Ex-dividend Arbitrage in Option Markets

Jia Hao
Wayne State University

Avner Kalay
Tel Aviv University and University of Utah

Stewart Mayhew
U.S. Securities and Exchange Commission

We examine the behavior of call options surrounding the underlying stock's ex-dividend date. The evidence is inconsistent with the predictions of a rational exercise policy; a significant fraction of the open interest remains unexercised, resulting in a windfall gain to option writers. This triggers a sophisticated trading scheme that enables short-term traders to receive a significant fraction of the gains. The trading scheme inflates reported volume and distorts its traditional relations to liquidity. The dramatic increases in the volume of trade on the last cum-dividend day are facilitated by limitations on transaction costs passed by the various option exchanges. (*JEL* G13, G14, G18)

Holders of equity call options are not entitled to receive the cash dividend paid to owners of the underlying stock, unless they exercise the calls prior to the ex-dividend date. Consequently, some owners of American-style call options have an incentive to exercise immediately before the stock goes ex dividend (Roll 1977). Essentially, exercise on the last cum-dividend day will be optimal if the value of the dividend exceeds the "time value" remaining in the option after the dividend. This is most likely to be the case for deep-in-the-money and short-term call options.

Previous research indicates that option owners do not always follow the optimal exercise strategy. Numerous studies (e.g., Finucane 1997; Kalay and Subrahmanyam 1984; Poteshman and Serbin 2003) document that option holders sometimes err in exercising when they should not, as well as failing to

---

We wish to thank Yakov Amihud, Peter Demarzo, David Dubofsky, Michael Gordy, Bjorn Jorgensen, Elizabeth King, Matt Spiegel, and participants of the finance seminar at the Indian School of Business, IDC, the Hebrew University, the University of New South Wales, Queensland University of Technology, the University of Utah, Tel Aviv University, the Securities and Exchange Commission, the University of Kansas, Skinance 2007, CICF 2008, and WFA 2008 for many helpful comments. Part of this research was undertaken while Kalay was visiting academic scholar at the Securities and Exchange Commission in 2004. The Securities and Exchange Commission, as a matter of policy, disclaims responsibility for any private publication or statement by any of its employees. The views expressed herein are those of the authors and do not necessarily reflect the views of the Commission or of the author's colleagues on the staff of the Commission. All errors and omissions are the sole responsibility of the authors. Send correspondence to Jia Hao, School of Business Administration, Wayne State University, Prentis Building Room 313, 5201 Cass Ave., Detroit, MI 48202; telephone: 313-577-5059; fax: 313-577-5486. E-mail: jia.hao@wayne.edu.

doi:10.1093/rfs/hhp038 Advance Access publication May 21, 2009

exercise when they should. If some option owners fail to exercise optimally on the last cum-dividend day, this generates a windfall benefit, accruing to those option writers who are not assigned an exercise. In the exchange-traded equity option market in the United States for the period 1996–2006, we estimate that more than 40% of the call options that should have been exercised remain unexercised.

In this article, we describe a trading scheme known as a "dividend play," through which market makers or other arbitrageurs extract these gains, diverting them away from the original option writers. The strategy exploits the mechanics of the allocation algorithm employed by the clearinghouse to assign option exercises, and involves two parties executing large offsetting buy and sell call option transactions on the last cum-dividend day. Our findings indicate that this activity generates substantial trading volume in call options, in some cases large enough to significantly affect statistics such as average daily volume, put-call ratios, and exchange market share.

The dividend play strategy tends to be most profitable where there is an in-the-money, short-term call option series on a stock that pays a high dividend (so that early exercise is optimal), and where the series has a significant unexercised open interest on the last cum-dividend date. In its simplest form, the strategy is executed as follows. Trader 1 buys a large number of calls from trader 2, who immediately buys an offsetting position back from trader 1. Because the two trades are exactly offsetting and executed at the same price, the initial position has zero risk and requires no capital. The two traders then exercise all their long positions.

At the end of the day when the clearinghouse adjusts its accounts to reflect the day's activity, purchases are processed before exercises are assigned, but sales that close out open positions are processed after assignment. Thus, when the clearinghouse assigns the day's exercises across all option writers, the resulting assignments will close out a large portion of the two traders' new short position, but will also close out a large portion of the preexisting short positions. The larger the positions taken by the two traders, compared to the preexisting open interest, the higher the proportion of preexisting shorts that will get forced out of their position, and the greater the proportion of the benefit extracted by the traders.[1]

Footprints of this activity are readily evident in large spikes in call volume observed immediately prior to ex-dividend days on stocks that pay large dividends. As an example, Figure 1 shows daily contract volume on Altria options, surrounding the $0.68 dividend on 11 March 2004, based on data reported by the Options Clearing Corporation (OCC). The contract volume on 10 March

[1] Appendix A contains background information on the process through which option trades are cleared at the Options Clearing Corporation (OCC).

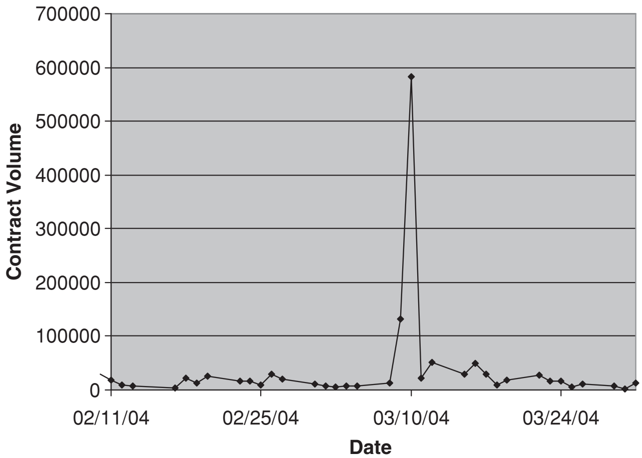

Figure 1
Example of excess trading volume resulting from dividend arbitrage
This figure graphs daily contract volume, as reported by the OCC, in Altria (formerly Philip Morris) options, from 5 February 2004 to 31 March 2004. The reported volume number is aggregated across all Altria series (excluding LEAPS) and across option exchanges. 11 March 2004 was the ex date for a dividend of $0.68 per share.

was roughly thirty times the typical level. Overwhelmingly, this volume spike was composed of trading in call options by market makers.[2]

The ability of arbitrageurs to capture these gains is limited only by the marginal cost of executing and clearing trades. During our sample period, in the year 2003, two exchanges adopted new rules that put a cap on total transaction fees, effectively dropping the marginal transaction fee to zero for active market makers. Because these rule changes significantly lowered the cost of implementing the strategy, we would expect to see more dividend play activity on these two exchanges after the rule change.

This article explores the economics underlying the dividend play strategy, and empirically investigates the extent to which the dividend play strategy is employed, the magnitude of the potential profits, the degree to which the transaction fee reductions facilitated the strategy, and the extent to which the reported volume numbers are affected. We show that if the marginal cost for some market participants is low, the optimal size of a dividend play arbitrage trade can easily be multiple times higher than the open interest prior to the ex-dividend date. We find that abnormal option trading volume on the last cum-dividend date is strongly related to our estimate of the potential profitability of the dividend play. We also find evidence that the rule changes decreasing the marginal transaction fee had a large impact on the volume of dividend

[2] As reported by the OCC, trading volume on that day was composed of 97% call options and 3% put options; market makers accounted for 89% of the volume, non-market-maker firms accounted for 6%, and customer transactions accounted for 5%.

play activity. Finally, we find evidence that dividend play volume is sufficiently large and it has a significant impact on trading volume numbers, making market participants more difficult to evaluate the true level of liquidity available on a certain exchange, or on certain option classes or series.

Our results have implications for the ongoing debate about whether and when irrational behavior by market participants is likely to influence prices. We find clear evidence that option owners make systematic mistakes that generate windfall benefits for option writers. One might surmise that such mistakes could lead to a lower equilibrium option price that would reflect these mistakes. However, our evidence indicates that arbitrageurs take aggressive action to capture these benefits, so that only a small fraction of the benefit accrues to parties who sold options prior to the last cum-dividend day. Thus, even though a significant proportion of option buyers make mistakes, the option price should reflect close to rational exercise. In addition, our results suggest that when transaction costs or other limits to arbitrage prevent traders from fully exploiting the mistakes of other investors, competitive pressure will push financial intermediaries to develop new ways to lower these costs.

Our results also have implications for the design and regulation of clearinghouses. An apparently innocuous allocation rule that appears designed to place all option writers on equal footing actually enables certain participants to skew the process in their favor, and creates a huge amount of extraneous trading activity in the process. This lesson may be relevant in other contexts where clearinghouses must allocate assignments.[3]

Additionally, our results demonstrate that when the marginal cost of trading is low for certain intermediaries, they will be willing to transact a huge volume of trades in order to capture a relatively modest benefit. The resulting volume spikes are merely an artifact of the settlement process—inasmuch as these spikes represent prearranged, offsetting trades between arbitrageurs, the higher volume is not indicative of the greater liquidity available to other market participants. The activity described in this article represents a significant source of short-term variation in trading volume that is unrelated to short-term variation in transaction costs, information, divergence of opinion, or other drivers of trading volume.

Our analysis proceeds as follows. Section 1 contains a detailed explanation of the dividend play trading strategy. Section 2 describes our hypotheses and our experimental design. Section 3 describes our data and presents descriptive statistics indicating that a large fraction of the open interest on the opening of trade on the last cum-dividend day remains open on the closing of trade. The evidence indicates that option owners fail to learn through time as the unexercised portion of the open interest remains stable during our sample period (1996–2006). Section 4 presents the results of our analysis. The strategy

[3] As an example, in the equity market, when sellers fail to deliver securities on time, the clearinghouse must assign the delivery failures across clearing members (see Boni 2006).

is widely employed as evidenced by the excessive trading volume that at the extreme is two thousand times the volume of trade in a randomly selected day. The two exchanges that imposed transaction fee caps experienced an increase in their fraction of the potential profits. Section 4 also contains evidence of a differential impact on the reported volume of trade for different exchanges at different months of the year 2003. As expected, the bigger effects are found for the exchanges that enacted a fee reduction rule. Section 5 concludes the article.

1. Analysis

1.1 Ex-dividend day behavior of call options

The last day that investors can buy a stock with the rights to receive an earlier declared dividend is referred to as the "last cum-dividend day." After this point, the stock is referred to as trading "ex dividend." At the opening of trade on the ex-dividend day, the ownership of stocks is separated from the ownership of the dividend. Consequently, no profit opportunities in the stock market imply an ex-dividend stock price drop in an amount equal to the dividend per share, adjusted for possible tax considerations, risk, and transaction costs.

In contrast to stocks, options are not entitled to cash dividends. If the options are American style, however, the owner can exercise and thus convert the option to stock in time to receive the dividend. In cases where it is not optimal to exercise the option, the assumption of no profit opportunities in the stock market implies that the option price should not respond in a predictable way to dividends. In other words, a predictable drop in the price of the option implies profit opportunities to the short.

Let

S_c = the stock price cum (with) dividend;

S_e = the stock price ex (without) dividends;

$E(S_e)$ = the expected stock price on ex-dividend day;

D^* = the dividend per share (tax-adjusted value perceived by marginal investor);

$C_c(X)$ = the price on the last cum-dividend day of a call option with strike price X;

$C_e(X)$ = the price on the ex-dividend day of a call option with strike price X; and

$C_e(X) - \max(0, S_e - X)$ = the time value of the option on the ex-dividend day.

For American-style options in states of the world where early exercise is not optimal, the following conditions must hold to rule out profit opportunities in the stock market and the option market:

$$S_c - E(S_e) = D^* \tag{1}$$

$$C_c = C_e \tag{2}$$

$$C_c \geq S_c - X \tag{3}$$

$$C_e \geq S_e - X. \tag{4}$$

The first condition simply states that the stock price is expected to fall by the amount of the dividend, possibly adjusted for tax effects. The second condition rules out the profits arising from predictable changes in the option price. The third and fourth conditions restate the well-known lower bound arbitrage properties of American options, which must hold both before and after the ex dividend. If conditions (3) and (4) are violated, arbitrageurs could profit by buying options and exercising immediately.

If the expected time value of the call option on the opening of trade on the ex-dividend day is smaller than the dividend (after adjusting for any tax consequences), the above four conditions cannot be simultaneously met. In this case, one would expect to observe a drop in the call price between cum to ex dividend, so that condition (2) would not hold. Yet, if all the call option holders exercise their options on the last cum-dividend day, there would be no profit opportunities. Indeed, if the time value of the call option on the opening of the ex-dividend day is lower than the dividend, it pays to exercise early. The cost of early exercise (the forgone time value) is lower than the benefit (receiving the dividend). As demonstrated by Roll (1977) and Kalay and Subrahmanyam (1984), early exercise of American-style calls will be optimal immediately prior to the underlying stock going ex dividend if the call is sufficiently far in the money. The following numerical example illustrates the point.

Example 1. *The expected time value of the call option on the opening of the ex-dividend day is smaller than the dividend.*

Assume that $D^* = \$1$, $S_c = \$10$, $X = \$7$, $\pi = \$0.25$. On the morning of ex dividend, the stock price is expected to open at \$9. Consequently, the call option is expected to be \$2 in the money and have a time value of \$0.25, so its total market value is expected to equal \$2.25. But on the cum-dividend day, this American call option can be purchased and exercised immediately, to generate a payoff of \$3, which establishes a lower bound on its value. The option price is expected to drop by \$0.75 from cum to ex. However, this does not represent an exploitable profit opportunity to the short seller. The entire open interest of the call option should disappear on the cum-dividend day. If all the call option holders exercise optimally, we should observe the open interest drop to zero on the last cum-dividend day. In other words, even though call option price decline is predictable, there is no profitable trading opportunity if the buyers all exercise rationally. Any attempt to exploit the price change would require writing the option and maintaining the written position until after the stock goes ex dividend, but this attempt would be frustrated when the option owners exercise.

The above argument assumes that all owners exercise optimally. If it is expected that some portion of the option holders will fail to exercise optimally, writing an option immediately prior to ex dividend can be a profitable strategy. If some holders fail to exercise, this of course implies that some writers will not be assigned. These unassigned writers will benefit when the option price drops. This does not mean, however, that competition among writers will push down the option price and eliminate the profit opportunity. Just prior to ex dividend, the option should be trading at intrinsic value (cum-dividend stock price minus the exercise price). If the option price falls below intrinsic value, a different arbitrage opportunity would arise, as condition (3) would be violated. As the empirical evidence presented below reveals, a significant fraction of the call option holders are not exercising their options when they should. The suboptimal exercise behavior of the option holders constitutes profits to the option seller.[4]

1.2 Ex-dividend arbitrage

Suboptimal exercise by option holders represents money "lying on the table," a benefit to those option writers who remain with intact written positions after all exercises have been assigned. The dividend play strategy flushes out preexisting written positions, leaving the arbitrageurs holding open written positions after the assignment. This section explains the mechanics of the strategy.

Consider an in-the-money call with strike price X, immediately prior to a dividend. As before, let D^* represent the amount of the dividend, adjusted for any possible tax effects that would induce the arbitrageur to value a dividend differently from a capital gain, and let $\pi_e = C_e - (S_e - X)$ represent the time value of the call option immediately after the stock goes ex dividend. Assume that $D^* > \pi_e$. In this case, if an option can be written on the last cum-dividend day at the intrinsic value $(S_c - X)$, the writer would unwind the position the next day, gaining the difference $D^* - \pi_e$. The extent to which the option writer will make profits depends upon the fraction of the open interest that remains unexercised.

Let H denote the open interest going into the last cum-dividend day. Suppose that some proportion of preexisting option holders, denoted by α, fail to exercise

[4] The algorithm used by the OCC to assign exercises has an element of randomness. Thus, it is impossible to predict in advance which option writers will be assigned. In this respect, writing an option immediately prior to ex dividend is like entering a lottery with some chance of a benefit.

optimally. The total money left lying on the table is given by the expression

$$\textit{Losses to Holders} = \textit{Gains to Sellers} = \alpha\ H(D^* - \pi_e). \tag{5}$$

Thus, the potential profits from writing options at the intrinsic values on the last cum-dividend day are larger the bigger α and D^* are, and the smaller π_e is.

Expression (5) describes the total value transferred from option buyers to option writers as a result of the buyer's suboptimal exercise strategy, but does not say anything about how the gains are distributed across the option writers. Arbitrageurs can devise a strategy ("dividend play") which enables them to extract almost all of the expected gains.

The strategy is executed as follows. On the last cum-dividend date, trader 1 buys calls on $Q/2$ shares from trader 2, and trader 2 buys calls on $Q/2$ shares from trader 1, so that together, the new trades establish new option positions on Q shares. The variable Q will be selected by the arbitrageurs to maximize the profitability of the strategy. As stated before, the two trades are executed at the same price, so no cash is required to settle these trades. The traders then exercise all Q shares of their long positions. At the end of the day, the clearinghouse allocates all exercises across written positions. This assignment process occurs prior to any netting of long and short positions within the same account, so the clearinghouse will allocate $(1-\alpha)\,H + Q$ exercises across $H + Q$ written positions. These assignments will be allocated across the preexisting writers, who hold a proportion $H/(H+Q)$ of all written positions, and the two new traders, who hold a proportion $Q/(H+Q)$ of the written positions.

After the assignment, the two traders would, on average, remain with $Q/(H+Q)$ of the unassigned written positions, and expect to capture that proportion of the total gains. Because the initial trades are exactly offsetting, no initial cash flow is required. The profits of the strategy are given by the total gain captured minus transaction and clearing fees. Denoting the per-unit transaction and clearing fees by C, the total expected profit from executing the strategy is thus

$$\Pi = \left(\frac{Q}{H+Q}\right)\alpha H(D^* - \pi_e) - CQ. \tag{6}$$

One unresolved question is how the traders would unwind their position after executing the strategy. When the assigned exercises settle, the arbitrageurs will deliver cash, receive a long position in the stock, and remain with a written position in the option. Any transaction cost associated with liquidating this position would eat into the arbitrage profits. If the options are short term and deep in the money, perhaps an easier way for the traders to exit their position would be to simply wait until expiration, at which point the options will be exercised and the traders will deliver their stock and get back the cash. This requires some short-term capital outlay, and may involve some risk that the stock price will fall below the strike price before expiration.

In Appendix B, we formulate expressions for the arbitrageur's profit and solve for the profit-maximizing scale of the dividend play arbitrage, under the assumption of a monopolist arbitrageur, and then in a purely competitive environment. In both cases, we provide expressions for the equilibrium amount of dividend arbitrage, as a function of the preexisting open interest, the proportion of option holders who fail to exercise optimally, the dividend size, the ex-dividend time value, and the marginal clearing cost. When calibrated with realistic parameter values, we find that at the optimal scale, we would expect to see dividend play volume multiple times higher than the cum-dividend open interest, even in the monopolist case. Under the competitive equilibrium, dividend play volume would be even higher.

2. Hypotheses and Experimental Design

2.1 Conditions for the dividend play

If market participants are engaged in dividend arbitrage, we would expect to see a flurry of trading activity on the last cum-dividend day, and would expect the activity to be confined to those series for which the strategy is potentially profitable. Within those series, we would expect arbitrageurs to be more aggressive at seeking to capture these profits when the potential profits are greater.

The dividend play strategy will be profitable on the series for which (1) the option's expected time value is lower than the amount of the dividend, and (2) the open interest going into the ex-dividend day is positive. The first condition implies that the option should be exercised early. The second implies that some holders errantly choose not to exercise and leave the money on the table. These are the two necessary conditions for the dividend play to be a possibility. Evaluating condition (1) requires an option-pricing model to estimate time value, and any model may be subject to error. To help ensure that our classification identifies only observations where early exercise is optimal, we add another requirement: (3) the open interest on the close of the last cum-dividend day is smaller than that on the previous day.

We wish to test the hypothesis that trading volume prior to a dividend date is abnormally high when these three conditions hold, and that volume is higher when the potential profits are higher. We test this by identifying an "arbitrage sample" consisting of those observations when these three conditions hold. The relationship between profit possibilities and excess trading activity may be modeled as follows:

$$Volume_i = \alpha + \beta_1 PROFIT_i + \beta_2 ARB_i + \beta_3 PROFIT\_ARB_i + MVOL_i + \varepsilon_i, \qquad (7)$$

where *Volume* represents trading activity on the last cum-dividend day, *PROFIT* represents the per-unit profitability of dividend arbitrage, *ARB* is an indicator variable corresponding to inclusion in the arbitrage sample, and *MVOL* is the average trading volume over the prior thirty trading days. The variable *PROFIT* is defined as the difference between the dividend and the expected time value of the option on the opening of the ex-dividend day. The third independent variable estimates the sensitivity of the excess trading activity to the potential profits within the arbitrage activity sample.

The option's expected time value (π_e) is estimated by

$$\pi_e = \mathrm{E}[C] - \mathrm{E}[S] + X, \tag{8}$$

where E[C] is the expected ex-dividend price of the call, E[S] is the expected ex-dividend price of the stock, and X is the strike price. We assume that the expected ex-dividend stock price is simply the closing price on the last cum-dividend day minus the cash dividend per share. The expected ex-dividend option price is estimated using the Black–Scholes–Merton model,[5] evaluated at the expected ex-dividend stock price. For the volatility parameter, we use the implied volatility provided in OptionMetrics for the specific option. Out of our sample of 438,269 observations, we find implied volatilities for 402,341 (91.80%) while they are missing for 35,928 (8.2%) observations. For these cases, we estimate the implied volatility by the linear combination of the implied volatilities of the two closest options with higher and lower exercise prices and the same time to expiration. To control for the potential effects of heteroscedasticity, we report test statistics based on White (1980) adjusted standard errors.

The amount of money left on the table (or the potential profit that can be extracted through dividend plays) will be equal to the value of early exercise times the open interest going into ex dividend. We would expect the abnormal trading activity on the last cum-dividend day to increase with respect to this measure.

2.2 Dividend plays and the fee structure

As the dividend play strategy can generate extremely high trading volume, one might expect that the strategy would incur great transaction costs. As a practical matter, the strategy would not be feasible for investors who have to pay commissions, or who could not arrange to cross trade at a single price inside the spread.

Market makers are likely to have the lowest marginal cost of implementing a dividend play strategy. In general, option market makers face two components of trading costs: a transaction fee, paid to the exchange, and a clearing fee, paid to the clearinghouse. However, several exchanges have implemented rules that place caps on the total monthly transaction fees paid by a given market maker. In some cases, these caps apply specifically to dividend capture strategies. Once the market maker reaches the cap, their incremental transaction fee is reduced to

[5] Black and Scholes (1973). For options that pay more dividends prior to expiration, we adjust the calculation by converting discrete dividends into equivalent continuous yields and using Merton's (1973) formula.

zero. Fee capping rules were introduced on the Pacific Exchange (PCX) on 19 August 2003, and on the Philadelphia Stock Exchange (PHLX) on 8 September 2003. The American Stock Exchange (AMEX) implemented fee caps in early 2004, but the fee change was applied retroactively back through July 2003. The Chicago Board Option Exchange implemented fee caps for dividend strategies in 2004.[6]

As stated by the Pacific Exchange (PCX) in the rule filing:

> From time to time, market participants engage in financing strategies known as option strategy plays for the purpose of reducing risk. These transactions include reversals and conversions, dividend spreads, and box spreads. Because the referenced options strategy transactions are generally executed by professionals whose profit margins are generally narrow, the PCX proposes to cap the transaction fees associated with such executions at $2000. The PCX believes that, by keeping fees low, it will be able to attract liquidity by accommodating these transactions.

Once the transaction fee is capped, the marginal cost of executing the dividend capture would consist only of the clearing fee. In the time period of our analysis, the clearing fee was $0.09 per contract ($0.0009 per underlying share) for trades up to five hundred contracts, and lower than this for larger trades.[7] We hypothesize that the capping rule facilitates dividend capture behavior by market makers. Abnormal trading activity should be more sensitive to the presence of dividend play opportunities on exchanges that implement fee caps and after the fee cap rule was enacted.

Our model predicts that abnormal trading activity on the exchanges that implement fee caps is more sensitive to the presence of dividend play opportunities (as measured by money left on the table). To test this, we estimate the following regression equation:

$$Volume_{\text{ij}} = \sum_{j=1}^{5} \alpha_j Exchange_i^j + \sum_{j=1}^{5} \beta_j Exchange\_PROFIT_i^j + \sum_{j=1}^{5} \beta_j Exchange\_MVOL_i^j + \varepsilon_i, \qquad (9)$$

where volume represents trading activity on the last cum-dividend day on a particular exchange; $Exchange^j$, $j = 1, \ldots, 5$, represents a set of indicator

6 Securities and Exchange Commission, Release No. 34-48363; File No. SR-PCX-2003-39 (19 August 2003); Release No. 34-48459; File No. SR-Phlx-2003-61 (8 September 2003); Release No. 34-49358; File No. SR-Amex-2004-09 (10 March 2004); Release No. 34-50175; File No. SR-CBOE-2004-38 (17 August 2004).

7 The actual clearing fee was lower than this because at the end of each year, the OCC refunds all revenues in excess of operating expenses, in proportion to fees paid. According to the OCC annual report, they refunded approximately 30% of clearing fees in 2003.

variables for the five exchanges,[8] and each variable takes the value of 1 if option i is trading in exchange j, and 0 otherwise; $Exchange\_PROFIT^j$ is the interaction between the exchange indicator and the expected profit on the dividend plays; and *Exchange_MVOL* is the interaction between the exchange indicator and the average trading volume over the prior thirty trading days. Again, hypothesis tests are based on White's heteroscedasticity-adjusted standard errors.

As a second approach to testing whether the fee cap rule facilitated dividend arbitrage, we regress the volume of trading on an indicator variable corresponding to the time period after the rule went into effect. Using data from the arbitrage sample only, we estimate the following model to test whether trading activity increased after the rule passed:

$$\begin{aligned} Volume_i = {} & \alpha + \beta_1 PROFIT_i + \beta_2 CAP_i + \beta_3 PROFIT\_CAP_i \\ & + MVOL_i + \varepsilon_i, \end{aligned} \tag{10}$$

where *PROFIT* is the expected profit (the product of open interest and potential profits per contract) on the ex-dividend day, *CAP* is an indicator variable taking the value of 1 in the period following the fee-capping rule and 0 otherwise, *PROFIT_CAP* is the interaction between the expected profit and the *CAP* indicator, and *MVOL* is the average trading volume over the prior thirty trading days.

Relevant data separated into the five different exchanges are available to us only for the year 2003. Hence both specifications (Equations (9) and (10)) were tested using data from 2003 only. We estimate Equation (10) separately for each exchange. For the two exchanges that implemented fee caps, the indicator variable is defined by the effective date of the rule. For the exchanges that did not have a rule change during our sample period, we set the indicator variables equal to 1 after 1 August 2003. The regressions on these exchanges act as a control that should capture the effects of any other factors that may have influenced dividend play activity. In addition, we can look into the coefficients in these other three exchanges to test whether dividend play activity migrated away from those exchanges when the fees were reduced on competing exchanges.

2.3 Impact on reported volume

One of the most important features of well-functioning capital markets is liquidity, which is related to the ease with which an asset can be traded without significantly affecting the price. Similar to other markets (such as the stock market), option markets employ market makers who stand ready to provide liquidity and implement some rules in order to attract liquidity. A dividend play involves coordination between two parties, most likely between two market makers. Because the trades are exactly offsetting, dividend play trades create trading volume without adding any liquidity to the market.

[8] The exchanges are AMEX: American Stock Exchange, CBOE: Chicago Board Options Exchange, ISE: International Securities Exchange, PCX: Pacific Stock Exchange, PHLX: Philadelphia Stock Exchange.

As the marginal transaction cost approaches zero, the optimal size of the arbitrage trade grows very large. According to the 2008 annual report of the OCC, clearing fees are now down to less than 1½ cents per contract. If the marginal transaction cost falls to one penny, we anticipate that for particularly large dividends, the profit maximizing dividend play might increase to fifty times the open interest, or more. Such uncontrolled growth of reported trade volume might make it more difficult for investors, academics, and regulators to infer information about true liquidity from volume of trade. Extreme trading volume associated with dividend plays can significantly affect statistics, such as average daily trading volume, that are often used as proxies for the liquidity of the option markets as a whole, the liquidity of individual options, or relative liquidity across options. For example, dividend play activity will make it appear that options on high-dividend stocks are more liquid and have significantly lower put/call volume ratios than low-dividend stocks, and that short-term in-the-money calls are more liquid than other kinds of options. Also, it will boost the market share of those exchanges where dividend plays are most prevalent, perhaps giving investors the impression that more liquidity is available on those exchanges than is really the case. For these reasons, it is important to see the effect of dividend plays on trading volume, and the extent to which volume figures have been affected on each exchange.

We would like to test whether trading volume on exchanges after the transaction fee capping rule was enacted was sufficiently large to have a substantial effect on the reported monthly volume in those exchanges. We obtain equity option contract volume by exchange and by month from CBOE 2003 market statistics.[9] We compare the last cum day trading volume on the call options in the dividend play to the monthly call trading volume[10] of all equity options listed on the corresponding exchange.

3. Data and Summary Statistics

We first identify a set of exchange-traded options with underlying stocks paying cash dividends. Our sample consists of dividend-paying stocks, identified from the Center for Research in Security Prices (CRSP) over the period January 1996 to December 2006, which also have option trades reported by OptionMetrics. For the year 2003, we obtain exchange specific information provided by the Options Price Reporting Authority (OPRA).

We then construct a sample of 438,269 observations, each corresponding to a particular call option series and a particular dividend event extracted from OptionMetrics. Inclusion in the sample is based on the following selection

[9] http://www.cboe.com/data/marketstats-2003.pdf.

[10] As we consider only call options in our analysis, based on the detail trading volume of calls and puts in CBOE in each month, the trading volume on calls is about twice that of the respective puts. We assume that all the other exchanges have the same relative open interest of calls and puts as does the CBOE, i.e., the monthly call trading volume of all equity options are two-thirds of the total equity option contract volume.

criteria: (a) the underlying stock paid at least one cash dividend during the time period from January 1996 to December 2006, and (b) the option series has reported volume and open interest information on the date before the last cum-dividend day, the last cum-dividend day, and the ex-dividend day.

To be included in the arbitrage activity sample, the observation has to satisfy the following additional selection criteria: (c) the cash dividend is larger than the option's expected time value, estimated by the Black–Scholes–Merton model, (d) the open interest going into the ex-dividend day is positive, which implies that some holders errantly choose not to exercise and leave the money on the table, and (e) there is a reduction of open interest from the last cum day going into the ex-dividend day. The last criterion serves as an additional control. Market makers may use a better option valuation model than the one we employ, and hence may classify observations differently. To align our classification with theirs, we examine only observations for which we have no difference of opinion—the open interest indeed drops. We essentially assume that the market makers as a whole have a better estimate of the value of early exercise than our valuation. Thus, if our estimation indicates value for early exercise and it is not met by a reduction in the open interest, we eliminate the observation from the arbitrage sample.

We construct a third sample consisting of observations that should have been exercised by our estimates but do not meet condition (e)—the open interest did not necessary go down. We refer to this sample as the optimal early exercise sample. This is a larger sample (that includes the arbitrage activity sample) consisting of 46,646 observations. This sample controls for a potential built-in relationship between the arbitrage activity sample and reduction in the open interest around the ex-dividend day.

Table 1 contains descriptive statistics for the total sample and for the arbitrage activity sample. In particular, it provides information about the distribution within the sample of time to expiration, dividend size, option moneyness, and the value of early exercise. Table 1 also reports average trading volume on the last cum-dividend date (*Cum_vol*) and the average trading volume over the prior thirty days (*MVOL*) for the entire sample and various partitions. Panel A partitions the sample by time to expiration relative to the ex-dividend day. Time to expiration is classified as ≤ 1 when the option expires during the ex-dividend month, greater than 1 but ≤ 2 when the option expires in the next month, and so on. Cumulatively, 87.47% of the options in the arbitrage activity sample expire within three months of the ex-dividend day.

Panel B of Table 1 partitions the sample by the size of the cash dividend per share. Cumulatively, 70% of the options in the arbitrage activity sample pay cash dividends higher than 18 cents. Panel C partitions the sample based on the option's depth in the money on the last cum-dividend day. We use two definitions of depth in the money. Depth in the money reported in Table 1 is the simple difference between the stock price and the exercise price. Later on we use an alternative definition of depth in the money—the difference between the underlying stock's market price and the option's exercise (strike) price on

Table 1
Descriptive statistics

Panel A: Partitioned by time to expiration

Months to expiration		≤1	≤2	≤3	≤4	≤5	≤6	≤7	≤8	≤9	Total
Whole sample	No. of options	81,764	89,333	52,429	49,453	52,328	44,512	34,525	26,377	7548	438,269
	Percent (%)	18.66	20.38	11.96	11.28	11.94	10.16	7.88	6.02	1.72	100
	Cum_Vol	1089	308	129	47	49	40	30	46	117	
	MVOL	99	55	46	26	24	30	17	23	41	
Arbitrage activity sample	No. of options	14,734	5678	2432	1431	912	507	261	91	71	26,117
	Percent (%)	56.42	21.74	9.31	5.48	3.49	1.94	1.00	0.35	0.27	100
	Cum_Vol	4866	3447	1625	533	650	712	678	1264	2363	
	MVOL	62	23	15	22	14	12	15	12	4	

Panel B: Partitioned by dividend size

Deciles	1	2	3	4	5	6	7	8	9	10
Whole sample (438,269 observations)										
Cash dividend	0.02	0.05	0.07	0.10	0.14	0.18	0.23	0.29	0.38	0.71
Cum_Vol	75	46	40	54	63	128	290	360	671	1322
MVOL	64	46	35	44	40	57	46	40	52	47
Arbitrage sample (26,117 observations)										
Cash dividend	0.04	0.09	0.13	0.18	0.22	0.27	0.33	0.39	0.51	1.11
Cum_Vol	43	109	351	1101	1863	4401	3353	6654	5564	1,3768
MVOL	49	40	34	38	34	41	38	50	43	63

Panel C: Partitioned by depth in the money

Deciles	1	2	3	4	5	6	7	8	9	10
Whole sample (438,269 observations)										
Cash dividend	−18.47	−7.33	−4.13	−2.04	−0.37	1.27	3.20	5.74	9.97	25.85
Cum_Vol	24	44	67	101	129	255	609	561	649	601
MVOL	30	46	59	77	85	67	54	28	18	10
Arbitrage sample (26,117 observations)										
Cash dividend	1.77	3.41	4.73	6.07	7.64	9.48	11.76	15.02	20.44	41.88
Cum_Vol	4075	5189	3614	4148	3671	2886	4540	3500	3177	2491
MVOL	142	82	52	41	30	24	21	17	14	10

(*continued overleaf*)

Table 1
(Continued)

Panel D: Partitioned by value of early exercise

Deciles	1	2	3	4	5	6	7	8	9	10
Whole sample (438,269 observations)										
Value of early ex	−5.04	−2.28	−1.47	−0.97	−0.63	−0.40	−0.17	−0.03	0.08	0.35
Cum_Vol	47	59	66	80	94	116	181	151	250	1996
MVOL	39	45	51	54	56	57	59	42	36	35
Arbitrage sample (26,117 observations)										
Diff	0.01	0.04	0.07	0.10	0.14	0.17	0.23	0.30	0.40	0.84
Cum_Vol	848	435	629	1490	1873	3168	4943	5252	6697	11,967
MVOL	65	41	50	38	40	42	35	36	36	50

This table reports trading volume on the last cum-dividend date (*Cum_Vol*) and mean volume over the prior thirty days (*MVOL*), for various samples and subsamples. The whole sample includes 438,269 cash dividend events during the period 1996–2006 for stocks with listed options, where call option volume was positive on the ex-dividend day and on the previous two days. The "arbitrage activity" sample includes 26,117 cases for which the dividend paid on the underlying stock is larger than our estimate of the expected time value of the option, the option is in the money on the cum-dividend day, and the open interest in the beginning of the last cum-dividend day is positive and higher than the open interest at the end of the last cum-dividend day. The sample is further partitioned by time to expiration (panel A), deciles of dividend size (panel B), deciles of depth in the money (panel C), and deciles of the value of early exercise, or arbitrage profit per share (panel D).

the last cum-dividend day standardized by the stock price times the implied volatility of the option. Cumulatively, 70% of the options in the arbitrage activity sample have depth in the money higher than $6.07. Panel D partitions the sample by the value of exercise immediately prior to the ex-dividend date, which equals the potential arbitrage profits at the end of the last cum-dividend day. For the whole sample, the potential profits are negative in all but the two highest deciles. For the arbitrage activity sample, the mean potential profits are positive for each one of the deciles, as expected given the definition of the arbitrage sample.

Table 2 provides summary statistics on exercise patterns surrounding the ex-dividend date. Panel A reports the open interest going into the last cum-dividend date and going into the ex-dividend date, separately for the control and the arbitrage samples, and broken down by the open interest decile. As expected, the table documents significant reductions in the open interest in the arbitrage activity sample, and very small and insignificant changes in the open interest during the same time period for the control sample. More importantly, the table documents that a large fraction of the open interest is not exercised prior to the ex-dividend day. Across deciles, the fraction of unexercised options varies from a low of 40.71% to a high of 54.33%. We find no systematic relationship between the amount of open interest going into the last cum-dividend day and the fraction that remains unexercised.

Panel B of Table 2 reports the same numbers broken down by deciles of per-unit arbitrage profit, for the arbitrage sample and for the optimal early exercise sample (consisting of 46,646 observations). The percentage of contracts that remain unexercised in the arbitrage sample clearly decreases as the benefits of early exercise increase. In the lowest decile, we estimate the average benefit of early exercise to be about one penny per share, and here we find that 75% of the contracts are not exercised. Much of this can be perfectly rational, given that there are likely to be some transaction costs associated with exercising the option. Our evidence strongly suggests, however, that failure to exercise is not only due to transaction costs. In the top decile, where the average benefit of early exercise is $0.84/share, we still find that more than 32% of contracts are not exercised. The reductions in the open interest documented for the arbitrage sample are not spurious. Similar numbers are reported for the optimal early exercise sample. For the lowest decile where the profit per share is one cent, 86.88% of the open interest remains unexercised. The number drops to 39.34% for the highest decile where the profits per share are 72 cents. Our conclusion that the failure to exercise cannot be explained by transaction costs is further verified by Pool, Stoll, and Whaley (2008), who in parallel research examine option exercises on ex-dividend days using a methodology that explicitly incorporates a range of estimates of transaction costs.

Panel C of Table 2 presents the proportion of unexercised options sorted into deciles by depth in the money for the two samples (arbitrage and optimal early exercise). As the depth in the money of the options grows, the fraction of the

Table 2
Descriptive statistics on ex-dividend option exercises

Panel A: Exercise patterns by open interest decile

Open interest decile	1	2	3	4	5	6	7	8	9	10
Arbitrage sample (26,117 observations)										
Cum-dividend open interest	15	40	73	120	193	307	508	895	1810	8377
Ex-dividend open interest	7	18	31	49	79	125	211	408	897	4551
Percent unexercised	46.67	45.00	42.47	40.83	40.93	40.71	41.54	45.59	49.56	54.33
No arbitrage sample (412,152 observations)										
Cum-dividend open interest	4	14	30	57	102	180	327	646	1517	8817
Ex-dividend open interest	7	17	34	61	107	186	334	659	1537	8862

Panel B: Exercise patterns by profit decile

Profit decile	1	2	3	4	5	6	7	8	9	10
Average profit per share	0.01	0.04	0.07	0.10	0.14	0.17	0.23	0.30	0.40	0.84
Cum-dividend open interest	1549	1195	1287	1202	1246	1333	1253	1256	1260	1428
Ex-dividend open interest	1162	766	737	639	672	634	571	585	527	471
Percent unexercised	75.02	64.10	57.26	53.16	53.93	47.56	45.57	46.58	41.83	32.98
Optimal early exercise sample (46,646 observations)										
Average profit per share	0.01	0.03	0.05	0.08	0.11	0.14	0.19	0.25	0.35	0.72
Cum-dividend open interest	831	839	763	828	886	924	916	990	1100	1037
Ex-dividend open interest	722	669	554	547	598	594	542	567	629	408
Percent unexercised	86.88	79.74	72.61	66.06	67.49	64.29	59.17	57.27	57.18	39.34

Panel C: Exercise patterns by moneyness decile

Arbitrage sample (26,117 observations)										
Moneyness per share	1.32	2.58	3.46	4.41	5.61	7.45	10.17	14.90	24.66	70.54
Cum-dividend open interest	1874	1524	1518	1264	1157	1129	1034	969	914	957
Ex-dividend open interest	1522	879	730	614	529	545	501	410	350	296
Percent unexercised	81.22	57.68	48.09	48.58	45.72	48.27	48.45	42.31	38.29	30.93
Optimal early exercise sample (46,646 observations)										
Moneyness per share	0.62	1.87	2.66	3.48	4.46	5.78	7.87	11.39	18.82	57.99
Cum-dividend open interest	1453	1037	1010	991	869	799	773	746	681	754
Ex-dividend open interest	1407	872	682	568	488	426	426	398	304	267
Percent unexercised	96.83	84.09	67.52	57.32	56.16	53.32	55.11	53.35	44.64	35.41

This table contains a description of the open interest on the last cum-dividend day and the first ex-dividend day, for various groups of options. There are 438,269 observations of option series on ex-dividend days, of which 46,646 observations constitute the "Optimal Early Exercise Sample," 26,117 observations constitute the "arbitrage" sample, and 412,152 constitute the "no arbitrage" sample. The optimal early exercise sample includes those series for which the dividend paid on the underlying stock is larger than the estimate of the expected time value of the option and the option is in the money. The arbitrage sample includes those series for which early exercise is optimal but some holders fail to exercise. In panel A, the samples are sorted separately and divided into deciles based on the amount of open interest on the beginning of the last cum-dividend day. In panel B, the samples are sorted and divided into deciles based on the per share benefit of early exercise, measured as the dividend per share minus the time value of the option. In panel C, the samples are sorted and divided into deciles based on the moneyness, measured as $\frac{(S_t-K)}{S_t\sigma_{\text{imp}}(T-t)}$, where S_t is the stock price on ex, K is the exercise price, σ_{imp} is the implied volatility, and $T-t$ is the time to expiration. The sample period is from 1996 to 2006.

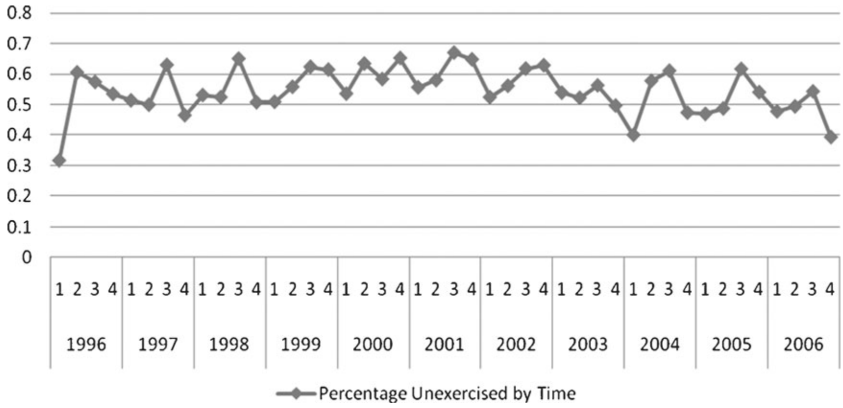

Figure 2
Fraction of unexercised options in the arbitrage activity sample
The fraction of unexercised options on the ex-dividend day is measured as the ratio of the open interest at the end of the last cum-dividend day (the opening of the ex-dividend day) and the open interest of the option at the beginning of the last cum-dividend day. Options are divided into groups based on the calendar quarters. There are 26,117 observations constituting the "arbitrage" sample from 1996 to 2006. The arbitrage sample includes those series for which early exercise is optimal but some holders fail to exercise.

unexercised options decreases for the arbitrage sample from a high of 81.22% to a low of 30.93%. Similar numbers are reported for the optimal early exercise sample. The fractions drop from a high of 96.83% to a low of 35.41% as the depth in the money increases.

Figure 2 presents the time-series mean fraction of the option in the arbitrage sample that remains unexercised during our sample period. We estimate the mean unexercised portion for each quarter during the period 1996–2006. The numbers presented in Figure 2 depict a fairly stable fraction of unexercised options. The empirical evidence seems to indicate that option holders fail to learn over the years as the fraction of unexercised options of those that should have been exercised does not decline.

4. Empirical Evidence

4.1 Potential profits and excess trading

Table 3 contains the results of regression (7). For the control sample, we find no association between the trading activity and the potential profits that can be made in a dividend play. The coefficient on *ARB* is significantly positive, indicating that there is substantially higher trading volume on the last cum-dividend day for the arbitrage activity sample. Further, the interaction coefficient of *ARB* and *PROFIT* is positive and significant. As hypothesized, the volume of trade of options providing profit opportunities on the last cum-dividend day increases with the potential profit. In other words, the volume of trade is positively correlated with the profit potential around the ex-dividend day. We document an

Table 3
Dividend plays in the arbitrage activity sample

Independent variable	Coefficient	t-stat
INTERCEPT	58.2857	0.84
PROFIT	−0.01438	−0.49
ARB	1446.1671	2.61
PROFIT ARB*	19.2936	3.71
MVOL	0.2514	0.11
Adj R^2	0.1748	

This table reports regression results testing whether trading activity on the last cum-dividend date is related to the expected profits from a dividend play strategy. The regression takes the form

$$Volume_{\mathrm{i}} = \alpha + \beta_1 PROFIT_i + \beta_2 ARB_i + \beta_3 PROFIT\_ARB_i + MVOL_i + \varepsilon_i,$$

where $Volume_i$ is a particular option's trading volume on the last cum-dividend day; *PROFIT* is the potential profit on the dividend play (the product of the open interest on the last cum-dividend day and the difference between the cash dividend paid on the underlying stock and the expected time value of the option); *ARB* is an indicator variable that takes the value of 1 if the option belongs to our defined arbitrage sample and 0 otherwise; *PROFIT_ARB* is the interaction between the potential profit and the arbitrage dummy; and *MVOL* is the average trading volume over the previous thirty business days. Regression coefficients and t-statistics computed using White (1980) standard errors are reported. The number of observations is 438,269, and the sample period is from 1996 to 2006.

insignificant coefficient on *MVOL* (the mean trading volume during the last thirty days). The adjusted R-square is 0.1748.

The evidence presented in Table 3 indicates that within the arbitrage sample, the excess trading activity increases as the potential profits are larger. To further explore this relation, we sort the data based on the potential profit per option and divide it into deciles. For each decile, we compute the excess trading activity in numbers of contracts and in ratios of expected to average volume. As is shown in Figure 3, the estimated profit positively correlates with both measures of abnormal trading activity. In the highest profit decile, when the mean estimated profit is $10,000, the mean excess trading volume is about 120,000 contracts and the mean ratio is 1400 (which is 1400 times the option's expected trading volume). Figure 3 depicts a nonlinear relationship between the excess trading volume in contracts and in ratio with the potential profits. Trading activity increases at an increasing rate.

4.2 Potential profits and choice of venue

The potential profits resulting from the suboptimal exercise decisions of the call option holders create an interesting connection across the various option exchanges. As we have seen, different exchanges have different marginal transaction costs for executing dividend plays—the strategy is likely to be more prevalent on some exchanges than others. However, an option writer cannot reduce the probability of being subjected to dividend plays by selecting an exchange where dividend plays are less common. The options traded on all of the exchanges are cleared against each other at a single clearinghouse, the OCC. Thus, investors who write options on one exchange are equally subject to the effects of dividend plays executed on other exchanges. To that extent, an

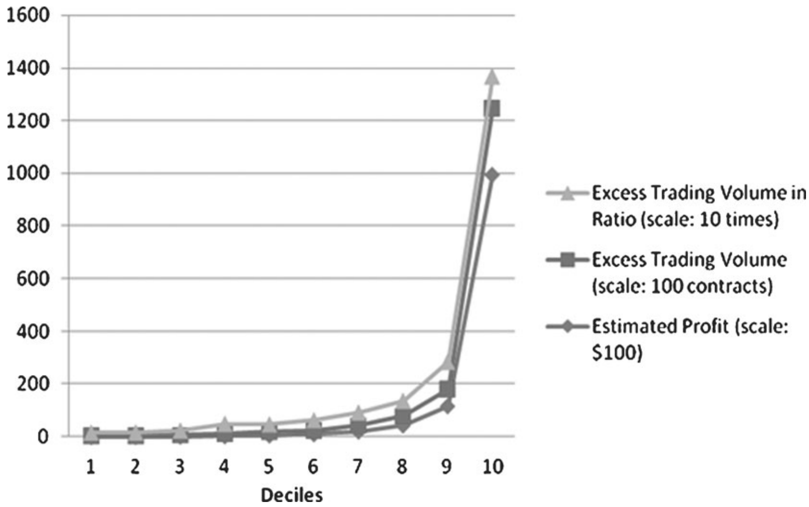

Figure 3
Potential profit and excess trading volume in the arbitrage activity sample
The potential profit on the dividend play is measured as the product of the open interest at the end of the last cum-dividend day and the difference between the cash dividend paid on the underlying stock and the expected time value of the option on the opening of the ex-dividend day. Options are divided into ten groups based on the value of the potential profit, and excess trading volumes are calculated within each group. Excess trading volume is calculated as the difference between the actual trading volume on the last cum-dividend day and the average trading volume during the previous thirty business days. The excess trading volume in ratio is the ratio of the actual trading volume on the last cum-dividend day to the average trading volume during the previous thirty business days. If the average trading volume is larger than the trading volume on the last cum-dividend day, we define the ratio as 0. The minimum value of the average historical trading volume is set equal to 1. There are 26,117 observations that constitute the "arbitrage" sample from 1996 to 2006. The arbitrage sample includes those series for which early exercise is optimal but some holders fail to exercise.

exchange cannot choose to stay out of the game. Investors trading in it are exposed to the consequences of dividend plays in other exchanges. It is therefore interesting to explore how the traders on different exchanges behave.

The evidence presented in Table 4 reveals significant differences across exchanges. Because different exchanges have somewhat different sets of listed options, the number of events in our arbitrage sample differs across exchanges. Panel A of Table 4 details the number of observations in the arbitrage activity sample in each of the five exchanges. The number varies from a low of 3192 in ISE to a high of 4299 in AMEX.

Panel B of Table 4 reports the estimation results of the regression in Equation (9). The coefficients on the exchange indicators are significantly positive for all the exchanges but ISE. Similarly, for these four exchanges, the coefficient on the interaction of the estimated profit and the exchange indicator is significantly positive. In other words, in these exchanges the market makers are responsive to the potential profit opportunities presented by the suboptimal exercise decisions of the option holders.

Table 4
Dividend play activity on different exchanges

Panel A: Number of observations in the various exchanges

Exchange	No. of observations
AMEX	4299
CBOE	3689
ISE	3192
PCX	3663
PHLX	3960

Panel B: Regression analysis—dependent variable:arbitrage activity

Independent variable	Coefficient	*t*-statistics
AMEX	65.4849	3.27
CBOE	156.6259	10.78
ISE	−2.3184	−0.76
PCX	317.4987	7.91
PHLX	612.5151	6.21
AMEX * *PROFIT*	1.6669	3.97
CBOE * *PROFIT*	1.8384	6.80
ISE * *PROFIT*	0.0947	1.84
PCX * *PROFIT*	5.4627	5.96
PHLX * *PROFIT*	15.6582	6.87
AMEX * *MVOL*	3.0303	0.97
CBOE * *MVOL*	0.5127	0.51
ISE * *MVOL*	1.1175	3.04
PCX * *MVOL*	−8.6256	−0.91
PHLX * *MVOL*	−9.8165	−0.66
Adj R^2	0.4103	

This table reports regression results testing whether cum-dividend trading volume on each of five exchanges is related to the expected profits from a dividend play trading strategy. The regression takes the form

$$Volume_{ij} = \sum_{j=1}^{5} \alpha_j Exchange_i^j + \sum_{j=1}^{5} \beta_j Exchange\_PROFIT_i^j + \sum_{j=1}^{5} \beta_j Exchange\_MVOL_i^j + \varepsilon_i,$$

where volume is the trading volume on the last cum-dividend day for a particular option on a particular exchange; $Exchange^j$, $j = 1,\ldots,5$, represents a set of dummy variables for five exchanges, and each variable takes the value of 1 if option i is trading in exchange j, and 0 otherwise. $Exchange\_PROFIT^j$ is the interaction between exchange dummy and the expected profit on the dividend plays. $Exchange\_MVOL^j$ is the interaction between exchange dummy and the average trading volume during the previous thirty business days. The sample period is January to December 2003, and the number of observations is 18,803. The number of observations in each exchange is shown in panel A. Regression coefficients and *t*-statistics based on White (1980) standard errors are shown in panel B.

Panel A of Table 5 describes the behavior of the excess trading volume in contracts and in ratios of the volume on the last cum-dividend day to the expected volume before and after the implementation of the change in the fees on PCX and PHLX, and before and after 1 August 2003 for the other exchanges. Only for PCX and PHLX do we observe a significant increase in the excess trading volume measured in contracts and for PCX, PHLX, and AMEX when measured in ratios. The evidence indicates that the volume of trade on the last cum-dividend day on exchanges that enacted a fee cap increases following the rule.

Panel B of Table 5 reports the result of regression (10), estimated separately for the arbitrage sample in each of the five exchanges. The coefficient on

Table 5
Excess trading volume across exchanges before and after the change of fee structure

Panel A: Mean excess trading volume

Exchange	Date	In ratio Before	In ratio After	In ratio *t*-stat	In contracts Before	In contracts After	In contracts *t*-stat
AMEX	1 August 2003	46	107	4.86	156	194	1.18
CBOE*	1 August 2003	81	104	1.76	240	275	0.92
ISE*	1 August 2003	0.8	1	1.76	14	1	−1.46
PCX	19 August 2003	116	512	8.70	295	938	6.53
PHLX	8 September 2003	418	1158	5.51	953	2238	5.31

Panel B: Dependent variable:trading volume on the last cum-dividend day

Exchange	AMEX		CBOE*		ISE*		PCX		PHLX	
No. of observations	4299		3689		3192		3663		3960	
	Coef.	*t*-stat	Coef.	*t*-stat	Coef.	*t*-stat	Coef.	*t*-stat	Coef.	*t*-stat
INTERCEPT	1.61	0.04	133.88	5.54	−0.51	−0.07	95.38	1.84	510.32	3.43
PROFIT	2.60	3.02	2.11	3.40	0.20	1.38	3.67	2.47	8.79	2.92
CAP	98.35	2.33	31.62	1.04	−3.41	−0.46	433.71	5.23	341.16	1.73
PROFIT_CAP	−1.22	−1.21	−0.34	−0.50	−0.14	−0.90	2.65	1.40	11.09	2.60
MVOL	2.38	0.87	0.49	0.48	1.06	2.97	−2.37	−0.23	−6.43	−0.31
Adj R^2	0.3341		0.2522		0.2195		0.2959		0.4631	

This table reports results testing whether trading activity related to dividend arbitrage increased after rule changes putting caps on transaction fees on three exchanges (PCX, PHLX, and AMEX) in 2003. For benchmark purposes, results before and after 1 August 1 2003 are reported for the exchanges that did not experience fee changes during the sample period. The experiment is limited to observations from the arbitrage activity sample, and to individual exchanges. Panel A reports average excess trading volume before and after the change in fee structure in each exchange. Panel B reports results from a regression, estimated separately for each exchange, taking the form

$$Volume_i = \alpha + \beta_1 PROFIT_i + \beta_2 CAP_i + \beta_3 PROFIT\_CAP_i + MVOL_i + \varepsilon_i,$$

where $Volume_i$ represents trading volume on the last cum-dividend day for a particular option; *PROFIT* is the potential profit on the dividend play (the product of the open interest on the last cum-dividend day and the difference between the cash dividend paid on the underlying stock and the expected time value of the option); *CAP* is an indicator variable that takes the value of 1 if the trading date is after the fee cap was introduced on the exchange, and 0 otherwise; *PROFIT_CAP* is the interaction between the potential profit and the fee cap indicator; and *MVOL* is the average trading volume during the previous thirty business days. Exchanges marked with an asterisk (*) did not introduce fee changes, and are included as a control sample. The sample period is January to December 2003.

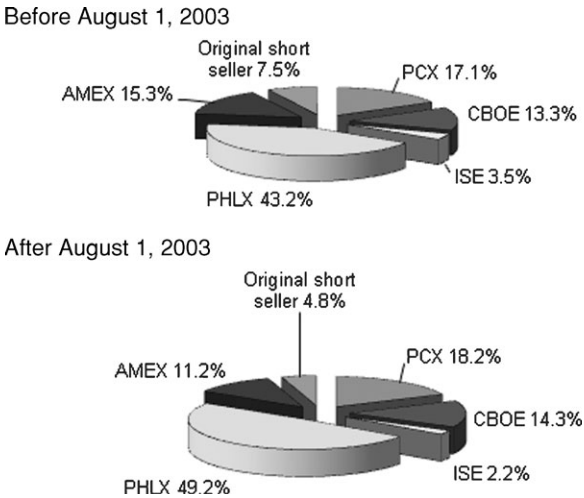

Figure 4
Division of potential profits among arbitrageurs on different exchanges and the original option writers
For each option, we compute the excess trading volume, i.e., the difference between the last cum-dividend day trading volume and the mean trading volume in the previous thirty business days in each exchange. This excess trading volume is our estimate of the fraction of the profits that the market makers capture, and the sum of the mean trading volume during the previous thirty business days in all the exchanges is the estimate of the fraction of the profits the original short sellers capture. We compute the division of the profits among the six parties before and after 1 August 2003. The division during the first part of 2003 is presented on top and the division after the reductions in fees is presented at the bottom. There are 18,803 observations and the sample period is January to December 2003.

PROFIT is significantly positive in all exchanges except for ISE, suggesting the presence of dividend play activity in four out of the five exchanges. The fee cap rule indicator is significantly positive only at the AMEX and PCX; it is marginally positively significant for PHLX. Finally, the interaction variable of CAP and PROFIT is significantly positive only for PHLX. The point estimator of the coefficient on PCX is positive but insignificant. The sensitivity of the excess trading volume to potential profits is larger after the reduction in fees at the PHLX and seems larger at the PCX.

Figure 4 details the division of the potential profits among the market makers in different exchanges and the original short sellers. For each option, we compute the excess trading volume, i.e., the difference between the last cum-dividend day trading volume and the mean trading volume in the previous thirty business days in each exchange. This excess trading volume is our estimate of the fraction of the profits that the market makers capture, and the sum of the mean trading volume during the previous thirty business days in all the exchanges is the estimate of the fraction of the profits the original short sellers capture. We compute the division of the profits among the six parties before and after 1 August 2003. The division during the first part of 2003 is presented on top and the division after the reductions in fees is presented at the bottom.

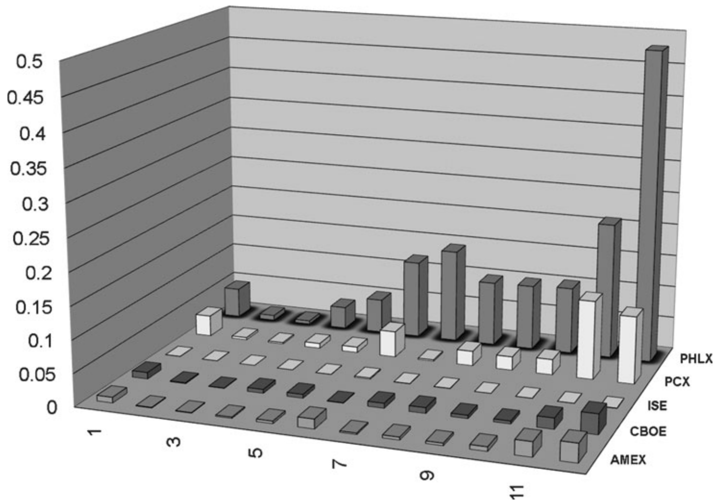

Figure 5
Impact on the reported volume of trade by month and by exchange
The impact on reported trading volume is measured as the fraction of the trading volume of call options belonging to the arbitrage sample on the last cum-dividend day in one exchange to the total trading volume in call options during the same calendar month in the same exchange. We estimate the total trading volume in equity (puts and calls) options by exchange for the month as reported in CBOE market statistics. The estimated trading volume in call options is two-thirds of these numbers.[11] The sample period is January 2003 to December 2003 and the number of observations is 18,803.

As expected, we observe a change in the composition of the profits following the implementations of the reduction in the fees. We observe a reduction of the fraction in the profits of the original short sellers from 7.5% before the change in the fee structure to 4.8% after the change and the corresponding increase in the share of PHLX from 43.2% before the reduction to 49.2% after it.

As discussed above, dividend play activity increases trading volume without increasing liquidity. Exchanges executing a large amount of dividend play trading volume might convey an incorrect impression to market participants about the level of liquidity available on that exchange. Differential quality in the reported volume is especially troubling as it can lead to an inefficient allocation of trading decisions for investors. Figure 5 depicts the extent to which volume figures have been affected by dividend play volume in each exchange. The effect on reported trading volume is measured as the fraction of the trading volume of call options belonging to the arbitrage sample on the last cum-dividend day in one exchange to the total trading volume in call options during the same calendar month in the same exchange. We estimate the total trading volume in equity (puts and calls) options by exchange for the month

[11] Information in this section was obtained from OCC Rules, OCC Rule change filings, the OCC's website, the Options Industry Council's website, and conversations with the OCC's staff.

Table 6
Windfall profits accruing to option writers

Deciles	1	2	3	4	5	6	7	8	9	10
No. of observations	2611	2611	2611	2612	2612	2612	2612	2612	2612	2612
Mean profit per option	20.76	75.51	167.7	321.6	592.3	1091.6	2113.0	4449.35	11739.63	99,539.28
Max profit per option	43.02	115.2	229.6	430.7	795.0	1461.7	2937.5	6576.87	20,679.30	6,664,400.10
Min profit per option	0.01	43.07	115.3	229.6	431.0	795.00	1462.5	2937.75	6577.43	20,680.22

This table contains a description of the profit of the writers of the call options that should have been exercised but were not. There are 26,117 observations constituting the "arbitrage" sample during the period 1996 to 2006. The arbitrage sample includes those series for which early exercise is optimal but some holders fail to exercise. The potential profit for a specific option is the product of the open interest remaining unexercised on the last cum-dividend day and the per share benefit of early exercise, which is the difference between the cash dividend paid on the underlying stock and the expected ex-dividend time value of the option. The potential profit is measured in dollars. The sample is sorted and divided into deciles based on the potential profit per option at the end of the last cum-dividend day.

as reported in CBOE market statistics. The estimated trading volume in call options is roughly two-thirds of these numbers.

Figure 5 describes substantial variation in the bias to reported volume by exchange and by calendar month. The variation in the biases across the various exchanges is problematic as it can mislead potential traders to divert their trades to a less liquid exchange.

Finally, Table 6 describes the magnitude of the profits to be divided in these dividend plays. This table contains a description of the profit of the writers of the call options that should have been exercised but were not during the period 1996 to 2006. The potential dollar profits are estimated for each of the 26,117 observations constituting the "arbitrage" sample. The sample is sorted and divided into deciles based on the potential profit per option at the end of the last cum-dividend day. The mean profit per option varies from a low of \$20.76 for the lowest decile to a high of \$4449.35 for the highest decile.

Figure 6 contains a description of the total profit to the option writers and the fraction captured by the market maker during the period 1996–2006 associated with the options in the arbitrage sample. The total potential profit in each quarter is the product of the potential profit per option (the product of the open interest remaining unexercised on the last cum-dividend day and the per share benefit of early exercise) and the number of options in the arbitrage sample. The market makers' share in the total profits is estimated by the ratio of the trading volume on the last cum-dividend day to the mean trading volume in the option during the most recent thirty business days. As is revealed in Figure 6, we find a nonmonotonic but dramatic increase in the profit opportunities associated with dividend plays and of those captured by the market maker during our sample period. Our estimate of the fraction captured by the market makers fluctuates from a low of 27.98% to a high of 68.74% but exhibits no trend. The fraction

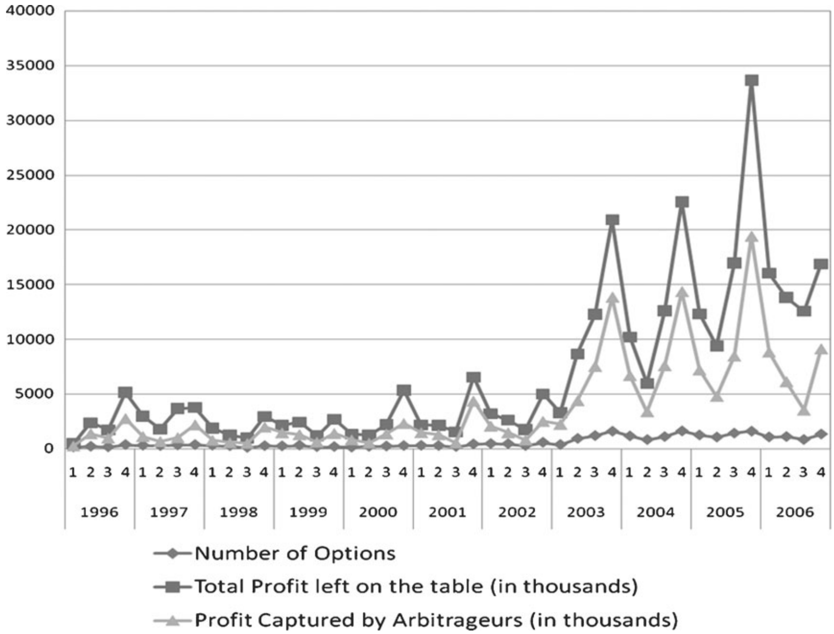

Figure 6
Profitability of dividend arbitrage
This figure shows time-series estimates of the total potential profit associated with dividend plays and the portion captured by arbitrageurs in each quarter. Results are based on the "arbitrage sample" including 26,117 observations from 1996 to 2006 for which early exercise is optimal but some holders fail to exercise. Potential profit per event is estimated as the product of the open interest remaining unexercised on the last cum-dividend day and the per share benefit of early exercise.

captured during the first quarter of 1996 is 57.54% and during the last quarter of 2006 it is 54.05%.

4.3 Implications for market efficiency

In this section, we briefly explore the implications of our results for market efficiency. In particular, we consider (1) whether suboptimal exercise behavior by the longs is likely to affect equilibrium option prices prior to the ex-dividend date, and (2) the extent to which dividend play activity keeps option prices consistent with the assumption of fully rational exercise.

To reiterate the evidence presented in Table 3, a substantial portion of open interest for in-the-money calls is not exercised before the stock goes ex dividend. Even for the most extreme decile where early exercise generates a benefit of about eighty-four cents per share, more than 32% of open interest is not exercised.

In the absence of arbitrage activity, option writers in this group might anticipate a greater than 32% chance of receiving a windfall gain of $0.84 per share. If the options were to trade at a price reflecting fully rational exercise, traders would be eager to write options. Competition among writers seeking a windfall benefit might exert downward pressure on option prices. However, as soon as prices fall below a price consistent with fully rational exercise, there will be an incentive for low-cost arbitrageurs to enter and follow the rational exercise strategy. Their actions will dilute the windfall benefit accruing to the option writers, and relieve the downward pressure on prices. Thus, the extent to which suboptimal exercise affects option prices depends on the ability of low-cost arbitrageurs to enter aggressively when the option is undervalued. In the extreme case where there are almost no limits to arbitrage, competition between rational sellers would not allow prices to fall below the fully rational price, even if a significant portion of long option holders follow an irrational exercise strategy. We would expect this to be the case for deep in-the-money calls found immediately prior to the ex-dividend date. These options have essentially no time value, and typically will be trading at intrinsic value. Any downward price pressure would cause their price to fall below the lower arbitrage bound, at which point there is an immediate, risk-free arbitrage strategy of buying and exercising immediately. Earlier in the option's life, when significant time value remains, limits to arbitrage may become somewhat more relevant and option prices may reflect suboptimal exercise.

The dividend play strategy described in this article is very low cost and is essentially a pure, risk-free arbitrage, at least for some market participants. Because the strategy has the effect of diverting the windfall gains away from the original option writers, these original writers should no longer price these gains into the option. Thus, the effect of dividend arbitrage is to eliminate pressures that might cause option prices to reflect suboptimal exercise. Our empirical evidence suggests that the degree to which option prices are consistent with the assumption of rational exercise is sensitive to market maker transaction and clearing fees.

The theoretical model presented in Appendix B also provides some guidance as to how much benefit we should expect to accrue to preexisting option writers. In this model, the expected benefit accruing to the marginal option writer prior to the last ex-dividend date is given by the probability of nonassignment times the benefit of nonassignment. This can be computed in the model by plugging in the equilibrium amount of dividend-play activity. This exercise confirms that the marginal benefit accruing to option owners is quite sensitive to the market makers' fee, and also to whether the equilibrium is monopolistic or competitive. If the dividend is large, the incidence of suboptimal exercise is high, the transaction cost for arbitrageurs is significant, and competition between arbitrageurs is not fully competitive, the marginal benefit to early option writers might be large enough to have a noticeable impact on option prices. As a simple example, consider the case where the dividend is 0.50 cents and 30% of preexisting option longs fail to exercise. A transaction cost of 0.35 cents per contract corresponds to a marginal benefit of about 0.023 cents per share under the monopolist equilibrium or 0.0035 cents in the competitive equilibrium. A transaction cost of 0.09 cents per contract corresponds to a marginal benefit of about 0.012 cents per share in the monopolist equilibrium or 0.0009 cents per share under the competitive equilibrium.

5. Conclusion

This article joins prior research in documenting that option holders often seem to follow a suboptimal exercise policy. We find that a significant fraction of the open interest of options that should have been exercised prior to the ex-dividend date remains unexercised, creating a benefit for option writers. The design of the exercise allocation mechanism enables certain market participants, such as market makers, to capture this benefit, through an arbitrage trading strategy that involves huge offsetting trades. Recent decreases in transaction and clearing fees for market makers have significantly increased the optimal scale of this strategy, and as a result, trading volume appears to have become less informative as a measure of liquidity.

One might surmise that this activity would be eliminated if the clearinghouse were to net out purchases and sales within each account before processing exercises. On the surface, it might appear that this would make the arbitrage infeasible. However, the strategy could still be implemented by multiple market makers coordinating their positions across accounts. Moreover, the implementation of the dividend arbitrage described in this article is just one simple variant of the strategy, but other possibilities exist. For example, if there are multiple in-the-money strike prices with open interest, two parties can implement the arbitrage by entering into opposite sides of a spread transaction, rather than executing the entire strategy on a single strike price.

The total amount of money left on the table is small, relative to the huge aggregate positions taken in this market in order to capture these gains. Because of the aggressive actions of arbitrageurs with a low marginal trading cost, the irrational behavior on the part of option holders does not generate any substantial benefit for parties who write options prior to the last cum-dividend day. When marginal trading costs are low for market makers, systematic suboptimal exercise by option holders should have no appreciable impact on the value of the option. Thus, despite evidence of irrational exercise, our evidence indicates that it is a good approximation to value options under the assumption of a rational exercise policy.

Appendix A: Institutional Background

Equity options traded on U.S. exchanges are cleared through the Options Clearing Corporation (OCC).[12] The OCC currently has approximately 130 clearing members, including registered broker-dealers and foreign securities firms. For every trade negotiated on a U.S. option exchange,

[12] The trading volume in call options is about two-thirds of the trading volume in all the equity options (call and put) in the market, according to CBOE 2003 market statistics.

each side is intermediated by a clearing member. That is, both the buyer and seller must either be a clearing member, have an account with a clearing member, or trade through a broker that has an account with a clearing member.

The OCC tracks long and short option positions at the clearing member level, and also provides clearing firms with separate customer, firm, and market maker accounts. If multiple market makers clear through the same firm, each market maker's transactions will clear through a subaccount. If the same individual acts as a market maker in some classes but also trades in other classes not as a market maker, the two types of trades will be cleared through different accounts.

A key aspect of the clearing system is that the OCC becomes the buyer (seller) to the seller (buyer) for matched trades reported by the option exchanges to the OCC. By standing as counterparty to both sides of the transaction, the OCC guarantees performance on the contract, and assumes the counterparty default risk. Once the trade is cleared, there is no longer any direct link between the buyer and seller. Thus, when a long party chooses to exercise, the process for assigning the exercise does not in any way take into account which among the shorts happened to be counterparty to the original trade.

When option trades are reported to the OCC, each party marks its side as an "opening" or "closing" transaction—that is, a transaction that opens a new option position or that closes out an old one. Market maker trades, however, are not marked as their positions are carried net at the OCC. At the end of each business day, the OCC's clearing system updates all accounts to reflect all that day's trading activity and exercises, and then assigns the day's exercises. These positions are updated in the following sequence. For market maker accounts, purchases are processed first, then exercises, and then sales. For other accounts, opening buys are processed first, then opening sells, then closing buys, then exercises, then closing sells.

Because closing buys are processed before assignments, a party who closes out a short position is not at risk of having the closed position assigned on the same day. Because purchases are processed before exercises, it is possible for a party to buy and exercise the same day. Because sales are not processed until after exercises, a party can buy, exercise, and sell the same series on the same day.

After all the accounts have been adjusted to reflect trading and exercises for each series, the clearinghouse assigns all that day's exercises, allocating these assignments across the clearing members who represent short positions in such series. These assignments are made according to the "standard algorithm."[13] Under this algorithm, each position account (including market maker subaccounts) is given a unique identification code. All the short positions in a given series are arranged on a virtual "wheel," in order of this identification code, with larger positions taking up more "space" on the wheel in proportion to the size of the short position. Then, a random starting place is selected on the wheel, and the first twenty-five contracts are assigned. A certain "skip interval" is then calculated, which is a function of the total number of short positions on the wheel and the total number of exercises. The algorithm then rotates around the wheel, skipping a number of positions indicated by the skip interval calculations, and then assigning another twenty-five contracts. The algorithm repeats like this, skipping around the wheel and assigning exercises in blocks of twenty-five contracts, until all the exercises have been assigned.

Under this algorithm, the assignment has an important random element, inasmuch as the starting point on the wheel is selected at random. This makes it impossible to predict whether any given small position will be assigned, or to predict the exact proportion of a large position that will be assigned. However, by rotating through all positions in this fashion, the algorithm ensures that the realized assignment is not too far from a *pro rata* allocation, in proportion to the size of the positions. Because the OCC has separate position subaccounts for individual market makers, the allocation of exercises to market makers is determined entirely by the OCC algorithm. On the other hand, a clearing member's customers' positions are generally cleared in a single omnibus account.

[13] The algorithm is described in OCC Rule change filing, Securities and Exchange Commission Release 34-46735, and is also published in the Federal Register, 5 November 2002.

In this case, the OCC would allocate assignments to a clearing member's omnibus account (i.e., its customers' account on the OCC's books), but allocating the exercises among the individual customers would be done at the clearing member level. The rules of the option exchanges[14] dictate that their member firms must establish fixed procedures for assigning exercises, but give the brokers flexibility to use "first in, first out" or random selection.

Appendix B: Optimal Scope of Dividend Arbitrage for Monopolist and Competitive Traders

We consider an option for which early exercise is optimal. As before, let H represent the open interest on the last cum-dividend date before the arbitrageurs enter, and let α represent the fraction of options that are not exercised. In the absence of any arbitrage activity, the clearinghouse will assign the $(1-\alpha)\,H$ exercises across the H written option positions, using their standard allocation algorithm. Those option writers lucky enough to avoid assignment avoid having to pay the value of the dividend, but remain with the liability of a written option.

B.1 The case of a monopolist

On the last cum-dividend date, trader 1 buys calls on $Q/2$ shares from trader 2, and trader 2 buys calls on $Q/2$ shares from trader 1, so that together, the new trades establish new option positions on Q shares. The variable Q will be selected by the arbitrageurs to maximize the profitability of the strategy. As stated before, the two trades are executed at the same price, so no cash is required to settle these trades. The traders then exercise all Q shares of their long positions. The clearinghouse now must allocate the $(1-\alpha)\,H+Q$ exercises across $H+Q$ written positions. These assignments will be allocated across the preexisting writers, who hold a proportion $H/(H+Q)$ of all written positions, and the two new traders, who hold a proportion $Q/(H+Q)$ of the written positions.

After the assignment, the two traders would, on average, remain with $Q/(H+Q)$ of the unassigned written positions, and expect to capture that proportion of the total gains. Because the initial trades are exactly offsetting, no initial cash flow is required. The profits of the strategy are given by the total gain captured, minus transaction and clearing fees. Denoting the per-unit transaction and clearing fees by C, the total profit from executing the strategy is thus

$$\Pi = \left(\frac{Q}{H+Q}\right)\alpha H(D^{*}-\pi_e) - CQ. \tag{B1}$$

Monopolist arbitrageurs would choose Q to maximize profits. The first-order condition for profit maximization is

$$\frac{d\Pi}{dQ} = \frac{H}{(H+Q)^2} + \alpha H(D^{*}-\pi_e) - C = 0 \tag{B2}$$

or, solving for the profit maximizing size,

$$Q = H\left(\sqrt{\frac{\alpha(D^{*}-\pi_e)}{C}} - 1\right). \tag{B3}$$

Thus, in a market of monopolistic pair of traders, we would expect to see the volume of dividend play activity proportional to the open interest going in to the last cum-dividend day, positively related to the proportion of option holders failing to exercise optimally and the per-unit early exercise benefit, and negatively related to transaction costs. Calibrated using reasonable numbers,

[14] Chicago Board Options Exchange Rule 11.2; International Securities Exchange Rule 1101; American Stock Exchange Rule 981; NYSE Arca Rule 6.25; Philadelphia Stock Exchange Rule 1043; Boston Option Exchange VII-2(a).

this formula predicts that trading volume on the last cum-dividend date should be multiple times higher than preexisting open interest. Also, the formula predicts that trading volume should be quite sensitive to a reduction in the marginal transaction cost.

For example, consider the case where the dividend is \$0.50 per share; the option is a deep in-the-money option and close to expiration, so its ex-dividend time value is essentially zero, and 30% of option owners fail to exercise. Equation (B3) indicates that if the marginal transaction and clearing fees are \$0.35 per contract, the optimal dividend play trade would be roughly 5½ times the preexisting open interest. If the marginal cost drops to \$0.09 per contract, the optimal dividend play would be nearly 12 times the open interest. The change in these fee numbers roughly corresponds to the marginal fee levels before and after the transaction capping fee rules implemented during our sample.

B.2 Competitive equilibrium

Suppose there are an unlimited number of arbitrageurs, each with a marginal transaction cost of C. Consider the decision facing the marginal arbitrageur, who comes to the market not having any open option position. If the marginal arbitrageur implements a dividend play for one unit, he/she expects a share of $1/(Q + H)$ of the total money lying on the table. Thus, the expected marginal benefit is

$$\left(\frac{1}{H+Q}\right)\alpha H(D^* - \pi_e). \tag{B4}$$

In the competitive environment, we would expect new arbitrageurs to enter until the marginal benefit of entering equals the marginal cost. So, the competitive equilibrium would be where

$$\left(\frac{1}{H+Q}\right)\alpha H(D^* - \pi_e) = C \tag{B5}$$

or

$$Q = H\left(\frac{\alpha(D^* - \pi_e)}{C} - 1\right). \tag{B6}$$

Compare this with the monopolist quantity derived earlier. The competitive equilibrium quantity is proportional to $\alpha(D^* - \pi_e)/C$, while the monopolist equilibrium is proportional to the square root of $\alpha(D^* - \pi_e)/C$. Expected trading volume is substantially higher under the competitive equilibrium. For the parameter values described above, the competitive equilibrium would generate trading volume more than forty times higher than the open interest for a fee of 0.35 cents per contract, or a trading volume more than 165 times higher than the open interest for a fee of 0.09 cents per contract.

References

Black, F., and M. Scholes. 1973. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* 81:637–54.

Boni, L. 2006. Strategic Delivery Failures in U.S. Equity Markets. *Journal of Financial Markets* 9:1–26.

Finucane, T. 1997. An Empirical Analysis of Common Stock Call Exercise: A Note. *Journal of Banking & Finance* 21:563–71.

Kalay, A., and M. G. Subrahmanyam. 1984. The Ex-dividend Day Behavior of Option Prices. *Journal of Business* 57:113–28.

Merton, R. 1973. Theory of Rational Option Pricing. *Bell Journal of Economics & Management Science* 4:141–83.

Pool, V. K., H. R. Stoll, and R. E. Whaley. 2008. Failure to Exercise Call Options: An Anomaly and a Trading Game. *Journal of Financial Markets* 11:1–35.

Poteshman, A. M., and V. Serbin. 2003. Clearly Irrational Financial Market Behavior: Evidence from the Early Exercise of Exchange Traded Stock Options. *Journal of Finance* 58:37–70.

Roll, R. 1977. An Analytic Valuation Formula for Unprotected American Call Options on Stocks with Known Dividends. *Journal of Financial Economics* 5:251–58.

White, H. 1980. A Heteroscedasticity-consistent Covariance Estimator and a Direct Test for Heteroscedasticity. *Econometrica* 48:817–38.