Box Spread Strategies and Arbitrage Opportunities

URI BENZION, SHMUEL DANAN, AND JOSEPH YAGIL

URI BENZION
is a professor of finance at Ben-Gurion University in Beer-Sheva, Israel.

SHMUEL DANAN
is a P.h.D candidate at The Technion and Haifa University in Haifa, Israel.

JOSEPH YAGIL
is an associate professor of finance in the School of Business at Haifa University in Haifa, Israel.
yagil@gsb.haifa.ac.il

Methodological problems have so far complicated attempts to examine the box spread strategy. The fully computerized trading system on the Tel-Aviv Stock Exchange and a special computer program that we devised now enable the detection of profitable arbitrage opportunities in real time.

It takes about one second for an arbitrage gain to vanish, but it takes less time than that to detect and exploit the opportunities. Yet since the arbitrage gain is relatively small, shrinks substantially with transaction fees, and disappears quickly with time, the Tel Aviv options market is highly efficient.

Our research represents the first attempt to examine the box spread (BS) strategy with very high precision. The BS arbitrage strategy uses European options with identical underlying assets and expiration dates. Two pairs of call options and put options are combined. Each pair has an identical exercise price but different from the exercise price of the other pair of options. Maintaining a certain position in each of the four options simultaneously produces a synthetic bond.

To achieve arbitrage opportunities by exploiting the difference between the return on the synthetic bond and the risk-free rate of interest (r), the strategy uses one of two types of transactions: 1) a long box spread, which is buying the synthetic bond and borrowing when the synthetic bond return is higher than r, and 2) a short BS, which entails shorting the synthetic bond and lending at r when the synthetic bond return is lower than r.

The box spread strategy is often used to test the efficiency of options markets. Since the strategy involves only options and the risk-free asset, it is particularly appropriate for testing the efficiency of options markets when the underlying asset is not traded.

Research on box spreads has suffered from methodological problems, particularly with regard to the conditions that are necessary to construct the BS strategy: 1) the trades examined were not simultaneous; 2) in several cases, the options studied were American rather than European and the early exercise feature results in a risky BS strategy; and 3) there has been nothing to guarantee that the prices used to construct the strategy were the real prices at which the strategy could have been executed.

In overcoming these three problems, we test the BS strategy more precisely. To obtain the intraday bid-ask prices required for this research, we wrote a special program to install on the computer of a brokerage firm. As the computer received market bid-ask quotes in real time, our program sampled data every four seconds for recording in our database. This allows us to detect arbitrage opportunities in real time and estimate the gain resulting for various scenarios of transaction fees and time delays associated with executing the BS strategy.

Our work is different from earlier studies in that it not only detects the arbitrage opportunities and calculates the resulting gain, but also analyzes the characteristics of the opportunities. Identifying a pattern can help shorten search time and increase the potential gain.

I. LITERATURE REVIEW

Although the box spread strategy offers a simple way to test the efficiency of options markets, only a few articles on the strategy have appeared since options were first traded on an organized exchange in 1973. One reason may be the absence of appropriate data. Furthermore, what research there is suffers from problems associated with computing the arbitrage gain on the strategy.

One of the earliest studies of the BS strategy is Ronn and Ronn [1989]. They use intraday bid-ask prices of Chicago Board Options Exchange American stock options. Their sample consists of eight trading days—one day for each of the eight years between 1977 and 1984. To overcome the problem associated with the early exercise of American options, they examine only long BS strategies. Despite a one-minute time difference between price quotations, only the last price quotes were sampled for each observation. Therefore, there is no guarantee that the strategy could have been executed at those prices. Thus, it is hard to determine whether the gain computed could have been realized.

Ronn and Ronn discover some gain-signaling opportunities, but the gain is relatively small and disappears when transaction fees are considered. Profits on a delayed execution of the strategy can be gained only if the time delay is less than five minutes. Ronn and Ronn find that market efficiency improves over the sample time period.

Marchand, Lindley, and Followill [1994] (MLF) also examine intraday bid-ask prices of CME American S&P 500 futures options. Their time period is 1983-1992; their sampling interval varies but is shorter than five minutes.

Like Ronn and Ronn [1989], MLF also use only the long and not the short BS strategy in order to avoid the early exercise problem associated with American options. Unlike Ronn and Ronn, MLF attempt to resolve the non-synchronization problem in the data using Bhattacharya's [1983] procedure.[1]

Under this procedure, the option price changes determine the boundaries of the option bid-ask prices. If the sampled bid-ask prices do not fall within the boundaries, they are suspected to be non-synchronized and thus will not be used for the gain from the strategy. The problems with this procedure are that, on the one hand, it may ignore profit opportunities that exist, and, on the other hand, some of the profit opportunities that are included may consist of non-synchronized prices.

MLF's findings indicate that the box spread average gain is negative, and that the longer it takes to complete the strategy, the more the profit varies. This result emphasizes the importance of synchronized prices for obtaining unbiased results in empirical tests of the BS strategy.

Hemler and Miller [1997] use the BS strategy to investigate the effect of the October 1987 stock market crash on market efficiency. Their sample consists of the bid-ask prices of S&P 500 European options during the months surrounding October 1987. Despite a 30-second period between observations, Hemler and Miller's work also suffers from the problem of non-synchronized prices. As a result, the reported profits might not have been achievable.

Hemler and Miller report a substantial increase in profit-signaling opportunities post-crash over pre-crash. Abnormal gains are also higher after the crash. For the \$2.50 to \$10.00 range they use for transaction costs, the gain disappears after only a one-minute delay in executing the short BS strategy and after five minutes for the long BS strategy (although these results might stem from invalid prices). Our suspicions are strengthened in light of their other results, which indicate that the more recent the proposed prices in the strategy, the smaller the profit. For prices not ten minutes old, the average gain becomes negative.

Ackert and Tian [2001] use a box spread arbitrage strategy to examine the efficiency of the market for S&P 500 index options. Their sample consists of daily closing option prices from February 1992 to January 1993. The option prices in their sample, however, are closing rather than bid-ask prices, as the strategy requires. To solve this problem, they take the usual spread in options as a substitute for the bid-ask prices. In this case, they find frequent violations of the arbitrage conditions. They also try a different method to derive substitutes for the bid-ask prices using the time to maturity of the option and the exercise price distance from the index level, but their results remain much the same.

Bharadwaj and Wiggins [2001] test box spread arbitrage conditions in the S&P 500 LEAPS market over 1994-1996. Their sample consists of the last bid-ask options prices offered in the market at 12:30 p.m. (The mean time difference between the quotes in their sample is more than one and a half hours.) Assuming no transaction fees, they find very few low-profit arbitrage opportunities.

We attempt to overcome the various shortcomings in methodologies that offer no guarantee the prices used are prices at which the strategy could actually have been executed. The special program we devised and installed on a broker's computer system together with the fully computerized simultaneous trading on the options exchange enable us to detect arbitrage opportunities in real time.

II. THEORETICAL FRAMEWORK

To construct the box spread strategy, one needs four European options characterized by 1) the same underlying asset; 2) the same expiration date; and 3) two different (low and high) exercise prices (EXPRC). The strategy is constructed as follows:

1. Long in low-EXPRC call option.
2. Short in high-EXPRC call option.
3. Long in high-EXPRC put option.
4. Short in low-EXPRC put option.

Because of the negative (positive) effect of the EXPRC on the call (put) price, the value of the options sold in the BS strategy is always lower than the value of the options purchased. Thus, the strategy always requires an initial investment.

The notation is as follows:

t = present time;
T = expiration date of the option;
$T - t$ = time to expiration of the option;
Z_t = present price of the underlying asset;
X = exercise prices;
N = number of exercise prices in the market;
$X = \{X_1, X_2, \ldots, X_N : X_1 < X_2 <, \ldots, < X_N\}$;
X_L (X_H) = lower (higher) of two exercise prices;
C_{xn} (P_{xn}) = price of a call (put) option whose exercise price is X_n; and
r = risk-free rate of interest.

The payoff at expiration from a BS strategy for each of the three possible price ranges of the underlying asset at expiration is shown in Exhibit 1. The payoff equals $(X_H - X_L)$ in each state of nature. Thus, purchasing BS is equivalent to purchasing a synthetic risk-free bond.

The BS strategy, as depicted by Exhibit 2, is a combination of a bull spread and a bear spread, where the bull

EXHIBIT 1
Payoffs from BS Strategy

Price Range For Z	Payoff from Long C_{XL}	Payoff from Short C_{XH}	Payoff from Long P_{XH}	Payoff from Short P_{XL}	Total Payoff
$Z_T \leq X_L$	0	0	$(X_H - Z_T)$	$(Z_T - X_L)$	$(X_H - X_L)$
$X_H > Z_T > X_L$	$(Z_T - X_L)$	0	$(X_H - Z_T)$	0	$(X_H - X_L)$
$Z_T \geq X_H$	$(Z_T - X_L)$	$(X_H - Z_T)$	0	0	$(X_H - X_L)$

Payoff is identical across the three possible ranges for the price of the underlying asset at expiration.

(bear) spread consists of both long and short positions in two identical calls (puts) with different exercise prices. The payoff from the bull and bear spreads is represented by the dashed lines. Combining these two payoff lines vertically yields the payoff from the BS strategy, which is represented by the solid line. The payoff is constant and independent of the value of the underlying asset at expiration.

To examine the potential for arbitrage opportunities, consider two portfolios: (A) the BS strategy portfolio, and (B) investing in a risk-free asset the sum of $(X_H - X_L)e^{-r(T-t)}$. At expiration, the terminal values of both portfolios are identical and equal to $(X_H - X_L)$. Thus, both portfolios should have the same initial value. That is:

$$C_{XL} - C_{XH} + P_{XH} - P_{XL} = (X_H - X_L)e^{-r(T-t)} \qquad (1)$$

for any $X_L < X_H$.

Equation (1) is known as box spread parity. When it is violated, arbitrage opportunities can arise as follows:

1. If Portfolio A [the left-hand side of Equation (1)] is more valuable than Portfolio B [the right-hand side of Equation (1), an arbitrage profit can be made by selling Portfolio A (which is a synthetic bond). Part of the proceeds will be used to buy a risk-free bond, and the remainder is the arbitrage profit. We denote this arbitrage strategy as S_1.
2. If Portfolio B is more valuable than Portfolio A, we can borrow and use part of the proceeds to buy Portfolio A, and the balance is an arbitrage profit. We denote this arbitrage strategy as S_2.

Underlying this analysis are the assumptions that 1) at each time there is only one price for buying and selling an option; and 2) the risk-free rate of interest, r, is the same for lending and borrowing. To simulate reality, we will assume that an option can be purchased at the ask price

EXHIBIT 2
Box Spread as Combination of Bull Spread and Bear Spread

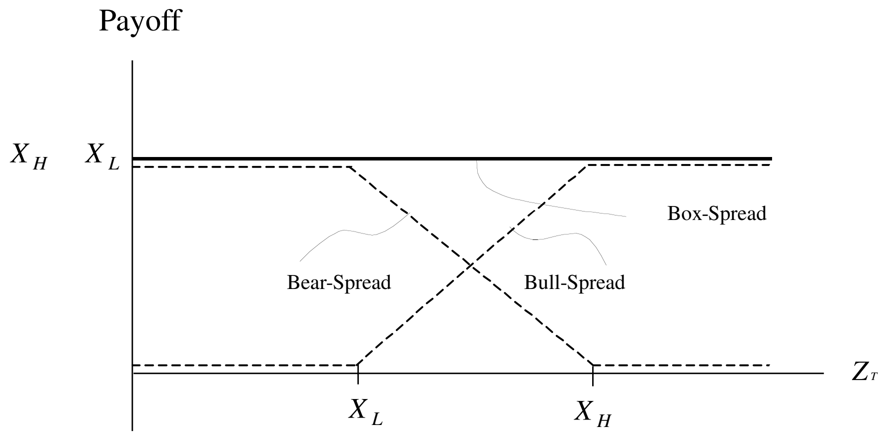

and sold at the bid price, and that the rate of interest for borrowing (r_b) is higher than the rate for lending (r_l). Thus, Strategy S_1 will be profitable only if the proceeds from selling Portfolio A exceed the cost of purchasing Portfolio B; or, formally, if:

$$C^{bid}_{XL} - C^{ask}_{XH} + P^{bid}_{XH} - P^{ask}_{XL} - (X_H - X_L)e^{-rl(T-t)} > 0 \quad (2)$$

Similarly, for S_2, there is an arbitrage profit only if:

$$(X_H - X_L)e^{-rb(T-t)} - (C^{ask}_{XL} - C^{bid}_{XH} + P^{ask}_{XH} - P^{bid}_{XL}) > 0 \quad (3)$$

Our empirical study involves a search for combinations or option price quotations that fulfill the conditions given by Equations (2) and (3).

III. METHODOLOGY AND SAMPLE

To detect arbitrage opportunities, we have to compute the gain from the strategy at each time for different combinations of exercise prices. Let us compute the number of combinations produced by N exercise prices. For constructing Strategy S_1 or S_2 when the exercise price X_1 is the lower one, the combinations should be selected from pairs of exercise prices as follows:

$$\{\{X_1, X_2\}, \{X_1, X_3\}, \{X_1, X_4\}, \ldots, \{X_1, X_N\}\}$$

which total N – 1 pairs. When X_2 is designated as the lower price, the combinations are based on the (N – 2) pairs:

$$\{\{X_2, X_3\}, \{X_2, X_4\}, \{X_2, X_5\}, \ldots, \{X_2, X_N\}\}$$

Consequently, when each and every one of the exercise prices is designated as the lower price, the total number of combinations (n) resulting from N exercise prices is:

$$n = (N - 1) + (N - 2) + (N - 3) + \ldots + 2 + 1 = N(N - 1)/2 \quad (4)$$

For each combination of exercise prices at a given time (t), the gain (G) from Strategy S_1, $G_{S1}(X_L, X_H, t)$, will be given by:

$$G_{S1}(X_L, X_H, t) = (C^{bid}_{XL} - C^{ask}_{XH} + P^{bid}_{XH} - P^{ask}_{XL}) - (X_H - X_L)e^{-rl(T-t)} \quad (5)$$

For Strategy S_2, the gain is:

$$G_{S2}(X_L, X_H, t) = (X_H - X_L)e^{-rb(T-t)} - (C^{ask}_{XL} - C^{bid}_{XH} + P^{ask}_{XH} - P^{bid}_{XL}) \quad (6)$$

Given these definitions of the gain, the next step is to scan all possible combinations for the presence of arbitrage opportunities.[2] How many opportunities there are depends on the definition of an arbitrage opportunity. There are two definitions, and the number of opportunities can be different, depending on the definition.

Exhibit 3 describes a segment of the time series of the gain resulting from Strategy S_1. One possible definition for an arbitrage opportunity is a position involving positive gain. By this definition, Exhibit 3 indicates 12 opportunities.

EXHIBIT 3
Gain on Arbitrage Opportunities over Time

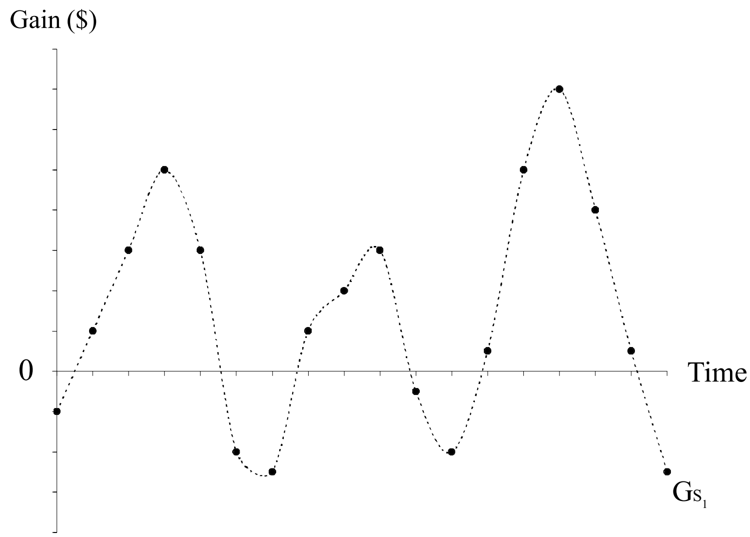

12 observations of positive gain but only three arbitrage opportunities.

By another definition, an arbitrage opportunity starts with the first positive-gain observation after a negative-gain observation, and ends with the last positive-gain observation before a negative-gain observation. According to this definition, the example in Exhibit 3 shows only three arbitrage opportunities.

We use this second definition in our empirical investigation. The rationale is that exploiting the first positive-gain observation can change the prices of the options constituting the strategy. The resulting price change can eliminate the subsequent positive-gain observations once seen.

The data set consists of index options traded on the Tel-Aviv Stock Exchange (TASE). The underlying asset is the Tel-Aviv 25 Stock Index (TA25), which is a weighted average of the 25 Israeli companies with the highest market value in the TASE. Trading in the TA25 options started in 1993.[3]

The weight of each share in the formulation of the index is determined according to its market value or 9.5%, whichever is lower. Weights are updated daily according to market value changes. Directors of the TASE revise the composition of the TA25 "basket" twice a year, in January and July.

The call and put options are European, and they are issued every month for a term of three months. Options are issued at exercise prices in multiples of 10 points in distance from the index.[4]

In 1999, the trading system in options changed from bilateral-sequential, where options are traded one after another by floor traders, to a simultaneous and fully automated trading system carried out by a computer system. In this system, bid-ask quotes are valid as long as they are not withdrawn (by the brokers) or as long as a transaction has not yet taken place. In other words, the quotes are valid as long as they appear on traders' and brokers' computer screens. Thus, to test the BS strategy, all that is needed is to obtain the bid-ask quotes that appear on the computer trading system.

To obtain these data, we wrote a computer program and installed it on the computer of a brokerage firm that is a member of the TASE.[5] The computer received real-time market bid-ask quotes from the TASE; the program sampled the data every four seconds and recorded the result in the database.

The fully computerized trading system has two components that are important for our research: 1) simultaneous trading, and 2) certainty regarding the bid-ask quotes that appear in the system. These elements, in conjunction with the computer program, facilitate this first attempt to test the BS strategy more precisely than ever before.

The sample time period is June and July 2000, which includes 43 trading days, of which we sampled 35.[6] We include in the sample only the option contracts with the nearest expiration dates for two reasons. First, options with the nearest expiration date account for most of the volume in options trading on the TASE. Second, including options with more time to expiration would slow the computer operation and affect the accuracy of sampling the quotes simultaneously.

Constrained by infrequent sampling, the earlier empirical studies had to use options that were around the money (at or slightly in or out). The assumption underlying this sample choice was that only for near-the-money options do price quotes stay valid for some reasonable length of time during which trades at those quotes will take place. Deep out-of-the-money options, however, are less liquid, making it less probable that transactions will be executed at bid-ask prices quoted earlier in the trading day.

This constraint does not apply to our sample because of the TASE computerized trading system. Our sample includes the whole range of exercise prices (distant from the index in multiples of 10 points), including those for deep in- or out-of-the-money options. It thus includes all call and put options with the nearest expiration date, totaling 19 calls and 19 puts.

For each of these 38 options, we collect six variables:

1. Settle price of the last transaction.
2. First bid price.
3. Quantity associated with the first bid price.
4. First ask price.
5. Quantity associated with the first ask price.
6. Total daily number of transactions in the option.

Consequently each sampling comprises 228 data points (6 variables times 38 options). With 4.136 seconds between one sampling and another, the number of samplings per hour totals 870.4 (3,600/4.136). During our two-month period we sample 250.46 hours, bringing the total number of samplings for the two-month period to 218,000 (870.40 × 250.46). With 228 data points in each sampling, the total number of data points in the data set amounts to 49,704,000 (228 × 218,000).

The risk-free rate of interest is measured as the yield to maturity on a three-month Treasury bill. The mean T-bill rate for the sample period was 9.10% with a standard deviation of 0.16%. Accordingly, the lending rate (r_1) associated with time deposits during the sample period was 8.5%. The interest rate for borrowing (r_b) is estimated as the published prime rate for the sample period of 10.8% per year. This was also the rate during the sample period that financial institutions charged investors for borrowed funds used to finance investment in options.

The transaction fee paid by a member brokerage firm is 0.6 new Israeli shekels (NIS) per buy or sell transaction. Brokerage firms charge their clients 2.5 to 10.0 NIS per transaction.[7]

IV. RESULTS

To identify arbitrage opportunities, we compute the gain for each of the exercise-price combinations produced by the 19 exercise prices in the database every four seconds. As indicated in Equation (4), the number of combinations is 171, (19(19 – 1)/2). For each of the two strategies, we have 171 time series that lasted for 35 days, each series characterized by a given combination of strike prices.

We first analyze the pattern of the arbitrage opportunities identified. We then estimate the gain resulting from the arbitrage opportunities. Finally, we investigate the sensitivity of the gain findings to execution delay and transaction costs.

Pattern of Arbitrage Opportunities

The combinations in the sample produce 4,505 profit-signaling arbitrage opportunities, 2,137 by Strategy S_1 and 2,368 by Strategy S_2. The specific pattern revealed by the signaled opportunities in Exhibit 4 has some implications both for the data that should be gathered for testing the BS strategy and for the method of computing the gain from the signals.

Exhibit 4 arranges the distribution of the arbitrage opportunities by the distance of the exercise prices characterizing the opportunities from the TA25 index. If an opportunity is characterized by a high exercise price of 530 points and a low exercise price of 490 points, for instance, and the relevant index level is 521.3 points when the opportunity is detected, the exercise price that is farther (from the index) is less than 40 points away. This opportunity would be characterized as 40 points in the exhibit.

The results in Exhibit 4 indicate that 68% of the opportunities in Strategy S_1 and 56% in S_2 are within 20- to 30-point distance of the index level; 13% to 14% are in the 10-point distance range (note that there is only

EXHIBIT 4

Distribution of Arbitrage Opportunities by Distance of Exercise Price from TA25 Index

Distance from Index (Points)	Distance from Index (%)	Strategy S_1 Number of Oppor-tunities	Percent of total (%)	Cumu-lative (%)	Strategy S_2 Number of Oppor-tunities	Percent of Total (%)	Cumu-lative (%)
10	1.8	288	13	13	337	14	14
20	3.6	983	46	59	804	34	48
30	5.5	461	22	81	523	22	70
40	7.3	247	12	93	290	12	82
50	9.2	100	5	98	172	7	89
60	11.0	32	1	99	108	5	94
other	-	26	1	100	134	6	100
Total	-	2,137	100	-	2,368	100	-

Strategy S_1: Sell higher-priced synthetic bond (options in the BS strategy), buy (lower-priced) risk-free bond.
Strategy S_2: Buy (lower-priced) synthetic bond, and sell (higher-priced) risk-free bond (borrow).

one combination of strike prices at this distance), and the rest lie farther away. Also, merely 1% of the opportunities in S_1 and 6% in S_2 lie in a range exceeding 60 points away from the index level. These results suggest that the signaled opportunities involve exercise prices that are at the money, and these occur less often for options that are deep out of the money.

This property implies that, to identify arbitrage opportunities, there is not much need to construct strategies involving options that are deep out of the money. One can shorten the search time by focusing on options at the money.

For our sample, instead of recording 19 exercise prices that imply 171 (19(19 – 1)/2) combinations, it is sufficient from a practical point of view to focus on 12 exercise prices (within 60 points of the index level) that result in only 66 (12(12 – 1)/2) combinations. This sharp reduction in the number of possible combinations can substantially shorten the search time for an arbitrage opportunity.

In trying to refine this pattern, we find that for the opportunities that emerge at a given time there is a certain structure of strike prices that occurs because the opportunities are produced as a result of a change in the bid or ask price of one option only. Consequently, acting on one of these opportunities can change the offer that caused the emergence of these opportunities in the first place and, as a result, make the other simultaneous opportunities disappear.

In other words, except for the executed opportunities, other opportunities that also appear simultaneously might not be real or realistic. That is, even if the reaction time of the arbitrageur approaches zero, these other opportunities might not be executable. Thus, in computing the arbitrage profit, it is erroneous to assume that one can exploit all opportunities that appear simultaneously.

To characterize these conditions we use terms as follows. A *bunch* is a set of opportunities appearing at the same time; a *focus* is an exercise price that is common to all opportunities in the bunch; and *surrounders* are the exercise prices of the opportunities in the bunch that are not the focus of this bunch. The pattern in our sample is that almost all bunches have a focus; the surrounders are clustered on one side of the focus; and the position of the bunch in relation to the relevant value of the index takes one of two forms:

1. The focus is beneath the index level, and the surrounder farthest from the focus is above the index level. This position of the bunch with respect to the index level will be called Category A.
2. The focus is above the index level, and the surrounder farthest from the focus is beneath the index level. This position will be called Category B.

As is evident in Exhibit 5, Categories A and B are just two of six possible categories of a bunch with respect to the index. The exhibit describes focused bunches with surrounders clustered on one side of the focus. The arrows represent a bunch whose base is the focus and head is the surrounder farthest from the focus. The location of the arrow with respect to the horizontal axis represents the position of the bunch with respect to the index at a given time, I_t.

For example, in Category 2, the focus of the bunch is located beneath the index, while the farthest surrounder is located above the index. Thus, Category 2 in Exhibit 5 in fact represents Category A as defined above, while the opposite holds true for Category 5 in Exhibit 5; it represents Category B as defined above.

Exhibit 6 presents the distribution of bunches by focus, surrounders, and category for Strategies S_1 and S_2. Starting with S_1, of the 2,137 opportunities for Strategy S_1, somewhere under one-half of them (1,186) appeared by themselves, while the other half appeared with at least another opportunity; namely, in bunches. The vast majority of these bunches are focused (95%), meaning that they have an exercise price that is common to all opportunities in the bunch.

Furthermore, almost all the focused bunches have clustered surrounders; i.e., the exercise prices that are not the focus are clustered on one side of the focus. Slightly more than one-half of these clustered surrounder bunches are in Category A, and the remainder are in Category B. Similar results are obtained for Strategy S_2.

Let us look, for example, at bunch size 2 for Strategy S_1 in Exhibit 6. At this bunch size, 246 bunches appeared of which 240 were focused, and for 239 of them the surrounders were clustered. Of these 239 bunches, 112 bunches are in Category A and 85 in Category B. The remaining 42 bunches are not in Category A or B.

What accounts for the pattern described by Exhibit 6, where bunches emerge as a result of a change in the bid-ask price of one option? Suppose that bunches emerge as a result of a change in the bid or ask prices of the options (in a given strategy) that have a different exercise price. A large enough change in the bid-ask prices of the options will in this case generate more opportunities with different combinations of exercise prices. As a result, bunches with two common exercise prices—each common to only part of the opportunities in the bunch—will appear.

The results in Exhibit 6, however, indicate that a vast majority of bunches are focused (namely, have an exercise price that is common to all opportunities in the bunch). Might bunches appear as a result of a change in both bid-ask prices of the options in the strategy that has the same exercise price?

If a bunch belongs to Category A, its focus is the lower exercise price (X_L), and it probably appears as the result of a change in the bid-ask price of at least one of two options whose exercise price is X_L. Given that for Category A, X_L is under the index level, there are greater changes in the bid-ask prices of the in-the-money call C_{XL} in the strategy than in the out-of-the-money put P_{XL}. Thus, it is likely that the bunch emerged as the result of the change in C_{XL} (C^{bid}_{XL} for Strategy S_1, and C^{ask}_{XL} for S_2). If, however, the bunch belongs to Category B, its focus is the higher exercise price (X_H), and it probably emerges because of the change in the bid-ask price of at least one of two options whose exercise price is X_H. Since X_H is above the index, the changes in Px_H are greater than the changes in C_{XH}. Thus, the bunch probably appeared due to the change in P_{XH} (P^{bid}_{XH} for S_1 and P^{ask}_{XH} for S_2).

This argument suggests two hypotheses:

H1: Category A bunches emerge due to a change in the bid-ask price of the lower-exercise price call option (C_{XL}).

H2: Category B bunches emerge due to a change in the bid-ask price of the higher-exercise price put option (P_{XH}).

To test Hypotheses H1 and H2, we identify the opportunity with the greatest gain for every bunch. For this opportunity we recompute the gain under the assumption that the bid-ask prices of C_{XL} and P_{XL} (in H1 and H2, respectively) have not changed for both the moment preceding the time the opportunity appeared and the moment the opportunity appeared. The test results are presented in the last two columns of Exhibit 6.

Most Category A or B bunches would have disappeared, had the change in C_{XL} or P_{XH} not occurred. For example, of the 112 Category A bunches in bunch size 2 for Strategy S_1, 90 bunches have disappeared without a price change in the lower-exercise price call option. Similarly, of the 85 bunches in Category B, 61 bunches have disappeared without a price change in the higher-exercise price put option. The results presented in Exhibit 6 support H1 and H2.

Identifying the option whose price change created the bunch gives the arbitrageur an important advantage over competitors who would attempt to exploit the strategy. Here is the reason. Once the arbitrageur's computer program identifies the strike prices that characterize the bunch, the program has to determine two items: 1) of all opportunities that appear in a given moment, which one is worth exploiting, and 2) in what sequence should

EXHIBIT 5
Bunch Locations Relative to Index

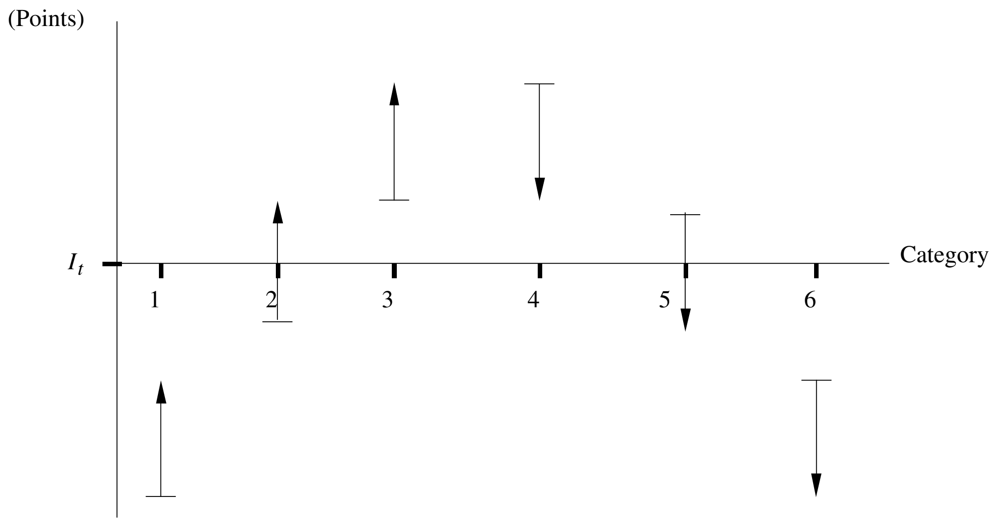

Six possible categories of focused bunches with respect to the index. Arrows describe a bunch whose base is the focus and whose head is the surrounder farthest from the focus. Categories 2 and 5 are equivalent to Categories A and B analyzed.

market buy and sell orders be launched.

Identifying the option whose price change created the bunch is critical because the first market order launched is the one relating to the option whose price change created the bunch.

Arbitrage Gain Findings

In computing the arbitrage gain, we make three assumptions.

Assumption 1. Every exploited opportunity is executed at the signaled prices of the opportunity (namely, the bid-ask prices at the time the opportunity emerges), and in an amount equal to the maximum quantity available at the time of the signal. To elaborate, the strategy combines the purchase (sale) of one unit from each option in the strategy. Thus, the maximum number of portfolios that can be constructed when the signal appears is equal to the quantity offered of the option for which the offered quantity is lowest among the options in the strategy.

Let us denote the maximum number of portfolios that can be constructed from a certain combination of exercise prices by q^{max} (S, X_L , $X_{H,}$ t), or, in short q^{max} , where S is either Strategy S_1 or S_2. The gain associated with this maximum portfolio strategy, $G_S{}^{max}$, will be

$$G_{S1}{}^{max} = q^{max}G_{S1} \quad (7)$$

and

$$G_{S2}{}^{max} = q^{max}G_{S2} \quad (8)$$

where G_{S1} and G_{S2} are expressed by Equations (5) and (6).

Assumption 2. The arbitrageur elects to exploit only one opportunity of all opportunities available at a given moment. This opportunity is the one offering the maximum gain for executing only one portfolio. We have noted that it would be incorrect to assume that all opportunities for a given moment could be executed. This, however, does not imply that only one opportunity can be executed. To elaborate, let us first denote by H^* the option whose change in price leads to the appearance of the bunch, and by q^* the quantity of Option H^* offered at the time that the signal from H^* appears. Consider two cases:

1. q^* is equal to q^{max} of the opportunity in the bunch that has the highest gain per portfolio execution. In this case, fully executing the opportunity in the bunch that has the highest gain per one-portfolio execution (by using the total number of options available) will make the other opportunities in the bunch disappear (assuming the next bid-ask price of Option H^* is such that the opportunity disappears).
2. q^* is greater than q^{max} of the opportunity in the bunch that has the highest gain per one-portfolio execution. In this case, fully executing the oppor-

EXHIBIT 6

Distribution of Bunches by Focus, Surrounders, and Category for Strategies S_1 and S_2

Bunch Size	Number of Bunches	Number of Focused Bunches	Number of Bunches with Clustered Surrounders	Number of Category A Bunches	Number of Category B Bunches	Number of TCD Category A Bunches	Number of TCD Category B Bunches
Strategy S_1							
1	1,186						
2	246	240	239	112	85	90	61
3	74	68	68	30	36	25	31
4	24	21	21	12	7	12	7
5	15	14	14	5	9	5	9
6	6	5	5	2	3	2	3
Over 6	3	2	2	1	1	1	1
Total	368	350	349	162	141	135	112
Strategy S_2							
1	1,180						
2	263	259	256	98	104	47	82
3	82	73	73	42	23	22	22
4	24	22	22	11	10	10	10
5	23	22	22	12	10	10	9
6	12	12	12	8	4	7	4
Over 6	16	13	13	10	3	10	3
Total	420	401	398	181	154	106	130

Bunch = set of opportunities appearing at the same time; Bunch size = number of arbitrage opportunities in a bunch; Focus = exercise price common to all opportunities in the bunch; Surrounders = exercise prices of opportunities in the bunch that are not the focus; Category A (B) Bunch = a bunch where focus is beneath (above) the index level, and the surrounder farthest from the focus is above (beneath) the index level; Clustered surrounders = surrounders located on one side of the focus; TCD = "that-can-disappear" bunches.

tunity will leave q^{max} of the opportunity in the bunch that has the highest gain per one-portfolio execution, less q^* units from Option H^*, that can be used for other opportunities in the bunch. Since executing the opportunities in this case complicates the strategy and takes time (which may result in losses), we decided to adopt the conservative assumption that at a given moment one can execute only one opportunity out of all the opportunities that emerge at the same moment.

Assumption 3. The third assumption applies to the transaction fee of executing the strategy. Two alternative scenarios are assumed:

1. No transaction fee. In this case the gain is computed using Equations (7) and (8).
2. A fixed strategy fee of 20 NIS (4 options × 5 NIS per option), paid when the strategy is executed. In this case the gain will be calculated by subtracting $5 \times q^{max}$ from Equations (7) and (8).

Exhibit 7 presents the number of arbitrage opportunities and the gain from them (with and without transaction fees) for Strategies S_1 and S_2 on each trading day. Of a total of 4,605 signaled arbitrage opportunities (2,368 for S_1 and 2,137 for S_2), 3,154 opportunities are found when we use Assumptions 1–3 in computing the arbitrage gain (1,554 for S_1 and 1,600 for S_2).

The last row in Exhibit 7 shows the total arbitrage gain on the 3,154 opportunities ranges between 146,499 NIS (62,313 NIS for S_1 and 84,186 NIS for S_2) when no transaction costs are paid to 27,073 NIS (10,801 NIS for S_1 and 16,272 NIS for S_2) when transaction fees of 20 NIS per opportunity are paid. This fee level substantially reduces both the number of opportunities (from 3,154 to only 252) and the total arbitrage gain (from 146,499 NIS to 27,073 NIS). The percentage reduction is much greater for the number of opportunities (252/3,154 – 1 = 92%) than for the gain (27,073/146,999 – 1 = 82%).

This result implies that the vast majority of the opportunities (2,902/3,154 = 92%) involve an arbitrage gain that is exceeded by the fee level assumed. Furthermore, when taking into account all 4,505 signaled opportunities, the percentage of positive-gain opportunities (net of transaction fees) amounts to 5.6% only (252/4,505).

Sensitivity of Results to Execution Delay

The results in Exhibit 7 indicate that profitable arbitrage opportunities do exist. An important question is whether it is possible to take advantage of them. Assumption 1 is that there is no delay in executing the options strategy; namely, the realized gain is the signaled gain. In reality, however, there will always be some time (as short as it might be) between the time an opportunity is detected and the time it is executed.

To investigate the extent to which the opportunities signaled could have been exploited, we compute the arbitrage gain from the BS strategy under various scenarios of execution time delay ranging from one to four time delays, and compare each with the no-delay gain. Recall that we sample the data every 4.136 seconds, so that one time delay is equivalent to 4.136 seconds. For each level of execution time delay, we compute the gain including transaction fees. The fee structure will start with zero and increase at a constant sum of 0.50 NIS per portfolio until no profitable opportunities are left.

We should note a new problem that arises when the execution time delay is positive. By Assumption 2, the number of portfolios used in exploiting the chosen opportunity is the maximum available. This assumption works when there is no time delay, but when there is a delay, there can be fewer portfolios at the execution time than the maximum number during the signaled time. In this case, a certain number of portfolios will be executed at the first bid-ask price, and the remainder at the subsequent bid-ask prices. We cannot compute the gain in this case, because our database includes only the first bid-ask price. To resolve this problem, we will assume instead that only one portfolio (instead of the maximum available) can be used to exploit the opportunity at the prices prevailing during the execution delay time assumed.

In order to find the gain for varying transaction fees and time delays, we employ an algorithm as follows:

1. The initial position involves no transaction fee.
2. We identify the first gain observation of all opportunities that appear for the fee level assumed.
3. For each opportunity in (2) above, we compute the gain from the strategy executed in time delays of 0, 1, 2, 3, and 4.
4. We sum the gain in (3) above for all opportunities and for each of the delay values assumed.
5. We increase the transaction fee above by 0.5 NIS per portfolio and repeat steps (1) – (4) above until we reach a fee level in (1) where no future profitable opportunities appear.

EXHIBIT 7

Number of Arbitrage Opportunities and Gain (with and without transaction costs) for Strategies S_1 and S_2 on Each Trading Day

			Strategy S_1				Strategy S_2			
Date	Day	T	n_a	n_b	G_a	G_b	n_a	n_b	G_a	G_b
05/28	S	32	13	3	838	194	25	4	708	141
05/29	M	31	7	1	266	89	12	1	181	23
05/30	T	30	19	0	310	0	25	5	796	242
05/31	W	29	33	8	2,213	886	46	6	1,966	796
06/1	R	28	26	3	1,580	158	37	7	2,140	386
06/4	S	25	34	6	2,739	1,023	28	12	2,629	1,430
06/5	M	24	17	1	926	256	23	2	713	143
06/6	T	23	32	4	1,425	301	30	2	1,281	62
06/7	W	22	45	1	1,026	21	51	2	1,762	143
06/11	S	18	36	2	1,979	161	18	5	1,694	661
06/12	M	17	70	4	3,006	116	46	11	4,997	2,542
06/13	T	16	47	6	2,119	220	60	12	4,253	1,010
06/14	W	15	64	9	3,370	991	50	5	3,047	1,023
06/21	W	8	55	4	1,942	309	68	3	3,734	263
06/22	R	7	91	6	4,688	1,682	69	4	3,518	745
06/25	S	4	65	0	1,400	0	120	2	4,670	161
06/26	M	3	94	2	2,189	204	56	2	2,493	241
06/27	T	2	60	2	667	3	65	3	2,744	285
06/28	W	1	161	7	5,595	390	142	6	9,684	223
06/29	R	28	21	3	1,320	356	36	4	1,366	251
07/2	S	25	13	1	270	0	18	0	274	0
07/3	M	24	29	3	906	352	25	2	779	35
07/4	T	23	45	3	1,726	85	23	2	779	164
07/5	W	22	4	0	43	0	2	0	42	0
07/6	R	21	28	0	861	0	2	1	1,062	38
07/9	S	18	45	3	1,720	139	64	11	3,205	609
07/10	M	17	104	8	3,891	958	48	6	2,600	717
07/11	T	16	22	2	882	177	28	2	987	329
07/12	W	15	23	1	551	1	28	0	1,213	0
07/13	R	14	38	3	2,321	766	55	2	1,853	227
07/16	S	11	23	0	384	0	24	1	693	121
07/17	M	10	19	1	1,229	94	13	1	588	24
07/19	T	8	58	8	2,918	708	74	17	7,066	3,235
07/25	W	2	61	2	2,222	10	74	0	3,435	0
07/26	R	1	52	2	2,792	151	79	0	5,235	0
Total			1,554	109	62,313	10,801	1,600	143	84,186	16,272

T = days to expiration of options; n = number of arbitrage opportunities; G = daily gain from opportunities in NIS; subscripts a and b denote without and with transaction fees, respectively.

Exhibit 8 depicts the number of arbitrage opportunities for different levels of transaction fees and execution delays. The time delay (d) varies from zero to four. For a zero delay, the vertical axis in Exhibit 8 indicates 1,554 and 1,600 opportunities for S_1 and S_2, respectively. These are the values found in Exhibit 7.

For a given delay value, the number of opportunities drops at a declining rate, implying many low-gain opportunities and a few high-gain opportunities. Also, for a given fee level, the number of opportunities drops at a declining rate with the time delay (the vertical gap between the curves in Exhibit 8 narrows with the time

EXHIBIT 8

Number of Arbitrage Opportunities for Different Transaction Fees and Execution Delays

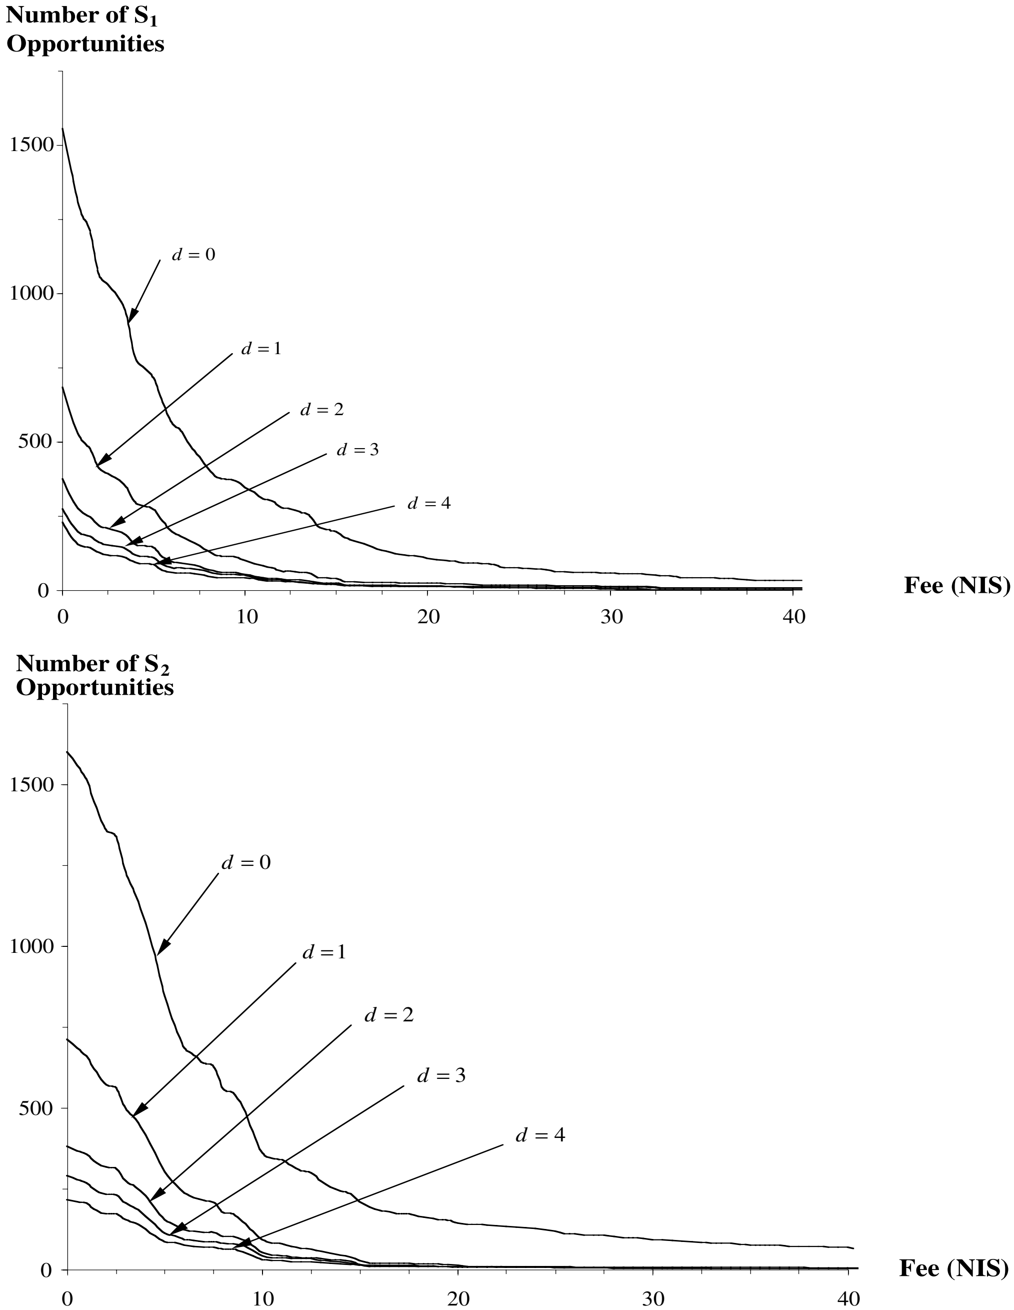

delay). This finding implies that there are many opportunities that disappear quickly and only a few that last longer.

The arbitrage gain, as a function of the transaction fee and execution delay, is depicted in Exhibit 9. The plots demonstrate that the gain is less sensitive to changes in the fee level when there is no delay; that is, even with a relatively high fee, the gain remains positive.

A notable result that Exhibit 9 makes clear is that the gain is very sensitive to the execution delay. A single delay is sufficient to cause the total gain to be negative for any fee level, even for a zero-fee case. A comparison of the two strategies, S_1 and S_2, in Exhibit 9 indicates that for any fee level the no-delay gain for S_2 exceeds the gain for S_1, but the loss resulting from the execution delay is higher

E X H I B I T 9

Arbitrage Gain as Function of Transaction Fee and Execution Delay

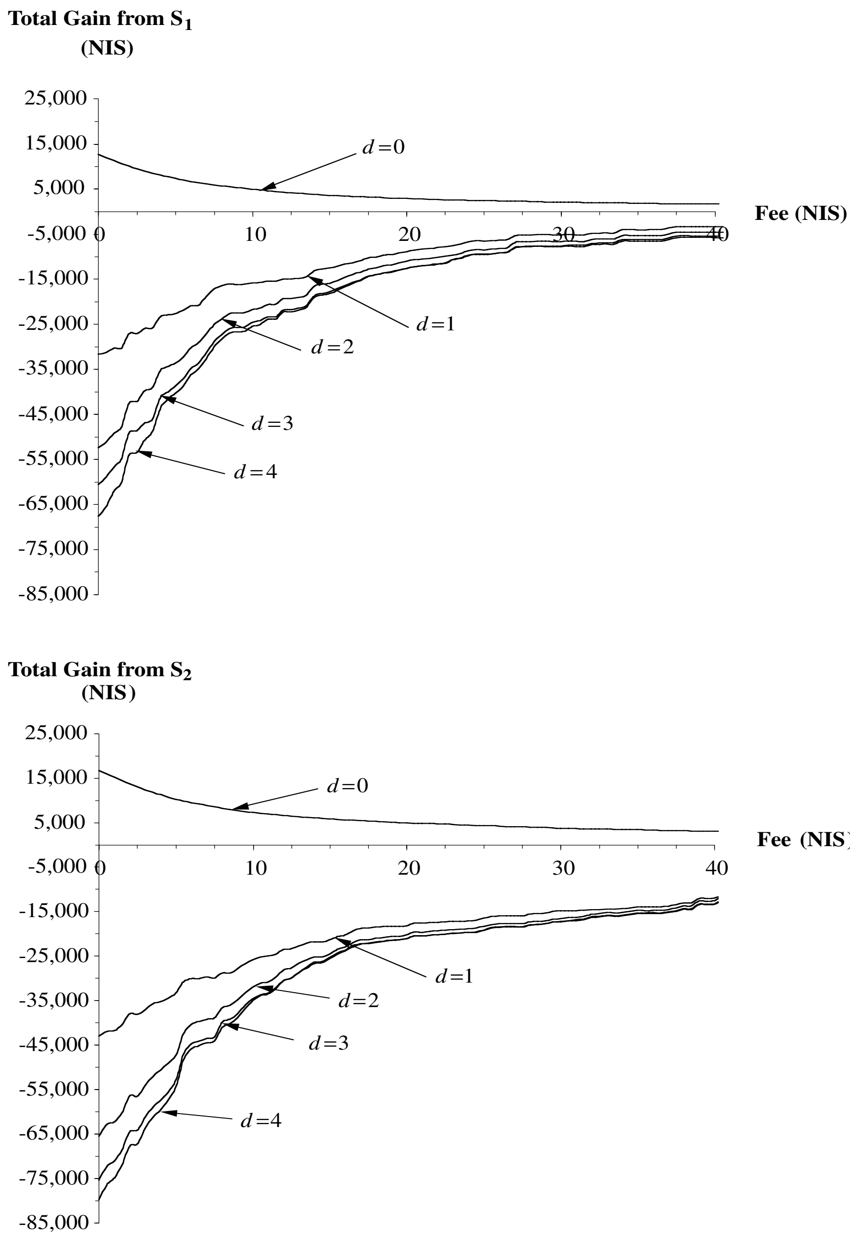

for S_2 than for S_1. To put this differently, the reduced gain in S_1 is offset by the reduced risk of a greater loss when execution delay time increases.

Exhibit 9 presents another interesting phenomenon as well. For the no-delay case, the gain declines monotonically with the transaction fee, indicated by the falling smooth curve. For the delay cases, however, the loss is diminished with the fee, but the reduction, as seen in the non-smooth increasing curves in Exhibit 9, is not monotonically smooth.

This phenomenon occurs because the fee has two opposing effects on the gain when the time delay is positive. First, an increase in the fee eliminates opportunities whose time delay causes a loss. As a result, the total gain increases, or, equivalently, the total loss is diminished. Second, the opportunities still present following the increase in the fee now produce reduced gains, and the resulting total gain falls. Thus, the net change in the total gain resulting from an increase in the fee depends on which of the two effects is stronger. For an immediate (no-delay) execution, only the second effect is present, which means the total gain declines with the cost.

Although our sampling intervals are much shorter than in other research, we cannot determine how long it takes the total arbitrage gain to vanish. The findings indicate only that with no execution delay the gain is positive, while with even one delay the gain becomes negative.

Exhibit 10 posits the (time-delay) speed at which the gain falls until it vanishes. The horizontal axis in Exhibit 10 now is the delay time rather than the transaction fee as in Exhibit 9. The exhibit depicts the gain as a function of the delay for five fee levels (F). The delay variable also consists of five values.

The gain curves in Exhibit 10 indicate that it takes roughly one-quarter of the time between two delay observations for the profit to vanish, or about one second (given the four seconds between samplings). The computer we used can scan the 66 combinations in order to detect profitable opportunities and to exploit one of them in under one second. This implies it is possible to obtain arbitrage gains.[8]

V. SUMMARY

We have used the fully computerized trading system on the Tel-Aviv Stock Exchange and a special program installed on a broker's computer trading system to examine the box spread (BS) strategy with a high degree of precision. This allows us to detect arbitrage opportunities in real time and compute gains for various scenarios of transaction fees and time delays. The sample of index options traded on the TA25 stock index (TA25) in June-July 2000 includes not only near but also deep in- or out-of-the-money options.

To detect arbitrage opportunities using two BS strategies, we sampled bid-ask prices every 4.136 seconds, producing close to 50 million data points. Most of the opportunities were at-the-money options; there are less often opportunities for options that are deep out of the money.

About half of the opportunities appear in groups with an exercise price that is common to all opportunities in the group as well as exercise prices clustered to one side. Exploiting one opportunity in a group can automatically eliminate the other opportunities. Or, trying to exploit all the opportunities that appear at a given time can lead to gains in only one and inevitable losses in the others. This finding implies that some so-called arbitrage opportunities are merely an illusion. The option whose bid-price change created the group is the option whose trade order should be executed first.

Our time delay sensitivity analysis demonstrates that the arbitrage gain becomes negative for even a four-second delay. A combined delay and fee analysis shows that as fee and delay increase, many low-gain opportunities dissipate quickly, but high-gain opportunities last longer. Because arbitrage opportunities are generated mainly for at-the-money options, the search for opportunities can be shortened substantially.

The arbitrage gain vanishes in about one second, longer than it takes our computer program to detect these profitable opportunities. Given that execution time of option trading on the TASE is about 0.04 seconds, it should be possible to exploit the profit-signal opportunities. Even if they pay high transaction fees, small investors can profit if they act quickly.

Our findings do not imply that the TASE option market is inefficient. First, our 31,000 profitable arbitrage opportunities represent merely 1.4% of the samplings. Second, most of these opportunities produce small gains, and can therefore be exploited only by low-fee traders. Third, the overall total gain found shrinks substantially with transaction fees and disappears quickly with time delays. Thus, while arbitrage profits can be made, they hardly pose a serious challenge to the efficiency of the options market in Tel Aviv.

EXHIBIT 10
Arbitrage Gain for Given Transaction Fee as Function of Execution Delay

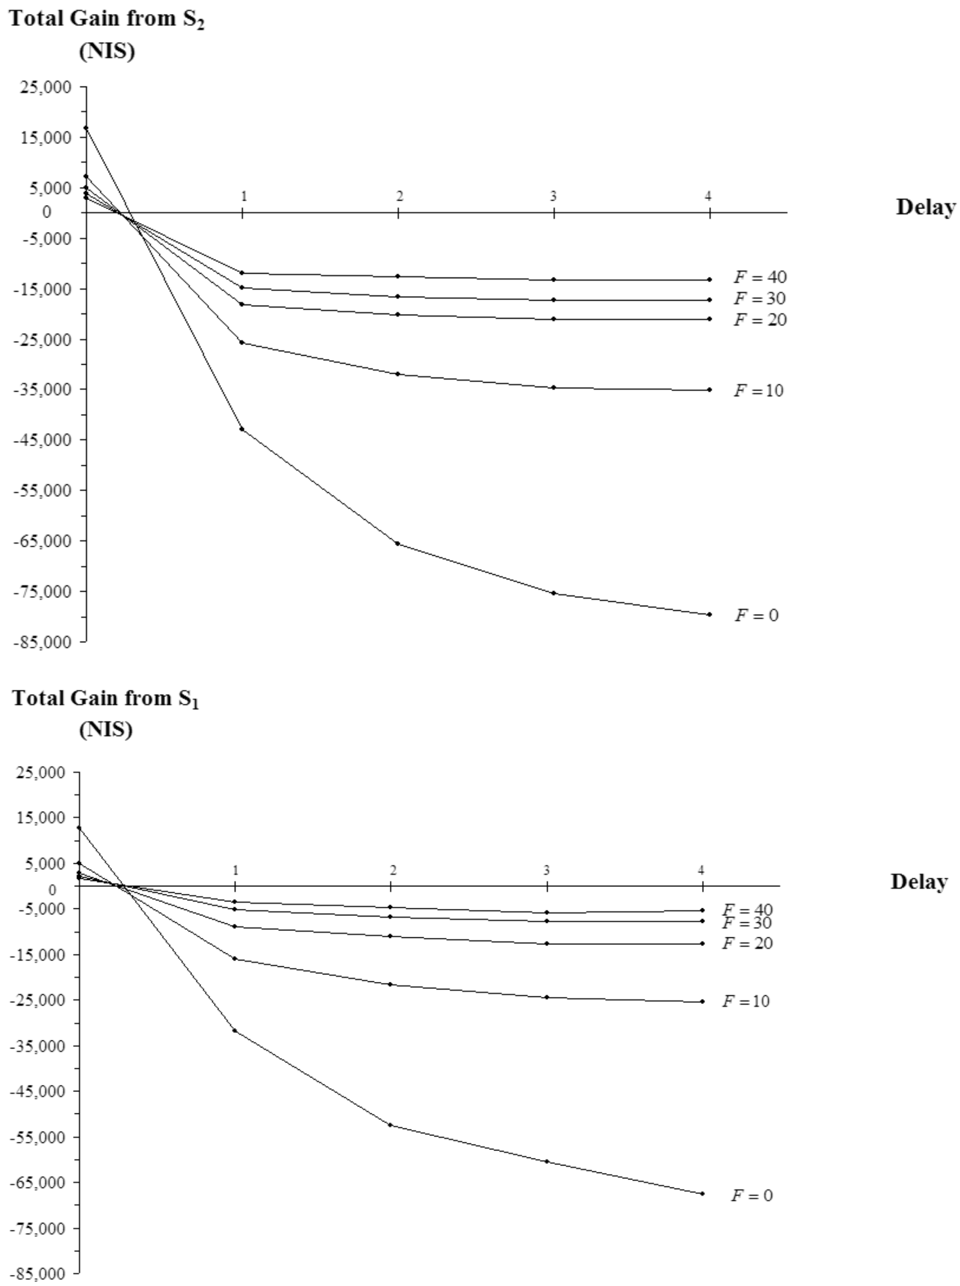

ENDNOTES

[1]A non-synchronization problem arises when option bid-ask prices used in the arbitrage strategy are no longer valid at the time the strategy is executed. Because there is a lag between the time an opportunity is detected and the time it is executed, bid-ask prices can change.

[2]The gain is certain net present value (NPV), and is "above-normal" when it is positive.

[3]During our sample period (June-July 2000) the TA25 ranged between 514 and 566 points with a mean of 546 points and a standard deviation of 10 points.

[4]For exercise prices very close to the prevailing index level, the multiple of exercise prices is five points. There were

only two such options during our sample period, and we excluded them in order to maintain the speed of recording the data.

[5]We thank Dani Gonen, manager of the Haifa branch of Ilanot-Batuha Brokerage Firm (a TASE member), for letting us install the computer program on the firm's trading computer system (which is directly connected with the TASE).

[6]In 8 of the total 43 trading days, the computer with our program did not get turned on.

[7]During the sample period, the exchange rate between the U.S. dollar and the new Israeli shekel ranged between 4.07 to 4.17 NIS, with a mean of 4.10 and a standard deviation of 0.03.

[8]Exhibit 3 might give the impression that restricting trading to just the first observation of the arbitrage opportunity may reduce profitability (the first observation in all the arbitrage opportunities in Exhibit 3 involves lower profits than in subsequent observations). To test whether "waking up" the system when the first observation of an opportunity appears, but trading only on the second one may yield more profit than jumping at the first chance, we examine all the arbitrage opportunities detected.

We find no arbitrage opportunities for which the profit in the first observation was less than the profit in one of the subsequent observations. In other words, restricting trading to the first observation of the arbitrage opportunity is better than using the first observation of the opportunity as a signal to act on one of the subsequent observations. This result holds however quickly the strategy is executed.

REFERENCES

Ackert, Lucy, and Yisong Tian. "Efficiency in Index Options Markets and Baskets." *Journal of Banking and Finance,* 25(9) (September 2001), pp. 1607-1634.

Bharadwaj, Anu, and James B. Wiggins. "Box Spread and Put-Call Parity Tests for the S&P 500 Index LEAPS Market." *The Journal of Derivatives,* 8(4) (Summer 2001), pp. 62-71.

Bhattacharya, M. "Transactions Data Tests of Efficiency of the Chicago Board Options Exchange." *Journal of Financial Economics,* 12 (December 1983), pp. 161-185.

Hemler, M.L., and T.W. Miller. "Box Spread Arbitrage Profits Following the 1987 Market Crash: Real or Illusory?" *Journal of Financial and Quantitative Analysis,* 32(1), (March 1997), pp. 71-90.

Marchand, P., A. Lindley, and R.A. Followill. "Further Evidence on Parity Relationships in Options on S&P 500 Index Futures." *Journal of Futures Markets,* 14(6) (September 1994), pp. 757-771.

Ronn, A.G., and E.I. Ronn. "The Box-Spread Arbitrage: Theory, Tests and Investment Strategies." *Review of Financial Studies*, 2(1) (Spring 1989), pp. 91-108.