THE JOURNAL OF FINANCE • VOL. XLIV, NO. 5 • DECEMBER 1989

Options Arbitrage in Imperfect Markets

STEPHEN FIGLEWSKI*

ABSTRACT

Option valuation models are based on an arbitrage strategy—hedging the option against the underlying asset and rebalancing continuously until expiration—that is only possible in a frictionless market. This paper simulates the impact of market imperfections and other problems with the "standard" arbitrage trade, including uncertain volatility, transactions costs, indivisibilities, and rebalancing only at discrete intervals. We find that, in an actual market such as that for stock index options, the standard arbitrage is exposed to such large risk and transactions costs that it can only establish very wide bounds on equilibrium options prices. This has important implications for price determination in options markets, as well as for testing of valuation models.

AMONG ALL THEORIES IN finance, the Black-Scholes option pricing model has perhaps had the biggest impact on the real world of securities trading. Virtually all market participants are aware of the model and use it in their decision making. Academics regularly test the model's valuation on actual market prices and typically conclude that, while not every feature is accounted for, the model works very well in explaining observed option prices.[1]

Most option valuation models are based on an arbitrage argument. Under the assumptions of the model, the option can be combined with the underlying asset into a hedged position that is riskless for local changes in the asset's price and in time and must therefore earn the riskless interest rate. This leads to a theoretical value for the option such that profitable arbitrage is ruled out.

However, while virtually all options traders are aware of option pricing theory and most use it in some way, the arbitrage mechanism assumed in deriving the theory cannot work in a real options market in the same way that it does in a frictionless market. The disparity between options arbitrage in theory and in practice is the subject of this paper.

Some of the important assumptions made in deriving the Black-Scholes model are the following.

- The price of the underlying asset follows a logarithmic diffusion process that can be written

$$dP/P = R\ dt + v\ dz, \tag{1}$$

where R is the drift of the price per unit time, dt denotes an infinitesimal

* Stern School of Business, New York University. The author would like to thank John Merrick, Roni Michaely, William Silber, and the referee, Mark Rubinstein, for helpful comments.

[1] Empirical studies of the option pricing model include Black and Scholes (1972), Galai (1977), and Macbeth and Merville (1979), among many others. Galai (1983b) provides a review of the literature on testing option models.

time increment, v is the volatility of proportional price change per unit time, and dz represents Brownian motion, the realization of a random variable distributed as normal with mean 0 dt and variance 1 dt.
- The volatility v is known.
- There are no indivisibilities.
- There are no transactions costs.
- Markets are "perfect" in other ways. There is no limit on borrowing or lending at the same riskless interest rate, and there are no taxes or constraints on short selling with full use of the proceeds.

In fact, none of these assumptions is true of real financial markets, and the arbitrage by which the theoretical pricing relation is supposed to be enforced, i.e., forming a riskless hedge, rebalancing continuously, and holding until option expiration, cannot actually be done with real options. For example, prices do not follow a continuous diffusion when the market is closed. Between trading sessions, prices can make nonlocal changes from one trade to the next with no possibility of rebalancing a hedged portfolio in between.[2]

One of the biggest problems in real world options trading is determining the volatility of the underlying asset's price. This is not a known constant parameter; nor is there even general agreement on the best procedure for estimating it. On the contrary, actual volatility, and also the market's volatility estimate, appear to vary randomly over a wide range.[3] Moreover, even if the true value for volatility is known, the *realized* volatility in a given (finite) series of prices will differ from the true value due to sampling variation.

Errors in predicting volatility lead to two kinds of errors in trading options. Most important is the error in evaluating the fair price for the option. Too low (high) a volatility estimate gives a model value that is also too low (high), and for options that are not deep-in-the-money the price is quite sensitive to small changes in the volatility parameter. An investor who has a more accurate volatility estimate than the market can, in theory, form a fully hedged position earning a return higher than the riskless rate. However, such a trade is complicated by the fact that the unknown volatility is also a determinant of the hedge ratio. The arbitrage position will not earn its excess return risklessly if an incorrect hedge ratio is used. Thus, unknown volatility affects both the return and the risk in options arbitrage.

Securities are also indivisible. Most traders would prefer to trade stock, for

[2] Rebalancing at discrete intervals rather than continuously has been examined by several authors. Galai (1983a) finds that there is little impact on the mean return earned by a hedged position but that its variance is increased. Boyle and Emanuel (1980) observe that the probability distribution of hedge returns is affected, and therefore so is the methodology required for empirical tests on options. Interestingly, both Brennan (1979) and Rubinstein (1976) show that continuous rebalancing is not necessary for the Black-Scholes model to hold. With the right combination of security price distribution and investor utility function, the equilibrium option price will be the Black-Scholes value even if rebalancing is impossible. An important assumption needed for this result is that aggregation conditions hold, so that the market behaves as if there were a single "representative" investor.

[3] Time variation in volatility of stock prices has been discussed by a number of authors, including Black (1976), Beckers (1981), and Christie (1982). Volatility estimates implied by the option pricing model also are highly variable, as shown by Latane and Rendleman (1976) and Rubinstein (1985).

example, in round lots of 100 shares. The effect of indivisibilities is greater when futures contracts are used to hedge an options position, because contract size is large. For instance, suppose stock index options with market exposure equal to $500,000 worth of stock are to be hedged using Standard and Poor's 500 futures contracts. At current (January 1989) prices, the value of one S&P contract is about $140,000. This allows one to construct only positions with hedge ratios of 0.28, 0.56, or 0.84 by selling one, two, or three contracts. Obviously, the indivisibility of the futures contract will lead to hedging inaccuracy.

Probably the most important "imperfection" of real financial markets is the existence of transactions costs. Arbitrage relationships that hold in theory are always affected by transactions costs in practice. Broadly speaking, transactions costs create bounds around the theoretical price within which the market price may fall without giving rise to a *profitable* arbitrage opportunity large enough to cover the cost of exploiting it.[4]

Arbitrage bounds on options prices cannot be easily computed. Because of the dynamic nature of the hedging strategy, the total transaction cost in a particular arbitrage trade will depend on how much the position has to be rebalanced. That is a function of the actual path taken by prices, so the trader cannot know in advance how large these costs will be. Nor can a researcher testing an option pricing model on market data know how big a deviation has to be before it represents a large enough after-cost profit to be a true "mispricing."

It is obvious that the other perfect markets assumptions, such as unlimited borrowing at the riskless interest rate, short sales with full use of the proceeds, absence of taxes, and so forth, do not hold any better in real markets than the ones we have already mentioned.

How can we find out how much options arbitrage is affected by market imperfections? One possibility would be to attempt to incorporate the imperfections directly into our theoretical valuation models. For some cases this is possible. For example, we can compute the amount of mispricing that arises when the wrong volatility estimate is used, and Leland (1985) is able to make some headway in determining theoretically the effect of proportional transactions costs. However, for the most part, the mathematical problems raised by treating realistic market imperfections are too complex to be tractable theoretically.

A second solution is to simulate trading strategies on historical option price data and to tabulate the results.[5] Analyzing historical data has always been the standard approach for testing option models, but there are several problems with it. The researcher is limited to examining a single set of data that may not be very long and over which he or she has no control. The researcher cannot know, for example, what the true ex ante distribution for prices was. Other difficulties

[4] Most published empirical tests of option pricing models find some mispricing in the market relative to the models' prescriptions, but they also find that these potential profit opportunities disappear when some estimate of the transactions costs involved is considered. Phillips and Smith (1980) document the costs of setting up and unwinding an options hedge and then show that these outweigh the possible profits uncovered by many earlier studies.

[5] Garcia and Gould (1987) simulate the performance of portfolio insurance, an application of option arbitrage, on historical data with realistic transactions costs and rebalancing. They find substantial deviations between actual results and theoretical estimates.

are that actual volatilities change over time, realized volatilities may differ considerably from the ex ante values, market prices may at times be distinctly nonlognormal, and so on.

The approach we will take in this paper is to simulate the performance of options hedge strategies on *simulated* price data. We specify values for the drift and volatility, R and v, and construct 250 randomly drawn price series and then do Monte Carlo simulations to determine the effect of some of the market imperfections discussed above.[6] This procedure has several things to commend it. First, it is relatively easy to do. Second, we know exactly what the true parameters of the price-generating process are. The results we obtain empirically, therefore, are directly comparable to those we would obtain by using the same parameters in a correctly specified theoretical model. Third, like other numerical methods, our simulations can be made arbitrarily accurate simply by using (i.e., creating) more observations.

In the next section we describe the experimental design, how the prices are generated, and how the hedged positions are constructed. Then we examine summary statistics on the realized returns and risk on the securities in the sample. These results will show that rebalancing the hedge daily rather than continuously has a considerable impact on its risk. Section II looks at other market imperfections which predominantly affect the level of risk in a hedged option position. These are the use of an incorrect volatility estimate in computing the hedge ratio and indivisibilities.

Section III analyzes transactions costs. We compute the after-costs returns and risk on hedged positions using approximately the transactions cost structure that currently applies to market makers and retail traders in stock index options. We also look at the performance of alternative trading strategies designed to reduce costs by rebalancing less frequently.

In Section IV we consider the arbitrage bounds that this cost structure would imply and compare them to typical market bid-ask spreads. We find that hedging an option with the underlying asset dynamically and rebalancing the position once a day until expiration would be exposed to such large transactions costs and risk in actual markets that it is impractical even for an options market maker. Rebalancing less frequently can reduce costs, but risk increases. Thus, the "standard" arbitrage can establish only very wide bounds on real option prices. This has important implications for price determination in options markets, as well as for testing valuation models.

The simulations we look at in the paper cover the standard arbitrage trade with one-month options. In Section V we present some results and discussion relating the analysis to longer maturity options and to other option replication strategies, such as portfolio insurance and the trading of actual market makers.

The final section summarizes our findings in more detail.

[6] In an early paper, Boyle (1977) proposes Monte Carlo simulation as a procedure for valuing options. Etzioni (1986) uses a simulation strategy to examine alternative rebalancing procedures for portfolio insurance. We will discuss the particular case of portfolio insurance in more detail below.

I. Experimental Design

If an option is priced at its arbitrage-based value in the market, the strategy of buying the option, forming the hedged portfolio, and carrying it, rebalancing continuously, until expiration will return exactly the riskless rate of interest.[7] To see how this strategy behaves with market imperfections such as are present in actual options markets, we begin by constructing a set of price series and option values to use as the basic data in the hedge tests.

The choice of parameters such as volatility is open. Throughout the paper we use parameter values that have been typical for actual options on broad-based stock indices. For convenience we will often refer to the underlying asset as the "stock," but our results will of course apply to other kinds of options as well.

Two hundred fifty price series of 25 observations each (indexed as $t = 0, \cdots, 24$) were constructed, each starting at the initial value of $P_0 = 100$. This corresponds to a hedge period of about one month.[8] Notice that, in doing this, we have already departed from a world in which continuous rebalancing is possible. The hedge is rebalanced at most once a day.

For each series, subsequent prices are computed according to equation (2). This process for prices is implied by the assumed returns equation (1):

$$P_{t+1} = P_t e^{R+vz}, \tag{2}$$

where R is the mean rate of price change per day, v is the daily volatility, and z is a random draw from a standardized normal probability distribution.

For this study we have set R and v to be the daily equivalent of an annual 15% and 0.15, respectively. That is,

$$R = (\log\ (1.15))/260 = 0.000538,$$

$$v = 0.15/260^{0.5} = 0.00930.$$

Next, corresponding series of option prices and theoretical hedge ratios were constructed using the Black-Scholes model, for call options with four different strike prices: 97, 100, 103, and 105. This gives us an in-the-money, an at-the-money, an out-of-the-money, and a deep-out-of-the-money option. In pricing the options, the riskless interest rate is assumed to be 5.0 percent and the volatility is taken to be the true volatility, 0.15 at an annual rate.

Table I shows summary statistics for the constructed price series in our sample. The first line describes the stock price series. The initial price was 100 for each series, and the mean terminal price was 101.43, with a standard deviation across series of 4.85. The mean of the annualized percentage rate of return was 15.52, calculated from $100 \times (P_{24}/P_0 - 1)$ and annualized at simple interest. The standard deviation was annualized by multiplying by the square root of 260/24.

[7] Note that, since the value of the funds invested in the portfolio changes as it is rebalanced, earning the riskless rate implies earning a continuous cash flow at that rate on the current, time-varying value of the portfolio.

[8] There are between 250 and 260 trading days per year, so the typical month has 21 to 22 trading days. We will use 260 in converting between daily and annualized parameter values.

Table I

Summary Statistics for Simulation Data Sample

The table shows summary statistics from 250 simulated series of 25 daily prices each. Stock prices are generated with an annual drift of 15% and volatility of 0.15. Option prices are computed from the underlying stock prices using the Black-Scholes model, assuming volatility of 0.15 and riskless interest of 5.0 percent. "Mean" and "Std Dev" figures refer to sample means and standard deviations across the 250 series. Returns are annualized by multiplying by 260/24; volatilities by $(260/24)^{0.5}$.

Stock

Initial Price	Final Price		Percent Return		Volatility	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
100	101.43	4.85	15.52	15.95	0.150	0.022

Calls

Strike	Initial Price	Final Value		In the Money	
		Mean	Std. Dev.	Number	Avg. Amount
97	4.01	4.91	4.13	204	6.02
100	2.05	2.75	3.31	152	4.52
103	0.84	1.25	2.29	92	3.40
105	0.41	0.65	1.63	60	2.72

Within each series of 25 prices, the standard deviation of the log price relatives produces a realized volatility. The mean of these sample volatilities in our constructed data turned out to be 0.150, as it should. This suggests that the series are representative.

It is very interesting to see that the standard deviation of the realized volatilities across series was as high as 0.022. In other words, knowing that the true volatility is 0.15 only tells us that there is about ⅔ probability that a given month's volatility will be somewhere between 0.128 and 0.172.[9] To options market makers, a difference of that size between volatility estimates is considered very large. What this shows is that, *even if the true volatility is known*, there is a sizable standard error of forecast in predicting what volatility will actually be experienced in daily prices over a month.

The second part of the table summarizes the returns to buying call options "naked." For example, the at-the-money call was priced initially at 2.05. Its end-of-period value averaged 2.75, with a standard deviation of 3.31. Of course, the distribution of these values is highly skewed: of 250 options, 152 ended up in the money, by an average amount of 4.52. The remainder expired worthless.

Table II shows the trading strategies and expected arbitrage profits that a trader would be able to earn under the perfect markets assumptions of the Black-Scholes model, assuming that the market had mispriced these options.

Consider the at-the-money calls. Knowing that the true price volatility of the underlying asset was 0.15, the trader would value the 100 strike price calls at 2.05. If the market's volatility estimate was 0.10, those calls would be selling for 1.45. The indicated arbitrage would then be to buy the underpriced calls and

[9] Due to the effect of the Central Limit Theorem, the distribution of these realized volatilities is very close to Gaussian in our sample.

Table II

Theoretical Arbitrage Profit When True Volatility is 0.15

The table shows the standard option arbitrage trade from the perspective of an arbitrageur who believes the true underlying volatility to be 0.15. Each line gives the market's volatility, the market option price, and hedge ratio and analyzes the arbitrage position that would be taken.

Strike Price	Market's Volatility Estimate	Call Price	Hedge Ratio	Trading Strategy: Call	Trading Strategy: Stock	Cost of Riskless Position	Excess Return: $ amt.	Excess Return: Annual %
97	0.10	3.62	0.878	Buy	Sell	−74.90	0.385	5.57
	0.13	3.84	0.817	Buy	Sell	−74.69	0.170	2.46
	0.15	4.01	0.785	No Trade		74.52	0.000	0.00
	0.17	4.19	0.759	Sell	Buy	74.33	0.184	2.68
	0.20	4.48	0.728	Sell	Buy	74.04	0.478	6.99
100	0.10	1.45	0.565	Buy	Sell	−53.39	0.600	12.18
	0.13	1.81	0.553	Buy	Sell	−53.03	0.240	4.91
	0.15	2.05	0.548	No Trade		52.79	0.000	0.00
	0.17	2.29	0.545	Sell	Buy	52.55	0.241	4.96
	0.20	2.65	0.542	Sell	Buy	52.19	0.602	12.50
103	0.10	0.35	0.209	Buy	Sell	−29.56	0.490	17.95
	0.13	0.64	0.269	Buy	Sell	−29.27	0.206	7.63
	0.15	0.84	0.299	No Trade		29.07	0.000	0.00
	0.17	1.06	0.323	Sell	Buy	28.85	0.215	8.06
	0.20	1.39	0.351	Sell	Buy	28.52	0.548	20.81
105	0.10	0.10	0.075	Buy	Sell	−17.03	0.309	19.67
	0.13	0.27	0.135	Buy	Sell	−16.87	0.144	9.24
	0.15	0.41	0.171	No Trade		16.72	0.000	0.00
	0.17	0.57	0.203	Sell	Buy	16.56	0.163	10.69
	0.20	0.84	0.242	Sell	Buy	16.29	0.435	28.96

short the stock using a hedge ratio of 0.548 shares per call. The initial cost of this position would be negative, meaning there would be a net cash inflow of 53.39.

As time elapsed and the stock price changed, this position would be rebalanced by adjusting the amount of stock sold short, always computing the new hedge ratio using 0.15 as the volatility estimate. Since the trade brings in cash at the beginning, "earning the riskless rate of return" would correspond to a continuous *loss* equal to paying the riskless interest rate on the net amount of funds remaining in the position. Regardless of the actual course of prices during the 25-day hedge period, the initial mispricing of the option would yield an excess return, (i.e., a reduction in the cost of obtaining funds) of $0.600. As a percentage of the initial value of the hedge portfolio, this would be an annualized 12.18 percent.

On the other hand, if the market's volatility estimate were 0.20 instead of 0.10, the arbitrageur would write calls at 2.65, buy stock, and create a hedged portfolio costing 52.19. Over time this would earn the riskless rate plus the initial option overpricing of $0.602. The excess return would be 12.50 percent.

The table shows two important properties of this arbitrage trade. One is that the dollar value of the arbitrage portfolio tends to be large compared to the amount of mispricing, even for a large difference in volatility estimates. In the example we just described, it was necessary to take a position in over 50 dollars worth of securities and to be prepared to manage the position carefully for a month in order to earn an excess return of 60 cents. This problem is less severe for the out-of-the-money calls.

Also apparent in these figures is the fact that using an incorrect volatility estimate makes a bigger difference in the value of a call than in its hedge ratio, or delta. The delta for the at-the-money calls, in particular, is hardly affected at all by changes in volatility over a wide range. By far the largest impact is on the deep-out-of-the-money calls when the implied volatility is too low. In that case, the market price and the delta both are much too close to zero. (It should be understood that these properties come from the Black-Scholes equation—they are not produced by our simulation.)

II. Market Imperfections and Risk

This section begins to examine hedge strategies. In every case we consider the arbitrage trade of buying a call option at the market price and selling the number of shares indicated by the hedge ratio. Thereafter, each day the hedge ratio is recalculated and the hedged portfolio is rebalanced by buying or selling the underlying asset. The day's excess return is calculated using equation (3):

$$ER_t = (C_t - C_{t-1}) - h_{t-1}(P_t - P_{t-1}) - r(C_{t-1} - h_{t-1}P_{t-1}), \tag{3}$$

where ER is the excess return, C is the call price, P is the price of the underlying asset, h is the hedge ratio called for by the particular trading strategy, and r is the one-day riskless interest rate, based on an annual rate of 5.00 percent.

The daily excess return figures are then cumulated, leading to 250 total excess return amounts, one from each price series, for each combination of hedge strategy and strike price. The sample mean and standard deviation statistics for the excess return totals show how introducing different market imperfections into the system affects the expected return and the risk of an options hedge.

As a base for comparison, we begin by analyzing a "Base Case" without imperfections. As mentioned above, it is not the Black-Scholes case exactly, since the position is rebalanced only once a day. It does, however, correspond to the typical methodology used in empirical tests of option pricing models.

The first line in Table III gives the Base Case results on the standard deviation of excess returns across the 250 price series. The figures are shown as annualized percentage rates. Thus, although buying the 100 strike calls at the theoretical value of 2.05 and selling short 0.548 shares per option would yield zero excess return with a standard deviation of zero if it were possible to rebalance the position continuously, daily rebalancing leads to a risky hedge whose annualized rate of return over the holding period has a standard deviation of 6.52 percent. The mean return is also nonzero in the sample but not statistically significant. (Since none of the means for the strategies examined in Table III was significantly different from zero, we do not report them.)

Before going on to look at market imperfections, it is worth reflecting briefly

Table III

Effects of Market Imperfections on Hedge Standard Deviation (Annualized Percent Standard Deviation)

The table shows the standard deviation across 250 price series of the annualized cumulative excess returns to option expiration on hedged positions that are long the call option and short the underlying stock. Positions are rebalanced daily. The Base Case uses the exact hedge ratio computed from the true volatility of 0.15. Incorrect Volatility results show the effect of computing the hedge ratio with an incorrect v. Indivisibilities results show the effect of rounding the correct (v = 0.15) hedge ratio to the nearest integer multiple of K.

	$X = 97$	$X = 100$	$X = 103$	$X = 105$
Base Case	3.34	6.52	11.18	16.70
Incorrect Volatility				
$v = 0.10$	4.89	8.18	24.07	64.31
$v = 0.13$	3.59	6.70	13.56	24.15
$v = 0.15$	3.34	6.52	11.18	16.70
$v = 0.17$	3.68	6.84	10.76	14.88
$v = 0.20$	4.85	7.84	11.80	16.01
Indivisibilities				
$K = 0.02$	3.42	6.62	11.25	16.05
$K = 0.05$	3.25	6.64	11.38	19.68
$K = 0.10$	3.82	7.82	12.09	16.46
$K = 0.25$	5.38	9.64	19.06	16.80
$K = 0,1$	11.68	17.10	1680.91	2868.26

on the meaning of these standard deviations. The riskless interest rate has been assumed to be five percent. In comparison, a standard deviation of more than six percent makes the position quite risky. It is apparent that, simply by rebalancing discretely instead of continuously, we have departed markedly from the theoretical world of Black-Scholes.

In the second section of Table III we look at hedging when the volatility of the underlying asset is not known. When a trader uses a volatility estimate that is not equal to the volatility that is actually experienced during the option's lifetime, both the expected return and the risk of the trader's arbitrage portfolio are affected.

For example, we saw in Table II that, if the volatility is 0.15, the true value of the 100 strike call is 2.05. Suppose the trader uses an incorrect volatility estimate of 0.10. (This could be due to an incorrect estimation procedure, or it might be the stock's true ex ante volatility, but the realized prices during the option's lifetime could have a *sample* volatility of 0.15.)[10] The call value at a 0.10 volatility is only 1.45. If the trader were to write the option at that price and hedge it by buying the stock, he or she would have sold it for 0.60 below its true value. A perfect hedge would then lock that mispricing in as a certain loss on the trader's position.

The second problem caused by an inaccurate volatility estimate is that the

[10] Note that the sample volatility can differ from the true volatility only in discrete time, such as in the daily price series we are considering. If prices generated by a diffusion process could be followed continuously, the realized volatility over any finite time interval must equal the true volatility with probability 1.0.

hedge ratio used in forming the arbitrage portfolio will be wrong. In this case, correct hedging (i.e., using $v = 0.15$) would lead to (approximately) a 0.60 loss; incorrect hedging with a volatility of 0.10 would also induce a standard deviation in the annualized hedge return of 8.18 percent.

For the in-the-money and the at-the-money options, the standard deviation of returns on a hedged position does not increase much compared to the risk level that is already present in the Base Case. However, for the out-of-the-money and especially the deep-out-of-the-money calls, the impact is substantial. There is also an interesting asymmetry between overestimating and underestimating the volatility. Hedging with too high a volatility estimate does not seem to increase hedge risk much at all. However, underestimating the volatility leads to a considerably larger standard deviation. These results suggest that, in trying to cope with uncertainty about the volatility, it might be appropriate to compute the hedge ratio for out-of-the-money options using a higher volatility than what the trader expects to prevail in the future, on the grounds that it is less costly to err on the side of overestimating than underestimating volatility for these options.

The other imperfection we consider in Table III is the indivisibility of the underlying asset. When the number of options to be traded is small or the underlying asset is like a futures contract that must be traded in large units, the hedge ratio cannot be set to the exact value dictated by the valuation model and the position will necessarily be slightly over or under hedged at all times. How much additional risk does this cause?

The third section of Table III shows the effect on hedge risk when the set of possible hedge ratios is constrained by the indivisibility of the underlying asset. Consider a hedged position consisting of an option on a single share of stock. Regardless of the delta produced by the valuation model, the hedge ratio can take only values of zero or one, depending on whether a share is sold. The hedger attempting to use the Black-Scholes model in this case might sell a share if the delta were greater than 0.5 and remain unhedged if it were less than 0.5.

With option contracts on N units of the underlying asset, the hedge ratio can only take on values that are an integral multiple of $K = 1/N$. Table III shows the standard deviation of hedge returns for several values of K. For the most part, the results bear out the expectation that the larger the value of K (i.e., the less accurate the hedge can be because of indivisibility), the greater will be the risk. However, it is interesting that the effect of rounding the hedge ratio to the nearest K is not very great for the out-of-the-money options even though a given K leads to a relatively larger inaccuracy for them due to their smaller deltas. But by the time K is 0.25, there is a sizable degradation in the effectiveness of the hedge for all but the 105 strike calls.

The final line in Table III shows the strategy of either no hedge or a full hedge ($h = 0$ or 1.0), depending on whether the theoretical delta is less than or greater than 0.5. This is close to a strategy of hedging only options that are in the money and leaving out-of-the-money options unhedged. The impact on hedge risk is substantial for in- and at-the-money calls. For the out-of-the-money options, the initial "hedged" position contains only the naked call because the deltas are both below 0.5. The standard deviation is expressed here as a percentage of the initial position value, which is very small; hence, the numbers in the table become very large.

III. Transactions Costs

We now turn our attention to the impact of transactions costs on options arbitrage. The market imperfections discussed in the last section induced risk but no bias toward higher or lower returns, but transactions costs unambiguously reduce the profitability of every trading strategy. The complication is that the transactions costs that must be paid in hedging an options position depend on the realized path taken by prices.

The transactions cost structure also varies considerably among different classes of traders. Commissions paid by retail investors, in particular, are much larger than those paid by market makers, and they depend on the brokerage firm (whether a "discount" or a "full service" broker) as well as on the size of the trade (with discounts for larger volume). We will look at two cost structures, one corresponding to the trading costs borne by a typical options market maker, and the other to the costs that would be paid by a retail trader dealing with a discount broker. These are representative of the costs applying to trading in stock index options in early 1987.

An options market maker is assumed to pay an exchange fee of $1 per option contract traded (i.e., $0.01 per underlying share). There is no charge for exercise of options finishing in the money, and it is assumed that the market maker trades options without having to pay the bid-ask spread. For hedging transactions in the underlying stock, the market maker pays $0.05 per share plus one half of the bid-ask spread, which we assume to be ⅛. Transactions costs are paid on both the initial sale and subsequent repurchase of the shares.

Since retail customers pay commissions that vary with quantity, we assume that three option contracts are traded, that is, calls on 300 shares. The cost is $8 per contract plus 1.5% of the dollar amount of the transaction. A similar fee must be paid to exercise options expiring in the money. The commission on a stock trade of $35 plus 0.5% of the dollar amount traded, plus half of the bid-ask spread of ⅛. Again, commissions are paid on both opening and closing trades.

Table IV shows the effect of transactions costs on the option hedges we have been considering. As before, we begin with an analysis of the Base Case. We then look at alternative strategies designed to limit transactions costs by reducing the number of rebalancing transactions.

Consider the Base Case results for the 100 strike calls. The arbitrage trade is to buy the call at its theoretical value of 2.05 and to sell 0.548 shares short, rebalancing daily. At expiration, options finishing in the money are exercised and the remainder expire worthless. (We assume cash settlement, so that one does not have additional costs to dispose of stock acquired through exercise.) The hedge position in the stock that remains on the expiration day is liquidated in the stock market.

The mean number of stock trades across all 250 series was 24.6, and the mean total number of shares traded was 2.64 (per call option on one share). Thus, on

Table IV

Comparison of Mean Transactions Costs

The table shows the mean and standard deviation across 250 series for arbitrage excess returns including transactions costs under different rebalancing strategies. The cost structure and arbitrage strategies are described in the text. Returns are in dollars per option on one share. Trades is the average number of days with a transaction in the stock. Shares traded is the average total number of shares traded in hedging the option for 25 days.

			Excess Return Including Transaction Costs					
			No Costs		Market Maker		Retail	
	Trades	Shares Traded	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
			Strike 97					
Base Case	24.2	2.64	0.037	0.230	−0.269	0.223	−10.104	0.684
No Rebalance	2.0	1.57	0.126	1.022	−0.060	1.022	−1.638	1.019
Rebalance Every K Days								
$K = 2$	12.8	2.30	0.043	0.349	−0.226	0.341	−5.934	0.465
$K = 5$	6.0	2.03	0.028	0.505	−0.210	0.496	−3.414	0.532
Rebalance When h Changes by K								
$K = 0.10$	6.1	2.21	0.039	0.334	−0.220	0.326	−3.508	1.175
$K = 0.25$	2.8	1.73	0.128	0.720	−0.076	0.718	−1.951	0.879
			Strike 100					
Base Case	24.6	2.64	0.038	0.318	−0.269	0.314	−10.144	0.664
No Rebalance	2.0	1.10	0.160	1.420	0.026	1.420	−1.175	1.417
Rebalance Every K Days								
$K = 2$	12.9	2.20	0.043	0.457	−0.214	0.447	−5.785	0.607
$K = 5$	6.0	1.76	0.051	0.687	−0.157	0.672	−3.100	0.710
Rebalance When h Changes by K								
$K = 0.10$	8.3	2.16	0.040	0.391	−0.213	0.380	−4.142	1.163
$K = 0.25$	3.7	1.64	0.049	0.641	−0.145	0.633	−2.193	0.778
			Strike 103					
Base Case	24.4	2.11	0.043	0.300	−0.204	0.294	−9.694	0.870
No Rebalance	2.0	0.60	0.113	1.445	0.036	1.445	−0.801	1.443
Rebalance Every K Days								
$K = 2$	12.8	1.67	0.049	0.461	−0.149	0.440	−5.350	0.672
$K = 5$	5.9	1.23	0.048	0.695	−0.101	0.674	−2.668	0.713
Rebalance When h Changes by K								
$K = 0.10$	7.9	1.67	0.050	0.354	−0.148	0.345	−3.607	1.509
$K = 0.25$	3.7	1.21	0.025	0.592	−0.121	0.581	−1.880	0.819
			Strike 105					
Base Case	23.9	1.51	0.033	0.258	−0.147	0.254	−9.198	0.985
No Rebalance	2.0	0.34	0.076	1.216	0.027	1.216	−0.629	1.215
Rebalance Every K Days								
$K = 2$	12.6	1.18	0.042	0.415	−0.101	0.395	−4.982	0.684
$K = 5$	5.9	0.83	0.032	0.570	−0.072	0.548	−2.385	0.599
Rebalance When h Changes by K								
$K = 0.10$	6.3	1.14	0.022	0.332	−0.116	0.318	−2.704	1.621
$K = 0.25$	3.0	0.82	−0.004	0.618	−0.107	0.605	−1.411	0.956

average, the dynamic hedging strategy required approximately five share transactions for every share in the initial hedge position. This can lead to a very costly hedge for retail traders who pay a minimum charge per stock trade no matter how few shares are involved.

For comparison, the next two columns show the sample means and standard deviations of excess returns when there are no transactions costs. The effect of trading costs can then be seen in the deviations from these values. Thus, without taking account of transactions costs, the sample mean excess return for the 100 strike hedges was \$0.038, with a standard deviation of \$0.318. Including transactions costs, a market maker would have experienced a mean of \$−0.269, with a standard deviation of \$0.314, meaning that the net impact of transactions costs was to reduce mean excess returns by (−0.269 − 0.038) = \$−0.307.

A retail trader would have huge costs if he or she attempted to trade the option arbitrage according to the dictates of the Black-Scholes model, losing more than \$10 on average when the option's initial value was only about \$2. Due to the variation across price series in the amount of rebalancing, and therefore trading costs, the retail trader's standard deviation was also substantially higher than that borne by the market maker. However, this is obviously of lesser importance than the effect on the mean return.

The patterns exhibited by the at-the-money calls were also present in the others. The reduction in mean for the market maker varied from −0.180 to −0.307, while the standard deviation was only slightly affected. The retail trader took losses far greater than the initial price of the option in each case.

It is clear from these figures that transactions costs make a substantial difference in the outcome of an options arbitrage, even when done by a market maker. We will see this in more detail in the next table. It suggests that trading strategies that economize on the number of transactions or the number of shares traded might be worth pursuing, even though one would expect risk to increase when the hedge is not rebalanced as often as possible.

At the opposite pole from continuous rebalancing is no rebalancing at all. That strategy is examined in the second line of results for each strike price. The trader takes an options position at the outset, hedges it according to the theoretical hedge ratio, and holds it until expiration without any further trading of the stock. This reduces the number of transactions to two: an opening and a closing trade, and the total number of shares traded is just twice the initial hedge ratio.

For the 100 strike call, the no-rebalancing strategy increases the mean excess return without transactions costs to \$0.160, but the standard deviation has more than quadrupled, to \$1.420. Commissions paid by a market maker reduce the mean to \$0.026, a net cost of \$0.134 on average instead of the previous \$0.307. A retail trader's cost is cut by nearly a factor of ten, but it remains high enough that the trade is still very unattractive.

What an arbitrageur wants is an intermediate strategy that limits both trading costs and risk. Two possibilities are commonly suggested. One is to rebalance less frequently than every day, perhaps every two days or every week. This limits the number of trades but allows for the possibility that the hedge proportions can get far out of line in between rebalancing points. A second approach is to monitor the discrepancy between the actual and the theoretical hedge ratios daily, but only to rebalance when the hedge gets too far from the correct value. This leads to frequent rebalancing in some cases and little in others, depending on the actual stock price path. The strategy allows more variation in the number of trades among hedges but keeps the hedge ratio close to the theoretical value at all times.

The table shows the results of two such strategies of each type. Lines three and four for each strike examine rebalancing the hedge every two days and once a week (i.e., five trading days). The following two lines relate to the strategy of rebalancing only when the actual hedge ratio differs from the theoretical by at least 0.10 and 0.25, respectively.

Rebalancing only every two days or every week reduces the number of trades substantially, with the latter leading to an average number of shares traded that is about halfway between the two polar cases of rebalancing daily and not at all. Once again, the retail trader has such heavy costs that he or she would not follow any of these strategies. There is clearly no point in discussing the returns to retail traders any further.

Rebalancing only when the actual hedge ratio gets too far away from the theoretical value is a logical way to reduce the amount of trading of the underlying asset. The two values for the maximum permitted deviation, $K = 0.10$ and $K = 0.25$, were chosen because the mean number of shares traded and the average reduction in excess return were comparable to the equivalent figures for the previous strategies.

Although these results can be analyzed carefully to try to determine which strategy performed better for which options, for none of them is the resulting combination of risk and return very favorable. In all cases, costs remain substantial and risk levels increase quickly as the frequency of rebalancing is reduced.

One of the most apparent conclusions to be drawn from Table IV is that, even for a market maker in options, the transactions costs entailed by the arbitrage strategy underlying the Black-Scholes model are quite large. For example, Table II shows that, if the market were pricing the 100 calls on a volatility of 0.13 while the market maker believed the true value was 0.15, the model indicates an arbitrage profit of $0.240. However, the transactions costs involved in trying to capture that excess return would be $0.307, more than enough to wipe out all of the profit.

IV. Arbitrage Bounds Based on the Standard Arbitrage

In a frictionless market, the force of arbitrage drives the price of an option exactly to its Black-Scholes value. With costly arbitrage, there will be bounds around the theoretical option price within which the market price may fluctuate freely, because the potential arbitrage profit would be outweighed by the cost of trying to capture it. The results of the previous section allow us to analyze these bounds.

If arbitrageurs derive prices at which they will enter the market to buy or sell calls by calculating the expected cost of the standard arbitrage, the figures shown in Table IV can be used to compute the bid and ask prices required for them to break even or to achieve any specified probability of earning a profit. The results of this calculation are displayed in Table V. We assume that arbitrageurs face the market maker cost structure described above.

Consider the Base Case for the at-the-money calls. With a 0.15 volatility, the option's theoretical value is 2.05. In Table IV, we saw that the arbitrage trade would entail an average total transactions cost of $0.307. If the arbitrageur's strategy is to trade the option and then hedge it and hold to expiration, he or she must buy the option at no more than 1.74 or sell it no cheaper than 2.35 in order to expect to break even. The width of the no-arbitrage band is (at least) 2 × 0.307 = $0.61.

Traders often think about option pricing in the market not in terms of prices but in terms of implied volatilities. For instance, if they thought the true volatility was 0.15, in this case they would be willing to "buy the option on a .125 volatility and sell it on a .176 volatility." For each case in Table V we show the arbitrage boundary bid and ask prices in dollars and, on the following line, the implied volatilities corresponding to those prices.

The break-even calculation involves only the expected cost of the arbitrage trade: the risk does not enter. However, risk does come into the calculation if

Table V

Arbitrage Bounds with Transactions Costs

The table shows the arbitrage bounds on option prices and implied volatilities at which an arbitrageur would bid for and offer options and have a 50 percent or 75 percent probability of covering costs.

Rebalancing Strategy		Market Maker Cost Structure: Breakeven Bid	Breakeven Ask	75% Profit Prob. Bid	75% Profit Prob. Ask
Strike 97—Model Value = 4.01					
Base Case	Price	3.70	4.31	3.55	4.46
	Vol.	0.112	0.183	0.087	0.198
No Rebalance	Price	3.82	4.19	3.13	4.88
	Vol.	0.128	0.170	Negative	0.239
Strike 100—Model Value = 2.05					
Base Case	Price	1.74	2.35	1.53	2.57
	Vol.	0.125	0.176	0.107	0.193
No Rebalance	Price	1.91	2.18	0.96	3.14
	Vol.	0.139	0.161	0.058	0.241
Strike 103—Model Value = 0.84					
Base Case	Price	0.60	1.09	0.40	1.29
	Vol.	0.126	0.173	0.105	0.191
No Rebalance	Price	0.77	0.92	No Bid	1.90
	Vol.	0.143	0.157	Negative	0.244
Strike 105—Model Value = 0.41					
Base Case	Price	0.23	0.59	0.06	0.76
	Vol.	0.125	0.172	0.089	0.191
No Rebalance	Price	0.36	0.46	No Bid	1.28
	Vol.	0.144	0.156	Negative	0.244

arbitrageurs require more than a 50 percent probability of covering their costs. How large a profit probability traders will require to engage in the standard arbitrage will depend on several factors, including their level of risk aversion and the degree of competition among them. For illustration, we calculate the bid-ask spreads that would produce a 75 percent probability of covering costs.

The 75th percentile of the normal distribution occurs at 0.67 standard deviations, so an arbitrageur has a 75 percent chance of covering costs if he or she quotes a bid lower than the model value and an offer above it by an amount equal to the expected transactions cost plus 0.67 standard deviations. In the Base Case for the 100 strike call, that would be a bid of 1.53 and an offer of 2.57 (implied volatilities of 0.107 and 0.193, respectively).

It is appropriate to note here that bid-ask spreads in actual options markets are much smaller than these figures. The typical market bid-ask spread on a one-month at-the-money stock index call option selling at about 2 would be no more than ¼ point, and normally less. That is, the price would be quoted as 1⅞ bid, offered at 2⅛, or better.

The second case examined for each strike price is no rebalancing. This is the strategy with the lowest average transactions cost but the highest standard deviation. For the at-the-money option, an arbitrageur could expect to break even bidding 1.91 and offering at 2.18, for a bid-ask spread of just about ¼ point. However, the risk involved in the arbitrage trade is so great that, to be 75 percent sure of covering costs, the bid should be no more than 0.96 and the ask at least 3.14, a much larger spread than in the Base Case.

For the options with different strike prices, the risk of the No-Rebalance strategy is sufficiently great that the appropriate bid price violates the option's boundary conditions. A bid price of 3.13 on the 97 call is less than the current stock price minus the present value of the exercise price, so the implied volatility would be negative. For the out-of-the-money options, there is no positive bid price that would allow a 75 percent profit probability.

In comparing results in this table for the different strategies and strike prices, several significant features are visible. First, only, the No-Rebalance break-even strategy leads to bid-ask spreads that are comparable to those observed in actual options markets. Typical spreads for these options would be ¼ to ⅜ for the 97 strike calls, about ⅛ to ¼ for the 100 strikes, and about ⅛ for the 103s and 105s. In no case would the spread indicated in Table V be close to the observed value if market makers only did the standard arbitrage and required as much as a 75 percent probability of making a profit on their trades.[11]

We do not report the results for the alternative strategies examined in Table IV since they were not particularly effective. In all cases, the Base Case results indicated the narrowest or almost the narrowest spread when a 75 percent profitability hurdle was imposed, because the effect of the lower mean cost for the other strategies was offset by their increased variability of returns.

[11] Market makers, of course, are able to rebalance their positions more frequently than every day, which would allow them to eliminate more of the risk than this table shows. However, more rebalancing also means higher transactions costs. In the limit, as Leland (1985) observes, it is a mathematical property of a logarithmic diffusion process that rebalancing continuously would require an infinite number of transactions and would involve trading an infinite number of shares.

V. Other Option Maturities and Trading Strategies

The results we have developed so far apply to call options that are not too far in or out of the money and have about one month to expiration. How badly market imperfections impact the standard arbitrage trade depends on the amount of rebalancing required to keep the portfolio hedged. This in turn is a function of how much the hedge ratio changes as the stock price moves—in other words, the option's "gamma."[12] Gamma is greatest for options that are close to expiration and at the money, exactly the ones we are looking at.

For example, our 100 strike calls have a gamma of 0.082. Because delta changes, if the stock price rises from 100 to 101 with no rebalancing in between, the value of the "hedged" portfolio will change by \$0.043. In other words, there would be "replication error" of more than four cents. However, if this were a one-year option, gamma would be only 0.024 and the same price change would only produce a \$0.012 replication error.

For exchange-traded options, the greatest trading volume and open interest is nearly always in the contracts for which our results are representative, those close to the money, with less than three months to expiration. However, arbitrage involving long dated traded options, warrants, and other optional contracts is not uncommon. One example of an arbitrage-like strategy that attempts to replicate longer maturity options is portfolio insurance, which we will discuss in detail below.

A. Longer Maturity Options

To see how much our results would change with longer dated options, we created a new set of stock price series and call option values with 75 trading days, following the same procedures as above. The only difference was that we limited the sample to 100 price series. Table VI displays summary results for the standard arbitrage with three-month calls. We have essentially computed the most relevant results from each of the earlier tables with these new data.

Each stock price series started at 100 and then followed a logarithmic random walk with an annualized drift of 15 percent and volatility of 0.15. Options prices and hedge ratios for calls with strike prices of 97, 100, 103, and 105 were computed from the Black-Scholes model.

The first three lines of the table show the theoretical option values and hedge ratios at the outset, as well as the cost of the arbitrage portfolios. As before, the arbitrage portfolio is long the option and short the underlying asset, so it produces a net cash inflow (negative "cost"). The option theoretical values are higher than those in Tables I and II, the deltas are closer to 0.5, and the values of the hedge portfolios still are very large compared to a typical amount of mispricing of the options.

Comparing the standard deviations of annualized holding-period returns on the hedge portfolios to those shown for the Base Case in Table III reveals a distinct decrease for the longer dated options. This is partly an artifact from annualizing the returns. Multiplying N-day cumulative excess returns by $260/N$

[12] Gamma is defined as the derivative of the hedge ratio with respect to the stock price or, alternatively, the second derivative of the option value with respect to the stock.

Table VI
Summary of Results for Three-Month Options

The table summarizes results from applying the procedures reported in the previous tables to a new sample consisting of 100 price series of 75 trading days each. See the text and earlier tables for a full description.

Strike Price	97	100	103	105
Theoretical				
Call Value	5.76	3.91	2.49	1.78
Hedge Ratio	0.724	0.585	0.438	0.346
Arbitrage Portfolio Cost	−66.62	−54.55	−41.32	−32.82
Hedge Portfolio				
Base Case Percent Standard Deviation	1.29	2.07	3.24	3.70
Transactions Costs for Market Maker (in $)				
Mean Cost	−0.406	−0.431	−0.422	−0.391
Std. Dev. of Hedge Portfolio	0.255	0.322	0.373	0.345
Arbitrage Band (Ask − Bid)				
Breakeven	0.82	0.86	0.84	0.78
75% Profit Probability	1.16	1.30	1.34	1.24

to annualize them multiplies their sample standard deviation by the same factor. However, if daily excess returns are serially independent, tripling the number of days in each series should only increase the standard deviation of the cumulative total by the *square root* of three. Thus, even if the risk *per day* of the arbitrage portfolios for one-month and three-month options were the same, the standard deviations of the annualized holding-period returns reported in Table VI would appear to drop by a factor of about 1/1.73.

However, the annualized standard deviations for three-month calls are substantially lower than can be accounted for in this way, particularly for the out-of-the-money options. It does seem that the arbitrage is considerably less risky per day for three-month than for one-month options.

Transactions costs for longer holding periods, on the other hand, can be expected to cumulate. At the outset, the lower gammas for three-month options may yield smaller trading costs per day than for one-month options. However, three-month calls eventually become one-month calls, since the standard arbitrage requires the position to be held until expiration. We might therefore expect the total transactions cost for the standard arbitrage to rise monotonically with option maturity. Table VI bears this out. The increase in mean transactions costs ranges from about 30 percent for the 97 strike calls up to 117 percent for the 105s. The standard deviations of hedge returns including costs increase also.

The combination of increased mean hedging cost and increased standard deviation leads to substantially wider arbitrage bands for three-month than for one-month options. For example, to break even on the at-the-money call, an arbitrageur would bid no more than 0.307 below the Black-Scholes price for a one-month option, but he or she would only pay the model prices less 0.431 for a three-month call. The comparable figures to have a 75 percent chance of covering costs would be 0.52 and 0.65, respectively.

The results in Table VI indicate that the standard arbitrage trade becomes less risky when longer dated options are involved. However, trading costs increase with option maturity, leading to wider arbitrage bounds.

B. Market Making and Price Determination in the Options Market

Since a rational options market maker who based his or her trading on the standard arbitrage strategy or any of the variants described in Table IV would insist on much wider bid-ask spreads than are observed in the marketplace, a reasonable conclusion is that actual market makers should probably not follow these strategies. Indeed, observation reveals that they do not.

A typical market maker does not buy an options contract with the expectation that it will be held in inventory and hedged until expiration. Rather, he or she buys it at his or her bid price, anticipating that he or she will sell it again fairly quickly at his or her offer price. More precisely, he or she buys on one implied volatility, hedges the position, and tries to sell as soon as possible on a higher implied volatility.

Options positions that are not turned over immediately are hedged, but not necessarily by setting up the standard arbitrage against the underlying asset. An option may be hedged with another option, or with a related futures contract, rather than with the underlying asset. Normally, a market maker's entire portfolio of options on a given asset is aggregated, with the result that considerable netting out of market exposure may occur. The resulting option position is evaluated not only for its delta, but also for its gamma, and probably its theta (rate of time decay) and its kappa or "vega" (sensitivity to volatility movements), as well. Each of these represents exposure to a type of risk, and one that can potentially be hedged with other options.

This trading strategy embodies important deviations from the model of market making implied by the Black-Scholes and other models based on the standard arbitrage.

First, since arbitrage is not riskless, trading so as to profit from a theoretical mispricing of the option relative to its underlying asset will not be the dominant strategy, as it is in a frictionless world. The supply of arbitrage services to a real market will not be perfectly elastic. As with other trading strategies, traders will take limited positions and carefully weigh the expected profit against the risk.[13] Under the right circumstances, options prices may be allowed to deviate very far from their model values without inducing a large amount of arbitrage trading to push them back into line.

Second, when a trader takes on an options position with the expectation that it will be unwound quickly in the market, the important thing is how the *market* will *price* options in the immediate future. The trader has less interest in the true volatility of the underlying asset than in the option's future *implied* volatility, regardless of whether this is a very good estimate of how volatile the stock price

[13] Figlewski (1988) examines the impact of incomplete arbitrage in the market for NYSE index options. Mispricing of options relative to the underlying index is found to be associated with the use of an alternative (risky) trading strategy involving NYSE index futures.

will actually be over the option's lifetime. Thus, it is perfectly appropriate for a market maker to buy an option that he or she believes is currently overpriced relative to its long run value, if he or she expects that the market will continue to overprice such options for a while and he or she will be able to earn a quick profit by reselling it at his or her ask price.[14]

All of these factors imply that, while the Black-Scholes model may give a great deal of guidance about how one option should be priced to be consistent with other options on the same stock, the force of arbitrage driving options prices to their theoretical values relative to the underlying asset based on the market's best estimate of its true volatility is severely blunted.

C. Option Replication in Portfolio Insurance

Our results appear to contrast markedly with those from studies of portfolio insurance, a well-known application of the same kind of option replication we are examining. Etzioni (1986), for example, uses a methodology similar to ours with daily rebalancing and finds that a portfolio insurance program replicating a one-year, at-the-money put option on a $100 million stock portfolio would cost $745.5 thousand and would have a replication error of only $199.1 thousand.

Several things account for this apparent discrepancy. A major one is that portfolio insurance transactions are generally done in stock index futures contracts rather than the underlying stocks, at approximately one tenth of the cost. Thus, Etzioni's results considerably understate what it would have cost to replicate the put with stock transactions.

A second factor is that, as we have mentioned, portfolio insurance programs try to replicate much longer dated options than we have been examining, which reduces replication error. Moreover, they are often set up to reduce the problem of high gamma near maturity, by targeting a final payoff pattern that is smooth rather than kinked at the strike price.

Finally, it is important to recognize that we are comparing the costs and replication errors of options arbitrage to the option's price, and more specifically to the amount by which it may be mispriced. This is what is relevant for pricing traded options contracts in the market and evaluating a market maker's arbitrage strategy. However, in the context of portfolio insurance, the cost and risk of the program are expressed relative to the total value of the insured portfolio, so they naturally appear much smaller. Thus, the replication error of $199 thousand found by Etzioni seems insignificant relative to the $100 million portfolio being insured, even though it is pretty large compared with the $750 thousand total value of the put options being replicated. Indivisibilities are also not a problem when such a large portfolio is being hedged.

Thus, there is no inconsistency between our results and those from portfolio insurance studies.

[14] Brennan and Schwartz (1988) have taken a first step in modeling the short run optimal trading behavior of an arbitrageur with limited capital in a stock index futures market. Their approach offers a useful starting point for analyzing the more complex option market maker's problem.

VI. Conclusions

In this paper we have addressed a number of issues involved in applying arbitrage-based option valuation models to actual, imperfect, markets. Since these questions are sufficiently complex that a general theoretical treatment is infeasible, we have adopted a simulation approach that allows us to derive accurate answers for specific values of the underlying parameters of the market system. Our results, therefore, are precise but not completely general. We have tried, as far as is possible, to look at cases that accurately reflect the realities of exchange traded stock index options contracts. We would anticipate that other parameter values would lead to qualitatively similar, though quantitatively different, results.

The following are the major conclusions indicated by our results.

- The volatility of the underlying asset is an extremely important determinant of option value, but sampling error makes the ex post volatility in daily closing prices hard to predict accurately even when the true underlying volatility is known. Mistakes in forecasting volatility cause both option values and hedge ratios to be wrong. However, the impact on hedging accuracy is relatively slight, except for out-of-the-money options. For them, the impact is asymmetrical, so that it is substantially worse to underestimate volatility than to overestimate it.
- Indivisibilities, which make it impossible to achieve exactly the right hedge ratio, increase risk in a hedged position but do not have a large effect on expected return. Hedges involving futures contracts that are relatively large will be most affected. If the hedge ratio can be set to the nearest 0.10 (per share) or better, the impact of indivisibilities is limited. However, except for far out-of-the-money options, limiting possible hedge ratios to multiples of 0.25 increased the standard deviation of the hedge return by more than half. A "yes or no" hedge (i.e., $h = 0$ or 1) is highly risky.
- Transactions costs to do the standard arbitrage trade upon which the Black-Scholes model is based are large, even for a market maker. For a retail trader, they are prohibitive. Strategies for reducing transactions costs by rebalancing the hedged position less frequently do not help much: there is a substantial reduction in cost only at the expense of a substantial increase in risk. The tradeoff appears to be slightly more favorable for a strategy of rebalancing only as the hedge ratio moves far enough away from its correct value, rather than rebalancing only after a fixed number of days.
- The normal transactions costs for the standard arbitrage induce arbitrage bounds around the theoretical option value that are substantially wider than the bid-ask spreads that are observed in practice. Partly based on direct observation, we suggest that market makers and others engaged in arbitrage of exchange traded options follow different strategies. They try to achieve quick turnover, thus reducing costs, but this is not a riskless strategy. Quick turnover also implies that traders will be more interested in forecasting the option's implied volatility for the immediate future than the true volatility of the underlying asset.

- The standard arbitrage with longer dated options is exposed to less risk per day than with the one-month calls we examined in the bulk of the paper. However, the total transactions cost to maintain a hedged position through option expiration increases with option maturity, so that the arbitrage bounds on the market price become wider.

One general conclusion suggested by this research is that, while empirical research has shown that option valuation theory plays a very important role in determining prices in real options markets, the impact of market imperfections is also large, and probably larger than many researchers have realized. (Nor have we covered all major imperfections, having left out margin requirements, nonlognormal price paths, and taxes, to name a few.) Under these conditions, the standard arbitrage cited in the literature as the basis of valuation models becomes a weak force to drive actual option prices toward their theoretical values.

We do not currently have a model of option pricing in a market populated by arbitrageurs who engage almost exclusively in non-"standard", short-term, incompletely hedged, arbitrage-like strategies. The standard arbitrage eliminates price expectations and risk aversion from option pricing in a frictionless market. However, within the wide bounds on prices that are all that can be established by the standard arbitrage in an actual, imperfect options market, there is certainly room for these and many other factors to have an influence.

REFERENCES

Beckers, S., 1981, Standard deviations implied in option prices as predictors of future stock price variability, *Journal of Banking and Finance* 5, 363–382.

Black, Fischer, 1976, Studies of stock price volatility changes, in *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economics Statistics Section.*

——— and Myron Scholes, 1972, The valuation of option contracts and a test of market efficiency, *Journal of Finance* 27, 399–418.

Boyle, Phelim, 1977, Options: A Monte Carlo approach, *Journal of Financial Economics* 4, 323–338.

——— and David Emanuel, 1980, Discretely adjusted option hedges, *Journal of Financial Economics* 8, 259–282.

Brennan, Michael J., 1979, The pricing of contingent claims in discrete time models, *Journal of Finance* 34, 53–68.

——— and Eduardo Schwartz, 1988, Arbitrage in stock index futures, Working Paper, UCLA Graduate School of Management.

Christie, Andrew, 1982, The stochastic behavior of common stock variances, *Journal of Financial Economics* 10, 407–432.

Etzioni, Eytan, 1986, Rebalance disciplines for portfolio insurance, *Journal of Portfolio Management* 13, 59–62.

Figlewski, Stephen, 1988, Arbitrage-based pricing of stock index options, *Review of Research in Futures Markets* 7, 250–270.

Galai, Dan, 1977, Tests of market efficiency of the Chicago Board Options Exchange, *Journal of Business* 50, 167–197.

———, 1983a, The components of the return from hedging options against stock, *Journal of Business* 56, 45–54.

———, 1983b, A survey of empirical tests of option pricing models, in M. Brenner, ed.: *Option Pricing* (D. C. Heath, Lexington, MA).

Garcia, C. B. and F. J. Gould, 1987, An empirical study of portfolio insurance, *Financial Analysts Journal* 44–54.

Latane, Henry and Richard Rendleman, 1976, Standard deviations of stock price ratios implied in option prices, *Journal of Finance* 31, 369–382.

Leland, Hayne, 1985, Options pricing and replication with transactions costs, *Journal of Finance* 40, 1283–1301.

Macbeth, James and Larry Merville, 1979, An empirical examination of the Black-Scholes call option pricing model, *Journal of Finance* 34, 1173–1186.

Phillips, Susan and Clifford Smith, 1980, Trading costs for listed options: The implications for market efficiency, *Journal of Financial Economics* 8, 179–201.

Rubinstein, Mark, 1976, The valuation of uncertain income streams and the pricing of options, *Bell Journal of Economics* 7, 407–425.

———, 1985, 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.