Proceedings of the 2014 Winter Simulation Conference
A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, eds.

ON THE SENSITIVITY OF GREEK KERNEL ESTIMATORS TO BANDWIDTH PARAMETERS

Marie Chau

Department of Mathematics
University of Maryland
1305 Mathematics Building
College Park, MD 20742, USA

Michael C. Fu

Robert H. Smith School of Business & ISR
University of Maryland
4305 Van Munching Hall
College Park, MD 20742, USA

ABSTRACT

The Greeks measure the rate of change of (financial) derivative prices with respect to underlying market parameters, which is essential in financial risk management. This paper focuses on a modified pathwise method that overcomes the difficulty of estimating Greeks with discontinuous payoffs as well as second-order Greeks and involves a kernel estimator whose accuracy/performance relies on a smoothing parameter (bandwidth). We explore the accuracy of the Greek delta, vega, and theta estimators of Asian digital options and up-and-out barrier call options with varying bandwidths. In addition, we investigate the sensitivity of a proposed iterative scheme that generates the "optimal" bandwidth. Our numerical experiments indicate that the Greek estimators are quite sensitive to the bandwidth choice, and the "optimal" bandwidth generated is sensitive to input parameters.

1 INTRODUCTION

Financial derivatives have been in existence for over 300 years and have become an integral part of risk management. Therefore, accessing and quantifying the risk of these assets have sparked great interest in the finance industry. The Greeks are derivatives of (financial) derivative prices with respect to an underlying market parameter. Each Greek measures a different dimension of market risk. For instance, "delta," "vega," and "theta" are first-order derivatives of the expected payoff with respect to price, volatility, and time, respectively, and "gamma" is the second-order derivative with respect to price.

One type of financial derivative used for hedging is an option, which gives the security holder the right to exercise the option once certain conditions are satisfied, which is dependent on the option type. For example, the exercise condition for an Asian digital option depends on depends on the average price over a pre-specified period, and it pays a lump sum if the option is "in the money" and zero otherwise. Up-and-out barrier call options can be exercised once the price reaches the strike price but does not cross the upper barrier. Both of these options have discontinuous payoff functions, and the option price can be written as

$$E[g(S)\cdot I_{\{h(S)\geq 0\}}], \tag{1}$$

where the expectation is taken under the risk-neutral measure, S is the price modeled by a stochastic process, $I_{\{\cdot\}}$ is the indicator function, and $g(S)$ is the discounted payoff when the exercise condition, $h(S)\geq 0$, is satisfied. The price S depends on the market parameters $\theta \in \Theta$, where Θ is an open set. For simplicity, assume θ is one-dimensional, since we can focus on one market parameter and hold all else constant. The derivative of the payoff function (1) with respect to θ is a first-order Greek, and the derivative of the first-order Greek with respect to θ is a second-order Greek, both of which we wish to estimate. Some well-known derivative estimation methods include finite difference approximation, perturbation analysis (also known as the pathwise method), the score function or likelihood ratio (SF/LR) method, and the weak derivative (WD) method. Each method has its advantages and drawbacks.

The most straightforward approach to estimate Greeks is the finite difference method. Although this method is easy to implement, the resulting estimators are biased and generally have a high mean-squared error (MSE) (Glasserman 2004). In addition, a perturbation sequence $\{\Delta\theta\}$ must be selected, which affects the bias/variance tradeoff, although symmetric differences can be used to reduce the order of bias (Fu 2008). In addition, smaller perturbations could lead to more noise because of the stochasticity. The pathwise method, generally known as infinitesimal perturbation analysis (IPA) in the simulation community, is easy to implement and often performs well in comparison with other methods, but unfortunately, it is not always applicable. For instance, it is not suitable for discontinuous functions nor second-order Greeks (Liu and Hong 2011). Smoothed perturbation analysis (SPA) is able to overcome the discontinuity issue, and Fu and Hu (1997) suggest using SPA to smooth out the discontinuity by conditioning on appropriate random variables, e.g., Wang, Fu, and Marcus (2009). The key, which is the most difficult part of this method, is choosing the random variable on which to condition. On the contrary, the weak derivative (WD) method is applicable to discontinuous functions, but may require a high number of simulations (Fu 2008). The likelihood ratio method can also approximate the gradient of discontinuous functions, but usually leads to high variance. For a comprehensive review of gradient estimation methods, refer to Glasserman (2004), Fu (2006), and Asmussen and Glynn (2007).

Liu and Hong (2011) introduce a modified pathwise method to estimate first- and second-order Greeks for options with discontinuous payoffs. It is a generalization of the classical pathwise method, which extends to options with discontinuous payoffs (1) by rewriting the Greek as a sum of an expectation and a derivative(s) with respect to an auxiliary parameter. The general pathwise method and kernel method are applied to the expectation and derivative terms, respectively. The accuracy of the kernel estimator depends/hinges on the chosen kernel function, which relies on a bandwidth/smoothing parameter, described in more detail in Section 3.3. A pilot simulation proposed in Liu and Hong (2011) generates an "optimal" bandwidth parameter prior to applying the modified pathwise method. The pilot simulation involves various input parameters, which affect the bandwidth output. The modified pathwise method has been applied to estimate delta, vega, theta and gamma of the Asian digital option and up-and-out barrier call option.

In this paper, we examine the accuracy of the Greeks estimated using a modified pathwise method for an Asian digital option and up-and-out barrier call option by conducting two sets of numerical experiments. The first set investigates the sensitivity of the Greek estimators to the bandwidth parameter(s), and the second explores the sensitivity of a proposed method used to generate "optimal" bandwidths for the modified pathwise method to various input parameters. Our numerical results show that the Greek estimators are quite sensitive to the bandwidth, so the performance of the pilot simulation is critical.

The paper is organized as follows. Section 2 provides a brief overview of the generalized pathwise method for the first- and second-order Greeks from Liu and Hong (2011) and kernels. Section 3 details the implementation of the Greeks for the Asian digital option and the up-and-out barrier option. Section 4 derives the the optimal bandwidth parameter and summarizes the pilot simulation given in the e-companion of Liu and Hong (2011) for the first-order Greek estimators. In Section 5, we describe our two sets of numerical experiments and corresponding results of 1) the sensitivity of Greek kernel estimators to bandwidth and 2) the sensitivity of the bandwidth to various input parameters in the pilot simulation. Finally, we conclude in Section 6.

2 PROBLEM SETTING

We focus on options with a discontinuous payoff function that can be written in the form in (1), where $g(\cdot)$ and $h(\cdot)$ are differentiable functions, and S is a vector of discretized prices. Let S_{t_i} denote the price of the security at time $t_i \geq 0$, where $t_i = iT/k$ for $i = 0,1,2,\ldots,k$, which are evenly spaced time points between 0 and T. For simplicity, denote S_{t_i} by S_i, and $S = (S_0, S_1, \ldots, S_k)^{'}$. Let $p(\theta) = E[g(S) \cdot I_{\{h(S) \geq 0\}}]$, where $p^{'}(\theta) = \partial p(\theta)/\partial\theta$ is the first-order Greek with respect to θ and $p^{''}(\theta) = \partial^2 p(\theta)/\partial\theta^2$ is the

Table 1: Parameters.

	r	σ	b	μ	S_0	K	T	U	k
Asian digital option	0.05	0.30	0.2	98	100	100	1	-	10, 20, 50
up-and-out barrier call option	0.05	0.20	-	-	100	100	1	120	20, 50

second-order Greek with respect to θ. If $\theta = S_0$, $\theta = \sigma$, and $\theta = \tau$, then $p^{'}(\theta)$ is called delta, vega, and theta, respectively, and if $\theta = S_0$, then $p^{''}(\theta)$ is called gamma.

We explore Asian digital options and up-and-out barrier call options, both of which have discontinuous payoffs. For the Asian digital option, the discounted payoff function $g(S) = e^{-rT}$, where r is the risk-free rate and T is the duration of time considered, and the exercise criterion is based on an average of k discretized prices $\overline{S} = \frac{1}{k}\sum_{i=1}^{k} S_i$ exceeding a pre-specified threshold K over a predetermined period of time, i.e., whether or not $h(S) = \overline{S} - K \geq 0$. The up-and-out barrier call option can be exercised if the final price S_k is greater than or equal to a threshold K, but the maximum price $S_{max} = \max\{S_1, \ldots, S_k\}$ never exceeds an upper barrier U, i.e., whether or not $h(S) = \min\{S_k - K, U - S_{max}\} \geq 0$. The payoff function is the amount by which the final price S_k exceeds the strike price K discounted by e^{-rT}, i.e., $g(S) = e^{-rT}(S_k - K)$.

In the numerical experiments, the prices considered for the Asian digital option and the up-and-out barrier call option follow an Ornstein-Uhlenbeck process and a geometric Brownian motion generated, respectively, using

$$\begin{aligned} S_i &= S_{i-1}e^{-b\tau} + \mu(1 - e^{-b\tau}) + \sigma\sqrt{(1 - e^{-2b\tau})/(2b)}Z_i \quad \text{for } i = 1, 2, ..., k, \\ S_t &= S_0 e^{(r - \sigma^2/2)t + \sigma B_t} \quad \text{for } t = T/k, 2T/k, ..., T, \end{aligned}$$

where b, σ, μ, r, k, and τ are the mean reversion rate, volatility, mean return, risk-free interest rate, number of discretized intervals, and incremental time step (T/k), respectively. The variables Z_i and B_t are the independent standard normal random variables and standard Brownian motion, respectively. The values considered are listed in Table 1.

3 GENERALIZED PATHWISE METHOD

The pathwise method is in general not applicable to second-order Greeks nor discontinuous functions, but the modified kernel version is one method that circumvents this issue.

3.1 First-Order Greeks

The first-order Greek estimator based on the modified pathwise method is based on the following theoretical result from Liu and Hong (2011).

Assumption 1 For any $\theta \in \Theta$, $g(S)$ and $h(S)$ are differentiable with respect to θ with probability 1 (w.p.1), and there exist random variables K_g and K_h with finite second moments that may depend on θ, such that $|g(S(\theta + \Delta\theta)) - g(S(\theta))| \leq K_g|\Delta\theta|$ and $|h(S(\theta + \Delta\theta)) - h(S(\theta))| \leq K_h|\Delta\theta|$ when $|\Delta\theta|$ is sufficiently small.

Assumption 2 For any $\theta \in \Theta$, $\partial_\theta \phi(\theta, y)$ exists and is continuous at (θ, y) with $\phi(\theta, y) = E[g(S) \cdot 1_{\{h(S) \geq y\}}]$.

Theorem 1 Suppose $g(\cdot)$ and $h(\cdot)$ satisfy Assumptions 1 and 2, and $E[|g(S)|^2] < +\infty$ and $E[|h(S)|^2] < +\infty$, then

$$p^{'}(\theta) = E[\partial_\theta g(S) \cdot 1_{\{h(S) \geq 0\}}] - \partial_y E[g(S)\partial_\theta h(S) \cdot 1_{\{h(S) \geq y\}}]|_{y=0} \tag{2}$$

The two terms on the right hand side of (2) can be estimated easily, because the first term is an ordinary expectation, where $g(S)$ is differentiable with respect to θ, and the second term is a derivative with respect to y, which is independent of $g(S)$ and $h(S)$. Thus, the first term can be estimated using the law of large numbers, and finite differences can be used to estimate the second term based on

$$
\begin{aligned}
-\partial_y E\left[g(S)\partial_\theta h(S)\cdot 1_{\{h(S)\geq y\}}\right]|_{y=0} &= -\lim_{\delta\to 0}\frac{1}{\delta}\{E\left[g(S)\partial_\theta h(S)\cdot 1_{\{h(S)\geq \delta/2\}}\right] \\
&\qquad - E\left[g(S)\partial_\theta h(S)\cdot 1_{\{h(S)\geq -\delta/2\}}\right]\} \\
&= \lim_{\delta\to 0}\frac{1}{\delta}E[g(S)\partial_\theta h(S)\cdot 1_{\{-\delta/2\leq h(S)\leq \delta/2\}}] \\
&= \lim_{\delta\to 0}\frac{1}{\delta}E\left[g(S)\partial_\theta h(S)\cdot Z\left(\frac{h(S)}{\delta}\right)\right], \qquad (3)
\end{aligned}
$$

where $Z(u) = 1_{\{-1/2\leq u\leq 1/2\}}$, which is known as the uniform or naive kernel because $u\sim U(-1/2,1/2)$. It can be replaced with any other kernel, which we elaborate on in Section 3.3. In practice, Assumptions 1 and 2 are typically satisfied. Refer to Liu and Hong (2011) for the proof and a discussion of the assumptions.

3.2 Second-Order Greeks

The second-order Greeks are estimated using a theorem from Liu and Hong (2011) based on the following two assumptions:

Assumption 3 For any $(\theta_1,\theta_2)\in\Theta$, $g(S)$ and $h(S)$ are differentiable with respect to θ_1 with probability 1 (w.p.1), and $\partial_{\theta_1}g(S)$ and $\partial_{\theta_1}h(S)$ are differentiable with respect to θ_2 with probability 1 (w.p.1), and there exists variables K_g, K_h, L_g, and L_c with finite fourth moments, which may depend on (θ_1,θ_2), such that $|g(S(\theta_1+\Delta\theta_1,\theta_2)) - g(S(\theta_1,\theta_2))|\leq K_g|\Delta\theta_1|$, $|g(S(\theta_1,\theta_2+\Delta\theta_2,)) - g(S(\theta_1,\theta_2))|\leq K_g|\Delta\theta_2|$, $|h(S(\theta_1+\Delta\theta_1,\theta_2)) - h(S(\theta_1,\theta_2))|\leq K_h|\Delta\theta_1|$, $|h(S(\theta_1,\theta_2+\Delta\theta_2,)) - h(S(\theta_1,\theta_2))|\leq K_h|\Delta\theta_2|$, $|\partial_{\theta_1}g(S(\theta_1+\Delta\theta_1,\theta_2)) - \partial_{\theta_1}g(S(\theta_1,\theta_2))|\leq L_g|\Delta\theta_2|$, and $|g(S(\theta_1,\theta_2+\Delta\theta_2,))\partial_{\theta_1}h(S(\theta_1,\theta_2+\Delta\theta_2,)) - g(S(\theta_1,\theta_2))\partial_{\theta_1}h(S(\theta_1,\theta_2))|\leq L_c|\Delta\theta_2|$, when $|\Delta\theta_1|$ and $|\Delta\theta_2|$ are sufficiently small.

Assumption 4 For any $\theta\in\Theta$, $\partial_{\theta_1}\partial_{\theta_2}\phi(\theta_1,\theta_2,y)$ exists and is continuous at $(\theta_1,\theta_2,0)$, where $\phi(\theta_1,\theta_2,y) = E[g(S)\cdot 1_{\{h(S)\geq y\}}]$.

Theorem 2 Suppose $g(\cdot)$ and $h(\cdot)$ satisfy Assumptions 3 and 4, and $E[|g(S)|^4] < +\infty$ and $E[|h(S)|^4] < +\infty$, then

$$
\begin{aligned}
\partial_{\theta_1}\partial_{\theta_2}p(\theta_1,\theta_2) &= E[\partial_{\theta_1}\partial_{\theta_2}g(S)\cdot 1_{\{h(S)\geq 0\}}] \\
&\quad -\partial_y E[(g(S)\partial_{\theta_1}\partial_{\theta_2}h(S) + \partial_{\theta_1}g(S)\partial_{\theta_2}h(S) + \partial_{\theta_2}g(S)\partial_{\theta_1}h(S))\cdot 1_{\{h(S)\geq y\}}]|_{y=0} \\
&\quad +\partial_y^2 E[g(S)\partial_{\theta_1}h(S)\partial_{\theta_2}h(S)1_{\{h(S)\geq y\}}]|_{y=0}.
\end{aligned}
$$

For the purposes of this paper, we focus on the second-order Greek gamma, where $\theta = S_0$; therefore,

$$
\begin{aligned}
\partial_\theta^2 p(\theta) &= E[\partial_\theta^2 g(S)\cdot 1_{\{h(S)\geq 0\}}] \\
&\quad -\partial_y E[(g(S)\partial_\theta^2 h(S) + 2\cdot\partial_\theta g(S)\partial_\theta h(S))\cdot 1_{\{h(S)\geq y\}}]|_{y=0} \\
&\quad +\partial_y^2 E[g(S)\partial_\theta h(S)\partial_\theta h(S)1_{\{h(S)\geq y\}}]|_{y=0}.
\end{aligned}
$$

Again, the first term on the right hand side can be estimated using using a sample mean. The finite differences derivation for the second term is similar to (3), and for the finite difference representation of the third term, observe that

$$
-\partial_y E[g(S)(\partial_\theta h(S))^2\cdot 1_{\{h(S)\geq y\}}]|_{y=0} = \lim_{\delta\to 0}\frac{1}{\delta}E\left[g(S)(\partial_\theta h(S))^2\cdot Z\left(\frac{h(S)-y}{\delta}\right)\right]\Big|_{y=0},
$$

so

$$\partial_y^2 E\left[g(S)(\partial_\theta h(S))^2 \cdot 1_{\{h(S)\geq y\}}\right]|_{y=0} \quad = \quad \lim_{\delta\to 0}\frac{1}{\delta^2}E\left[g(S)(\partial_\theta h(S))^2 \cdot Z'\left(\frac{h(S)}{\delta}\right)\right], \tag{4}$$

where $Z(\cdot)$ is the kernel function. The approximation of the second-order auxiliary term (4) along with the first-order term (3) can be seen in the estimator (6).

3.3 Kernel

A d-dimensional kernel, $K:\mathbb{R}^d \to \mathbb{R}$, is a bounded symmetric density with respect to the the Lebesgue measure with $\lim_{\|u\|^d\to\infty} \| u \| K(u) = 0$ and $\int_{\mathbb{R}^d} \| u \|^2 K(u)du < \infty$ where $\|\cdot\|$ is any norm on $\mathbb{R}^d$ (Bosq 1998). A kernel K has the form

$$K(u) = \frac{1}{h}K\left(\frac{u}{h}\right),\ u\in\mathbb{R}^d,$$

where h is the bandwidth also known as the smoothing parameter (Bosq 1998). For this paper, we focus on the case for $d = 1$. According to Bosq (1998), a reasonable kernel does not affect the asymptotic behavior of the estimator; however, the bandwidth significantly influences the accuracy of the estimator. Large bandwidths decrease the variance but increase the bias, whereas small bandwidths lead to the exact opposite; hence, there is a tradeoff between variance and bias. Therefore, careful attention should be given to the selection process. Although there are other methods to quantify the accuracy of the kernel estimator, Liu and Hong (2011) measure it based on the MSE. Consequently, they attempt to select the bandwidth that maximizes the accuracy using an iterative pilot simulation.

3.4 First- and Second-Order Greek Estimators

Based on Theorem 1 and 2, the following estimators are used to approximate the first- and second-order Greeks, respectively:

$$\overline{G}_n \quad = \quad \frac{1}{n}\sum_{l=1}^{n} g_l^{'}\cdot 1_{[h_l\geq 0]} + \frac{1}{n\delta_n}\sum_{l=1}^{n} g_l\cdot h_l^{'}\cdot Z\left(\frac{h_l}{\delta_n}\right), \tag{5}$$

$$\overline{H}_n \quad = \quad \frac{1}{n}\sum_{l=1}^{n} g_l^{''}\cdot 1_{[h_l\leq 0]} + \frac{1}{n\delta_n}\sum_{l=1}^{n}\{g_l\cdot h_l^{''}+2\cdot g_l^{'}\cdot h_l^{'}\}\cdot Z\left(\frac{h_l}{\delta_n}\right) + \frac{1}{n\gamma_n^2}\sum_{l=1}^{n} g_l\cdot (h_l^{'})^2\cdot Z^{'}\left(\frac{h_l}{\gamma_n}\right),$$

where $(g_l,h_l,g_l',h_l',g_l'',h_l'')$ is the lth observation of $(g(S),h(S),\partial_\theta g(S),\partial_\theta h(S),\partial_\theta^2 g(S),\partial_\theta^2 h(S),\ Z(\cdot)$ is a kernel, and δ_n and γ_n are constant bandwidth parameters generated using a pilot simulation. The second-order estimator (6) is a special case when both derivatives are taken with respect to the same parameter.

4 PILOT SIMULATION

Ideally, the optimal bandwidth parameter would minimize the error between the Greek estimator and its true value, and the pilot simulation proposed in Liu and Hong (2011) employs an iterative method that generates an "optimal" bandwidth that minimize the asymptotic mean-squared error (MSE). The MSE of the first-order estimator can be expressed as

$$\begin{aligned}\mathrm{MSE}(\overline{G}_n) \quad &= \quad [E(\overline{G}_n) - \partial_\theta E[g(S)\cdot 1_{\{h(S)\geq 0\}}]]^2 + \mathrm{Var}(\overline{G}_n)\\ &= \quad \delta_n^4\left[\frac{\psi^{''}(0)}{2}\int_{-\infty}^{\infty}u^2 Z(u)du + \varepsilon_n\right]^2 + \frac{1}{n\delta_n}(\sigma_1^2+\xi_n),\end{aligned}$$

where $\psi(u) = f_h(u) \cdot E[g(S) \cdot \partial_\theta h(S) | h(S) = u]$, $\psi_\beta(u) = f_h(u) \cdot E[|g(S) \cdot \partial_\theta h(S)|^\beta | h(S) = u]$, and $\sigma_1^2 = \psi_2(0) \int_{-\infty}^{\infty} Z^2(u) du$. Since $\varepsilon_n \to 0$ and $\xi_n \to 0$ as $n \to \infty$, then the optimal bandwidth choice is $c \cdot n^{-1/5}$, where

$$c = \left[\frac{\sigma_1^2}{(\psi''(0) \int_{-\infty}^{\infty} u^2 Z(u) du)^2} \right]^{1/5}.$$

This constant c is estimated by

$$\hat{c} = \left[\frac{\overline{V}_n}{(\widehat{\psi}''(0) \int_{-\infty}^{\infty} u^2 Z(u) du)^2} \right]^{1/5},$$

where $\overline{V}_n$ is the sample variance of $\overline{G}_n$ and $\widehat{\psi}^{''}(0)$ is the finite difference approximation of $\psi^{''}(0)$.

$$\widehat{\psi}^{''}(0) = \frac{\overline{G}_n(s) + \overline{G}_n(-s) - 2\overline{G}_n(0)}{s^2},$$

where s is a sufficiently small step size and

$$\overline{G}_n(u) = \frac{1}{n\delta_n} \sum_{l=1}^{n} g_l \cdot h_l^{'} \cdot Z\left(\frac{h(S) - u}{\delta_n} \right).$$

Therefore, the optimal bandwidth is approximated by $\delta_n^\star = \hat{c} \cdot n^{-1/5}$. The pilot simulation begins with $\hat{c} = 1$, and is updated after each iteration using the same sample. More specifically, at each iteration, $\overline{V}_n$ and $\widehat{\psi}^{''}(0)$ are simulated using the updated $\delta_n^\star$ with $n = 500$. This process continues for a pre-specified number of iterations, and the bandwidth selected is the one generated after 30 pilot iterations, as in Liu and Hong (2011). The pilot simulation algorithm does not specify the variance sample size n_v to generate $\overline{V}_n$, step size s to estimate $\psi''(0)$, and the number of iterations for the pilot simulation n_p, which are all predetermined. Refer to the e-companion of Liu and Hong (2011) for the algorithm details and the pilot simulation for the second-order Greeks.

5 NUMERICAL EXPERIMENTS

We conduct two sets of numerical experiments to test the robustness of a modified pathwise method and a pilot simulation used to generate an input parameter for the modified method. The first is a sensitivity analysis of Greek estimators for both the Asian digital option and up-and-out barrier call option to bandwidth parameters, and the second explores the effect of the input parameters to the bandwidth generated from the pilot simulation. We considered identical settings as in Liu and Hong (2011), which are shown in Table 1 and described in Section 2. However, the sample size considered here is 100 as opposed to 1000 in Liu and Hong (2011). From our preliminary results, the increase in sample size only tightens the confidence band and smooths out the curves, but the overall behavior is similar across estimators, so we opt to use a smaller sample size. Furthermore, we also consider the standard normal density as the kernel to increase the robustness of the resulting estimator.

5.1 Sensitivity of Greeks to Bandwidth

In our experiment for each Greek, we generate 100 estimators for incrementally increasing bandwidths and observe the effect on the relative root mean-squared error (RRMSE) and the 95% confidence band (CB). Figures 1, 2, 3, and 4 plot the 95% confidence band for the Greek estimator generated using 1000 sample paths for Asian Greeks, 10000 sample paths for barrier Greeks, and sample size of 100 for incrementally increasing $\hat{c}$. The black curve in between is the mean, and the horizontal dashed line represents the exact

Table 2: Asian delta.

n	k	c^*	RRMSE	lower	upper	range	δ_n^*	δ_n for s 0.1	0.01	0.001
500	10	0.64	0.028	0.915	1.021	0.106	0.185	0.14	0.06	0.06
	20	0.46	0.038	0.975	1.132	0.156	0.133	0.15	0.05	0.05
	50	0.54	0.033	0.998	1.130	0.132	0.156	0.14	0.06	0.06
1000	10	0.74	0.020	0.931	1.005	0.074	0.186	0.14	0.05	0.04
	20	0.54	0.028	0.996	1.104	0.108	0.136	0.13	0.05	0.05
	50	0.60	0.024	1.019	1.117	0.099	0.151	0.13	0.05	0.05
2000	10	0.84	0.014	0.943	0.996	0.053	0.184	0.13	0.03	0.03
	20	0.56	0.021	1.012	1.097	0.084	0.122	0.13	0.03	0.02
	50	0.68	0.017	1.034	1.105	0.071	0.149	0.13	0.03	0.02
5000	10	1.00	0.009	0.954	0.988	0.035	0.182	0.15	0.06	0.06
	20	0.66	0.014	1.027	1.083	0.055	0.120	0.16	0.07	0.07
	50	0.80	0.011	1.050	1.094	0.044	0.146	0.16	0.06	0.06
10000	10	1.14	0.007	0.960	0.985	0.020	0.181	0.15	0.03	0.02
	20	0.72	0.011	1.036	1.076	0.040	0.114	0.12	0.03	0.02
	50	0.92	0.008	1.056	1.088	0.032	0.146	0.13	0.03	0.02

Greek value. The three vertical lines in Figures 1, 2, and 3 represent the 95% confidence interval for the $\hat{c}$ generated using the pilot simulation with $n_p = 30, n_v = 100$, which we will elaborate on in the next section. In the same figures, the second row of graphs plots the RRMSE of the 100 generated Greeks for incrementally increasing $\hat{c}$, where the additional vertical dashed line denotes the c^* that minimizes the RRMSE. The only difference when $n = 1000$ increases to $n = 10000$ for the Asian Greek estimators is the increased in smoothness of the curves and tighter confidence bands, but again, the overall behavior is the same. Liu and Hong (2011) report results for $n = 10^3, 10^4, 10^5$ for the Asian Greeks, but we only test the case for $n = 10^3$.

Not surprisingly, the bandwidth affects the accuracy and precision of the Greek estimator. Generally, as the bandwidth decreases, the variance increases, which can be seen in the widening of the 95% confidence band as $\hat{c} \to 0$ and is observable in all of the estimators. Similarly, the 95% confidence band decreases as the bandwidth increases, but eventually the confidence band no longer contains the exact Greek value. Figures 2, 3, and 4 for the Asian vega, Asian gamma, and barrier theta estimators illustrate this behavior, respectively. In addition, the minimum RRMSE of these estimates occur in one particular region for all discretization levels. Figure 2 shows that the Asian vega estimator is rather close to the exact value for a range of $\hat{c}$ values less than 0.1. The barrier theta estimator behaves similarly but for a different range of $\hat{c}$ values as seen in Figure 4. In contrast, the Asian gamma estimator is extremely sensitive to $\hat{c}$, and the RRMSE increases significantly with even slight deviations from the optimal bandwidth. Moreover, the 95% confidence interval (CI) for the bandwidth parameter is very narrow and does not contain the exact Asian gamma value.

The numerical results for the Asian delta are quite different as illustrated in Figure 1. The mean of the delta estimator intersects with the exact/analytical value at least twice for $k = 10$ and $k = 50$. For the case $k = 10$, the RRMSE has a minimum and a nearly flat region. The minimum of the RRMSE for each discretization case occurs at different points for a given number of sample paths n, unlike the Asian vega, Asian gamma, and barrier theta. The exact $\hat{c}$ values at which the minimums occurs, the resulting RRMSE, and 95% confidence interval of the Asian delta estimators generated using the $\hat{c}$ are listed in Table 2.

Another interesting case is the gamma estimator for the up-and-out barrier call option, which relies on two bandwidth parameters δ_n and γ_n. The Asian gamma is also a second-order Greek, but the second term

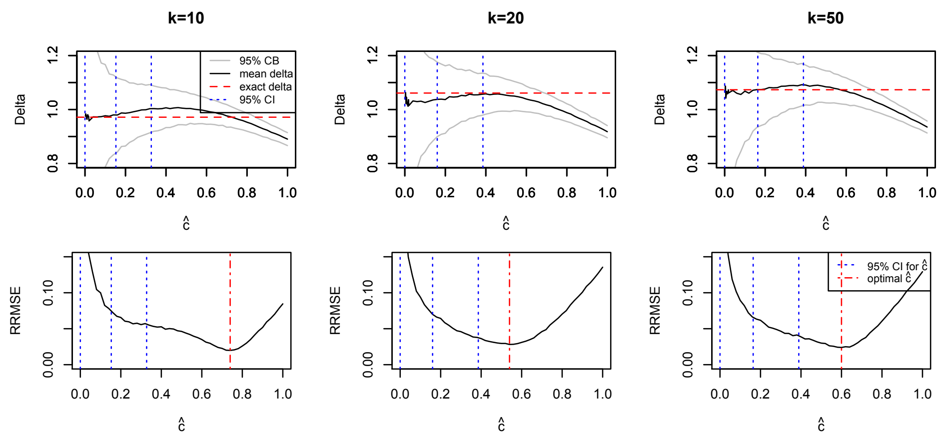

Figure 1: 95% confidence band for Asian delta (solid curves) for $n = 1000$ with w/95% confidence interval for bandwidth (vertical dashed lines) generated using $s = 0.1$ and RRMSE for 100 sample paths.

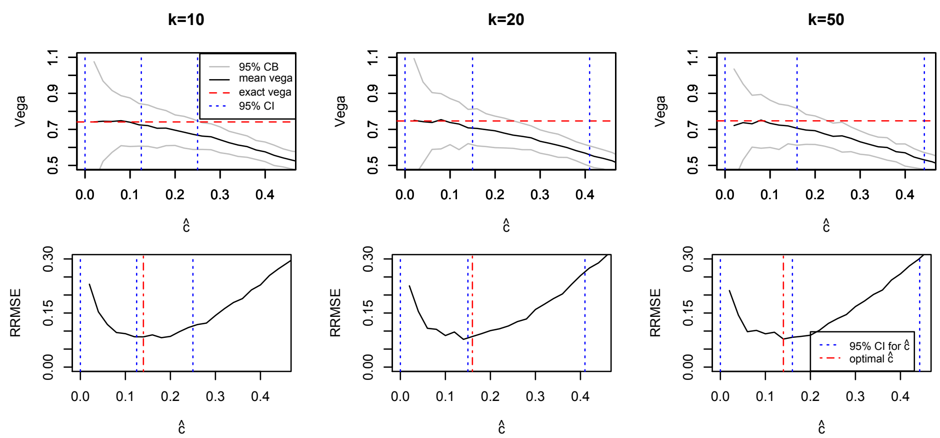

Figure 2: 95% confidence band for Asian vega (solid curves) for $n = 1000$ w/95% confidence interval for bandwidth (vertical dashed lines) generated using $s = 0.001$ and RRMSE for 100 sample paths.

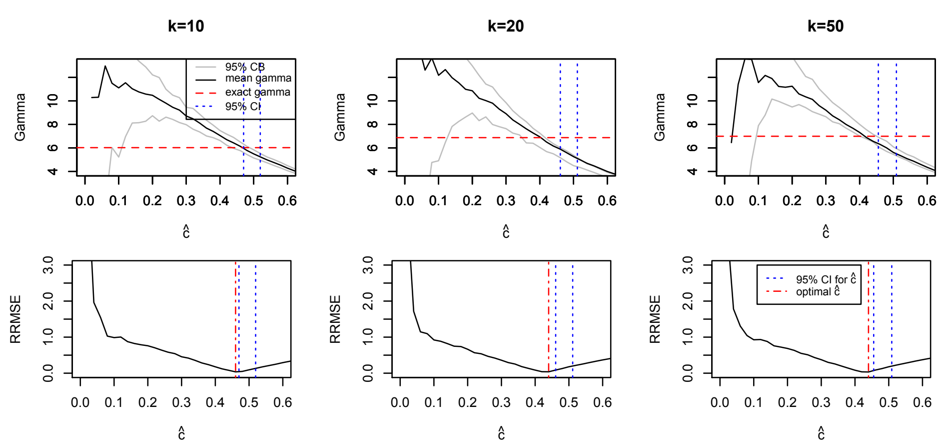

Figure 3: 95% confidence band for Asian gamma (solid curves) for $n = 1000$ w/95% confidence interval for bandwidth (vertical dashed lines) generated using $s = 0.01$ and RRMSE for 100 sample paths.

in (5) containing δ_n disappears with the given settings. Figure 5 plots the barrier gamma as a function of the bandwidth parameters δ_n (visible horizontal axis) and γ_n (not visible), where the vertical planes represent the approximately exact gamma for $k = 20, 50$, respectively. It is clear that gamma is insensitive to γ_n, but extremely sensitive to δ_n, especially when $\delta_n < .03$. Unfortunately, this is the region in which the estimator attains the value closest to the true value. Since the analytical expression is not available for comparison, the true values are approximated using finite differences with a large sample of 10^9. In this case, the steepness of the gamma estimator near the "exact" value prompts the need for a very accurate and precise δ_n; otherwise, the estimate could be drastically different from its true value. As expected, the margin of error decreases as the number of sample paths n increases. For each discretization level k, $\hat{c}$ increases with n, but the resulting δ_n values are similar for each k regardless of n.

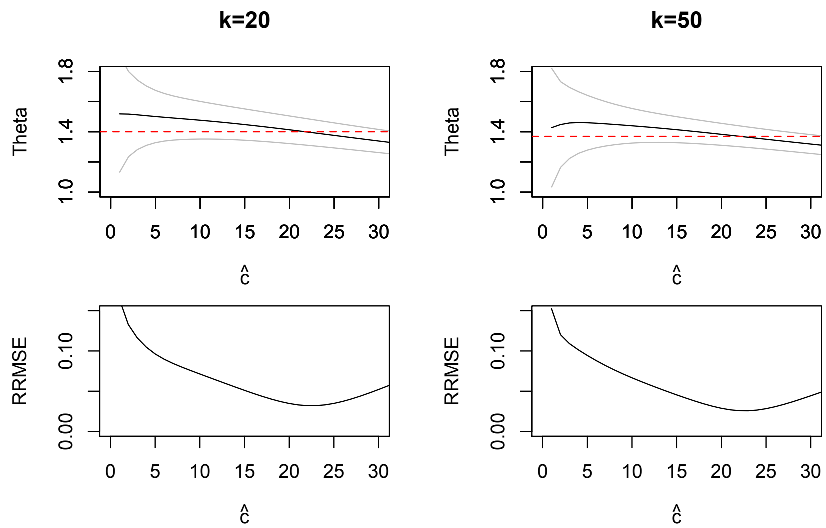

Figure 4: 95% confidence band for barrier theta for $n = 10000$ and RRMSE for 100 sample paths.

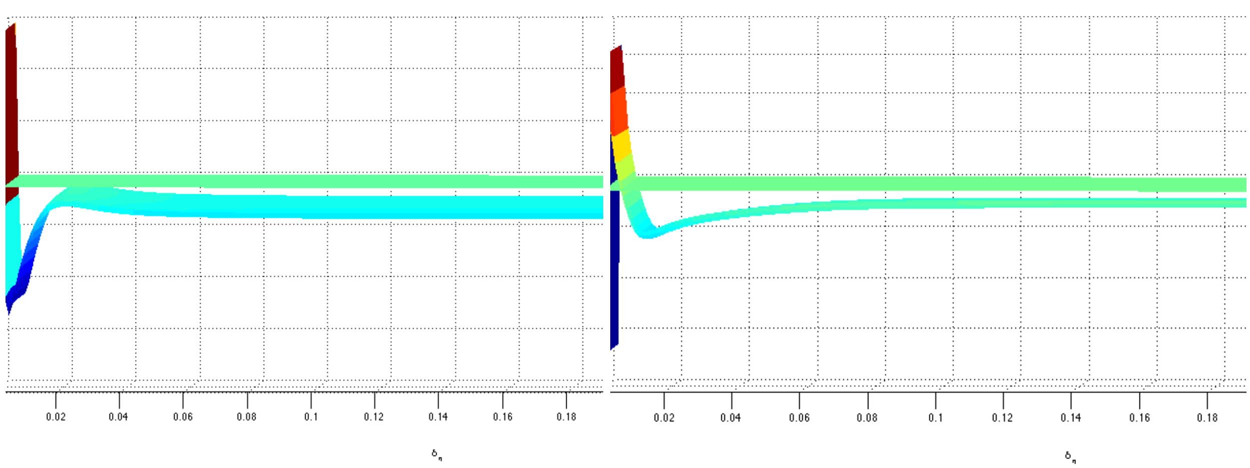

Figure 5: Barrier gamma estimator for $n = 10000$, $k = 20$ (left), $k = 50$ (right) for 100 sample paths.

The conclusions in this section are representative of all of the estimators tested. For instance, the Asian delta is the only estimator whose mean matches the exact Greek value for two different bandwidth values and the barrier gamma is the only estimator that requires the input of two bandwidth parameters. The Asian vega, in general, has a decreasing 95% confidence band whose mean also decreases with $\hat{c}$ and matches the exact Greek value for one particular $\hat{c}$, and there is one global minimum RRMSE, which represents the results of the remaining Greek estimators.

5.2 Sensitivity of Bandwidth to Pilot Simulation

Our sensitivity analysis of the modified pathwise method applied to Greeks concludes that it is, in general, very sensitive to the bandwidth parameter(s), so the key to a good estimator is in the bandwidth selection process. Therefore, we investigate the sensitivity of the pilot simulation to the various input parameters such as variance sample size n_v, perturbation size s, number of sample paths n, and number of pilot iterations n_p, as described in Section 4. We considered the following settings: $n_v \in \{50, 100, 500, 1000\}$, $s \in \{0.001, 0.01, 0.1\}$, $n \in \{500, 1000, 2000, 5000\}$, and $n_p \in \{1, \ldots, 100\}$. For each setting, we compute the 95% confidence interval for the bandwidth output and compare it against the optimal bandwidth that minimizes the RRMSE in Table 2. The pilot simulation experiments consider a variance sample size of 100 (i.e., $n_v = 100$) and 30 pilot iterations (i.e, $n_p = 30$), unless stated otherwise. Over an extensive set of parameter settings, the results are quite similar, so we limit our discussion to representative cases.

First, we discuss the results of the pilot simulation for the Asian delta. For any fixed perturbation size, the average $\hat{c}$ under the Asian delta estimator is quite similar across sample paths and discretization level k, with a maximum range of 0.04. As n increases, the average $\hat{c}$ decreases ever so slightly. One would think that variance of $\hat{c}$ decreases with n, i.e., the confidence band would decrease with the n; however, the perturbation size s appears to be the determining factor from Figure 6. In this case, the smaller perturbation

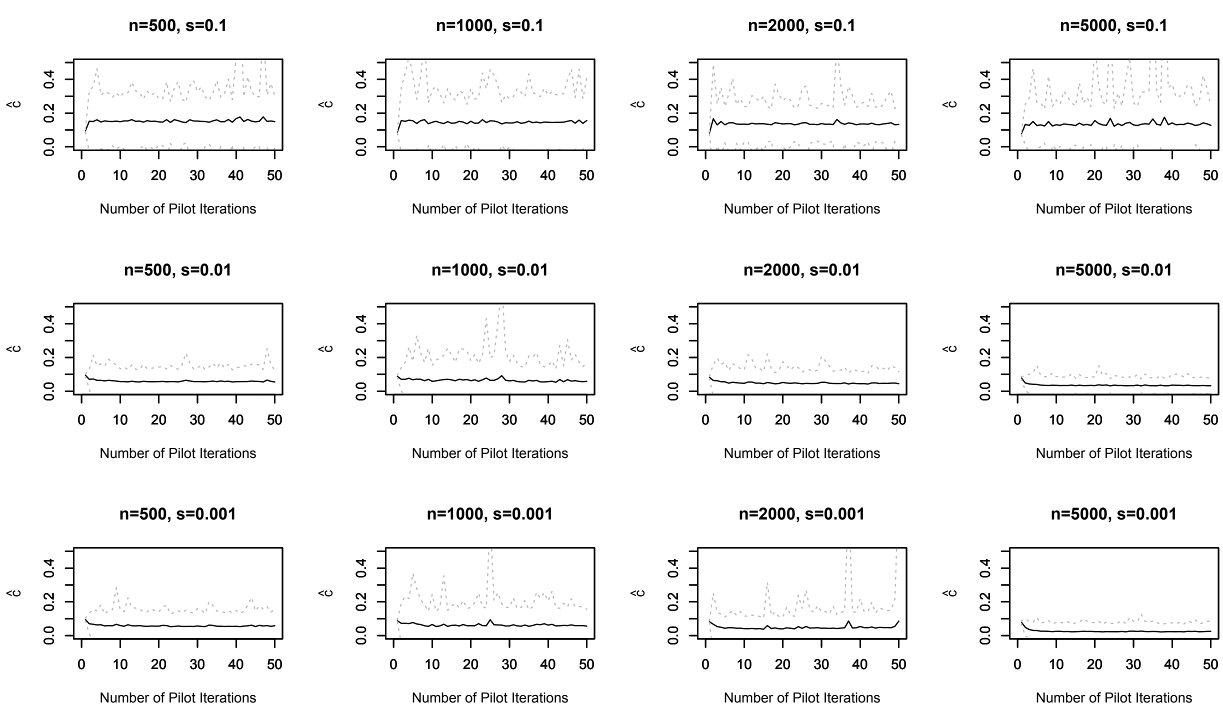

Figure 6: Asian delta pilot 95% confidence interval for $\hat{c}$ for $k = 10$.

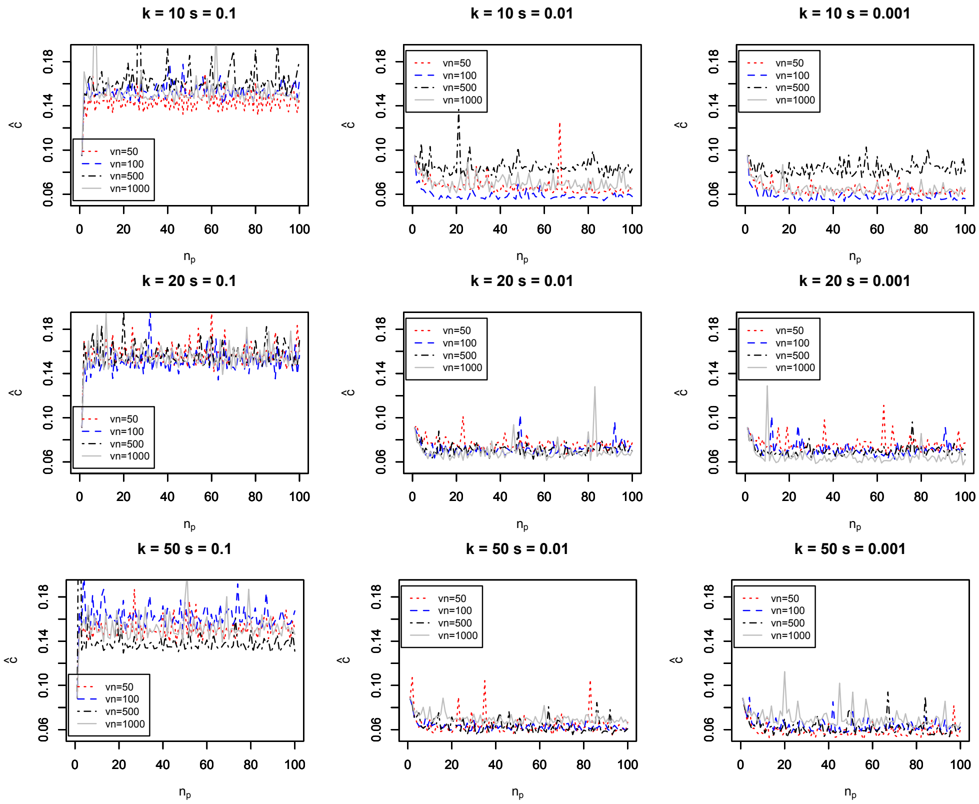

Figure 7: Asian delta pilot, $n = 500$, sample size $= 100$.

size $s = 0.1$ leads to wider confidence intervals, which is interesting, because smaller perturbations tend to generate noisier estimates. Therefore, it could be advantageous to consider an average or a window average, as opposed to selecting the $\hat{c}$ generated after a certain number of iterations. For instance, the average $\hat{c}$ fluctuates, especially in the case of $s = 0.1$, so taking a further average of multiple pilot iterations at the later stage, could generate a parameter that is more stable.

Another input parameter that was not specified is the sample size n_v used to generate the sample variance $\overline{V}_n$. So we investigate the sensitivity of the bandwidth parameter to n_v for each of the perturbation sizes s. Again, we focus on the Asian delta case for the discretization case $k = 10$, sample path $n = 1000$, vary $n_v \in \{50, 100, 500, 1000\}$, and $s \in \{0.001, 0.01, 0.1\}$. Figure 7 plots the generated $\hat{c}$ against the number of pilot simulations n_p for a sample size of 100. The results of the sensitivity to n_v are inconclusive, since the behavior is inconsistent across different parameter settings. For this particular example, the larger perturbation size $s = 0.1$ leads to a higher bandwidth parameter, which can also be seen in Figure 6.

Clearly, the input parameters have a significant impact on the bandwidth parameter generated from the pilot simulation, and we investigate further to determine whether or not the confidence interval of the bandwidth parameter captures the bandwidth that minimizes the RRMSE for parameters $n = 1000$, $n_v = 100$, $n_p = 30$, sample size 100, $s = 0.1$ for the delta, $s = 0.01$ for the gamma, and $s = 0.001$ for both vega and theta Asian Greek estimators. Values for the perturbation parameters were not specified in Liu and Hong (2011), but were obtained from the authors via personal communication. The sizes were chosen relative to the parameter of interest. Our results vary significantly, from successfully capturing the optimal bandwidth to completely missing it. Figures 1, 2, and 3 contain vertical dotted lines representing the 95% confidence interval for the optimal c^* with the above specified parameters, the dotted line denotes the mean, and the dashed line is the c^* that minimizes the RRMSE from our first set of numerical experiments. None of the 95% confidence intervals of $\hat{c}$ capture the minimum RRMSE for the Asian delta Greek estimator. Instead, the CI contains mean $\hat{c}^*$ The confidence interval for the $\hat{c}$ generated from the pilot simulation for the Asian vega, covers the minimum RRMSE bandwidth, which can be seen in Figure 2.

6 CONCLUSION

In this paper, we analyzed the sensitivity of the Greek kernel estimators of Liu and Hong (2011) to the bandwidth parameters and the sensitivity and performance of the pilot simulation in selecting the "optimal" bandwidth parameters to input parameters. The simulated experiments show that the kernel estimators are quite sensitive to the bandwidth choice, and the pilot simulation is sensitive to the input parameters. Of the input parameters, the perturbation step size s significantly impacts the generated bandwidth. If s is too small relative to the parameter of interest, then the pilot simulation could generate bandwidths that lead to poor gradient estimates. However, even if the s is chosen proportional to the parameter of interest, the 95% confidence interval for the $\hat{c}$ might not capture the c^* that minimizes the RRMSE. The input parameters play a vital role in the performance of the estimator, so the input selection process of the pilot simulation and other bandwidth selection methods are important future research areas.

REFERENCES

Asmussen, S., and P. Glynn. 2007. *Stochastic Simulation: Algorithms and Analysis*. New York: Springer.

Bosq, D. 1998. *Nonparametric Statistics for Stochastic Processes*. Second ed. Springer.

Fu, M. 2006. "Gradient Estimation". In *Handbooks in Operations Research and Management Science: Simulation*, edited by S. G. Henderson and B. L. Nelson, Chapter 19, 575–616. Elsevier.

Fu, M., and J.-Q. Hu. 1997. *Conditional Monte Carlo, Gradient Estimation and Optimization Applications*. Boston: Kluwer Academic Publishers.

Fu, M. C. 2008. "What You Should Know About Simulation and Derivatives". *Naval Research Logistics* 55 (8): 723 – 736.

Glasserman, P. 2004. *Monte Carlo Methods in Financial Engineering*. New York, NY: Springer.

Liu, G., and L. Hong. 2011. "Kernel estimation of the Greeks for options with discontinuous payoffs". *Operations Research* 59 (1): 96–108.

Wang, Y., M. Fu, and S. Marcus. 2009. "Sensitivity analysis for barrier options". In *Proceedings of the 2009 Winter Simulation Conference*, edited by M. Rossetti, R. Hill, B. Johansson, and R. Ingalls, 1272 – 1282. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

MARIE CHAU is a Ph.D. candidate in Applied Mathematics, Statistics, and Scientific Computation at the University of Maryland. She has a M.S. in applied math and degrees in math, finance and economics from UMCP. Her research interests lie in stochastic approximation, stochastic optimization and financial engineering. Her email address is mchau@math.umd.edu.

MICHAEL C. FU is Ralph J. Tyser Professor of Management Science in the Robert H. Smith School of Business, with a joint appointment in the Institute for Systems Research and affiliate faculty appointment in the Department of Electrical and Computer Engineering, all at the University of Maryland. His research interests include simulation optimization and applied probability, with applications in supply chain management and financial engineering. He has a Ph.D. in applied math from Harvard and degrees in math and EECS from MIT. He served as WSC2011 Program Chair, NSF Operations Research Program Director, *Management Science* Stochastic Models and Simulation Department Editor, and *Operations Research* Simulation Area Editor. He is a Fellow of **INFORMS** and **IEEE**. His email address is mfu@umd.edu.