Discussion of:
Pricing currency risks
Mikhail Chernov, Magnus Dahlquist, and Lars Lochstoer
Kent Daniel[†]
[†]Columbia Business School & NBER
2021 Western Finance Association Meetings
17 June 2021
Columbia Business School

10 minute discussion.

1. Very cool paper.
2. Interesting idea; beautifully empirical analysis
3. I'm going to make a few comments, and relate this a bit to the literature on equity factors that Lars mentioned.
4. There is a lot of really good analysis that I won't cover.
   - Multi-Horizon Regression (MHR); Currency denomination analysis

Discussion of: Pricing currency risks

Mikhail Chernov, Magnus Dahlquist, and Lars Lochstoer

Kent Daniel[†]

[†]Columbia Business School & NBER

2021 Western Finance Association Meetings
17 June 2021

Columbia Business School

Basic Idea
Key Questions:
What matters in determining currency expected returns (μ)?
What determines the currency covariance structure (Σ)?
What is the link between μ and Σ?
Answers:
Three characteristics determine expected currency returns.
interest rate differentials, 1 yr. momentum/trend, and ∼5 yr. mean reversion in PPP differentials
A UMVE based on a linear-characteristic model explains the returns of nine standard strategies extremely well.
The converse is not true.
μ does not line up with the first few PCs of Σ
for Σs for currency returns and for strategy returns.
At least 85% of the risk in standard currency strategies ("factor-portfolios") is unpriced.
Suggests that extant explanations for currency premia are problematic
Priced component is correlated with consumption growth.

- Three characteristics are interest rate differentials, momentum/trend, and longer-term (∼5 year) mean-reversion in PPP differentials.
- (Answers 2) Note that no combination of strategies can price the UMVE portfolio.

Basic Idea

Key Questions:

- What matters in determining currency expected returns (μ)?
- What determines the currency covariance structure (Σ)?
- What is the link between μ and Σ?

Answers:

- Three characteristics determine expected currency returns.
  - interest rate differentials, 1 yr. momentum/trend, and ∼5 yr. mean reversion in PPP differentials
- A UMVE based on a linear-characteristics model explains the returns of nine standard strategies extremely well.
  - The converse is not true.
- μ does not line up with the first few PCs of Σ
  - for Σs for currency returns and for strategy returns.
- At least 85% of the risk in standard currency strategies ("factor-portfolios") is unpriced.
  - Suggests that extant explanations for currency premia are problematic
  - Priced component is correlated with consumption growth.

Basic Idea

Key Questions:

- What matters in determining currency expected returns (μ)?
- What determines the currency covariance structure (Σ)?
- What is the link between μ and Σ?

Answers:

- Three characteristics determine expected currency returns.
  - interest rate differentials, 1 yr. momentum/trend, and ∼5 yr. mean reversion in PPP differentials
- A UMVE based on a linear-characteristics model explains the returns of nine standard strategies extremely well.
  - The converse is not true.
- μ does not line up with the first few PCs of Σ
  - for Σs for currency returns and for strategy returns.
- At least 85% of the risk in standard currency strategies (“factor-portfolios”) is unpriced.
  - Suggests that extant explanations for currency premia are problematic
  - Priced component is correlated with consumption growth.

- Three characteristics are interest rate differentials, momentum/trend, and longer-term (∼5 year) mean-reversion in PPP differentials.
- (Answers 2) Note that no combination of strategies can price the UMVE portfolio.

Basic Idea

Key Questions:

- What matters in determining currency expected returns (μ)?
- What determines the currency covariance structure (Σ)?
- What is the link between μ and Σ?

Answers:

- Three characteristics determine expected currency returns.
  - interest rate differentials, 1 yr. momentum/trend, and ∼5 yr. mean reversion in PPP differentials
- A UMVE based on a linear-characteristics model explains the returns of nine standard strategies extremely well.
  - The converse is not true.
- μ does not line up with the first few PCs of Σ
  - for Σs for currency returns and for strategy returns.
- At least 85% of the risk in standard currency strategies (“factor-portfolios”) is unpriced.
  - Suggests that extant explanations for currency premia are problematic
  - Priced component is correlated with consumption growth.

Notation & Basics
Definitions:
- $\mathbf{r}_{t+1}$ is the N_t-vector of realized excess currency returns from $t \to t+1$.
- Expected return; covariance matrix:
- Conditional MVE portfolio:
- Pricing equation:
where
- Pricing Kernel formulation (Hansen and Richard, 1987):

- Everything bold is a vector or matrix; non-bold is a scalar
- $\mathbf{w}_t^C$ can of course be scaled up or down
  - It is the portfolio that captures the premium in this universe of securities with minimum variance.
  - Thus, everything orthogonal to this portfolio earns zero premium
  - Thus, you can decompose any security's return into some MVE portfolio, and a residual which earns zero premium.
- Note that the denominator of the projection coefficient ($\boldsymbol{\mu}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}$) is both the expected return and the squared Sharpe-ratio of the CMVE portfolio.

Notation & Basics

Definitions:

- $\mathbf{r}_{t+1}$ is the N_t-vector of realized excess currency returns from $t \to t+1$.
- Expected return; covariance matrix:

$$\begin{aligned}\boldsymbol{\mu}_t &= \mathbb{E}_t[\mathbf{r}_{t+1}] \\ \boldsymbol{\Sigma}_t &= \mathbb{E}_t[(\mathbf{r}_{t+1} - \boldsymbol{\mu}_t)(\mathbf{r}_{t+1} - \boldsymbol{\mu}_t)']\end{aligned}$$

- Conditional MVE portfolio:

$$\begin{aligned}\mathbf{w}_t^C &= \boldsymbol{\Sigma}_t^{-1}\boldsymbol{\mu}_t \\ r_{t+1}^C &= \mathbf{w}_t^{C\prime}\,\mathbf{r}_{t+1}\end{aligned}$$

- Pricing equation:

$$\mathbf{r}_{t+1} = \boldsymbol{\beta}_t r_{t+1}^C + \boldsymbol{\epsilon}_{t+1}, \quad \mathbb{E}_t[\boldsymbol{\epsilon}_{t+1}] = \mathbf{0}$$

where $$\boldsymbol{\beta}_t = \frac{\boldsymbol{\mu}}{\boldsymbol{\mu}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}}$$

- Pricing Kernel formulation (Hansen and Richard, 1987):

$$m_{t+1} = 1 - (r_{t+1}^C - \mu_t^C) \;\Rightarrow\; \mathbb{E}_t[m_{t+1}\mathbf{r}_{t+1}] = \mathbf{0}$$

- For the pricing equation on the last page to work, the returns on the right side have to give you exposure to the sources of premium with minimum risk.
- If you form the factor portfolios using just the characteristics (which capture μ), and not *Sigma*, you won't generally get exposure with minimum variance.

How much unpriced risk could there be?

- Let's look at what contributes to the variance of Fama and French's HML portfolio.

Factor Pricing Models

- In the last few decades, the standard has become the use of "factor" models, based on characteristic-sorts.
- Such models dominate the asset pricing literature devoted to pricing equities, but are also used in currency pricing models.
  - See, e.g., Fama and French (1993, 2015); Carhart (1997); Daniel and Moskowitz (2016); Pástor and Stambaugh (2003); Daniel and Titman (2006); Lustig, Roussanov, and Verdelhan (2011); Asness, Moskowitz, and Pedersen (2013)
- The models will only price the full cross-section if the resulting set of factor-portfolios span the MVE portfolio.
  - They may, however, successfully price other characteristic-sorted portfolios. See Lewellen, Nagel, and Shanken (2010) and Daniel and Titman (2012).
- Specifically, the problem is that these factor-portfolios ignore information about the covariance structure.
  - Thus, the resulting portfolios will contain both *priced* and *unpriced* risk.

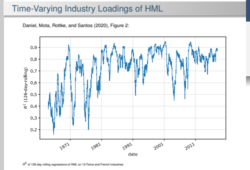

Time-Varying Industry Loadings of HML

Daniel, Mota, Rottke, and Santos (2020), Figure 2:

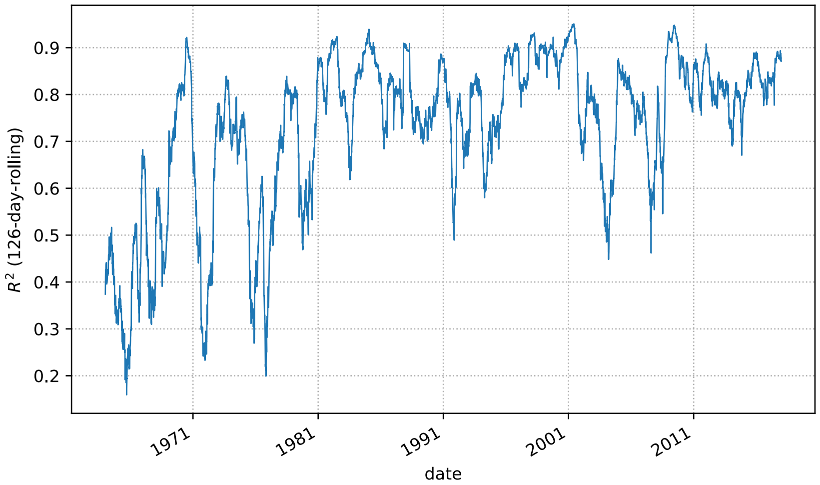

R^2 of 126-day rolling regressions of HML on 12 Fama and French industries

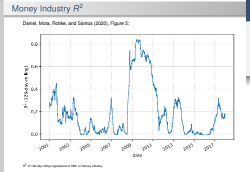

- I should show the R^2 from regressing HML on both Money and BusEq; this is above 90% in 2010.
- Note that the adjusted-R^2 is about zero before the start of the financial crisis.
- However, with the onset of the financial crisis, many large financial firms fall in price and become value firms
  - in a value-weighted portfolio.
- In addition, the portfolio of financial firms become really volatile.
- This leads to the correlation of HML with Money jumping above 90% for a period of time.
- So, HML gives you exposure to the *priced* value factor, but a lot of the risk in HML can be hedged-out without affecting the return
  - It is *unpriced*.
- To separate out the priced and unpriced risk, you need information both about the expected returns (from the characteristics), and the covariance/correlation structure.

Money Industry R^2

Daniel, Mota, Rottke, and Santos (2020), Figure 5:

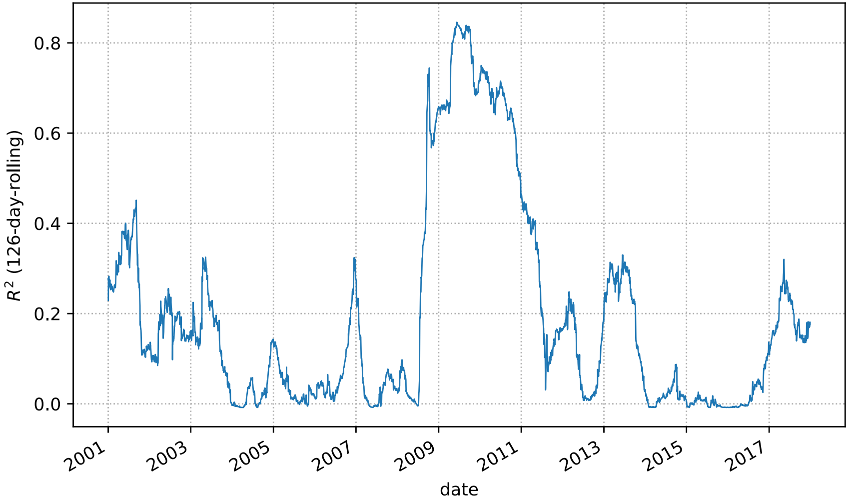

R^2 of 126-day rolling regressions of HML on Money industry

- What we do in our RFS paper is to show how to hedge out the unpriced risk via an iterative procedure.
- Lars, Mike and Magnus instead build an MVE portfolio directly using an estimated covariance matrix, in addition to the characteristics.
- The estimated CMVE portfolio is just based on the estimated $\hat{\Sigma}$ and $\hat{\mu}$.
- The Unconditional MVE portolio (UMVE), just levers this up or down based on the (estimated) squared-Sharpe ratio of the MVE portfolio.

Building the CMVE and UMVE Portfolios

$$\hat{\mathbf{w}}_t^C = \hat{\Sigma}_t^{-1}\hat{\boldsymbol{\mu}}_t, \qquad r_{t+1}^C = \hat{\mathbf{w}}_t^{C\prime}\mathbf{r}_{t+1}$$

$$r_{t+1}^U = \frac{r_{t+1}^C}{1 + \hat{\boldsymbol{\mu}}_t'\hat{\Sigma}_t^{-1}\hat{\boldsymbol{\mu}}_t}$$

- Building a good MVE portfolio based on an estimated $\hat{\Sigma}$ is tricky.
  - e.g., use of a sample covariance matrix will result in an estimated MVE which lines up with small eigenvalues of $\hat{\Sigma}$
  - There are some techniques, starting with Black and Litterman (1991), designed to resolve these problems.
- However here, with a much smaller universe, the problem is not as severe.
  - $\hat{\Sigma}_t^{-1}$ is estimated using an exponential-weighted moving average applied to daily data with shrinkage, following Ledoit and Wolf (2020).
  - This seems to work pretty well in forecasting the covariance structure.
    - *Would it be useful to use different decay rates for correlations and volatilities?*

Building the CMVE and UMVE Portfolios

- What we do in our RFS paper is to show how to hedge out the unpriced risk via an iterative procedure.
- Lars, Mike and Magnus instead build an MVE portfolio directly using an estimated covariance matrix, in addition to the characteristics.
- The estimated CMVE portfolio is just based on the estimated $\hat{\Sigma}$ and $\hat{\mu}$.
- The Unconditional MVE portolio (UMVE), just levers this up or down based on the (estimated) squared-Sharpe ratio of the MVE portfolio.

Building the CMVE and UMVE Portfolios

$$\hat{\mathbf{w}}_t^C = \hat{\Sigma}_t^{-1}\hat{\mu}_t, \qquad r_{t+1}^C = \hat{\mathbf{w}}_t^{C\prime}\mathbf{r}_{t+1}$$

$$r_{t+1}^U = \frac{r_{t+1}^C}{1 + \hat{\mu}_t'\hat{\Sigma}_t^{-1}\hat{\mu}_t}$$

- Building a good MVE portfolio based on an estimated $\hat{\Sigma}$ is tricky.
  - e.g., use of a sample covariance matrix will result in an estimated MVE which lines up with small eigenvalues of $\hat{\Sigma}$
  - There are some techniques, starting with Black and Litterman (1991), designed to resolve these problems.
- However here, with a much smaller universe, the problem is not as severe.
  - $\hat{\Sigma}_t^{-1}$ is estimated using an exponential-weighted moving average applied to daily data with shrinkage, following Ledoit and Wolf (2020).
  - This seems to work pretty well in forecasting the covariance structure.
    - *Would it be useful to use different decay rates for correlations and volatilities?*

Building the CMVE and UMVE Portfolios

$$\hat{\mathbf{w}}_t^C = \hat{\Sigma}_t^{-1}\hat{\mu}_t, \quad r_{t+1}^C = \hat{\mathbf{w}}_t^{C\prime}\mathbf{r}_{t+1}$$

$$r_{t+1}^U = \frac{r_{t+1}^C}{1+\hat{\mu}_t'\hat{\Sigma}_t^{-1}\hat{\mu}_t}$$

- Building a good MVE portfolio based on an estimated $\hat{\Sigma}$ is tricky.
  - e.g., use of a sample covariance matrix will result in an estimated MVE which lines up with small eigenvalues of $\hat{\Sigma}$
  - There are some techniques, starting with Black and Litterman (1991), designed to resolve these problems.
- However here, with a much smaller universe, the problem is not as severe.
  - $\hat{\Sigma}_t^{-1}$ is estimated using an exponential-weighted moving average applied to daily data with shrinkage, following Ledoit and Wolf (2020).
  - This seems to work pretty well in forecasting the covariance structure.
    - *Would it be useful to use different decay rates for correlations and volatilities?*

- What we do in our RFS paper is to show how to hedge out the unpriced risk via an iterative procedure.
- Lars, Mike and Magnus instead build an MVE portfolio directly using an estimated covariance matrix, in addition to the characteristics.
- The estimated CMVE portfolio is just based on the estimated $\hat{\Sigma}$ and $\hat{\mu}$.
- The Unconditional MVE portolio (UMVE), just levers this up or down based on the (estimated) squared-Sharpe ratio of the MVE portfolio.

Building the CMVE and UMVE Portfolios

$$\hat{\mathbf{w}}_t^C = \hat{\Sigma}_t^{-1}\hat{\mu}_t, \quad r_{t+1}^C = \hat{\mathbf{w}}_t^{C\prime}\mathbf{r}_{t+1}$$

$$r_{t+1}^U = \frac{r_{t+1}^C}{1+\hat{\mu}_t'\hat{\Sigma}_t^{-1}\hat{\mu}_t}$$

- Building a good MVE portfolio based on an estimated $\hat{\Sigma}$ is tricky.
  - e.g., use of a sample covariance matrix will result in an estimated MVE which lines up with small eigenvalues of $\hat{\Sigma}$
  - There are some techniques, starting with Black and Litterman (1991), designed to resolve these problems.
- However here, with a much smaller universe, the problem is not as severe.
  - $\hat{\Sigma}_t^{-1}$ is estimated using an exponential-weighted moving average applied to daily data with shrinkage, following Ledoit and Wolf (2020).
  - This seems to work pretty well in forecasting the covariance structure.
    - *Would it be useful to use different decay rates for correlations and volatilities?*

Linear Characteristics Model ($\hat{\boldsymbol{\mu}}_t$)

$$\hat{\boldsymbol{\mu}}_t = \gamma_t \cdot \left(\frac{\mathbf{S}_t - \mathbf{F}_t}{\mathbf{F}_t}\right) + \delta_t \cdot \mathbf{z}_t^Q + \phi_t \cdot \left(\frac{\mathbf{S}_t}{\mathbf{S}_{t-12}}\right)$$

- $\hat{\boldsymbol{\mu}}_t$ is based on a linear characteristics model:
  - The RWH coefficient $\gamma_t = 1$
  - The coefficients on the RER and momentum signals, δ_t and ϕ_t, are estimated with an expanding window, and after the first few years are fairly stable.
- Note that CDL are very much "hanging their hat" on this linear-characteristic/constant-coefficient model specification.
  - Note that any evidence of factor timing here arises as a result of time variation in the (1) characteristic loadings of the factor portfolios and (2) the covariance structure.
  - *Should there be any timing instruments?*
  - *Should the coefficients be a function of the currencies?*

The three terms are:

1. The Random Walk Hypothesis (i.e., the expected change in the spot rate is zero).
   – If the spot rate is 2% higher than the future, then the expected return is 2%.
2. High (real) prices adjust mean revert via exchange rate depreciation.
   – relative to 5 years before
3. Currencies that have appreciated over the last year keep appreciating.

To these suggestions:

- A simple model is robust. This is about as simple as it gets.
- You might expect the "model" to change over time, to work better in come conditions than others, etc.
- I'm guessing that this should work "better" for emerging currencies than developed. Suggests that coefficients should vary with currency characteristics.

Linear Characteristics Model ($\hat{\boldsymbol{\mu}}_t$)

$$\hat{\boldsymbol{\mu}}_t = \gamma_t \cdot \left(\frac{\mathbf{S}_t - \mathbf{F}_t}{\mathbf{F}_t}\right) + \delta_t \cdot \mathbf{z}_t^Q + \phi_t \cdot \left(\frac{\mathbf{S}_t}{\mathbf{S}_{t-12}}\right)$$

- $\hat{\boldsymbol{\mu}}_t$ is based on a linear characteristics model:
  - The RWH coefficient $\gamma_t = 1$
  - The coefficients on the RER and momentum signals, δ_t and ϕ_t, are estimated with an expanding window, and after the first few years are fairly stable.
- Note that CDL are very much "hanging their hat" on this linear-characteristic/constant-coefficient model specification.
  - Note that any evidence of factor timing here arises as a result of time variation in the (1) characteristic loadings of the factor portfolios and (2) the covariance structure.
  - *Should there be any timing instruments?*
  - *Should the coefficients be a function of the currencies?*

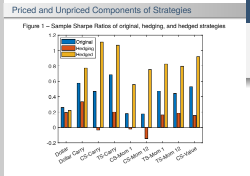

- Blue bars are the Sharpe ratios of each of the 9 strategies
- To hedge, you “pull out” the part of the strategy returns that is unpriced based on their model (orange, labeled “hedging”)
  - In theory, this should have an SR of zero
- What is left should have a higher SR.
  - This is bigger in sample, for 9/9 strategies.

Priced and Unpriced Components of Strategies

Figure 1 – Sample Sharpe Ratios of original, hedging, and hedged strategies

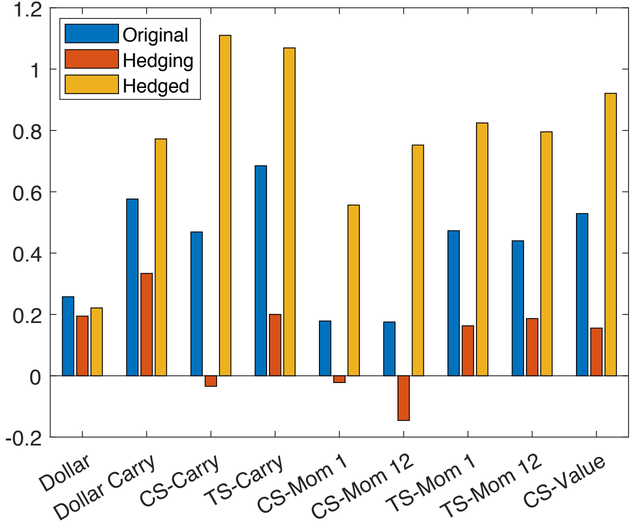

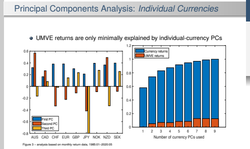

- Note that the return of the CMVE portfolio lies in the space spanned by the returns of the nine currencies
- However, this chart shows that the UMVE returns are only minimally explained by the nine PCs.
  - The reason is that the weights on the nine currencies are not static, they are constantly rotating depending on $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$.

Principal Components Analysis: *Individual Currencies*

- UMVE returns are only minimally explained by individual-currency PCs

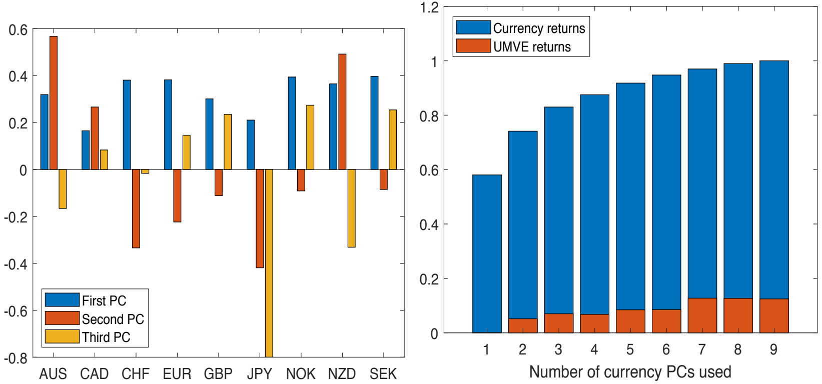

Figure 3 – analysis based on monthly return data, 1985:01–2020:05

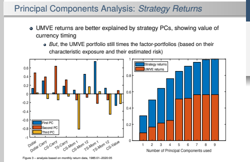

- The UMVE portfolio lines up much better with the strategy returns
  - This suggests that the UMVE is pretty well explained by a portfolio with static weights on the nine-strategies.
  - However, the R^2 is still far from 100%.

Principal Components Analysis: *Strategy Returns*

- UMVE returns are better explained by strategy PCs, showing value of currency timing
  - *But*, the UMVE portfolio still times the factor-portfolios (based on their characteristic exposure and their estimated risk)

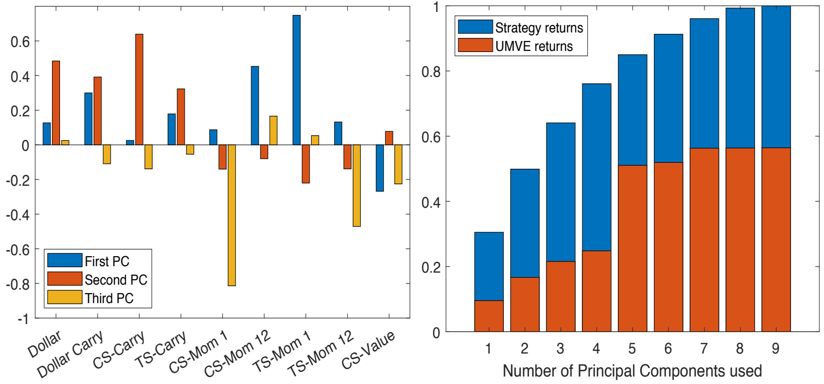

Figure 3 – analysis based on monthly return data, 1985:01–2020:05

Pushing the Empirical Analysis

- It is really interesting that the correlations of carry with intermediary capital and global volatility results from a correlation with the *unpriced* component.
- CDL find that the priced component of currency returns is correlated with consumption growth.
  - The R^2 for quarterly Δc is about 3%; for three-year Δc it is about 7%.
  - t-statistics are 2.5 and 3.2, respectively.
  - *What should we expect here?*
- Are there additional instruments that could enhance timing?
- *What factors drive* Σ_t
  - CDL show that it is not the factors that explain expected returns.
    - But what is it?
  - Is it regional factors? Is it related to trade?

Pushing the Empirical Analysis

- It is really interesting that the correlations of carry with intermediary capital and global volatility results from a correlation with the *unpriced* component.
- CDL find that the priced component of currency returns is correlated with consumption growth.
  - The R^2 for quarterly Δc is about 3%; for three-year Δc it is about 7%.
  - t-statistics are 2.5 and 3.2, respectively.
  - *What should we expect here?*
- Are there additional instruments that could enhance timing?
- *What factors drive* Σ_t
  - CDL show that it is not the factors that explain expected returns.
    - But what is it?
  - Is it regional factors? Is it related to trade?

References I

Asness, Clifford S., Toby J. Moskowitz, and Lasse Heje Pedersen, 2013, Value and momentum everywhere, *The Journal of Finance* 58, 929–895.

Black, Fischer, and Robert Litterman, 1991, Global asset allocation with equities, bonds, and currencies, Goldman, Sachs & Co. Fixed Income Research Report.

Carhart, Mark M., 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57–82.

Daniel, Kent, Lira Mota, Simon Rottke, and Tano Santos, 2020, The cross section of risk and return, *The Review of Financial Studies* 33, 1927–1979.

Daniel, Kent D., and Tobias J. Moskowitz, 2016, Momentum crashes, *Journal of Financial Economics* 122, 221–247.

Daniel, Kent D., and Sheridan Titman, 2006, Market reactions to tangible and intangible information, *Journal of Finance* 61, 1605–1643.

———, 2012, Testing factor-model explanations of market anomalies, *Critical Finance Review* 1, 103–139.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

———, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1–22.

Hansen, Lars Peter, and Scott F. Richard, 1987, The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models, *Econometrica* 55, 587–613.

Ledoit, Olivier, and Michael Wolf, 2020, Analytical nonlinear shrinkage of large-dimensional covariance matrices, *Annals of Statistics* 48, 3043–3065.

Lewellen, Jonathan, Stefan Nagel, and Jay Shanken, 2010, A skeptical appraisal of asset pricing tests, *Journal of Financial Economics* 96, 175–194.

Lustig, Hanno N., Nikolai L. Roussanov, and Adrien Verdelhan, 2011, Common risk factors in currency markets, *Review of Financial Studies* 24, 3731–3777.

Pástor, Ľuboš, and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642–685.