NERA
ECONOMIC CONSULTING

July 2016

# Estimating Equity Betas: What Can a Time-Varying Approach Add?

## A Comparison of Ordinary Least Squares and the Kalman Filter

By **Giulio Renzi-Ricci**[1]

## Introduction

The estimation of systematic risk coefficients (i.e., equity betas or simply betas) has several applications in economics and finance. Equity betas are at the centre of finance theory, being embedded in Modern Portfolio Theory (MPT) and the Capital Asset Pricing Model (CAPM). In particular, betas represent the component of risk to holders of an asset or investment related to general market dynamics, as opposed to idiosyncratic factors. They are important to financial practitioners because they measure the risk that cannot be reduced through diversification.[2]

Betas are also used for purposes other than portfolio theory or asset management. Market models and the estimation of sensitivities of returns to different risk factors are applied, for instance, in risk management, in corporate finance for estimating weighted average cost of capital (WACC), and in counterfactual analysis for assessing damages.

Regardless of the application, it is crucial to estimate reliable betas. The dynamics of financial markets and businesses can lead to structural changes over time in the relationship between the market's performance and that of individual investments. In this context, the standard assumption of a time-invariant beta may potentially be inadequate, as a constant estimate fails to capture the changes over time in systematic (i.e., non-diversifiable) risk. When this limitation is likely to be material, a technique that is often used to take into account changes in the systematic risk of an asset is a rolling-window Ordinary Least Squares (OLS) regression.[3]

This paper reviews a relatively less common method for the dynamic estimation of betas—the Kalman filter. The Kalman filter is a recursive process[4] that refines the model's estimates over time by taking into account the new information it receives. In certain conditions, it provides a useful check on results obtained using rolling-window OLS.

In the second section of this paper, I provide a brief and intuitive description of the Kalman filter, its functioning, and its properties. In the third section, I compare the Kalman filter and OLS regression beta estimates for synthetic time series data. The results show how, under certain circumstances, the Kalman filter compares to standard OLS techniques in detecting a form of structural break in the beta. In the fourth section, I provide an overview of how the Kalman filter can be used in counterfactual analysis and regulatory contexts. Finally, in the last section, I summarise the results, highlighting the strengths and weaknesses of the Kalman filter approach for beta prediction.

## The Intuition Behind the Kalman Filter

The Kalman filter has been extensively used in fields that involve modelling dynamic elements exposed to measurement error, such as control system engineering. More recently, the filter has been applied in economics and finance.

The Kalman filter is a recursive algorithm, i.e., one based on a repeated, updating procedure. It is used to form estimates of an unknown variable[5]—in this case, the beta—that varies over time. The Kalman filter process updates the beta estimates by using the new observed information at each point in time and by measuring a prediction error. Under certain conditions, this technique surmounts certain limitations that OLS techniques have in measuring time-varying betas.

The Kalman filter is based on two building blocks: 1) a measurement equation and 2) a transition equation.

- The measurement equation relates the unknown variable (i.e., beta) to observable variables (i.e., the stock and market returns, $R_s$ and $R_m$, respectively). In this case, the measurement equation is a standard market model.[6]

- The transition equation allows the beta to change over time through an autoregressive process.[7] In other words, the transition equation relates the beta in any period ($t$) to the beta in the previous period ($t - 1$). This relationship is captured by the parameter $T$. A second element ($\theta_t$), which is an error term, constitutes the random component of the change in the beta.

For the purposes of this paper, the mathematical representation that forms the measurement and transition equations can be specified as:

*Measurement Equation*: $R_{s,t} = \alpha + \beta_t R_{m,t} + \varepsilon_t$

*Transition Equation*: $\beta_{t+1} = T_t \beta_t + \theta_t$

where $\varepsilon_t \sim \mathrm{N}(0, h_t)$ and $\theta_t \sim \mathrm{N}(0, q_t)$.[8]

More specifically, the Kalman filter algorithm can be summarised as:

- Step 1: The filter starts the procedure with the current estimate of the unknown variable ($\beta_t$) in the transition equation as well as some initial "guess" values for $T_t$ and the accompanying errors in the transition and measurement equations.[9] Given these initial values, the filter calculates the best *ex ante* estimate of beta for the following period ($\widehat{\beta_{t+1}}$).

- Step 2: Given $\widehat{\beta_{t+1}}$, the filter estimates the predicted stock return for the following period ($\widehat{R_{s,t+1}}$) by plugging $\widehat{\beta_{t+1}}$ in the measurement equation and using the known observed value of the market return ($R_{m,t+1}$).

- Step 3: Following an observation of the actual stock return ($R_{s,t+1}$), the model calculates the prediction error, which is defined as the difference between the observed and the predicted stock return ($R_{s,t+1} - \widehat{R_{s,t+1}}$).

- Step 4: Finally, the model adjusts the beta prediction by allowing part of the prediction error to feed through in the adjusted beta ($\widehat{\beta_{t+1,\,Adj}}$).[10] The adjusted beta is then used in Step 1 as $\beta_{t+1}$, and the process starts over again.

The algorithm is then solved recursively via a Maximum-Likelihood Estimation (MLE),[11] which identifies the values of the unknown parameters that minimise the prediction error.

## Application to Synthetic Data

In this section, I discuss a synthetically created example in order to compare the predictive power and stability of three different techniques for dynamic beta estimation:

1. Entire-period OLS: an OLS regression where the model's coefficients are estimated using all available observations in the synthetic data sample up to a specific point in time (e.g., the beta estimate at time 30 is estimated using all observations from 1 to 30, the beta estimate at time 31 is estimated using all observations from 1 to 31, and so on);

2. Rolling-window OLS: an OLS regression where the model's coefficients are estimated using overlapping rolling windows consisting of the same number of observations (e.g., a 30-data points rolling window); and

3. Kalman filter: as described in section 2.

I generated a sample pair of 1,000 market ($R_{m,t}$) and asset ($R_{s,t}$) returns with a known relationship, where for the first 500 observations, the true beta was assigned to be equal to 3 and for the subsequent 500 observations, it was assigned to be equal to 6.[12] In summary, the specified model is the following:[13]

$$\begin{cases} R_{s,t} = \beta R_{m,t} + \varepsilon_t \text{ for } t \leq 500, \text{ where } \beta = 3 \\ R_{s,t} = \beta R_{m,t} + \zeta_t \text{ for } t > 500, \text{ where } \beta = 6 \end{cases}$$

where $\varepsilon_t$ and $\zeta_t$ were generated as two vectors of normally distributed random numbers.

Thus, I have created a hypothetical market model where I know what the *actual* beta is and how it changes over time. In fact, I constructed the market model such that there is a single and persistent shift in the beta coefficient (i.e., a jump in level) precisely in the middle of the time period of the synthetic data sample.

After setting up the model, I ran the Kalman filter algorithm, the entire-period OLS regression, and the rolling-window OLS regression using a 30-trading day (i.e., 30 observations) rolling window. At each point in time (i.e., data observation), I recorded the estimated beta coefficient of the market model. Figure 1 shows the resulting betas derived over time, across the three methods.[14]

Figure 1. **Estimated Market Model Betas**

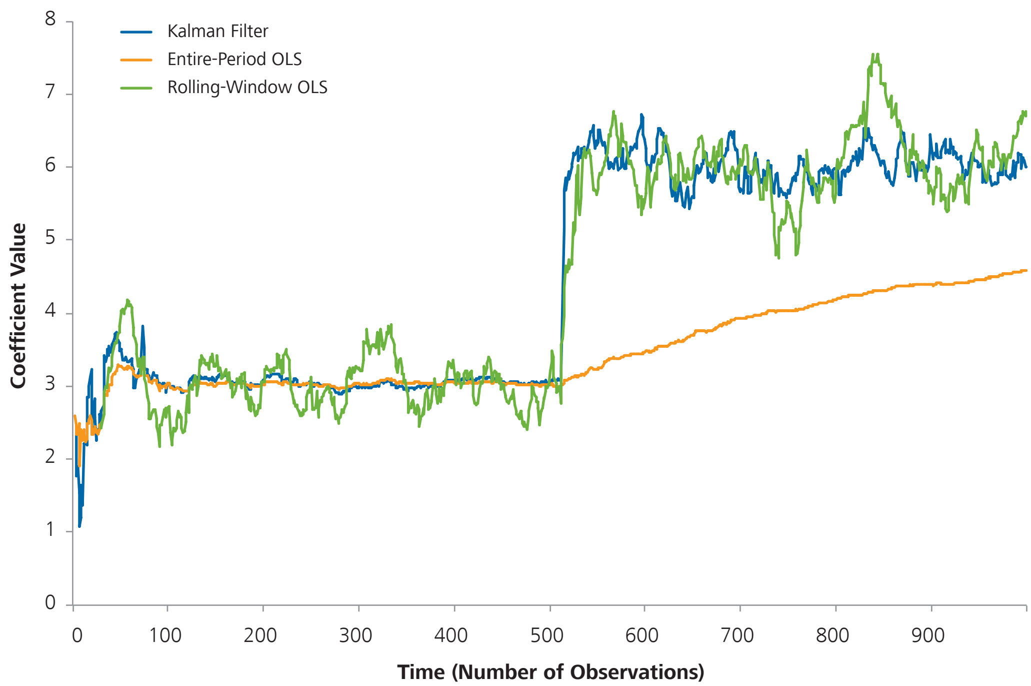


As depicted in Figure 1, the entire-period OLS estimated beta does not converge to the actual value. The estimated coefficient does react to the change in the systematic risk (i.e., the orange line starts to increase after observation 500)—however, this happens very slowly. Even after 500 additional observations (i.e., at 1,000 observations), the entire-period OLS beta estimate has still not converged to the actual beta of six. This is due to the fact that the entire-period OLS includes all observations that are previous to the change in the beta. In contrast, both the Kalman filter and the rolling-window OLS's estimates react quickly to the structural change in the beta coefficient. However, the Kalman filter's estimate adjusts more quickly and appears to also be more stable (i.e., lower level of volatility) than the estimate derived through the rolling-window OLS.[15]

This is confirmed by Figure 2, where I have plotted the Mean-Squared Error (MSE) of the market model's beta estimated using the Kalman filter and the rolling-window OLS.[16] In this case, the MSE compares over time the estimated beta with the actual known parameter coefficient. A lower MSE value would indicate that the model provides, on average, a more precise prediction of the estimated beta (i.e., with a lower level of dispersion around its real value). As shown by Figure 2, aside from an initial phase in which the Kalman filter is still adjusting its predictions, the MSE estimated through the Kalman filter is lower than the MSE of the beta estimated using a 30-observations rolling-window OLS regression.[17]

Figure 2. **Mean-Squared Error of Market Model Betas**

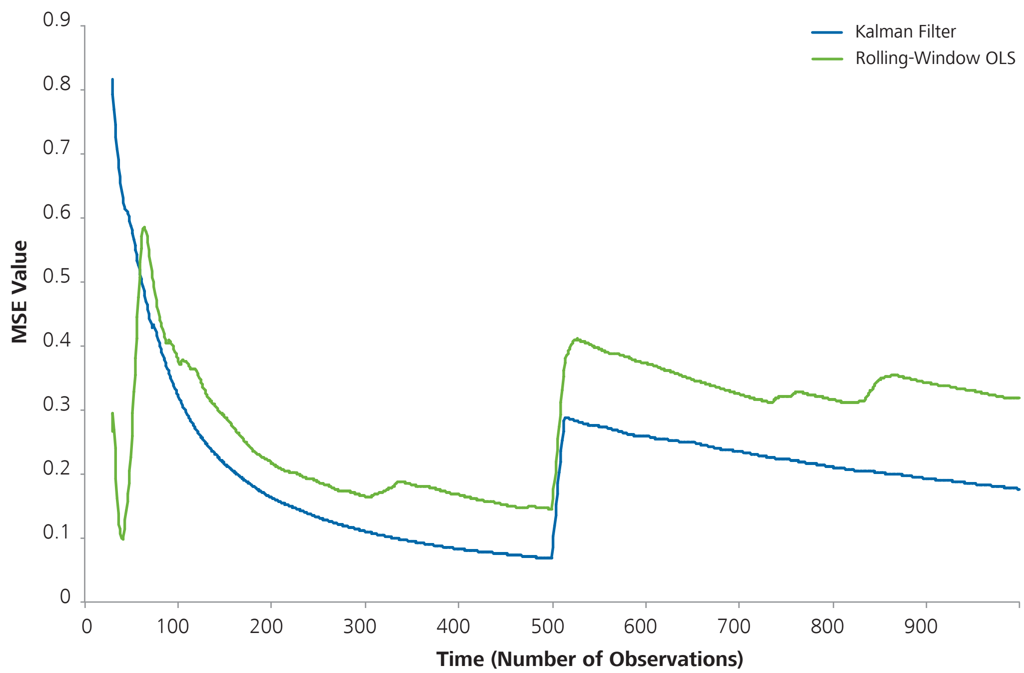


In summary, once possible shifts in the systematic risk of an asset are taken into account, the Kalman filter algorithm seems to provide—at least in a controlled environment using one particular form of a shift—more robust estimates for market model coefficients. In fact, the Kalman filter's beta estimates converge more quickly and with more precision to the actual values, compared to estimates derived using classical OLS regression techniques for the particular shift examined.[18]

## Application to Financial Disputes and Regulation

As mentioned in the previous sections, a reliable estimate of betas is critically important to the construction of any market model. Market models are widely used by experts in economics, finance-related disputes, and regulatory finance for various purposes including, but not limited to:

- Event studies[19] for collective actions by investors, estimating the but-for performance of the share or other asset price. For example, an expert might use the Kalman filter for assessing loss causation and materiality of a stock price drop, following the disclosure of a patent infringement;
- Assessing the damages related to loss of profits in a counterfactual analysis. For example, the Kalman filter could be employed for the counterfactual estimation of a fund’s returns assuming that a programming error in the investment algorithm of an asset manager had not occurred;
- Analysis of Key Performance Indicators (KPIs) of a fund, including: counterfactual assets under management, management and performance fees, and capital inflows;
- Estimation of cost of credit of bidding banks in order to assess potential market manipulation; and
- Estimation of equity betas and related cost of capital of a company for regulatory enquiries.

## Conclusions

The Kalman filter can be a valuable supplement to standard econometric techniques (such as rolling-window OLS) for detecting structural changes in systematic risk and estimating equity betas.

An advantage of the Kalman filter is its ability to adjust quickly to new information. More specifically, at least in the one type of shift examined here, the Kalman filter can adjust to the actual beta value more quickly than standard OLS approaches. This advantage is even more pronounced when compared to “long-term” OLS estimates (e.g., one-year or two-year data window). In addition, this responsiveness does not necessarily come at a cost in the form of higher uncertainty or volatility of the estimated beta, as is generally the case with “short-term” OLS estimates (e.g., 30-day data window).

However, the Kalman filter does present some limitations. First, it can be sensitive to initial parametrisation and take some time to converge when the dataset is small. Second, the Kalman filter betas can be more volatile than estimates derived using long-term OLS regressions. Third, the technique is also more computationally complex and takes more processing time than standard regression techniques.

I think that the Kalman filter approach to beta estimation should not be used in isolation. The Kalman filter does, however, provide economic experts a valuable cross-check to OLS beta estimates, especially when the context suggests changes in systematic risk during the relevant period.

## Bibliography


1. Alexander, Carol, *Market Models: A Guide to Financial Data Analysis*, John Wiley & Sons, 2001.
2. Arnold, Tom, Bertus, Mark J., and Godbey, Jonathan, "A Simplified Approach to Understanding the Kalman Filter Technique," *The Engineering Economist*, 53 (2), December 2007, p. 140-155.
3. Berardi, Andrea, Corradin, Stefano, and Sommacampagna, Cristina, *Estimating Value at Risk with the Kalman Filter*, Working Paper, 2002, available at: http://quantlabs.net/academy/download/free_quant_instituitional_books_/Estimating_Value_at_Risk_with_Kalman_Filter_Berardi_Corradin_Sommacampagna.pdf.
4. Brooks, Chris, *Introductory Econometrics for Finance, 2d ed*., Cambridge University Press, 2008, Chap. 2.
5. Callaghan, Joseph, Fu, Liang, and Liu, Jing, "Systematic Risk Estimation: OLS v. State-Space Methods," *International Research Journal of Applied Finance*, 3 (7), July 2012, p. 940-953.
6. Commandeur, Jacques and Koopman, Siem Jan, *An Introduction to State Space Time Series Analysis*, Oxford University Press, 2007.
7. Das, Atanu and Ghoshal, Tapan Kumar, "Market Risk Beta Estimation Using Adaptive Kalman Filter," *International Journal of Engineering Science and Technology*, 2 (6), 2010, p. 1923-1934.
8. Elton, Edwin, Gruber, Martin, Brown, Stephen, and Goetzmann, William, *Modern Portfolio Theory and Investment Analysis, 6th ed*., John Wiley & Sons, 2003.
9. Faff, Robert W., Hillier, David, and Hillier Joseph, "Time Varying Beta Risk: An Analysis of Alternative Modelling Techniques," *Journal of Business Finance & Accounting*, 27 (5) and 27 (6), June/July 2000, p. 523-554.
10. Faragher, Ramsey, "Understanding the Basis of the Kalman Filter Via a Simple and Intuitive Derivation," *IEEE Signal Processing Magazine*, September 2012, p. 128-132, available at: http://www.cl.cam.ac.uk/~rmf25/papers/Understanding%20the%20Basis%20of%20the%20Kalman%20Filter.pdf.
11. Hamilton, James D., "State-Space Models," *Handbook of Econometrics, Volume IV*, ed. R.F. Engle and D.L. McFadden, Elsevier Science B.V., 1994.
12. Hern, Richard and Zalewska, Ania, "Recent Evidence on Beta and the Cost of Capital for UK Electricity Companies," *National Economic Research Associates – Topics*, 25, October 2001.
13. Lie, Frida, Brooks, Robert, and Faff, Robert, "Modelling the Equity Beta Risk of Australian Financial Sector Companies," *Australian Economics Papers*, 39 (3), 2002, p. 301-311.
14. Park, Kwang Woo, "Time-Series Analysis of Return and Beta in U.S.," *Journal of Academy of Business and Economics*, 3, 2004, p. 43-54.
15. Tabak, David and Dunbar, Frederick C., *Materiality and Magnitude: Event Studies in the Courtroom*, NERA Working Paper No. 34, April 1999, available at: http://www.nera.com/content/dam/nera/publications/archive1/3841.pdf.
16. Wells, Curt, "Variable Betas on the Stockholm Exchange 1971-1989," *Applied Financial Economics*, 4, 1994, p. 75-92.

## References

[1] Mr. Renzi-Ricci is a Consultant at NERA in London. The author would like to thank David Tabak, Robert Patton, Timothy McKenna, and Erin McHugh for valuable feedback on earlier drafts.

[2] In the CAPM, beta captures the only type of risk that should be rewarded with an expected return higher than the risk-free rate. The market portfolio of all assets (e.g., market benchmark) has a beta equal to one. For an individual security, a beta lower than one indicates that market-wide events tend to cause the asset to move less than the market portfolio. Conversely, a beta greater than one means that market-wide events will tend to cause the asset to move more than the market portfolio. For further details on the CAPM betas, see Elton et al. (2003).

[3] A rolling-window regression for time series data implies using different overlapping (or non-overlapping) data ranges referring to different time intervals. For example, at each point in time (i.e., data observation), the estimation might be performed over the prior 30 observations. For an introduction to OLS regression and its applications to finance, see Brooks (2008).

[4] The Kalman filter applies the same rule and algorithm for each new data observation in order to update its best estimate. See section 2 for further details.

[5] It is considered "unknown" because I am making an estimate of what it is most likely to be.

[6] In this paper, I will be estimating the beta of a market model. The market model is a factor model extensively used by practitioners in regulatory and financial litigation contexts. A market model establishes a linear relationship between the return on an asset and the return on the market portfolio. For further details, refer to Alexander (2001). A similar approach can be used with other measurement equations that refer to different asset pricing models, such as the CAPM or multifactor models.

[7] An autoregressive process is one where the current value of a variable depends only on its own past values and an error component. For further details, see Brooks (2008).

[8] These are normally distributed error terms with zero mean and variances equal to $h_t$ and $q_t$, respectively.

[9] $T_t$ varies over time such that the Maximum-Likelihood Estimation process is optimised. For further details, see Hamilton (1994).

[10] The fraction of the prediction error that is used to estimate the adjusted beta is sometimes called "Kalman Gain."

[11] The Maximum-Likelihood Estimation is a statistical technique that selects those values, for some given model's parameters, that maximise the probability of obtaining the observed data.

[12] These two coefficients have been chosen for illustrative purposes only, and they are not meant to be representative of typical equity market betas. However, a single and persistent structural shift in the beta coefficient, as in the example presented in this paper, can happen in financial markets. For instance, a change in the systematic risk of a company could be related to a change in the level of competition in its industry, a change in the level of its financial leverage, or entry into a new market.

[13] This model is the one presented in the measurement equation above. Here, for illustrative purposes, I assumed the constant component ($\alpha$) to be equal to zero.

[14] All figures and results in this paper refer to one particular simulation run of the hypothetical model specified above. I also ran several hundred Monte Carlo simulations of synthetic data series based on the same parametrisation. The sample result presented is representative of the average result of all simulations.

[15] The rolling-window OLS can be made to react more quickly by reducing the number of observations in the window. However, this will tend to increase the volatility of the rolling-window OLS estimates.

[16] The MSE is a statistical measure assessing the quality of an estimator. It is often used for comparing the level of accuracy of different models; see, for instance, Faff et al. (2000) and Callaghan et al. (2012) for further details.

[17] The rolling-window OLS method is static since it implicitly gives —within the time period window—the same weight to each observation. Hence, the estimated value may significantly depend on the length of the data series that is used. The Kalman filter technique adjusts for this limitation by upgrading model estimates using new information at each point in time. In this regard, the Kalman filter is similar to a weighted OLS that implicitly gives more weight to observations that are more informative.

[18] My results and conclusions might not apply to other types of structural breaks. In fact, a one-off persistent shift in the beta, like the one analysed in this paper, is only one of a multitude of different types of structural breaks that could happen to actual (i.e., non-synthetic) time series data. Beta coefficients could, for instance, be subject to trend changes or mean-reverting processes where the Kalman filter might perform similarly to, better than, or worse than more standard OLS techniques.

[19] An event study is a statistical technique assessing the impact of a specific event on an asset price. An event study separates the "abnormal" component of the price movement (i.e., the part that is security-specific and thus may be attributable to the event being studied) from the component that relates to market fluctuations. For further information on the how event studies can be applied by economic experts in dispute resolution, see Tabak et al. (1999).

NERA
ECONOMIC CONSULTING

## About NERA

NERA Economic Consulting (**www.nera.com**) is a global firm of experts dedicated to applying economic, finance, and quantitative principles to complex business and legal challenges. For over half a century, NERA's economists have been creating strategies, studies, reports, expert testimony, and policy recommendations for government authorities and the world's leading law firms and corporations. We bring academic rigor, objectivity, and real world industry experience to bear on issues arising from competition, regulation, public policy, strategy, finance, and litigation.

NERA's clients value our ability to apply and communicate state-of-the-art approaches clearly and convincingly, our commitment to deliver unbiased findings, and our reputation for quality and independence. Our clients rely on the integrity and skills of our unparalleled team of economists and other experts backed by the resources and reliability of one of the world's largest economic consultancies. With its main office in New York City, NERA serves clients from more than 25 offices across North America, Europe, and Asia Pacific.

## Contact

For further information and questions, please contact the author:

**Giulio Renzi-Ricci**
Consultant
+44 20 7659 8546
giulio.renzi.ricci@nera.com