A Simplified Approach to Understanding the Kalman Filter Technique

Tom Arnold
(Contact Author)
The Robins School of Business
Department of Finance
University of Richmond
Richmond, VA 23173
O: 804-287-6399
F: 804-289-8878
tarnold@richmond.edu

Mark Bertus
Department of Finance
Lowder Business Building
415 West Magnolia Avenue, Suite 303
Auburn University, AL 36849
O: 334-844-3004
F: 334-844-4960
bertumj@auburn.edu

and

Jonathan Godbey
Department of Finance
MSC 0203
James Madison University
Harrisonburg, VA 22807
O: 540-568-3074
F: 540-568-3017
godbeyjm@jmu.edu

December 21, 2007

Preliminary, do not cite without permission

A Simplified Approach to Understanding the Kalman Filter Technique

The Kalman Filter is a time series estimation algorithm that is applied extensively in the field of engineering and recently (relative to engineering) in the field of finance and economics. However, presentations of the technique are somewhat intimidating despite the relative ease of generating the algorithm. This paper presents the Kalman Filter in a simplified manner and produces an example of an application of the algorithm in Excel. This scaled down version of the Kalman filter can be introduced in the (advanced) undergraduate classroom as well as the graduate classroom.

1

INTRODUCTION:

Many models in economics and finance depend on data that are not observable. These unobserved data are usually in a context in which it is desirable for a model to predict future events. The Kalman Filter has been used to estimate an unobservable source of jumps in stock returns, unobservable noise in equity index levels, unobservable parameters and state variables in commodity futures prices, unobservable inflation expectations, unobservable stock betas, and unobservable hedge ratios across interest rate contracts[1]. In the field of engineering a Kalman Filter (Kalman, 1960) is employed for similar problems involving physical phenomena. The technique is appearing more frequently in the fields of finance and economics. However, understanding the technique can be very difficult given the available resource material.

When viewing chapter thirteen of Hamilton's *Times Series Analysis* text (1994), one can understand why the topic of Kalman Filters is generally reserved for the graduate classroom. However, as we will demonstrate, the technique is not quite as difficult as one may perceive initially and has similarities to standard linear regression analysis. Consequently, if placed in the correct context, it is accessible to the undergraduate student. In order to make the Kalman Filter more accessible, an Excel application is developed in this paper to work the student through the mechanics of the process.

In the first section, a derivation of the Kalman Filter algorithm is presented in a univariate context and a connection is made between the algorithm and linear regression. In the second section, the Kalman Filter is combined with Maximum Likelihood Estimation (MLE) to create an iterative process for parameter estimation. In the third

---

[1] See Bertus, Beyer, Godbey and Hinkelmann (2006), Bertus, Denny, Godbey and Hinkelmann (2006), Burmeister and Wall (1982), Burmeister, Wall, and Hamilton (1986), Faff, Hillier, and Hillier (2000), Fink, Fink, and Lange (2005), Godbey and Hilliard (forthcoming), and Schwartz (1997).

section, an Excel application/example of using the Kalman Filter/MLE iterative routine is performed.

SECTION 1: DEVELOPING THE KALMAN FILTER ALGORITHM

There are two basic building blocks of a Kalman Filter, the measurement equation and the transition equation. The measurement equation relates an unobserved variable (X_t) to an observable variable (Y_t). In general, the measurement equation is of the form:

$$Y_t = m_t * X_t + b_t + \varepsilon_t \tag{1}$$

To simplify the exposition, assume the constant "b_t" is zero and "m_t" remains constant through time eliminating the need for a subscript. Further, "ε_t" has a mean of zero and a variance of "r_t". Equation (1) becomes:

$$Y_t = m * X_t + \varepsilon_t \tag{2}$$

The transition equation is based on a model that allows the unobserved variable to change through time. In general, the transition equation is of the form:

$$X_{t+1} = a_t * X_t + g_t + \theta_t \tag{3}$$

Again, to simplify the exposition, assume the constant "g_t" is zero and "a_t" remains constant through time eliminating the need for a subscript. Further, "θ_t" has a mean of zero and a variance of "q_t". Equation (3) becomes:

$$X_{t+1} = a * X_t + \theta_t \tag{4}$$

To begin deriving the Kalman Filter algorithm, insert an initial value, "X_0" into equation (4) (the transition equation) for "X_t". "X_0" has a mean of "μ_0" and a standard deviation of "σ_0". It should be noted: "ε_t", "θ_t", and "X_0" are uncorrelated (Note: these variables are also uncorrelated relative to lagged variables). Equation (4) becomes:

$$X_{1P} = a * X_0 + \theta_0 \tag{5}$$

3

where, "X_{1P}" is the predicted value for "X_1"

"X_{1P}" is inserted into equation (2) (the measurement equation) to get a predicted value for "Y_1", call it "Y_{1P}":

$$Y_{1P} = m * X_{1P} + \varepsilon_1 = m * [a * X_0 + \theta_0] + \varepsilon_1 \qquad (6)$$

When "Y_1" actually occurs, the error, "Y_{1E}", is computed by subtracting "Y_{1P}" from "Y_1":

$$Y_{1E} = Y_1 - Y_{1P} \qquad (7)$$

The error can now be incorporated into the prediction for "X_1". To distinguish the adjusted predicted value of "X_1" from the predicted value of "X_1" in equation (5), the adjusted predicted value is called "$X_{1P\text{-}ADJ}$":

$$\begin{aligned}
X_{1P-ADJ} &= X_{1P} + k_1 * Y_{1E} \\
&= X_{1P} + k_1 [Y_1 - Y_{1P}] \\
&= X_{1P} + k_1 [Y_1 - m * X_{1P} - \varepsilon_1] \\
&= X_{1P} [1 - m * k_1] + k_1 * Y_1 - k_1 * \varepsilon_1
\end{aligned} \qquad (8)$$

where "k_1" is the Kalman gain, which will be determined shortly

The Kalman gain variable is determined by taking the partial derivative of the variance of "$X_{1P\text{-}ADJ}$" relative to "k_1" in order to minimize the variance based on "k_1" (i.e. the partial derivative is set to zero and then one finds a solution for "k_1"). For ease of exposition, let "p_1" be the variance of "X_{1P}" (technically, "p_1" equals: $(a * \sigma_0)^2 + q_0$). The solution for the Kalman gain is as follows (see Joseph (2007) for a numerical example):

$$Var(X_{1P-ADJ}) = p_1 * [1 - m * k_1]^2 + k_1^2 * r_1 \qquad (9)$$

$$\frac{\partial Var(X_{1P-ADJ})}{\partial k_1} = -2m * [1 - m * k_1] * p_1 + 2 * k_1 * r_1 = 0 \qquad (10)$$

4

$$\therefore k_1 = \frac{p_1 * m}{(p_1 * m^2 + r_1)} = \frac{Cov(X_{1P}, Y_{1P})}{Var(Y_{1P})} \tag{11}$$

Notice, the Kalman gain is equivalent to a β-coefficient from a linear regression with "X_{1P}" as the dependent variable and "Y_{1P}" as the independent variable. Not that one would have a sufficient set of data to perform such a regression, but the idea that a β-coefficient is set to reduce error in a regression is equivalent to the idea of the Kalman gain being set to reduce variance in the adjusted predicted value for "X_1".

The next step is to use "$X_{1P\text{-}ADJ}$" in the transition equation (equation (4)) for "X_t" and start the process over again to find equivalent values when t = 2. However, before ending this section, it is important to note the advantages of "$X_{1P\text{-}ADJ}$" over "X_{1P}". Recall, the variance for "X_{1P}" is "p_1". Substituting equation (11) into equation (9), the variance of "$X_{1P\text{-}ADJ}$" is:

$$Var(X_{1P\text{-}ADJ}) = p_1 * \left[1 - \frac{1}{\left(1 + \frac{r_1}{p_1 * m^2} \right)} \right]^2 + k_1^2 * r_1 \tag{12}$$

The portion of the equation that pertains to the variance of "X_{1P}", i.e. "p_1", has a bracketed term that is less than one (and is further reduced because the "less than one quantity" is squared). This means the portion of the variance attributed to estimating "X_1" has been reduced by using "$X_{1P\text{-}ADJ}$" instead of "X_{1P}".

For reference, the Kalman Filter algorithm is summarized in the table below:

(INSERT TABLE 1 HERE)

In the next section, it will be necessary to use the mean and variance of "X_{1P-ADJ}" and of "Y_{1P}". Although, some of these quantities have already been calculated, all are presented below for reference purposes with the time index variable "t" incorporated ($t = 1$ to T) and the adjusted predicted values for "X_t" incorporated into "Y_{tP}":

$$E\left[X_{tP-ADJ}\right] = E\left[X_{tP} + k_t * Y_{tE}\right] = E\left[X_{tP}\right] + k_t * \left(Y_t - E\left[Y_{tP}\right]\right) \tag{13}$$

$$Var\left[X_{tP-ADJ}\right] = p_t * \left[1 - \frac{1}{\left(1 + \frac{r_t}{p_t * m^2}\right)}\right]^2 + k_t^2 * r_t \tag{14}$$

$$E\left[Y_{tP}\right] = E\left[m * \left(X_{tP-ADJ}\right) + \varepsilon_t\right] = m * E\left[X_{tP-ADJ}\right] \tag{15}$$

$$Var\left[Y_{tP}\right] = Var\left[X_{tP-ADJ}\right] * m^2 + r_t \tag{16}$$

Note: ε_t technically appears within equation (13) within the Y_{tE} term and within equation (15), however, these error terms are independent of each other. In other words, equation (15) and (16) refer to an "updated" or "adjusted" version of the Y_{tP} term in equations (13) and (14). Consequently, the error terms corresponding to Y_{tP} within the two sets of equations are uncorrelated.

In the classroom setting, it is important to keep the application in a univariate setting initially to allow the student to follow the logic of the filter. Further, it is suggested that the instructor reinforce the logic of the algorithm using Table 1 in conjunction with an assignment (such as the assignment developed in section three of this paper) or a quiz. Because this presentation is not reliant on many expectation calculations and only one variance calculation, it is a more palatable introduction of the Kalman Filter than what many texts present. Consequently, this presentation works best as an introduction to the technique which can eventually lead to the more sophisticated presentations available in many time series texts. If the instructor only requires an

introduction to the Kalman Filter technique with the ability to create an assignment then this presentation of the algorithm will be sufficient without a text.

SECTION 2: APPLYING MAXIMUM LIKELIHOOD ESTIMATION TO THE KALMAN FILTER

The Kalman Filter provides output throughout the time series in the form of estimated values for an unobservable variable: "$X_{tP\text{-}ADJ}$" with a mean and a variance defined in equations (13) and (14). Further, the observable variable has a time series of values and a distribution based on its predicted value, "Y_{tP}", which has a mean and variance defined by equations (15) and (16). What the Kalman Filter cannot determine are unknown model parameters in the measurement equation, "ε_t", in equation (2) (note: "m" is a constant and assumed known) and unknown parameters in the transition equation, "a" and "θ_t", in equation (4). Consequently, it is necessary to have a means of estimating these parameters and when estimated, allow the Kalman Filter to generate the time series of the unobservable variable that is desired.

If we assume that the distribution for each "Y_{tP}" is serially independent and normally distributed with a mean and variance as defined by equations (15) and (16); note: the mean and variance both incorporate the mean and variance of the unobservable variable "$X_{tP\text{-}ADJ}$", a joint likelihood function emerges:

$$\prod_{t=1}^{t=T}\left\{\left[\frac{1}{\sqrt{2\pi * Var[Y_{tP}]}}\right]^{T} e^{-\frac{\sum_{t=1}^{T}(Y_t - E[Y_{tP}])^2}{2 * Var[Y_{tP}]}}\right\} \tag{17}$$

The idea behind the likelihood function is that the observable data emerges from this jointly normal distribution. Consequently, the parameters to be estimated within the distribution are chosen in a manner that maximizes the value of the likelihood function (i.e. provides the highest probability that the observed data actually occur). To simplify calculations using the likelihood function, it is common to use the natural logarithm of the likelihood function (i.e. the log-likelihood function):

$$-\frac{T*\ln(2\pi)}{2} - \frac{1}{2}\sum_{t=1}^{T}\ln[Var[Y_{tP}]] - \frac{1}{2}\sum_{t=1}^{T}\frac{(Y_t - E[Y_{tP}])^2}{Var[Y_{tP}]} \qquad (18)$$

As mentioned previously, the parameters of interest are "ε_t", "a", and "θ_t" which may be constants or defined by a distribution (the parameters of the distribution that generates the variable then become the parameters of interest instead of the variable). Further simplifying assumptions may be employed, for example, the variance of "ε_t" and "θ_t" will be constant through time (i.e. "q_t" = "q" and "r_t" = "r"). The partial derivative of the log-likelihood function with respect to each parameter is calculated and set to zero in order to maximize the log-likelihood function.

After a set of parameters is estimated (these are called maximum likelihood estimates or MLEs), the Kalman Filter algorithm is applied again which will produce new time series of "Y_{tP}" and "X_{tP-ADJ}" with associated distributions. The likelihood estimation is then performed again producing new MLEs which will again enter into the Kalman Filter. This iterative process will continue until the value of equation (18) does not improve by a significant amount (say 0.0001). In this context, equation (18) is often referred to as "the score". The use of maximum likelihood estimation in conjunction with the Kalman Filter in an iterative fashion is referred to as the Expectation Maximization (EM) algorithm (see Brockwell and Davis (2002)).

Notice, the EM algorithm is critical in the final estimation of the unobserved data, however, it is not essential to understanding the Kalman Filter process. Consequently, one can choose to treat the EM algorithm in a cursory fashion depending on the level of the class. In the next section, the Kalman Filter with the EM algorithm are applied together in Excel. The application allows the student to work through the Kalman Filter process, define the MLE equation, and then execute the entire system using Excel's "Solver" function. The "Solver" function performs the EM algorithm with minimal input from the student.

SECTION 3: A NUMERICAL EXAMPLE IN EXCEL

The remaining portion of this presentation is a simple example of the EM algorithm. This example presents an iterative computation of the Kalman Filter and the maximum likelihood estimation when the observations can be viewed as incomplete data. In particular, to illustrate the usefulness of the Kalman filter, we analyze its application in a pricing model framework for a commodity spot and futures market.

3.1   Description of the Example

Consider an agent who participates in the oil market. This market participant may buy and sell oil in two different markets; the spot market and the futures market. When buying (selling) oil in the spot market, the trader is looking to take (lose) immediate possession of oil. Alternatively, if the agent does not have an immediate need for oil, but does at sometime in the future, the trader may arrange today to take ownership of oil at the later date by purchasing a futures contract today. The value of this futures contract today of course will then depend on the current spot price of oil, the time period of the

agreement and some time value of money factor (we simplify the time value of money factor to only incorporate the risk-free rate). That is, the relationship between the spot price (S_t) and futures price (F_t) is given by

$$F_t = S_t \exp(r\tau).$$ (19)

where, "r" is the annual risk-free rate and "τ" is the time to maturity of the contract measured in years

Equation (19) has an important practical function for the crude oil market. Spot market crude oil does not have a single organized trading floor and therefore does not have an observable spot price. Crude oil futures, however, do actively trade on an organized exchange and are observable. Given the relationship between the spot price and the futures price in equation (19), traders can use the Kalman Filter to accurately infer spot price levels of crude oil. To estimate these unobserved spot prices we need the pricing relation from equation (19) and the underlying dynamics of the spot prices.

For simplicity, assume the spot price follows geometric Brownian motion:

$$dS_t = \mu S_0 dt + \sigma S_0 dZ_t$$ (20)

$$dZ_t \sim N[0, dt]$$ (21)

where, "dt" is an infinitesimally small step forward in time

Because "dZ_t" is distributed normally with a zero mean and a variance of "dt", "dS_t" also follows a normal distribution:

$$dS_t \sim N\left[\mu S_0 dt, (\sigma S_0)^2 dt\right]$$ (22)

Although correct in its present form, it is much easier to utilize a linear form of the relationship by taking the natural logarithm of both sides of equation (19) and adding

an error term (ε_t is an error term with $E[\varepsilon_t] = 0$ *and* $Var[\varepsilon_t] = q_t$). We now have the

measurement equation for the Kalman Filter

$$\ln(F_t) = \ln(S_t) + r\tau + \varepsilon_t \tag{23}$$

For notational ease let $X_t \equiv \ln(S_t)$ and $Y_t \equiv \ln(F_t)$ where "t" indicates a point in

time. The measurement equation is similar to equation (1) with "Y_t" equal to "$\ln(F_t)$",

"X_t" equal to "$\ln(S_t)$", "b_t" equal to "$r\tau$", and a similarly defined error term. By Ito's

lemma (if *dX = a*dt + b*dZ* and *Q = f(X)*, then *dQ = [a*Q_X + 0.5*b²*Q_XX + Q_t]*dt +*

b*Q_X*dZ, subscripts indicate partial derivatives; see Arnold and Henry (2003) for a more

expansive explanation of Ito's lemma in the context of asset prices):

$$dX_t = \left(\mu - 0.5\sigma^2\right)dt + \sigma dZ_t \tag{24}$$

Equation (24) implies that $S_t = S_0 \exp\left\{\left(\mu - 0.5\sigma^2\right)\tau + \sigma\int_0^t dZ_t\right\}$ where $\tau = t - 0$. By

taking equation (24) and changing "*dt*" to discrete time, "*Δt*", the transition equation for

the Kalman Filter is defined.

$$X_t = X_{t-1} + \left(\mu - 0.5 * \sigma^2\right) * \Delta t + \theta_t \tag{25}$$

where $E[\theta_t] = 0$ *and* $Var[\theta_t] = \sigma^2 \Delta t$

This is similar to equation (3) with "a_t" equal to one and "g_t" equal to $\left(\mu - 0.5 * \sigma^2\right) * \Delta t$.

To perform the Kalman Filter algorithm, we only need initial values for "X_0", "μ", and

"σ" within the transition equation along with a time series of the observable futures

prices.

3.2   Monte Carlo Simulation

To illustrate the estimation ability of the Kalman filter, we conduct a Monte Carlo experiment. To begin, we will produce a spot price time series for crude oil using a random number generator, and parameter values for equation (25). Next, these spot prices will be used in equation (19), along with parameter values for the risk free rate and the time to maturity, to construct a time series of futures prices. Once these futures prices are obtained, we will use only these futures prices along with equations (19) and (25) to estimate the simulated spot prices using the Kalman filter. Lastly, we will compare the Kalman estimated spot prices with the simulated spot prices to show how accurate the Kalman filter estimates the unobservable (or latent) variable.

To produce a numerical example in Excel, begin the with an initial spot price of $50.00. The price moves forward in time by the process:

$$S_t = S_{t-1} \exp\{(\mu - 0.5\sigma^2)\Delta t + \sigma\Delta Z_t\}. \tag{26}$$

To generate a random number for S_t in Excel, use the following command: =NORMINV(RAND(), (μ-0.5*σ^2)*Δt, σ*SQRT(Δt)) with applicable values for "μ", "σ", and "Δt". Once these values are obtained, substitute these values into equation (26) to produce a time series for the spot price, S_t . After generating this series of prices, we may obtain the futures prices by multiplying each spot price by ($e^{(r*\tau)}$).   Figure 1 illustrates the steps described above with $\mu = 10\%$ annually, $\sigma = 25\%$ annually, $\Delta t = 1/52$, $r = 4\%$ annually, and $\tau = 1$.

(INSERT FIGURE 1 HERE)

To set up the Kalman Filter, it is necessary to understand what is actually known: $ln(F_0) = 3.9520$, $F_0 = \$52.04$, $r = 4\%$, $\Delta t = 0.01923$, and $\tau = 1$. Take the expectation of

the measurement equation, $E[ln(F_{t=0})] = E[ln(S_{t=0})] + r*\tau$, and solve for $E[ln(S_{t=0})]$

based on the known parameters (i.e. $E[ln(S_0)] = 3.9520 - 4\%*1 = 3.9120$). The variance

of $ln(S_0)$ is assumed to be zero.

It is helpful to rewrite the measurement and transition equations with known

values to determine what parameters still need to be estimated.

$$\ln(F_t) = \ln(S_t) + 4\%*1.00 + \varepsilon_t \quad where \ E[\varepsilon_t] = 0 \ and \ Var[\varepsilon_t] = q_t \tag{27}$$

Consequently, q_t, needs to be estimated in the measurement equation.

$$\ln(S_t) = \ln(S_{t-1}) + (\mu - 0.5*\sigma^2)*0.01923 + \theta_t \quad where \ E[\theta_t] = 0 \ and \tag{28}$$

$$Var[\theta_t] = \sigma^2 * 0.01923$$

Consequently, μ and σ need to be estimated in the transition equation.

The selection of initial values for these parameters to be estimated can be

performed somewhat strategically depending on one's knowledge of the system. In

theory, the values can be any set of numbers consistent with the numerical attributes of

the parameters (e.g. variance parameters should not be negative). However, an extensive

discussion of this issue is not presented here, but is certainly a worthy topic of discussion

in the classroom. To perform the Excel example, the initial parameter estimates are $\mu =$

15%, $\sigma = 32\%$, and $q_t = 10\%$. Figure 2 extends the spreadsheet in Figure 1 to

demonstrate the Kalman Filter application.

(INSERT FIGURE 2 HERE)

With the Kalman Filter entered, the EM Algorithm for the maximum likelihood estimation requires two addition columns to calculate the equivalent of equation (18) (cell M1 of Figure 3). Figure 3 illustrates the calculations assuming 100 observations. (INSERT FIGURE 3 HERE)

Next, the "Solver" function is implemented to maximize the log-likelihood function. The "Solver" function is a selection within the Excel "Tools" menu. Should "Solver" not be available, it can be loaded by selecting "Solver" under the "Add-In" menu which is within the "Tools" menu (Note: when loading "Solver", the original Excel compact disk will be requested). Within the "Solver" application, the goal is to maximize the log-likelihood function (cell M1) by adjusting the unknown parameters (cells G2, G5, and G6) while maintaining the constraint that any variance parameters remain positive (cells G2 and G6).

By implementing "Solver" with the above conditions, Excel will iterate between the Kalman Filter solutions and the maximization of the log-likelihood function (the EM Algorithm). The solutions for the particular set of data generated for this paper are $\mu =$ 8.9835%, $\sigma = 24.1354\%$, and $q_t = 0.0002\%$ with a log-likelihood function value of 197.8224. The parameter values are close to the actual parameter values used to generate the data: $\mu = 10.00\%$, $\sigma = 25.00\%$, and $q_t = 0.00\%$. One should be aware that different sets of randomly generated data produce different solutions. The solutions can vary greatly and many times depend on how well the random number generator performs at a given time (this sentiment has been echoed by others who have used this example in a classroom setting).

Consequently, it is advisable for the instructor to test the randomly generated data prior to giving it to the student. Because this is an exercise for the student to understand the Kalman Filter and not an exercise about data or modeling issues, it is best for the student's estimated dataset of the unobserved variable (Column F) to match up well with the technically "unobserved" dataset (Column D). There are two advantages to this: 1) the student can now judge via a benchmark how well their estimation of the unobserved time series performs (Column D can be made available to the student prior to or after the estimation of the Kalman Filter depending on the instructor's objectives) and 2) the student gains confidence in executing the technique and gains confidence in the technique itself assuming a correct model. In reality, one never actually compares the estimated "unobserved data" with actual "unobserved data". However, because the controlled environment developed here allows such a comparison, the instructor should take advantage of it. The exercise illustrated in Figures 1 through 3 is available from the authors upon request.

The Kalman Filter series in this exercise matched the actual generated series so well, that a graph comparing the two series is not very meaningful for the purposes of this paper as the two series simply overlay on top of each other. However, when used in the classroom, students find a graph illustrating the near perfect overlay of predicted data over actual (technically unobserved) data very compelling. The mean error between the two price series is $0.00005 with a standard deviation of $0.00341. Further, the observable data, the futures prices, when compared to the Kalman Filter predicted futures price have an average error of -$0.00186 with a standard deviation of $1.91290.

Although the topic of this paper is to only present the Kalman Filter technique, it is necessary to mention how the Kalman Filter is actually applied empirically. The Kalman Filter provides a testing environment for different model specifications for the unobserved variables to be compared. Some of the issues that emerge include:

- testing whether model parameters remain constant throughout the time series

- if the parameters do change throughout the time series, in what manner do the parameters change

- does the model actually do an adequate job at forecasting

- do particular time series elements emerge, such as, autocorrelation

Ultimately, the best (and hopefully correct) model fits the observable data with the least amount of error. The exercise presented in this paper can be as simple or complex as the instructor desires. Although considered a time series technique which places the exercise in the realm of econometrics, the exercise is suitable for a course devoted to estimating asset pricing models in finance or simply as an exercise in an Excel modeling course.

CONCLUSION:

The existing presentations of the Kalman Filter technique are daunting despite the relative ease in which the filter is implemented. Part of the problem is the matrix notation (avoided in this presentation), but equally to blame, is that most of the "scaled down" examples are applied in engineering and not in terms of economic/financial analysis. By providing an accessible example in Excel, the Kalman Filter becomes a powerful analytical tool.

The example in this paper has been presented successfully as a "stand alone" assignment in which the students follow the paper and produce their own Monte Carlo simulation to be estimated with a Kalman Filter. The example in this paper has also been used in conjunction with other material to simply illustrate the Kalman Filter technique.

Echoing some of the feedback received regarding this paper, to implement as a lecture, it is suggested that the instructor generate the Monte Carlo simulation prior to the lecture. It is important, particularly for a student's initial introduction to this material, that the technique perform well and not be subject to problems with the randomly generated set of data. A comparison between the filter generated unobservable variables and the "actual" unobservable variables is instructive, but can be omitted at the instructor's discretion.

Although this paper has been successfully assigned as a "stand-alone" assignment, when giving an assignment, we suggest providing the student with only the observable generated data (after it has been pre-tested to make certain the generated data leads to appropriate solutions). If desired, multiple sets of the generated data with different parameter figures can be created to make individual assignments. As part of the solutions to the assignment, the instructor can make the "actual" unobservable data available to the student to see how well the filter performed.

The programming for the filter is minimal, however, the ability to grasp the idea of a latent variable is the truly novel part of the classroom presentation. Consequently, this material is best suited for advanced undergraduate (econometric or financial modeling) classes because of the introduction of a latent variable. Alternatively, the material can simply be used as an Excel programming assignment by allowing the

student to program the filter in an effort to estimate the unobservable data (which the instructor can provide in this instance) through the observable data. The exercise allows the student to program the filter and provides a suitable context for using Excel's "Solver" function.

At the graduate level (drawing from feedback received from graduate students who downloaded earlier versions of the paper), the material is suitable for a course on time series analysis or advanced financial modeling. In fact, the paper can be assigned as background reading in a doctoral course for understanding the Kalman Filter prior to reading empirical research based on the Kalman Filter. In this context, the paper is not very extensive because it does not address many of the econometric issues associated with time series analysis. Yet, it is still a useful resource particularly for students who are unfamiliar with the topic.

REFERENCES:

Arnold, T. and S. Henry. (2003) Visualizing the stochastic calculus of option pricing with Excel and VBA. *Journal of Applied Finance* 13 (1): 56 – 65.

Bertus, M., S. Beyer, J. Godbey and C. Hinkelmann. (2006) Market behavior and equity prices: what can the S&P 500 index derivatives markets tell us? working paper.

Bertus, M., T. Denny, J. Godbey and C. Hinkelmann. (2006) Noise and equity prices: evidence from the stock index futures market. working paper.

Brockwell, P. and R. Davis. (2002) *Introduction to Time Series and Forecasting (2nd Edition)* (New York: Springer-Verlag).

Burmeister, E. and K. Wall. (1982) Kalman filtering estimation of unobserved rational expectations with an application to the German hyperinflation. *Journal of Econometrics* 20: 255-284.

Burmeister, E., K. Wall, and J Hamilton (1986) Estimation of unobserved expected monthly inflation using Kalman Filtering. *Journal of Business and Economic Statistics* 4, 147-160

Faff, R.W., D. Hillier, and J. Hillier. (2000) Time varying beta risk: an analysis of alternative modeling techniques. *Journal of Business Finance & Accounting* 27(5): 523-554.

Fink, J., K.E. Fink, and S. Lange. (2005) The use of term structure information in the hedging of mortgage-backed securities. *The Journal of Futures Markets* 25(7): 661-678.

Godbey, J. and J.E. Hilliard. (forthcoming) Hedging long-term commitments under stochastic convenience yield: a minimum variance approach. *Quantitative Finance.*

Hamilton, J. (1994) *Time Series Analysis* (Princeton, New Jersey: Princeton University Press).

Joseph, P. (2007) *Kalman Filters*, available on the internet: http://ourworld.compuserve.com/homepages/PDJoseph/kalman.htm

Kalman, R. (1960) A new approach to linear filtering and prediction problems. *Journal of Basic Engineering* 82: 34 – 45.

Schwartz, E. (1997) The stochastic behavior of commodity prices implications for valuations and hedging. *Journal of Finance* 52, 923-973.

Predict future unobserved variable (X_{t+1}) based on the current estimate of the unobserved variable, call it $X_{(t+1)P}$:

$$X_{(t+1)P} = a_t * X_{tP-ADJ} + g_t + \theta_t$$

Note: $X_{0P-ADJ} = X_0$ which is N(μ_0, σ_0^2)

Use the predicted unobserved variable to predict the future observable variable (Y_{t+1}), call it $Y_{(t+1)P}$:

$$Y_{(t+1)P} = m_t * X_{(t+1)P} + b_{t+1} + \varepsilon_{t+1}$$

When the future observable variable actually occurs, calculate the error in the prediction:

$$Y_{(t+1)E} = Y_{(t+1)} - Y_{(t+1)P}$$

Generate a better estimate of the unobserved variable at time (t + 1) and start the process over for time (t + 2):

$$X_{(t+1)P-ADJ} = X_{(t+1)P} + k_{(t+1)} * Y_{(t+1)E}$$

Note: k_{t+1} is the "Kalman gain" and is set to minimize the variance of $X_{(t+1)P-ADJ}$; p_{t+1} is the variance of $X_{(t+1)P}$:

$$k_{t+1} = \frac{p_{t+1} * m}{\left(p_{t+1} * m^2 + r_t\right)} = \frac{Cov\left(X_{(1+1)P}, Y_{(t+1)P}\right)}{Var\left(Y_{(t+1)P}\right)}$$

These equations are based on the more general measurement and transition equations, equations (1) and (3) respectively.

FIGURE 1: Monte Carlo Generation of the Observed and Unobserved Time Series in Excel

	A	B	C	D	E
1	Actual Parameters:				
2	Mean:	10%	Annually		
3	Volatility:	25%	Annually		
4	Risk-free rate:	4%	Annually		
5	Maturity of Futures:	1.00	Years		
6	Time Increment	0.01923[a]	Years		
7	Current Spot Price:	$50.00			
8					
9	Time (in weeks):	Security Price:	Futures Price:	Ln(Security):	Ln(Futures):
10	0	$50.00[b]	$52.04[c]	3.9120[d]	3.9520[e]
11	1	51.58[f]	$ 53.68	3.9431	3.9831
12	2	52.97	$ 55.13	3.9697	4.0097
13	3	56.04	$ 58.32	4.0260	4.0660
14	4	55.25	$ 57.50	4.0118	4.0518

[a] Cell Formula: =1/52   [b] Cell Formula: =B7
[c] Cell Formula: =B10*EXP(B4*B5)   [d] Cell Formula: =LN(B10)
[e] Cell Formula: =LN(C10)
[f] Cell Formula: =B10*EXP(NORMINV(RAND(),(B2-0.5*B3^2)*B6,B3*SQRT(B6)))
Note: Column D contains the unobserved time series and Column E contains the observed time series. Further, it may be necessary to copy the Monte Carlo data in column B over itself. Highlight the data, target the data over itself, and then use the menu sequence: Edit/Paste Special/Values. This will save the Monte Carlo data without having the simulation update itself every time a new Excel command is executed (this is an Excel default setting).

FIGURE 2: Kalman Filter Application for the Monte Carlo Data

	E	F	G	H	I	J	K
1		Measure Eq					
2		h(t):	10%				
3							
4		Trans Eq					
5		Mu:	15%				
6		Sigma:	32%				
7							
8							
9	Ln(Futures):	Pred. ln(S):	Pred. ln(F):	Error:	P(t):	K(t):	Var(t):[a]
10	3.9520	3.9120[b]					0.0000
11	3.9831	3.9145[c]	3.9539[d]	0.0292[e]	0.0020[f]	0.0193[g]	0.0019[h]
12	4.0097	3.9184	3.9564	0.0533	0.0039	0.0375	0.0038
13	4.0660	3.9260	3.9603	0.1057	0.0057	0.0541	0.0054
14	4.0518	3.9337	3.9679	0.0839	0.0074	0.0688	0.0069

[a] This is the variance of the predicted natural logarithm of the spot price. It is set at zero for t = 0.

[b] Cell Formula: =E10 - B4*B5

[c] Cell Formula: =F10 + (G5 – 0.5*G6^2)*B6 + J11*H11

[d] Cell Formula: =F10 + (G5 – 0.5*G6^2)*B6 + B4*B5

[e] Cell Formula: =E11 – G11   [f] Cell Formula =K10 + G6^2*B6

[g] Cell Formula: =I11/(I11+G2)   [h] Cell Formula: =I11*(1-J11)

Note: Cells B4, B5, and B6 refer to the spreadsheet in Figure 1.

FIGURE 3: Applying Maximum Likelihood Estimation to the Kalman Filter

	F	G	H	I	…	L	M
1	Measure Eq					log-like:	15.4068[a]
2	h(t):	10%					
3							
4	Trans Eq						
5	Mu:	15%					
6	Sigma:	32%					
7							
8							
9	Pred. ln(S):	Pred. ln(F):	Error:	P(t):	…	MLE(1):	MLE(2):
10	3.9120						
11	3.9145	3.9539	0.0292	0.0020	…	1.1415[b]	-0.0042[c]
12	3.9184	3.9564	0.0533	0.0039	…	1.1322	-0.0137
13	3.9260	3.9603	0.1057	0.0057	…	1.1235	-0.0529
14	3.9337	3.9679	0.0839	0.0074	…	1.1157	-0.0328

[a] Cell Formula: = -100*LN(2*PI())/2 + SUM(L11:L110) + SUM(M11:M110)
[b] Cell Formula: = -LN(I11 + G2)/2   [c] Cell Formula: = -(H11^2/(I11 + G2))/2