Currency Risk Management

FX Bootcamp
EPFL

Dr. Urban Ulrych[1, 2, 3]
urban.ulrych@epfl.ch

[1]College of Management of Technology,
École Polytechnique Fédérale de Lausanne

[2]Department of Mathematics,
ETH Zurich

[3]Swiss Finance Institute

Swiss Franc Unpeg

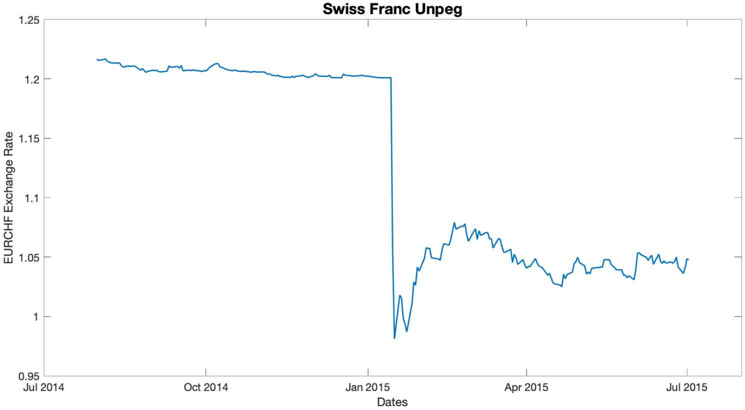

On January 15th 2015, the Swiss National Bank unexpectedly removed the peg of 1.20 francs per euro. In the initial reaction to the news, the Swiss franc rallied a massive 30% versus the euro and 25% against the US dollar.

Foreign Asset Holdings

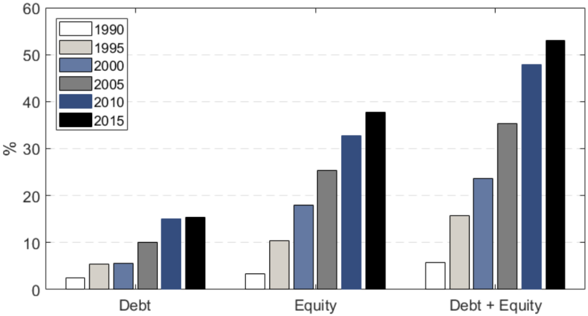

Figure 1: U.S.-Resident Holdings of Foreign Debt and Equity (% GDP)
The figure presents year-end U.S.-resident holdings of foreign debt and equity as a percentage of Gross Domestic Product (GDP). Data is collected from Philip Lane's website (see Lane and Milesi-Ferretti (2001, 2007) for further details).

Source: Opie and Riddiough, Global Currency Hedging with Common Risk Factors, 2018.

Currency Risk Management

FX Market and Interest Rate Parities

What Is the FX Market?

- *Foreign Exchange (FX)* is the global marketplace for trading national currencies.
- With around *7–8 trillion USD* traded daily, FX is the *largest financial market* in the world.
- Trades occur *24 hours a day*, five days a week, across major financial centers.
- The market is highly *liquid*, resulting in tight bid-ask *spreads* and *low transaction costs*.
- The FX market is largely *decentralized*, operating over-the-counter (OTC) rather than on a central exchange.

Who Trades in the FX Market?

- FX participants span across sectors, with different *motives*:
  - **Corporations:** Manage foreign revenues and costs (e.g., exports/imports).
  - **Asset Managers:** Rebalance global portfolios, currency overlay.
  - **Hedge Funds:** Speculative trading on macro views.
  - **Central Banks:** Currency stability, monetary policy.
  - **Banks/Dealers:** Market making and arbitrage.
- Activities range from *hedging* to *speculation*.
- Instruments: Spot, Forward, Swap, Option, Non-Deliverable Forward (NDF).

Why Does Currency Risk Matter?

- Currency fluctuations affect:
  - *Investment returns* in foreign assets.
  - *Corporate earnings* and competitiveness.
  - *Balance sheets* of firms and sovereigns.
- Even passive investors face *unhedged FX exposure* in global portfolios.
- FX risk management helps reduce volatility and protect against *unexpected exchange rate movements*.

Covered Interest Rate Parity

- *Covered interest rate parity* is a *no-arbitrage condition* representing an equilibrium state under given interest rates available on bank deposits in two countries.
- Investors should not earn arbitrage profits by *borrowing* in a country with a lower interest rate, *exchanging* for foreign currency, and *investing* in a foreign country with a higher interest rate, due to gains or losses from *exchanging back* to their domestic currency at maturity.
- Consider a *domestic investor* who invests into a foreign money or bond market.
- Let $r_{t,k}^{d}$ denote the k-period cumulative discrete interest (with compounding) earned on the domestic currency at time t.
- Let $r_{t,k}^{f}$ denote the k-period cumulative discrete interest (with compounding) earned on the foreign currency at time t.

Covered Interest Rate Parity

- Let S_t denote the *spot exchange rate* expressed as the number of units of domestic currency (e.g., CHF) per unit of foreign currency (e.g., EUR), for example, 0.95 CHF per EUR.
- Let $F_{t,k}$ denote the *forward foreign exchange rate*, i.e., the FX rate one can agree to today for a foreign currency transaction with delivery in k periods. If one sells the foreign currency forward, she will receive $F_{t,k}$ units of the domestic currency per unit of foreign currency at time $t + k$.
- By *no arbitrage*, the forward rate is given by:

$$F_{t,k} = S_t \frac{1 + r^d_{t,k}}{1 + r^f_{t,k}}.$$

Covered Interest Rate Parity

- This result is known as *covered interest parity* (CIP). It says that the forward rate must be such that the return on a riskless investment in the domestic currency $(1 + r^d_{t,k})$ is identical to that of a hedged investment in the foreign currency, $F_{t,k}(1 + r^f_{t,k})/S_t$.

- Put differently, borrowing at home and lending abroad or doing the reverse earns *zero return* if the FX risk is *hedged*.

Uncovered Interest Rate Parity

- *Uncovered interest rate parity* considers the return from an unhedged investment in foreign currency.
- It states that the *expected return* from an investment in *foreign* currency should be the same as that of an investment in *domestic* currency – *on average*, the FX appreciation/depreciation should offset the interest rate differential.
- The uncovered interest rate parity (UIP) is:

$$(1 + r^d_{t,k}) = \frac{\mathbb{E}_t[S_{t+k}]}{S_t}(1 + r^f_{t,k}).$$

- For example, if $r^f_{t,k} > r^d_{t,k}$, UIP says that the foreign currency should, on average, depreciate at a rate matching the interest rate differential in order to make the domestic and foreign investments *equally profitable*. If you *gain* on the the *interest rate differential* you tend to *lose* on the *exchange rate move*.

Uncovered Interest Rate Parity

- Thus, UIP claims that $F_{t,k} = \mathbb{E}_t[S_{t+k}]$ – the *forward rate* is an *unbiased predictor* of the *future spot rate*.

- By contrast with CIP, UIP is *not* an *arbitrage relationship*. It is a condition based on equilibrium reasoning that *may* or *may not hold*.

- One case where UIP should hold is that where investors are *risk-neutral*.

- There are actually two statements in UIP:
  - There is no risk premium from holding the foreign currency – the appreciation/depreciation of the foreign currency offsets the interest rate differential.
  - Currency excess returns are unpredictable.

Uncovered Interest Rate Parity

- *Empirically*, UIP does *not hold*: there is a currency risk premium and foreign exchange returns are predictable.
- *On average*, countries with high interest rates also experience an appreciation of their currency.
- This means that there are *profitable opportunities*, i.e., one can capture a positive currency risk premium by borrowing in low interest rate countries and investing in high interest rate countries. However, such a strategy is subject to *steep drawdowns*.

Currency Carry Trade

- The *carry trade* is a popular *trading strategy* that exploits deviations from UIP (if UIP would hold empirically, carry trade would not be possible).
- Carry traders *borrow* in *low interest rate currencies* and *lend* in *high interest rate currencies*.
- Let $r^f_{t,k} > r^d_{t,k}$, then the carry trade involves going long the foreign currency and short domestic (and vice versa for $r^f_{t,k} < r^d_{t,k}$).
- The return $R_{t,k}$ from borrowing in the domestic currency and taking a long position in the foreign currency is:

$$R_{t,k} = (1 + r^f_{t,k})\frac{S_{t+k}}{S_t} - (1 + r^d_{t,k}).$$

- A trader stands to make a profit of the difference in the interest rates of the two countries as long as the *exchange rate* between the currencies *does not change*.

Currency Carry Trade

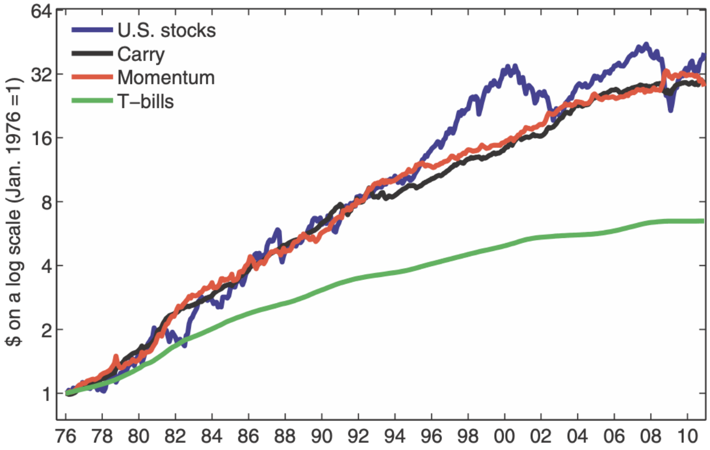

The cumulative returns to the carry (and momentum) portfolios are almost as high as the cumulative return to investing in stocks!

Source: Burnside et al., Carry Trade and Momentum in Currency Markets, 2011.

Currency Hedging and Risk Management

International Asset Allocation

- *International asset allocation* is a natural way to improve risk-adjusted portfolio performance $\rightarrow$ *diversification*.

- One of the main challenges for global asset allocation is the *currency risk*.

- Foreign currency exposure adds risk and full hedging is theoretically and empirically not optimal $\rightarrow$ we can do better by taking into account the correlations $\rightarrow$ *currency overlay* portfolios.

- The *return* in *domestic currency* of an investment in an *asset denominated* in *foreign currency* has two components:
  - The return on the *investment* expressed in *foreign currency*; and
  - The return on the *foreign currency* expressed in *domestic currency*.

Hedged Portfolio Return

- Variables at time t:
  - $P_{i,t}$: Value of asset i in local currency.
  - $\tilde{R}_{i,t+1} = (P_{i,t+1} - P_{i,t})/P_{i,t}$: Simple return of asset i.
  - $S_{c,t}$: Spot foreign exchange rate in domestic terms per foreign currency c.
  - $e_{c,t+1} = (S_{c,t+1} - S_{c,t})/S_{c,t}$: Return of the foreign exchange rate.

- Return of an *unhedged* investment in *domestic* terms $\bar{R}_{i,t+1}$:

$$\begin{aligned}\bar{R}_{i,t+1} &= \frac{P_{i,t+1}S_{c,t+1}}{P_{i,t}S_{c,t}} - 1 = (1 + \tilde{R}_{i,t+1})(1 + e_{c,t+1}) - 1 = \\ &= \tilde{R}_{i,t+1} + e_{c,t+1} + \tilde{R}_{i,t+1}e_{c,t+1}.\end{aligned}$$

Hedged Portfolio Return

- Denote with $x_{i,t}$ a fraction of wealth invested in asset i at time t.
- Consider an investor with an arbitrary domestic currency and a portfolio consisting of N assets.
- The *unhedged portfolio return* is given by

$$R_{\mathcal{P},t+1} = \sum_{i=1}^{N} x_{i,t}\bar{R}_{i,t+1}.$$

Hedged Portfolio Return

- Variables at time t:
  - $F_{c,t}$: Forward exchange rate in domestic currency per foreign currency c.
  - $f_{c,t} = (F_{c,t} - S_{c,t})/S_{c,t}$: The forward premium.
  - $\phi_{c,t}$: Relative notional value of the forward contract position (short for $\phi_{c,t} > 0$) in currency c, expressed as a percentage of total portfolio value.
- Suppose that there is a universe of K possible foreign currencies which can be traded on the market. The *return* of a *portfolio* with currency *forwards* is equal to

$$R^h_{\mathcal{P},t+1} = R_{\mathcal{P},t+1} + \sum_{c=2}^{K+1} \phi_{c,t}(f_{c,t} - e_{c,t+1}). \tag{1}$$

- For the domestic currency $c = 1$, set $\phi_{1,t} = 1 - \sum_{c=2}^{K+1} \phi_{c,t}$.

Hedged Portfolio Return

- Denote with $\mathcal{A}_{c,t}$ a set of all assets denominated in currency c. Then, $w_{c,t} := \sum_{j\in\mathcal{A}_{c,t}} x_{j,t}$ is a fraction of wealth invested in assets denominated in currency c.

- Define a *hedge ratio* as $h_{c,t} := \phi_{c,t}/w_{c,t}$, for $w_{c,t} \neq 0$.

- Return of a *fully hedged portfolio*, where $h_{c,t} = 1$ for $c = 2, \dots, k+1$ and $\phi_{c,t} = 0$ for $c = k+2, \dots, n+1$, is equal to

$$R^{fh}_{\mathcal{P},t+1} = \sum_{i=1}^{N} x_{i,t}\tilde{R}_{i,t+1} + \sum_{c=2}^{K+1} w_{c,t} f_{c,t} + \sum_{i=1}^{N} x_{i,t}\tilde{R}_{i,t+1} e_{c,t+1}.$$

Hedged Portfolio Return

- Define an *exposure* to currency c as $\psi_{c,t} := w_{c,t} - \phi_{c,t}$.
- *Hedged portfolio return* from (1) can be rewritten as

$$R^{h}_{\mathcal{P},t+1} = R^{fh}_{\mathcal{P},t+1} + \sum_{c=2}^{K+1} \psi_{c,t}(e_{c,t+1} - f_{c,t}). \tag{2}$$

- Domestic currency exposure becomes $\psi_{1,t} = -\sum_{c=2}^{n+1} \psi_{c,t}$, which implies that the currency portfolio is a *zero investment portfolio*.
- Note that these expressions are *model-free*! No underlying dynamics for asset or currency returns assumed.

- Optimal *currency overlay* (i.e., currency hedging portfolio) usually refers to the *optimal exposure* $\boldsymbol{\Psi}^*_{t,risk}$ that *minimizes* the *variance* of the *hedged portfolio returns* from (2):

$$\boldsymbol{\Psi}^*_{t,risk} = \arg\min_{\boldsymbol{\Psi}_t} \left\{ \mathrm{Var}(R^h_{\mathcal{P},t+1}) \right\} = -\mathrm{Var}(\mathbf{e}_{t+1} - \mathbf{f}_t)^{-1} \cdot \mathrm{Cov}(R^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t), \tag{3}$$

  where the $(K \times K)$ matrix $\mathrm{Var}^{-1}(\mathbf{e}_{t+1} - \mathbf{f}_t)$ is the inverse of the covariance matrix of the excess currency returns, and the $(K \times 1)$ vector $\mathrm{Cov}(R^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t)$ represents the covariances between the fully hedged portfolio returns and the excess currency returns.

- *Derivation*: Take the vector derivative of $\mathrm{Var}(R^h_{\mathcal{P},t+1})$ w.r.t. $\boldsymbol{\Psi}_t$, apply first- and second-order optimality conditions.

Currency Hedging and Risk Management

- Intuition: *Hedge more* against currencies with *high variance* and *positive correlation* to (fully hedged) portfolio returns.

- For easier interpretation, assume $K = 1$:
  - In case of zero correlation (covariance) between the asset portfolio and currency excess returns $\text{Cov}(R^h_{\mathcal{P},t+1}, e_{t+1} - f_t) = 0$, it is optimal to fully hedge (have zero currency exposure). In this case currency exposure only adds risk to the investor's portfolio.
  - If $\text{Cov}(R^h_{\mathcal{P},t+1}, e_{t+1} - f_t) > 0$ an investor can decrease risk by over-hedging.
  - If $-\text{Var}(e_{t+1} - f_t) < \text{Cov}(R^h_{\mathcal{P},t+1}, e_{t+1} - f_t) < 0$ partial hedging is optimal.
  - If $\text{Cov}(R^h_{\mathcal{P},t+1}, e_{t+1} - f_t) < -\text{Var}(e_{t+1} - f_t)$ it is optimal to under-hedge (increase exposure to such risk reducing currency).

Currency Hedging and Risk Management

- Let's compute the optimal currency exposure $\boldsymbol{\Psi}^*_{t,mv}$ that maximizes the *mean-variance objective*:

$$
\begin{aligned}
\boldsymbol{\Psi}^*_{t,mv} &= \underset{\boldsymbol{\Psi}_t}{\arg\max} \left\{ \mathrm{E}[R^h_{\mathcal{P},t+1}] - \frac{\lambda}{2}\mathrm{Var}(R^h_{\mathcal{P},t+1}) \right\} \\
&= -\mathrm{Var}(\mathbf{e}_{t+1} - \mathbf{f}_t)^{-1} \cdot \left[\mathrm{Cov}(R^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t) - \frac{1}{\lambda}\mathrm{E}[\mathbf{e}_{t+1} - \mathbf{f}_t] \right] \\
&= \boldsymbol{\Psi}^*_{t,risk} + \frac{1}{\lambda}\mathrm{Var}(\mathbf{e}_{t+1} - \mathbf{f}_t)^{-1}\mathrm{E}[\mathbf{e}_{t+1} - \mathbf{f}_t].
\end{aligned}
$$

- Observe that $\boldsymbol{\Psi}^*_{t,mv}$ can be interpreted as a sum of: $\boldsymbol{\Psi}^*_{t,risk}$ and the *market price of currency risk*, which represents the *trade-off* between the expected excess return on currencies and the variance of these returns.

- In the presence of constraints on $\boldsymbol{\Psi}^*_{t,risk}$ or $\boldsymbol{\Psi}^*_{t,mv}$, the solution is obtained via *quadratic programming*.

Currency Hedging and Risk Management

- The optimal *variance-minimizing currency exposure* $\mathbf{\Psi}^*_{t,risk}$ can equivalently be obtained from a *multivariate OLS regression*, see Campbell et al. (2010).
- Regress the demeaned fully hedged portfolio return $R^{fh}_{\mathcal{P},t+1}$ on the demeaned excess currency returns $(\mathbf{e}_{t+1} - \mathbf{f}_t)$, *without an intercept*, and swap the sign of the regression coefficients:

$$R^{fh}_{\mathcal{P},t+1} = (-\mathbf{\Psi}_t^\top)(\mathbf{e}_{t+1} - \mathbf{f}_t) + \varepsilon_{t+1}.$$

- The estimated regression coefficients $\mathbf{\Psi}^*_{t,risk}$ correspond to the *optimal currency exposures* that minimize the variance of the hedged portfolio returns.
- This OLS interpretation provides an intuitive view: hedge more against currencies that *co-move* with the fully hedged portfolio return.

Replication of Campbell et al. (2010)

- I replicate some of the key results from *Global Currency Hedging* (Campbell, Serfaty-de Medeiros, and Viceira, 2010) using more recent data.
- *Table 1:* A domestic investor chooses *one foreign currency at a time* to minimize portfolio return variance. Reported currency positions are the amounts of dollars invested in foreign currency per dollar in the portfolio.
- *Table 2:* The investor chooses a *vector of positions across all available currencies* simultaneously. Results reflect optimal variance-minimizing exposure per currency.
- In both tables, the investor holds a domestic equity portfolio and takes foreign currency positions purely for risk reduction.

Replication of Campbell et al. (2010)

Table 1: Optimal Single-Currency Exposure for Single-Country Portfolios

	Currency						
Stock Market	Australia	Canada	Switzerland	Eurozone	UK	Japan	USA
Australia		0.49**	0.61***	0.65***	0.39***	0.49***	0.51***
		(0.21)	(0.13)	(0.19)	(0.13)	(0.05)	(0.10)
Canada	−0.70**		0.60***	0.52***	0.23	0.61***	0.94***
	(0.28)		(0.16)	(0.15)	(0.18)	(0.14)	(0.23)
Switzerland	−0.76***	−0.64***		−0.81***	−0.69***	0.25	−0.20
	(0.11)	(0.14)		(0.30)	(0.16)	(0.19)	(0.19)
Eurozone	−0.76**	0.08	1.11**		−0.02	0.87***	1.10***
	(0.31)	(0.21)	(0.57)		(0.32)	(0.17)	(0.26)
UK	−0.70***	−0.49***	0.33*	0.14		0.34***	0.34
	(0.13)	(0.14)	(0.18)	(0.16)		(0.12)	(0.26)
Japan	−0.70***	−0.92***	−0.50	−0.52*	−0.74***		−0.90***
	(0.13)	(0.16)	(0.35)	(0.31)	(0.17)		(0.25)
USA	−0.67***	−0.90***	−0.02	−0.23	−0.58*	0.41**	
	(0.12)	(0.22)	(0.20)	(0.25)	(0.30)	(0.19)	

Optimal Currency Allocations

- Table 1 shows that optimal demands for foreign currency are large, positive, and statistically significant for two stock markets (rows of the table), those of Australia and Canada.
- Investors in the Australian and Canadian stock markets are keen to hold foreign currency, regardless of the particular currency under consideration, because the Australian and Canadian dollars tend to depreciate against all currencies when their stock markets fall.
- Thus, any foreign currency serves as a hedge against fluctuations in these stock markets.
- The long positions in euros, Swiss francs, or US dollars are particularly large and statistically significant.

Optimal Currency Allocations

- At the opposite extreme, it is optimal for investors in the Swiss and Japanese stock markets to hold economically and statistically large short positions (exposures) in almost all currencies.
- This implies that the Swiss franc and Japanese yen tend to appreciate against (almost) all currencies when the Swiss and Japanese stock markets fall.
- Results are similar for the US dollar.
- The Euroland stock market generates large positive demand for the Swiss franc, Japanese yen and US dollar, and negative or zero demands for all other currencies.

Replication of Campbell et al. (2010)

Table 2: Optimal Multiple-Currencies Exposure for Single-Country Portfolios

	Currency						
Stock Market	Australia	Canada	Switzerland	Eurozone	UK	Japan	USA
Australia	-0.40^{***}	-0.29^{**}	0.04	0.34	-0.31^{**}	0.36^{***}	0.25^{*}
	(0.10)	(0.13)	(0.23)	(0.27)	(0.15)	(0.09)	(0.14)
Canada	-0.77^{***}	-0.28	0.14	0.59^{**}	-0.34^{**}	0.22^{**}	0.45^{***}
	(0.14)	(0.18)	(0.20)	(0.23)	(0.13)	(0.10)	(0.13)
Switzerland	-0.58^{***}	-0.17	0.61^{***}	0.10	-0.26^{*}	0.31^{***}	-0.02
	(0.14)	(0.18)	(0.20)	(0.21)	(0.15)	(0.12)	(0.17)
Eurozone	-0.68^{***}	-0.24	0.08	-0.18	-0.52^{***}	0.30^{**}	1.25^{***}
	(0.15)	(0.19)	(0.36)	(0.40)	(0.19)	(0.14)	(0.16)
UK	-0.76^{***}	-0.27^{**}	0.12	0.31	0.15	0.30^{***}	0.15
	(0.12)	(0.13)	(0.21)	(0.24)	(0.12)	(0.09)	(0.14)
Japan	-0.40^{**}	-0.61^{**}	0.15	0.73^{*}	-0.55^{**}	0.81^{***}	-0.14
	(0.19)	(0.28)	(0.34)	(0.39)	(0.26)	(0.31)	(0.31)
USA	-0.72^{***}	-0.33^{**}	0.30^{*}	0.42^{**}	-0.35^{**}	0.32^{***}	0.36^{**}
	(0.15)	(0.13)	(0.17)	(0.21)	(0.14)	(0.12)	(0.16)

Optimal Currency Allocations

- When single-country stock market investors consider investing in all currencies simultaneously, they almost always choose positive exposures to the Swiss franc, Euro, Japanese yen and US dollar, and negative exposures to the Australian dollar, Canadian dollar, and British pound.

- Relative to Table 1, the optimal currency demands are generally larger and statistically more significant for the US dollar, and less statistically significant for the euro and the Swiss franc.

Optimal Currency Allocations

- ▷ This reflects two features of the multiple-currency analysis:
  - First: A position that is long the US dollar and short the Canadian dollar is an effective hedge against stock market declines.
  - Thus, allowing investors to use both North American currencies increases the risk management demand for the US dollar.
  - Second: The euro and Swiss franc are both good hedges but they are highly correlated; thus, the demand for each currency is less precisely estimated when investors are allowed to take positions in both currencies.
  - In this sense the euro and the Swiss franc are substitutes for one another.
- ▷ Note: Risk-minimizing currency demands for internationally diversified bond market investors are generally very small and not statistically significantly different from zero $\Rightarrow$ full hedging is optimal.

Global Currency Hedging with Ambiguity

Ulrych & Vasiljević (2025), Journal of Banking & Finance

Link: https://doi.org/10.1016/j.jbankfin.2024.107366

- We study optimal currency exposure for an international investor who is both *risk-* and *ambiguity-averse*.
- Methodologically, we show that ambiguity leads to a *generalized ridge regression* solution for currency hedging.
- Ambiguity aversion strengthens hedging demand and improves the *stability* of out-of-sample allocations.
- Empirically, we find that accounting for ambiguity reduces portfolio *volatility net of transaction costs* and narrows the confidence intervals for optimal currency positions.

Optimal Currency Exposure with Ambiguity

▷ *Model uncertainty*: The situation in which an investor is *uncertain* about the *true* probabilistic *model* governing the occurrence of different states.

▷ Notation:

- λ: positive coefficient representing risk aversion,
- θ: positive coefficient representing ambiguity aversion,
- $\mathbb{Q}$: particular model (probability measure) from the set of possible models $\mathcal{Q}$,
- μ: agent's prior probability on the space Δ of possible models from $\mathcal{Q}$,
- $\bar{\mathbb{Q}}$: reduced probability $\int_{\Delta} \mathbb{Q}\, d\mu(\mathbb{Q})$ induced by the prior μ.

▷ A risk-and-ambiguity-averse international investor *maximizes* the robust mean-variance *utility*

$$\max_{\boldsymbol{\Psi}_t} U(R^h_{\mathcal{P},t+1}) = \max_{\boldsymbol{\Psi}_t} \left\{ \mathrm{E}_{\bar{\mathbb{Q}}}[R^h_{\mathcal{P},t+1}] - \frac{\lambda}{2} \mathrm{Var}_{\bar{\mathbb{Q}}}(R^h_{\mathcal{P},t+1}) - \frac{\theta}{2} \mathrm{Var}_{\mu}(\mathrm{E}_{\mathbb{Q}}[R^h_{\mathcal{P},t+1}]) \right\}.$$

Optimal Currency Exposure with Ambiguity

- The argument $\boldsymbol{\Psi}_t^*$ that maximizes the above expression is the *optimal currency exposure* for a *risk-and-ambiguity-averse* international investor

$$\boldsymbol{\Psi}_t^* = -\left[\lambda \mathrm{Var}_{\bar{\mathbb{Q}}}(\mathbf{e}_{t+1} - \mathbf{f}_t) + \theta \mathrm{Var}_\mu \left(\mathrm{E}_{\mathbb{Q}}[\mathbf{e}_{t+1} - \mathbf{f}_t]\right)\right]^{-1} \cdot \left[\lambda \mathrm{Cov}_{\bar{\mathbb{Q}}}(R^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t) + \theta \mathrm{Cov}_\mu(\mathrm{E}_{\mathbb{Q}}[R^{fh}_{\mathcal{P},t+1}], \mathrm{E}_{\mathbb{Q}}[\mathbf{e}_{t+1} - \mathbf{f}_t]) - \mathrm{E}_{\bar{\mathbb{Q}}}[\mathbf{e}_{t+1} - \mathbf{f}_t]\right].$$

- Special case I: *Infinite risk aversion* (i.e., $\lambda \to \infty$)

$$\boldsymbol{\Psi}_{t,risk}^* := \lim_{\lambda\to\infty} \boldsymbol{\Psi}_t^* = -\mathrm{Var}_{\bar{\mathbb{Q}}}[\mathbf{e}_{t+1} - \mathbf{f}_t]^{-1} \cdot \mathrm{Cov}_{\bar{\mathbb{Q}}}[R^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t].$$

- Special case II: *Infinite ambiguity aversion* (i.e., $\theta \to \infty$)

$$\boldsymbol{\Psi}_{t,amb}^* := \lim_{\theta\to\infty} \boldsymbol{\Psi}_t^* = -\mathrm{Var}_\mu[\mathrm{E}_{\mathbb{Q}}[\mathbf{e}_{t+1} - \mathbf{f}_t]]^{-1} \cdot \mathrm{Cov}_\mu[\mathrm{E}_{\mathbb{Q}}[R^{fh}_{\mathcal{P},t+1}], \mathrm{E}_{\mathbb{Q}}[\mathbf{e}_{t+1} - \mathbf{f}_t]].$$

Optimal Currency Exposure with Ambiguity

▷ Special case III: *Risky assets and ambiguous currencies* (domestic cash is risk-free, fully hedged assets are purely risky, and foreign assets and currencies are ambiguous)

$$\boldsymbol{\Psi}^*_{t,fh} = -\Big(\mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,] + \frac{\theta}{\lambda}\mathrm{Var}_{\mu}[\,\mathrm{E}_{\mathbb{Q}}[\mathbf{e}_{t+1} - \mathbf{f}_t]\,]\Big)^{-1}$$
$$\cdot\Big(\mathrm{Cov}_{\bar{\mathbb{Q}}}[\,R^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t\,] - \frac{1}{\lambda}\mathrm{E}_{\bar{\mathbb{Q}}}[\mathbf{e}_{t+1} - \mathbf{f}_t]\Big).$$

- Observe:

  When $\lambda \to \infty$, the optimal currency exposure converges to $\boldsymbol{\Psi}^*_{t,risk}$.

  When $\theta \to \infty$, the optimal currency exposure converges to zero (i.e., *full hedging*).

Global Currency Hedging with Ambiguity: Empirical Analysis

- Empirical study of the *impact* of *risk* and *ambiguity* on the *optimal currency exposure*. Focus on analyzing the (ambiguity-adjusted) *hedging demand* (i.e., no speculative demand) $\Longrightarrow$ *ordinary* ridge regression.
- The *empirical analysis* employs the *data* of: spot and forward currency exchange rates, equity broad market indices, and fixed income total return indices.
- The data is obtained from Refinitiv Datastream and covers the period from January 1999 (the euro's inception) to December 2019 for seven *developed economies*: Australia, Canada, Switzerland, Eurozone, United Kingdom, Japan, and United States.
- We assume a diversified portfolio consisting of equities and government bonds, with a ratio of three-to-one in favor of equities, equally weighted across the seven economies.

In-Sample Analysis: The Impact of Ambiguity Aversion

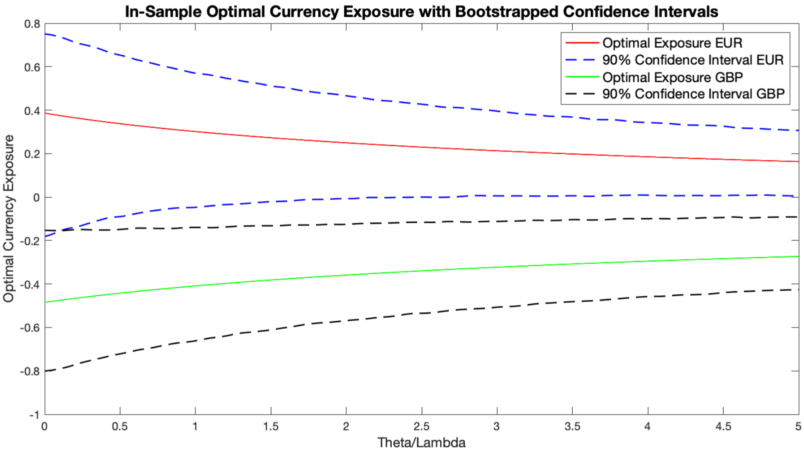

Figure 1: Optimal hedging demand in EUR and GBP for a USD-based investor and the corresponding 90% confidence intervals estimated with the non-parametric bootstrap, as a function of θ/λ.

In-Sample Analysis: The Impact of Ambiguity Aversion

Optimal In-Sample Currency Exposure

Base Currency	**Model**	AUD	CAD	CHF	EUR	GBP	JPY	USD
CHF	RO	-0.47^{***}	-0.30^{*}	0.22	0.27	-0.33^{**}	0.26^{***}	0.35^{**}
		(0.11)	(0.15)	(0.20)	(0.22)	(0.15)	(0.08)	(0.13)
	AA	-0.45^{***}	-0.25^{**}	0.26	0.17	-0.26^{**}	0.27^{***}	0.27^{**}
		(0.09)	(0.11)	(0.18)	(0.16)	(0.11)	(0.08)	(0.11)
EUR	RO	-0.47^{***}	-0.29^{*}	0.21	0.28	-0.34^{**}	0.26^{***}	0.35^{**}
		(0.11)	(0.15)	(0.20)	(0.22)	(0.15)	(0.08)	(0.13)
	AA	-0.46^{***}	-0.25^{**}	0.17	0.27	-0.28^{**}	0.27^{***}	0.28^{***}
		(0.09)	(0.11)	(0.15)	(0.18)	(0.12)	(0.07)	(0.10)
USD	RO	-0.47^{***}	-0.32^{**}	0.27	0.25	-0.34^{**}	0.25^{***}	0.35^{**}
		(0.11)	(0.15)	(0.20)	(0.21)	(0.14)	(0.08)	(0.13)
	AA	-0.44^{***}	-0.29^{**}	0.24^{*}	0.20	-0.28^{**}	0.25^{***}	0.33^{***}
		(0.10)	(0.12)	(0.14)	(0.14)	(0.11)	(0.08)	(0.12)

Table 1: The estimated optimal hedging demand and the corresponding standard deviations are reported for risk-only (RO) and ambiguity-adjusted (AA) investors with $\theta = 0$ and $\theta = \lambda$.

Global Currency Hedging with Ambiguity: Out-of-Sample Backtest

- Analyze the *out-of-sample performance* of model based hedging (RO and AA hedge) in comparison to the naive *benchmarks* (zero, half, and full hedge).
- Currency hedging is implemented using forward contracts which are rolled over *quarterly*.
- We assume that *asset positions* are also *rebalanced* quarterly such that initial portfolio weights are preserved over the backtest.
- The *constrained* strategies RO-C and AA-C restrict the currency exposures to the interval $[-\frac{2}{7}, \frac{3}{7}]$ ensuring the *symmetric* treatment of over- and under-hedging.
- All results are presented *net of transaction* (hedging) *costs*, which are assumed to be 10 basis points relative to the notional of the entered currency forward contract.

Global Currency Hedging with Ambiguity: Out-of-Sample Backtest

	Naive Hedge			Model Hedge			
Panel A: CHF	Zero	Half	Full	RO	RO-C	AA	AA-C
Average (%)	5.31	5.14	4.98	5.20	5.11	4.46	4.51
Volatility (%)	12.87	10.70	9.34	9.09	7.38	7.23	7.19
Sharpe Ratio	0.41	0.48	0.53	0.57	0.69	0.62	0.63
Sortino Ratio	0.56	0.66	0.73	0.81	0.98	0.86	0.87
Turnover	0	0.43	0.86	2.36	1.82	1.33	1.33
Panel B: EUR	Zero	Half	Full	RO	RO-Con	AA	AA-Con
Average (%)	7.21	6.39	5.56	6.20	6.39	5.33	5.37
Volatility (%)	10.10	9.41	9.35	9.00	7.35	7.17	7.17
Sharpe Ratio	0.71	0.68	0.59	0.69	0.87	0.74	0.75
Sortino Ratio	0.99	0.93	0.81	1.00	1.26	1.04	1.04
Turnover	0	0.43	0.86	2.42	1.79	1.21	1.21
Panel C: USD	Zero	Half	Full	RO	RO-Con	AA	AA-Con
Average (%)	6.75	6.67	6.59	7.18	7.04	6.29	6.29
Volatility (%)	12.82	10.78	9.34	9.20	7.27	7.04	7.04
Sharpe Ratio	0.53	0.62	0.71	0.78	0.97	0.89	0.89
Sortino Ratio	0.74	0.86	0.97	1.14	1.39	1.26	1.25
Turnover	0	0.43	0.86	2.56	1.82	1.26	1.26

Global Currency Hedging with Ambiguity: Out-of-Sample Backtest

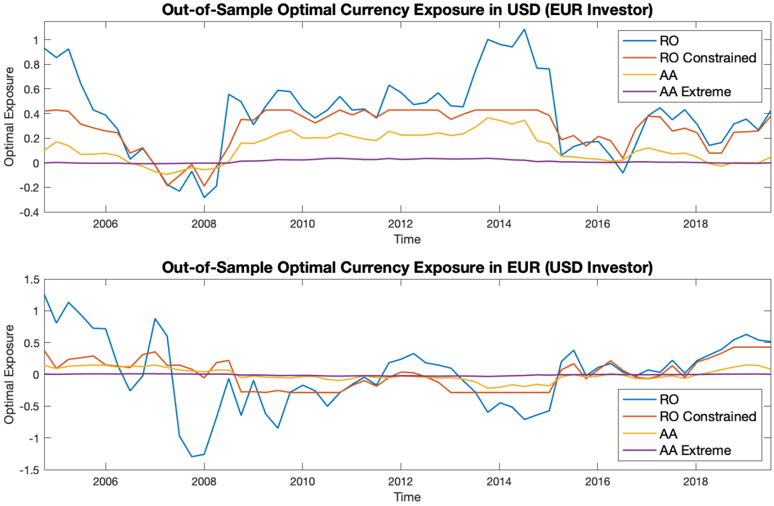

Figure 2: Temporal dynamics of the optimal currency exposures for the model-based strategies, including an additional AA-based strategy called 'AA Extreme' with $\theta/\lambda = 40$.

Global Currency Hedging with Ambiguity: Conclusion

- *Closed-form* expressions of the *optimal currency exposure* for a *risk-and-ambiguity-averse* robust mean-variance investor are derived.
- The optimal currency exposure capturing investor's dislike for *risk* and *ambiguity* is characterized by a *generalized ridge regression* $\Longrightarrow$ the *penalty term* corresponds to the *utility loss* arising from *model uncertainty* and induces *shrinkage*.
- Elevated *ambiguity aversion* can lead to *insufficient currency diversification* of global investors.
- Accounting for *ambiguity*
  - corresponds to an *increase* in *bias* and a simultaneous *shrinkage* of the *confidence intervals* of the sample efficient optimal currency exposure;
  - *enhances* the *stability* of the estimated optimal currency exposures and significantly *reduces* the portfolio *volatility* in- and out-of-sample.

Dynamic Currency Hedging with Non-Gaussianity and Ambiguity

Polak & Ulrych (2024), Quantitative Finance

https://doi.org/10.1080/14697688.2023.2301419

- We propose a *dynamic currency hedging strategy* under *non-Gaussian return dynamics* and investor-specific *ambiguity*.
- Ambiguity is *parametrized* from market data using a *continuous mean-variance mixture* representation.
- We provide a *closed-form* solution for mean-variance hedging and a *numerical algorithm* for optimizing arbitrary risk measures (e.g., expected shortfall) via *generalized filtered historical simulation*.
- An out-of-sample *backtest* shows the proposed strategy is *robust, risk-reductive*, and *outperforms* both constant and Gaussian benchmarks *net of transaction costs*.

Joint Optimization of Assets and Currencies

Sparse and Stable International Portfolio Optimization

Burkhardt & Ulrych (2023), Journal of International Money and Finance
https://www.sciencedirect.com/science/article/pii/S026156062300150X

- We propose a *joint optimization framework* for international asset and currency allocation with *regularization* to improve out-of-sample performance.
- Regularized optimization yields *sparse and stable portfolio weights*, reducing sensitivity to parameter uncertainty.
- The *joint* optimization approach outperforms standard *currency overlay* strategies and *equally-weighted* fully hedged portfolios, *net of transaction costs*.
- The paper challenges the *industry practice* of separate hedging (i.e., currency overlay), showing that *joint optimization* leads to improved risk-adjusted performance.

International Portfolio Optimization: Hedged Portfolio Return

- Return in domestic currency:

$$\tilde{R}^u_{i,t+1} = \frac{P_{i,t+1} S_{c_i,t+1}}{P_{i,t} S_{c_i,t}} - 1 = R_{i,t+1} + e_{c_i,t+1} + R_{i,t+1} e_{c_i,t+1}.$$

- Portfolio return:

$$\tilde{R}^u_{\mathcal{P},t+1} = \sum_{i=1}^{N} x_{i,t} \tilde{R}^u_{i,t+1}.$$

- Hedged portfolio return:

$$\tilde{R}^h_{\mathcal{P},t+1} = \tilde{R}^u_{\mathcal{P},t+1} + \sum_{c=2}^{M+1} \phi_{c,t} (f_{c,t} - e_{c,t+1}),$$

with

$$f_{c,t} = \frac{(F_{c,t} - S_{c,t})}{S_{c,t}}.$$

International Portfolio Optimization Problem

- Separate optimization problem:

$$\begin{aligned}
\boldsymbol{x}_t^* &= \arg\max_{\boldsymbol{x}_t} \left\{ \mathbb{E}(\tilde{R}^{fh}_{\mathcal{P},t+1}) - \frac{\gamma}{2}\mathrm{Var}(\tilde{R}^{fh}_{\mathcal{P},t+1}) \right\} \\
&= \arg\max_{\boldsymbol{x}_t} \left\{ \boldsymbol{x}_t^T \boldsymbol{\mu}^a - \frac{\gamma}{2}\boldsymbol{x}_t^T \boldsymbol{\Sigma}^a \boldsymbol{x}_t \right\} \\
&\quad \text{subject to} \qquad \boldsymbol{x}_t^T \mathbf{1}_N = 1,
\end{aligned}$$

$$\begin{aligned}
\phi_t^* \mid \boldsymbol{x}_t^* &= \arg\max_{\phi_t} \left\{ \mathbb{E}(\tilde{R}^{h}_{\mathcal{P},t+1}) - \frac{\gamma}{2}\mathrm{Var}(\tilde{R}^{h}_{\mathcal{P},t+1}) \mid \boldsymbol{x}_t^* \right\} \\
&= \arg\max_{\phi_t} \left\{ \phi_t^T \boldsymbol{\mu}^c - \frac{\gamma}{2}\phi_t^T \boldsymbol{\Sigma}^c \phi_t - \gamma \boldsymbol{x}_t^{*T} \boldsymbol{\Sigma}^{ac} \phi_t \ \mid \boldsymbol{x}_t^* \right\}.
\end{aligned}$$

- Joint optimization problem:

$$\boldsymbol{\theta}_t := (x_{1,t}, \ldots, x_{N,t}, \phi_{2,t}, \ldots, \phi_{M+1,t}),$$

$$\boldsymbol{r}_{t+1} := (\tilde{R}^u_{1,t+1}, \ldots, \tilde{R}^u_{N,t+1}, (f_{2,t} - e_{2,t+1}), \ldots, (f_{M+1,t} - e_{M+1,t+1})),$$

$$\begin{aligned}
\boldsymbol{\theta}^*_t &= \arg\max_{\boldsymbol{\theta}_t} \left\{ \mathbb{E}(\tilde{R}^h_{\mathcal{P},t+1}) - \frac{\gamma}{2}\mathrm{Var}(\tilde{R}^h_{\mathcal{P},t+1}) \right\} \\
&= \arg\max_{\boldsymbol{\theta}_t} \left\{ \boldsymbol{\theta}_t^T \boldsymbol{\mu} - \frac{\gamma}{2}\boldsymbol{\theta}_t^T \boldsymbol{\Sigma} \boldsymbol{\theta}_t \right\} \\
&\quad \text{subject to} \qquad \boldsymbol{\theta}_t^T \mathbf{q}_{N,M} = 1,
\end{aligned}$$

with

$$\boldsymbol{\mu} = \mathbb{E}(\boldsymbol{r}_{t+1}) \quad \text{and} \quad \boldsymbol{\Sigma} = \mathrm{Var}(\boldsymbol{r}_{t+1}),$$

and where $\mathbf{q}_{N,M}$ denotes an $(N+M) \times 1$ vector with the first N elements equal to one and the rest M elements equal to zero.

Sparse Multi-Currency Portfolio

- Separate optimization:

$$\boldsymbol{x}_t^* = \arg\min_{\boldsymbol{x}_t} \left\{ \frac{\gamma}{2}\boldsymbol{x}_t^T \hat{\boldsymbol{\Sigma}}^a \boldsymbol{x}_t - \boldsymbol{x}_t^T \hat{\boldsymbol{\mu}}^a + \lambda_1^a \|\boldsymbol{x}_t\|_1 \right\}$$
$$\text{subject to} \qquad \boldsymbol{x}_t^T \mathbf{1}_N = 1,$$

$$\phi_t^* \mid \boldsymbol{x}_t^* = \arg\min_{\phi_t} \left\{ \frac{\gamma}{2}\phi_t^T \hat{\boldsymbol{\Sigma}}^c \phi_t - \phi_t^T \hat{\boldsymbol{\mu}}^c + \gamma \boldsymbol{x}_t^{*T} \hat{\boldsymbol{\Sigma}}^{ac} \phi_t + \lambda_1^c \|\phi_t\|_1 \mid \boldsymbol{x}_t^* \right\}.$$

- Joint optimization:

$$\boldsymbol{\theta}_t^* = \arg\min_{\boldsymbol{\theta}_t} \left\{ \frac{\gamma}{2}\boldsymbol{\theta}_t^T \hat{\boldsymbol{\Sigma}} \boldsymbol{\theta}_t - \boldsymbol{\theta}_t^T \hat{\boldsymbol{\mu}} + \|\boldsymbol{S}\boldsymbol{\theta}_t\|_1 \right\}$$
$$\text{subject to} \qquad \boldsymbol{\theta}_t^T \mathbf{q}_{N,M} = 1.$$

- Denoted by NC1 in the empirical analysis.

Stable Multi-Currency Portfolio

- *Shrinkage* estimator for the covariance matrix:

$$\hat{\mathbf{\Sigma}}_s = v\hat{\mathbf{\Sigma}}_g + (1 - v)\hat{\mathbf{\Sigma}}.$$

- Corresponding joint optimization problem:

$$\boldsymbol{\theta}_t^* = \arg\min_{\boldsymbol{\theta}_t} \left\{ \frac{\gamma}{2}\boldsymbol{\theta}_t^T \hat{\mathbf{\Sigma}}_s \boldsymbol{\theta}_t - \boldsymbol{\theta}_t^T \hat{\boldsymbol{\mu}} \right\},$$
$$\text{subject to} \qquad \boldsymbol{\theta}_t^T \mathbf{q}_{N,M} = 1,$$

  where in the empirical analysis, $\hat{\mathbf{\Sigma}}_g$ is set to a constant correlation matrix, denoted by CC, as introduced by Ledoit and Wolf (2004).

- *Constraining* the *$\mathcal{L}^2$-norm* is equivalent to shrinking the sample covariance matrix towards the identity matrix. In the empirical analysis, this is denoted by NC2.

- The *sparse and stable* portfolio is determined by combining both sparsity (i.e., $\mathcal{L}^1$-regularization) and stability (i.e., covariance shrinkage or $\mathcal{L}^2$-regularization).

International Portfolio Optimization: Empirical Analysis

- Daily data ranging from 1999 to 2019 and encompassing 21 equity broad market indices denoted in 10 currencies.
- Spot and forward exchange rate data for 10 currencies: Australian dollar, Canadian dollar, Swiss franc, euro, British pound, Japanese yen, US dollar, Norwegian krone, New Zealand dollar, and Singapore dollar.
- The investor rebalances her portfolio every *month* and hedging is only possible one period ahead.
- All returns are computed *net of transaction costs* – 20 bp for assets and 2 bp for currency forwards.
- Parameters are estimated on a 12-month rolling estimation window and the penalization parameters are determined via *cross-validation*.

International Portfolio Optimization: In-Sample Analysis

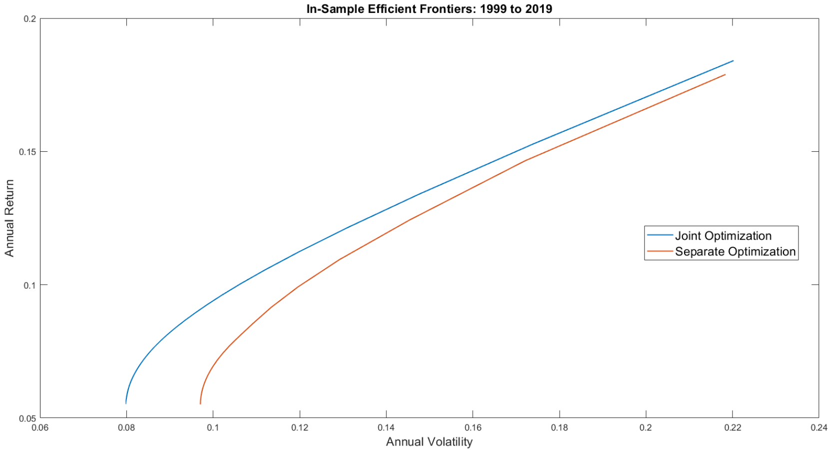

Figure 3: In-sample efficient frontiers for the dataset of 21 equity indices and 10 currencies from 1999 to 2019. The curves were derived by varying the risk aversion parameter γ in both steps of the separate optimization as well as in the joint optimization.

International Portfolio Optimization: Out-of-Sample Analysis

	Sharpe ratio		Sortino ratio		Certainty equivalent		Volatility	
	Sep.	**Joint**	**Sep.**	**Joint**	**Sep.**	**Joint**	**Sep.**	**Joint**
MV	0.3481	0.4625	0.3939	0.5314	0.0217	0.0322	0.2167	0.2701
MV-CB	0.2488	0.4119	0.2779	0.4919	0.0091	0.0390	0.1924	0.2005
NC1	0.4586	0.5035	0.5474	0.5702	0.0489	0.0580	0.1091	0.1425
NC1-CB	0.4496	0.5048	0.5193	0.6020	0.0473	0.0552	0.1046	0.1170
NC2	0.5107	0.5600	0.6142	0.6620	0.0582	0.0641	0.1343	0.1296
NC2-CB	0.5198	0.5171	0.6319	0.6246	0.0583	0.0573	0.1255	0.1212
NC1-CC	0.3699	0.4542	0.4307	0.5092	0.0394	0.0510	0.1150	0.1443
NC1-CC-CB	0.3880	0.5168	0.4397	0.6125	0.0413	0.0569	0.1110	0.1182
NC1-NC2	0.4962	0.6172	0.6305	0.7646	0.0488	0.0659	0.0883	0.1082
NC1-NC2-CB	0.5532	**0.6810**	0.6882	**0.8314**	0.0538	**0.0715**	**0.0881**	0.1046
1/N	0.3413		0.4071		0.0357		0.1302	

International Portfolio Optimization: Out-of-Sample Analysis

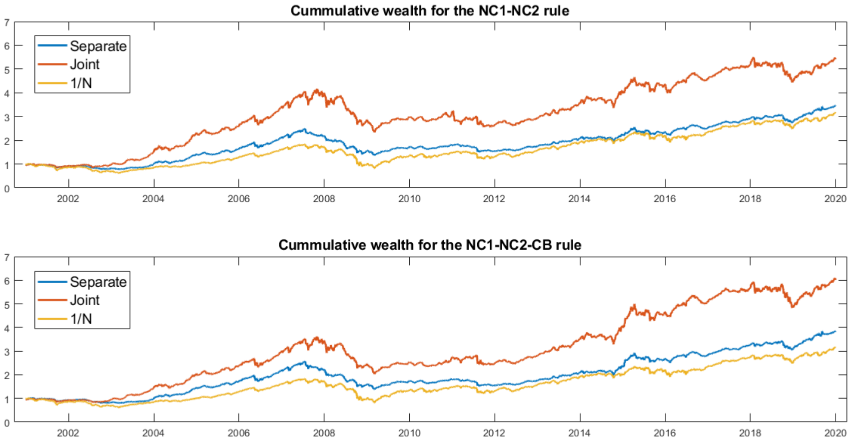

Figure 4: This figure portrays the cumulative wealth evolution for the two sparse and stable portfolio rules, i.e., the NC1-NC2 portfolio (top) and the NC1-NC2-CB portfolio (bottom), with an initial wealth of $1. All reported values are adjusted for transaction costs.

International Portfolio Optimization: Conclusion

- This paper proposes a *sparse and stable international asset allocation framework* where asset and currency weights are determined in a regularized fashion $\Longrightarrow$ *contribution* to the literature on multi-currency asset allocation.

- The introduced sparse and stable *joint optimization* outperforms the equivalent approaches based on *separate* (i.e., currency overlay) optimization as well as the *equally weighted* fully hedged portfolio *net of transaction costs*.

- Employing *currency overlay* strategies is *suboptimal* not only in sample but also from the *out-of-sample* perspective $\Longrightarrow$ improvement via *regularized joint optimization*.

Regularized Multi-Currency Expected Shortfall Portfolios

Lucescu & Ulrych (2024), Work in Progress

- We develop a regularized optimization framework for *multi-currency portfolio allocation* under *expected shortfall (ES)*. Two approaches are considered: a *separate currency overlay* and a *joint asset-currency optimization*.
- To address instability in ES minimization, we apply *$\mathcal{L}^1$ and $\mathcal{L}^2$ regularization* and a leverage constraint to promote *sparsity and numerical stability*.
- *Out-of-sample backtests* with transaction costs show improved *risk-adjusted returns* and *lower tail risk* compared to unregularized and naive benchmarks.
- *Joint optimization* yields better results in a *small asset universe*, where asset-currency interactions are more effectively captured, while the *overlay* approach performs *better* in a *large universe* due to its *dimensionality reduction*.

Expected Shortfall Optimization

- Optimal *expected shortfall* portfolio is given by

$$\boldsymbol{x}_t^* = \arg\min_{\boldsymbol{x}_t} \mathbb{ES}_{t,\beta}[R_{\mathcal{P},t+1}], \text{ with } \boldsymbol{x}_t^T \mathbf{1}_N = 1. \tag{4}$$

- As shown by Rockafellar and Uryasev (2000), (4) can be recast as:

$$\boldsymbol{x}_t^* = \arg\min_{\boldsymbol{x}_t,\alpha} \alpha + (1-\beta)^{-1}\mathbb{E}_t[(-\boldsymbol{x}_t^T \boldsymbol{R}_{t+1} - \alpha)^+]. \tag{5}$$

- Sampling the historical return distribution, (5) can be *linearized*:

$$\boldsymbol{x}_t^* = \arg\min_{\boldsymbol{x}_t,\alpha} \alpha + \frac{1}{q(1-\beta)} \sum_{k=1}^{q} u_k,$$

subject to the linear constraints

$$\sum_{i=1}^{N} x_i = 1,\ \alpha > 0,\ u_k \geq 0,\ \text{and } \boldsymbol{x}_t^T \boldsymbol{R}_k + \alpha + u_k \geq 0 \text{ for } k = 1, ..., q.$$

Regularized Expected Shortfall Portfolios

- ES minimization is prone to *numerical instability*. This is particularly the case in *high-dimensional problems*.
- $\mathcal{L}^1$-norm penalty:
  - Defined as $\|\boldsymbol{x}\|_1 = \sum_{i=1}^N |x_i|$.
  - *Limits sensitivity to colinearities*. Promotes *sparse solutions* that are practically appealing.
- $\mathcal{L}^2$-norm penalty:
  - Defined as $\|\boldsymbol{x}\|_2^2 = \sum_{i=1}^N x_i^2$.
  - Directly *removes instability* in ES minimization. Related to *shrinkage methods* in covariance estimation.
- Assets leverage constraint:
  - Defined as $\|\boldsymbol{x}_t\|_1 \leq \ell$ for $\ell \geq 1$. Practical *control* over portfolio exposure.

Regularized Expected Shortfall: Optimization Setup

- The *return* of a *portfolio* $\mathcal{P}$ hedged with *currency forwards* is equal to:

$$R^{h}_{\mathcal{P},t+1} = \sum_{i=1}^{N} x_{i,t} R^{fh}_{i,t+1} + \sum_{c=1}^{M} \psi_{c,t}(e_{c,t+1} - f_{c,t}).$$

- Variables at time t:
  - $\boldsymbol{\mu}^{fh}$: Expected fully hedged asset return vector $\mathbb{E}[\boldsymbol{R}^{fh}_{t+1}] \in \mathbb{R}^N$.
  - $\boldsymbol{\mu}^{c}$: Expected excess currency return vector $\mathbb{E}[\boldsymbol{e}_{t+1} - \boldsymbol{f}_t] \in \mathbb{R}^M$.
  - γ: Investor's relative risk aversion parameter, with $\gamma > 0$.
  - β: Confidence level at which the ES is evaluated, with $\beta \in (0, 1)$.
  - ℓ: Maximum leverage allowed for the asset portfolio, with $\ell \geq 1$.

Regularized Expected Shortfall: Separate Optimization

- Asset-level optimization problem:

$$\boldsymbol{x}_t^* = \arg\min_{\boldsymbol{x}_t} \left\{ -\boldsymbol{x}_t^T \boldsymbol{\mu}^{fh} + \gamma \mathbb{ES}_\beta(\boldsymbol{x}_t^T \boldsymbol{R}_{t+1}^{fh}) + \lambda_1^a \|\boldsymbol{x}_t\|_1 + \lambda_2^a \|\boldsymbol{x}_t\|_2^2 \right\},$$
$$\text{subject to} \quad \boldsymbol{x}_t^T \mathbb{1}_N = 1, \quad \|\boldsymbol{x}_t\|_1 \leq \ell.$$

- Second-step currency overlay problem:

$$\boldsymbol{\psi}_t^* \mid \boldsymbol{x}_t^* = \arg\min_{\boldsymbol{\psi}_t} \left\{ -\boldsymbol{\psi}_t^T \boldsymbol{\mu}^c + \gamma \mathbb{ES}_\beta \left(\boldsymbol{x}_t^{*T} \boldsymbol{R}_{t+1}^{fh} + \boldsymbol{\psi}_t^T (\boldsymbol{e}_{t+1} - \boldsymbol{f}_t) \right) + \lambda_1^c \|\boldsymbol{\psi}_t\|_1 + \lambda_2^c \|\boldsymbol{\psi}_t\|_2^2 \mid \boldsymbol{x}_t^* \right\}.$$

- *Separate penalizing terms* for assets (λ_1^a, λ_2^a) and currencies (λ_1^c, λ_2^c) reflecting different *budget constraints* and *transaction costs*.

Regularized Expected Shortfall: Joint Optimization

- Variables at time t:
  - $\boldsymbol{\theta}_t$: Combined decision vector $(\boldsymbol{x}_t, \boldsymbol{\psi}_t)^T \in \mathbb{R}^{N+M}$.
  - $\boldsymbol{R}_{t+1}$: Combined return vector $(\boldsymbol{R}^{fh}_{t+1}, \boldsymbol{e}_{t+1} - \boldsymbol{f}_t)^T \in \mathbb{R}^{N+M}$.
  - $\boldsymbol{\mu}$: Joint expected return vector $\mathbb{E}[\boldsymbol{R}_{t+1}]$.
- Joint optimization:

$$\boldsymbol{\theta}_t^* = \arg\min_{\boldsymbol{\theta}_t} \left\{ -\boldsymbol{\theta}_t^T \boldsymbol{\mu} + \gamma \mathbb{ES}_\beta(\boldsymbol{\theta}_t^T \boldsymbol{R}_{t+1}) + \|\boldsymbol{S}_1 \boldsymbol{\theta}_t\|_1 + \|\boldsymbol{S}_2 \boldsymbol{\theta}_t\|_2^2 \right\},$$

$$\text{subject to} \quad \boldsymbol{\theta}_t^T \boldsymbol{q}_{N,M} = 1, \quad \|\boldsymbol{L}\boldsymbol{\theta}_t\|_1 \le \ell,$$

  where vector $\boldsymbol{q}_{N,M}$ controls the budget constraint, vector $\boldsymbol{L}$ enforces the leverage constraint, and diagonal matrices $\boldsymbol{S}_1$ and $\boldsymbol{S}_2$ control regularization.
- Both optimization problems remain *tractable*, and can be expressed as *quadratic programs with linear constraints*.

Regularized Expected Shortfall: Empirical Analysis

- Daily data ranging from 1990 to 2023, encompassing 14 major economies and individual assets denominated in their corresponding currencies.
- Two investment universes:
  - Small Universe: Largest 10 companies by market capitalization per economy.
  - Large Universe: Largest 50 companies by market capitalization per economy.
- The investor rebalances their portfolio every *quarter* and hedging is only possible one period ahead.
- Empirical analysis is conducted assuming *infinite risk aversion* $\gamma = \infty$, eliminating noisy mean estimates.
- All returns are presented *net of transaction costs*.

Regularized Expected Shortfall: Empirical Analysis

- Parameters are estimated on a 10-year rolling estimation window and the penalization parameters are determined via *temporal validation*.
- We evaluate ES at the 80% confidence level, striking a *balance* between downside risk and out-of-sample performance.
- *Out-of-sample backtest* of the following strategies:
  - **1/N**: The fully hedged equally-weighted portfolio.
  - **OVERLAY**: The unregularized two-step (currency overlay) portfolio.
  - **OVERLAY-REG**: The regularized overlay portfolio with $\mathcal{L}^1$ and $\mathcal{L}^2$ penalties.
  - **JOINT**: The unregularized joint optimization portfolio.
  - **JOINT-REG**: The regularized joint optimization portfolio with $\mathcal{L}^1$ and $\mathcal{L}^2$ penalties.

Regularized Expected Shortfall: Out-of-Sample Analysis

Panel A: Small Universe					
	1/N	**OVERLAY**	**JOINT**	**OVERLAY-REG**	**JOINT-REG**
ES (%)	*18.06*	*10.86*	*10.48*	*10.44*	*10.28*
Volatility (%)	13.76	8.19	7.96	7.93	7.81
Max Draw (%)	53.42	30.01	28.06	28.87	28.55
ES Ratio	0.12	0.07	0.13	0.12	0.17
Sharpe Ratio	*0.15*	*0.10*	*0.17*	*0.15*	*0.22*
Sortino Ratio	0.19	0.12	0.22	0.19	0.28
Turnover[a]	0.24	1.29	1.18	0.96	1.02
Turnover[c]	0.92	2.36	2.97	1.70	1.71

Panel B: Large Universe					
	1/N	**OVERLAY**	**JOINT**	**OVERLAY-REG**	**JOINT-REG**
ES (%)	*16.34*	*7.85*	*7.81*	*6.76*	*6.69*
Volatility (%)	12.39	6.16	6.09	5.90	5.78
Max Draw (%)	54.86	39.29	39.56	42.53	41.02
ES Ratio	0.25	0.32	0.31	0.53	0.51
Sharpe Ratio	*0.33*	*0.41*	*0.40*	*0.61*	*0.59*
Sortino Ratio	0.40	0.51	0.51	0.72	0.71
Turnover[a]	0.24	3.36	3.28	1.45	1.42
Turnover[c]	0.92	0.95	1.75	1.25	1.19

Regularized Expected Shortfall: Cumulative Returns

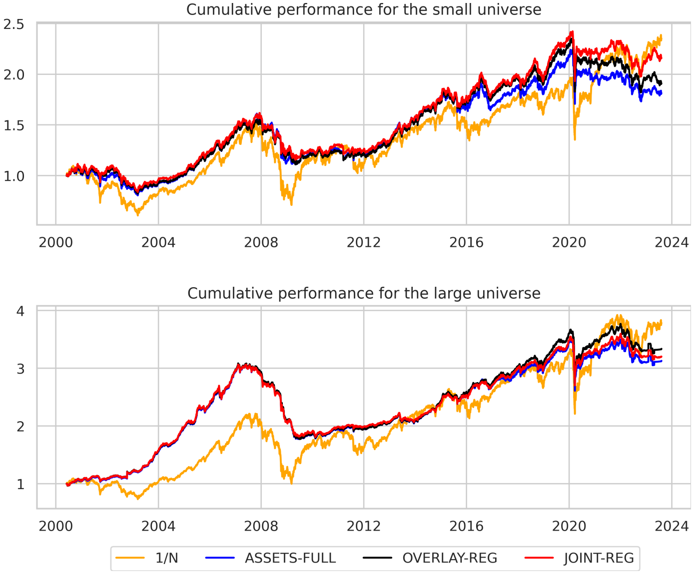

Regularized Expected Shortfall: Rolling ES

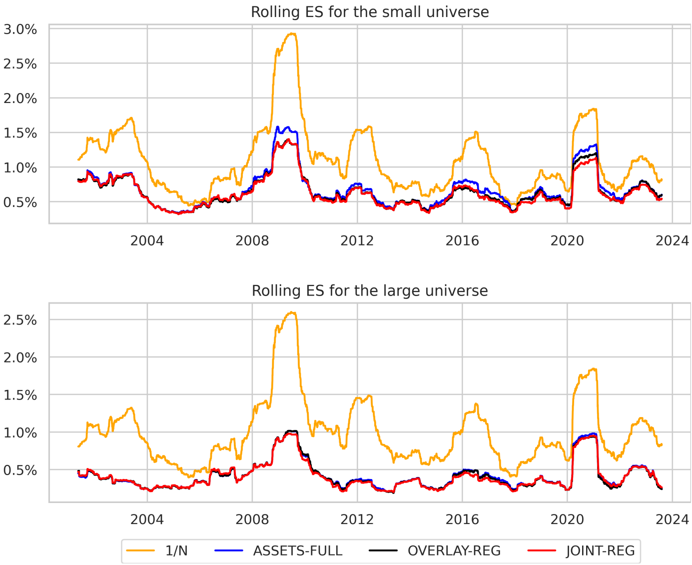

Regularized Expected Shortfall: Conclusion

- This paper addresses known challenges of downside risk optimization in *high-dimensional international setting*.
- We propose a *regularized optimization framework* for multi-currency portfolio allocation under *expected shortfall*.
- *Regularized portfolios* outperform *unregularized* and *naïve* benchmarks in terms of realized risk and risk-adjusted returns *net of transaction costs*.
- The *joint optimization* delivers superior performance when the investment universe is relatively *small*, while the *separate overlay* approach proves more effective in larger universes due to its *dimensionality reduction*.

Currency Risk Management: Summary and Implications

Key Takeaways on Currency Risk

- *Currency risk* is a major component of international investments, and effective management is crucial for *portfolio stability*.
- Currency exposure is not only a source of risk — it can also be a *strategic return driver* when appropriately managed.
- Traditional approaches often separate asset and currency decisions (*currency overlay*), but this may ignore important interactions.
- *Hedging strategies* based on return variance or expected shortfall require *robust modeling* to avoid instability and poor out-of-sample performance.

Insights from My Research

- *Joint optimization* of assets and currencies generally outperforms separate overlay approaches in *smaller asset universes*, where asset-currency dependencies can be estimated more *reliably*.
- *Currency overlay* performs better in *larger universes*, where its *dimensionality reduction* makes it more robust and scalable.
- *Regularization techniques* (sparsity, shrinkage) improve numerical stability and enhance out-of-sample risk-adjusted returns.
- *Ambiguity-aware models* can induce economically-informed shrinkage, stabilizing currency exposure and reducing estimation error.

Practical Implications

- *Integrating asset and currency decisions* within a single optimization framework can streamline portfolio construction and reduce operational complexity.
- Portfolio managers should carefully choose between *overlay* vs *joint optimization* depending on the portfolio size and performance objective.
- Accounting for downside risk and *transaction costs* ensures portfolio strategies remain *relevant and feasible* under real-world conditions.
- Combining risk modeling with statistical tools (e.g., regularization, robust estimation) can help *bridge theory* and *practice*.

Open Questions and Future Work

- How can currency strategies adapt to *regime shifts*, structural breaks, or geopolitical shocks?
- Can *machine learning* improve forecasts of risk premia or hedging performance without sacrificing *interpretability*?
- How can *robustness to model uncertainty* improve portfolio stability under changing market conditions?

Thank you for your attention!

References

U. Ulrych and N. Vasiljević. Global Currency Hedging with Ambiguity. *Journal of Banking & Finance*, 172(107366), 2025.
Available at: https://doi.org/10.1016/j.jbankfin.2024.107366

P. Polak and U. Ulrych. Dynamic Currency Hedging with Non-Gaussianity and Ambiguity. *Quantitative Finance*, 24(2): 305-327, 2024.
Available at: https://doi.org/10.1080/14697688.2023.2301419

R. Burkhardt and U. Ulrych. Sparse and Stable International Portfolio Optimization and Currency Risk Management. *Journal of International Money and Finance*, 139(102949), 2023.
Available at:
https://www.sciencedirect.com/science/article/pii/S026156062300150X

References

Global Currency Hedging, John Y. Campbell, Karine Serfaty-de Medeiros, and Luis M. Viceira. Journal of Finance, 65(1):87–121, 2010.

Alpha as ambiguity: Robust mean–variance portfolio analysis, F. Maccheroni, M. Marinacci, and D. Ruffino. Econometrica, 81(3):1075–1113, 2013.

An Introduction to Statistical Learning, Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. Springer, 2013.

A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms, V. DeMiguel, L. Garlappi, F. J. Nogales, and R. Uppal. Management science, 55(5):798–812, 2009.

References

Optimization of conditional value-at-risk, R. T. Rockafellar and S. Uryasev. Journal of risk, 2:21–42, 2000.

Global currency hedging with common risk factors, W. Opie and S. J. Riddiough. Journal of Financial Economics, 136(3):780–805, 2020.

Honey, I Shrunk the Sample Covariance Matrix, Olivier Ledoit and Michael Wolf. The Journal of Portfolio Management, 30(4):110–119, 2004.

Carry trade and momentum in currency markets, C. Burnside, M. Eichenbaum, and S. Rebelo. Annual Review of Financial Economics, 3(1):511–535, 2011.

Hedging with an edge: parametric currency overlay, P. Barroso, J. A. Reichenecker, and M. J. Menichetti. Management Science, 68(1):669–689, 2022.

Machine Learning: The Future of Financial Risk Management?

The Big Data Opportunity

- New technologies have made available vast quantities of digital data.
  - 90% of all data in existence has been created in the past two years!
- There is a paradigm shift from machine programming to machine learning.
  - In conventional programming, tell computer what to do, breaking big problems into many small, precisely defined tasks.
  - Learn (estimate) from observational data, instead of requiring pre-specified logic, for decision making and problem solving.

Can Machines Learn Finance?

- Machines capable of many things (playing chess, speech recognition, translation, self-driving cars etc.), but can they learn finance?
- Finance is different!
  - Low signal-to-noise ratio.
  - Non-stationary, evolving markets.
  - Competition.
- Market efficiency: Returns must be dominated by news in well-functioning markets. Low signal-to-noise ratio is not a coincidence. Market efficiency reinforces it!
- Despite the huge hype around machines and bots in asset and risk management, understanding the true value of machine learning in finance is still in its early stages.

What Is Machine Learning?

- Goal: Find a model that is flexible enough to accommodate important patterns but not so flexible that it overspecializes to specific data set.
  - Most of the modern methods concern with high dimensional models: N observations, P parameters, and $N \approx P$, or $N < P$.

- Supervised learning
  - Want to predict a target variable Y with input variables X.
  - Also called: "predictive analytics".

- Unsupervised learning
  - Want to find structure within a set of variables X.
  - Also called: "exploratory data analysis", "fancy descriptive statistics".

What Is Machine Learning?

The definition of "machine learning" is inchoate and is often context specific. We use the term to describe:

i) A diverse collection of high-dimensional models for statistical prediction, combined with
ii) Regularization methods for model selection and mitigation of overfit, and
iii) Efficient algorithms for searching among a vast number of potential model specifications.

Statistical Learning vs Machine Learning

- Machine learning arose as a subfield of Artificial Intelligence.
- Statistical learning arose as a subfield of Statistics.
- There is much overlap — both fields focus on supervised and unsupervised problems:
  - Machine learning has a greater emphasis on distribution free approaches, large scale applications and prediction accuracy.
  - Statistical learning emphasizes models and their interpretability, and precision and uncertainty.
- But the distinction has become more and more blurred, and there is a great deal of "cross-fertilization".
- Machine learning has the upper hand in Marketing!

Predictive Modelling

- We generically refer to a *response* or *target* that we wish to predict as Y. *Features*, or *inputs*, or *predictors* are named as X_1, X_2, etc.
- We can refer to the *input vector* collectively as

$$X = \begin{bmatrix} X_1 \\ X_2 \\ X_3 \end{bmatrix}.$$

- Predictive relationship is characterized by a model that links a target with input:

$$Y = f(X) + \epsilon,$$

  where ϵ captures measurement errors and other discrepancies.
- Note: $f(X)$ represents the *signal* and ϵ the *noise*.

What Is $f(X)$ Good For?

Prediction vs Inference:

- With a good f we can make *predictions* of Y at new points $X = x$.
- We can understand which components of $X = (X_1, X_2, \ldots, X_p)$ are *important* in *explaining* Y , and which are irrelevant, e.g. Seniority and Years of Education have a big impact on Income, but Marital Status typically does not.
- Depending on the complexity of f, we may be able to *understand* how each component X_j of X *affects* Y.

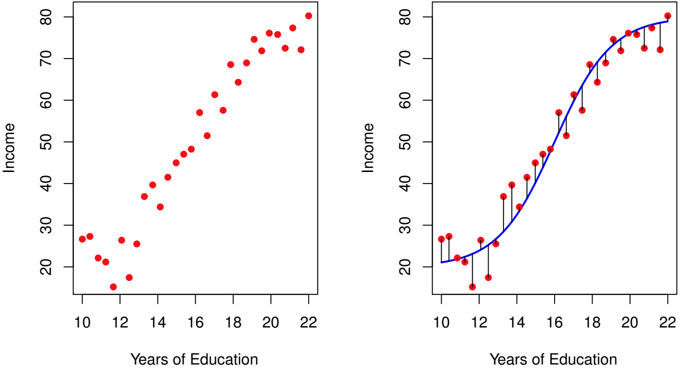

What is a good value for $f(X)$ at any selected value of X, say $X = 18$? There can be many possible Y values at $X = 18$. A good value is

$$f(18) = \mathbb{E}(Y|X = 18),$$

which means expected value (average) of Y given $X = 18$. This ideal $f(x) = \mathbb{E}(Y|X = x)$ is called the *regression function*.

The Regression Function $f(x)$

- It is also defined for vector X; e.g. $f(x) = f(x_1, x_2, x_3) = E(Y \mid X_1 = x_1, X_2 = x_2, X_3 = x_3)$
- It is the *optimal* predictor of Y with regard to mean-squared prediction error: $f(x) = \mathbb{E}(Y \mid X = x)$ is the function that minimizes $\mathbb{E}[(Y - g(X))^2 \mid X = x]$ over all functions g at all points $X = x$.
- $\epsilon = Y - f(x)$ is the *irreducible* error — i.e. even if we knew $f(x)$, we would still make errors in prediction, since at each $X = x$ there is typically a distribution of possible Y values.
- For any estimate $\hat{f}(x)$ of $f(x)$, we have

$$\mathbb{E}[(Y - \hat{f}(X))^2 \mid X = x] = [f(x) - \hat{f}(x)]^2 + \mathrm{Var}(\epsilon),$$

where we assumed $\hat{f}$ and X are fixed and $[f(x) - \hat{f}(x)]^2$ represents the *reducible* error and $\mathrm{Var}(\epsilon)$ the *irreducible* error.

Parametric and Structured Models

The *linear* model is an important example of a parametric model:

$$f_L(X) = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \cdots + \beta_p X_p.$$

- A linear model is specified in terms of $p+1$ parameters $\beta_0, \beta_1, \ldots, \beta_p$.
- We estimate the parameters by fitting the model to training data.
- Although it is *almost never correct*, a linear model often serves as a good and *interpretable approximation* to the unknown true function $f(X)$.

A linear model $\hat{f}_L(X) = \hat{\beta}_0 + \hat{\beta}_1 X$ gives a reasonable fit here.

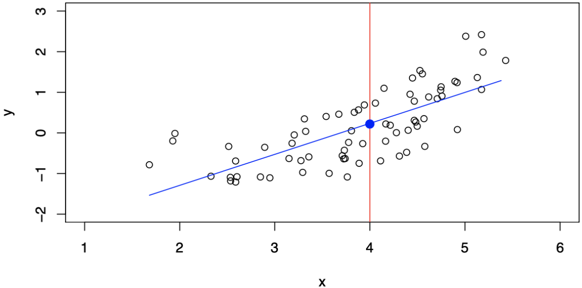

A quadratic model $\hat{f}_Q(X) = \hat{\beta}_0 + \hat{\beta}_1 X + \hat{\beta}_2 X^2$ fits slightly better.

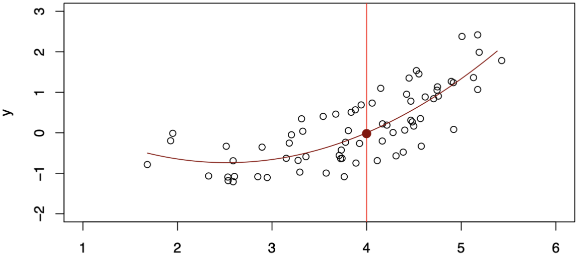

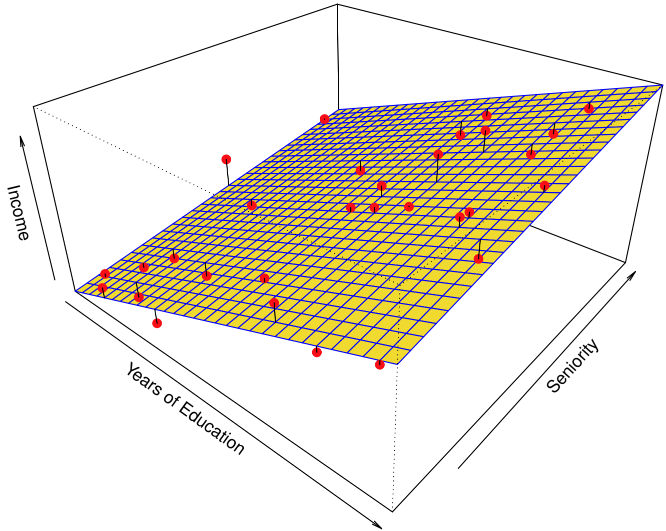

Linear regression model fit to the simulated data.

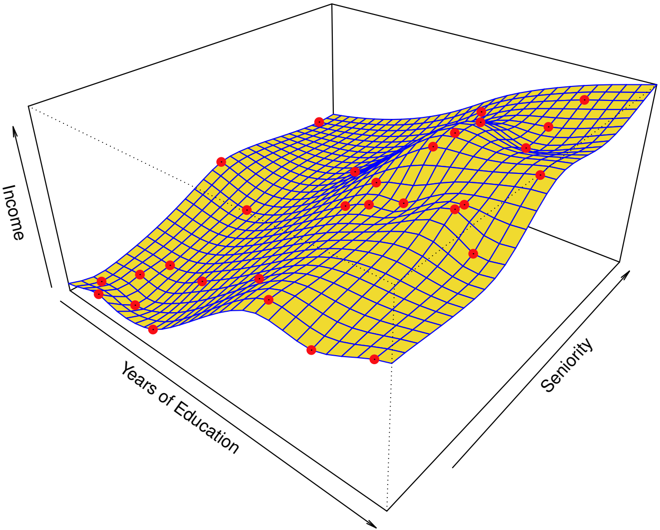

Very flexible spline regression model fit to the simulated data. Here the fitted model makes no errors on the training data! Also known as *overfitting*.

Assessing Model Accuracy

Suppose we fit a model $f(x)$ to some training data $\mathrm{Tr} = \{x_i, y_i\}$ for $i = 1, \ldots, N$, and we wish to see how well it performs.

- We could compute the average squared prediction error over Tr:

$$\mathrm{MSE}_{\mathrm{Tr}} = \mathrm{Ave}_{i \in \mathrm{Tr}}[y_i - \hat{f}(x_i)]^2.$$

This may be biased toward more overfit models.

- Instead we should, if possible, compute it using fresh test data $\mathrm{Te} = \{x_i, y_i\}$ for $i = 1, \ldots, M$:

$$\mathrm{MSE}_{\mathrm{Te}} = \mathrm{Ave}_{i \in \mathrm{Te}}[y_i - \hat{f}(x_i)]^2.$$

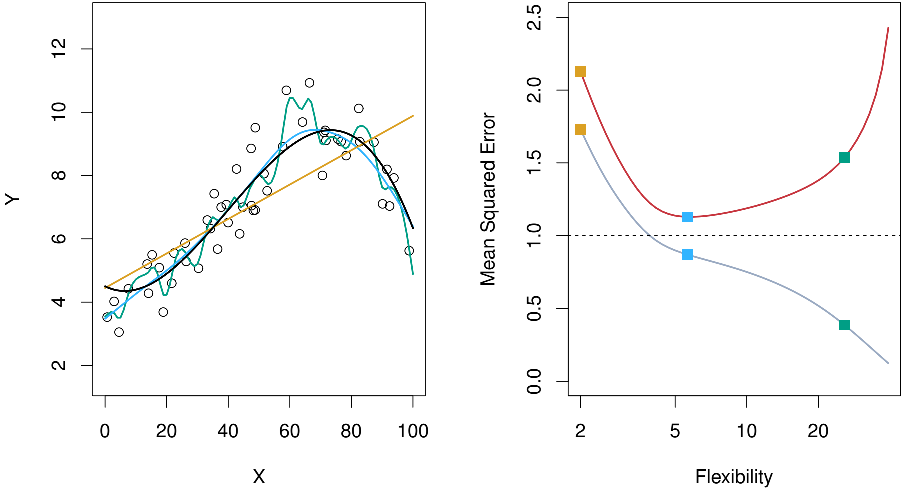

Black curve is the truth. Red curve on right is $\mathrm{MSE}_{\mathrm{Te}}$, grey curve is $\mathrm{MSE}_{\mathrm{Tr}}$. Orange, blue and green curves/squares correspond to fits of different flexibility.

Bias-Variance Trade-Off

- Suppose we have fit a model $f(x)$ to some training data Tr, and let (x_0, y_0) be a test observation drawn from the population. If the true model is $Y = f(X) + \epsilon$ (with $f(x) = \mathbb{E}[Y \mid X = x]$), then

$$\mathbb{E}[y_0 - \hat{f}(x_0)]^2 = \mathrm{Var}(\hat{f}(x_0)) + [\mathrm{Bias}(\hat{f}(x_0))]^2 + \mathrm{Var}(\epsilon).$$

- The expectation averages over the variability of y_0 as well as the variability in Tr. Note that $\mathrm{Bias}(\hat{f}(x_0)) = \mathbb{E}[\hat{f}(x_0)] - f(x_0)$.

- Typically as the *flexibility* of $\hat{f}$ increases, its variance increases, and its bias decreases. So choosing the flexibility based on average test error amounts to a *bias-variance trade-off*.

Bias-Variance Trade-Off

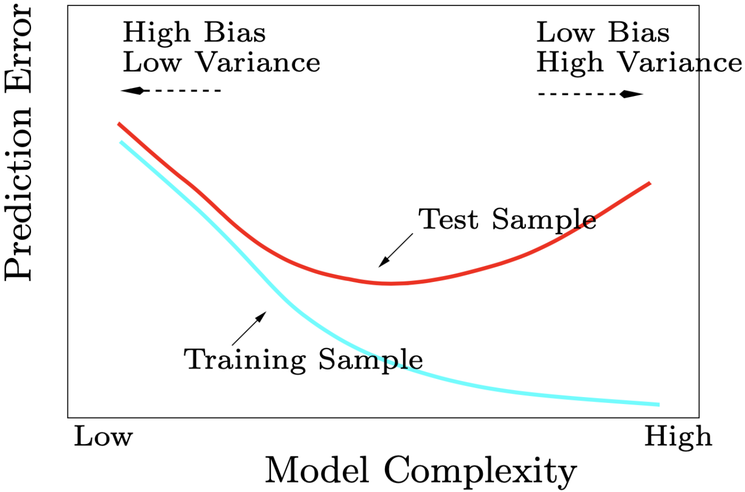

Training- vs test-set performance (in terms of prediction error) given model complexity.

Bias-Variance Trade-Off

- *Trade-off* ability to *extract signal* (bias) with *overfitting* (variance).
    - More complicated models tend to approximate f better, but run into the risk of fitting ϵ - specific to the *sample* you have.
    - Always leave some room for "noise" - the effect of all remaining covariates (not present the model, usually also non-observable).

- Difficulty hinges on the signal-to-noise ratio, dimensionality, and sample size.
    - Lower signal-to-noise ratio or higher dimensionality result in an elevated risk of overfitting: simpler models tend to forecast better!
    - Imposing assumptions/structures (regularization, shrinkage, sparsity, priors) helps and is often necessary. Otherwise, one needs lots of data!

Why Applying Machine Learning to Risk Management?

- ▷ *Reason 1*: Measurement of an asset's *expected* (excess) *return* (=risk premium) is fundamentally a problem of prediction!
  - A risk premium is a conditional expectation of a future realized (excess) return.
  - Machine learning methods are largely specialized for prediction tasks, thus, ideally suited to risk premium measurement.
- ▷ *Reason 2*: *Functional form* is *unknown* and likely complex.
  - Theoretical literature offers little guidance of winning lists of conditioning variables and functional forms.
  - Machine learning methods are designed to approximate complex non-linear associations. They cast a wide net in model search. Parameter penalization and conservative model selection criteria help avoid overfit and false discovery.

Why Applying Machine Learning to Risk Management?

▷ *Reason 3*: The collection of *candidate* conditioning *variables* is large and messy!

- Researchers and practitioners have accumulated a staggering list of return predictors (stock level predictive characteristics, macroeconomic predictors of the aggregate market).
- They are often close cousins and highly correlated.
- Historical data in finance are typically not large in size (barely 10s of years of data + lack of stationarity).
- With the emphasis of variable selection and dimensionality reduction, machine learning is well suited for such challenging empirical issues.

Have You Already Used Machine Learning?

▷ *Yes*, probably without knowing or emphasizing it.

- *Dimension reduction*: use portfolios instead of individual assets.
- *Variable selection* and *factor analysis*: CAPM, Fama-French factors.
- *Nonlinearity*: sort returns by characteristics.
- *Priors* (regularization): use economic intuition/theory.
- . . .

Caveats of Machine Learning

- *Data snooping*: find patterns when there are none.
  - Machine learning (ML) makes data snooping even easier.
- *Prediction* vs *Causal Inference*
  - ML aims at better prediction instead of valid inference (hot area: ML + econometrics).
  - Of course, the usual caveat still applies: correlation is not causation.
- *Biased data*: discrimination, inequality.
  - Artificial Intelligence is only human.
- *Interpretability*: black-boxes, tricks without theoretical justification.
  - Investors, clients, regulators, etc., demand transparency.
  - Black-box algorithms may not necessarily beat white-box ones.
  - Need more research or trade-off performance with interpretability.