Time-dependent Black-Litterman

An ORTEC research paper by
Martin van der Schans and Hens Steehouwer

Discussion by Kris Boudt

Vrije Universiteit Brussel, Amsterdam

Why I (*think I*) am the discussant?

- My own experience with Black-Litterman: Reverse engineering of expected returns

Implied Expected Returns and the Choice of a Mean–Variance Efficient Portfolio Proxy

David Ardia and Kris Boudt

The Journal of Portfolio Management
Summer 2015, Vol. 41, No. 4: pp. 68-81
DOI: 10.3905/jpm.2015.41.4.068

which is not what the authors do in this paper.

- Research and practical work in the field of time series modelling and portfolio allocation. Algorithmic.

Structure of the discussion

- My waow-feeling when reading the paper;
- How?
- Hungry for more (OR/TEC+FIN).

Waow!

Figure 3: Portfolio weights (on the y-axis) of the allocation with the highest expected return divided by standard deviation for different investment horizons (on the x-axis in number of years since December 2008).

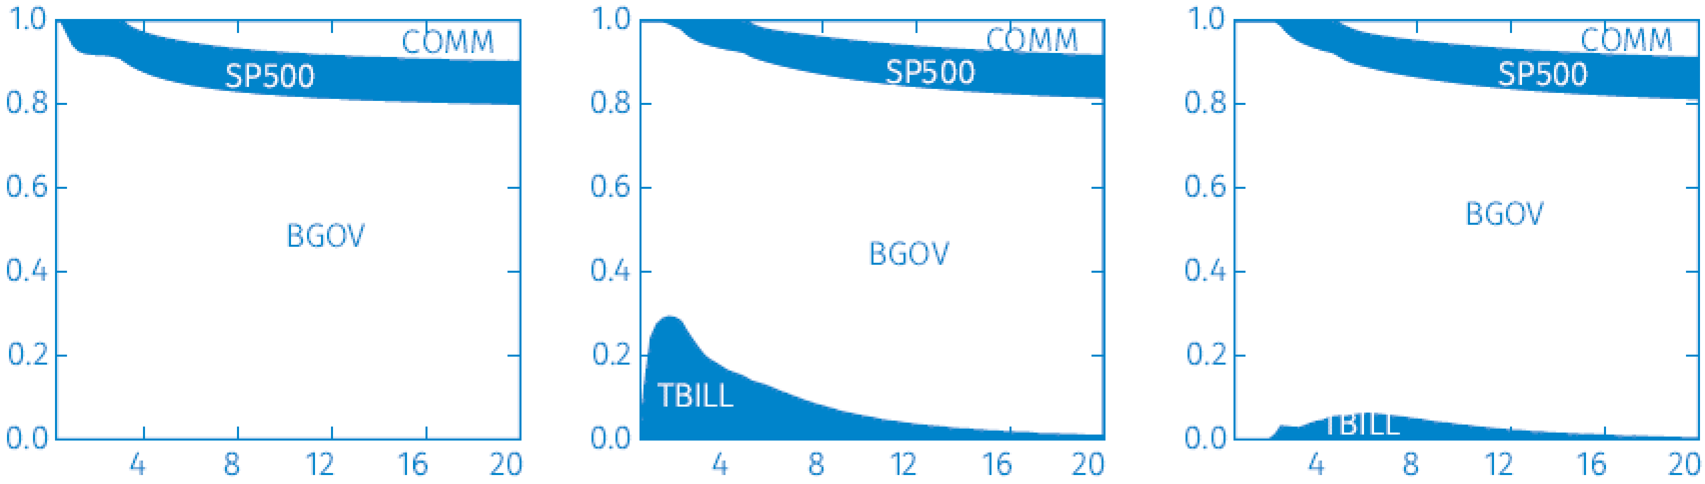

(a) No view, based on unadjusted forecast.

(b) With view, based on adjusted forecast using sampling approach with unconditional views.

(c) With view, based on adjusted forecast using sampling approach with conditional views.

How?

- Combining two major challenges in forecasting the return density:
  1. Going from a 1-period time series model to a h-period forecast = h-period investment horizon:
     - 1-period requirement often follows from the data: trade-off between timeliness of the data used (robustness to eg structural break) and sufficient number of observation. Eg use of weekly data for monthly investment horizon. Or monthly data for yearly investment horizon.
  2. Tilting the pure data-driven return density forecast (leading to prior estimates $\hat{\mu}_{t|t-1}$ and $\hat{\Sigma}_{t|t-1}$) into a return density forecast that respects the **views** of the expert portfolio manager about linear combinations of expected returns:
     - Views could be more easily formulated over shorter horizons?
     - Views can be conditional: if this happens, then that will happen in the next period (scenarios).

How?

1. Going from a 1-period time series model to a h-period forecast:
   - Mean and volatility? Because of time series dependence “annualization” by multiplying expected log-returns with h and volatilities by square root of h is not precise. Iterative calculation.

   See also: Ghysels, E.; Rubia, A. & Valkanov, R. (2009), 'Multi-period Forecasts of Volatility: Direct, Iterated and Mixed-data Approaches', Working paper.
   - Density? For (non-linear) time series models, explicit results are rarely available; Aggregational gaussianity? Use of simulation analysis?
   - Authors take a novel approach using a type of bayesian updating:
     - Based on time series model: Forecast: $\hat{\mu}_{t|t-1}$ and $\hat{\Sigma}_{t|t-1}$
     - Update based on new information of some state variable: $\hat{\mu}_{t|t}$ and $\hat{\Sigma}_{t|t}$
     - Then iterate for t+1,...

How?

2. Tilting the pure data-driven return density forecast (leading to: $\hat{\mu}_{t|t-1}$ and $\hat{\Sigma}_{t|t-1}$) into a return density forecast that respects the **views** of the expert portfolio manager about linear combinations of expected returns:

$$v_t \; = P_t \; \mu_t \; + \xi_t$$

where P_t is a matrix and $\xi_t \sim N(0, \Omega_t)$.

- Includes views on outperformance of (groups) of assets versus other (groups) of assets;
- This is achieved through the Black-Litterman approach (under the assumption of normality).

How?

Flexibility of the proposed technology: It can deal with everal approaches:

- Bayesian and frequentist;
- Unconditional and conditional view:

In practice, however, views can be constructed conditionally, e.g., an investor might expect a stock market crash next year and a recovery the year after. The recovery, of course, is only expected to happen conditional on the crash having taken place. Mathematically, we interpret

- Tree? Feasible when h becomes large or many views? Artificial (?) solution in the paper:

where $\mu_{t-1} = m_{t-1}$ is the event that should have taken place for the view to hold. In this paper, we investigate a specific conditional view and assume the investor constructs his view conditional on the posterior returns at time $t-1$ equal their expected value, i.e., conditional on $\mu_{t-1|t-1} = \mathrm{E}\,\mu_{t-1|t-1}$ (and $\varepsilon_{t-1} = 0$):

$$v_t = \left[\left(P_t \mu_{t|t-1} + \xi_t \right) \middle| \left(\mu_{t-1|t-1} = \mathrm{E}\,\mu_{t-1|t-1} \right) \right]. \quad (11)$$

How?

- Time series model used for the prior fully data-driven denisty forecast: dynamic factor model of Stock and Watson.

A well known dynamic factor model is the one presented in Stock and Watson (2002, 2011). There, the historical time series of the factors are estimated with a principal component analysis on a dataset containing 108 economic and financial time series for the United States. Selecting the first four to ten principal components is sufficient to explain most of the variance and correlation structure of this dataset. As in Stock and Watson (2012b), we consider a dynamic factor model of the form:

$$F_t = C_F + AF_{t-1} + \phi_t^F, \tag{25a}$$

$$R_t = C_R + BF_t + \phi_t^R. \tag{25b}$$

Here, $t = 1, \dots T$, F_t is a vector of factors, R_t is a vector of economic and financial variables regressed on the factors, C_F and C_R are vectors with regression constants, A and B are matrices with regression coefficients, $\phi_t^F \sim N(0, \Phi_F)$, $\phi_t^R \sim N(0, \Phi_R)$ (independent of η_t) and Φ_R is assumed to be diagonal.

Resulting forecasts for the expected return at different horizons

	GDP		CPI	
	IMF	Model	IMF	Model
2009	-2.8 %	2.4 %	-0.9 %	-7.1 %
2010	-0.1 %	4.7 %	-0.1 %	-2.4 %
2011	3.5 %	3.1 %	0.7 %	0.2 %
2012	3.6 %	2.7 %	1.7 %	1.0 %
2013	3.3 %	2.6 %	2.1 %	1.3 %
2014	2.5 %	2.6 %	2.2 %	1.6 %

view *view*

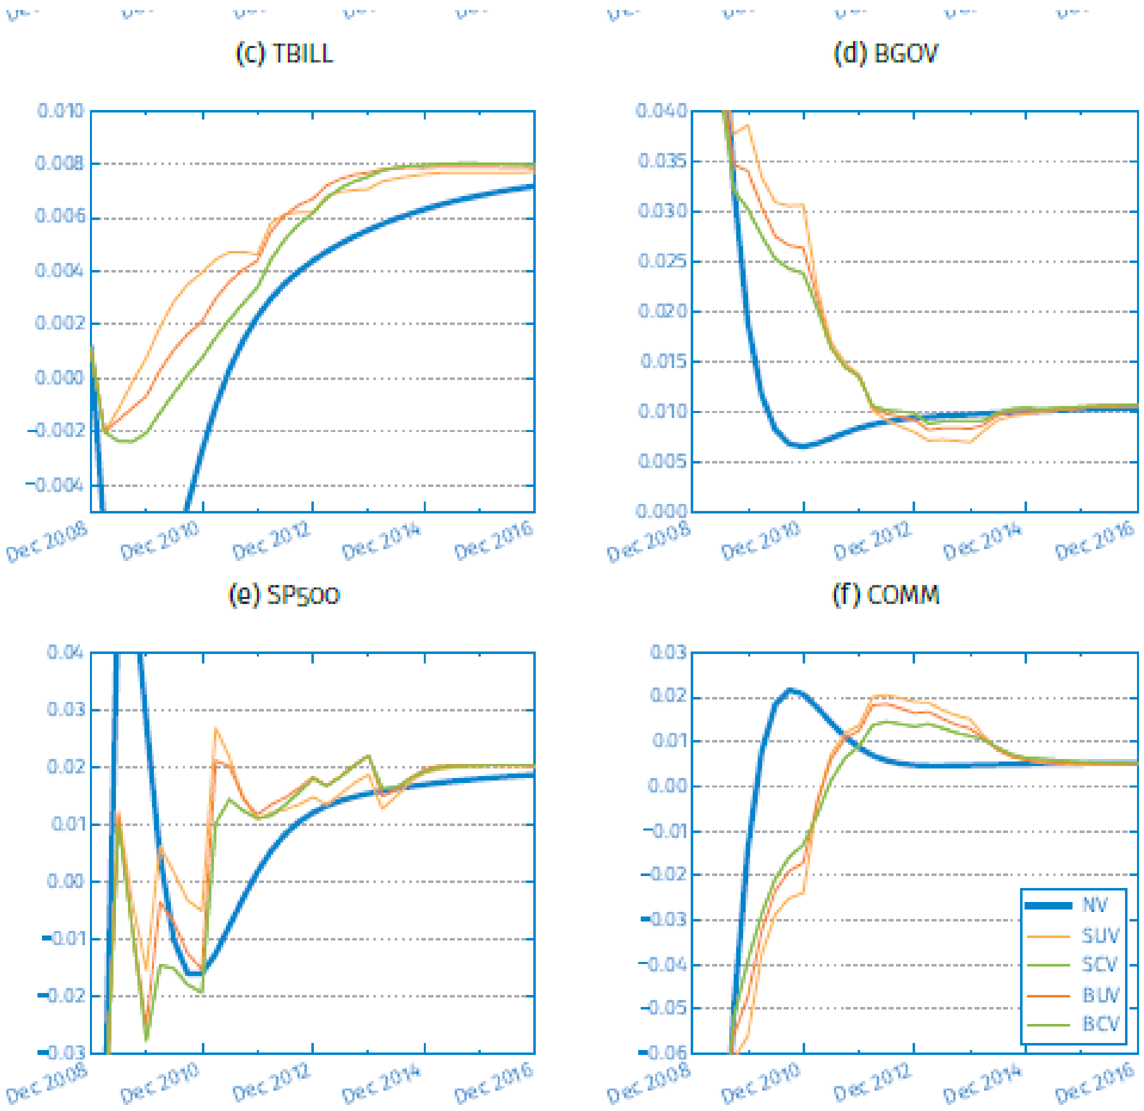

Waow! But hungry for more (OR/TEC,FIN)

Figure 3: Portfolio weights (on the y-axis) of the allocation with the highest expected return divided by standard deviation for different investment horizons (on the x-axis in number of years since December 2008).

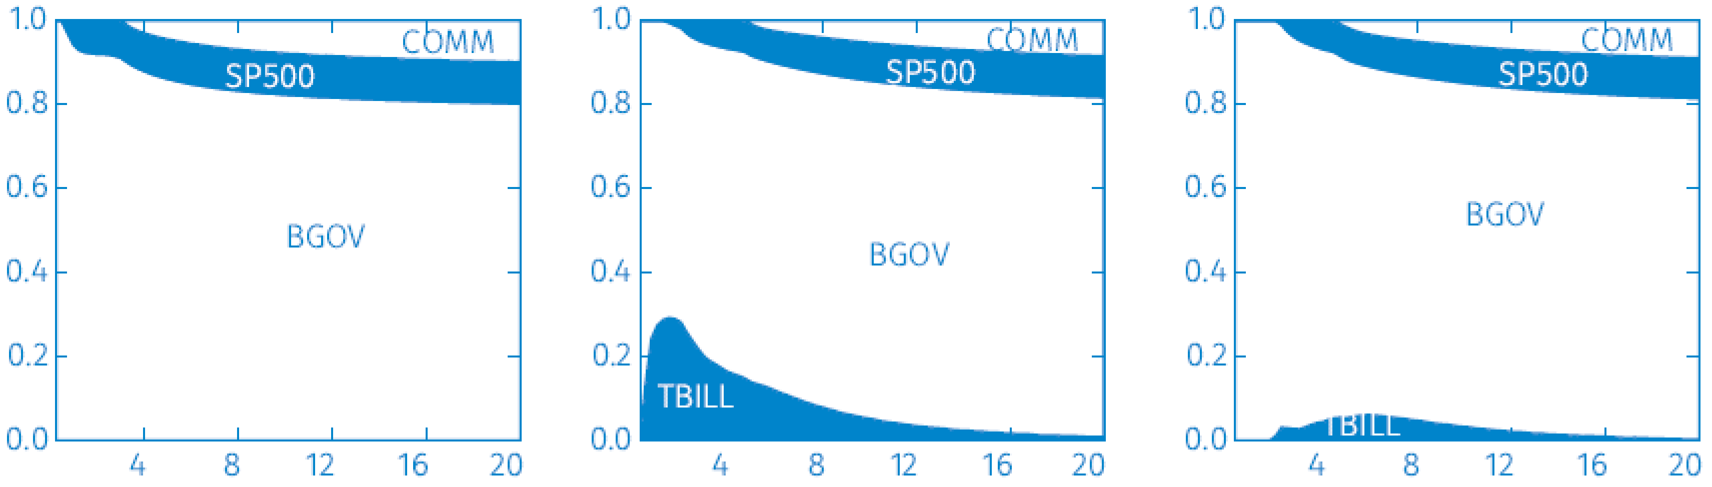

(a) No view, based on unadjusted forecast.

(b) With view, based on adjusted forecast using sampling approach with unconditional views.

(c) With view, based on adjusted forecast using sampling approach with conditional views.

Waow! But hungry for more (OR/TEC)

- I’d like to play with the code
  - (paper is well written) Study the implementation and understand better the technicalities in the various conditioning and how the iteration affect the forecasts $t|t-1$, $t|t$, $t+1|t$,$t+1|t+1$,...
  - And see how sensitive the results are:
    - To the views and their uncertainty;
    - To assumptions like normality;
- Understand better the conditional view mechanism and whether there is a curse of dimensionality with forecast horizon h.

Waow! But hungry for more (FIN)

- Understand better this allocation and how it changes when the portfolio objectives and constraints change;
- The “term structure” of out-of-sample performance of
  - The return density forecasts (statistical evaluation)
  - The resulting portfolio decision (economic evaluation)
  - Compared to:
    - alternative approaches to obtaining long term density forecasts (eg using simulation);
    - Alternative approaches of combining views with prior distributions (eg using entropy pooling as in David Ardia and Attilio Meucci).
  - Term structure: I’d be interested to see how the results change with the forecast horizon.

Conclusion

- Optimizing portfolios with a longer investment horizon than the horizon of the data is a common and challenging problem;
- A useful paper, interesting for both
  - OR/TEC
  - FIN ;
  - Needs a thorough out-of-sample analysis and comparison with other approaches to be a complete paper.