ELSEVIER

Journal of Financial Economics 55 (2000) 173–204

JOURNAL OF Financial ECONOMICS

www.elsevier.com/locate/econbase

# When is time continuous?☆

Dimitris Bertsimas[a], Leonid Kogan[b], Andrew W. Lo[a,*]

[a]*MIT Sloan School of Management, 50 Memorial Drive, Cambridge, MA 02142-1347, USA*
[b]*Department of Finance, Wharton School, University of Pennsylvania, Philadelphia, PA 19104-6367, USA*

**Abstract**

Continuous-time stochastic processes are approximations to physically realizable phenomena. We quantify one aspect of the approximation errors by characterizing the asymptotic distribution of the replication errors that arise from delta-hedging derivative securities in discrete time, and introducing the notion of *temporal granularity* which measures the extent to which discrete-time implementations of continuous-time models can track the payoff of a derivative security. We show that granularity is a particular function of a derivative contract's terms and the parameters of the underlying stochastic process. Explicit expressions for the granularity of geometric Brownian motion and an Ornstein–Uhlenbeck process for call and put options are derived, and we perform Monte Carlo simulations to illustrate the empirical properties of granularity. © 2000 Elsevier Science S.A. All rights reserved.

*JEL classification:* G13

*Keywords:* Derivatives; Delta hedging; Continuous-time models

☆This research was partially supported by the MIT Laboratory for Financial Engineering, the National Science Foundation (Grant No. SBR-9709976) and a Presidential Young Investigator Award (Grant No. DDM-9158118) with matching funds from Draper Laboratory. We thank Michael Brandt, George Martin, Matthew Richardson, Bill Schwert, Jiang Wang, an anonymous referee, and seminar participants at Columbia, the Courant Institute, the Federal Reserve Board, the London School of Economics, the MIT finance lunch group, the MIT LIDS Colloquium, the NBER Asset Pricing Program, NYU, Renaissance Technologies, the Wharton School, the University of Chicago, the University of Massachusetts at Amherst, the University of Southern California, and the 1997 INFORMS Applied Probability Conference for valuable discussions and comments.

* Corresponding author. Tel.: + 1-617-2530920; fax: + 1-617-2585727.

*E-mail address:* alo@mit.edu (A.W. Lo)


PII: S0304-405X(99)00049-5

## 1. Introduction

Since Wiener's (1923) pioneering construction of Brownian motion and Itô's (1951) theory of stochastic integrals, continuous-time stochastic processes have become indispensable to many disciplines ranging from chemistry and physics to engineering to biology to financial economics. In fact, the application of Brownian motion to financial markets (Bachelier, 1900) pre-dates Wiener's contribution by almost a quarter century, and Merton's (1973) seminal derivation of the Black and Scholes (1973) option-pricing formula in continuous time – and, more importantly, his notion of delta hedging and dynamic replication – is often cited as the foundation of today's multitrillion-dollar derivatives industry.

Indeed, the mathematics and statistics of Brownian motion have become so intertwined with so many scientific theories that we often forget the fact that continuous-time processes are only approximations to physically realizable phenomena. In fact, for the more theoretically inclined, Brownian motion may seem more "real" than discrete-time discrete-valued processes. Of course, whether time is continuous or discrete is a theological question best left for philosophers. But a more practical question remains: under what conditions are continuous-time models good approximations to specific physical phenomena, i.e., when does time seem "continuous" and when does it seem "discrete"? In this paper, we provide a concrete answer to this question in the context of continuous-time derivative-pricing models, e.g., Merton (1973), by characterizing the replication errors that arise from delta hedging derivatives in discrete time.

Delta-hedging strategies play a central role in the theory of derivatives and in our understanding of dynamic notions of spanning and market completeness. In particular, delta-hedging strategies are recipes for replicating the payoff of a complex security by sophisticated dynamic trading of simpler securities. When markets are dynamically complete (e.g., Harrison and Kreps, 1979; Duffie and Huang, 1985) and continuous trading is feasible, it is possible to replicate certain derivative securities perfectly. However, when markets are not complete or when continuous trading is not feasible, e.g., when there are trading frictions or periodic market closings, perfect replication is not possible and the usual delta-hedging strategies exhibit *tracking* errors. These tracking errors are the focus of our attention.

Specifically, we characterize the asymptotic distribution of the tracking errors of delta-hedging strategies using continuous-record asymptotics, i.e., we implement these strategies in discrete time and let the number of time periods increase while holding the time span fixed. Since the delta-hedging strategies we consider are those implied by continuous-time models like Merton (1973), it is not surprising that tracking errors arise when such strategies are implemented in discrete time, nor is it surprising that these errors disappear in the limit of continuous time. However, by focusing on the continuous-record asymptotics of the tracking error, we can quantify the discrepancy between the discrete-time hedging strategy and its continuous-time limit, answering the question "When is time continuous?" in the context of replicating derivative securities.

We show that the normalized tracking error converges weakly to a particular stochastic integral and that the root-mean-squared tracking error is of order $N^{-1/2}$ where $N$ is the number of discrete time periods over which the delta hedging is performed. This provides a natural definition for *temporal granularity*: it is the coefficient that corresponds to the $\mathbf{O}(N^{-1/2})$ term. We derive a closed-form expression for the temporal granularity of a diffusion process paired with a derivative security, and propose this as a measure of the "continuity" of time. The fact that granularity is defined with respect to a derivative-security/price-process pair underscores the obvious: there is a need for specificity in quantifying the approximation errors of continuous-time processes. It is impossible to tell how good an approximation a continuous-time process is to a physical process without specifying the nature of the physical process.

In addition to the general usefulness of a measure of temporal granularity for continuous-time stochastic processes, our results have other, more immediate applications. For example, for a broad class of derivative securities and price processes, our measure of granularity provides a simple method for determining the approximate number of hedging intervals $N^*$ needed to achieve a target root-mean-squared error $\delta$: $N^* \approx g^2/\delta^2$ where $g$ is the granularity coefficient of the derivative-security/price-process pair. This expression shows that to halve the root-mean-squared error of a typical delta-hedging strategy, the number of hedging intervals must be increased approximately fourfold.

Moreover, for some special cases, e.g., the Black–Scholes case, the granularity coefficient can be obtained in closed form, and these cases shed considerable light on several aspects of derivatives replication. For example, in the Black–Scholes case, does an increase in volatility make it easier or more difficult to replicate a simple call option? Common intuition suggests that the tracking error increases with volatility, but the closed-form expression (3.2) for granularity shows that it achieves a maximum as a function of $\sigma$ and that beyond this point, granularity becomes a decreasing function of $\sigma$. The correct intuition is that at lower levels of volatility, tracking error is an increasing function of volatility because an increase in volatility implies more price movements and a greater likelihood of hedging errors in each hedging interval. But at higher levels of volatility, price movements are so extreme that an increase in volatility in this case implies that prices are less likely to fluctuate near the strike price where delta-hedging errors are the largest, hence granularity is a decreasing function of $\sigma$. In other words, at sufficiently high levels of volatility, the nonlinear payoff function of a call option "looks" approximately linear and is therefore easier to hedge. Similar insights can be gleaned from other closed-form expressions of granularity (see, for example, Section 3.2).

In Section 2, we provide a complete characterization of the asymptotic behavior of the tracking error for delta hedging an arbitrary derivative security, and formally introduce the notion of granularity. To illustrate the practical relevance of granularity, in Section 3 we obtain closed-form expressions for granularity in two specific cases: call options under geometric Brownian motion, and under a mean-reverting process. In Section 4 we check the accuracy of our continuous-record asymptotic approximations by presenting Monte Carlo simulation experiments for the two examples of Section 3 and comparing them to the corresponding analytical expressions. We present other extensions and generalizations in Section 5, including a characterization of the sample-path properties of tracking errors, the joint distributions of tracking errors and prices, a PDE characterization of the tracking error, and more general loss functions than root-mean-squared tracking error. We conclude in Section 6.

## 2. Defining temporal granularity

The relation between continuous-time and discrete-time models in economics and finance has been explored in a number of studies. One of the earliest examples is Merton (1969), in which the continuous-time limit of the budget equation of a dynamic portfolio choice problem is carefully derived from discrete-time considerations (see also Merton, 1975, 1982). Foley's (1975) analyis of 'beginning-of-period' versus 'end-of-period' models in macroeconomics is similar in spirit, though quite different in substance.

More recent interest in this issue stems primarily from two sources. On the one hand, it is widely recognized that continuous-time models are useful and tractable approximations to more realistic discrete-time models. Therefore, it is important to establish that key economic characteristics of discrete-time models converge properly to the characteristics of their continuous-time counterparts. A review of recent research along these lines can be found in Duffie and Protter (1992).

On the other hand, while discrete-time and discrete-state models such as those based on binomial and multinomial trees, e.g., Cox et al. (1979), He (1990, 1991), and Rubinstein (1994), may not be realistic models of actual markets, nevertheless they are convenient computational devices for analyzing continuous-time models. Willinger and Taqqu (1991) formalize this notion and provide a review of this literature.

For derivative-pricing applications, the distinction between discrete-time and continuous-time models is a more serious one. For all practical purposes, trading takes place at discrete intervals, and a discrete-time implementation of Merton's (1973) continuous-time delta-hedging strategy cannot perfectly replicate an option's payoff. The tracking error that arises from implementing a continuous-time hedging strategy in discrete time has been studied by several authors. One of the first studies was conducted by Boyle and Emanuel (1980), who consider the statistical properties of "local" tracking errors. At the beginning of a sufficiently small time interval, they form a hedging portfolio of options and stock according to the continuous-time Black–Scholes/Merton delta-hedging formula. The composition of this hedging portfolio is held fixed during this time interval, which gives rise to a tracking error (in continuous time, the composition of this portfolio would be adjusted continuously to keep its dollar value equal to zero). The dollar value of this portfolio at the end of the interval is then used to quantify the tracking error.

More recently, Toft (1996) shows that a closed-form expression for the variance of the cash flow from a discrete-time delta-hedging strategy can be obtained for a call or put option in the special case of geometric Brownian motion. However, he observes that this expression is likely to span several pages and is therefore quite difficult to analyze.

But perhaps the most relevant literature for our purposes is Leland's (1985) investigation of discrete-time delta-hedging strategies motivated by the presence of transactions costs, an obvious but important motivation (why else would one trade discretely?) that spurred a series of studies on option pricing with transactions costs, e.g., Figlewski (1989), Hodges and Neuberger (1989), Bensaid et al. (1992), Boyle and Vorst (1992), Edirisinghe et al. (1993), Henrotte (1993), Avellaneda and Paras (1994), Neuberger (1994), and Grannan and Swindle (1996). This strand of the literature provides compelling economic motivation for discrete delta-hedging – trading continuously would generate infinite transactions costs. However, the focus of these studies is primarily the *tradeoff* between the magnitude of tracking errors and the cost of replication. Since we focus on only one of these two issues – the approximation errors that arise from applying continuous-time models discretely – we are able to characterize the statistical behavior of tracking errors much more generally, i.e., for large classes of price processes and payoff functions.

Specifically, we investigate the discrete-time implementation of continuous-time delta-hedging strategies and derive the asymptotic distribution of the tracking error in considerable generality by appealing to continuous-record asymptotics. We introduce the notion of temporal granularity which is central to the issue of when time may be considered continuous, i.e., when continuous-time models are good approximations to discrete-time phenomena. In Section 2.1, we describe the framework in which our delta-hedging strategy will be implemented and define tracking error and related quantities. In Section 2.2, we characterize the continuous-record asymptotic behavior of the tracking error and define the notion of temporal granularity. We provide an interpretation of granularity in Section 2.3 and discuss its implications.

### 2.1. Delta hedging in complete markets

We begin by specifying the market environment. For simplicity, we assume that there are only two traded securities: a riskless asset (bond) and a risky asset (stock). Time $t$ is normalized to the unit interval so that trading takes place from $t = 0$ to $t = 1$. In addition, we assume the following:

(A1) markets are frictionless, i.e., there are no taxes, transactions costs, short-sales restrictions, or borrowing restrictions;
(A2) the riskless borrowing and lending rate is zero; and
(A3) the price $P_t$ of the risky asset follows a diffusion process

$$\frac{\mathrm{d}P_t}{P_t} = \mu(t, P_t)\,\mathrm{d}t + \sigma(t, P_t)\,\mathrm{d}W_t, \quad \sigma(t, P_t) \geqslant \sigma_0 > 0 \tag{2.1}$$

where the coefficients $\mu$ and $\sigma$ satisfy standard regularity conditions that guarantee existence and uniqueness of the strong solution of (2.1) and market completeness (see Duffie, 1996).

Note that Assumption (A1) entails little loss of generality since we can always renormalize all prices by the price of a zero-coupon bond with maturity at $t = 1$ (e.g., Harrison and Kreps, 1979). However, this assumption does rule out the case of a stochastic interest rate.

We now introduce a European derivative security on the stock that pays $F(P_1)$ dollars at time $t = 1$. We will call $F(\cdot)$ the payoff function of the derivative. The equilibrium price of the derivative, $H(t, P_t)$, satisfies the following partial differential equation (PDE) (e.g., Cox et al., 1985):

$$\frac{\partial H(t,x)}{\partial t} + \frac{1}{2}\sigma^2(t,x)x^2\frac{\partial^2 H(t,x)}{\partial x^2} = 0 \tag{2.2}$$

with the boundary condition

$$H(1,x) = F(x). \tag{2.3}$$

This is a generalization of the standard Black–Scholes model which can be obtained as a special case when the coefficients of the diffusion process (2.1) are constant, i.e., $\mu(t, P_t) = \mu$, $\sigma(t, P_t) = \sigma$, and the payoff function $F(P_1)$ is given by $\mathrm{Max}[P_1 - K, 0]$ or $\mathrm{Max}[K - P_1, 0]$.

The delta-hedging strategy was introduced by Black and Scholes (1973) and Merton (1973) and when implemented continuously on $t \in [0, 1]$, the payoff of the derivative at expiration can be replicated perfectly by a portfolio of stocks and riskless bonds. This strategy consists of forming a portfolio at time $t = 0$ containing only stocks and bonds with an initial investment of $H(0, P_0)$ and rebalancing it continuously in a *self-financing* manner – all long positions are financed by short positions and no money is withdrawn or added to the portfolio – so that at all times $t \in [0, 1]$ the portfolio contains $\partial H(t, P_t)/\partial P_t$ shares of the stock. The value of such a portfolio at time $t = 1$ is exactly equal to the payoff, $F(P_1)$, of the derivative. Therefore, the price, $H(t, P_t)$, of the derivative can also be considered the production cost of replicating the derivative's payoff $F(P_1)$ starting at time $t$.

Such an interpretation becomes important when continuous-time trading is not feasible. In this case, $H(t, P_t)$ can no longer be viewed as the equilibrium price of the derivative. However, the function $H(t, P_t)$, defined formally as a solution of (2.2)–(2.3), can still be viewed as the production cost $H(0, P_0)$ of an approximate replication of the derivative's payoff, and can be used to define the production process itself (we formally define a discrete-time delta-hedging strategy below). The term "approximate replication" indicates the fact that when continuous trading is not feasible, the difference between the payoff of the derivative and the end-of-period dollar value of the replicating portfolio will not, in general, be zero; Bertsimas et al. (1997) discuss derivative replication in discrete time and the distinction between production cost and equilibrium price. Accordingly, when we refer to $H(t, P_t)$ as the derivative's "price" below, we shall have in mind this more robust interpretation of production cost and approximate replication strategy.[1]

More formally, we assume:

(A4) trading takes place only at $N$ regularly spaced times $t_i$, $i = 1, \ldots, N$, where

$$t_i \in \left\{0, \frac{1}{N}, \frac{2}{N}, \ldots, \frac{N-1}{N}\right\}.$$

Under (A4), the difference between the payoff of the derivative and the end-of-period dollar value of the replicating portfolio – the *tracking error* – will be nonzero.

Following Hutchinson et al. (1994), let $V_{t_i}^{(N)}$ be the value of the replicating portfolio at time $t_i$. Since the replicating portfolio consists of shares of the stock and the bond, we can express $V_{t_i}^{(N)}$ as

$$V_{t_i}^{(N)} = V_{S,t_i}^{(N)} + V_{B,t_i}^{(N)} \tag{2.4}$$

where $V_{S,t_i}^{(N)}$ and $V_{B,t_i}^{(N)}$ denote the dollar amount invested in the stock and the bond, respectively, in the replicating portfolio at time $t_i$. At time $t = 0$ the total

[1] Alternatively, we can conduct the following equivalent thought experiment: while some market participants can trade costlessly and continuously in time and thus ensure that the price of the derivative is given by the solution of (2.2) and (2.3), we will focus our attention on other market participants who can trade only a finite number of times.

value of the replicating portfolio is equal to the price (production cost) of the derivative

$$V_0^{(N)} = H(0, P_0) \tag{2.5}$$

and its composition is given by

$$V_{S,0}^{(N)} = \left.\frac{\partial H(t,P_t)}{\partial P_t}\right|_{t=0} P_0, \quad V_{B,0}^{(N)} = V_0^{(N)} - V_{S,0}^{(N)}, \tag{2.6}$$

hence the portfolio contains $\partial H(t,P_t)/\partial P_t|_{t=0}$ shares of stock. The replicating portfolio is rebalanced at time periods $t_i$ so that

$$V_{S,t_i}^{(N)} = \left.\frac{\partial H(t,P_t)}{\partial P_t}\right|_{t=t_i} P_{t_i}, \quad V_{B,t_i}^{(N)} = V_{t_i}^{(N)} - V_{S,t_i}^{(N)}. \tag{2.7}$$

Between time periods $t_i$ and $t_{i+1}$, the portfolio composition remains unchanged. This gives rise to nonzero tracking errors $\varepsilon_{t_i}^{(N)}$:

$$\varepsilon_{t_i}^{(N)} \equiv H(t_i, P_{t_i}) - V_{t_i}^{(N)}. \tag{2.8}$$

The value of the replicating portfolio at time $t = 1$ is denoted by $V_1^{(N)}$ and the end-of-period tracking error is denoted by $\varepsilon_1^{(N)}$.

The sequence of tracking errors contains a great deal of information about the approximation errors of implementing a continuous-time hedging strategy in discrete time, and in Sections 2.2 and 5 we provide a complete characterization of the continuous-time limiting distribution of $\varepsilon_1^{(N)}$ and $\{\varepsilon_{t_i}\}$. However, because tracking errors also contain noise, we investigate the properties of the root-mean-squared error (RMSE) of the end-of-period tracking error $\varepsilon_1^{(N)}$:

$$\mathrm{RMSE}^{(N)} = \sqrt{\mathbf{E}_0[(\varepsilon_1^{(N)})^2]}, \tag{2.9}$$

where $\mathbf{E}_0[\cdot]$ denotes the conditional expectation, conditional on information available at time $t = 0$. Whenever exact replication of the derivative's payoff is impossible, $\mathrm{RMSE}^{(N)}$ is positive.

Of course, root-mean-squared error is only one of many possible summary statistics of the tracking error; Hutchinson et al. (1994) suggest other alternatives. A more general specification is the expected loss of the tracking error

$$\mathbf{E}_0[U(\varepsilon_1^{(N)})],$$

where $U(\cdot)$ is a general *loss function*, and we consider this case explicitly in Section 5.4.

## *2.2. Asymptotic behavior of the tracking error and RMSE*

We characterize analytically the asymptotic behavior of the tracking error and RMSE by appealing to continuous-record asymptotics, i.e., by letting the number of trading periods $N$ increase without bound while holding the time span fixed. This characterization provides several important insights into the behavior of the tracking error of general European derivative securities that previous studies have hinted at only indirectly and only for simple put and call options (e.g., Boyle and Emanuel, 1980; Hutchinson et al., 1994; Leland, 1985; Toft, 1996). A byproduct of this characterization is a useful definition for the temporal granularity of a continuous-time stochastic process (relative to a specific derivative security).

We begin with the case of smooth payoff functions $F(P_1)$.

**Theorem 1.** *Let the derivative's payoff function* $F(x)$ *in* (2.3) *be six times continuously differentiable and all of its derivatives be bounded, and suppose there exists a positive constant K such that functions* $\mu(\tau, x)$ *and* $\sigma(\tau, x)$ *in* (2.1) *satisfy*

$$\left|\frac{\partial^{\beta+\gamma}}{\partial\tau^{\beta}\partial x^{\gamma}}\mu(\tau,x)\right| + \left|\frac{\partial^{\beta+\gamma}}{\partial\tau^{\beta}\partial x^{\gamma}}\sigma(\tau,x)\right| + \left|\frac{\partial^{\alpha}}{\partial x^{\alpha}}(x\sigma(\tau,x))\right| \leqslant K, \tag{2.10}$$

*where* $(\tau, x) \in [0,\ 1] \times [0, \infty)$, $1 \leqslant \alpha \leqslant 6$, $0 \leqslant \beta \leqslant 1$, $0 \leqslant \gamma \leqslant 3$, *and all partial derivatives are continuous. Then under assumptions* (A1)–(A4),

(a) *the* RMSE *of the discrete-time delta-hedging strategy* (2.7) *satisfies*

$$\mathrm{RMSE}^{(N)} = \mathbf{O}\left(\frac{1}{\sqrt{N}}\right), \tag{2.11}$$

(b) *the normalized tracking error satisfies*

$$\sqrt{N}\varepsilon_1^{(N)} \Rightarrow G$$

*where*

$$G \equiv \frac{1}{\sqrt{2}}\int_0^1 \sigma^2(t, P_t)P_t^2 \frac{\partial^2 H(t, P_t)}{\partial P_t^2}\, \mathrm{d}W_t' \tag{2.12}$$

($W_t'$ *is a Wiener process independent of* $W_t$, *and* "$\Rightarrow$" *denotes convergence in distribution*), *and*

(c) *the* RMSE *of the discrete-time delta-hedging strategy* (2.7) *satisfies*

$$\mathrm{RMSE}^{(N)} = \frac{g}{\sqrt{N}} + \mathbf{O}\left(\frac{1}{N}\right), \tag{2.13}$$

*where*

$$g = \sqrt{\mathbf{E}_0[\mathscr{R}]}, \tag{2.14}$$

$$\mathscr{R} = \frac{1}{2}\int_0^1 \left(\sigma^2(t, P_t)P_t^2 \frac{\partial^2 H(t, P_t)}{\partial P_t^2}\right)^2 \mathrm{d}t. \tag{2.15}$$

**Proof.** See the appendix.

Theorem 1 shows that the tracking error is asymptotically equal in distribution to $G/\sqrt{N}$ (up to $\mathbf{O}(N^{-1})$ terms), where $G$ is a random variable given by (2.12). The expected value of $G$ is zero by the martingale property of stochastic integrals. Moreover, the independence of the Wiener processes $W'_t$ and $W_t$ implies that the asymptotic distribution of the normalized tracking error is symmetric, i.e., in the limit of frequent trading, positive values of the normalized tracking error are just as likely as negative values of the same magnitude.

This result might seem somewhat counterintuitive at first, especially in light of Boyle and Emanuel's (1980) finding that in the Black–Scholes framework the distribution of the local tracking error over a short trading interval is significantly skewed. However, Theorem 1(b) describes the asymptotic distribution of the tracking error over the *entire life* of the derivative, not over short intervals. Such an aggregation of local errors leads to a symmetric asymptotic distribution, just as a normalized sum of random variables will have a Gaussian distribution asymptotically under certain conditions, e.g., the conditions for a functional central limit theorem to hold.

Note that Theorem 1 applies to a wide class of diffusion processes (2.1) and to a variety of derivative payoff functions $F(P_1)$. In particular, it holds when the stock price follows a diffusion process with constant coefficients, as in Black and Scholes (1973).[2] However, the requirement that the payoff function $F(P_1)$ is smooth – six times differentiable with bounded derivatives – is violated by the most common derivatives of all, simple puts and calls. In the next theorem, we extend our results to cover this most basic set of payoff functions.

**Theorem 2.** *Let the payoff function* $F(P_1)$ *be continuous and piecewise linear, and suppose* (2.10) *holds. In addition, let*

$$\left| x^2 \frac{\partial^\alpha \sigma(\tau, x)}{\partial x^\alpha} \right| \leqslant K_2 \tag{2.16}$$

[2] For the Black–Scholes case, the formula for the RMSE (2.14)–(2.15) was first derived by Grannan and Swindle (1996). Our results provide a more complete characterization of the tracking error in their framework – we derive the asymptotic distribution – and our analysis applies to more general trading strategies than theirs, e.g., they consider strategies obtained by deterministic time deformations; our framework can accommodate deterministic and stochastic time deformations.

*for* $(\tau, x) \in [0, 1] \times [0, \infty)$, $2 \leqslant \alpha \leqslant 6$, *and some positive constant* $K_2$. *Then under assumptions* (A1)–(A4),

(a) *the* RMSE *of the discrete-time delta-hedging strategy* (2.7) *satisfies*

$$\mathrm{RMSE}^{(N)} = \frac{g}{\sqrt{N}} + \mathbf{o}(1/\sqrt{N}),$$

*where* $g$ *is given by* (2.14)–(2.15), *and*

(b) *the normalized tracking error satisfies*

$$\sqrt{N}\,\varepsilon_1^{(N)} \Rightarrow \frac{1}{\sqrt{2}} \int_0^1 \sigma^2(t, P_t) P_t^2 \frac{\partial^2 H(t, P_t)}{\partial P_t^2}\, \mathrm{d}W'_t, \tag{2.17}$$

*where* $W'_t$ *is a Wiener process independent of* $W_t$.

**Proof.** Available from the authors upon request.

By imposing an additional smoothness condition (2.16) on the diffusion coefficient $\sigma(\tau, x)$, Theorem 2 assures us that the conclusions of Theorem 1 also hold for the most common types of derivatives, those with piecewise linear payoff functions. Theorems 1 and 2 allow us to define the coefficient of *temporal granularity* $g$ for any combination of continuous-time process $\{P_t\}$ and derivative payoff function $F(P_1)$ as the constant associated with the leading term of the RMSE's continuous-record asymptotic expansion:

$$g \equiv \sqrt{\frac{1}{2}\mathbf{E}_0\left[\int_0^1 \left(\sigma^2(t, P_t) P_t^2 \frac{\partial^2 H(t, P_t)}{\partial P_t^2}\right)^2 \mathrm{d}t\right]} \tag{2.18}$$

where $H(t, P_t)$ satisfies (2.2) and (2.3).

### 2.3. Interpretation of granularity

The interpretation for temporal granularity is clear: it is a measure of the approximation errors that arise from implementing a continuous-time delta-hedging strategy in discrete time. A derivative-pricing model – recall that this consists of a payoff function $F(P_1)$ and a continuous-time stochastic process for $P_t$ – with high granularity requires a larger number of trading periods to achieve the same level of tracking error as a derivative-pricing model with low granularity. In the former case, time is 'grainier', calling for more frequent hedging activity than the latter case. More formally, according to Theorems 1 and 2, to a first-order approximation the RMSE of an $N$-trade delta-hedging strategy is

$g/\sqrt{N}$. Therefore, if we desire the RMSE to be within some small value $\delta$, we require

$$N \approx \frac{g^2}{\delta^2}$$

trades in the unit interval. For a fixed error $\delta$, the number of trades needed to reduce the RMSE to within $\delta$ grows quadratically with granularity. If one derivative-pricing model has twice the granularity of another, it would require four times as many delta-hedging transactions to achieve the same RMSE tracking error.

From (2.18) is it clear that granularity depends on the derivative-pricing formula $H(t, P_t)$ and the price dynamics $P_t$ in natural ways. Eq. (2.18) formalizes the intuition that derivatives with higher volatility and higher 'gamma' risk (large second derivative with respect to stock price) are more difficult to hedge, since these cases imply larger values for the integrand in (2.18). Accordingly, derivatives on less volatile stocks are easier to hedge. Consider a stock price process which is almost deterministic, i.e., $\sigma(t, P_t)$ is very small. This implies a very small value for $g$, hence derivatives on such a stock can be replicated almost perfectly, even if continuous trading is not feasible. Alternatively, such derivatives require relatively few rebalancing periods $N$ to maintain small tracking errors.

Also, a derivative with a particularly simple payoff function should be easier to hedge than derivatives on the same stock with more complicated payoffs. For example, consider a derivative with the payoff function $F(P_1) = P_1$. This derivative is identical to the underlying stock, and can always be replicated perfectly by buying a unit of the underlying stock at time $t = 0$ and holding it until expiration. The tracking error for this derivative is always equal to zero, no matter how volatile the underlying stock is. This intuition is made precise by Theorem 1, which describes exactly how the error depends on the properties of the stock price process and the payoff function of the derivative: it is determined by the behavior of the integral $\mathscr{R}$, which tends to be large when stock prices 'spend more time' in regions of the domain that imply high volatility and high convexity or gamma of the derivative.

We will investigate the sensitivity of $g$ to the specification of the stock price process in Sections 3 and 4.

## 3. Applications

To develop further intuition for our measure of temporal granularity, in this section we derive closed-form expressions for $g$ in two important special cases: the Black–Scholes option pricing model with geometric Brownian motion, and the Black–Scholes model with a mean-reverting (Ornstein–Uhlenbeck) process.

### 3.1. Granularity of geometric Brownian motion

Suppose that stock price dynamics are given by

$$\frac{\mathrm{d}P_t}{P_t} = \mu\,\mathrm{d}t + \sigma\,\mathrm{d}W_t, \tag{3.1}$$

where $\mu$ and $\sigma$ are constants. Under this assumption we obtain an explicit characterization of the granularity $g$.

**Theorem 3.** *Under Assumptions* (A1)–(A4), *stock price dynamics* (3.1), *and the payoff function of simple call and put options*, *the granularity* $g$ *in* (2.13) *is given by*

$$g = K\sigma\left(\int_0^1 \exp\left[-\frac{[\mu t + \ln(P_0/K) - \sigma^2/2]^2}{\sigma^2(1+t)}\right]/(4\pi\sqrt{1-t^2})\,\mathrm{d}t\right)^{1/2}, \tag{3.2}$$

*where K is the option's strike price.*

**Proof.** Available from the authors upon request.

It is easy to see that $g = 0$ if $\sigma = 0$ and $g$ increases with $\sigma$ in the neighborhood of zero. When $\sigma$ increases without bound, the granularity $g$ decays to zero, which means that it has at least one local maximum as a function of $\sigma$. The granularity $g$ also decays to zero when $P_0/K$ approaches zero or infinity. In the important special case of $\mu = 0$, we conclude by direct computation that $g$ is a unimodal function of $P_0/K$ that achieves its maximum at $P_0/K = \exp(\sigma^2/2)$.

The fact that granularity is not monotone increasing in $\sigma$ may seem counterintuitive at first – after all, how can delta-hedging errors become smaller for larger values of $\sigma$? The intuition follows from the fact that at small levels of $\sigma$, an increase in $\sigma$ leads to larger granularity because there is a greater chance that the stock price will fluctuate around regions of high gamma, i.e., near the money where $\partial^2 H(t, P_t)/\partial P_t^2$ is large, leading to greater tracking errors. However, at very high levels of $\sigma$, prices fluctuate so wildly that an increase in $\sigma$ will decrease the probability that the stock price stays in regions of high gamma for very long. In these extreme cases, the payoff function 'looks' approximately linear, hence granularity becomes a decreasing function of $\sigma$.

Also, we show below that $g$ is not very sensitive to changes in $\mu$ when $\sigma$ is sufficiently large. This implies that, for an empirically relevant range of parameter values, $g$, as a function of the initial stock price, achieves its maximum close to the strike price, i.e., at $P_0/K \approx 1$. These observations are consistent with the behavior of the tracking error for finite values of $N$ that we see in the Monte Carlo simulations of Section 4.

When stock prices follow a geometric Brownian motion, expressions similar to (3.2) can be obtained for derivatives other than simple puts and calls. For example, for a 'straddle', consisting of one put and one call option with the same strike price $K$, the constant $g$ is twice as large as for the put or call option alone.

### 3.2. Granularity of a mean-reverting process

Let $p_t \equiv \ln(P_t)$ and suppose

$$\mathrm{d}p_t = (-\gamma(p_t - (\alpha + \beta t)) + \beta)\,\mathrm{d}t + \sigma\,\mathrm{d}W_t, \tag{3.3}$$

where $\beta = \mu - \sigma^2/2$ and $\alpha$ is a constant. This is an Ornstein–Uhlenbeck process with a linear time trend, and the solution of (3.3) is given by

$$p_t = (p_0 - \alpha)\mathrm{e}^{-\gamma t} + (\alpha + \beta t) + \sigma \int_0^t \mathrm{e}^{-\gamma(t-s)}\,\mathrm{d}W_s. \tag{3.4}$$

**Theorem 4.** *Under assumptions* (A1)–(A4), *stock price dynamics* (3.3), *and the payoff function of simple call and put options, the granularity g in* (2.13) *is given by*

$$g = K\sigma\left(\int_0^1 \frac{\sqrt{\gamma}\,\exp\left[-\frac{\gamma[\alpha+\mu t+(\ln(P_0/K)-\alpha)\exp(-\gamma t)-\sigma^2/2]^2}{\sigma^2[\gamma(1-t)+1-\exp(-2\gamma t)]}\right]}{4\pi\sqrt{1-t}\sqrt{\gamma(1-t)+1-\exp(-2\gamma t)}}\,\mathrm{d}t\right)^{1/2}, \tag{3.5}$$

*where K is the option's strike price.*

**Proof.** Available from the authors upon request.

Expression (3.5) is a direct generalization of (3.2): when the mean-reversion parameter $\gamma$ is set to zero, the process (3.3) becomes a geometric Brownian motion and (3.5) reduces to (3.2). Theorem 4 has some interesting qualitative implications for the behavior of the tracking error in presence of mean-reversion. We will discuss them in detail in the next section.

## 4. Monte Carlo analysis

Since our analysis of granularity is based entirely on continuous-record asymptotics, we must check the quality of these approximations by performing Monte Carlo simulation experiments for various values of $N$. The results of these Monte Carlo simulations are reported in Section 4.1. We also use Monte Carlo simulations to explore the qualitative behavior of the RMSE for various parameter values of the stock price process, and these simulations are reported in Section 4.2.

### 4.1. Accuracy of the asymptotics

We begin by investigating the distribution of the tracking error $\varepsilon^{(N)}$ for various values of $N$. We do this by simulating the hedging strategy of Section 2.2 for call and put options assuming that price dynamics are given by a geometric Brownian motion (3.1). According to Theorem 1, the asymptotic expressions for the tracking error and the RMSE are the same for put and call options since these options have the same second partial derivative of the option price with respect to the current stock price. Moreover, it is easy to verify, using the put-call parity relation, that these options give rise to identical tracking errors. When the stock price process $P_t$ follows a geometric Brownian motion, the stock price at time $t_{i+1}$ is distributed (conditional on the stock price at time $t_i$) as $P_{t_i}\exp((\mu - \sigma^2/2)\,\Delta t + \sigma\sqrt{\Delta t}\,\eta)$, where $\eta \sim \mathcal{N}(0, 1)$. We use this relation to simulate the delta-hedging strategy. We set the parameters of the stock price process to $\mu = 0.1$, $\sigma = 0.3$, and $P_0 = 1.0$, and let the strike price be $K = 1$. We consider $N = 10$, 20, 50, and 100, and simulate the hedging process 250,000 times for each value of $N$.

Fig. 1a shows the empirical probability density function (PDF) of $\varepsilon_1^{(N)}$ for each $N$. As expected, the distribution of the tracking error becomes tighter as the trading frequency increases. It is also apparent that the tracking error can be significant even for $N = 100$. Fig. 1b contains the empirical PDFs of the normalized tracking error, $\sqrt{N}\varepsilon_1^{(N)}$, for the same values of $N$. These PDFs are

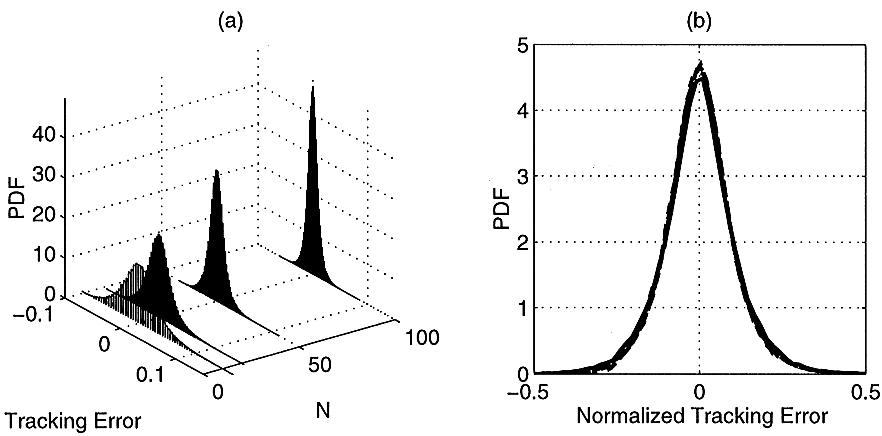


Fig. 1. Empirical probability density functions of (a) the tracking error and (b) the normalized tracking error (dashed line) are plotted for different values of the trading frequency $N$. (b) Also, shows the empirical probability density function of the asymptotic distribution (2.17) (solid line). The stock price process is given by (3.1) with parameters $\mu = 0.1$, $\sigma = 0.3$, and $P_0 = 1.0$. The option is a European call (put) option with strike price $K = 1$.

compared to the PDF of the asymptotic distribution (2.17), which is estimated by approximating the integral in (2.17) using a first-order Euler scheme. The functions in Fig. 1b are practically identical and indistinguishable, which suggests that the asymptotic expression for the distribution of $\sqrt{N}\varepsilon_1^{(N)}$ in Theorem 1(b) is an excellent approximation to the finite-sample PDF for values of $N$ as small as ten.

To evaluate the accuracy of the asymptotic expression $g/\sqrt{N}$ for finite values of $N$, we compare $g/\sqrt{N}$ to the actual RMSE from Monte Carlo simulations of the delta-hedging strategy of Section 2.2. Specifically, we simulate the delta-

Table 1
The sensitivity of the RMSE as a function of the initial price $P_0$. The RMSE is estimated using Monte Carlo simulation. Options are European calls and puts with strike price $K = 1$. 250,000 simulations are performed for every set of parameter values. The stock price follows a geometric Brownian motion (3.1). The drift and diffusion coefficients of the stock price process are $\mu = 0.1$ and $\sigma = 0.3$. $\mathrm{RMSE}^{(N)}$ is compared to the asymptotic approximation $gN^{-1/2}$ in (2.13)–(3.2). The relative error (RE) of the asymptotic approximation is defined as $|gN^{-1/2} - \mathrm{RMSE}^{(N)}|/\mathrm{RMSE}^{(N)} \times 100\%$.

| Parameters | | $gN^{-1/2}$ | $\mathrm{RMSE}^{(N)}$ | R.E. (%) | Call option | | Put option | |
|---|---|---|---|---|---|---|---|---|
| $N$ | $P_0$ | | | | $H(0,P_0)$ | $\frac{\mathrm{RMSE}^{(N)}}{H}$ | $H(0,P_0)$ | $\frac{\mathrm{RMSE}^{(N)}}{H}$ |
| 10 | 0.50 | 0.0078 | 0.0071 | 8.9 | 7E-4 | 9.64 | 0.501 | 0.014 |
| 20 | 0.50 | 0.0055 | 0.0052 | 7.9 | 7E-4 | 6.88 | 0.501 | 0.010 |
| 50 | 0.50 | 0.0035 | 0.0033 | 5.9 | 7E-4 | 4.43 | 0.501 | 0.007 |
| 100 | 0.50 | 0.0025 | 0.0024 | 3.1 | 7E-4 | 3.22 | 0.501 | 0.005 |
| 10 | 0.75 | 0.0259 | 0.0248 | 3.8 | 0.023 | 1.08 | 0.273 | 0.091 |
| 20 | 0.75 | 0.0183 | 0.0177 | 2.9 | 0.023 | 0.760 | 0.273 | 0.065 |
| 50 | 0.75 | 0.0116 | 0.0113 | 2.6 | 0.023 | 0.490 | 0.273 | 0.041 |
| 100 | 0.75 | 0.0082 | 0.0082 | 2.3 | 0.023 | 0.345 | 0.273 | 0.029 |
| 10 | 1.00 | 0.0334 | 0.0327 | 4.1 | 0.119 | 0.269 | 0.119 | 0.269 |
| 20 | 1.00 | 0.0236 | 0.0227 | 3.4 | 0.119 | 0.192 | 0.119 | 0.192 |
| 50 | 1.00 | 0.0149 | 0.0145 | 2.3 | 0.119 | 0.122 | 0.119 | 0.122 |
| 100 | 1.00 | 0.0106 | 0.0104 | 1.9 | 0.119 | 0.087 | 0.119 | 0.087 |
| 10 | 1.25 | 0.0275 | 0.0263 | 5.8 | 0.294 | 0.088 | 0.044 | 0.588 |
| 20 | 1.25 | 0.0194 | 0.0187 | 3.9 | 0.294 | 0.064 | 0.044 | 0.423 |
| 50 | 1.25 | 0.0123 | 0.0120 | 2.7 | 0.294 | 0.041 | 0.044 | 0.271 |
| 100 | 1.25 | 0.0087 | 0.0087 | 1.8 | 0.294 | 0.029 | 0.044 | 0.195 |
| 10 | 1.50 | 0.0181 | 0.0169 | 7.7 | 0.515 | 0.033 | 0.015 | 1.130 |
| 20 | 1.50 | 0.0128 | 0.0122 | 5.3 | 0.515 | 0.024 | 0.015 | 0.816 |
| 50 | 1.50 | 0.0081 | 0.0076 | 2.9 | 0.515 | 0.015 | 0.015 | 0.528 |
| 100 | 1.50 | 0.0057 | 0.0056 | 3.0 | 0.515 | 0.011 | 0.015 | 0.373 |

hedging strategy for a set of European put and call options with strike price $K = 1$ under geometric Brownian motion (3.1) with different sets of parameter values for ($\sigma$, $\mu$, and $P_0$). The tracking error is tabulated as a function of these parameters and the results are summarized in Tables 1–3.

Tables 1–3 show that $g/\sqrt{N}$ is an excellent approximation to the RMSE across a wide range of parameter values for ($\mu$, $\sigma, P_0$), even for as few as $N = 10$ delta-hedging periods.

### 4.2. Qualitative behavior of the RMSE

The Monte Carlo simulations of Section 4.1 show that the RMSE increases with the diffusion coefficient $\sigma$ in an empirically relevant range of parameter values (see Table 2), and that the RMSE is not very sensitive to the drift rate $\mu$ of the stock price process when $\sigma$ is sufficiently large (see Table 3). These properties are illustrated in Figs. 2a and 3. Fig. 2a plots the logarithm of the RMSE against the logarithm of trading periods $N$ for $\sigma = 0.1$, 0.2, and 0.3 – as $\sigma$ increases, the

Table 2
The sensitivity of the RMSE as a function of volatility $\sigma$. The RMSE is estimated using Monte Carlo simulation. Options are European calls and puts with strike price $K = 1$. 250,000 simulations are performed for every set of parameter values. The stock price follows a geometric Brownian motion (3.1). The drift coefficient of the stock price process is $\mu = 0.1$, and the initial stock price is $P_0 = 1.0$. $\text{RMSE}^{(N)}$ is compared to the asymptotic approximation $gN^{-1/2}$ in (2.13)–(3.2). The relative error (RE) of the asymptotic approximation is defined as $|gN^{-1/2} - \text{RMSE}^{(N)}|/\text{RMSE}^{(N)} \times 100\%$.

| Parameters | | $gN^{-1/2}$ | Call and put options | | | |
|---|---|---|---|---|---|---|
| $N$ | $\sigma$ | | $\text{RMSE}^{(N)}$ | R.E. (%) | $H(0, P_0)$ | $\frac{\text{RMSE}^{(N)}}{H}$ |
| 10 | 0.3 | 0.0334 | 0.0327 | 4.1 | 0.119 | 0.269 |
| 20 | 0.3 | 0.0236 | 0.0227 | 3.4 | 0.119 | 0.192 |
| 50 | 0.3 | 0.0149 | 0.0145 | 2.3 | 0.119 | 0.122 |
| 100 | 0.3 | 0.0106 | 0.0104 | 1.9 | 0.119 | 0.087 |
| 10 | 0.2 | 0.0219 | 0.0212 | 3.4 | 0.080 | 0.266 |
| 20 | 0.2 | 0.0155 | 0.0151 | 3.0 | 0.080 | 0.189 |
| 50 | 0.2 | 0.0098 | 0.0096 | 2.1 | 0.080 | 0.121 |
| 100 | 0.2 | 0.0069 | 0.0068 | 1.7 | 0.080 | 0.086 |
| 10 | 0.1 | 0.0100 | 0.0102 | 1.6 | 0.040 | 0.255 |
| 20 | 0.1 | 0.0071 | 0.0071 | 0.04 | 0.040 | 0.177 |
| 50 | 0.1 | 0.0045 | 0.0044 | 1.1 | 0.040 | 0.111 |
| 100 | 0.1 | 0.0032 | 0.0031 | 0.9 | 0.040 | 0.078 |

Table 3
The sensitivity of the RMSE as a function of the drift $\mu$. The RMSE is estimated using Monte Carlo simulation. Options are European calls and puts with strike price $K = 1$. 250,000 simulations are performed for every set of parameter values. The stock price follows a geometric Brownian motion (3.1). The diffusion coefficient of the stock price process is $\sigma = 0.3$, the initial stock price is $P_0 = 1.0$, and the number of trading periods is $N = 20$. $\mathrm{RMSE}^{(N)}$ is compared to the asymptotic approximation $gN^{-1/2}$ in (2.13)–(3.2). The relative error (RE) of the asymptotic approximation is defined as $|gN^{-1/2} - \mathrm{RMSE}^{(N)}|/\mathrm{RMSE}^{(N)} \times 100\%$.

| Parameters | $gN^{-1/2}$ | Call and put options | | | |
|---|---|---|---|---|---|
| $\mu$ | | $\mathrm{RMSE}^{(N)}$ | R.E. (%) | $H(0, P_0)$ | $\frac{\mathrm{RMSE}^{(N)}}{H}$ |
| 0.0 | 0.0235 | 0.0226 | 4.3 | 0.119 | 0.189 |
| 0.1 | 0.0236 | 0.0229 | 3.4 | 0.119 | 0.192 |
| 0.2 | 0.0230 | 0.0226 | 1.7 | 0.119 | 0.190 |
| 0.3 | 0.0218 | 0.0220 | 1.0 | 0.119 | 0.184 |

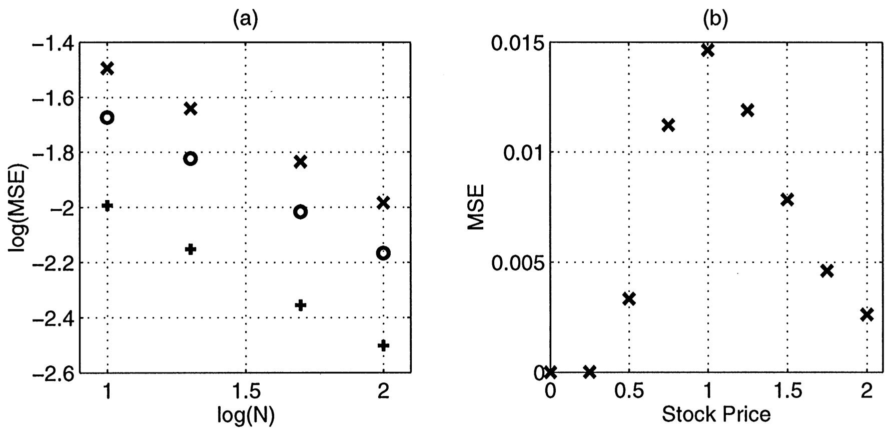


Fig. 2. (a) The logarithm of the root-mean-squared error $\log_{10}(\mathrm{RMSE}^{(N)})$ is plotted as a function of the logarithm of the trading frequency $\log_{10}(N)$. The option is a European call (put) option with the strike price $K = 1$. The stock price process is given by (3.1) with parameters $\mu = 0.1$ and $P_0 = 1.0$. The diffusion coefficient of the stock price process takes values $\sigma = 0.3$ (x's), $\sigma = 0.2$ (o's) and $\sigma = 0.1$ ( + 's). (b) The root-mean-squared error RMSE is plotted as a function of the initial stock price $P_0$. The option is a European put option with the strike price $K = 1$. The parameters of the stock price process are $\mu = 0.1$ and $\sigma = 0.3$.

locus of points shifts upward. Fig. 3 shows that granularity $g$ is not a monotone function of $\sigma$ and goes to zero as $\sigma$ increases without bound.

Fig. 2b plots the RMSE as a function of the initial stock price $P_0$. RMSE is a unimodal function of $P_0/K$ (recall that the strike price has been normalized to

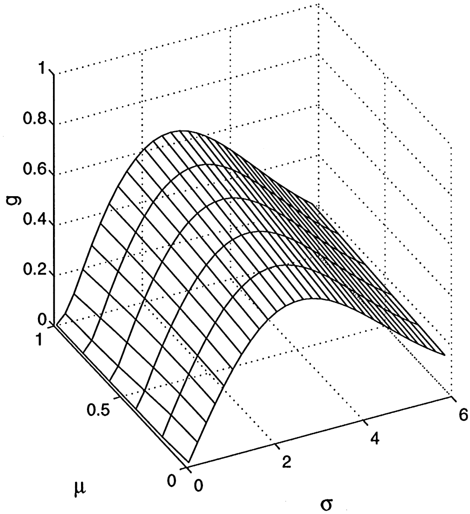


Fig. 3. The granularity $g$ is plotted as a function of $\sigma$ and $\mu$. The option is a European call (put) option with strike price $K = 1$. The stock price process is geometric Brownian motion and initial stock price $P_0 = 1$.

$K = 1$ in all our calculations), achieving its maximum around one and decaying to zero as $P_0/K$ approaches zero or infinity (see Table 1). This confirms the common intuition that close-to-the-money options are the most difficult to hedge (they exhibit the largest RMSE).

Finally, the relative importance of the RMSE can be measured by the ratio of the RMSE to the option price: $\text{RMSE}^{(N)}/H(0, P_0)$. This quantity is the root-mean-squared error per dollar invested in the option. Table 1 shows that this ratio is highest for out-of-the-money options, despite the fact that the RMSE is highest for close-to-the-money options. This is due to the fact that the option price decreases faster than the RMSE as the stock moves away from the strike.

Now consider the case of mean-reverting stock price dynamics (3.3). Recall that under these dynamics, the Black–Scholes formula still holds, although the numerical value for $\sigma$ can be different than that of a geometric Brownian motion because the presence of mean-reversion can affect conditional volatility, holding unconditional volatility fixed; see Lo and Wang (1995) for further discussion. Nevertheless, the behavior of granularity and RMSE is quite different in this case. Fig. 4 plots the granularity $g$ of call and put options for the Ornstein–Uhlenbeck process (3.3) as a function of $\alpha$ and $P_0$. Fig. 4a assumes a value of 0.1 for the mean-reversion parameter $\gamma$ and Fig. 4b assumes a value of 3.0. It is clear from these two plots that the degree of mean reversion $\gamma$ has an enormous impact on granularity. When $\gamma$ is small, Fig. 4a shows that the RMSE is highest

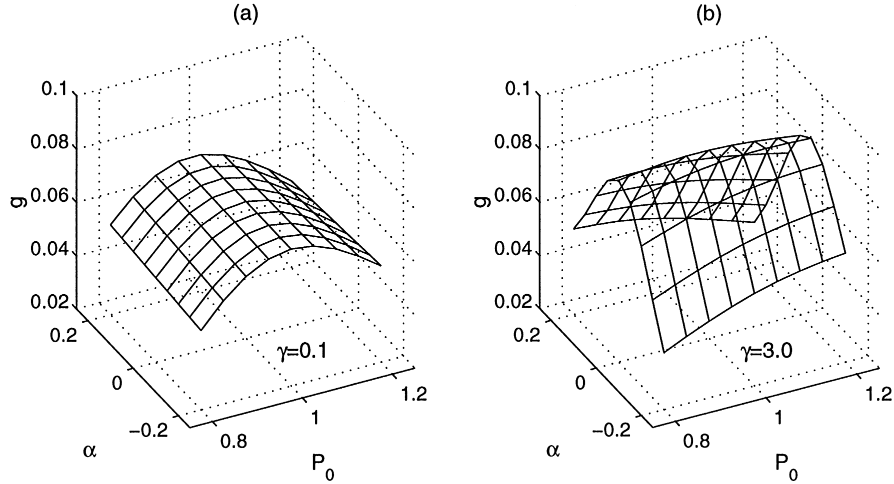


Fig. 4. Granularity $g$ is plotted as a function of $P_0$ and $\alpha$. The option is a European call (put) option with the strike price $K = 1$. The parameters of the stock price process are $\sigma = 0.2$ and $\mu = 0.05$. The stock price process is given by (3.4). Mean-reversion parameter $\gamma$ takes two values: (a) $\gamma = 0.1$ and (b) $\gamma = 3.0$.

when $P_0$ is close to the strike price and is not sensitive to $\alpha$. But when $\gamma$ is large, Fig. 4b suggests that the RMSE is highest when $\exp(\alpha)$ is close to the strike price and is not sensitive to $P_0$.

The influence of $\gamma$ on granularity can be understood by recalling that granularity is closely related to the option's gamma (see Section 2.3). When $\gamma$ is small, the stock price is more likely to spend time in the neighborhood of the strike price – the region with the highest gamma or $\partial^2 H(t, P_t)/\partial P_t^2$ – when $P_0$ is close to $K$. However, when $\gamma$ is large, the stock price is more likely to spend time in a neighborhood of $\exp(\alpha)$, thus $g$ is highest when $\exp(\alpha)$ is close to $K$.

## 5. Extensions and generalizations

The analysis of Section 2 can be extended in a number of directions, and we briefly outline four of the most important of these extensions here. In Section 5.1, we show that the normalized tracking error converges in a much stronger sense than simply in distribution, and that this stronger 'sample-path' notion of convergence – called, ironically, 'weak' convergence – can be used to analyze the tracking error of American-style derivative securities. In Section 5.2 we characterize the asymptotic joint distributions of the normalized tracking error and asset prices, a particularly important extension for investigating the tracking error of delta hedging a portfolio of derivatives. In Section 5.3, we provide another characterization of the tracking error, one that relies on PDEs, that offers important computational advantages. And in Section 5.4, we consider alternatives to mean-squared error loss functions and show that for quite general loss functions, the behavior of the expected loss of the tracking error is characterized by the same stochastic integral (2.17) as in the mean-squared-error case.

## 5.1. Sample-path properties of tracking errors

Recall that the normalized tracking error *process* is defined as

$$\sqrt{N}\varepsilon_t^{(N)} = \sqrt{N}(H(t,P_t) - V_t^{(N)}), \quad t \in [0, 1].$$

It can be shown that $\sqrt{N}\varepsilon_t^{(N)}$ converges weakly to the stochastic process $G_t$, characterized by the stochastic integral in (2.12) as a function of its upper limit:

$$G_t = \frac{1}{\sqrt{2}} \int_0^t \sigma^2(t,P_s) P_s^2 \frac{\partial^2 H(s,P_s)}{\partial P_s^2} \, \mathrm{d}W'_s.$$

The proof of this result consists of two steps. The first step is to establish that the sequence of measures induced by $\sqrt{N}\varepsilon_t^{(N)}$ is tight (relatively compact). This can be done by verifying local inequalities for the moments of processes $\sqrt{N}\varepsilon_t^{(N)}$ using the machinery developed in the proof of Theorem 2 (we must use Burkholder's inequality instead of the isometric property and Hölder's inequality instead of Schwarz's inequality throughout – details are available from the authors upon request). The second step is to characterize the limiting process. Such a characterization follows from the proof of Theorem 1(b) in the appendix and the fact that the results in Duffie and Protter (1992) guarantee weak convergence of stochastic processes, not just convergence of their one-dimensional marginal distributions.

This stronger notion of convergence yields stronger versions of Theorems 1 and 2 that can be used to analyze a number of sample-path properties of the tracking error by appealing to the Continuous Mapping Theorem (Billingsley, 1986). This well-known result shows that the asymptotic distribution of any continuous functional $\xi(\cdot)$ of the normalized tracking error is given by $\xi(G_t)$. For example, the maximum of the normalized tracking error over the entire life of the derivative security, $\max_t \sqrt{N}\varepsilon_t^{(N)}$, is distributed as $\max_t G_t$ asymptotically.

These results can be applied to the normalized tracking errors of American-style derivatives in a straightforward manner. Such derivatives differ from European derivatives in one respect: they can be exercised prematurely. Therefore, the valuation of these derivatives consists of computing both the derivative price function $H(t, P_t)$ and the optimal exercise schedule, which can be represented as a stopping time $\tau$. Then the tracking error at the moment when the derivative is exercised behaves asymptotically as $G_\tau/\sqrt{N}$; (some technical regularity conditions, e.g., the smoothness of the exercise boundary, are required to ensure convergence; see, e.g., Kushner and Dupuis, 1992).

The tracking error, conditional on the derivative not being exercised prematurely, is distributed asymptotically as $(G_1/\sqrt{N}|\tau = 1)$.

### 5.2. Joint distributions of tracking errors and prices

Theorems 1 and 2 provide a complete characterization of the tracking error and RMSE for individual derivatives, but what is often of more practical interest is the behavior of a portfolio of derivatives. Delta hedging a portfolio of derivatives is typically easier because of the effects of diversification. As long as tracking errors are not perfectly correlated across derivatives, the portfolio tracking error will be less volatile than the tracking error of individual derivatives.

To address portfolio issues, we require the joint distribution of tracking errors for multiple stocks, as well as the joint distribution of tracking errors and prices. Consider another stock with price $P_t^{(2)}$ governed by the diffusion equation

$$\frac{\mathrm{d}P_t^{(2)}}{P_t^{(2)}} = \mu^{(2)}(t, P_t^{(2)})\,\mathrm{d}t + \sigma^{(2)}(t, P_t^{(2)})\,\mathrm{d}W_t^{(2)} \tag{5.1}$$

where $W_t^{(2)}$ can be correlated with $W_t$. According to the proof of Theorem 1(b) in the Appendix, the random variables $(W_{t_{i+1}} - W_{t_i})^2 - (t_{i+1} - t_i)$ and $W_{t_{i+1}}^{(2)} - W_{t_i}^{(2)}$ are uncorrelated. This follows from the fact that, for every pair of standard normal random variables $X$ and $Y$ with correlation $\rho$, $X = \rho Y + \sqrt{1 + \rho^2} Z$, where $Z$ is a standard normal random variable, independent of $Y$. Thus $X$ and $Y^2 - 1$ are uncorrelated. It follows that the Wiener processes $W_t'$ and $W_t^{(2)}$ are independent. Therefore, as $N$ increases without bound, the pair of random variables $(\sqrt{N}\varepsilon_1^{(N)}, P_1^{(2)})$ converges in distribution to

$$(\sqrt{N}\varepsilon_1^{(N)}, P_1^{(2)}) \Rightarrow \left(\frac{1}{\sqrt{2}}\int_0^1 \sigma^2(t, P_t) P_t^2 \frac{\partial^2 H(t, P_t)}{\partial P_t^2}\,\mathrm{d}W_t', P_1^{(2)}\right) \tag{5.2}$$

where $W_t'$ is independent of $W_t$ and $W_t^{(2)}$.

An immediate corollary of this result is that the normalized tracking error is uncorrelated with any asset in the economy. This follows easily from (5.2) since, conditional on the realization of $P_t$ and $P_t^{(2)}$, $t \in [0, 1]$, the normalized tracking error has zero expected value asymptotically. However, this does not imply that the asymptotic joint distribution of $(\sqrt{N}\varepsilon_1^{(N)}, P_1^{(2)})$ does not depend on the correlation between $W_t$ and $W_t^{(2)}$ – it does, since this correlation determines the joint distribution of $P_t$ and $P_1^{(2)}$.

The above argument applies without change when the price of the second stock follows a diffusion process different from (5.1), and can also easily be extended to the case of multiple stocks.

To derive the joint distribution of the normalized tracking errors for multiple stocks, we consider the case of two stocks since the generalization to multiple stocks is obvious. Let $W_t$ and $W_t^{(2)}$ have mutual variation $\mathrm{d}W_t\,\mathrm{d}W_t^{(2)} = \rho(t, P_t, P_t^{(2)})\,\mathrm{d}t$, where $\rho(\cdot)$ is a continuously differentiable function with bounded first-order partial derivatives. We have already established that the asymptotic distribution of the tracking error is characterized by the stochastic integral (2.12). To describe the asymptotic joint distribution of two normalized tracking errors, it is sufficient to find the mutual variation of the Wiener processes in the corresponding stochastic integrals. According to the proof of Theorem 1(b) in the appendix, this amounts to computing the expected value of the product

$$((W_{t_{i+1}} - W_{t_i})^2 - (t_{i+1} - t_i))\,((W_{t_{i+1}}^{(2)} - W_{t_i}^{(2)})^2 - (t_{i+1} - t_i)).$$

Using Itô's formula, it is easy to show that the expected value of the above expression is equal to

$$\mathbf{E}_0[2\rho^2(t, P_{t_i}, P_{t_i}^{(2)})](\Delta t)^2 + \mathbf{O}((\Delta t)^{5/2}).$$

This implies that $\rho^2(t, P_t, P_t^{(2)})$ is the mutual variation of the two Wiener processes in the stochastic integrals (2.12) that describe the asymptotic distributions of the normalized tracking errors of the two stocks. Together with Theorem 1(b), this completely determines the asymptotic joint distribution of the two normalized tracking errors, and is a generalization of the results of Boyle and Emanuel (1980).

Note that the correlation of two Wiener processes describing the asymptotic behavior of two normalized tracking errors is always nonnegative, regardless of the sign of the mutual variation of the original Wiener processes $W_t$ and $W_t^{(2)}$. In particular, when two derivatives have convex price functions, this means that even if the returns on the two stocks are negatively correlated, the tracking errors resulting from delta hedging derivatives on these stocks are asymptotically positively correlated.

### *5.3. A PDE characterization of the tracking error*

It is possible to derive an alternative characterization of the tracking error using the intimate relation between diffusion processes and PDEs. Although this may seem superfluous given the analytical results of Theorems 1 and 2, the numerical implementation of a PDE representation is often computationally more efficient.

To illustrate our approach, we begin with the RMSE. According to Theorem 1(c), the RMSE can be completely characterized asymptotically if $g$ is known.

Using the Feynman-Kac representation of the solutions of PDEs (Karatzas and Shreve, 1991, Proposition 4.2.), we conclude that $g^2 = u(0, P_0)$, where $u(t,x)$ solves the following:

$$\left[\frac{\partial}{\partial t} + \mu(t,x)x\frac{\partial}{\partial x} + \frac{1}{2}\sigma^2(t,x)x^2\frac{\partial^2}{\partial x^2}\right]u(t,x) + \frac{1}{2}\left(\sigma^2(t,x)x^2\frac{\partial^2 H(t,x)}{\partial x^2}\right)^2 = 0 \tag{5.3}$$

$$u(1,x) = 0, \quad \forall x. \tag{5.4}$$

The PDE (5.3)–(5.4) is of the same degree of difficulty as the fundamental PDE (2.2)–(2.3) that must be solved to obtain the derivative-pricing function $H(t, P_t)$. This new representation of the RMSE can be used to implement an efficient numerical procedure for calculating RMSE without resorting to Monte Carlo simulation. Results from some preliminary numerical experiments provide encouraging evidence of the practical value of this new representation.

Summary measures of the tracking error with general loss functions can also be computed numerically along the same lines, using the Kolmogorov backward equation. The probability density function of the normalized tracking error $\sqrt{N}\varepsilon_1^{(N)}$ can be determined numerically as a solution of the Kolmogorov forward equation (see, e.g., Karatzas and Shreve, 1991, pp. 368–369).

### *5.4. Alternative measures of the tracking error*

As we observed in Section 2.2, the root-mean-squared error is only one of many possible summary measures of the tracking error. An obvious alternative is the $L_p$-norm:

$$\mathbf{E}_0[|\varepsilon_1^{(N)}|^p]^{1/p} \tag{5.5}$$

where $p$ is chosen so that the expectation is finite (otherwise the measure will not be particularly informative). More generally, the tracking error can be summarized by

$$\mathbf{E}_0[U(\varepsilon_1^{(N)})] \tag{5.6}$$

where $U(\cdot)$ is an arbitrary loss function.

Consider the set of measures (5.5) first and assume for simplicity that $p \in [1, 2]$. From (2.17), it follows that

$$\mathbf{E}_0[|\varepsilon_1^{(N)}|^p]^{1/p} \sim N^{-1/2}\mathbf{E}_0\left[\left|\frac{1}{\sqrt{2}}\int_0^1 \sigma^2(t,P_t)P_t^2\frac{\partial^2 H(t,P_t)}{\partial P_t^2}\mathrm{d}W_t'\right|^p\right]^{1/p} \tag{5.7}$$

hence, the moments of the stochastic integral in (2.17) describe the asymptotic behavior of the moments of the tracking error. Conditional on the realization of $\{P_t\}$, $t \in [0, 1]$, the stochastic integral on the right side of (5.7) is normally distributed with zero mean and variance

$$\int_0^1 \left( \sigma^2(t, P_t) P_t^2 \frac{\partial^2 H(t, P_t)}{\partial P_t^2} \right)^2 \mathrm{d}t$$

which follows from Hull and White (1987). The intuition is that, conditional on the realization of the integrand, the stochastic integral behaves as an integral of a deterministic function with respect to the Wiener process which is a normal random variable. Now let $m_p$ denote an $L_p$-norm of the standard normal random variable. If $X$ is a standard normal random variable, then $m_p = \mathbf{E}_0[|X|^p]^{1/p}$, hence (5.7) can be rewritten as:

$$\mathbf{E}_0[|\varepsilon_1^{(N)}|^p]^{1/p} \sim \frac{m_p}{\sqrt{N}} \mathbf{E}_0[\mathscr{R}^{p/2}]^{1/p} \tag{5.8}$$

where $\mathscr{R}$ is given by (2.15).

As in the case of a quadratic loss function, $\mathscr{R}$ plays a fundamental role here in describing the behavior of the tracking error. When $p = 2$, $\mathscr{R}$ enters (5.8) linearly and closed-form expressions can be derived for special cases. However, even when $p \neq 2$, the qualitative impact of $\mathscr{R}$ on the tracking error is the same as for $p = 2$ and our discussion of the qualitative behavior of the tracking error applies to this case as well.

For general loss functions $U(\cdot)$ that satisfy certain growth conditions and are sufficiently smooth near the origin, the delta method can be applied and we obtain

$$\mathbf{E}_0[U(\varepsilon_1^{(N)})] \sim \frac{1}{N}|U''(0)|g^2 = \frac{1}{N}|U''(0)|\mathbf{E}_0[\mathscr{R}]. \tag{5.9}$$

When $U(\cdot)$ is not differentiable at zero, the delta method cannot be used. However, we can use the same strategy as in our analysis of $L_p$-norms to tackle this case. Suppose that $U(\cdot)$ is dominated by a quadratic function. Then

$$\mathbf{E}_0[U(\varepsilon_1^{(N)})] \approx \mathbf{E}_0\left[ U\left( \frac{1}{\sqrt{2N}} \int_0^1 \sigma^2(t, P_t) P_t^2 \frac{\partial^2 H(t, P_t)}{\partial P_t^2} \, \mathrm{d}W'_t \right) \right]. \tag{5.10}$$

Now let

$$m_U(x) = \mathbf{E}_0[U(x\eta)], \quad \eta \sim \mathscr{N}(0, 1).$$

Then

$$\mathbf{E}_0[U(\varepsilon_1^{(N)})] \approx \mathbf{E}_0[m_U(\sqrt{\mathscr{R}/N})]. \tag{5.11}$$

When the loss function $U(\cdot)$ is convex, $m_U(\cdot)$ is an increasing function (by second-order stochastic dominance). Therefore, the qualitative behavior of the measure (5.6) is also determined by $\mathscr{R}$ and is the same as that of the RMSE.

## 6. Conclusions

We have argued that continuous-time models are meant to be approximations to physical phenomena, and as such, their approximation errors should be better understood. In the specific context of continuous-time models of derivative securities, we have quantified the approximation error through our definition of temporal granularity. The combination of a specific derivative security and a stochastic process for the underlying asset's price dynamics can be associated with a measure of how 'grainy' the passage of time is. This measure is related to the ability to replicate the derivative security through a delta-hedging strategy implemented in discrete time. Time is said to be very granular if the replication strategy does not work well; in such cases, time is not continuous. If, however, the replication strategy is very effective, time is said to be very smooth or continuous.

Under the assumption of general Markov diffusion price dynamics, we show that the tracking errors for derivatives with sufficiently smooth or continuous piecewise linear payoff functions behave asymptotically (in distribution) as $G/\sqrt{N}$. We characterize the distribution of the random variable $G$ as a stochastic integral, and also obtain the joint distribution of $G$ with prices of other assets and with other tracking errors. We demonstrate that the root-mean-squared error behaves asymptotically as $g/\sqrt{N}$, where the constant $g$ is what we call the coefficient of *temporal granularity*. For two special cases – call or put options on geometric Brownian motion and on an Ornstein–Uhlenbeck process – we are able to evaluate the coefficient of granularity explicitly.

We also consider a number of extensions of our analysis, including an extension to alternative loss functions, a demonstration of the weak convergence of the tracking error process, a derivation of the joint distribution of tracking errors and prices, and an alternative characterization of the tracking error in terms of PDEs that can be used for efficient numerical implementation.

Because these results depend so heavily on continuous-record asymptotics, we perform Monte Carlo simulations to check the quality of our asymptotics. For the case of European puts and calls with geometric Brownian motion price dynamics, our asymptotic approximations are excellent, providing extremely accurate inferences over the range of empirically relevant parameter values, even with a small number of trading periods.

Of course, our definition of granularity is not invariant to the derivative security, the underlying asset's price dynamics, and other variables. But we regard this as a positive feature of our approach, not a drawback. After all, any plausible definition of granularity must be a relative one, balancing the coarseness of changes in the time domain against the coarseness of changes in the 'space' or price domain. Although the title of this paper suggests that time is the main focus of our analysis, it is really the relation between time and price that determines whether or not continuous-time models are good approximations to physical phenomena. It is our hope that the definition of granularity proposed in this paper is one useful way of tackling this very complex issue.

## Appendix A

The essence of these proofs involves the relation between the delta-hedging strategy and mean-square approximations of solutions of systems of stochastic differential equations described in Milstein (1974, 1987, 1995). Readers interested in additional details and intuition should consult these references directly. We present the proof of Theorem 1 only; the proofs for the other theorems are available from the authors upon request.

### *A.1. Proof of Theorem 1(a)*

First we observe that the regularity conditions (2.10) imply the existence of a positive constant $K_1$ such that

$$\left|\frac{\partial^{\beta+\gamma}}{\partial\tau^{\beta}\partial x^{\gamma}}H(\tau,x)\right| \leqslant K_1 \tag{A.1}$$

for $(\tau, x)\in[0,\ 1]\times[0,\infty),0\leqslant\beta\leqslant 1$, and $1\leqslant\gamma\leqslant 4$, and all partial derivatives are continuous. Since the price of the derivative $H(\tau, x)$ is defined as a solution of (2.2), it is equal to the expectation of $F(P_1)$ with respect to the equivalent martingale measure (Duffie, 1996), i.e.,

$$H(\tau,x) = \mathbf{E}_{(t=\tau,P_t^*=x)}[F(P_1^*)], \tag{A.2}$$

where

$$\frac{\mathrm{d}P_t^*}{P_t^*} = \sigma(t,P_t^*)\,\mathrm{d}W_t^*. \tag{A.3}$$

and $W_t^*$ is a Brownian motion under the equivalent martingale measure. Eq. (A.1) now follows from Friedman (1975, Theorems 5.4 and 5.5, p. 122). The same line of reasoning is followed in He (1989, p. 68). Of course, one could derive (A.1) using purely analytic methods, e.g. Friedman (1964; Theorem 10, p. 72; Theorem 11, p. 24; and Theorem 12, p. 25). Next, by Itô's formula,

$$H(1,P_1) = H(0,P_0) + \int_0^1\left(\frac{\partial H(t,P_t)}{\partial t} + \frac{1}{2}\sigma^2(t,P_t)P_t^2\frac{\partial^2 H(t,P_t)}{\partial P_t^2}\right)\mathrm{d}t + \int_0^1\frac{\partial H(t,P_t)}{\partial P_t}\,\mathrm{d}P_t. \tag{A.4}$$

According to (2.2), the first integral on the right-hand side of (A.4) is equal to zero. Thus,

$$H(1,P_1) = H(0,P_0) + \int_0^1 \frac{\partial H(t,P_t)}{\partial P_t}\,\mathrm{d}P_t \tag{A.5}$$

which implies that $H(t,P_t)$ can be characterized as a solution of the system of stochastic differential equations

$$\mathrm{d}X_t = \frac{\partial H(t,P_t)}{\partial P_t}\mu(t,P_t)P_t\,\mathrm{d}t + \frac{\partial H(t,P_t)}{\partial P_t}\sigma(t,P_t)P_t\,\mathrm{d}W_t,$$

$$\mathrm{d}P_t = \mu(t,P_t)P_t\,\mathrm{d}t + \sigma(t,P_t)P_t\,\mathrm{d}W_t. \tag{A.6}$$

At the same time, $V_1^{(N)}$ is given by

$$V_1^{(N)} = H(0,P_0) + \sum_{i=0}^{N-1} \frac{\partial H(t,P_t)_{t=t_i}}{\partial P_t}(P_{t_{i+1}} - P_{t_i}), \tag{A.7}$$

which can be interpreted as a solution of the following approximation scheme of (A.6) as defined in Milstein (1987):

$$\bar{X}_{t_{i+1}} - X_{t_i} = \frac{\partial H(t,P_t)_{t=t_i}}{\partial P_t}(P_{t_{i+1}} - P_{t_i}),$$

$$\bar{P}_{t_{i+1}} - P_{t_i} = P_{t_{i+1}} - P_{t_i}, \tag{A.8}$$

where $\bar{X}$ and $\bar{P}$ denote approximations to $X$ and $P$, respectively. We now compare (A.8) to the Euler approximation scheme in Milstein (1995):

$$\bar{X}_{t_{i+1}} - X_{t_i} = \frac{\partial H(t,P_t)_{t=t_i}}{\partial P_t}\mu(t_i,P_{t_i})P_{t_i}(t_{i+1} - t_i)$$

$$+ \frac{\partial H(t,P_t)_{t=t_i}}{\partial P_t}\sigma(t_i,P_{t_i})P_{t_i}(W_{t_{i+1}} - W_{t_i}),$$

$$\bar{P}_{t_{i+1}} - P_{t_i} = \mu(t_i,P_{t_i})(t_{i+1} - t_i) + \sigma(t_i,P_{t_i})(W_{t_{i+1}} - W_{t_i}). \tag{A.9}$$

Regularity conditions (2.10) and (A.1) allow us to conclude (see Milstein, 1995, Theorem 2.1) that a one-step version of the approximation scheme (A.9) has order-of-accuracy two in expected deviation and order-of-accuracy one in mean-squared deviation; see Milstein (1987, 1995) for definitions and further discussion. It is easy to check that the approximation scheme (A.8) exhibits this same property. Milstein (1995, Theorem 1.1) relates the one-step order-of-accuracy of the approximation scheme to its order-of-accuracy on the whole interval (see also Milstein, 1987). We use this theorem to conclude that (A.8) has mean-square order-of-accuracy 1/2, i.e.,

$$\sqrt{\mathbf{E}_0[(X(1,P_1) - \bar{X}(1,P_1))^2]} = \mathbf{O}\left(\frac{1}{\sqrt{N}}\right). \tag{A.10}$$

We now recall that $X(t,P_t)=H(t,P_t)$ and $\bar{X}(1,P_1)=V_1^{(N)}$ and conclude that

$$\sqrt{\mathbf{E}_0[(H(1,P_1)-V_1^{(N)})^2]}=\mathbf{O}\left(\frac{1}{\sqrt{N}}\right) \tag{A.11}$$

which completes the proof. □

*A.2. Proof of Theorem 1(b)*

We follow the same line of reasoning as in the proof of Theorem 1(a), but we use the Milstein approximation scheme for (A.6) instead of the Euler scheme:

$$\begin{aligned}
\bar{X}_{t_{i+1}}-X_{t_i} &= \frac{\partial H(t,P_t)_{t=t_i}}{\partial P_t}\mu(t_i,P_{t_i})P_{t_i}(t_{i+1}-t_i) \\
&\quad+\frac{\partial H(t,P_t)_{t=t_i}}{\partial P_t}\sigma(t_i,P_{t_i})P_{t_i}(W_{t_{i+1}}-W_{t_i}) \\
&\quad+\left(\frac{\partial^2 H(t,P_t)_{t=t_i}}{\partial P_t^2}\sigma(t_i,P_{t_i})P_{t_i}+\frac{\partial H(t,P_t)_{t=t_i}}{\partial P_t}\frac{\partial(\sigma(t,P_t)P_t)_{t=t_i}}{\partial P_t}\right) \\
&\quad\times\tfrac{1}{2}\sigma(t_i,P_{t_i})P_{t_i}((W_{t_{i+1}}-W_{t_i})^2-(t_{i+1}-t_i)) \\
\bar{P}_{t_{i+1}}-P_{t_i} &= \mu(t_i,P_{t_i})P_{t_i}(t_{i+1}-t_i)+\sigma(t_i,P_{t_i})P_{t_i}(W_{t_{i+1}}-W_{t_i}) \\
&\quad+\tfrac{1}{2}\sigma(t_i,P_{t_i})P_{t_i}\frac{\partial(\sigma(t,P_t)P_t)_{t=t_i}}{\partial P_t}((W_{t_{i+1}}-W_{t_i})^2-(t_{i+1}-t_i)).
\end{aligned} \tag{A.12}$$

According to Milstein (1974) (see also Milstein, 1995, Theorem 2.1), this one-step scheme has order-of-accuracy two in expected deviation and 1.5 in mean-squared deviation. It is easy to check by comparison that the scheme

$$\begin{aligned}
\bar{X}_{t_{i+1}}-X_{t_i} &= \frac{\partial H(t,P_t)_{t=t_i}}{\partial P_t}(P_{t_{i+1}}-P_{t_i}) \\
&\quad+\tfrac{1}{2}\sigma(t_i,P_{t_i})^2P_{t_i}^2\frac{\partial^2 H(t,P_t)_{t=t_i}}{\partial P_t^2}((W_{t_{i+1}}-W_{t_i})^2-(t_{i+1}-t_i)) \\
\bar{P}_{t_{i+1}}-P_{t_i} &= P_{t_{i+1}}-P_{t_i}
\end{aligned} \tag{A.13}$$

has the same property. We now use Milstein (1995, Theorem 1.1) to conclude that

$$\begin{aligned}
H(1,P_1)-V_1^{(N)} &= \sum_{i=0}^{N-1}\tfrac{1}{2}\sigma(t_i,P_{t_i})^2P_{t_i}^2\frac{\partial^2 H(t,P_t)_{t=t_i}}{\partial P_t^2} \\
&\quad\times((W_{t_{i+1}}-W_{t_i})^2-(t_{i+1}-t_i))+\mathbf{O}\left(\frac{1}{N}\right)
\end{aligned} \tag{A.14}$$

where $f = \mathbf{O}(\frac{1}{N})$ means that $\lim_{N\to\infty} N\sqrt{\mathbf{E}_{t=0}[f^2]} < \infty$. By Slutsky's theorem, we can ignore the $\mathbf{O}(\frac{1}{N})$ term in considering the convergence in distribution of $\sqrt{N}\ (H(1,P_1) - V_1^{(N)})$, since $\sqrt{N}\ \mathbf{O}(\frac{1}{N})$ converges to zero in mean-squared and, therefore, also in probability. Observe now that, since $W_{t_{i+1}} - W_{t_i}$ and $W_{t_{j+1}} - W_{t_j}$ are independent for $i \neq j$, $(W_{t_{i+1}} - W_{t_i})^2$ and $W_{t_{i+1}} - W_{t_i}$ are uncorrelated, $\mathbf{E}_0[(W_{t_{i+1}} - W_{t_i})^2 - (t_{i+1} - t_i)] = 0$, and $\mathbf{E}_0[((W_{t_{i+1}} - W_{t_i})^2 - (t_{i+1} - t_i))^2] = 2/(t_{i+1} - t_i)^2$, by the functional central limit theorem (see Ethier and Kurtz, 1986), a piecewise constant martingale

$$\sqrt{N/2} \sum_{i=0}^{[Nt]-1} ((W_{t_{i+1}} - W_{t_i})^2 - (t_{i+1} - t_i)) \tag{A.15}$$

converges weakly on [0, 1] to a standard Brownian motion $W'_t$, which is independent of $W_t$. The notation $[Nt]$ denotes the integer part of $Nt$ and we use the convention $\sum_0^{-1} = 0$. We complete the proof by applying Duffie and Protter (1992, Lemma 5.1 and Corollary 5.1).

### *A.3. Proof of Theorem 1(c)*

Eq. (2.13) follows immediately from Theorem 1(a) and the proof of Theorem 1(b). Relation (A.14), established as a part of the proof of Theorem 1(b), guarantees that convergence in (2.12) occurs not only in distribution, but also in mean-square. Combined with Theorem 1(b), (2.13) implies that

$$g = \sqrt{\frac{1}{2}\mathbf{E}_0\left[\left(\int_0^1 \sigma^2(t,P_t)P_t^2 \frac{\partial^2 H(t,P_t)}{\partial P_t^2}\,\mathrm{d}W'_t\right)^2\right]}. \tag{A.16}$$

Eq. (2.14) follows from (A.16) using the isometric property of stochastic integrals.

## References

Avellaneda, M., Paras, A., 1994. Dynamic hedging portfolios for derivative securities in the presence of large transactions costs. Applied Mathematical Finance 1, 165–193.

Bachelier, L., 1900. Theory of speculation. In: Cootner, P. (Ed.), 1964. The Random Character of Stock Market Prices. MIT Press, Cambridge, MA.

Bensaid, B., Lesne, J., Pages, H., Scheinkman, J., 1992. Derivative asset pricing with transaction costs. Mathematical Finance 2, 63–86.

Bertsimas, D., Kogan, L., Lo, A., 1997. Pricing and hedging derivative securities in incomplete markets: an ε-arbitrage approach. Working paper No. LFE–1027–97, MIT Laboratory for Financial Engineering.

Billingsley, P., 1986. Probability and Measure, 2nd Edition. Wiley, New York.

Black, F., Scholes, M., 1973. Pricing of options and corporate liabilities. Journal of Political Economy 81, 637–654.

Boyle, P., Emanuel, D., 1980. Discretely adjusted option hedges. Journal of Financial Economics 8, 259–282.
Boyle, P., Vorst, T., 1992. Option replication in discrete time with transaction costs. Journal of Finance 47, 271–294.
Cox, J., Ingersoll, J., Ross, S., 1985. An intertemporal general equilibrium model of asset prices. Econometrica 36, 363–384.
Cox, J., Ross, S., Rubinstein, M., 1979. Option pricing: a simplified approach. Journal of Financial Economics 7, 229–263.
Duffie, D., 1996, Dynamic Asset Pricing Theory, 2nd Edition. Princeton University Press, Princeton, NJ.
Duffie, D., Huang, C., 1985. Implementing Arrow–Debreu equilibria by continuous trading of few long–lived securities. Econometrica 53, 1337–1356.
Duffie, D., Protter, P., 1992. From discrete-to continuous-time finance: weak convergence of the financial gain process, Mathematical Finance. 2 1–15.
Edirisinghe, C., Naik, V., Uppal, R., 1993. Optimal replication of options with transaction costs and trading restrictions. Journal of Financial and Quantitative Analysis 28, 117–138.
Ethier, S., Kurtz, T., 1986, Markov Processes: Characterization and Convergence. Wiley, New York.
Figlewski, S., 1989. Options arbitrage in imperfect markets. Journal of Finance 44, 1289–1311.
Foley, D., 1975. On two specifications of asset equilibrium in macroeconomic models. Journal of Political Economy 83, 303–324.
Friedman, A., 1964. Partial Differential Equations of Parabolic Type. Prentice-Hall, Englewood Cliffs, NJ.
Friedman, A., 1975. Stochastic Differential Equations and Applications, Vol. 1. Academic Press, New York.
Grannan, E., Swindle, G., 1996. Minimizing transaction costs of option hedging strategies. Mathematical Finance 6, 341–364.
Harrison, J., Kreps, D., 1979. Martingales and arbitrage in multiperiod securities markets. Journal of Economic Theory 20, 381–408.
He, H., 1989. Essays in dynamic portfolio optimization and diffusion estimations. Ph.D. Thesis, MIT Press, Cambridge, MA.
He, H., 1990. Convergence from discrete to continuous-time contingent claim prices. Review of Financial Studies 3, 523–546.
He, H., 1991. Optimal consumption-portfolio policies: a convergence from discrete to continuous time models. Journal of Economic Theory 55, 340–363.
Henrotte, P., 1993. Transactions costs and duplication strategies. Unpublished working paper, Stanford University.
Hodges, S., Neuberger, A., 1989. Optimal replication of contingent claims under transactions costs. unpublished paper. Financial Options Research Center, University of Warwick.
Hull, J., White, A., 1987. The pricing of options on assets with stochastic volatilities. Journal of Finance 42, 281–300.
Hutchinson, J., Lo, A., Poggio, T., 1994. A nonparametric approach to the pricing and hedging of derivative securities via learning networks. Journal of Finance 49, 851–889.
Itô, K., 1951. On stochastic differential equations. Memoirs of the American Mathematical Society 4, 1–51.
Karatzas, I., Shreve, S., 1991. Brownian Motion and Stochastic Calculus. Springer, New York.
Kushner, H., Dupuis, P., 1992. Numerical Methods for Stochastic Control Problems in Continuous Time. Springer, New York.
Leland, H., 1985. Option pricing and replication with transaction costs. Journal of Finance 40, 1283–1301.
Lo, A., Wang, J., 1995. Implementing option pricing models when asset returns are predictable. Journal of Finance 50, 87–129.

Merton, R., 1969. Lifetime portfolio selection under uncertainty: the continuous-time case. Review of Economics and Statistics 51, 247–257.

Merton, R., 1973. Theory of rational option pricing. Bell Journal of Economics and Management Science 4, 141–183.

Merton, R., 1975. Theory of finance from the perspective of continuous time. Journal of Financial and Quantitative Analysis 10, 659–674.

Merton, R., 1982. On the mathematics and economic assumptions of continuous-time models. In: Sharpe, W., Cootner, C. (Eds.). Financial Economics: Essays in Honor of Paul Cootner. Prentice-Hall, Englewood Cliffs, NJ.

Milstein, G., 1974. Approximate integration of stochastic differential equations. Theory of Probability and Its Applications 19, 557–563.

Milstein, G., 1987. A theorem on the order of convergence of mean-square approximations of solutions of systems of stochastic differential equations. Theory of Probability and Its Applications 32, 738–741.

Milstein, G., 1995. Numerical Integration of Stochastic Differential Equations. Kluwer Academic Publishers, Boston, MA.

Neuberger, A., 1994. Option replication with transaction costs: an exact solution for the pure jump process. In: Chance, D., Trippi, R. (Eds.). Advances in Futures and Options Research, Vol. 7. JAI Press, Greenwich, CT.

Rubinstein, M., 1994. Implied binomial trees. Journal of Finance 49, 771–818.

Toft, K., 1996. On the mean-variance tradeoff in option replication with transactions costs. Journal of Financial and Quantitative Analysis 31, 233–263.

Wiener, N., 1923. Differential space, Journal of Mathematical Physics. 2 131–174.

Willinger, W., Taqqu, M., 1991. Towards a convergence theory for continuous stochastic securities market models. Mathematical Finance 1, 55–99.