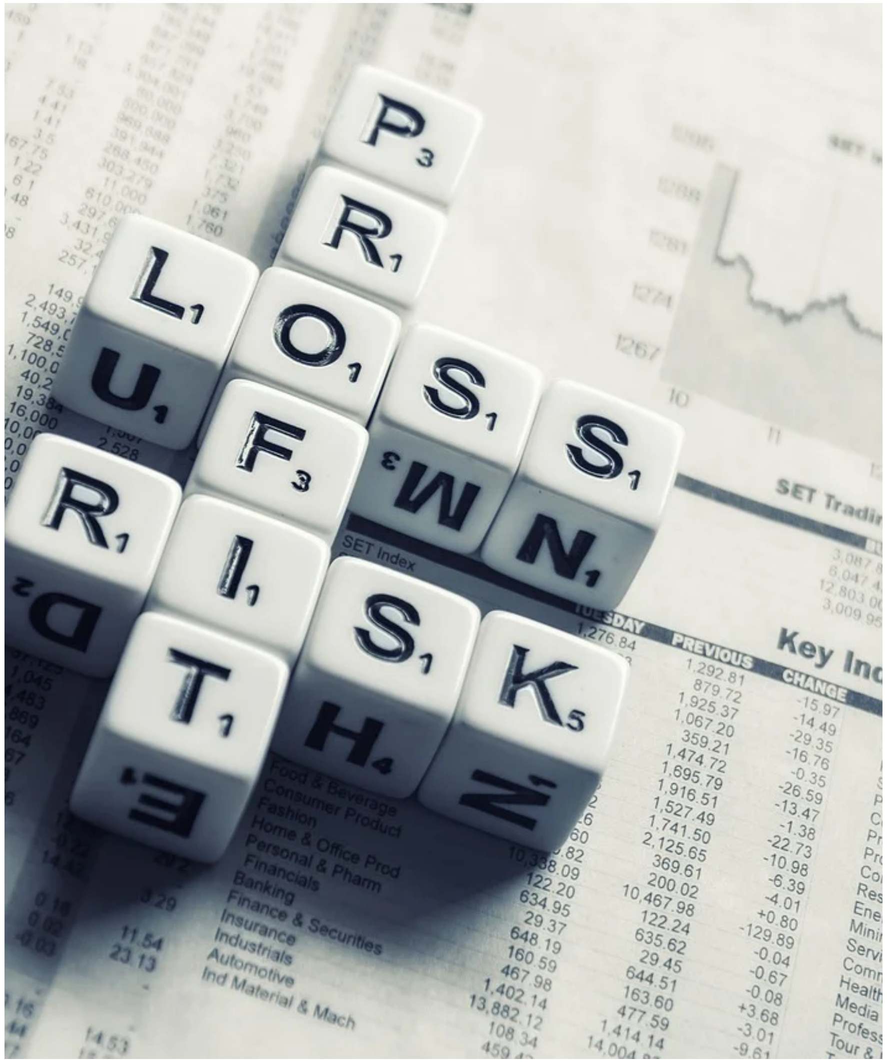

Photo courtesy of Pixabay

Case 1: We begin with a simple model set in a Black-Scholes-Merton world

- Consider an investor:
  - looking to construct a portfolio to fund her retirement in T years.
  - initial wealth \$$v$, and her
  - degree of risk sensitivity is $\theta > 0$.
  - can invest her wealth in
    - a stock index fund S and
    - a risk-free money market instrument B.

The Black-Scholes-Merton World

- Brownian motion $W(t)$;
- Stock $S(t)$ with price dynamics:

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW_t, \quad S(0) = s;$$

- Risk-free instrument $B(t)$ with price dynamics

$$\frac{dB_t}{B_t} = r dt, \quad B(0) = 1.$$

Case 1: RSIM in the Black-Scholes-Merton World

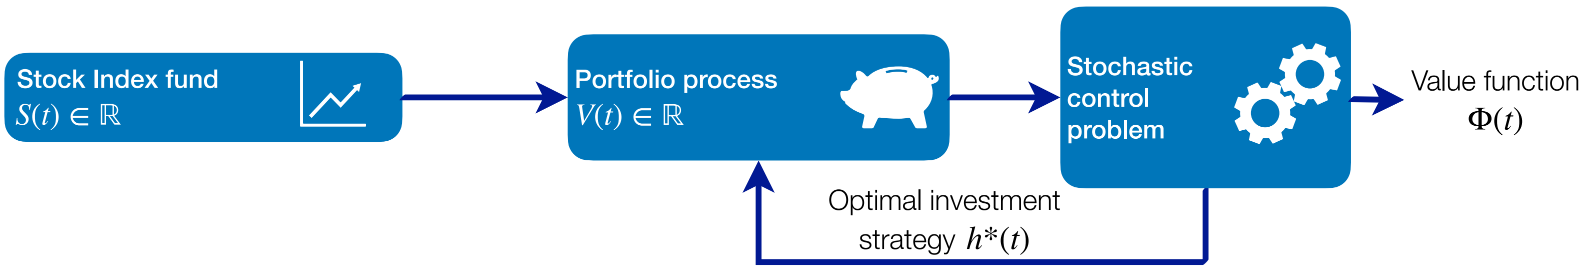

Mathematically,

Discounted asset prices $\tilde{S}(t) := S(t)/B(t)$ are modelled as geometric Brownian motions

$$\frac{d\tilde{S}_t}{\tilde{S}_t} = (\mu - r)dt + \sigma dW_t, \ \tilde{S}(0) = s$$

W_t is a $\mathscr{F}_t$-Brownian motion on a filtered complete probability space $\left(\Omega, \mathscr{F}, (\mathscr{F})_{t=0}^{T}, \mathbb{P}\right)$

The wealth process is a controlled geometric process

$$\frac{dV(t)}{V(t)} = h(t)(\mu - r)dt + h(t)\Sigma dW(t)$$
$$V(0) = v$$

Risk-sensitive criterion

$$J(t, x, h; \theta) := -\frac{1}{\theta} \ln \mathbb{E}\left[e^{-\theta R(t)}\right]$$

Control Variable

- The control variable $h(t)$ represents the **proportion of the investor's total wealth in the stock index fund** at time $0 \leq t \leq T$.
  - When $h(t) > 0$, the investor is long the stock;
  - $h(t) < 0$ is short the stock;
  - $h(t) > 1$ implies leverage, funded by shorting the money market instrument.
- Technically, the control variable $h(t)$ is a $\mathscr{F}_t$-adapted and progressively measurable stochastic process.
- We say that $h(t)$ is **admissible**, or in class $\mathscr{H}$, if it also satisfies the technical condition $P\left(\int_0^T |h_t|^2 dt < +\infty\right) = 1.$
  - **Main objective of this condition**: to prevent unbounded leverage.

Understanding the objective function

- We define the **stochastic reward $r(T)$ as the logarithmic excess return $\ln \frac{V(T)}{v}$** that the portfolio earns on top of the risk-free rate over the investment horizon.

- The **risk-sensitive objective function** is:

$$J(h) := -\frac{1}{\theta} \ln \mathbb{E}\left[e^{-\theta \ln \frac{V(T)}{v}}\right] = \ln v - \frac{1}{\theta} \ln \mathbb{E}\left[e^{-\theta \ln V(T)}\right],$$

for $\theta \in (-1,0) \cup (0,\infty)$.

  - The initial wealth $\$v$ is simply an additive constant, so it does not affect the control problem.

  - For simplicity and without loss of generality, we take $v = \$1$.

- **Objective function is intuitive**: the investor seeks to maximize the log excess return of their portfolio over the risk-free rate consistently with their risk aversion.

Building Connections

- This objective function connects neatly to **utility theory**.
  - The term $e^{-\theta \ln V(T)} = V^{-\theta}(t)$ inside the expectation acts as a **power utility function**.
  - The $-\frac{1}{\theta} \ln$ outside the expectation normalizes the criterion to the same unit as the reward.
- The Taylor expansion $J(h,t) \approx \mathbb{E}\left[\ln V(t)\right] - \frac{\theta}{2}\text{Var}\left[\ln V(t)\right],$ is tantamount to a `**dynamic Markowitz**.'
- When we take the limit of the objective function as $\theta \to 0$, we recover the logarithmic utility, which also corresponds to the **Kelly criterion**:

$$K(h) := \lim_{\theta \to 0} J(h) = \lim_{\theta \to 0} -\frac{1}{\theta} \ln \mathbb{E}\left[e^{-\theta \ln V(T)}\right] = \mathbb{E}\left[\ln V(T)\right].$$

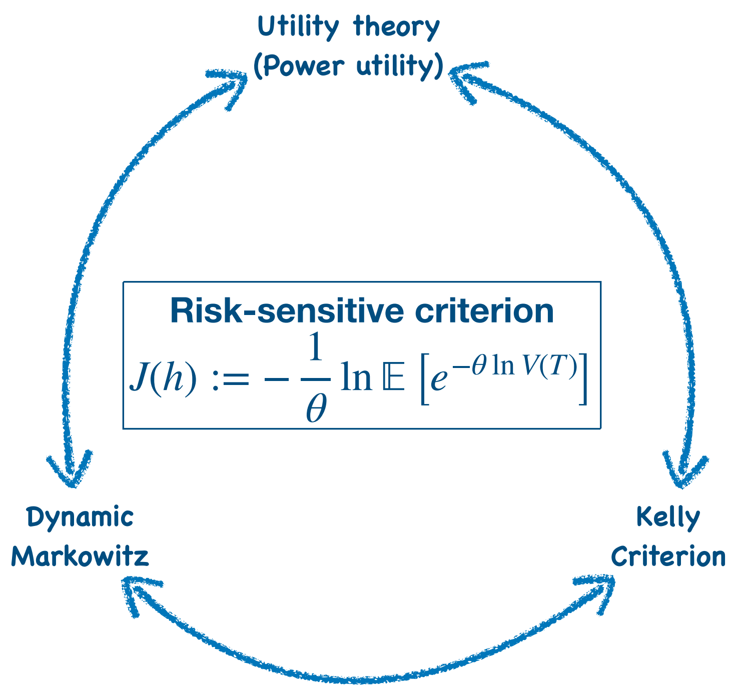

Solving the RSIM problem

- **The investor chooses $h(t)$ to maximize the objective function.**
- Define the value function $\Phi(t)$ as

$$\Phi(t) := \sup_{h \in \mathcal{H}} J(h) = \sup_{h \in \mathcal{H}} -\frac{1}{\theta} \ln \mathbb{E}\left[e^{-\theta \ln V(T)}\right],$$

where $\mathcal{H}$ is the class of admissible controls.

- The cleanest and most direct solution is to perform a **change of probability measure**.

- Focus on the term $e^{-\theta \ln V(T)}$ inside the criterion $J(h)$.
- **Apply Itô to get the following dynamics for $\ln V(t)$:**

$$d \ln V(t) = \left[\left(\mu - r \right) h(t) - \frac{1}{2} \sigma^2 h^2(t) \right] dt + V(t) h(t) \sigma dW(t)$$

- Thus,

$$e^{-\theta \ln V(T)} = \exp \left\{ -\theta \int_0^T \left[\left(\mu - r \right) h(t) - \frac{1}{2} \sigma^2 h^2(t) \right] dt - \theta \int_0^T h(t) \sigma dW(t) \right\}$$

is an exponential process with randomness driven by $\exp \left\{ -\theta \int_0^T h(t) \sigma dW(t) \right\}$.

- Before performing a change of measure, we need to `**complete the exponential martingale.**'

- To do so, multiply and divide $e^{-\theta \ln V(T)}$ by $\exp\left\{-\frac{1}{2}\theta^2\int_0^T h^2(t)\sigma^2 dt\right\}$ to obtain

$$e^{-\theta \ln V(T)} := \exp\left\{\theta\int_0^T g(h(t))dt\right\}\chi^h(T),$$

where

- $g(h) := \frac{1}{2}(\theta+1)\sigma^2h^2 - \left(\mu - r\right) h(t)$

- $\chi^h(t) := \exp\left\{-\frac{1}{2}\theta^2\int_0^t h^2(s)\sigma^2 ds - \theta\int_0^t h(s)\sigma dW(s)\right\}.$

- For now, assume that $h(t)$ is such that $\chi^h(t)$ is an exponential martingale.

- Then we can define a **new measure** $\mathbb{P}_h$ on $(\Omega, \mathcal{F}_T)$ via the **Radon-Nikodym derivative**

$$\frac{d\mathbb{P}_h}{d\mathbb{P}} := \chi_T^h$$

- Taking the expectation of $e^{-\theta \ln V(T)}$ and applying the change of measure, we obtain

$$\mathbb{E}\left[e^{-\theta \ln V(T)}\right] = \mathbb{E}\left[\exp\left\{-\theta \int_0^T g(h)dt\right\}\chi^h(T)\right] = \mathbb{E}^h\left[\exp\left\{-\theta \int_0^T g(h(t))dt\right\}\right],$$

where $\mathbb{E}^h[\ \cdot\ ]$ denotes the exponential under the measure $\mathbb{P}_h$.

- Therefore,

$$J^h(h) = -\frac{1}{\theta}\ln \mathbb{E}^h\left[v^{-\theta}e^{-\theta \ln V(T)}\right] = -\frac{1}{\theta}\ln \mathbb{E}^h\left[\exp\left\{-\theta\int_0^T g(h(t))dt\right\}\right]$$

$$\Phi(t) = \sup_{h\in\mathscr{H}} J^h(h) = \sup_{h\in\mathscr{H}} -\frac{1}{\theta}\ln \mathbb{E}^h\left[\exp\left\{-\theta\int_0^T g(h(t))dt\right\}\right].$$

- Finally, **it suffices to maximize the function g pointwise to achieve the supremum in this expression**.

- The function g is quadratic in h so it achieves a **global maximum**
  - at $g^* = \frac{1}{2}\frac{1}{\theta+1}\left(\frac{\mu-r}{\sigma}\right)^2$,
  - for $h^* = \frac{1}{\theta+1}\cdot\frac{\mu-r}{\sigma^2}$.

- The coefficients r, μ, σ are constant, so the optimal asset allocation h^* and optimal value g^* are constant.
  - Thus, **$\chi^h(T)$ is an exponential martingale for our choice of h^***.

Case 2: RSIM with factors and benchmark
What's new?

This case improves the Black-Scholes-Merton formulation 3 ways:

1. consider a financial market with m **risky securities**;

2. generalize the dynamics of the risky assets by introducing a **factor dependence in their drift**.

- Crucial because the Black-Scholes-Merton setup assumes a constant risk premium $\mu - r$ while empirical evidence suggests that the risk premium is stochastic.

- Introducing factors is also necessary to capture into our model the rapid development of the literature on **empirical asset pricing** and **factor investing**.

3. Include an **investment benchmark**, such as a financial index or a bespoke portfolio, to explore active and passive management simultaneously.

- This formulation recognizes that most professional asset managers are tasked with replicating or outperforming a benchmark.

- When no benchmark is specified, we retrieve the risk-sensitive asset management criterion as a special case.

The original surveyor's bench-mark!

Case 2: RSIM with factors and benchmark

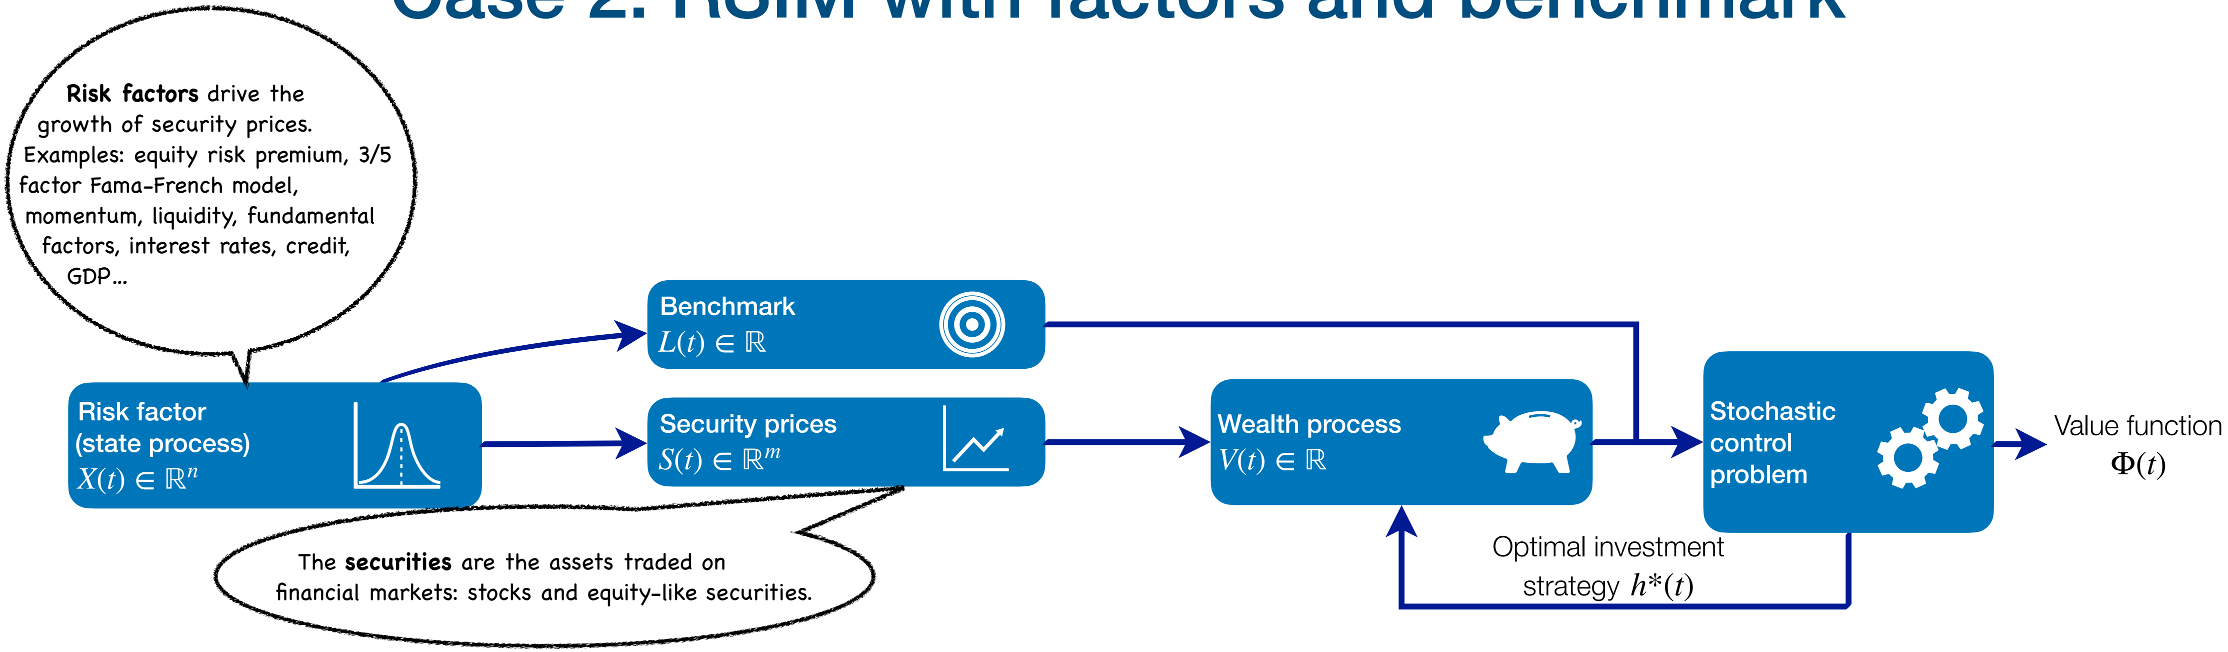

Mathematically,

Here $S_i(t)$ is discounted price of asset $i, i = 1, \ldots, m$

$$dX(t) = (b + BX(t))dt + \Lambda dW(t),$$
$$X(0) = x_0$$

$$\frac{dS_t^i}{S_t^i} = \left(a + AX(t)\right)_i dt + \sum_{j=1}^{d} \sigma_{ij} dW_t^j,$$
$$S_i(0) = s_i, i = 1, \ldots, m,$$

$$\frac{dL_t}{L_t} = \left(c + CX(t)\right) dt + \Xi dW_t, \quad L(0) = l$$

W_t is a $\mathbb{R}^d$-valued $\mathcal{F}_t$-Brownian motion on a filtered complete probability space $\left(\Omega, \mathcal{F}, \left(\mathcal{F}\right)_{t=0}^{T}, \mathbb{P}\right)$, with $d = n + m + 1$

$$\frac{dV(t)}{V(t)} = h'(t)\left(a + AX(t)\right) dt + h'(t)\Sigma dW(t)$$
$$V(0) = v$$

Log excess return process $R(t) := \ln \dfrac{V(t)}{L(t)}$.

$$dR(t) = \left[\left(-\frac{1}{2}h'(t)\Sigma'\Sigma h(t) + h'(t)a + \frac{1}{2}\Xi'\Xi - c\right) + \left(h'(t)A - C\right) X(t)\right] dt + \left(h'(t)\Sigma - \Xi\right) dW(t)$$

$$R(0) = \ln \frac{v}{l}.$$

Risk-sensitive criterion:

$$J(t, x, h; \theta) := -\frac{1}{\theta} \ln \mathbb{E}\left[e^{-\theta R(t)}\right]$$

Benchmark categorisation (for Case 2 and Case 3)

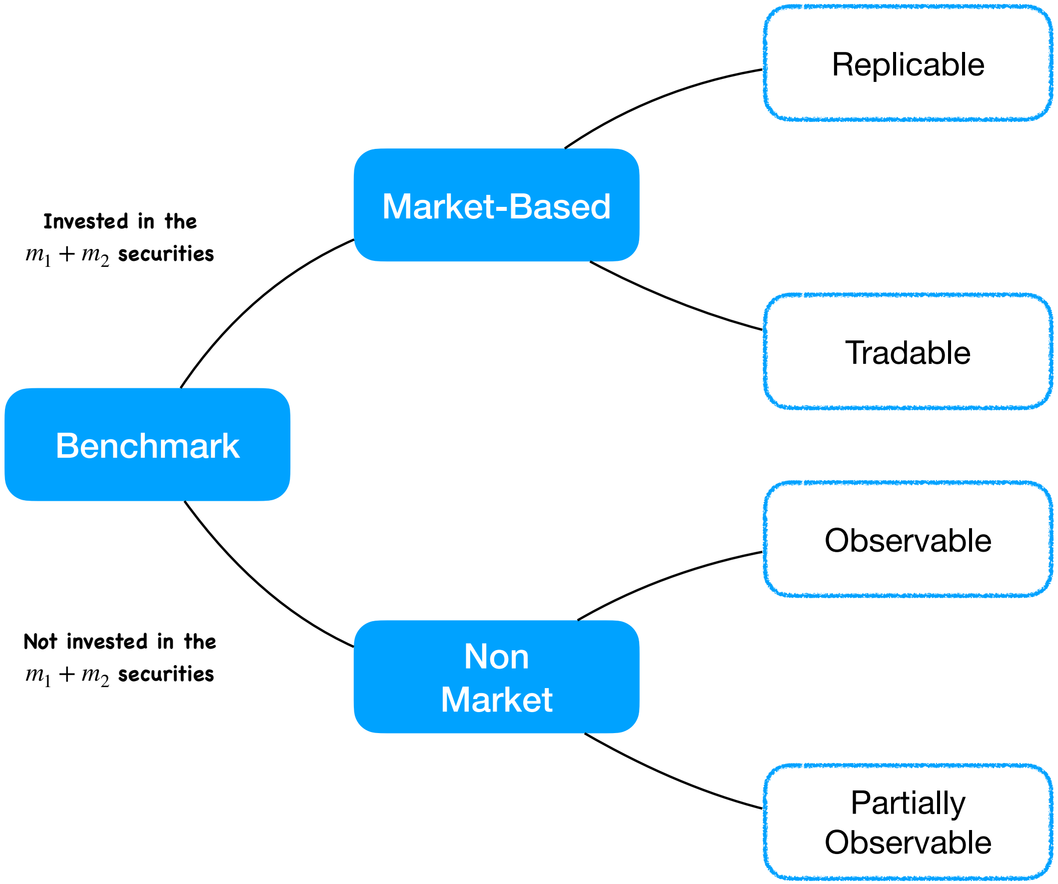

Replicable: The benchmark is a combination of the m_1 securities in the manager's investment universe.

Example: think of S&P500 for a large cap equity manager on the US market.

Tradable: The benchmark is a combination of the $m_1 + m_2$ securities traded on the financial market, but the manager might not be allowed to invest in all of them.

Example: think of Wilshire 5000 for a large cap equity manager.

Observable: The benchmark is observable but its individual constituents might not be observable individually.

Example: real estate property index.

Partially Observable: The benchmark is based directly on a factor or a combination of factors, so it cannot be observed directly.

Example: the return benchmark of an endowment fund set at the inflation rate plus 400 basis points. Because inflation is a partially observable factor, the manager needs to estimate the current level of the benchmark dynamically.

Constructing the benchmark criterion

- The discounted wealth process $V(t)$ is the market value of the self-financing investment portfolio subject to the investment strategy $h(t)$. It solves the SDE:

$$\frac{dV_t}{V_t} = \sum_{i=1}^{m} h_i(t) \frac{dS_i(t)}{S_i(t)} = h'(t)\big(a(t) + A(t)X(t)\big)\, dt + h'(t)\Sigma(t) dW_t, \qquad V_0 = v\,.$$

- The **log excess return** $R_t := \ln \frac{V_t}{L_t}$ tracks the portfolio's performance relative to its benchmark. Its dynamics is:

$$dR(t) = \left[\left(-\frac{1}{2}h'(t)\Sigma(t)\Sigma'(t)h(t) + h'(t)a(t) + \frac{1}{2}\Xi(t)\Xi'(t) - c(t)\right) + \big(h'(t)A(t) - C(t)\big)\, X(t)\right] dt + \big(h'(t)\Sigma(t) - \Xi(t)\big)\, dW(t), \quad R(0) = \ln \frac{v}{l}$$

- Without loss of generality, we index the benchmark's initial level on the investor's starting wealth v by setting $l := v$.

- The **risk-sensitive benchmarked criterion** J is

$$\begin{aligned} J(h;\theta) &:= -\frac{1}{\theta} \ln \mathbb{E}\left[e^{-\theta R(T)}\right] \\ &= -\frac{1}{\theta} \ln \mathbb{E}\left[\exp\left\{-\theta \int_0^T \left(-\frac{1}{2}h'(t)\Sigma(t)\Sigma'(t)h(t) + h'(t)a(t) + \frac{1}{2}\Xi(t)\Xi'(t) - c(t)\right) + \big(h'(t)A(t) - C(t)\big)\, X(t)dt - \theta \int_0^T \big(h'(t)\Sigma(t) - \Xi(t)\big)\, dW(t)\right\}\right]. \end{aligned}$$

- The Taylor expansion becomes $J(h, t) \approx \mathbb{E}\left[R(t)\right] - \frac{\theta}{2}\text{Var}\left[R(t)\right]$, which is again a \`**Dynamic Markowitz**' applied to the excess log return over the benchmark.

To solve the risk-sensitive benchmarked problem, we proceed with a change of measure, as in the previous case.

- We `**complete the exponential martingale**' to express the risk-sensitive criterion as

$$J(h;\theta) := -\frac{1}{\theta}\ln\mathbb{E}\left[\exp\left\{\theta\int_0^T g(t,X(t),h(t);\theta)dt\right\}\chi_T^h\right],$$

where

$$g(s,x,h;\theta) = \frac{1}{2}(\theta+1)\,h'\Sigma(s)\Sigma'(s)h - h'(a(s)+A(s)x) - \theta h'\Sigma(s)\Xi'(s) + c(s) + C(s)x - \frac{1}{2}(\theta-1)\,\Xi(s)\Xi'(s),$$

and

$$\chi_T^h := \exp\left\{-\theta\int_0^T\left(h(t)'\Sigma(t)-\Xi(t)\right)dW(t) - \frac{1}{2}\theta^2\int_0^T\left(h'(t)\Sigma(t)-\Xi(t)\right)\left(\Sigma'(t)h(t)-\Xi'(t)\right)dt\right\}.$$

- We also assume that the investment strategy $h(t)$ is in class $\mathscr{A}(T)$.

- **Definition 2.5 (Class $\mathscr{A}(T)$)**

A $\mathbb{R}^m$-valued control process $h(t)$ is in class $\mathscr{A}(T)$ if the following conditions are satisfied:

(i) $h(t)$ is progressively measurable with respect to $\left\{\mathscr{B}([0,t]) \otimes \mathscr{F}_t^Y\right\}_{t\geq 0}$ and is càdlàg;

(ii) $P\left(\int_0^T \left|h(s)\right|^2 ds < +\infty\right) = 1;$

(iii) the Doléans exponential χ_T^h is an exponential martingale, thus $\mathbb{E}\left[\chi_T^h\right] = 1.$

- Let $\mathbb{P}_h$ be the measure on $(\Omega, \mathcal{F}_T)$ defined via the **Radon-Nikodym derivative** $\dfrac{d\mathbb{P}_h}{d\mathbb{P}} := \chi_T^h$.

- Under the **measure** $\mathbb{P}_h$,

$$W^h(t) := W(t) + \theta \int_0^t \left(\Sigma'(s)h(s) - \Xi'(s)\right) ds$$

  is a standard Brownian motion for $h \in \mathscr{A}(T)$ and the risk-sensitive control criterion is

$$J^h(h;\theta) = -\frac{1}{\theta} \ln \mathbb{E}^h \left[\exp \left\{ \theta \int_0^T g(t, X(t), h(t); \theta) ds \right\} \right],$$

  where $\mathbb{E}^h [\,\cdot\,]$ denotes the expectation taken with respect to the measure $\mathbb{P}_h$.

- The **dynamics of the factors** $X(t)$ under the new measure,

$$dX(t) = \left[b(t) + B(t)X(t) - \theta\Lambda(t)\left(\Sigma'(t)h(t) - \Xi'(t)\right) \right] dt + \Lambda(t)dW^h(t),\ t \in [0,T],$$

  is a **controlled diffusion process** that depends on the investment strategy $h(t)$.

- We cannot conclude directly.
  - The function g depends on the stochastic process $X(t)$, so the solution is no longer deterministic.
- However,
  - the change of measure has expressed our risk-sensitive investment problem as a standard linear-quadratic-Gaussian risk-sensitive control problem, with

    a controlled state process $X(t)$ that is a linear and Gaussian and

    a reward function g quadratic in its h argument and linear in its x argument.
  - This problem is a special case of Jacobson's **LEQG problem** (Jacobson, 1973; Bensoussan, 1992), which we can solve efficiently using dynamic programming methods.

The Hamilton-Jacobi-Bellman equation

- Let $\Phi(t,x) := \sup_{h \in \mathscr{A}(T)} J^h(t,x;h;T,\theta)$ be the **value function** for the control problem, with associated **Hamilton-Jacobi-Bellman partial differential equation**:

$$\frac{\partial \Phi}{\partial t}(t,x) + \sup_{h \in \mathbb{R}^m} L_t^h(t,x,D\Phi,D^2\Phi) = 0,$$

where $D\Phi = \left(\frac{\partial \Phi}{\partial x_1}, \ldots, \frac{\partial \Phi}{\partial x_i}, \ldots, \frac{\partial \Phi}{\partial x_n}\right)'$, $D^2\Phi = \left[\frac{\partial^2 \Phi}{\partial x_i \partial x_j}\right]$, $i,j = 1,\ldots,n$, and

$$L_t^h(t,x,p,M) = \left(b(t) + B(t)x - \theta\Lambda(t)\Sigma'(t)h\right)' p + \frac{1}{2}\mathrm{tr}\left(\Lambda(t)\Lambda(t)'M\right)$$
$$-\frac{\theta}{2}p'\Lambda(t)\Lambda(t)'p - g(t,x,h;\theta),$$

for $p \in \mathbb{R}^n$ and subject to terminal condition $\Phi(T,x) = 0$.

- The term inside the $\sup$ is quadratic in h. Its unique maximizer corresponds to the candidate optimal control

$$\hat{h}(t,x,p) = \frac{1}{\theta+1}\left(\Sigma(t)\Sigma'(t)\right)^{-1}\left[a(t) + A(t)x + \theta\Sigma(t)\left(\Xi'(t) - \Lambda'(t)p\right)\right],$$

where (t,x,p) stands in for $\left(t, X(t), D\Phi(t, X(t))\right)$.

- **Moreover, the value function $\Phi(t,x) = \frac{1}{2}x'Q(t)x + x'q(t) + k(t)$, where $Q(t)$ is the unique symmetric non-negative solution to the matrix Riccati equation, $q(t)$ solves a linear ODE, and $k(t)$ is found by integration.** Specifically,

- $Q(t)$ solves

$$\dot{Q}(t) - Q(t)K_0(t)Q(t) + K_1'(t)Q(t) + Q(t)K_1(t) + \frac{1}{\theta+1}A'(t)\big(\Sigma(t)\Sigma'(t)\big)^{-1}A(t) = 0,$$

where $K_0(t) = \theta\left[\Lambda(t)\left(I - \frac{\theta}{\theta+1}\Sigma'(t)\big(\Sigma(t)\Sigma'(t)\big)^{-1}\Sigma(t)\right)\Lambda'(t)\right], K_1(t) = B(t) - \frac{\theta}{\theta+1}\Lambda(t)\Sigma'(t)\big(\Sigma(t)\Sigma'(t)\big)^{-1}A(t),$

and I is the $n \times n$ identity matrix.

- The vector-valued function $q(t)$solves

$$\dot{q}(t) + \left(K_1'(t) - Q(t)K_0(t)\right)q(t) + Q(t)\big(b + \theta\Lambda(t)\Xi'(t)\big) + \frac{1}{\theta+1}\left(A'(t) - \theta Q(t)\Lambda(t)\Sigma(t)\right)\left(\Sigma(t)\Sigma'(t)\right)^{-1}\left(a + \theta\Sigma(t)\Xi'(t)\right) - C(t) = 0,$$

- and $k(t) = \int_t^T \ell(t)dt$, where

$$\ell(s) = \frac{1}{2}\text{tr}\left(\Lambda(t)\Lambda'(t)Q(t)\right) - \frac{\theta}{2}q'(t)\Lambda(t)\Lambda'(t)q(t) + b'(t)q(t) + \frac{1}{2}\frac{1}{\theta+1}a'(t)\big(\Sigma(t)\Sigma'(t)\big)^{-1}a(t) + \frac{1}{2}\frac{\theta^2}{\theta+1}q'(t)\Lambda(t)\Sigma'(t)\big(\Sigma(t)\Sigma'(t)\big)^{-1}\Sigma(t)\Lambda'(t)q(t)$$
$$-\frac{\theta}{\theta+1}q'(t)\Lambda(t)\hat{\Sigma}'(t)\big(\Sigma(t)\Sigma'(t)\big)^{-1}a - \frac{\theta^2}{\theta+1}q'(t)\Lambda(t)\Sigma'(t)\big(\Sigma(t)\Sigma'(t)\big)^{-1}\Sigma(t)\Xi'(t) + \theta\Xi(t)\Lambda'(t)q(t) - \frac{1}{2}(\theta-1)\,\Xi(t)\Xi'(t) + \frac{\theta}{\theta+1}a'(t)\big(\Sigma(t)\Sigma'(t)\big)^{-1}\Sigma(t)\Xi'(t)$$
$$+\frac{1}{2}\frac{\theta^2}{\theta+1}\Xi(t)\Sigma'(t)\big(\Sigma(t)\Sigma'(t)\big)^{-1}\Sigma(t)\Xi'(t)\,.$$

- A standard **verification argument** completes the resolution of the stochastic control problem. The following Theorem summarizes these results.

- **Theorem (Risk-Sensitive Benchmarked Asset Management)**

  (i) The value function Φ is the unique $C^{1,2}$ solution to associated HJB PDE. It has the form $\Phi(t,x) = \frac{1}{2}x'Q(t)x + x'q(t) + k(t)$.

  (ii) There is a unique Borel measurable maximiser $\hat{h}(t,x,p)$ for $(t,x,p) \in [0,T] \times \mathbb{R}^n \times \mathbb{R}^n$ given by
  $$\hat{h}(t,x,p) = \frac{1}{\theta+1}\left(\Sigma(t)\Sigma'(t)\right)^{-1}\left[a(t) + A(t)x + \theta\Sigma(t)\left(\Xi'(t) - \Lambda'(t)p\right)\right].$$

  (iii) The maximizer is optimal, meaning $h^*(t,X(t)) = \hat{h}\left(t,X(t),D\Phi(t,X(t))\right)$.

Proposition: Fractional Kelly Strategy (FKS)

- The optimal investment strategy $h^*(t, \hat{X}(t))$ consists of an allocation between three funds: h^K, h^{Bench}, and h^{PIHP}.

  (i) The fund h^K is a **Kelly portfolio** with factor-dependent allocation

  $$h^K(t, X(t)) = (\Sigma(t)\Sigma'(t))^{-1}\big(a(t) + A(t)X(t)\big)\,.$$

  (ii) The fund h^{Bench} a **benchmark-tracking portfolio** with deterministic allocation

  $$h^{\text{Bench}}(t) = (\Sigma(t)\Sigma'(t))^{-1}\Sigma(t)\Xi'(t)\,.$$

  (iii) The fund h^{IHP} is an **Intertemporal Hedging Portfolio (IHP)** with factor-dependent allocation

  $$h^{IHP}(t, X(t)) = (\Sigma(t)\Sigma'(t))^{-1}\Sigma(t)\Lambda(t)'\big(q(t) + Q(t)X(t)\big)\,.$$

- The relative allocation of each fund is constant at $f := \dfrac{1}{\theta + 1}$ for h^K, $1 - f$ for h^{Bench}, and $f - 1$ for h^{IHP}.

The Active ↔ Passive Continuum Summarized

Optimal Asset Allocation

$$h^*(t) = \frac{\theta}{\theta+1}\underbrace{(\Sigma\Sigma')^{-1}\Sigma\Xi(t)}_{\text{Benchmark tracking portfolio}} - \frac{\theta}{\theta+1}\underbrace{(\Sigma\Sigma')^{-1}\Sigma\Lambda'(t)D\Phi(t,X(t))}_{\text{Intertemporal hedging portfolio (IHP)}} + \frac{1}{\theta+1}\underbrace{(\Sigma\Sigma')^{-1}(t)\big(a(t)+A(t)X(t)\big)}_{\text{Kelly portfolio}}.$$

A deeper look

Intuition: projection / OLS regression of the benchmark on the securities space.

Objective: hedges future consumption against changes in factor levels.

$(\Sigma\Sigma')^{-1}(t)(a(t)+A(t)X(t))$ ➜ Kelly portfolio

$= (\Sigma\Sigma')^{-1}(t)a(t)$ ➜ Security selection

$+(\Sigma_1\Sigma_1')^{-1}(t)A(t)X(t)$ ➜ factor tilting

Degree of active risk

Passive → No active risk $\theta \to \infty$

Active → High active risk $\theta \to 0$

Overbetting $\theta \in (-1,0)$

Strategy

- **Index fund**: Replicate the benchmark.
- **Core-Satellite**: Mix of benchmark and Kelly portfolio
- **Kelly investor**: Maximize the growth rate of wealth.
- **Overbetting**: "Sell your skill and bet on luck."

Benchmark-tracking portfolio

Choice of benchmark is crucial: provides risk signature and return.

Benchmark plays no role.

PIHP

Short position to hedge the optimal utility against changes in factor estimates.

PIHP plays no role.

The Active ↔ Passive Continuum Summarized

Optimal Asset Allocation

$$h^*(t) = \frac{\theta}{\theta+1}\underbrace{(\Sigma\Sigma')^{-1}\Sigma\Xi(t)}_{\text{Benchmark tracking portfolio}} - \frac{\theta}{\theta+1}\underbrace{(\Sigma\Sigma')^{-1}\Sigma\Lambda'(t)D\Phi(t,X(t))}_{\text{Intertemporal hedging portfolio (IHP)}} + \frac{1}{\theta+1}\underbrace{(\Sigma\Sigma')^{-1}(t)\big(a(t)+A(t)X(t)\big)}_{\text{Kelly portfolio}}.$$

A deeper look

Intuition: projection / OLS regression of the benchmark on the securities space.

Objective: hedges future consumption against changes in factor levels.

$(\Sigma\Sigma')^{-1}(t)(a(t)+A(t)X(t))$ ➜ Kelly portfolio

$= (\Sigma\Sigma')^{-1}(t)a(t)$ ➜ Security selection

$+(\Sigma_1\Sigma_1')^{-1}(t)A(t)X(t)$ ➜ factor tilting

Degree of active risk

Passive → No active risk
$\theta \to \infty$

Active → High active risk
$\theta \to 0$

Overbetting
$\theta \in (-1,0)$

Strategy

Index fund
Replicate the benchmark.

Core-Satellite
Mix of benchmark and Kelly portfolio

Kelly investor
Maximize the growth rate of wealth.

Overbetting
"Sell your skill and bet on luck."

Benchmark-tracking portfolio

Choice of benchmark is crucial: provides risk signature and return.

Benchmark plays no role.

PIHP

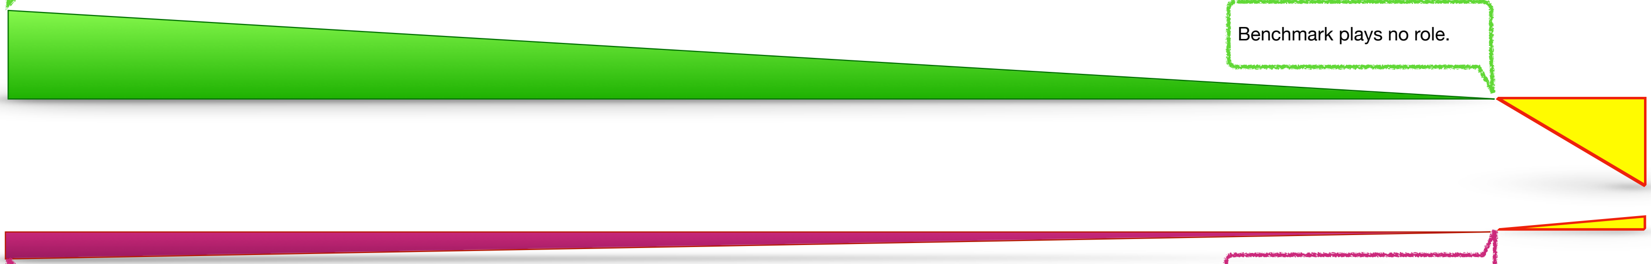

PIHP plays no role.

Short position to hedge the optimal utility against changes in factor estimates.

Case 3: RSIM with unobservable factors and expert opinions

- We remove the assumption that the factors are observable in real-time because in reality **factors may not be observable**.
  - **Statistical variables**, such as latent variables, are usually unobservable.
  - **Macroeconomic variables** are typically monthly or quarterly, with a lag and, possibly, revisions.
  - **Empirical asset pricing factors** are often constructed *ex-post* based on portfolio performance.
- However, removing this assumption raises two fundamental questions.
  - How are we going to estimate the factors?
  - What is the effect of unobservable factors on the stochastic control problem?

Filter Setup

- In dynamical systems, **filtering techniques** provide a natural way to
  - estimate the current value of a set of variables, called the **state variables**,
  - from another set of related but noisy variables, called the **observation variables**.


Photo courtesy of Pixabay

Filtering to Estimate the Factor Process

- In Case 3, the **state variables are the factors**, and **observations can come from**:
  - **Asset prices** depend on factor values, so they are relevant to estimating X. We can observe their prices directly on the financial market. While asset prices are a natural start, relying exclusively on them has the **downside of favoring momentum strategies**.
  - **Expert forecasts** and opinions offer another popular source of observations to complement asset prices. Experts include financial analysts, economists, policy experts, and nowcasting models.
    - This approach produces a dynamical model in the spirit of Black and Litterman (1992).
    - Expert opinions may exhibit **behavioral biases**. Davis and Lleo (2016, 2020) show how to identify and mitigate behavioral biases.
  - **Alternative data** are time series constructed from structured and unstructured data. Some examples include usage trends, product review trends, and sentiment indexes.
    - Davis and Lleo (2022) propose a risk-sensitive benchmarked model that combines asset prices, expert forecasts, and alternative data as observations.
    - Alternative data often feature non-Gaussian noise, so they are best modeled using jump-diffusion processes which are outside of the scope of this article.

Photo courtesy of Pixabay

Effect of unobservability on the stochastic control problem

- The kind of stochastic control problem we need to solve is a **partial observation problem**.
- Partial observation problems are **notoriously difficult to solve because they require simultaneously estimating the state and optimizing the system**.
- So, can we perform these two tasks separately?
- **Separation** is a good idea in practice, but it may produce a suboptimal solution.

- Fortunately, **risk-sensitive investment problems are separable**, so we do not lose optimality by estimating first and then optimizing (see Lleo and Runggaldier, 2023).

- This is great news because:

    - We can **implement the filter directly**, as discussed above.

    - We can **reuse the reasoning developed in the previous Case** to solve the control problem.

    - We can **use filtering to produce portable estimates** that we can use in other aspects of asset management business businesses, such as trading and risk management.

Photo courtesy of Pixabay

Case 3: RSIM with unobservable factors and expert opinions

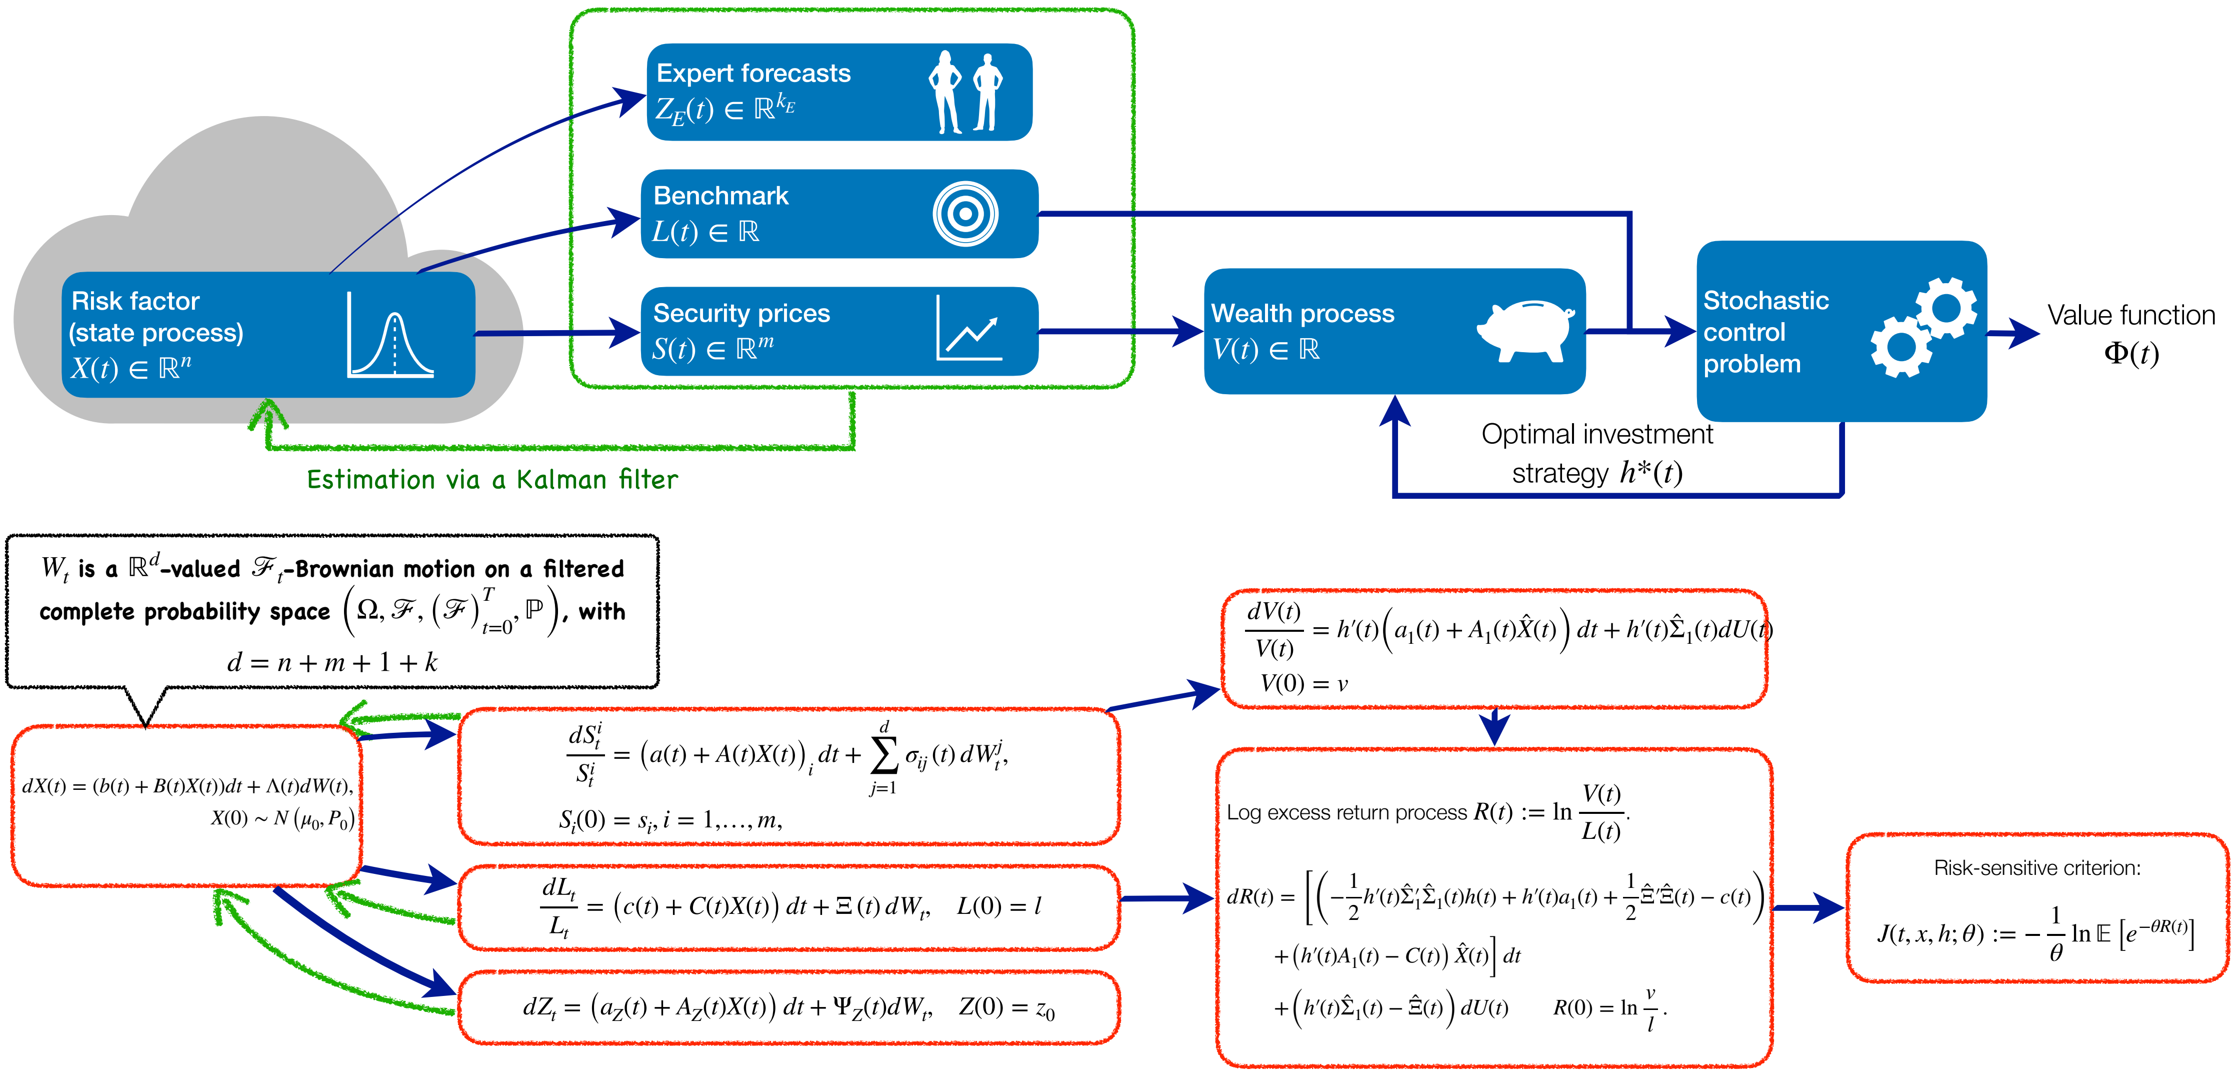

Tradable and Non-Tradable Assets

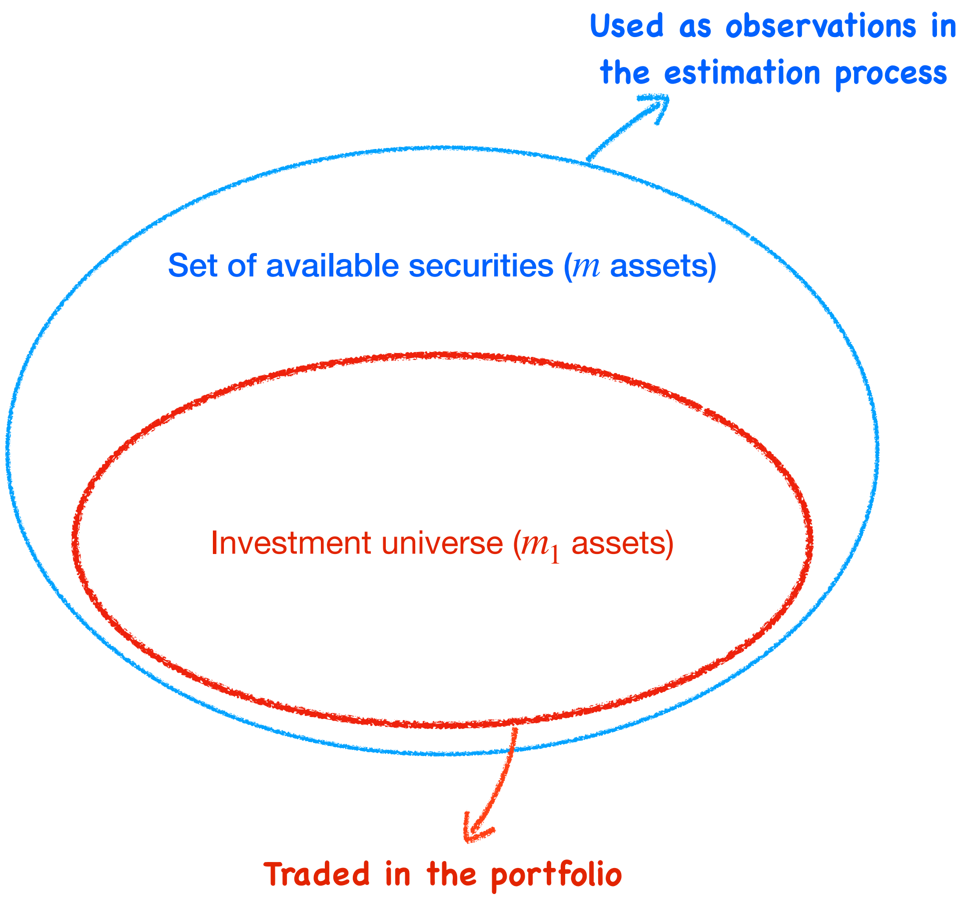

- In view of the estimation task, we adopt a broader view of the financial market.
- We include all the securities that can help estimate the factors, whether the investor can trade them or not.
- To distinguish between tradable and non-tradable assets, we split the m risky assets into an investment universe of $0 < m_1 \leq m$ assets that the investor is allowed to trade, and the remaining $m_2 = m - m_1 \geq 0$ assets that the investor can only observe.
- Accordingly, we express the securities price vector as $S_t := \left(S_t^{(1)'} \quad S_t^{(2)'}\right)'$, where $S_t^{(1)}$ is the m_1-vector process of tradable securities prices and $S^{(2)}$ is the m_2-vector process of untradable, but observable, securities prices.
- We perform a similar decomposition for the vector- and matrix-valued functions a, A and Σ. We denote by $a^{(1)}, A^{(1)}$ and $\Sigma^{(1)}$ the subvector and submatrix corresponding to the m_1 tradable assets, with an analogous definition for $a^{(2)}, A^{(2)}$ and $\Sigma^{(2)}$.

Filter Setup

- We gather all these observable processes to construct the observation process $Y(t)$.
  - Discounted financial securities prices and benchmark levels have a geometric dynamics, so they are not suitable observations for the (linear) Kalman filter.
  - We replace them with their excess log returns $\mathfrak{s}(t)$ as $\mathfrak{s}_i(t) = \ln(S_i(t)),\ i = 1,\ldots,m$, with affine-Gaussian dynamics

$$d\mathfrak{s}(t) = \left[\left(a(t) - \frac{1}{2}d_\Sigma(t)\right) + A(t)X(t)\right]dt + \Sigma(t)dW(t), \mathfrak{s}(0) = \ln(s),$$

where

$$d_\Sigma(t) = \left((\Sigma\Sigma')_{11}(t) \quad (\Sigma\Sigma')_{22}(t) \quad \ldots \quad (\Sigma\Sigma')_{mm}(t)\right)'.$$

  - Similarly, the excess log return vector $\mathfrak{l}(t) := \ln L(t)$ solves the SDE:

$$d\mathfrak{l}(t) = \left[\left(c(t) - \frac{1}{2}\Xi\Xi'(t)\right) + C(t)X(t)\right]dt + \Xi(t)dW(t), \quad \mathfrak{l}(0) = \ln(l).$$

Photo courtesy of Pixabay

- Accordingly, the **observation vector** $Y(t) := \left(\mathfrak{z}(t) \quad Z(t) \quad \mathfrak{l}(t)\right)'$ is affine in the state with Gaussian noise:

- $dY(t) = (a_Y(t) + A_Y(t)X(t))dt + \Gamma(t)dW(t), \; Y(0) = y_0,$

  where

$$a_Y(t) = \begin{pmatrix} a(t) - \frac{1}{2}d_\Sigma(t) \\ a_Z(t) \\ c(t) - \frac{1}{2}\Xi'\Xi(t) \end{pmatrix}, \quad A_Y(t) = \begin{pmatrix} A(t) \\ A_Z(t) \\ C(t) \end{pmatrix}, \quad \Gamma(t) = \begin{pmatrix} \Sigma(t) \\ \Psi_Z(t) \\ \Xi'(t) \end{pmatrix}.$$

- Now, let $\mathscr{F}_t^Y = \sigma\{Y(u), 0 \le u \le t\}$ **be the filtration generated by the observation process only**.

- The conditional distribution of the factor process $X(t)$ is normal $N(\hat{X}(t), P(t))$ where

  - $\hat{X}(t) = \mathbb{E}[X(t) \,|\, \mathscr{F}_t^Y]$ satisfies the **Kalman filter equation** and

  - $P(t)$ is a deterministic matrix-valued function.

$\mathscr{F}_t$ is the full set of information.

It includes everything:
- factors,
- securities,
- benchmark, and
- expert forecasts.

$\mathscr{F}_t^Y = \sigma\{Y(u), 0 \le u \le t\}$ is the subset of information that we get from observing $Y(t)$ only, that is, the
- securities,
- benchmark, and
- expert forecasts.

Kalman Filter:
Filtering equations (Davis 1977, Davis and Lleo 2011, 2020)

- The **Kalman estimate** $\hat{X}(t)$ is the unique solution of the SDE:

$$d\hat{X}(t) = (b(t) + B(t)\hat{X}(t))dt + \hat{\Lambda}(t)dU(t), \qquad \hat{X}(0) = \mu_0,$$

where $\hat{\Lambda}(t) = \left(\Lambda\Gamma(t)' + P(t)A_Y'\right)\left(\Gamma(t)\Gamma(t)'\right)^{-1/2}$.

- The **variance** $P(t)$ is the unique non-negative definite symmetric solution of the matrix Riccati equation

$$\begin{aligned}\dot{P}(t) = \Lambda\Upsilon^{\perp}\Lambda' - P(t)A_Y'\left(\Gamma(t)\Gamma(t)'\right)^{-1}A_Y P(t) + \left(B(t) - \Lambda\Gamma(t)'\left(\Gamma(t)\Gamma(t)'\right)^{-1}A_Y\right)P(t) \\ + P(t)\left(B(t)' - A_Y'\left(\Gamma(t)\Gamma(t)'\right)^{-1}\Gamma(t)\Lambda'\right), \qquad P(0) = P_0,\end{aligned}$$

with $\Upsilon^{\perp} := I - \Gamma(t)'\left(\Gamma(t)\Gamma(t)'\right)^{-1}\Gamma(t)$.

- The **innovation** process $U(t)$ defined by

$$dU(t) = \left(\Gamma(t)\Gamma(t)'\right)^{-1/2}(dY^1(t) - A_Y\hat{X}(t)dt), \qquad U(0) = 0$$

is a $\mathbb{R}^{m+K}$-valued $\left(\mathcal{F}_t^Y\right)$-Brownian motion on $\left(\Omega, \mathcal{F}, \left(\mathcal{F}_t\right)_{t=0}^{T}, \mathbb{P}\right)$.

- As a final step in the application of the filter, we decompose the $(m+k+1) \times (m+k+1)$ matrix $(\Gamma(t)\Gamma'(t))^{1/2}$ as $(\Gamma(t)\Gamma'(t))^{1/2} := \begin{pmatrix} \hat{\Sigma}'(t) & \hat{\Psi}'_Z(t) & \hat{\Xi}'(t) \end{pmatrix}'$, where

  - $\hat{\Sigma}(t)$ is a $m \times (m+k+1)$ matrix such that $\hat{\Sigma}\hat{\Sigma}'(t) = \Sigma\Sigma'(t)$,

  - $\hat{\Psi}_Z(t)$ is a $k \times (m+k+1)$ matrix such that $\hat{\Psi}_Z\hat{\Psi}_Z(t)' = \Psi_Z\Psi_Z(t)'$, and

  - $\hat{\Xi}(t)$ is a $(m+k+1)$-element vector such that $\hat{\Xi}\hat{\Xi}'(t) = \Xi\Xi'(t)$

  - and where we are using the notational shortcut $\Gamma\Gamma(t)'$ for $\Gamma(t)\Gamma(t)'$, $\Sigma\Sigma(t)'$ for $\Sigma(t)\Sigma(t)'$, etc.

- Having completed the estimation task, we move to the optimisation.

- The stochastic control problem is separable, so we simply `substitute' the Kalman filter estimate $\hat{X}(t)$ for the unobservable factor value $X(t)$ in the risk-sensitive criterion:

$$J(h;\theta) := -\frac{1}{\theta}\ln\mathbb{E}\left[\exp\left\{-\theta\int_0^T\left(-\frac{1}{2}h'(t)\hat{\Sigma}_1\hat{\Sigma}_1'(t)h(t) + h'(t)a_1(t) + \frac{1}{2}\hat{\Xi}\hat{\Xi}'(t) - c(t)\right)\right.\right.$$
$$\left.\left. + \left(h'(t)A_1(t) - C(t)\right)\hat{X}(t)dt - \theta\int_0^T\left(h'(t)\hat{\Sigma}_1(t) - \hat{\Xi}(t)\right)dU(t)\right\}\right],$$

where the dynamics of $\hat{X}(t)$ is given by the Kalman filter.

- A standard verification argument completes the resolution of the stochastic control problem. The following Theorem summarizes these results.

- **Theorem (Risk-Sensitive Benchmarked Asset Management)**

  (i) The value function Φ is the unique $C^{1,2}$ solution to associated HJB PDE. It has the form $\Phi(t,\hat{x}) = \frac{1}{2}\hat{x}'Q(t)\hat{x} + \hat{x}'q(t) + k(t)$.

  (ii) There is a unique Borel measurable maximiser $\hat{h}(t,x,p)$ for $(t,x,p) \in [0,T] \times \mathbb{R}^n \times \mathbb{R}^n$ given by
  $\hat{h}(t,\hat{x},p) = \frac{1}{\theta+1}\left(\hat{\Sigma}_1\hat{\Sigma}_1'(t)\right)^{-1}\left[a_1(t) + A_1(t)\hat{x} + \theta\hat{\Sigma}_1(t)\left(\hat{\Xi}'(t) - \hat{\Lambda}'(t)p\right)\right]$.

  (iii) The maximizer is optimal, meaning $h^*(t,\hat{X}(t)) = \hat{h}\left(t,\hat{X}(t),D\Phi(t,\hat{X}(t))\right)$.

Proposition: Fractional Kelly Strategy (PFKS)

- The optimal investment strategy $h^*(t, \hat{X}(t))$ consists of an allocation between three funds: h^K, h^{Bench}, and h^{PIHP}.

  (i) The fund h^K is a **personal Kelly portfolio** with factor-dependent allocation

$$h^K(t, \hat{X}(t)) = (\hat{\Sigma}_1 \hat{\Sigma}_1'(t))^{-1} \Big(a_1(t) + A_1(t)\hat{X}(t) \Big) .$$

  (ii) The fund h^{Bench} a **benchmark-tracking portfolio** with deterministic allocation

$$h^{\text{Bench}}(t) = (\hat{\Sigma}_1 \hat{\Sigma}_1'(t))^{-1} \hat{\Sigma}_1(t) \hat{\Xi}' .$$

  (iii) The fund h^{IHP} is an **Intertemporal Hedging Portfolio (IHP)** with factor-dependent allocation

$$h^{PIHP}(t, \hat{X}(t)) = (\hat{\Sigma}_1 \hat{\Sigma}_1'(t))^{-1} \hat{\Sigma}_1(t) \hat{\Lambda}(t)' \Big(q(t) + Q(t)\hat{X}(t) \Big) .$$

- The relative allocation of each fund is constant at $f := \dfrac{1}{\theta + 1}$ for h^K, $1 - f$ for h^{Bench}, and $f - 1$ for h^{IHP}.

The Active ↔ Passive Continuum Summarized

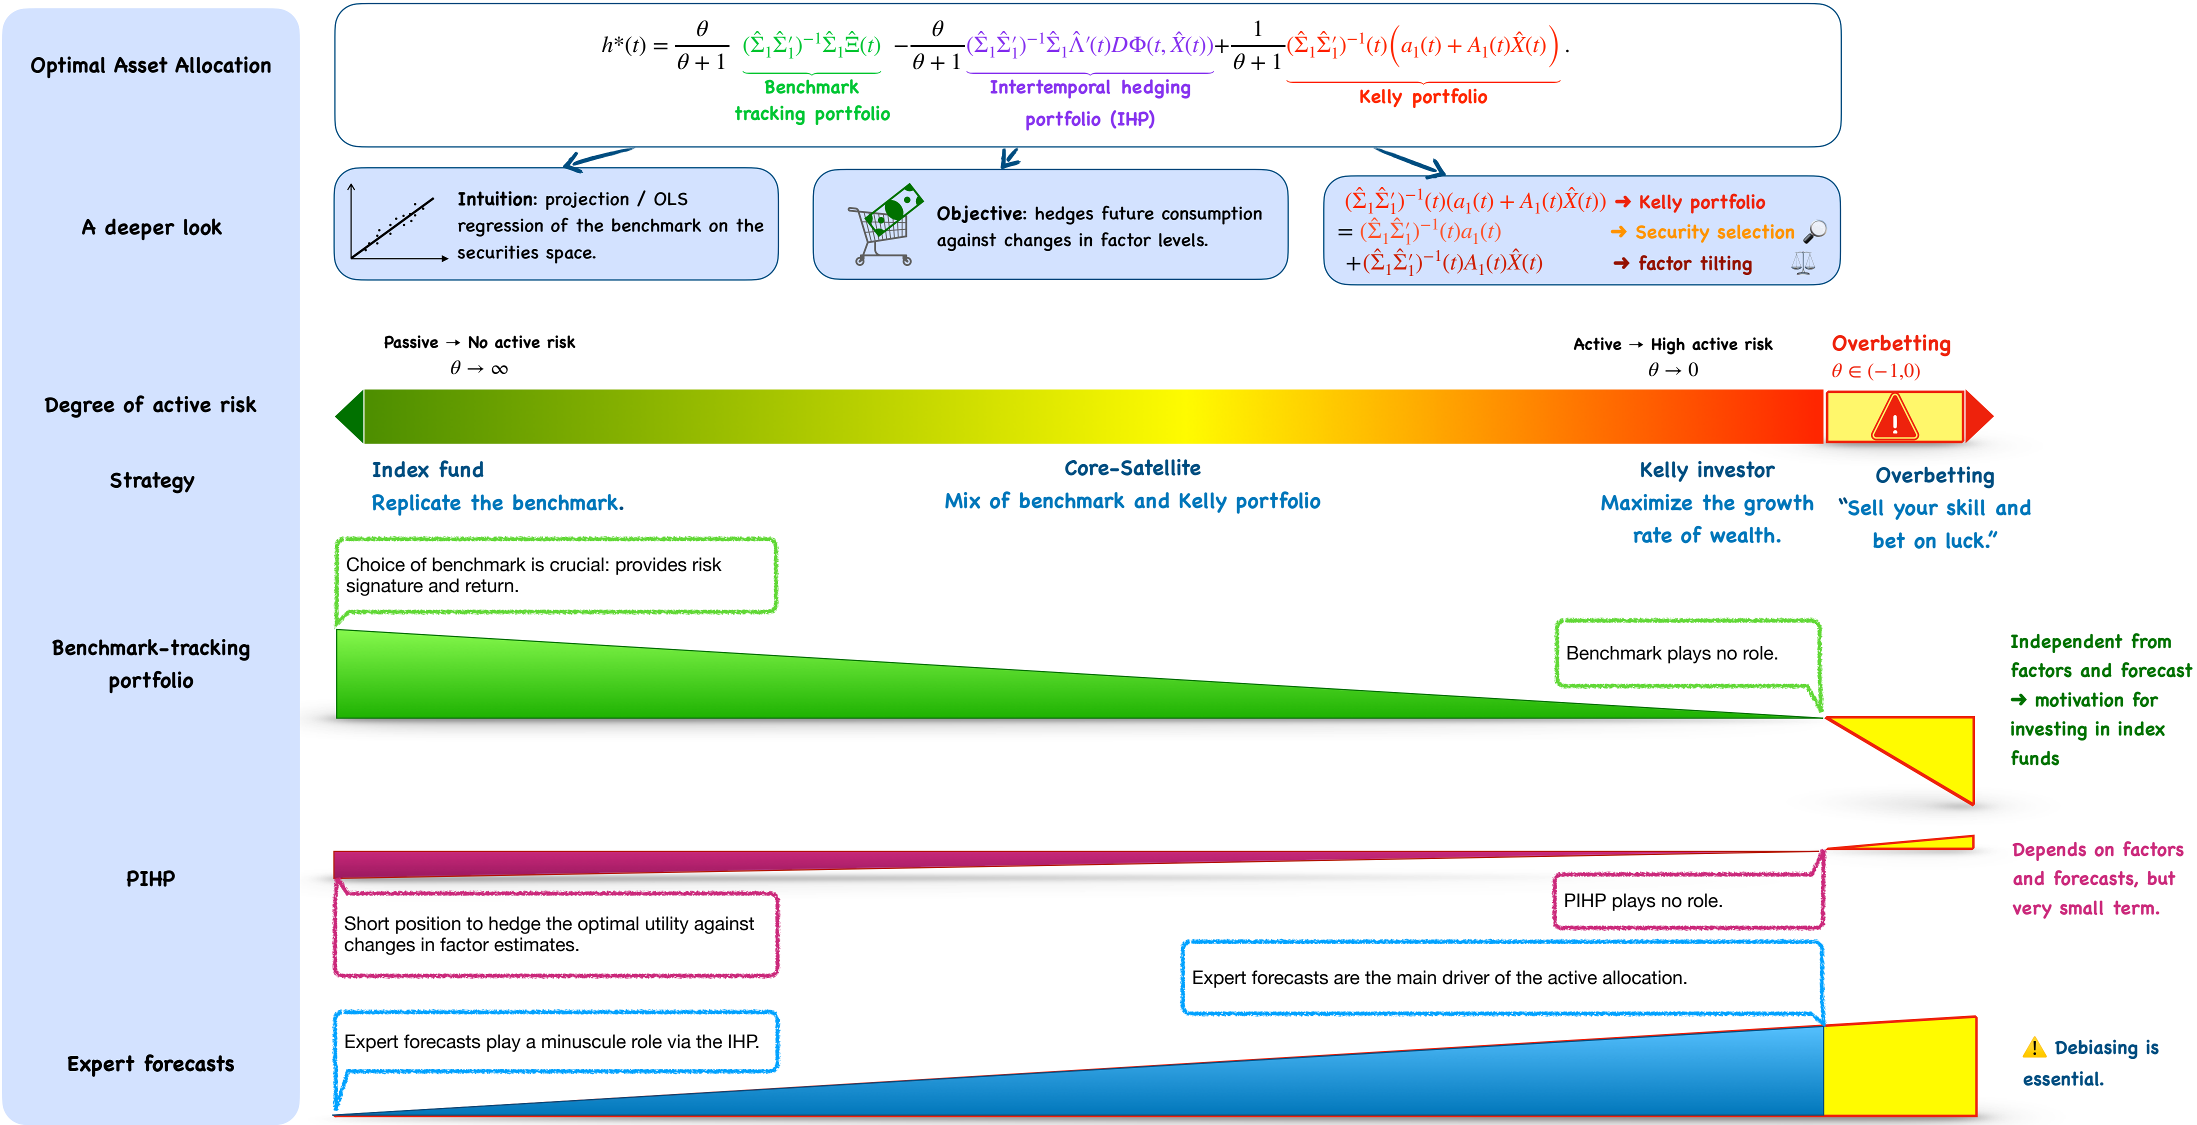

A Word of Thanks

I extend my heartfelt gratitude to

✓Paul Wilmott,

✓Riaz Ahmad,

✓Paul Shaw and 7City,

✓Randeep Gug and the team at FitchLearning.

Being part of the CQF in 2003 changed my life.

To all my CQF students, past, present, and future, in the $\mathbb{P}$- or any other measure.

To an absent friend of our community

Professor William T. Ziemba, a.k.a. Doctor Z.

Happy Birthday,
CQF & Wilmott Mag!!!