The long-run behavior of firms' stock returns: Evidence and interpretations

Gregory R. Duffee[*]

Haas School of Business
University of California – Berkeley

First version: June 2002
This Draft: August 28, 2002

ABSTRACT

The distribution of long-horizon idiosyncratic returns to individual stocks is strongly asymmetric, in contrast to both the distribution of shorter-horizon returns and the intuition of the Central Limit Theorem. The lower tail of the distribution is much fatter than the upper tail. One reason for this asymmetry is a modest, extremely persistent negative relation between a stock's return and its future idiosyncratic return volatility. The long-run relation between returns and volatility is driven by two balance-sheet effects that work in opposite directions. One is the leverage effect: a higher asset/equity ratio corresponds to both greater idiosyncratic volatility and greater exposure to common risks. The other appears to be an asset-mix effect. A higher book/market ratio corresponds to both *less* idiosyncratic volatility and *less* exposure to common risks. This result is consistent with the idea that book/market is correlated with the proportion of a firm's assets that are lower-risk assets-in-place instead of higher-risk intangible assets such as growth options. The leverage effect induces a negative relation between returns and volatility, while the asset-mix effect induces a positive relation.

[*]Voice 510-642-1435, fax 510-643-1420, email duffee@haas.berkeley.edu. Address correspondence to 545 Student Services Building #1900, Berkeley, CA 94720. I thank my colleagues at Berkeley for helpful comments. The most recent version of this paper is at http://faculty.haas.berkeley.edu/duffee/.

1 Introduction

Extensive research documents the properties of daily and monthly returns to individual stocks. Thus we know, for example, whether the distributions of these returns are typically symmetric and whether the tails of the densities look Gaussian.[1] Researchers have also studied key features of the short-run dynamics of these returns, such as the persistence in volatility and the correlation between returns and future market betas.

However, we know little about the properties of returns to individual stocks over horizons of a year or more. Our ignorance is partially a consequence of limited data. It is hard to say much about the tails of the density of a particular firm's stock return over, say, a two-year horizon when no security in the Center for Research in Security Prices files has more than 40 observations of nonoverlapping two-year returns. This gap in our knowledge is unfortunate because dynamic asset pricing theories often have strong implications for distribution of long-horizon stock returns. For example, structural contingent-claim models for pricing credit-risky instruments (e.g., Merton (1974)) posit joint dynamics of firm value and debt that largely determine the distribution of long-horizon stock returns. Models of dynamic project choice by firms (e.g., Berk, Green, and Naik (1999)) also generate distributions of long-horizon stock returns, as well as implications for the relation between returns and future exposure to market risks. More generally, models that characterize long-run dynamics of short-horizon returns have implications for the behavior of long-horizon returns. To take the simplest case, the assumption that short-horizon log returns are stationary and independent implies (using the intuition of the Central Limit Theorem) that the distribution of long-horizon log returns is symmetric.

Although data limitations prevent us from estimating with precision the behavior of long-horizon returns to particular stocks or the market as a whole, in this paper I exploit the information in a broad panel of stock returns (1.6 million monthly observations) to study the behavior of long-horizon returns to a typical stock. These long-horizon returns consist of common and idiosyncratic components. Because the panel has no information about the behavior of common components that cannot be gleaned from studying portfolio returns, I focus on two particular issues here. First, what is the shape of the distribution of idiosyncratic multi-year returns? Second, what is the relation between yearly returns to stocks and the stocks' future aggregate risk exposures (betas)?

I find that the distribution of idiosyncratic five-year returns is wildly inconsistent with the intuition of the Central Limit Theorem. The lower tail of the distribution is much fatter than

[1]At the daily horizon returns to individual stocks are positively skewed. At the monthly horizon returns are close to symmetric and have fatter tails than the Gaussian distribution.

the upper tail. For example, the lower and upper one percentiles of five-year idiosyncratic log returns for a typical firm listed on the NYSE are about −350 and 240 percent, respectively. This asymmetry contrasts sharply with the nearly symmetric distribution of idiosyncratic monthly returns.

This asymmetric return distribution is, in part, a consequence of an asymmetric relation between returns and idiosyncratic return volatility. I find that when a stock's price falls, the volatility of idiosyncratic returns rises modestly for the next few years. This persistent negative relation conflicts with what we know about asymmetric volatility in other contexts. At the aggregate level, volatility rises sharply after a decline in stock prices but quickly returns to its earlier level. At the firm level, the relation between returns and volatility is positive in the very short run.

Asymmetric volatility is simply a statistical property; the important economic question is what drives it. I find that changes in a firm's capital structure are closely associated with changes in volatility. There are two competing balance-sheet effects at work. The first is the leverage effect. As others have documented, a decrease in a firm's stock price increases its debt-equity ratio, and changes in debt/equity are positively associated with changes in return volatility. New to this paper is evidence of a book-to-market effect. A decrease in a firm's stock price increases its book-to-market ratio, and changes in a firm's book/market are *negatively* associated with changes in return volatility. This latter relation is consistent with an asset-mix effect, in which assets-in-place (proxied by book equity) are less volatile than intangible assets such as real options.

These empirical links between idiosyncratic volatility and balance sheets have their natural counterparts in links between betas and balance sheets. An increase in a firm's leverage corresponds to an increase in its future stock return volatility owing to common factors, while an increase in its book-market ratio corresponds to a decrease in this common component of volatility. The net relation between a firm's stock return and its exposure to common risks is weak. This weak relation appears inconsistent with earlier evidence that "winners" subsequently have lower risk premia than do "losers." However, in the context of the three-factor of Fama and French (1993), there is no contradiction. A winner's exposure to common risks shifts from factors with high risk premia (small-cap and high book-to-market factors) to a factor with a low risk premium (a large-cap factor).

The outline of the remainder of this paper is as follows. Section 2 discusses how this paper fits in with the earlier literature. Section 3 explains how I construct idiosyncratic returns across different horizons and then discusses the empirical distributions of these returns. Section 4 investigates the relation between stock returns and future return volatility. Concluding comments are in the final section.

2 Related research

To date there is no direct evidence on the behavior of long-horizon returns to individual stocks. There is, however, a large literature on short-horizon returns–daily and monthly. If we had an accurate descriptions of the distribution and dynamics of short-horizon returns we could infer the behavior of long-horizon returns.

Since the 1960s, researchers have studied the probability distributions of short-horizon returns. Textbook treatments of empirical distributions include Fama (1976). He argued that distributions of monthly returns to firms' stocks are approximately Gaussian, while distributions of daily returns are skewed and highly leptokurtotic. Subsequent research has expanded our knowledge of these distributions, but not altered the way we interpret the contrast highlighted by Fama.[2] Campbell, Lo, and MacKinlay (1997) state the standard intuition: "Since all moments are finite, the Central Limit Theorem applies and long-horizon returns will tend to be closer to the normal distribution." (p. 19) In other words, tail events of short-horizon returns have only a modest impact on longer-horizon returns.

Taken to its extreme, this Central Limit Theorem (CLT) argument implies that log returns over multi-year horizons are Gaussian. But the CLT intuition applies to situations in which random variables are (sufficiently close to) independent. There is no question that short-horizon returns are not independent, although it is not clear whether the degree of dependence is enough to overturn the CLT. Aside from the conclusion that short-horizon returns exhibit strong volatility persistence, there is no generally-accepted description of their dynamics. An important stumbling block is that key features of these dynamics have signs that depend on the return horizon. Cho and Engle (1999) and Duffee (1995) find that at the daily horizon, return volatility rises after a firm's stock price rises. At the monthly horizon, Braun, Nelson, and Sunier (1995) and Duffee (1995) find no clear pattern between a firm's stock returns and future volatility. At the daily horizon, Cho and Engle (1999) find a negative relation between a firm's idiosyncratic stock return and its future market beta. The relation disappears at the monthly horizon (Braun et al. (1995)). The role of financial leverage in the short-run dynamics of firm-level stock returns has been extensively studied,[3] but the link, if any, between leverage and long-horizon returns is unknown.

Prior research into the behavior of long-horizon returns has focused on portfolio-level returns. There is evidence of time-variation in both means and covariances with common factors. Analyses of conditional means follow DeBondt and Thaler (1985) and Jegadeesh

[2]Much of this research considered whether stock returns are better described by a finite-variance or infinite-variance distribution. See Chapter 1 in Campbell et al. (1997) for references. Simkowitz and Beedles (1978, 1980) characterize more precisely the statistical description of firms' monthly stock returns.

[3]Recent studies include Bekaert and Wu (2000) and Figlewski and Wang (2000).

and Titman (1993), who document horizon-specific serial correlation patterns of returns. Attempts to understand these patterns emphasize time-varying covariances with common factors. Chan (1988) and Ball and Kothari (1989) find that stocks that were "winners" during a multiyear "ranking period" had higher betas in that ranking period than they did during the subsequent "performance period." Grundy and Martin (2001) refine these results by considering multifactor models and by conditioning on the factor returns. Related evidence is in Fama and French (1996). They find that risk premia, as measured by their three-factor model, are higher for stocks that have decreased in price over the past few years. In addition, there is a large literature on time-variation in betas that focuses on the variation of betas with aggregate variables.[4]

Option prices provide a little indirect evidence on the distribution of long-horizon stock returns. Lauterbach and Schultz (1990) find that prices of warrants (long-dated options) are more consistent with a constant elasticity of variance pricing model than with the Black and Scholes (1973) model. This result suggests that the distribution of long-horizon stock returns is asymmetric, with the lower tail fatter than the upper tail. In principle, prices of corporate debt should contain substantial information about the distribution of long-horizon stock returns. However, the empirical accuracy of existing structural models is too low to trust the models' implications for this distribution.[5]

In a nutshell, our knowledge of the behavior of long-horizon returns to individual stocks is fairly sparse. In the next section I begin to fill in the gaps by examining the distribution of long-horizon returns.

3 The distribution of long-horizon returns

3.1 Theory

To motivate the empirical analysis of long-horizon returns it is useful to look at a simple model. Assume that we can write the log excess return to stock i from the end of period $t-1$ to the end of period t as

$$r_i(t) = \alpha_{it} + \sum_{j=1}^{J} \beta_{jit} r_j(t) + \epsilon_i(t), \quad \mathrm{Var}_t[\epsilon_i(t)] = \sigma_{it}^2. \tag{1}$$

[4]See, for example, Jagannathan and Wang (1996) and its references.

[5]See Eom, Helwege, and Huang (2002) for an empirical analysis of a variety of structural models. They conclude "... all of the models have substantial [credit] spread prediction errors."

In (1), there are J common components to stock returns $r_j(t)$. Think of them as returns to mimicking portfolios for some underlying factors. Time-varying parameters are given time subscripts. A subscript of t indicates a variable known at the end of $t-1$. The time dimension of returns is indicated with parentheses; the return $r_i(t)$ is not known until the end of t. The log excess return to this stock over T periods is denoted $r_i^T(t)$ and is the sum of short-horizon returns from t to $t+T-1$:

$$r_i^T(t) \equiv \sum_{k=0}^{T-1} \left[\alpha_{i(t+k)} + \sum_{j=1}^{J} \beta_{ji(t+j)} r_j(t+k) \right] + \epsilon_i^T(t), \quad \epsilon_i^T(t) \equiv \sum_{k=0}^{T-1} \epsilon_i(t+k). \tag{2}$$

Because the use of log returns creates some complications later, it is useful to recall the motivation behind their use. The intuition of the CLT is based on the behavior of sums of random variables. Unlike log returns, long-horizon simple net returns are not the sum of short-horizon simple net returns. Thus there is no reason to believe that the density of longer-horizon simple returns should be symmetric. In fact, it cannot be; a stock price cannot fall below zero, but is unbounded above.

If we faced no data limitations we could determine the statistical distribution of long-horizon returns $r_i^T(t)$ by constructing a histogram from a large number of nonoverlapping draws of these returns. However, for T greater than a few months, the available time series are too short to construct a reliable histogram for $r_i^T(t)$ for a given stock i.

The approach I adopt is to construct a histogram across both stocks and time. This panel approach gives us additional information about the distribution of long-horizon stock-specific shocks but cannot tell us anything new about long-horizon market returns. Therefore I focus on the distribution of $\epsilon_i^T(t)$. To construct this distribution the common components must be removed from $r_i^T(t)$, which requires specifying the common factors and estimating each stock's loadings on the factors. In the empirical work that follows I use the three Fama-French factors as the common components. Since these factors are not the universe of common components to stocks, the use of the term "idiosyncratic" is a little strained, but it is easier to use than "the part of stock returns that is not captured by the Fama-French factors."

The focus on idiosyncratic returns instead of total returns is not an important limitation because idiosyncratic components dominate the volatility of individual stock returns. For example, the standard deviation of log returns to a typical stock on the NYSE is about 13 percent per month, while the standard deviation of the idiosyncratic component of this return (i.e., after removing Fama-French factor loadings) is about 11 percent.

There are two critical choices to make when constructing the empirical distribution of long-horizon returns. First, we must decide how to combine information about returns to stocks that differ in their return variances. Second, we must decide how to treat truncated returns—returns to stocks that do not have a complete history of returns over the T-period horizon. I address these two choices in order.

3.1.1 Standardizing returns

The idiosyncratic returns $\epsilon_i(t)$ have different variances across stocks i. If the returns are not adjusted for these differences, the resulting distributions of $\epsilon_i(t)$ and $\epsilon_i^T(t)$ are completely unconditional distributions: we do not condition on either i or t. These are well-defined distributions, but they are not particularly interesting given the goals of this paper. One objective is to see if the intuition of the CLT holds when we compare long-horizon return distributions to short-horizon distributions. The theorem breaks down when there is sufficient dependence between successive short-horizon returns. If we do not scale $\epsilon_i(t)$, there is an obvious reason for dependence: successive returns are generated by the distribution for a given stock i. Put differently, the distribution of long-horizon returns will have fat tails simply because there are some stocks with large return variances.

Cross-sectional differences in variances could be removed if we could scale $\epsilon_i(t)$ by its standard deviation conditioned on i. However, we do not observe these stock-specific standard deviations, and scaling by estimates of these standard deviations turns out to have highly undesirable consequences. (I discuss these consequences in Section 3.2.) In addition, scaling by stock i's unconditional standard deviation (more precisely, the standard deviation conditioned on i but not on t) does not adjust for variations over time in the volatility of stock i's return. Although these variations are much smaller than cross-sectional variations in volatility, they have the same qualitative effect on the distribution of long-horizon returns: the distribution will have fat tails because there are some time periods in which return variances are large.

If we could observe the conditional standard deviations σ_{it}, we could construct standardized returns

$$z_i(t) = \epsilon_i(t)/\sigma_{it}, \quad z_i^T(t) = \epsilon_i^T(t)/\sigma_{it}.$$

I adopt this method (using an estimate of σ_{it}), but it is an imperfect solution when there is mean reversion in σ_{it}^2. Although $\text{Var}_t[z_i(t)] = 1$ for all i and t, $\text{Var}_t[z_i^T(t)]$ depends on whether σ_{it} is above or below its mean (conditioned on i). If, say, the volatility of stock i's return is temporarily low, we expect it to rise in the future and therefore the conditional variance $\text{Var}_t[\epsilon_i(\tau)/\sigma_{it}], \tau > t$, is greater than one. Thus the variance of long-horizon scaled returns will tend to be high when σ_{it} is low. The resulting distribution of long-horizon returns will still be a mixture of high-variance returns and low-variance returns, although the mixture will not be as extreme as it would be if we scaled by unconditional standard deviations.

3.1.2 The treatment of truncated returns

The second critical choice is how to treat long-horizon returns that are truncated. The most common reasons for truncation are merger, delisting, or liquidation. There is no obvious way to consistently treat truncated and non-truncated returns. One method is to treat them identically. In other words, we can construct the distribution of idiosyncratic returns over the holding period from t until either the end of the return horizon T or until the stock price is unavailable:

$$\epsilon_i^T(t) \equiv \sum_{j=0}^{\tau-1} \epsilon_i(t+j), \; \tau = \min(T, \text{index of first missing obs in}[t+1, t+T]). \tag{3}$$

This corresponds to the distribution of returns to an investment strategy that borrows at t to buy a stock, hedge out the market return component, and hold the position until the end of the return horizon or until the stock disappears. (It is essential to include the stock's delisting return in this strategy, which leads to a set of problems that are discussed below.)

However, I am more interested in understanding the dynamics of equity valuation than in returns to investor strategies. When a firm's stock is no longer publicly traded, the equity value does not necessarily disappear. Instead, it is unobserved because it trades privately or is incorporated in the equity of a different firm. In such situations the distribution of (3) is not the distribution of long-horizon returns to equity. To illustrate this point, consider the two-period return to a stock. Each period's return is normal with variance one. However, at the end of period one, the firm's stock may stop publicly trading because the firm is merged into another firm. This event occurs with probability one-half. If mergers are independent of the period one returns, the distribution of two-period returns using (3) is nonnormal (it is a mixture of normals). But in another sense, the distribution of two-period returns is normal with variance two, and observed two-period returns are censored. We can recover the latter distribution of two-period returns by scaling up the truncated returns by the square root of two.

More generally, we could construct long-horizon returns by scaling (3) to account for the shorter horizons of the truncated returns:

$$\epsilon_i^T(t) \equiv \sqrt{T/\tau} \; \sum_{j=0}^{\tau-1} \epsilon_i(t+j), \; \tau = \min(T, \text{index of first missing obs in}[t+1, t+T]). \tag{4}$$

Scaling by $\sqrt{T/\tau}$ implies that that the unobserved portion of the truncated long-horizon return is drawn from the same distribution as the observed portion of the truncated return. Therefore this scaling is appropriate if we assume that the events that lead to truncation are independent of past returns. Of course, this assumption is false. Stocks often disappear precisely because of their history of returns; bankruptcy is an obvious example. As we will see in the empirical work later in this paper, truncated long-horizon returns have higher variances and are more negatively skewed than non-truncated long-horizon returns. Therefore scaling by $\sqrt{T/\tau}$ exaggerates the negative tail of long-horizon returns.

Yet not scaling truncated returns underestimates the variance of long-horizon returns because it effectively assumes that idiosyncratic returns after truncation are identically zero. In addition, not scaling underestimates the size of the negative tail of long-horizon returns. Since returns tend to be truncated after stock prices fall to some barrier, observations of stock prices below the barrier are censored. In short, (3) underestimates and (4) overestimates the variance and lower tail of long-horizon returns. We cannot say more without putting some structure on the joint dynamics of returns and the event of truncation. Such modeling is beyond the scope of the present paper, thus I report information about both distributions in the paper.

The main points of Sections 3.1.1 and 3.1.2 are (1) returns should be scaled to adjust for differences in conditional variances across stocks and time, and (2) any method used to scale long-horizon returns based on the length of time over which they are observed is problematic. In the next subsection I discuss the mechanics of constructing these returns.

3.2 The construction of returns

3.2.1 Estimation of factor loadings

I use monthly stock returns from CRSP. If the stock has been delisted I augment the CRSP return series with a final month's return equal to the delisting return, if available. I then form returns in excess of the riskfree rate by subtracting from the log monthly return the continuously-compounded yield on a three-month Treasury bill as of the end of month $t-1$. A few of the delisting prices on CRSP are zero.[6] (Of course this is not a transaction price.) To handle these prices, I adopt an ad hoc rule that replaces a minus infinity log monthly return with a -400 percent log monthly return. This has little effect on the results emphasized here because there were too few zero prices in my sample to affect the tail behavior of long-horizon stock returns at the percentiles that I examine (0.05, 0.01, and 0.005).

[6]Delisting returns on CRSP are subject to much more error than returns based on exchange transactions. See Shumway (1997) and Shumway and Warther (1999) for details on problems associated with CRSP's delisting returns.

The next step is to remove common components from firm i's stock monthly return. I assume that these common components are captured by the three Fama-French factors.[7] I use rolling regressions to estimate each stock's one-month-ahead factor loadings. (The use of rolling regressions is motivated later.) The regression is

$$r_i(t) = b_{0it} + b_{1it} R_M(t) + b_{2it} SMB(t) + b_{3it} HML(t) + e_i(t), \quad (5)$$

where $r_i(t)$ is the excess log return for stock i, $R_M(t)$ is the Fama-French market return deflated by the riskfree rate, $SMB(t)$ is the small-stock factor, and $HML(t)$ is the book-to-market factor. The regression used to construct the month-t residual is estimated on data from month $t-60$ to month $t-1$. The regression is not estimated if $r_i(t)$ is missing or if fewer than 55 observations are available to estimate this regression. A convenient byproduct of the use of rolling regressions is that the results are unaffected by the treatment of delisting stock prices that are reported as zero. Since the regression is estimated for stock i and month t only when a valid return exists for month t, we know the stock price did not fall to zero prior to month t.

3.2.2 Estimation of mean returns

Given these rolling regressions, perhaps the most obvious way to construct idiosyncratic returns is to form one-month-ahead residuals with $\hat{e}_i(t) = r_i(t) - \hat{b}_{0it} - \hat{b}_{1it} R_M(t) - \hat{b}_{2it} SMB(t) - b_{3it} HML(t)$. However, the noise in $\hat{b}_{0it}$ accumulates over the T months in the long-horizon return and thus adds substantial noise to this long-horizon return.[8] My approach is to assume that exposure to the common factors are the only priced risks, so that theory can be used to determine the value of the constant. Because we are working with log returns, this value is not zero.

Under the assumption of joint log-normality of short-horizon returns and the stochastic discount factor, standard arguments imply that

$$E_t(r_i(t+1)) = -\frac{1}{2}\mathrm{Var}_t(r_i(t+1)) - \mathrm{Cov}_t(m(t+1), r_i(t+1)). \quad (6)$$

where $m(t+1)$ is the log of the stochastic discount factor. The Jensen's inequality term on the right side of (6) plays a nontrivial role in the results of this paper.[9] It implies that if

[7] The factors were introduced in Fama and French (1993). Thanks to Ken French for making these data available.

[8] In practice, the resulting distribution of long-horizon idiosyncratic stock returns has a very large standard deviation relative to the distribution of monthly idiosyncratic stock returns.

[9] The precise form of the Jensen's inequality term in (6) relies on log-normality. As we will see, the empirical distribution of monthly log returns is reasonably close to normal, thus its use is appropriate here.

two stocks have identical covariance risk but differ in their idiosyncratic volatility, the stock with the larger idiosyncratic volatility will have a lower expected log return—investors get more of their expected return in the form of volatility. If we ignore this fact and assign the two stocks equal expected returns, we will introduce a spurious source of negative skewness in stock returns. High-volatility stocks will tend to have negative returns, adding weight to the lower tail of the distribution.

At short horizons, correcting for Jensen's inequality has only a small effect on the distribution of returns. A typical standard deviation of monthly returns is around 13 percent. This corresponds to a Jensen's inequality term of 85 basis points per month, which is small relative to the monthly standard deviation. At long horizons, the correction is important because the Jensen's inequality term grows linearly with the horizon while the standard deviation grows with the square root of the horizon. A standard deviation of 13 percent per month corresponds to a Jensen's inequality term of 51 percent over a five-year horizon, which is over one-half of the corresponding standard deviation of 101 percent.

The difficulty of correcting for Jensen's inequality is that we do not observe the conditional variances of returns. I use sample variances from a rolling 60-month window to proxy for these unobserved conditional variances. Therefore the idiosyncratic return to stock i in month t is constructed as

$$\hat{e}_i(t) = r_i(t) + c + \frac{1}{2}\widehat{\mathrm{Var}_t(r_i}(t)) - \hat{b}_{1it}R_M(t) - \hat{b}_{2it}SMB(t) - \hat{b}_{3it}HML(t). \tag{7}$$

(No Jensen's inequality terms are needed for the Fama-French factor returns because they are not log returns.) The residuals $\hat{e}_i(t)$ are the empirical counterpart to the idiosyncratic shocks $\epsilon_i(t)$ in (1). The constant term c has no effect on the shape of the distribution of idiosyncratic returns. It is constant across stocks and time, and is chosen to set the grand mean idiosyncratic return (across stocks and time) equal to zero.[10] Its only role is to simplify the interpretation of the results.

The use in (7) of rolling sample covariances instead of true (unobserved) conditional covariances biases the shape of the distribution of long-horizon returns. The sign of the bias depends on whether the rolling covariances overreact or underreact to absolute return shocks. The extreme case of overreaction is when covariances is constant over time, so that variations in sample covariances are entirely noise. Then the bias works to reduce the lower tail of the return distribution relative to the upper tail. For example, if $\hat{e}_i(t)$ is large in absolute value, sample conditional variances for the next 60 months will tend to be high. The Jensen's inequality correction will bias up subsequent monthly returns. Therefore long-

[10]Depending on the sample, this constant ranges from about 1 to −7 basis points per month.

horizon returns that are large in absolute value will tend to have positively biased returns.

At the other extreme, a large absolute $\hat{e}_i(t)$ may correspond to a correspondingly large $\text{Var}_t(e_i(t{+}1))$. The rolling sample covariances adjust sluggishly to this large absolute idiosyncratic return, hence subsequent monthly returns will be biased down. Therefore long-horizon returns that are large in absolute value will tend to have negatively biased returns. The net effect of these offsetting biases is unknown.

3.2.3 Scaling

The construction of idiosyncratic returns is completed by scaling the residuals. Monthly return residuals are scaled by the standard error of the estimate from the rolling regression (5):

$$\hat{z}_i(t) = \frac{\hat{e}_i(t)}{\widehat{SE}_{it}}.$$

Idiosyncratic returns over T months are denoted $\hat{e}_i^T(t)$ and are the sum of the T residuals $\hat{e}_i(t+j), j = t, T-1$. If any of the returns $r_i(t+1), \ldots, r_i(t+T-1)$ are missing (say, the stock disappears from the CRSP tape after month $t+2$), the sum is truncated at month τ immediately preceding the first missing value. Standardized long-horizon returns, denoted $\hat{z}_i^T(t)$, are transformations of $\hat{e}_i^T(t)$. The transformation depends on whether I scale truncated returns.

$$\text{no scaling of truncated returns:} \qquad \hat{z}_i^T(t) = \frac{\hat{e}_i^T(t)}{\sqrt{T}\widehat{SE}_{it}}, \tag{8}$$

$$\text{scaling of truncated returns:} \quad \hat{z}_i^T(t) = \sqrt{T/\tau}\frac{\hat{e}_i^T(t)}{\sqrt{T}\widehat{SE}_{it}}. \tag{9}$$

Divison by $\sqrt{T}$ in (8) and (9) facilitates comparison of the distribution of long-horizon returns to the distribution of short-horizon returns.

3.2.4 The motivation for rolling regressions

This rolling regression procedure may seem overelaborate. A much simpler approach is to estimate (5) over the entire sample of a stock's monthly returns, then standardize the in-sample residuals by the standard error of the estimate. There are two major problems with this approach. First, factor loadings vary through time. Second, and more important for the purposes of this paper, this methodology biases long-horizon idiosyncratic returns to zero. By construction, in-sample residuals sum to zero, thus long-horizon idiosyncratic returns constructed with these residuals will be close to zero. An alternative is to estimate (5) over the entire sample but without the constant term. The residuals can be standardized by the standard error of the estimate, and T-period returns can be constructed by summing the monthly standardized residuals. This may appear to be similar to scaling by the unconditional standard deviation of stock i's return, which is one of the potential scaling methods suggested in Section 3.1. However, this scaling method is also flawed. I illustrate this point with a hypothetical example.

Consider two stocks that, as of time t, have the same price and identical conditional return distributions that are independent of each other. Then the returns realized on these two stocks over, say, the next five years are independent draws from the same distribution of five-year returns. After time t, the two stock prices will diverge. Because the volatility of returns is inversely related to past returns, the stock with the lower price will have higher realized volatility along its sample path. As we will see, this inverse relation between returns and volatility is the reason why the density of long-horizon returns that has a fat lower tail relative to its upper tail. But if we standardize the realized monthly returns over this five-year period by the realized volatility in this same period, we will erase the cross-sectional evidence of fat lower tails. Put differently, this standardization improperly implies that the two stock returns did not have identical conditional return distributions as of time t. The moral of this example is that returns need to be standardized by a volatility measure that is not contaminated by future returns.

3.3 Sample selection

Stock return data are from the 2001 version of the CRSP monthly NYSE/Amex/Nasdaq file. The analysis is restricted to common stocks of domestic firms. (These are securities with CRSP sharecodes of 10 or 11 over their entire sample.) The first date in the sample is July 1927, which is the first month for which the Fama-French factors are availabe. The number of listed firms has grown dramatically over the sample period, especially since the incorporation in the early 1970s of thousands of small-capitalization stocks traded on Nasdaq. It is possible that the distribution of long-horizon returns of small-cap firms differs from the corresponding distribution for high-cap firms. To allow for this possibility, I report results for two samples of stocks. The complete sample includes all stocks. The restricted sample drops stocks with market capitalization below the 20th percentile of NYSE-listed stocks.

I produce empirical distributions of both monthly and five-year returns ($T = 60$ months). The complete sample of monthly returns is the set of all non-missing $\hat{z}_i(t)$. The restricted sample is the set of non-missing $\hat{z}_i(t)$ for which the market capitalization of stock i at the end of month $t-1$ is at least as large as the 20th percentile of market capitalizations of NYSE-listed stocks as of that month-end. The sample of five-year returns is the set of non-overlapping returns, and is constructed as follows.

Let month t_0 be the first month for which stock i has a non-missing observation of $\hat{z}_i(t)$. (For the restricted sample, this month must also satisfy the requirement that at month-end $t_0 - 1$, the market capitalization of the stock was not below the 20th percentile of market capitalizations of NYSE-listed stocks for that month.) Then $\hat{z}_i^T(t_0)$ is included in the distribution of five-year returns. The next potential non-overlapping draw of a five-year return for stock i is $\hat{z}_i^T(t_0 + 60)$. If $z_i(t_0 + 60)$ is nonmissing (and, for the restricted sample, if the market cap at month-end $t_0 - 59$ is not below the 20th percentile), then $\hat{z}_i^T(t_0 + 60)$ is also included in the empirical distribution. This procedure is followed until stock i exits the CRSP tape or until January 1997, which is the last month for which a five-year return can be calculated.

3.4 The empirical distributions

3.4.1 Monthly idiosyncratic returns

I first examine the distribution of monthly idiosyncratic stock returns. Table 1 reports summary statistics and Fig. 1 displays the empirical densities. The main conclusions to draw from this information is that monthly returns are nonnormal, but the tails of the densities are close to symmetric. Before I get to the evidence supporting these conclusions it is helpful to discuss certain features of the data in more detail.

The sample of idiosyncratic returns to all stocks has more than 1.6 million observations across nearly 13,000 stocks. The sample of returns to large-cap stocks has about half the number of observations and stocks. The table reports that the mean monthly standard deviation of idiosyncratic returns, as estimated by the rolling regressions, is 13 percent for the entire sample and 11 percent for the large-cap sample. Sample moments for demeaned standardized returns are also reported. The term ‘demeaned’ requires some explanation. By construction, the mean monthly idiosyncratic return is zero (this is the role of c in (7)). Therefore the mean standardized return depends on the relation between the conditional standard deviation $\widehat{SE_{it}}$ and the idiosyncratic return $\hat{e}_i(t)$. They are slightly negatively correlated in these samples, resulting in mean standardized returns of 0.011 and 0.002, respectively. To simplify interpretation of the percentiles of the distributions I remove these means from the standardized returns.

The standard deviations of standardized returns are a little larger than one, thus $\widehat{SE_{it}}$ underestimates the standard deviation of $\hat{e}_i(t)$. There are two potential explanations for this. First, the conditional standard deviation $\widehat{SE_{it}}$, which is based on returns from $t-60$ to $t-1$, is estimated only for survivors; stocks that exist at time t. Survivorship probably imparts a negative bias to the conditional standard deviation. (It also imparts an upward bias to the mean return, but any such bias is removed when returns are demeaned.) Second, there has been a steady increase over time in the variance of idiosyncratic returns, as documented in Campbell, Lettau, Malkiel, and Xu (2001). The higher moments in the table, which are standardized by powers of the sample standard deviation, are reported for completeness. Because they are strongly influenced by extreme outliers I do not focus on them.

The shape of the empirical return distributions is evident from the selected percentiles of the distributions reported in the table and the densities plotted in Fig. 1. To aid in the interpretation of these densities, densities of normal distributions (dashes) and t distributions (dots) are also plotted.[11] It is clear from the figure that the distributions are nonnormal. The t distribution is a better description of the data. The tails are approximately symmetric, in line with the view of Fama (1976). For example, the largest difference between the lower-tail absolute percentiles and upper-tail percentiles occurs at the one percentile standardized return for the entire sample of stocks. The lower tail is -2.95 and the upper tail is 3.13. Evaluated at the mean standard deviation of monthly returns of 13 percent, these standardized returns correspond to stock returns of about -38 and 41 percent, respectively.

The symmetry apparent in Fig. 1 appears to contradict the negative skewness coefficients reported in the table. However, given the extremely high kurtosis of the distributions, these coefficients are simply describing a few observations. We would like a formal statistical test of whether the tails of the distribution of monthly returns are symmetric. Statistical tests of symmetry are quite sensitive to the choice of the null symmetric distribution used to construct them. Rather than picking a particular theoretical distribution, I assume that the true distribution is symmetric around zero and that the true density of absolute returns is identical to the empirical density. I then use a bootstrap approach to construct confidence intervals for the percentiles reported in Table 1.

Consider, for example, the distribution of returns to the entire sample of stocks. I randomly draw (with replacement) 1,633,649 absolute returns from the empirical density of absolute returns. I then randomly assign each of these absolute returns a sign (with equal probabilities for positive and negative returns) and determine various percentiles of the resulting distribution of returns. I repeat this exercise 10,000 times, producing a distribution of percentiles. The standard deviations of the bootstrapped 0.005, 0.01, and 0.05 percentiles are reported in parentheses in the table, underneath their respective percentiles. By construction, these are identical to the standard deviations of the bootstrapped 0.995, 0.99,

[11] These distributions are transformed to have the same standard deviations as the empirical distributions. The t distribution in the top panel has four degrees of freedom and the t distribution in the bottom panel has five degrees of freedom.

and 0.95 percentiles. I also calculate the standard deviation of the difference between the absolute lower-tail percentiles and the upper-tail percentiles. These standard deviations are reported in parentheses underneath their respective upper-tail percentiles.

The bootstrapped standard errors allow us to reject the hypothesis that the distribution of monthly idiosyncratic returns to all stocks is symmetric. For example, the one percentile values of -2.95 and 3.13 differ by 0.18. The standard error is 0.014, thus the difference is nearly 13 standard deviations away from zero. The distribution of monthly idiosyncratic returns to large-cap stocks is statistically indistinguishable from a symmetric distribution.

3.4.2 Five-year idiosyncratic returns

Table 2 reports selected statistics for distributions of five-year idiosyncratic returns. There are 29,000 observations for the entire sample of stocks and slightly more than half that number for the sample of large-cap stocks. As with monthly returns, the means of these standardized returns are slightly different from zero, so I demeaned the series before calculating the statistics. The fraction of these returns that were truncated prior to five years ranges from 20 percent (large-cap stocks) to 30 percent (all stocks). The first six rows in the table are based on returns defined by (8), which means that truncated returns are not scaled up to account for their shorter horizons. For expositional purposes I refer to these as "unscaled" returns, as indicated in the second column in the table. The final two rows are based on "scaled" returns, which use (9) to scale up the variance of truncated returns. Because the behavior of truncated returns drives the interesting features of the distribution of five-year returns, the table also breaks the results down by truncated versus non-truncated returns.

There are two main conclusions to draw from these results. First, the distributions of five-year returns are strongly asymmetric: the lower tails are much fatter than the upper tails. Second, most of the asymmetry is driven by the returns to stocks that disappear before the end of the five-year period. I discuss these conclusions in turn.

The best way to see the asymmetry is to compare the upper and lower percentiles of the empirical distributions. For example, we see from the first and fourth rows of the table that the lower one percentiles of the unscaled distributions are about 1.5 times the size (in absolute value) of their corresponding upper one percentile values. As discussed in Section 3.1, the asymmetry in the distribution of unscaled long-horizon returns underestimates the asymmetry in the true distribution of long-horizon returns. From the seventh and eight rows, we see that for scaled returns the ratio is even larger. In every case, the hypothesis that these upper and lower percentiles are equal in absolute value is overwhelmingly rejected. (The bootstrapped standard deviations are produced using the same methodology described for monthly returns.)

It is perhaps more intuitive to discuss this asymmetry in terms of returns rather than standardized returns. To illustrate the asymmetry of returns, consider returns to large-cap stocks. Using the percentiles of the unscaled distribution and an 11 percent monthly standard deviation of returns, the one percentile lower tail corresponds to a return of −328 percent ($-3.85\sqrt{60} \times 11$). The one percentile upper tail corrresponds to a return of 231 percent. For scaled returns to large-cap stocks, the one percentile five-year returns are −366 percent and 245 percent.

Fig. 2 and Fig. 3 plot the densities of standardized five-year unscaled and scaled returns, respectively. Densitites of t distributions with four degrees of freedom are also plotted.[12] The difference between these two figures arise from the treatment of truncated returns. The sharp peaks in Fig. 2 are attributable to the returns that are truncated within a few months. These returns tend to be small in absolute value. When scaled up, these returns contribute to the flattening of the overall distribution of returns seen in Fig. 3.

The treatment of truncated returns has a large effect on the overall distribution of returns because most of the asymmetry in these five-year returns is attributable to truncated returns. This conclusion is obvious from a glance at the empirical densities for truncated and non-truncated returns in Fig. 4 (all stocks) and Fig. 5 (large market-cap stocks). In these figures the truncated returns are unscaled. Conditional on a stock surviving through the five-year period, the distribution of its return is closer to a symmetric distribution than is the unconditional distribution of returns. This is verified by the percentiles reported in Table 2. Conditional on truncation, the return is highly volatile and strongly asymmetric. For the entire sample of stocks, the mean truncated return is negative. None of these facts are particularly surprising. What is a little surprising is that the mean truncated return for the sample of large-cap stocks is slightly positive.

It is clear from a comparison of the distributions of monthly returns and five-year returns that the simple intuition of the CLT does not fit the facts. Since this intuition relies on the independence of short-horizon returns, a natural place to look for an explanation is in potential dependencies among returns. Given the asymmetry in the distribution of long-horizon returns, a promising source of dependence is asymmetric volatility: The tendency of volatility to vary with past returns.

[12]These distributions are transformed to have the same standard deviations as the empirical distributions.

4 Aymmetric volatility

Asymmetric volatility is best known in the context of aggregate stock returns. A large literature (Black (1976) and French, Schwert, and Stambaugh (1987) are early contributions) documents a strong inverse relation between returns and future volatility. As discussed in Section 2, the relation at the firm level is more ambiguous. A negative relation between returns and future volatility will result in negatively-skewed long-horizon returns even if short-horizon returns are symmetric. If period t's return is positive and large, the probability of another large positive return declines because the conditional variance of returns declines. If period t's return is negative and large, the probability of a large negative return rises.

The most obvious source of asymmetric volatility is a firm's balance sheet. Balance-sheet effects can induce either a positive or negative relation between returns and future volatility. Rubinstein (1983) shows that a positive relation is produced when firms hold assets with different levels of risk. Christie (1982) formalizes the negative relation produced by financial leverage. A simple theoretical model illustrates these two effects.

4.1 A model of asymmetric volatility

Assume that a firm has a risky asset-in-place with market value $F(t)$. It also holds an intantible asset with market value $G(t)$. We can think of this informally as a growth option, although careful modeling of growth options is beyond the scope of this paper. The asset values evolve according to the continuous-time processes

$$\frac{dF}{F} = \mu_F dt + v_F^I dZ_f(t) + v_F^C dZ_c(t),$$

$$\frac{dG}{G} = \mu_G dt + v_G^I dZ_g(t) + v_G^C dZ_c(t)$$

where $dZ_f(t)$ and $dZ_g(t)$ are independent shocks to the two asset values, v_F^I and v_G^I are the responses of the log asset values to these shocks, $dZ_c(t)$ is a shock common to all firms' assets' values, and v_F^C, v_G^C are the responses of log asset values to this common shock. The working hypothesis here is that the volatility of the return to the intangible asset exceeds that of the return to the assets in place, or $v_G^I = kv_F^I$, $v_G^C = kv_F^C$, $k > 1$.

The firm has issued default-free floating rate zero-coupon debt with market value $D(t)$. The dynamics of $D(t)$ are

$$\frac{dD}{D} = rdt.$$

The market value of the firm's equity is $S(t) = F(t) + G(t) - D(t)$. (The firm pays no dividends.) Equity dynamics are

$$\frac{dS}{S} = \left[\frac{F(t)}{S(t)}\mu_F + \frac{G(t)}{S(t)}\mu_G - \frac{D(t)}{S(t)}r\right] dt + v_{FS}^I(t)dZ_f(t) + v_{GS}^I(t)dZ_g(t) + v_S^C(t)dZ_c(t)$$

where equity volatilities are

$$v_{FS}^I(t) = \frac{F(t)}{S(t)}v_F^I \tag{10}$$

$$v_{GS}^I(t) = \frac{kG(t)}{S(t)}v_F^I \tag{11}$$

$$v_S^C(t) = \frac{F(t)+kG(t)}{S(t)}v_F^C. \tag{12}$$

We are interested in how the firm's stock return volatility varies with its stock price. It is fairly easy to see that the sign of this relation is ambiguous. There are three effects at work. They are the standard leverage effect (inducing a negative relation), the asset-mix effect (inducing a positive relation), and diversification (ambiguous; positive if the higher-risk asset is less valuable than the lower-risk asset, negative otherwise.) To clarify these effects, consider a stock-price increase. The increase lowers the leverage of the stock and thus lowers its return volatility. However, the increase in price is more likely the result of an increase in the value of the high-volatility asset than in the value of the low-volatility asset. Thus on average, an increase in stock price corresponds to an increase in the proportion of the total assets of the firm that are high-volatility, raising the stock's return volatility. If, however, the high-volatility asset is a small contributor to the firm's overall asset volatility, an increase in its value can decrease the stock's volatility through greater diversification. (Of course, the opposite can hold; if the high-volatility asset is the major contributor to overall asset volatility, an increase in its value lowers diversification.)

Diversification has no effect on exposure to the common component. Therefore the asset-mix effect implies that an increase in the value of the intangible asset relative to the value of the tangible asset unambiguously raises the sensitivity of the stock's return to the common shock. The effect of this relative change on the stock return's idiosyncratic volatility can be greater or less than its effect on exposure to the common component, depending on whether the change decreases or increases diversification.

This model implies that if the ratio of tangible to intangible assets is held constant, an increase in financial leverage corresponds to higher idiosyncratic volatility and higher common-factor exposure. In addition, holding leverage constant, an increase in the ratio of intangible to tangible assets raises common-factor exposure and (assuming no strong diversification effects) also raises idiosyncratic volatility. The net relation between returns and either idiosyncratic volatility or common-factor exposure is unsigned. However, if the relation is negative, it suggests that the leverage effect outweights the asset-mix effect.

I test these hypotheses using annual measures of stock returns, idiosyncratic stock return volatility, and balance sheet data. The data are summarized in the next subsection.

4.2 Data construction and summary

Balance sheet information is from Compustat annual files. The available data limit the analysis to fiscal years from 1961 through 2000. Firm i's book value of equity for year-end t, denoted $BE_i(t)$, is constructed using the definition in Table I of Fama and French (1995). Book values less than zero are ignored. Book value of debt, $D_i(t)$, is the sum of short-term liabilities, long-term liabilities, and preferred stock. I do not require that firms have December fiscal year-ends, therefore balance sheet information for fiscal year t can be realized as late as the end of May in calendar year $t + 1$. Accordingly, I define the market value of equity for fiscal year t, $ME_i(t)$, as the market capitalization of firm i's stock as of the end of June in calendar year $t + 1$.

The accounting variables that I use are (1) assets/equity measured using market value of equity; (2) assets/equity measured using book value of equity; and (3) book equity/market equity. Because the extremely high asset/equity ratios of financial firms might skew the results, I dropped financial firms (SIC codes 6000-6999) from the sample. I trimmed outliers of book equity as follows. Consider the universe of observations of assets/book equity available on Compustat across all non-financial firms and all years. If the observation for firm i and year t is in the upper 1/2 percent of this distribution, I set the observation of book equity to a missing value.

I also construct yearly excess returns, yearly estimates of the variance of idiosyncratic returns, and yearly estimates of Fama-French factor loadings for each year from 1961 through 2001. For consistency with the dating of market capitalization, I define the excess return to stock i during fiscal year t as as the sum of monthly excess returns to the stock from the end of June in calendar year t to the end of June in calendar year $t + 1$.[13] Therefore the accounting ratios for year t are realized by the beginning of the period over which year $t+1$'s return is measured. If there any months without valid stock returns I set this annual return to a missing value. [14] The annual return is denoted $R_i(t)$.

[13]This is the sum of monthly returns in levels, not logs. This definition of yearly returns follows Grundy and Martin (2001).

[14]The year 2001 begins at month-end June 2001 and ends with month-end June 2002. Since I use the 2001 CRSP tape, stock returns for the final six months are unavailable. Therefore the 2001 "year" of stock

Year t's variance of annual idiosyncratic returns is estimated by the sum of squared monthly idiosyncratic returns. The monthly idiosyncratic returns are produced using the methodology of Section 3. This estimate is denoted $V_i(t)$. Note that this is a measure of realized volatility, not a conditional estimate formed in year $(t-1)$ of the volatility of year t's return. If any of these monthly returns are missing I set the variance to a missing value. Annual estimates of loadings on the Fama-French factors are constructed with regressions. For each stock and each year, I estimate a regression of the firm's monthly excess stock return on the factors using the twelve observations in the year. If any of the monthly returns is missing I set the estimates of the loadings to missing values. As with stock return variances, these are measures of realized factor loadings, not conditional factor loadings.

Because I use a panel data approach with fixed effects, the explanatory power in the data lies in the time-series variation for a given stock. Thus I drop all stocks for which there are fewer than 15 yearly observations available to estimate a regression of volatility on the previous year's stock return and balance sheet ratios. This filter restricts the sample to stocks of 1,620 firms. These firms have, on average larger market capitalizations than the typical firm. More importantly, the sample of firms is largely drawn from survivors. As we saw in Section 3.4, long-run stock returns that are not truncated (i.e., they are the returns to surviving stocks) exhibit much less extreme behavior than do long-run stock returns that are truncated. Thus the selection criteria focuses our attention on firms for which asymmetries are muted. More generally, the focus on survivors limits our ability to say much about the dynamics of stock return behavior of firms that get into trouble.

Table 3 presents summary information for the stock return variables and balance-sheet ratios. For example, the mean annual excess return to a stock is 10 percent with a standard deviation of 36 percent. This standard deviation, like all the standard deviations reported in this column, is calculated by first removing both time and cross-sectional fixed effects. Annual returns exhibit a small amount of negative serial correlation. The AR(1) coefficient is calculated after removing both sets of fixed effects, thus these coefficients are subject to a fairly strong small-sample bias. (Effectively the AR coefficient is a cross-sectional mean of individual AR coefficients, which are based on an average of 30 annual observations per stock.)

The table reports that the annual volatility estimates are moderately persistent; the AR(1) coefficient is 0.32. This coefficient is downward biased both because of the small-sample bias and because of the errors-in-variables problem associated with using an estimate of annual volatility on the right-hand-side instead of true annual volatility. This errors-in-variables problem is especially apparent in the annual estimates of factor loadings. These

returns consists of six monthly observations.

estimates are very noisy because they are taken from a regression of 12 observations using four explanatory variables (a constant and the factor returns). The substantial noise in the estimates shows up in the small AR(1) coefficients.

The mean balance-sheet ratios indicate that the typical firm has less debt than equity and higher market equity than book equity. All the ratios have mean AR(1) coefficients of about 0.7. Again, these coefficients are downward biased because of the small sample sizes.

4.3 The relation between returns and idiosyncratic volatility

In this subsection I test for the presence of asymmetric volatility. The regression is

$$\begin{aligned}\log V_i^{1/2}(t) - \log V_i^{1/2}(t-4) &= \Big[b_1 + b_2 R_i(t-2) + b_3 R_i(t-3)\Big] R_i(t-1) \\ &+ \Big[b_4 + b_5 R_i(t-3)\Big] R_i(t-2) + b_6 R_i(t-3) + g_i(t). \quad (13)\end{aligned}$$

The log change in realized idiosyncratic volatility from year $t-4$ to year t is forecasted by the annual returns in the years $t-1, t-2$, and $t-3$. Volatility in year $t-4$ plays the role of the baseline. The predicted values of $V_i(t)$ can be thought of as year $t-1$ conditional variances for year t idiosyncratic returns. The cross terms are suggested by the theoretical model. To see this clearly, drop the intangible asset from the model so that it is a model of pure financial leverage. Then the regression coefficient on year $t-1$'s return should equal the negative of the firm's debt/asset ratio as of year $t-1$. Holding all else constant, if year $t-2$'s stock return was positive, the year $t-1$ debt/asset ratio will be lower, thus the coefficient on year $t-1$'s return should be closer to zero. This logic implies that the coefficients b_1, b_4, and b_6 should be negative, while the coefficients b_3, b_4, and b_5 should be positive.

There are 41 years of data for use in estimating (13). Therefore t ranges from 1965 through 2001. I estimate the regressions jointly for each stock using a panel data approach. Fixed effects are included in the regression, which explains the lack of a constant term. The standard errors are calculated assuming that the residuals of the volatility regressions are uncorrelated across stocks and time. In practice, this assumption is too strong, even when cross-sectional and yearly fixed effects are included in the regressions. For example, industry and size effects will produce correlated residuals for firms that are similar to each other. Therefore the results are subject to the caveat that the standard errors underestimate the true variability in the parameter estimates.

The results are in Table 4. Four sets of results are presented, depending on whether the cross-terms are included and whether firm-specific fixed effects are included. (All regressions include annual fixed effects.) The latter results are included for completeness, but the inclusion of cross-sectional fixed effects is problematic. First, because the dependent variable is a change in volatility rather than a level, the firm-level fixed effects are not particularly useful in this regression. Second, their use can magnify a predictive regressions bias (Stambaugh (1999)). Because returns and changes in idiosyncratic volatility are positively contemporaneously correlated (see Panel B of Table 3), there is a finite-sample negative bias in a regression of changes in volatility on lagged returns. With firm-level fixed effects, the regression equation is effectively estimated on a smaller sample, magnifying the bias. Third, firm-level fixed effects produce the same problem discussed in Section 3.2.4. Stocks that were winners (losers) over the sample period tended to exhibit declining (increasing) volatility over the period. Including firm-level fixed effects eliminates this information from the sample.

There are two main conclusions to draw from these results. The first conclusion is that there is a modest (but statistically strong) and persistent negative relation between returns and future idiosyncratic volatility. The results indicate that if a firm's stock price rises in a given year, idiosyncratic volatility will be higher for at least the next three years. To get a sense of the magnitude, compare the implications for volatility of a year t return of either 36 or -36 percent. (This is one standard deviation.) Using the results of regression [1], the implied standard deviation of idiosyncratic returns in year $t+1$ given a return of -36 percent is only 1.026 times the corresponding standard deviation given a return of 36 percent. ($\exp[(-0.36 \times -0.036) - (0.36 \times -0.036)]$). The implied standard deviations in years $t+1$ and $t+2$ are more sensitive to this return, but the magnitudes still seem quite small. This behavior strongly contrasts with asymmetric volatility at the aggregate level. The aggregate relation is negative, very strong at short horizons, and dissipates quickly.[15]

The second conclusion is that the interaction between returns is in the direction implied by the leverage hypothesis. Volatility in year t is more sensitive to year $t-1$'s stock return if the stock fell in value in years $t-2$ and $t-3$. For example, if $R_i(t-2) = R_i(t-3) = -0.36$, the responsiveness of year t's log volatility to year $t-1$'s return implied by the estimates of regression [2] is -0.077. If the returns in years $t-2$ and $t-3$ were 36 percent instead of -36 percent, the implied responsiveness is only -0.013.

Since Black (1976), researchers have proposed a host of theories that link returns and volatility.[16] It is almost certain that there are multiple effects at work. Without multiple effects, it seems impossible to explain why there exist simultaneously (a) a large and short-lived inverse relation between aggregate returns and aggregate return volatility; (b) a positive

[15]The short-lived nature of the aggregate relation is documented in Figlewski and Wang (2000) and Duffee (2002).

[16]I do not attempt a review here. Duffee (2002) discusses some of this literature.

and short-lived relation between short-horizon returns and idiosyncratic return volatility (documented in Duffee (1995, 2002)); and (c) a small and long-lived inverse relation between returns and idiosyncratic return volatility documented in this paper. The long persistence of the latter relation is suggestive of balance-sheet effects. The link between the balance sheet and volatility is explored next.

4.4 Idiosyncratic volatility and balance sheets

The basic approach here is to forecast the change in realized idiosyncratic volatility from year $t-2$ to year t using variables realized in year $t-1$. The main regression is

$$\begin{aligned} \log V_i^{1/2}(t) - \log V_i^{1/2}(t-2) &= b_1 R_i(t-1) + b_2 \Delta \log \frac{D_i + ME_i}{ME_i}(t-1) \\ &+ b_3 \Delta \log \frac{BE_i}{ME_i}(t-1) + g_i(t) \end{aligned} \tag{14}$$

where $\Delta x(t-1) = x(t-1) - x(t-2)$. As with (13), the predicted values of $V_i(t)$ are year $t-1$ conditional variances for year t idiosyncratic returns. The asset/equity ratio is included to pick up variations in financial leverage and the book/market ratio is designed to pick up variations in the ratio of assets-in-place to intangible assets such as growth options. The model of Section 4.1 assumes that these intangible assets are riskier than tangible assets. If the model is correct, the negative relation between returns and future idiosyncratic volatility documented in Section 4.3 is a consequence of the leverage effect outweighing the asset-mix effect.

The range of t is 1963 through 2001. I estimate (14) across all stocks using annual fixed effects. In results not reported here, I also included cross-sectional fixed effects. They had little effect on the results, which is likely a consequence of the fact that all of the balance-sheet variables are in logs and first-differenced.

Estimation results are in Table 5. There are two main conclusions to draw from these results. The first conclusion is that conditional volatilities depend on leverage. Consider, for example, regression [2] in the table, which excludes the book/market ratio. The elasticity of idiosyncratic volatility with respect to the asset/equity ratio is about 0.2, and is strongly statistically significant. This result is consistent with earlier research that finds a positive relation between leverage and volatility at shorter horizons. A typical result for short-horizon returns is that changes in leverage cannot completely explain the inverse relation between returns and volatility.[17] But at the annual horizon examined here, changes in the asset/equity ratio capture the negative relation between returns and future volatility: the coefficient on

[17] See, e.g., Bekaert and Wu (2000) and Figlewski and Wang (2000).

returns in [2] is actually positive.

A well-known criticism of regressions like (14) is that most of the action in leverage ratios over time is in the denominator. Since we know that volatility rises after stock prices fall, any variable included on the right-hand-side that has the stock price in the denominator is likely to be positively associated with volatility. Therefore it is possible that financial leverage matters only to the extent that it allows the regression to pick up a nonlinear relation between returns and future volatility. To investigate this possibility, I replaced the market value of equity in the numerator and denominator of asset/equity with the book value of equity. This measure of asset/equity emphasizes fluctuations in leverage owing to changes in the firm's debt load over time. The results for regression [3] in the table indicate that this proxy for the asset/equity ratio also forecasts realized volatility. The additional noise in this proxy results in a smaller estimate (an elasticity of 0.13), but the estimate remains strongly statistically significant.

The second conclusion to draw from Table 5 is that conditional volatilities depend on book/market. When book/market rises, future volatility declines. Regression [4] in the table reports the estimated elasticity as -0.08. Since returns and book/market are negatively contemporaneously correlated, this effect runs counter to the overall negative relation between returns and future volatility. Put differently, the negative relation between returns and future volatility is stronger when book/market is held constant, as in regression [4].

The results of regression [5], which includes both assets/equity and book/market, can be interpreted as follows. An increase in a firm's stock price in year $t-1$ implies a contemporaneous decrease in both its leverage and its book/market ratio. (See Panel B of Table 3.) The decrease in leverage corresponds to lower volatility of year t's idiosyncratic return, while the decrease in book/market corresponds to higher volatility. These effects largely offset each other; the coefficient on the lagged stock return in [5] is almost identical to the coefficient on the lagged stock return in [1], which excludes the balance-sheet variables.

These results provide support for the model sketched in Section 4.1. If the model is correct, we should see the same kind of patterns in exposure to marketwide risks. Lower leverage should be associated with less exposure to common factors, while lower book/market should be associated with greater exposure to these factors.

4.5 Factor loadings and balance sheets

The regression approach here is the same I used above. For example, a stock's conditional year-t loading on the Fama-French market factor is expressed as

$$\begin{aligned}\beta_{i,M}(t) - \beta_{i,M}(t-2) &= b_1 R_i(t-1) + b_2 \Delta \log \frac{D_i + ME_i}{ME_i}(t-1) \\ &+ b_3 \Delta \log \frac{BE_i}{ME_i}(t-1) + g_i(t) \end{aligned} \tag{15}$$

I also estimate versions of (15) with different dependent variables. I use the SMB factor loading, the HML factor loading, and a measure of the total return volatility owing to the Fama-French factors. This measure is $\log(\beta_{it}' \Sigma \beta_{it})^{1/2}$, where Σ is the observed variance-covariance matrix of monthly returns to the factors from January 1961 through December 2001. The final dependent variable I consider is a measure of the total risk premium owing to these factors. This measure is $\beta_{it}' P$, where P is the observed mean return vector to the Fama-French factors from January 1961 through December 2001. To conserve space, I report results using only two combinations of explanatory variables: the year $t-1$ stock return by itself and all three variables. The results are displayed in Table 6.

Perhaps the most surprising result in the table is that stock returns are (weakly) positively associated with future volatility owing to common factors. The coefficient of 0.018 is statistically indistinguishable from zero, but the news is that the coefficient is not strongly negative. Chan (1988) and Ball and Kothari (1989) find that at the portfolio level, betas of "winners" tend to decline, while betas of "winners" tend to increase. Part of the explanation for the divergence in results is that Chan and Ball-Kothari estimate their "before" betas over the same period for which they define winners and losers. In results I do not report in detail here, I find that betas and returns are positively contemporaneously correlated. A stock that is a winner in year $t-1$ has a higher exposure to market risks in year $t-1$ than it had in year $t-2$; in year t, this exposure returns to its original level.

Although a stock's total common-factor exposure does not fall after its price rises, its risk premium (as measured by F-F exposure) does fall. The reason is that its factor exposure shifts away from the high-premium SMB and HML factors to the low-premium market factors. This is consistent with the results of Fama and French (1996) discussed in Section 2. Put differently, the results here do not resolve the momentum puzzle.

For the purposes of evaluating the balance sheet model of volatility, the important regression is the regression of changes in total common-factor exposure on the three explanatory variables. The results of this regression are close to the results for changes in idiosyncratic volatility. The elasticity of total common-factor exposure with respect to the asset/equity ratio is about 0.2, while the elasticity with respect to the book/market ratio is about -0.1. Controlling for these balance sheet variables, the relation between stock returns and future total factor exposure is statistically indistinguishable from zero. In other words, total factor exposure and idiosyncratic return volatility respond to these variables in almost the same way.

5 Concluding comments

In this paper I make three contributions to our knowledge of the long-run behavior of stock returns. First, I show that long-horizon log idiosyncratic returns to a typical stock have much fatter lower tails than upper tails. This result is surprising, because distributions of shorter-horizon idiosyncratic returns are nearly symmestric (and to the extent that they are not, they have slightly fatter upper tails than lower tails). Second, I show that yearly idiosyncratic stock returns exhibit highly persistent asymmetric volatility. Third, I link changes in firms' stock return volatility to changes in their balance sheets. Theory tells us that both the leverage effect and the asset-mix effect should affect the return–volatility relation, and the theory is supported here. Evidence of the link to leverage reinforces earlier work, while evidence of the link to book/market (proxying for variations in asset mix)is new.

These results naturally lead to a large question that I do not attempt to answer here. Is the magnitude of asymmetric volatility observed in long-horizon returns sufficient to explain their asymmetric distributions? A careful answer to this question requires estimation of a realistic model of firm-level asset dynamics that combines leverage and asset-mix effects. The goal of the current paper is to document features of long-horizon returns without imposing much overidentifying structure. Therefore I defer this question to later research.

References

Ball, Ray, and S.P. Kothari, 1989, Nonstationary expected returns: Implications for tests of market efficiency and serial correlation in returns, *Journal of Financial Economics* 25, 51-74.

Bekaert, Geert, and Goujun Wu, 2000, Asymmetric volatility and risk in equity markets, *Review of Financial Studies* 13, 1-42.

Berk, Jonathan B., Richard C. Green, and Vasant Naik, 1999, Optimal investmetn, growth options, and security returns, *Journal of Finance* 54, 1553-1607.

Black, Fischer, 1976, Studies of stock price volatility changes, *Proceedings of the 1976 Meetings of the Business and Economics Statistics Section*, 177-181.

Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637-659.

Braun, Phillip A., Daniel B. Nelson, and Alain M. Sunier, 1995, Good news, bad news, volatility, and betas, *Journal of Finance* 50, 1575-1603.

Campbell, John Y., Martin Lettau, Burton G. Malkiel, and Yexiao Xu, 2001, Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk, *Journal of Finance* 56, 1-43.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, *The Econometrics of Financial Markets*, Princeton University Press, Princeton, NJ.

Chan, K.C., 1988, On the contrarian investment strategy, *Journal of Business* 61, 147-163.

Cheung, Yin-Wong, and Lilian Ng, 1992, Stock price dynamics and firm size: An empirical investigation, *Journal of Finance* 47, 1985-1997.

Cho, Young-Hye, and Robert F. Engle, 1999, Time-varying betas and asymmetric effects of news: Empirical analysis of blue chip stocks, NBER working paper 7330.

Christie, Andrew A., 1982, The stochastic behavior of common stock variances, *Journal of Financial Economics* 10, 407-432.

DeBondt, Werner F.M., and Richard H. Thaler, 1985, Does the stock market overreact? *Journal of Finance* 42, 557-581.

Duffee, Gregory R., 1995, Stock returns and volatility, *Journal of Financial Economics* 37, 399–420.

Duffee, Gregory R., 2002, Balance sheet explanations for asymmetric volatility, working paper, Haas School, University of California–Berkeley.

Eom, Young Ho, Jean Helwege, and Jing-zhi Huang, 2002, Structural models of corporate bond pricing: An empirical analysis, working paper, College of Business, Ohio State University.

Fama, Eugene F., 1976, *Foundations of Finance*, Basic Books, New York, NY.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3-56.

Fama, Eugene F., and Kenneth R. French, 1995, Size and book-to-market factors in earnings and returns, *Journal of Finance* 50, 131-155.

Fama, Eugene F., and Kenneth R. French, 1996, Multifactor explanations of asset pricing anomalies, *Journal of Finance* 51, 55-84.

Figlewski, Stephen, and Xiaozu Wang, 2000, Is the "leverage effect" a leverage effect?, Working paper, NYU Stern School.

French, Kenneth R., G. William Schwert, and Robert Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.

Grundy, Bruce D., and J. Spencer Martin, 2001, Understanding the nature of the risks and the source of the rewards to momentum investing, *Review of Financial Studies* 14, 29-78.

Jagannathan, Ravi, and Zhenyu Wang, 1996, The conditional CAPM and the cross-section of expected returns, *Journal of Finance* 51, 3-53.

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65-91.

Lauterbach, Beni, and Paul Schultz, 1990, Pricing warrants: An empirical study of the Black-Scholes model and its alternatives, *Journal of Finance* 45, 1181-1209.

Merton, Robert C., 1974, On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance* 29, 449-470.

Rubinstein, Mark, 1983, Displaced diffusion option pricing, *Journal of Finance* 38, 213-217.

Schwert, G. William, 1989, Why does stock market volatility change over time, *Journal of Finance* 44, 1115-1153.

Shumway, Tyler, 1997, The delisting bias in CRSP data, *Journal of Finance* 52, 327-340.

Shumway, Tyler, and Vincent A. Warther, 1999, The delisting bias in CRSP's Nasdaq data and its implications for the size effect, *Journal of Finance* 54, 2361-2379.

Simkowitz, Michael A., and William L. Beedles, 1978, Diversification in a three-moment world, *Journal of Financial and Quantitative Analysis* 13, 927-941.

Simkowitz, Michael A., and William L. Beedles, 1980, Asymmetric stable distributed security returns, *Journal of the American Statistical Association* 75, 306-312.

Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375-421.

						Percentiles					
Sample	Obs	$\widehat{SE}_{it}$	Std. Dev.	Skew	Kurtosis	0.005	0.01	0.05	0.95	0.99	0.995
All stocks	1633649	0.133	1.195	−1.409	51.153	−3.61	−2.95	−1.75	1.85	3.13	3.75
						(0.013)[a]	(0.008)[a]	(0.003)[a]	(0.005)[b]	(0.014)[b]	(0.022)[b]
Large-cap stocks	861908	0.108	1.131	−0.387	20.755	−3.46	−2.89	−1.75	1.79	2.90	3.42
						(0.016)[a]	(0.010)[a]	(0.004)[a]	(0.006)[b]	(0.016)[b]	(0.027)[b]

Table 1: Summary statistics for standardized monthly stock returns

Rolling regressions of 60 months are used to estimate stock return betas and corresponding standard deviations of idiosyncratic returns for stocks in the given samples. The parameter estimates are used to construct standardized one-month-ahead idiosyncratic returns. The table reports the grand means (across stocks and months) of the parameter estimates from these regressions. It also reports the first four moments and selected percentiles for the distribution of demeaned one-month-ahead standardized returns.

[a] Bootstrapped standard errors for the given percentile, assuming symmetry in the distribution.

[b] Bootstrapped standard errors for the difference between the given percentile and the absolute value of the corresponding lower-tail percentile, assuming symmetry in the distribution.

Sample	Scaled	Obs	Mean	Std. Dev.	Skew	Kurtosis	Percentiles					
							0.005	0.01	0.05	0.95	0.99	0.995
All stocks	N	29066	0.000	1.273	−1.577	13.491	−5.44	−4.08	−2.05	1.76	2.70	3.06
							(0.153)[a]	(0.072)[a]	(0.027)[a]	(0.044)[b]	(0.117)[b]	(0.248)[b]
Truncated rets	N	8224	−0.156	1.491	−2.719	18.342	−7.37	−6.10	−2.82	1.51	2.34	2.77
Non-truncated rets	N	20842	0.062	1.170	−0.513	6.119	−3.87	−3.19	−1.82	1.84	2.76	3.17
Large-cap stocks	N	16764	0.000	1.246	−1.695	16.160	−5.10	−3.85	−1.97	1.72	2.71	3.08
							(0.163)[a]	(0.093)[a]	(0.029)[a]	(0.049)[b]	(0.152)[b]	(0.262)[b]
Truncated rets	N	3092	0.025	1.543	−3.837	28.806	−8.08	−6.60	−2.67	1.55	2.28	2.65
Non-truncated rets	N	13672	−0.006	1.169	−0.535	6.001	−3.88	−3.32	−1.90	1.77	2.75	3.17
All stocks	Y	29066	0.000	1.604	−3.583	53.587	−7.61	−5.78	−2.32	2.01	2.99	3.42
							(0.254)[a]	(0.141)[a]	(0.031)[a]	(0.051)[b]	(0.229)[b]	(0.408)[b]
Large-cap stocks	Y	16764	0.000	1.433	−4.440	103.602	−6.25	−4.30	−2.04	1.89	2.88	3.27
							(0.284)[a]	(0.093)[a]	(0.033)[a]	(0.055)[b]	(0.151)[b]	(0.459)[b]

Table 2: Summary statistics for standardized five-year idiosyncratic stock returns

Non-overlapping five-year idiosyncratic stock returns for stock i are constructed by summing monthly idiosyncratic returns for 60 months. If any of these returns is missing the return is truncated at the month τ preceding the first missing observation. The returns are divided by the product of the conditional standard deviation of the first month's idiosyncratic return and either $\sqrt{60}$ (unscaled) or $\sqrt{\min(60,\tau)}$ (scaled). This table reports the first four moments and selected percentiles for the entire distribution of five-year returns across all stocks in the sample.

[a] Bootstrapped standard errors for the given percentile, assuming symmetry in the distribution.
[b] Bootstrapped standard errors for the difference between the given percentile and the absolute value of the corresponding lower-tail percentile, assuming symmetry in the distribution.

Panel A. Univariate statistics

Variable	Obs.	Mean	Median	Std. Dev.	AR(1)
Excess returns	48306	0.100	0.077	0.362	−0.027
SD of idiosyncratic returns	42721	0.310	0.270	0.124	0.319
β (market)	48306	0.849	0.899	1.120	−0.012
β (SMB)	48306	0.681	0.502	1.800	0.000
β (HML)	48306	0.245	0.243	1.969	0.013
log (D+ME)/ME	48308	0.436	0.309	0.293	0.765
log (D+BE)/BE	48000	0.511	0.423	0.308	0.760
log BE/ME	45528	−0.391	−0.325	0.534	0.695

Panel B. Correlations of first differences (returns in levels)

	Excess returns	log idio SD	log (D+ME)/ME	log BE/ME
Excess returns	1.00			
log idio SD	0.11	1.00		
log (D+ME)/ME	−0.57	0.02	1.00	
log BE/ME	−0.79	−0.07	0.47	1.00

Table 3: Summary statistics for annual stock returns and balance sheet ratios

The sample consists of 1620 publicly-traded firms with at least 15 years of selected annual balance sheet ratios during 1961 through 2001. Annual excess stock returns, yearly estimates of the stocks' loadings on the Fama-French factors, and the standard deviation of annual idiosyncratic returns (i.e., after removing the Fama-French factor components) are calculated from monthly data. Balance-sheet ratios are (debt+equity)/equity measured with both market equity ME and book equity BE, and book equity/market equity. The table reports the number of observations (across years and stocks) for each variable, their unweighted means and medians, standard deviations, and AR(1) coefficients of the variables. The standard deviations and AR(1) coefficients are calculated after removing annual and stock-specific fixed effects. These fixed effects are also removed before calculating correlations.

Regression coef	Regression number [1]	[2]	[3]	[4]
b_1	-0.036	-0.045	-0.031	-0.040
	(0.007)	(0.007)	(0.007)	(0.007)
b_2	-	0.049	-	0.050
		(0.011)		(0.011)
b_3	-	0.041	-	0.043
		(0.012)		(0.012)
b_4	-0.066	-0.078	-0.060	-0.073
	(0.007)	(0.007)	(0.007)	(0.007)
b_5	-	0.066		0.065
		(0.011)		(0.011)
b_6	-0.072	-0.084	-0.063	-0.075
	(0.007)	(0.007)	(0.007)	(0.007)

Table 4: Conditioning volatilities of annual idiosyncratic stock returns on past returns, 1965 through 2001

The realized variance of firm i's idiosyncratic stock return in year t, denoted $V_i(t)$, is estimated by the sum of squared monthy idiosyncratic returns in year t. The change in this variance from year $t-4$ to year t is explained by excess annual returns to the stock in years $t-1, t-2$, and $t-3$.

$$\begin{aligned}\log V_i^{1/2}(t) - \log V_i^{1/2}(t-4) &= \Big[b_1 + b_2 R_i(t-2) + b_3 R_i(t-3)\Big] R_i(t-1) \\ &+ \Big[b_4 + b_5 R_i(t-3)\Big] R_i(t-2) + b_6 R_i(t-3) + g_i(t).\end{aligned}$$

A panel data approach is used to estimate the regressions across 1620 firms and 41 years (including lags), for a total of 35539 non-missing observations. Regressions [1] and [2] includes only annual fixed effects. Regressions [3] and [4] include both annual and cross-sectional fixed effects. Standard errors are in parentheses.

Lagged explanatory variable	Regression number [1] (39080 obs.)	[2] (38864 obs.)	[3] (36761 obs.)	[4] (36801 obs.)	[5] (36761 obs.)
Stock return	-0.034	0.037	-0.019	-0.093	-0.036
	(0.006)	(0.007)	(0.006)	(0.010)	(0.010)
Δ log asset/equity (market value)	-	0.219	-	-	0.209
		(0.012)			(0.014)
Δ log asset/equity (book value)	-	-	0.133	-	-
			(0.011)		
Δ log book/market	-	-	-	-0.077	-0.086
				(0.009)	(0.009)

Table 5: Conditioning volatilities of annual idiosyncratic stock returns on past returns and balance sheets, 1963 through 2001

The realized standard deviation of firm i's idiosyncratic stock return in year t is estimated by the square root of the sum of squared monthy idiosyncratic returns in year t. The log change in this standard deviation from year $t-2$ to year t is explained by excess annual stock returns to the stock in year $t-1$, the change in the log asset/equity ratio from year $t-2$ to year $t-1$, and the change in log (book equity)/(market equity) from year $t-2$ to year $t-1$. The asset/equity ratio is measured using either market equity or book equity. A panel data approach is used to estimate the regressions across 1620 firms and 41 years (including lags). The regressions include annual fixed effects. Standard errors are in parentheses.

Lagged explanatory variable	Dependent variable				
	Market beta	SMB beta	HML beta	log SD	Premium
First regression					
Stock return (only)	0.089	−0.258	−0.019	0.018	−0.026
	(0.022)	(0.033)	(0.006)	(0.009)	(0.003)
Second regression					
Stock return	0.206	−0.294	−0.529	0.002	−0.015
	(0.039)	(0.007)	(0.036)	(0.016)	(0.005)
Δ log asset/equity	0.130	0.401	0.199	0.216	0.026
	(0.052)	(0.079)	(0.086)	(0.022)	(0.006)
Δ log book/market	0.068	−0.159	0.048	−0.103	0.002
	(0.033)	(0.050)	(0.055)	(0.014)	(0.004)

Table 6: Conditioning betas on past returns and balance sheets, 1963 through 2001

Year-t loadings of the return to stock i on the three Fama-French factors are estimated using monthly returns during year t. An estimate of the log monthly standard deviation of returns owing to these common factors is $\log(\beta_{it}'\Sigma\beta_{it})^{1/2}$, where the loadings are stacked in the vector β_{it} and Σ is the unconditional variance-covariance matrix of monthly returns to the factors from January 1961 through December 2001. The implied annual risk premium is $\beta_{it}'P$, where P is the vector of mean returns to the factors from January 1961 through December 2001.

For each of these variables, two regressions are estimated. The first is a regression of the change in the variable from year $t-2$ to year t on the stock's excess return in year $t-1$. The second adds as explanatory variables the change in the log asset/equity ratio from year $t-2$ to year $t-1$ and the change in log (book equity)/(market equity) from year $t-2$ to year $t-1$. A panel data approach is used to estimate the regressions across 1620 firms and 41 years (including lags). The regressions include annual fixed effects. Standard errors are in parentheses.

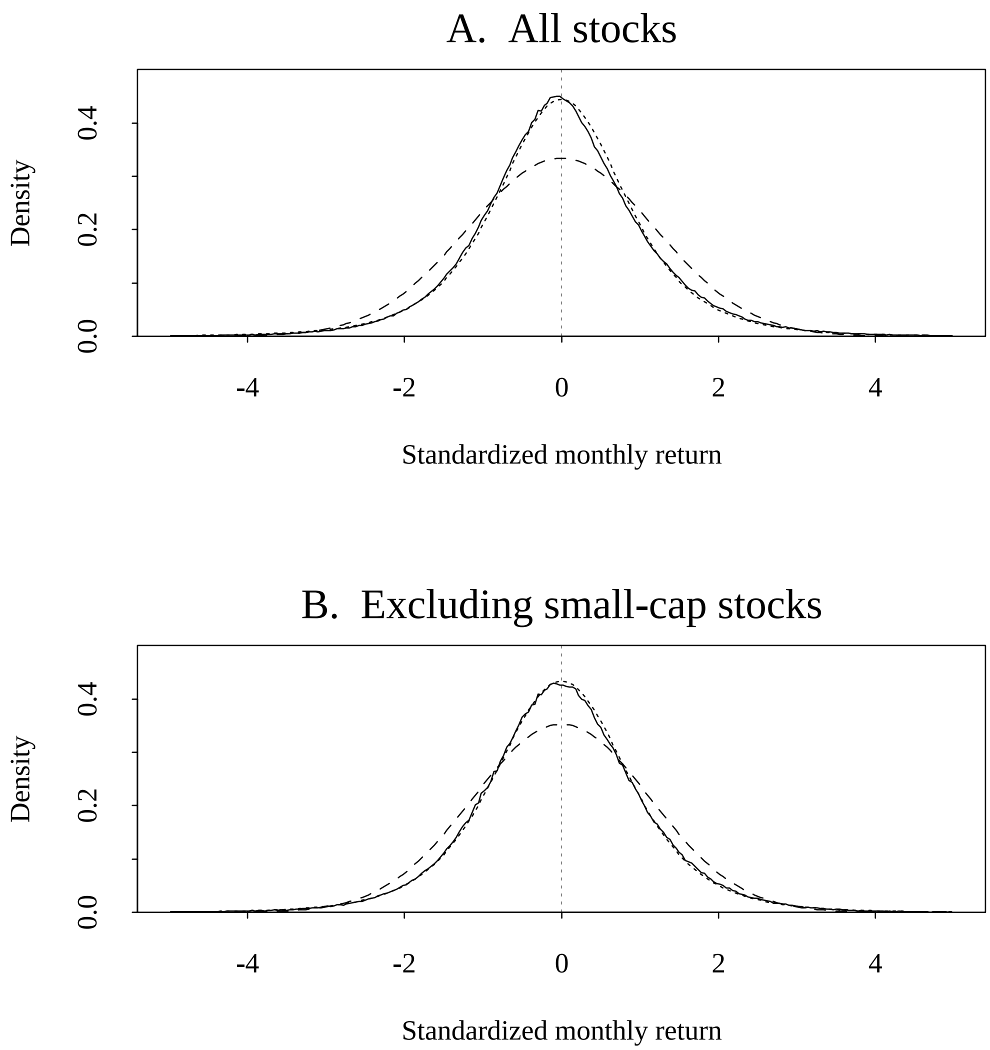

Figure 1: The density of standardized monthly idiosyncratic stock returns

Idiosyncratic monthly returns are constructed for each stock and month, then standardized by estimates of their standard deviations. Empirical densities of the returns are calculated across all stocks and months (top panel) and across all stocks with market capitalizations at or above the 20th percentile of NYSE-listed stocks (bottom panel). The dashed and dotted lines are densities for the normal and t distributions, respectively. They are adjusted to have the same standard deviations as the empirical distributions. The t distribution in the upper panel has four degrees of freedom and the t distribution in the bottom panel has five degrees of freedom, adjusted to have the same standard deviations as the empirical distributions.

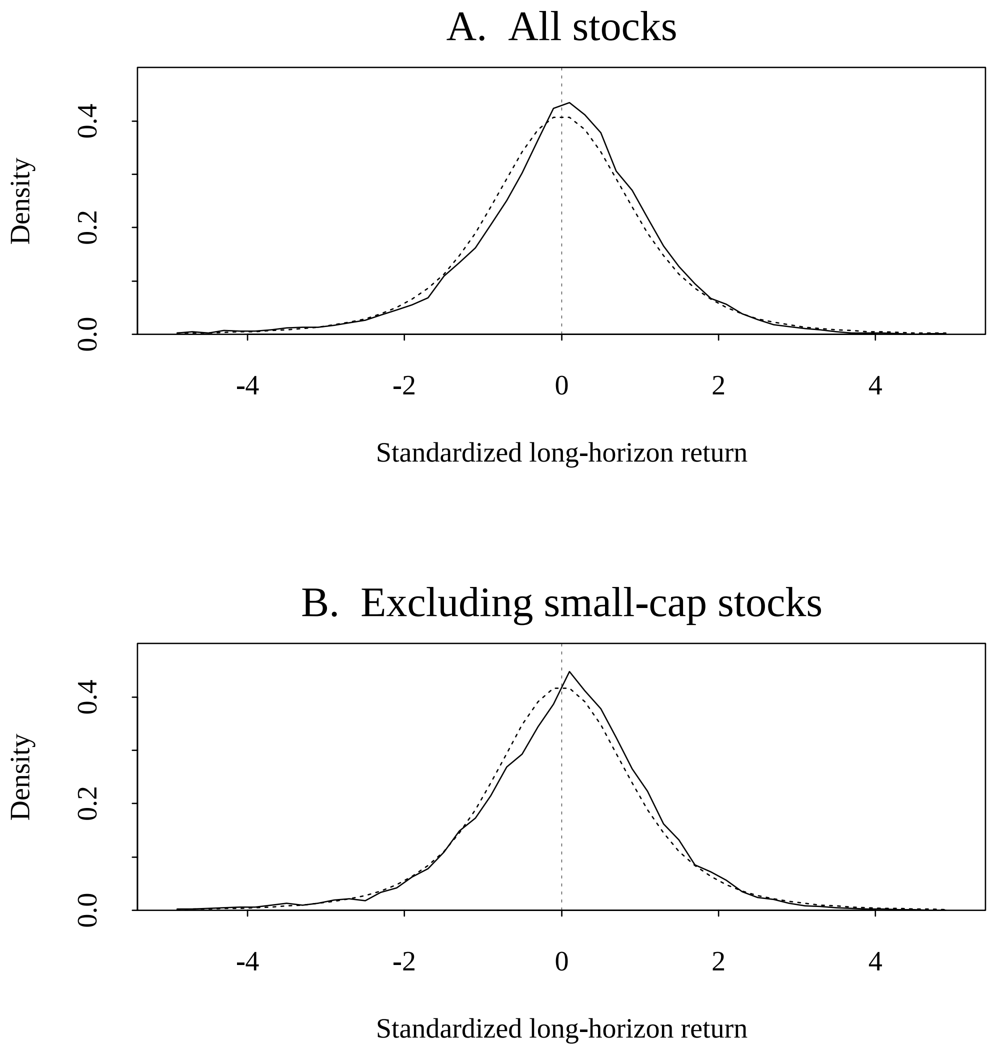

Figure 2: The density of long-horizon idiosyncratic stock returns, standardized by a five-year horizon

Idiosyncratic monthly returns are constructed for each stock and month. The long-horizon return for a given stock and month is the sum of the next five years of monthly idiosyncratic returns unless any monthly stock returns in this period are missing (typically because of delisting). In this case, the horizon is the time to the first missing observation. Long-horizon returns are standardized by the conditional standard deviation of the first month's return times the square root of sixty. The solid lines display the empirical densities. The bottom panel excludes returns of stocks with market capitalizations below the 20th percentile of NYSE-listed stocks. The dotted lines are densities of t distributions with four degrees of freedom, adjusted to have the same standard deviations as the empirical distributions.

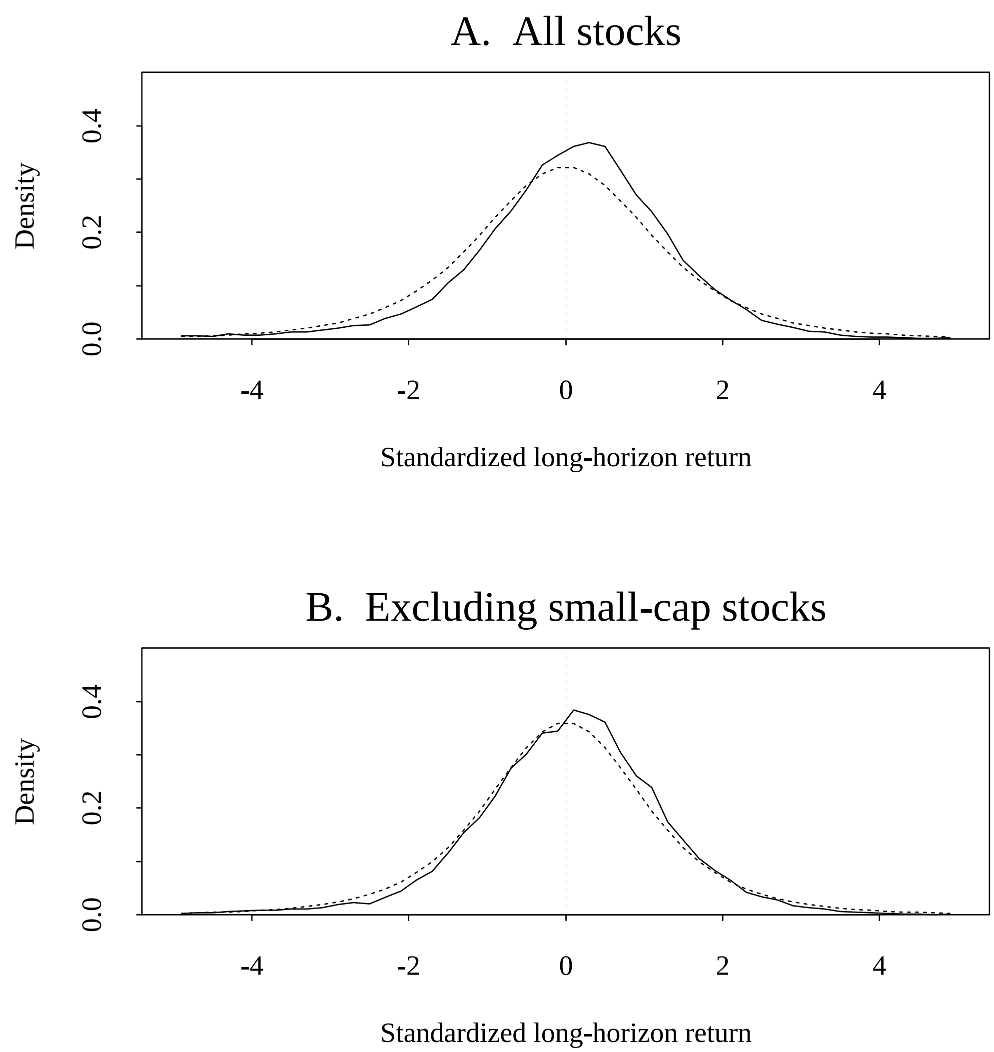

Figure 3: The density of long-horizon idiosyncratic stock returns, standardized by their actual horizons

Idiosyncratic monthly returns are constructed for each stock and month. The long-horizon return for a given stock and month is the sum of the next five years of monthly idiosyncratic returns unless any monthly stock returns in this period are missing (typically because of delisting). In this case, the horizon is the time to the first missing observation. Long-horizon returns are standardized by the conditional standard deviation of the first month's return times the square root of the horizon of the return. The solid lines display the empirical densities. The bottom panel excludes returns of stocks with market capitalizations below the 20th percentile of NYSE-listed stocks. The dotted lines are densities of t distributions with four degrees of freedom, adjusted to have the same standard deviations as the empirical distributions.

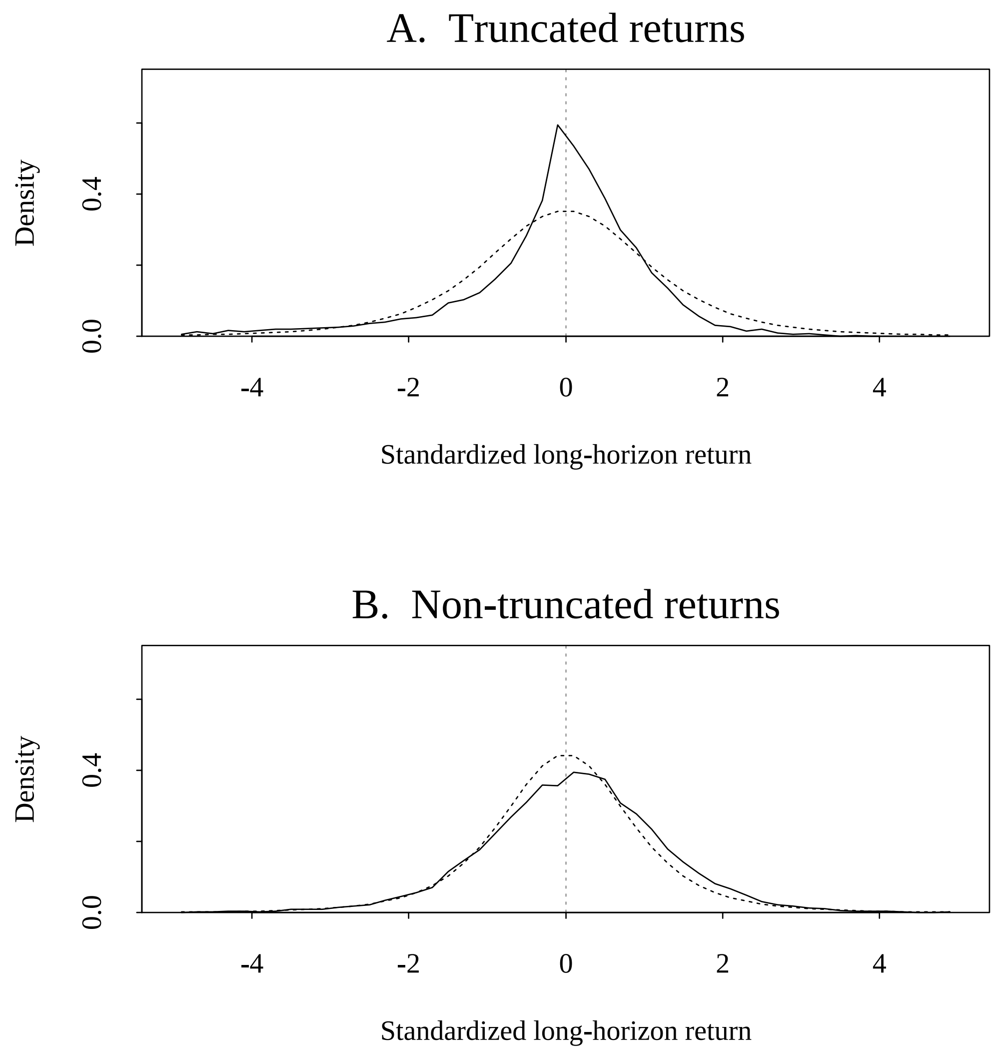

Figure 4: The density of long-horizon idiosyncratic stock returns, standardized by a five-year horizon

Idiosyncratic monthly returns are constructed for each stock and month. The long-horizon return for a given stock and month is the sum of the next five years of monthly idiosyncratic returns unless any monthly stock returns in this period are missing (typically because of delisting). In this case, the horizon is the time to the first missing observation and the return is called a "truncated" return. Long-horizon returns are standardized by the conditional standard deviation of the first month's return times the square root of sixty. The top panel reports the distribution of truncated returns and the bottom panel reports the distribution of non-truncated returns. The solid lines display the empirical densities. The dotted lines are densities of t distributions with four degrees of freedom, adjusted to have the same standard deviations as the empirical distributions.

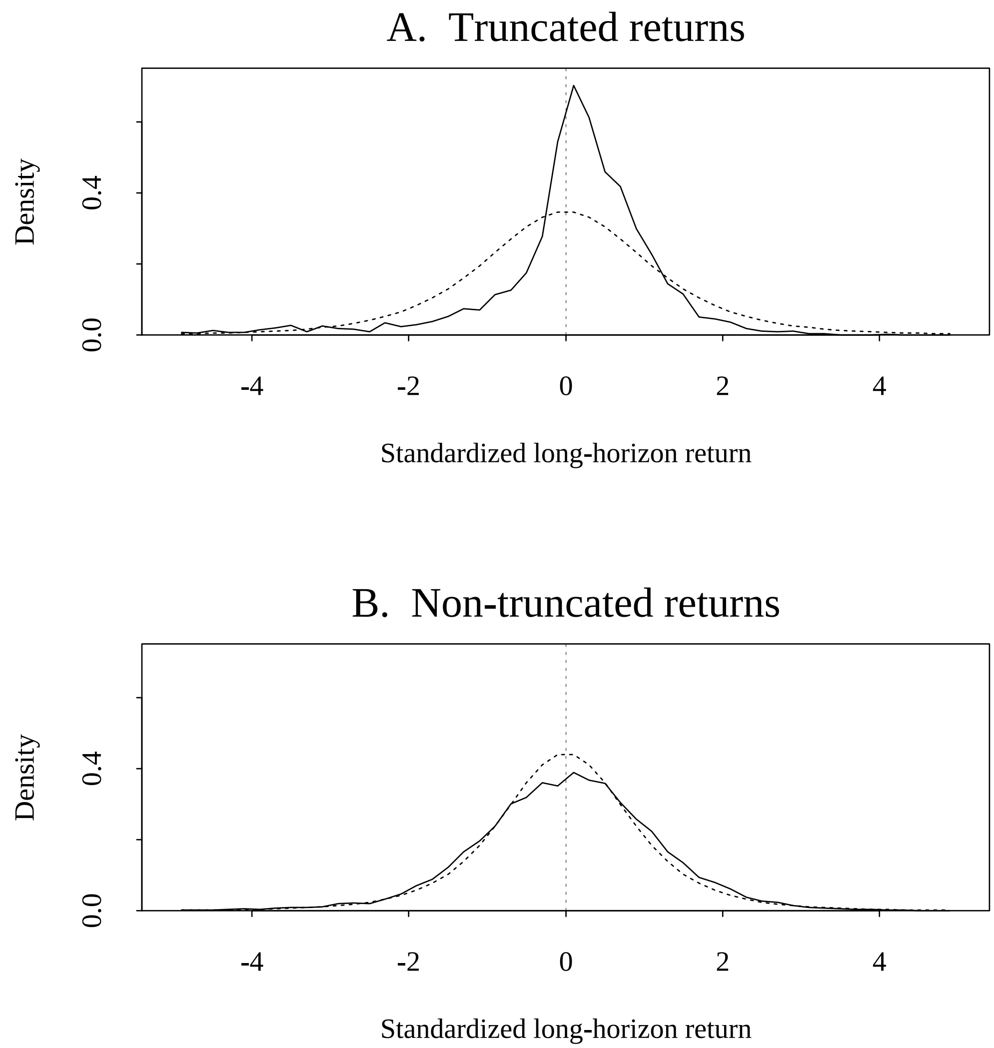

Figure 5: The density of long-horizon idiosyncratic stock returns to large-cap stocks, standardized by a five-year horizon

Idiosyncratic monthly returns are constructed for each stock and month. The long-horizon return for a given stock and month is the sum of the next five years of monthly idiosyncratic returns unless any monthly stock returns in this period are missing (typically because of delisting). In this case, the horizon is the time to the first missing observation and the return is called a "truncated" return. Long-horizon returns are standardized by the conditional standard deviation of the first month's return times the square root of sixty. The top panel reports the distribution of truncated returns and the bottom panel reports the distribution of non-truncated returns. Returns of stocks with market capitalizations below the 20th percentile of NYSE-listed stocks are excluded. The solid lines display the empirical densities. The dotted lines are densities of t distributions with four degrees of freedom, adjusted to have the same standard deviations as the empirical distributions.