

**PRICING AND HEDGING DERIVATIVE
SECURITIES IN INCOMPLETE
MARKETS: AN ϵ -ARBITRAGE APPROACH**

**Dimitris Bertsimas
Leonid Kogan
Andrew W. Lo**

Working Paper 6250

NBER WORKING PAPER SERIES

PRICING AND HEDGING DERIVATIVE
SECURITIES IN INCOMPLETE
MARKETS: AN ϵ -ARBITRAGE APPROACH

Dimitris Bertsimas
Leonid Kogan
Andrew W. Lo

Working Paper 6250
<http://www.nber.org/papers/w6250>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 1997

This research was partially supported by the MIT Laboratory for Financial Engineering, a Presidential Young Investigator Award DDM-9158118 with matching funds from Draper Laboratory and a grant from the National Science Foundation to the National Bureau of Economic Research. We thank Chi-fu Huang and Jiang Wang for helpful discussions and seminar participants at MIT, NYU, the 1997 Spring INFORMS Conference, Fudan University, and Tsinghua University for comments. This paper is part of NBER's research program in Asset Pricing. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

© 1997 by Dimitris Bertsimas, Leonid Kogan and Andrew W. Lo. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Pricing and Hedging Derivative Securities in
Incomplete Markets: An ϵ -Arbitrage Approach
Dimitris Bertsimas, Leonid Kogan and Andrew W. Lo
NBER Working Paper No. 6250
November 1997
JEL No. G13
Asset Pricing

ABSTRACT

Given a European derivative security with an arbitrary payoff function and a corresponding set of underlying securities on which the derivative security is based, we solve the dynamic replication problem: find a self-financing dynamic portfolio strategy—involving only the underlying securities—that most closely approximates the payoff function at maturity. By applying stochastic dynamic programming to the minimization of a mean-squared-error loss function under Markov state-dynamics, we derive recursive expressions for the optimal-replication strategy that are readily implemented in practice. The approximation error or “ ϵ ” of the optimal-replication strategy is also given recursively and may be used to quantify the “degree” of market incompleteness. To investigate the practical significance of these ϵ -arbitrage strategies, we consider several numerical examples including path-dependent options and options on assets with stochastic volatility and jumps.

Dimitris Bertsimas
Sloan School of Management, E52-359
Massachusetts Institute of Technology
50 Memorial Drive
Cambridge, MA 02142-1347

Leonid Kogan
Sloan School of Management, E52-450
Massachusetts Institute of Technology
50 Memorial Drive
Cambridge, MA 02142-1347

Andrew W. Lo
Sloan School of Management, E52-432
Massachusetts Institute of Technology
50 Memorial Drive
Cambridge, MA 02142-1347
and NBER
alo@mit.edu

Contents

1	Introduction	1
2	ϵ-Arbitrage Strategies	5
2.1	The Dynamic Replication Problem	5
2.1.1	Examples	9
2.2	ϵ -Arbitrage in Discrete Time	10
2.3	ϵ -Arbitrage in Continuous Time	13
2.4	Interpreting ϵ^* and V_0^*	15
2.4.1	V_0^* Is Not a Price	15
2.4.2	Why Mean-Squared Error?	16
3	Risk-Neutralized ϵ-Arbitrage Strategies	17
3.1	Equilibrium Pricing Models and ϵ -Arbitrage	17
3.2	How to Obtain ν^*	19
3.2.1	Theoretical Methods	19
3.2.2	Empirical Methods	21
4	Illustrative Examples	22
4.1	State-Independent Returns	22
4.2	Geometric Brownian Motion	23
4.3	Jump-Diffusion Models	25
4.3.1	The Continuous-Time Limit	27
4.3.2	Perturbation Analysis with Small Jump Amplitudes	28
4.4	Stochastic Volatility	30
4.4.1	The Continuous-Time Solution	30
5	Numerical Analysis	32
5.1	The Numerical Procedure	32
5.2	Geometric Brownian Motion	33
5.3	Jump-Diffusion Models	35
5.4	Stochastic Volatility	38
5.5	Stochastic Volatility Under The Risk-Neutral Measure	44
5.6	Path-Dependent Options	45
6	Specification Analysis of Replication Errors	48
7	Conclusion	53
A	Appendix	55
A.1	Proof of Theorem 1	55
A.2	Proof of Theorem 2	55

1 Introduction

One of the most important breakthroughs in modern financial economics is Merton's (1973) insight that under certain conditions the frequent trading of a small number of long-lived securities can create new investment opportunities that would otherwise be unavailable to investors. These conditions—now known collectively as *dynamic spanning* or *dynamically complete markets*—and the corresponding asset-pricing models on which they are based have generated a rich literature and an even richer industry in which complex financial securities are synthetically replicated by sophisticated trading strategies involving considerably simpler instruments.¹ This approach is the basis of the celebrated Black and Scholes (1973) and Merton (1973) option-pricing formula, the arbitrage-free method of pricing and, more importantly, hedging other derivative securities, and the martingale characterization of prices and dynamic equilibria.

The essence of dynamic spanning is the ability to replicate exactly the payoff of a complex security by a dynamic portfolio strategy of simpler securities which is *self-financing*, i.e., no cash inflows or outflows except at the start and at the end. If such a *dynamic-hedging* strategy exists, then the initial cost of the portfolio must equal the price of the complex security, otherwise an arbitrage opportunity exists. For example, under the assumptions of Black and Scholes (1973) and Merton (1973), the payoff of a European call-option on a non-dividend-paying stock can be replicated exactly by a dynamic-hedging strategy involving only stocks and riskless borrowing and lending.

But the conditions that guarantee dynamic spanning are nontrivial restrictions on market structure and price dynamics (see, for example, Duffie and Huang [1985]), hence there are situations in which exact replication is impossible.² These instances of *market incompleteness* are often attributable to institutional rigidities and market frictions—transactions costs, periodic market closures, and discreteness in trading opportunities and prices—and while the pricing of complex securities can still be accomplished in some cases via equilib-

¹In addition to Merton's seminal paper, several other important contributions to the finance literature are responsible for our current understanding of dynamic spanning. In particular, see Cox and Ross (1976), Duffie (1985), Duffie and Huang (1985), Harrison and Kreps (1979), and Huang (1985a,b).

²Suppose, for example, that stock price volatility σ in the Black and Scholes (1973) framework is stochastic.

rium arguments,³ this still leaves the question of dynamic replication unanswered. Perfect replication is impossible in dynamically incomplete markets, but how close can one come, and what is does the optimal-replication strategy look like?

In this paper we answer these questions by applying optimal control techniques to the dynamic replication problem: given an arbitrary payoff function and a set of fundamental securities, find a self-financing dynamic portfolio strategy involving only the fundamental securities that most closely approximates the payoff. The initial cost of such an optimal strategy can be viewed as a proxy for the price of the security—it is the cost of the best dynamic approximation to the payoff function given the set of fundamental securities traded, i.e., the minimum “production cost” of the option.⁴ Such an interpretation is more than a figment of economic imagination: the ability to synthesize options via dynamic trading strategies has fueled the growth of the multi-trillion-dollar over-the-counter derivatives market.⁵

Of course, the nature of the optimal-replication strategy depends intimately on how we measure the closeness of the payoff and its approximation. For tractability and other reasons (see Section 2.4), we choose a mean-squared-error loss function and we denote by ϵ the root-mean-squared-error. In a dynamically complete market, the approximation error ϵ is identically zero, but when the market is incomplete ϵ can be used to quantify the “degree” of incompleteness. Although from a theoretical point of view dynamic spanning either holds or does not hold, a gradient for market completeness seems more natural from an empirical and a practical point of view. We provide examples of stochastic processes that imply dynamically incomplete markets, e.g., stochastic volatility, and yet still admit ϵ -arbitrage strategies for replicating options to within ϵ where ϵ can be evaluated numerically.

More importantly, we introduce a slight modification of the mean-squared-error loss func-

³Examples of continuous-time incomplete-markets models include Duffie (1987), Duffie and Shafer (1985, 1986), Föllmer and Sonderman (1986), and He and Pearson (1991). Examples of discrete-time incomplete-markets models include Aiyagari (1994), Aiyagari and Gertler (1991), He and Modest (1995), Heaton and Lucas (1992, 1996), Lucas (1994), Scheinkman and Weiss (1986), Telmer (1993), and Weil (1992).

⁴This minimum production cost of the optimal-replication strategy cannot be interpreted as a price because we have not specified a set of preferences and market-clearing conditions that supports such a strategy. In particular, agents that differ in preferences may well value the optimal strategy differently. See Section 2.4.1 for further discussion.

⁵In contrast to exchange-traded options such as equity puts and calls, over-the-counter derivatives are considerably more illiquid. If investment houses were unable to synthesize them via dynamic trading strategies, they would have to take the other size of every option position that their clients' wish to take (net of offsetting positions among the clients themselves). Such risk exposure would dramatically curtail the scope of the derivatives business, limiting both the size and type of contracts available to end users.

tion which does allow us to interpret the minimum production cost as a price: we replace the probability measure of mean-squared-error loss function with the *equivalent martingale measure*. Optimizing this loss function yields a minimum production cost that must equal the equilibrium price of the option, hence under the equivalent martingale measure our ϵ -arbitrage strategies have a deeper economic motivation.

In this respect, our contribution extends the results of Schweizer (1992, 1995) in which the dynamic replication problem is also solved for a mean-squared-error loss function but under the probability measure of the original price process, not the equivalent martingale measure. Also, Schweizer considers more general stochastic processes than we do—we focus only on Markov price processes—and uses variational principles to characterize the optimal-replication strategy. Although our approach can be viewed as a special case of his, the Markov assumption allows us to obtain considerably sharper results and yields an easily implementable numerical procedure (via dynamic programming) for determining the optimal-replication strategy and the replication error ϵ in practice.

To demonstrate the practical relevance of our optimal-replication strategy, even in the simplest case of the Black and Scholes (1973) model where an explicit dynamic-replication strategy is available, Table 1 presents a comparison of our optimal-replication strategy with the standard Black-Scholes “delta-hedging” strategy for replicating an at-the-money put option on 1,000 shares of a \$40-stock over 25 trading periods for two simulated sample paths of a geometric Brownian motion with drift $\mu = 0.07$ and diffusion coefficient $\sigma = 0.13$ (rounded to the nearest \$0.125).

V_t^* denotes the period- t value of the optimal replicating portfolio, θ_t^* denotes the number of shares of stock held in that portfolio, and V_t^{BS} and θ_t^{BS} are defined similarly for the Black-Scholes strategy.

Despite the fact that both sample paths are simulated geometric Brownian motions with identical parameters, the optimal-replication strategy has a higher replication error than the Black-Scholes strategy for path A and a lower replication error than Black-Scholes for path B.⁶ That the optimal-replication strategy underperforms the Black-Scholes strategy for path

⁶Specifically, $V_{25}^* - 1000 \times \text{Max}[0, \$40 - P_{25}] = \$199.1$ and $V_{25}^{\text{BS}} - 1000 \times \text{Max}[0, \$40 - P_{25}] = \$172.3$ for path A, and $V_{25}^* - 1000 \times \text{Max}[0, \$40 - P_{25}] = -\$40.3$ and $V_{25}^{\text{BS}} - 1000 \times \text{Max}[0, \$40 - P_{25}] = -\$299.2$ for path B.

A is not surprising since the optimal-replication strategy is optimal only in a mean-squared sense (see Section 2.1), not path by path.⁷ That the Black-Scholes strategy underperforms the optimal-replication strategy for path B is also not surprising since the former is designed to replicate the option with continuous trading whereas the optimal-replication strategy is designed to replicate the option with 25 trading periods.

Table 1: Comparison of optimal-replication strategy and Black-Scholes delta-hedging strategy for replicating an at-the-money put option on 1,000 shares of a \$40-stock over 25 trading periods for two simulated sample paths of a geometric Brownian motion with parameters $\mu = 0.07$ and $\sigma = 0.13$.

Period t	Sample Path A					Sample Path B				
	P_t	V_t^*	θ_t^*	V_t^{BS}	θ_t^{BS}	P_t	V_t^*	θ_t^*	V_t^{BS}	θ_t^{BS}
0	40.000	1461.0	-474.8	1466.4	-481.7	40.000	1461.0	-474.8	1466.4	-481.7
1	40.750	1104.9	-394.2	1105.1	-400.8	39.875	1520.4	-491.8	1526.6	-495.9
2	42.125	562.9	-252.1	554.0	-264.0	40.250	1336.0	-452.4	1340.6	-454.3
3	41.375	751.9	-318.6	752.0	-331.8	36.500	3032.3	-762.4	3044.3	-845.8
4	42.000	552.8	-256.1	544.7	-267.2	36.875	2746.4	-744.1	2727.1	-822.1
5	43.125	264.7	-157.8	244.0	-169.5	36.500	3025.4	-777.9	3035.4	-858.3
6	43.250	245.0	-145.6	222.8	-155.1	37.000	2636.5	-750.1	2606.3	-824.5
7	42.250	390.6	-216.7	377.9	-229.5	39.875	479.9	-407.8	235.8	-500.5
8	43.000	228.2	-149.7	205.8	-160.6	39.875	479.9	-412.1	235.8	-501.4
9	41.750	415.2	-250.3	406.6	-267.9	40.125	376.9	-384.8	110.4	-468.4
10	42.000	352.7	-221.1	339.6	-235.5	39.500	617.3	-466.9	403.2	-556.1
11	42.625	214.5	-157.2	192.4	-168.9	41.250	-199.7	-227.2	-570.0	-315.0
12	41.750	352.1	-233.8	340.2	-248.5	40.625	-57.7	-300.8	-373.1	-394.7
13	41.500	410.5	-258.4	402.3	-271.0	39.875	167.9	-403.7	-77.0	-506.9
14	42.625	119.8	-128.5	97.5	-141.7	39.375	369.8	-482.4	176.4	-590.1
15	42.875	87.7	-100.5	62.1	-110.7	39.625	249.2	-452.1	28.9	-552.9
16	42.875	87.7	-91.5	62.1	-99.2	39.750	192.7	-439.2	-40.2	-534.3
17	43.125	64.8	-65.5	37.3	-70.4	39.250	412.3	-533.4	226.9	-632.3
18	43.000	73.0	-62.3	46.1	-65.4	39.500	278.9	-500.5	68.8	-592.6
19	43.000	73.0	-50.8	46.1	-51.7	39.750	153.8	-461.4	-79.3	-546.5
20	41.875	130.2	-121.8	104.3	-128.2	39.750	153.8	-472.5	-79.3	-552.5
21	41.125	221.5	-209.8	200.4	-219.8	39.875	94.8	-452.8	-148.4	-526.6
22	41.375	169.1	-137.5	145.5	-140.7	39.625	208.0	-538.7	-16.7	-610.2
23	40.625	272.2	-263.5	251.0	-271.2	39.875	73.3	-476.8	-169.3	-542.8
24	40.000	436.9	-475.7	420.5	-496.3	40.000	13.7	-432.2	-237.1	-496.3
25	40.500	199.1	0.0	172.3	0.0	40.125	-40.3	0.0	-299.2	0.0

⁷These two sample paths were chosen to be illustrative, not conclusive. In a more extensive simulation study in which 250,000 sample paths were generated, the average replication error of the Black-Scholes strategy is \$248.0 and the average error of the optimal-replication strategy is \$241.2.

For sample path A, the differences between the optimal-replication strategy and the Black-Scholes are not great— V_t^* and θ_t^* are fairly close to their Black-Scholes counterparts. However, for sample path B, where there are two large price movements, the differences between the two replication strategies and the replication errors are substantial. Even in such an idealized setting, the optimal-replication strategy can still play an important role in the dynamic hedging of risks.

In Section 2 we introduce the dynamic replication problem and propose a solution based on stochastic dynamic programming. In Section 3 we recast the dynamic replication problem under the equivalent martingale measure, which generalizes the typical derivative pricing and hedging results to dynamically incomplete markets. The scope of the ϵ -arbitrage approach is illustrated in Sections 4 and 5 analytically and numerically for several examples including path-dependent options and options on assets with mixed jump-diffusion and stochastic-volatility price dynamics. The sensitivity of the replication error to price dynamics is studied in Section 6, and we conclude in Section 7.

2 ϵ -Arbitrage Strategies

In this section, we formulate and propose a solution approach for the problem of option pricing in incomplete markets. In Section 2.1 we introduce the replication problem and the principle of ϵ -arbitrage. In Sections 2.2 and 2.3 we propose stochastic dynamic programming algorithms in discrete and continuous time, respectively.

2.1 The Dynamic Replication Problem

Consider an asset with price P_t at time t where $0 \leq t \leq T$ and let $F(P_T, \mathbf{Z}_T)$ denote the payoff of some European derivative security at maturity date T which is a function of P_T and other variables \mathbf{Z}_T (see below). For expositional convenience, we shall refer to the asset as a *stock* and the derivative security as an *option* on that stock, but our results are considerably more general.

As suggested by Merton's (1973) derivation of the Black-Scholes formula, the dynamic replication problem is to find a dynamic portfolio strategy—purchases and sales of stock and riskless borrowing and lending—on $[0, T]$ that is self-financing and comes as close as possible

to the payoff $F(P_T, \mathbf{Z}_T)$ at T . To formulate the dynamic replication problem more precisely, we begin with the following assumptions:

(A1) *Markets are frictionless, i.e., there are no taxes, transactions costs, shortsales restrictions, and borrowing restrictions.*

(A2) *The riskless borrowing and lending rate is 0.*⁸

(A3) *There exists a finite-dimensional vector \mathbf{Z}_t of state variables whose components are not perfectly correlated with the prices of any traded securities, and $[P_t \ \mathbf{Z}_t]'$ is a vector Markov process*

(A4) *Trading takes place at known fixed times $t \in \mathcal{T}$. If $\mathcal{T} = \{t_0, t_1, \dots, t_N\}$, trading is said to be discrete. If $\mathcal{T} = [0, T]$, trading is said to be continuous.*

At time 0 consider forming a portfolio of stocks and riskless bonds at a cost V_0 and as time progresses, let θ_t , B_t , and V_t denote the number of shares of the stock held, the dollar value of bonds held, and the market value of the portfolio at time t , respectively, $t \in \mathcal{T}$, hence:

$$V_t = \theta_t P_t + B_t . \tag{2.1}$$

In addition, we impose the condition that after time 0, the portfolio is *self-financing*, i.e., all long positions in one asset are completely financed by short positions in the other asset so that the portfolio experiences no cash inflows or outflows:

$$P_{t_{i+1}}(\theta_{t_{i+1}} - \theta_{t_i}) + B_{t_{i+1}} - B_{t_i} = 0 \quad , \quad 0 < t_i < t_{i+1} \leq T . \tag{2.2}$$

This implies that:

$$V_{t_{i+1}} - V_{t_i} = \theta_{t_i}(P_{t_{i+1}} - P_{t_i}) \tag{2.3}$$

⁸This entails no loss of generality since we can always renormalize all prices by the price of a zero-coupon bond with maturity at time T (see, for example, Harrison and Kreps [1979]).

and, in continuous time,

$$dV_t = \theta_t dP_t . \quad (2.4)$$

We seek a self-financing portfolio strategy $\{\theta_t\}$, $t \in \mathcal{T}$, such that the terminal value V_T of the portfolio is as close as possible to the option's payoff $F(P_T, \mathbf{Z}_T)$. Of course, there are many ways of measuring "closeness", each giving rise to a different dynamic replication problem. For reasons that will become clear shortly (see Sections 2.4 and 3), we choose a mean-squared-error loss function, hence our version of the dynamic replication problem is:⁹

$$\min_{\{\theta_t\}} E^\nu \left[V_T - F(P_T, \mathbf{Z}_T) \right]^2 \quad (2.5)$$

subject to self-financing condition (2.3) or (2.4), the dynamics of $[P_t \ \mathbf{Z}_t]'$, and the initial wealth V_0 , where the expectation E^ν is taken with respect to a probability measure ν that represents the randomness of the difference $V_T - F(P_T, \mathbf{Z}_T)$, conditional on information at time 0.¹⁰

A natural measure of the success of the optimal-replication strategy is the square root of the mean-squared replication error (2.5) evaluated at the optimal $\{\theta_t\}$, hence we define

$$\epsilon(V_0) \equiv \sqrt{\min_{\{\theta_t\}} E^\nu \left\{ [V_T - F(P_T, \mathbf{Z}_T)]^2 \right\}} . \quad (2.6)$$

We shall show below that $\epsilon(V_0)$ can be minimized with respect to the initial wealth V_0 to yield the *least-cost* optimal-replication strategy and a corresponding measure of the minimum

⁹Other recent examples of the use of mean-squared-error loss functions in related dynamic-trading problems include Duffie and Jackson (1990), Duffie and Richardson (1991), Schäl (1994), and Schweizer (1992, 1995).

¹⁰Note that we have placed no constraints on $\{\theta_t\}$, hence it is conceivable that for certain replication strategies, V_T is negative with positive probability. Imposing constraints on $\{\theta_t\}$ to ensure the non-negativity of V_T would render the dynamic replication problem (2.5) intractable. However, negative values for V_T is not nearly as problematic in the context of the dynamic replication problem as it is for the optimal consumption and portfolio problem of, for example, Merton (1971). In particular, V_T does not correspond to an individual's wealth, but is the terminal value of a portfolio designed to replicate a particular payoff function. See Dybvig and Huang (1988) and Merton (1992, Chapter 6) for further discussion.

replication error ϵ^* :

$$\epsilon^* \equiv \min_{\{V_0\}} \epsilon(V_0) . \quad (2.7)$$

In the case of Black and Scholes (1973) and Merton (1973), there exists dynamic replication strategies for which $\epsilon^* = 0$, hence we say that *perfect* arbitrage pricing holds.

But there are situations—dynamically incomplete markets, for example—where perfect arbitrage pricing does not hold. In particular, assumption (A3), the presence of state variables \mathbf{Z}_t that are not perfectly correlated with the prices of any traded securities, is the source of market incompleteness in our framework. While this captures only one potential source of incompleteness—and does so only in a “reduced-form” sense—nevertheless, it is a particularly relevant source of incompleteness in financial markets. Of course, we recognize that the precise nature of incompleteness, e.g., institutional rigidities, transactions costs, technological constraints, will affect the pricing and hedging of derivative securities in complex ways.¹¹ Nevertheless, how well one security can be replicated by sophisticated trading in other securities does provide one measure of the degree of market incompleteness even if it does not completely characterize it. In much the same way that the Black and Scholes (1973) and Merton (1973) models focus on the *relative* pricing of options—relative to the exogenously specified price dynamics for the underlying asset—we hope to capture the degree of *relative* incompleteness, relative to an exogenously specified set of Markov state variables that are not completely hedgeable.

In some of these cases, we shall show in Sections 2.2 and 2.3 that ϵ -arbitrage pricing is possible, i.e., it is possible to derive a mean-square-optimal dynamic replication strategy that is able to approximate the terminal payoff $F(P_T, \mathbf{Z}_T)$ of an option to within ϵ^* . But before turning to the solution of the dynamic replication problem, we provide several illustrative examples that delineate the scope of our framework.

¹¹For more “structural” models in which institutional sources of market incompleteness are studied, e.g., transactions costs, shortsales constraints, undiversifiable labor income, see Aiyagari (1994), Aiyagari and Gertler (1991), He and Modest (1995), Heaton and Lucas (1992, 1996), Lucas (1994), Scheinkman and Weiss (1986), Telmer (1993), and Weil (1992). See Magill and Quinzii (1996) for a comprehensive analysis of market incompleteness.

2.1.1 Examples

In particular, despite the restrictions imposed by assumptions (A1)–(A4), our framework can accommodate many kinds of market incompleteness as the following examples illustrate:

- (a) **Stochastic Volatility.** Consider a stock price process that follows a diffusion process with stochastic volatility, e.g., Hull and White (1987) and Wiggins (1987). The stock price and stock-price volatility are assumed to be governed by the following pair of stochastic differential equations:

$$\begin{aligned} dP_t &= \mu P_t dt + \sigma_t P_t dW_{P_t} \\ d\sigma_t &= g(\sigma_t) dt + \kappa \sigma_t dW_{\sigma_t} \end{aligned}$$

where W_{P_t} and W_{σ_t} are Brownian motions with mutual variation $dW_{P_t} dW_{\sigma_t} = \rho dt$. This stochastic volatility model is included in our framework by defining $Z_t = \sigma_t$. Then, clearly the vector process $[P_t \ Z_t]'$ is Markov.

- (b) **Options on the Maximum.** In this and the next two examples we assume that $\mathcal{T} = \{t_0, t_1, \dots, t_N\}$ and that the stock price P_t process is Markov for expositional simplicity. The payoff of the option on the maximum stock price is given by

$$F\left(\max_{i=0, \dots, N} P_{t_i}\right). \quad (2.8)$$

Define the state variable

$$Z_{t_i} \equiv \max_{k=0, \dots, i} P_{t_k}.$$

The process $[P_{t_i} \ Z_{t_i}]'$ is Markov since the distribution of $P_{t_{i+1}}$ depends only on P_{t_i} and

$$Z_{t_{i+1}} = \max(Z_{t_i}, P_{t_{i+1}}) \quad , \quad Z_0 = P_0.$$

The payoff of the option can be expressed in terms of the terminal value of the state variables (P_T, Z_T) as $F(Z_T)$.

(c) **Asian Options.** The payoff of “Asian” or “average-rate” options is given by

$$F\left(\frac{1}{N+1}\sum_{i=0}^N P_{t_i}\right).$$

Let Z_{t_i} be the following state variable

$$Z_{t_i} \equiv \frac{1}{i+1}\sum_{k=0}^i P_{t_k}$$

and observe that the process $[P_{t_i} \ Z_{t_i}]'$ is Markov since the distribution of $P_{t_{i+1}}$ depends only on P_{t_i} and

$$Z_{t_{i+1}} = \frac{Z_{t_i}(i+1) + P_{t_{i+1}}}{(i+2)}, \quad Z_0 = P_0.$$

As before, the payoff of the option can be written as $F(Z_T)$.

(d) **Knock-Out Options.** Given a knock-out price \bar{P} , the payoff of a knock-out option is $\beta_T h(P_T)$, where $h(\cdot)$ is a function of the terminal stock price and

$$\beta_T = \begin{cases} 1 & \text{if } \max_{i=0,\dots,N} P_{t_i} \leq \bar{P} \\ 0 & \text{if } \max_{i=0,\dots,N} P_{t_i} > \bar{P}. \end{cases}$$

Define the state variable Z_t :

$$Z_0 = \begin{cases} 1 & \text{if } P_0 \leq \bar{P} \\ 0 & \text{if } P_0 > \bar{P} \end{cases}$$

$$Z_{t_{i+1}} = \begin{cases} 1 & \text{if } P_{t_{i+1}} \leq \bar{P} \text{ and } Z_{t_i} = 1, \\ 0 & \text{otherwise} \end{cases}.$$

It is easy to see that resulting process $[P_{t_i} \ Z_{t_i}]'$ is Markov, $Z_T = \beta_T$. The payoff of the option is given by $F(P_T, Z_T) = Z_T h(P_T)$.

2.2 ϵ -Arbitrage in Discrete Time

In this section, we propose a solution for the dynamic replication problem (2.5) in discrete time via stochastic dynamic programming. To simplify notation, we adopt the following

convention for discrete-time quantities: time subscripts t_i are replaced by i , e.g., the stock price P_{t_i} will be denoted as P_i and so on. Under this convention, we can define the usual *cost-to-go* or *value* function J_i as:

$$J_i(V_i, P_i, \mathbf{Z}_i) \equiv \min_{\theta^{(k, V_k, P_k, \mathbf{Z}_k)}, i \leq k \leq N-1} \mathbb{E}^\nu \left[[V_N - F(P_N, \mathbf{Z}_N)]^2 \mid V_i, P_i, \mathbf{Z}_i \right] \quad (2.9)$$

where V_i , P_i , and \mathbf{Z}_i comprise the state variables, θ_i is the control variable, and the self-financing condition (2.3) and the Markov property (A3) comprise the law of motion for the state variables. By applying Bellman's principle of optimality recursively (see, for example, Bertsekas [1995]):

$$J_N(V_N, P_N, \mathbf{Z}_N) = [V_N - F(P_N, \mathbf{Z}_N)]^2 \quad (2.10)$$

$$J_i(V_i, P_i, \mathbf{Z}_i) = \min_{\theta^{(i, V_i, P_i, \mathbf{Z}_i)}} \mathbb{E}^\nu \left[J_{i+1}(V_{i+1}, P_{i+1}, \mathbf{Z}_{i+1}) \mid V_i, P_i, \mathbf{Z}_i \right] \quad (2.11)$$

$i = 0, \dots, N-1$

the optimal-replication strategy $\theta^*(i, V_i, P_i, \mathbf{Z}_i)$ can be characterized and computed. In particular, we have:¹²

Theorem 1 *Under Assumptions (A1)–(A4) and (2.3), the solution of the dynamic replication problem (2.5) for $\mathcal{T} = \{t_0, t_1, \dots, t_N\}$ is characterized by the following:*

(a) *The value function $J_i(V_i, P_i, \mathbf{Z}_i)$ is quadratic in V_i , i.e., there are functions $a_i(P_i, \mathbf{Z}_i)$, $b_i(P_i, \mathbf{Z}_i)$, and $c_i(P_i, \mathbf{Z}_i)$ such that*

$$J_i(V_i, P_i, \mathbf{Z}_i) = a_i(P_i, \mathbf{Z}_i) \cdot [V_i - b_i(P_i, \mathbf{Z}_i)]^2 + c_i(P_i, \mathbf{Z}_i), \quad i = 0, \dots, N. \quad (2.12)$$

(b) *The optimal control $\theta^*(i, V_i, P_i, \mathbf{Z}_i)$ is linear in V_i , i.e.,*

$$\theta^*(i, V_i, P_i, \mathbf{Z}_i) = p_i(P_i, \mathbf{Z}_i) - V_i q_i(P_i, \mathbf{Z}_i) \quad (2.13)$$

¹²Proofs are relegated to the Appendix.

(c) The functions $a_i(\cdot)$, $b_i(\cdot)$, $c_i(\cdot)$, $p_i(\cdot)$, and $q_i(\cdot)$, are defined recursively as

$$a_N(P_N, \mathbf{Z}_N) = 1 \quad (2.14)$$

$$b_N(P_N, \mathbf{Z}_N) = F(P_N, \mathbf{Z}_N) \quad (2.15)$$

$$c_i(P_N, \mathbf{Z}_N) = 0 \quad (2.16)$$

and for $i = N-1, \dots, 0$

$$p_i(P_i, \mathbf{Z}_i) = \frac{\mathbb{E}^\nu[a_{i+1}(P_{i+1}, \mathbf{Z}_{i+1}) \cdot b_{i+1}(P_{i+1}, \mathbf{Z}_{i+1}) \cdot (P_{i+1} - P_i) | P_i, \mathbf{Z}_i]}{\mathbb{E}^\nu[a_{i+1}(P_{i+1}, \mathbf{Z}_{i+1}) \cdot (P_{i+1} - P_i)^2 | P_i, \mathbf{Z}_i]} \quad (2.17)$$

$$q_i(P_i, \mathbf{Z}_i) = \frac{\mathbb{E}^\nu[a_{i+1}(P_{i+1}, \mathbf{Z}_{i+1}) \cdot (P_{i+1} - P_i) | P_i, \mathbf{Z}_i]}{\mathbb{E}^\nu[a_{i+1}(P_{i+1}, \mathbf{Z}_{i+1}) \cdot (P_{i+1} - P_i)^2 | P_i, \mathbf{Z}_i]} \quad (2.18)$$

$$a_i(P_i, \mathbf{Z}_i) = \mathbb{E}^\nu[a_{i+1}(P_{i+1}, \mathbf{Z}_{i+1}) \cdot (1 - q_i(P_i, \mathbf{Z}_i)(P_{i+1} - P_i))^2 | P_i, \mathbf{Z}_i] \quad (2.19)$$

$$b_i(P_i, \mathbf{Z}_i) = \frac{1}{a_i(P_i, \mathbf{Z}_i)} \mathbb{E}^\nu[a_{i+1}(P_{i+1}, \mathbf{Z}_{i+1}) \cdot (b_{i+1}(P_{i+1}, \mathbf{Z}_{i+1}) - p_i(P_i, \mathbf{Z}_i)(P_{i+1} - P_i)) \cdot (1 - q_i(P_i, \mathbf{Z}_i)(P_{i+1} - P_i)) | P_i, \mathbf{Z}_i] \quad (2.20)$$

$$c_i(P_i, \mathbf{Z}_i) = \mathbb{E}^\nu[c_{i+1}(P_{i+1}, \mathbf{Z}_{i+1}) | P_i, \mathbf{Z}_i] + \mathbb{E}^\nu[a_{i+1}(P_{i+1}, \mathbf{Z}_{i+1}) \cdot (b_{i+1}(P_{i+1}, \mathbf{Z}_{i+1}) - p_i(P_i, \mathbf{Z}_i)(P_{i+1} - P_i))^2 | P_i, \mathbf{Z}_i] - a_i(P_i, \mathbf{Z}_i) \cdot b_i(P_i, \mathbf{Z}_i)^2 \quad (2.21)$$

(d) Under the optimal-replication strategy θ^* , the minimum replication error as a function of the initial wealth V_0 is

$$J_0(V_0, P_0, \mathbf{Z}_0) = a_0(P_0, \mathbf{Z}_0) \cdot [V_0 - b_0(P_0, \mathbf{Z}_0)]^2 + c_0(P_0, \mathbf{Z}_0) . \quad (2.22)$$

hence the initial wealth that minimizes the replication error is $V_0^* = b_0(P_0, \mathbf{Z}_0)$, the least-cost optimal-replication strategy is the $\{\theta^*(i, V_i, P_i, \mathbf{Z}_i)\}$ that corresponds to this initial wealth, and the minimum replication error over all V_0 is:¹³

$$\epsilon^* = \sqrt{c_0(P_0, \mathbf{Z}_0)} . \quad (2.23)$$

¹³It is simple to show by induction that $a_i(P_i, \mathbf{Z}_i) > 0$ and $c_i(P_i, \mathbf{Z}_i) \geq 0$.

2.3 ϵ -Arbitrage in Continuous Time

For the continuous-time case $\mathcal{T} = [0, T]$, let $[P_t \ \mathbf{Z}_t]'$ follow a vector Markov diffusion process

$$dP_t = \mu_0(t, P_t, \mathbf{Z}_t)P_t dt + \sigma_0(t, P_t, \mathbf{Z}_t)P_t dW_{0t} \quad (2.24)$$

$$dZ_{jt} = \mu_j(t, P_t, \mathbf{Z}_t)Z_{jt} dt + \sigma_j(t, P_t, \mathbf{Z}_t)Z_{jt}dW_{jt}, \quad j = 1, \dots, J \quad (2.25)$$

where W_{jt} , $j = 0, \dots, J$ are Wiener processes with mutual variation

$$dW_{jt}dW_{kt} = \rho_{jk}(t, P_t, \mathbf{Z}_t) dt .$$

The continuous-time counterpart of the Bellman recursion is the Hamilton-Jacobi-Bellman equation (see, for example, Fleming and Rishel [1975]), and this yields the following:

Theorem 2 *Under Assumptions (A1)–(A4) and (2.4), the solution of the dynamic replication problem (2.5) for $\mathcal{T} = [0, T]$ is characterized by the following:*

- (a) *The value function $J(t, V_t, P_t, \mathbf{Z}_t)$ is quadratic in V_t , i.e., there are functions $a(t, P_t, \mathbf{Z}_t)$, $b(t, P_t, \mathbf{Z}_t)$, and $c(t, P_t, \mathbf{Z}_t)$ such that*

$$J(t, V_t, P_t, \mathbf{Z}_t) = a(t, P_t, \mathbf{Z}_t) \cdot [V_t - b(t, P_t, \mathbf{Z}_t)]^2 + c(t, P_t, \mathbf{Z}_t), \quad 0 \leq t \leq T . \quad (2.26)$$

- (b) *For $t \in [0, T]$ the functions $a(t, P_t, \mathbf{Z}_t)$, $b(t, P_t, \mathbf{Z}_t)$, and $c(t, P_t, \mathbf{Z}_t)$ satisfy the following system of partial differential equations:¹⁴*

$$\begin{aligned} \frac{\partial a}{\partial t} + \sum_{j=0}^J \mu_j Z_j \frac{\partial a}{\partial Z_j} + \frac{1}{2} \sum_{i,j=0}^J \sigma_i \sigma_j Z_i Z_j \rho_{ij} \frac{\partial^2 a}{\partial Z_i \partial Z_j} = \\ \left(\frac{\mu_0}{\sigma_0}\right)^2 a + 2 \sum_{j=0}^J \frac{\sigma_j}{\sigma_0} \mu_0 \rho_{0j} Z_j \frac{\partial a}{\partial Z_j} + \\ \frac{1}{a} \sum_{i,j=0}^J \sigma_i \sigma_j Z_i Z_j \rho_{0i} \rho_{0j} \frac{\partial a}{\partial Z_i} \frac{\partial a}{\partial Z_j} \end{aligned} \quad (2.27)$$

¹⁴We omit the arguments of $a(\cdot)$, $b(\cdot)$, and $c(\cdot)$ in (2.27)–(2.29) to economize on notation.

$$\begin{aligned} \frac{\partial b}{\partial t} + \sum_{j=0}^J \mu_j Z_j \frac{\partial b}{\partial Z_j} + \frac{1}{2} \sum_{i,j=0}^J \sigma_i \sigma_j Z_i Z_j \rho_{ij} \frac{\partial^2 b}{\partial Z_i \partial Z_j} = \\ \sum_{j=0}^J \frac{\partial b}{\partial Z_j} Z_j \left(\frac{\sigma_j}{\sigma_0} \mu_0 \rho_{0j} - \frac{1}{2} \sum_{i=0}^J \frac{\sigma_i \sigma_j Z_i}{a} (\rho_{0i} \rho_{0j} - \rho_{ij}) \frac{\partial a}{\partial Z_i} \right) \end{aligned} \quad (2.28)$$

$$\begin{aligned} \frac{\partial c}{\partial t} + \sum_{j=0}^J \mu_j Z_j \frac{\partial c}{\partial Z_j} + \frac{1}{2} \sum_{i,j=0}^J \sigma_i \sigma_j Z_i Z_j \rho_{ij} \frac{\partial^2 c}{\partial Z_i \partial Z_j} = \\ a \sum_{i,j=0}^J \sigma_i \sigma_j Z_i Z_j \frac{\partial b}{\partial Z_i} \frac{\partial b}{\partial Z_j} (\rho_{0i} \rho_{0j} - \rho_{ij}) . \end{aligned} \quad (2.29)$$

with boundary conditions:

$$a(T, P_T, \mathbf{Z}_T) = 1 \quad , \quad b(T, P_T, \mathbf{Z}_T) = F(P_T, \mathbf{Z}_T) \quad , \quad c(T, P_T, \mathbf{Z}_T) = 0 \quad (2.30)$$

where Z_i denotes the i -th component of \mathbf{Z}_t and $Z_0 \equiv P_t$.

(c) The optimal control $\theta^*(t, V_t, P_t, \mathbf{Z}_t)$ is linear in V_t and is given by:

$$\begin{aligned} \theta^*(t, V_t, P_t, \mathbf{Z}_t) = \sum_{j=0}^J \frac{\sigma_j Z_j}{\sigma_0 Z_0} \rho_{0j} \frac{\partial b}{\partial Z_j} - \frac{V_t - b}{a} \sum_{j=0}^J \frac{\sigma_j Z_j}{\sigma_0 Z_0} \rho_{0j} \frac{\partial a}{\partial Z_j} - \\ \frac{(V_t - b) \mu_0}{\sigma_0^2 Z_0} . \end{aligned} \quad (2.31)$$

(d) Under the optimal-replication strategy θ^* , the minimum replication error as a function of the initial wealth V_0 is

$$J(0, V_0, P_0, \mathbf{Z}_0) = a(0, P_0, \mathbf{Z}_0) [V_0 - b(0, P_0, \mathbf{Z}_0)]^2 + c(0, P_0, \mathbf{Z}_0) \quad (2.32)$$

hence the initial wealth that minimizes the replication error is $V_0^* = b(0, P_0, \mathbf{Z}_0)$, the least-cost optimal-replication strategy is the $\{\theta^*(t, V_t, P_t, \mathbf{Z}_t)\}$ that corresponds to this initial wealth, and the minimum replication error over all V_0 is:¹⁵

$$\epsilon^* = \sqrt{c(0, P_0, \mathbf{Z}_0)} . \quad (2.33)$$

¹⁵It can be shown that $a(t, P_t, \mathbf{Z}_t) > 0$ and $c(t, P_t, \mathbf{Z}_t) \geq 0$.

2.4 Interpreting ϵ^* and V_0^*

Theorems 1 and 2 show that the dynamic replication problem (2.5) can be solved for a mean-squared-error measure of replication error under Markov state dynamics. In particular, the optimal-replication strategy $\theta^*(\cdot)$ is a dynamic trading strategy that yields the minimum mean-squared replication error $\epsilon(V_0)$ for an initial wealth V_0 . The fact that $\epsilon(V_0)$ depends on V_0 should come as no surprise, and the fact that $\epsilon(V_0)$ is quadratic in V_0 emphasizes the fact that delta-hedging strategies can be under- or over-capitalized, i.e., there exists a unique V_0^* that minimizes the mean-squared replication error. One attractive feature of our approach is the ability to quantify the impact of capitalization V_0 on the replication error $\epsilon(V_0)$.

2.4.1 V_0^* Is Not a Price

In this sense, V_0^* may be viewed as the minimum production-cost of replicating the payoff $F(P_T, \mathbf{Z}_T)$ as closely as possible, to within ϵ^* . However, because we have assumed that markets are dynamically incomplete (otherwise ϵ^* is 0 and perfect replication is possible), V_0^* cannot be interpreted as the price of a derivative security with payoff $F(P_T, \mathbf{Z}_T)$ unless additional economic structure is imposed. In particular, in dynamically incomplete markets derivatives cannot be priced by arbitrage considerations alone—we must resort to an equilibrium model in which the prices of all traded assets are determined by supply and demand.

To see why V_0^* cannot be interpreted as a price, observe that two investors with different risk preferences may value $F(P_T, \mathbf{Z}_T)$ quite differently, and will therefore place different valuations on the replication error ϵ^* . While both investors may agree that V_0^* is the minimum cost for the optimal-replication strategy $\theta^*(\cdot)$, they may differ in their willingness to pay such a cost for achieving the replication error ϵ^* .¹⁶ Moreover, some investors' preferences may not be consistent with a symmetric loss function, e.g., they may value negative replication errors quite differently than positive replication errors.

More to the point, an asset's price is the outcome of a market equilibrium in which investors' preferences, budget dynamics, and information structure interact through the imposition of market-clearing conditions, i.e., supply equals demand. In contrast, V_0^* is the

¹⁶See Duffie and Jackson (1990) and Duffie and Richardson (1991) for examples of replication strategies under specific preference assumptions.

solution to a simple dynamic optimization problem that does not typically incorporate any notion of economic equilibrium. However, in Section 3 we modify the dynamic optimization problem to account for such equilibrium considerations, and the V_0^* that solves this modified optimal-replication problem does correspond to the equilibrium price of the derivative security (see Theorem 3).

2.4.2 Why Mean-Squared Error?

In fact, there are many possible loss functions, each giving rise to a different set of dynamic replication strategies, hence a natural question to ask in interpreting Theorems 1 and 2 is why use mean-squared error?

The first reason is, of course, tractability. We showed in Sections 2.2 and 2.3 that the dynamic replication problem can be solved via stochastic dynamic programming for a mean-squared-error loss function and Markov state dynamics, and that the solution can be implemented as an exact and efficient recursive algorithm. In Sections 4 and 5, we apply this algorithm to a variety of derivative securities in incomplete markets and demonstrate its practical relevance analytically and numerically.

The second reason is that a symmetric loss function is the most natural choice when we have no prior information about whether the derivative to be replicated is being purchased or sold. In such cases, asymmetric loss functions are inappropriate since positive replication errors for a long position become negative replication errors for the short position. Indeed, when a derivatives broker is asked by a client to provide a price quote, the client does not reveal whether he is a buyer or seller until after he receives both bid and offer prices. Therefore, it is in the interest of the broker to provide as “tight” a spread as possible, i.e., to minimize mean-squared error.

Of course, in more structured applications such as Duffie and Jackson (1990) in which investors’ preferences, budget dynamics, and information sets are specified, it is not apparent that mean-squared-error optimal-replication strategies are optimal from a particular investor’s point of view. However, even in these cases, a slight modification of the mean-squared-error loss function yields optimal-replication strategies that have natural economic interpretations. In particular, we show in the next section that by defining mean-squared

error with respect to equivalent martingale measure, the minimum production cost V_0^* associated with this loss function can be interpreted as an equilibrium market price which, by definition, incorporates all aspects of the economic environment in which the derivative security is traded. This is the third and perhaps the most compelling motivation for a mean-squared error loss function.

3 Risk-Neutralized ϵ -Arbitrage Strategies

To relate the minimum production cost V_0^* of the optimal-replication strategy to the market price of a derivative security with payoff $F(P_T, \mathbf{Z}_T)$, in this section we propose a minor but important modification to the dynamic replication problem (2.5) of Section 2. The modification consists of evaluating the mean-squared replication error with respect to an adjusted probability measure ν^* —the *risk-neutralized* or *equivalent martingale* measure—hence the risk-neutralized dynamic replication problem becomes:

$$\min_{\{\theta_i\}} E^{\nu^*} \{ [V_T - F(P_T, \mathbf{Z}_T)]^2 \}. \quad (3.1)$$

Although (3.1) seems virtually identical to (2.5), the implications of using ν^* in place of ν are significant. In Section 3.1, we show that the minimum production cost V_0^* associated with (3.1) does correspond to the equilibrium price of a derivative security with payoff $F(P_T, \mathbf{Z}_T)$, and is not merely a “proxy” for the price. As a consequence, the algorithm of Theorems 1 and 2 provides an explicit optimal dynamic replication strategy that corresponds to the equilibrium price of the derivative security, which complements the standard delta-hedging strategies and generalizes them to an incomplete-markets setting.

Of course, ν^* is not always readily observable and additional structure is needed to infer ν^* from existing market prices—we discuss this issue in Section 3.2.

3.1 Equilibrium Pricing Models and ϵ -Arbitrage

In Section 2.4 we argued that the minimum production cost V_0^* cannot be interpreted as the price of a derivative security with payoff $F(P_T, \mathbf{Z}_T)$ because $E^\nu [V_T - F(P_T, \mathbf{Z}_T)]^2$ does not necessarily reflect an investor’s preferences regarding the replication error. To derive

the equilibrium price of the derivative security, we require additional economic structure, i.e., investors' preferences, budget dynamics, information structure, and the imposition of market-clearing conditions.

Such economic structure is summarized by the equivalent martingale measure ν^* , also known as the *state-price density* or the *risk-neutral density*. Cox and Ross (1976) and Harrison and Kreps (1979) show that under certain regularity conditions, the equilibrium prices of all traded securities must be martingales under this adjusted probability measure, hence the price of any security can be determined simply as the expectation of its payoff, where the expectation is evaluated with respect to ν^* .¹⁷ Therefore, if $H(0, P_0, \mathbf{Z}_0)$ denotes the equilibrium market price of the option at time 0, then

$$H(0, P_0, \mathbf{Z}_0) = E^{\nu^*}[F(P_T, \mathbf{Z}_T)] . \quad (3.2)$$

But observe that an implication of minimizing mean-squared error (3.1) is that:¹⁸

$$E^{\nu^*}[V_T^* - F(P_T, \mathbf{Z}_T)] = 0 \quad (3.3)$$

where V_T^* denotes the terminal value of the portfolio under the optimal replicating strategy $\theta^*(\cdot)$. Since the optimal-replication strategy is self-financing so that there are no cash inflows or outflows during the interval $(0, T)$, it must be the case that $V_0^* = E^{\nu^*}[V_T^*]$. This in turn implies:

$$V_0^* = E^{\nu^*}[V_T^*] = E^{\nu^*}[F(P_T, \mathbf{Z}_T)] = H(0, P_0, \mathbf{Z}_0) . \quad (3.4)$$

Therefore, we have:

Theorem 3 *Under Assumptions (A1)–(A4) and (2.3) or (2.4), the minimum production cost V_0^* corresponding to the risk-neutralized dynamic replication problem (3.1) is the equilibrium price of a derivative security with payoff $F(P_T, \mathbf{Z}_T)$.*

Observe that Theorem 3 does not assume dynamically complete markets, unlike the

¹⁷See Duffie (1996), Duffie and Huang (1985), Huang (1987), Huang and Litzenberger (1988), and Merton (1992) for further details.

¹⁸See, for example, Schweizer (1995).

standard arbitrage-based pricing model of Merton (1973). Therefore, when ν^* is substituted for ν in Theorems 1 and 2, the recursive algorithms outlined in those two theorems yield *risk-neutralized* optimal-replication strategies that generalize the standard Merton (1973) delta-hedging strategies to dynamically incomplete markets. Hereafter, we shall refer to $\theta^*(\cdot)$ under ν^* as a *generalized delta-hedging strategy*.

3.2 How to Obtain ν^*

While Theorem 3 seems to suggest that the generalization of delta-hedging to dynamically incomplete markets is straightforward—substitute ν^* for ν —obtaining ν^* can often be quite a challenge. In a dynamic equilibrium model such as Lucas (1978) and Rubinstein (1976), the equivalent martingale measure is a weighted average of the probability measure ν , where the weighting function is the equilibrium marginal rate of substitution of the representative agent. This dynamic equilibrium interpretation illustrates the enormous information content of ν^* and the enormous information reduction that the equivalent martingale measure affords. Indeed, from a pricing perspective, ν^* is a “sufficient statistic” in the sense that it contains all relevant information about preferences and business conditions for purposes of pricing financial securities.

But as a practical matter, how does one obtain ν^* to make Theorems 1–3 operational? There are at least two possible approaches to this challenge—theoretical and empirical—and we shall describe each of these in the next two sections.

3.2.1 Theoretical Methods

The theoretical approach is to provide sufficient economic structure, i.e., a fully articulated dynamic equilibrium model as in Cox, Ingersoll, and Ross (1985), to yield a unique ν^* . In such an environment, using ν^* leads to a significant simplification of the dynamic replication problem. For example, consider a continuous-time model in which stock prices and the state variables \mathbf{Z}_t are described by the following system of Markov diffusion processes:

$$\begin{aligned} dP_t &= \mu_0(t, P_t, \mathbf{Z}_t)P_t dt + \sigma_0(t, P_t, \mathbf{Z}_t)P_t dW_{0t} \\ dZ_{jt} &= \mu_j(t, P_t, \mathbf{Z}_t)Z_{jt} dt + \sigma_j(t, P_t, \mathbf{Z}_t)Z_{jt} dW_{jt} \quad , \quad j = 1, \dots, J . \end{aligned}$$

Under the risk-neutral measure ν^* , the risk-neutralized drift rate $\mu_0^*(t, P_t, \mathbf{Z}_t)$ becomes:

$$\mu_0^*(t, P_t, \mathbf{Z}_t) = 0$$

while the diffusion coefficients remain unchanged. This simplifies the system of PDE's (2.27)–(2.29) to:

$$a(t, P_t, \mathbf{Z}_t) = 1 \quad (3.5)$$

$$\frac{\partial b(t, P_t, \mathbf{Z}_t)}{\partial t} + \sum_{j=1}^J \mu_j^* Z_j \frac{\partial b(t, P_t, \mathbf{Z}_t)}{\partial Z_j} + \frac{1}{2} \sum_{i,j=0}^J \sigma_i \sigma_j Z_i Z_j \rho_{ij} \frac{\partial^2 b(t, P_t, \mathbf{Z}_t)}{\partial Z_i \partial Z_j} = 0 \quad (3.6)$$

$$\begin{aligned} \frac{\partial c(t, P_t, \mathbf{Z}_t)}{\partial t} + \sum_{j=1}^J \mu_j^* Z_j \frac{\partial c(t, P_t, \mathbf{Z}_t)}{\partial Z_j} + \frac{1}{2} \sum_{i,j=0}^J \sigma_i \sigma_j Z_i Z_j \rho_{ij} \frac{\partial^2 c(t, P_t, \mathbf{Z}_t)}{\partial Z_i \partial Z_j} = \\ \sum_{i,j=0}^J \sigma_i \sigma_j Z_i Z_j \frac{\partial b(t, P_t, \mathbf{Z}_t)}{\partial Z_i} \frac{\partial b(t, P_t, \mathbf{Z}_t)}{\partial Z_j} (\rho_{0i} \rho_{0j} - \rho_{ij}) \end{aligned} \quad (3.7)$$

and the optimal number of shares in the replicating portfolio is given by:¹⁹

$$\theta^*(t, V_t, P_t, \mathbf{Z}_t) = \frac{\partial b(t, P_t, \mathbf{Z}_t)}{\partial Z_0} + \sum_{j=1}^J \frac{\sigma_j Z_j}{\sigma_0 Z_0} \rho_{0j} \frac{\partial b(t, P_t, \mathbf{Z}_t)}{\partial Z_j}. \quad (3.8)$$

Interestingly, the optimal replicating strategy $\theta^*(t, V_t, P_t, \mathbf{Z}_t)$ depends on the equivalent martingale measure ν^* only indirectly, through the option price $H(t, P_t, \mathbf{Z}_t) = b(t, P_t, \mathbf{Z}_t)$. In other words, all information necessary for the construction of the optimal replicating strategy is contained in the option price itself. Moreover, the optimal replicating strategy is independent of the value of the portfolio V_t . These are properties of the delta-hedging strategies of arbitrage-based models such as Merton (1973), and they carry through to the generalized delta-hedging strategies of ϵ -arbitrage models as well.

Indeed, in addition to the term $\partial b(t, P_t, \mathbf{Z}_t)/\partial Z_0$ which is the well-known Black-Scholes hedge ratio, the generalized delta-hedging strategy $\theta^*(t, V_t, P_t, \mathbf{Z}_t)$ contains additional terms

¹⁹(3.8) follows from (4.4) in Schweizer (1992), which was obtained under a set of assumptions different from those adopted in the present paper. The results in Schweizer (1992) do not apply to the case when the objective function is defined using the original probability measure ν , but they are applicable here since the drift rate of the stock price process under the equivalent martingale measure is equal to zero.

of the form $\partial b(t, P_t, \mathbf{Z}_t)/\partial Z_j$, that use changes in the stock price to hedge against changes in the non-traded state variables Z_j . These terms are weighted by the correlation coefficients $(\sigma_j Z_j/\sigma_0 Z_0)\rho_{0j}$, that determine the degree to which such a hedging is effective. To see that (3.8) is a direct generalization of the Black and Scholes (1973) and Merton (1973) delta-hedging formula, observe that it reduces to the Black and Scholes (1973) and Merton (1973) option delta when the state variables Z_j are instantaneously uncorrelated with the stock price.

3.2.2 Empirical Methods

An alternative to developing a fully articulated dynamic equilibrium model is to estimate ν^* from the prices of existing financial securities. This is the approach taken in several recent papers, including Ait-Sahalia and Lo (1996), Derman and Kani (1994), Dumas et al. (1995), Dupire (1994), Hutchinson et al. (1994), Jarrow and Rudd (1982), Longstaff (1992, 1995), Rady (1994), Rubinstein (1994), and Shimko (1993).

For example, Ait-Sahalia and Lo (1996) propose a nonparametric method for estimating ν^* : construct a nonparametric estimator of a call-option pricing formula using market prices, then take the second derivative of the estimated pricing formula with respect to the strike price. Banz and Miller (1978), Breeden and Litzenberger (1978), and Ross (1976) show that this second derivative is ν^* .

Another method for determining ν^* empirically is Rubinstein's (1994) *implied binomial tree*, in which the risk-neutral probabilities $\{\pi_i^*\}$ associated with the binomial terminal stock price P_T are estimated by minimizing the sum of squared deviations between $\{\pi_i^*\}$ and a set of *prior* risk-neutral probabilities $\{\tilde{\pi}_i^*\}$, subject to the restrictions that $\{\pi_i^*\}$ correctly price an existing set of options and the underlying stock [in the sense that the optimal risk-neutral probabilities yield prices that lie within the bid-ask spreads of the options and the stock]. This approach is similar in spirit to Jarrow and Rudd (1982) and Longstaff's (1992, 1995) method of fitting risk-neutral density functions using a four-parameter Edgeworth expansion.²⁰

²⁰However, Rubinstein (1994) points out several important limitations of Longstaff's method when extended to a binomial model, including the possibility of negative probabilities. See, also, Derman and Kani (1994), Shimko (1993), Dupire (1994), and Dumas et al. (1995)].

Once ν^* has been estimated, a numerical implementation of the recursive algorithms in Theorems 1 and 2 can be undertaken. We hope to examine the properties of such procedures in future research.

4 Illustrative Examples

To illustrate the scope and power of the ϵ -arbitrage approach to the dynamic replication problem, we apply the results of Section 2 to four specific cases for the return-generating process: state-independent returns (Section 4.1), geometric Brownian motion (Section 4.2), a jump-diffusion model (Section 4.3), and a stochastic volatility model (4.4).

4.1 State-Independent Returns

Suppose that stock returns are state-independent so that

$$P_i = P_{i-1}(1 + \phi_{i-1}) \quad (4.1)$$

where ϕ_{i-1} is independent of the current stock price and all other state variables. This, together with the Markov assumption (A3) implies that returns are statistically independent (but not necessarily identically distributed) through time. Also, let the payoff of the derivative security $F(P_T)$ depend only on the price of the risky asset at time T .

In this case, there is no need for additional state variables \mathbf{Z}_i and the expressions in Theorem 1 simplify to:

$$a_N = 1 \quad , \quad b_N(P_N) = F(P_N) \quad , \quad c_N(P_N) = 0 \quad (4.2)$$

and for $i = N-1, \dots, 0$,

$$a_i = a_{i+1} \frac{\sigma_i^2}{\sigma_i^2 + \mu_i^2} \quad (4.3)$$

$$b_i(P_i) = E^\nu [b_{i+1}(P_i(1 + \phi_i))|P_i] - \frac{\mu_i}{\sigma_i^2} \text{Cov}^\nu [\phi_i, b_{i+1}(P_i(1 + \phi_i))|P_i] \quad (4.4)$$

$$c_i(P_i) = E^\nu [c_{i+1}(P_i(1 + \phi_i))|P_i] + \frac{a_{i+1}}{\sigma_i^2} \left\{ \sigma_i^2 \text{Var}^\nu [b_{i+1}(P_i(1 + \phi_i))|P_i] - \right.$$

$$\left. \text{Cov}^\nu [\phi_i, b_{i+1}(P_i(1 + \phi_i)) | P_i]^2 \right\} \quad (4.5)$$

$$p_i(P_i) = \frac{\text{E}^\nu [\phi_i b_{i+1}(P_i(1 + \phi_i)) | P_i]}{(\sigma_i^2 + \mu_i^2) P_i} \quad (4.6)$$

$$q_i(P_i) = \frac{\mu_i}{(\sigma_i^2 + \mu_i^2) P_i} \quad (4.7)$$

where $\mu_i = \text{E}^\nu[\phi_i]$ and $\sigma_i^2 = \text{Var}^\nu[\phi_i]$.

4.2 Geometric Brownian Motion

Let the stock price process follow the geometric Brownian motion of Black and Scholes (1973) and Merton (1973). We show that the ϵ -arbitrage approach yields the Black-Scholes/Merton results in the limit of continuous time, but in discrete time there are important differences between the optimal-replication strategy of Theorem 1 and the standard Black-Scholes/Merton delta-hedging strategy.

For notational convenience, let all discrete time intervals $[t_i, t_{i+1})$ be of equal length $t_{i+1} - t_i = \Delta t$. The assumption of geometric Brownian motion then implies:

$$P_{i+1} = P_i \cdot (1 + \phi_i) \quad (4.8)$$

$$\log(1 + \phi_i) = \left(\mu - \frac{\sigma^2}{2}\right)\Delta t + \sigma\sqrt{\Delta t}z_i \quad (4.9)$$

$$z_i \sim \mathcal{N}(0, 1) . \quad (4.10)$$

Recall that for $\Delta t \ll 1$ (a large number of time increments on $[0, T]$), the following approximation holds (see, for example, Merton [1992, Chapter 3]):

$$\phi_i \sim \mathcal{N}(\mu\Delta t, \sigma^2\Delta t) + O(\Delta t^{3/2})$$

This, and Taylor's theorem, imply the following approximations for the recursive relations (4.3)–(4.5) of Section 4.1:

$$\text{Var}^\nu [b_{i+1}(P_i(1 + \phi_i)) | P_i] = b'_{i+1}(P_i)\sigma^2 P_i^2 \Delta t + O(\Delta t^2)$$

$$\begin{aligned}
\text{Cov}^\nu[\phi_i, b_{i+1}(P_i(1 + \phi_i))|P_i] &= b'_{i+1}(P_i)\sigma^2 P_i \Delta t + O(\Delta t^2) \\
\mathbb{E}^\nu[b_{i+1}(P_i(1 + \phi_i))|P_i] &= b_{i+1}(P_i) + b'_{i+1}(P_{i-1})\mu P_i \Delta t + b''_{i+1}(P_i)\frac{\sigma^2}{2}P_i^2 \Delta t + O(\Delta t^2) \\
\mathbb{E}^\nu[c_{i+1}(P_i(1 + \phi_i))|P_i] &= c_i(P_i) + c'_{i+1}(P_{i-1})\mu P_i \Delta t + c''_{i+1}(P_i)\frac{\sigma^2}{2}P_i^2 \Delta t + O(\Delta t^2).
\end{aligned}$$

We can then rewrite (4.4)–(4.5) as

$$\begin{aligned}
b_i(P_i) &= b_{i+1}(P_i) + b''_{i+1}(P_i)\frac{\sigma^2}{2}P_i^2 \Delta t + O(\Delta t^2) \\
c_i(P_i) &= c_{i+1}(P_i) + c'_{i+1}(P_{i-1})\mu P_i \Delta t + c''_{i+1}(P_i)\frac{\sigma^2}{2}P_i^2 \Delta t + O(\Delta t^2)
\end{aligned}$$

and conclude that the system (4.4)–(4.5) approximates the following system of PDE's

$$\frac{\partial b(t, P)}{\partial t} = -\frac{\sigma^2 P^2}{2} \frac{\partial^2 b(t, P)}{\partial P^2} \tag{4.11}$$

$$\frac{\partial c(t, P)}{\partial t} = -\mu P \frac{\partial c(t, P)}{\partial P} - \frac{\sigma^2 P^2}{2} \frac{\partial^2 c(t, P)}{\partial P^2} \tag{4.12}$$

up to $O(\Delta t)$ terms. But (4.11) is the Black and Scholes (1973) PDE, hence we see that in the limit of continuous trading, i.e., as $N \rightarrow \infty$ and $\Delta t \rightarrow 0$ for a fixed $T \equiv N\Delta t$, the discrete-time optimal-replication strategy of Theorem 1 characterizes the Black and Scholes (1973) and Merton (1973) models.

Moreover, the equation for $c(t, P)$, (4.12) is homogeneous, hence $c(t, P) \equiv 0$ due to the boundary condition $c(T, \cdot) = 0$. This is consistent with the fact that in the Black-Scholes case it is possible to replicate the option exactly, so that the replication error vanishes in the continuous-time limit.

The continuous-time limit of the optimal-replication strategy $\theta^*(\cdot)$ is given by:

$$\theta^*(t, V_t, P_t) = \frac{\partial b(t, P_t)}{\partial P_t} - \frac{\mu}{\sigma^2 P_t} [V_t - b(t, P_t)] . \tag{4.13}$$

At time $t = 0$, and for the minimum production-cost initial wealth V_0^* , this becomes

$$\theta^*(0, V_0^*, P_0) = \frac{\partial b(0, P_0)}{\partial P_0}$$

since $V_0^* = b(0, P_0)$. Since exact replication is possible in this case, the value of the replicating portfolio is always equal to $b(t, P_t)$, for every realization of the stock price process, i.e.,

$$V_t = b(t, P_t)$$

for all $t \in [0, T]$, which implies that

$$\theta^*(t, V_t, P_t) = \frac{\partial b(t, P_t)}{\partial P_t}. \quad (4.14)$$

As expected, for every realization of the stock price process the optimal replication strategy coincides with the delta-hedging strategy given by the Black-Scholes hedge ratio. However, note that the functional form of (4.13) is different from the Black-Scholes hedging formula—the optimal-replication strategy depends explicitly on its value V_t .

4.3 Jump-Diffusion Models

In this section, we apply results of Section 2 to the replication and pricing of options on a stock with mixed jump-diffusion price dynamics. As before, we assume that all time intervals $t_{i+1} - t_i = \Delta t$ are regularly spaced. Following Merton (1976), we assume the following model for the stock price process:

$$P_{i+1} = P_i(1 + \phi_i) \quad (4.15)$$

$$\log(1 + \phi_i) = \left(\mu - \lambda k - \frac{\sigma^2}{2}\right)\Delta t + \sigma\sqrt{\Delta t}z_i + \sum_{j=0}^{n_i} \log Y_j \quad (4.16)$$

$$z_i \sim \mathcal{N}(0, 1) \quad (4.17)$$

$$k = E^\nu[Y_j - 1] \quad (4.18)$$

$$\text{Prob}(n_i = m) = e^{-\lambda\Delta t} \frac{(\lambda\Delta t)^m}{m!} \quad (4.19)$$

where the jump magnitudes $\{Y_j\}$ are independently and identically distributed random variables and jump arrivals follow a Poisson process with constant arrival rate λ .

We consider two types of: jumps of deterministic magnitude and jumps with lognormally

distributed jump magnitudes. In the first case:

$$Y_j = 1 + \delta . \tag{4.20}$$

If we set $\sigma = 0$ in (4.15), this model corresponds to the continuous-time jump process considered by Cox and Ross (1976). In the second case:

$$\log Y_i \sim \mathcal{N}(0, \delta^2) . \tag{4.21}$$

There are two methods of calculating the optimal-replication strategy for the mixed jump-diffusion model. One method is to begin with the solutions of the dynamic programming problem given in Sections 2.2 and 2.3, derive a limiting system of partial differential equations as in Section 4.2, and solve this system numerically, using one of the standard finite difference schemes.

The second method is to implement the solution of the dynamic programming problem directly, without the intermediate step of reducing it to a system of PDE's.

The advantage of the second method is that it treats a variety of problems in a uniform fashion, the only problem-dependent part of the approach being the specification of the stochastic process. On the other hand, the first approach yields a representation of the solution as a system of PDE's, which can often provide some information about the qualitative properties of the solution even before a numerical solution is obtained.

With these considerations in mind, we shall derive a limiting system of PDE's for the deterministic-jump-magnitude specification (4.20) and use it to find conditions on the parameters of the stochastic process which allow exact replication of the option's payoff, or, equivalently, arbitrage pricing. For the lognormal-jump-magnitude specification (4.21), we shall obtain numerical solutions directly from the dynamic programming algorithm of Theorem 1.

4.3.1 The Continuous-Time Limit

To derive the continuous-time limit of (4.3)–(4.5) we follow the same procedure as in Section 4.2 which yields the following system of PDE's:

$$\begin{aligned} \frac{\partial b(t, P)}{\partial t} = & -\lambda [b(t, P(1 + \delta)) - b(t, P)] + \lambda \delta P \frac{\partial b(t, P)}{\partial P} - \frac{\sigma^2 P^2}{2} \frac{\partial^2 b(t, P)}{\partial P^2} - \\ & \frac{\mu \lambda \delta}{\lambda \delta^2 + \sigma^2} \left[\delta P \frac{\partial b(t, P)}{\partial P} - [b(t, P(1 + \delta)) - b(t, P)] \right] \end{aligned} \quad (4.22)$$

$$\begin{aligned} \frac{\partial c(t, P)}{\partial t} = & -\lambda [c(t, P(1 + \delta)) - c(t, P)] - (\mu - \lambda \delta) P \frac{\partial c(t, P)}{\partial P} - \frac{\sigma^2 P^2}{2} \frac{\partial^2 c(t, P)}{\partial P^2} - \\ & a(t) \frac{\lambda \sigma^2}{\lambda \delta^2 + \sigma^2} \left[\delta P \frac{\partial b(t, P)}{\partial P} - [b(t, P(1 + \delta)) - b(t, P)] \right]^2 \end{aligned} \quad (4.23)$$

$$\frac{da(t)}{dt} = \frac{\mu^2}{\lambda \delta^2 + \sigma^2} a(t) \quad (4.24)$$

with boundary conditions:

$$a(T) = 1 \quad (4.25)$$

$$c(T, P) = 0 \quad (4.26)$$

$$b(T, P) = F(P) . \quad (4.27)$$

We can use the boundary conditions to solve (4.24):

$$a(t) = \exp \left[\frac{\mu^2}{\lambda \delta^2 + \sigma^2} (t - T) \right] . \quad (4.28)$$

The optimal-replication strategy is given by:

$$\begin{aligned} \theta^*(t, V_t, P_t) = & \frac{\partial b(t, P_t)}{\partial P_t} - \frac{\mu}{(\lambda \delta^2 + \sigma^2) P_t} [V_t - b(t, P_t)] - \frac{\lambda \delta^2}{\lambda \delta^2 + \sigma^2} \frac{\partial b(t, P_t)}{\partial P_t} + \\ & \frac{\lambda \delta}{(\lambda \delta^2 + \sigma^2) P_t} [b(t, P_t(1 + \delta)) - b(t, P_t)] . \end{aligned} \quad (4.29)$$

For exact replication to be possible, $c(t, P) \equiv 0$ must be a solution of (4.23). This implies

that (4.23) is homogeneous, i.e.,

$$\frac{\lambda\sigma^2}{\lambda\delta^2 + \sigma^2} \left\{ \delta P \frac{\partial b(t, P)}{\partial P} - [b(t, P(1 + \delta)) - b(t, P)] \right\}^2 = 0 \quad (4.30)$$

for all $b(t, P)$ satisfying (4.22), which is equivalent to

$$\lambda\delta\sigma^2 = 0. \quad (4.31)$$

Condition (4.31) is satisfied if at least one of the following is true:

- Jumps occur with zero probability.
- Jumps have zero magnitude.
- The diffusion coefficient is equal to zero, i.e., stock price follows a pure jump process.

But these are precisely the conditions for the arbitrage-pricing of options on mixed jump-diffusion assets, e.g., Merton (1976).

4.3.2 Perturbation Analysis with Small Jump Amplitudes

Consider the behavior of $b(t, P)$ and $c(t, P)$ when the jump magnitude is small, i.e., $\delta \ll 1$. In this case the market is “almost complete” and solution of the option replication problem is obtained as a perturbation of the complete-markets solution of Black and Scholes (1973) and Merton (1973). In particular, we treat the amplitude of stock price jumps as a small parameter and look for a solution of (4.22)–(4.27) of the following form:

$$b(t, P) = b_0(t, P) + \delta b_1(t, P) + \delta^2 b_2(t, P) + \dots \quad (4.32)$$

$$c(t, P) = c_0(t, P) + \delta c_1(t, P) + \delta^2 c_2(t, P) + \delta^3 c_3(t, P) + \delta^4 c_4(t, P) + \dots \quad (4.33)$$

After substituting this expansion into (4.23)–(4.27), it is apparent that the functions $b_0(t, P)$, $b_2(t, P)$, and $c_4(t, P)$ must satisfy the following system of partial differential equations:

$$\frac{\partial b_0(t, P)}{\partial t} = - \frac{\sigma^2 P^2}{2} \frac{\partial^2 b_0(t, P)}{\partial P^2} \quad (4.34)$$

$$\frac{\partial b_2(t, P)}{\partial t} = - \frac{\lambda P^2}{2} \frac{\partial^2 b_0(t, P)}{\partial P^2} \quad (4.35)$$

$$\frac{\partial c_4(t, P)}{\partial t} = - \mu P \frac{\partial c_4(t, P)}{\partial P} - \frac{\sigma^2 P^2}{2} \frac{\partial^2 c_4(t, P)}{\partial P^2} - a(t) \frac{\lambda P^4}{4} \left(\frac{\partial^2 b_0(t, P)}{\partial P^2} \right)^2 \quad (4.36)$$

with boundary conditions:

$$b_0(T, P) = F(P) \quad (4.37)$$

$$b_2(T, P) = 0 \quad (4.38)$$

$$c_4(T, P) = 0 \quad (4.39)$$

and

$$b_1 = c_1 = c_2 = c_3 = 0 .$$

System (4.34)–(4.39) can be solved to yield:

$$b(t, P) = b_0(t, P) + \frac{\lambda \delta^2}{\sigma^2} [b_0(t, P) - F(P)] + O(\delta^3) \quad (4.40)$$

where $b_0(t, P)$ is the option price in the absence of a jump component, i.e., the Black-Scholes formula in the case of put and call options. Observe that for an option with a convex payoff function $b_0(t, P) \geq F(P)$, which implies that $b(t, P) \geq b_0(t, P)$, i.e., the addition of a small jump component to geometric Brownian motion increases the price of the option. This qualitative behavior of the option price is consistent with the results in Merton (1976) which were obtained with equilibrium arguments.

The optimal-replication strategy (4.29) is given by:

$$\begin{aligned} \theta^*(t, V_t, P_t) = & \frac{\partial b_0(t, P_t)}{\partial P_t} + \frac{\mu}{\sigma^2 P_t} [b_0(t, P_t) - V_t] + \\ & \frac{\lambda \delta^2}{\sigma^2} \left[\frac{\partial b_0(t, P_t)}{\partial P_t} - \frac{\partial F(P_t)}{\partial P_t} + V_t - F(P_t) \right] + O(\delta^3) . \end{aligned} \quad (4.41)$$

and the corresponding replication error is:

$$c(t, P) = \delta^4 c_4(t, P) + O(\delta^6) = O(\delta^4) \quad (4.42)$$

where $c_4(t, P)$ solves (4.36) and (4.39).

Equations (4.40) and (4.41) provide closed-form expressions for the replication cost and the optimal-replication strategy when the amplitude of jumps is small, i.e., when markets are almost complete, and (4.42) describes the dependence of the replication error on the jump magnitude.

4.4 Stochastic Volatility

Let stock prices follow a diffusion process with stochastic volatility as in Hull and White (1987) and Wiggins (1987):

$$dP_t = \mu P_t dt + \sigma_t P_t dW_{Pt} \quad (4.43)$$

$$d\sigma_t = g(\sigma_t) dt + \kappa \sigma_t dW_{\sigma t} \quad (4.44)$$

where W_{Pt} and $W_{\sigma t}$ are Brownian motions with mutual variation $dW_{Pt}dW_{\sigma t} = \rho dt$.

4.4.1 The Continuous-Time Solution

Although applying the results of Section 2 to (4.43)–(4.44) is conceptually straightforward, the algebraic manipulations are quite involved in this case. A simpler alternative to deriving a system of PDE's as the continuous-time limit of the solution in Theorem 1, we formulate the problem in continuous time at the outset and solve it using continuous-time stochastic control methods. This approach simplifies the calculations considerably.

Specifically, the pair of stochastic processes (P_t, σ_t) satisfies assumptions of Section 2.3, therefore results of this section can be used to derive the optimal-replication strategy, the minimum production-cost of optimal replication, and the replication error. In particular, the application of the results of Section 2.3 to (4.43)–(4.44) yields the following system of

PDE's:

$$\begin{aligned} \frac{\partial a(t, \sigma)}{\partial t} = & \frac{\mu^2}{\sigma^2} a(t, \sigma) - (g(\sigma) + 2\rho\kappa\mu) \frac{\partial a(t, \sigma)}{\partial \sigma} + \\ & \frac{1}{a(t, \sigma)} \left(\rho\kappa\sigma \frac{\partial a(t, \sigma)}{\partial \sigma} \right)^2 - \frac{1}{2} \kappa^2 \sigma^2 \frac{\partial^2 a(t, \sigma)}{\partial \sigma^2} \end{aligned} \quad (4.45)$$

$$\begin{aligned} \frac{\partial b(t, P, \sigma)}{\partial t} = & - \left(g(\sigma) - \rho\mu\kappa \right) \frac{\partial b(t, P, \sigma)}{\partial \sigma} - \frac{\kappa^2 \sigma^2}{2} \frac{\partial^2 b(t, P, \sigma)}{\partial \sigma^2} - \\ & \frac{\sigma^2 P^2}{2} \frac{\partial^2 b(t, P, \sigma)}{\partial P^2} - \rho\kappa\sigma^2 P \frac{\partial^2 b(t, P, \sigma)}{\partial P \partial \sigma} - \\ & (1 - \rho^2) \frac{\kappa^2 \sigma^2}{a(t, \sigma)} \frac{\partial a(t, \sigma)}{\partial \sigma} \frac{\partial b(t, P, \sigma)}{\partial \sigma} \end{aligned} \quad (4.46)$$

$$\begin{aligned} \frac{\partial c(t, P, \sigma)}{\partial t} = & - g(\sigma) \frac{\partial c(t, P, \sigma)}{\partial \sigma} - \mu P \frac{\partial c(t, P, \sigma)}{\partial P} - \frac{\kappa^2 \sigma^2}{2} \frac{\partial^2 c(t, P, \sigma)}{\partial \sigma^2} - \\ & \rho\kappa\sigma^2 P \frac{\partial^2 c(t, P, \sigma)}{\partial P \partial \sigma} - \frac{\sigma^2 P^2}{2} \frac{\partial^2 c(t, P, \sigma)}{\partial P^2} + \\ & a(t, \sigma) \kappa^2 \sigma^2 (\rho^2 - 1) \left(\frac{\partial b(t, P, \sigma)}{\partial \sigma} \right)^2 \end{aligned} \quad (4.47)$$

with boundary conditions:

$$a(T, \sigma) = 1 \quad , \quad b(T, P, \sigma) = F(P, \sigma) \quad , \quad c(T, P, \sigma_T) = 0 \quad .$$

The optimal-replication strategy is given by:

$$\begin{aligned} \theta^*(t, V_t, P_t, \sigma_t) = & \frac{\partial b(t, P_t, \sigma_t)}{\partial P_t} + \frac{\rho\kappa}{P_t} \frac{\partial b(t, P_t, \sigma_t)}{\partial \sigma_t} - \frac{V_t - b(t, P_t, \sigma_t)}{a(t, \sigma_t)} \times \\ & \frac{\rho\kappa}{P_t} \frac{\partial a(t, \sigma_t)}{\partial \sigma_t} - \frac{\mu}{\sigma_t^2 P_t} [V_t - b(t, P_t, \sigma_t)] \quad . \end{aligned} \quad (4.48)$$

Exact replication is possible when the following equation is satisfied:

$$\kappa^2 (\rho^2 - 1) = 0 \quad .$$

and this corresponds to the following special cases:

- Volatility is a deterministic function of time.

- The Brownian motions driving stock prices and volatility are perfectly correlated.

Both of these conditions yield well-known special cases where arbitrage-pricing is possible (see, for example, Geske [1979] and Rubinstein [1983]). If we set $\kappa = g(\sigma) = 0$, (4.46) reduces to the Black and Scholes (1973) PDE.

5 Numerical Analysis

The essence of the ϵ -arbitrage approach to the dynamic replication problem is the recognition that although perfect replication may not be possible in some situations, the optimal-replication strategy of Theorem 1 may come very close. How close is, of course, an empirical matter hence in this section we present several numerical examples that complement the theoretical calculations of Section 4.

In Section 5.1 we describe our numerical procedure and implement it for following examples: geometric Brownian motion (Section 5.2), a mixed jump-diffusion model with a lognormal jump magnitude (Section 5.3), and a stochastic volatility model (Section 5.4). In addition, we also implement our numerical solution algorithm for a stochastic volatility model under an equivalent martingale measure in Section 5.5. Finally, in Section 5.6 we apply our algorithm to the path-dependent option to “sell at the high”.

5.1 The Numerical Procedure

To implement the solution (2.17)–(2.21) of the dynamic replication problem numerically, we begin by representing the functions $a_i(P, \mathbf{Z})$, $b_i(P, \mathbf{Z})$, and $c_i(P, \mathbf{Z})$ by their values over a spatial grid $\{(P^j, \mathbf{Z}^k) : j = 1, \dots, J, k = 1, \dots, K\}$. For any given (P, \mathbf{Z}) , values $a_i(P, \mathbf{Z})$, $b_i(P, \mathbf{Z})$, and $c_i(P, \mathbf{Z})$ are obtained from $a_i(P^j, \mathbf{Z}^k)$, $b_i(P^j, \mathbf{Z}^k)$, and $c_i(P^j, \mathbf{Z}^k)$ using a piecewise quadratic interpolation. This procedure provides an accurate representation of $a_i(P, \mathbf{Z})$, $b_i(P, \mathbf{Z})$, and $c_i(P, \mathbf{Z})$ with a reasonably small number of sample points. The values $a_i(P^j, \mathbf{Z}^k)$, $b_i(P^j, \mathbf{Z}^k)$, and $c_i(P^j, \mathbf{Z}^k)$ are updated according to the recursive procedure (2.17)–(2.19).

We evaluate the expectations in (2.17)–(2.19) by replacing them with the corresponding integrals. For all the models considered in this paper, these integrals involve Gaussian kernels. We use Gauss-Hermite quadrature formulas (see, for example, Stroud [1971]) to obtain efficient numerical approximations of these integrals.

In all cases except for the path-dependent options, we perform numerical computations for a European put option with a unit strike price ($K = 1$), i.e., $F(P_T) = \max(0, K - P_T)$, and a six-month maturity. It is apparent from (2.17)–(2.21) that for a call option with the same strike price K , the replication error $c_i(\cdot)$ is the same as that of a put option, and the replication cost $b_i(\cdot)$ satisfies the put-call parity relation. We assume 25 trading periods, defined by $t_0 = 0, t_{i+1} - t_i = \Delta t = 1/50$.




Figure 1: The difference between the replication cost and the intrinsic value of a six-month maturity European put option, plotted as a function of the initial stock price. The stock price follows a geometric Brownian motion with parameter values $\mu = 0.07$ and $\sigma = 0.13$ corresponding to the solid line. In Panel (a), μ is varied and σ is fixed; in Panel (b), σ is varied and μ is fixed. In both cases, the variation in each parameter is obtained by multiplying its original value by 1.25 (dashed-dotted line), 1.5 (dots), 0.75 (dashed line) and 0.5 (pluses).

5.2 Geometric Brownian Motion

Let stock prices follow a geometric Brownian motion, which implies that returns are lognormally distributed as in (4.8)–(4.10). We set $\mu = 0.07$ and $\sigma = 0.13$, and to cover a range of empirically plausible parameter values, we vary each parameter by increasing and decreasing them by 25% and 50% while holding the values of other parameter fixed. Figure 1 displays the minimum replication cost V_0^* minus the intrinsic value $F(P_0)$, for the above range of parameter values, as a function of the stock price at time 0.

Figure 1 shows that V_0^* is not sensitive to changes in μ and increases with σ . This is not surprising given that V_0^* approximates the Black-Scholes option pricing formula.

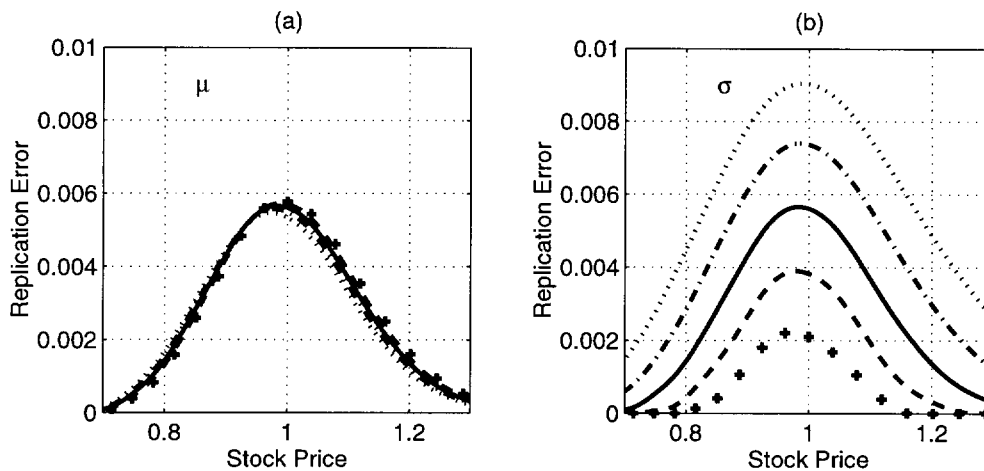


Figure 2: The replication error of a six-month maturity European put option, plotted as a function of the initial stock price. The stock price follows a geometric Brownian motion with parameter values $\mu = 0.07$ and $\sigma = 0.13$ corresponding to the solid line. In Panel (a), μ is varied and σ is fixed; in Panel (b), σ is varied and μ is fixed. In both cases, the variation in each parameter is obtained by multiplying its original value by 1.25 (dashed-dotted line), 1.5 (dots), 0.75 (dashed line) and 0.5 (pluses).

Figure 2 shows the dependence of the replication error ϵ^* on the initial stock price. Again we observe low sensitivity to the drift μ but, as in Figure 1, the replication error tends to increase with the volatility. We also observe that the replication error is highest when the stock price is close to the strike price.

Another important characteristic of the replication process is the ratio of the replication error to the replication cost ϵ^*/V_0^* , which we call the relative replication error. This ratio is more informative than the replication error itself since it describes the replication error per dollar spent, as opposed to the error of replicating a single option contract.

The dependence of the relative replication error on the initial stock price is displayed in Figure 3. This figure shows that the relative replication error is an increasing function of the initial stock price, i.e., it is higher for out-of-the-money options. Also, the relative replication error decreases with volatility for out-of-the-money options. This is not surprising given that it was defined as a ratio of the replication error to the hedging cost, both of which are

increasing functions of volatility. According to this definition, the dependence of the relative replication error :

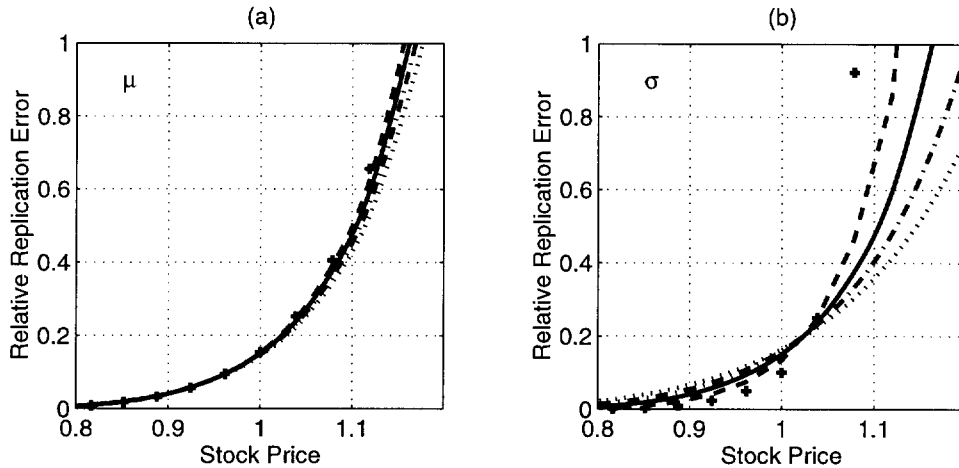


Figure 3: The relative replication error of a six-month maturity European put option (relative to the replication cost), plotted as a function of the initial stock price. The stock price follows a geometric Brownian motion with parameter values $\mu = 0.07$ and $\sigma = 0.13$ corresponding to the solid line. In Panel (a), μ is varied and σ is fixed; in Panel (b), σ is varied and μ is fixed. In both cases, the variation in each parameter is obtained by multiplying its original value by 1.25 (dashed-dotted line), 1.5 (dots), 0.75 (dashed line) and 0.5 (pluses).

5.3 Jump-Diffusion Models

Our numerical computations are based on the model (4.15)–(4.19), (4.21). In our numerical implementation we restrict the number of jumps over a single time interval to be no more than three, which amounts to modifying the distribution of n_i in (4.16), originally given by (4.19).²¹ Specifically, we replace (4.19) with

$$\text{Prob}[n_i = m] = e^{-\lambda\Delta t} \frac{(\lambda\Delta t)^m}{m!}, \quad m = 1, 2, 3 \quad (5.1)$$

$$\text{Prob}[n_i = 0] = 1 - \sum_{m=1}^3 \text{Prob}[n_i = m]. \quad (5.2)$$

²¹This “truncation problem” is a necessary evil in the estimation of jump-diffusion models. See Ball and Torous (1985) for further discussion.

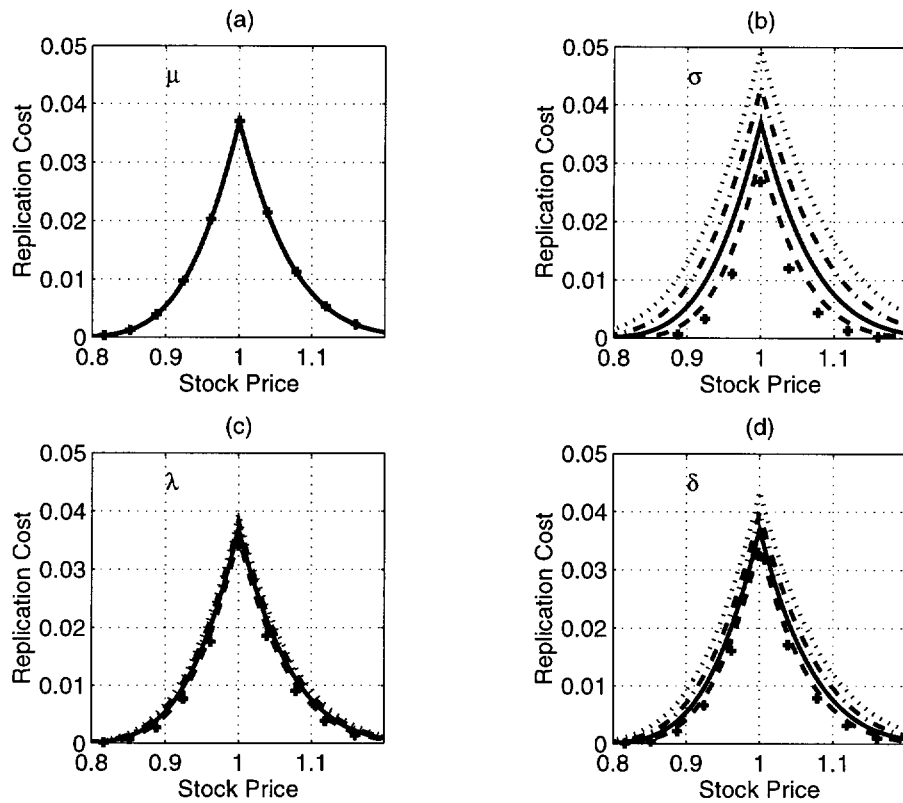


Figure 4: The difference between the replication cost and the intrinsic value of a six-month maturity European put option, plotted as a function of the initial stock price. The stock price follows the mixed jump-diffusion process given in (4.15)–(4.18), (4.21), (5.1), and (5.2) with parameter values $\mu = 0.07$, $\sigma = 0.106$, $\lambda = 25$, and $\delta = 0.015$ corresponding to the solid line. In Panels (a)–(d), μ , σ , λ , and δ are each varied, respectively, while the other parameter values are held fixed. The variation in each parameter is obtained by multiplying its original value by 1.25 (dashed-dotted line), 1.5 (dots), 0.75 (dashed line), and 0.5 (pluses).

Besides this adjustment in the distribution of returns, our numerical procedure is exactly the same as in Section 4.2. We start with the following parameter values:

$$\mu = 0.07 \quad , \quad \sigma = 0.106 \quad , \quad \lambda = 25 \quad , \quad \delta = 0.015 \quad .$$

Then we study sensitivity of the solution to the parameter values by increasing and decreasing them by 25% and 50% while holding the other parameter values fixed. Our numerical results are summarized in Figures 4, 5, 6.

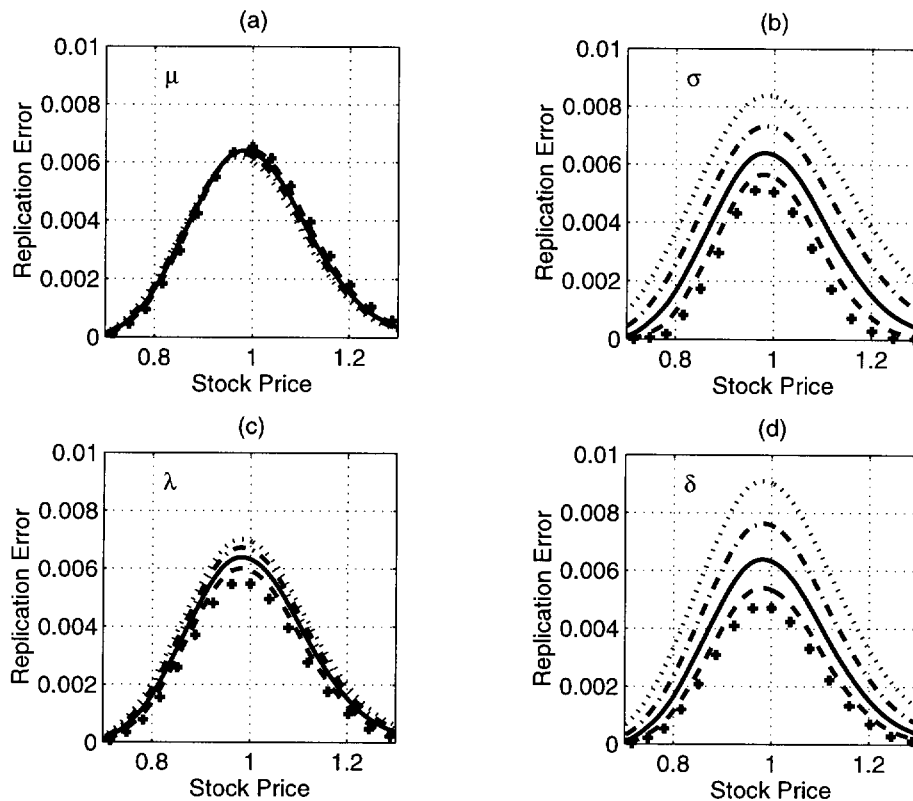


Figure 5: The replication error of a six-month maturity European put option, plotted as a function of the initial stock price. The stock price follows the mixed jump-diffusion process given in (4.15)–(4.18), (4.21), (5.1), and (5.2) with parameter values $\mu = 0.07$, $\sigma = 0.106$, $\lambda = 25$, and $\delta = 0.015$ corresponding to the solid line. In Panels (a)–(d), μ , σ , λ , and δ are each varied, respectively, while the other parameter values are held fixed. The variation in each parameter is obtained by multiplying its original value by 1.25 (dashed-dotted line), 1.5 (dots), 0.75 (dashed line), and 0.5 (pluses).

Figure 4 shows that the replication cost V_0^* is not sensitive to the drift rate μ and is increasing in volatility σ , the jump intensity λ , and the standard deviation δ of the jump

magnitude. It is most sensitive to σ . According to Figure 5, the replication error ϵ^* is not sensitive to μ and increases with all other parameters, with the highest sensitivity to δ . Finally, Figure 6 shows that the relative replication error ϵ^*/V_0^* is sensitive only to σ and it decreases as a function of σ for out-of-the-money options.




Figure 6: The relative replication error of a six-month maturity European put option (relative to the replication cost), plotted as a function of the initial stock price. The stock price follows the mixed jump-diffusion process (4.15)–(4.18), (4.21), (5.1), and (5.2) with parameter values $\mu = 0.07$, $\sigma = 0.106$, $\lambda = 25$, and $\delta = 0.015$ corresponding to the solid line. In Panels (a)–(d), μ , σ , λ , and δ are each varied, respectively, while the other parameter values are held fixed. The variation in each parameter is obtained by multiplying its original value by 1.25 (dashed-dotted line), 1.5 (dots), 0.75 (dashed line), and 0.5 (pluses).

5.4 Stochastic Volatility

We begin by assuming a particular functional form for $g(\sigma)$ in (4.44):

$$g(\sigma) = -\delta\sigma(\sigma - \zeta).$$

We also assume that the Brownian motions driving the stock price and volatility are uncorrelated. Since the closed-form expressions for the transition probability density of the diffusion process with stochastic volatility are not available, we base our computations on the discrete-time approximations of this process.²² The dynamics of stock prices and volatil-

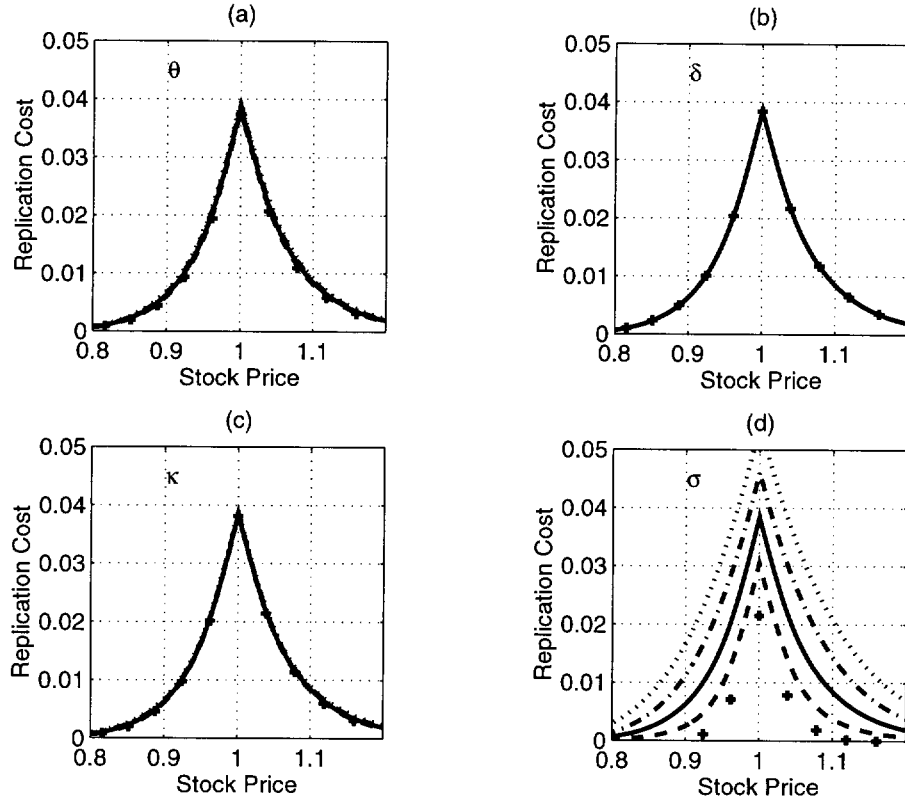


Figure 7: The difference between the replication cost and the intrinsic value of a six-month maturity European put option, plotted as a function of the initial stock price. The stock price follows the with stochastic volatility model (5.3)–(5.4) with parameter values $\mu = 0.07$, $\zeta = 0.153$, $\delta = 2$, $\kappa = 0.4$, and $\sigma_0 = 0.13$ corresponding to the solid line. In Panels (a)–(d), ζ , δ , κ , and σ_0 are each varied, respectively, while the other parameter values are held fixed. The variation in each parameter is obtained by multiplying its original value by 1.25 (dashed-dotted line), 1.5 (dots), 0.75 (dashed line), and 0.5 (pluses).

²²This is done mostly for convenience, since we could approximate the transition probability density using Monte Carlo simulations. It should be pointed out that, while the discrete-time approximations lead to significantly more efficient numerical algorithms, they are also consistent with many estimation procedures, replacing continuous-time processes with their discrete-time approximations (see, for example, Ball and Torous [1985] and Wiggins [1987]).

ity are described by

$$P_{i+1} = P_i \exp\left((\mu - \sigma_i^2/2)\Delta t + \sigma_i\sqrt{\Delta t}z_{P_i}\right) \quad (5.3)$$

$$\sigma_{i+1} = \sigma_i \exp\left((-\delta(\sigma_i - \zeta) - \kappa^2/2)\Delta t + \kappa\sqrt{\Delta t}z_{\sigma_i}\right) \quad (5.4)$$

where $z_{P_i}, z_{\sigma_i} \sim \mathcal{N}(0, 1)$ and $E^\nu[z_{P_i}z_{\sigma_i}] = 0$. The parameters of the model are chosen to be

$$\mu = 0.07 \quad , \quad \zeta = 0.153 \quad , \quad \delta = 2 \quad , \quad \kappa = 0.4 \quad . \quad (5.5)$$

We also assume that at time $t = 0$ volatility σ_0 is equal to 0.13. As before, we study sensitivity of the solution to parameter values. Our findings are summarized in Figures 7, 8, 9.

We do not display the dependence on μ in these figures since the sensitivity to this parameter is so low. Figure 7 shows that the replication cost is sensitive only to the initial value of volatility σ_0 and, as expected, the replication cost increases with σ_0 . Figure 8 shows that the replication error is sensitive to κ and σ_0 and is increasing in both of these parameters. According to Figure 9, the relative replication error is increasing in κ . It also increases in σ_0 for in-the-money options and decreases for out-of-the-money options.

In addition its empirical relevance, the stochastic volatility model (4.43)–(4.44) also provides a clear illustration of the use of ϵ^* as a quantitative measure of dynamic market-incompleteness. Table 2 reports the results of Monte Carlo experiments in which the optimal-replication strategy is implemented for six sets of parameter values for the stochastic volatility model, including the set that yields geometric Brownian motion.

For each set of parameter values, 1,000 independent sample paths of the stock price are simulated, each sample path containing 25 observations, and for each path the optimal-replication strategy is implemented. The averages (over the 1,000 sample paths) of the minimum production cost V_0^* , the realized replication error $\hat{\epsilon}^*$, the initial optimal stock holdings θ_0^* , and the average optimal stock holdings $\bar{\theta}^*$ (over the 25 periods), is reported in each row. For comparison, the theoretical replication error ϵ^* is also reported.

Since stochastic volatility implies dynamically incomplete markets whereas geometric Brownian motion implies the opposite, these six sets of simulations comprise a sequence of

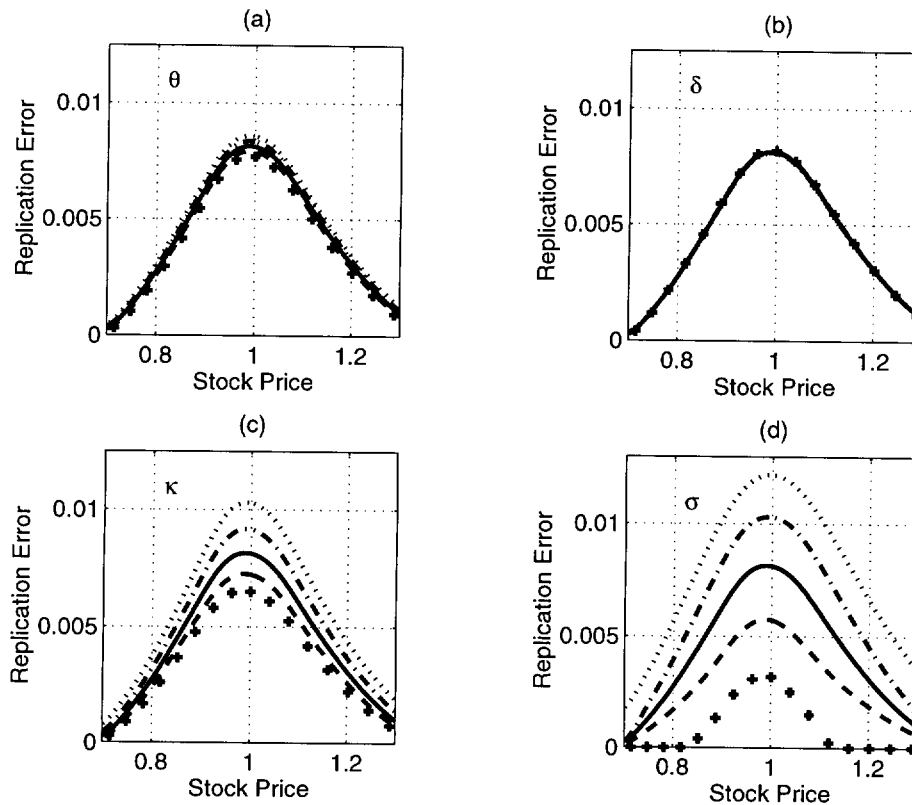


Figure 8: The replication error of a six-month maturity European put option, plotted as a function of the initial stock price. The stock price follows the with stochastic volatility model (5.3)–(5.4) with parameter values $\mu = 0.07$, $\zeta = 0.153$, $\delta = 2$, $\kappa = 0.4$, and $\sigma_0 = 0.13$ corresponding to the solid line. In Panels (a)–(d), ζ , δ , κ , and σ_0 are each varied, respectively, while the other parameter values are held fixed. The variation in each parameter is obtained by multiplying its original value by 1.25 (dashed-dotted line), 1.5 (dots), 0.75 (dashed line), and 0.5 (pluses).




Figure 9: The relative replication error of a six-month maturity European put option (relative to the replication cost), plotted as a function of the initial stock price. The stock price follows the with stochastic volatility model (5.3)–(5.4) with parameter values $\mu = 0.07$, $\zeta = 0.153$, $\delta = 2$, $\kappa = 0.4$, and $\sigma_0 = 0.13$ corresponding to the solid line. In Panels (a)–(d), ζ , δ , κ , and σ_0 are each varied, respectively, while the other parameter values are held fixed. The variation in each parameter is obtained by multiplying its original value by 1.25 (dashed-dotted line), 1.5 (dots), 0.75 (dashed line), and 0.5 (pluses).




Figure 10: The difference between the replication cost and the intrinsic value of a six-month maturity European put option, plotted as a function of the initial stock price. The stock price follows the with stochastic volatility model (5.3)–(5.4) with parameter values: $\mu = 0.07$, $\zeta = 0.153$, $\delta = 2$, $\kappa = 0.4$, $\sigma_0 = 0.13$ (solid line); $\mu = 0$, $\zeta = 0.153$, $\delta = 2$, $\kappa = 0.4$, $\sigma_0 = 0.13$ (dashed line); and $\mu = 0$, $\zeta = 0$, $\delta = 2$, $\kappa = 0.4$, $\sigma_0 = 0.13$ (dashed-dotted line).

Table 2: Monte Carlo simulation of the optimal-replication strategy θ^* for replicating a six-month at-the-money European put-option, for six sets of parameter values of the stochastic volatility model (5.3)–(5.4), including the set of parameter values that yields a geometric Brownian motion (last row). For each set of parameter values, 1,000 independent sample paths were simulated, each path containing 25 periods, and $P_0 = 1$.

Model				Performance of Optimal Replication Strategy				
σ_0	ζ	δ	κ	V_0^*	ϵ^*	$\hat{\epsilon}^*$	θ_0^*	$\bar{\theta}^*$
0.13	0.153	2.00	0.400	0.0374	0.0084	0.0086	-0.474	-0.44
0.13	0.137	1.50	0.200	0.0367	0.0067	0.0068	-0.475	-0.43
0.13	0.133	1.00	0.100	0.0366	0.0062	0.0063	-0.475	-0.42
0.13	0.131	0.50	0.050	0.0366	0.0060	0.0059	-0.475	-0.41
0.13	0.130	0.25	0.025	0.0365	0.0060	0.0060	-0.475	-0.43
0.13	0.130	0.00	0.000	0.0365	0.0060	0.0062	-0.475	-0.42

models that illustrate the fact that market completeness need not be a binary characteristic. In particular, Table 2 shows that as the parameter values move closer to geometric Brownian motion, the average replication error ϵ^* decreases from 0.0086 to 0.0060. Moreover, the decrease between the first and second rows is considerably larger than the decrease between the second and third rows—the second and third rows imply price processes that are closer to each other in their degree of market completeness than that of the first row. Such specific rank orderings and sharp numerical comparisons are simply unavailable from standard dynamic equilibrium models that have been used to model market incompleteness.

Of course, ϵ^* is only one of many possible measures of market incompleteness—a canonical measure seems unlikely to emerge from the current literature—nevertheless it is an extremely useful measure given the practical implications that it contains for dynamically hedging risks.

5.5 Stochastic Volatility Under The Risk-Neutral Measure

To study the effects of changing the original probability measure ν to the risk-neutral probability measure ν^* , we consider the case when the stock price follows the stochastic volatility process (5.3)–(5.4). Under the original probability measure ν , we assume that parameters of the model are:

$$\mu = 0.07 \quad , \quad \zeta = 0.153 \quad , \quad \delta = 2 \quad , \quad \kappa = 0.4 \quad , \quad \sigma_0 = 0.13 \quad . \quad (5.6)$$

Under the risk-neutral probability measure ν^* the drift rate of the stock price is zero. To define the process completely we need to specify the drift rate of the volatility under the new measure. We consider two cases. In the first case the new (risk-adjusted) drift rate is equal to the original drift rate, i.e., parameters δ and ζ remain unchanged. In the second case we assume that there exists a risk premium on the risk associated with the volatility process, so that the risk-adjusted drift rate of this process is different from its original value. To be specific, we leave δ unchanged and set $\zeta = 0$. Our results are summarized in Figure 10.

Figure 10 shows that the replication cost obtained by using the original probability measure ν is almost identical to the price of the option in case when volatility risk is not priced by the market. On the other hand, when there exists a risk premium on the volatility risk, the replication cost is different from exact price. The sign and magnitude of the difference

obviously depend on the exact form of the risk premium.

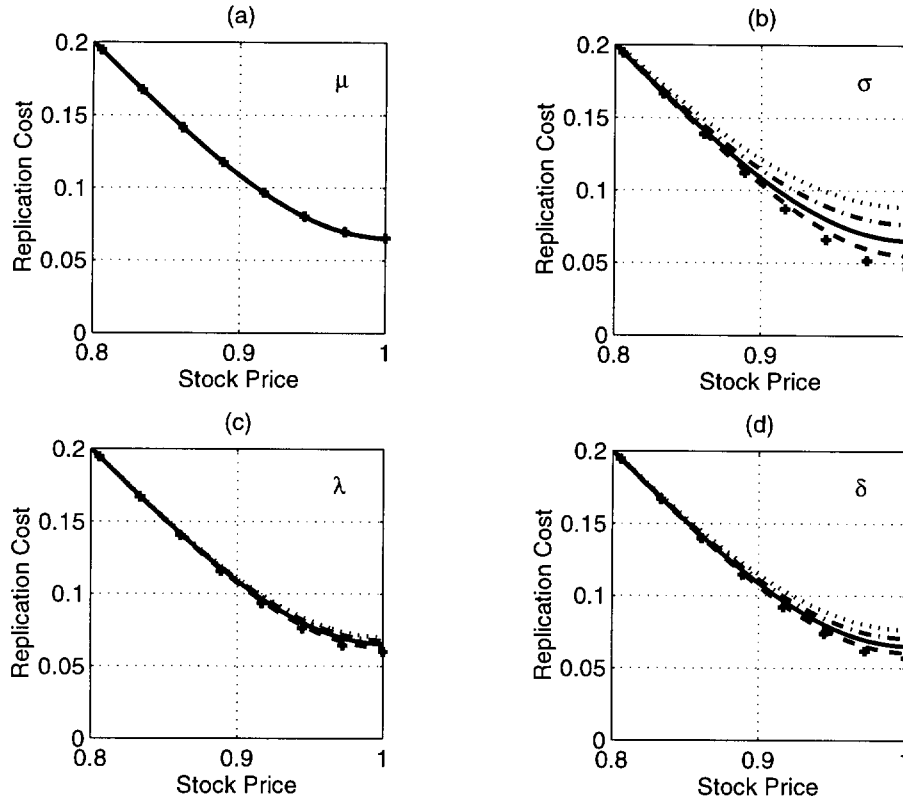


Figure 11: The replication cost of a six-month maturity European option to “sell at the high”, plotted as a function of the initial stock price. The stock price follows the mixed jump-diffusion process (4.15)–(4.18), (4.21), (5.1), and (5.2) with parameter values $m = 1$, $\mu = 0.07$, $\sigma = 0.106$, $\lambda = 25$, and $\delta = 0.015$ corresponding to the solid line. In Panels (a)–(d), μ , σ , λ , and δ are each varied, respectively, while the other parameter values are held fixed. The variation in each parameter is obtained by multiplying its original value by 1.25 (dashed-dotted line), 1.5 (dots), 0.75 (dashed line), and 0.5 (pluses).

5.6 Path-Dependent Options

We consider the option to “sell at the high” as described by Goldman et al. (1979), under the assumption that the stock price follows the mixed jump-diffusion process (4.15)–(4.18), (4.21), (5.1), (5.2). We define the state variable Z :

$$\begin{aligned} Z_0 &= m \geq P_0 \\ Z_{i+1} &= \max(Z_i, P_{i+1}) . \end{aligned}$$

According to this definition, Z_i is the running maximum of the stock price process at time t_i . The initial value of Z_i is m , i.e., we assume that at time 0 the running maximum is equal to m .

The payoff of the option is given by

$$F(P_T, Z_T) = Z_T - P_T .$$

In our numerical analysis we set $m = 1$ as a convenient normalization. Note that this convention is just a change of scale and does not lead to any loss of generality.

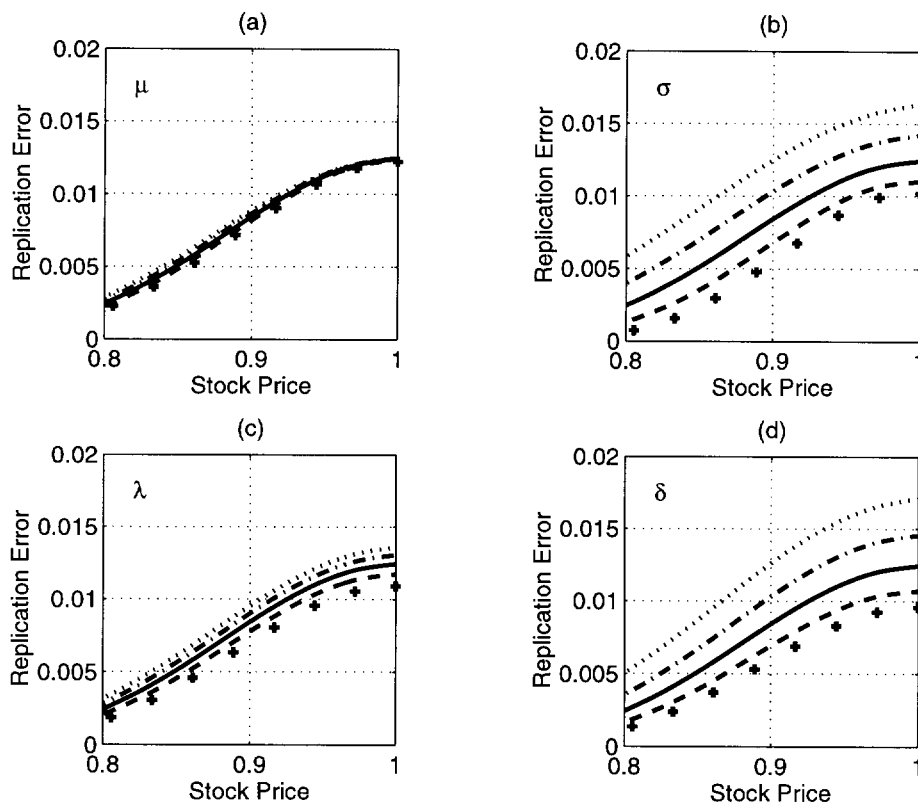


Figure 12: The replication error of a six-month maturity European option to “sell at the high”, plotted as a function of the initial stock price. The stock price follows the mixed jump-diffusion process (4.15)–(4.18), (4.21), (5.1), and (5.2) with parameter values $m = 1$, $\mu = 0.07$, $\sigma = 0.106$, $\lambda = 25$, and $\delta = 0.015$ corresponding to the solid line. In Panels (a)–(d), μ , σ , λ , and δ are each varied, respectively, while the other parameter values are held fixed. The variation in each parameter is obtained by multiplying its original value by 1.25 (dashed-dotted line), 1.5 (dots), 0.75 (dashed line), and 0.5 (pluses).

The parameters for the stock price process are taken to be the same as in Section (4.3).

The sensitivity of the replication cost and replication error on the initial stock price and parameters of the stock price process are reported in Figures 11, 12, 13.




Figure 13: The relative replication error of a six-month maturity European option to “sell at the high” (relative to the replication cost), plotted as a function of the initial stock price. The stock price follows the mixed jump-diffusion process (4.15)–(4.18), (4.21), (5.1), and (5.2) with parameter values $m = 1$, $\mu = 0.07$, $\sigma = 0.106$, $\lambda = 25$, and $\delta = 0.015$ corresponding to the solid line. In Panels (a)–(d), μ , σ , λ , and δ are each varied, respectively, while the other parameter values are held fixed. The variation in each parameter is obtained by multiplying its original value by 1.25 (dashed-dotted line), 1.5 (dots), 0.75 (dashed line), and 0.5 (pluses).

The qualitative behavior of the replication cost as a function of the initial stock price is similar to that of the option price as described in Goldman et al. (1979).²³ Figure 11 shows that the replication cost V_0 is not sensitive to the drift rate μ and is increasing in volatility σ , the jump intensity jumps λ , and the standard deviation δ of the jump magnitude. It is most sensitive to σ . These observations are consistent with the behavior of the replication error

²³The difference between our model and that in Goldman et al. (1979) is that the latter assumes that the stock price follows a geometric Brownian motion and that continuous-time trading is allowed. Also the running maximum of the stock price process is calculated continuously, not over a discrete set of time moments, as in our case.

of the European put option in Section (4.3). According to Figure 12, the replication error ϵ^* is not sensitive to μ and is increasing in all other parameters with the highest sensitivity to δ and σ . Figure 6 shows that the relative replication error ϵ^*/V_0 is sensitive to σ and δ . It is an increasing function of δ , while the sign of the change of ϵ^*/V_0 with σ depends on the initial stock price P_0 .

6 Specification Analysis of Replication Errors

In this section, we explore the sensitivity of the replication error and the replication cost of a particular option contract to the specification of the stock-price dynamics. Specifically, we compare the following models: geometric Brownian motion, a mixed jump-diffusion process, and a diffusion process with stochastic volatility. The parameters of these models are calibrated to give rise to identical values of the expected instantaneous rate of return and volatility, hence we can view these three models as competing specifications of the same data-generating process.

We consider a European put option with a unit strike price ($K = 1$) and a six-month maturity, i.e., $F(P_T) = \max(0, K - P_T)$. There are 25 trading periods, defined by $t_{i+1} - t_i = \Delta t = 1/50$. Since the closed-form expressions for the transition probability density of the mixed jump-diffusion process and the process with stochastic volatility are not available, we base our computations on the discrete-time approximations of these processes. The model specifications and corresponding parameter values are:

1. **Geometric Brownian Motion.** Returns on the stock are lognormal, given by (4.8)–(4.10). We use the following parameter values:

$$\mu = 0.07 \quad , \quad \sigma = 0.13 \quad . \quad (6.1)$$

2. **Mixed Jump-Diffusion.** The distribution of returns on the stock is given by (4.15)–(4.18), (4.21), (5.1), and (5.2). We use the following parameter values:

$$\mu = 0.07 \quad , \quad \sigma = 0.106 \quad , \quad \lambda = 25 \quad , \quad \delta = 0.015 \quad . \quad (6.2)$$




Figure 14: The difference between the replication cost and the intrinsic value of a six-month maturity European put option, plotted as a function of the initial stock price. Several processes for the stock price are plotted: geometric Brownian motion (4.8)–(4.10) (solid line); the mixed jump-diffusion model (4.15)–(4.18), (4.21), (5.1), and (5.2) (dashed line); and the stochastic volatility model (5.3)–(5.4) (dashed-dotted line). The parameter values are given by (6.1), (6.2), and (6.3).

3. **Diffusion with Stochastic Volatility.** Stock-price and volatility dynamics are given by (5.3)–(5.4), and the parameters are:

$$\mu = 0.07 \quad , \quad \zeta = 0.153 \quad , \quad \delta = 2 \quad , \quad \kappa = 0.4 \quad . \quad (6.3)$$

We assume that at time $t = 0$, volatility σ_0 is equal to 0.13.




Figure 15: The replication error of a six-month maturity European put option, plotted as a function of the initial stock price. Several processes for the stock price are plotted: geometric Brownian motion (4.8)–(4.10) (solid line); the mixed jump-diffusion model (4.15)–(4.18), (4.21), (5.1), and (5.2) (dashed line); and the stochastic volatility model (5.3)–(5.4) (dashed-dotted line). The parameter values are given by (6.1), (6.2), and (6.3).

Figures 14–16 and Table 3 summarize our numerical results. Figure 14 presents the replication cost V_0^* minus the intrinsic value $F(P_0)$ for the three models as a function of the stock price at time $t = 0$. The hedging costs for the first two models are practically identical, while the stochastic volatility model can give rise to a significantly higher hedging costs for a deep-out-of-the-money option. Figure 15 and Table 3 shows the dependence of the replication error ϵ^* on the initial stock price.

All three models exhibit qualitatively similar behavior: the replication error is highest close to the strike price. For our choice of parameter values the replication error is highest

Table 3: Comparison of replication costs and errors of the optimal replication strategy for replicating a six-month European put option under competing specifications of price dynamics: geometric Brownian motion (4.8)–(4.10); the mixed jump-diffusion model (4.15)–(4.18), (4.21), (5.1), and (5.2); and the stochastic volatility model (5.3)–(5.4). The parameter values are given by (6.1), (6.2), and (6.3).

Price Dynamics	Initial Stock Price P_0				
	0.90	0.95	1.00	1.05	1.10
Replication Cost Minus Intrinsic Value ($V_0^* - F(P_T)$)					
Geometric Brownian Motion	0.0054	0.0161	0.0365	0.0176	0.0074
Jump/Diffusion	0.0053	0.0159	0.0364	0.0175	0.0074
Stochastic Volatility	0.0060	0.0168	0.0374	0.0185	0.0082
Replication Error (ϵ^*)					
Geometric Brownian Motion	0.0046	0.0058	0.0060	0.0052	0.0038
Jump/Diffusion	0.0051	0.0066	0.0068	0.0059	0.0043
Stochastic Volatility	0.0061	0.0079	0.0084	0.0074	0.0056
Relative Replication Error (ϵ^*/V_0^*)					
Geometric Brownian Motion	0.043	0.088	0.164	0.292	0.509
Jump/Diffusion	0.049	0.100	0.187	0.335	0.585
Stochastic Volatility	0.058	0.119	0.226	0.400	0.679

for the stochastic volatility model and lowest for geometric Brownian motion. However, this need not hold in general. As we demonstrate in Section 4.3, the replication error of the mixed jump-diffusion process depends critically on δ and λ in (4.19, 4.21), thus, by varying these parameters, one can reverse the order of the curves in Figure 15 without changing the annualized volatility of the mixed jump-diffusion process.




Figure 16: The relative replication error of a six-month maturity European put option (relative to the replication cost), plotted as a function of the initial stock price. Several processes for the stock price are plotted: geometric Brownian motion (4.8)–(4.10) (solid line); the mixed jump-diffusion model (4.15)–(4.18), (4.21), (5.1), and (5.2) (dashed line); and the stochastic volatility model (5.3)–(5.4) (dashed-dotted line). The parameter values are given by (6.1), (6.2), and (6.3).

The dependence of the relative replication error on the initial stock price is captured in Figure 16. As in Figure 14, the relative replication error for the first two models are practically identical, while the stochastic volatility model can exhibit considerably higher errors. Also, while the relative replication error can be significant, particularly for an out-of-the-money option, the variation across the models is not as significant as one would expect. When continuous-time trading is allowed, the replication error for the geometric Brownian motion model is zero, while the other two models give rise to strictly positive replication errors. This is an implication of the fact that the first model describes a dynamically complete market, while the other two correspond to markets which are dynamically incomplete (due

to the absence of a sufficient number of traded instruments).

Nevertheless, as Figure 16 illustrates, the transition from continuous- to discrete-time trading can smear the differences between these models, leading to relative replication errors of comparable magnitude. This shows that impossibility of continuous-time trading is just as important as a source of market incompleteness, as the absence of traded instruments.

7 Conclusion

We have proposed a method for replicating derivative securities in dynamically incomplete markets. Using stochastic dynamic programming, we construct a self-financing dynamic portfolio strategy that best approximates an arbitrary payoff function in a mean-squared sense. When markets are dynamically complete, as in the Black and Scholes (1973) and Merton (1973) models, our optimal-replication strategy coincides with the delta-hedging strategies of arbitrage-based models. Moreover, we provide an explicit algorithm for computing such strategies, which can be a formidable challenge in spite of market completeness, e.g., path-dependent derivatives such as “look-back” options.

When markets are not dynamically complete, as in the case of options on assets with stochastic volatility or with jump components, our approach yields the minimum production cost of a self-financing portfolio strategy with a terminal value that comes as close as possible (in a mean-squared sense) to the option’s payoff. This is the essence of the ϵ -arbitrage approach to synthetically replicating a derivative security.

More importantly, we show that if the mean-squared loss function is evaluated with respect to the equivalent martingale measure, the minimum production cost associated with the optimal-replication strategy is the equilibrium price of the derivative security. Therefore, the ϵ -arbitrage approach can also be viewed as a method for pricing, not just replicating, derivative securities in dynamically incomplete markets. Of course, obtaining the equivalent martingale measure is the most difficult aspect of asset pricing in incomplete markets, and we discuss theoretical and empirical approaches to confronting this challenge.

We also argue that the replication error of the optimal-replication strategy can be used as a quantitative measure for the degree of market incompleteness. Despite the difficulties in making welfare comparisons between markets with different types of incompleteness (see,

for example, Duffie [1987], Duffie and Shafer [1985, 1986], and Hart [1974]), the minimum replication error of an ϵ -arbitrage strategy does provide one practical metric by which market completeness can be judged. After all, if it is possible to replicate the payoff of a derivative security to within some small error ϵ , the market for that security may be considered complete for all practical purposes even if ϵ is not zero.

Of course, this is only one of many possible measures of market completeness and we make no claims of generality here. Instead, we hope to have shown that Merton's (1973) seminal idea of dynamic replication has far broader implications than the dynamically-complete-markets setting in which it was originally developed. We plan to explore other implications in future research.

A Appendix

The proofs of Theorems 2 and 1 are conceptually straightforward but notationally quite cumbersome. Therefore, we present only a brief sketch of the proofs below—interested readers can contact the authors for the more detailed mathematical appendix.

A.1 Proof of Theorem 1

The proof of Theorem 1 follows from dynamic programming. For $i = N$, (2.14)–(2.16) are clearly true, given (2.10). We now show that (2.17)–(2.21) describe the solution of the optimization problem in (2.9). First, as we observed in Section 2.2, the functions $a_i(\cdot, \cdot)$ are positive. Together with (2.3) this implies that

$$E^\nu \left[J_{i+1} \left(V_i + \theta_i(P_{i+1} - P_i), P_{i+1}, \mathbf{Z}_{i+1} \right) \middle| V_i, P_i, \mathbf{Z}_i \right]$$

is a convex function of θ_i . Therefore, we can use the first-order condition to solve the optimization problem in (2.11):

$$\frac{d}{d\theta_i} E^\nu \left[J_{i+1} \left(V_i + \theta_i(P_{i+1} - P_i), P_{i+1}, \mathbf{Z}_{i+1} \right) \middle| V_i, P_i, \mathbf{Z}_i \right] = 0, \quad (\text{A.1})$$

where $J_{i+1}(\cdot, \cdot, \cdot)$ is given by (2.12). Equation (A.1) is a linear equation in θ_i and it is straightforward to check that its solution, $\theta^*(i, V_i, P_i, \mathbf{Z}_i)$, is given by (2.13), (2.17), and (2.18). We now substitute (2.13) into (2.3) and use (2.11) to calculate

$$J_i(V_i, P_i, \mathbf{Z}_i) = E^\nu \left[J_{i+1} \left(V_i + \theta^*(i, V_i, P_i, \mathbf{Z}_i) \cdot (P_{i+1} - P_i), P_{i+1}, \mathbf{Z}_{i+1} \right) \middle| V_i, P_i, \mathbf{Z}_i \right]. \quad (\text{A.2})$$

Equations (2.19)–(2.21) are obtained by rearranging terms in (A.2).

A.2 Proof of Theorem 2

The more tedious algebraic manipulations of this proof were carried out using the symbolic algebra program *Maple*. Therefore, we shall outline the main ideas of the proof without

reporting all of the details.

The cost-to-go function $J(t, V_t, P_t, \mathbf{Z}_t)$ satisfies the dynamic programming equation

$$\frac{\partial J}{\partial t} + \min_{\theta_t} \left\{ \left[\sum_{j=0}^J \mu_j Z_j \frac{\partial}{\partial Z_j} + \frac{1}{2} \sum_{i,j=0}^J \sigma_i \sigma_j Z_i Z_j \rho_{ij} \frac{\partial^2}{\partial Z_i \partial Z_j} + \theta_t \mu_0 Z_0 \frac{\partial}{\partial W} + \frac{1}{2} (\theta_t \sigma_0 Z_0)^2 \frac{\partial^2}{\partial W^2} + \theta_t \sum_{j=0}^J \sigma_0 \sigma_j \rho_{0j} Z_0 Z_j \frac{\partial^2}{\partial W \partial Z_j} \right] J \right\} = 0 \quad (\text{A.3})$$

with boundary condition:

$$J(T, V_T, P_T, \mathbf{Z}_T) = [V_T - F(P_T, \mathbf{Z}_T)]^2 \quad (\text{A.4})$$

where some of the functional dependencies were omitted to simplify the notation.

We must now check that the function $J(t, V_t, P_t, \mathbf{Z}_t)$, given by (2.26), (2.27)–(2.30), and the optimal control (2.31), satisfies (A.3)–(A.4). Boundary conditions (2.30) immediately imply (A.4). Next we substitute (2.26) into (A.3). It is easy to check, using equation (2.27), that function $a(\cdot)$ is positive. Therefore, the first-order condition is sufficient for the minimum in (A.3). This condition is a linear equation in θ_t which is solved by (2.31). It is now straightforward to verify that, whenever functions $a(\cdot)$, $b(\cdot)$, $c(\cdot)$ satisfy (2.27)–(2.29), (A.3) is satisfied as well.

References

- Aït-Sahalia, Y. and A. Lo, 1996, "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices", to appear in *Journal of Finance*.
- Aiyagari, R., 1994, "Uninsured Idiosyncratic Risk and Aggregate Saving", *Quarterly Journal of Economics* 109, 659–684.
- Aiyagari, R. and M. Gertler, 1991, "Asset Returns with Transactions Costs and Uninsured Individual Risk", *Journal of Monetary Economics* 27, 311–331.
- Amin, K., 1993, "Jump Diffusion Option Valuation in Discrete Time", *Journal of Finance* 48, 1833–1863.
- Amin, K. and V. Ng, 1993, "Option Valuation With Systematic Stochastic Volatility", *Journal of Finance* 48, 881–910.
- Ball, C. and W. Torous, 1985, "On Jumps in Common Stock Prices and Their Impact on Call Option Pricing", *Journal of Finance* 40, 155–173.
- Banz, R. and M. Miller, 1978, "Prices for State-Contingent Claims: Some Estimates and Applications", *Journal of Business* 51, 653–672.
- Bertsekas, D., 1995, *Dynamic Programming and Optimal Control*, Vol. I. Belmont, MA: Athena Scientific.
- Bick, A., 1995, "Quadratic-Variation-Based Dynamic Strategies", *Management Science*, 41(4), 722–732.
- Black, F. and M. Scholes, 1973, "Pricing of Options and Corporate Liabilities", *Journal of Political Economy* 81, 637–654.
- Boyle, P. and D. Emanuel, 1980, "Discretely Adjusted Option Hedges", *Journal of Financial Economics* 8, 259–282.
- Breeden, D. 1979, "An Intertemporal Capital Pricing Model with Stochastic Investment Opportunities", *Journal of Financial Economics* 7, 265–296.
- Breeden, D. and R. Litzenberger, 1978, "State contingent Prices Implicit in Option Prices", *Journal of Business* 51, 621–651.
- Brennan, M., 1979, "The Pricing of Contingent claims in Discrete-Time Models", *Journal of Finance* 34, 53–68.
- Cox, J. and S. Ross, 1976, "The Valuation of Options for Alternative Stochastic Processes", *Journal of Financial Economics* 3, 145–166.
- Cox, J., J. Ingersoll and S. Ross, 1985, "An Intertemporal General Equilibrium Model of Asset Prices", *Econometrica* 53, 363–384.
- Derman, E. and I. Kani, 1994, "Riding on the Smile", *RISK* 7 (February), 32–39.
- Duffie, D., 1996, *Dynamic Asset Pricing Theory*, Second Edition. Princeton, NJ: Princeton University Press.

- Duffie, D., 1987, "Stochastic Equilibria with Incomplete Financial Markets", *Journal of Economic Theory* 41, 405–416.
- Duffie, D. and C. Huang, 1985, "Implementing Arrow-Debreu Equilibria by Continuous Trading of Few Long-Lived Securities", *Econometrica* 53, 1337–1356.
- Duffie, D. and M. Jackson, 1990, "Optimal Hedging and Equilibrium in a Dynamic Futures Market", *Journal of Economic Dynamics and Control* 14, 21–33.
- Duffie, D. and W. Shafer, 1985, "Equilibrium in Incomplete Markets I: A Basic Model of Generic Existence", *Journal of Mathematical Economics* 14, 285–300.
- Duffie, D. and W. Shafer, 1986, "Equilibrium in Incomplete Markets II: Generic Existence in Stochastic Economies", *Journal of Mathematical Economics* 15, 199–216.
- Duffie, D. and M. Richardson, 1991, "Mean-Variance Hedging in Continuous Time", *Annals of Applied Probability* 1, 1–15.
- Duffie, D. and W. Zam, 1987, "The Consumption-Based Capital Asset Pricing Model", GSB Research Paper #922, Stanford University.
- Dybvig, P. and C. Huang, 1988, "Nonnegative Wealth, Absence of Arbitrage, and Feasible Consumption Plans", *Review of Financial Studies* 1, 377–401.
- Dumas, B., J. Flemming and R.E. Whaley, 1995, "Implied Volatility Functions: Empirical Tests", working paper, Fuqua School of Business, Duke University.
- Dupire, B., 1994, "Pricing with a Smile", *RISK* 7 (January), 18–20.
- Fleming W., R. Rishel, 1975, *Deterministic and Stochastic Optimal Control*. New York: Springer-Verlag.
- Föllmer, H. and D. Sonderman, 1986, "Hedging of Non-Redundant Contingent-Claims", in W. Hildebrand and A. Mas-Colell, eds., *Contributions to Mathematical Economics, in Honor of Gérard Debreu*. Amsterdam: North-Holland.
- Garman, M., 1976, "A General Theory of Asset Valuation under Diffusion State Processes", Working Paper No. 50, University of California Berkeley.
- Geske, R., 1979, "The Valuation of Compound Options", *Journal of Financial Economics* 7, 63–81.
- Harrison, J. and D. Kreps, 1979, "Martingales and Arbitrage in Multiperiod Securities Markets", *Journal of Economic Theory* 20, 381–408.
- Hart, O., 1974, "On the Existence of Equilibrium in a Securities Model", *Journal of Economic Theory* 9, 293–311.
- He, H. and D. Modest, 1995, "Market Frictions and Consumption-Based Asset Pricing", *Journal of Political Economy* 103, 94–117.
- Heaton, J. and D. Lucas, 1992, "The Effects of Incomplete Insurance Markets and Trading Costs in a Consumption-Based Asset Pricing Model", *Journal of Economic Dynamics and Control* 16, 601–620.
- Heaton, J. and D. Lucas, 1996, "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing", *Journal of Political Economy* 104, 443–487.

- He, H. and N. Pearson, 1991, "Consumption and Portfolio Policies with Incomplete Markets and Short-Sale Constraints", *Journal of Economic Theory* 54, 259–304.
- Huang, C., 1987, "An Intertemporal General Equilibrium Asset Pricing Model: The Case of Diffusion Information", *Econometrica* 55, 117–142.
- Hull, J. and A. White, 1987, "The Pricing of Options on Assets With Stochastic Volatilities", *Journal of Finance* 42, 281–300.
- Hutchinson, J., Lo, A. and T. Poggio, 1994, "A Nonparametric Approach to the Pricing and Hedging of Derivative Securities Via Learning Networks", *Journal of Finance* 49, 851–889.
- Jarrow, R. and A. Rudd, 1982, "Approximate Option Valuation for Arbitrary Stochastic Processes", *Journal of Financial Economics* 10, 347–369.
- Johnson H. and D. Shanno, 1987, "Option Pricing When the Variance is Changing", *Journal of Financial and Quantitative Analysis* 22, 143–151.
- Longstaff, F., 1992, "An Empirical Examination of the Risk-Neutral Valuation Model", working paper, College of Business, Ohio State University, and the Anderson Graduate School of Management, UCLA.
- Longstaff, F., 1995, "Option Pricing and the Martingale Restriction", *Review of Financial Studies* 8, 1091–1124.
- Lucas, R., 1978, "Asset Prices in an Exchange Economy", *Econometrica* 46, 1426–1446.
- Lucas, D., 1994, "Asset Pricing With Undiversifiable Income Risk and Short Sales Constraints: Deepening the Equity Premium Puzzle", *Journal of Monetary Economics* 34, 325–341.
- Magill, M. and M. Quinzii, 1996, *Theory of Incomplete Markets*. Cambridge, MA: MIT Press.
- Merton, R., 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model", *Journal of Economic Theory* 3, 373–413.
- Merton, R., 1973, "An Intertemporal Capital Asset Pricing Model", *Econometrica* 41, 867–887.
- Merton, R., 1976, "Option Pricing when Underlying Stock Returns are Discontinuous", *Journal of Financial Economics* 3, 125–144.
- Merton, R., 1992, *Continuous-Time Finance*. Cambridge, MA: Blackwell Publishers.
- Rady, S., 1994, "State Prices Implicit in Valuation Formulae for Derivative Securities: A Martingale Approach", Discussion Paper No. 181, LSE Financial Markets Group, London, UK.
- Rubinstein, M., 1976, "The Valuation of Uncertain Income Streams and the Pricing of Options", *Bell Journal of Economics* 7, 407–425.
- Rubinstein, M., 1983, "Displaced Diffusion Option Pricing", *Journal of Finance* 38, 213–217.
- Rubinstein, M., 1994, "Implied Binomial Trees", *Journal of Finance* 49, 771–818.

- Schäl, M., 1994, "On Quadratic Cost Criteria for Option Hedging" *Mathematics of Operations Research* 19, 121–131.
- Scheinkman, J. and L. Weiss, 1986, "Borrowing Constraints and Aggregate Economic Activity", *Econometrica* 54, 23–45.
- Schweizer, M., 1992, "Mean-Variance Hedging for General Claims", *Annals of Applied Probability* 2, 171–179.
- Schweizer, M., 1995, "Variance-Optimal Hedging in Discrete Time", *Mathematics of Operations Research* 20, 1–31.
- Shimko, D., 1993, "Bounds of Probability", *RISK* 6, 33–37.
- Stroud, A., 1971, *Approximate Calculation of Multiple Integrals*. Englewood Cliffs, NJ: Prentice-Hall.
- Stutzer, M., 1995, "A Simple Nonparametric Approach to Derivative Security Valuation", forthcoming in the *Journal of Finance*.
- Telmer, C., 1993, "Asset-Pricing Puzzles and Incomplete Markets", *Journal of Finance* 48, 1803–1832.
- Toft, K., 1996, "On the Mean-Variance Tradeoff in Option Replication with Transactions Costs", *Journal of Financial and Quantitative Analysis* 31, 233–263.
- Weil, P., 1992, "Equilibrium Asset Prices with Undiversifiable Labor Income Risk", *Journal of Economic Dynamics and Control* 16, 769–790.
- Wiggins, J., 1987, "Option Values Under Stochastic Volatility: Theory and Empirical Estimates", *Journal of Financial Economics* 5, 351–372.