NBER WORKING PAPER SERIES

PREDICTABLE PRICE PRESSURE

Samuel M. Hartzmark
David H. Solomon

Working Paper 30688
http://www.nber.org/papers/w30688

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2022

We have nothing to disclose. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Predictable Price Pressure
Samuel M. Hartzmark and David H. Solomon
NBER Working Paper No. 30688
November 2022
JEL No. G12,G14,G4

ABSTRACT

We demonstrate that predictable uninformed cash flows forecast market and individual stock returns. Buying pressure from dividend payments (announced weeks prior) predicts higher value-weighted market returns, with returns for the top quintile of payment days four times higher than the lowest. This holds internationally, and increases when reinvestment is high and market liquidity is low. High stock expense firms have lower returns from selling pressure after blackout periods, by 117 b.p. in four days. We estimate market-level price multipliers of 1.5 to 2.3. These results suggest price pressure is a widespread result of flows, not an anomaly.

Samuel M. Hartzmark
Boston College
Carroll School of Management
140 Commonwealth Ave
Chestnut Hill, MA 02467
and NBER
samuel.hartzmark@bc.edu

David H. Solomon
Boston College
140 Commonwealth Ave
Fulton Hall Room 330
Chestnut Hill, MA 02467
david.solomon@bc.edu

Do trades in and of themselves generally move equity prices? This question, while seemingly straightforward, does not have a readily agreed upon answer in the academic literature. Market efficiency, upon which most modern asset pricing theories are built, posits that prices move because of information - either about cashflows, or the discount rates applied to them. Trades themselves, unless they convey information, should not move prices.[1] In tension with this assumption, various papers document that uninformed flows can shift prices. To rule out the possibility that the trading in question represents information, the focus is generally on one-off or surprising events for individual stocks where trade occurs for non-fundamental reasons, like being added to an index (e.g. Shleifer 1986), or mutual fund fire sales (e.g. Coval and Stafford 2007).

One limitation of this approach is that the unusual nature of events, which allows for the identification of uninformed flows, makes it difficult to ascertain whether price pressure is an aberration limited to periods of unusual market dislocation, or a plausible null hypothesis for how trades affect prices in common settings. This leads to the curious outcome whereby academics who *study* price pressure view it as well established, but most others view it as economically unimportant.[2]

Even in models that rationalize the existence of price pressure (e.g. Vayanos and Vila (2021), Gabaix and Koijen (2021)), the prediction is that prices mostly move because of information about a trade, not the trade itself. For most trades, news of the trade occurs at the same time as the trade, so a price change cannot be separately ascribed to the news or the trade. However, if the trade were announced in advance, these models predict the price movement should occur largely when the market learns the news that the trade is coming, with minimal movement at the trade itself. The idea that the trade might still significantly move prices, even if known in advance, represents an extension that is less studied empirically, and less understood from a theory perspective.

We provide evidence for economically sizable price pressure, in a common and recurring setting,

[1]Or only move prices a very small amount. The field has moved past the basic efficient markets hypothesis intuition (e.g. that of Sharpe 1964, Ross 1976, and many others), but the general equilibrium models that follow yield similar intuition as they predict economically negligible responses to price pressure. See Gabaix and Koijen (2021) for quantification of the the demand curve slope in various models.

[2]Gabaix and Koijen (2021) surveyed asset pricing academics as to their beliefs about the money multiplier (the price increase from each dollar of equity purchases), and found a median answer of zero, and among those with a positive answer, the median was an economically tiny 0.01.

for the deepest and most liquid assets, when all of the information about the flows is known weeks in advance. We examine the impact of the predictable reinvestment of dividend payments in creating price pressure on the value-weighted market portfolio. Our identification is simple. When investors receive a dividend payment, some of the money is predictably invested into the market.[3] While the actual purchases (which we do not observe) may contain information when they occur, all of the information necessary to understand the predictable component of flows from dividend reinvestment is revealed weeks prior on the announcement day. By using dividend payments as a proxy for the predictable component of flows due to dividend reinvestment, our analysis examines flows that are predictable in both timing and amount (as in Hartzmark and Solomon (2013)) and occur almost every day. As a result, market participants should understand that these flows are uninformed, and the recurring nature of the event gives liquidity providers the best opportunity to recognize the situation and counteract it. The existence of price pressure in such a setting suggests that trades influencing prices is a general and ubiquitous property of financial markets

We find that when daily dividend payouts are higher, daily value-weighted market returns are significantly higher. The magnitude of this effect is considerable. A one standard deviation increase in payout is associated with higher returns of 3.2 basis points (compared to an unconditional mean return of 4.1 b.p.). Value-weighted market returns are four times higher in the highest quintile of dividend payment days than in the lowest. This is not due to shifts in the denominator of stock market capitalization. Similar effects are obtained with year-by-month fixed effects, or with a measure of abnormal dividends based on the difference from the longer–term average daily dividend. These effects survive controlling for day-of-the-week effects, turn-of-the-month effects (e.g. Lakonishok and Smidt 1988), Fed announcements (Lucca and Moench 2015) and other macroeconomic announcements (Savor and Wilson 2013). Daily value-weighted returns are 11.6 b.p. higher if dividends are in the top 10 days of the past year, and 7.1 b.p. higher if they are in the top quarter. Days with an above average dividend payment account for roughly 70% of the equity premium. We next show that this result holds on average in 58 international markets, which allows us to control

[3]Though clearly not all. For example, some dividends are consumed (e.g. Hartzmark and Solomon 2019, Di Maggio, Kermani, and Majlesi 2020, Baker, Nagel, and Wurgler 2006).

for worldwide daily patterns in returns. If a country has higher dividends on a given day relative to other countries, it tends to have higher returns relative to other countries on that day.

Consistent with price pressure from uninformed demand, market predictability is related not just to the amount of dividends, but their likelihood of reinvestment. Asset managers typically make their largest payouts in the fourth quarter, making them less likely to reinvest dividends at such times (as they need to shortly pay out the cash to their own investors). During the fourth quarter, dividend payments do not predict market returns. During the first quarter, when dividend investment is likely highest (as managers have a long time before needing to pay out again) we find the largest effect. A one standard deviation increase in dividend payments in the first quarter is associated with an increase in market returns of 6.4 b.p. Evidence from the time series is also consistent with the importance of asset managers. The impact of dividend price pressure has increased since roughly 1990, as mutual funds, ETFs and other asset managers have become a larger component of equity holdings. Consistent with liquidity being important for price pressure, we find larger effects when proxies of market liquidity are lower.

Up to this point, we have been agnostic over whether the price impact is permanent or temporary. Examining cumulative returns, we find no evidence of a reversal over the month after dividend payment. If anything, the data suggests a continuation, consistent with some investors reinvesting dividends with a delay. Extending out further adds significant noise to our estimates, and one quarter after payment we cannot reject either a permanent effect or a full reversal. While our setting offers a number of advantages, it does not allow precise quantification of the total extent and the timing of a reversal in the time series.

Our baseline strategy uses stale information and predicted purchases, rather than actual purchases, to ensure that payment amounts do not have any information released on the payment date itself. A potential concern is that our results reflect a delayed reaction to this stale information. To demonstrate this is unlikely to account for our results, we approximate a day's dividend payment using dividend payment from the same calendar day in prior years. After showing this rule explains a significant portion of the variation in daily dividend payment, we show that much of the returns can be predicted using data as old as 10 years prior. Thus the returns are unlikely to reflect a delayed reaction to news from recent dividend announcements.

The analysis of price pressure so far has considered buying pressure, simply due to the clean identification offered by dividend payments. To identify evidence of predictable selling pressure, we examine stock compensation. Some firms predictably have high stock expenses due to paying a significant portion of employee compensation using equity. These employees have incentives to sell this as soon as possible. Many firms allow employees to trade only within a limited window, typically after earnings are announced. This should lead to concentrated selling pressure after earnings announcements for high stock expense firms, but it is predictable and uninformed selling pressure that the market should easily understand and offset. To demonstrate the expense is easily predictable and not related to news in the current announcement, we focus on stock expenses announced the prior quarter. We find that firms with higher stock expense have more negative returns. Firms in the top 5% of stock expense have returns of more than -117 b.p. (t-statistic of 5.50) in the four days after their earnings announcement.

Finally, we perform a back-of-the-envelope calculation to estimate the multiplier of prices with respect to reinvested dividends, similar to Gabaix and Koijen (2021). This represents the magnitude of a market capitalization increase for a dollar of dividends reinvested (i.e. the inverse of elasticity). Our baseline estimate of 0.67 represents a lower bound on the multiplier (assuming 100% reinvestment), and demonstrates that standard models and survey respondents greatly underestimate its magnitude. Estimating reinvestment rates for different investor groups, and combining this with our results gives a multiplier of 1.5 to 2.3. This is somewhat lower than that in Gabaix and Koijen (2021), showing potentially that the size of the multiplier estimated may be sensitive to methodology and which flows are used for estimation. Nonetheless, given the information about the flows in our setting is announced in advance, our results are puzzlingly large under their model.

This paper fits into several strands of literature showing the existence of price pressure. A number of papers have explored reactions to one-off events. When companies are added to an index, this results in an increase in price (Shleifer 1986, Harris and Gurel 1986, Wurgler and Zhuravskaya

2002, Chang, Hong, and Liskovich 2015, Patel and Welch 2017). A related literature on fire sales explores how flows in mutual funds can induce price pressure from predictable purchases and sales of underlying assets (Coval and Stafford 2007, Lou 2012, Frazzini and Lamont 2008). Ben-David et al. (2020) argue that flows from Morningstar ratings changes induce price pressure to individual stocks and influenced the dynamics of cross-sectional style returns. Various papers have analyzed stock-level effects of dividend payouts. Hartzmark and Solomon (2013) show price pressure from dividend-seeking investors in the lead up to a dividend ex-day, and negative returns thereafter, an event that is also both predictable and recurring. Other papers examine stock-level effects when the dividend is actually paid. Berkman and Koch (2017) show positive payment date returns for dividend-paying stocks that have dividend reinvestment plans. Other papers show cross-sectional patterns where stocks owned by mutual funds receive spillover price pressure when the funds receive a dividend or a merger payment (Schmickler 2021, Chen 2022, Kvamvold and Lindset 2018). Most of this evidence posits effects through overlapping holdings (Chen 2022, Kvamvold and Lindset 2018) and/or industries (Schmickler 2021).[4] While these cross-sectional patterns raise the question of whether there may be aggregate time-series patterns, the existence of such a pattern is far from obvious from these results.[5]

Evidence of price pressure for *entire markets* is harder to find. While this paper focuses on equity markets, other papers have explored the influence of supply and demand in the market for bonds (Vayanos and Vila 2021, Greenwood and Vayanos 2014, Greenwood and Hanson 2013, Lou, Yan, and Zhang 2013, DAmico and King 2013), mortgage backed securities, (Gabaix, Krishnamurthy, and Vigneron 2007), and options (Garleanu, Pedersen, and Poteshman 2008). Closest to our paper in the examination of equity markets is Gabaix and Koijen (2021), who use granular instrumental variables to identify shocks to institutions to estimate the multiplier of flows on prices. Parker,

[4]Kvamvold and Lindset 2018 perhaps come the closest to our results. While they show daily price pressure in stocks connected by mutual fund holdings, they show inconsistent effects in S&P 500 stocks (compared with stocks that funds were holding). Given their emphasis on cross-sectional predictions based on mutual funds minimizing tracking error, they make a number of specification choices, which make it not obvious that the value-weighted market return is predicted by daily dividend payouts.

[5]In the return space the average cross-sectional predictor is not actually a good time-series predictor (Engelberg et al. (2021)).

Schoar, and Sun (2020) show mechanical re-balancing by target date funds influences the cross-section of stock returns and overall market dynamics. Da et al. (2018) finds market-wide price pressure in Chile for stocks and bonds following recommendations for pension allocations and Li, Pearson, and Zhang (2020) show that flows based on IPO regulations influence the Chinese market. Evidence on the effects of predictable flows at the market level has not, to our knowledge, been studied. Our paper complements this work by providing evidence of predictable price pressure through a simpler identification channel where the argument for the flows being uninformed is strong, where the flows are predictable in both timing and amount, and the event is a normal day-to-day occurrence. This reinforces the conclusion that price pressure is ubiquitous even at the asset class level, even when the trades are known in advance, for the largest and most liquid stock market in the world.

This notion has significant implications for financial markets. It provides a parsimonious framework for understanding the large price rises of cryptocurrencies, and of "meme" stocks like Gamestop in 2021. It provides simple and intuitive explanations for why order flow matters for currencies, and why the Federal Reserve can influence asset prices in low interest rate environments. It predicts that demographic shifts will affect prices. We discuss these ideas more at the end of the paper.

The results point to the existence of price pressure based only on predictable flows, in both buying and selling directions, at the broad level of the market and for individual firms. This suggests that price pressure should be the plausible default assumption in any scenario when considering how changes in buy and sell orders affect prices. This raises the difficult question of understanding why we observe such buying or selling in the first place. The conceptual difficulty of this task, however, does not justify assuming away the problem by positing that liquidity is infinite.

I. Framework

A. Price Pressure

If more people attempt to exchange twenty dollar bills for strawberries at their current price, textbook microeconomic theory suggests that the price of strawberries will rise. By contrast, if more people attempt to exchange twenty dollar bills for hundred dollar bills, except in unusual circumstances, the price of hundred dollar bills (denominated in twenty dollar bills) will not rise. What, then, should the prediction be when more people attempt to convert their twenty dollar bills into stocks? Is the stock market merely a stream of agreed-upon future cash flows, a sequence of hundred dollar bills of well-understood value? Or is it like strawberries, where the simple act of more people arriving to the market trying to buy is enough to influence its price?

When explaining microeconomics to undergrads, market participants are partitioned into consumers of product who make up the downward-sloping demand curve, and producers who make up the upward-sloping supply curve. The price is given by the intersection of those curves, so an increase in buyers (an increase in demand) or a reduction in sellers (a decrease in supply) leads to a price increase. In the case of stocks, the most common assumption is that they are instead like the hundred dollar bill, and the modeling is rather different. Supply is usually described as the number of shares outstanding, fixed in the short term. This means that the supply curve is a vertical line (see Shleifer (1986) for a discussion of the standard intuition). The (net) demand curve, of buyers minus sellers, is assumed to be horizontal.

This formulation yields intuition about market dynamics from a representative agent under the efficient market hypothesis, namely that the market will bear a net demand of any quantity at the prevailing price, with quantity fixed by the number of shares. But it thereby obscures the effect of different traders' beliefs or liquidity needs. The day-to-day trading of stocks is not represented by interactions of supply and demand, but occurs *within* the net demand curve. This framework muddies the question of whether price pressure is even a coherent concept. Not infrequently, one hears objections to the effect that "one investor is buying, one investor is selling, and the quantity of shares is fixed. How can investors in aggregate have price pressure? Pressure on what?"

This confusion seems to stem more from the textbook model, not the plausibility of the price pressure concept. For example, standard terms like "a shift in demand" are unclear for horizontal demand. For downward-sloping demand, it makes no difference if the demand curve is considered to be transposed *vertically* (i.e. a willingness to pay a higher price for each quantity) or *horizontally* (i.e. a willingness to buy a greater quantity at each price), because the final effect is the same. Indeed, the expressions "demand shifts up", "demand shifts to the right", and "demand shifts up and to the right" are generally used as *synonyms*, meaning that many economists likely have not considered which version they have in mind.

For horizontal demand, this equivalence is removed. An increase in quantity demanded given the price (a horizontal transposition) has no influence. It leads to an identical line, as unconstrained net demand is infinite at the market price, and is zero at a price of one penny more. However, if a shift in demand is interpreted as a shift in the willingness to pay (a vertical transposition), this leads to a shift in price. These two different predictions lack a distinct vocabulary, making it challenging to grasp how the model is meant to work.

An alternative conception, and we feel a more useful one, is to consider demand and supply in terms of people bringing trades to market. At a specific moment in time, this maps closely to the concept of the limit order book (aggregated over all trading venues). For most economic applications, the instantaneous elasticity is not the length of time over which price impact is evaluated, so it is worth extending the concept over the relevant horizon, such as a day or a month, etc. There is no real world analogue to this, but the theoretical goal is to capture investors in the current limit order book along with all those who would enter the market when prices shift. In such a scenario, supply is no longer the number of shares outstanding. Rather, it is the number of shares of a stock that would arrive as sell orders over the horizon of interest if the price shifted. We feel that this formulation yields significantly more intuition, as the number of shares outstanding is not generally a constraining factor in the number of sell limit orders that can be placed.[6] Similarly, demand is

[6] Short sellers can generally borrow shares and generate greater numbers of sell limit orders. Even the company itself can also be reduced to a potential seller or purchaser of shares, rather than the main determinant of supply.

the number of shares sought in buy orders under the same circumstances.

Price pressure in this framework is easily understood, as intuitions about elasticity map into orders submitted by traders. When a sufficiently large liquidity-demanding order is placed, it will eat through the limit orders in the book, and eventually will affect the last trade price. The question is what happens next. If aggregate net demand is perfectly elastic in the standard framework, this is equivalent to saying that this theoretical limit order book over some period of time has *infinite depth at both the bid and the ask.* Put this way, this assumption is far from innocuous. At short horizons we can observe a finite depth of the current limit order book for all securities.

Thus, to motivate perfectly elastic demand, there must be an arbitrage mechanism functioning beyond the current limit order book. The standard assumption is that investors will realize that the price has deviated from fundamentals, and place large trades in the opposite direction to restore the price to its previous level. While this could be the case, perhaps they won't, either due to lacking capital, or thinking the trade isn't worth it, or being uncertain whether the counter-party knows something they don't, or trading in the opposite direction thereby exacerbating rather than correcting the mispricing (e.g. Brunnermeier and Nagel 2004). In any case, the idea of offsetting limit orders to replace price-affecting trades (even if economically intuitive in certain contexts) is an additional assumption layered on top of the mechanics of the limit order book. In other words, price pressure is a general and largely *mechanical* prediction, whereas offsetting trades are an economic prediction arising from particular models of how investors trade.

One of the notable aspects of this framework relative to others considering price pressure is the focus on trading itself as a driver of price movements, separate from the announcement of that trading. This distinction is possible in our settings because the flows we examine are predictable. For example, the index inclusion literature mostly finds larger price effects on the announcement rather than the inclusion itself. For many other shocks, like institutional trades in Gabaix and Koijen (2021), it is not possible to distinguish announcements from trades. Finding that pre-announced events have price impact when they eventually occur provides evidence that the market is unlikely to perfectly offset the impact of trade, and that trade is important even when it is pre-announced.

This simple framework allows us to derive a number of testable predictions. First, buy and sell orders should have price pressure. More people attempting to buy a stock (or the market) will lead to price increases, and more people attempting to sell a stock will lead to price decreases. Second, the magnitude of these effects will increase with the desired amount of stock being purchased or sold. Finally, because price pressure is fundamentally about limited liquidity, we predict that the effects of a given amount of buying or selling will be greater in periods of lower liquidity.

B. Predictability

The key testable prediction of this framework relative to the textbook framework is that predictable uninformed flows are predicted to shift prices. In order to test this prediction a measure of flows is needed that should be considered uninformed under standard theories. Such flows ought to not only *be* uninformed, but should be easily *recognized* as being uninformed by a rational investor.

One way to do so is to focus on the component of flows that are ex-ante predictable long before actual trading occurs. The rationale is analogous to why asset pricing tests use information known prior to an event as a measure of ex-ante predictable returns. Ex-post realized flows likely contain information that could not be known before the trade occurs. By conditioning on the the ex-ante component of information, we do not include the total information known only ex-post that is included in actual flows. In doing so we are conditioning on the information the market had before trading occurred, making it a good proxy for the predictable component of flows that a textbook market would understand to be uninformed.

To illustrate the intuition using our dividend setting, assume there is an announced dividend that will be paid on a future date t. It is likely that the difference between predicted flows (based on the information from the dividend announcement) and ex-post realized flows occuring on date t contains information that shifts prices. For example, assume that when the payment is made more money is reinvested than would have been expected at the time of the announcement. This behavior could indicate positive news which in a rational market could lead to a price increase of the market. Similarly, assume the payout is used to purchase certain stocks and other stocks are not purchased.

This could release positive news for the purchased stocks and negative news for those that were not, leading to cross-sectional differences in returns.

The difference in prediction between our framework and the textbook framework comes from the influence of predictable flows, not realized flows. A test focusing on a realized flow could indeed reflect news in a rational market that was unknowable before the flow occurred. By using only predictable flows based on information released long before trade occurs, the flow cannot reflect news released when trading finally happens and thus in any textbook framework is not predicted to meaningfully influence prices upon payment.

II. Data

Stock return data is from the Center for Research in Securities Prices (CRSP), including returns, dividend amounts and payment dates. We limit our return data to ordinary common shares (codes 10 or 11) listed on the NYSE, NASDAQ or AMEX. For dividend amounts we examine ordinary cash dividends. We examine data from 1926 through 2018. Monthly mutual fund holdings come from CRSP. Stock compensation expenses come from Compustat. International stock returns data comes from Compustat Global.

Unless otherwise noted, when we refer to the timing of a dividend, we mean the payment date. For example, when we refer to a dividend date or a dividend yield, we are referring to date or yield based on the payment date. The major exception to this is when we examine mutual fund dividends, as there is not payment date information for fund dividends, so the ex-date is used.

III. Results

A. Aggregate Dividend Payments and Market Returns

Our first set of tests examines the impact of aggregate dividend payments on market returns. Aggregate dividend payments represent an interesting source of price pressure because their timing is completely predictable well in advance of the date the payments actually occur. The fact that the payments are known in advance, and result in predictable reinvestment, offers clean identification of whether uninformed flows into the market impact prices.

While there is uncertainty in the timing and amount of dividends that companies will declare, this is resolved when the dividend is announced. Once the dividend is declared, companies have a legal obligation to pay it.[7] Moreover, there is usually some delay, on average 21 days, between the announcement of the dividend and the ex-date (the first day when those who buy the share will not receive the dividend). The ex-date is when any implications of receiving a dividend payment, such as tax consequences, are resolved. We focus on the payment date, which on average occurs 22 days after the ex-date, and so is on average 43 days after the initial announcement. The payment date is the date when cash is disbursed, and lacks economically meaningful news or tax implications.

As a consequence, dividend payment amounts are predictable well before they occur. Market-wide dividend payments also exhibit substantial day-to-day variation due to idiosyncratic differences between individual companies' payment dates, as well as large differences in firm size. Further, while there is some seasonality in dividend payments, over 90% of trading days involve a dividend payment.

An increase in cash must end up somewhere - it can be consumed, left in cash or the money market, or invested into risky securities. We focus on the latter as a potential driver of returns. Hartzmark and Solomon (2019) show that investors of any type, including institutional investors and mutual funds, do not generally reinvest dividends into the securities from which they came. Given the paucity of same-stock reinvestment, this raises the question of whether the remainder of the dividend amounts are invested into other stocks, which may be inducing price pressure. To consider the broadest and most liquid set of test assets, we take as our main dependent variable the value-weighted market portfolio from CRSP.

The final question is the timeline under which people use a dividend payment to purchase securities. Dividend reinvestment likely occurs with different timing for institutional reasons, differences in attention, and differences in strategy. If investors trade once their cash balances update, differences in reinvestment timing could be due to differences in when the cash appears in an account.

[7]Even in the event of corporate bankruptcy before the payment date, shareholders are unsecured creditors for the amount of the dividend.

For example, some institutions update balances throughout the day (allowing for trade on t=0 and after), though many banks clear each day's deposits that night (allowing for trade on t=1 and after).[8] Further, not every attentive investor has a strategy that trades every day. In addition, some investors may not be attentive to a dividend payment, and reinvest only when the higher cash balance comes to their attention. We lack details for the myriad of institutions, time periods and strategies involved, so while we predict there will be reinvestment on the payment date and the days immediately following it, we are agnostic as to which specific day should have the largest effect.

Table I examines the effects of daily payment yield on value-weighted market returns. The independent variable is the dividend payment yield on a payment day t, which is the daily total dividend payments, divided by the previous day's total market capitalization. If a dividend payment is invested on the payment date and induces price pressure, the market return on date t would be predicted by the dividend payment on date t, while if it is invested the day after the payment date, the market return on date t would be predicted by the dividend payment on date t-1. In column 1, we consider dividend payments at various lags from zero days to four days. This shows a coefficient on payment yield of 55.76 for day t (t-statistic of 1.74), 60.04 for day t-1 (t-statistic of 2.73). Days t-2 and t-3 are smaller and statistically insignificant, though day t-4 is also larger and marginally significant (t-statistic of 1.89). In terms of magnitudes, a one standard deviation increase in payout yield (.0003176) predicts higher returns of 1.8 b.p. on day t, and 1.9 b.p. on day t+1 .

Payout yield uses the market level as a denominator, so we wish to make sure that the effects capture variation in dividends, not variation in the market level. We deal with this problem in several ways. As a first step, in column 2 we include a year-by-month fixed effect. If the concern is that the price-to-dividend ratio predicts long-term periods when the market has a higher or lower return, a year-by-month fixed effect will control for general differences in such an average. This also controls for specific calendar effects (e.g. the January effect in Rozeff and Kinney Jr (1976) and Thaler (1987)). The inclusion of this fixed effect increases the magnitude and significance of the

[8]There is likely heterogeneity across institutions as to when in a day a payment appears in an account, even if the payment appears on the same day. Further, we expect that in the early period of our sample there was likely more heterogeneity and longer lags for a payment to clear, though we lack documentation from this period.

coefficients, to 74.85 on day t (t-statistic of 2.26), 71.98 on day t-1 (t-statistic of 3.10), and 66.66 (2.32) on t-4 (with days t-2 and t-3 also being slightly larger, though insignificant).

Given this daily pattern, in columns 3 and 4 we focus on the main variable we will use, the cumulative dividend payment yield on the payment date and the day before.[9] Column 3 finds similar effects to column 1 - dividend payout from day t and day t-1 positively and significantly predicts value-weighted market returns on day t (coefficient of 59.50 and t-statistic of 3.32). When year-by-month fixed effects are added in column 4, the results are again stronger, with coefficients increased to 67.07 and a t-statistic of 3.47. A one standard deviation increase in payout yield over days t and t-1 (.0004711) predicts higher daily returns of 3.2 b.p. on day t.

In columns 7 and 8, we perform the same tests using the equal-weighted market portfolio. In general, the effects are slightly larger. These results seem to reflect a broadly similar picture to that using the value-weighted market.[10]

If the returns reflect price pressure from investing a dividend payment, then there is less reason to expect positive return predictability before the payment is made.[11] Thus the period directly prior to dividend payment offers a placebo where price pressure does not make strong predictions about returns. To examine this period, Table I columns 5 and 6 examines the influence of future dividends (on day t+1 and t+2) on current (day t) returns. The coefficients are economically small and insignificant. This is consistent with dividend payments only influencing returns after they occur.

While year-by-month fixed effects remove slow-moving trends in the market price, it is possible that the effects are driven by day-to-day fluctuations in the denominator. We employ a number of

[9]We find the t-4 effect likely reflects delays in receiving the dividend in the early period of our sample, so we mostly focus on the first two days where the predictions are clearest. Running the Column 2 specification using data prior to 1960 yields a coefficient on Mkt Div Pay[t-4] of 79.9 (t-statistic of 1.83), post 1960 this coefficient is 45.73 (t statistic of 1.58), and post 1980 is 32.23 (t-statistic of 0.73).

[10]For the remainder of the paper we focus on value-weighted returns, but similar results are obtained using equal-weighted.

[11]In principle, an attentive investor could trade before the payment date, since cash is only required at settlement. The settlement date occurs two or three days after the purchase date, depending on the sample period, so one could buy earlier and have the cash arrive before settlement. An investor with access to leverage could borrow against a future dividend and trade as soon as it was announced. They would need a processes in place to trigger a trade at this earlier date, so while this outcome is possible, it is considerably less likely.

empirical strategies to ensure that this is not the case. First, we scale the dividend payment amount not by the prior day market value, but with further lags. Table I Panel B repeats the analysis using the prior month's, quarter's and year's market capitalization. All of the regressions yield positive and significant coefficients with t-statistics over 3 and a similar economic magnitude.

While day-to-day fluctuations in the price cannot explain the pattern, there may be an unaddressed concern with using price in any manner. To address such a concern, in Table II Panel A we examine a measure of dividend payment that does not employ market prices at all. Specifically, we construct an abnormal dividend yield which is dividends on days t and t-1 divided by the average daily dividend paid over the prior year.[12] When market returns are regressed on abnormal dividends, the results are positive and significant at the 1% level, when returns are value-weighted (columns 1 and 2) or equal-weighted (columns 3 and 4), and regardless of whether year-by-month fixed effects are included. Taking the value-weighted estimate in column 2, a one standard deviation increase in the dividend yield is associated with an increase in expected returns of 2.7 b.p., similar to the 3.2 b.p. found from the same specification normalizing by market capitalization.

Figure 1 examines how value-weighted returns vary with this abnormal dividend.[13] We split days into quintiles based on the abnormal dividend and take the average of the value-weighted market returns in each quintile. The graph shows a monotonically increasing relationship between payment yield and returns. The returns in the lowest two quintiles are about 2 basis points. For the next two quintiles this value increases to roughly 4 basis points. The top quintile of abnormal dividends is associated with value-weighted market returns of about 8 basis points.

The graph suggests that the largest effect is concentrated in more extreme dividend payment dates. To test this, and to use a simple rule that is easy to interpret and ex-ante tradable, we examine whether or not a given day's dividend payment is high relative to the payments made in the previous 252 days. Specifically, we regress market returns on dummy variables equal to one if the cumulative dividend payment on day t and t-1 is in the top 2 weeks (10 days), quarter (63

[12] We calculate this average using the trading days from t-20 to t-272. The average year contains 252 trading days and we skip about a month (20 trading days) to ensure the denominator excludes any recent market information.

[13] In addition to not being influenced by daily variation in the market price, the abnormal dividend is also largely independent of long run changes to the market dividend yield (e.g., Fama and French 2001).

days), third (84 days) or half (126 days) of days in the past year, and equal to zero otherwise. These regressions display the magnitude of the return effect from the regression coefficient, as it represents how different the return on these extreme dividend payment days is relative to the other days.

Column 1 of Table II Panel B shows that with no controls, the top ten days have higher value-weighted returns by 9.8 b.p., with a t-statistic of 2.39. In column 2, with a year-by-month fixed effect, this increases slightly to 11.6 b.p. In the specifications including year-by-month fixed effects, the top quarter of days experienced returns 7.1 b.p. higher (column 4, with a t-statistic of 4.27) and the top third of days experienced returns 5.7 b.p. higher (column 6, with a t-statistic of 3.71). Column 8 shows that days with above median dividend payments experience returns 4.7 b.p. higher with a t-statistic of 3.15. The constant indicates that below average days earn returns of 1.7 b.p. This means that above median days earned returns of 6.4 b.p., nearly four times higher than the below median days.

Figure 2 demonstrates the cumulative magnitude of these effects. We plot the cumulative returns to a \$1 investment in 1926 for a strategy that invests in the value-weighted market portfolio only on days where the dividend payment yield on day t or t-1 is above the median level (red line) compared with the same investment only on days below the median (blue line). Despite the fact that the blue line and the red line have about the same number of days of exposure to the market (and the associated equity premium), the high payment day red line cumulated to \$512 by the end of the sample, compared to only \$10 for the low payment blue line.

An alternative way to quantify the magnitude of the effect is to ask how much of the equity premium is earned on days with high dividend payments. We explore this question in Table III. The first row shows the average annual return to investing on days with high dividend payments, the "Top" columns, and the returns to investing on days with low dividend payments, the "Not" columns. Column 1 shows that investing in the top 2 weeks earns an average annual return of 1.4%, while investing on all other days earns a return of 10.2%. This means that the top 2 weeks per year (3.8% of trading days) earn about 12% of the equity premium. This ratio of 3.13 indicates that the top 2 weeks account for roughly triple the amount of the equity premium compared to the number of days they are in the market. The final two columns, examine being above or below the median dividend payment. The above median days experience average annual returns of 7.7% while the below median days experience average returns of 3.5%. This indicates that the above median days experience roughly 69% of the equity premium, even though they only account for 50% of the days.

A question that naturally arises is whether the effects represent price pressure in the companies that paid the dividend, or elsewhere in the market. We note that Hartzmark and Solomon (2019) document that, when looking at mutual funds and institutions, dividends are rarely reinvested into the stocks that they came from. If this is the case, we predict that price pressure should not be limited to the firms that actually paid the dividend.[14]

In Table IV, we consider how the effects are different for the firms paying dividends that day versus other firms, and find the strongest effect for firms that do not pay the dividend. When examining non-dividend-payers (columns 1-4), the results are similar in magnitude and significance to Table I. When only firms payingdividends on those days are included, the results are smaller - around two thirds as large for value-weighted returns (columns 5 and 6), and one-quarter to one-third as large, and statistically insignificant, for equal-weighted returns (columns 7 and 8).

Another potential confounding effect is that returns vary over the week (with the lowest returns on Monday and the highest on Friday), as well as over the month (with the highest returns at the turn-of-the-month, as in Lakonishok and Smidt 1988). If dividend payments are more likely on certain days of the week or periods of the month, payment yield might proxy for such effects.

Table V Panel A examines such effects and finds they are unlikely to account for the results. Columns 1 and 2 include day-of-the-week fixed effects, columns 3 and 4 include turn-of-the-month fixed effects and columns 5 and 6 include both.[15] All of the coefficients are positive and significant, suggesting that these calendar patterns do not account for our result.

Another predictor of daily market returns is macroeconomic announcements. Perhaps the most

[14] It is difficult to predict the relative magnitudes of effects for paying and non-paying firms. Even if the majority of money is invested outside paying firms, there are many more non-paying firms. Thus we predict *some* price pressure for non-paying firms, but the relative amount of price pressure for the two groups is ultimately an empirical question.

[15] Following Lakonishok and Smidt 1988, our turn-of-the-month dummy variable is equal to one for the last trading date of the month and the first three trading dates of the month.

attention has been given to Federal Open Market Committee (FOMC) announcements, when Lucca and Moench 2015 argue a large fraction of the equity premium is earned. Savor and Wilson 2013 initially documented this FOMC effect, as well as similar positive returns coinciding with other macroeconomic announcements. If dividend payments coincide with FOMC announcements, or macroeconomic announcements more broadly, such an effect could account for the results.

Table V Panel B examines FOMC announcements. Our FOMC announcement data runs from 1988 through 2019, so the first two columns repeat our baseline analysis over this time period, and show stronger results than for the whole sample.[16] Columns 3 and 4 include dummy variables for FOMC announcement days, and finds coefficients that are roughly unchanged. This suggests that the FOMC announcement effect is largely unrelated to the effects we document.

Table V Panel C extends to other macroeconomic announcements including CPI, PPI, Initial Claims, Employment and GDP. This data runs from 1994-2018, so the first two columns repeat our baseline in this period, and show stronger results than using the whole sample. Columns 3 and 4 include dummy variables for these days, while columns five and six also include FOMC announcements. The results are similar across these columns, which suggests that dividend payments are not proxying for the influence of macroeconomic announcements on returns.

B. Delayed Reaction to Announcement

The key assumption underlying our analysis is that the information from the dividend is announced early enough that an efficient market should easily incorporate information from the predictable component of flows prior to the payment. The gap between the announcement and payment averages over 40 days, which should be long enough to incorporate this news. One potential concern is that there could be exceptions that are important in predicting returns where the time is too short for the market to respond before payment. Such instances are rare, as the 1st percentile of the time between announcement and payment is 17 days and the 25th percentile is 29 days.

To show that the effect is driven by standard observations with a significant period of time

[16] Data on all of the macroeconomic announcements are the same as used in Neuhierl and Weber 2019.

between the announcement and the payment, we repeat the analysis restricting the sample to observations where the gap is more than 17 days or 29 days.[17] Table VI Panel A shows similar results to our baseline with these restrictions.

Another potential concern is that an issue related to the ex-date is contaminating the payment date effect. Time gaps between the days are again fairly long, with a 1st percentile gap of 9 days and a 25th percentile of 17 days.[18] When we include only the days with gaps above these cutoffs, in Table VI Panel B, results are again similar to our baseline analysis. These results together with those in Panel A suggest that any information released on the announcement date (or any effects related to the ex-date) should have been incorporated into market prices long before the payment date, and should not be able to predict payment date returns under standard theories.

Thus, if the payment date return patterns we find are related to news, it must be that there is news released upon the dividend announcement, but this news is ignored until dividend payment. Given this is a standard event and there is a significant amount of time between the announcement and the payment date, behavioral explanations such as limited attention would be needed to explain such an occurrence. This is complicated by the fact that there tends to be significant attention given to a dividend announcement and to the ex-date, but minimal attention given to a firm on its payment date. Thus ex-ante it seems unlikely that under a limited attention explanation, the payment date is when investors finally pay attention to the previously released news. Nonetheless, the prior results do not rule out such a possibility.

To test whether such an explanation could account for our results, we predict dividend payments using data from long before the most recent dividend announcement. This means that the predicted dividend amount that we are sorting on lacks *any* news contained in the most recent dividend announcement. In general, firm level dividends are sticky in terms of both amount and timing at annual frequencies (Hartzmark and Solomon 2013). This means that past years' dividend information is likely a reasonable proxy for current information. To proxy for the current day's

[17] Announcement date information is only included in CRSP starting during 1962, so this analysis is restricted to 1963 and later when the data is available.

[18] Ex-date information is available for the full sample, so this analysis does not restrict by year.

dividend yield, we use the dividend yield from a prior year on the same calendar date. We examine lags between 1 and 10 years. For example, with a one year lag, if we are examining returns on March 25, 1982 we use the dividend payment yield from March 25, 1981. We do so all the way to 10 years which would predict returns from March 25, 1982 using the payment yield from March 25, 1972.

Before examining whether the measure predicts returns, it is worth asking if such a crude proxy is a meaningful predictor of the current year's dividend payment. To do so, we regress the current dividend payment yield from year y on the lagged payment yield from the same date in year y-L. To capture any longer term predictability we also include a year by month fixed effect. We repeat this exercise using a one year lag and continue to a 10 year lag.

Table VII Panel A shows that this simple rule yields significant predictability. Examining the one year lag in Column 1 gives a coefficient of about 0.8 with a t-statistic of over 200. In terms of evaluating the predictability though, the important statistic is the within R^2, which indicates how much of the variation this lagged value predicts after removing the influence of the fixed effects. We find that this simple rule using data that is a full year stale yields an adjusted R^2 of 0.719. Given this simple proxy is able to account for such a large fraction of the day-to-day variability in the dividend payment, we feel confident using it as a daily dividend proxy in the analysis that follows.[19]

Given the strong predictability of the lagged dividend measure, in Panel B we regress returns on day t in year y on the two day dividend payment yield from day t year y-L, where L is the number of years lagged indicated by each column. Examining Column 1 we see a coefficient of 69.69 with a t-statistic of 3.64. Days without a payment yield, which occur if the day on the prior year occurred on a day without trading, such as a weekend, are dropped. Thus Panel C conducts our baseline analysis on the subsample of days with dividend information for comparison to the results in Panel B. For a 1 year lag the sample loses about 6,000 days and on this subsample our baseline

[19] We note that this is an intentionally simple rule which could be improved upon in a variety of dimensions. It likely would be possible to go to the company level and better predict patterns, such as firms that prefer certain days of the week, or which corrects for weekends, or any number of improvements. The goal of this analysis is to show how predictable these values are using the most simple rule possible. Thus the coarse measure here offers the best illustration of the predictability of these flows, rather than a better, but significantly more complex measure.

regression has a coefficient of 92.24 with a t-statistic of 3.79. Given we are using a noisy proxy of daily dividend payment we would expect some attenuation bias. Nine of the 10 the point estimates in Panel B are lower than the point estimate in Panel C, consistent with this prediction.

We find strong results using data as much as 10 years old. With a 10-year lag we see a point estimate of 58.85 with a t-statistic of 4.19. This suggests that these dividends are highly predictable and that a rational market should have no problem understanding the flows. It also makes it quite unlikely that the returns reflect a response to information from a recent dividend announcement.

C. International Evidence on Dividend Payments and Market Returns

Table VIII switches to an international setting, and considers how daily dividend payments predict market returns across 58 international markets.[20] These tests examine predictable price pressure while abstracting away from a number of possible institutional details that are unique to the US. We again focus on the dividend payment on day t and t-1. While the payment mechanisms will vary from country to country, there is a reasonable assumption that, payment dates and/or the day after should roughly correspond to when investors can use the payment to purchase securities. To the extent that institutional details lead to mismeasurement of reinvestment, or that the data is less accurate, we expect this to bias against finding a result.

In column 1 of Table VIII, with no controls, we find that a higher dividend payment yield is associated with higher returns, with a coefficient of 20.30 and a t-statistic of 2.19. Column 2 adds country-by-year-by-month fixed effects, and the coefficient increases to 30.21 with a t-statistic of 3.59. The economic magnitudes of these effects are somewhat smaller than in the US. A one standard deviation increase in payment (0.000389) predicts higher returns in column 2 by 1.2 b.p.

This panel setting means that we can compare returns across countries on the same day based

[20]These 58 markets are the subset with sufficient data from the 125 markets in Compustat Global. A given market and day is included when it contains at least 100 stocks with prices. The countries in the analysis are: Australia, Austria, Belgium, Bangladesh, Bulgaria, Bermuda, Brazil, the Cayman Islands Chile, China, Croatia, Cyprus, Czechia, Denmark, Egypt, Finland, France, Great Britain, Germany, Greece, Guernsey, Hong Kong, India, Indonesia, Ireland, Israel, Italy, Japan, Jordan, Korea, Kuwait, Luxembourg, Malaysia, Mexico, the Netherlands, Nigeria, New Zealand, Norway, Pakistan, Peru, the Philippines, Poland, Romania, Russia, Saudi Arabia, Singapore, Spain, Sri Lanka, South Africa, Sweden, Switzerland, Thailand, Turkey, Taiwan, Ukraine, United Arab Emirates, and Vietnam. Data cover 1986-2017.

on the difference in countries' dividend payments. Column 3 adds a day fixed effect, and finds similar magnitudes to column 1, but with a larger t-statistic. Column 4 includes date and country-by-year-by-month fixed effects, and the significance increases again, to a t-statistic of 4.50.

D. Returns based on reinvestment intensity

Prior research and results in Section III.H suggest that retail investors rarely reinvest dividends. Thus, it is likely that the returns we document are largely related to the reinvestment rate of professional investors. If a fund is about to send cash to investors and receives a dividend from its holdings, it may not be worth it to pay the transaction costs to hold the new shares only for a short time. As a consequence, reinvestment rates are likely to be lower as the fund's own dividend payment draws near. In the US, funds and ETFs are subject to the pass-through rule, whereby to avoid paying corporate income tax at the fund or ETF level, all dividends and realized capital gains (minus fund expenses) must be distributed to the fund's investors by the end of the year. This leads many funds to wait until the end of the year to pay out most of their dividends.

Figure 3 Panel A graphs the average proportion of fund annual payouts distributed each quarter.[21] Specifically, the sum of mutual fund payouts that quarter is divided by the sum of mutual fund payouts that calendar year and averaged across years. Consistent with the intuition above, the graphs indicate that a significant fraction of mutual fund payouts, nearly two thirds, occur during the fourth quarter. The next panel graphs the payouts by month, and shows that much of this pattern is driven by December payouts, which constitute roughly 60% of the annual payout.

If money managers are less likely to reinvest when they are about to send cash to investors, and if they are responsible for a significant portion of the price effects, then the effects should diminish prior to these large payouts. In other words, a given dollar of dividends paid will lead to less reinvestment, and less price pressure. Given the pattern in dividend payments, we would expect the lowest reinvestment in prior to the large December payouts. This suggests the lowest price effects in Q4 and December in particular. On the other hand, once the next year arrives, given the long

[21]Mutual fund data does not have information on payment dates, so this is based on ex-dates. These payouts include both capital gains and dividends combined, though dividend payments alone display a similar pattern.

time until the December payouts, we expect the most reinvestment at the beginning of the year.

Table IX explores whether this is the case. It repeats our baseline analysis, but examines how the coefficient on dividend yield varies by quarter and includes a year-by-month fixed effect to ensure the regressions are not picking up variation in the level of returns in a given calendar month. Column 1 shows that the magnitudes of the coefficients are consistent with the price pressure predictions. The weakest effect is in the fourth quarter when the largest payouts occur. We find insignificant coefficients that are close to zero during this time. Also consistent with the predictions, the largest effect is during the first quarter, which is the farthest away from the large end of year payments. While the graph of mutual fund dividend payments suggests a strong quarterly pattern, the monthly graph suggests that much of this pattern is specifically due to December. To test this, we split the effects of dividend yield into December versus all other months. Consistent with this, we find no significant effect of dividend payment on market returns in December.

In addition to patterns within the year based on payout frequency, there have been large time series shifts in the popularity of investment vehicles like ETFs and mutual funds. Such products seem particularly likely to reinvest dividends, so one might expect the patterns we observe to become more pronounced as these products become more popular. Figure 4 suggests a significant increase in the effect coinciding with the rise of mutual funds and ETFs. The figure shows regression coefficients of the market return on the dividend payment yield on day t and t-1, conducted separately each decade. For example, the point graphed for 1970 is the regression coefficient using daily data from 1970-1979. Prior to the 1990s, the dividend payment coefficient was always positive, with a mean of about 50 and a maximum of 83. The three decades post 1990 represent the three largest coefficients in our sample, each greater than 100. These results further support the interpretation of the main results being due to price pressure, as the size of the effect from a dollar of dividends varies with the reinvestment rate. Meanwhile, it also confirms the importance of financial intermediaries like mutual funds, who are responsible for much of the actual dividend reinvestment in the market.

E. Returns based on market liquidity

An additional prediction of price pressure is that its impact should be larger when there is less liquidity. Similar to Nagel (2012), we proxy for liquidity using measures of market volatility based on the VIX index, which uses implied volatility from S&P 500 options.[22] The VIX data we examine begins in 1990, so we also use the news VIX index developed in Manela and Moreira (2017) which is available at the beginning of our sample. VIX is a noisy proxy for marketwide liquidity that likely captures other market dynamics. To demonstrate robustness using an alternative we also construct a marketwide measure of daily Amihud illiquidity (Amihud 2002).[23] While marketwide liquidity is an elusive concept that is difficult to measure, we think that showing the results using two separate measures, each with their own shortcomings, should assuage concerns that idiosyncracies of the measure drive the result. With that said, given the difficulty of concretely measuring marketwide liquidity, we think the results should be taken as suggestive, but not conclusive.

Table X examines how the impact of dividend payments interacts with levels of the VIX and News VIX in predicting market returns. We augment our baseline regression by including VIX and also an interaction of VIX with the dividend payment yield. Column 1 examines the coefficient on this interaction term and shows a significant coefficient of about 30. Column 2 shows a similar coefficient of about 24 after controlling for year-by-month fixed effects. A one standard deviation increase in the VIX of 7.8, thus leads to an increase in the sensitivity of dividend payment of 187. This suggests that in times when VIX is higher, the impact of dividend reinvestment is also higher.

A potential concern with these results is that they only cover the period from 1990 onwards and thus may be driven by a relatively short sample. In columns 3 and 4 we use the news VIX index to explore the entire sample. All of the major patterns are the same. The 11.5 coefficient in column 4 suggests that a one standard deviation increase in the news VIX of 4.6 coincides with an increased

[22]Theoretically the VIX index should be strongly correlated with the returns to liquidity provision (e.g. Nagel (2012); Gromb and Vayanos 2002; Brunnermeier and Pedersen 2009; Adrian and Shin 2010; Ang, Gorovyy, and Van Inwegen 2011; Ben-David et al. 2011).

[23]We follow the steps in Amihud (2002) and Amihud and Mendelson (2015) in constructing the measure. Specifically, we take the measure at the individual stock level, drop the highest 1% of observations per day, and value weight the remaining observations. We multiply the measure by 10^7.

sensitivity of dividend payment of about 53. This is about 80% of our baseline coefficient of 67.

Columns 5 and 6 finds similar results using Amihud illiquidity. The coefficient on the interaction term of 18 implies that a one standard deviation increase in illiquidity of 3.1 is associated with a higher coefficient on dividend payment of about 56. This is nearly identical to the estimate using the news VIX. These results further support the interpretation of price pressure, with each dollar of dividends having more effect in periods when there is less liquidity.

F. Long Horizon Returns

Returns around dividend payment dates offer sharp tests of the price pressure prediction because of the clean variation provided by daily payments. Nonetheless, these may not reflect the full extent of price pressure, as some investors may reinvest dividends with a delay. Longer horizons also test whether these returns quickly reverse afterwards.

While the effects are not as tightly identified from timing, we consider how daily payment yields predict returns at horizons out to a quarter. To do so, we regress cumulative returns at different horizons on a dummy for whether day zero was above the median dividend payment, plus a constant.[24] We conduct this regression separately for each horizon, beginning with just the payment date alone, and extending to the cumulative return from the payment date to 60 trading days after the payment date, and graph the coefficient on the dummy variable. Thus each point can be interpreted as how much higher the cumulative return was by that point after a day with above median payments. We run versions of this regression with no controls, as well as versions including year fixed effects to control for time trends and daily future dividend payments (during the period of the cumulative return) to ensure we are not capturing price pressure from future payments.

Figure 5 Panel A graphs the results without controls and Panel B includes controls. Each point is the regression coefficient on the payment yield using the cumulative return from t=0 to the the x-axis value. The first week shows similar results of strong price pressure of about 10 b.p.. After this, the pattern is noisier. With that said, the effect increases over roughly the first month,

[24]We examine abnormal dividend yield to define the dummy variable so as to not induce a positive relationship between scaled price measures and longer period returns in the specifications without controls.

peaking around $t+20$. Taking the point estimates seriously suggests that reinvestment continues at a lower rate in weeks two through four. This may arise if some investors are inattentive about the timing of dividend payments (consistent with the observation in Hartzmark and Solomon (2020) that dividend payments are often not flagged in brokerage platforms). The noisiness of these results makes it difficult to distinguish between additional buying pressure, or a lack of meaningful returns.

Perhaps most interestingly, the results show no evidence of a short-term reversal in the first month, even in point estimates. Before seeing these results, a plausible hypothesis might be that the price pressure we observe quickly reverses, having minimal impact on any long term measure. None of the regressions present evidence consistent with a short-term reversal.

Extending further in the future the noise increases substantially. By the end of roughly a quarter (60 days), we are unable to reject either the hypothesis that the effect has totally reversed, or the hypothesis that it is permanent. Our point estimates suggest that this is not an effect that quickly reverses. Unfortunately, we lack the statistical power to measure whether the effect has a permanent impact or reverses over longer horizons. While cross-sectional reversals have been studied in a variety of contexts, much less is known about the speed and extent one should expect entire markets to reverse price pressure. We think this is an important avenue of study for future research.

G. Stock Compensation Expense and Predictable Selling Pressure

The role of price pressure is not conceptually limited to the arrival of buy orders, but ought to apply to predictable selling pressure as well. To demonstrate this we switch to the perspective of selling pressure at the individual stock level.

In order to examine a setting where there is a predictable and sizable increase in selling that does not reveal information, we examine stock compensation. Many companies pay considerable stock-based compensation to employees - not only senior executives, but often lower level employees as well. This is a core component of how these firms conduct their business, so such firms predictably expense large amounts of stock. To prevent insider trading, companies often have blackout periods during which employees cannot trade the company's shares. Employees are usually allowed to trade immediately after the earnings announcement, when presumably all (or most) material information has been disclosed to the market, and insiders are (hopefully) not at an advantage.[25]

For employees receiving a significant portion of compensation through stock grants, standard portfolio theory suggests that they should diversify to other assets to minimize idiosyncratic risk from large holdings in a single firm, and because they already have a considerable exposure to the firm's returns through their labor income (Malmendier and Tate 2005, Cohen 2009). Both effects mean that employees should want to sell shares they receive as quickly as possible, which predicts that when these blackout periods end many employees will predictably sell shares. If such sales create predictable price pressure, returns should be negative just after earnings announcements for these high stock expense firms

We test this by examining returns around earnings announcements from 2002 to 2020.[26] The main variable of interest is stock expense as a proportion of market capitalization one day before the announcement. Similar to DellaVigna and Pollet (2009), we define "announcement day," which we refer to as t=0, as the first trading day when the announcement can affect prices.[27] We examine characteristic-adjusted returns which take a firm's return and subtracts the returns of a portfolio matched on quintiles of market capitalization, book-to-market and momentum (Daniel et al. 1997).

In Figure 6 Panel A, we examine how returns after the earnings announcement vary with the level of stock expense. We graph a local linear plot of the characteristic adjusted returns from t+1 to t+3 after the announcement on the y-axis (skipping t=0 to avoid reactions to announcement news for now), relative to the level of stock expense on the x-axis. The graph illustrates a strong negative relation between the two. Companies with expenses of about 0.1% of market cap experienced negative returns of about 25 b.p., those with expenses of 0.2% experienced about -80 b.p., those with expenses of 0.3% experienced about -120 b.p., and those with expenses of 0.4% experienced

[25]Bettis, Coles, and Lemmon (2000) report that, as of 2000, 78% of companies had explicit blackout periods. Most policies only allowing trades shortly after earnings announcements. Jagolinzer, Larcker, and Taylor (2011) reports that the mean blackout period is lifted 0.8 calendar days after the earnings announcement, with a 25th percentile of zero days and a 75th percentile of one day.

[26]We combine information from both Compustat and IBES. 2002 is the first year with meaningful stock compensation data. There is sparse data in 2001 (results including 2001 are similar) and no data prior.

[27]For example, if a company announces earnings before market open, t=0 is defined as that day. If a company announces earnings after market close, t=0 is defined as the next trading day.

returns of nearly -150 b.p. Consistent with the predictions of price pressure, the negative returns experienced by firms with stock expenses increases with the level of stock expense.

Figure 6 Panel B explores the cumulative characteristic-adjusted returns based on the level of stock expense for the announcement day through twenty days after the announcement. By t=3, companies above the 95th percentile in stock expenses experienced abnormal returns of about -120 b.p., firms in the 90th to 95th percentile experienced about -90 b.p. returns and firms in the 80th to 90th percentile experienced -40 b.p. returns. Firms with any expenses underperform relative to firms with none who experience positive returns (the earnings announcement premium (Frazzini and Lamont 2007)). Looking further into the future there is evidence of a partial reversal with moderate groups reversing after about two weeks. For the two highest groups the underperformance lessens, but is still around -50 b.p. four weeks after the announcement.

A potential concern is that these results do not reflect predictable selling pressure, but rather capture correlations with earnings announcement news. This could reflect fundamental information from the announcement, be it earnings information, information from the amount of stock expensed, or some other information in the report. However, if our measures capture characteristics of how firms generally conduct their business, then most of the information needed to predict price pressure should be available well before the announcement itself.

To demonstrate the extent to which stock expense represents predictable price pressure and to rule out such information effects, in Table XI we examine this pattern using stale data available before the current earnings announcement. The table shows regressions of characteristic-adjusted returns over different time periods after earnings announcements on dummy variables for the level of stock expense. In Panel A we sort on stock expenses announced in the prior quarter's earnings announcement.[28] If the relevant information was released months beforehand, it becomes difficult to explain returns as any kind of response to the current earnings information.

Column 1 shows returns from t=0 to t+3 for each category of stock expense relative to firms

[28]For this lagged measure we use annual expense announced the prior quarter. While quarterly stock expense (which we use in the contemporaneous analysis) is the best measure of the likely amount of selling for that particular announcement, the annual stock expense provides a better proxy for the average quarter. Results are substantially similar if we use quarterly stock expense lagged four quarters.

with no stock expenses. Firms with the highest stock expense (above the 95th percentile) had -117 b.p. returns (t-statistic of -5.50), firms with expenses between the 90 and 95th percentile had returns of -77 b.p. (t-statistic of -4.83) and firms in the 80th to 90th percentile experienced returns of -33 b.p. (t-statistic of -2.93). Firms with positive, but lower, stock expenses (below the 80th percentile) still underperformed firms announcing earnings by 18 b.p. (t-statistic of -3.69).

Another potential concern is that the results represent a delayed reaction to earnings news or a behavioral response to information in the announcement that could have been predicted prior (e.g. the post-earnings announcement drift in Bernard and Thomas 1989). To demonstrate the results are not driven by announcement day effects, Column 2 examines returns from $t{+}1$ through $t{+}3$ thereby skipping the announcement day. We expect $t{=}0$ to also reflect price pressure, particularly since our $t{=}0$ definition often coincides with the calendar date after the announcement, thus results should be smaller beginning at $t{=}1$ under a price pressure explanation. To control for a delayed reaction to the response to earnings news, we include as a control the $t{=}0$ characteristic-adjusted return, split into positive and negative returns. The results in column 2 are smaller (consistent with missing some selling pressure), but display a similar pattern for firms with significant stock expense.

Next we examine whether there is a reversal. Columns 3 and 4 repeat the analysis for returns on days $t{+}4$ to $t{+}10$, and $t{+}4$ to $t{+}20$, respectively. For returns up to $t{+}10$, there are significant positive point estimates for firms above the 80th percentile of expense, consistent with some reversal. Extending to $t{+}20$ in Column 4, each estimate has a positive estimate, with significant estimates for firms above the 90th percentile of expenses. In each instance the point estimate is below the returns from $t{=}0$ to $t{+}3$, suggesting a partial reversal of the effect over the next 20 days.

These results use lagged data to demonstrate the extent of expense predictability, but the actual contemporaneous stock expense (which is publicly known for the periods examined) is a better reflection of the amount of stock currently sold. Panel B repeats the analysis using the current expenses. As predicted, there are modest increases in the point estimates. For example, the 95th percentile of expense moves from -117 b.p. (t-statistic of 5.50) using the lagged variable to -127 b.p. (t-statistic of -5.65) when using contemporaneous data. This underscores that the majority of the selling pressure represents a predictable component of certain firm's business strategy that lacks fundamental information, yet induces large price impact.

These results reinforce the flip side of the dividend results - predictable price pressure is not limited to purchases, but also holds for sales. While the sales we examine are announced before the returns we examine, the economic effects we find are large. This is consistent with markets generally not being fully liquid, and so predictable purchases and sales lead to predictable price changes.

H. Price Multipliers

The previous analysis of dividend payments and market returns considered economic magnitudes primarily in terms of the observed variation in the dependent variable - that is, how much larger are market returns with higher dividends. Another important metric of economic magnitude is the multiplier of market price with respect to a dollar of predictable investment. This is the inverse of the price elasticity (the change in quantity demanded as a function of price changes). In order to estimate such a figure, we need an estimate of what fraction of a dollar of dividends is typically reinvested into the market. That is to say, we need to approximate the number of dollars spent on purchases, not the number of dollars received. While data limits the ability to provide exact figures, this section provides a range of such estimates and what they imply for the market's demand elasticity.

One potential (though implausible) benchmark is to assume that all funds are reinvested on either the payment date or the following day. Under this assumption, the coefficient of 67 from Table I implies a multiplier of 0.67. This suggests that a $1 reinvestment leads to an increase in market value of $0.67. Table XIII lists the implied multiplier based on the reinvestment rate with the first row showing the estimate of 0.67 assuming that all investors invest all dividends.

The estimate of 0.67 is an implausible lower bound for the multiplier, but is already too high to be explained by any standard asset pricing model. 100% reinvestment is intuitively unlikely and inconsistent with the behavior of some investors. For example, retail investors consume dividends at a high rate (e.g., Baker, Nagel, and Wurgler 2006; Di Maggio, Kermani, and Majlesi 2020) meaning they cannot be reinvesting them. Further, these investors typically trade about once a year (e.g. Odean 1998) and are less likely to trade if they hold positions that pay dividends (Hartzmark and Solomon 2019). Both behaviors are associated with lower short-term reinvestment rates.

With that said, it could be there is a subset of retail investors that reinvest dividends. To examine whether this is likely, we look at the propensity of retail investors trading from their own accounts from 1991-1996 (as in Barber and Odean 2000) to buy a position (either open a new position or expand an old position) based on receiving a dividend payment in Table XII Panel A.[29] We regress a dummy variable equal to one if a position is purchased that day on a dummy variable for whether a position in the portfolio received a dividend payment on the current or prior day.

The results suggest dividends do not induce a meaningful increase in buying for individual traders. The constant in column 1 shows trading activity is low, as only 0.18% of days include a buy trade. The dividend payment dummy variable is significant, and of a similar magnitude to the constant, suggesting that on days when dividends are received the probability of buying doubles to roughly 0.36%. The magnitude of this coefficient suggests that dividend reinvestment is not common. Further, the odds of receiving a dividend vary with portfolio characteristics and trading behavior. To partially control for this, column 2 adds an account fixed effect which decreases the influence of receiving a dividend payment to 0.03%. While there may be some increase in buying when a dividend is received, the small economic magnitude of the effects, and the prior literature, suggest that most payments are not quickly reinvested. Thus we think a reasonable approximation of the reinvestment rate by retail investors of single name stock holdings is zero.[30]

The assumption that retail investors do not reinvest dividends affects the estimated multiplier. According to the Financial Accounts of the United States release Z.1 from March 2020, retail investors hold 34% of assets. Taking this estimate and assuming that retail investors never reinvest, and all other investors reinvest 100% of dividends, suggests that 66% of dividends are reinvested

[29] All days that an investor could have traded are included and the analysis is conducted at the portfolio level. See Hartzmark 2015 for a description of restrictions on this dataset.

[30] We note that this direct evidence is for retail investors in the 1990s, and it is unclear to what extent it holds in other datasets or more recent time periods where internet-based trading and lower trading costs may increase such portfolio reinvestment. We have no direct evidence on these issues, so we leave it to future research.

which yields an estimate of the multiplier of 1.02 in the second row of Table XIII.

The remaining 66% of non-retail investors generally represent more professional investors who likely reinvest dividends at higher rates. With that said, it is implausible that reinvestment is 100%. For example, if an asset manager needs to increase his cash buffer or send cash to his clients, he likely would not always reinvest a dividend and may instead use it for these needs. Ultimately we think it is an empirical question as to what this rate is.

While the data is fairly limited to calculate such a number, we use data on mutual funds to estimate their reinvestment rate. Mutual funds represent a large class of professional investors (22% of the US equity market in March 2020) that publicly disclose enough information to estimate a reinvestment rate. While the ideal data set would involve information on daily holdings, dividends received and trades, we do not have access to such a dataset.

Instead, the analysis focuses on monthly holdings data from CRSP which reports the positions a mutual fund holds at the end of the month. To calculate the dividends received we take the positions at the end of a given month and use them to measure dividends received the next month. To calculate the fraction of TNA of a fund accounted for by its listed holdings, we sum a variable that lists the proportion of TNA of each position.[31]

When a dividend payment is received, the value of fund TNA increases by the amount of the dividend payment. If that money is used to buy equity securities that get listed in the monthly reports, then the dividend payment will have no impact on the measure of the fund's TNA explained by the holdings data. If the payment is not reinvested, for example if it remains in cash, then the fund's TNA will increase but the value of equity holdings will not. Thus, absent reinvestment, the fund's fraction of TNA measured in listed holdings will decrease by the dividend payment amount.

An additional complexity arises if the ex-dividend date occurs within the same calendar month as the payment. When a stock goes ex-dividend, the price will drop by the dividend amount, absent taxes and frictions (Miller and Modigliani 1961). In this scenario, the TNA in equity holdings drops by the dividend amount on the ex-date, but if the fund engages in full reinvestment, TNA is

[31]To ensure that these positions represent actual holdings we limit the sample to holdings with non-missing CUSIP variables. We focus on this variable as it seems the least prone to erroneous entries and data errors.

replenished by the same amount on the payment date, leading to no change in monthly measures if both occur in the same month. Similarly, zero reinvestment would be associated with a decrease by the dividend amount. In practice, prices on the ex-day drop by less than the amount of the dividend (Elton and Gruber 1970, Hartzmark and Solomon 2013). Full reinvestment with such positive ex-day returns will lead to position values greater than the original holding. Predictions are much less clear though, because the amount of the price drop (and thus the predicted change in TNA under full reinvestment) varies with a number of factors.[32] This means that our predictions are cleanest for payments where the ex-date occurs during an earlier calendar month. Since these observations also represent about two thirds of such payments, we focus on them for our estimates.[33]

To explore reinvestment behavior, we first graph the change in TNA from the prior month relative to the dividend payments received that month in Figure 7. The figures are a bin-scatter plot splitting the data into 20 bins based on the magnitude of the dividend yield (represented by its location on the x-axis). The y-axis shows the average change in percentage holdings in that month for funds in that bin. The red line is the outcome of a regression of the change in TNA on the dividend payment. The figures are value-weighted by the fund's prior month TNA.

Figure 7 Panel A shows the relation with no controls. The first thing to notice is a strong, positive, roughly linear relation. This is consistent with a general reinvestment of dividend payments that are received. Given that the x-axis and y-axis are on the same scale, perfect reinvestment would be represented by a 45 degree line. The line is substantially flatter than this. Thus the figure suggests that mutual funds reinvest dividends, but do so at a rate significantly below 100%. Such results could be explained by differences over certain time periods, or by heterogeneity across mutual funds. Panel B repeats the analysis but plots the data as the residuals after controlling for fixed effects for each fund and time period. The pattern is similar suggesting some reinvestment, but not 100%.

For more precise estimates of reinvestment, we regress the change in a fund's fraction of holdings

[32]Differences in ex-day returns are affected by dividend yield, liquidity, recessions, VIX, etc (Hartzmark and Solomon (2013))

[33]We add a number of additional filters to the mutual fund database to attempt to avoid the noise inherent in this data. We focus on funds with a TNA above $10 million in the prior month that report holding at least 10 positions. We focus on funds where the estimate of the fraction of TNA accounted for by holdings is between 50% and 105% in the current and prior month.

on the dividend payment received that month. We value-weight by the prior month TNA to get a better estimate of the dollar-weighted reinvestment amount. With full reinvestment, this regression should yield a coefficient of one. With zero reinvestment this regression should yield a coefficient of zero. With partial reinvestment, that amount should be captured by the regression coefficient.

Table XII Panel B reports the results of these regressions. The first column does not include controls and shows a coefficient of 0.53 with a t-statistic of 4.39, suggesting 53% of dividends received in a given month are reinvested that month. The next 3 columns include year-by-month fixed effects, fund fixed effects and both together. This yields estimates of dividend reinvestment from 45% to 69%.

We take the high (69%) and low (45%) of these estimates for our our back-of-the-envelope calculations of multipliers.[34] In March 2020, mutual funds held 22% of the US equity market and ETFs held 6%. We think a reasonable assumption is that ETFs behave similarly to mutual funds with respect to reinvestment, so we apply our estimates to 28% of the market.[35] Taking the high estimate of 69% reinvestment suggests an aggregate reinvestment rate of 57% (retaining the assumption that retail investors do not reinvest and the rest of the market reinvests at a rate of 100%). This implies a multiplier of 1.17. Taking the low estimate of 45% reinvestment rate from the value-weighted estimate, we find a multiplier estimate of 1.33.

These elasticities are based on the assumption that the remaining investors who are not retail, mutual funds or ETFs reinvest at 100%, which again seems implausibly high. This group includes foreign investors (16%), pensions (11%), business holdings (4%), hedge funds (3%) and other investors (3%). There is clearly variation across these groups, but they largely represent professional investors who also sometimes pay out money to constituents, similar to a mutual fund. Given that we lack the data to estimate the actual reinvestment rate, we think a reasonable approximate as-

[34]Taking this monthly estimate and applying it to our baseline regression coefficient implies that when a dividend is received it is reinvested within two days. While we think it is plausible that mutual funds will reinvest quickly if they are planning to reinvest at all, this is a conjecture the data does not allow an explicit test of. We note that fund managers often are worried underperforming benchmarks (even going as far as to mis-specify benchmarks in order to beat them - see Sensoy (2009)), and thus have an incentive to quickly redeploy capital rather than having an additional performance drag from holding excess cash.

[35]We lack ETF holdings data to directly perform the same calculation for them.

sumption of their reinvestment rate is to take the mutual fund rates and apply it to these investors. Taking the estimate of 69% reinvestment for these other investors yields an aggregate reinvestment rate of 46% and a multiplier of 1.47, while 45% reinvestment yields a multiplier estimate of 2.26. Thus we think a reasonable back-of-the-envelope multiplier estimate is likely in the 1.5 to 2.3 range.

IV. Discussion

The existence of price pressure from predictable flows means there is a credible basis to believe that price pressure is a plausible null hypothesis *in general*. Any finite limit order book will have price pressure at some point. The extent of such pressure depends on the ability and willingness of liquidity-providing traders to engage in offsetting trades, which should be easier for traders to do with more advance notice of the timing. In this sense, finding price pressure for the most liquid assets when flows are predictable makes it more likely that price pressure exists when flows are not predictable and assets are illiquid.

The implications of such a null hypothesis are wide-ranging. Most asset pricing models start with fundamental value from risk-adjusted discounted future cash flows. While they sometimes add frictions or behavioral factors to explain deviations from the rational benchmark, the actual buying and selling is usually relegated to the background. However, beginning with price pressure, and taking fundamental value or psychology as inputs into this trading, can help explain pricing anomalies especially in situations where fundamental value is difficult to model or entirely absent.

Perhaps the strongest example of this phenomenon is cryptocurrencies, a setting in which standard finance models struggle to make any meaningful predictions. Economists have long puzzled over why assets like Bitcoin should have non-zero prices, given they have no underlying cash-flows or clear value proposition other than the ability to sell them to other investors.[36] Informal justifi-

[36] Which, in equilibrium, becomes entirely circular, and potentially applies to any asset whatsoever. Indeed, finance has had a lot more success in modeling the usefulness of tokens, where the digital asset is an input into a real-world project, and the sale is a means of raising funds for such a project. See, for instance Li and Mann (2018), Cong, Li, and Wang (2021), Catalini and Gans (2018), and others. There are relatively fewer models of coins that exist purely as objects of exchange, even though these assets (especially Bitcoin) are not only the original cryptocurrencies, but still the largest by market capitalization.

cations often appeal to the idea of it being a "hedge asset," conditional on it already being widely traded (e.g. Dyhrberg 2016). But this sidesteps the question of why cryptocurrencies have a non-zero price in the first place. In a world of price pressure, the somewhat tautological answer is that these assets have a non-zero price because people keep wanting to buy them. Because Bitcoin is in a fixed total supply, with a small and decreasing amount released through mining, if enough people continue to buy Bitcoin, this is predicted to keep raising prices.

This characterization highlights both the strengths and weaknesses of price pressure as an asset pricing framework. It is only a partial theory, because it does not take a stand on what is causing the underlying demand.[37] There is also a sense in which price pressure seems almost tautological, in that enough people buying an asset will push up its price. The rejoinder, however, is that starting with price pressure makes it intuitive to take the existing demand for Bitcoin as a basic fact, and something over which it is acceptable to be agnostic as to the cause or likely duration. While it is unable to identify the source of demand, it still makes a simple and powerful prediction. If demand has been broadly increasing for 12 years, for whatever reason, then prices will rise even for an asset that, by design, produces no cash flows or inherent utility. By contrast, worldviews that start with fundamental value tend to predict that not only should the price of Bitcoin be zero, but that it should *always* have been zero, which suggests that going to zero in the future is somewhere between "highly likely" and "inevitable". This makes it difficult to know how to update expectations for each successive year when prices are not zero, and not obviously heading towards zero.

While cryptocurrencies provide the starkest example where predictions based on fundamental value are insufficient, price pressure can also provide a relatively parsimonious explanation for a number of other results where modeling deviations from fundamental value is either complicated, or incomplete. Price pressure readily accounts for market episodes like the "meme stocks" of Gamestop, Nokia and others during early 2021, that experienced large and surprisingly long-lived price increases

[37]Credible contributors seem likely to be some combination of true believers in the project, people desiring to transact outside the formal banking system, and extrapolative traders. While not directly about cryptocurrencies, Van Wesep and Waters (2021) present a model of unstable valuations due to leveraged optimists, with a similar spirit to some of this discussion. But these components of optimism and leverage are largely conceptually orthogonal to the price pressure component.

as a result of attention on internet forums such as Reddit's r/WallStreetBets. Coordinating lots of people to buy in this manner is predicted to increase prices (even if fundamentals are unchanged), and the findings from our dividend reinvestment results suggest that this need not quickly revert.

Price pressure also provides a way of simplifying other results in the literature where standard asset pricing explanations are plausible but somewhat convoluted. In international finance, it has long been known that macroeconomic models have difficulty explaining exchange rate movements at horizons less than a year (Meese and Rogoff 1983, Meese 1990). Evans and Lyons (2002) shows empirically that order flow matters a great deal for exchange rate movements, and explains this using a model where order flow reveals fundamental decentralized economic information such as future interest rate differentials. Relatedly, Jiang, Krishnamurthy, and Lustig (2021) explain the high valuation of the US dollar, and the importance of shifting supply and demand of dollars for exchange rates, by positing a time-varying convenience yield to holding US dollar assets. In both cases, price pressure provides the very simple prediction that order flow inherently matters for prices, regardless of its source, without the need for additional assumptions as to what such price pressure represents. In other words, it is *possible* that such flows derive from fundamental information about the macroeconomy or expressions of convenience yields, but this need not be the case.

A price pressure framework provides a simple explanation of how the Federal Reserve is able to influence prices in low interest rate environments. Namely, quantitative easing matters, because the almost unlimited budget constraint of the Fed means its purchases directly affect prices (and thus yields) in the assets it buys. While there is considerable evidence that quantitative easing has important effects on markets (Hamilton and Wu 2012, Wu and Xia 2016, Gilchrist, López-Salido, and Zakrajšek 2015, Rebucci, Hartley, and Jiménez 2020) the mechanism driving this is less clear. As Bernanke and Reinhart (2004) describe, close to the zero lower bound the Fed can influence prices by either altering investor expectations, by changing the composition of the Fed's balance sheet, or by changing the quantity of the balance sheet. These latter two options function because of price pressure, though they are often described using different terms.[38] While some of

[38]For example, Bernanke and Reinhart (2004) describe the influence by stating "changes in relative demands by a large purchaser have the potential to alter relative security prices" if "investors do not treat all securities as perfect

the evidence is consistent with either an expectations or a price pressure channel,[39] some episodes illustrate the importance of price pressure. For example, Vissing-Jørgensen (2021) documents that, during Covid-19, rates spiked due to selling pressure from March 9th-18th and then reversed due to buying pressure by the Fed on purchase days (not announcement days). Price pressure represents a parsimonious explanation for the Fed's influence on prices, requiring only that the Fed's budget constraint is huge, and a common understanding that the Fed is willing to use it.[40]

Price pressure also makes predictions for where other asset pricing puzzles may be found. To take one example, in a world of price pressure, demographics matter. More baby boomers selling equity for retirement is predicted to lower prices (holding fundamental news constant). This and other related outcomes are intuitive places to look when the underlying question is "what phenomena are likely to drive significant trade in the same direction?" and may not be as obvious if the initial inquiry is instead "what phenomena will influence the macroeconomy or psychology of investors?".

The idea of price pressure does not imply that fundamental value is unimportant for asset pricing, but the explanation for *why* it matters is different from standard theories. Under a price pressure explanation, fundamental value matters when it affects the prices at which traders are willing to place orders. This means that its impact likely varies according to the context. For example, if investors generally view prices to be anchored to fundamental value and think other investors share this view (as in the beauty contest in Keynes 1936), then willingness to trade, and thereby the price, will reflect beliefs about fundamental value. On the other hand, if investors believe prices are decoupled from fundamentals (such as in the meme stock example above) or that convergence to

substitutes". The disconnect between a macroeconomics literature that assumes Treasury securities with different duration are *not* perfect substitutes, and a finance literature that assumes entirely different companies *are* perfect substitutes, is curious, to say the least. Price pressure also explains why the evidence for purely compositional changes affecting prices may be "contentious" (as Bernanke and Reinhart (2004) characterize it) - without an overall balance sheet expansion, at least one asset has to be sold, which pushes down prices for the asset the Fed previously held.

[39]Empirically, prices of assets that the Fed announces it will buy tend to increase on announcement days (e.g. Rebucci, Hartley, and Jiménez 2020). This is consistent with both expectation based and price pressure based explanations, as announcing future Fed buying may induce other traders to front run those trades.

[40]Price pressure offers an indirect explanation for the Fed's role in asset classes that it does not directly purchase. For example, (Cieslak and Vissing-Jorgensen, 2020) describe the "Fed Put", that monetary interventions will be used to prevent large stock market declines, satirized in popular discourse as the monetary policy of "big line go up." If Fed actions in one market, such as using buying pressure to lower the interest rate, makes equities appear more attractive, perhaps due to reaching for yield (Lian, Ma, and Wang, 2019) or demand for dividends (Hartzmark and Solomon, 2019), this could induce equity buying pressure and increase the price of the market.

fundamentals is a long-term phenomenon (e.g. Royal Dutch / Shell (Rosenthal and Young 1990)), than fundamental value is unlikely to drive trading decisions and will be less relevant for prices.[41]

The results in this paper suggest that an important course of future study is to pair the base hypothesis of price pressure together with an understanding of the context-specific shifts in buying or selling pressure in order to better understand asset price movements.

V. Conclusion

In this paper, we show that predictable increases in the number of people buying or selling shares leads to predictable changes in price for those shares. Put in such stark terms, such a finding may seem obvious. And yet many of the standard models in finance assume this result does not hold - that investors are simply exchanging twenty dollar bills for hundred dollar bills, and the price for such an exchange is constant regardless of volume. Even when trades are considered to move prices, the general assumption is that the main force is information about the trade, not the trade itself. We consider price changes in the largest and most liquid assets (the value-weighted market portfolio), where the flows in question ought to be known to not result from private information, and where the timing and amount of such flows is known in advance. Even with these aspects stacking the deck against the possibility of price pressure, price pressure is evident.

One interesting question is whether one would expect price pressure to be more pronounced at the individual stock level or at the market level. Academic finance generally predicts that the larger and more liquid the asset, the more likely that the asset behaves in accordance with the efficient markets hypothesis. It is unclear if this is the appropriate null hypothesis though, since while individual stocks may be close substitutes for each other, there are relatively few good substitutes for the market as a whole, an argument made by Samuelson (as reported in Jung and Shiller 2005).

[41] Deviations from fundamental value also gives rise to longer-term trading strategies based on economic fundamentals. If a stock's price is low relative to its discounted future cash flows, this represents an opportunity for a long-term investor to buy the stock and receive its cash flows. Conversely, if an asset is overpriced, it may be profitable to issue more equity, or to found a new company and sell equity at the overvalued price. These strategies likely provide weaker constraints on prices, because they take considerable time to implement, and have different risks than betting on mispricing correcting itself.

Another interesting question is whether potential arbitrageurs are aware of the pattern that we document. Given that dividend payments are obvious and recurring, the standard prediction is that it is likely that the pattern is understood by informed investors. However, our results suggest that if arbitrageurs know about the pattern, they are unwilling or unable to correct it. This would suggest that a key component of arbitrageur behavior is absent from standard models.

While arbitrageurs perhaps *ought* to know about the findings in this paper, it is not clear that they actually *do*. The implicit assumption in the idea that informed traders know about a newly documented mispricing pattern is that academics are the last to know about it. This assumption is far from innocuous and appears at odds with the empirical evidence (e.g., McLean and Pontiff 2016). Mundane daily predictable basis point fluctuations may be the type of setting that people focus less attention on, making it less likely that arbitrageurs engage in offsetting trades, leading to price changes unrelated to information flows (similar to Da, Gurun, and Warachka 2014; Chang et al. 2016). With that said, the space of possible flows with price impact is large, and the setting we study seems relatively easy to understand. Thus, if investors are unaware of this pattern, this suggests liquidity provision by arbitrageurs is unlikely to occur in the presence of less obvious flows. This would imply such liquidity provision is present only in certain obvious situations and is not the ubiquitous force typically assumed. Either of these mechanisms are consistent with the importance of price pressure in a way that is incompatible with standard models, so directly testing these channels is an important avenue of future research.

One way of interpreting our results is that, for whatever reason, investors have strong views as to which assets they want to hold, and changes in the price of such stocks do not greatly change their willingness to purchase them. The question of *why* this is the case is one of great interest.

REFERENCES

Adrian, Tobias and Hyun Song Shin, 2010, Liquidity and leverage, *Journal of financial intermediation* 19, 418–437.

Amihud, Yakov, 2002, Illiquidity and stock returns: cross-section and time-series effects, *Journal of financial markets* 5, 31–56.

Amihud, Yakov and Haim Mendelson, 2015, The pricing of illiquidity as a characteristic and as risk, *Multinational Finance Journal* 19, 149–168.

Ang, Andrew, Sergiy Gorovyy, and Gregory B Van Inwegen, 2011, Hedge fund leverage, *Journal of Financial Economics* 102, 102–126.

Baker, Malcolm, Stefan Nagel, and Jeffrey Wurgler, 2006, The effect of dividends on consumption, Tech. rep., National Bureau of Economic Research.

Barber, Brad M and Terrance Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *The journal of Finance* 55, 773–806.

Ben-David, Itzhak, Francesco Franzoni, Rabih Moussawi, et al., 2011, *ETFs, arbitrage, and contagion* (University of Geneva Geneva).

Ben-David, Itzhak, Jiacui Li, Andrea Rossi, and Yang Song, 2020, Advice-driven demand and systematic price fluctuations, *NBER working paper* .

Berkman, Henk and Paul D Koch, 2017, Drips and the dividend pay date effect, *Journal of Financial and Quantitative Analysis* 52, 1765–1795.

Bernanke, Ben S and Vincent R Reinhart, 2004, Conducting monetary policy at very low short-term interest rates, *American Economic Review* 94, 85–90.

Bernard, Victor L. and Jacob K. Thomas, 1989, Post-earnings-announcement drift: Delayed price response or risk premium, *Journal of Accounting Research* 1–36.

Bettis, J Carr, Jeffrey L Coles, and Michael L Lemmon, 2000, Corporate policies restricting trading by insiders, *Journal of financial economics* 57, 191–220.

Brunnermeier, Markus K. and Stefan Nagel, 2004, Hedge funds and the technology bubble, *Journal of Finance* 59, 2013–2040.

Brunnermeier, Markus K and Lasse Heje Pedersen, 2009, Market liquidity and funding liquidity, *The review of financial studies* 22, 2201–2238.

Catalini, Christian and Joshua S Gans, 2018, Initial coin offerings and the value of crypto tokens, Tech. rep., National Bureau of Economic Research.

Chang, Tom Y., Samuel M. Hartzmark, David H. Solomon, and Eugene F. Soltes, 2016, Being surprised by the unsurprising: Earnings seasonality and stock returns, *The Review of Financial Studies* 30, 281–323.

Chang, Yen-Cheng, Harrison Hong, and Inessa Liskovich, 2015, Regression discontinuity and the price effects of stock market indexing, *The Review of Financial Studies* 28, 212–246.

Chen, Huaizhi, 2022, Cash induced demand, *Journal of Financial and Quantitative Analysis* .

Cieslak, Anna and Annette Vissing-Jorgensen, 2020, The economics of the fed put, Tech. rep., National Bureau of Economic Research.

Cohen, Lauren, 2009, Loyalty-based portfolio choice, *The Review of Financial Studies* 22, 1213–1245.

Cong, Lin William, Ye Li, and Neng Wang, 2021, Tokenomics: Dynamic adoption and valuation, *The Review of Financial Studies* 34, 1105–1155.

Coval, Joshua and Erik Stafford, 2007, Asset fire sales (and purchases) in equity markets, *Journal of Financial Economics* 86, 479–512.

Da, Zhi, Umit G Gurun, and Mitch Warachka, 2014, Frog in the pan: Continuous information and momentum, *The review of financial studies* 27, 2171–2218.

Da, Zhi, Borja Larrain, Clemens Sialm, and José Tessada, 2018, Destabilizing financial advice: Evidence from pension fund reallocations, *The Review of Financial Studies* 31, 3720–3755.

DAmico, Stefania and Thomas B King, 2013, Flow and stock effects of large-scale treasury purchases: Evidence on the importance of local supply, *Journal of Financial Economics* 108, 425–448.

Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers, 1997, Measuring mutual fund performance with characteristic-based benchmarks, *Journal of Finance* 52, 1035–1058.

DellaVigna, Stefano and Joshua M. Pollet, 2009, Investor inattention and friday earnings announcements, *Journal of Finance* 64, 709–749.

Di Maggio, Marco, Amir Kermani, and Kaveh Majlesi, 2020, Stock market returns and consumption, *The Journal of Finance* 75, 3175–3219.

Dyhrberg, Anne Haubo, 2016, Hedging capabilities of bitcoin. is it the virtual gold?, *Finance Research Letters* 16, 139–144.

Elton, Edwin J. and Martin J. Gruber, 1970, Marginal stockholder tax rates and the clientele effect, *Review of Economics and Statistics* 52, 68–74.

Engelberg, Joseph, R David McLean, Jeffrey Pontiff, and Matthew C Ringgenberg, 2021, Do cross-sectional predictors contain systematic information?, *Journal of Financial and Quantitative Analysis* 1–45.

Evans, Martin DD and Richard K Lyons, 2002, Order flow and exchange rate dynamics, *Journal of political economy* 110, 170–180.

Fama, Eugene F and Kenneth R French, 2001, Disappearing dividends: changing firm characteristics or lower propensity to pay?, *Journal of Financial Economics* 60, 43.

Frazzini, Andrea and Owen A. Lamont, 2007, The earnings announcement premium and trading volume, Working paper, Harvard University.

Frazzini, Andrea and Owen A Lamont, 2008, Dumb money: Mutual fund flows and the cross-section of stock returns, *Journal of financial economics* 88, 299–322.

Gabaix, Xavier and Ralph SJ Koijen, 2021, In search of the origins of financial fluctuations: The inelastic markets hypothesis, *Available at SSRN 3686935* .

Gabaix, Xavier, Arvind Krishnamurthy, and Olivier Vigneron, 2007, Limits of arbitrage: Theory and evidence from the mortgage-backed securities market, *The Journal of Finance* 62, 557–595.

Garleanu, Nicolae, Lasse Heje Pedersen, and Allen M Poteshman, 2008, Demand-based option pricing, *The Review of Financial Studies* 22, 4259–4299.

Gilchrist, Simon, David López-Salido, and Egon Zakrajšek, 2015, Monetary policy and real borrowing costs at the zero lower bound, *American Economic Journal: Macroeconomics* 7, 77–109.

Greenwood, Robin and Samuel G Hanson, 2013, Issuer quality and corporate bond returns, *The Review of Financial Studies* 26, 1483–1525.

Greenwood, Robin and Dimitri Vayanos, 2014, Bond supply and excess bond returns, *The Review of Financial Studies* 27, 663–713.

Gromb, Denis and Dimitri Vayanos, 2002, Equilibrium and welfare in markets with financially constrained arbitrageurs, *Journal of financial Economics* 66, 361–407.

Hamilton, James D and Jing Cynthia Wu, 2012, The effectiveness of alternative monetary policy tools in a zero lower bound environment, *Journal of Money, Credit and Banking* 44, 3–46.

Harris, Lawrence and Eitan Gurel, 1986, Price and volume effects associated with changes in the s&p 500 list: New evidence for the existence of price pressures, *the Journal of Finance* 41, 815–829.

Hartzmark, Samuel M., 2015, The worst, the best, ignoring all the rest: The rank effect and trading behavior, *Review of Financial Studies* 28, 1024–1059.

Hartzmark, Samuel M and David H Solomon, 2013, The dividend month premium, *Journal of Financial Economics* 109, 640–660.

Hartzmark, Samuel M and David H Solomon, 2019, The dividend disconnect, *The Journal of Finance* 74, 2153–2199.

Hartzmark, Samuel M and David H Solomon, 2020, Reconsidering returns, *Review of Financial Studies* Forthcoming.

Jagolinzer, Alan D, David F Larcker, and Daniel J Taylor, 2011, Corporate governance and the information content of insider trades, *Journal of Accounting Research* 49, 1249–1274.

Jiang, Zhengyang, Arvind Krishnamurthy, and Hanno Lustig, 2021, Foreign safe asset demand and the dollar exchange rate, *The Journal of Finance* 76, 1049–1089.

Jung, Jeeman and Robert J Shiller, 2005, Samuelson's dictum and the stock market, *Economic Inquiry* 43, 221–228.

Keynes, John Maynard, 1936, *The general theory of employment, interest, and money.*

Kvamvold, Joakim and Snorre Lindset, 2018, Do dividend flows affect stock returns?, *Journal of Financial Research* 41, 149–174.

Lakonishok, Josef and Seymour Smidt, 1988, Are seasonal anomalies real? a ninety-year perspective, *The review of financial studies* 1, 403–425.

Li, Jennifer Jie, Neil D Pearson, and Qi Zhang, 2020, Impact of demand shocks on the stock market: Evidence from chinese ipos, *Neil D. and Zhang, Qi, Impact of Demand Shocks on the Stock Market: Evidence from Chinese IPOs (November 15, 2020)* .

Li, Jiasun and William Mann, 2018, Initial coin offerings and platform building .

Lian, Chen, Yueran Ma, and Carmen Wang, 2019, Low interest rates and risk-taking: Evidence from individual investment decisions, *The Review of Financial Studies* 32, 2107–2148.

Lou, Dong, 2012, A flow-based explanation for return predictability, *The Review of Financial Studies* 25, 3457–3489.

Lou, Dong, Hongjun Yan, and Jinfan Zhang, 2013, Anticipated and repeated shocks in liquid markets, *The Review of Financial Studies* 26, 1891–1912.

Lucca, David O and Emanuel Moench, 2015, The pre-fomc announcement drift, *The Journal of finance* 70, 329–371.

Malmendier, Ulrike and Geoffrey Tate, 2005, Ceo overconfidence and corporate investment, *The journal of finance* 60, 2661–2700.

Manela, Asaf and Alan Moreira, 2017, News implied volatility and disaster concerns, *Journal of Financial Economics* 123, 137–162.

McLean, R David and Jeffrey Pontiff, 2016, Does academic research destroy stock return predictability?, *The Journal of Finance* 71, 5–32.

Meese, Richard, 1990, Currency fluctuations in the post-bretton woods era, *Journal of Economic Perspectives* 4, 117–134.

Meese, Richard A and Kenneth Rogoff, 1983, Empirical exchange rate models of the seventies: Do they fit out of sample?, *Journal of international economics* 14, 3–24.

Miller, Merton H and Franco Modigliani, 1961, Dividend policy, growth, and the valuation of shares, *the Journal of Business* 34, 411–433.

Nagel, Stefan, 2012, Evaporating liquidity, *Review of Financial Studies* 25, 2005–2039.

Neuhierl, Andreas and Michael Weber, 2019, Monetary policy communication, policy slope, and the stock market, *Journal of Monetary Economics* 108, 140–155.

Odean, Terrance, 1998, Are investors reluctant to realize their losses, *Journal of Finance* 53, 1775–1798.

Parker, Jonathan A, Antoinette Schoar, and Yang Sun, 2020, Retail financial innovation and stock market dynamics: The case of target date funds, Tech. rep., National Bureau of Economic Research.

Patel, Nimesh and Ivo Welch, 2017, Extended stock returns in response to s&p 500 index changes, *The Review of Asset Pricing Studies* 7, 172–208.

Rebucci, Alessandro, Jonathan S Hartley, and Daniel Jiménez, 2020, An event study of covid-19 central bank quantitative easing in advanced and emerging economies, Tech. rep., National Bureau of Economic Research.

Rosenthal, Leonard and Colin Young, 1990, The seemingly anomalous price behavior of royal dutch/shell and unilever nv/plc, *Journal of Financial Economics* 26, 123–141.

Ross, Stephen A., 1976, The arbitrage theory of capital asset pricing, *The journal of economic theory* 13, 341–360.

Rozeff, Michael S and William R Kinney Jr, 1976, Capital market seasonality: The case of stock returns, *Journal of financial economics* 3, 379–402.

Savor, Pavel and Mungo Wilson, 2013, How much do investors care about macroeconomic risk? evidence from scheduled economic announcements, *Journal of Financial and Quantitative Analysis* 343–375.

Schmickler, Simon, 2021, Demand pressure spillover effects of corporate payouts .

Sensoy, Berk A, 2009, Performance evaluation and self-designated benchmark indexes in the mutual fund industry, *Journal of Financial Economics* 92, 25–39.

Sharpe, William F, 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, *The journal of finance* 19, 425–442.

Shleifer, Andrei, 1986, Do demand curves for stocks slope down?, *The Journal of Finance* 41, 579–590.

Thaler, Richard H, 1987, Anomalies: the january effect, *Journal of economic perspectives* 1, 197–201.

Van Wesep, Edward Dickersin and Brian Waters, 2021, The sky's the limit: Asset prices can be indeterminate when margin traders are all in, *Available at SSRN* .

Vayanos, Dimitri and Jean-Luc Vila, 2021, A preferred-habitat model of the term structure of interest rates, *Econometrica* 89, 77–112.

Vissing-Jørgensen, Annette, 2021, The treasury market in spring 2020 and the response of the federal reserve .

Wu, Jing Cynthia and Fan Dora Xia, 2016, Measuring the macroeconomic impact of monetary policy at the zero lower bound, *Journal of Money, Credit and Banking* 48, 253–291.

Wurgler, Jeffrey and Ekaterina Zhuravskaya, 2002, Does arbitrage flatten demand curves for stocks?, *The Journal of Business* 75, 583–608.

Figure 1. Daily Market Returns Based on Dividend Payments

This graph shows the the daily value-weighted market return (in %) on day t based on the quintile of the abnormal dividend payment on day t and $t-1$. Red bars represent 95% confidence intervals.

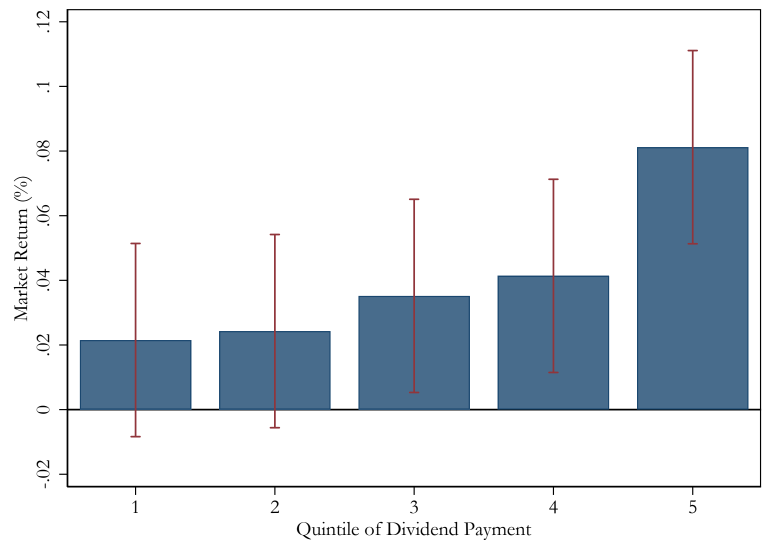

Figure 2. Cumulative Value-Weighted Returns from Dividend Payment Strategies

This graph shows the cumulative performance of a $1 investment in the market portfolio based on the dividend payout yield. The maroon line invests in the market if the dividend payment today or yesterday is above the median of the prior 252 days and earns zero returns otherwise. The navy line invests in the market when the dividend payment today or yesterday is less than or equal to the median of the prior 252 trading days and earns zero returns otherwise.

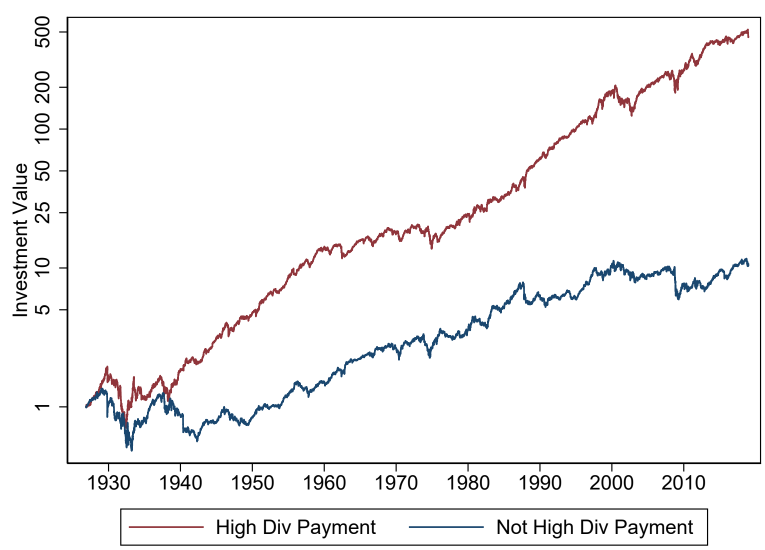

Figure 3. Mutual Fund Distributions

This graph shows the average fraction of the annual total distribution (including dividends and capital gains) by month and by quarter. In Panel A, the sum of all distributions in a given quarter is divided by the sum of all distributions in that calendar year. This value is averaged across years and graphed for each quarter. Thus each bar is the average fraction of distributions paid in a year that occur in a specific quarter. Panel B repeats the exercise using month instead of quarter.

Panel A: Average Distribution by Quarter

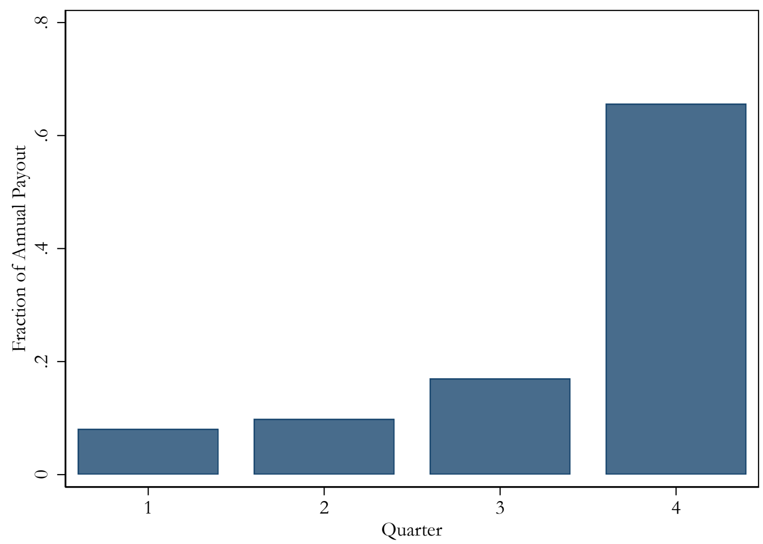

Panel B: Average Distribution by Month

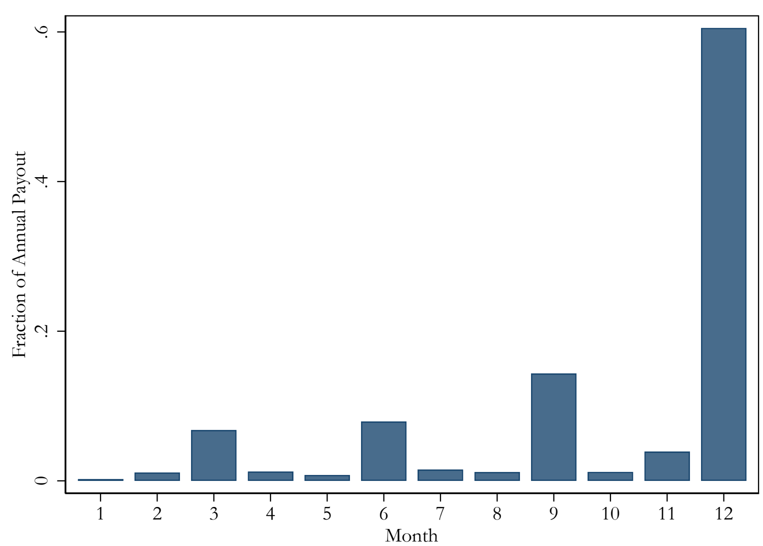

Figure 4. Return Sensitivity to Divided Payment by Decade

This graph shows the coefficient of value weighted market returns regressed on Market Div Pay[t-1,t] conducted separately for each decade. Each data point represents a decade, for example, the data point for 1930 is the regression coefficient when the regression is conducted on all data occurring from 1930 through 1939.

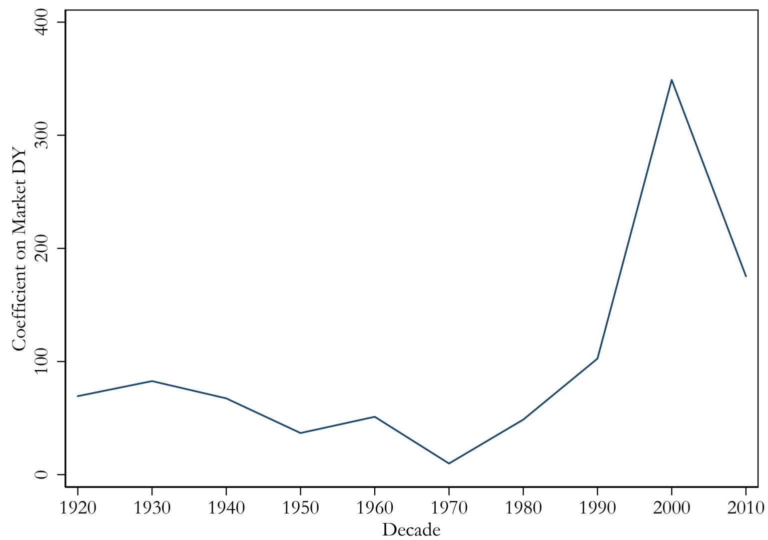

Figure 5. Longer Term Returns Based on Dividend Payment

These graphs shows the difference over time between returns on days with an above-median dividend payment and returns on days with a below-median dividend payment. Specifically, the cumulative return from t=0 through the number of days indicated on the x-axis is regressed on a dummy variable equal to one if the day is above the median dividend payment and a constant. This coefficient is graphed in blue and its 95% confidence interval is shown by the gray dashed line. The second graph includes controls for year fixed effects, as well as future dividend payments included up to the number of days on the x-axis.

Panel A: No Controls

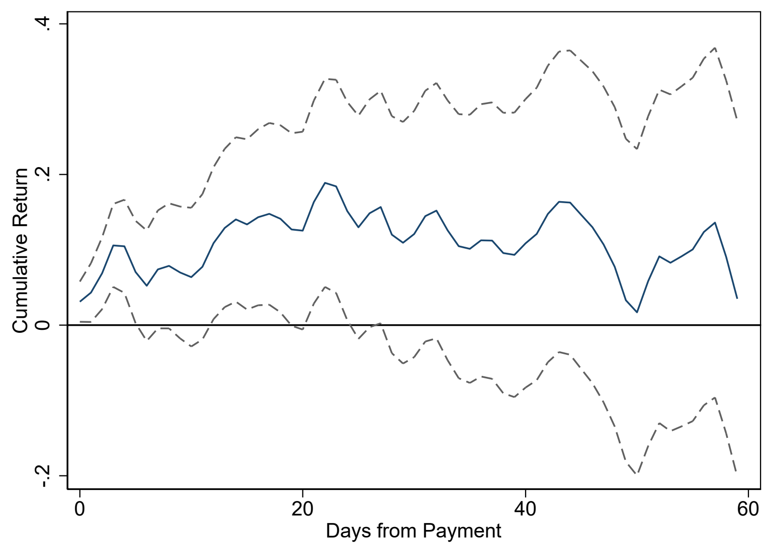

Panel B: Controlling for future dividends and year fixed effects

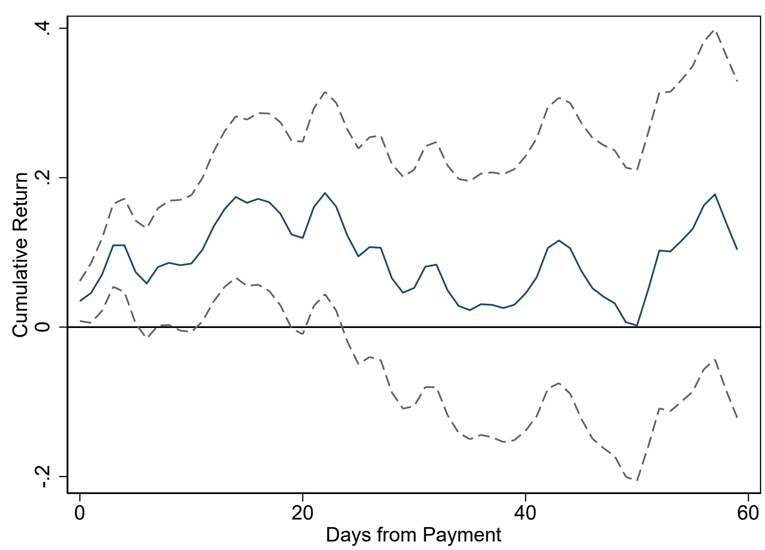

Figure 6. Stock Expense and Returns After Earnings Announcements

These figures show how returns after earnings announcements vary based on the level of stock expense. Panel A show a local linear plot of the cumulative returns from t+1 to t+3 based on the level of stock expense (winsorized at the 99th percentile). The gray dotted lines show the indicated percentiles of stock expense, among firms with non-zero stock expense. The gray shaded area represents the 95% confidence interval. Panel B shows the cumulative characteristic-adjusted returns from the announcement date (t=0) until 20 days after the announcement. Returns are shown separately based on the percentile of stock expense.

Panel A: Returns by Stock Expense

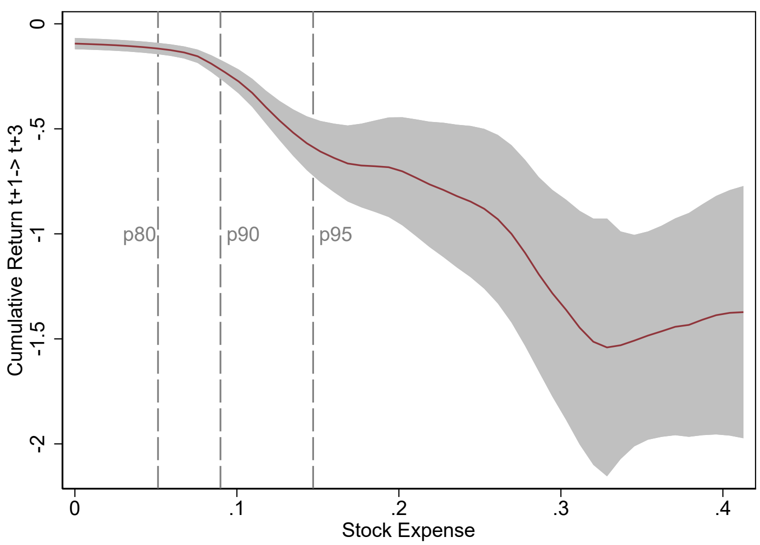

Panel B: Returns after Announcement

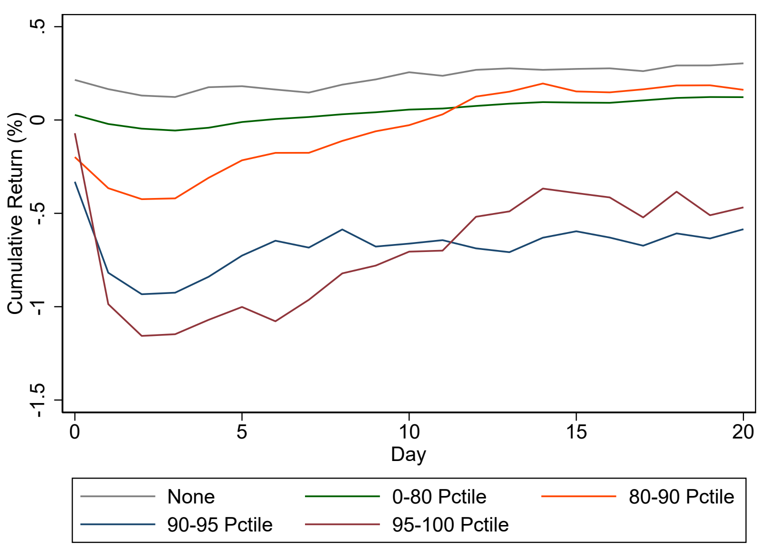

Figure 7. Mutual Fund Reinvestment

This graph shows a bin scatter plot of the change in fraction of TNA covered by holdings (y-axis) relative to the level of dividends received from those holdings (x-axis). Figures are value-weighted by prior month TNA. Panel A does not include controls. Panel B removes a fixed effect for fund and year-by-month. The analysis includes observations with a prior month TNA above $10 million, with at least 10 holdings, and where the fraction of TNA covered by holdings is between 50% and 105% in the current and prior month.

Panel A: No Controls

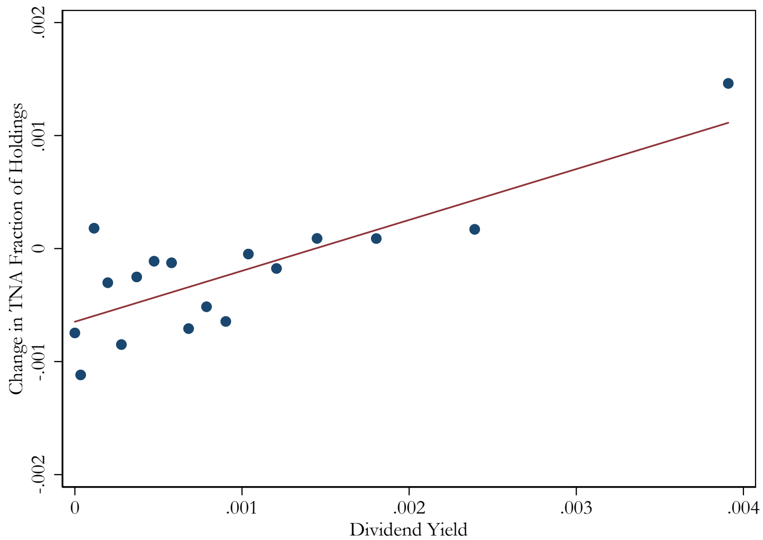

Panel B: Fund and Year-by-Month Fixed Effects

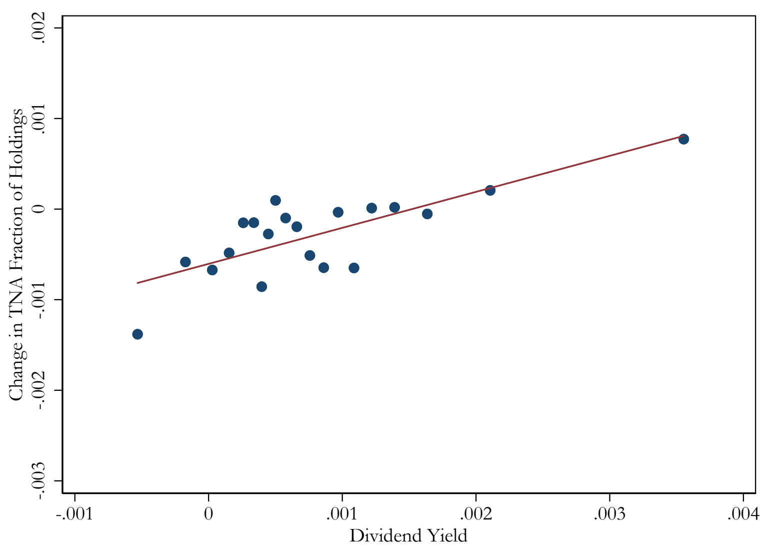

Table I

Market Returns Based on Dividend Payment

This table shows how the market return varies with the dividend payment yield. In Panel A, market return is measured in percent as either the CRSP value-weighted market index or the CRSP equal-weighted index. Market dividend payment yield is measured as the cumulative payment yield on day t and t-1 (Market Div Pay[t-1,t]), or the payment yield on the indicated day scaled by total market cap from the prior day. Panel B repeats the regressions, but changes the date of the market cap used for scaling. The first two columns use market cap from a month (t-20) prior, the next two columns use market cap from a quarter (t-60) prior, while the final two columns use market cap from a year (t-252) prior. All regressions are based on CRSP value weighted market returns. Even-numbered columns contain year-by-month fixed effects. t-statistics based on heteroskedasticity-robust standard errors are in parentheses.

Panel A: Baseline Results

	Value Weighted						Equal Weighted	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mkt Div Pay[t-1,t]			59.50***	67.07***			69.18***	72.35***
			(3.32)	(3.47)			(3.48)	(3.50)
Mkt Div Pay[t]	55.76*	74.85**						
	(1.74)	(2.26)						
Mkt Div Pay[t-1]	60.04***	71.98***						
	(2.73)	(3.10)						
Mkt Div Pay[t-2]	23.59	35.98						
	(0.98)	(1.43)						
Mkt Div Pay[t-3]	14.56	25.73						
	(0.50)	(0.85)						
Mkt Div Pay[t-4]	53.96*	66.66**						
	(1.89)	(2.32)						
Mkt Div Pay[t+1]					14.57	26.48		
					(0.59)	(0.99)		
Mkt Div Pay[t+2]					-25.23	-12.73		
					(-1.07)	(-0.51)		
YM FE	No	Yes	No	Yes	No	Yes	No	Yes
R^2	0.00105	0.0503	0.000700	0.0498	0.0000697	0.0491	0.000964	0.0858
Observations	24534	24534	24537	24537	24536	24536	24537	24537

Panel B: Alternative Market Scaling

	Month Lag		Quarter Lag		Year Lag	
	(1)	(2)	(3)	(4)	(5)	(6)
Mkt Div Pay[t-1,t]	59.26***	60.56***	59.49***	62.62***	54.41***	55.16***
	(3.48)	(3.29)	(3.63)	(3.51)	(3.54)	(3.29)
YM FE	No	Yes	No	Yes	No	Yes
R^2	0.000688	0.0497	0.000706	0.0498	0.000644	0.0497
Observations	24518	24518	24478	24478	24286	24286

Table II

Market Returns Based on Dividend Payment: Alternative Dividend Normalizations

This table shows how the market return varies with the dividend payment yield using different dividend normalizations and examining stocks by dividend-paying status. In Panel A, the dividend payment yield is measured as the sum of dividends paid on days t and t-1 divided by the average daily dividend payment on days t-20 through t-272. In Panel B, Top Days shows the coefficient on a dummy variable if the dummy variable listed in the column header is equal to 1. The 2 Weeks column include a dummy variable equal to one if the payment yield is in the top 10 of the prior 252 trading days, Quarter includes a dummy variable equal to one if the payment yield is in the top 63 of the prior 252 trading days, Third includes a dummy variable equal to one if the payment yield is in the top 84 of the prior 252 trading days, and Median includes a dummy variable equal to one if the payment yield is in the top 126 of the prior 252 trading days. Even-numbered columns contain year-by-month fixed effects. t-statistics based on heteroskedasticity-robust standard errors are in parentheses.

Panel A: Abnormal Dividend

	Value Weighted		Equal Weighted	
	(1)	(2)	(3)	(4)
Mkt Abnormal Div[t-1,t]	0.00698***	0.00786***	0.00637***	0.00790***
	(3.39)	(3.58)	(2.89)	(3.43)
YM FE	No	Yes	No	Yes
R^2	0.000499	0.0497	0.000424	0.0857
Observations	24266	24266	24266	24266

Panel B: Top Dividend Payment Days

	2 Weeks		Quarter		Third		Median	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Top Days	0.0983**	0.116***	0.0568***	0.0709***	0.0443***	0.0573***	0.0311**	0.0473***
	(2.39)	(2.75)	(3.56)	(4.27)	(3.04)	(3.71)	(2.28)	(3.15)
Constant	0.0370***	0.0363***	0.0266***	0.0231***	0.0260***	0.0217**	0.0252***	0.0171*
	(5.36)	(5.22)	(3.41)	(2.92)	(3.15)	(2.57)	(2.63)	(1.69)
YM FE	No	Yes	No	Yes	No	Yes	No	Yes
R^2	0.000315	0.0495	0.000535	0.0498	0.000387	0.0497	0.000215	0.0495
Observations	24286	24286	24286	24286	24286	24286	24286	24286

Table III

Equity Premium Based on Dividend Payment

This table shows how the annual market return and equity premium varies with the dividend payment yield. Columns labeled "Top" show values when the column header dummy variable is equal to 1 and rows labeled "Not" show values when the dummy variable is equal to 0. The first row shows the average cumulative return for each calendar year. The next row shows the fraction of the total (the sum of the Top and Not first row) the average in the first row represents. The third row shows the fraction of days each dummy variable is equal to 1. The final row shows how much of the equity premium was earned compared to the number of days the dummy variable was equal to 1 (i.e. row 2 divided by row 3). "2 Weeks" is a dummy variable equal to one if the payment yield is in the top 10 of the prior 252 trading days, "Quarter" is a dummy variable equal to one if the payment yield is in the top 63 of the prior 252 trading days, "Third" is a dummy variable equal to one if the payment yield is in the top 84 of the prior 252 trading days, and "Median" is a dummy variable equal to one if the payment yield is in the top 126 of the prior 252 trading days.

	2 Weeks		Quarter		Third		Median	
	Top	Not	Top	Not	Top	Not	Top	Not
Avg. Annual Return	0.014	0.102	0.058	0.053	0.065	0.047	0.077	0.035
% of Equity Premium	11.9	88.1	51.6	48.4	56.8	43.2	69.3	30.7
% of Days	3.8	96.2	24.8	75.2	33.2	66.8	50.0	50.0
% Premium/% Days	3.13	0.92	2.08	0.64	1.71	0.65	1.39	0.61

Table IV

Market Returns Based on Dividend Payment: Dividend-Paying vs Non-Dividend-Paying Stocks

This table shows how the market return varies with the dividend payment yield. The return is restricted to stocks that did not pay a dividend on days t and t-1 in the first four columns. The return is restricted to stocks that did pay a dividend on either t or t-1 in the last four columns. Even-numbered columns contain year-by-month fixed effects. t-statistics based on heteroskedasticity-robust standard errors are in parentheses.

	No Div Payment				Div Payment			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mkt Div Pay[t-1,t]	57.39***	65.12***	70.20***	74.39***	40.83**	45.91*	24.27	16.09
	(3.16)	(3.34)	(3.46)	(3.52)	(2.06)	(1.94)	(1.20)	(0.62)
YM FE	No	Yes	No	Yes	No	Yes	No	Yes
Value Weight	Yes	Yes	No	No	Yes	Yes	No	No
Equal Weight	No	No	Yes	Yes	No	No	Yes	Yes
R^2	0.000644	0.0494	0.000942	0.0855	0.000142	0.0486	0.0000404	0.0601
Observations	24537	24537	24537	24537	24355	24355	24355	24355

Table V

Market Returns Based on Dividend Payment: Calendar Effects and Macroeconomic Announcements

This table shows how the market return varies with the dividend payment yield controlling for patterns in returns related to calendar events and macroeconomic announcements. The first two columns of Panel A include dummy variables for the day of the week. The next two columns include a turn-of-the-month dummy variable equal to one if it is the last day of the month or the first three days of the month. The final two columns include both dummy variables. Panel B explores FOMC announcements. The announcement data ranges from 1988-2019, so the first two columns show the baseline analysis conducted over this period. The next two columns include dummy variables for days with FOMC announcements. Panel C examines macroeconomic announcements. The announcement dates range from 1994-2019, so the first two columns repeat the baseline analysis for this period. Columns 3 and 4 include dummy variables for CPI, PPI, employment and GDP announcements. Columns 5 and 6 also include dummy variables for FOMC announcements. Market return is measured in percent as the CRSP value-weighted market index. Market payment yield is the measured as the cumulative payment yield on day t and t-1 (Market Div Pay[t-1,t]). Even-numbered columns contain year-by-month fixed effects. t-statistics based on heteroskedasticity-robust standard errors are in parentheses.

Panel A: Day of Week and Turn of Month

	Day of Week		Turn of Month		Both	
	(1)	(2)	(3)	(4)	(5)	(6)
Mkt Div Pay[t-1,t]	74.46***	85.52***	32.01**	35.42**	48.00***	55.71***
	(5.16)	(5.54)	(2.14)	(2.19)	(3.19)	(3.43)
YM FE	No	Yes	No	Yes	No	Yes
R^2	0.00595	0.0552	0.00234	0.0514	0.00743	0.0565
Observations	24537	24537	24537	24537	24537	24537

Panel B: FOMC Announcements (1988-2019)

	(1)	(2)	(3)	(4)
Mkt Div Pay[t-1,t]	161.7***	174.8***	164.5***	179.8***
	(3.08)	(2.94)	(2.93)	(3.02)
FOMC	No	No	Yes	Yes
YM FE	No	Yes	No	Yes
R^2	0.00106	0.0348	0.00350	0.0370
Observations	7812	7812	7812	7812

Panel C: Macroeconomic Announcements (1994-2019)

	(1)	(2)	(3)	(4)	(5)	(6)
Mkt Div Pay[t-1,t]	223.5**	251.1***	228.0**	256.4***	230.8***	264.3***
	(2.25)	(2.67)	(2.56)	(2.71)	(2.60)	(2.80)
Macro	No	No	Yes	Yes	Yes	Yes
FOMC	No	No	No	No	Yes	Yes
YM FE	No	Yes	No	Yes	No	Yes
R^2	0.00101	0.0332	0.00140	0.0336	0.00426	0.0364
Observations	6294	6294	6294	6294	6294	6294

Table VI

Market Returns Based on Dividend Payment: Restricting to :Longer Times Between Payment and Announcement or Ex-Date

This table shows how the market return varies with the dividend payment yield after restricting to dividends with a larger number of days between announcement date and payment date (Panel A) and days between ex-date and payment date (Panel B). Market return is measured in percent as the CRSP value-weighted market index. Dividends are excluded if the number of days between the given date and the payment date is less than the first percentile of the measure (columns 1 and 2) or less than the 25th percentile of the measure (columns 3 and 4). The remaining dividends are used as a measure of the cumulative payment yield on day t and t-1 (Market Div Pay[t-1,t]). Even-numbered columns contain year-by-month fixed effects. t-statistics based on heteroskedasticity-robust standard errors are in parentheses.

Panel A: Days between Announcement Date and Payment Date

	1st Pctile (17 days)		25th Pctile (29 days)	
	(1)	(2)	(3)	(4)
Mkt Div Pay[t-1,t]	58.43***	66.42***	60.21***	67.45***
	(3.06)	(3.25)	(3.01)	(3.16)
YM FE	No	Yes	No	Yes
R^2	0.000524	0.0455	0.000505	0.0454
Observations	14096	14096	14096	14096

Panel B: Days between Ex-date and Payment Date

	1st Pctile (9 days)		25th Pctile (17 days)	
	(1)	(2)	(3)	(4)
Mkt Div Pay[t-1,t]	59.61***	67.04***	64.74***	73.44***
	(3.30)	(3.45)	(3.39)	(3.58)
YM FE	No	Yes	No	Yes
R^2	0.000696	0.0498	0.000666	0.0498
Observations	24537	24537	24537	24537

Table VII

Market Returns Based on Lagged Dividend Payment

This table shows how the current days' dividend is predicted by dividends from that day in prior years (Panel A) and how market returns vary with prior year's dividends (Panel B). Each column uses dividend payments of the indicated lag. For example, if the left hand side variable is from day D, month M and year Y, then the column labeled "X Yr Lag" indicates the dividend payment yield matched to that observation is from day D, month M, and year Y-X. Panel C performs the baseline regressions of contemporaneous measures of dividend payment on returns on the sample of observation included for a given lag in panel B. All regressions contain year-by-month fixed effects. Panel A reports the within R^2 which is the R^2 after accounting for the influence of the fixed effects. t-statistics based on heteroskedasticity-robust standard errors are in parentheses.

Panel A: Current Dividends predicted by dividends in prior years

	1 Yr Lag	2 Yr Lag	3 Yr Lag	4 Yr Lag	5 Yr Lag	6 Yr Lag	7 Yr Lag	8 Yr Lag	9 Yr Lag	10 Yr Lag
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Mkt Div Pay	0.789***	0.646***	0.640***	0.757***	0.822***	0.758***	0.582***	0.525***	0.576***	0.679***
	(210.45)	(116.03)	(100.31)	(119.40)	(193.88)	(198.36)	(113.91)	(84.23)	(88.18)	(121.34)
Within R^2	0.719	0.483	0.417	0.501	0.670	0.672	0.474	0.350	0.373	0.498
R^2	0.756	0.558	0.513	0.581	0.716	0.717	0.553	0.457	0.486	0.580
Observations	18449	15485	15149	15273	19597	20249	15446	14210	14073	15824

Panel B: Current Returns predicted by dividends in prior years

	1 Yr Lag	2 Yr Lag	3 Yr Lag	4 Yr Lag	5 Yr Lag	6 Yr Lag	7 Yr Lag	8 Yr Lag	9 Yr Lag	10 Yr Lag
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Mkt Div Pay	69.69***	62.79***	88.79***	51.29***	34.47**	73.16***	55.83***	73.58***	81.95***	58.85***
	(3.64)	(3.45)	(5.02)	(2.73)	(2.06)	(4.49)	(3.54)	(4.87)	(5.42)	(4.19)
Observations	18449	15485	15149	15273	19597	20249	15446	14210	14073	15824

Panel C: Current Returns predicted by current dividends restricting sample to Panel B Observations

	1 Yr Lag	2 Yr Lag	3 Yr Lag	4 Yr Lag	5 Yr Lag	6 Yr Lag	7 Yr Lag	8 Yr Lag	9 Yr Lag	10 Yr Lag
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Mkt Div Pay	92.24***	59.58**	90.76***	53.23**	54.73**	94.76***	69.73***	80.44***	85.68***	59.51***
	(3.79)	(2.51)	(3.41)	(2.20)	(2.51)	(4.52)	(2.60)	(4.53)	(4.88)	(3.68)
Observations	18449	15485	15149	15273	19597	20249	15446	14210	14073	15824

Table VIII

International Market Returns Based on Dividend Payment

This table shows how market returns vary based on the dividend payment yield across 58 international markets. Data is in a panel format with data from a market/day combination included when the market reports information for at least 100 stocks. The market return is calculated as the value-weighted averages of those stocks. Market payment yield is measured as the cumulative payment yield on day t and t-1 in that market. Columns 2 and 4 include month-by-market fixed effects. Columns 3 and 4 include date fixed effects. Standard errors are clustered by date and market, with t-statistics in parentheses

	(1)	(2)	(3)	(4)
Mkt Div Pay[t-1,t]	20.30**	30.21***	22.65***	26.89***
	(2.19)	(3.59)	(3.02)	(4.50)
Country YM FE	No	Yes	No	Yes
Date	No	No	Yes	Yes
R^2	0.0000330	0.0611	0.268	0.309
Observations	237185	237070	236775	236660

Table IX

Market Returns Based on Dividend Payment by Quarter and Month

This table shows how the market return varies with the dividend payment yield based on the calendar quarter or month. Market return is regressed on Market Div Pay[t-1,t] interacted with time period dummy variables. Columns 1 and 3 include quarter-of-year dummy variables and columns 2 and 4 include dummy variables for whether the calendar month is December, or other months of the year. All columns contain year-by-month fixed effects. t-statistics based on heteroskedasticity-robust standard errors are in parentheses.

	Value Weighted		Equal Weighted	
	(1)	(2)	(3)	(4)
Q1*Mkt Div Pay	134.8**		228.8***	
	(2.40)		(3.66)	
Q2*Mkt Div Pay	54.92*		31.60	
	(1.68)		(0.92)	
Q3*Mkt Div Pay	73.78**		53.16*	
	(2.44)		(1.91)	
Q4*Mkt Div Pay	12.01		-6.124	
	(0.35)		(-0.17)	
Not December*Mkt Div Pay		82.91***		89.99***
		(3.89)		(3.96)
December*Mkt Div Pay		-48.92		-56.87
		(-1.25)		(-1.30)
YM FE	Yes	Yes	Yes	Yes
R^2	0.0501	0.0501	0.0872	0.0862
Observations	24537	24537	24537	24537

Table X

Market Returns Based on Dividend Payment and Liquidity Measures

This table shows how the market return varies with the dividend payment yield interacted with measures of liquidity. Liquidity (LIQ) is proxied for by the VIX (Columns 1-2), the News VIX (Columns 3-4) and Amihud Illiquidity (Columns 5-6). Value-weighted market returns are regressed on the liquidity measure,along with Market Div Pay[t-1,t] and the interaction of the two. . Columns 2 and 6 include year-by-month fixed effects and Column 4 includes year fixed effects. t-statistics based on heteroskedasticity-robust standard errors are in parentheses.

	VIX		News VIX		Amihud	
	(1)	(2)	(3)	(4)	(5)	(6)
LIQ*Mkt Div Pay	30.08***	24.12***	10.74***	11.53***	17.72***	18.29***
	(3.79)	(2.95)	(2.89)	(3.05)	(5.88)	(5.79)
Mkt Div Pay	-416.5**	-240.6	-210.1**	-227.0**	0.250	0.579
	(-2.47)	(-1.40)	(-2.22)	(-2.36)	(0.01)	(0.03)
YM FE	No	Yes	No	No	No	Yes
Year FE	No	No	No	Yes	Yes	Yes
R^2	0.0191	0.122	0.00198	0.00690	0.00355	0.0562
Observations	7304	7304	23822	23822	24537	24537

Table XI

Stock Expense and Post Earnings Announcement Returns

This table shows characteristic-adjusted returns around earnings announcements as a function of the firm's stock expense. The dependent variable is stock returns at various horizons relative to the earnings announcement date (t=0), minus the returns on a portfolio matched on size, book-to-market ratio and momentum (Daniel et al. (1997)). Panel A measures stock expense based on the annual stock expense announced the prior quarter, scaled by market capitalization. Panel B measures stock expense using the quarterly stock expense in the current quarter divided by market capitalization one day before the earnings announcement. Dummy variables are created based on the percentile of stock expense. Values are equal to one if stock expense is in the range indicated in a row, with the omitted group being zero stock expense. Regressions other than column 1 include a control for the t=0 return, allowing for separate effect of positive and negative returns. Standard errors are clustered by date and firm, with t-statistics in parentheses.

Panel A: Prior Quarter Expense

	(1) (t=0, t+3)	(2) (t+1, t+3)	(3) (t+4, t+10)	(4) (t+4, t+20)
Stock Expense >0-80 Pctile	-0.184***	-0.00182	-0.0112	0.0301
	(-3.69)	(-0.05)	(-0.27)	(0.50)
Stock Expense 80-90 Pctile	-0.332***	-0.101	0.238***	0.0965
	(-2.93)	(-1.36)	(2.70)	(0.76)
Stock Expense 90-95 Pctile	-0.774***	-0.236**	0.152	0.477**
	(-4.83)	(-2.04)	(1.19)	(2.37)
Stock Expense 95-100 Pctile	-1.165***	-0.875***	0.455**	0.924**
	(-5.50)	(-5.77)	(2.06)	(2.46)
Constant	0.0987**	-0.0500	0.0819**	0.0465
	(2.24)	(-1.34)	(2.03)	(0.79)
t=0 Return	No	Yes	Yes	Yes
R^2	0.000453	0.000656	0.000289	0.000458
Observations	300456	300473	299789	298274

Panel B: Current Expense

	(1) (t=0, t+3)	(2) (t+1, t+3)	(3) (t+4, t+10)	(4) (t+4,t+20)
Stock Expense >0-80 Pctile	-0.173***	0.00492	-0.0259	0.00993
	(-3.42)	(0.14)	(-0.58)	(0.15)
Stock Expense 80-90 Pctile	-0.543***	-0.123	0.218***	0.359***
	(-5.06)	(-1.64)	(2.59)	(2.89)
Stock Expense 90-95 Pctile	-1.048***	-0.493***	0.151	0.176
	(-6.56)	(-4.24)	(1.19)	(0.76)
Stock Expense 95-100 Pctile	-1.271***	-0.884***	0.222	0.453
	(-5.65)	(-5.81)	(1.02)	(1.31)
Constant	0.123***	-0.0450	0.0973**	0.0619
	(2.73)	(-1.17)	(2.23)	(0.99)
t=0 Return	Yes	Yes	Yes	No
R^2	0.000682	0.000777	0.000233	0.000364
Observations	300456	300473	299789	298274

Table XII

Retail and Mutual Fund Dividend Reinvestment

This table examines how often individual investors and mutual funds reinvest dividends in their portfolio. Panel A considers observations of individual investors at the portfolio level on each day they have holdings, from an anonymous discount brokerage between 1991 and 1996. The dependent variable is a dummy equal to one if the investor made a purchase in any security that day. The independent variable is a dummy equal to one if the investor received a dividend that day or the day before. Column 2 adds account fixed effects, and column 3 adds both account and date fixed effects. Panel B examines monthly mutual fund portfolio changes from CRSP. The dependent variable is the monthly change in the fraction of TNA covered by holdings. The independent variable is the dividend yield those holdings received. Columns 2 and 4 include year-by-month fixed effects and column 4 includes fund fixed effects. Regressions are value-weighted by prior month TNA. The analysis includes observations with a prior month TNA above $10 million, with at least 10 holdings, and where the fraction of TNA covered by holdings is between 50% and 105% in the current and prior month. t-statistics are in parentheses, and are clustered by account and date in Panel A, fund and month in Panel B.

Panel A: Individual Investors

	(1)	(2)	(3)
Div Pay[t-1,t]	0.00176***	0.000307***	0.000370***
	(14.02)	(4.79)	(6.35)
Constant	0.00178***		
	(59.83)		
Account FE	No	Yes	Yes
Date FE	No	No	Yes
R^2	0.0000572	0.0165	0.0168
Observations	46000677	46000623	46000623

Panel B: Mutual Funds

	(1)	(2)	(3)	(4)
Div Received	0.530***	0.446***	0.692***	0.641***
	(4.39)	(3.88)	(4.95)	(4.59)
YM FE	No	Yes	No	Yes
Portno FE	No	No	Yes	Yes
R^2	198141	198139	197716	197714

Table XIII

Price Multiplier Estimates

This table presents estimates of price multipliers under different assumptions about reinvestment rates. Estimates are based on the coefficient of 67.07 from Table I Panel A Column 2. The first column presents multiplier estimates based on the assumption for aggregate reinvestment rate in Column 2. The multiplier estimate is 67.07 divided by the aggregate reinvestment percentage. The aggregate reinvestment rate is based on the assumptions about reinvestment rates for retail investors (column 3), mutual funds and ETFs (column 4) and all other investors (column 5). The aggregate reinvestment % implied by these numbers is the average of columns 3 through 5, weighted by the proportion of each investor type in the market, which is listed in the column headings and is based on the March 2020 Financial Accounts of the United States release Z.1.

Multiplier	Aggregate Reinvestment %	Retail (34%)	Mutual Funds & ETFs (28%)	Other (38%)
0.67	100%	100%	100%	100%
1.02	66%	0%	100%	100%
1.17	57.3%	0%	69%	100%
1.33	50.6%	0%	45%	100%
1.47	45.5%	0%	69%	69%
2.26	29.7%	0%	45%	45%