Strategic Arbitrage in Segmented Markets

Svetlana Bryzgalova,* Anna Pavlova,† and Taisiya Sikorskaya‡

August 31, 2023

Abstract

We propose a model in which arbitrageurs act strategically in markets with entry costs. In a repeated game, arbitrageurs choose to specialize in some markets, which leads to the highest combined profits. We present evidence consistent with our theory from the options market, in which suboptimally unexercised options create arbitrage opportunities for intermediaries. Using transaction-level data, we identify the corresponding arbitrage trades. Consistent with the model, only 53% of these opportunities attract entry by arbitrageurs. Of those that do, 56% attract only one arbitrageur. Finally, our paper details how market participants circumvent a regulation devised to curtail this arbitrage strategy.

JEL Classification: G4, G5, G11, G12

Keywords: Arbitrage, repeated game, entry cost, dividend play, retail investors

*We thank Brad Barber, Robert Battalio, Jean-Pierre Benoit, Anne Haubo Dyhrberg, Terry Hendershott, Su Li, Emre Ozdenoren, Dmitriy Muravyev, Chester Spatt, Dimitri Vayanos, Ingrid Werner, and seminar participants at Indiana University, London Business School, Stockholm Business School, the University of Central Florida, and the 2023 Women in Microstructure meeting for their thoughtful comments, Miguel Chumbo for excellent research assistance, and the AQR Asset Management Institute at London Business School for generous support. A part of this paper was previously circulated under the title, "Profiting from Investor Mistakes: Evidence from Suboptimal Option Exercise."

*London Business School and CEPR, sbryzgalova@london.edu

†London Business School and CEPR, apavlova@london.edu

‡London Business School, tsikorskaya@london.edu

1 Introduction

The dominant paradigm in the literature studying limits to arbitrage and arbitrage in general is the competitive arbitrageur paradigm, whereby arbitrageurs are atomistic. Yet in practice, we often observe just a handful of large arbitrageurs operating in some markets, because of high fixed costs of entry or owing to economies of scale. In many applications, strategic arbitrageurs interact repeatedly, which gives rise to repeated game considerations, currently absent from the existing literature. An environment incorporating these features may offer a better description of arbitrageur interaction in markets that are more opaque and concentrated.

In this paper, we depart from the competitive arbitrageur paradigm and introduce strategic interactions. Intermediaries, who act as arbitrageurs, face fixed costs of market entry. We consider both one-shot and repeated games, and show that the latter give rise to the highest arbitrageur profits. In repeated game equilibria, arbitrageurs do not necessarily enter all possible markets, as they would in a one-shot Nash equilibrium. Instead, it is more profitable for the intermediaries to specialize in harvesting arbitrage opportunities in their "natural" markets, in which their fixed costs are low or zero, and enter other markets only when no (or few) other intermediaries conduct arbitrage activities in those markets. This is a "tacit collusion" equilibrium, which arises in repeated games.

As a laboratory for the study of arbitrageur interactions, we use the so-called dividend play, a strategy available to the intermediaries in the U.S. options market which produces (virtually) riskless arbitrage profits. These profits derive from call options left suboptimally unexercised before the underlying stock goes ex-dividend. Exploiting transaction-level data, we are able to accurately identify arbitrageur trades in the dividend play strategy in each option contract. We document that only 53% of arbitrage opportunities – pertinent call option contracts – attract entry of arbitrageurs, and these arbitrage opportunities are not the most profitable. This is consistent with our model's prediction that the remaining seemingly unexploited arbitrage profits are fully harvested by incumbent natural intermediaries and in the data we should not observe other arbitrageurs' entry. This lack of entry by other arbitrageurs cannot be explained by their costs or constraints. We further document that 56% of opportunities that attract entry, are exploited by only one arbitrageur. This is consistent with the model's implication that, for markets (contracts) in which there is no natural intermediary, it is profitable for arbitrageurs to specialize and to harvest arbitrage profits only in a subset of available contracts.

Here is how the model works. There are three markets with one arbitrage opportunity of fixed value each. There are two risk-neutral intermediaries who exploit these opportunities.

The first (second) market is the natural market for the first (second) intermediary. A natural market is defined as a market with zero fixed costs of entry. Examples of natural markets include securities in which an intermediary is already a market maker or a major dealer and thus has a natural advantage in monitoring and accessing potential arbitrage opportunities in these securities. More generally, economies of scale could be creating natural markets for intermediaries. In all other markets, intermediaries in our model face (per-period) fixed costs of entry. Intermediaries act strategically, choosing which markets to enter and, conditional on entry, how much resource to commit to each market. The share of arbitrage opportunities that is captured by them in each market is proportional to their market share. The game is symmetric in payoffs and we focus on symmetric equilibria.

We first consider a one-shot game. The only Nash equilibrium in that game is one in which the intermediaries enter all markets and equally split the arbitrage opportunities in them. This split, however, does not maximize the joint welfare of the intermediaries. In a welfare-maximizing allocation, arbitrageurs specialize in a subset of markets and thus reduce the total fixed cost paid, while still harvesting the same arbitrage gains. A cooperative equilibrium of this type is not a Nash equilibrium of a one-shot game but it becomes a subgame perfect equilibrium of an infinitely-repeated game, supported by standard trigger strategies. This equilibrium is known in the industrial organization literature as a tacit collusion equilibrium (Tirole (1988)). The term "tacit" conveys that the equilibrium does not require any explicit agreements between intermediaries. In the tacit collusion equilibrium in our model, intermediaries do not enter each others' natural markets and alternate their entry in the third market. This results in the minimum per-period fixed costs. The testable predictions of the model is that in the data we should observe no market entry into intermediaries' natural markets and observe entry of a limited number or arbitrageurs – just one in our model – in markets that are neither intermediaries' natural markets.

The U.S. options market offers a unique laboratory to test our theory. We are able to identify an arbitrage opportunity, for which we observe arbitrageur entry and the amount of resources they commit to exploiting it with high accuracy. Most of related literature on limits to arbitrage has focused on the effects of arbitrage on prices, normally because quantities are not observed in non-proprietary data. Our setting is an exception. In fact, the effects on prices are not relevant for our application. An arbitrage opportunity, as in our model, is best thought of as a windfall gain, which accrues to arbitrageurs in proportion to their market shares. The opportunity results from a failure of investors to exercise in-the-money call options before the underlying stock or exchange-traded fund (ETF) goes ex-dividend when it is optimal to do so.[1] To benefit from it, arbitrageurs engage in "dividend play,"

[1]We note that sometimes call options may be purchased as part of any strategy that involves holding multiple

an arbitrage strategy that diverts windfall gains from the original writer of the option that was suboptimally left unexercised by dormant investors. Option writers receive this gain in expectation, as the Options Clearing Corporation (OCC) randomly assigns options writers who receive an exercise notice and must deliver the underlying. The writers who are not assigned receive the windfall capital gain. Dividend play is a dilution strategy in which on the last cum-dividend date, arbitrageurs write an extremely large number of call option contracts, maximizing the probability that the gain from exercise mistakes is assigned to them rather than to the original option writer. To hedge the large short position required for the strategy, arbitrageurs simultaneously establish an offsetting large long position in a similar contract. Dividend play trades are usually prearranged by pairs of arbitrageurs, who serve as counterparties to each others' long and short positions, and in our sample, actual transaction prices are close to the midpoint of the bid-ask spread. So price impact is not a concern in our application; the windfall gain is fixed and it is effectively divided pro-rata between arbitrageurs who participate in the dividend play strategy, as in our model.

The dividend play strategy is normally executed on a physical exchange floor. This is because some exchanges facilitate this trade by capping the fees for this strategy for floor participants,[2] and because the floor offers pairs of arbitrageurs an opportunity to execute (large) bilateral trades on agreed terms, without exposing them to other market makers' quotes. Our empirical strategy for classifying dividend play trades relies on the flags for floor trades in the transaction-level Options Price Reporting Authority (OPRA) data. With most of the options trading being electronic, exchange floors are more akin to a library than to a bustling trading floor of the past. So, like in our model, there is usually a small number of participants in the dividend play arbitrage, who interact repeatedly with each other on pre-scheduled (last cum-dividend) dates.

Since the dividend play strategy requires opening extremely large positions, the daily trading volume on the last cum-dividend dates in in-the-money call options targeted by the strategy often exceeds trading volume on the remaining dates by several orders of magnitude. Even for the largest S&P 500 ETF, SPY—the ticker with the most actively traded options in 2021—the last cum-dividend date trading volume is typically 14–53 times larger.[3] In

option contracts. In those circumstances, or whenever transaction costs outweigh profits from early exercise, exercising an option may not be optimal.

[2]For example, PHLX imposes daily fee caps for floor market makers and other floor traders engaging in dividend play. See, e.g., https://listingcenter.nasdaq.com/rulebook/phlx/rules/phlx-options-7, accessed July 26, 2023, for the dividend strategy fee caps imposed by the PHLX options exchange. More than two-thirds of dividend play transactions in our sample are executed on PHLX. Dividend play is not the only strategy that benefits from fee caps. For example, PHLX offers similar fee caps for five other arbitrage strategies.

[3]The lower bound compares the average last cum-dividend date dollar trading volume in call options to an average across all days in our sample, while the upper bound compares to the average volume in a week

this paper, we often refer to the last cum-dividend dates on which dividend play trade takes place as dividend play dates.

During the recent retail investor boom in options (documented in e.g., Bryzgalova, Pavlova, and Sikorskaya (2023)), expected profits from the dividend play trades have been growing rapidly. Most of this profit derives from the sheer increase in open interest due to investor inflow, coupled with a higher fraction of options that are left unexercised on dividend play dates. For the latter, we document that in the cross-section, arbitrage activity is higher in contracts with a larger share of retail activity in the preceding week. Overall, in our sample, dividend play participants behave like unconstrained arbitrageurs in that they are able to establish extremely large long-short positions in pertinent contracts.

There are, however, four striking patterns that emerge from our examination of dividend play transactions. First, only 53% of arbitrage opportunities (profitable option contracts) attract arbitrageur entry, which implies that about 50% of potential profits appear to be left of the table.[4] Second, if arbitrageurs do decide to participate in dividend play in a given contract, they harvest close to all available profits in that contract. Third, around 42% of contracts in the top profitability tercile do not attract entry by arbitrageurs. More generally, even among very similar profitable contracts, arbitrageurs often enter only a few. In light of the standard explanations of limits to arbitrage, the first and third patterns are quite puzzling. Arbitrageurs' daily fee on dividend play trades is capped by most exchanges on which dividend play trades take place, so additional fees are not a plausible explanation for leaving money on the table. Trading costs are quite low because dividend play transactions are executed close to the midpoint of the bid-ask spread. Capital constraints are also unlikely to explain the findings because market participant exposure is normally computed at a ticker level, and so the large long and short positions in contracts on the same ticker, required for the dividend play trade, are netted to zero.

The fourth pattern concerns the number of arbitrageurs that engage simultaneously in dividend play in a given contract. We take advantage of the granularity of our data to make inferences about this number. Strikingly, in 56% of the contracts, in which we detect arbitrageur entry, we observe that the entire profit is captured by a single arbitrageur. Across all contracts, the number of arbitrageurs simultaneously participating in dividend play in a given contract exceeds 3 in only 6% of the cases and exceeds 7 in less than 1%. This is surprising because there are no obvious restrictions on free entry of other arbitrageurs who already engage in dividend play in other contracts. Partly owing to their sheer size and

prior to the last cum-dividend date.

[4]Table 4 in the Appendix quantifies forgone profits of market makers in the top 40 most popular underlying stocks and ETFs for the dividend play strategy in our sample.

distinct execution patterns, dividend play trades are easy to detect in transaction-level data (e.g., OPRA). Arbitrageurs are therefore aware of the history of other arbitrageurs' dividend play activity in a given contract before they decide to engage in dividend play in it.

The implications of our model are consistent with all of the above patterns. One may interpret the contracts in which we observe no entry as natural markets of some arbitrageurs. Many arbitrageurs in this strategy are also market makers, who write options that retail investors buy (and fail to exercise optimally). Using an additional data source covering a key subset of exchanges in our dataset,[5] we are able to show that market makers are indeed net short on average in the top 100 most mentioned tickers on `WallStreetBets`, an investing forum popular with retail investors (Figure 5 in the Appendix). These market makers then do not incur any fixed costs and simply pocket the windfall gains on the corresponding written options already in their inventory, unless other arbitrageurs enter and compete away those gains. This fits exactly our definition of a natural market for an intermediary. Our empirical tests reveal that arbitrageurs are less likely to enter profitable contracts that have positive retail buy imbalances (and hence market makers likely hold short positions in these contracts in their inventory). In contrast, in written contracts that neither intermediary has in their inventory, we should expect arbitrageur entry so as to harvest the gains. This entry, however, should be limited. The costs of arbitrage are fixed because of the exchanges' daily fee caps on dividend play. Our model implies that arbitrageurs should therefore specialize and enter only a subset of markets, with the tacit collusion equilibrium corresponding to only one arbitrageur entering a market (contract) at a time.

There are other potential explanations of the empirical patterns we document. It is possible that the reluctance of some firms to engage in the dividend play arbitrage could be explained by the operational risk of the trade.[6] Another possibility is the presence of potentially prohibitive clearing fees, faced by market participants. Unfortunately, information on clearing fees is not public, and so we have not been able to verify this potential explanation. It is also possible that there is a stigma associated with the dividend play strategy, since it is frowned upon by the U.S. Securities and Exchange Commission (SEC).

Finally, we discuss the difficulties of devising effective regulation in the derivatives market. Concerned about the impact of dividend play trades on the orderly functioning of the market, in 2014 the SEC issued a new rule designed to make the strategy impractical,[7]

[5]We thank Nasdaq Global Data Products for generously sharing their historical PHOTO and NOTO End-of-Day data with us. PHOTO covers the PHLX exchange, in which we observe the majority of the trading volume associated with dividend play.

[6]For example, a human error in the dividend play strategy inflicted a $10 million loss on Bank of America Merrill Lynch. See https://www.reuters.com/article/us-usa-options-apple-idUSKBN0IQ2FA20141106.

[7]See https://www.sec.gov/rules/sro/occ/2014/34-73438.pdf.

which resulted in much lower trading volumes on dividend play dates. However, the recent dramatic increase in options trading by inexperienced retail investors appears to have led to a resurgence of the strategy, whereby arbitrageurs have found a way to circumvent the barriers created by the SEC rule. We propose a rule change that may curtail the strategy.

Our paper is related to the limits-to-arbitrage literature, such as Gromb and Vayanos (2002) and Shleifer and Vishny (1997). With the exception of Fardeau (2021), arbitrageurs in this literature behave competitively. In contrast to the strategic arbitrageurs in Fardeau (2021) for whom price impact is a primary consideration, in our setting, price impact is not a concern. Instead, our primary focus is not on prices but on quantities of arbitrage and on multi-market arbitrageur interactions and specialization in a repeated-game setting. Arbitrage opportunities are in fact fully exploited in our case, and limits to arbitrage should be understood as (endogenous) limits on the number of arbitrageurs in a given market. Our model is consistent with recent evidence of Siriwardane, Sunderam, and Wallen (2022) (see also Boyarchenko, Eisenbach, Gupta, Shachar, and Tassel (2020)), that strongly suggests that arbitrage is segmented and arbitrageurs specialize in a small set of arbitrage opportunities.

The closest paper considering a repeated game in financial markets is a tacit collusion model of Dutta and Madhavan (1997). As in the classical industrial organization literature, the authors consider pricing strategies that give rise to prices above competitive levels. Their model explains the puzzling tendency of NASDAQ market makers in equities to avoid odd-eighth quotes, uncovered in an influential paper by Christie and Schultz (1994). There is of course a vast literature in industrial organization on tacit collusion (Tirole (1988)), but it normally studies the behavior of price-setting firms. Our mechanism is novel in that the welfare gain to the players from tacit collusion is generated by reduced fixed costs as opposed to higher prices. The application to arbitrage in financial markets in also novel.

The empirical application, dividend play, has been considered in Hao, Kalay, and Mayhew (2009) and Pool, Stoll, and Whaley (2008), who were the first to describe this strategy. Unlike us, they do not compute the number of arbitrageurs in a given contract and do not document the widespread selective entry of arbitrageurs in profitable contracts, which is our primary focus. Our paper also relates to the literature on optimal option exercise by investors. It has been previously documented that not all American options are exercised rationally (e.g., Poteshman and Serbin (2003)). Battalio, Figlewski, and Neal (2020), Cosma, Galluccio, Pederzoli, and Scaillet (2020), Jensen and Pedersen (2016), and Barraclough and Whaley (2012) focus on early exercise decisions and show in more recent data that a fraction of investors still fail to exercise their options optimally. More generally, the literature has documented a variety of mistakes that retail investors make (see, e.g., Barber and Odean (2001), Calvet, Campbell, and Sodini (2007), and Barber and Odean (2013)).

2 Model

In this section, we develop a game-theoretic model of arbitrage in segmented markets with a small number of arbitrageurs, who repeatedly interact with each other. In the empirical part of the paper, we focus on a specific application, which is likely to satisfy the model's assumptions – dividend play – and show that the patterns of arbitrage we observe in the data are consistent with our model's predictions. The model, however, is applicable more broadly, to environments which feature strategic arbitrage with repeated arbitrageur interactions.

2.1 Economic environment

Time is discrete and the horizon is infinite. There are three markets, M_1, M_2, and M_3, each with one arbitrage opportunity. For simplicity, the per-period value of the arbitrage opportunities in the first two markets is the same, and we denote it by A. The per-period value of the arbitrage opportunity in the third market is A^*.

There are two special intermediaries, 1 and 2, who are the natural arbitrageurs in markets M_1 and M_2, respectively. That is, Intermediary 1 has a zero cost of harvesting the arbitrage opportunity in market M_1 and Intermediary 2 in market M_2. Neither of the two intermediaries is a natural arbitrageur in market M_3. Both intermediaries are risk-neutral. In what follows, we use the terms "arbitrageur" and "intermediary" interchangeably.

Each intermediary has 1 unit of resource that she allocates across the three markets. Intermediary 1's per-period costs of exploiting the arbitrage opportunities in markets M_1, M_2, and M_3 are 0, f, and f^*, respectively, and those of Intermediary 2 are f, 0, and f^*. The fixed cost is incurred only if an intermediary enters a market, i.e., commits a non-zero resource to the market. If an intermediary does not enter a market, its payoff in that market is 0.

Intermediary 1's per-period payoff is

$$\pi_1(k_1^1,\, k_1^2,\, k_1^3) \equiv A\frac{k_1^1}{k_1^1+k_2^1} + A\frac{k_1^2}{k_1^2+k_2^2} + A^*\frac{k_1^3}{k_1^3+k_2^3} - f\,1_{k_1^2>0} - f^*\,1_{k_1^3>0}, \tag{1}$$

where resources allocated to each of the three markets $(k_1^1,\, k_1^2,\, k_1^3)$ are choice variables and $k_1^1 + k_1^2 + k_1^3 \leq 1$. Analogously, Intermediary 2's per-period payoff is

$$\pi_2(k_2^1,\, k_2^2,\, k_2^3) \equiv A\frac{k_2^1}{k_1^1+k_2^1} + A\frac{k_2^2}{k_1^2+k_2^2} + A^*\frac{k_2^3}{k_1^3+k_2^3} - f\,1_{k_2^1>0} - f^*\,1_{k_2^3>0}, \tag{2}$$

where $k_2^1 + k_2^2 + k_2^3 \leq 1$. The intermediaries therefore harvest the arbitrage opportunities in

proportion to their market share of resources invested in a given market. In the special case of $k_2^1 = 0$, Intermediary 1 harvests the entire arbitrage opportunity in market M_1, even if she commits no resources to that market. Likewise, if Intermediary 1 refrains from entry in market M_2, the entire arbitrage profit A accrues to Intermediary 2, regardless of the amount of resource in that market. In none of the intermediaries enter market M_3, the arbitrage opportunity in that market remains unexploited.

Our main focus is on repeated games. The time discount factor is δ. The risk-neutral intermediaries act strategically and choose their best response per-period resource allocations $(k_{it}^1, k_{it}^2, k_{it}^3)$ to maximize

$$\sum_{t=0}^{\infty} \delta^t \pi_i(k_{it}^1, k_{it}^2, k_{it}^3),$$

taking the resource allocation of the other intermediary as given. We will also consider Pareto efficient equilibrium allocations, which maximize social welfare, i.e.,

$$\mu_1 \sum_{t=0}^{\infty} \delta^t \pi_1(k_{1t}^1, k_{1t}^2, k_{1t}^3) + \mu_2 \sum_{t=0}^{\infty} \delta^t \pi_i(k_{2t}^1, k_{2t}^2, k_{2t}^3), \tag{3}$$

subject to the resource constraints $k_{it}^1 + k_{it}^2 + k_{it}^3 \leq 1$, $i \in \{1, 2\}$, where μ_1 and μ_2 are the Pareto weights. Social welfare here is understood in the game-theoretic sense: it takes into account welfare of the two players (the intermediaries), but not any other agents in the economy. With a restriction on the time discount factor δ, Pareto efficient equilibria can be supported as non-cooperative (subgame perfect) equilibria of a repeated game, as we demonstrate below.

The concept of arbitrageurs' natural markets is novel to this paper. We think of natural markets as the ones in which an arbitrageur incurs lower costs than the competition. This can happen if a arbitrageur is an intermediary that is close to the order flow in the market (e.g., a market maker or a prominent high-frequency trading firm) or has a structure of its balance sheet that gives it a natural advantage in exploiting a particular arbitrage opportunity. One example is the options market arbitrage from the trading floor that we use as a laboratory in the next section. In other markets, natural hedges may create a similar environment. For example, if an intermediary is a bank and has structured products on its balance sheet, it may be "naturally" exposed to dividend risk Mixon and Onur (2017).

Fixed costs of market access are an important ingredient of our model. Recent limits-to-arbitrage literature (Siriwardane, Sunderam, and Wallen (2022)) stresses that in practice arbitrage is segmented, whereby arbitrageurs specialize in one market or a small handful of markets. Such segmentation would naturally happen in the presence of fixed costs of arbitrage. We model fixed costs in the natural markets to be zero for simplicity, but what is

key to our implications is that fixed costs in a given market are lower for natural arbitrageurs in that market.

For the emergence of cooperative or tacit collusion equilibria as subgame perfect equilibria in our model, it is important that we consider a repeated game. Repeated interaction of arbitrageurs commonly occurs in practice. If arbitrage is indeed segmented, it is practical for the same arbitrageurs to specialize in a given market over time, as opposed to switching to a different market every period.

Finally, the focus of the model is on the markets with a fixed number arbitrageurs who act strategically. An important macroeconomic trend of the past few decades is the rise in industry concentration. (See Grullon, Larkin, and Michaely (2019), Autor, Dorn, Katz, Patterson, and Van Reenen (2020), and De Loecker, Eeckhout, and Unger (2020).) Financial intermediation and market making and is no exception. We believe that it is important to relax the standard competitive arbitrageur paradigm and to consider strategic arbitrageur interactions, at least in some more opaque and in more concentrated markets.

2.2 Pareto efficient and Nash equilibria of a stage game

We first consider a one-shot game and derive its Nash equilibrium. Intermediary 1's best response to Intermediary 2's actions is the solution to the following optimization problem, taking $(k_2^1,\, k_2^2,\, k_2^3)$ as given:

$$\max_{k_1^1,k_1^2,k_1^3} A\frac{k_1^1}{k_1^1+k_2^1} + A\frac{k_i^2}{k_1^2+k_2^2} + A^*\frac{k_1^3}{k_1^3+k_2^3} - f\,1_{k_1^2>0} - f^*\,1_{k_1^3>0}, \tag{4}$$

$$s.t.\ k_1^1+k_1^2+k_1^3 \leq 1. \tag{5}$$

Analogously, Intermediary's 2 best response is the solution to

$$\max_{k_2^1,k_2^2,k_2^3} A\frac{k_2^1}{k_1^1+k_2^1} + A\frac{k_2^2}{k_1^2+k_2^2} + A^*\frac{k_2^3}{k_1^3+k_2^3} - f\,1_{k_2^1>0} - f^*\,1_{k_2^3>0}, \tag{6}$$

$$s.t.\ k_2^1+k_2^2+k_2^3 \leq 1, \tag{7}$$

where she takes the resource allocation of Intermediary 1, $(k_1^1,\, k_1^2,\, k_1^3)$, as given. The best responses are presented in Lemmas 1 and 2 in the Appendix. The intersection of the best responses is a Nash equilibrium of the one-shot game. Of specific interest to us is the interior equilibrium, in which the intermediaries enter all markets.

Proposition 1 (Nash equilibrium) *Suppose that a Nash equilibrium in which* $k_i^j > 0$, $i \in \{1,2\}$, $j \in \{1,2,3\}$, *exists. Then it is given by a solution to the following system of*

equations:

$$k_1^1 = \left[\Lambda_1 \left(\frac{A}{k_2^1}\right)^{1/2} - 1\right] k_2^1, \quad k_1^2 = \left[\Lambda_1 \left(\frac{A}{k_2^2}\right)^{1/2} - 1\right] k_2^2, \quad k_1^3 = \left[\Lambda_1 \left(\frac{A^*}{k_2^3}\right)^{1/2} - 1\right] k_2^3, \tag{8}$$

$$k_2^1 = \left[\Lambda_2 \left(\frac{A}{k_1^1}\right)^{1/2} - 1\right] k_1^1, \quad k_2^2 = \left[\Lambda_2 \left(\frac{A}{k_1^2}\right)^{1/2} - 1\right] k_1^2, \quad k_2^3 = \left[\Lambda_2 \left(\frac{A^*}{k_1^3}\right)^{1/2} - 1\right] k_1^3, \tag{9}$$

$$\Lambda_1 = 2\left[\left(Ak_2^1\right)^{1/2} + \left(Ak_2^2\right)^{1/2} + \left(A^*k_2^3\right)^{1/2}\right]^{-1}, \tag{10}$$

$$\Lambda_2 = 2\left[\left(Ak_1^1\right)^{1/2} + \left(Ak_1^2\right)^{1/2} + \left(A^*k_1^3\right)^{1/2}\right]^{-1}. \tag{11}$$

In this paper, we focus on symmetric Nash equilibria. Substituting $k_1^1 = k_2^1$, $k_1^2 = k_2^2$, and $k_1^3 = k_2^3$ in (8)–(11), we arrive at the following corollary to Proposition 1.

Corollary 1 (Symmetric Nash equilibrium) *In a symmetric Nash equilibrium,*

(i) the resource allocations to markets M_1 and M_2 are

$$k_1^1 = k_2^1 = k_1^2 = k_2^2 = \frac{A}{2A + A^*}$$

and to market M_3 are

$$k_1^3 = k_2^3 = \frac{A^*}{2A + A^*}.$$

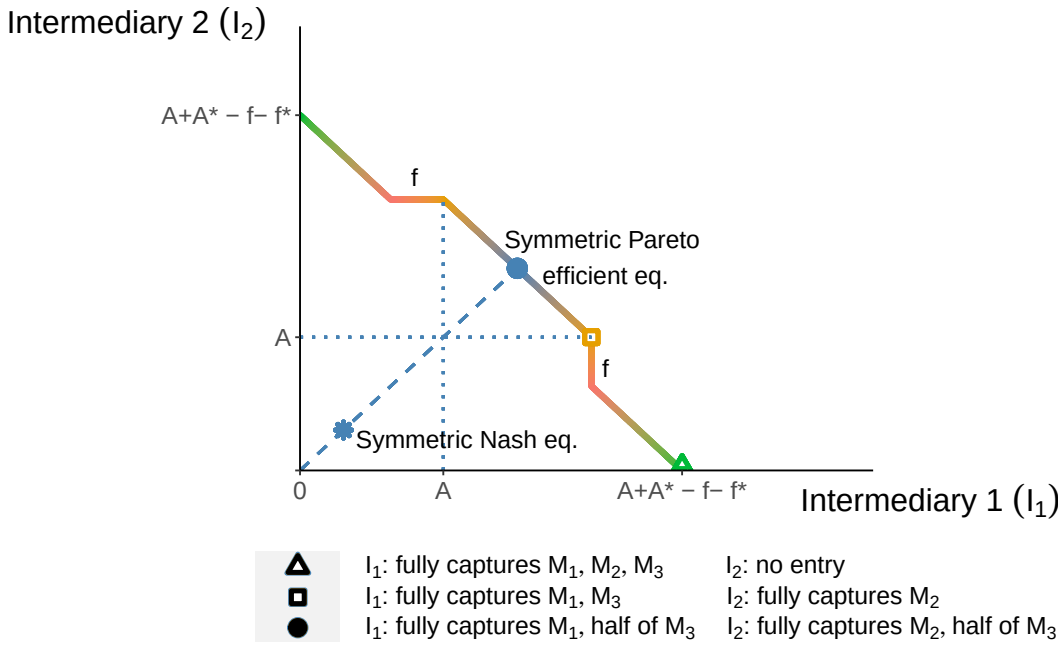
Furthermore, $k_1^1 < k_1^3$ iff $A < A^*$.

(ii) each intermediary's payoff is $A + A^*/2 - f - f^*$. The symmetric Nash equilibrium exists if $A/2 \geq f$ and $A^*/2 \geq f^*$.

In what follows, we maintain the assumption that $A + A^*/2 - f - f^* \geq 0$ and hence the symmetric Nash equilibrium reported in Corollary 1 exists. Since the market shares of the two intermediaries in each market are the same (1), each arbitrage opportunity is split equally between the two of them. Moreover, the larger the arbitrage opportunity, the more resources an arbitrageur is willing to commit to exploiting it. This implication is intuitive. In equilibrium, however, the market shares in all three markets end up being the same because the best response to the opponent who commits more resources to a market is to commit more resources herself.

In the symmetric Nash equilibrium, arbitrageurs enter all three markets, jointly paying the maximum per-period fixed cost. This allocation is clearly not Pareto efficient. For

Figure 1: Pareto frontier



The Pareto frontier of the stage game is depicted by the solid line. The symmetric Nash equilibrium is denoted by the star symbol (✱), while the symmetric Pareto efficient equilibrium by a circle (●).

example, an allocation whereby both arbitrageurs enter their natural markets but not the natural markets of their opponents (while still equally splitting the third market), yields exactly the same gain from harvesting the arbitrage opportunities as in Corollary 1, but at a lower individual fixed cost. To generalize this reasoning, we turn to exploring the set of payoffs that maximize social welfare in a stage game.

Figure 1 depicts the Pareto frontier and the symmetric Nash equilibrium of this game. The Pareto frontier, which is the boundary of the feasible set of payoffs, is depicted by the solid line. To understand its shape, let us start from its rightmost point, denoted by a triangle (△), which corresponds to the highest possible payoff to Intermediary 1 and zero payoff to Intermediary 2. The highest payoff achievable by Intermediary 1 occurs when she enters all three markets and Intermediary 2 enters none. As Intermediary 2 enters its natural market, M_2 (and no other market), the share of the arbitrage opportunity in market M_2, A, captured by Intermediary 1 decreases, while the value captured by Intermediary 2 increases proportionately to its market share of that market. The slanted segment on the Pareto frontier originating from its rightmost point (△) corresponds to this case. When Intermediary 1 exists market M_2 (while remaining in markets M_1 and M_3), the Pareto frontier jumps up by the value of the fixed cost f, because of the fixed cost savings when market 2 is left to the natural arbitrageur in that market, Intermediary 2. This allocation

is denoted by a square (□). The slanted segment right after the jump corresponds to the case in which Intermediary 1 is operating in markets M_1 and M_3 and Intermediary 2 in markets M_2 and M_3. To maximize the total surplus, it is inefficient for both intermediaries to be in present in market M_3 at the same time, both paying a fixed cost of operating in that market. Pareto efficient equilibria that lie on the segment that contains the symmetric Pareto efficient equilibrium therefore involve randomization, in which Intermediary 1 enters market M_3 half of the time and Intermediary 2 the remaining half (denoted by a circle, ●). Since (with the exception of the natural markets) the problem in symmetric, the remaining part of the Pareto frontier is the mirror image of the right part of the frontier.

We are going to focus on the Pareto efficient equilibrium that is symmetric in payoffs, denoted by a circle (●) in Figure 1. The following proposition formally derives and characterizes this equilibrium.

Proposition 2 (Symmetric Pareto efficient equilibrium) *In the Pareto efficient equilibrium that is symmetric in payoffs,*

(i) the two intermediaries agree to allocate no resources to markets M_1 and M_2. That is, $(k_1^1, k_1^2) = (0, 0)$, and $(k_1^1, k_1^2) = (0, 0)$, and hence the intermediaries fully harvest the arbitrage opportunities in their natural markets.

In market M_3, the intermediaries randomize their market entry, and agree that only one intermediary enters at a time. Intermediary 1 plays $k_1^3 \in (0, 1]$ and captures the entire market M_3 with probability $\frac{1}{2}$ and stays away from the market ($k_1^3 = 0$) with probability $\frac{1}{2}$. Intermediary 2 plays the mirror image of this strategy.

(ii) each intermediary's expected payoff is $A + A^*/2 - f^*/2$.

The Pareto frontier contains a set of *cooperative* equilibria in this game. Comparing the intermediaries' payoffs in the symmetric Nash and Pareto efficient equilibria (Corollary 1 and Proposition 2), denoted by a star (✻) and a circle (●) in Figure 1, we see that the latter is clearly higher. The Pareto efficient equilibrium is a cooperative equilibrium that involves explicit agreements between players. It is not achievable in a one-shot game if the intermediaries are playing non-cooperatively. This is because a deviation from playing the Pareto efficient equilibrium is profitable for each intermediary. Indeed, by entering the other intermediary's natural market, the player captures the entire natural market of the opponent, increasing its payoff in the stage game by $A - f$. The unique symmetric equilibrium in a non-cooperative game is the Nash equilibrium.

Achieving a Pareto efficient equilibrium in a one-shot game is possible only if the intermediaries are making explicit agreements with each other. In an infinitely repeated

game, however, a Pareto efficient equilibrium becomes one of the possible equilibria of the stage game, and supporting it does not require any explicit agreements between arbitrageurs.

2.3 Tacit collusion in an infinitely repeated game

We now move from the stage game to the infinitely repeated game. Our objective is to support the symmetric Pareto efficient equilibrium of the stage game, analyzed in the previous subsection. In the infinitely repeated game, the set set of possible strategies expands considerably, and in particular it now includes strategies that depend on the other intermediary's actions in previous periods. This yields a significant enlargement of the set of subgame perfect equilibria, if the intermediaries are sufficiently patient. Randomization each period in the third market is also no longer required: the intermediaries can simply implicitly agree to alternate their entry in the market, i.e., one of them enters in odd periods while the other in even. Infinite horizon is required because if the horizon were finite, the intermediaries have an incentive to deviate from the Pareto efficient equilibrium in the last period, which by backward induction implies that they deviate in all previous periods and therefore the Pareto efficient equilibrium is not subgame perfect and the only symmetric non-cooperative equilibrium of this game is the Nash equilibrium of the stage game, derived above, played every period.

To support the symmetric Pareto efficient equilibrium as a subgame perfect equilibrium, we focus on the following *trigger strategy:* cooperate and play the Pareto efficient equilibrium if the other player cooperates and punish any deviation from it by the opponent by playing the non-cooperative stage game strategies thereafter. Considering again the symmetric case, in the post-deviation phase the intermediaries therefore play the stage-game Nash equilibrium computed in Corollary 1 in every stage game forever. The following proposition makes these statements more precise.

Proposition 3 (Symmetric tacit collusion equilibrium) *Consider the following symmetric trigger strategy:*

Cooperative phase: *Cooperate and*

(i) allocate no resources to markets M_1 and M_2. That is, $(k_{1t}^1, k_{1t}^2) = (0, 0)$, and $(k_{1t}^1, k_{1t}^2) = (0, 0)$, so that the natural intermediaries capture all of arbitrage opportunities in their natural markets;

(ii) alternate entry in market M_3, with Intermediary i entering the market and harvesting the entire arbitrage opportunity A in odd periods ($k_{it}^3 > 0$, t is odd), and Intermediary j, $j \neq i$ in even periods ($k_{jt}^3 > 0$, t is even).

Punishment for deviation: *If the opponent deviated from the above cooperative strategy in the previous period, enter all markets (M_1, M_2, M_3) and play purely competitively forever, which results in a Nash equilibrium in each stage game thereafter (Proposition 1 and Corollary 1).*

If both players play the trigger strategy above and $\delta \geq \frac{A-f-f^*/2}{A}$, the symmetric Pareto optimal (or tacit collusion) equilibrium is a subgame perfect equilibrium of the repeated game.

The result of Proposition 3 is consistent with the Folk Theorem of repeated games (see Fudenberg and Maskin (1986)) that for the time discount factor sufficiently close to 1, a cooperative equilibrium of the stage game can be supported via a purely non-cooperative mechanism in a repeated game. An equilibrium of this type is known as a tacit collusion equilibrium in the industrial organization literature. In contrast to the industrial organization literature, where players charge higher prices in the cooperative stage, in our model cooperation involves staying away from a subset of markets so as to achieve the highest the social welfare (3) by paying the minimum combined fixed cost. This is the novelty of our tacit collusion mechanism.

We stress that the trigger strategy played by the intermediaries and the resulting tacit collusion equilibrium does not involve any explicit agreements every period. All that is required for the implementation of the strategy is that the intermediaries observe that there was market entry by other intermediaries in the previous period. In our application that follows, we observe this information in our non-proprietary dataset.

To better fit our empirical analysis, our model requires a straightforward modification. Instead of there being one market that is natural for neither intermediary, M_3, we need two markets, M_3 and M_4, with an arbitrage opportunity A^* each. The fixed costs of market entry into markets M_3 and M_4 are the same and are equal to f^*. In this game, there is a symmetric tacit collusion equilibrium in which (i) neither intermediary enters markets M_1 and M_2, and therefore the arbitrage opportunity accrues to the natural arbitrageur in that market, and (ii) only one of the intermediaries enters market M_3 and the other intermediary enters market M_4. This equilibrium is an equilibrium in an infinitely repeated game, in which reversion to Nash equilibrium upon a deviation acts as a threat. Proposition 4 in the Appendix formalizes this discussion.

We conclude the section by presenting several testable predictions of our model.

Testable Prediction 1: The natural markets are fully exploited by the incumbent natural arbitrageur(s) and in the data we should not observe entry into those markets.

Testable Prediction 2: If a market is neither arbitrageur's natural market, we should observe only one arbitrageur entering that market.

Testable Prediction 3: If the arbitrageurs are playing one-shot-game Nash equilibrium strategies, they enter all markets. The higher the value of the arbitrage opportunity in a market, the more resources they allocate to that market.

3 Dividend play

In this section, we describe a specific arbitrage strategy, dividend play, for which we can accurately identify trades of arbitrageurs and examine the properties of these trades. The environment in which this strategy is executed closely resembles the setup of our model.

3.1 Dataset

Our options data comes from two main sources: OPRA and OptionMetrics. Transaction-level data from OPRA LiveVol is provided by CBOE. The data ranges from November 4, 2019, to July 31, 2021. It contains all transactions in index, ETF, and equity options on 16 U.S. options exchanges. In our analysis, we exclude index options and focus on single-name options on equities and ETFs. We use daily option price, volume, and open interest data from OptionMetrics, available at a contract level for the period between January 4, 1996, and June 30, 2021. We lag open interest for all the data after November 28, 2000, to have a series of consistent open interest as of the end of day.[8]

To compute trade imbalances, we follow the method described in Muravyev (2016), which classifies trades with prices above (below) the midpoint as “buy" (“sell") trades and trades at midpoint according to the quote rule on the exchange where the trade took place.[9]

Our data-cleaning procedure is as follows. Following the literature, we remove the first 15 and last 10 minutes in the day, canceled trades, trades with nonpositive size or price, and negative spread (the difference between best ask and best bid). We only keep trades for which trade price is above (best bid minus spread) and below (best ask plus spread). We aggregate trades of the same contract with the same quote time, exchange ID, trade price, and trade type into one line.

We complement this data by the daily options trading volume data from the OCC, which is available from the OCC website in open access for the 24 months preceding its

[8]The lag is due to the change in the reporting format of OptionMetrics. This implies that end-of-day open interest is measured after option exercises.

[9]Our results hold if we use the Lee and Ready (1991) algorithm (that is, applying the tick rule to classify trades at midpoint instead of excluding them). The ticker-level imbalances produced by the two methods have a correlation of 94%.

publication date. In our tests, we restrict our sample to the same time period, November 4, 2019, to July 31, 2021, as our OPRA sample.

We also use data from Nasdaq Options Trade Outline (NOTO) and the PHLX Options Trade Outline (PHOTO) End-of-Day files with order classification by the originating counterparty (customers, professional customers, market makers, firms, or broker/dealers). The data ranges from November 1, 2019, until December 31, 2022.

For all stock data, we use CRSP. We obtain dividend history, stock prices and returns, outstanding shares, and rolling monthly volatility of daily returns. We rely on the SecId-PERMNO crosswalk provided by WRDS to link CRSP with OptionMetrics.

3.2 Resurgence of dividend play

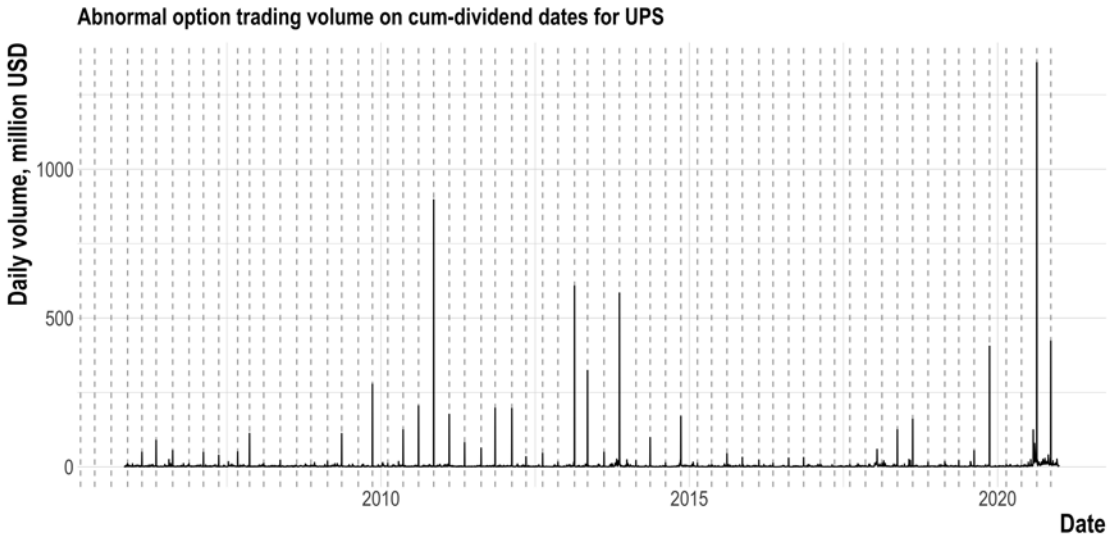
Daily trading volume in options on high-dividend stocks in the U.S. exhibits an intriguing seasonality, illustrated in Figure 2 for the UPS case. The spikes in trading volume apparent from the figure occur every quarter, on the last cum-dividend date, that is, the day before UPS pays a dividend. The average daily traded notional for UPS is $125.3 million on the last cum-dividend dates and only $2.5 million on the remaining dates. This pattern is common for options on high dividend paying stocks. Appendix B.3 presents more examples.

On the last cum-dividend dates, arbitrageurs engage in an arbitrage trade known as the dividend play. This strategy is available only for transactions originating from the floor of the exchange,[10] or, in other words, only to the market participants who must be physically located on the trading floor. The strategy involves establishing long and short options positions that are so large that an operational error may potentially destabilize the market. Concerned about the impact of dividend play trades on the orderly functioning of the market, in 2014 the SEC issued a new rule designed to make the strategy impractical (see footnote 7), which resulted in much lower trading volumes on dividend play dates. However, the recent dramatic increase in options trading by inexperienced retail investors appears to have led to a resurgence of the strategy, despite the barriers created by the SEC rule.

The goal of the dividend play strategy is to take advantage of inattentive investors who fail to exercise their call options on dividend-paying stocks when it is optimal to do so. It is optimal to exercise a call option if the value of exercising it on the last cum-dividend date and collecting a dividend exceeds the value of the call the next day, when the stock goes ex-dividend. Computing option values involves an application of the Black-Scholes-Merton formula or a more sophisticated option pricing method, which is typically difficult for novice

[10] In fact, dividend play could be organized off the exchange floor, but it would then not qualify for transaction fee caps. In our data, most abnormal volume on cum-dividend dates goes through floor trades on two exchanges, PHLX and BOX, as we discuss below.

Figure 2: Abnormal trading volume on last cum-dividend dates for UPS



This figure plots daily trading volume for all call option contracts on UPS, in millions of U.S. dollars, as reported in OptionMetrics. The dashed lines indicate the last cum-dividend dates.

retail investors. Alternatively, some retail investors may be unaware of the possibility of early exercise or are simply inattentive.[11] Since a fraction of in-the-money call options remains suboptimally unexercised, the writers of these options would not be asked to deliver the stock and would therefore receive a windfall capital gain on their position. It is a zero-sum game.

If all in-the-money call option contracts on a stock have been exercised on or before the last cum-dividend date, *all* holders of short positions in the same contracts receive a request to deliver the stock. If some contracts are left unexercised, however, the OCC randomly "assigns" short positions that must deliver the stock. The unassigned holders simply hold on to their options and profit from a capital gain. Arbitrageurs can divert this capital gain to themselves by simultaneously buying and selling a large number of in-the-money call options on the same underlying.[12] They exercise all long positions and deliver on all assigned short positions. Since some fraction of the options remains unassigned (owing to suboptimal exercise by investors), arbitrageurs then capture dividends on their net long stock positions while staying fully hedged. Usually, two arbitrageurs agree on a dividend play trade in advance and serve as a counterparty to each other on their arbitrage positions.

[11] There might be other reasons why investors do not exercise, such as costs of unwinding more complex strategies. Hao, Kalay, and Mayhew (2009) show that dividend play profits outweigh such costs in most cases.

[12] The current SEC rule prohibits simultaneous buying and selling of the same contract. See footnote 7.

Table 1: Dividend play: An example

	OI_{t-1}	New positions(t)	Available for ex.	No. exercised	Prob. non-assign. orig. option writer	Prob. non-assign. arbitrageur	Gain per share	Expected gain orig. option writer	Expected gain arbitrageur
	(A)	(B)	(C)	(D)	(E)	(F)	(G)	(E*G*100)	(F*G*100)
Case 1. Optimal exercise									
Customer	1	0	1	1	0		5	0	
Case 2. Suboptimal exercise									
Case 2.1. Without dividend play									
Customer	1	0	1	0	1		5	**500**	
Case 2.2. With dividend play									
Customer	1	0	1	0	1/101		5	**5**	
Arbitrageur	0	100	100	100		100/101	5		**495**
Total	1	100	101	100					

This table illustrates the dividend play strategy. Date t refers to the last cum-dividend date, and OI_t stands for the open interest on date t.

Table 1 illustrates the mechanics of the dividend play strategy by means of an example. Suppose there is 1 call option contract outstanding, and it is optimal to exercise it.[13] Case 1 corresponds to the case when the option is exercised, the holder of the short position is assigned to deliver the underlying, and there is no profit for the dividend play strategy to harvest. Case 2 describes what happens if the contract is left unexercised. Without arbitrageur involvement, the short position in the contract does not get assigned, and the option writer receives a windfall gain of \$500 for sure.[14] Now consider the entry of an arbitrageur. The arbitrageur attempts to pocket most of the potentially harvestable profit of \$500. To do so, the arbitrageur buys and simultaneously sells 100 contracts and exercises all of their long positions. The probability of assignment increases, but, because of the OCC's random assignment, with probability 100/101, the arbitrageur holds the short position that does not get assigned and hence yields a gain. For the original option writer, this probability is now only 1/101. Hence, the expected gain of the arbitrageur is \$495 out of the total gain of \$500 and that of the original option writer drops to \$5. A dividend play strategy, therefore, dilutes the share of the gain that accrues to the original option writer.

In what follows, we detect dividend play activity at a contract level in the full sample and characterize its importance relative to the overall trading volume on the last cum-dividend dates.

[13] Appendix B.4 provides another example, in which there are multiple contracts outstanding, some of which are exercised optimally and some are not.

[14] We assume that, as in the data, each option contract is for 100 shares of underlying.

3.3 Arbitrageur activity in dividend play strategy

We first present our novel measure of arbitrageur activity in the dividend play strategy. Through fee caps, exchanges incentivize dividend play strategies to originate from the physical floor. To construct our measure of dividend play activity, we therefore exploit OPRA new detailed trade type flags to isolate option transactions that are executed on the floor. The trade types that cover most of the dividend play transactions are SLFT and MLFT, which are single-leg and multi-leg floor trades, respectively. (See Appendix B.2 for a more detailed description.) Other floor trade types, used infrequently in our sample are MLCT, MSFL, SLCN, TLFT, and TLFT. We analyze this measure in detail below and show that it accurately captures dividend play trades in our sample. To our knowledge, this is the most precise measure of arbitrageur activity in the dividend play strategy in the literature, which typically uses trading volume on the last cum-dividend dates in excess of the past average volume.[15]

In our data, we see bursts of simultaneous buy and sell activity in neighboring-strike call option contracts, executed usually within several seconds, all coming from the floor. In an effort to reduce operationally risky dividend play trades, in 2014 the SEC began to forbid arbitrageurs from simultaneously buying and selling the same contract. Market participants have adjusted their trading strategies, and they now simultaneously buy and sell neighboring contracts, which ultimately achieves the same objective. The trades are typically prearranged by pairs of arbitrageurs. We see no similar bursts of simultaneous buy and sell activity in call option contracts in any other OPRA trade types, which assures us that our measure accurately captures arbitrageur activity in the dividend play strategy.

Table 2 presents descriptive statistics of trading activity on the last cum-dividend compared to any other dates for dividend-paying stocks and ETFs. We see a considerable difference in floor trading volume and volume of large trades on the last cum-dividend dates relative to other dates. Moreover, on the last cum-dividend dates we see colossal spikes in volume on two exchanges that cap fees for the dividend play strategy: PHLX and BOX. Breaking the trades by moneyness, we see that the primary increase in volume comes from trading deep-in-the-money calls (which are more likely to be optimal to exercise). This pattern is a signature of the dividend play strategy. The utter size of the dividend play positions is astonishing. Clearly, market participants have been able to circumvent the regulators' attempt to curtail this strategy, and profit-driven incentives of arbitrageurs have pushed the trading volume in dividend play to new highs.

[15] Strictly speaking, we should define our measure of arbitrageur activity in dividend play as *abnormal* floor volume on the last cum-dividend dates, but our measure is simpler and our results stay virtually the same if we use the abnormal floor volume measure instead.

Table 2: Characteristics of activity on last cum-dividend dates

	Average ticker dollar volume ($ million) on		Total market dollar volume share (%) on	
	last cum-dividend date	any other date	last cum-dividend date	any other date
	(1)	(2)	(3)	(4)
Panel A. Option type				
Call	24.9	1.5	92.8	54.3
Put	2.1	1.4	7.2	45.7
Panel B. Moneyness				
In the money	26.8	0.8	80.6	18.3
At the money	4.8	2.0	17.6	70.2
Out of the money	0.6	0.4	1.8	11.4
Panel C. Trade size				
Small	1.5	0.7	5.6	27.5
Large	31.0	2.8	94.4	72.5
Panel D. Floor trade				
Yes	48.2	0.9	76.9	6.5
No	6.1	2.4	23.1	93.5
Panel E. Exchange				
PHLX or BOX	24.6	0.5	79.7	15.0
Any other	5.5	2.3	20.3	85.0

This table compares option trading activity for dividend-paying tickers (2,153 stocks and ETFs) on the last cum-dividend date with any other date. The average volume in columns (1) and (2) is computed at ticker-day level, and the volume share in columns (3) and (4) is for the entire market. In Panel B, we define "in the money" as $(Midpoint\,Price - Strike)/Strike > 0.1$ for call options and $(Midpoint\,Price - Strike)/Strike < -0.1$ for put options. "At the money" are contracts for which this value is between -0.1 and 0.1, and "out of the money" are all other contracts. In Panel C, we define trade as 'small' if the trade size is at or below 10 contracts. In Panel D, we define floor trades as trades with SLFT and MLFT OPRA trade types.

3.4 Dividend play profits

We now derive exploitable profits from the dividend play strategy. Some of these profits come from an increase in the open interest, some from investors' failure to exercise, and some from the value of early exercise of each contract. With an inflow of inexperienced investors in the options market, we expect the first two components to increase. We therefore find it useful to decompose the exploitable profit from a contract into three parts: (i) open interest, (ii) fraction unexercised, and (iii) early exercise value.

The exploitable dividend play profit on all the interest for each contract is defined as

$$\pi_t = OI_{t-1} \times f_t \times EEV_t, \tag{12}$$

where $t-1$ is the day before the last cum-dividend date, OI_{t-1} denotes open interest on that date (measured after all trades, exercises, and assignments on that date), $f_t \equiv OI_t/OI_{t-1}$ is the fraction unexercised, and EEV_t is the early exercise value, computed below. Note

that the fraction unexercised reflects the fraction of open interest in an option contract that remains outstanding after the last cum-dividend date (after all trades, exercises, and assignments on that date). Both EEV_t and f_t are estimated quantities. Open interest as of the day before the last cum-dividend day (OI_{t-1}) and fraction not exercised (f_t) are available from OptionMetrics. In rational and frictionless markets, we expect that $f_t = 0$ if $EEV > 0$.

The early exercise value is model-based, and we rely on the Black-Scholes-Merton option pricing formula to compute it.[16] Denote the expected ex-dividend price of an option by c_{ex}, its strike by K, and the current (cum-dividend) underlying stock price by S. The expected option ex-dividend price represents the expected time value of the option. *Early exercise value (EEV)* is therefore the difference between the current stock price, strike, and this expected time value of the option: $S - K - c_{ex}$.[17] The details of the computation of c_{ex} are in Appendix B.5.

In the following analyses, we restrict our sample to call option contracts that are optimal to exercise on cum-dates and refer to it as the *dividend play sample*. Further details related to its construction are provided in Appendix B.6, and Table 10 in the Appendix presents the descriptive statistics for our dividend play sample.

4 Selective entry by arbitrageurs

In this section, we show that arbitrageurs engaging in the dividend play strategy seemingly leave money on the table by failing to capture arbitrage profits in some call option contracts. We explore the determinants of this puzzling behavior and show that it is consistent with our model.

4.1 Case study

November 11, 2020, was the last cum-dividend date for UPS, a high-dividend paying stock, and a number of call options on UPS were deeply in-the-money and optimal to exercise on that day. Table 3 focuses on a pair of such contracts, both expiring on November 20, 2020. We first compare the trading volume in the contracts on November 11, 2020. Notice that the trading volume in Contract 2 exceeds that in Contract 1 by two orders of magnitude.

[16] To ensure our results are robust to the choice of the underlying pricing model, we considered the sample of broad-index ETFs and computed their corresponding option prices with the Merton and Bates models, following Bakshi, Cao, and Chen (1997) and Cosma, Galluccio, Pederzoli, and Scaillet (2020). Options on these ETFs represent over 10% of contracts in our dividend play sample and 55% of potential dividend play profits. All our results hold in that sample and are available upon request.

[17] Note that this definition is from Pool, Stoll, and Whaley (2008), and it is equivalent to the definition in Hao, Kalay, and Mayhew (2009). The latter uses dividend instead: $Dividend - c_{ex} + S_{ex} - K$.

Table 3: Case study of arbitrageur activity: Two UPS call options on the last cum-dividend date

	Strike	EEV	OI (t-1)	Moneyness	Spread	Fraction unexercised	Cum-date volume	Floor share
Contract 1	160	0.29	1,945	3.15	0.045	0.76	45	0.000
Contract 2	155	0.43	2,487	4.62	0.039	0.47	3,255	0.998

Notice also that Contract 2 has most of the orders from the trading floor on that day, while Contract 1 has zero. We also see characteristic bursts of floor orders in the transaction-level data for Contract 2. This means that arbitrageurs engaging in a dividend play trade exploited Contract 2 but not Contract 1.

Why did the arbitrageurs leave money on the table in Contract 1? The contract had a high EEV and a large fraction unexercised. Using equation (12) to compute the arbitrageur's forgone profits from not participating in Contract 1, we arrive at $1,945 \times 0.76 \times 0.29 \times 100 \approx \$42,900$, a significant sum.[18]

Trading costs do not explain the market participants' reluctance to trade Contract 1. First, exchanges offer daily fee caps for the dividend play strategy, and so if market makers (or other arbitrageurs) entered in Contract 2, they should have also entered in Contract 1. Second, contract bid-ask spreads in Table 4 are very similar. In the regression framework that follows, we further control for the option contract's liquidity and show that trading costs do not explain why arbitrageurs forgo profitable opportunities.

It is puzzling that arbitrageurs fully exploited the arbitrage opportunity in Contract 2 but not in Contract 1. In the following section, we show that this is the general pattern in our sample. The unexploited profit in Contract 1 accrued to the writer of this contract, who could be a market maker or perhaps a retail investor. The latter is less likely because retail brokerages take an automated action to close short positions that have dividend risk on behalf of their clients.[19] Appendix B.9 presents an excerpt from Robinhood's Terms and Conditions to provide an example of such automated action. It is therefore more likely that the writer of the contract who received the windfall gain was a market maker. The market maker who is a writer of the contract of course has no incentive to engage in a dividend play strategy in this contract because this would mean sacrificing his own profit. But it is puzzling that other market makers or arbitrageurs would not wish to exploit Contract 1 and reap arbitrage profits.

[18] Each option contract in our sample is for 100 shares of the underlying stock or ETF.

[19] Since each option contract is for delivery of 100 shares of the underlying, for small retail investors the cash outlay needed for purchases of the underlying stock and delivering it could be quite significant. A brokerage would therefore close a short position if there are not enough funds in the account to buy and deliver the underlying.

4.2 Lack of entry by arbitrageurs

Table 4 expands our case study and reports forgone profits by ticker for the top 40 underlying stocks and ETFs sorted by the total size of forgone profits in our sample. We aggregate our data to the ticker level and report the number of profitable individual contracts per ticker. The total amount of harvested profits in the top 40 tickers in our sample is around \$55 million, whereas the total amount of forgone profit stands at \$67 million. For a virtually riskless arbitrage strategy, the amount of money left on the table is striking.

In the full sample, the values of harvested and forgone profits are \$84 million and \$81 million, respectively; hence, the total potential arbitrageur profit from the dividend play strategy is around \$170 million. To put the total amount of profits from this strategy into perspective, it is useful to compare these numbers to arbitrageurs' profits from other strategies that exploit retail investor mistakes. de Silva, Smith, and So (2022) study retail investor trading in options around earnings announcements and find that this trading amounts to a wealth transfer of \$360 million from retail investors to market makers in their sample that is seven times longer than ours, running from 2010 to 2021.

Table 4 does not reveal any particular pattern in harvested and forgone profits: There is a large variation in arbitrageur participation across and within tickers. In what follows, we examine possible explanations why some contracts attract entry by arbitrageurs and some do not.

To examine drivers of arbitrage activity, we start by contrasting potential and harvested profits from the dividend play strategy on the last cum-dividend dates in our sample. Figure 3 presents potential profit of the dividend play strategy in all outstanding contracts, computed using equation (12), and profits harvested by floor traders. It emerges from Panel (a), that a large fraction of potential profit, about 50%, remains unharvested. If we restrict the sample, however, to the contracts with non-zero floor trading volume—that is, contracts in which we detect dividend play activity—most of the potential profit resulting from the failure of investors to exercise their options on cum-dividend dates is harvested. In other words, arbitrageurs selectively exploit profitable contracts, capturing almost 100% of exploitable gains (more precicely, 90.5%), but forgo arbitrage profits in contracts they do not enter.

We have now established the following two empirical facts regarding arbitrageur activity on dividend play dates in profitable contracts.

Fact 1: *Only 53% of profitable dividend play arbitrage opportunities (contracts) attract arbitrageur entry. About 50% of potential profits appear to be left of the table.*

Table 4: Dividend play profits by ticker

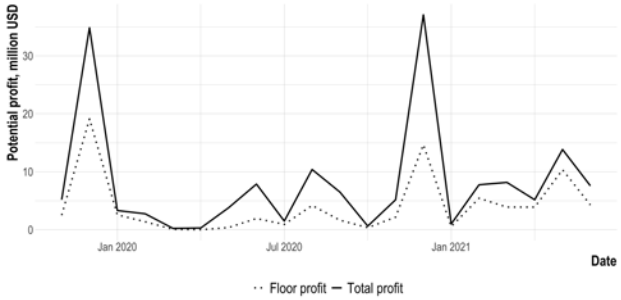
Ticker	Profit, USD		No. contracts			Traded volume (contracts)
	Harvested	Forgone	Fully harvested	Partly harvested	Forgone	
Ticker	(1)	(2)	(3)	(4)	(5)	(6)
SPY	2,807,822.0	23,664,548.0	270	48	1357	1,071,808
AAPL	3,834,985.0	9,900,164.0	343	210	393	1,344,226
EEM	11,200,000.0	4,761,947.0	247	8	51	6,056,682
IWM	1,485,149.0	3,550,559.0	55	4	218	521,360
XLE	1,818,708.0	2,840,936.0	207	12	112	522,140
VALE	1,919,459.0	2,741,013.0	80	7	26	1,965,400
QQQ	109,689.3	2,052,260.0	24	2	274	29,950
EFA	4,546,631.0	1,952,792.0	202	8	39	2,573,722
EWZ	3,448,678.0	1,330,618.0	157	2	52	1,596,842
KO	448,687.9	1,095,654.0	82	29	103	329,065
HYG	38,107.7	912,239.3	15	4	86	69,660
XOM	8,376,230.0	754,073.5	277	157	192	2,124,907
SAN	-	753,507.4	0	0	12	-
HD	984,543.6	714,045.3	143	45	190	305,970
COST	1,648.3	676,289.4	12	6	69	1,672
QCOM	861,147.2	628,717.5	132	51	144	559,989
XLF	482,727.7	620,491.9	64	5	63	351,070
IBM	843,184.9	615,173.6	169	183	151	493,061
DIA	166,536.8	583,968.0	89	9	229	27,054
BHP	175,917.7	553,663.6	45	6	14	60,570
ET	674,788.8	495,248.8	60	38	53	613,655
GOLD	314,559.5	456,460.8	32	4	38	68,864
VIAC	1,774,224.0	425,584.1	100	1	65	467,210
XLI	10,878.1	401,835.9	13	1	26	17,423
RIO	27,423.2	378,598.4	18	9	14	73,662
XLP	131,354.4	370,952.7	19	1	49	22,370
T	2,706,384.0	368,861.1	174	54	82	2,538,321
JPM	927,727.1	346,714.9	107	44	106	1,149,259
CVX	632,718.6	327,097.1	298	122	178	524,514
GILD	362,869.9	314,560.7	83	36	92	299,828
FXI	885,153.6	312,136.5	103	6	45	1,371,416
MRO	-	307,658.0	0	0	26	-
NVDA	-	306,058.0	0	0	91	-
BP	378,855.8	296,584.8	131	61	106	310,435
DIS	873,927.8	269,175.9	51	17	14	514,155
MSFT	562,511.9	251,602.3	201	35	55	518,632
MPC	595,727.9	246,978.6	127	44	124	477,051
TGT	102,606.3	244,211.7	120	44	115	151,552
AVGO	748,359.4	235,800.2	150	86	85	257,670
DOW	144,568.3	228,715.5	49	17	105	99,844
Total	**55,404,491.7**	**67,287,496.5**	**4,449**	**1,416**	**5,244**	**29,481,009**

This table reports the top 40 tickers in terms of dividend play profits forgone by floor traders in our sample. Values are aggregated across all contracts within a ticker in 11/2019–06/2021. Total dividend play potential profits are computed as in equation (12). To compute Harvested profits, we multiply the total profits by the floor volume share on the last cum-dividend date and attribute the residual to Forgone profits. No. of Fully harvested contracts in column (3) is the number of contracts with floor share above 90%, and in column (5) – with zero floor share.[a] Traded volume in column (6) is the total floor trading volume in all contracts.

[a]The average floor share is over 99% in Fully harvested contracts and 69% in Partly harvested contracts.

Figure 3: Total and floor trader profit from dividend play strategy

(a) All profitable contracts



(b) Contracts with non-zero floor share



This figure illustrates the implied share of potential dividend play profits captured by arbitrageurs on the trading floor. The solid plot depicts the potential profit from the dividend play strategy, and the dashed plot depicts the profit harvested by floor traders (arbitrageurs).

Fact 2: *If an option contract attracts arbitrageur entry, arbitrageurs extract close to 100% of dividend play profits in that contract.*

The selective entry of arbitrageurs in profitable dividend play opportunities is puzzling. In what follows, we try to understand the features on the contracts into which market participants are likely to enter.

4.3 Retail investor presence and arbitrageur activity

The resurgence of dividend play in our sample is likely to be linked to the pandemic-related retail trading boom, with retail investors making more early exercise mistakes than other market participants and hence boosting arbitrageur profits in the dividend play strategy. Moreover, the sheer size of the increase in open interest due to retail investor inflow is likely to have contributed to dividend play profits (see equation (12)). BPS document that over 60% of contract trading volume in the U.S. options market is now retail. We will formally establish below that contracts with higher retail participation are more profitable to exploit for dividend play arbitrageurs. Yet, despite their higher profitability, some of these contracts do not attract arbitrageur entry. As we show below, retail investor buying activity over the past week is a strong predictor of this selective entry.

We use two measures of retail trading activity in the options market, constructed from the transaction-level OPRA data. Our first measure, often used in the industry, is the daily volume share of small trades (up to 10 contracts), *Small Share*. One could compute it as a frequency share and as a trading volume share. We adopt the latter definition, as it would be more relevant for assessing the influence of retail traders on asset prices. Our second measure

is from BPS (Ernst and Spatt (2022) and Hendershott, Khan, and Riordan (2022) propose the same methodology.) They take advantage of the new trade type codes, introduced by OPRA on November 4, 2019, to construct a new measure based on transactions that are likely to be intermediated by wholesalers.[20] Wholesalers, such as Citadel, Susquehanna, and Wolverine, intermediate the overwhelming majority of retail order flow in the U.S. and they often use the so-called price improvement mechanisms to execute the orders (see BPS for more details). We borrow their acronym SLIM to denote retail trades and use the volume share of these trades, *SLIM Share*, as a measure of retail activity.

How do retail trading trends relate to the last cum-dividend date exercise rates? To answer this question, we run the following regression:

$$Y_{c,t} = \beta_1 \times share_{c,t}^{SLIM} + \beta_2 \times share_{c,t}^{small} + \boldsymbol{\gamma}' X_{c,t} + \alpha_{i,t} + \varepsilon_{c,t} \tag{13}$$

where, for each contract c on the last cum-date t, we consider two dependent variables, $Y_{c,t}$: (a) fraction of open interest not exercised by ex-dividend date, and (b) potential profits from dividend play strategy as defined in Equation (12). Our vector of controls $X_{c,t}$ includes the following contract-level variables: log OI, EEV, log dollar trading volume, relative spread, implied volatility, moneyness, and days to expiration.[21] Our specification also includes the ticker by date fixed effects $\alpha_{i,t}$. Our measures of retail investor trading are computed over one trading week before the last cum-date, because more recent data on retail trading is likely to be also more informative about the arbitrageur (retail counterparty) positions in the underlying contracts. Furthermore, we get identical results by considering trade imbalance over the three-day period, immediately preceding the cum-dividend date.

Panel A of Table 5 reports the results of the regression in (13), with the fraction of open interest unexercised as the outcome variable. We find that there is a strong positive relationship between retail investor trading and the fraction of options that were suboptimally left unexercised on the last cum-dividend day. While we measure retail investor trading in two different ways, both variables prove to be strong predictors of failures to exercise the option. A one standard deviation increase in the share of SLIM or small trades in the contract in the week preceding the last cum-date increases the fraction unexercised by approximately 1 percentage point, depending on the specification. This result is robust, and the magnitudes of the coefficients of interest do not change much as we relax the specification of fixed effects and use ticker-level controls instead (see columns (1)–(3) and (5)).

[20] Specifically, we use the same OPRA type SLAN, which stands for single-leg non-ISO price improvement auctions. See Appendix B.2 for a description.

[21] Since log OI and EEV are components of potential dividend play profits, we do not include them in the specification in Panel B below.

Table 5: Dividend play profits and additional measures of retail investor popularity

	Dividend play profitability feature				
	(1)	(2)	(3)	(4)	(5)
Panel A. Fraction of OI not exercised, %					
SLIM Share	4.939***		4.835***	3.773***	5.085***
	(5.14)		(5.03)	(3.76)	(5.27)
Small Share		4.155***	3.999***		
		(3.67)	(3.53)		
D(SLIM buy imbalance)				2.040***	
				(5.43)	
Observations	33,562	33,562	33,562	33,562	33,562
Adjusted R-squared	0.209	0.209	0.209	0.210	0.186
Panel B. Potential profits, log U.S. dollar					
SLIM Share	1.624***		1.616***	1.110***	1.707***
	(11.10)		(11.18)	(7.81)	(11.40)
Small Share		0.428*	0.389*		
		(1.80)	(1.67)		
D(SLIM buy imbalance)				0.862***	
				(16.99)	
Observations	33,562	33,562	33,562	33,562	33,562
Adjusted R-squared	0.256	0.250	0.256	0.265	0.234
FE	Ticker*Date	Ticker*Date	Ticker*Date	Ticker*Date	Ticker and Date
Contract controls	Y	Y	Y	Y	Y
Ticker controls	N	N	N	N	Y

This table reports estimates of (13) and (14) in our dividend play sample. SLIM Share and Small Share are the contract-level volume shares of SLIM and small trades, respectively, averaged over one trading week before the last cum-dividend date. D(SLIM buy imbalance) is the dummy variable of the SLIM buy imbalance, which is equal to one, if there was on average a positive buy imbalance in SLIM trading volume over the previous trading week, and zero otherwise. In Panel B, contract controls include log dollar trading volume, relative spread, IV, moneyness, and days to expiration. In Panel A, they additionally include log OI and EEV. Ticker controls include underlying price, underlying volatility, underlying relative bid-ask spread, and underlying market cap. Standard errors are clustered by ticker and date. Robust t-statistics are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

So far our results demonstrated the link between the overall retail presence in a given contract and dividend play profit features. Their connection becomes even more apparent in one considers the direction of retail trading. To do this, we create a dummy variable of the SLIM buy imbalance, which is equal to one, if there was on average a positive buy imbalance in SLIM trading volume over the previous trading week, and zero otherwise. We then run

the following regression:

$$Y_{c,t} = \beta_1 \times share^{SLIM}_{c,t} + \beta_2 \times D(SLIM\ buy\ imbalance)_{c,t} + \boldsymbol{\gamma}' X_{c,t} + \alpha_{i,t} + \varepsilon_{c,t} \qquad (14)$$

where, as in equation (13), for each contract c on the last cum-date t, we consider two dependent variables, $Y_{c,t}$: (a) fraction of open interest not exercised by ex-dividend date, and (b) potential profits from dividend play strategy as defined in Equation (12). We also include the same set of controls $X_{c,t}$ and fixed effects.

Our findings are presented in Table 5, column (4). Both SLIM Share and SLIM buy imbalance have a significant and positive effect on the fraction of options left unexercised. This further supports our hypothesis that it is indeed retail investors who recently *bought* call options, who are unlikely to optimally exercise them on the cum dividend date. In particular, we find that if there was a positive buy imbalance in SLIM over the previous trading week, even conditional on the overall level of retail trading, the fraction of open interest left unexercised is 2 percentage points higher.

Panel B of Table 5 considers the regression in (13), but with potential profits as the outcome variable. It reveals that the coefficients on all measures of retail activity are positive and significant: A one standard deviation increase in SLIM Share corresponds to around \$4,200 higher profit *per contract*. In other words, the higher the retail activity in a contract in a week preceding the last cum-dividend date, the more profitable it is for arbitrageurs to engage in a dividend play in the contract. Higher profits come from both (a) higher fraction unexercised (documented in Panel A) and (b) higher open interest in the contracts popular with retail investors. In Panel B, column (4), we further extend our specification to include SLIM buy imbalance. Similar to Panel A, we find that a positive buy imbalance of retail trading increases the potential profits from the dividend play strategy.

We are now ready to test whether arbitrageurs or aware of and exploit retail investor mistakes in their dividend play arbitrage. We estimate the following regression in the sample of contracts that should optimally be exercised on cum-date:

$$share^{floor}_{c,t} = \beta_1 \times share^{SLIM}_{c,t} + \beta_2 \times D(SLIM\ buy\ imbalance)_{c,t} + \boldsymbol{\gamma}' X_{c,t} + \alpha_{i,t} + \varepsilon_{c,t}, \quad (15)$$

where the main regressors are $share^{SLIM}_{c,t}$, the average dollar volume share in SLIM trades over one trading week before the last cum-dividend date t, and the dummy variable indicating a positive average buy imbalance in SLIM trades over the same period. $X_{c,t}$ includes the following contract-level variables: log OI, EEV, log dollar trading volume, relative spread, implied volatility, moneyness, and days to expiration.[22] Our specification also includes the

[22]Since log OI and EEV are components of potential dividend play profits, we do not include them in the

Table 6: Arbitrageur activity and retail investor popularity

	Floor trading measure			
	(1)	(2)	(3)	(4)
Panel A: Floor trading share on last cum-date				
SLIM Share	0.029*		0.042***	0.106***
	(1.80)		(2.66)	(3.62)
D(SLIM buy imbalance)		-0.019***	-0.023***	-0.030***
		(-3.80)	(-4.90)	(-3.48)
Observations	33,562	33,562	33,562	10,650
Adjusted R-squared	0.419	0.419	0.419	0.399
Panel B: D(floor share > 0)				
SLIM Share	0.027*		0.039**	0.091***
	(1.65)		(2.47)	(2.89)
D(SLIM buy imbalance)		-0.018***	-0.022***	-0.032***
		(-3.58)	(-4.59)	(-3.57)
Observations	33,562	33,562	33,562	10,650
Adjusted R-squared	0.412	0.412	0.412	0.381
Panel C: Floor trading volume, log				
SLIM Share	0.401***		0.524***	1.132***
	(2.85)		(3.85)	(3.96)
D(SLIM buy imbalance)		-0.163***	-0.215***	-0.351***
		(-3.63)	(-5.16)	(-4.03)
Observations	33,562	33,562	33,562	10,650
Adjusted R-squared	0.502	0.502	0.502	0.423
Sample	All	All	All	Top profit tercile

This table reports estimates of (15) in our dividend play sample. Floor trading share on cum-date is the contract-level volume share of trades executed on the traded floor in the total traded volume on the last cum-dividend date. SLIM Share and is the contract-level volume share of SLIM trades, averaged over one trading week before the last cum-dividend date. D(SLIM buy imbalance) = 1 if there was a buy imbalance in SLIM volume over the previous trading week, on average, and 0 otherwise. All regressions include ticker-by-date fixed effects and contract controls (log OI, EEV, log dollar trading volume, relative spread, IV, moneyness, and days to expiration). Top profit tercile includes contracts in the top tercile of total potential profits. t-statistics are based on standard errors clustered by ticker and date (in parentheses). *** p<0.01, ** p<0.05, * p<0.1.

ticker by date fixed effects $\alpha_{i,t}$. Table 6 reports the results of estimating the regression.

specification in Panel B below.

Table 6 reveals that arbitrageur activity, as measured by floor trading share, volume, or entry, is positively related to SLIM Share, the measure of retail investor trading over the preceding week (Panels A–C). One would expect this behavior from profit-maximizing arbitrageurs. However, strikingly, arbitrageur activity in a contract is *negatively* related to the positive SLIM buy imbalance, and this result is highly statistically significant across all specifications in Table 6, irrespective of the measure of arbitrageur activity (Panels A–C).

Why would arbitrageurs turn away from contracts that are more likely to remain unexercised by inattentive retail investors (as we have established in Table 5)? Some dividend play arbitrageurs could be options market makers. These market participants have access to the exchanges' trading floors and they may find the strategy attractive. In their role as market makers, they intermediate retail order flow and in particular write options that retail investors buy. As option writers, these market makers are then set to receive a windfall capital gain at zero cost if retail investors do not exercise these option contracts optimally. If the contracts remain in their inventory, they represent the "natural markets" for arbitrage activities of these intermediaries, as defined in our model. (We elaborate on this in Section 4.5 and present evidence in Figure 5 in the Appendix.) The negative sign on SLIM buy imbalance in Table 6 is consistent with our model's implications. The test in Panel B is particularly aligned with the model, indicating that arbitrageurs are less likely to enter natural markets or other intermediaries.

The specification in column (4) of Table 6 restricts the sample to the top tercile of most profitable contracts (by total profits) for dividend play arbitrageurs. Standard theories of limits-to-arbitrage imply that we should see more arbitrageur activity in those contracts. This is not what our estimation reveals. While the coefficient on SLIM Share does indeed go up relative to the full-sample regressions, we still see a strong negative effect of SLIM buy imbalance. The latter is inconsistent with standard limits-to-arbitrage theories but consistent with the hypothesis that arbitrageurs selectively enter different contracts, even the most profitable ones, avoiding those that represent natural markets for market makers.

To confirm the robustness of our results, we also pursued an alternative empirical strategy, based on propensity score matching. Matching is a natural strategy in our setup because the set of characteristics on which one should match options to keep the expected profitability – the main driver of arbitrageur entry in the limits-to-arbitrage literature – constant is well understood. In particular, we match contracts based on open interest, early exercise value, and moneyness (we also considered relative spread and underlying price). Furthermore, such empirical exercise brings us closer to the case study presented in Table 3. In Appendix B.10, we show that across the whole propensity score spectrum, there exist contracts with both zero and positive floor volume. As shown before, in the latter case, floor

traders represent almost 100% of trading, so they seem to exhaust most of the potential profits. This result again suggests that profitability characteristics do not predict entry well.

We further highlight arbitrageurs' selective entry by exploring their participation rate in the most profitable contracts. We find that arbitrageurs do not enter 42% of contracts in the top EEV tercile, 47% in the top moneyness tercile, and 36% in the top total potential profit tercile. The average across these three measures of the attractiveness of an arbitrage opportunity is 42%. Therefore, in the dividend play arbitrage there is no evidence of the expected pecking order, in which arbitrageurs exploit the most profitable contracts first, followed by the second most profitable, etc.

In sum, the analysis in this subsection establishes the link between retail investor presence and dividend play arbitrage and uncovers the following novel fact about arbitrage activity.

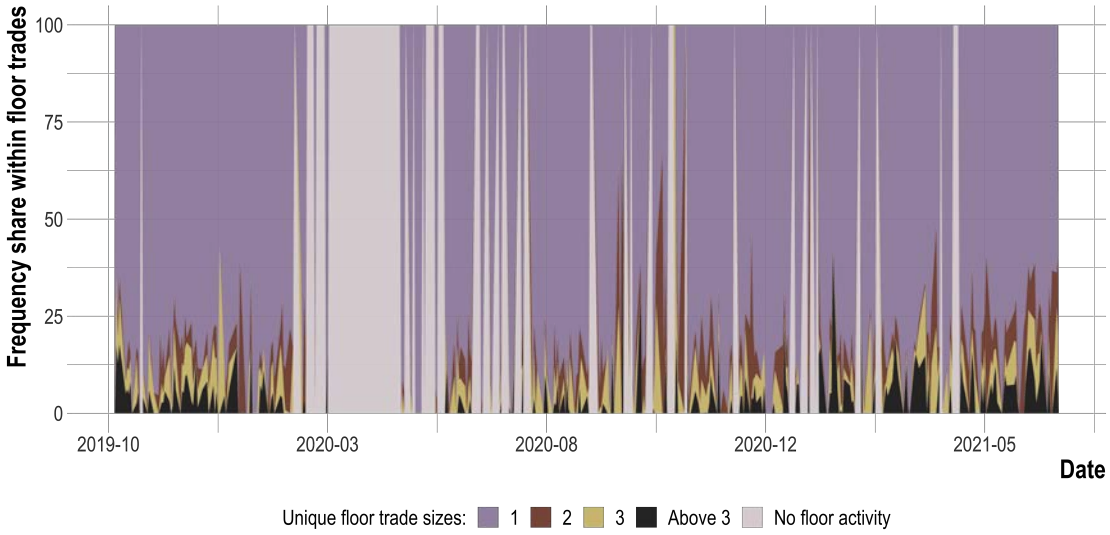
Fact 3: *More profitable arbitrage opportunities do not necessarily attract more arbitrage activity.*

4.4 How many arbitrageurs engage in dividend play?

Earlier in this section, we have documented the facts that arbitrageurs do not enter some profitable contracts, and if they do enter, they exploit all available profits in a contract. In this subsection, we examine the contracts that do attract arbitrageur entry and provide suggestive evidence for the number of arbitrageurs simultaneously engaging in the dividend play strategy in each contract.

An advantage of our granular, transaction-level data is that we can estimate the number of arbitrageurs exploiting each contract, which prior literature has not done. As mentioned previously, dividend play trades in our data appear as a sequence of trades (normally five or more), executed within several seconds from the floor of an options exchange. There are two simultaneously executed legs of this trade by each arbitrageur: One leg is a sequence of buys of one call option contract and the other a sequence of sells of a contract with a neighboring strike. Trade sizes within each sequence of trades are typically the same, but they differ across sequences. Differences in trade sizes across dividend play trade sequences are likely due to execution preferences of individual arbitrageurs. We therefore use the number of unique trade sizes within each sequence of trades that we classify as dividend play in a given contract as a proxy for the number of arbitrageurs engaging in dividend play in that contract. The dividend play trade is normally executed by a pair of arbitrageurs, simultaneously establishing long and short positions in the contract. The latter collects (a share of) the windfall gain and the former facilitates the trade by serving as a counterparty.

Figure 4: Number of arbitrageurs in dividend play trades



This figures depicts percentage split of trades executed on exchange floor by the number of unique trade sizes, i.e., the number of arbitrageurs. We include only contracts in our dividend play sample. The largest gray shaded area corresponds to the period of floor closures on all exchanges.

Our measure identifies the number of such pairs, and we count each pair as one arbitrageur (because only one party receives the gain).

Figure 4 plots a percentage split of dividend play trades by the number of arbitrageurs exploiting each contract. It is striking that in the majority of contracts we observe only one arbitrageur. Only in a small fraction of contracts is the number of arbitrageurs above three. Since there are no obvious impediments to free entry of other arbitrageurs who regularly engage in dividend play, it is puzzling that so few arbitrageurs participate in each contract.

Dividend play trades are easy to detect in transaction-level data, which is available to arbitrageurs from standard datafeeds (e.g., OPRA or those provided by exchanges). There may be some delays with reporting floor trades, but these delays are hardly significant for our application. Arbitrageurs are therefore fully aware of the presence of other arbitrageurs in a given contract before they decide to engage in a dividend play trade in it. According to textbook theories, however, this should not deter them from entering and competing for profits in the same contract.

The largest gray shaded area in Figure 4 corresponds to the closure of all exchange floors in the U.S. due to the COVID-19 pandemic. Our measure of floor trading is indeed zero over this period. Furthermore, the total trading volume on dividend play dates during the closures is the same as on any other day, which provides additional validation of the measure. Even when PHLX floor was closed but ARCA and BOX floors were open, the

mean trading volume on dividend play was an order of magnitude lower.

Table 7: Number of arbitrageurs in contracts that attract entry

				Percentile				
	Mean	**Median**	**Std. Dev.**	**1%**	**10%**	**90%**	**99%**	**Observations**
Number of arbitrageurs								
- full sample	1.77	1	1.43	1	1	3	7	19,425
- top EEV tercile	1.83	1	1.49	1	1	3	7	7,145
- top moneyness tercile	1.71	1	1.27	1	1	3	7	6,466
- top profit tercile	2.08	2	1.81	1	1	4	9	7,768

This table reports the distribution of the number of arbirtageurs for all contracts in the dividend play sample that attract arbitrageur entry.

It is clearly also a possibility that our measure of the number or arbitrageurs is inaccurate and that multiple arbitrageurs use trades of identical size. This caveat aside, we have established the following stylized fact in our dividend play sample.

Fact 4: *Of all arbitrage opportunities (option contracts) that attract arbitrageur entry, 56% are exploited by only one arbitrageur.*

Table 7 demonstrates that the number of arbitrageurs engaging in a given contract is generally small. Across all contracts that attract arbitrageur entry, in only 6% of those contracts, the number of arbitrageurs exceeds 3.

4.5 Possible explanations of selective arbitrageur participation

The evidence summarized in Facts 1–4 is consistent with our model. We could interpret the contracts that do not attract any arbitrageur entry as natural markets of some intermediaries. Some of the intermediaries in our application are market makers, who normally have access to the trading floors and who might be interested in the dividend play strategy. Due to significant economies of scale in market making, this segment is served by a small number of firms.[23] Market makers intermediate retail investor transactions, and in particular write options that retail investors buy. Prior to the dividend play date, they may accumulate an inventory of such written options. Because of early exercise mistakes, they are then set to benefit from the windfall capital gains, unless another arbitrageur engages in dividend play in these contracts. This parallels exactly our model's definition of a natural

[23] See, e.g., https://www.barrons.com/articles/market-makers-in-equity-options-are-vanishing-1496459364. In a recent paper, Ernst and Spatt (2022) list market makers currently operating in the U.S. options market, and their number is small.

market. Based on the data from Nasdaq exchanges (which included PHLX), Figure 5 in the Appendix confirms that market makers are indeed net sellers on average of call option contracts in our dividend play sample on the 100 most mentioned tickers on the `WallStreetBets` over the past week, which is a proxy for retail investor popularity.

Owing to daily fee caps on dividend play on options exchanges, the entry costs are fixed, as in our model. Our model predicts that we should therefore observe no entry in the natural markets of options market makers, and our Fact 1 is consistent with that. Moreover, our empirical tests show that arbitrageurs are less likely to enter profitable contracts that have positive retail buy imbalances (and, hence, are natural markets for market makers who are the likely option writers for these contracts). For profitable contracts that are in neither intermediary's inventory, we should observe entry of only one arbitrageur. Consistent with the model, Fact 4 indeed documents that, conditional on entry, in the majority of the contracts we observe entry by only one arbitrageur. The version of the model that is closest to the data is the one discussed at the end of Section 2, in which intermediaries enter only one of the last two markets, leaving the other to the other intermediary. In the data, arbitrageurs operate in pairs, each pocketing a gain in one contract, while facilitating the other arbitrageur's gain in another contract.

We note that sometimes we observe entry by two or more arbitrageurs in a given contract. We interpret this not as a reversion to Nash equilibrium (punishment) but as occasional "misfires." Tacit collusion does not require any explicit agreements, and so occasional misfires are likely to happen. This is especially true in our application because there are many contracts that may be profitable to exploit on a dividend play date (the median in our sample is 54). Also, unlike in our model, there is some noise in the data. For example, floor transactions are likely to be reported late and so an arbitrageur may not realize that another arbitrageur has already entered that contract.

Fact 3 is also consistent with our model and is inconsistent with standard theories of arbitrage. Unless the arbitrageurs are playing a cooperative strategy, there should be a pecking order in which they enter profitable contracts, with most profitable being exploited first.

We note that the repeated game assumption of the model hardly requires any defense in our application. Most of the dividend play transactions in our sample occur on the floor on only two exchanges, PHLX and BOX (Table 2). Given that most of trade in options is electronic, modern-day exchange floors are populated by a small number of market participants, who specialize in floor trading and frequently interact with each other. Finally, while our assumption is that dividend play profits are observed with certainly does not hold in the data, with risk-neutral intermediaries it is sufficient to replace certain profits by

expected profits in our model for the results to go through.

Other theories could potentially rationalize the phenomena that we document. In the rest of this section, we explore alternative explanations of the selective entry by arbitrageurs, based on the existing limits-to-arbitrage literature. Trading costs is the most natural one. First, per-contract exchange fees might prevent arbitrageurs to fully exploit the available arbitrage opportunities. However, there exist dividend-play specific fee caps on PHLX and, more recently, BOX exchanges, in which most of the dividend play activity takes place (Table 2).[24] Those fee caps limit the total costs paid by a market maker or other arbitrageur on a particular day at the options class level: Harvesting the profit from an additional contract would not increase payments to the exchanges once the limit is reached. Second, given that dividend play usually requires two participating parties, it is highly likely that they would agree on the transaction price that allows for mutually beneficial profit sharing. There is no clear reason why they would omit any particular contract from their agreement due to its otherwise lower liquidity. Finally, in the analysis above, we always control for contract liquidity or match on contract relative spread. It is therefore unlikely that the contracts in which arbitrageurs do not engage in dividend play are systematically less liquid.

Another potential explanation is that arbitrageurs expose themselves to risk (or hedging cost) if the delta of the options they buy and sell is not 1. So less deep in-the-money options on more volatile underlyings would be riskier or costlier. However, options in our dividend play sample are normally deep in-the-money. Our sample of arbitrage opportunities has a mean (median) moneyness of over 6 (4) and the top tercile by profitability has mean (median) moneyness close to 8 (5), and yet we still do not observe arbitrage activity in many of those contracts. Moreover, we observe in our data that traders simultaneously establish offsetting long and short positions, and so they are fully hedged. There can also be financing costs, which would likely be increasing in the time to expiration. It is also possible that the arbitrageurs' capital constraints bind. However, most regulatory requirements typically involve netted positions, which are relatively low given the symmetric and fully hedged nature of the strategy. It is therefore not clear why capital constraints may bind unless they bind due to the arbitrageurs' internal risk management guidelines. Relatedly, such large trades are associated with high operational risks. According to SIFMA, Bank of America Merrill Lynch incurred a $10 million loss due to a human error when executing the dividend play strategy.[25] Still, such explanations cannot produce the variation in floor trader activity within and across tickers that we document: Table 4 illustrates that there

[24] See the PHLX pricing schedule: https://listingcenter.nasdaq.com/rulebook/phlx/rules/phlx-options-7 and BOX fee schedule: https://boxoptions.com/regulatory/fee-schedule/.

[25] See https://www.reuters.com/article/us-usa-options-apple-idUSKBN0IQ2FA20141106.

are many profitable contracts in which floor traders do not participate at all.

Finally, it has been documented that even sophisticated market players exhibit limits to attention (Kacperczyk, Nieuwerburgh, and Veldkamp (2016)). Indeed, there may be hundreds of potentially profitable contracts available to dividend play on each cum-dividend day (thousands, in the case of SPY). Perhaps traders simply cannot evaluate all relevant pricing parameters, enter into an agreement with each other, and process the necessary number of trades? First, it is not clear why other exchange members do not enter to reap arbitrage profits if such limits exist. However, we proceeded to test this hypothesis more formally. To do so, we used the number of stock-level EPS (Hirshleifer, Lim, and Teoh (2009)) and macroeconomic announcements (Savor and Wilson (2014)) as proxies for limits to attention and did not find that those mattered for floor trader activity. These additional results are available upon request.

One alternative explanation that we cannot rule out is that some profits are left unexploited because of the stigma and reputational costs associated with the dividend play strategy. The SEC has clearly signaled its disapproval of the strategy in its 2014 rule aimed at making the strategy impractical. (See footnote 7.) Reputational costs could explain the lack of entry of new arbitrageurs. However, they cannot explain why arbitrageurs who regularly engage in this strategy, and hence are willing to incur reputation costs, still leave money on the table. Finally, we cannot rule out an explanation based on clearing fees. Clearing fees are negotiated bilaterally and they are not available in the public domain.

5 Discussion and policy implications

The dividend play arbitrage does not serve any particularly useful purpose in financial markets. It neither improves liquidity nor aids price discovery. It simply reallocates profits from the original option writers to arbitrageurs. In the process of doing so, however, it dramatically inflates trading volume and hence skews various important statistics that market participants rely on, while exposing arbitrageurs and potentially exchanges to large operational risk. With its 2014 regulation, the SEC attempted to put an end to this strategy.

Our paper demonstrates that market participants found a way around the 2014 SEC regulation, leading to a resurgence of the strategy, and highlights the difficulties of devising effective regulation in financial markets. Two previous papers on dividend play, Hao, Kalay, and Mayhew (2009) and Pool, Stoll, and Whaley (2008), written before the SEC regulation, argue that arbitrageurs engage in dividend play by simultaneously opening long and short positions in the *same* call option contract. The 2014 regulation affected clearing rules, and it was no longer beneficial for arbitrageurs to long and short the same contract. In our

transaction-level data, we clearly see that after the regulation, dividend play arbitrageurs simultaneously long and short contracts with *neighboring* strikes. For the purposes of computing regulatory risk exposures, these positions are simply netted to zero. There is some economic risk exposure resulting from this trade because the deltas of the long and short positions are not exactly the same, but the short-horizon nature of the trade and the ability of arbitrageurs to delta hedge their economic exposure, if desired, still makes it a worthwhile strategy, leading to its resurgence in the data.

We propose an alternative way to curtail the strategy. The problem lies in the random assignment by the OCC of options writers who are required to deliver the underlying. Dividend play is a play on this random assignment. An alternative would be the First-in, First-out (FIFO) rule, similar to a regulation imposed by the National Futures Association (NFA) in the U.S. on trading international currencies and currency derivatives (forex). It requires that currency traders close out their oldest positions first, before closing out more recent positions. This rule would favor the original option writers over the dividend play arbitrageurs.[26]

Finally, it is apparent from our analysis that the new generation of investors is still lacking in financial education that is required to trade options. Their mistakes in sophisticated financial decisions such as early exercise of an option generate transfers to arbitrageurs or market makers. It is not clear whether retail investing platforms have the appropriate incentives to prevent their customers from making trading mistakes. The question of optimal options exercise requires knowledge of option pricing models, which retail investors are likely to lack. One possibility would be to require retail brokerages to report options' early exercise values to investors. The early exercise value could be computed from the Black-Scholes model. Another possibility is to make *automatic* early exercise on last cum-dividend dates, when it is optimal to do so, a default option for investors, from which they can opt out if they wish. This simple enhancement would prevent retail investor losses of around $170 million in our sample.[27] More generally, the literature on investor protection has long been concerned about complex investment products and incentives of intermediaries.[28] The complexity of options contracts from the viewpoint of an average retail investor and the potentially misaligned incentives of intermediaries call for enhancements to investor protection on retail trading platforms.

[26] We thank Terry Hendershott for suggesting this idea to us.

[27] While $170 million is the total potential profit to be exploited by dividend play arbitrageurs, a large fraction of it is likely to come from retail investors, as sophisticated investors are likely to exercise their options optimally.

[28] See e.g., Barbu (2022), Bhattacharya, Illanes, and Padi (2019), Egan (2019), Heimer and Simsek (2019), Célérier and Vallée (2017), and Campbell, Jackson, Madrian, and Tufano (2011).

6 Conclusion

In this paper we contribute to the literature on arbitrage in financial markets by introducing strategic arbitrageurs. Arbitrageurs are large and their number is limited because of economies of scale and entry costs. Furthermore, arbitrageurs in our model interact repeatedly with each other. Repeated game considerations enlarge the set of possible equilibrium outcomes, and in particular give rise to tacit collusion. Arbitrageurs enter into only a subset of markets, which reduces their combined entry costs and thus increases social welfare.

Our empirical application mimics the setup of the model. We consider dividend play, a specific arbitrage strategy in the options market for which we can accurately track arbitrageur activity. Dividend play arbitrageurs execute their trades on the floor of the exchange, which is typically populated by a limited number of players, and they interact repeatedly. Needless to say, there are entry costs associated with floor trading, although these costs are lower for intermediaries who are already present in the options markets, such as for example market makers. We document that arbitrageurs routinely do not take part in dividend play in a large number of exploitable options contracts, even the most profitable ones. Moreover, if arbitrageurs do decide to engage in a dividend play in a given contract, their number is usually limited to only one arbitrageur. These implications are consistent with our model. We also consider alternative explanations for the selective entry of arbitrageurs, many of which we are able to rule out.

Dividend play arbitrage is not the only setting that satisfies our model's assumptions. For example, in the options market there are five further arbitrage strategies, facilitated by exchanges via fee caps, which are also executed on the floor. More broadly, financial markets with entry costs and natural economies of scale, intermediated by a limited number of players, may present other applications of our model, as strategic interactions cannot be ruled out in such markets. More opaque, for example, over-the-counter markets may also present suitable applications, as repeated game considerations are more likely to present themselves in more opaque markets.

References

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen, 2020, The Fall of the Labor Share and the Rise of Superstar Firms, *The Quarterly Journal of Economics* 135, 645–709.

Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 1997, Empirical Performance of Alternative Option Pricing Models, *The Journal of Finance* 52, 2003–2049.

Barber, Brad, and Terrance Odean, 2013, The Behavior of Individual Investors, vol. 2 . chap. 22, pp. 1533–1570 (Elsevier).

Barber, Brad M., and Terrance Odean, 2001, Boys will be Boys: Gender, Overconfidence, and Common Stock Investment, *Quarterly Journal of Economics* 116, 261–292.

Barbu, Alexandru, 2022, Ex-Post Loss Sharing In Consumer Financial Markets, working paper, London Business School.

Barraclough, Kathryn, and Robert E. Whaley, 2012, Early Exercise of Put Options on Stocks, *The Journal of Finance* 67, 1423–1456.

Battalio, Robert, Stephen Figlewski, and Robert Neal, 2020, Option Investor Rationality Revisited: The Role of Exercise Boundary Violations, *Financial Analyst Journal* 76, 82–99.

Bhattacharya, Vivek, Gaston Illanes, and Manisha Padi, 2019, Fiduciary Duty and the Market for Financial Advice, working paper, available at SSRN: https://ssrn.com/abstract=3281345.

Boyarchenko, Nina, Thomas M. Eisenbach, Pooja Gupta, Or Shachar, and Peter Van Tassel, 2020, Bank-Intermediated Arbitrage, Federal Reserve Bank of New York Staff Report 858.

Bryzgalova, Svetlana, Anna Pavlova, and Taisiya Sikorskaya, 2023, Retail Trading in Options and the Rise of the Big Three Wholesalers, *The Journal of Finance,* forthcoming.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini, 2007, Down or out: Assessing the welfare costs of household investment mistakes, *Journal of Political Economy* 115, 707–747.

Campbell, John Y, Howell E. Jackson, Brigitte C. Madrian, and Peter Tufano, 2011, Consumer Financial Protection, *Journal of Economic Perspectives* 25, 91–114.

Célérier, Claire, and Boris Vallée, 2017, Catering to Investors Through Security Design: Headline Rate and Complexity, *The Quarterly Journal of Economics* 132, 1469–1508.

Christie, William G., and Paul H. Schultz, 1994, Why do NASDAQ Market Makers Avoid Odd-Eighth Quotes?, *The Journal of Finance* 49, 1813–1840.

Cosma, Antonio, Stefano Galluccio, Paola Pederzoli, and Olivier Scaillet, 2020, Early Exer-

cise Decision in American Options with Dividends, Stochastic Volatility, and Jumps, *Journal of Financial and Quantitative Analysis* 55, 331–356.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger, 2020, The Rise of Market Power and the Macroeconomic Implications, *The Quarterly Journal of Economics* 135, 561–644.

de Silva, Tim, Kevin Smith, and Eric C. So, 2022, Losing is Optional: Retail Option Trading and Earnings Announcement Volatility, working paper, available at SSRN: https://ssrn.com/abstract=4050165.

Dutta, Prajit K., and Ananth Madhavan, 1997, Competition and collusion in dealer markets, *The Journal of Finance* 52, 245–276.

Egan, Mark, 2019, Brokers versus Retail Investors: Conflicting Interests and Dominated Products, *The Journal of Finance* 74, 1217–1260.

Ernst, Thomas, and Chester S. Spatt, 2022, Payment for Order Flow and Asset Choice, *NBER Working Paper 29883*.

Fardeau, Vincent, 2021, Arbitrage with Financial Constraints and Market Power, working paper, available at SSRN: https://ssrn.com/abstract_id=3630914.

Fudenberg, Drew, and Eric Maskin, 1986, The Folk Theorem in Repeated Games with Discounting or with Incomplete Information, *Econometrica* 54, 533–554.

Gromb, Denis, and Dimitri Vayanos, 2002, Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs, *Journal of Financial Economics* 66, 361–407.

Grullon, Gustavo, Yelena Larkin, and Roni Michaely, 2019, Are US Industries Becoming More Concentrated?, *Review of Finance* 23, 697–743.

Hao, Jia, Avner Kalay, and Stewart Mayhew, 2009, Ex-dividend Arbitrage in Option Markets, *The Review of Financial Studies* 23, 271–303.

Heimer, Rawley, and Alp Simsek, 2019, Should Retail Investors' Leverage be Limited?, *Journal of Financial Economics* 132, 1–21.

Hendershott, Terrence J., Saad Khan, and Ryan Riordan, 2022, Option Auctions, working paper, available at SSRN: https://ssrn.com/abstract=4110516.

Hirshleifer, David, Sonya Seongyeon Lim, and Siew Hong Teoh, 2009, Driven to Distraction: Extraneous Events and Underreaction to Earnings News, *The Journal of Finance* 64, 2289–2325.

Jensen, Mads V., and Lasse H. Pedersen, 2016, Early Option Exercise: Never Say Never, *Journal of Financial Economics* 121, 278–299.

Kacperczyk, Marcin, Stijn Van Nieuwerburgh, and Laura Veldkamp, 2016, A Rational Theory of Mutual Funds' Attention Allocation, *Econometrica* 84, 571–626.

Lee, Charles M. C., and Mark J. Ready, 1991, Inferring Trade Direction from Intraday Data, *The Journal of Finance* 46, 733–746.

Mixon, Scott, and Esen Onur, 2017, Dividend Swaps and Dividend Futures: State of Play, *The Journal of Alternative Investments* 19, 27–39.

Muravyev, Dmitriy, 2016, Order Flow and Expected Option Returns, *The Journal of Finance* 71, 673–708.

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

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

Savor, Pavel, and Mungo Wilson, 2014, Asset Pricing: A Tale of Two Days, *Journal of Financial Economics* 113, 171–201.

Shleifer, Andrei, and Robert W. Vishny, 1997, The Limits of Arbitrage, *The Journal of Finance* 52, 35–55.

Siriwardane, Emil, Adi Sunderam, and Jonathan L. Wallen, 2022, Segmented Arbitrage, NBER Working Paper 30561.

Tirole, Jean, 1988, *The Theory of Industrial Organization* . MIT Press Books (The MIT Press).

A Appendix: Proofs

Lemma 1 *Suppose that the best response of Intermediary 1 is to enter all three markets, i.e., $k_1^j > 0$, $j = 1, 2, 3$. The best response of Intermediary 1 is then given by*

$$k_1^1 = \left[\Lambda_1 \left(\frac{A}{k_2^1}\right)^{1/2} - 1\right] k_2^1, \quad k_1^2 = \left[\Lambda_1 \left(\frac{A}{k_2^2}\right)^{1/2} - 1\right] k_2^2, \quad \textit{and} \quad k_1^3 = \left[\Lambda_1 \left(\frac{A^*}{k_2^3}\right)^{1/2} - 1\right] k_2^3,$$

where $\Lambda_1 = 2\left[(Ak_2^1)^{1/2} + (Ak_2^2)^{1/2} + (A^*k_2^3)^{1/2}\right]^{-1}$.

Lemma 2 *Suppose that the best response of Intermediary 1 is to enter all three markets, i.e., $k_1^j > 0$, $j = 1, 2, 3$. The best response of Intermediary 2 is then given by*

$$k_2^1 = \left[\Lambda_2 \left(\frac{A}{k_1^1}\right)^{1/2} - 1\right] k_1^1, \quad k_2^2 = \left[\Lambda_2 \left(\frac{A}{k_1^2}\right)^{1/2} - 1\right] k_1^2, \quad \textit{and} \quad k_2^3 = \left[\Lambda_2 \left(\frac{A^*}{k_1^3}\right)^{1/2} - 1\right] k_1^3,$$

where $\Lambda_2 = 2\left[(Ak_1^1)^{1/2} + (Ak_1^2)^{1/2} + (A^*k_1^3)^{1/2}\right]^{-1}$.

Proof of Lemmas 1 and 2.

We first prove Lemma 1.The problem defined in (4)–(5) is a convex constrained optimization problem, and the (linear) resource constraint (5) will be satisfied with equality. The first-order conditions are:

$$\begin{aligned}
A\frac{1}{(1 + k_1^1/k_2^1)^2} &= \lambda k_2^1, \\
A\frac{1}{(1 + k_1^2/k_2^2)^2} &= \lambda k_2^2, \\
A^*\frac{1}{(1 + k_1^3/k_2^3)^2} &= \lambda k_2^3,
\end{aligned}$$

where λ is the Lagrange multiplier on the resource constraint (5). The first-order conditions can be rewritten as

$$k_1^1 = \left[A\left(\frac{1}{\lambda k_2^1}\right)^{1/2} - 1\right] k_2^1, \tag{16}$$

$$k_1^2 = \left[A\left(\frac{1}{\lambda k_2^2}\right)^{1/2} - 1\right] k_2^2, \tag{17}$$

$$k_1^3 = \left[A^*\left(\frac{1}{\lambda k_2^3}\right)^{1/2} - 1\right] k_2^3. \tag{18}$$

Substituting (16)–(18) into the resource constraint (5) and simplifying, we arrive at

$$\left(\frac{1}{\lambda}\right)^{1/2}\left[(Ak_2^1)^{1/2}+(Ak_2^2)^{1/2}+(A^*k_2^3)^{1/2}\right]-(k_2^1+k_2^2+k_2^3)=1.$$

Using the resource constraint of Intermediary 2 (satisfied with equality), (6), defining $\Lambda_1 \equiv \left(\frac{1}{\lambda}\right)$ and solving the above equation for Λ_1, we arrive at

$$\Lambda_1 \equiv \left(\frac{1}{\lambda}\right)^{1/2} = 2\left[\left(Ak_2^1\right)^{1/2}+\left(Ak_2^2\right)^{1/2}+\left(A^*k_2^3\right)^{1/2}\right]^{-1}.$$

Substituting Λ_1 for $\left(\frac{1}{\lambda}\right)^{1/2}$ in (16)–(18) and rearranging terms, we arrive at the statement in the lemma. A proof of Lemma 2 is analogous. □

Proof of Proposition 1. Nash equilibria are given by the intersection of the intermediaries' best responses. Combining the best responses from Lemmas 1 and 2, we arrive at the system of equations (8)–(11) in the statement of the proposition. □

Proof of Corollary 1. Given symmetry of the intermediaries' payoffs in the first two markets, we look for a Nash equilibrium in which $k_1^1 = k_2^1 = k_1^2 = k_2^2$. Similarly, given the symmetry of the intermediaries' payoffs in market M_3, we look for a Nash equilibrium in which, additionally, $k_1^3 = k_2^3$. Under these conditions, the system of equations in Proposition 1 simplifies significantly, leading to the expressions in part (i) of the corollary.

In the symmetric Nash equilibrium, the intermediaries enter all three markets and they equally split arbitrage opportunities in each of the three markets. The total fixed cost that each of them pays is only $f + f^*$ because the fixed cost of entry into the intermediaries' natural market is zero.

The equilibrium presented in part (i) of the corollary exists if the intermediaries' fixed costs are not too high so that entry in all three markets is profitable for them. The condition specified in the corollary guarantee that. □

Proof of Proposition 2. In a stage game, the Pareto efficient equilibrium that is symmetric in payoffs is an allocation resulting from maximizing social welfare subject to the resource constraints:

$$\max_{k_i^1,k_i^2,k_i^3,\, i\in\{1,2\}} 0.5\pi_1(k_1^1,\, k_1^2,\, k_1^3) + 0.5\pi_2(k_2^1,\, k_2^2,\, k_2^3), \tag{19}$$

$$\text{s.t. } k_i^1 + k_i^2 + k_i^3 \leq 1,\ i \in \{1,2\}, \tag{20}$$

where we have assigned the intermediaries identical Pareto weights $\mu_1 = \mu_2 = 0.5$. Since the

fixed costs of market entry in markets M_1 and M_2 can be reduced to zero, it is straightforward to see that in the symmetric Pareto optimal allocation only Intermediary 1 should enter M_1 and only Intermediary 2 market M_2. The harvested profits from these markets are identical and are equal to A. Similarly, it is inefficient from the social welfare perspective for both intermediaries to pay the fixed cost of entry into market M_3. It is therefore welfare maximizing that they agree to flip a fair coin to decide which one of the two intermediaries enters the market and which one stays out. Their payoff from this strategy is the same in expectation, and they each earn $A^*/2 - f^*/2$ in market M_3. Note that this payoff is higher than that achieved with mixed strategies whereby the intermediaries enter market M_3 with some probability and otherwise stay away, because under these strategies sometimes both intermediaries end up entering the market, each paying a fixed cost of entry, and sometimes none of them would enter and therefore the arbitrage opportunity A^* would be left unharvested. Under the strategies described in the proposition, in all three markets the intermediaries earn $A + A^*/2 - f^*/2$ each in expectation, which is the maximum profit possible in a symmetric solution to (19)–(20).

In markets in which the intermediaries are present, the allocation of resources k_i^j, $i, j \in \{1, 2\}$ are indeterminate, as long as $k_i^3 > 0$ for the intermediary who is chosen to enter market M_3. We focus on the solution in which the intermediaries allocate zero resources to their natural markets, since they do not need to commit any resources to harvest the arbitrage opportunity there. Any positive amount of resource committed to market M_3 by the (only) entrant $k_i^3 \in (0, 1]$ guarantees the that that intermediary captures the entire arbitrage opportunity. □

Proof of Proposition 3. To prove that the proposed equilbrium is indeed a subgame perfect equilbrium of the game, let us explore a deviation from it and show that it is not profitable in an infinitely repeated game. If the intermediaries play cooperatively, each intermediary's payoff is

$$\begin{aligned} &\sum_{t=0}^{\infty} \delta^t \left(A + \frac{A^*}{2} - \frac{f^*}{2} \right) \\ &= \frac{1}{1-\delta} \left(A + \frac{A^*}{2} - \frac{f^*}{2} \right) \end{aligned} \tag{21}$$

in expectation (Proposition 2). If an intermediary deviates, its payoff is

$$\begin{aligned} &2A + \frac{A^*}{2} - f - f^* + \sum_{t=1}^{\infty} \delta^t \left(A + \frac{A^*}{2} - f - f^* \right) \\ &= 2A + \frac{A^*}{2} - f - f^* + \frac{\delta}{1-\delta} \left(A + \frac{A^*}{2} - f - f^* \right), \end{aligned} \tag{22}$$

where the term before the infinite sum is the payoff from a deviating and the infinite sum is the payoff in the ensuing punishment phase. The payoff in (21) is greater or equal to the payoff in (22) if and only if

$$\delta \geq \frac{A - f - \frac{f^*}{2}}{A},$$

which establishes the result in the proposition. □

Proposition 4 (Tacit collusion equilibrium with four markets) *Consider the following symmetric trigger strategy:*

Cooperative phase: *Cooperate and*

(i) allocate no resources to markets M_1 and M_2. That is, $(k_{1t}^1, k_{1t}^2) = (0, 0)$, and $(k_{1t}^1, k_{1t}^2) = (0, 0)$, so that the natural intermediaries capture all of arbitrage opportunities in their natural markets;

(ii) Intermediary i enters market M_3, committing $k_{it}^3 > 0$ and harvesting the entire arbitrage opportunity A^*, and does not enter market M_4.

(iii) Intermediary j, $j \neq i$, enters market M_4, committing $k_{jt}^4 > 0$ and harvesting the entire arbitrage opportunity A^*, and does not enter market M_3.

Punishment for deviation: *If the opponent deviated from the above cooperative strategy in the previous period, enter all markets (M_1, M_2, M_3, and M_4) and play purely competitively forever, which results in a Nash equilibrium in each stage game thereafter.*

If both players play the trigger strategy above and $\delta \geq \frac{A+A^*/2-f-f^*}{A+A^*/2}$, the symmetric Pareto optimal (or tacit collusion) equilibrium is a subgame perfect equilibrium of the repeated game.

Proof of Proposition 4. To prove that the proposed equilbrium is indeed a subgame perfect equilbrium of the game, let us explore a deviation from it and show that it is not profitable in an infinitely repeated game. If the intermediaries play cooperatively, each

intermediary's payoff is

$$\begin{aligned}&\sum_{t=0}^{\infty}\delta^t\left(A+A^*-f^*\right)\\&=\frac{1}{1-\delta}\left(A+A^*-f^*\right).\end{aligned}\tag{23}$$

If an intermediary deviates, its payoff is

$$\begin{aligned}&2A+\frac{3}{2}A^*-f-2f^*+\sum_{t=1}^{\infty}\delta^t\left(A+A^*-f-2f^*\right)\\&=2A+\frac{3}{2}A^*-f-2f^*+\frac{\delta}{1-\delta}\left(A+A^*-f-2f^*\right),\end{aligned}\tag{24}$$

where the term before the infinite sum is the payoff from a deviating and the infinite sum is the payoff in the ensuing punishment phase. The payoff in (21) is greater or equal to the payoff in (22) if and only if

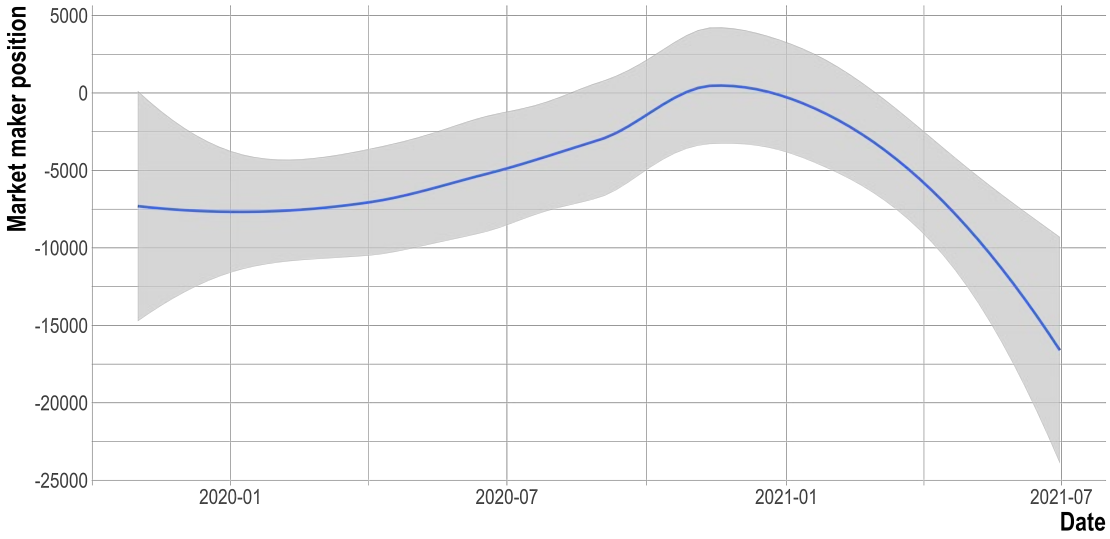
$$\delta\geq\frac{A+\frac{A^*}{2}-f-f^*}{A+\frac{A^*}{2}},$$

which establishes the result in the proposition. □

B Appendix: Empirical Analysis

B.1 Market makers' position in the top 100 most mentioned tickers on `WallStreetBets`

Figure 5: Market makers' position in the top 100 most mentioned tickers on `WallStreetBets` on the PHLX exchange



This figures depicts the smoothed market makers' trade position in the top 100 most mentioned tickers on `WallStreetBets` on the PHLX exchange. The market makers' trade position in option contract c on date t is computed using the End-of-Day Nasdaq PHOTO data as ΔMarket maker position$_{c,t}$ = $\sum_{s=t-5}^{t}$(Market maker opening buys$_{c,s}$ − Market maker opening sells$_{c,s}$ + Market maker closing buys$_{c,s}$ − Market maker closing sells$_{c,s}$). We aggregate the positions across contracts on the top 100 most mentioned tickers on `WallStreetBets`. To smooth the data, we fit a local second-order polynomial on 3/4 of observations.

B.2 OPRA trade types

The following table presents OPRA trade types, together with their descriptions, implemented on November 4, 2019. We also include the corresponding Trade Condition IDs from LiveVol, our data provider.

Table 8: OPRA trade types for transactions in U.S. options exchanges

OPRA Type Description	**OPRA Message Type**	**LiveVol Trade Condition ID**	**OPRA Condition Description**
AUTO		18	Transaction was executed electronically. Prefix appears solely for information; process as a regular transaction.
CANC		40	Transaction previously reported (other than as the last or opening report for the particular option contract) is now to be cancelled.
CBMO	Multi Leg Floor Trade of Proprietary Products	133	Transaction represents execution of a proprietary product non-electronic multi leg order with at least 3 legs. The trade price may be outside the current NBBO.
CNCL		41	Transaction is the last reported for the particular option contract and is now cancelled.
CNCO		42	Transaction was the first one (opening) reported this day for the particular option contract. Although later transactions have been reported, this transaction is now to be cancelled.
CNOL		43	Transaction was the only one reported this day for the particular option contract and is now to be cancelled.
ISOI		95	Transaction was the execution of an order identified as an Intermarket Sweep Order. Process like normal transaction.
LATE		13	Transaction is being reported late, but is in the correct sequence; i.e., no later transactions have been reported for the particular option contract.
MASL	Multi Leg Auction against single leg(s)	125	Transaction was the execution of an electronic multi leg order which was "stopped" at a price and traded in a two sided auction mechanism that goes through an exposure period and trades against single leg orders/ quotes. Such auctions mechanisms include and not limited to Price Improvement, Facilitation or Solicitation Mechanism.
MESL	Multi Leg auto-electronic trade against single leg(s)	123	Transaction represents an electronic execution of a multi Leg order traded against single leg orders/ quotes.
MLAT	Multi Leg Auction	120	Transaction was the execution of an electronic multi leg order which was "stopped" at a price and traded in a two sided auction mechanism that goes through an exposure period in a complex order book. Such auctions mechanisms include and not limited to Price Improvement, Facilitation or Solicitation Mechanism.
MLET	Multi Leg auto-electronic trade	119	Transaction represents an electronic execution of a multi leg order traded in a complex order book.

continuation on the next page

Table 8: OPRA trade types for transactions in U.S. options exchanges (cont.)

MLCT	Multi Leg Cross	121	Transaction was the execution of an electronic multi leg order which was "stopped" at a price and traded in a two sided crossing mechanism that does not go through an exposure period. Such crossing mechanisms include and not limited to Customer to Customer Cross and QCC with two or more options legs.
MLFT	Multi Leg floor trade	122	Transaction represents a non-electronic multi leg order trade executed against other multi-leg order(s) on a trading floor. Execution of Paired and Non-Paired Auctions and Cross orders on an exchange floor are also included in this category.
MSFL	Multi Leg floor trade against single leg(s)	126	Transaction represents a non-electronic multi leg order trade executed on a trading floor against single leg orders/ quotes. Execution of Paired and Non-Paired Auctions on an exchange floor are also included in this category.
OPEN		6	Transaction is a late report of the opening trade and is out of sequence; i.e., other transactions have been reported for the particular option contract.
OPNL		7	Transaction is a late report of the opening trade, but is in the correct sequence; i.e., no other transactions have been reported for the particular option contract.
OSEQ		2	Transaction is being reported late and is out of sequence; i.e., later transactions have been reported for the particular option contract.
REOP		21	Transaction is a reopening of an option contract in which trading has been previously halted. Prefix appears solely for information; process as a regular transaction.
SCLI	Single Leg Cross ISO	117	Transaction was the execution of an Intermarket Sweep electronic order which was "stopped" at a price and traded in a two sided crossing mechanism that does not go through an exposure period. Such crossing mechanisms include and not limited to Customer to Customer Cross.
SLAI	Single Leg Auction ISO	115	Transaction was the execution of an Intermarket Sweep electronic order which was "stopped" at a price and traded in a two sided auction mechanism that goes through an exposure period. Such auctions mechanisms include and not limited to Price Improvement, Facilitation or Solicitation Mechanism marked as ISO.
SLAN	Single Leg Auction Non ISO	114	Transaction was the execution of an electronic order which was "stopped" at a price and traded in a two sided auction mechanism that goes through an exposure period. Such auctions mechanisms include and not limited to Price Improvement, Facilitation or Soliciation Mechanism.
SLCN	Single Leg Cross Non ISO	116	Transaction was the execution of an electronic order which was "stopped" at a price and traded in a two sided crossing mechanism that does not go through an exposure period. Such crossing mechanisms include and not limited to Customer to Customer Cross and QCC with a single option leg.

continuation on the next page

Table 8: OPRA trade types for transactions in U.S. options exchanges (cont.)

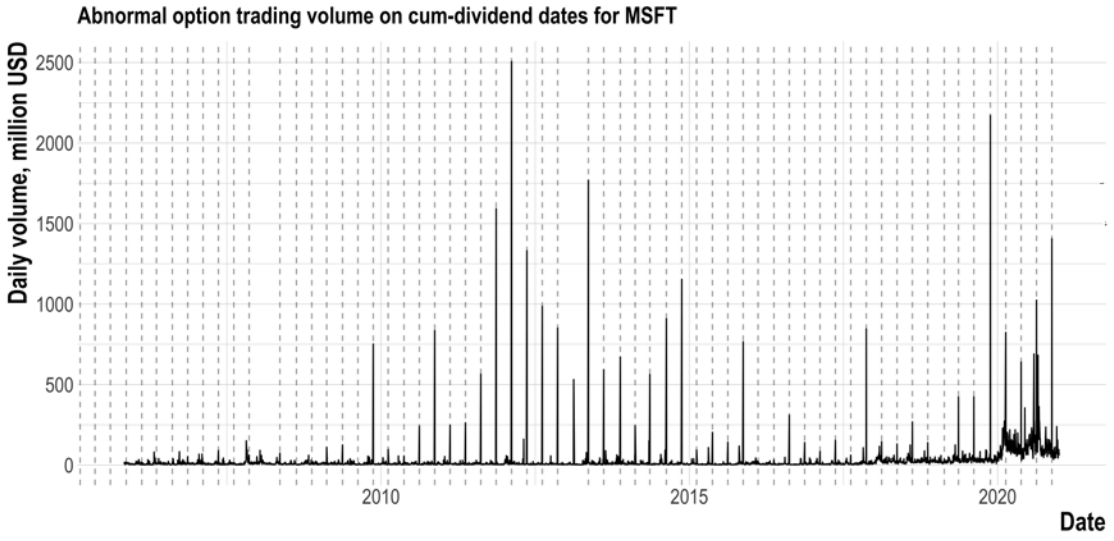
MLCT	Multi Leg Cross	121	Transaction was the execution of an electronic multi leg order which was "stopped" at a price and traded in a two sided crossing mechanism that does not go through an exposure period. Such crossing mechanisms include and not limited to Customer to Customer Cross and QCC with two or more options legs.
SLFT	Single Leg Floor Trade	118	Transaction represents a non-electronic trade executed on a trading floor. Execution of Paired and Non-Paired Auctions and Cross orders on an exchange floor are also included in this category.
TASL	Stock Options Auction against single leg(s)	131	Transaction was the execution of an electronic multi leg stock/options order which was "stopped" at a price and traded in a two sided auction mechanism that goes through an exposure period and trades against single leg orders/ quotes. Such auctions mechanisms include and not limited to Price Improvement, Facilitation or Solicitation Mechanism.
TESL	Stock Options auto-electronic trade against single leg(s)	130	Transaction represents an electronic execution of a multi Leg stock/options order traded against single leg orders/ quotes.
TFSL	Stock Options floor trade against single leg(s)	132	Transaction represents a non-electronic multi leg stock/options order trade executed on a trading floor against single leg orders/ quotes. Execution of Paired and Non-Paired Auctions on an exchange floor are also included in this category.
TLAT	Stock Options Auction	124	Transaction was the execution of an electronic multi leg stock/options order which was "stopped" at a price and traded in a two sided auction mechanism that goes through an exposure period in a complex order book. Such auctions mechanisms include and not limited to Price Improvement, Facilitation or Solicitation Mechanism.
TLCT	Stock Options Cross	128	Transaction was the execution of an electronic multi leg stock/options order which was "stopped" at a price and traded in a two sided crossing mechanism that does not go through an exposure period. Such crossing mechanisms include and not limited to Customer to Customer Cross.
TLET	Stock Options auto-electronic trade	127	Transaction represents an electronic execution of a multi leg stock/options order traded in a complex order book.
TLFT	Stock Options floor trade	129	Transaction represents a non-electronic multi leg order stock/options trade executed on a trading floor in a Complex order book. Execution of Paired and Non-Paired Auctions and Cross orders on an exchange floor are also included in this category.

This table reports OPRA trade types and their descriptions. The type of each transaction in U.S. options exchanges has to be classified using a type description from the table and reported to OPRA. This reporting requirement was implemented on November 4, 2019.

B.3 Abnormal trading volume on cum-dividend dates: Further examples

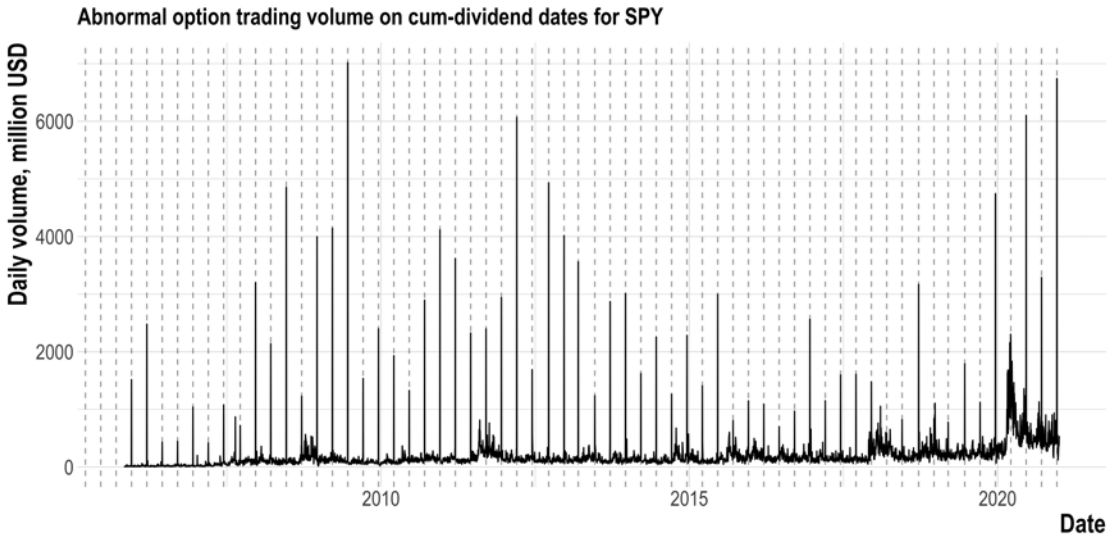
This appendix contains two further examples of abnormal trading volume on cum-dividend dates. The figures below plot daily trading volume of options on Microsoft, MSFT, and on the largest S&P 500 ETF, SPY.

Figure 6: Abnormal trading volume on last cum-dividend dates for MSFT



This figure plots daily trading volume for all call option contracts on MSFT, in millions of U.S. dollars, as reported in OptionMetrics. The dashed lines indicate the last cum-dividend dates.

Figure 7: Abnormal trading volume on last cum-dividend dates for SPY



This figure plots daily trading volume for all call option contracts on SPY, in millions of U.S. dollars, as reported in OptionMetrics. The dashed lines indicate the last cum-dividend dates.

B.4 Dividend play: Another example

Table 9 provides an additional example illustrating the mechanics of the dividend play strategy. Case 1 corresponds to the case when all 1,000 outstanding contracts are exercised and all 1,000 short positions get assigned, so there is no profit for a dividend play strategy to harvest. Case 2 describes what happens if 500 of 1,000 outstanding contracts are left unexercised. Without arbitrageur involvement, half of the short positions in the contract get assigned; the remaining positions deliver a gain of \$0.5 per share and \$25,000 in total for the unassigned short positions, a gain to the original customers with short positions. Now consider the entry of market makers. The market makers attempt to recover most of the potentially harvestable profit of \$25,000. To do so, they buy and simultaneously sell 5,000 contracts and exercise all their long positions. The probability of assignment increases, but, because of the OCC's random assignment, some of the market makers' short positions remain unassigned and hence yield a gain. In our example, market makers harvest \$20,850 out of the total gain of \$25,000. To divert a larger fraction of the total gain from the original customers with short positions, market makers simply increase the number of contracts they buy and sell.

Table 9: Dividend play: Another Example

	OI_{t-1}	New positions(t)	Available for ex.	No. exercised	Prob. Assign.	No. assign.	No. not assign.	Gain per share	Total gain on unassign. positions	OI_t	Fraction unex.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Case 1. Optimal exercise											
Customer	1000	0	1000	1000	100%	1000	0.00	0.5	0	0.00	
Case 2. Suboptimal exercise											
Case 2.1. Without dividend play											
Customer	1000	0	1000	500	50%	500	500	0.5	25000	500	0.5
Case 2.2. With dividend play											
Customer	1000	0	1000	500		916.7	83.33	0.5	4166.7		
Arbitrageurs	0	5000	5000	5000		4583.3	416.67	0.5	20833.3		
Total	1000	5000	6000	5500	92%	5500	500		25000	500	0.5

This table illustrates the dividend play strategy. Date t refers to the last cum-dividend date and OI_t stands for the open interest on date t. This table is similar to Table 1 in Pool, Stoll, and Whaley (2008).

B.5 Dividend play: Technical details

We compute the expected call option ex-dividend price using the Black-Scholes-Merton formula, as follows:

$$
\begin{aligned}
c_{ex} &= S_{ex}e^{-y(T-t)}N(d_1) - Ke^{-r(T-t)}N(d_2), \\
d_1 &= \frac{1}{\sigma\sqrt{T-t}} ln\left(\frac{S_{ex}}{K} + \left[r - y + \frac{\sigma^2}{2}\right](T-t)\right), \\
d_2 &= d_1 - \sigma\sqrt{T-t}, \\
y &= Dividend_{ex}/S_{ex},
\end{aligned}
$$

where S_{ex} is the expected price after the stock goes ex-dividend, that is, price at close on the last cum-dividend day minus expected dividend; $T-t$ is time to maturity in years, that is, difference in the expiration date and the current date in days divided by 360; K is the contract strike; σ^2 is the annualized implied volatility[29]; r is the interpolated maturity-specific interest rate provided by OptionMetrics (annualized %); and $Dividend_{ex}$ is the expected dividend after the ex-date.[30]

[29] We use the daily contract-level implied volatility from OptionMetrics. If it is missing, we interpolate it from the neighboring strikes.

[30] We assume that its size is equal to the current dividend if the stock pays one more dividend after the current dividend until the option expires and 0 otherwise.

B.6 Dividend play sample: Data filters and calculated variables

We use our dataset described in Section 3.1 together with the following filters to arrive at the final dividend play sample. We include all call option contracts with $EEV > 0$. Furthermore, since our valuation might be imperfect, we add a market-based filter of the optimality of exercise: We keep only contracts with a decline in open interest on the last cum-dividend date.[31] By implication, we have only contracts with non-zero open interest on the last cum-dividend date and on the day before that.

Following the early papers on dividend play, we remove contracts with no trading volume on the last cum-dividend date. Additionally, we remove contracts expiring immediately after the ex-dividend date.[32]

To measure arbitrageur activity, we use floor trading share, defined as the total volume in transactions of OPRA types SLFT and MLFT, divided by the total volume on the last cum-dividend date.[33] For both SLIM and Small Share, we compute a one-week moving average and use its lagged value on the cum-dividend date. We use the same rolling measures for the retail activity variables described in the main text, as well as volume, spread, and implied volatility controls.

We compute relative spread quoted at the time of each option trade as $2(best\,ask - best\,bid)/(best\,ask + best\,bid)$ (relative to the midpoint price). We compute moneyness of the trade as $0.5(underlying\,bid + underlying\,ask)/strike - 1$.[34]

[31] This is consitent with Hao, Kalay, and Mayhew (2009).
[32] The last filter does not change results significantly.
[33] In unreported tests, we confirm that using dollar volume-based measures instead yields similar results.
[34] In the absence of TAQ data, we use the underlying bid-ask midpoint as a high-frequency price.

Table 10: Dividend play sample descriptive statistics

	Mean	Median	St. Dev.	p1	p99
Fraction of OI unexercised, %	19.46	2.92	29.41	0.00	98.97
Floor trades volume share on last cum-date	0.49	0.58	0.47	0.00	1.00
D(floor share > 0)	0.53	1.00	0.50	0.00	1.00
SLIM Share	0.14	0.07	0.18	0.00	0.83
D(SLIM buy imbalance)	0.22	0.00	0.41	0.00	1.00
Small Share	0.85	1.00	0.21	0.27	1.00
Internalized volume in underlying (share)	0.17	0.16	0.05	0.06	0.300
WSB mentions	12.45	0.17	44.80	0.00	310.17
OI, log	4.21	4.11	2.20	0.00	9.48
Early exercise value (EEV), $	0.50	0.32	0.60	0.00	2.90
Market EEV, $	0.08	0.02	0.38	-0.43	1.07
Dollar potential profit	4735.93	61.53	46,502.69	0.00	76,546.38
Dollar volume, log	1.71	1.46	1.23	0.00	5.82
Relative spread	0.08	0.05	0.12	0.00	0.59
Implied volatility, annualized	0.43	0.36	0.41	0.00	1.62
Moneyness	6.49	4.32	7.45	0.49	37.69
Days to expiry	62	20	117	3	618

This table reports descriptive statistics for all contracts in the dividend play sample (34,696 observations). SLIM Share and Small Share are the contract-level volume shares of SLIM and small trades, respectively, averaged over one trading week before the last cum-dividend date. D(SLIM buy imbalance)=1 if there was an average buy imbalance in SLIM trading volume over one trading week before the cum-dividend date, and 0 otherwise. The share of internalized volume in the total trading volume in the underlying is computed as a ticker-level non-ATS OTC share of volume executed in the non-ATS OTC space relative to the total trading volume, averaged over one trading week before the last cum-dividend date. WSB mentions is the number of underlying ticker mentions on the `WallStreetBets` forum, averaged over one trading week before the last cum-dividend date. Relative spread is the option contract quoted spread at the time of the trade relative to the midpoint price. Implied volatility is as reported in LiveVol, interpolated using nearest strikes if missing. Moneyness is measured as $(Midpoint\,Price - Strike)/Strike$.

B.7 OTC trading volume, by venue

Table 11: Top 15 venues in the non-ATS OTC market in the U.S.

Firm	**OTC volume, billion shares**	**Venue share in total volume, %**	**Cumulative share, %**
CITADEL SECURITIES	477.82	44.31	44.31
VIRTU	357.61	33.16	77.47
SUSQUEHANNA	119.10	11.04	88.52
TWO SIGMA	48.50	4.50	93.01
JANE STREET CAPITAL	28.49	2.64	95.66
UBS	25.35	2.35	98.01
WOLVERINE	7.29	0.68	98.68
COMHAR CAPITAL MARKETS	3.84	0.36	99.04
HRT EXECUTION SERVICES	3.46	0.32	99.36
LEK SECURITIES CORPORATION	2.27	0.21	99.57
GOLDMAN	2.20	0.20	99.77
ACS EXECUTION SERVICES	0.44	0.04	99.81
IMC	0.32	0.03	99.84
MORGAN STANLEY	0.29	0.03	99.87
COWEN	0.28	0.03	99.90

This table reports the top 15 venues in terms of their total OTC non-ATS volume in 11/2019–06/2021. Based on FINRA OTC Transparency data.

B.8 Dividend play profits and additional measures of retail investor popularity

In addition to the measures of retail trading in options, Small Share and SLIM Share, described in the main text, we use several measures of popularity of the *underlying stock* among retail investors. These measures have been suggested in the previous literature. Although the measures are all for stocks, we hypothesize that cash-constrained, risk-loving retail investors who bet on a stock may also trade options on this stock, which is a leveraged bet on the stock.

Our first measure of stock popularity, introduced in BPS, is a fraction of stock trading volume that was internalized on private trading platforms, off lit exchanges. We utilize FINRA OTC Transparency data, specifically stock-level non-ATS trading volume, which is internalized trading volume. To construct a measure of internalized volume in the underlying, we rely on FINRA OTC Transparency data, which reports stock-level trading volume on automated trading systems (ATS), typically referred to as dark pools, and non-ATS, with the latter representing internalized trades. This data is available from April 2016, by security

and venue. Securities are split into NMS Tier 1, Tier 2, and OTCE. Details are on the FINRA website: https://otctransparency.finra.org/otctransparency/AtsIssueData.[35]

Our second measure of stock popularity is a ticker-level *WallStreetBets mentions*, which is based on the number of mentions of a ticker on `WallStreetBets` subreddit of `reddit.com` popular with retail investors. For our WallStreetBets count measure, we download all comments from Daily Discussion and What Are Your Moves Tomorrow threads on the `WallStreetBets` subreddit of `reddit.com`. This data is collected via PRAW, which is a Python API toolkit to access `reddit.com` for the period from October 1, 2019, to June 30, 2021. In particular, we download all the comments (original posts and reactions to them) for each daily Daily Discussion and What Are Your Moves Tomorrow thread.[36]

We extend our baseline regression (13) in the main text to include the additional popularity variables.

$$Y_{c,t} = \beta_1 \times share_{c,t}^{SLIM} + \beta_2 \times share_{c,t}^{small} + \beta_3' \times Popularity_{c,t}\boldsymbol{\gamma}' X_{c,t} + \alpha_{i,t} + \varepsilon_{c,t} \qquad (25)$$

where, for each contract c on cum-date t, we consider two dependent variables, $Y_{c,t}$: Fraction of open interest not exercised by ex-dividend date and potential profits from dividend play strategy as defined in Equation (12). Slim Share and Small Share are as defined in the main text. The vector *Popularity* contains the two additional popularity measures, both available only at a ticker level. The first measure of retail investor popularity is *Internalized volume in underlying*, measured by the Non-ATS OTC share, which is the share of Non-ATS OTC (internalized) volume in the total trading volume of ticker i in the week of date t. The second measure is *WSB mentions, log*, which is the logarithm of the number of times ticker i was mentioned on `WallStreetBets` forum on date t. Our vector of controls X_t is the same as in the main and it includes: log OI, EEV, log dollar trading volume, relative spread, implied volatility, moneyness, days to expiry. Our specification also includes the ticker by date fixed effects $\alpha_{i,t}$. Our measures of retail investor trading are computed over one trading week before the cum-date because the new generation of retail investors have a strong preference for options expiring within a week (as documented in BPS). Table **??** in Appendix **??** presents an alternative specification in which we measure retail trading over two weeks preceeding a cum-dividend date.

[35] See also FINRA's regulatory notice requiring reporting of this data: https://www.finra.org/rules-guidance/notices/15-48.

[36] Some dates are missing due to retrieval limitations on `reddit.com`. We interpolate between the neighboring dates to fill in those values. We search for each unique historical ticker from CRSP in all the comments, and then sum the number of mentions for each date. We search only for capitalized tickers, which the reddit audience typically uses. Since we might omit any lowercase mentions, and since we do not cover other threads of the forum (such as occasional megathreads), our measure provides a lower bound for ticker popularity.

Column (1) of Table 12 reports the results of the regression in (25), with the fraction of open interest unexercised as the outcome variable. We find that the share of internalized volume in the underlying strongly predicts the fraction left unexercised. While this is a measure of retail investor trading in the underlying stock, it is correlated with retail investor trading in options on that stock. A call option is a leveraged position in a stock, and we hypothesize that some risk-loving retail investors may wish to trade options rather than the stock. We find that a one standard deviation increase in the internalized volume share increases the fraction of options left unexercised by 1.6 percentage points. Using WSB mentions, we also find a positive and significant relationship between retail investor interest and the fraction unexercised. Column (2) shows that the two additional retail measures positively affect potential profits from dividend play. Finally, in both columns Small Share and SLIM share remain statistically significant, and the point estimates in the regression are close to those in Table 5 in the main text.

Table 12: Suboptimal exercise and retail investor popularity

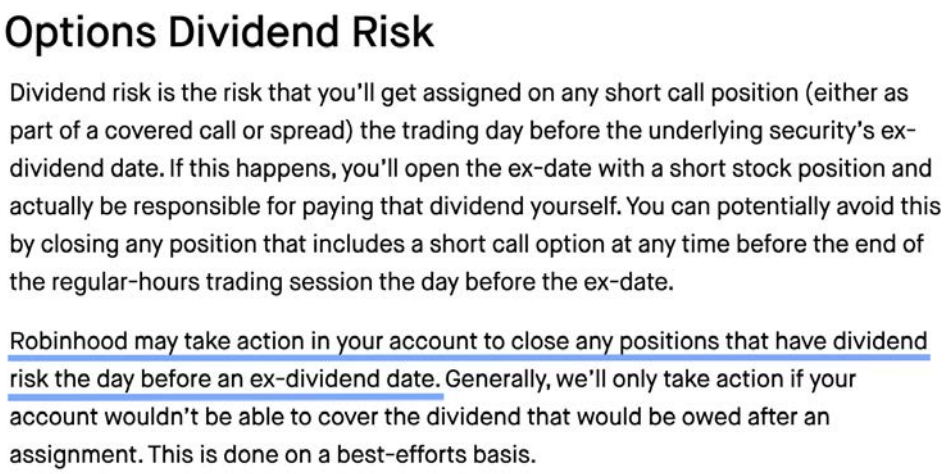
	Dividend play profitability feature	
	Fraction of OI not exercised, %	Potential profits, log U.S. dollar
	(1)	(2)
SLIM Share	5.188***	1.725***
	(5.23)	(11.34)
Small Share	3.932***	0.463**
	(3.57)	(2.05)
Internalized volume in underlying	31.800***	4.093***
	(3.22)	(2.78)
WSB mentions, log	0.562**	0.024
	(2.33)	(0.79)
Observations	32,336	32,336
Adjusted R-squared	0.189	0.235

This table reports estimates of (13) in our dividend play sample. SLIM Share and Small Share are the contract-level volume shares of SLIM and small trades, respectively, averaged over one trading week before the last cum-dividend date. Internalized volume in underlying is the share of non-ATS OTC (i.e., internalized) volume in the total trading volume in the underlying stock or ETF, averaged over one trading week before the cum-dividend date. WSB mentions, log, is the logarithm of total mentions of the ticker on the `WallStreetBets` forum. In column (2), contract controls include log dollar trading volume, relative spread, IV, moneyness, and days to expiration. In column (1), they additionally include log OI and EEV. Ticker controls include underlying price, underlying volatility, underlying relative bid-ask spread, and underlying market cap. Both columns include ticker and date fixed effects. Standard errors are clustered by ticker and date. Robust t-statistics are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

B.9 Dividend risk and automatic actions of retail brokerages

This appendix presents an example of an automatic action to close short positions exposed to dividend risk on the last cum-dividend dates undertaken by retail brokerages. The example is from the Robinhood Terms and Conditions.

Figure 8: Excerpt from Robinhood's Terms and Conditions



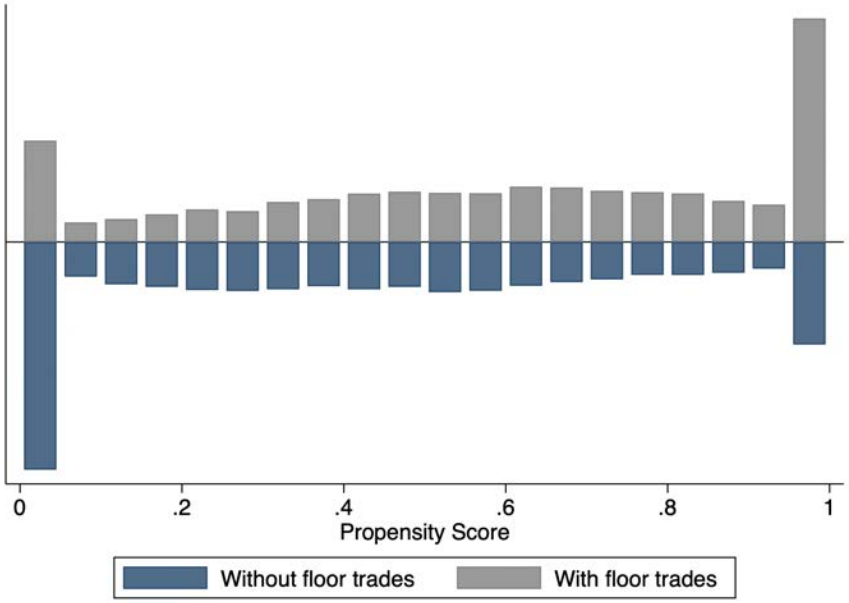
Options Dividend Risk

Dividend risk is the risk that you'll get assigned on any short call position (either as part of a covered call or spread) the trading day before the underlying security's ex-dividend date. If this happens, you'll open the ex-date with a short stock position and actually be responsible for paying that dividend yourself. You can potentially avoid this by closing any position that includes a short call option at any time before the end of the regular-hours trading session the day before the ex-date.

Robinhood may take action in your account to close any positions that have dividend risk the day before an ex-dividend date. Generally, we'll only take action if your account wouldn't be able to cover the dividend that would be owed after an assignment. This is done on a best-efforts basis.

B.10 Selective participation in matched contracts

Figure 9: Floor traders' entry across propensity score levels



This figure depicts the number of contracts with and without floor trades across the scores of propensity to have floor trades. The propensity scores are based on the full set of controls: log OI, EEV, log trading volume, relative spread, IV, moneyness, days to expiry, underlying price, underlying volatility, underlying relative bid-ask spread, and underlying market cap. We report the balance tests in Appendix **??**.

Figure 10: Covariate balance for Floor share in Figure 9

