
OPTIMAL NO-ARBITRAGE BOUNDS ON S&P 500 INDEX OPTIONS AND THE VOLATILITY SMILE

PATRICK J. DENNIS

This article shows that the volatility smile is not necessarily inconsistent with the Black–Scholes analysis. Specifically, when transaction costs are present, the absence of arbitrage opportunities does not dictate that there exists a unique price for an option. Rather, there exists a range of prices within which the option's price may fall and still be consistent with the Black–Scholes arbitrage pricing argument. This article uses a linear program (LP) cast in a binomial framework to determine the smallest possible range of prices for Standard & Poor's 500 Index options that are consistent with no arbitrage in the presence of transaction costs. The LP method employs dynamic trading in the underlying and risk-free assets as well as fixed positions in other options that trade on the same underlying security. One-way transaction-cost levels on the index, inclusive of the bid–ask spread, would have to be below six basis points for deviations from Black–Scholes pricing to present an arbitrage opportunity. Monte Carlo simulations are employed to assess the hedging error induced with a

This article has benefited substantially from the helpful suggestions of Richard J. Rendleman, Jr., Robert A. Connolly, Jennifer S. Conrad, Jacky So, Robert E. Whaley, and Marc Zenner. This article has also been improved by suggestions from seminar participants at the College of William and Mary, the University of Massachusetts, the University of North Carolina at Chapel Hill, Purdue University, and the University of Virginia. The author thanks the McIntire School of Commerce at the University of Virginia for financial support.

For correspondence, Patrick J. Dennis, McIntire School of Commerce, University of Virginia, Monroe Hall, Charlottesville, Virginia; e-mail: pjd9v@virginia.edu

Received April 2000; Accepted May 2001

■ *Patrick J. Dennis is an Assistant Professor in the Finance Area at the McIntire School of Commerce at the University of Virginia in Charlottesville, Virginia.*

12-period binomial model to approximate a continuous-time geometric Brownian motion. Once the risk caused by the hedging error is accounted for, transaction costs have to be well below three basis points for the arbitrage opportunity to be profitable two times out of five. This analysis indicates that market prices that deviate from those given by a constant-volatility option model, such as the Black–Scholes model, can be consistent with the absence of arbitrage in the presence of transaction costs. © 2001 John Wiley & Sons, Inc. *Jrl Fut Mark* 21:1151–1179, 2001

INTRODUCTION

This article explores deviations from Black–Scholes pricing, commonly known as the volatility smile within the context of efficient pricing in the presence of transaction costs. By establishing no-arbitrage bounds with pricing information from related securities and the costs of trading these securities, I show that the volatility smile may be consistent with constant volatility in that it does not present arbitrage opportunities when transaction costs are present. A linear program (LP) model is used to obtain the tightest possible no-arbitrage pricing bounds for traded options in the presence of transaction costs. The optimization model, originally proposed by Edirisinghe, Naik, and Uppal (1993) and cast in an LP framework by Dennis and Rendleman (1995), incorporates both dynamic trading in the underlying and risk-free asset and fixed positions in other options that trade on the same underlying security. The solution to the LP provides the optimal mix of securities and optimal rebalancing strategy for replicating long and short option positions. I find that the proportional transaction cost on the underlying asset would have to be less than three basis points for the deviations from Black–Scholes pricing to represent a reasonably low-risk arbitrage opportunity to a trader. At transaction costs above this level, no arbitrage opportunities are present.

Volatility Smile

Black and Scholes (1973) were the first to develop the strategy of synthetic option replication. Under various assumptions, including perfect markets and no transaction costs, they show that the payoffs to a continuously revised portfolio of stock and risk-free securities can replicate those of a put or call option.

If the underlying asset on which several options are written has constant volatility and the options are priced according to the Black–Scholes model, the volatilities that are implied from the market prices of the options with the Black–Scholes model should all be equal to the

volatility of the underlying asset. Several studies, including those of MacBeth and Merville (1979) and Rubinstein (1985), show not only that the implied volatilities are not equal but also that the deviations are systematic and form a smile or sneer shape. Thus, this variation in implied volatility, known as the volatility smile, must be due to a violation of one or more of the Black–Scholes assumptions.

To explain why the volatility smile exists, several competing theories have been developed, each of which relaxes one or more of the Black–Scholes assumptions. The first set of these theories, such as those of Rubinstein (1994) and Derman and Kani (1994), relax the constant volatility assumption and allow for a deterministic time- and state-dependent volatility function. The time- and state-dependent volatility function allows the returns of the underlying security to have distributions that deviate from the constant volatility normal distribution. In these models, the time- and state-dependent volatility function is determined by the fitting of the model to the observed market prices of traded options. In this respect, these models are similar to the stochastic term structure models of Ho and Lee (1986) and Heath, Jarrow, and Morton (1992), in that the model can be adjusted so that the model prices equal the market prices.¹ Dumas, Fleming, and Whaley (1998) provided an empirical test of these time- and state-dependent volatility models and found that there was a large mean square error between the model and market prices. They concluded that the time- and state-dependent volatility approach did not do a good job of explaining observed option prices.

The second set of these models, such as those of Hull and White (1987) and Heston (1993), are based on stochastic volatility and can also explain the volatility smile. In contrast with the models of Rubinstein (1994) and Derman and Kani (1994), the volatility of the underlying asset is not deterministic but stochastic. If the volatility has a negative correlation with the underlying security, then when the price of the underlying is low, volatility will be high, the probability of large price movements will be high, and the left tail will be fatter than in the case of constant volatility. Conversely, when the price of the underlying is high, volatility will be low, the probability of large price movements will be low, and the right tail will be thinner than in the case of constant volatility. Using this type of distribution to price options will result in relatively high prices for out-of-the-money puts and relatively low prices for in-the-money puts.

¹For this to hold in Rubinstein (1994), the time- and state-dependent volatility function has to be defined with as many free parameters as there are options to price so that there is an exactly identified system of equations.

Although stochastic volatility models would seem to offer a good explanation for the volatility smile, they do present a problem. To be able to price options with no-arbitrage principles, one must be able to trade a claim on volatility or volatility must be uncorrelated to aggregate consumption.

Option Replication with Transaction Costs

In addition to assuming that the volatility of an asset is known and nonstochastic, Black and Scholes also assumed that one can trade securities with no transaction costs. Of course, markets are not perfect, and transaction costs do exist. Instead of examining whether deviations of option prices from the Black–Scholes model occur because of a nonconstant volatility, I investigate whether the deviations are consistent with the absence of arbitrage opportunities in the stock, bond, and options markets, given that transaction costs exist. If one can show that the volatility smile is sandwiched between the long and short no-arbitrage bounds, the smile may be consistent with the constant volatility Black–Scholes model, once transaction costs have been accounted for.

One of the first studies to examine the impact of proportional transaction costs on the price of an option was conducted by Leland (1985). Leland used a discrete-time stock–bond replication strategy and demonstrated that the payoff of this replication scheme converges almost surely to that of a European call option as the revision interval becomes infinitesimally small. The price of a European call option is determined with the Black–Scholes formula with an inflated variance, $\hat{\sigma}^2 = \sigma^2(1 + \frac{\sqrt{2/\pi\kappa}}{\sigma\sqrt{\Delta t}})$, when the transaction cost, κ , is large or the time interval between successive portfolio rebalancings, Δt , is small; the adjusted volatility is large; and the price of the option is relatively high. Although letting large amounts of time lapse between portfolio rebalancings reduces the cost of the option, it also increases the chance that the value of the replicating portfolio will significantly deviate from the desired payoff. Although Leland's approach is limited to derivative securities that have convex payoffs, Whaley and Wilmont (1993) presented an analysis similar to that of Leland that can be used to replicate concave and convex payoffs. Toft (1996) extended Leland's analysis by examining the mean-variance tradeoff between the variance of the error, which is reduced by more frequent rebalancing, and cost, which is increased by more frequent rebalancing.

Although Leland (1985) and Whaley and Wilmont (1993) arrived at a closed-form solution for the price of a call option, Merton (1990) took

a slightly different approach and modeled the stock price as a two-period binomial process. Allowing for proportional transaction costs when trading the stock, Merton then solved for the dynamic trading strategy that exactly replicated the payoff to a call option. Merton's analysis was extended to binomial trees of arbitrary length by Boyle and Vorst (1992). Using a binomial tree approach is useful for pricing American options and exotics such as Asian options.

Option replication with transaction costs was approached differently by Constantinides (1993). Constantinides modeled the stock price as a continuous-time geometric Brownian motion, used a proportional transaction cost when purchasing or selling stock, and assumed that the investor has either constant-relative-risk-aversion or hyperbolic-absolute-risk-aversion preferences. He then derived the investor's reservation buy or sell price for a call option on the stock. This analysis was generalized by Constantinides (1996) and Constantinides and Zariphopoulou (1999), for whom the reservation write or purchase price depended only on an investor's utility functions being monotonic and concave, not on a specific form of the utility function.

The price of the call option obtained with the models of Merton (1990), Boyle and Vorst (1992), Leland (1985), Constantinides (1996), Constantinides and Zariphopoulou (1999), and Whaley and Wilmont (1993) is suboptimal in three respects. First, the stock-bond replicating portfolio has to be rebalanced at each revision date. Because of the frequent rebalancing, the replication cost grows without bound as the revision interval becomes smaller. When transaction costs are present, it is sometimes optimal to not rebalance the replicating portfolio at every single revision date. Second, when transaction costs are present, it is sometimes less expensive to produce a higher payoff than is needed to exactly replicate the call option. In a constrained optimization framework, one can think of this approach as changing the constraint that the replicating portfolio exactly equals the payoff of the call option to a constraint in which the replicating portfolio equals or exceeds the payoff of the call option. Finally, these models do not allow for the use of other derivative securities that may be helpful in replicating the call option's payoff. As shown later, using traded puts and calls in the replicating portfolio, in addition to the stock and the bond, helps to lower the cost of replication. The option-pricing methodology described in the Methodology section overcomes these three problems.

In contrast to these models, Edirisinghe et al. (1993) and Bensaid, Lesne, Pagès, and Scheinkman (1992) showed that the cost of the replicating portfolio may be reduced with a replicating strategy of dominating

an option's payoff instead of exactly replicating the payoff. Payoff domination is a less expensive replication strategy when transaction costs are high because it may be cheaper to overhedge at certain times and not to rebalance the stock–bond portfolio at other times.²

Although options can be priced with a dynamic stock–bond replicating portfolio and a no-arbitrage argument, stock options can also be priced in a now and then economy if markets are complete. Although complete markets are a convenient theoretical construction, they do not exist in practice. In incomplete markets, where the number of states of nature exceeds the number of securities, state prices cannot be implied from observed prices, and stock options cannot be priced with static replication alone. Nevertheless, Ritchken (1985) used an LP to show that the upper and lower no-arbitrage price bounds can be computed for an option with price information from the underlying asset and a risk-free bond, assuming that there are no transaction costs and that the positions in the two securities are maintained without revision until the option's maturity date.³ In a related work, Cochrane and Saá-Requejo (2000) provided no-arbitrage bounds for option pricing in incomplete markets by choosing a stochastic discount factor to maximize and minimize the price of a call subject to the constraints that the stochastic discount factor is positive and that its volatility is less than or equal to some prespecified limit.

Prior empirical work on testing no-arbitrage bounds has focused on put–call parity, which is a special case of static replication. For example, Klemkosky and Resnick (1979) tested put–call parity and discovered some violations, but they did not account for transaction costs. Nisbit (1992) tested put–call parity and found that when bid–ask spreads and transaction costs are accounted for, there are no violations. None of these articles, however, has addressed the issue of whether the volatility smile is consistent with the absence of arbitrage opportunities within the context of a static–dynamic replication model. An exception is a recent study by Peña, Rubio, and Serna (1999), who provided some evidence that the shape of the volatility smile is related to transaction costs.

Given that markets are neither perfect nor complete, bounds for stock option prices can be obtained with either a dynamic stock–bond replication argument in the presence of transactions costs, as in Boyle and

²Bensaid et al. (1992) presented a simple two-period binomial example that nicely illustrated this point.

³In addition to the price information from the stock and risk-free bond, Ritchken's method can be easily extended to use the price information contained in other options.

Vorst (1992), or an incomplete market static replication argument, as in Ritchken (1985), Cochrane and Saá-Requejo (2000), and Garman (1976). Edirisinghe et al. (1993) and Dennis and Rendleman (1995) combined both these concepts in an LP framework and showed that it was possible to obtain even tighter no-arbitrage bounds for the option's price by the simultaneous use of both dynamic and static replication. This article uses the LP methodology to determine upper and lower no-arbitrage prices for the Standard & Poor's (S&P) 500 Index and examine whether the empirically documented volatility smile is consistent with the absence of arbitrage opportunities in the presence of transaction costs.

METHODOLOGY

Intuition Behind the LP

This section describes the LP technique used to formulate an optimal option replication strategy. The objective of the LP is to determine the lowest cost portfolio, consisting of stock, a risk-free bond, and traded options, that replicates the payoffs to a nontraded option. The LP combines both dynamic and static replication techniques. An example is provided here to convey the intuition. A detailed description of the model is provided in the appendix and in Edirisinghe et al. (1993) and Dennis and Rendleman (1995).

Suppose that there exists a risky asset that has a current price of \$100 and a volatility of 20% per annum and our goal is to replicate a 1-year call with a strike price of \$100. The evolution of the stock price is modeled as a two-period binomial tree with 0.5 year per period and is shown in Figure 1.




FIGURE 1

Two-period binomial tree on which the LP is based.

In addition to the stock, there also exists a risk-free asset that has a current value of \$1 and that earns a continuous return of 10% per annum. Its value at the end of the first period is $e^{(0.10)(0.5)} = 1.0513$, and its value at the end of the second period is $e^{(0.10)(1.0)} = 1.1052$. Finally, there is a 1-year call option available with a striking price of \$80, which currently costs \$28. A 1% transaction cost is incurred whenever the stock, the risk-free asset, or the call are traded.

Our objective is to minimize the up-front portfolio cost, which, if use Equation 16 and account for the 1% transaction cost, can be expressed as

$$101\Delta_{0,0}^{LA} - 99\Delta_{0,0}^{SA} + 1.01\Delta_{0,0}^{LB} - 0.99\Delta_{0,0}^{SB} + 28.28\Delta_{0,0}^{LO_1} - 27.72\Delta_{0,0}^{SO_1} \quad (1)$$

where $\Delta_{i,j}^{Lx}$ represents the change in the quantity of asset x held long (L) or short (S) after the stock has moved up i times and down j times.

The first set of constraints that must be specified are the linking constraints, which link the total position of the stock and risk-free asset in each state to the changes in the positions of these assets between adjacent states. For this problem, these constraints are

$$\Phi_{0,0}^{LA} - \Phi_{0,0}^{SA} = \Delta_{0,0}^{LA} - \Delta_{0,0}^{SA} \quad (2)$$

$$\Phi_{0,0}^{LB} - \Phi_{0,0}^{SB} = \Delta_{0,0}^{LB} - \Delta_{0,0}^{SB} \quad (3)$$

$$\Phi_{0,0}^{LO_1} - \Phi_{0,0}^{SO_1} = \Delta_{0,0}^{LO_1} - \Delta_{0,0}^{SO_1} \quad (4)$$

$$\Phi_{1,0}^{LA} - \Phi_{1,0}^{SA} = \Phi_{0,0}^{LA} - \Phi_{0,0}^{SA} + \Delta_{1,0}^{LA+} - \Delta_{1,0}^{SA+} \quad (5)$$

$$\Phi_{1,0}^{LB} - \Phi_{1,0}^{SB} = \Phi_{0,0}^{LB} - \Phi_{0,0}^{SB} + \Delta_{1,0}^{LB+} - \Delta_{1,0}^{SB+} \quad (6)$$

$$\Phi_{0,1}^{LA} - \Phi_{0,1}^{SA} = \Phi_{0,0}^{LA} - \Phi_{0,0}^{SA} + \Delta_{0,1}^{LA-} - \Delta_{0,1}^{SA-} \quad (7)$$

$$\Phi_{0,1}^{LB} - \Phi_{0,1}^{SB} = \Phi_{0,0}^{LB} - \Phi_{0,0}^{SB} + \Delta_{0,1}^{LB-} - \Delta_{0,1}^{SB-} \quad (8)$$

Constraint 2 sets the total initial position of the stock, $\Phi_{0,0}^{LA} - \Phi_{0,0}^{SA}$, equal to the amount of stock initially purchased, $\Delta_{0,0}^{LA} - \Delta_{0,0}^{SA}$. These two quantities must be the same because there was no stock already owned at the beginning of the replication program. Constraints 3 and 4 apply to the risk-free asset and traded option, respectively, and parallel Constraint 2. Constraint 5 sets the total amount of stock owned in state $\{1, 0\}$, $\Phi_{1,0}^{LA} - \Phi_{1,0}^{SA}$, equal to the amount of stock owned in state $\{0, 0\}$, $\Phi_{0,0}^{LA} - \Phi_{0,0}^{SA}$, plus the amount purchased in state $\{1, 0\}$, $\Delta_{1,0}^{LA+} - \Delta_{1,0}^{SA+}$. Constraint 6 is identical to Constraint 5, except that it applies to the risk-free asset. Constraints 7 and 8 for state $\{0, 1\}$ parallel Constraints 5 and 6 for state $\{1, 0\}$.

In addition to the linking constraints, the portfolio must be self-financing in state $\{1, 0\}$ and state $\{0, 1\}$. Before we can define the self-financing constraints, we must remember to account for transaction costs. For example, the cost of purchasing the stock in state $\{1, 0\}$ is $(115.19)(1.01) = 116.34$, and the cash inflow from shorting the stock is $(115.19)(0.99) = 114.04$. The following two constraints ensure that the replicating portfolio is self-financing in both the up and down states at time 1:

$$116.34\Delta_{1,0}^{LA+} - 114.04\Delta_{1,0}^{SA+} + 1.0618\Delta_{1,0}^{LB+} - 1.0408\Delta_{1,0}^{SB+} + S_{1,0}^+ = 0 \quad (9)$$

$$87.68\Delta_{0,1}^{LA-} - 85.94\Delta_{0,1}^{SA-} + 1.0618\Delta_{0,1}^{LB-} - 1.0408\Delta_{0,1}^{SB-} + S_{0,1}^- = 0 \quad (10)$$

Equation 9 constrains the replicating portfolio to be self-financing in state $\{1, 0\}$. The constraint ensures that the cost of purchasing stock at 116.34, less the cash inflow generated from shorting stock at 114.04, plus the cost of purchasing the risk-free asset at 1.0618, less the cash inflow of shorting the asset at 1.0408, is less than or equal to zero. Constraint 10, which applies to state $\{0, 1\}$, parallels Constraint 9, which applies to state $\{1, 0\}$. The variables $S_{1,0}^+$ and $S_{1,0}^-$ are nonnegative slack variables that transform Constraints 9 and 10 from inequalities into equalities.

To complete the LP, we must specify the constraints that ensure that, at maturity, the payoffs of the replicating portfolio meet or exceed those of the target. The following four constraints accomplish this:

$$131.36\Phi_{1,0}^{LA} - 134.02\Phi_{1,0}^{SA} + 1.0941\Phi_{1,0}^{LB} - 1.1162\Phi_{1,0}^{SB} + 52.16\Phi_{1,0}^{LO_1} - 53.22\Phi_{1,0}^{SO_1} - S_{2,0}^+ = 32.69 \quad (11)$$

$$99.00\Phi_{1,0}^{LA} - 101.00\Phi_{1,0}^{SA} + 1.0941\Phi_{1,0}^{LB} - 1.1162\Phi_{1,0}^{SB} + 19.80\Phi_{1,0}^{LO_1} - 20.20\Phi_{1,0}^{SO_1} - S_{1,1}^- = 0 \quad (12)$$

$$99.00\Phi_{0,1}^{LA} - 101.00\Phi_{0,1}^{SA} + 1.0941\Phi_{0,1}^{LB} - 1.1162\Phi_{0,1}^{SB} + 19.80\Phi_{0,1}^{LO_1} - 20.20\Phi_{0,1}^{SO_1} - S_{1,1}^+ = 0 \quad (13)$$

$$74.61\Phi_{0,1}^{LA} - 76.12\Phi_{0,1}^{SA} + 1.0941\Phi_{0,1}^{LB} - 1.1162\Phi_{0,1}^{SB} + 0\Phi_{0,1}^{LO_1} - 0\Phi_{0,1}^{SO_1} - S_{0,2}^- = 0 \quad (14)$$

Constraint 11 applies to state $\{2, 0\}$. The positive terms in 11 represent the after-transaction-cost cash inflow that is generated at time 2

from liquidating the long position in the stock at \$131.36, the risk-free asset at \$1.0941, and the call option at \$52.16. The negative terms in 11 represent the after-transaction-cost cash outflows from covering the short position in the stock at \$134.02, the risk-free asset at \$1.1162, and the call option at \$53.52. Because $S_{2,0}^+$ is a nonnegative slack variable, the total cash inflows must exceed the total cash outflows by at least \$32.69, which is the target replicating portfolio value in state $\{2, 0\}$. Although Constraint 11 ensures that the target value will be met or exceeded when the stock increases from state $\{1, 0\}$ to state $\{2, 0\}$, Constraints 12, 13, and 14 ensure that the target value will be met or exceeded when the stock price moves from state $\{1, 0\}$ to $\{1, 1\}$, state $\{0, 1\}$ to $\{1, 1\}$, and state $\{0, 1\}$ to $\{0, 2\}$, respectively.

Objective Function 1, along with the Constraints 2–14, constitutes an LP whose solution represents the optimal cost and optimal trading strategy for replicating a call option. In this particular case, the objective function is \$14.16 at the optimum.

DATA

The primary data set used is a subset of the Berkeley Options Data Base, which contains both quotes and trades recorded on the floor of the Chicago Board of Exchange from January 4, 1993 to April 30, 1993. Each record is time-stamped to the nearest second and contains the bid–ask quote, ticker symbol, date, time, type (put or call), strike price, and latest value of the underlying security. The options used are S&P 500 Index options, which are cash-settled and European in nature. The term structure at each point in time is constructed with bid and ask Treasury bill (T-bill) rates collected from *The Wall Street Journal*.

There are two concerns that must be addressed when the subset of options to be used for the static component of the replication is constructed: the depth of the market and the synchronicity of the observed option prices. First, it is important to have enough market depth so that trading in the options does not move the price significantly. Because most of the trading in index options occurs in the near-the-money options, only the three closest-to-the-money calls and puts are used in the replicating portfolio, for a maximum of six options. For ease of exposition, let Z represent the set of six options that meets the liquidity criteria described previously.

In addition to the liquidity criteria, it is also important that the option prices are observed at the same point in time. Although it would

be ideal to observe all option prices at the same point in time, t^* , in practice the prices are usually quoted at different points in time, and this nonsynchronicity can introduce error. To minimize this error, only those options in set Z that are closest to time t^* , without being more than 5 min away, are used in the static component of the replication.

Once the set of options to be used for the static component of the replication has been selected, the next step is to construct the binomial tree that will be used to determine the dynamic part of the replication. When the binomial tree is constructed, there are three issues that must be addressed: the treatment of dividends, the determination of the initial index level, and the determination of the volatility of the index.

Dividends on the S&P 500 Index must be treated with care. Harvey and Whaley (1992) showed that there exist strong seasonal patterns in the dividend payouts for the S&P 100 Index, with February, May, August, and November having the highest dividend payouts. They also showed that there is a strong day-of-the-week effect, with Monday having the highest dividend payout and Wednesday having the lowest. Because of these effects, dividends should not be treated as continuous but rather as discrete. Daily cash dividends for the S&P 500 Index were obtained from the S&P 500 Information Bulletin.⁴ Although the dividends are discrete, the binomial tree was adjusted for discrete dividends in such a way that it recombines.⁵

To model the evolution of the index value with a binomial tree, both the initial level of the index and the volatility of the index must be known. Although each record in the Berkeley Options Data Base contains the current index level, it may not represent the true level of the underlying index. This discrepancy arises because the index is a composite of prices of 500 stocks, some of which do not trade that frequently. Because of this stale quote problem, the quoted index level may not reflect the true market consensus of what the index value is. To mitigate this stale index problem, both the volatility and the level of the index are jointly implied from those options that meet both the liquidity and synchronicity criteria described previously by minimization of the sum-of-squared deviations between the market prices of the options and their model prices.⁶

⁴Robert E. Whaley graciously provided these data.

⁵See Hull (1997, pp. 354–356) for a discussion of this technique.

⁶The model prices are computed with the dividend-adjusted binomial tree and the term structure implied by T-bill prices from *The Wall Street Journal*.

RESULTS

Volatility Smile

Although full empirical tests are discussed in the next subsection, this subsection gives an example of how the volatility smile can be consistent with the absence of arbitrage opportunities in the presence of transactions costs. As long as the volatility for the traded option falls between that of the synthetic long and short positions, there will be no arbitrage opportunities involving the target security and any combination of securities in the replicating portfolio. If the volatility smile is sandwiched between the implied volatility bounds formed by the long and short synthetic positions, we can conclude that the volatility smile is consistent with the absence of arbitrage opportunities, once we have accounted for transaction costs.

Table I shows the market prices on January 4, 1993 at 2:00 p.m. A proportional transaction cost of 0.5% is used for the underlying index, and no transaction cost, other than the bid-ask spread, is used for the index options and risk-free asset. The transaction cost of 0.5% was chosen simply to illustrate this example. In the empirical tests in the next section, the transaction cost is continuously lowered until an arbitrage is possible. This level of transaction cost can then be compared to actual costs that traders incur to see if arbitrage is feasible at realistic transaction-cost levels.

TABLE I
Data Used in the Replication Program Example

Annual risk-free rate	2.97% per year			
Volatility of index	10.76% per year			
Index level	434.50			
Number of periods for binomial tree	8			
Time to expiration	46 days			
<i>Option Number</i>	<i>Bid</i>	<i>Ask</i>	<i>Strike</i>	<i>Type</i>
1	10.25	10.75	430.00	Call
2	7.13	7.25	435.00	Call
3	4.63	5.00	440.00	Call
4	7.25	7.75	435.00	Put
5	9.63	10.13	440.00	Put

Note. This table contains the data used to generate the results in Table II. The upper panel specifies the data necessary to construct the binomial lattice. The lower panel contains prices of near-the-money Standard & Poor's 500 Index options at 2:00 p.m. on January 4, 1993. The options mature in 46 days.

Table II shows the no-arbitrage volatility bounds computed for February calls with striking prices ranging from 420 to 450. The prices for the synthetic calls in the table are constructed with the data for the stock, bond, and options 1–5 in Table I. The upper and lower panels show replicating costs and implied volatilities for long and short positions, respectively. In each panel, the first column lists the strike price from 420 to 450. The second column contains the synthetic price of a long call with the stock, bond, and options 1–5 in Table I. The third column contains the bid price for the option that was quoted closest to 2:00 p.m. The fourth and fifth columns contain the implied volatility for the prices in the second and third columns.

There are several observations that can be made about the upper panel of Table II. First, the volatility smile is evident in the observed bid prices. The implied volatility is 11.25% at a strike of 420 and declines to a value of 10.33% at a strike of 450. Second, at most strike prices, the cost to synthetically create a long position is greater than the cash inflow

TABLE II
Synthetic Costs and Implied Volatilities as a Function of the Strike Price

Target Strike	Costs		Implied Volatilities (%)	
	Long Synthetic Call	Observed Bid	Long Synthetic Call	Observed Bid
420	18.222	17.63	12.70	11.25
425	14.400	14.00	12.56	11.77
430	10.781	10.25	12.03	11.11
435	7.793	7.13	11.79	10.71
440	5.148	4.63	11.21	10.35
445	2.870	2.88	10.23	10.25
450	1.958	1.75	10.80	10.33

Target Strike	Costs		Implied Volatilities (%)	
	Short Synthetic Call	Observed Ask	Short Synthetic Call	Observed Ask
420	16.691	18.38	8.44	13.08
425	13.292	14.63	10.32	13.01
430	9.433	10.75	9.67	11.98
435	6.751	7.25	10.09	10.90
440	4.281	5.00	9.77	10.96
445	2.010	3.25	8.57	10.94
450	0.903	2.00	8.18	10.90

Note. Long and short prices of synthetic options with various strike prices having 46 days to maturity are computed with the data in Table I. The linear program uses both a dynamic strategy with the stock and bond and a static strategy with options 1–5. Observed bid and ask prices for the corresponding traded options are also shown. Implied volatilities are obtained with the Black–Scholes formula.

that would be realized by shorting a traded option at the bid price. For example, it would cost \$18.222 (implied volatility of 12.70%) to synthesize a long call struck at 420, but only \$17.63 (implied volatility of 11.25%) could be generated by shorting the option in the market. Hence, there is no arbitrage opportunity. The one exception to this is the case where the strike price is 445. Here the synthetic cost of a long call is 2.870, but the call has a bid price of $2 \frac{7}{8}$ (2.875). This implies that even after a transaction cost of 0.5%, an arbitrage profit is possible. This may not be a true arbitrage opportunity, however, for two reasons. First, the option is so far out of the money that one may not be able to short the option at $2 \frac{7}{8}$ because of a thin market. Second, even if one can short the option at $2 \frac{7}{8}$ there will be error introduced into the hedge because the binomial model is only an approximation to the actual stochastic process followed by the stock. As long as the stock's price evolves according to the binomial model, the LP replication will be exact. If, however, the stock's price path deviates from the binomial lattice at some future date, hedging error will be introduced. The impact of this hedging error on the transaction cost necessary for arbitrage opportunities to exist is discussed in the Hedging Effectiveness subsection.

The lower panel of Table II is similar to the upper panel, except that it contains the costs of synthetic short positions and the ask prices of the traded options. As in the upper panel, the volatility smile is evident in the observed ask prices. Also, at each strike price, the cost to purchase a long option in the market is greater than the cash inflow that would be realized by the synthetic creation of a short position. Hence, the volatility smile is bounded below by the implied volatilities of the synthetic short positions. Because the volatility smile is bounded above by the implied volatilities of the synthetic long positions (except for the case in which the strike is 445) and is bounded below by the implied volatilities of the synthetic short positions, the volatility smile is consistent with the absence of arbitrage opportunities, once transaction costs have been accounted for.

Transaction Costs Necessary to Permit Arbitrage

This subsection discusses the minimum level of transaction costs such that the deviations from Black–Scholes model are still consistent with the constant volatility model and the absence of arbitrage. To determine these costs, the analysis outlined in the previous subsection could be repeated at successively lower transaction costs until an arbitrage opportunity was detected. This would be extremely time-consuming, as

the LP would have to be rerun several times at each level of transaction costs to recreate the bounds on the smile. An alternative approach to detecting arbitrage opportunities is to reformulate the LP to minimize the up-front cost of the replicating portfolio subject to the constraint that a nonnegative payoff be generated at each state at the maturity of the replicating portfolio. If a portfolio can be found that satisfies the non-negativity constraint at the replicating portfolio's maturity and has a negative price today (i.e., a cash inflow), an arbitrage opportunity has been detected. At high levels of transaction costs, no arbitrage opportunities exist. As transaction costs fall, a point will be reached at which an arbitrage opportunity is possible. This method is more efficient because the LP only has to be rerun once at each transaction-cost level.

Table III shows the results from 20 replicating portfolios that were constructed at 1 p.m. on each trading day during January 1993 and the

TABLE III
Transaction Costs Needed to Permit Arbitrage: January 1993

<i>Date of Option Quote</i>	<i>Days to Expiration</i>	<i>Break-Even Transaction Cost</i>	<i>M Hedging Error</i>	<i>SD of Hedging Error</i>	<i>Ratio of M to SD</i>
930104	11	0.0003777	10	5,240	0.00
930105	10	0.0001236	731	5,181	0.14
930106	44	0.0000206	322	1,046	0.31
930107	43	0.0008995	77	4,491	0.02
930108	42	0.0007141	524	4,614	0.11
930111	39	0.0016479	234	1,792	0.13
930112	38	0.0005630	900	3,584	0.25
930113	37	0.0005630	405	3,433	0.12
930114	36	0.0010506	291	2,438	0.12
930115	35	0.0010025	290	3,252	0.09
930118	32	0.0000961	-131	1,631	-0.08
930119	31	0.0007484	143	1,207	0.12
930120	30	0.0010506	272	2,694	0.10
930121	29	0.0007690	275	2,273	0.12
930122	28	0.0009132	2,115	21,503	0.10
930125	25	0.0000893	254	1,940	0.13
930126	24	0.0008446	46	1,568	0.03
930127	23	0.0004875	435	5,834	0.07
930128	22	0.0001099	146	1,081	0.14
930129	21	0.0000412	151	1,162	0.13
<i>M</i>		0.0005756	374	3,798	0.11

Note. This table presents the levels of transaction costs necessary to permit arbitrage for portfolios with cross sections of Standard & Poor's 500 Index options formed daily at 1 p.m. during January 1993. Arbitrage is possible when a cash inflow can be produced today with no future liability. The transaction costs presented in the table are the proportional one-way costs on the underlying asset. There are no transaction costs on the risk-free asset and traded options other than the bid-ask spread.

transaction-cost levels at which arbitrage opportunities exist. The options in the replicating portfolio were selected according to the liquidity and synchronicity criteria outlined in the Data section. There is no transaction cost placed on the risk-free asset, other than the bid-ask spread obtained from *The Wall Street Journal*, and there is no transaction cost placed on the traded options, other than the bid-ask spread obtained from the Berkeley Options Data Base. The level of transaction cost at which arbitrage opportunities exist is determined by the setting of the initial transaction cost on the index to a very high level and the lowering of the cost until an arbitrage opportunity is detected.⁷ This transaction cost is then compared to the bid-ask spread and commission costs for trading the index.

The first three columns in Table III list the date that the replicating portfolio was constructed, the number of days to maturity for the replicating portfolio, and the break-even transaction cost. The break-even transaction cost represents the largest cost at which arbitrage is feasible within the context of the binomial model. For example, if, on January 4, 1993, a trader incurred a one-way transaction cost of 3.78 basis points or lower, there would be an arbitrage opportunity. The average transaction cost at which arbitrage is feasible is roughly six basis points. It is useful to compare these costs to costs that would actually be incurred by traders. The typical bid-ask spread was around two to three ticks (0.10–0.15) on S&P 500 Index futures in 1993,⁸ and the commission per contract was roughly \$30. This gives a one-way transaction cost of around $\frac{(500)(0.05)(2.5/2) + 30}{(430)(500)}$, or roughly three basis points per trade. To see if the results are robust through time, I repeated the arbitrage analysis previously presented for observations during January 1995. The results are presented in Table IV. The results for January 1995 indicate that the average transaction cost at which arbitrage is feasible is roughly five basis points, which is on the same order of magnitude as the transaction cost for January 1993.

Although it appears from these data that arbitrage is economically feasible, a trader may not be able to take advantage of these opportunities. Because the model is a binomial approximation to the continuous-time process being used, these arbitrage positions may not be entirely risk-free. As discussed in the next subsection, the transaction costs have to be much lower than three basis points to have a reasonably low-risk arbitrage opportunity.

⁷Varying the transaction cost on all three types of securities complicates matters because there are three choice variables instead of one.

⁸This figure was obtained from the Chicago Mercantile Exchange.

TABLE IV
Transaction Costs Needed to Permit Arbitrage: January 1995

<i>Date of Option Quote</i>	<i>Days to Expiration</i>	<i>Break-Even Transaction Cost</i>
950103	17	0.0006592
950104	16	0.0008720
950105	15	0.0011055
950106	14	0.0007278
950109	11	0.0006798
950110	10	0.0000755
950111	37	0.0019089
950113	35	0.0005219
950116	32	0.0001373
950117	31	0.0004532
950118	30	0.0003708
950119	29	0.0005424
950120	28	0.0006386
950123	25	0.0007280
950124	24	0.0002609
950125	23	0.0001991
950126	22	0.0001785
950127	21	0.0000481
950130	18	0.0000412
950131	17	0.0000961
<i>M</i>		0.0005122

Note. This table presents the levels of transaction costs necessary to permit arbitrage for portfolios with cross sections of Standard & Poor's 500 Index options formed daily at 1 P.M. during January 1995. Arbitrage is possible when a cash inflow can be produced today with no future liability. The transaction costs presented in the table are the proportional one-way costs on the underlying asset. There are no transaction costs on the risk-free asset and traded options other than the bid-ask spread.

Hedging Effectiveness

In the previous subsection, it was shown that arbitrage opportunities could be detected if we constrain the replicating portfolio to produce a nonnegative payoff in each state and then gradually lower the transaction-cost levels to discover at what point the up-front replicating portfolio has a negative price. One of the problems with this is that the LP model is necessarily based on the assumption that the price of the underlying stock follows a binomial process when, in fact, the binomial model is being used as an approximation to a continuous lognormal distribution. The arbitrage trading is guaranteed to work perfectly if the actual stock price outcomes fall exactly on the nodes of the binomial tree. If, however, the stock price outcomes do not fall exactly on the binomial tree, as is the case if the true stock price is lognormally distributed, any arbitrage opportunities that are detected may not be risk-free. In this subsection, Monte Carlo

simulations are used to determine the magnitude of the risk introduced by the binomial approximation.

It is important to recognize that the LP solution based on the binomial model determines not only the initial set of portfolio holdings for the replicating portfolio but also the optimal holdings in each subsequent binomial state. In the simulations, it is assumed that the investor expects stock returns to follow binomial outcomes and also expects to resolve the LP at the end of each binomial period on the basis of future anticipated binomial outcomes. In fact, this procedure is implicit in the optimal binomial-based LP solution. Actually, however, stock return outcomes are generated from a normal distribution. At the end of the first binomial period, the actual stock price is likely to be different from either of the two binomial outcomes that had been originally anticipated and reflected in the initial LP solution. As such, the LP, based on $T - 1$ remaining periods, must be resolved with the simulated stock price realized at the end of the first binomial period and with the security positions carried forward from time zero. Because the binomial model is being used as an approximation of a continuous-time stochastic process, most of the time the random return will not fall exactly on the up or down nodes of the binomial tree. Hence, the replicating portfolio will most likely not be self-financing, and cash will have to be added to, or withdrawn from, the replicating portfolio at the end of the first binomial period.

Having solved for the optimal replicating portfolio at the end of the first period, we then realize a random return for the stock at the end of the second binomial period, a $T - 2$ period binomial tree is constructed, and the process is repeated. This algorithm is continued until the maturity of the option being replicated is reached. The sum of the future value of the cash infusions or withdrawals at each time represents a hedging error for some price path realization from the continuous-time stochastic process. By drawing M random sample paths from the true stochastic process and computing the hedging error for each sample path, we can determine the hedging precision of the LP model, defined as the standard deviation of the future value of the M hedging errors. To generate a series of normally distributed random returns for the stock, uniform random deviates are generated with the RAN2 routine from Press, Teukolsky, Vetterling, and Flannery (1992). This random number generation routine provides random numbers that have a periodicity of about 2.3×10^{18} . The uniform random deviates are transformed into normal random numbers via the Box-Muller transformation by the GASDEV routine from Press et al. (1992).

The initial replicating portfolio, which may consist of stock, bonds, and options, is constrained to produce a zero payoff in each of the 13 possible states at the end of the binomial lattice. If the LP is used as it is formulated in the Intuition Behind the LP section, an arbitrage opportunity will result in an unbounded objective function.⁹ If the purpose of the LP is only to detect the presence of an arbitrage opportunity, as in the previous subsection, the fact that the solution is unbounded does not matter. If one is interested in the composition of the optimal arbitrage portfolio, however, the solution to LP cannot be unbounded. This is prevented by the addition of a constraint to the LP to ensure that the value of the objective function cannot be lower than -1.00 , thereby ensuring that the solution to the LP will be bounded. This constraint also standardizes the initial arbitrage profit to be $\$1.00$ for various levels of transaction costs because, when an arbitrage opportunity is present, the constraint will be binding.

We can generate lognormally distributed stock price outcomes by starting with the initial value of the asset and generating subsequent values for the asset with $S_{t+1} = s_t e^{(r_f + \sigma \tilde{z} \sqrt{\tau})}$, where $\tilde{z} \sim N(0,1)$ and r_f is the risk-free rate determined from the midpoint of the T-bill bid-ask spread observed at time t . The volatility used for the asset's returns is the implied volatility described in the Data section. The transaction cost on the underlying asset is taken from Column 3 in Table III, which is the largest cost where arbitrage is still possible.

The fourth and fifth columns of Table III contain the mean and standard deviation of the hedging error after 100 simulated price paths. The magnitude of the standard deviation of the hedging error is quite large. To understand this, one must consider that to realize $\$1$ of profit at the largest transaction-cost level at which arbitrage is still possible, very large positions must be taken in the underlying securities. This is because transaction costs consume most of the profit. Typical initial positions are shown in Table V. The first column shows the positions that must be taken when the transaction costs are just below the break-even point. As can be seen from the first column, the magnitudes of the positions are quite large. Because a volume of 10,000 for an S&P 500 Index option is pretty heavy, there is not even enough liquidity to make some of the positions feasible at the observed prices.

⁹To see this, let us say that there exists some portfolio that satisfies the constraints and has a value of $-\pi$ today. Doubling the portfolio positions will result in a value today of -2π yet still satisfy all the constraints. Because the objective is to minimize the up-front portfolio cost, the value of the objective function at the optimum will be $-\infty$.

TABLE V
Initial Replicating Portfolio Positions: January 4, 1993

	<i>At Break-Even Transaction Cost</i>	<i>10% Below Break-Even Transaction Cost</i>	<i>50% Below Break-Even Transaction Cost</i>
Total stock	-7247	-81	-59
Total bond	3,146,657	-36,628	26,042
Option 1 value	0	308	22
Option 2 value	4102	0	0
Option 3 value	-128,263	-4,326	1,082
Option 4 value	61,426	2,607	500
Option 5 value	41	0	0

Note. This table presents the magnitude of the initial positions in the stock, bond, and options needed to realize the arbitrage opportunity at various levels of transaction costs. A negative sign indicates a short position in that particular security.

TABLE VI
Summary Statistics for Total Hedging Errors Associated with 12-Period
Binomial Approximation

<i>Transaction-Cost Level</i>	<i>Average Transaction Cost</i>	<i>M Error (\$)</i>	<i>SD of Error (\$)</i>	<i>Percent Negative</i>	<i>Average Loss if Negative (\$)</i>
AT BETC	0.0005756	374	3,798	51	-1,629
10% below BETC	0.0005181	101	1,394	49	-876
50% below BETC	0.0002878	85	623	43	-257

Note. This table contains summary statistics for the Monte Carlo simulation results for the 20 trading days in January 1993. For each day, 100 price paths were generated at the break-even transaction cost (BETC) and at costs of 10 and 50% below this level. The percent negative column represents the percentage of time that the trading strategy lost money (a negative hedging error), and the average loss column is the dollar loss conditional on the hedging error being negative.

The last column of Table III contains the ratio of the mean error to the standard deviation. Although the replicating strategy produces a positive mean hedging error, the magnitude of the mean hedging error is small compared with the magnitude of the standard deviation, indicating that there is a good chance of large losses. The first row of Table VI summarizes the data in Table III and contains two additional statistics. The first is the percentage of time that the hedging error is negative. In the case in which the transaction cost is just below the break-even point, the arbitrage strategy loses money roughly half of the time. The second is the dollar loss conditional on the hedging error being negative. At the break-even transaction-cost level, the average dollar loss is \$1,629. This is a large risk to make \$1 in arbitrage profit.

It is natural to ask if the situation can be improved if a trader incurs a lower transaction cost. Two observations can be made. First, the size of the

positions required for arbitrage is lower at the reduced transaction cost. Table V shows the positions required for arbitrage on a typical day when transaction costs are 10 and 50% below the break-even level. Although relatively smaller positions are required at these lower transaction-cost levels, they are still large compared with the typical daily volume in the S&P 500 Index options market. Second, although the arbitrage transactions are relatively less risky at these lower transaction-cost levels, they are still quite risky on an absolute basis. The second row of Table VI shows the summary statistics for 100 price path simulations on each day of January 1993. The results were obtained in the same manner as those in the first row of the table, except that the transaction cost used for each of the simulations is 10% less than the break-even transaction cost for that particular day. Although the standard deviation of the error is reduced dramatically, it is still much larger than the mean. The percentage of time that a trader loses money does not change and is still around 50%. Furthermore, the dollar loss conditional on the hedging error being negative is \$867, which is still quite large compared with the \$1 of arbitrage profit. The last row contains the results when the transaction cost is lowered by 50% from the break-even value to three basis points. Even at this level of transaction costs, there is still a large risk, with a loss occurring 43% of the time and averaging \$257.

This analysis suggests that although arbitrage may at first appear feasible, it is by no means risk-free. Although arbitrage opportunities were detected at a transaction-cost level averaging six basis points, the positions in the replicating portfolio have to be so large to realize \$1 in profit that a small amount of model misspecification can induce very large errors. Even as the transaction costs are lowered to 50% less than the break-even value, any attempt at arbitrage is still very risky.

CONCLUSION

This article shows that the deviations of option prices from those given by a constant-volatility option model, such as the Black-Scholes model, can be consistent with the absence of arbitrage in the presence of transaction costs. Specifically, an LP is used to determine the optimal dynamic and static replication method to synthesize long and short positions for an S&P 500 Index option. As long as the bid price is above the synthetic short price and the ask price is below the synthetic long price, there is no way to make an arbitrage profit. If, however, either of these conditions is violated, it is possible to make an arbitrage profit, although the arbitrage opportunity may not be entirely risk-free. The use of both dynamic and

static replication allows for stronger tests of the existence of arbitrage opportunities than either one alone.

With a 12-period binomial model used to capture the price dynamics of a continuous-time process, one-way transaction costs, inclusive of the bid-ask spread, would have to be lower than six basis points for the volatility smile to be inconsistent with the absence of arbitrage. Although an arbitrage opportunity may be detected at this transaction-cost level, it may not be risk-free because of model misspecification. Monte Carlo simulations are used to assess the risk induced with a binomial tree to approximate a continuous-time process. Although an arbitrage opportunity may be present at a transaction-cost level of six basis points, such large positions have to be taken to realize a \$1 profit that the arbitrage is very risky. Even at transaction-cost levels lower than three basis points, which is roughly what traders would incur, there are no reasonably risk-free arbitrage opportunities.

APPENDIX: LP DETAILS

This appendix describes the LP technique used to formulate an optimal option replication strategy. The objective of the LP is to determine the lowest cost portfolio consisting of stock, a risk-free bond, and traded options that replicates the payoffs to a nontraded option. The LP combines both dynamic and static replication techniques.¹⁰

The uncertainty in the future stock price is represented with a binomial model as in Cox, Ross, and Rubinstein (1979) and Rendleman and Bartter (1979). Figure 2 shows a three-period binomial price process. Let $A_{0,0}$ denote the current price of the underlying stock and $A_{i,j}$ denote the price of the underlying stock after it has increased in value i times and decreased in value j times. Let τ represent the length of one binomial period in years, T represent the total number of binomial periods, and σ represent the annual volatility of the stock's returns. The future price of the stock after i ups and j downs is given by $A_{i,j} = A_{0,0}e^{\sigma(i-j)\sqrt{\tau}}$. Assume that a risk-free asset exists that has an annual continuously compounded return of r per year. Let its value at time zero be denoted $B_{0,0}$ and its value in state i,j be denoted $B_{i,j}$. The future value of the price of the risk-free asset is given by $B_{i,j} = B_{0,0}e^{r(i+j)\tau}$.

Although the LP methodology can accommodate both fixed and proportional transaction costs, the model described here will focus on

¹⁰See Edirisinghe et al. (1993) and Dennis and Rendleman (1995).




FIGURE 2

Example of a three-period binomial model of the stock price evolution.

proportional costs only. Define ρ_A as the proportional cost of trading the underlying stock, ρ_B as the proportional cost of trading the risk-free asset, and N_O as the proportional cost of trading an option on the underlying stock. Assume that there are N_O options that trade in secondary markets available for inclusion in the replicating portfolio. Let the symbol $\kappa_{i,j}^{LA}(\kappa_{i,j}^{SA})$ represent the total cash outflow (inflow), after all transaction costs, of going long (short) one share of the stock in state $\{i, j\}$. The cash outflow (inflow) of going long (short) one share of the risk-free asset is represented by $\kappa_{i,j}^{LB}(\kappa_{i,j}^{SB})$, and the cash outflow (inflow) of going long (short) option k , which costs $O_{i,j}^k$, is represented by $k_{i,j}^{LO_k}(k_{i,j}^{SO_k})$, where each option is written on one share of stock.

In the LP, it is assumed that options that are in the replicating portfolio can be traded only at time zero and at their maturity (T_k). At these two times, the prices of the options are known with certainty. If any options are mispriced, however, there is no way to know in advance how they will be priced in any binomial state prior to maturity. Therefore, to avoid needing to estimate the option's price at intermediate times and needing to take into account convergence to equilibrium pricing, we assume that all option positions are maintained until maturity without revision.

To account for this aspect of trading in the LP, let $\theta_{i,j}^k$ be an indicator variable that takes on a value of one when option k is available for trade at time zero and T_k and zero otherwise:

$$\theta_{i,j}^k = \begin{cases} 1 & \text{if } i + j = 0 \text{ or } i + j = T_k \\ 0 & \text{otherwise} \end{cases}$$

By multiplying $\theta_{i,j}^k$ by the cost of going long or short an option, we can write the LP compactly; when $\theta_{i,j}^k$ is multiplied by the after-transaction-cost

option price, the option will not appear in the LP at any time other than time zero and T_k .

We now define the cash outflow (inflow), the net of transaction costs, of going long (short) one share of the stock, one share of the risk-free asset, or one option of type k after the underlying stock has increased in value i times and decreased in value j times:

$$\begin{aligned}
 k_{i,j}^{LA} &= A_{i,j}(1 + \rho_A) \\
 k_{i,j}^{SA} &= A_{i,j}(1 - \rho_A) \\
 k_{i,j}^{LB} &= B_{i,j}(1 + \rho_B) \\
 k_{i,j}^{SB} &= B_{i,j}(1 - \rho_B) \\
 k_{i,j}^{LO_k} &= \theta_{i,j}^k 0_{i,j}^k (1 + \rho_O) \\
 k_{i,j}^{SO_k} &= \theta_{i,j}^k 0_{i,j}^k (1 - \rho_O)
 \end{aligned} \tag{15}$$

In the LP, separate decision variables are needed to denote long and short positions in each asset because LP requires that all decision variables be nonnegative.

The decision variables $\Phi_{i,j}^{LA}$ and $\Phi_{i,j}^{SA}$ denote the total number of shares of the risky asset held long (short) in state $\{i,j\}$. The variables $\Phi_{i,j}^{LB}$ and $\Phi_{i,j}^{SB}$ are identical to $\Phi_{i,j}^{LA}$ and $\Phi_{i,j}^{SA}$ except that they apply to the risk-free asset. Similarly, the variables $\Phi_{i,j}^{LO_k}$ and $\Phi_{i,j}^{SO_k}$ denote the total number of the k th option held long (short). In addition to determining the levels of each asset in each binomial state, the LP must determine the changes in the positions of each asset between binomial states. The decision variables $\Delta_{i,j}^{LA}$ and $\Delta_{i,j}^{SA}$ represent an increase in the long (short) position in the stock in state $\{i,j\}$, $\Delta_{i,j}^{LB}$ and $\Delta_{i,j}^{SB}$ represent an increase in the long (short) position in the risk-free asset, and $\Delta_{i,j}^{LO_k}$ and $\Delta_{i,j}^{SO_k}$ represent an increase in the long (short) position of the k th option. An additional superscript of $+$ or $-$ is used in connection with all variables for all states for which $i + j > 0$ to indicate that the last stock price move was an increase ($+$) or a decrease ($-$). For example, $\Delta_{i,j}^{LA+}$ denotes the increase in the long position of the stock in state $\{2, 1\}$ where the last move was an up. This implies that the prior state had to have been $\{1, 1\}$, not $\{2, 0\}$.

In the LP, the objective is to minimize the initial cost of the replicating portfolio. The relevant decision variables for establishing the

up-front cost of the replicating portfolio are the change variables (Δ). Therefore, the objective function can be stated as

MINIMIZE

$$\begin{aligned} & \Delta_{0,0}^{LA} \kappa_{0,0}^{LA} - \Delta_{0,0}^{SA} \kappa_{0,0}^{SA} + \Delta_{0,0}^{LB} \kappa_{0,0}^{LB} - \Delta_{0,0}^{SB} \kappa_{0,0}^{SB} \\ & + \sum_{\kappa=1}^{N_O} (\Delta_{0,0}^{LO_k} \kappa_{0,0}^{LO_k} - \Delta_{0,0}^{SO_k} \kappa_{0,0}^{SO_k}) \end{aligned} \quad (16)$$

The first term in the objective function is the cost of establishing a long position in the underlying stock. The second term represents the cash inflow from establishing a short position in the underlying stock. The third and fourth terms are the costs and cash inflows of establishing long and short positions in the risk-free asset, respectively. The terms in the summation represent the cost of establishing positions in the N_O options. The cost of establishing the replicating portfolio is simply the sum of the size of the long (short) position of each security in the portfolio multiplied by the appropriate unit cost (cash inflow) of acquiring the position.

There are three types of constraints that the solution to the LP must satisfy. First, the portfolio must be self-financing, so that when the portfolio is rebalanced, there is enough income from the sale of existing securities held in the portfolio at time t to cover both the purchase of new securities at time $t + 1$ and all transaction costs. The following two constraints ensure that this is the case:

$$\begin{aligned} & \Delta_{i,j}^{LA+} \kappa_{i,j}^{LA} - \Delta_{i,j}^{SA+} \kappa_{i,j}^{SA} + \Delta_{i,j}^{LB+} \kappa_{i,j}^{LB} - \Delta_{i,j}^{SB+} \kappa_{i,j}^{SB} \\ & - \sum_{k=1}^{N_O} (\Delta_{i,j}^{LO_k} \kappa_{i,j}^{SO_k} - \Delta_{i,j}^{SO_k} \kappa_{i,j}^{LO_k}) + S_{i,j}^+ = 0 \end{aligned} \quad (17)$$

$$\begin{aligned} & \Delta_{i,j}^{LA-} \kappa_{i,j}^{LA} - \Delta_{i,j}^{SA-} \kappa_{i,j}^{SA} + \Delta_{i,j}^{LB-} \kappa_{i,j}^{LB} - \Delta_{i,j}^{SB-} \kappa_{i,j}^{SB} \\ & - \sum_{k=1}^{N_O} (\Delta_{i,j}^{LO_k} \kappa_{i,j}^{SO_k} - \Delta_{i,j}^{SO_k} \kappa_{i,j}^{LO_k}) + S_{i,j}^+ = 0 \end{aligned} \quad (18)$$

The first term in Constraint 17 represents the total cost of financing a change in the long position of the underlying stock. Note that terms that are greater than zero represent a cash outflow, whereas terms less than zero represent cash inflow. The second term represents the cash inflow generated by shorting $\Delta_{i,j}^{SA+}$ shares of the underlying stock. The third (fourth) term represents the cash outflow (inflow) of increasing the long (short) position in the risk-free asset. The terms inside the summation represent the unwinding of option positions at maturity that were

established at time zero. Long options provide income of κ^{LO_k} per option, whereas short options have to be covered at a cost of κ^{SO_k} per option. The final term, $S_{i,j}^+$, represents the amount by which cash inflow exceeds cost and converts the inequality to an equality. Constraint 17 applies to the case in which the stock increased in value from time $i + j - 1$ to time $i + j$. Constraint 18 is identical to 17 except that it pertains to the case in which the stock's value decreased from time $i + j - 1$ to time $i + j$. Constraints 17 and 18 must hold for all $i, j \mid i + j < T$ and $i + j > 0$.

The second set of constraints are called linking constraints because they link changes in the positions of each security between consecutive trading times (Δ) to the total holdings of that security at one particular time (Φ). These constraints simply state that the total number of shares of a security at time $t = i + j$ must be the total number of shares in the previous state at time $t = i + j - 1$ plus any changes made to that position at time $t = i + j$. Linking constraints are not needed for the options because no changes are made in the option positions after time zero.

The linking constraints for the LP are

$$\Phi_{0,0}^{LA} - \Phi_{0,0}^{SA} = \Delta_{0,0}^{LA} - \Delta_{0,0}^{SA} + I_{0,0}^A \quad (19)$$

$$\Phi_{0,0}^{LB} - \Phi_{0,0}^{SB} = \Delta_{0,0}^{LB} - \Delta_{0,0}^{SB} + I_{0,0}^B \quad (20)$$

$$\Phi_{0,0}^{LO_k} - \Phi_{0,0}^{SO_k} = \Delta_{0,0}^{LO_k} - \Delta_{0,0}^{SO_k} + I_{0,0}^k \quad (21)$$

$$\Phi_{i,j}^{LA} - \Phi_{i,j}^{SA} = \Phi_{i-1,j}^{LA} - \Phi_{i-1,j}^{SA} + \Delta_{i,j}^{LA+} - \Delta_{i,j}^{SA+} \quad (22)$$

$$\Phi_{i,j}^{LB} - \Phi_{i,j}^{SB} = \Phi_{i-1,j}^{LB} - \Phi_{i-1,j}^{SB} + \Delta_{i,j}^{LB+} - \Delta_{i,j}^{SB+} \quad (23)$$

$$\Phi_{i,j}^{LA} - \Phi_{i,j}^{SA} = \Phi_{i,j-1}^{LA} - \Phi_{i,j-1}^{SA} + \Delta_{i,j}^{LA-} - \Delta_{i,j}^{SA-} \quad (24)$$

$$\Phi_{i,j}^{LB} - \Phi_{i,j}^{SB} = \Phi_{i,j-1}^{LB} - \Phi_{i,j-1}^{SB} + \Delta_{i,j}^{LB-} - \Delta_{i,j}^{SB-} \quad (25)$$

The left-hand side of Constraint 19 represents the initial net position of the underlying stock, and the right-hand side represents the net number of shares purchased. The parameter $I_{0,0}^A$ is the number of shares already owned at time zero and, therefore, is not a decision variable for the LP. It is necessary to include the number of shares already owned if one uses the LP to optimally rebalance a replicating portfolio when the portfolio already contains some amount of the underlying stock and risk-free bonds.

Constraint 20, which applies to the risk-free asset, and Constraint 21, which applies to the k th traded option, parallel Constraint 19, which applies to the risky asset. Constraint 22 applies to an upward move in the stock price. The left-hand side of Constraint 22 represents the net position

in the underlying stock in state $\{i, j\}$, which must be equal to the number of shares in the previous state, $\Phi_{i-1,j}^{LA} - \Phi_{i-1,j}^{SA}$, plus the net change in position due to trading in state $\{i, j\}$, $\Delta_{i-1,j}^{LA+} - \Delta_{i-1,j}^{SA+}$. Constraint 23 for the risk-free asset parallels Constraint 22 for the risky (underlying) stock. Constraints 24 and 25 are identical to 22 and 23 except that they link the total positions in adjacent states when the stock price moves down in price. Constraints 22–25 must hold for all $i, j \mid i + j < T$ and $i + j > 0$.

The third constraint that the LP needs to satisfy is that the payoffs of the replicating portfolio must equal or exceed the payoffs of the security being replicated (the target). Let $G_{i,j}$ represent the value of the target security in state $\{i, j\} \mid i + j = T$. Constraint 26 applies to the case in which the stock's value increases from time $i + j - 1$ to time $i + j$. The first (second) term in 26 is the cash flow from liquidating (covering) the total long (short) position in stock A. The third (fourth) term in 26 is the cash flow from liquidating (covering) the total long (short) position in risk-free security B. The term inside the summation represents the cash flow from liquidating the k th option. The final term, $S_{i,j}^+$, represents the amount by which the value of the replicating portfolio exceeds that of the target in state $\{i, j\}$. Constraint 27 is identical to 26, except that it pertains to the case in which the stock's value decreases from time $i + j - 1$ to time $i + j$. Both Constraints 26 and 27 must hold for all $i, j \mid i + j = T$:

$$\begin{aligned} & \Phi_{i-1,j}^{LA} \kappa_{i,j}^{SA} - \Phi_{i-1,j}^{SA} \kappa_{i,j}^{LA} + \Phi_{i-1,j}^{LB} \kappa_{i,j}^{SB} - \Phi_{i-1,j}^{SB} \kappa_{i,j}^{LB} \\ & + \sum_{k=1}^{N_O} (\Phi_{i,j}^{LO_k} \kappa_{i,j}^{SO_k} - \Phi_{i,j}^{SO_k} \kappa_{i,j}^{LO_k}) - S_{i,j}^+ = G_{i,j} \end{aligned} \quad (26)$$

$$\begin{aligned} & \Phi_{i,j-1}^{LA} \kappa_{i,j}^{SA} - \Phi_{i,j-1}^{SA} \kappa_{i,j}^{LA} + \Phi_{i,j-1}^{LB} \kappa_{i,j}^{SB} - \Phi_{i,j-1}^{SB} \kappa_{i,j}^{LB} \\ & + \sum_{k=1}^{N_O} (\Phi_{i,j}^{LO_k} \kappa_{i,j}^{SO_k} - \Phi_{i,j}^{SO_k} \kappa_{i,j}^{LO_k}) - S_{i,j}^- = G_{i,j} \end{aligned} \quad (27)$$

The objective function, along with Constraints 19–27, constitutes an LP whose solution provides the lowest cost strategy for replicating the payoffs to an option.

BIBLIOGRAPHY

- Bensaid, B., Lesne, J.-P., Pages, H., & Scheinkman, J. (1992). Derivative asset pricing with transaction costs. *Mathematical Finance*, 2, 63–86.
- Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81, 637–654.

- Boyle, P. P., & Vorst, T. (1992). Option replication in discrete time with transaction costs. *Journal of Finance*, 47, 271–293.
- Cochrane, J. H., & Saá-Requejo, J. (2000). Beyond arbitrage: "Good-deal" asset price bounds in incomplete markets. *Journal of Political Economy*, 108, 79–119.
- Constantinides, G. M. (1993). Option pricing bounds with transaction costs (University of Chicago Working Paper 384). Chicago: University of Chicago.
- Constantinides, G. M. (1996). Transaction costs and the implied volatility smile (University of Chicago Working Paper 442). Chicago: University of Chicago.
- Constantinides, G. M., & Zariphopoulou, T. (1999). Bounds on prices of contingent claims in an intertemporal economy with proportional transactions costs and general preferences. *Finance and Stochastics*, 3, 345–369.
- Cox, J., Ross, S., & Rubinstein, M. (1979). Option pricing: A simplified approach. *Journal of Financial Economics*, 7, 1979, 229–263.
- Dennis, P., & Rendleman, R. J., Jr. (1995). An LP approach to synthetic option replication with transactions costs and multiple security selection. *Advances in Futures and Options Research*, 8, 53–84.
- Derman, E., & Kani, I. (1994). Riding on a smile. *Risk*, 7, 32–39.
- Dumas, B., Fleming, J., & Whaley, R. E. (1998). Implied volatility functions: Empirical tests. *Journal of Finance*, 53, 2059–2106.
- Edirisinghe, C., Naik, V., & Uppal, R. (1993). Optimal replication of options with transactions costs and trading restrictions. *Journal of Financial and Quantitative Analysis*, 28, 117–138.
- Garman, M. (1976). An algebra for evaluating hedge portfolios. *Journal of Financial Economics*, 3, 403–427.
- Harvey, C. R., & Whaley, R. E. (1992). Dividends and S&P 100 Index option valuation. *Journal of Futures Markets*, 12, 123–137.
- Heath, D., Jarrow, R., & Morton, A. (1992). Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica*, 60, 77–105.
- Heston, S. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6, 327–344.
- Ho, T., & Lee, S.-B. (1986). Term structure movements and pricing interest rate contingent claims. *Journal of Finance*, 5, 1011–1029.
- Hull, J. C. (1997). *Options, futures, and other derivatives* (3rd ed.). Upper Saddle River, NJ: Prentice Hall.
- Hull, J. C., & White, A. (1987). The pricing of options on assets with stochastic volatilities. *Journal of Finance*, 42, 281–300.
- Klemkosky, R. C., & Resnick, B. G. (1979). Put call parity and market efficiency. *Journal of Finance*, 34, 1141–1155.
- Leland, H. E. (1985). Option pricing and replication with transaction costs. *Journal of Finance*, 40, 1215–1221.
- MacBeth, J. D., & Merville, L. J. (1979). An empirical examination of the Black–Scholes call option pricing model. *Journal of Finance*, 34, 1173–1186.

- Merton, R. (1990). *Continuous time finance*. Oxford: Blackwell.
- Nisbit, M. (1992). Put-call parity and an empirical test of the efficiency of the London traded options market. *Journal of Banking and Finance*, 16, 381-403.
- Peña, I., Rubio, G., & Serna, G. (1999). Why do we smile? On the determinants of the implied volatility function. *Journal of Banking and Finance*, 16, 23, 1151-1179.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (1992). *Numerical recipes in C* (2nd ed.). New York: Cambridge University Press.
- Rendleman, R. J., & Bartter, B. (1979). Two-state option pricing. *Journal of Finance*, 34, 1093-1110.
- Ritchken, P. H. (1985). On option pricing bounds. *Journal of Finance*, 40, 1219-1233.
- Rubinstein, M. (1978). Nonparametric tests of alternative option pricing models using all reported trades and quotes on the 30 most active CBOE option classes from August 23, 1976 through August 31, 1978. *Journal of Finance*, 40, 455-480.
- Rubinstein, M. (1994). Implied binomial trees. *Journal of Finance*, 49, 771-818.
- Toft, K. (1996). On the mean-variance tradeoff in option replication with transaction costs. *Journal of Financial and Quantitative Analysis*, 31, 233-263.
- Whaley, E., & Wilmont, P. (1993). Counting the costs. *Risk*, 6, 59-65.

Copyright of Journal of Futures Markets is the property of John Wiley & Sons, Inc. / Business and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.