

Measuring Mutual Fund Flow Pressure as Shock to Stock Returns

MALCOLM WARDLAW*

ABSTRACT

A large and rapidly growing literature examines the impact of misvaluation on firm policies by using mutual fund outflow-induced price pressure to isolate nonfundamental price variation. I demonstrate that the standard approach to computing outflow-induced price pressure produces a measure that is inadvertently a direct function of a stock's actual realized return during the outflow quarter, raising doubts about its orthogonality to fundamentals. After removing these direct measurements of return, outflows generate a fairly negligible quarterly decline in returns, with no subsequent reversal, and many established results in this literature no longer hold. I provide suggestions for future analysis.

FINANCIAL ECONOMICS HAS LONG BEEN interested in whether nonfundamental movement in stock prices impacts corporate decision making. Empirically identifying such impact is challenging, however, as it requires an independent shock to stock prices that is both fully observable to the econometrician and completely orthogonal to firm fundamentals. Over the last 10 years, a rapidly expanding literature has used the investor flows to and from mutual funds as a source of exogenous price pressure. The idea behind this approach is that large investor redemptions may place pressure on mutual funds to sell the stocks they hold. If the required sales are sufficiently large, the funds' liquidity needs may put downward pressure on prices that is unrelated to the fundamental value of the underlying stocks. If cleanly identified, this price pressure then creates a laboratory for studying market feedback effects by breaking the endogenous relationship between prices and fundamental firm value.

*Malcolm Wardlaw is with University of Georgia, Terry School of Business. Special thanks to Elizabeth Berger, Nicholas Crain, Jonathan Cohn, Alex Edmans, Itay Goldstein, Wei Jiang, Uday Rajan, Breno Schmidt, Sheridan Titman, Wei Wei, Toni Whited, Steven Xiao, and Alex Young for comments and suggestions. Thanks also to seminar participants at the University of Michigan, University of Illinois at Chicago, University of Warwick, University of South Carolina, University of Georgia, Chinese University of Hong Kong, University of New South Wales, the 2019 Midwest Finance Association, and the 2019 Western Finance Association annual meetings. The author has no conflicts of interest as identified by *The Journal of Finance's* disclosure policy.

Correspondence: Malcolm Wardlaw, Terry School of Business, University of Georgia, 620 South Lumpkin Street, B332 Amos Hall, Athens, GA 30602; telephone: (706) 542-0356; e-mail: malcolm.wardlaw@uga.edu.

DOI: 10.1111/jofi.12962

© 2020 the American Finance Association

Interest in this approach began with Coval and Stafford (2007), who provide suggestive evidence of this flow pressure using observed mutual fund sales. However, measuring sell pressure with this approach is of limited usefulness in cleanly identifying nonfundamental shocks. Since the buy and sell actions of mutual funds are measured directly, these actions reflect of the information used in that decision. Edmans, Goldstein, and Jiang (2012), EGJ, propose solving this problem by estimating not the sells themselves, but a measure of the quarterly outflows to each fund scaled by the proportion of each stock that makes up the mutual fund's portfolio. Essentially, this measure estimates what the total sales of each stock in each fund would be if each stock was sold in proportion to the fund's initial beginning-of-quarter holding of that stock. Since the measure abstracts from which stocks are sold in the quarter and the information about firm fundamentals that those sales might contain, it potentially satisfies the exclusion restriction for an instrumental variables approach. Exposure to the measure also appears to result in large quarterly price declines followed by a reversion over the subsequent two years. This observation suggests that the price declines are solely due to the exposure to flow pressure in the event quarter and hence are unrelated to fundamentals before or during the quarter. EGJ use this measure to instrument for nonfundamental declines in market value and the effect that these declines have on takeovers.

Since the publication of EGJ, a large number of papers published in high-quality finance and business journals have used this measure and basic econometric approach to identify declines in corporate equity value that are unrelated to fundamentals and the impact of these declines on a wide array of corporate decisions. These papers include Derrien, Kecskés, and Thesmar (2013), who examine the impact on payout; Phillips and Zhdanov (2013), who examine the impact on R&D; Norli, Ostergaard, and Schindele (2015), who examine the impact on shareholder activism; Zuo (2016), who examines the impact on managerial earnings forecasts; Lee and So (2017), who examine changes in analyst coverage; Bonaime et al. (2018), who examine the impact on mergers and acquisitions; Eckbo, Makaew, and Thorburn (2018), who examine stock-financed takeovers; Lou and Wang (2018) who examine the impact on corporate investment; and Dessaint et al. (2018), who examine the intraindustry cross-firm impact on corporate investment.¹

In this paper, I demonstrate that this approach misidentifies the primary source of these declines in stock price. This misidentification occurs not

¹ Each of these articles is published or forthcoming in a high-quality finance or accounting journal such as *Journal of Finance*, *Journal of Financial Economics*, *Review of Financial Studies*, *Journal of Finance and Quantitative Analysis*, and *Journal of Accounting and Economics*. A number of papers such as Acharya et al. (2014), Deng, Hung, and Qiao (2018), and Bilinski et al. (2018) use this measure as a control in their tests. In addition, a large and growing number of current working papers also utilize this measure for identification, including but not limited to Agarwal and Zhao (2016), Badertscher, Shanthikumar, and Teoh (2017), Chang et al. (2017), Dong, Hirshleifer, and Teoh (2018), Gredil, Kapadia, and Lee (2018), Henning, Oesch, and Schmid (2015), Honkanen and Schmidt (2017), and Sun (2017). Recent working papers by Dessaint et al. (2019) and Gredil, Kapadia, and Lee (2019) have partially responded to the issues raised in an earlier working version of this paper.

because mutual fund flows or holdings may be related to fundamentals for economic reasons, but because large variation in returns is inadvertently introduced into the measure by construction. Because this large, unintended variation in returns is introduced directly, it is likely that other shocks affecting asset prices contaminate the returns. The contamination of the instrument can make it mechanically correlated with what it instruments for (the price discount), weakening the identification. In addition, I demonstrate that the apparent reversal simply reflects the well-documented size effect and does not represent an actual reversal.

Examination of this measure of price pressure, typically denoted by *MFFlow*, shows that it can be decomposed into three multiplicative terms. Of these three terms, only one is related to mutual fund flows. The remaining two are monotonic, nonlinear transformations of the reported return itself and the reported share volume. I demonstrate that the majority of the correlation between *MFFlow* and stock returns is due to independent variation in these directly measured quantities rather than variation in exposure to outflows.

Since the construction of *MFFlow* involves scaling by the *dollar* volume in each stock, the denominator contains the end-of-quarter price (P_t). The construction of *MFFlow* places the beginning-of-quarter price (P_{t-1}) in the numerator to calculate the dollar holdings of each stock by a group of mutual funds. Since the measure is weakly negative by construction, this function of prices represents a monotonic, increasing transformation of the realized quarterly return itself. This direct effect of the inclusion of the quarterly gross return in *MFFlow* is the primary determinant of the correlation between *MFFlow* and equity returns during the event quarter.

The impact of volume is slightly more subtle, but just as important. Abnormally high stock returns lead to a significant increase in volume for several quarters.² This effect is independent of mutual fund sell pressure, and represents an information or behavioral response by the market as a whole. Consequently, any measure that is scaled by share volume will naturally sort stocks on past and contemporaneous stock returns. This phenomenon is especially problematic when summing the measure up over multiple quarters or when examining models in which returns have a long-term effect on the outcome.

These two direct measures independently induce the majority of the correlation between quarterly returns and *MFFlow*. If some of the variation in returns and volume is driven by shocks from fund flows, the full relationship between *MFFlow* and stock returns may not be entirely mechanical. However, any measure of flow pressure that is constructed with these terms will vary significantly with shocks from all other sources, making it impossible to cleanly identify a nonfundamental change in returns.

While the direct inclusion of returns and volume creates a serious contamination problem for identification, a separate issue creates a problem for

² See Statman, Thorley, and Vorkink (2006), Griffin, Nardari, and Stulz (2007), Chordia, Huh, and Subrahmanyam (2007), and Lou and Shu (2017).

inference about the reversal used to justify the claim that *MFFlow* captures nonfundamental information. In particular, nearly all mutual funds have a minimum scale at which they invest in a given stock. A direct implication of this minimum scale is that the position of any *individual* fund, as a percentage of market capitalization, is greater for small-cap stocks than for large-cap stocks. The construction of *MFFlow* assumes that stocks are sold in proportion to their holdings. As a result, a disproportionate number of stocks in the extreme decile of *MFFlow* will be small-cap stocks by construction.

The construction of any measure of extreme flow pressure inherently requires the researcher to limit the number of exposed funds. As a result, a stock's exposure to flow pressure is often the result of only one or two funds. Consequently, the relative weighting of flow pressure more closely mirrors the holdings of *individual* mutual funds rather than the aggregate holdings of all mutual funds collectively.³

This sorting effect is *not* primarily a function of stock selection by the affected mutual funds. Rather, it holds for the stocks held by any randomly selected subset of mutual funds, and it is more pronounced when more funds are selectively omitted. Market-adjusted returns for these stocks would have increased, on average, even in the absence of a decrease in price pressure-driven outflow, simply because they are predominately small-cap stocks exposed to the size premium. In the absence of the size premium, this increase would not occur. This result removes one of the pillars supporting the claim that *MFFlow* captures nonfundamental variation in stocks prices, namely, that this variation is transitory rather than permanent.

I next propose several alternatives to the standard fund flow price pressure measure. Each of these alternative measures attempts to remove the unintended variation in returns that is not attributable to fund outflow pressure. For each of these constructions, I highlight the fundamental tradeoff between clean identification and power. I replicate several key tests from papers in various areas, and I examine how the results change when the direct market-wide measure of returns is excluded. In the bulk of these replications, the alternative measures usually fail to produce significant results, with the evidence indicating that existing results are likely contaminated by unintended fundamental variation in returns.

The results suggest that future research should treat existing evidence on the real effects of flow-induced mispricing with skepticism and seek other ways to reestablish these results. Additionally, future research should firmly reestablish that any measure of fire-sale pressure exhibits three important characteristics: It is free from unintended contamination by the scaling variables, exposed stocks exhibit abnormal return behavior around the event that is consistent for all quantiles of the measure rather than just the most extreme quantile, and the baseline cumulative average abnormal return (CAAR) fully

³ This effect will also occur in measures that measure the sells directly, since the magnitude of stock sales are fundamentally tied to the size of the holdings.

reflects known return anomalies that continuously impact the average abnormal return.⁴

It is worth emphasizing that EGJ's theoretical motivation and basis for the construction of *MFFlow* is reasonable, and the analysis of the approach initially passes most sensible hurdles for exogenous identifying variation. The contamination is subtle and quite difficult to spot without significant additional graphical and numerical analysis. This is evinced by the large number of researchers who have examined and used their instrument without recognizing the problem. Further, this paper is *not* advocating the use of fund sales directly as in Coval and Stafford (2007), which EGJ rightfully point out is contaminated by an inherent selection effect.

This paper's results complement those of a recent working paper by Berger (2019), but the methodology and scope of this paper are distinct. In her study, Berger (2019) first demonstrates that firms with extreme values of *MFFlow* are different along many observable dimensions such as leverage, cash flows, and stock price volatility. She then shows that firms with extreme values of *MFFlow* have similar conditional average values of corporate investment, issuance, and payout when compared with a control group matched on these observables.

This paper takes a different approach, decomposing the measure itself and identifying the sources of variation in detail. In doing so, I explain why the apparent sample selection arises in the first place. This decomposition is important because differences in observable characteristics do not all exist for the same reason and do not affect identification in the same way. More importantly, the primary source of sample selection that drives the key results in the literature arises because *MFFlow* directly sorts firms by their realized return. This fact makes sample matching on observables an ineffective solution since the factors driving the quarterly stock returns are largely unobservable, and addressing this unobservability problem is the point of the whole exercise.

The paper proceeds as follows. Section I provides the formal definition of the primary measure and the mathematical decompositions used in the paper. Section II replicates several key results in the literature and examines how the measure's construction impacts economic interpretation. Section III provides guidance on how to interpret evidence of a postevent reversal. Section IV concludes.

I. Measuring Mutual Fund Fire-Sale Pressure

Building on the initial work of Coval and Stafford (2007), who examine the return patterns of stocks that are sold by mutual funds with large outflows,

⁴ Note that while this paper proposes a potential fix to *MFFlow* by purging the original measure of its unintended contamination by returns and volume, this fix should not be used naively. The adjusted measure, termed the Flow-to-Stock, does not exhibit the type of decline and reversal that should be required by a nonfundamental shock to returns. It may, however, provide a starting point for a more sophisticated approach.

EGJ propose a measure of mutual fund fire-sale pressure to be constructed in such a way as to exclude any potential information effects implied by the act of selling by funds. In particular, they propose using the extreme outflows of a group of mutual funds scaled by the percentage of the mutual fund portfolio represented by each stock. They then sum the scaled flow measure over all mutual funds that experience large outflows and scale the price pressure by the dollar volume of the stock over the quarter. This measure, denoted by *MFFlow*, effectively captures the total dollar amount of each stock sold by these funds, scaled by its dollar volume, if all of the funds in question were to sell their stocks in proportion to their initial holdings.

The components of this measure calculated for each fund j , stock i , and quarter t are the net dollar flow to each mutual fund in the quarter ($F_{j,t}$), the percentage of the value of each fund j made up by each stock i at the end of the previous quarter ($s_{i,j,t-1}$), the dollar volume of each stock over the quarter ($VOL_{i,t}$), the shares held by each fund at the end of the last quarter ($SHARES_{i,j,t-1}$), the price of the stock at the end of the last quarter ($PRC_{i,t-1}$), and the total asset value of each fund at the end of the last quarter ($TA_{j,t-1}$). Formally, *MFFlow* is given as follows:

$$MFFlow_{i,t} = \sum_j^m \frac{F_{j,t} s_{i,j,t-1}}{VOL_{i,t}} \quad (1)$$

$$s_{i,j,t-1} = \frac{SHARES_{i,j,t-1} \times PRC_{i,t-1}}{TA_{j,t-1}} \quad (2)$$

$$\Rightarrow MFFlow_{i,t} = \sum_j^m \frac{F_{j,t} \times SHARES_{i,j,t-1} \times PRC_{i,t-1}}{TA_{j,t-1} \times VOL_{i,t}}, \quad (3)$$

conditional on the *outflow* of fund j being greater than 5% of total assets. That is, $\frac{F_{j,t}}{TA_{j,t-1}} < -5\%$.

The measure of EDJ thus calculates the percentage holdings of each fund at the beginning of the quarter (or, more precisely, the end of the previous quarter $t - 1$) and multiplies by the flow over quarter t scaled by the dollar volume over quarter t . In the calculation, only extreme outflows of 5% or more are considered because these are the funds most likely to be forced into a “fire sale” of their holdings. Because fund flows are measured as a net change, flow $F_{j,t}$ is always negative, and thus *MFFlow* is also always negative by construction. As a result, stocks with a more negative value of *MFFlow* experience greater outflow pressure and should see a larger sell-pressure-induced decrease in stock returns over the quarter.

A. Construction and Decomposition of *MFFlow*

Because the most extreme flows F are always negative, *MFFlow* is always weakly negative by construction. Therefore, it is helpful to recast the equation as a function of an always-positive outflow $|F| = -F \geq 0$:

$$MFFlow_{i,t} = (-1) \sum_j^m \left(\frac{|F_{j,t}| \times SHARES_{i,j,t-1} \times PRC_{i,t-1}}{TA_{j,t-1} \times VOL_{i,t}} \right). \quad (4)$$

It is important to note that $VOL_{i,t}$ is the *dollar* volume of stock i , not the share volume. Because F is denominated in dollars, when the measure is constructed this way, volume will also be in dollars. Specifically, $VOL_t \equiv SHARE_VOL_{i,t} \times PRC_{i,t}$. Note that because CRSP does not provide a measure of dollar volume, researchers manually calculate it this way directly from the data.⁵ Plugging this into equation (4) and rearranging terms gives

$$\begin{aligned} MFFlow_{i,t} &= (-1) \sum_j^m \left(\frac{|F_{j,t}| \times SHARES_{i,j,t-1} \times PRC_{i,t-1}}{TA_{j,t-1} \times SHARE_VOL_{i,t} \times PRC_{i,t}} \right) \\ &= (-1) \sum_j^m \left(\frac{|F_{j,t}| \times SHARES_{i,j,t-1}}{TA_{j,t-1}} \right) \left(\frac{1}{SHARE_VOL_{i,t}} \right) \left(\frac{PRC_{i,t-1}}{PRC_{i,t}} \right). \end{aligned} \quad (5)$$

This yields three grouped terms. Note that the last term, $\frac{PRC_{i,t-1}}{PRC_{i,t}}$, is the inverse of the gross return,

$$\frac{PRC_{t-1}}{PRC_t} = 1 / \left(\frac{PRC_t}{PRC_{t-1}} \right) = \frac{1}{1 + r_t}. \quad (6)$$

Because the gross return is weakly positive by construction, equation (6) is a monotonically decreasing transformation of the quarterly stock return, bounded between zero and ∞ . Multiplying by (-1) yields a monotonically *increasing* transformation bounded between $-\infty$ and zero.

It is worth stopping for a moment to emphasize the following point: because *MFFlow* always multiplies by lagged prices in the numerator whereas end-of-quarter prices are in the denominator, whatever is calculated in the rest of the equation is always multiplied by the gross return itself. Because the gross return is always strictly positive, the negative of the inverse of the gross return will always be a strict monotonic transformation of the quarterly net return. Stocks with very negative quarterly returns will produce very negative values of $-\frac{1}{1+r_t}$ and thus very negative values of *MFFlow*, while highly positive quarterly returns will produce negative values of $-\frac{1}{1+r_t}$ closer to zero and thus

⁵ See Internet Appendix, Section I, for more details on alternative constructions. The Internet Appendix may be found in the online version of this article.

negative values of $MFFlow$ that are also closer to zero, the upper bound of $MFFlow$.⁶

To properly scale the first two terms, I define share turnover as a function of share volume ($SHARE_VOL$) and total shares outstanding ($SHROUT$),

$$TURNOVER_{i,t} = \frac{SHARE_VOL_{i,t}}{SHROUT_{i,t-1}}.$$

I then rewrite equation (5), pulling all of the non- j terms out of the summation, as follows:

$$\begin{aligned} MFFlow_{i,t} &= (-1) \left(\frac{PRC_{i,t-1}}{PRC_{i,t}} \right) \left(\frac{1}{SHARE_VOL_{i,t}} \right) \sum_j^m \left(\frac{|F_{j,t}| \times SHARES_{i,j,t-1}}{TA_{j,t-1}} \right) \\ &= (-1) \left(\frac{PRC_{i,t-1}}{PRC_{i,t}} \right) \left(\frac{1}{SHARE_VOL_{i,t}} \right) \left(\frac{SHROUT_{i,t-1}}{SHROUT_{i,t-1}} \right) \sum_j^m \\ &\quad \left(\frac{|F_{j,t}| \times SHARES_{i,j,t-1}}{TA_{j,t-1}} \right) \\ &= (-1) \left(\frac{PRC_{i,t-1}}{PRC_{i,t}} \right) \left(\frac{SHROUT_{i,t-1}}{SHARE_VOL_{i,t}} \right) \left(\frac{1}{SHROUT_{i,t-1}} \right) \sum_j^m \\ &\quad \left(\frac{|F_{j,t}| \times SHARES_{i,j,t-1}}{TA_{j,t-1}} \right) \\ &= (-1) \left[\left(\frac{1}{1+r_{i,t}} \right) \left(\frac{1}{TURNOVER_{i,t}} \right) \left(\sum_j^m \frac{|F_{j,t}|}{TA_{j,t-1}} \times \frac{SHARES_{i,j,t-1}}{SHROUT_{i,t-1}} \right) \right]. \quad (7) \end{aligned}$$

In equation (7) we have three distinct terms, all of which are ultimately multiplied by -1 . Each of these terms is weakly positive by construction, and thus always have a proportional scaling effect on $MFFlow$ as a whole. The first term is the inverse quarterly gross return. The second term is the inverse of quarterly turnover. The third term contains the relative flow pressure on each stock as a percentage of the stock owned by each fund. This third term is now free of the direct mechanical effect of returns and turnover and contains only

⁶ One could choose to define dollar volume as a summation of $PRC \times SHARE_VOL$ over three separate months (or even over 63 trading days). This would make $PRC_{i,t}$ a more complicated weighted average over the end of the three previous months rather than a single end-of-quarter value. This slightly alters the form of r_t but not the mechanical effect, because the calculation still involves lagged prices over current prices. Because this construction precludes a closed-form algebraic decomposition of each component, the simpler method will be used. From a practical standpoint this makes very little difference. After winsorizing outliers, the two construction methods have a correlation of 0.995, and all of the graphs and large-sample statistical properties are essentially the same. See Internet Appendix Section I for more details.

the components that are mutual fund specific scaled by shares outstanding. I formally define the three terms in equation (7) as follows:

$$\begin{aligned}
 \text{Inverse Gross Return : } & \left(\frac{1}{1+r_{i,t}} \right) = \left(\frac{PRC_{i,t-1}}{PRC_{i,t}} \right) \\
 \text{Inverse Turnover : } & \left(\frac{1}{\overline{TURNOVER}_{i,t}} \right) = \left(\frac{SHROUT_{i,t-1}}{SHARE_VOL_{i,t}} \right) \\
 \text{Flow-to-Stock : } & \left(\sum_j^m \frac{|F_{j,t}|}{TA_{j,t-1}} \times \frac{SHARES_{i,j,t-1}}{SHROUT_{i,t-1}} \right).
 \end{aligned} \tag{8}$$

The relative empirical variation in each of these terms potentially affects the way in which *MFFlow* as a whole varies with returns. The construction of *MFFlow* is intended to proxy for *nonfundamental* movement in stock returns, but the measure is always highly correlated with returns primarily because (i) it is multiplied by a direct calculation of the quarterly return itself and (ii) it is scaled by overall volume, which is correlated with past returns independent of mutual fund participation. As I show below, variation in the direct return itself and contemporaneous volume are by far the most significant drivers of the effect this measure has in spreading quarterly returns, while mutual fund flows and holdings have little such effect.

Before moving to the numerical analysis, I define a fourth alternative measure, the Flow-to-Volume, as follows:

$$\text{Flow-to-Volume : } \left(\sum_j^m \frac{|F_{j,t}|}{TA_{j,t-1}} \times \frac{SHARES_{i,j,t-1}}{SHARE_VOL_{i,t}} \right). \tag{9}$$

Note that both Flow-to-Stock and Flow-to-Volume have a natural interpretation that is in line with to the original goal of the *MFFlow* construction. Equation (5) reduces to Flow-to-Volume if we simply change $PRC_{i,t}$ in the denominator to $PRC_{i,t-1}$. In this case, each mutual fund can be thought of as selling a fixed proportion of its shares in a stock and then scaling those sales by either the available number of shares or the available share volume.

B. Empirical Analysis of MFFlow

To analyze *MFFlow* and its components, I replicate the construction and empirical methodology given by EGJ. I obtain holdings data from CDA Spectrum/Thomson Financial and mutual fund flow and individual stock return data from CRSP. In replicating the main results, I obtain M&A data from Securities Data Company (SDC) and basic accounting data from Compustat following the instructions and time window given in the original paper. All data are pulled from their original source as of the beginning of 2019.

While the replication work in this paper follows the instructions exactly as given in the original paper, some details regarding the construction of *MFFlow* are not explicitly stated. For instance, while the original paper describes “remov[ing] funds that specialize in a single industry,” the mutual fund data do not contain a single flag for this specialization and a collection of flags must be

identified over time to satisfy this requirement. The paper also does not fully describe the somewhat complicated process through which fund flow data are matched to the 13F holdings data, which contain a number of data inconsistencies that must be addressed. Problematically, subsequent changes made to improve the quality of the holdings data may have also made this process more difficult. Where the original paper is unclear, I select from the set of possible criteria and choose the one that produces results that most closely match the summary statistics and results of the original paper. Slight differences exist in the tabulated results but these differences are minor. Summary statistics and the primary coefficients and standard errors from the main regressions can be matched to within a rounding error. Furthermore, most alternative choices in variable construction do not change the direction, statistical significance, or economic magnitude of the results. Details on construction can be found in the codebase provided in the Data Appendix of this paper, which constructs all variables and results from the unmodified core data sets.

To analyze the relative magnitude of each component, I create portfolios of stocks formed on deciles of *MFFlow* and each of the three individual components. I then graph their CAARs. Typically, as in EGJ, only the portfolio CAAR of the most extreme decile is graphed. However, it is important to graph all deciles to fully understand how the measure works. Approximately 40% of the stocks in the sample are not held by any fund experiencing these extreme outflows and therefore have an *MFFlow* of zero. Consequently, only six deciles of *MFFlow* have nonzero values. To allow for direct comparison, I normalize each component to zero when *MFFlow* = 0, with Decile 1 representing the most extreme decile, effectively multiplying each component by -1 . This allows the graph of each component to be compared to the graph for the full measure. The cumulative abnormal return is calculated each month, for 39 months, relative to its characteristic-matched portfolio on size, value, and momentum as in Daniel et al. (1997).⁷

Figure 1, Panel A, plots the CAARs for all deciles of *MFFlow*. The graph demonstrates not only the large negative abnormal event return in Decile 1, but also the progressively more positive abnormal return for stocks with values closest to zero, a phenomenon that theoretically should not occur. Figure 1, Panel B, plots the CAARs for deciles of Inverse Gross Return. Each decile shows a predictable pattern in whereby flows are spread during the quarter, with no meaningful reversal after the quarter ends. It is worth noting that all of the exposed stocks show an adjusted CAAR of around 4% to 5% in the run-up to the event quarter, such that any decile based on alternative sorting within these exposed stocks will have a pre-event CAAR that is biased upward.

Figure 1, Panel C, plots the CAAR for similar portfolio deciles constructed from Inverse Turnover. Share turnover has a well-known lagged relationship

⁷ Using a characteristic-adjusted rather than an equal-weighted benchmark is not important for the decomposition, but it will be important below when examining evidence of a postevent reversal. Note also that the tests replicated in Section II do not incorporate *any* adjustment as part of the actual tests, so the choice of adjustment here is only for the purpose of illustration.




Figure 1. CAAR path for all deciles of *MFFlow* and its components. This figure presents CAAR paths for portfolios constructed from all six nonzero deciles (i.e., the six deciles for which outflow pressure is nonzero) of *MFFlow*, Inverse Gross Return, Inverse Turnover, Flow-to-Stock, and Flow-to-Volume, and the portfolio constructed from stocks for which outflow pressure is zero. The abnormal monthly return is calculated by subtracting the monthly Daniel et al. (1997) size-value-momentum portfolio return from the monthly stock return. The deciles are calculated such that Decile 1 is the most extreme decile. That is, by multiplying by -1 to facilitate comparison with previous work. (Color figure can be viewed at wileyonlinelibrary.com)

with returns, with large excess returns generally followed by periods of high turnover. We see that this relation holds here, as high turnover during the quarter follows a large change in lagged excess returns, especially in the “highest” decile where turnover is very high (and negative Inverse Turnover is close to zero). These results explain a large part of the pre-quarter spread in CAAR seen in Figure 1, Panel A.

In addition, turnover and the negative of inverse turnover are positively correlated with both lagged and contemporaneous returns. The highest and lowest deciles of turnover see a spread in returns not only leading up to the event quarter, but *within* the event quarter as well. This effect is completely independent of any impact of mutual fund flows. Forward-looking returns, post quarter, are largely unaffected by turnover in line with weak-form market efficiency.

Note that turnover is given as share volume scaled by shares outstanding, but the strong correlation with returns does not depend on expressing it this way. Sorting on raw, unscaled share volume would produce similar graphs in which both lagged and contemporaneous quarterly returns are higher for stocks with higher quarterly volume. In short, any measure that varies substantially with share volume will produce similar results.⁸

In Figure 1, Panel D, I construct the same graph for deciles of the Flow-to-Stock measure, which removes the mechanical effect of returns and turnover from the equation. The results demonstrate very little impact on abnormal returns during the event quarter. The most extreme decile appears to show a CAAR of around 1% to 2% within the first two months *prior* to the event quarter in which the outflows occur, followed by a decline of 1% to 2% during the first two months of the event quarter itself. A formal test of the abnormal returns reveals that only the abnormal returns for these four months. That is, the two months prior to and the first two months during the event quarter, have statistically significant conditional correlations with Flow-to-Stock.

Finally, as an additional test, I plot the CAAR deciles of Flow-to-Volume in Figure 1, Panel E. Note that while the bottom decile shows a decline, this decline is roughly similar in magnitude to the decline formed when sorting on raw turnover. Moreover, the increasing nonzero deciles exhibit the same pattern of *outperformance* rather than progressively less underperformance, extending at least 12 months prior to the event. Returns remain unaffected going forward, so it is reasonable to consider using lagged volume as a scaling variable. However, because the volume effect clearly exhibits a permanent rather than temporary effect on returns, it is important for test to isolate any change in value to *only* the event period.

The graphical analysis in Figure 1 provides a stark projection of the impact of each component on return outcomes. The effect of *MFFlow* in spreading returns is dominated by the direct measurement of returns and volume rather than the impact of flow pressure. I find further support for this

⁸ Some working papers have also chosen to scale this measure by market capitalization rather than dollar volume. While doing so removes the volume effect, it does not remove the direct price effect because market capitalization is share price times shares outstanding.

observation in numerical analysis in Internet Appendix Section III, where I calculate semielasticities in a composite construction. I also show that while each of these three components may spread returns individually, the interaction between volume and flow does not meaningfully affect the abnormal returns. Consequently, while volume is a natural scaling variable to account for differences in available liquidity, these differences do not actually affect the impact of flow pressure.

Although none of the tests above rules out the possibility that part of the variation in returns and volume may be a direct result of fundflow pressure, the primary goal of constructing *MFFlow* is to isolate *only* the variation due to fund outflows. Because the directly measured return and volume components of *MFFlow* generate most of the identifying variation, such variation is at the very least significantly contaminated by fundamentals-based changes in returns. Because the interaction between these variables adds little to identification, restricting the identifying variation in *MFFlow* to come solely from the price pressure itself is therefore preferable.

II. Analysis of Key Results

I now examine how some of the key results in the literature are affected by the construction of the measure. I begin by replicating several established results from well-published papers in the field. I then examine the baseline results in context of the construction of the measure. Finally, I analyze how interpretation changes, if at all, when we apply different parts of the decomposition.

A. Takeovers

I begin by replicating the core result from EGJ, who first introduced the measure of mutual funds fire-sale pressure and the core econometric approach. In their paper, the authors study how market feedback effects impact takeover probability as a function of the inherent “discount” at which a firm hypothetically trades. The idea is that most firms likely trade at a market value that is a discount of the shadow value they would be worth if they were purchased by another firm. This discount is abstracted, but it could result from poor management that would be fixed in a takeover or by Coase (1937)-Williamson (1990) type corporate boundaries that would be more efficiently drawn around the target and acquirer. To measure this discount, the market value of the firm, as captured for by Tobin’s *Q*, is subtracted from a predicted Tobin’s *Q* from a quantile regression. In the simplest case, the predicted *Q* is just the value of Tobin’s *Q* at the 80th percentile of each industry-year. Firms below the 80th percentile are assumed to be trading at a discount, and the lower the Tobin’s *Q* relative to the industry benchmark, the higher the discount. Thus, the variable of interest is the negative of Tobin’s *Q*, adjusted and scaled by some benchmark.

A fundamental challenge in trying to measure the impact of the implied discount on takeovers is that the market-based measure of firm value will

Table I

Determinants of Discount (Edmans, Goldstein, and Jiang (2012))

This table reports results of an OLS regression of Discount on *MFFlow*, Flow-to-Stock, and a number of controls. The dependent variable is Discount, the value of predicted Q (the 80th percentile industry-year Tobin's Q) minus firm Q scaled by predicted Q. Year fixed effects are included in all specifications. Standard errors are double-clustered at the year and firm levels, and the corresponding *t*-statistics are reported in parentheses below the coefficients. The column dy/ex presents the semi-elasticity, the effect of a 1% change in the independent variable on Discount. All additional controls included in EGJ are included here: SalesRank, ATO, MktShr, Growth, Beta, Leverage, Payout, HHIFirm, HHISIC3, Amihud, and year fixed effects. *, **, and *** denote significance at 10%, 5%, and 1%, respectively.

	(1) Coef	(2) dy/ex	(3) Coef	(4) dy/ex
MFFlow	−0.00829*** (−6.84)	0.00852		
Flow-to-Stock			−0.000813 (−0.62)	0.00156
Controls	Y		Y	
<i>R</i> ²	0.0578		0.0566	
<i>N</i>	102,231	102,231	102,231	102,231

inherently contain information about whether the firm is a potential takeover target. Thus, there are two countervailing forces at work that make identification difficult: a direct *trigger effect*, whereby lower market prices and higher discounts induce a takeover effect because the assets are priced more cheaply, and an *anticipation effect*, whereby the higher probability of a takeover pushes the market price up as traders anticipate the takeover, in which case any direct estimate of the effect of market-implied discounts on takeover probabilities will be biased back towards zero. To cleanly shock the market price, and by extension the implied discount, EGJ use mutual fund fire-sale pressure as an instrument for the discount. Instrumenting in this way should theoretically isolate variation in the discount that is unrelated to the *anticipation* effect because the latter is a function of the fact that price changes reflect underlying economic fundamentals. In the model, fire-sale pressure is unrelated to changes in stock price that occur for any reason other than liquidity-induced mispricing. The instrumented discount should therefore result in a larger, more economically significant effect on takeover probabilities.

EGJ first calculate the baseline impact of this discount measure (Discount) on the probability of being taken over in the subsequent year. The authors find a statistically significant but economically small effect in the baseline specification. They then instrument for Discount using *MFFlow* summed over the previous year. They find that *MFFlow* is strongly negatively correlated with the discount because lower (i.e., more negative) values of *MFFlow* lead to lower stock prices and thus lower values of Q and higher values of Discount.

In Table I, I replicate the OLS regressions of Discount on *MFFlow* and several control variables. Column (1) reports the main results from EGJ.

Table II
Volume and the Time-Series Persistence of Discount

This table presents results of multiple OLS and fixed-effects regressions of Discount on volume scaled by fund flow and turnover at multiple annual lags. Standard errors are double-clustered at the year and firm levels, and the corresponding *t*-statistics are reported in the parentheses below the coefficients. All additional controls included in EGJ are included here. *, **, and *** denote significance at 10%, 5%, and 1%, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Flow/Lag Volume	-0.00402*** (-2.79)				-0.000981 (-0.91)	0.000964 (1.14)
Inverse Turnover _{<i>t</i>-1}		-0.0334*** (-3.09)				
Inverse Turnover _{<i>t</i>-2}			-0.0341*** (-3.75)			
Inverse Turnover _{<i>t</i>-3}				-0.0345*** (-3.98)		
Discount _{<i>t</i>-1}					0.585*** (34.70)	
<i>R</i> ²	0.0568	0.0568	0.0568	0.0569	0.391	0.492
<i>N</i>	102,231	102,231	102,231	102,231	90,875	100,763
Other Controls	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Firm FE	N	N	N	N	N	Y

As predicted, *MFFlow* shows up as a highly significant determinant of the discount. In column (3), I replace *MFFlow* with the annual Flow-to-Stock measure, which strips out the quarterly return and turnover components. The coefficient is negative but not statistically significant. Economic significance inferred by examining the semielasticities in columns (2) and (4) appears to be negligible as well.⁹

Using Flow/Lagged Volume as an alternative measure could represent a potential fix as it purges the event-period contamination from *MFFlow*, but this measure cannot be directly applied to the original specification without creating additional problems. Because the original EGJ specification is not a differencing regression and contains no firm fixed effects, any correlation that this measure has with *past* price increases will also be incorporated into the measure. The volume component of Flow/Volume is correlated with large and persistent increases in prices that occur over previous years. As a consequence, stocks with high lagged volume will have substantially higher valuation ratios in the cross-section even if no price increases occurred in the event period.

I examine this issue in Table II. Naively applying Flow/Volume to the original specification produces a reasonably statistically significant relationship between the measure and the discount. However, the problem with this approach

⁹ In this test the dependent variable *y*, the discount, is already naturally expressed as a percentage of the “true” value, so the semielasticity yields a percentage impact. The mean value of both measures is negative by construction, so the semielasticity will be positive.

Table III
Probability of Takeover Probit Model (Edmans, Goldstein, and Jiang (2012))

This table presents results of a probit and instrumented-probit regressions. The dependent variable is Takeover. All standard errors are adjusted for heteroskedasticity and correlation clustered at the firm level. The column dPr/dX gives the marginal effect on takeover probability of a one unit (or 100 percentage point) change in each regressor. Year fixed effects are included in all specifications. All additional controls included in EGJ are included here. The corresponding *t*-statistics are reported in parentheses below the coefficients. *, **, and *** denote significance at 10%, 5%, and 1%, respectively.

	(1) Probit	(2) dy/ex	(3) IV-MFFlow	(4) dy/dx	(5) IV-FTS	(6) dy/dx
Discount	0.156*** (10.29)	0.0172	0.778*** (3.91)	0.116	1.589*** (79.46)	0.528
Controls	Y		Y		Y	
N	102,231	102,231	102,231	102,231	102,231	102,231
Weak instrument tests						
<i>F</i> (1, # obs)			154.04		0.33	
<i>p</i> -Value			0.00		0.57	

is illustrated in columns (2) through (4) using only the Inverse Turnover component lagged one, two, and three years. Because this pure turnover effect creates a highly persistent change in the valuation ratio, the measure needs to somehow control for previous run-ups to isolate the impact of fund flows for that year.¹⁰ This can be achieved by controlling for the previous period's discount as in column (5) or by using firm fixed effects as in column (6). In both cases, the effect of Flow/Lagged Volume on the innovation in the discount variable is statistically insignificant, and in the case of a fixed effects regression exhibits the wrong sign. Given that volume scaling creates these additional problems, and the interaction between volume and flowpressure appears to produce little additional impact on returns, it may be preferable to simply avoid volume scaling altogether.

In Table III, I replicate the probit and instruments-probit results from EGJ. The baseline probit regression in columns (1) and (2) is statistically significant with a very small economic effect, while the instrumented regression in columns (3) and (4) produces an effect that is nearly seven times as large. While this is predicted under the theoretical model, the prediction does not rule out a substantially larger coefficient for purely statistical reasons when the instrument is inadvertently driven by fundamentals. Instrumenting with Flow-to-Stock produces an estimate over 30 times as large, although it fails the

¹⁰ Note that it is not sufficient to separately control for turnover. Inverse Turnover represents both something that needs to be controlled for (i.e., any existing past price run-ups implied by turnover) and a multiplicative component that one is trying to retain (i.e., the component that makes flow comparisons across stocks comparable). As a consequence, controlling for Inverse Turnover is an imperfect proxy for what one wants to control for, namely, the component of Inverse Turnover due to past price run-ups.

weak instruments test by a large margin. However, naively using the inverse gross return as an instrument also produces an estimate of similar magnitude, and predictably produces a very strong first stage.

B. Additional Replications

I also replicate and examine the results from two other papers. Lee and So (2017) examine the relation between mutual fund price pressure and analyst forecasts. Lou and Wang (2018) examine the relationship between flow pressure and corporate investment decisions. The full results are presented and discussed in Internet Appendix Section IV. Analyst forecast tests produce some significant results using Flow-to-Stock, but fail when using Flow/Lagged Volume. In the corporate investment models, the corrected measures fail to produce significant second-stage results. I also highlight the additional problems that arise due to persistence in the effect of past returns on long-horizon outcomes. In general, the use of mutual fund flow pressure as a source of clean identification is problematic, and orthogonality to firm fundamentals should not be naively assumed.

III. Evidence of Postevent Reversals


A. Portfolio Event Return Paths

The primary argument for why mutual fund outflows represent a nonfundamental shock to stock returns is that the graph of the equal-weighted CAAR path for the most extreme portfolio appears to show a steady postevent increase over the course of the next 24 months. However, this abnormal return does not represent a true “reversal.” Instead, the abnormal return is the consequence of two separate effects that add a constant positive return to the extreme portfolio. These effects are unrelated to the actual outflows and are a consequence of the portfolio sorting that occurs in constructing *MFFlow*.

The nature of these persistent effects can be noticed by comparing the plots of the equal weighted CAAR and the characteristic-adjusted CAAR for each of the portfolio deciles, as shown in Figure 2. First, note that the difference in postevent CAAR between deciles is completely eliminated by the characteristic adjustment. Second, note that all of the deciles appear to show an identical increasing CAAR after the characteristic adjustment.

In EGJ and most subsequent papers, only Decile 1 benchmarked against the CRSP equal weighted average is plotted. However, this plot is misleading for two reasons. First, because the construction of the measure naturally sorts stocks on size, the stocks in the bottom decile are tilted heavily towards small stocks relative to the top nonzero decile. This results in a CAAR for this decile that permanently outperforms the CRSP equal-weighted average. This can easily be seen in Figure 3, Panel A, by extending the CAAR graph of the bottom decile past 24 months. Rather than slowly “reverting” to its true

Panel A. CRSP Equal-Weighted Adjustment



Panel B. DGTW Adjustment





Figure 2. Visual impact of equal weighted abnormal returns. This figure presents CAAR paths for portfolios constructed from all six nonzero deciles of *MFFlow* and the portfolio constructed from stocks for which *MFFlow* = 0. Abnormal monthly returns are calculated by subtracting either the CRSP equal-weighted average or the monthly Daniel et al. (1997), DGTW, size-value-momentum portfolio return from the monthly stock return. (Color figure can be viewed at wileyonlinelibrary.com)

Panel A. Decile 1 extended to t+48 months



Panel B. All in-sample stocks




Figure 3. Constant outperformance of CAAR path. This figure presents the CAAR path, adjusted by the CRSP equal-weighted average, for a portfolio of all stocks in the sample (i.e., those stocks which are present and have valid data in CRSP and Thompson) and for Decile 1 extended to t + 48 months instead of t + 24 months. (Color figure can be viewed at wileyonlinelibrary.com)

pre-event value, the CAAR increases continuously as a function of the outperformance of the small stocks that dominate the portfolio.

Second, for reasons that have not been widely studied, a portfolio that consists only of the stocks held by mutual funds in Thompson/CDA Spectrum outperforms the CRSP equal-weighted average by around 2% to 3% per year. This is illustrated in Figure 3, Panel B, which plots the CRSP equal-weighted CAAR for a portfolio of all stocks in the entire sample for each event period. The CAAR graph of any randomly created portfolio will therefore naturally exhibit an abnormal return of 2% to 3% per year, which can also be seen in the parallel increasing CAAR for all portfolios in Figure 2, Panel B. This result is somewhat striking, as it would seem to imply either that equal-weighted CRSP is not an investable benchmark or stocks held by mutual funds outperform

Table IV
Abnormal Return Tests

Panel A reports the results of a calendar-time Fama and French (1993) three-factor regression of the monthly excess return over the risk-free rate, for a portfolio of stocks that have been in the bottom decile of *MFFlow* between 1 and 23 months prior. The standard errors have not been adjusted for serial correlation so as to not be overly conservative. Adjusting for serial correlation via the Newey-West procedure for any number of lags increases the standard errors slightly and makes the *t*-statistics smaller. Panel B reports results of a means test of the buy-and-hold abnormal return (BHAR) for a portfolio constructed of stocks from the bottom decile of *MFFlow*. Returns are calculated over two years post event, adjusted by their corresponding size-value-momentum characteristic-matched portfolio return. Returns are calculated for portfolios in the bottom decile and portfolios in the unaffected decile. That is *MFFlow* = 0. As returns overlap for seven quarters by construction, Hansen-Hodrick moving average standard errors are computed over seven lagged quarters. The corresponding *t*-statistics are reported in parentheses below the coefficients.

	Monthly Return	
α		0.000635
		(0.76)
Mkt-Rf		1.016
		(51.72)
SMB		0.741
		(25.95)
HML		0.406
		(13.55)
Observations		441

	Bottom Decile	<i>MFFlow</i> = 0
BHAR	0.0497	0.0545
	(3.67)	(2.36)
Observations	140	140

CRSP. The latter is unlikely given the lack of outperformance by the average mutual fund. Regardless of the cause, this empirical fact must be accounted for when drawing inferences about postevent reversals in stocks held by mutual funds.

The lack of reversal can also be demonstrated more rigorously by examining the three-factor alpha for the bottom decile portfolio using a formal asset pricing model. Table IV presents results from a monthly calendar-time Fama and French (1993) three-factor regression for any stock exposed to the bottom decile of *MFFlow* in the previous 24 months. The monthly excess alpha is only nine basis points and statistically insignificant. The returns also have a significant loading on small-minus-big (SMB), confirming their exposure to the size effect over the post event period. A buy-and-hold return test appears to show outperformance in the extreme decile, but such tests are known to be problematic at long horizons (see Fama (1998)). This is evidenced by the fact that the unaffected stocks show almost identical outperformance.

B. Measure Construction and Size

Finally, it is important to understand why the small stocks appear to sort heavily into the extreme deciles of *MFFlow*. Berger (2019) documents that funds that have extreme outflows are more likely to have investment objectives that favor small stocks. She suggests that this difference in investment objectives drives a selection bias that weights *MFFlow* towards small stocks.

Although this selection by funds does have a small effect on the size of the stocks in each decile of *MFFlow*, it is not the primary cause. Instead, the primary cause of this sorting is a function of the inherent overweighting of small stocks by *all* funds and the fact that *MFFlow*, as with most measures of flow pressure, sums over only a fraction of the total number of available funds. Nearly all mutual funds have a minimum scale at which they invest in a given stock. A direct implication of this minimum scale is that the position of any *individual* fund, as a percentage of market capitalization, is greater for small-cap stocks than for large-cap stocks. Construction of *MFFlow* assumes that stocks are sold in proportion to their holdings. As a result, a disproportionate number of stocks in the extreme decile of *MFFlow* will be small-cap stocks by construction.

This effect depends heavily on the fact that the construction of any measure of “extreme” flow pressure inherently requires the researcher to limit the number of exposed funds. Because the construction of *MFFlow* purposefully includes only funds with extreme outflows, omits sector-based funds, and excludes many mutual funds simply due to data limitations, many exposure weights are created by the holdings of only one or two funds. Consequently, the relative weighting of flow pressure more closely mirrors the holdings of individual mutual funds than the aggregate holdings of all mutual funds collectively.

This effect can be illustrated by randomly assigning shocks to mutual funds in a simulation. In each quarter, I take the outflow as a percentage of total assets of each fund, and I randomly reassign it to another fund in the available universe of funds. I then recalculate the full measure, Flow-to-Volume, and Flow-to-Stock using this randomly assigned extreme outflow. I simulate this random assignment 1,000 times and calculate the mean size of the stocks in the most extreme decile and the remaining deciles for all of the simulations. The results are presented in Table V. For each construction, the extreme decile of flow pressure has an average market capitalization that is significantly smaller than that of the remaining nonzero deciles in both the actual measure and the simulation.

Because this size sort occurs regardless of the stock selection made by a given fund, any measure of mutual fund flow pressure must at least account for this automatic sorting before considering whether there is selection by the funds themselves. Further, researchers must be careful when trying to address potential selection by funds. A common approach to addressing concerns about selection by mutual funds is to simply exclude certain types of funds. However, this approach intrinsically reduces the number of funds that are used

Table V
Simulated Outflows and Stock Size

This table presents the average of the total market cap (in \$100 millions) of all stocks in the 10th (most extreme) decile of each of three mutual fund flow measures and the average log assets of all stocks in the remaining nonzero deciles. These constructions are scaled by dollar volume, share volume, and shares outstanding. The left column of each group presents the results when constructing each measure using actual fund outflows, while the right column presents the average of 1,000 simulations in which outflows are randomly reassigned to any available fund in each quarter.

	Dollar Volume		Share Volume		Shares Out	
	Actual	Sim	Actual	Sim	Actual	Sim
10 th Decile	6.1	11.6	6.5	12.1	9.4	15.1
Remaining	30.1	27.7	30.0	27.6	29.4	27.0

to construct the measure. Reducing the number of funds available automatically reduces the number of funds that expose stocks to outflow pressure. The smaller the number of funds, the more the measure will reflect the weighting of individual funds rather than the weighting of the mutual fund market in aggregate. This exacerbates the sorting mechanism and potentially worsens the problem.

IV. Conclusion

The large-scale use of mutual fund outflows as instruments and proxies for non-fundamental price changes deserves a thorough reevaluation. The question of stock market feedback effects is an important one, and the promise of a clean, straightforward proxy for such changes is appealing. However, mutual fund price pressure is perhaps less well understood as a phenomenon than it should be. The impact of flows on price pressure deserves further study, and a comprehensive model, perhaps appealing to microstructure data, could shed light on the extent to which these price pressures can be accurately measured. Once we know this, we may better understand overall magnitude of the effect, as well as when and where it matters. Trying to construct a general measure of sell pressure using noisy, low-frequency data, however, is fraught with peril.

The use of trading and flow-based measures of nonfundamental stock performance in a corporate context could also benefit from a more thorough examination from an asset pricing perspective. Significant consideration should be given to the magnitudes resulting from any given change and whether they are compatible with a framework in which markets are at least somewhat efficient. More generally, greater scrutiny should be applied to parse and examine the empirical determinants of all constructed measures. This is especially true for instrumental variables, where measures that ostensibly satisfy the exclusion restriction may have unknown and unexplained correlations with fundamentals that invalidate that restriction.

The questions of both market feedback effects and mutual fund externalities are fundamental to our understanding of financial markets and corporate behavior. It is therefore important that researchers have a thorough understanding of what is known to be true about these economic forces. This paper helps clarify our understanding and provides tools that can help advance research in this area on a solid footing.

Initial submission: October 15, 2018; Accepted: November 19, 2019

Editors: Stefan Nagel, Philip Bond, Amit Seru, and Wei Xiong

REFERENCES

- Acharya, Viral, Heitor Almeida, Filippo Ippolito, and Ander Perez, 2014, Credit lines as monitored liquidity insurance: Theory and evidence, *Journal of Financial Economics* 112, 287–319.
- Agarwal, Vikas, and Haibei Zhao, 2016, What are the actual effects of cash holdings? Evidence from the mutual fund industry, SSRN Working Paper 2901723, Social Science Research Network, Rochester, NY.
- Badertscher, Brad A., Devin M. Shanthikumar, and Siew Hong Teoh, 2017, Private firm investment and public peer misvaluation, SSRN Working Paper 2689880, Social Science Research Network, Rochester, NY.
- Berger, Elizabeth, 2019, Does stock mispricing drive firm policies? Mutual fund fire sales and selection bias, SSRN Working Paper 3011027, Social Science Research Network, Rochester, NY.
- Bilinski, Pawel, Douglas J. Cumming, Lars Helge Hass, Konstantinos Stathopoulos, and Martin Walker, 2018, Strategic distortions in analyst forecasts in the presence of short-term institutional investors, *Accounting and Business Research* 49, 305–341.
- Bonaime, Alice, Huseyin Gulen, and Mihai Ion, 2018, Does policy uncertainty affect mergers and acquisitions? *Journal of Financial Economics* 129, 531–558.
- Chang, Ching-Hsing, Chinghung (Henry) Chang, Pi Kun Hsu, and Sheng-Yung Yang, 2017, The catalytic effect of internationalization on innovation, SSRN Working Paper 3041135, Social Science Research Network, Rochester, NY.
- Chordia, Tarun, Sahn-Wook Huh, and Avanidhar Subrahmanyam, 2007, The cross-section of expected trading activity, *Review of Financial Studies* 20, 709–740.
- Coase, Ronald, 1937, The nature of the firm, *Economica* 4, 386–405.
- Coval, Joshua, and Erik Stafford, 2007, Asset fire sales (and purchases) in equity markets, *Journal of Financial Economics* 86, 479–512.
- Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers, 1997, Measuring mutual fund performance with characteristic-based benchmarks, *Journal of Finance* 52, 1035–1058.
- Deng, Xin, Shengmin Hung, and Zheng Qiao, 2018, Mutual fund herding and stock price crashes, *Journal of Banking & Finance* 94, 166–184.
- Derrien, François, Ambrus Kecskés, and David Thesmar, 2013, Investor horizons and corporate policies, *Journal of Financial and Quantitative Analysis* 48, 1755–1780.
- Dessaint, Olivier, Thierry Foucault, Laurent Frésard, and Adrien Matray, 2018, Noisy stock prices and corporate investment, *Review of Financial Studies* 32, 2625–2672.
- Dessaint, Olivier, Jacques Olivier, Clemens A. Otto, and David Thesmar, 2019, Capm-based company (mis)valuations, SSRN Working Paper 3050928, Social Science Research Network, Rochester, NY.
- Dong, Ming, David A. Hirshleifer, and Siew Hong Teoh, 2018, Stock market overvaluation, moon shots, and corporate innovation, SSRN Working Paper 2878418, Social Science Research Network, Rochester, NY.
- Eckbo, B. Espen, Tanakorn Makaew, and Karin S. Thorburn, 2018, Are stock-financed takeovers opportunistic? *Journal of Financial Economics* 128, 443–465.

- Edmans, Alex, Itay Goldstein, and Wei Jiang, 2012, The real effects of financial markets: The impact of prices on takeovers, *Journal of Finance* 67, 933–971.
- Fama, Eugene, 1998, Market efficiency, long-term returns, and behavioral finance, *Journal of Financial Economics* 49, 283–306.
- Fama, Eugene, and Kenneth French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- Gredil, Oleg, Nishad Kapadia, and Jung Hoon Lee, 2018, Do rating agencies deserve some credit? Evidence from transitory shocks to credit risk, SSRN Working Paper 2998242, Social Science Research Network, Rochester, NY.
- Gredil, Oleg, Nishad Kapadia, and Jung Hoon Lee, 2019, Do rating agencies deserve some credit? Evidence from transitory shocks to credit risk, SSRN Working Paper 2998242, Social Science Research Network, Rochester, NY.
- Griffin, John M., Federico Nardari, and René M. Stulz, 2007, Do investors trade more when stocks have performed well? Evidence from 46 countries, *Review of Financial Studies* 20, 905–951.
- Henning, Laura Sophie, David Oesch, and Markus Schmid, 2015, Stock underpricing and firm news disclosure, Working Paper, University of St. Gallen, St. Gallen, Switzerland.
- Honkanen, Pekka, and Daniel Schmidt, 2017, Price and liquidity spillovers during fire sale episodes, SSRN Working Paper 2980006, Social Science Research Network, Rochester, NY.
- Lee, Charles M. C., and Eric C. So, 2017, Uncovering expected returns: Information in analyst coverage proxies, *Journal of Financial Economics* 124, 331–348.
- Lou, Xiaoxia, and Tao Shu, 2017, Price impact or trading volume: Why is the Amihud (2002) measure priced? *Review of Financial Studies* 30, 4481–4520.
- Lou, Xiaoxia, and Albert Y. Wang, 2018, Flow-induced trading pressure and corporate investment, *Journal of Financial and Quantitative Analysis* 53, 171–201.
- Norli, Øyvind, Charlotte Ostergaard, and Ibolya Schindele, 2015, Liquidity and shareholder activism, *Review of Financial Studies* 28, 486–520.
- Phillips, Gordon M., and Alexei Zhdanov, 2013, R&D and the incentives from merger and acquisition activity, *Review of Financial Studies* 26, 34–78.
- Statman, Meir, Steven Thorley, and Keith Vorkink, 2006, Investor overconfidence and trading volume, *Review of Financial Studies* 19, 1531–1565.
- Sun, Sophia, 2017, Do customers learn from stock prices? SSRN Working Paper 2871154, Social Science Research Network, Rochester, NY.
- Williamson, Oliver E., 1990, *Economic Organization: Firms, Markets and Policy Control* (New York University Press, New York).
- Zuo, Luo, 2016, The informational feedback effect of stock prices on management forecasts, *Journal of Accounting and Economics* 61, 391–413.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.
Replication code.