A new unscented Kalman filter with higher order moment-matching

KSENIA PONOMAREVA, PARESH DATE AND ZIDONG WANG
Department of Mathematical Sciences, Brunel University,
Uxbridge, UB8 3PH, UK.

***Abstract*— This paper is concerned with filtering nonlinear multivariate time series. A new approximate Bayesian algorithm is proposed which generates sample points and corresponding probability weights that match exactly the predicted values of average marginal skewness and average marginal kurtosis of the unobserved state variables, in addition to matching their mean and the covariance matrix. The performance of the algorithm is illustrated by an empirical example of yield curve modelling with real financial market data. Results show an improvement in accuracy in comparison with extended Kalman filter (EKF) and traditional unscented Kalman filter (UKF).**

Keywords: State estimation, sigma point filters, nonlinear time series

I. INTRODUCTION

This paper is concerned with the problem of latent state estimation for a nonlinear time series in discrete time. Our analysis will focus on the general class of systems with the following state space form:

$$\mathcal{X}(k+1) = \mathbf{f}(\mathcal{X}(k)) + Q(\mathcal{X}(k))\mathbf{w}(k+1), \quad (1a)$$
$$\mathcal{Y}(k) = \mathbf{h}(\mathcal{X}(k)) + \mathbf{v}(k), \quad (1b)$$

where $\mathcal{X}(k)$ and $\mathcal{Y}(k)$ are the respective state vector and measurement vector at time t_k; $\mathbf{f}$,$\mathbf{h}$ are given vector-valued deterministic functions; Q is a matrix valued deterministic function; and $\mathbf{v}(k), \mathbf{w}(k)$ are vector-valued random variables. The time increment $t_k - t_{k-1}$ is assumed constant for all k. We wish to find an estimate of the random vector $\mathcal{X}(k)$, $k \geq 1$, based on the noisy time series data $\mathcal{Y}(1), \mathcal{Y}(2), ..., \mathcal{Y}(k)$.

In the special case when $\mathbf{f}$, $\mathbf{h}$ are affine in $\mathcal{X}(k)$, Q is an identity matrix and $\mathbf{v}(k), \mathbf{w}(k)$ are Gaussian, the optimal recursive solution to the state estimation problem is given by linear Kalman filter, as first outlined in [1]. However, nonlinear and non-Gaussian models are used to capture the dynamics of many phenomena occurring in the fields of radar navigation, climatology, financial modeling and econometrics, among others. The optimal recursive solution to the state estimation problem in nonlinear systems is usually not available in closed form. Current approaches to address the nonlinear filtering problems fall under one of the following approximate Bayesian filtering methods:

(a) Extended Kalman filter (EKF). Under this filter, equation (1) or its continuous time analogue is locally linearized resulting in a linear state space system. A Kalman filter is then employed to obtain the conditional state density of $\mathcal{X}(k)$. This approach is popular in engineering for more than three decades and standard textbooks such as [2] carry an extensive discussion of its theoretical underpinnings and implementation. If the system is approximately linear then EKF will work well. Nevertheless, such assumption is often not easy to validate.

(b) Sequential Monte Carlo filter or *particle filter* (PF). For this technique, the required conditional density function of $\mathcal{X}(k)$ given measurement $\mathcal{Y}(k)$ at time t_k is represented by a set of random samples (or *particles*) and associated probability weights. The particles and weights are updated recursively as new measurements become available; see [3]-[5] and references therein for more details on PF. PF can perform better than EKF for highly nonlinear systems. However, as large number of samples need to be generated at each time t_k, this type of techniques are computationally quite expensive to implement, especially for large state dimensions.

(c) Unscented Kalman filters (UKF). This class of filters provides an increasingly popular alternative to particle filters in signal processing. UKF may be viewed as a compromise between an EKF (in the sense that it uses a closed-form expression for updating the state estimate) and a PF (in the sense that it uses a set of particles - or *sigma points* - and weights to evaluate the terms in that expression). Several applications in communication, tracking and navigation are discussed in [6] and [7], among others. *Ensemble filter* (EF) used in climatology is closely related to UKF; see [8], [9] and references therein.

UKF suffers from one major disadvantage, especially for systems with significant noise terms in the transition equation (1). Even if the density of $\mathcal{X}(k|k)$ at time t_k is Gaussian, a nonlinear $\mathbf{f}$ will lead to a prediction $\mathcal{X}(k+1|k)$ whose density is non-Gaussian in general. Unscented filter assumes conditional Gaussianity throughout the filter recursions and may lead to misleading results in case the density departs too far from the assumed Gaussian density. We propose an algorithm to partially alleviate this problem while still working within UKF framework. Specifically, the sigma points and weights are modified at each time step to match exactly the predicted values of the average marginal skewness and average marginal kurtosis of the the unobserved state variables, in addition to matching their mean and the covariance matrix.

The rest of this paper are organized as follows. Section II reviews the traditional unscented Kalman filter and Section

III introduces the new algorithm for unscented filter with higher order moment matching. Section IV illustrates the utility of our method by comparing its performance in filtering nonlinear, multivariate time series with the performance of the UKF as well as the EKF. Section V summarizes the contributions of the paper and outlines directions for further research.

II. Unscented Filter

Consider a random n-vector $\mathcal{X}$ with mean $\hat{\mathcal{X}}$ and covariance P_{xx}. A nonlinear tranformation relates $\mathcal{X}$ to a second random vector $\mathcal{Y}$ through

$$\mathcal{Y}(k) = \mathbf{h}(\mathcal{X}(k)) + \mathbf{v}(k),$$

where $\mathbf{v}(k)$ is the zero-mean noise vector as in equation (1). Augmentation method (see, *e.g.* [10]) incorporates noise into the augmented random state vector and from here onwards we will assume $\mathbf{h}$ to be an augmented function. The problem now is to calculate the mean $\hat{\mathcal{Y}}$ and covariance P_{yy} of $\mathcal{Y}$.

In unscented filter, 2n+1 symmetric sigma points are chosen so that they have the same mean and covariance as $\mathcal{X}$.

$$\begin{aligned} X^{(0)} &= \hat{\mathcal{X}}, \\ X^{(i)} &= \hat{\mathcal{X}} \pm (\sqrt{(n+\kappa)P_{xx}})_i, \end{aligned}$$

where i=1,2,...,n, κ is a scaling parameter and $(\sqrt{P_{xx}})_i$ is the i^{th} column of the matrix square root of P_{xx}. Probability weights W_i associated with the i^{th} sigma point $X^{(i)}$ are defined as

$$W_0 = \frac{\kappa}{n+\kappa}, \tag{2}$$

$$W_i = \frac{1}{2(n+\kappa)}, \tag{3}$$

for i=1,2,...,2n. Given these sigma points and corresponding weights $\hat{\mathcal{Y}}$ and P_{yy} are calculated as follows.

(a) Propagate each point through the function to get the set of transformed sigma points

$$\mathcal{Y}^{(i)} = \mathbf{h}(X^{(i)}).$$

(b) The mean is given by the weighted average of transformed poins

$$\hat{\mathcal{Y}} = \sum_{i=0}^{2n} W_i \mathcal{Y}^{(i)}.$$

(c) The covariance is the weighted outer product of the transformed points

$$P_{yy} = \sum_{i=0}^{2n} W_i (\mathcal{Y}^{(i)} - \hat{\mathcal{Y}})(\mathcal{Y}^{(i)} - \hat{\mathcal{Y}})^T.$$

More details on this algorithm can be found in [6]. Note that, while matching first and second moment accurately, UKF does not propagate information on 3rd and 4th moments. Other suggested algorithms, which try to match higher moments, either require optimization or rely heavily on analytical solver, as in [11].

III. New Algorithm for Unscented Kalman Filtering

A. Sigma point generation

We want to extend Julier and Uhlmann method in [6] to asymmetric distributions by introducing additional variables α and β in order to capture 3rd and 4th moments of $\mathcal{X}(k+1|k)$ using augmented UKF. Here we consider an augmented method mentioned in section II, which has been shown in [10] to give more accurate results compared to non-augmented UKF in presence of significant noise terms.

Suppose we have $\mathcal{X}$, random n-vector with mean $\hat{\mathcal{X}}$ and covariance P_{xx}, and noise, m-vector with zero mean and covariance Q_{xx}, as in equation (1). Matrices $P > 0$, $Q > 0$ are such that P_{xx}=PP^T and Q_{xx}=QQ^T, where P^T is transpose of P and P_i is the ith column of matrix P. We create 2(n+m)+1 sigma points and corresponding weights as follows:

$$\begin{aligned} X^{(0)} &= \begin{pmatrix} \hat{\mathcal{X}} \\ \mathbf{0}_{m\times 1} \end{pmatrix} = \bar{\mathcal{X}}, \\ X^{(i)} &= \begin{pmatrix} \hat{\mathcal{X}} + \alpha\sqrt{N}P_i \\ \mathbf{0}_{m\times 1} \end{pmatrix}, i = 1, 2, ..., n, \\ &= \begin{pmatrix} \hat{\mathcal{X}} - \beta\sqrt{N}P_i \\ \mathbf{0}_{m\times 1} \end{pmatrix}, i = n+1, ..., 2n, \\ &= \begin{pmatrix} \hat{\mathcal{X}} \\ \sqrt{N}Q_{i-2n} \end{pmatrix}, i = 2n+1, ..., 2n+m, \\ &= \begin{pmatrix} \hat{\mathcal{X}} \\ -\sqrt{N}Q_{i-2n} \end{pmatrix}, i = 2n+m+1, ..., 2n+2m, \end{aligned}$$

where $N = n + m$ and j^{th} element of a sigma point $X^{(i)}$ will be denoted as $X_j^{(i)}$. Note that $\alpha \neq \beta$ would mean a set of points distributed asymmetrically about the mean $\hat{\mathcal{X}}$. The probability weight corresponding to the point $X^{(i)}$ is denoted as W_i. In order to exactly match the expected mean and the covariance matrix, the weights W_i have to satisfy the following conditions:

$$\sum_{i=0}^{2N} W_i = 1, \tag{4}$$

$$\sum_{i=0}^{2N} W_i X^{(i)} = \begin{pmatrix} \hat{\mathcal{X}} \\ \mathbf{0}_{m\times 1} \end{pmatrix}, \tag{5}$$

$$\sum_{i=0}^{2N} W_i (X^{(i)} - \bar{\mathcal{X}})(X^{(i)} - \bar{\mathcal{X}})^T = \begin{pmatrix} P_{xx} & \mathbf{0}_{n\times m} \\ \mathbf{0}_{m\times n} & Q_{xx} \end{pmatrix} \tag{6}$$

Note that $W_i \geq 0$ and $\sum_{i=0}^{2N} W_i = 1$ mean the set of probability weights and corresponding sigma points $\{W_i, X^{(i)}\}$ forms a valid probability distribution. This is *not* always the case in unscented Kalman filter, since κ in (2)-(3) is not restricted to be positive. Equations (4)-(6) give us explicit expressions for weights in terms of N, α and β:

$$\begin{aligned} W_i &= \frac{1}{\alpha(\alpha+\beta)N},\ i = 1, 2, ..., n, && (7)\\ &= \frac{1}{\beta(\alpha+\beta)N},\ i = n+1, ..., 2n, && (8)\\ &= (\frac{1}{2N}), \qquad i = 2n+1, ..., 2n+2m, && (9)\\ W_0 &= 1 - \sum_{i=1}^{2N} W_i. && (10)\end{aligned}$$

Define the marginal 3rd and 4th central moments as

$$\begin{aligned} \omega_j &= \sum_{i=0}^{2N} W_i (X_j^{(i)} - \bar{\mathcal{X}}_j)^3,\\ \psi_j &= \sum_{i=0}^{2N} W_i (X_j^{(i)} - \bar{\mathcal{X}}_j)^4.\end{aligned}$$

As we have only two degrees of freedom (*viz* α and β), we choose to match the *average* third and fourth marginal moments of the state vector alone (i.e. ignoring the moments of noise terms). Note that it is possible to match average marginal moments of the augmented state vector, although it is not done here. We have

$$\begin{aligned} \frac{1}{n}\sum_{j=1}^{n}\sum_{i=0}^{2N} W_i (X_j^{(i)} - \bar{\mathcal{X}}_j)^3 &= \frac{1}{n}\sum_{j=1}^{n}\omega_j,\\ \frac{1}{n}\sum_{j=1}^{n}\sum_{i=0}^{2N} W_i (X_j^{(i)} - \bar{\mathcal{X}}_j)^4 &= \frac{1}{n}\sum_{j=1}^{n}\psi_j.\end{aligned}$$

Substituting expressions for W_i from (7)-(10) we get

$$\alpha - \beta = \frac{\sum_{j=1}^{n}\omega_j}{\sqrt{N}\sum_{l=1}^{n}\sum_{k=1}^{n}P_{ij}^3} =: \phi_1, \qquad (11)$$

$$\alpha^2 - \alpha\beta + \beta^2 = \frac{\sum_{j=1}^{n}\psi_j}{N\sum_{l=1}^{n}\sum_{k=1}^{n}P_{ij}^4} =: \phi_2, \qquad (12)$$

where P_{ij} is entry in i^{th} row and j^{th} column of matrix P, so that ϕ_1 and ϕ_2 are known from data. Solving these yields

$$\begin{aligned} \alpha &= \frac{1}{2}\phi_1 \pm \frac{1}{2}\sqrt{4\phi_2 - 3\phi_1^2}, && (13)\\ \beta &= -\frac{1}{2}\phi_1 \pm \frac{1}{2}\sqrt{4\phi_2 - 3\phi_1^2}, && (14)\end{aligned}$$

where we take values of the same sign. Provided $\phi_2 \geq \frac{3}{4}\phi_1^2$, (which is trivially true for symmetric distributions), α and β allow us to capture and propagate the marginal skewness and marginal kurtosis.

Note that unscented filter in section II employs the same weights W_i for all sigma points $X^{(i)}$ for $i > 0$. In comparison, we have different expressions for probability weights depending on i in (7)-(9). A similar approach is taken in [11]. The paper proposes a method where a set of nonlinear algebraic equations is solved to find support points and probability weights to match a given set of moments. However, the probability weights found are held constant throughout the recursion and a closed form analytic solution is given only for the Gaussian case (in particular, with zero skewness).

B. Filtering Algorithm for Unscented Kalman Filtering

Once the sigma points $X^{(i)}$ and weights W_i are generated as in the previous section, we can find the one step prediction for mean and covariance matrix using

$$\mathbb{E}(\mathbf{x}|\mathbf{y}) = \mathbb{E}(\mathbf{x}) + \Sigma_{xy}\Sigma_{yy}^{-1}(\mathbf{y} - \mathbb{E}(\mathbf{y})), \quad (15)$$

$$\mathbb{E}(\mathbf{x} - \mathbb{E}(\mathbf{x}|\mathbf{y}))(\mathbf{x} - \mathbb{E}(\mathbf{x}|\mathbf{y}))^\top = \Sigma_{xx} - \Sigma_{xy}\Sigma_{yy}^{-1}\Sigma_{yx}, \quad (16)$$

for x and y jointly Gaussian.

The filtering algorithm can be described as follows.

(a) Given $\hat{\mathcal{X}}(k|k-1)$, $P_{xx}(k|k-1)$ use

$$\begin{aligned} \mathcal{X}^{(i)}(k+1|k) &= \mathbf{f}(X^{(i)}) + P_{xy}P_{yy}^{-1}(Y_k - \mathbf{h}(\mathcal{X}^{(i)}(k|k))),\\ \hat{\mathcal{X}}(k+1|k) &= \sum_{i=0}^{2N} W_i \mathcal{X}^{(i)}(k+1|k),\end{aligned}$$

where Y_k is the true measurement and $\mathcal{X}^{(i)}(k|k) = \mathbf{f}(X^{(i)})$. $P_{xx}(k+1|k)$ is computed similarly using (16). Covariance matrices $P_{xy}(k+1|k)$ and $P_{yy}(k+1|k)$ are worked out as

$$\begin{aligned} P_{xy}(k+1|k) &= \sum_{i=0}^{2N} W_i(\mathcal{X}^{(i)}(k|k) - \hat{\mathcal{X}}(k|k))v^{(i)}(k)^T,\\ P_{yy}(k+1|k) &= \sum_{i=0}^{2N} W_i v^{(i)}(k) v^{(i)}(k)^T,\end{aligned}$$

where $v^{(i)}(k) = \mathcal{Y}^{(i)}(k+1) - \hat{\mathcal{Y}}(k+1)$, $\mathcal{Y}^{(i)}(k+1) = \mathbf{h}(\mathcal{X}^{(i)}(k|k))$ and $\hat{\mathcal{Y}}(k+1) = \sum_{i=0}^{2N} W_i\mathcal{Y}^{(i)}(k+1)$.

(b) Calculating average marginal skewness and average marginal kurtosis of $\mathcal{X}^{(i)}(k+1|k)$ provides us with updated values for α and β via (13)-(14). Now we use these values to generate new set of sigma points and corresponding weights at time t_{k+1}.

IV. Numerical Example

To test the efficiency of the new algorithm for unscented filtering we consider discretisation of the multi-factor CIR model with nonlinear measurement equation. The state evolution is given as below. This is a multivariable extension of the model first proposed in [12]; see [13] for more details on the use of this model in filtering context.

$$\mathcal{X}_j(k+1) = \kappa_j\epsilon_j\theta_j + (1 - \kappa_j\epsilon_j)\mathcal{X}_j(k) + Q_j(k)w_j(k),$$

for $j = 1, 2$, where $w_j(k)$ are zero mean, unit variance and uncorrelated Gaussian random variables. The standard deviation Q is given by

$$Q_j(k) = \sigma_j\sqrt{\epsilon_j(\frac{1}{2}\theta_j(\kappa_j\epsilon_j + (1 - \kappa_j\epsilon_j)\mathcal{X}_j(k-1))},$$

where κ_j, σ_j and θ_j are constants and

$$\epsilon_j = \frac{(1 - e^{(-\kappa_j \Delta)})}{\kappa_j}.$$

The observable variables are exponential in the latent states and are given by

$$\mathcal{Y}_i(k) = \Pi_{j=1}^2 \left(A_{i,j} \exp(-\sum_{j=1}^{2} (B_{i,j} \mathcal{X}_j(k))) \right) + z_i(k),$$

where

$$A_{i,j} = \left(\frac{2\gamma_j \exp((\kappa_j + \gamma_j + \lambda_j)T_i/2)}{2\gamma_j + (\kappa_j + \lambda_j + \gamma_j)(exp(T_i\gamma_j) - 1)} \right)^{\frac{2\kappa_j\theta_j}{\sigma_j^2}},$$
$$B_{i,j} = \frac{2(exp(T_i\gamma_j) - 1)}{2\gamma_j + (\kappa_j + \lambda_j + \gamma_j)(exp(T_i\gamma_j) - 1)},$$
$$\gamma_j = \sqrt{(\kappa_j + \lambda_j)^2 + 2\sigma_j^2}.$$

$z_i(k)$ is observational noise with zero mean and a constant variance h^2 for each i and λ_i are constants. In practice, T_i represents time to maturity and $\mathcal{Y}_i(k)$ represents the price of a zero coupon bond with maturity $T_i + t_k$, at time t_k. Here we use three maturities, $T_1 = 1$, $T_2 = 2$ and $T_3 = 4$.

For numerical experiments, we use weekly data from February 2001 to July 2005 for 3 different UK government bond yields. Here 180 observations were used for calibration and 42 were used for out-of-sample validation. A 2-factor model was calibrated using the extended Kalman filter and the maximum likelihood method. In-built optimization routines from MATLAB were used for calibration. Table 1 reports the parameter values obtained as a result of calibration.

Table 1. Parameter values

θ_1	0.0254
θ_2	0.0175
σ_1	0.0710
σ_2	0.1870
κ_1	0.0978
κ_2	0.8035
λ_1	−0.0350
λ_2	−0.0490
h	0.001

After calibration, we use the sigma point generation method described in section III to generate sigma points at each t_k, with initial values for mean θ_j and diagonal elements of covariance as $\frac{\theta_j\sigma_j^2}{2\kappa_j}$. Eleven sigma points are generated at each t_k. Bearing in mind the nonnegativity restriction on state variables $\mathcal{X}_j(k) \geq 0$ we replace any negative element of state estimate $\mathcal{X}_j(k|k-1)$ with zero. These points are then used to construct $\hat{\mathcal{X}}_j(k+1|k)$, $j = 1, 2$ and the corresponding predictions of $\mathcal{Y}_i(k+1)$, $i = 1, 2, 3$. As a benchmark for comparison, we use the predictions made using the extended Kalman filter and traditional UKF described in section II.

To compare the performance of sigma point filters and extended Kalman filters, for each time to maturity τ_k we consider the sample mean of the relative absolute error (MRAE) defined as

$$MRAE_i = \frac{1}{M} \sum_{j=1}^{M} \frac{|Y_i(j) - \hat{\mathcal{Y}}_i(j)|}{Y_i(j)}.$$

This was computed over the relevant set of M observations of in-sample and out-of-sample data separately. Tables 2-3 list the errors computed for one step ahead prediction of yields for the extended Kalman filter (EKF), traditional unscented filter (UKF) and the higher order sigma point filter (HOSPF) proposed here. It can be seen that the new sigma point filter outperforms the extended Kalman filter and unscented filter, in-sample and especially out-of-sample and for all yields. In particular, the improvement for out-of-sample predictions achieved with HOSPF is over 10% for all the yields, as compared to UKF. This small improvement is obtained with very little extra computational effort.

Table 2. Relative absolute errors of 1-step ahead in-sample prediction

τ_k	EKF	UKF	HOSPF
$1Y$	0.00237	0.00079	0.00075
$2Y$	0.00388	0.00198	0.00190
$4Y$	0.00606	0.00365	0.00349

Table 3. Relative absolute errors of 1-step ahead out-of-sample prediction

τ_k	EKF	UKF	HOSPF
$1Y$	0.004114	0.00075	0.00066
$2Y$	0.006832	0.00144	0.00125
$4Y$	0.011320	0.00300	0.00268

V. Summary

In this paper, we have proposed a new algorithm in which the sigma points and weights are modified at each step to match exactly the predicted values of the average marginal skewness and the average marginal kurtosis, besides matching the mean and covariance matrix. For filtering high dimensional data, this algorithm is a very useful alternative to the extended Kalman filter (due to improved accuracy) and is computationally no more difficult than the standard UKF. A numerical example shows that our method ouperforms the traditional UKF and the EKF. Currently we are testing the sensitivity of this algorithm to the changes in model parameters and the theoretical accuracy using Taylor series expansion of the estimation error.

References

[1] Kalman, R.E. (1960) A new approach to linear filtering and prediction problems, *Journal of Basic Engineering*, **82**, 35-45.

[2] B. Anderson and J. Moore (1979) Optimal Filtering, Prentice-Hall

[3] G. KITAGAWA (1996) Monte Carlo filter and smoother for non-Gaussian nonlinear state space models, *Journal of Computational and Graphical Statistics*, **5**, 1-25.

[4] ARULAMPALAM, M.S., MASKELL, S., GORDON, N. & CLAPP, T. (2002) A tutorial on particle filters for online nonlinear /non-Gaussian Bayesian tracking, *IEEE Transactions on Signal Processing*, **50**, 174-188.

[5] DAUM, F. (2005) Nonlinear filters: beyond the Kalman filter, *IEEE Aerospace and Electronic Systems Magazine*, **20**, 57-69.

[6] S. JULIER AND J. UHLMANN (2004) Unscented filtering and nonlinear estimation, *Proceedings of the IEEE*, **92**, 401-422.

[7] A. FARINA, B. RISTIC, AND D. BENVENUTI (2002) Tracking a ballistic target: comparison of several nonlinear filters, *IEEE Transactions on Aerospace and Electronic Systems*, **38**, 854-867.

[8] H. MITCHELL AND P. HOTEKAMER (1998) Data assimilation using an ensemble Kalman filter technique, *Monthly Weather Review*, **126**, 796-811.

[9] EVENSEN, G. (1994) Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics, *Journal of Geography*,**10**, 143-162.

[10] YUANXIN WU, DEWEN HU, MEIPING WU, AND XIAOPING HU (2005) Unscented Kalman Filtering for Additive Noise Case: Augmented versus Non-augmented, *IEEE Signal Processing Letters*, **12**, 357-360.

[11] D. TENNE AND T. SINGH (2003) The higher order unscented filter, *Proceedings of American Control Conference*, **3**, 2441-2446.

[12] COX, J.C., INGERSOLL, J.E. AND ROSS, S.A. (1985) A theory of the term structure of interest rates, *Econometrica*, **53**, 385-407.

[13] GEYER, A. AND PICHLER, S. (1993) A state-space approach to estimate and test multifactor cox-ingersoll-ross models of the term structure, *Journal of Financial Research*, **22**, 107-130.