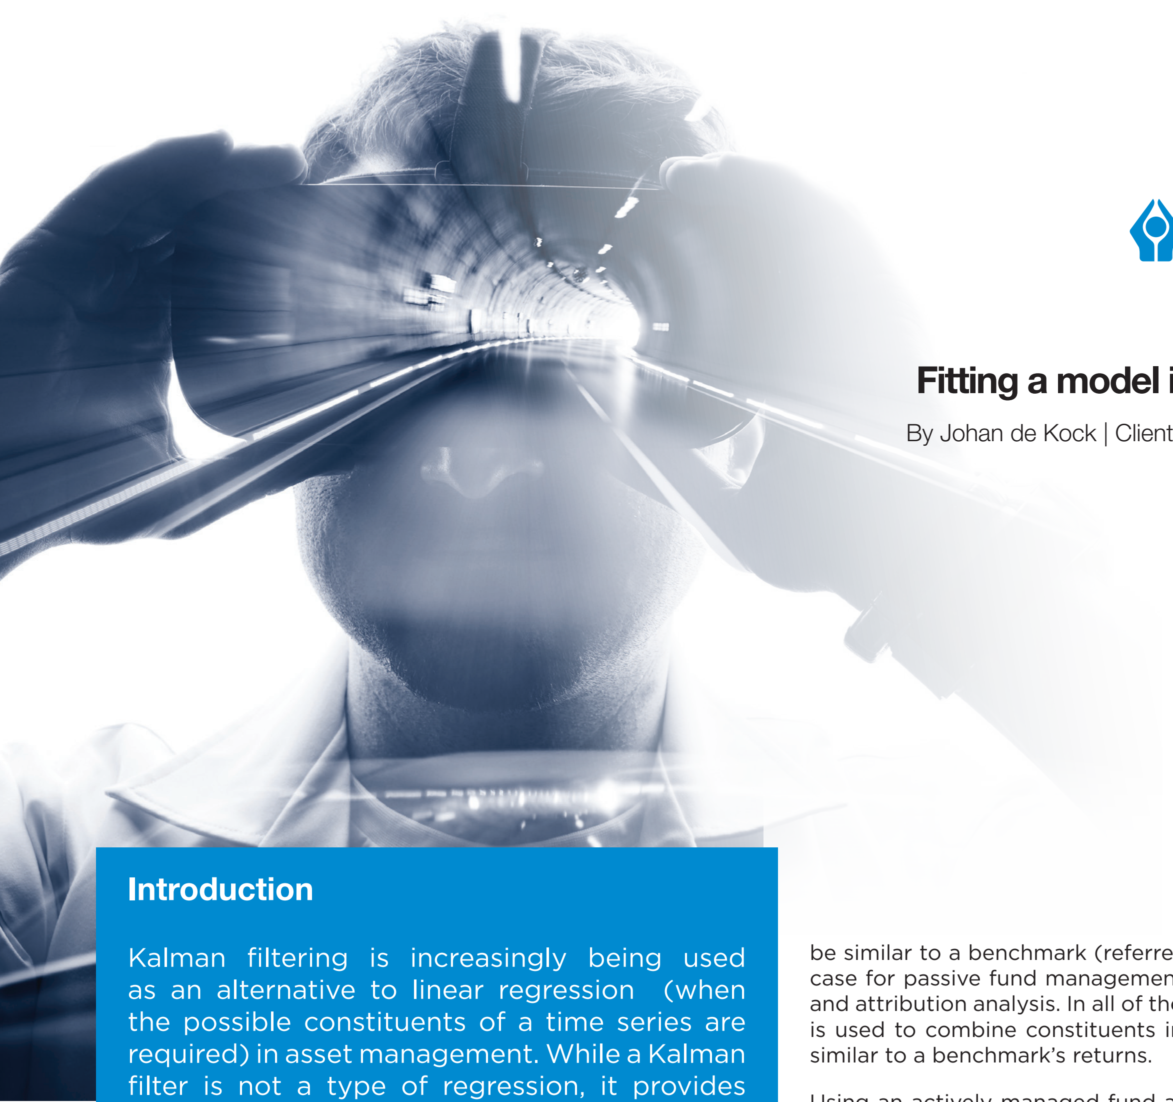

Kalman Filters

Fitting a model instead of a line

By Johan de Kock | Client Solutions and Research

22 January 2020

Investments

Introduction

Kalman filtering is increasingly being used as an alternative to linear regression (when the possible constituents of a time series are required) in asset management. While a Kalman filter is not a type of regression, it provides an entire class of models with rich sets of parameters. These parameters convey various insights into the real world scenario being modelled. A Kalman filter can also forecast whereas regression only attempts to fit a line through a given number of points and ignores the past before these points. A bigger issue with using regression in fund analysis is that it does not represent a model of the real world.

This article shows that Kalman filters can provide more accurate models of the real world in asset management applications and extremely smooth tracking weights.

Uses of Regression and Regression-Type Methods in Asset Management

A large part of asset management evolves around building an optimal portfolio given available **constituents**. These could be single assets like Alphabet shares, indices like the Capped SWIX (J433) or alternative asset classes like African credit. If this optimal portfolio represents an actively managed fund it has to outperform a selected benchmark, but does not need to have similar return rates or characteristics to the benchmark. Such active portfolios might be designed to maximise return and minimise risk, but there are many portfolios that are designed to be similar to a benchmark (referred to as the **target**). This is the case for passive fund management, tracking, and in factor/style and attribution analysis. In all of these applications some method is used to combine constituents in order to obtain return rates similar to a benchmark's returns.

Using an actively managed fund as an example, assume several local and foreign constituents (equity, bonds and property) make up the returns. Suppose monthly returns are available and the manager wants to combine the constituents (i.e. find the optimal portfolio **weights** of these constituents) that would yield returns similar to the **target**. A rolling window regression would fit a line through a set of returns called a window (typically 24 months) and use the resulting weights for the last point in the window. Combining the constituents in these weights yields the returns **achieved** by the portfolio. This process is repeated by moving the window on by one month. Firstly, due to the overlapping windows, the weights calculated for say December 2019 implicitly also hold for November 2019, yet they are ignored. The approach is therefore internally inconsistent. This model is appropriate if one is looking for the best way the constituents can be combined to match the target **on average**, but not exactly. Secondly, since regression tries to fit a line through the points, it is not really an accurate model of what is happening in asset allocation over time.

The distance between target returns and achieved returns (usually mean-squared error or R^2), might even be smaller in the case of regression than for Kalman filtering. This does not imply that the weights are correct, but means that they are more optimal locally. This local optimality can also lead to unstable weights.

Regression is a very useful statistical tool and is still the most appropriate way to obtain the average allocation over time and the appropriate model in this case. However, the article shows that Regression is simply the wrong model for the time-varying constituent analysis of a target, which a Kalman Filter is better suited to.

What is a Kalman Filter?

Kalman filters (similar to other statistical filtering techniques) always try to solve problems involving at least two event chains. The one process is observable and the other process is hidden. As an investor, one can only see the returns of the target (the observable process) and does not know exactly how the fund manager invests the money, i.e. the weightings of the constituents. Regulations require fund managers to disclose high-level asset allocation on a regular basis, but this information typically provided in MDDs does not truly enable an investor to calculate the exact holdings and their movements over time. Figure 1 is a graphical attempt to convey the reasoning behind Kalman filters, the next step of the observable process is provided based on a Kalman forecast of the hidden process. The forecasting accuracy improves with time. A single filtering step is described in figure 2. This step is repeated at each time-point.

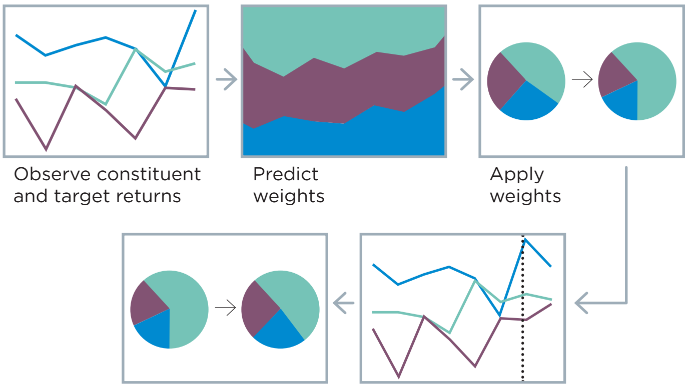

Figure 1

Why is a Kalman filter more appropriate for time-varying constituent analysis?

The first answer is that the weights are not calculated in isolation. Instead the weights evolve in time exactly as a fund manager might change allocations over time. Therefore, the Kalman filter models real world situations while regression does not. The second answer is that the Kalman filter has a proper forecasting mechanism. Using regression for forecasting is extremely questionable since it would imply that the "average of the past describes the immediate future".

Parallel to the Kalman filter, there are several Kalman smoothers, which work in conjunction with the filter and improve the result. The Rauch-Tung-Striebel smoother is used in this article's examples. There are also several extensions to the Kalman filter to deal with non-linearity.

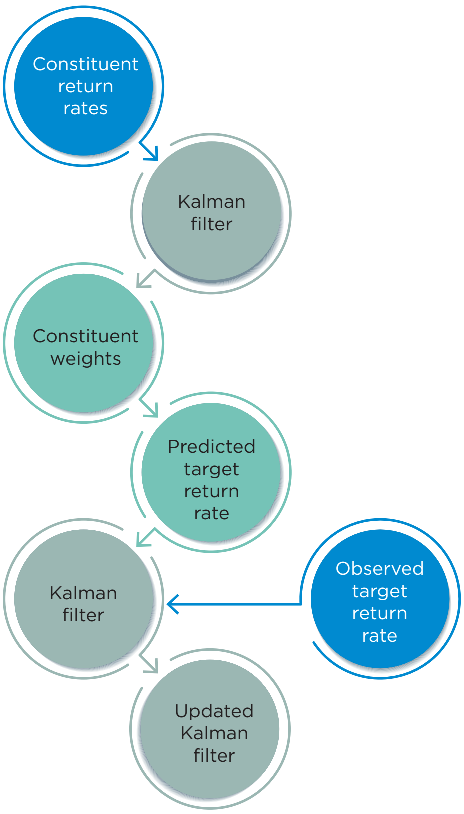

Figure 2: The filtering step

Example 1: Showcasing Kalman Filters

Here we demonstrate the use of the Kalman filter and compare it with regression by looking at the factor/style decomposition of a basket of equity funds. The basket is rebalanced quarterly and consists of nine popular South African equity funds:

- Absa Prime Equity Fund
- Allan Gray Equity Fund
- Ashburton Equity Fund
- Coronation Top 20 Fund
- Investec Equity Fund
- Kagiso Islamic Equity Fund
- Momentum Equity Fund
- SIM General Equity Fund
- SIM Top Choice Equity Fund

All time series are total returns and in South African Rand. The funds are net of fees which makes their returns comparable with the benchmark returns. Since this is a factor decomposition, one would not take trading costs in the tracking portfolio into account. However, this example shows the filter's ability to decompose and to track – thus we included trading costs of 10 bps for buying and selling. We use the factors and factor indices listed in table 1. Many equity funds have a cash or money market allocation, so an appropriate non-equity index should be included in order to accurately track the basket but our primary goal for this example is to demonstrate decomposition. Using a Kalman filter to track a multi-asset fund is similar, but instead the benchmarks would be replaced by tracker products like the Satrix Capped SWIX Fund and the Satrix MSCI World ETF. We use both a Kalman filter and regression to track the basket from October 2014 until December 2019. The regression window used is 24 months. The Kalman filter is a custom implementation in R (due to the many statistical packages available) by SI Client Solutions and Research. The regression was performed using the *quadprog* package from R.

Factor/style	Factor index
Value	J330
Growth	J331
Quality	S&P Quality SA
Momentum	S&P Momentum SA
Low volatility	S&P SA Low Volatility

Table 1: Equity factors

The tracking and performance measures are in tables 2 and 3. The performance plots (rebased to 1) are shown in figure 3. We have chosen this equity decomposition because it is a challenging problem. The Kalman filter can track a multi-asset basket, using the appropriate asset classes, far more accurately resulting in R^2 values of 94%. As mentioned before, one could improve the tracking error by including non-equity components.

Measure	Kalman Filter	Regression
Mean-Squared Error (MSE)	6.7×10^{-5}	5.9×10^{-5}
Tracking Error (TE)	2.9%	2.7%
R^2	92.1%	93.0%
Adjusted R^2	91.4%	92.4%

Table 2: Tracking measures

	Target	Kalman Filter	Regression
Annualised Rate of Return	5.8%	4.7%	4.2%
Annualised Volatility	10.2%	10.6%	10.6%

Table 3: Performance measures

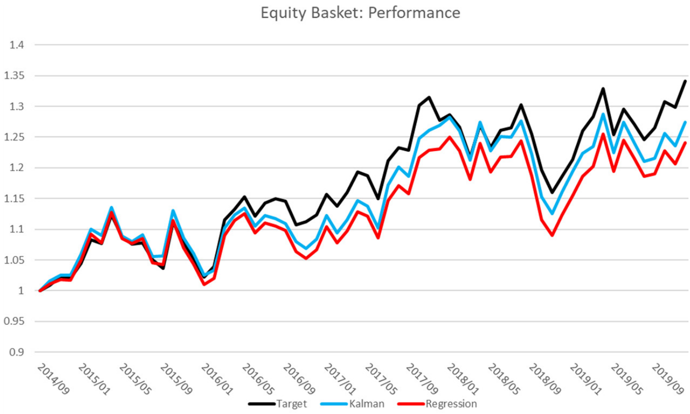

Figure 3: Tracking Performance for the equity basket

The results at this point in the analysis show that regression can also decompose and track. Our main criticism of using regression in this context thus far has been philosophical, but the next graphs depict a practical element: Kalman weights (figure 4) are smoother than regression weights (figure 5). The numbers in table 4 confirm this. This is expected, as the Kalman filter continually adapts weights, whereas regression optimises weights for different time-points in isolation.

The Kalman results presented here were are from a filter aimed at providing stable weights. Changing the forecasting parameter estimation would lead to weights, which are still less volatile than regression weights, but significantly reduce the tracking error.

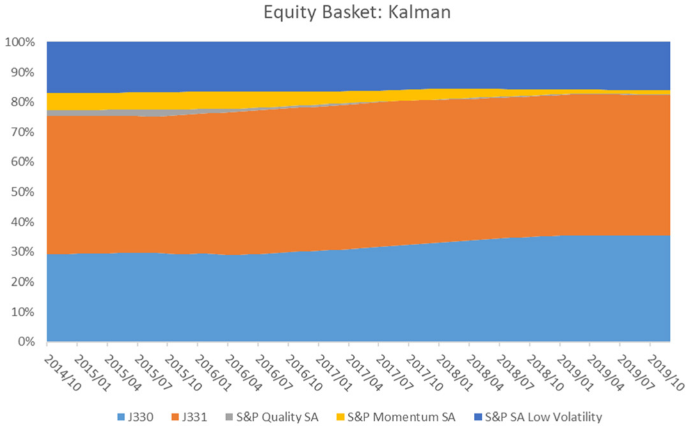

Figure 4: Kalman weights for the equity basket

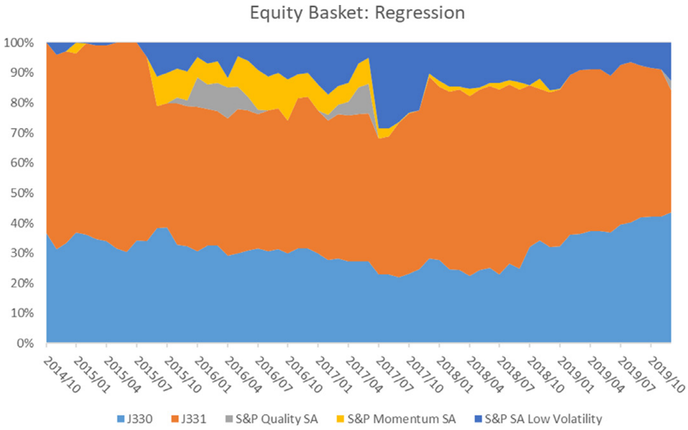

Figure 5: Regression weights for the equity basket

	Weight volatility		Average weight change	
Factor/style	**Kalman Filter**	**Regression**	**Kalman Filter**	**Regression**
Value	2.6%	5.5%	0.1%	1.7%
Growth	0.9%	7.4%	0.1%	2.3%
Quality	0.8%	3.0%	0.1%	0.8%
Momentum	1.7%	4.3%	0.1%	1.2%
Low volatility	0.5%	6.6%	0.0%	2.5%

Table 4: Weight stability

It is important to note what these smooth weights are NOT trying to convey. Firstly, the factor analysis was for a basket. Therefore, the individual investment philosophies play less of a role, leading to the stable weights. Secondly, the funds in the basket are not completely representative of the South African equity market. Doing the same analysis for the SWIX, for example, will lead to clearer cycles, ie less smooth weights.

Example 2: Extreme case – regression failure

Our second example looks at the opposite situation: the desired weights have to change significantly over time. Here the true weights (black line in figure 6) are known (by us not by the filter). Our example is the South African version of a challenging filtering test case from Roncalli & Teïletche (5). We refer the interest reader to this paper for a detailed discussion. The target fund consists of two constituents (Capped SWIX and MSCI World) both in ZAR. These are combined using the weights in figure 6 for Capped SWIX and 1 - these weights for the MSCI World.

Figure 6 shows the true Capped SWIX weights as well as the weights calculated by the Kalman filter and regression. Regression fails completely, whereas the Kalman result is extremely accurate.

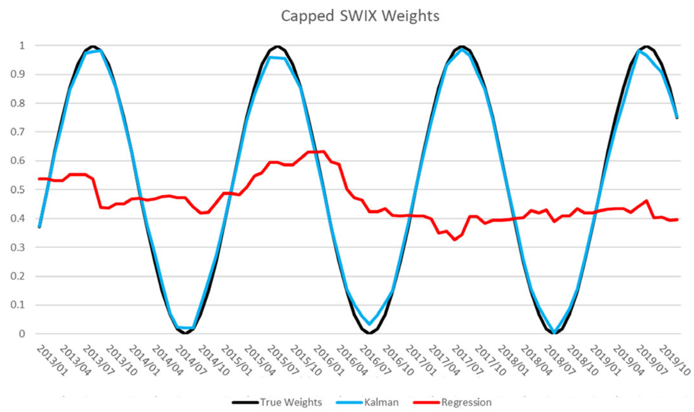

Figure 6

Future approach

Sanlam Investments has developed Kalman filter approaches for asset allocation, factor/style analysis of funds, asset class mappings and tracking. Currently SI is working on risk tools based on Kalman filtering as well as on novel approaches to improve filtering results.

References and further reading

1. Anderson, BDO and Moore, JB. *Optimal Filtering*. Dover Publications, 2005
2. Einicke, GA. *Optimum minimum variance fixed interval smoothing*. Proceedings of the 8th International Symposium on Signal Processing and its Applications, 2005
3. Jazwinski, AH. *Stochastic Processes and Filtering Theory*. Dover Publications, 2007
4. Kalman, RE. *A New Approach to Linear Filtering and Prediction Problems*. Journal of Basic Engineering, 1960
5. Roencalli, T and Teïletche, J. *An Alternative Approach to Alternative Beta*. URL: https://ssrn.com/abstract=1035521, 2007
6. Ruckdeschel, P. Optimally Robust Kalman Filtering. URL: http://publica.fraunhofer.de/dokumente/N-137186.html Berichte des Fraunhofer ITWM 185, 2010