Money Illusion in Asset Pricing

Kelly Shue and Richard R. Townsend*

February 2, 2018

Abstract

A form of "money illusion" in financial markets may cause investors to think that news should correspond to a dollar change in price rather than a percentage change in price, leading to return underreaction for high-priced stocks and overreaction for low-priced stocks. Consistent with a simple model of money illusion, we find that total volatility, idiosyncratic volatility, and absolute market beta are significantly higher for stocks with low share prices or negative past returns. To identify a causal effect of price, we show that volatility increases sharply following stock splits and drops following reverse stock splits. The economic magnitudes are large: money illusion can explain a significant portion of the "leverage effect" puzzle, in which volatility is negatively related to past returns, as well as the volatility-size and beta-size relations in the data. We also show that money illusion distorts investor reactions to news that is itself reported in nominal rather than real units. Investors react to nominal earnings per share surprises, after controlling for the earnings surprise scaled by price. The reaction to the nominal earnings surprise reverses in the long run, consistent with correction of mispricing.

*Kelly Shue: Yale University and NBER, kelly.shue@yale.edu. Richard Townsend: University of California San Diego, rrtownsend@ucsd.edu. We thank Huijun Sun for excellent research assistance and the International Center for Finance at the Yale School of Management for assistance. We thank James Choi, Bryan Kelly, and Stefano Giglio for helpful comments.

1 Introduction

Seminal studies of money illusion have shown that people tend to focus on nominal rather than real wages. For example, workers may perceive a 2% cut in nominal wages in a zero-inflation environment as unfair, but a 2% rise in nominal income in a 4%-inflation environment as fair (e.g., Fisher, 1928; Benartzi and Thaler, 1995).[1] While numerous studies have examined money illusion in terms of the nominal versus real value of money, we argue in this paper that money illusion may encompass a broader class of biases in which people think about an asset's value in the wrong units. In particular, rational agents should always think in terms of proportional rather than nominal price changes in financial markets. The nominal price level of any financial security has no real meaning; its price can easily be changed through stock splits or reverse splits. What matters for financial securities is returns, i.e., the proportional change in price. However, changes in the value of stocks are frequently reported in dollar units rather than or in addition to percentage returns. For example, the print version of the Wall Street Journal historically only displayed the daily dollar change in share prices and modern apps such as the Apple iPhone stock application display only the dollar change in prices as the default option. Given the emphasis on dollar changes in share prices in the financial media, we hypothesize that money illusion may lead investors to think that a given piece of news should correspond to a certain dollar change in price rather than a percentage change in price.

For example, consider two otherwise identical stocks, one trading at \$20/share and another trading at \$30/share. Money illusion may lead investors to think the same piece of good news should correspect to a dollar increase in price for both stocks. Thinking about this news in dollar

[1] A number of related papers argue that money illusion with respect to nominal and real dollars may contribute to important patterns in macroeconomics, such as the relation between inflation and unemployment, price stickiness, and the non-indexation of contracts to inflation (see e.g., Fehr and Tyran, 2001; Shafir, Diamond, and Tversky, 1997). In addition, Cannon and Cipriani (2006) show that money illusion with respect to changes in currency surrounding the adoption of the Euro led consumers to mistakenly perceive the prices of goods.

rather than return units leads to return underreaction for the high-priced stock at \$30/share and overreaction for the low-priced \$20/share stock. For a given sequence of news, money illusion would then lead to high return volatility for low-priced stocks and low return volatility for high-priced stocks. Similarly, money illusion may lead investors to overreact to relevant macro news for low-priced stocks, leading to higher absolute market beta for lower priced stocks.

Consistent with these predictions, we find that lower nominal share price is associated lower volatility, measured in three ways: total return volatility, idiosyncratic volatility, and absolute market beta. The economic magnitudes are large: a doubling in share price corresponds to a 20-30 percent reduction in these three measures of volatility. The results hold both in the cross section and in panel regressions that control for fixed characteristics of each stock. They also hold for stocks in the recent time period and among stocks with high share prices. Thus, the results cannot be driven by historical tick-size limitations.

Of course, the negative relation between volatility and nominal share prices could still be caused by other factors besides money illusion. In particular, it is widely known in the asset pricing literature that small-cap stocks tend to have higher total volatility, idiosyncratic volatility, and market beta, possibility because small-cap stocks are fundamentally more risky. Small-cap stocks also tend to have lower nominal share prices, so the price-volatility relation in the data could be driven by size. However, we find that the negative price-volatility relation remains equally strong after introducing flexible control variables for market capitalization. Moreover, the negative relation between size and volatility flattens by more than 80% after we introduce a single control variable for nominal share price. Thus, the results suggest that money illusion may explain the size-volatility relation rather than the reverse.

Aside from size, we remain concerned that other omitted factors may drive the negative relation between price and volatility. For example, low nominal share price is often the result of negative

past returns, and poor past performance may directly be associated with volatility and risk. To better account for this and other potential omitted factors, we conduct an event study around stock splits. Following a standard 2-for-1 stock split, the share price falls by half, while the fundamentals of the stock are unlikely to have dramatically changed. We find a sharp discontinuity around stock splits: the stock's return volatility, idiosyncratic volatility, and absolute market beta increase by approximately 30 percent immediately after the split. Further, the volatility remains high with a gradual monotonic decline over the course of the next twelve months. We further find sharp declines in volatility following reverse stock splits (e.g., when 2 shares become 1 share), in which the share price jumps up.

We also show that our results are unlikely to be caused by potentially endogenous determinants of stock splits. Firms that are doing well or that have experienced positive returns are more likely to engage in stock splits. However, we show that these trends are slow moving and cannot explain the sharp rise volatility immediately after stock splits. It's also possible that stock splits raise investor attention, which may directly impact volatility. However, we show that attention as proxied by volume quickly reverts to its low original level after a split. In contrast, volatility after stock splits remains high for over a year.

Our empirical results so far are consistent with a simple model of money illusion in which investors react to news with a reference point for a share price in mind. Investors observe the magnitude of the news and choose a dollar reaction to the news that approximately translates to the correct rational percentage price change to the news if the share price equaled the reference price. This reference price could the price of a typical stock in the market or the actual share price in the recent past or just before a stock split. The fact that volatility sharply rises and then gradually declines following stock splits is further consistent with a model in which investors gradually update the reference price toward the current stock price. To explore the rate at which investors update a

stock's reference price, we look at the relation between volatility and the stock's past returns over various return windows. By holding the total return over various time horizons fixed, we can vary the rate at which prices have changed. We find that the negative relation between past returns and subsequent realized volatility becomes weaker the farther back the return window is extended. In other words, a stock that has doubled in the last month is significantly more volatile than a stock that doubled in value over the last six months, which is in turn more volatile than a stock that has doubled in value over the the last year. These results suggest that investors gradually and incompletely update reference prices toward the current price level, implying that money-illusion-induced misreaction to news should be greater for stocks that have experienced recent large absolute returns. These results also show that money illusion may contribute to the well-known "leverage effect," in which volatility is negatively related to past returns (e.g., Black, 1976). While a number of papers (e.g., Christie, 1982) argue that the negative return-volatility relation may be due to leverage (as asset values decline and debt stays approximately constant, the equity becomes more leveraged and therefore more risky), other research (e.g., Figlewski and Wang, 2001) cast doubt on the leverage explanation for the leverage effect. We show that money illusion offers a compelling alternative explanation for this empirical pattern: as prices decline, volatility increases because investors react to news in dollar units based upon a higher reference price and thereby overreact in percentage units.

In the final part of the paper, we explore a related prediction relating to money illusion. We hypothesize that money illusion may distort investors' reactions to news if the news itself is reported in nominal rather than the appropriate proportional units. In the case of firm earnings announcements, the right measure of the news should the nominal value of the earnings surprise, *scaled* by the firm's price just before the news is released. For example, earnings news in which a firm beats analyst expectations by 10 cents per share is a greater positive surprise if the firm's share price

is \$20/share than if the firm's share price is \$30/share. However, investors may mistakenly react to the nominal earnings surprise because that is the value that is most commonly reported in the financial press. We find that investors react strongly to nominal earnings per share surprises, after controlling for the earnings surprise scaled by share price. If prices move toward fundamentals in the long run, we expect the initial return reaction to the nominal earnings surprise to reverse over time as the mispricing is corrected. Because investors react to the nominal earnings surprise, they also underreact to the scaled earnings surprise, so we expect the scaled earnings surprise to predict future drift in prices. Consistent with these predictions, we find that returns drift in the direction of the scaled earnings surprise and against the direction of the nominal earnings surprise in the long run.

Our results contribute to the literature in two main ways. First, we document a new form of money illusion and show that it can significantly affect financial market outcomes. To date, relatively few papers have explored the implications of thinking about value in the wrong units in financial markets. In related work, Shue and Townsend (2017) show that the tendency to think about executive option grants in terms of the number of options rather than the Black-Scholes value contributed to the dramatic rise in CEO pay starting in the late 1990s. Birru and Wang (2015) use option data to show that nominal price illusion causes investors to mistakenly believe that low-priced stocks have more "room to grow."

Second, our findings shed light on the potential origins of volatility and risk in financial markets. Since Shiller (1981), academics have explored the question of what factors determine volatility and risk. Our results suggest that money illusion may be an important part of the explanation and that well-known asset pricing facts such as the leverage effect and the size-volatility relation can be reinterpreted through the lens of money illusion.

2 A simple model

Consider a stock with current share price P. Let P_0 be the reference price for the stock in the minds of investors. P_0 could be the price of a typical stock in the stock market, the price of the stock in a previous time period, or the price of the stock prior to a stock split. Suppose news Z is released that contains information relevant for the valuation of the stock. If markets are fully efficient and rational, the release of news Z should imply a δ fractional change in the price of the stock, i.e., δ is the rational return reaction to the news. However, money illusion may lead investors to apply a heuristic and think that news Z should move prices by a nominal amount X. The dollar movement of X is such that it roughly equals the rational return reaction if the stock's price equaled the reference price P_o, i.e., $X = \delta P_0$. Thus, money illusion implies the return reaction to news Z is $\frac{X}{P} = \frac{\delta P_0}{P}$. If we allow investors to partially suffer from money illusion, the return reaction the news Z can be expressed as:

$$r = \theta \frac{\delta P_0}{P} + (1 - \theta)\delta \tag{1}$$

$\theta \in [0, 1]$ measures the extent to which investors suffer from money illusion. If investors are fully rational, $\theta = 0$, and the return reaction $r = \delta$. If investors fully suffer from money illusion, $\theta = 1$, and the return reaction to the news behaves as though the stock had an reference price P_0, leading to $r = \delta P_0/P$.

This simple framework delivers a number of testable predictions. First, whether investors under- or overreact to news will depend on the ratio of the reference price to the current price: P_0/P. If the stock's price is high relative to the reference price, then investors will underreact to the news, leading to $|r| < |\delta|$. If the stock's price is low relative to the reference price, then investors will overreact to the news, leading to $|r| > |\delta|$. Second, this initial under- or overreaction represents mispricing, which implies drift patterns if prices converge toward fundamentals in the long run.

Specifically, if the stock's price is high relative to the reference price, we expect continued drift to correct for the initial underreaction. If the stock's price is low relative to the reference price, we expect a long run reversal to correct for the initial overreaction. Third, for a given sequence of news over time, we expect the return volatility of the a stock to be higher when the stock's price is lower related to the reference price. The higher volatility arises from the return overreaction to each piece of news. Finally, we expect the absolute value of the market beta of a stock to be higher if its price is lower relative to the reference price, because the return for the stock will overreact to market-level news. Note that money illusion amplifies the absolute value of beta rather than beta. For example, if a stock's true beta is negative and the market news is positive, the stock's share price should drop and the share price should drop by more if investors overreact to the news.

To test these predictions, we examine cases in which P_0/P is likely to be low or high. First, P_0 may be a simple constant representing a typical share price in the market, e.g. \$25/share. If so, P_0/P is high for stocks with low nominal share price and low for stocks with high nominal share price. Second, P_0 may be the price of a stock just prior to a stock split event. After a 2-for-1 stock split $P_0/P = 0.5$, so we expect return overreaction, leading to higher return volatility and higher absolute beta. Looking at stock splits is also helpful for identifying a money illusion effect, because differences in return patterns before and after the stock split cannot be due to fundamental differences in returns patterns across different stocks. Further, while the general occurance of a stock split is likely to be endogenous to other factors that may affect returns (e.g., firms to engage in stock splits following good performance), the exact timing of when a stock split occurs is more likely to be exogenous, allowing us to estimate a regression discontinuity using return patterns immediately before and after the stock split. Finally, investors may think of each stock's reference price as its price at some period in the past. Therefore, stocks that have recently decreased in value will have $P_0/P > 1$, so we would again expect return overreaction, leading to higher return volatility and

higher absolute beta. Looking at how volatility and betas are related to past returns also helps us to identify how investors think about and update the reference price for each stock. We can compare the explanatory power of past returns measured over various time horizons.

To summarize, when prices are low (high) relative to a reference price, we expect:

1. Initial overreaction to news (initial underreaction to news)
2. Long run reversal (long run drift)
3. Higher volatility (lower volatility)
4. Higher beta (lower beta).

Because we don't always usually the arrival of specific pieces of news, we will focus our baseline analysis on testing the third and fourth predictions, which can be tested even if the news itself is not observed. In supplemental analysis, we attempt to isolate large news shocks and test for initial under- or over-reaction and subsequent corrections through either long-run drift or reversals.

We also hypothesize that money illusion may lead investors to exhibit biased reactions to news that is itself reported in nominal rather than real terms. We consider the case of earnings surprises, which are often reported in terms of the nominal surprise (the raw difference between actual earnings and analyst forecasts) rather than the scaled surprise (the nominal surprise divided by the share price just before the news is released). If investors are fully rational, they should only react to the the scaled surprise. However, if investors fixate on the nominal surprise, we predict that returns will also react to the nominal surprise. On net, this leads to initial overreaction to the nominal earnings surprise and underreaction to the scaled surprise. If prices converge toward fundamentals in the long run, then the initial mispricing should reverse. We predict continued drift in the direction of the scaled surprise and a reversal with respect to the nominal surprise.

3 Data

3.1 Stock market data

We obtain stock market data from CRSP, which offers information relating to returns, nominal share prices, stock splits, daily high and low, volume, and market capitalization. Data on the market excess return, risk-free rate, SMB, HML, UMD, and size category cutoffs come from the Kenneth French Data Library. We measure the return for day t as the return from market close on day $t-1$ to market close on day t.

The sample is restricted to stocks that are publicly traded on the NYSE, American Stock Exchange, or NASDAQ. We also restrict the sample to assets that are classified as common equity (CRSP share codes 10 and 11). To reduce the influence of outlier share prices, we exclude the top and bottom 1% of the sample in each year-month period in terms of 1-month lagged shared price from the analysis.

In our baseline tests, we measure firm i's volatility in month m as vol_{im}, equal to the annualized standard deviation of daily returns within each calendar month. In some robustness tests, we instead measure volatility as the mean of squared daily returns, again annualized. We measure each firm's monthly market beta as $beta_{im}$, equal to the covariance between daily firm excess returns and market excess returns divided by the variance of daily market excess returns within each calendar month. We measure each firm's idiosyncratic volatility as $ivol_{im}$, equal to the standard deviation of the firm's daily abnormal returns, where abnormal return is defined as the firm return minus $beta_{im}$ multiplied by the market return.

Our baseline tests uses observations at the firm-month level. To control for each firm's market capitalization, we match each firm's size of at the end the previous month to size vigintile categories as defined using the size cutoff data in the Kenneth French Data Library. To classify firms by

nominal share price, past returns, etc., we always use past information.

3.2 Earnings data

We use the I/B/E/S detail history file for data on analyst forecasts as well as the values and dates of earnings announcements. The sample is restricted to earnings announced on calendar dates when the market is open. Day t refers to the date of the earnings announcement listed in the I/B/E/S file.[2] We examine the quarterly forecasts of earnings per share.

The two key variables in our analysis are the nominal surprise for a given earnings announcement and the scaled surprise.[3] Broadly defined, the earnings surprise is the difference between announced earnings and the expectations of investors prior to the announcement. We follow a commonly-used method in the accounting and finance literature and measure expectations using analyst forecasts prior to announcement. This measure is available for a long time-series and does not require us to take a stand on specific modeling assumptions (for example, assuming that expectations follow random walk with drift as in Bernard, 1992).[4]

Following the methodology in Hartzmark and Shue (2018) and related studies of investor reactions to earnings announcement, we take each analyst's most recent forecast, thereby limiting the sample to one forecast per analyst, and then take the median of this number within a certain time window for each firm's earnings announcement. In our base specification, we take all analyst

[2]DellaVigna and Pollet (2009) highlight a potential concern regarding earnings announcement dates as reported in I/B/E/S. We address this issue by showing that our results continue to hold in a more recent time period in which there exists exact information on both the day and hour of each earnings announcement.

[3]We follow the literature on earnings announcements in characterizing earnings news as the surprise relative to expectations. We focus on surprise rather than levels because whether a given level of earnings is good or bad news depends on the level relative to investor expectations. In addition, stock prices should reflect current information—the stock market return response to an earnings announcement represents the change in a firm's valuation, which should depend on the difference in earnings relative to expectations. Moreover, the financial press typically reports earnings announcement news in terms of how much earnings beat or missed forecasts. Therefore, the earnings surprise is likely to be the measure of earnings news that is most salient to investors.

[4]Analysts are professionals who are paid to forecast future earnings. While there is some debate about how unbiased analysts are (e.g., McNichols and O'Brien, 1997; Lin and McNichols, 1998; Hong and Kubik, 2003; Lim, 2001; and So, 2013), our tests only require that such a bias is not correlated with the difference between the nominal and scaled earnings surprises.

forecasts made between two and thirty days prior to the announcement of earnings. We choose thirty days to avoid stale information and still retain a large sample of firms with analyst coverage. Our results remain qualitatively similar if we use alternative windows of 15 or 45 days prior to announcement.

We define the nominal earnings surprise as the dollar difference between actual earnings and the median analyst forecast:

$$nominal\,surprise_{it} = actual\,earnings_{it} - median\,estimate_{i,[t-30,t-2]}. \tag{2}$$

We define the scaled earnings surprise as the nominal earnings surprise divided by the share price of the firm three trading days prior to the announcement:

$$scaled\,surprise_{it} = \frac{\left(actual\,earnings_{it} - median\,estimate_{i,[t-30,t-2]}\right)}{price_{i,t-3}}. \tag{3}$$

Most of the existing literature exploring the return reactions to earnings surprises focus on the scaled surprise measure. Scaling by price accounts for the fact that a given level of earnings surprise implies different magnitudes of news shocks depending on the price per share. However, many media outlets report earnings surprises as the nominal (unscaled) difference between actual earnings and analyst forecasts, and investors may mistakenly pay attention to the nominal surprise. Therefore we compare how markets react to the nominal and scaled surprise measures. To reduce the influence of outliers, which may be relatively more problematic for the scaled surprise measure because it is measured as a ratio, we measure both the scaled and nominal surprise as percentile rank variables.

We then construct measures of returns over various event windows around the earnings announcement. We measure the direct short-term reaction to the earnings announcement as the firm's abnormal return in the window $[t-1, t+1]$, i.e., the firm's return from market close on

day $t-2$ to market close on $t+1$, minus the market return over the same period. In unreported results, we find similar patterns using characteristic-adjusted returns instead of returns in excess of the market. We can also test for long-run drift and reversals by examining the firm's abnormal returns over longer event windows.

4 Results

4.1 Prices, volatility, and beta

We begin by exploring how return volatility varies with the share price. Using data at the stock-month level, we estimate the following regression:

$$log\left(vol_{it}\right) = \beta_0 + \beta_1 log\left(price_{i,t-1}\right) + size_{i,t-1} + \tau_t + \nu_i + \epsilon_{it}. \tag{4}$$

We regress each stock i's volatility in month t on the stock's nominal share price at the end of the previous month, indicator variables for 20 size categories using the stock's market capitalization at the end of the previous month relative to other stocks in the same time period, calendar-year month fixed effects, and stock fixed effects. We measure volatility and nominal share price in logarithm form because a simple model of money illusion implies that volatility should change proportionately with share price. To account for correlated observations, we double-cluster standard errors by stock and year-month.

We present our baseline results in Table 1, which begins with the minimum set of control variables and builds up to the full specification as described above. Consistent with the predictions from a simple money illusion model, we find that higher nominal share price is associated with lower total return volatility. The negative coefficient on price remains highly significant and stable in magnitude as we introduce control variables for size (either as the log of lagged market capitalization or with

20 size bin indicators) and time fixed effects. The results hold in the cross section (without stock fixed effects) and in the panel (with stock fixed effects). The economic magnitudes are also quite large. With the full set of control variables in Column 4, a doubling in share price is associated with a 34% decline in volatility in the cross section and a 27% decline in volatility in the panel, i.e., within stock over time.

In Tables 2 and 3, we find very similar empirical patterns after replacing the dependent variable with idiosyncratic volatility and absolute market beta. The economic magnitudes are again large. With the full set of control variables in Column 4, a doubling in share price is associated with 35% decline in idiosyncratic volatility and a 31% decline in absolute market beta. In the panel, a doubling in price is associated with a 29% decline in idiosyncratic volatility and a 19% decline in absolute market beta. As discussed previously, we use the absolute value of market beta instead of the raw level of beta because money illusion should lead to overreaction to market news for low priced stocks, resulting in larger betas for positive-beta stocks and more negative betas for negative-beta stocks. However, one may be concerned that stocks with measured betas in the negative range may simply be stocks where beta is measured with error. To show that this is unlikely to drive our results, we present results in the Online Appendix that restrict the sample to observations with positive market betas and continue to find a very similar relation.

4.2 Size and risk

The empirical patterns shown so far are consistent with a money illusion story. However, share prices are not randomly assigned, so an omitted factor could determine both price and volatility. Our results can already reject one key alternative explanation involving size: It is a well-known in the asset pricing literature that small-caps, i.e., stocks with low market capitalization, tend to have higher return volatility, idiosyncratic volatility, and market beta. The size-volatility relations

in the data may even be viewed by some as unsurprising, given that it seems plausible that small stocks may be fundamentally more risky. Since small-cap stocks also tend to have low nominal share prices, size may simultaneously determine share price and volatility.

However, we showed in Tables 1-3 that the coefficient on lagged share price remains stable in magnitude and significant after controlling for the logarithm of lagged market capitalization or after controlling for 20 bins for market capitalization. We also see in Columns 2 and 3 of each table that, while size negatively predicts volatility if we do not control for price, the size-volatility relation flattens toward zero once we control for lagged nominal share price.

We explore the relation between size and volatility in more detail in Figure 1. Panel A shows the coefficients from a regression of log volatility on 20 bins for size (the largest size bin is the omitted category), after controlling for year-month fixed effects. As expected, we find a strong negative relation between size and volatility. In Panel B, we report the same set of coefficients for the 20 bins representing size, after adding a single control variable for the log of the lagged nominal share price to the regression. We see that the relation between size and volatility flattens dramatically. In the range between size categories 4 and 20, size continues to negatively predict volatility. However, the magnitude of the slope shrinks by more than 80 percent. Thus, the results suggest that money illusion may explain a significant portion of the well-known size-volatility and size-beta empirical relation rather than the reverse.

4.3 Stock splits

Aside from size, we remain concerned that other omitted factors may drive the negative relation between price and volatility. For example, low nominal share price is often the result of negative past returns. If bad times or poor performance causes an increase in risk and volatility, then negative past returns may directly increase our measures of volatility. To better account for this and other

potential omitted factors, we conduct an event study around stock splits. Following a standard 2-for-1 stock split, the share price falls by half, but the fundamentals of the stock are unlikely to have dramatically changed, an assumption we verify in later analysis. Therefore, we can more credibly attribute changes in volatility or beta after a stock split to the change in share price.

To explore how volatility changes after splits, we estimate the following regression:

$$log\left(vol_{it}\right) = \beta_0 + \beta_1 Post_{it} + \tau_t + \nu_i + \epsilon_{it}. \tag{5}$$

The regression sample uses observations at the stock-month level. We begin by considering any positive stock split in which one old share is converted into two or more new shares (the results remain qualitatively similar if we restrict the definition of an event to 2-for-1 stock splits, the most common type of split). The month of a stock split counts as event date 0, and the sample is restricted to observations in the window $[t-12, t+12]$ around the event date, conditional on the observation not being in the windows of another stock split for the same stock. The coefficient of interest is β_1, which measures the difference in volatility in the twelve months after the split relative to the twelve months before. If the drop in share prices following a stock split leads to an increase in volatility, we expect that $B_1 > 0$. To examine how volatility varies with event time in greater detail, we also present results in which the $Post_{it}$ indicator is replaced with event-month indicators. In all specifications, we control for year-month and stock fixed effects and double cluster standard errors by stock and time.

We find that volatility rises significantly after stock splits. Table 4shows that total volatility, idiosyncratic volatility, and absolute beta increase by approximately 20 percent in the 12-month period after the split relative to the 12-month period before the split. If we introduce event month indicators for the post-split period as in the even-numbered columns in Table 4, we find that

volatility sharply increases in the first month after the stock split and remains high relative to the pre-period for the next 12 months. There is also evidence of a monotonic decay over time: the initial jump in volatility in the first month after the stock split falls by approximately 35% over the next 12 months.

In Figure 2, we explore volatility around stock splits in more detail. We plot the coefficients for each month in event time, relative to the omitted category of 12 months prior to the split. We find that there is a slight pre-trend in that volatility rises in the six months leading up to the split. This pretrend is consistent with the view that splits are not random. Managers of firms choose to engage in stock splits following periods of good performance, which may coincide with small increases in volatility. However the direction and magnitude of the pre-trend in volatility cannot explain the sudden large jump in volatility after the split, nor the slow monotonic decay in volatility over the next 12 months.

In Figure 3, we explore patterns of volatility following reverse stock splits, defined as events in which two more stocks are converted to one stock. We expect volatility to decline following reverse stock splits because the nominal share price increases by two-fold or more following the reverse split. Consistent with these predictions, we find that volatility drops by more than 20 percent in the month following reverse stock splits and the drop remains persistent over the next 12 months. As with positive stock splits, we observe a positive pre-trend in volatility in the months leading up to the split. However, the direction and small absolute magnitude of the pre-trend in volatility cannot explain the sudden and persistent drop in volatility following the split.

These results related to reverse splits also help to rule out a potential alternative explanation related to increased investor attention. Stock splits typically receive media coverage, which may increase investor attention, leading to higher volatility in the period after the split. However, investor attention should also increase following reverse splits, and yet we find that volatility declines

following these reverse stock splits, consistent with a money illusion model. Moreover, we note that any changes in investor attention following stock splits are unlikely to persist over a multi-month horizon, and therefore cannot explain the persistent changes in volatility.

In our final set of tests relating to stock splits, we move from monthly stock data to daily stock data. The more granular daily stock data allows us to estimate a regression discontinuity around the date of the stock split. While stock splits are unlikely to be randomly assigned across firms, the fundamentals of each firm are unlikely to change exactly on the date of each stock split. Therefore, we can credibly attribute changes in volatility immediately after the split to the change in share price. For the regression discontinuity, we change our measure of volatility from the standard deviation of daily returns within each calendar month to the scaled intraday price range, defined that the difference between the intraday high and low, scaled by the share price at market close on the previous day. This scaled intraday price range allows us to proxy for daily volatility even in the absence of minute-by-minute pricing data.

In Figure 4 (regression results in Table 5), we find that the scaled intraday price range increases by 0.014 immediately after the split, an approximate 40 percent increase relative to the pre-split scaled intraday price range. These magnitudes are very similar regardless of the exact regression discontinuity method that we adopt. We fit local linear regressions on either side of the regression discontinuity, using either a triangular or Epanechnikov kernel, and the rule-of-thumb optimal bandwidth.

4.4 Past returns

In this section, we explore the relation between a stock's volatility and past returns. If a stock has experienced negative past returns, then it's current share price is more likely to be low relative to the reference price in the minds of investors. Therefore, we expect a negative relation between past

returns and volatility. Examining past returns over various windows also allows us to see if the evidence is consistent with a model in which investors use a stock's past share price as a reference price. In the analysis in the previous section, the fact that volatility sharply rises and then gradually declines following stock splits is suggestive of a model in which investors gradually, and incompletely, update the reference price toward the current price. To explore the rate at which investors update a stock's reference price level, we look at the relation between volatility and the stock's past returns over various return windows. By holding fixed the total return over various time horizons, we can vary the rate at which prices have changed.

Table 6 shows the regression results. We estimate the following regression:

$$log\left(vol_{it}\right) = \beta_0 + \beta_1 returns_{i,[t-x,t-1]} + \tau_t + \epsilon_{it}, \tag{6}$$

where $returns_{i,[t-x,t-1]}$ represents past returns over the past 2, 4, 6, 8, 10, and 12 month windows. Consistent with money illusion, we find a strong negative relation between past returns and volatility. We also find that the negative relation between past returns and subsequent realized volatility becomes weaker the farther back the return window is extended. In other words, a stock that has doubled in the last two months is significantly more volatile than a stock that doubled in value over the last six months, which is in turn more volatile than a stock that has doubled in value over the the last year. These results suggest that investors gradually and incompletely update reference prices toward the current share price over time.

These results also have important implications for how prices will react to news as a function of past returns. We expect that money-illusion-induced misreaction to news will be greater for stocks that have experienced recent changes in stock price levels. For example, stocks that have recently experienced strong negative returns will have current share prices that are farther below

their reference prices, leading to greater overreaction and volatility.

These results also show that money illusion may contribute to the well-known "leverage effect," in which volatility is negatively related to past returns (e.g., Black, 1976). While a number of papers including Christie (1982) argue that the negative relation between returns and volatility may be due to leverage (as asset values decline and debt stays approximately constant, the equity becomes more leveraged and therefore more risky), other research (e.g., Figlewski and Wang, 2001) cast doubt on the leverage explanation for the leverage effect. We show that money illusion offers a compelling alternative explanation for this empirical pattern: as prices decline, volatility increases because investors react to news in dollar units based upon a higher reference price and thereby overreact in percentage units.

4.5 Reactions to nominal earnings news

In the section, we explore a related prediction from a simple model of money illusion. We hypothesize that money illusion may distort investors' reactions to news if the news itself is reported in nominal rather than the appropriate proportional units. In the case of firm earnings announcements, the right measure of the magnitude of the news should the nominal value of the earnings surprise, scaled by the firm's price just before the news is released. For example, earnings news in which a firm beats analyst expectations by 10 cents per share is a greater positive surprise if the firm's share price is \$20/share than if the firm's share price is \$30/share.

However, the financial press commonly reports the nominal earnings surprise of 10 cents per share, without scaling by share price. Therefore, money illusion may lead investors to react to the nominal earnings surprise instead of, or in additional to, the scaled earnings surprise. To test this, we estimate the following regression:

$$CAR_{i,[t-1,t+1]} = \beta_0 + \beta_1 nominal\, surprise_{it} + \beta_2 scaled\, surprise_{it} + \epsilon_{it}$$

$CAR_{i,[t-1,t+1]}$ is the firm's cumulative abnormal return over the three-day window around the earnings announcement. If markets are fully rational, we expect to find that $\beta_2 > 0$ because bigger real surprises should correspond to larger return reactions, and β_1 to be close to zero, because the nominal surprise should contain no information beyond the scaled surprise.[5] On the other hand, money illusion predicts that $\beta_1 > 0$.

When estimating this regression, we measure the nominal surprise and the scaled surprise as percentile rankings for two reasons. First, percentile rankings allows for a direct comparison of the relative magnitudes of the return reaction to each type of earnings surprise. Second, percentile rankings reduce the potential influence of outliers, particularly for the measure of the scaled surprise which can take on very large values when the denominator approaches zero. Expressing earnings surprise in ranked form also follows the convention in the earnings literature, e.g., Dellavigna and Pollet (2009) and Hartzmark and Shue (2018).

Table 7 Panel A presents the results. We find that investors indeed react strongly to nominal surprises. If fact, the return reaction to the nominal surprise is slightly larger in magnitude and more statistically significant than the reaction to the scaled surprise. These results holds in the full sample, as well as the subsamples for large cap and small cap firms in Columns 3 and 4, respectively. This shows that these patterns are not driven only by small firms. However, the relative return reaction to the nominal surprise is much larger for the sample of small cap firms, consistent with

[5]If markets are rational and the scaled surprise variable is a sufficient statistics for news, it's still possible that the coefficient on nominal surprise would be non-zero. It could be that the rational relation between CAR and scaled surprise is non-linear. Since the nominal surprise is correlated with the scaled surprise, β_1 may be non-zero to pick up part of this nonlinear relation. To address this possibility, we focus on long run reversals, which provide more direct evidence of mispricing and a subsequent correction.

a story in which the investor base for small cap firms are less sophisticated or a story in which arbitrage frictions or shorting constraints are more likely to apply to small cap firms. In Columns 5 and 6, we show that the same patterns hold before and after the year 2001. In the more recent time period, I/B/E/S data records the announcement date and time of earnings announcement with greater accuracy. Finding similar results in the more recent sample period shows that these patterns cannot be explained by date recording errors.

If prices move toward fundamentals in the long run, we expect the initial return reaction to the nominal earnings surprise to reverse over time as the mispricing is corrected. Because investors react to nominal earnings surprise, they also underreact to the scaled earnings surprise, so we expect the scaled earnings surprise to predict future drift in returns. Consistent with this, we find that returns continue to drift in the direction of the scaled earnings surprise and against the direction of the nominal earnings surprise in the long run. Table 7 Panel B shows the long run return reactions to the scaled surprise and the nominal surprise. We find over the course of 100 trading days after earnings announcement that the return reaction to the nominal surprise converges toward zero and the return reaction to the scaled surprise increases in magnitude, consistent with a correction of the initial overreaction tot he nominal surprise and underreaction to the scaled surprise.

5 Robustness

A potential concern with our findings is that the negative price-volatility relation may be driven by tick-size limitations. A tick size is the minimum price movement for a financial security. Tick size as a fraction for share price is larger for stocks with lower nominal share price, which may artificially inflate the measured volatility of low-priced stocks.

In Figure 5, we show that tick size cannot explain our results. We plot the coefficients of a regression of volatility on 20 equally spaced bins in nominal share price, controlling for 20 size

category bins, and stock and time fixed effects. All plotted coefficients measure the difference in volatility within each share price bin relative to the omitted bin of 20 (the largest share price). We observe a strong monotonic negative relation between volatility and share price. The negative relation holds even in the range of very high nominal share price bins, when tick size limits should have minimal impact.

The strong monotonic pattern in this figure also shows that our findings of a negative relation between volatility and nominal share price are unlikely to be driven by a few outlier observations. Rather, the negative relation holds between any two adjacent nominal price bins.

6 Conclusion

We show that a form of money illusion in financial markets may cause investors to think that news should correspond to a dollar change in price rather than a percentage change in price, leading to return underreaction for high-priced stocks and overreaction for low-priced stocks. Consistent with a simple model of money illusion, we find that total volatility, idiosyncratic volatility, and absolute market beta are significantly higher for stocks with low share prices or negative past returns. To identify a causal effect of price, we show that volatility increases sharply following stock splits and drops following reverse stock splits. The economic magnitudes are large: money illusion can explain a significant portion of the "leverage effect" puzzle, in which volatility is negatively related to past returns, as well as the volatility-size and beta-size relations in the data. We also show that money illusion distorts investor reactions to news that is itself reported in nominal rather than real units. Investors react to nominal earnings per share surprises, after controlling for the earnings surprise scaled by price. The reaction to the nominal earnings surprise reverses in the long run, consistent with correction of mispricing.

References

Benartzi, Shlomo and Richard H. Thaler, 1995, Myopic loss aversion and the equity premium puzzle, *Quarterly Journal of Economics* 110, 73–92.

Bernard, Victor Lewis, 1992, *Stock Price Reactions to Earnings Announcements: A Summary of Recent Anomalous Evidence and Possible Explanations, in: R. Thaler, ed., Advances in Behavioral* (Russell Sage Foundation, New York, NY).

Birru, Justin and Baolian Wang, 2015, The nominal price premium, *Working Paper* .

Black, Fischer, 1976, Studies of stock price volatility changes, *Proceedings of the Business and Economics Section of the American Statistical Association* 177–181.

Cannon, Edmund and Giam Pietro Cipriani, 2006, Euro-illusion: A natural experiment, *Journal of Money, Credit and Banking* 38, 1391–1403.

Christie, Andrew A., 1982, The stochastic behavior of common stock variances: Value, leverage and interest rate effects, *Journal of Financial Economics* 10, 407 – 432.

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

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

Fehr, Ernst and Jean-Robert Tyran, 2001, Does money illusion matter?, *American Economic Review* 91, 1239–1262.

Figlewski, Stephen and Xiaozu Wang, 2001, Is the 'leverage effect' a leverage effect?, *Working Paper* .

Fisher, Irving, 1928, *The money illusion* (Adelphi Company, New York, NY).

Hartzmark, Samuel M. and Kelly Shue, 2018, A tough act to follow: Contrast effects in financial markets, *Journal of Finance, forthcoming* DOI: 10.3386/w23883.

Hong, Harrison and Jeffrey D Kubik, 2003, Analyzing the analysts: Career concerns and biased earnings forecasts, *Journal of Finance* 313–351.

Lim, Terence, 2001, Rationality and analysts' forecast bias, *The Journal of Finance* 56, 369–385.

Lin, Hsiou-wei and Maureen F McNichols, 1998, Underwriting relationships, analysts' earnings forecasts and investment recommendations, *Journal of Accounting and Economics* 25, 101–127.

McNichols, Maureen and Patricia C O'Brien, 1997, Self-selection and analyst coverage, *Journal of Accounting Research* 167–199.

Shafir, Eldar, Peter Diamond, and Amos Tversky, 1997, Money illusion, *The Quarterly Journal of Economics* 112, 341–374.

Shiller, Robert J., 1981, The use of volatility measures in assessing market efficiency, *The Journal of Finance* 36, 291–304.

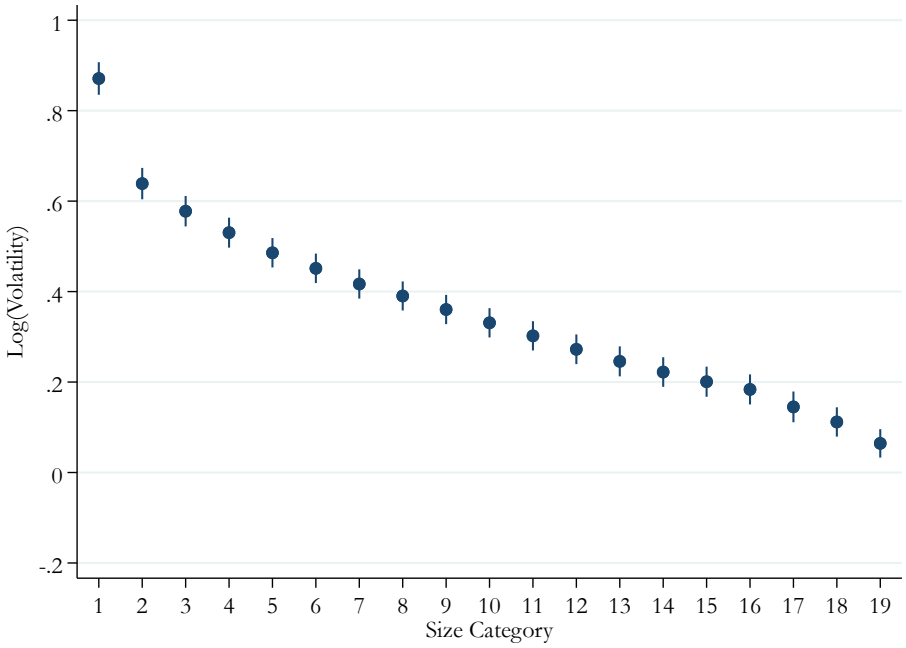
Shue, Kelly and Richard R. Townsend, 2017, Growth through rigidity: An explanation for the rise in CEO pay, *Journal of Financial Economics* 123, 1 – 21.

So, Eric C, 2013, A new approach to predicting analyst forecast errors: Do investors overweight analyst forecasts?, *Journal of Financial Economics* 108, 615–640.

Figure 1
Size and Volatility

This figure explores the relation between market capitalization and volatility. Panel A shows the coefficients from a regression of log volatility on 20 bins for size (the largest size bin is the omitted category), after controlling for year-month fixed effects. In Panel B, we report the same set of coefficients for the 20 bins representing size, after adding in a single additional control variable for the log of the lagged nominal share price. The dots represent the point estimates and the lines represent 95% confidence intervals. Standard errors are clustered by stock and month.

Panel A: Market Capitalization and Volatility, No Stock Price Control



Panel B: Market Capitalization and Volatility, Stock Price Control

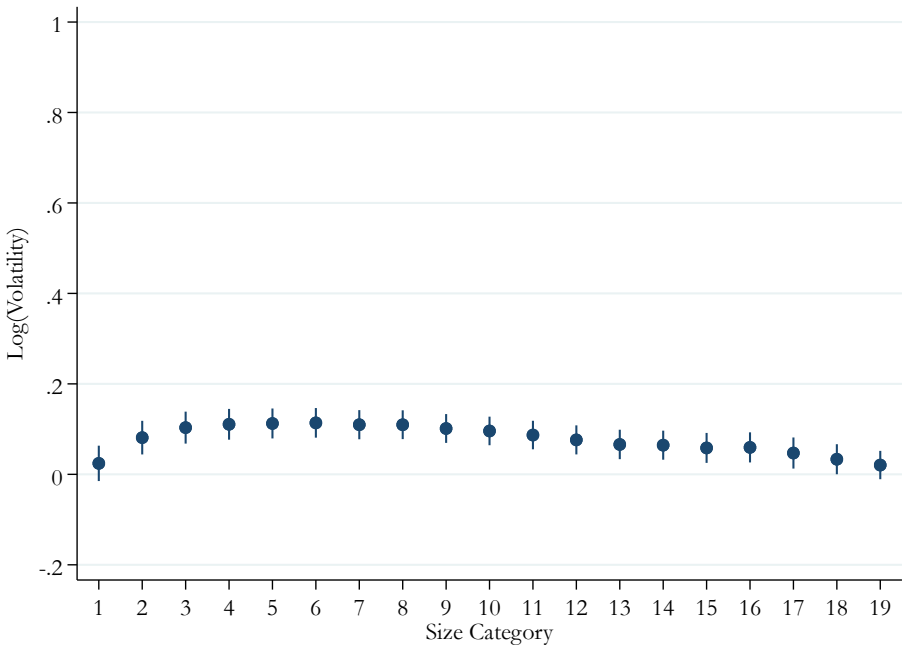


Figure 2
Volatility Around Stock Splits

In this figure we explore the pattern of volatility around 2-for-1 stock splits or greater (e.g., 3-for-1, 4-for-1, etc.). We plot the coefficients for each month in event time, relative to the omitted category of 12 months prior to the split. The regressions contain stock and month fixed effects. The dots represent the point estimates and the lines represent 95% confidence intervals. Standard errors are clustered by stock and month.

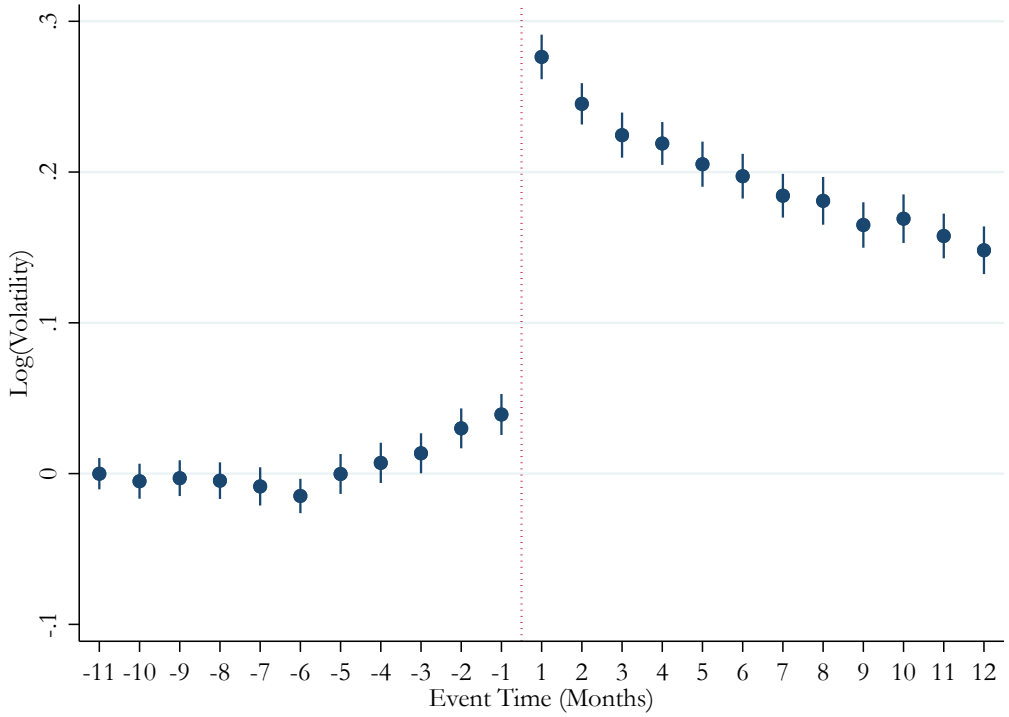


Figure 3
Volatility Around Reverse Stock Splits

In this figure we explore the pattern of volatility around 1-for-2 stock splits or greater (e.g., 1-for-3, 1-for-4, etc.). We plot the coefficients for each month in event time, relative to the omitted category of 12 months prior to the split. The regressions contain stock and month fixed effects. The dots represent the point estimates and the lines represent 95% confidence intervals. Standard errors are clustered by stock and month.

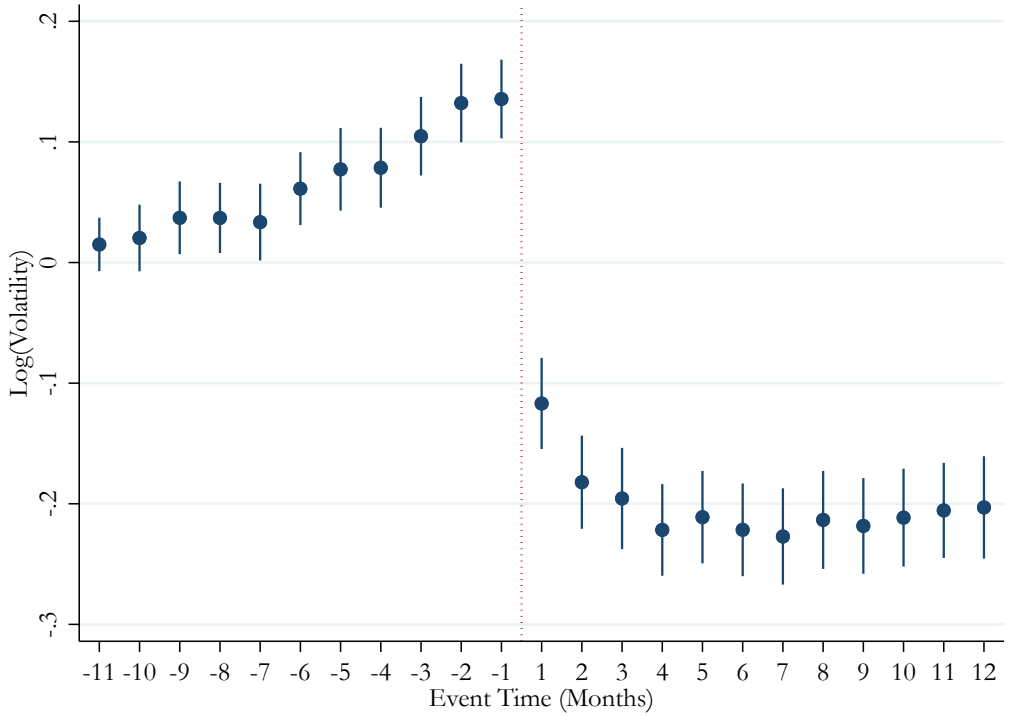


Figure 4

Regression Discontinuity: Scaled Intraday Price Range Around Stock Splits

In this figure we explore the pattern of volatility around 2-for-1 stock splits or greater (e.g., 3-for-1, 4-for-1, etc.). We examine 90 days before and after the split. The outcome, scaled intraday price range, is defined as difference between the intraday high and intraday low, normalized by the lagged closing price. The thick lines represent non-parametric estimates of the mean on a given day, estimated using local linear regression with a triangular kernal and MSE-optimal bandwidth. The thin lines represent 95% confidence intervals. The dot show raw means for each event day.

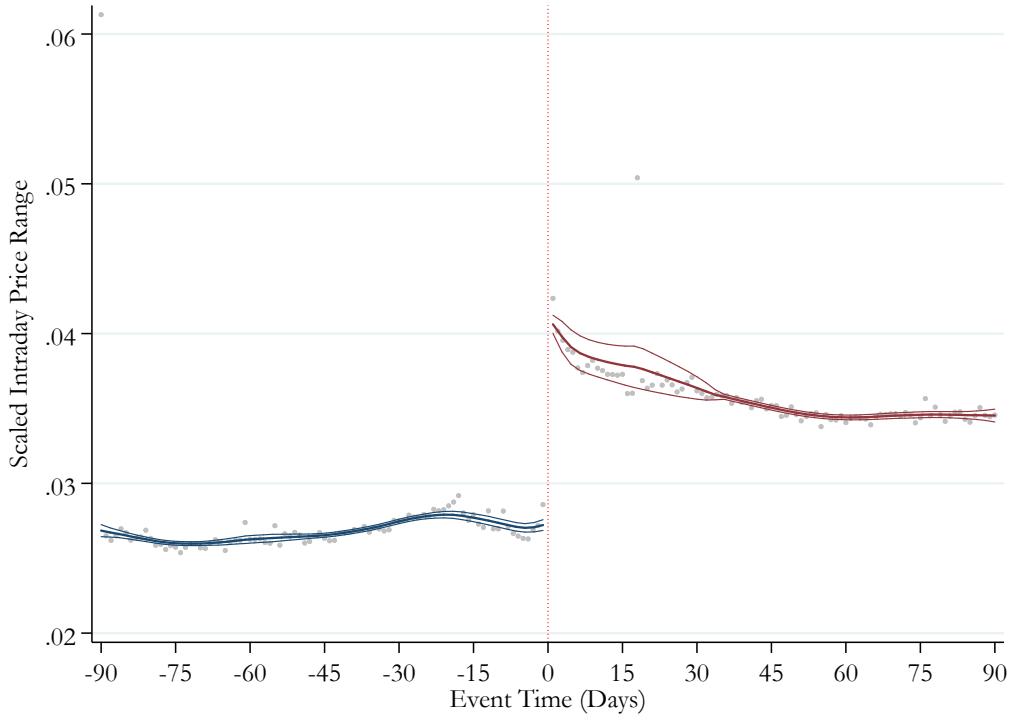


Figure 5
Shape of Volatility-Price Relation
This figure explores the shape of the volatility-price relation by binning lagged prices into vigintiles and repeating the regression from Panel A of Table 1, column 4, replacing the continuous *Log(Lagged Price)* variable with these vigintile dummies. The resulting coefficients are plotted with 95% confidence intervals. Category 20 is omitted.

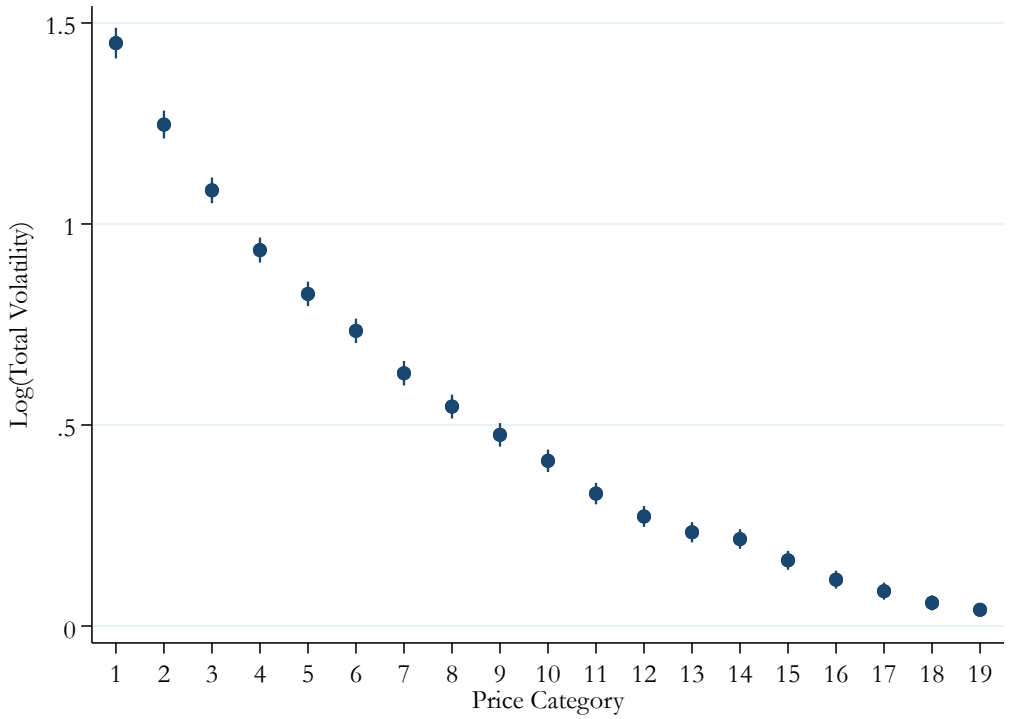


Table 1
Baseline Results: Total Volatility
This Table explores how return volatility varies with the share price. Using data at the stock-month level, we estimate the following regression:

$$log\left(vol_{it}\right) = \beta_0 + \beta_1 log\left(price_{i,t-1}\right) + Size_{i,t-1} + \tau_t + \nu_i + \epsilon_{it}.$$

We regress each stock i's volatility in month t on the stock's nominal share price at the end of the previous month, indicator variables for 20 size categories using the stock's market capitalization at the end of the previous month relative to other stocks in the same time period, calendar-year month fixed effects, and stock fixed effects. Volatility is estimated using daily returns from month t. The sample exclude observations with extreme lagged price (the bottom and top 1%). To account for correlated observations, we double-cluster standard errors by stock and year-month. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Panel A: Between Stocks

	Log(Total Volatility)			
	(1)	(2)	(3)	(4)
Log(Lagged Price)	-0.326***		-0.332***	-0.339***
	(0.00339)		(0.00446)	(0.00405)
Log(Lagged Market Cap)		-0.146***	0.00431	
		(0.00235)	(0.00311)	
Month FE	Yes	Yes	Yes	Yes
ME Breakpoint FE	No	No	No	Yes

Table 2
Baseline Results: Idiosyncratic Volatility

This Table repeats the analysis of Table 1, using idiosyncratic volatility as the outcome, where idiosyncratic volatility for a given month is estimated as the volatility of CAPM residuals that month using daily returns.

Panel A: Between Stocks

	Log(Idiosyncratic Volatility)			
	(1)	(2)	(3)	(4)
Log(Lagged Price)	-0.357***		-0.330***	-0.345***
	(0.00317)		(0.00432)	(0.00397)
Log(Lagged Market Cap)		-0.171***	-0.0211***	
		(0.00216)	(0.00308)	
Month FE	Yes	Yes	Yes	Yes
ME Breakpoint FE	No	No	No	Yes
R-squared	0.469	0.363	0.471	0.474
Observations	3,254,548	3,254,302	3,254,302	3,254,200

Panel B: Within Stocks

	Log(Idiosyncratic Volatility)			
	(1)	(2)	(3)	(4)
Log(Lagged Price)	-0.285***		-0.262***	-0.286***
	(0.00375)		(0.00461)	(0.00391)
Log(Lagged Market Cap)		-0.187***	-0.0258***	
		(0.00324)	(0.00377)	
Security FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
ME Breakpoint FE	No	No	No	Yes
R-squared	0.603	0.582	0.603	0.603
Observations	3,254,436	3,254,200	3,254,200	3,254,088

Table 3
Baseline Results: Market Beta

This Table repeats the analysis of Table 1, using Beta as the outcome, where Beta from a given month is estimated using daily returns.

Panel A: Between Stocks

	Log(|Beta|)			
	(1)	(2)	(3)	(4)
Log(Lagged Price)	-0.109***		-0.334***	-0.305***
	(0.00565)		(0.00642)	(0.00538)
Log(Lagged Market Cap)		0.0232***	0.174***	
		(0.00383)	(0.00468)	
Month FE	Yes	Yes	Yes	Yes
ME Breakpoint FE	No	No	No	Yes
R-squared	0.056	0.047	0.085	0.085
Observations	3,254,548	3,254,302	3,254,302	3,254,200

Panel B: Within Stocks

	Log(|Beta|)			
	(1)	(2)	(3)	(4)
Log(Lagged Price)	-0.0881***		-0.236***	-0.192***
	(0.00619)		(0.00698)	(0.00594)
Log(Lagged Market Cap)		0.0182***	0.164***	
		(0.00508)	(0.00584)	
Security FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
ME Breakpoint FE	No	No	No	Yes
R-squared	0.185	0.183	0.190	0.190
Observations	3,254,436	3,254,200	3,254,200	3,254,088

Table 4
Volatility Around Stock Splits

To explore how volatility changes after splits, we estimate the following regression:

$$log\left(vol_{it}\right) = \beta_0 + \beta_1 Post_{it} + \tau_t + \nu_i + \epsilon_{it}.$$

The regression sample uses observations at the stock-month level. We consider any positive stock split in which one old share is converted into two or more new shares. The month of a stock split counts as event date 0, and the sample is restricted to observations in the window $[t-12, t+12]$ around the event date, conditional on the observation not being in the windows of another stock split for the same stock. To examine how volatility varies with event time in greater detail, we also present results in which the $Post_{it}$ indicator is replaced with event-month indicators. In all specifications, we control for year-month and stock fixed effects and double cluster standard errors by stock and time. To account for correlated observations, we double-cluster standard errors by stock and year-month. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

	Log(Total Volatility)		Log(Idiosyncratic Volatility)		Log(|Beta|)	
	(1)	(2)	(3)	(4)	(5)	(6)
Post Split	0.197***		0.194***		0.220***	
	(0.00455)		(0.00464)		(0.00738)	
Post Split (0 Month)		0.234***		0.232***		0.255***
		(0.00575)		(0.00585)		(0.0133)
Post Split (1 Month)		0.271***		0.264***		0.318***
		(0.00598)		(0.00619)		(0.0128)
Post Split (2 Month)		0.240***		0.236***		0.278***
		(0.00561)		(0.00581)		(0.0126)
Post Split (3 Month)		0.220***		0.215***		0.255***
		(0.00604)		(0.00620)		(0.0129)
Post Split (4 Month)		0.214***		0.210***		0.237***
		(0.00565)		(0.00582)		(0.0123)
Post Split (5 Month)		0.200***		0.198***		0.217***
		(0.00613)		(0.00621)		(0.0140)
Post Split (6 Month)		0.192***		0.190***		0.205***
		(0.00604)		(0.00608)		(0.0141)
Post Split (7 Month)		0.179***		0.175***		0.214***
		(0.00621)		(0.00620)		(0.0146)
Post Split (8 Month)		0.176***		0.172***		0.189***
		(0.00632)		(0.00639)		(0.0147)
Post Split (9 Month)		0.160***		0.155***		0.184***
		(0.00626)		(0.00642)		(0.0139)
Post Split (10 Month)		0.164***		0.163***		0.175***
		(0.00636)		(0.00633)		(0.0144)
Post Split (11 Month)		0.152***		0.150***		0.164***
		(0.00594)		(0.00597)		(0.0137)
Post Split (12 Month)		0.143***		0.142***		0.144***
		(0.00664)		(0.00677)		(0.0153)
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Stock FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.609	0.611	0.598	0.600	0.258	0.259
Observations	167,938	167,938	167,938	167,938	167,938	167,938

Table 5

Regression Discontinuity: Scaled Intraday Price Range Around Stock Splits

In this figure we explore the pattern of volatility around 2-for-1 stock splits or greater (e.g., 3-for-1, 4-for-1, etc.). We examine 90 days before and after the split. The outcome, scaled intraday price range, is defined as difference between the intraday high and intraday low, normalized by the lagged closing price. Control functions on each side of the cutoff are estimated non-parametrically using local linear regression. Bandwidths are selected using one common MSE-optimal bandwidth selector, two different MSE-optimal bandwidth selectors on each side of the cutoff, or one common MSE-optimal bandwidth selector for the sum of regression estimates (as opposed to the difference thereof). The kernel is either Triangular or Epanechnikov as labelled. The estimated coefficient represents the size of the discontinuity at the split date, as illustrated in Figure 4.

	(Intraday High - Intraday Low) / Lagged Close					
	(1)	(2)	(3)	(4)	(5)	(6)
Discontinuity at Split	0.0138***	0.0139***	0.0139***	0.0137***	0.0138***	0.0139***
	(0.000600)	(0.000636)	(0.000673)	(0.000512)	(0.000607)	(0.000661)
Degree Local Poly	1	1	1	1	1	1
Bandwidth Left	15.39	14.62	13.47	16.49	12.82	12.32
Bandwidth Right	15.39	14.38	13.47	16.49	13.74	12.32
Kernel	Triangular	Triangular	Triangular	Epanechnikov	Epanechnikov	Epanechnikov
Observations	1,268,679	1,268,679	1,268,679	1,268,679	1,268,679	1,268,679

Table 6
Volatility and Past Returns

This table shows regressions of the form:

$$log\left(vol_{it}\right) = \beta_0 + \beta_1 returns_{i,[t-x,t-1]} + \tau_t + \epsilon_{it},$$

where $returns_{i,[t-x,t-1]}$ represents past returns over the recent months. Volatility is estimated using daily returns from month t. The sample excludes observations with extreme lagged priced (the bottom and top 1%). To account for correlated observations, we double-cluster standard errors by stock and year-month. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Panel A: Between Stocks

	Log(Total Volatility)					
	(1)	(2)	(3)	(4)	(5)	(6)
	2-Month	4-Month	6-Month	8-Month	10-Month	12-Month
Lagged Return	-0.125***	-0.103***	-0.0895***	-0.0784***	-0.0663***	-0.0557***
	(0.0249)	(0.0202)	(0.0158)	(0.0126)	(0.0106)	(0.00899)
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.179	0.181	0.183	0.184	0.185	0.186
Observations	3,227,112	3,173,006	3,119,908	3,067,813	3,016,552	2,966,212

Panel B: Within Stocks

	Log(Total Volatility)					
	(1)	(2)	(3)	(4)	(5)	(6)
	2-Month	4-Month	6-Month	8-Month	10-Month	12-Month
Lagged Return	-0.106***	-0.0845***	-0.0711***	-0.0613***	-0.0519***	-0.0438***
	(0.0115)	(0.00975)	(0.00771)	(0.00602)	(0.00500)	(0.00423)
Security FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.539	0.542	0.544	0.545	0.547	0.548
Observations	3,227,012	3,172,885	3,119,769	3,067,672	3,016,348	2,966,021

Table 7
Response to Nominal Earnings Surprises

Panel A of this table shows results of estimating equations of the form:

$$CAR_{i,[t-1,t+1]} = \beta_0 + \beta_1 nominal\, surprise_{it} + \beta_2 scaled\, surprise_{it} + \epsilon_{it}$$

$CAR_{i,[t-1,t+1]}$ is the firm's cumulative abnormal return over the three-day window around the earnings announcement. *Nominal Surprise* is defined as the difference between announced earnings per share and analyst consensus. *Scaled Surprise* is defined as *Nominal Surprise* divided by the lagged stock price (from 3 trading days prior to announcement). Analyst concensus is defined as the median analyst earnings forecast among analysts that make a forecast within 30 days of the announcement. We measure both the the nominal surprise and the scaled surprise as percentile rankings in Columns 2-6. Standard errors are clustered by date. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Panel A: Immediate Impact

	Cummulative Abnormal Return [-1,1]					
	(1) All	(2) All	(3) Large Cap	(4) Small Cap	(5) Year $\geq$ 2001	(6) Year $<$ 2001
Surprise Scaled	26.91***					
	(3.369)					
Surprise Nominal	7.734***					
	(0.233)					
Percentile Surprise Scaled		0.0346***	0.0295***	0.0187***	0.0400***	0.0277***
		(0.00203)	(0.00349)	(0.00252)	(0.00301)	(0.00266)
Percentile Surprise Nominal		0.0347***	0.0281***	0.0624***	0.0456***	0.0228***
		(0.00190)	(0.00288)	(0.00274)	(0.00278)	(0.00255)
R-squared	0.034	0.069	0.060	0.077	0.093	0.043
Observations	180,740	180,740	90,592	90,148	98,668	82,072

Table 7
(Continued)
Panel B: Long Term Reversal

	Cummulative Abnormal Return				
	(1) [-1,1]	(2) [-1,25]	(3) [-1,50]	(4) [-1,75]	(5) [-1,100]
Percentile Surprise Scaled	0.0346***	0.0557***	0.0714***	0.0794***	0.0875***
	(0.00203)	(0.00466)	(0.00645)	(0.00747)	(0.00852)
Percentile Surprise Nominal	0.0347***	0.0279***	0.0237***	0.0208***	0.0126
	(0.00190)	(0.00436)	(0.00606)	(0.00701)	(0.00798)
R-squared	0.069	0.028	0.019	0.015	0.011
Observations	180,740	180,153	178,120	177,196	176,155