ELSEVIER Journal of Financial Economics 61 (2001) 345–381

JOURNAL OF Financial ECONOMICS
www.elsevier.com/locate/econbase

Forecasting crashes: trading volume, past returns, and conditional skewness in stock prices ☆

Joseph Chen[a], Harrison Hong[a], Jeremy C. Stein[b,*]

[a] *Graduate School of Business, Stanford University, Stanford, CA 94305, USA*
[b] *Department of Economics, Harvard University, Cambridge, MA 02138, USA*

Received 6 January 2000; accepted 17 July 2000

Abstract

We develop a series of cross-sectional regression specifications to forecast skewness in the daily returns of individual stocks. Negative skewness is most pronounced in stocks that have experienced (1) an increase in trading volume relative to trend over the prior six months, consistent with the model of Hong and Stein (NBER Working Paper, 1999), and (2) positive returns over the prior 36 months, which fits with a number of theories, most notably Blanchard and Watson's (Crises in Economic and Financial Structure. Lexington Books, Lexington, MA, 1982, pp. 295–315) rendition of stock-price bubbles. Analogous results also obtain when we attempt to forecast the skewness of the aggregate stock market, though our statistical power in this case is limited. © 2001 Elsevier Science S.A. All rights reserved.

JEL classification: G12; G14

Keywords: Crashes; Trading volume; Skewness

☆ We are grateful to the National Science Foundation for research support, and to John Campbell, Kent Daniel, Ken Froot, Ravi Jagannathan, Phillipe Jorion, Chris Lamoreaux, Ken Singleton, an anonymous referee, and seminar participants at Arizona, Arizona State, Cornell, Harvard Business School, Northwestern, Maryland, Stanford, Texas, the UCLA Liquidity Conference, and the NBER for helpful comments and suggestions. Thanks also to Jun Pan for generously sharing her option-pricing software with us.

*Corresponding author.

E-mail address: jeremy_stein@harvard.edu (J.C. Stein).

PII: S0304-405X(01)00066-6

1. Introduction

Aggregate stock market returns are asymmetrically distributed. This asymmetry can be measured in several ways. First, and most simply, the very largest movements in the market are usually decreases, rather than increases – that is, the stock market is more prone to melt down than to melt up. For example, of the ten biggest one-day movements in the S&P 500 since 1947, nine were declines.[1] Second, a large literature documents that market returns exhibit negative skewness, or a closely related property, "asymmetric volatility" – a tendency for volatility to go up with negative returns.[2] Finally, since the crash of October 1987, the prices of stock index options have been strongly indicative of a negative asymmetry in returns, with the implied volatilities of out-of-the-money puts far exceeding those of out-of-the-money calls; this pattern has come to be known as the "smirk" in index-implied volatilities. (See, e.g., Bates, 1997; Bakshi et al., 1997; and Dumas et al., 1998.)

While the existence of negative asymmetries in market returns is generally not disputed, it is less clear what underlying economic mechanism these asymmetries reflect. Perhaps the most venerable theory is based on leverage effects (Black, 1976; Christie, 1982), whereby a drop in prices raises operating and financial leverage, and hence the volatility of subsequent returns. However, it appears that leverage effects are not of sufficient quantitative importance to explain the data (Schwert, 1989; Bekaert and Wu, 2000). This is especially true if one is interested in asymmetries at a relatively high frequency, e.g., in daily data. To explain these, one has to argue that *intraday* changes in leverage have a large impact on volatility – that a drop in prices on Monday morning leads to a large increase in leverage and hence in volatility by Monday afternoon, so that overall, the return for the full day Monday is negatively skewed.

An alternative theory is based on a "volatility feedback" mechanism. As developed by Pindyck (1984), French et al. (1987), Campbell and Hentschel (1992), and others, the idea is as follows: When a large piece of good news arrives, this signals that market volatility has increased, so the direct positive effect of the good news is partially offset by an increase in the risk premium. On the other hand, when a large piece of bad news arrives, the direct effect and the risk-premium effect now go in the same direction, so the impact of the news is amplified. While the volatility-feedback story is in some ways more attractive

[1] Moreover, the one increase – of 9.10% on October 21, 1987 – was right on the heels of the 20.47% decline on October 19, and arguably represented a correction of the microstructural distortions that arose on that chaotic day, rather than an independent price change.

[2] If, in a discrete-time setting, a negative return in period t raises volatility in period $t+1$ and thereafter, returns measured over multiple periods will be negatively skewed, even if single-period returns are not. The literature on these phenomena includes Pindyck (1984), French et al. (1987), Campbell and Hentschel (1992), Nelson (1991), Engle and Ng (1993), Glosten et al. (1993), Braun et al. (1995), Duffee (1995), Bekaert and Wu (2000), and Wu (2001).

than the leverage-effects story, there are again questions as to whether it has the quantitative kick that is needed to explain the data. The thrust of the critique, first articulated by Poterba and Summers (1986), is that shocks to market volatility are for the most part very short-lived, and hence cannot be expected to have a large impact on risk premiums.

A third explanation for asymmetries in stock market returns comes from stochastic bubble models of the sort pioneered by Blanchard and Watson (1982). The asymmetry here is due to the popping of the bubble – a low-probability event that produces large negative returns.

What the leverage-effects, volatility-feedback, and bubble theories all have in common is that they can be cast in a representative-investor framework. In contrast, a more recent explanation of return asymmetries, Hong and Stein (1999), argues that investor heterogeneity is central to the phenomenon. The Hong-Stein model rests on two key assumptions: (1) there are differences of opinion among investors as to fundamental value, and (2) some – though not all – investors face short-sales constraints. The constrained investors can be thought of as mutual funds, whose charters typically prohibit them from taking short positions; the unconstrained investors can be thought of as hedge funds or other arbitrageurs.

When differences of opinion are initially large, those bearish investors who are subject to the short-sales constraint will be forced to a corner solution, in which they sell all of their shares and just sit out of the market. As a consequence of being at a corner, their information is not fully incorporated into prices. For example, if the market-clearing price is \$100, and a particular investor is sitting out, it must be that his valuation is less than \$100, but one has no way of knowing by how much – it could be \$95, but it could also be much lower, say \$50.

However, if after this information is hidden, other, previously more-bullish investors have a change of heart and bail out of the market, the originally more-bearish group may become the marginal "support buyers" and hence more will be learned about their signals. In particular, if the investor who was sitting out at a price of \$100 jumps in and buys at \$95, this is good news relative to continuing to sit on the sidelines even as the price drops further. Thus, accumulated hidden information tends to come out during market declines, which is another way of saying that returns are negatively skewed.

With its focus on differences of opinion, the Hong-Stein model has distinctive empirical implications that are not shared by the representative-investor theories. In particular, the Hong-Stein model predicts that negative skewness in returns will be most pronounced around periods of heavy trading volume. This is because – like in many models with differences of opinion – trading volume proxies for the intensity of disagreement. (See Varian, 1989; Harris and Raviv, 1993; Kandel and Pearson, 1995; and Odean, 1998a for other models with this feature.)

When disagreement (and hence trading volume) is high, it is more likely that bearish investors will wind up at a corner, with their information incompletely revealed in prices. And it is precisely this hiding of information that sets the stage for negative skewness in subsequent rounds of trade, when the arrival of bad news to other, previously more-bullish investors can force the hidden information to come out.

In this paper, we undertake an empirical investigation that is motivated by this differences-of-opinion theory. We develop a series of cross-sectional regression specifications that attempt to forecast skewness in the daily returns to individual stocks. Thus, when we speak of "forecasting crashes" in the title of the paper, we are adopting a narrow and euphemistic definition of the word "crashes," associating it solely with the conditional skewness of the return distribution; we are not in the business of forecasting negative expected returns. This usage follows Bates (1991, 1997), who also interprets conditional skewness – in his case, inferred from options prices – as a measure of crash expectations.

One of our key forecasting variables is the recent deviation of turnover from its trend. For example, at the firm level, we ask whether the skewness in daily returns measured over a given six-month period (say, July 1–December 31, 1998) can be predicted from the detrended level of turnover over the prior six-month period (January 1–June 30, 1998). It turns out that firms that experience larger increases in turnover relative to trend are indeed predicted to have more negative skewness; moreover, the effect of turnover is strongly statistically and economically significant.

In an effort to isolate the effects of turnover, our specifications also include a number of control variables. These control variables can be divided into two categories. In the first category are those that, like detrended turnover, capture time-varying influences on skewness. The most significant variable in this category is past returns. We find that when past returns have been high, skewness is forecasted to become more negative. The predictive power is strongest for returns in the prior six months, but there is some ability to predict negative skewness based on returns as far back as 36 months. In a similar vein, glamour stocks – those with low ratios of book value to market value – are also forecasted to have more negative skewness. (Harvey and Siddique (2000) also examine how skewness varies with past returns and book-to-market.) These results can be rationalized in a number of ways, but they are perhaps most clearly suggested by models of stochastic bubbles. In the context of a bubble model, high past returns or a low book-to-market value imply that the bubble has been building up for a long time, so that there is a larger drop when it pops and prices fall back to fundamentals.

The second category of variables that help to explain skewness are those that appear to be picking up relatively fixed firm characteristics. For example, it has been documented by others (e.g., Damodaran, 1987; Harvey and Siddique, 2000) that skewness is more negative on average for large-cap firms – a pattern that also shows up strongly in our multivariate regressions. We are not aware of any theories that would have naturally led one to anticipate this finding. Rather, for our purposes a variable like size is best thought of as an atheoretic control – it is included in our regressions to help ensure that we do not mistakenly attribute explanatory power to turnover when it is actually proxying for some other firm characteristic. Such a control might be redundant to the extent that detrending the turnover variable already removes firm effects, but we keep it in to be safe.

In addition to running our cross-sectional regressions with the individual-firm data, we also experiment briefly with analogous time-series regressions for the U.S. stock market as a whole. Here, we attempt to forecast the skewness in the daily returns to the market using detrended market turnover and past market returns. Obviously, this pure time-series approach entails an enormous loss in statistical power – with data going back to 1962, we have less than 70 independent observations of market skewness measured at six-month intervals – which is why it is not the main focus of our analysis. Nevertheless, it is comforting to note that the qualitative results from the aggregate-market regressions closely parallel those from the cross-sectional regressions in that high values of both detrended turnover and past returns also forecast more negative market skewness. The coefficient estimates continue to imply economically meaningful effects, although that for detrended turnover is no longer statistically significant.

While both the cross-sectional and time-series results for turnover are broadly consistent with the theory we are interested in, we should stress that we do not at this point view them as a tight test. There are several reasons why one might wish to remain skeptical. First, beyond the effects of turnover, we document other strong influences on skewness, such as firm size, that are not easily rationalized within the context of the Hong-Stein model, and for which there are no other widely accepted explanations. Second, even if innovations to trading volume proxy for the intensity of disagreement among investors, they likely capture other factors as well – such as changes in trading costs – that we have not adequately controlled for. Finally, and most generally, our efforts to model the determinants of conditional skewness at the firm level are really quite exploratory in nature. Given how early it is in this game, we are naturally reluctant to declare an unqualified victory for any one theory.

The remainder of the paper is organized as follows. In Section 2, we review in more detail the theoretical work that motivates our empirical specification. In Section 3, we discuss our sample and the construction of our key variables. In Section 4, we present our baseline cross-sectional regressions, along with a variety of sensitivities and sample splits. In Section 5, we consider the analogous time-series regressions, in which we attempt to forecast the skewness in aggregate-market returns. In Section 6, we use an option-pricing metric to evaluate the economic significance of our results. Section 7 concludes.

2. Theoretical background

The model of Hong and Stein (1999), which provides the principal motivation for our empirical tests, begins with the assumption that there are two investors, A and B, each of whom receives a private signal about a stock's terminal payoff. As a matter of objective reality, each investor's signal contains some useful information. However, each of the two investors only pays attention to their own signal, even if that of the other investor is revealed to them. This deviation from full Bayesian rationality – which can be thought of as a form of overconfidence – leads to irreducible differences of opinion about the stock's value.

In addition to investors A and B, the model also incorporates a class of fully rational, risk-neutral arbitrageurs. These arbitrageurs recognize that the best estimate of the stock's true value is formed by averaging the signals of A and B. However, the arbitrageurs may not always get to see both of the signals, because A and B face short-sales constraints. Importantly, the arbitrageurs themselves are not short-sales constrained, so they can take infinitely large positive or negative positions. Perhaps the most natural interpretation of these assumptions is not to take the short-sales constraint literally – as an absolute technological impediment to trade – but rather to think of investors A and B as institutions like equity mutual funds, many of whom are precluded by their charters or operating policies from ever taking short positions.[3] In contrast, the arbitrageurs might be thought of as hedge funds who are not subject to such restrictions.

Even though investors A and B can be said to suffer from behavioral biases (i.e., overconfidence), the market as a whole is efficient, in the sense of there being no predictability in returns. This is because of the presence of the risk-neutral, unconstrained arbitrageurs. Hence, unlike most of the behavioral finance literature, which relies on limited arbitrage, the model's only implications are for the higher-order moments of the return distribution.

There are two trading dates. To see how the model can generate asymmetries, imagine that at time 1, investor B gets a pessimistic signal, so that B's valuation for the stock lies well below A's. Because of the short-sales constraint, B will simply sit out of the market, and the only trade will be between investor A and the arbitrageurs. The arbitrageurs are rational enough to figure out that B's signal is below A's, but they cannot know by how much.

[3] In fact, Almazan et al. (1999) document that roughly 70% of mutual funds explicitly state (in Form N-SAR that they file with the SEC) that they are not permitted to sell short. This is obviously a lower bound on the fraction of funds that never take short positions. Moreover, Koski and Pontiff (1999) find that 79% of equity mutual funds make no use whatsoever of derivatives (either futures or options), suggesting that funds are also not finding synthetic ways to take short positions.

Thus the market price at time 1 impounds A's prior information, but does not fully reflect B's time-1 signal.

Next, move to time 2, and suppose that A gets a new positive signal. In this case, A continues to be the more optimistic of the two, so A's new time-2 signal is incorporated into the price, while B's time-1 signal remains hidden. On the other hand, if A gets a bad signal at time 2, some of B's previously hidden information might come out. This is because as A bails out of the market at time 2, the arbitrageurs learn something by observing if and at what price B steps in and starts being willing to buy. In other words, there is information in how B responds to A's reduced demand for the stock – in whether or not B gets up off the sidelines and provides buying support. Thus more information comes out, and variance is greater, when the stock price is falling at time 2, as opposed to rising. This greater variance on the downside implies that time-2 returns will be negatively skewed.

However, this logic is not sufficient to establish that unconditional returns (i.e., the average across time 1 and time 2) are negatively skewed. There is a countervailing positive-skewness effect at time 1, since the most negative draws of B's signal are the ones that get hidden from the market at this time. When A's and B's priors are sufficiently close to one another, the positive time-1 skewness can actually overwhelm the negative time-2 skewness, so that returns are on average positively skewed. Nevertheless, Hong and Stein show that if the ex ante divergence of opinion (i.e., the difference in priors) between A and B is great enough, the time-2 effect dominates, and unconditional returns are negatively skewed. It is this unconditional skewness feature – driven by the short-sales constraint – that most clearly distinguishes the model of Hong and Stein from other related models in which pent-up information is revealed through the trading process (e.g., Grossman, 1988; Genotte and Leland, 1990; Jacklin et al., 1992; and Romer, 1993). In these other models, returns are on average symmetrically distributed, albeit potentially quite volatile.

Moreover, the ex ante divergence in priors between A and B – which Hong and Stein denote by H – not only governs the extent of negative skewness, it also governs trading volume. In particular, when H is large, trading volume is unusually high at times 1 and 2. This high trading volume is associated with a greater likelihood of B moving to the sidelines at time 1, and subsequently moving off the sidelines at time 2 – precisely the mechanism that generates negative skewness. Thus the comparative statics properties of the model with respect to the parameter H lead to the prediction that increases in trading volume should forecast more negative skewness. This comparative static result holds regardless of whether unconditional skewness (averaged across different values of H) is positive or negative, and it forms the basis for our empirical tests.

In order to isolate this particular theoretical effect, we need to be aware of other potentially confounding factors. For example, it is well known that trading volume is correlated with past returns (Shefrin and Statman, 1985; Lakonishok and Smidt, 1986; Odean, 1998b). And, as noted above, past returns might also help predict skewness, if there are stochastic bubbles of the sort described by Blanchard and Watson (1982).[4] Indeed, just such a pattern has been documented in recent work by Harvey and Siddique (2000). To control for this tendency, all of our regressions include a number of lags of past returns on the right-hand side.

In a similar vein, one might also worry about skewness being correlated with volatility. There are a number of models that can deliver such a correlation; in the volatility-feedback model of Campbell and Hentschel (1992), for example, higher levels of volatility are associated with more negative skewness. To the extent that such an effect is present in our data, we would like to know whether turnover is forecasting skewness directly – as it should, according to the Hong-Stein model – or whether it is really just forecasting volatility, which is in turn correlated with skewness. To address this concern, all of our regressions include some control for volatility, and we experiment with several ways of doing this control.

3. Data

To construct our variables, we begin with data on daily stock prices and monthly trading volume for all NYSE and AMEX firms, from the CRSP daily and monthly stock files. Our sample period begins in July 1962, which is as far back as we can get the trading volume data; because our regressions use many lags, we do not actually begin to forecast returns until December 1965. We do not include NASDAQ firms because we want to have a uniform and accurate measure of trading volume, and the dealer nature of the NASDAQ market is likely to render turnover in its stocks not directly comparable to that of NYSE and AMEX stocks. We also follow convention and exclude ADRs, REITs, closed-end funds, primes, and scores—i.e., stocks that do not have a CRSP share type code of 10 or 11.

For most of our analysis, we further truncate the sample by eliminating the very smallest stocks in the NYSE/AMEX universe – in particular, those with a market capitalization below the 20th percentile NYSE breakpoint. We do so because our goal is to use trading volume as a proxy for differences of opinion. Theoretical models that relate trading volume to differences of opinion typically assume that transactions costs are zero. In reality, variations in transactions costs are likely to be an important driver of trading volume, and

[4] In the model of Coval and Hirshleifer (1998), there is also conditional negative skewness after periods of positive returns, even though unconditional average skewness is zero.

more so for very small stocks. By eliminating the smallest stocks, we hope to raise the ratio of signal (differences of opinion) to noise (transactions costs) in our key explanatory variable. We also report some sensitivities in which the smallest stocks are analyzed separately (see Table 4 below), and as one would expect from this discussion, the coefficients on turnover for this subsample are noticeably smaller.

Our baseline measure of skewness, which we denote NCSKEW, for "negative coefficient of skewness," is calculated by taking the negative of (the sample analog to) the third moment of daily returns, and dividing it by (the sample analog to) the standard deviation of daily returns raised to the third power. Thus, for any stock i over any six-month period t, we have

$$\mathrm{NCSKEW}_{it} = -\left(n(n-1)^{3/2}\sum R_{it}^{3}\right)\bigg/\left((n-1)(n-2)\left(\sum R_{it}^{2}\right)^{3/2}\right), \quad (1)$$

where R_{it} represents the sequence of de-meaned daily returns to stock i during period t, and n is the number of observations on daily returns during the period. In calculating NCSKEW, as well as any other moments that rely on daily return data, we drop any firm that has more than five missing observations on daily returns in a given period. These daily "returns" are, more precisely, actually log changes in price. We use log changes as opposed to simple daily percentage returns because they allow for a natural benchmark – if stock returns were lognormally distributed, then an NCSKEW measure based on log changes should have a mean of zero. We have also redone everything with an NCSKEW measure based instead on simple daily percentage returns, and none of our main results are affected. Using simple percentage returns instead of log changes does have two (predictable) effects: (1) it makes returns look more positively skewed on average and (2) it induces a pronounced correlation between skewness and contemporaneously measured volatility. However, given that we control for volatility in all of our regression specifications, using simple percentage returns does not materially alter the coefficients on turnover and past returns.

Scaling the raw third moment by the standard deviation cubed allows for comparisons across stocks with different variances; this is the usual normalization for skewness statistics (Greene, 1993). By putting a minus sign in front of the third moment, we are adopting the convention that an increase in NCSKEW corresponds to a stock being more "crash prone" – i.e., having a more left-skewed distribution.

For most of our regressions, the daily firm-level returns that go into the calculation of the NCSKEW variable are market-adjusted returns – the log change in stock i less the log change in the value-weighted CRSP index for that day. However, we also run everything with variations of NCSKEW based on both (1) excess returns (the log change in stock i less the T-bill return) as well as

(2) beta-adjusted returns. As will be seen, these variations do not make much difference to our results with NCSKEW.

In addition to NCSKEW, we also work with a second measure of return asymmetries that does not involve third moments, and hence is less likely to be overly influenced by a handful of extreme days. This alternative measure, which we denote by DUVOL, for "down-to-up volatility," is computed as follows. For any stock i over any six-month period t, we separate all the days with returns below the period mean ("down" days) from those with returns above the period mean ("up" days), and compute the standard deviation for each of these subsamples separately. We then take the log of the ratio of (the sample analog to) the standard deviation on the down days to (the sample analog to) the standard deviation on the up days. Thus we have

$$\mathrm{DUVOL}_{it} = \log\left\{(n_\mathrm{u}-1)\sum_{\mathrm{DOWN}} R_{it}^2 \Big/ \left((n_\mathrm{d}-1)\sum_{\mathrm{UP}} R_{it}^2\right)\right\}, \tag{2}$$

where n_u and n_d are the number of up and down days, respectively. Again, the convention is that a higher value of this measure corresponds to a more left-skewed distribution. To preview, our results with NCSKEW and DUVOL are for the most part quite similar, so it does not appear that they depend on a particular parametric representation of return asymmetries.

In our regressions with firm-level data, we use nonoverlapping six-month observations on skewness. In particular, the NCSKEW and DUVOL measures are calculated using data from either January 1–June 30 or July 1–December 31 of each calendar year. We could alternatively use overlapping data, so that we would have a new skewness measure every month, but there is little payoff to doing so, since, as will become clear shortly, we already have more than enough statistical power as it is. We have, however, checked our results by re-running everything using different nonoverlapping intervals – e.g., February 1–July 31 and August 1–January 31, March 1–August 31 and September 1–February 28, etc. In all cases, the results are essentially identical. When we turn to the time-series regressions with aggregate-market data, statistical power becomes a real issue, and we use overlapping observations.

The choice of a six-month horizon for measuring skewness is admittedly somewhat arbitrary. In principle, the effects that we are interested in could be playing themselves out over a shorter horizon, so that trading volume on Monday forecasts skewness for the rest of the week, but has little predictive power beyond that. Unfortunately, the model of Hong and Stein does not give us much guidance in this regard. Lacking this theoretical guidance, our choice to use six months' worth of daily returns to estimate skewness is driven more by measurement concerns. For example, if we estimated skewness using only one month's worth of data, we would presumably have more measurement error; this is particularly relevant given that a higher-order moment like skewness is strongly influenced by outliers in the data. The important point to note, however, is that to the extent that our measurement horizon does not correspond well to the underlying theory, this should simply blur our ability to find what the theory predicts – i.e., it should make our tests too conservative.

Besides the skewness measures, the other variables that we use are very familiar and do not merit much discussion. SIGMA_{it} is the standard deviation of stock i's daily returns, measured over the six-month period t. RET_{it} is the cumulative return on stock i, also measured over the six-month period t. When we compute NCSKEW or DUVOL using either market-adjusted or beta-adjusted returns, SIGMA and RET are computed using market-adjusted returns. When we compute NCSKEW or DUVOL using excess returns, SIGMA and RET are based on excess returns as well.

LOGSIZE_{it} is the log of firm i's stock market capitalization at the end of period t. BK/MKT_{it} is firm i's book-to-market ratio at the end of period t. LOGCOVER_{it} is the log of one plus the number of analysts (from the I/B/E/S database) covering firm i at the end of period t. TURNOVER_{it} is the average monthly share turnover in stock i, defined as shares traded divided by shares outstanding over period t.

In our baseline specification, we work with detrended turnover, which we denote DTURNOVER. The detrending is done very simply, by subtracting from the TURNOVER variable a moving average of its value over the prior 18 months. Again, the rationale for doing this detrending is that, as a matter of conservatism, we want to eliminate any component of turnover that can be thought of as a relatively fixed firm characteristic. This detrending is roughly analogous to doing a fixed-effects specification in a shorter-lived panel. Since we have such a long time series, it makes little sense to require that firm effects be literally constant over the entire sample period. Instead, the detrending controls for firm characteristics that adjust gradually.

Table 1 presents a variety of summary statistics for our sample. Panel A shows the means and standard deviations of all of our variables for (1) the full sample of individual firms, (2) five size-based subsamples, and (3) the market as a whole, defined as the value-weighted NYSE/AMEX index. (When working with the market as a whole, all the variables are based on simple excess returns relative to T-bills.) Panels B and C look at contemporaneous correlations and autocorrelations, respectively, for the sample of individual firms. In Panels B and C, as in most of our subsequent regression analysis, we restrict the sample to those firms with a market capitalization above the 20th percentile NYSE breakpoint.

One interesting point that emerges from Panel A is that while there is negative skewness – i.e., positive mean values of NCSKEW and DUVOL – for the market as a whole, the opposite is true for individual stocks, which are positively skewed. This discrepancy can in principle be understood within the strict confines of the Hong-Stein model, since, as noted above, the model

Table 1
Summary statistics

The sample period is from July 1962 to December 1998, except for LOGCOVER_t, which is measured starting in December 1976. NCSKEW_t is the negative coefficient of (daily) skewness, measured using market-adjusted returns in the six-month period t. DUVOL_t is the log of the ratio of down-day to up-day standard deviation, measured using market-adjusted returns in the six-month period t. SIGMA_t is the standard deviation of (daily) market-adjusted returns measured in the six-month period t. LOGSIZE_t is the log of market capitalization measured at the end of period t. BK/MKT_t is the most recently available observation of the book-to-market ratio at the end of period t. LOGCOVER_t is the log of one plus the number of analysts covering the stock at the end of period t. DTURNOVER_t is average monthly turnover in the six-month period t, detrended by a moving average of turnover in the prior 18 months. TURNOVER_t is the average monthly turnover measured in the six-month period t. RET_t is the market-adjusted cumulative return in the six-month period t. Size quintiles are determined using NYSE breakpoints.

	All firms	Quintile 5 (largest) firms	Quintile 4 firms	Quintile 3 firms	Quintile 2 firms	Quintile 1 (smallest) firms	Market portfolio
Panel A: First and second moments							
NCSKEW_t							
Mean	−0.262	−0.139	−0.155	−0.198	−0.266	−0.362	0.268
Standard dev.	0.939	0.806	0.904	0.923	0.994	0.964	0.735
DUVOL_t							
Mean	−0.190	−0.128	−0.141	−0.171	−0.213	−0.224	0.172
Standard dev.	0.436	0.364	0.391	0.406	0.437	0.476	0.377

$SIGMA_t$							
Mean	0.025	0.015	0.017	0.020	0.023	0.034	0.008
Standard dev.	0.018	0.005	0.007	0.008	0.010	0.023	0.003
$LOGSIZE_t$							
Mean	5.177	8.249	6.860	5.924	4.984	3.121	N/A
Standard dev.	2.073	1.035	0.653	0.642	0.656	1.108	
BK/MKT_t							
Mean	0.983	0.667	0.782	0.824	0.935	1.275	N/A
Standard dev.	14.036	0.472	0.710	0.870	1.197	22.920	
$LOGCOVER_t$							
Mean	1.991	3.006	2.512	2.030	1.565	1.140	N/A
Standard dev.	0.840	0.431	0.503	0.563	0.564	0.508	
$DTURNOVER_t$							
Mean	0.001	0.000	0.002	0.002	0.001	−0.000	0.002
Standard dev.	0.066	0.039	0.040	0.042	0.046	0.095	0.005
$TURNOVER_t$							
Mean	0.050	0.051	0.056	0.055	0.054	0.043	0.037
Standard dev.	0.075	0.050	0.055	0.060	0.063	0.098	0.022
RET_t							
Mean	0.003	0.024	0.015	0.021	0.017	−0.019	0.029
Standard dev.	0.297	0.164	0.202	0.240	0.288	0.372	0.108
No. of obs.	100,898	13,988	14,291	14,727	16,651	41,241	421

Table 1 (*continued*)

	$NCSKEW_t$	$DUVOL_t$	$SIGMA_t$	$LOGSIZE_t$	BK/MKT_t	$LOGCOVER_t$	$DTURNOVER_t$	$TURNOVER_t$	RET_t
Panel B: Contemporaneous correlations (using only firms above 20th percentile in size)									
$NCSKEW_t$		0.875	0.008	0.038	0.311	0.081	0.007	0.028	−0.302
$DUVOL_t$			−0.076	0.045	0.068	0.100	−0.013	−0.042	−0.371
$SIGMA_t$				−0.307	−0.056	−0.238	0.130	0.398	0.034
$LOGSIZE_t$					−0.213	0.729	0.002	0.101	−0.014
BK/MKT_t						−0.026	0.026	0.006	−0.104
$LOGCOVER_t$							−0.013	0.158	−0.080
$DTURNOVER_t$								0.376	0.133
$TURNOVER_t$									0.061
RET_t									

	$NCSKEW_{t-1}$	$DUVOL_{t-1}$	$SIGMA_{t-1}$	$LOGSIZE_{t-1}$	BK/MKT_{t-1}	$LOGCOVER_{t-1}$	$DTURNOVER_{t-1}$	$TURNOVER_{t-1}$	RET_{t-1}
Panel C: Autocorrelations and cross-correlations (using only firms above 20th percentile in size)									
$NCSKEW_t$	0.047	0.059	−0.047	0.063	−0.030	0.056	0.022	0.032	0.043
$DUVOL_t$	0.061	0.090	−0.109	0.068	−0.011	0.066	0.016	−0.024	0.047
$SIGMA_t$	−0.008	−0.071	0.715	−0.292	−0.050	−0.218	0.042	0.318	−0.014
$LOGSIZE_t$	0.049	0.055	−0.342	0.976	−0.182	0.719	0.000	0.093	−0.011
BK/MKT_t	0.022	0.047	−0.067	−0.181	0.782	−0.027	0.022	0.017	−0.080
$LOGCOVER_t$	0.079	0.098	−0.257	0.736	−0.035	0.852	0.006	0.166	−0.079
$DTURNOVER_t$	−0.028	−0.028	−0.059	0.009	0.039	0.019	0.381	−0.130	0.119
$TURNOVER_t$	0.015	−0.052	0.294	0.104	0.029	0.179	0.195	0.781	0.086
RET_t	−0.002	0.006	−0.032	−0.042	0.051	−0.023	−0.013	−0.064	0.030

allows for either positive or negative unconditional skewness, depending on the degree of ex ante investor heterogeneity. In other words, if one is willing to assume that differences of opinion about the market are on average more pronounced than differences of opinion about individual stocks, the model can produce negative skewness for the latter and positive skewness for the former.

However, it is not clear that such an assumption is empirically defensible. An alternative interpretation of the data in Table 1A is that even if the Hong-Stein model provides a reasonable account of skewness in market returns, it must be missing something when it comes to explaining the mean skewness of individual stocks. For example, it might be that large positive events like hostile takeovers (which the theory ignores) impart an added degree of positive skewness to individual stocks but wash out across the market as a whole. This view does not imply that we cannot learn something about the theory by looking at firm-level data; the theory will certainly gain some credence if it does a good job of explaining cross-sectional variation in skewness, even if it cannot fit the mean skewness at the firm level. Nevertheless, it is worth emphasizing the caveat that, without further embellishments, the theory might not provide a convincing rationale for everything that is going on at the individual stock level.

The most noteworthy fact in Panel B of Table 1 is the contemporaneous correlation between our two skewness measures, NCSKEW and DUVOL, which is 0.88. While these two measures are quite different in their construction, they appear to be picking up much the same information. Also worth pointing out is that the correlation between NCSKEW and SIGMA is less than 0.01, and that between DUVOL and SIGMA is about −0.08; these low correlations lend some preliminary (and comforting) support to the notion that forecasting either of our skewness measures is a quite distinct exercise from forecasting volatility. Panel C documents that, unlike SIGMA – which has an autocorrelation coefficient of 0.72 – neither of our skewness measures has much persistence. For NCSKEW the autocorrelation is on the order of 0.05; for DUVOL it is 0.09.

4. Forecasting skewness in the cross-section

4.1. Baseline specification

Table 2 presents our baseline cross-sectional regression specification. We pool all the data (excluding firms with market capitalization below the 20th percentile NYSE breakpoint) and regress NCSKEW_{it+1} against its own lagged value, NCSKEW_{it}, as well as SIGMA_{it}, LOGSIZE_{it}, DTURNOVER_{it}, and six lags of past returns, $\text{RET}_{it} \ldots \text{RET}_{it-5}$. We also include dummy variables

Table 2
Forecasting skewness in the cross-section: pooled regressions

The sample period is from July 1962 to December 1998 and includes only those firms with market capitalization above the 20th percentile breakpoint of NYSE. The dependent variable is $NCSKEW_{t+1}$, the negative coefficient of (daily) skewness in the six-month period $t+1$. $NCSKEW_{t+1}$ is computed based on market-adjusted returns, beta-adjusted returns and simple excess returns in cols. 1–3, respectively. $SIGMA_t$ is the (daily) standard deviation of returns in the six-month period t. $LOGSIZE_t$ is the log of market capitalization at the end of period t. $DTURNOVER_t$ is average monthly turnover in the six-month period t, detrended by a moving average of turnover in the prior 18 months. $RET_t \ldots RET_{t-5}$ are returns in the six-month periods t through $t-5$ (these past returns are market adjusted in cols. 1−2 and excess in col. 3). All regressions also contain dummies for each time period (not shown); t-statistics, which are in parentheses, are adjusted for heteroskedasticity and serial correlation.

	(1) Base case: market-adjusted returns	(2) Beta-adjusted returns	(3) Excess returns
$NCSKEW_t$	0.053	0.051	0.052
	(7.778)	(7.441)	(7.920)
$SIGMA_t$	−4.566	−3.370	−2.701
	(−7.180)	(−5.242)	(−4.706)
$LOGSIZE_t$	0.037	0.046	0.059
	(11.129)	(13.465)	(19.110)
$DTURNOVER_t$	0.437	0.364	0.364
	(3.839)	(3.175)	(3.329)
RET_t	0.218	0.197	0.221
	(10.701)	(9.638)	(11.607)
RET_{t-1}	0.082	0.082	0.109
	(4.296)	(4.220)	(6.175)
RET_{t-2}	0.103	0.108	0.089
	(5.497)	(5.675)	(5.149)
RET_{t-3}	0.054	0.067	0.053
	(2.830)	(3.462)	(3.001)
RET_{t-4}	0.062	0.058	0.041
	(3.403)	(3.133)	(2.477)
RET_{t-5}	0.071	0.083	0.092
	(3.759)	(4.335)	(5.257)
No. of obs.	51,426	51,426	51,426
R^2	0.030	0.031	0.082

for each time period t. The regression can be interpreted as an effort to predict – based on information available at the end of period t – cross-sectional variation in skewness over period $t+1$.

In column 1, we use market-adjusted returns as the basis for computing the NCSKEW measure. In column 2 we use beta-adjusted returns, and in column 3 we use simple excess returns. The results are quite similar in all three cases. In particular, the coefficients on detrended turnover are positive and strongly statistically significant in each of the three columns, albeit somewhat larger (by about 20%) in magnitude when market-adjusted returns are used. (We expect lower coefficient estimates when using simple excess returns as compared to market-adjusted returns – after all, DTURNOVER is a firm-specific variable, so it should have more ability to explain skewness in the purely idiosyncratic component of stock returns.) The past return terms are also always positive and strongly significant. Thus stocks that have experienced either a surge in turnover or high past returns are predicted to have more negative skewness – i.e., to become more crash-prone, all else equal. The coefficient on size is also positive, suggesting that negative skewness is more likely in large-cap stocks.

As noted above, the findings for past returns and size run broadly parallel to previous work by Harvey and Siddique (2000). Nevertheless, there are several distinctions between our results and theirs. To begin, ours are couched in a multivariate regression framework, while theirs are based on univariate sorts. But more significantly, our measure of skewness is quite different from theirs, for two reasons. First, we look at daily returns, while they look at monthly returns. Second, we look at individual stocks, while they look at portfolios of stocks. The skewness of a portfolio of stocks is not the same thing as the average skewness of its component stocks, especially if, as Harvey and Siddique (2000) stress, coskewness varies systematically with firm characteristics.

We have done some detailed comparisons to make these latter points explicit. For 25 portfolios sorted on size and book-to-market, we have computed both (1) the skewness of monthly portfolio returns, as in Harvey and Siddique (2000), and (2) the average skewness of daily individual stock returns, a measure analogous to what we use here. We can then ask the following: Across the 25 portfolios, what is the correlation of the two skewness measures? The answer is about 0.22, a relatively low, albeit significantly positive, correlation. Thus, while it might have been reasonable to conjecture – based on the prior evidence in Harvey and Siddique (2000) – that our firm-level NCSKEW variable would also be related to past returns and size, such results were by no means a foregone conclusion.

As we have already stressed, the positive coefficient on size is not something one would have necessarily predicted ex ante based on the Hong-Stein model. Nevertheless, it is possible to come up with rationalizations after the fact. Suppose that managers can to some extent control the rate at which information about their firms gets out. It seems plausible that if they uncover good news, they will disclose all this good news right away. In contrast, if they are sitting on bad news, they may try to delay its release, with the result that the bad news dribbles out slowly. This behavior will tend to impart positive skewness to firm-level returns, and may explain why returns on individual stocks are on average positively skewed at the same time that market returns are negatively skewed. Moreover, if one adds the further assumption that it is easier for managers of small firms to temporarily hide bad news – since they face less scrutiny from outside analysts than do managers of large firms – the resulting positive skewness will be more pronounced for small firms. We return to this idea in Section 4.5 below, and use it to develop some additional testable implications.

4.2. Robustness

In Table 3 we conduct a number of further robustness checks. Everything is a variation on column 1 of Table 2, and uses an NCSKEW measure based on market-adjusted returns. First, in column 1 of Table 3, we truncate outliers of the NCSKEW variable, setting all observations that are more than three standard deviations from the mean in any period t to the three-standard-deviation tail values in that period. This has little impact on the results, suggesting that they are not driven by a handful of outlier observations.

In column 2, we replace the DTURNOVER variable with its un-detrended analog, TURNOVER. This means that we are now admitting into consideration differences in turnover across firms that are not merely temporary deviations from trend but more in the nature of long-run firm characteristics. In other words, we are essentially removing our fixed-effect control from the turnover variable. According to the theory, one might expect that long-run cross-firm variation in turnover would also predict skewness – some firms might be subject to persistently large differences in investor opinion, and these too should matter for return asymmetries. The coefficient estimate on TURNVOVER in column 2 confirms this notion, roughly doubling in magnitude from its base-case value. This implies that our fixed-effect approach of using DTURNOVER instead of TURNOVER everywhere else in the paper is quite conservative – in doing so, we are throwing out a dimension of the data that is strongly supportive of the theory.

In columns 3 and 4, we investigate whether our results are somehow tied to the way that we have controlled for volatility. Recall that the central issue here is whether $\mathrm{DTURNOVER}_{it}$ is really forecasting NCSKEW_{it+1} directly, or whether it is instead forecasting SIGMA_{it+1}, and showing up in the regression only because SIGMA_{it+1} is correlated with NCSKEW_{it+1}. Ideally, we would like to add a period-t control variable to the regression that is a good forecast of SIGMA_{it+1}, so that we can verify that $\mathrm{DTURNOVER}_{it}$ is still significant even after the inclusion of this control. Our use of SIGMA_{it} in the base-case specification can be motivated on the grounds that it is probably the best univariate predictor of $SIGMA_{it+1}$, given the very pronounced serial correlation in the SIGMA variable.

But using just one past lag is not necessarily the best way to forecast $SIGMA_{it+1}$. One can presumably do better by allowing for richer dynamics. In this spirit, we add in column 3 two further lags of SIGMA ($SIGMA_{it-1}$ and $SIGMA_{it-2}$) to the base-case specification. These two lags are insignificant, and hence our coefficients on $DTURNOVER_{it}$ as well as on the six RET terms are virtually unchanged. Motivated by the work of Glosten et al. (1993), who find that the effect of past volatility on future volatility depends on the sign of the past return, we also experiment with allowing two coefficients on $SIGMA_{it}$, one for positive past returns and one for negative past returns. This variation (not shown in the table) makes no difference to the results.

In column 4 we take our logic one step further. We create a fitted value of $SIGMA_{it+1}$, which we denote by $SIGMAHAT_{it+1}$, based on the following information set available in period t: $SIGMA_{it}$, $SIGMA_{it-1}$, $SIGMA_{it-2}$, $LOGSIZE_{it}$, $DTURNOVER_{it}$, and RET_{it} ... RET_{it-5}. We then replace $SIGMA_{it}$ in the base case with this fitted value of future volatility, $SIGMAHAT_{it+1}$. This is equivalent to an instrumental-variables regression in which future volatility $SIGMA_{it+1}$ is included on the right-hand side, but is instrumented using the information available in period t. As can be seen, this variation leads to almost exactly the same results as in the base case.

Overall, based on the evidence in columns 3 and 4 of Table 3, we conclude that it is highly unlikely that our base-case success in forecasting NCSKEW with the DTURNOVER and RET variables arises because these variables are able to forecast SIGMA. In other words, these variables really appear to be predicting cross-firm differences in the *asymmetry* of stock returns, rather than just differences in volatility.

In column 5, we add the book-to-market ratio, BK/MKT, to the base-case specification. This variable attracts a significant negative coefficient, which means that it tells the same story as the past-return terms: glamour stocks, like those with high past returns, are more crash-prone. However, the addition of BK/MKT has no impact on the DTURNOVER coefficient.

In column 6, we use the DUVOL measure of return asymmetry as the left-hand-side variable in place of NCSKEW. Although the difference in units precludes a direct comparison of the point estimates, the qualitative patterns are generally the same as in the corresponding specification in column 1 of Table 2. Indeed, the t-statistic on DTURNOVER is actually a bit higher (4.35 vs. 3.84) as is the R^2 of the regression (6.7% vs. 3.0%).

Finally, in an unreported sensitivity, we check to make sure our results are robust to how we have modeled the effect of the lagged skewness variable, $NCSKEW_{it}$. Instead of estimating just one coefficient on $NCSKEW_{it}$, we allow this effect to be a function of the realization of $NCSKEW_{it}$ itself. We implement this by interacting $NCSKEW_{it}$ with five dummy variables,

Table 3
Forecasting skewness in the cross-section: robustness checks

The sample period is from July 1962 to December 1998 and includes only those firms with market capitalization above the 20th percentile breakpoint of NYSE. In columns 1–5, the dependent variable is $NCSKEW_{t+1}$, the negative coefficient of (daily) skewness in the six-month period $t+1$. In column 6, the dependent variable is $DUVOL_{t+1}$, the log of the ratio of down-day to up-day standard deviation in the six-month period $t+1$. In all columns, returns are market-adjusted. $SIGMA_t$ is the standard deviation of (daily) returns in the six-month period t. $LOGSIZE_t$ is the log of market capitalization at the end of period t. BK/MKT_t is the most recently available observation of the book-to-market ratio at the end of period t. $DTURNOVER_t$ is average monthly turnover in the six-month period t, detrended by a moving average of turnover in the prior 18 months, except in column 3, where turnover is not detrended. $RET_t \dots RET_{t-5}$ are returns in the six-month periods t through $t-5$. $SIGMAHAT_{t+1}$ is the predicted value of $SIGMA_{t+1}$ calculated from a regression of $SIGMA_{t+1}$ on $SIGMA_t, \dots,$ $SIGMA_{t-2}$, $LOGSIZE_t$, $DTURNOVER_t$, and $RET_t \dots RET_{t-5}$. All regressions also contain dummies for each time period (not shown); t-statistics, which are in parentheses, are adjusted for heteroskedasticity and serial correlation.

	(1) Outliers truncated	(2) Turnover not detrended	(3) More lags of past volatility	(4) Fitted future volatility	(5) Book-to-market	(6) Using $DUVOL_{t+1}$
$NCSKEW_t$	0.050	0.053	0.053	0.051	0.054	0.096
($DUVOL_t$ in col. 6)	(8.675)	(7.837)	(7.663)	(7.454)	(7.750)	(16.627)
$SIGMAHAT_{t+1}$				−6.178		
				(−7.180)		
$SIGMA_t$	−4.994	−6.618	−3.953		−4.999	−4.956
	(−8.938)	(−9.822)	(−3.751)		(−7.552)	(−15.698)
$SIGMA_{t-1}$			−0.460			
			(−0.384)			
$SIGMA_{t-2}$			−0.367			
			(−0.353)			
$LOGSIZE_t$	0.035	0.033	0.037	0.034	0.035	0.014
	(12.047)	(9.980)	(10.898)	(9.351)	(10.095)	(9.572)
BK/MKT_t					−0.020	
					(−3.808)	

$DTURNOVER_t$	0.375	0.761	0.411	0.387	0.455	0.202
($TURNOVER_t$ in col. 2)	(3.729)	(7.685)	(3.459)	(3.410)	(3.848)	(4.346)
RET_t	0.206	0.217	0.218	0.208	0.213	0.142
	(11.787)	(10.887)	(10.761)	(10.249)	(10.071)	(15.810)
RET_{t-1}	0.075	0.071	0.083	0.084	0.081	0.014
	(4.587)	(3.828)	(4.329)	(4.428)	(4.087)	(1.671)
RET_{t-2}	0.100	0.088	0.104	0.106	0.098	0.045
	(6.273)	(4.734)	(5.472)	(5.621)	(5.038)	(5.587)
RET_{t-3}	0.049	0.033	0.054	0.056	0.056	0.009
	(3.030)	(1.727)	(2.819)	(2.943)	(2.816)	(1.131)
RET_{t-4}	0.048	0.041	0.060	0.064	0.051	0.014
	(3.084)	(2.287)	(3.337)	(3.523)	(2.722)	(1.808)
RET_{t-5}	0.057	0.054	0.072	0.073	0.066	0.014
	(3.580)	(2.923)	(3.789)	(3.820)	(3.324)	(1.705)
No. of obs.	51,426	52,229	51,393	51,426	48,630	51,426
R^2	0.039	0.031	0.030	0.030	0.030	0.067

one corresponding to each quintile of $NCSKEW_{it}$. In other words, we estimate five separate slope coefficients on lagged skewness, depending on the quintile that lagged skewness falls in. As it turns out, while there appear to be some modest nonlinearities in the effect of lagged skewness, these nonlinearities do not at all impact the coefficients on any of the other variables of interest.

4.3. Cuts on firm size

In Table 4, we disaggregate our base-case analysis by size. We take the specification from column 1 of Table 2 and run it separately for five size-based subsamples, corresponding to quintiles based on NYSE breakpoints. (Recall that in Tables 2 and 3, we omit the smallest of these five quintiles from our sample.) Two conclusions stand out. First, as suspected, the coefficient on DTURNOVER for the smallest category of firms is noticeably lower than for any other group, albeit still positive.[5] Again, this is probably because variation in turnover for these tiny firms is driven in large part by variation in trading costs, whereas our theory requires a good proxy for differences of opinion. Second, once one moves beyond the smallest quintile, the coefficients look reasonably stable. There is certainly no hint that the effects that we are interested in go away for larger firms. Indeed, the highest point estimate for the DTURNOVER coefficient comes from the next-to-largest quintile.

The fact that the coefficients on DTURNOVER are robust for large firms is not surprising in light of the underlying theory. As we have emphasized, the model of Hong and Stein is not predicated on impediments to arbitrage – it incorporates a class of fully risk-neutral arbitrageurs who can take infinite long or short positions. Thus, as long as some investors other than the arbitrageurs (e.g., mutual funds) continue to be short-sales constrained, the model does not have the feature that the key effects diminish as one moves to larger stocks, where arbitrage activity is presumably more efficient. This is in contrast to behavioral models based on limited arbitrage (e.g., DeLong et al., 1990) whose implications for return predictability are often thought of as applying more forcefully to small stocks.[6]

[5] Also for these smallest firms, the coefficient on SIGMA changes signs, and the coefficients on the past-return terms are smaller and much less significant. Given the potential distortions associated with infrequent trading and price discreteness for this group, we are reluctant to hazard an economic interpretation for these anomalies.

[6] Several recent papers find that predictability – based on either "momentum" or "value" strategies – is stronger in small-cap stocks (see, e.g., Fama, 1998; Hong et al., 2000; and Griffin and Lemmon, 1999).

Table 4
Forecasting skewness in the cross-section: cuts by firm size

The sample period is from July 1962 to December 1998. The dependent variable in all columns is $NCSKEW_{t+1}$, the negative coefficient of (daily) skewness in the six-month period $t+1$. In all columns, returns are market-adjusted. $SIGMA_t$ is the standard deviation of (daily) returns in the six-month period t. $LOGSIZE_t$ is the log of market capitalization at the end of period t. $DTURNOVER_t$ is average monthly turnover in the six-month period t, detrended by a moving average of turnover in the prior 18 months. $RET_t \ldots RET_{t-5}$ are returns in the six-month periods t through $t-5$. All regressions also contain dummies for each time period (not shown); t-statistics are adjusted for heteroskedasticity and serial correlation. Firm size cuts are based on NYSE breakpoints.

	Quintile 5 (largest) firms	Quintile 4 firms	Quintile 3 firms	Quintile 2 firms	Quintile 1 (smallest) firms
$NCSKEW_t$	0.053	0.059	0.054	0.043	0.045
	(3.758)	(3.653)	(4.341)	(3.690)	(5.431)
$SIGMA_t$	−3.043	−4.362	−4.409	−4.062	2.894
	(−1.243)	(−2.263)	(−3.771)	(−4.612)	(8.793)
$LOGSIZE_t$	0.009	0.057	0.049	0.105	0.066
	(1.021)	(1.855)	(1.590)	(3.639)	(8.800)
$DTURNOVER_t$	0.404	0.637	0.551	0.264	0.079
	(1.812)	(2.450)	(2.554)	(1.391)	(1.072)
RET_t	0.260	0.335	0.215	0.155	0.010
	(5.637)	(7.000)	(5.359)	(4.682)	(0.569)
RET_{t-1}	0.047	0.001	0.083	0.134	0.017
	(1.009)	(0.024)	(2.157)	(4.269)	(1.076)
RET_{t-2}	0.163	0.165	0.104	0.069	0.014
	(3.554)	(3.726)	(2.651)	(2.298)	(0.816)
RET_{t-3}	0.025	0.078	0.093	0.033	0.028
	(0.535)	(1.682)	(2.334)	(1.112)	(1.823)
RET_{t-4}	0.162	0.101	0.071	0.006	0.014
	(3.637)	(2.540)	(1.852)	(0.215)	(0.864)
RET_{t-5}	0.128	0.089	0.134	0.013	−0.010
	(2.906)	(1.801)	(3.503)	(0.465)	(−0.632)
No. of obs.	12,749	12,520	12,407	13,750	29,165
R^2	0.035	0.030	0.024	0.029	0.028

4.4. Stability over subperiods

In Table 5, we examine the intertemporal stability of our baseline regression, using a Fama-MacBeth (1973) approach. Specifically, we run a separate, purely cross-sectional variant of the regression in column 1 of Table 2 for every one of

Table 5
Forecasting skewness in the cross-section: Fama-MacBeth approach

The sample period is from July 1962 to December 1998 and includes only those firms with market capitalization above the 20th percentile breakpoint of NYSE. The dependent variable is $NCSKEW_{t+1}$, the negative coefficient of (daily) skewness in the six-month period $t+1$. In all cases, returns are market-adjusted. The specification is the same as in col. 1 of Table 2. $SIGMA_t$ is the standard deviation of (daily) returns in the six-month period t. $LOGSIZE_t$ is the log of market capitalization at the end of period t. $DTURNOVER_t$ is average monthly turnover in the six-month period t, detrended by a moving average of turnover in the prior 18 months. $RET_t \ldots RET_{t-5}$ are returns in the six-month periods t through $t-5$. Panel A reports only the coefficient on $DTURNOVER_t$ for each period. Panel B reports the mean coefficients for different subperiods, and the associated t-statistics, based on the time-series standard deviations of the coefficients, and adjusted for serial correlation.

1960s		1970s		1980s		1990s	
Panel A: Period-by-period regressions (12/1965 to 6/1998); coefficient on detrended turnover only							
12/1965	0.383	6/1970	0.129	6/1980	1.730	6/1990	1.780
6/1966	1.053	12/1970	0.973	12/1980	0.707	12/1990	−0.194
12/1966	0.248	6/1971	1.145	6/1981	−0.156	6/1991	1.065
6/1967	−0.081	12/1971	0.269	12/1981	−0.757	12/1991	0.058
12/1967	0.201	6/1972	0.955	6/1982	2.738	6/1992	0.835
6/1968	0.468	12/1972	−0.207	12/1982	0.373	12/1992	0.569
12/1968	1.218	6/1973	0.148	6/1983	2.314	6/1993	0.161
6/1969	1.101	12/1973	−0.904	12/1983	0.334	12/1993	0.803
12/1969	0.498	6/1974	2.257	6/1984	−0.751	6/1994	0.459
		12/1974	0.579	12/1984	0.545	12/1994	0.372
		6/1975	−0.363	6/1985	2.448	6/1995	1.026
		12/1975	−0.083	12/1985	−0.182	12/1995	−0.913
		6/1976	0.029	6/1986	−0.686	6/1996	−0.631
		12/1976	−0.016	12/1986	0.388	12/1996	1.981
		6/1977	0.876	6/1987	0.672	6/1997	0.643
		12/1977	1.901	12/1987	0.464	12/1997	0.062
		6/1978	0.918	6/1988	0.404	6/1998	0.381
		12/1978	1.512	12/1988	−0.941		
		6/1979	1.506	6/1989	0.121		
		12/1979	0.210	12/1989	−0.038		

	All periods	Late 1960s	1970s	1980s	1990s
Panel B: Average coefficients by subperiods					
$NCSKEW_t$	0.063	0.099	0.079	0.064	0.024
	(4.880)	(2.173)	(4.517)	(2.707)	(1.258)
$SIGMA_t$	−5.017	−11.577	−9.507	−3.884	2.407
	(−2.312)	(−2.614)	(−3.063)	(−1.061)	(0.288)
$LOGSIZE_t$	0.030	0.005	0.040	0.027	0.032
	(4.141)	(0.222)	(2.216)	(2.776)	(4.200)

Table 5 (*continued*)

	All periods	Late 1960s	1970s	1980s	1990s
DTURNOVER_t	0.532	0.565	0.592	0.486	0.497
	(3.981)	(2.280)	(2.549)	(1.372)	(2.326)
RET_t	0.249	0.335	0.234	0.229	0.242
	(6.614)	(1.807)	(3.909)	(3.663)	(2.312)
RET_{t-1}	0.099	0.229	0.026	0.085	0.132
	(3.287)	(1.684)	(0.427)	(1.838)	(2.711)
RET_{t-2}	0.139	0.100	0.222	0.132	0.071
	(4.357)	(1.098)	(3.452)	(2.323)	(1.387)
RET_{t-3}	0.082	0.057	0.139	0.017	0.104
	(2.555)	(0.645)	(2.596)	(0.341)	(1.513)
RET_{t-4}	0.081	0.045	0.091	0.044	0.133
	(2.887)	(0.390)	(1.524)	(0.917)	(2.453)
RET_{t-5}	0.082	0.139	0.056	0.036	0.136
	(1.967)	(1.767)	(1.014)	(0.193)	(1.492)
No. of obs.	66	9	20	20	17

the 66 six-month periods in our sample. We then take simple time-averages of the cross-sectional regression coefficients over various subperiods, and compute the associated *t*-statistics based on the time-series properties of the coefficients (and adjusting for serial correlation). In Panel A of Table 5, we display the coefficient on DTURNOVER from every one of the 66 regressions. In Panel B, we show time-averages of all the regression coefficients for the full sample and for each of four decade-based subperiods: the 1960s, the 1970s, the 1980s, and the 1990s.

The overriding conclusion that emerges from Table 5 is that our results are remarkably stable over time. For example, the coefficient on DTURNOVER – which averages 0.532 over the full sample period – reaches a low of 0.486 in the 1980s and a high of 0.592 in the 1970s. Moreover, even taken alone, three of the four decade-based subperiods produce a statistically significant result for DTURNOVER.

4.5. Why are small stocks more positively skewed?

One of the most striking patterns that we have documented is that small stocks are more positively skewed than large stocks. Given that this pattern is not clearly predicted by any existing theories (of which we are aware) we have had to come up with a new hypothesis after the fact in order to rationalize it. As described above, this hypothesis begins with the assumption that managers have some discretion over the disclosure of information, and prefer to announce good news immediately, while allowing bad news to dribble out slowly. This behavior tends to impart a degree of positive skewness to returns. Moreover, if there is more scope for such managerial discretion in small firms – perhaps because they face less scrutiny from security analysts – then the positive-skewness effect will be more pronounced in small stocks.

The one satisfying thing about this after-the-fact hypothesis is that it yields new testable predictions. Specifically, positive skewness ought to be greater in firms with fewer analysts, after controlling for size. Table 6 investigates this prediction, taking our baseline specifications for both NCSKEW and DUVOL, and in each case adding LOGCOVER, the log of one plus the number of analysts covering the stock. (The sample period in Table 6 is shorter, since analyst coverage is not available from I/B/E/S prior to December 1976.) The coefficients on LOGCOVER have the predicted positive sign, and are strongly statistically significant.[7] The coefficients on LOGSIZE go down a bit, but remain significant as well. Nothing else changes.

We do not mean to cast Table 6 as a definitive test of the discretionary-disclosure hypothesis; this idea is outside the main scope of the paper, and pursuing it more seriously would take us too far afield. Nevertheless, it is comforting to know that the most obvious auxiliary prediction of the hypothesis is borne out in the data, and that as a result, we at least have a plausible explanation for what would otherwise be a puzzling feature of our data.

5. Forecasting market skewness

We now turn to forecasting skewness in the returns to the aggregate market. While this is in many ways the more interesting exercise from an economic viewpoint, our statistical power is severely limited. Thus it may be asking too much to expect that the results here will be strongly statistically significant in their own right; rather, one might more reasonably hope that they look qualitatively similar to those from the cross-sectional regressions.

[7] After developing the discretionary-disclosure hypothesis, and running the regressions in Table 6, we became aware of a closely related working paper by Damodaran (1987). Using data from 1979 to 1983, he also finds that firms with fewer analysts have more positively skewed returns.

Table 6
Forecasting skewness in the cross-section: adding analyst coverage

The sample period is from December 1976 to December 1998 and includes only those firms with market capitalization above the 20th percentile breakpoint of the NYSE. The dependent variable in col. 1 is NCSKEW_{t+1}, the negative coefficient of skewness in the six-month period $t+1$, and in col. 2 it is DUVOL_{t+1}, the log of the ratio of down-day to up-day standard deviation in the six-month period $t+1$. SIGMA_t is the standard deviation of (daily) returns in the six-month period t. LOGSIZE_t is the log of market capitalization at the end of period t. LOGCOVER_t is the log of one plus the number of analysts covering the stock at the end of period t. DTURNOVER_t is average monthly turnover in the six-month period t, detrended by a moving average of turnover in the prior 18 months. $\text{RET}_t \ldots \text{RET}_{t-5}$ are returns in the six-month periods t through $t-5$. All regressions also contain dummies for each time period (not shown); t-statistics are adjusted for heteroskedasticity and serial correlation.

	(1) Using NCSKEW_{t+1} measure	(2) Using DUVOL_{t+1} measure
NCSKEW_t	0.049	0.090
(DUVOL_t in col.2)	(6.649)	(13.945)
SIGMA_t	−3.188	−4.022
	(−4.366)	(−11.586)
LOGSIZE_t	0.032	0.011
	(7.992)	(6.665)
DTURNOVER_t	0.504	0.207
	(3.504)	(3.648)
LOGCOVER_t	0.019	0.006
	(4.059)	(3.288)
RET_t	0.219	0.135
	(8.767)	(12.599)
RET_{t-1}	0.085	0.010
	(3.681)	(1.044)
RET_{t-2}	0.100	0.040
	(4.410)	(4.253)
RET_{t-3}	0.058	0.006
	(2.564)	(0.683)
RET_{t-4}	0.065	0.012
	(3.140)	(1.290)
RET_{t-5}	0.055	0.006
	(2.493)	(0.705)
No. of obs.	40,688	40,688
R^2	0.025	0.051

Our definition of the aggregate market is the value-weighted NYSE-AMEX index, and all returns are excess returns relative to T-bills. To avoid any temptation to further mine the data, we use essentially the same specification as in our baseline cross-sectional analysis.[8] Specifically, we use all the same right-hand-side variables, except for LOGSIZE and the time dummies. The DTURNOVER variable is constructed exactly as before, by detrending TURNOVER with its own moving average over the prior 18 months.

In an effort to get the most out of the little time-series data that we have, we now use monthly overlapping observations. (The t-statistics we report are adjusted accordingly.) This yields a total of 401 observations that can be used in the regressions. However, a new concern that arises with the time-series approach is the extent to which our inferences are dominated by the enormous daily movements during October 1987. To address this concern, we also re-run our regressions omitting October 1987. This brings us down to 371 observations.[9]

The results are summarized in Table 7. In columns 1 and 2 we use the NCSKEW measure of skewness, and run the regressions with and without October 1987, respectively. In columns 3 and 4 we use the DUVOL measure of skewness, and again run the regression with and without October 1987. The basic story is the same in all four columns. The six past-return terms are always positive, and many are individually statistically significant. In contrast, the coefficient on DTURNOVER, while positive in each of the four regressions, is never statistically significant. Dropping October 1987 seems to increase the precision of the DTURNOVER coefficient estimate somewhat, but the highest t-statistic across the four specifications is only 1.15.

Nevertheless, holding statistical significance aside, the point estimates suggest large quantitative effects relative to the cross-sectional regressions. Indeed, the coefficients on DTURNOVER and the RET terms are now on the order of ten times bigger than they were in the previous tables. Thus both turnover and past returns could well be very important for forecasting the skewness of market returns, but we lack the statistical power to assert these conclusions – particularly regarding turnover – with much confidence.

In light of this power problem, we obtained an alternative series on NYSE volume going back to 1928 from Gallant et al. (1992), who in turn take it from the S&P *Security Price Record*. The one drawback with these data is that we cannot use them to literally calculate turnover, since they give only the number of shares traded and not the number of shares outstanding. Thus we cannot

[8] Harvey and Siddique (1999) build an autoregressive conditional skewness model for aggregate-market returns; while their specification is very different from that here, it shares the common element that lagged skewness helps to forecast future skewness.

[9] We lose 30 observations because we do not allow any observation on NCSKEW, DUVOL, SIGMA, or DTURNOVER to enter the regression if it draws on data from October 1987. Because of the detrending, the DTURNOVER variable in any given month draws on 24 months' worth of data.

Table 7
Forecasting skewness in the market: time-series regressions

The sample period is from July 1962 to December 1998 and is based on market returns in excess of the risk-free rate, where the market is defined as the value-weighted portfolio of all NYSE/AMEX stocks. The dependent variable in cols. 1 and 2 is $NCSKEW_{t+1}$, the negative coefficient of skewness in the six-month period $t+1$, and in cols. 3 and 4 it is $DUVOL_{t+1}$, the log of the ratio of down-day to up-day standard deviation in the six-month period $t+1$. $SIGMA_t$ is the standard deviation of (daily) market returns in the six-month period t. $DTURNOVER_t$ is the average monthly turnover of the market portfolio in the six-month period t, detrended by a moving average of turnover in the prior 18 months. $RET_t \ldots RET_{t-5}$ are returns in the six-month periods t through $t-5$; t-statistics, which are in parentheses, are adjusted for heteroskedasticity and serial correlation.

	(1) Dep. variable is $NCSKEW_{t+1}$	(2) Dep. variable is $NCSKEW_{t+1}$, excluding 10/87	(3) Dep. variable is $DUVOL_{t+1}$	(4) Dep. variable is $DUVOL_{t+1}$, excluding 10/87
$NCSKEW_t$ ($DUVOL_t$ in col.3 and 4)	0.100	0.123	0.221	0.217
	(0.855)	(1.232)	(1.842)	(0.844)
$SIGMA_t$	18.183	13.708	1.196	−3.574
	(1.137)	(0.749)	(0.156)	(−0.300)
$DTURNOVER_t$	6.002	9.349	6.324	9.462
	(0.262)	(0.828)	(0.704)	(1.148)
RET_t	2.647	1.809	1.484	1.184
	(4.147)	(4.406)	(4.168)	(3.398)
RET_{t-1}	1.585	1.077	0.482	0.332
	(3.086)	(2.939)	(1.481)	(1.061)
RET_{t-2}	1.473	0.926	0.554	0.386
	(2.242)	(1.922)	(1.898)	(1.357)
RET_{t-3}	0.589	0.443	0.126	0.017
	(0.602)	(0.734)	(0.325)	(0.049)
RET_{t-4}	1.283	0.680	0.475	0.287
	(2.264)	(1.575)	(1.726)	(0.968)
RET_{t-5}	1.187	0.596	0.686	0.470
	(2.288)	(1.930)	(2.326)	(1.753)
No. of obs.	401	371	401	371
R^2	0.265	0.264	0.304	0.274

quite reproduce our baseline specifications for the longer post-1928 sample period. Nevertheless, using detrended values of raw trading volume to approximate detrended turnover, we get results for this sample period that are very similar to those reported in Table 7.

6. Economic significance of the results: an option-pricing metric

Thus far, we have focused on the statistical significance of our results, and have not really asked whether they imply magnitudes that are economically meaningful. Assessing economic significance in the current context is a bit tricky. The thought experiment that is typically undertaken is something like this: suppose that the right-hand-side variable of interest – in this case, DTURNOVER – is shocked by two standard deviations. How much does the left-hand-side variable – NCSKEW or DUVOL – move? What makes things difficult here is that most people have little sense for what would constitute an economically interesting change in NCSKEW or DUVOL.

To help frame things in a way that is hopefully more intuitive, we can translate statements about NCSKEW into statements about the prices of out-of-the-money put options. The idea behind our metric can be understood as follows. Imagine that you are pricing an out-of-the-money put on a stock whose returns you initially believe to be symmetrically distributed – i.e., a stock for which you believe that NCSKEW is equal to zero. Now the stock experiences a surge in turnover. As a result, you revise your forecast of NCSKEW, using the DTURNOVER coefficient estimate from our regressions. Given this new forecast of NCSKEW – but holding volatility fixed – by how much does the value of the put option increase?

To answer this sort of question precisely, we need to (1) find an option-pricing model that admits skewness in returns and (2) create a mapping from the parameters of this model to our NCSKEW variable. The model we use is the stochastic-volatility model of Das and Sundaram (1999), in which the dynamics of stock prices are summarized by the following two diffusion equations:

$$\mathrm{d}p_t = \alpha\,\mathrm{d}t + V_t^{1/2}\,\mathrm{d}z_1, \tag{3}$$

$$\mathrm{d}V_t = \kappa(V_0 - V_t)\mathrm{d}t + \eta V_t^{1/2}\,\mathrm{d}z_2, \tag{4}$$

where p_t is the log of the stock price, α is the expected return on the stock, V_t is the current variance, κ is the mean reversion parameter for the variance process, V_0 is the long-run mean level of variance, and η is the volatility of the variance process. The two Wiener processes $\mathrm{d}z_1$ and $\mathrm{d}z_2$ are instantaneously correlated, with a correlation coefficient of ρ. The parameter ρ is the one of central interest for our purposes, as it governs the skewness of stock returns: when $\rho = 0$, log returns are symmetrically distributed; when $\rho < 0$, log returns are negatively skewed.

In order to map the parameters of the option-pricing model into our NCSKEW variable, we draw on formulas given by Das and Sundaram that express the skewness in daily log returns as a function of the diffusion parameters. If we are willing to fix all the other parameters besides ρ, these formulas allow us to ask to what value of ρ a given value of NCSKEW corresponds. Once we have obtained the implied value of ρ in this way, we can calculate options prices and thereby see the impact of a given value of NCSKEW.

Table 8 illustrates the results of this exercise. Consider first Panel A, where the parameters are chosen so as to be reasonable for individual stocks: $\kappa=1$, $V_0=0.16$, $V_t=0.16$, and $\eta=0.4$. (Setting the variance V to 0.16 corresponds to an annual standard deviation of returns of 40%.) We also set the stock price $P=100$, and the riskless rate $r=0$. We begin with a hypothetical firm 1, which has symmetrically distributed returns – i.e., it has NCSKEW = 0. This is equivalent to a value of $\rho=0$. Next, we take firm 2, which is identical to firm 1, except that it has a two-standard-deviation higher value of DTURNOVER. The standard deviation of DTURNOVER (for firms above the 20th percentile NYSE breakpoint) is 0.042, and from Table 2, column 1, the coefficient on DTURNOVER is 0.437. Hence the value of NCSKEW for firm 2 is 0.037 ($2 \times 0.042 \times 0.437 = 0.037$). Using Eq. (21) in Das and Sundaram (1999, p. 223) this value of skewness in daily returns for firm 2 can be shown to imply $\rho=-0.38$, assuming all the other diffusion parameters stay fixed.

Panel A of Table 8 displays the impact of this change in ρ for the prices of six-month European put options. That is, it calculates put prices for both firm 1 (which has NCKSEW = 0 and thus $\rho=0$) and firm 2 (which has NCSKEW = 0.037 and thus $\rho=-0.38$). As can be seen, the impact on put prices is substantial, particularly if one goes relatively far out-of-the-money. For example, a put with a strike of 70 is worth 1.20 for firm 1 but 1.44 for firm 2, an increase of 20.14%. Or expressed in a different way, the firm 1 put has a Black-Scholes (1973) implied volatility of 40.33%, while the firm 2 put has an implied volatility of 42.50%.

Panel B undertakes a similar experiment to gauge the significance of our time-series results. We keep all the diffusion parameters the same as in Panel A, except that we now set $V_0 = V_t = 0.04$, corresponding to an annual standard deviation of returns of 20%. For the market as a whole, the standard deviation of DTURNOVER is 0.005 (see Table 1A). Using the coefficient estimate on DTURNOVER of 6.00 from Table 7, column 1, a two-standard-deviation shock to DTURNOVER translates into a movement of 0.06 in the NCSKEW variable. Given the other diffusion parameters, this value of 0.06 for NCSKEW is equivalent to $\rho=-0.33$.

Panel B then compares the prices of six-month European puts across two regimes, the first with $\rho=0$ and the second with $\rho=-0.33$. Once again, the differences appear to be meaningful. For example, a put with a strike price of 85 is worth 0.86 in regime 1 but 1.07 in regime 2, an increase of 24.66%. The corresponding implied volatilities are 20.36% and 21.84%, respectively. These results reinforce a point made above: while the time-series estimates are

Table 8
Economic significance of trading volume for skewness in stock returns: an option-pricing metric

Using the stochastic volatility option pricing model (and notation) of Das and Sundaram (1999) we consider what a two-standard-deviation shock in detrended trading volume implies for the prices of six-month European options.

Panel A: Options on individual stocks

The benchmark parameters are as follows: stock price $P=100$, interest rate $r=0$, annualized long-run variance $V_0=0.16$, current variance $V=0.16$, mean reversion in variance $\kappa=1$, and volatility of variance $\eta=0.4$. Firm 1 is assumed to have a value of $\rho=0$. Firm 2 is assumed to have a value of $\rho=-0.38$. These values of ρ imply that the difference in daily skewness between Firms 1 and 2 is equivalent to that created by a two-standard-deviation move in the DTURNOVER variable, using our baseline firm-level sample and coefficient estimates from Table 2, col.1.

	70	80	90	100	110	120	130
Firm 1: $\rho=0$							
Six-month European put price	1.197	3.044	6.287	11.082	17.325	24.748	33.044
Black-Scholes implied vol.	40.33%	39.79%	39.50%	39.41%	39.48%	39.67%	39.93%
Firm 2: $\rho=-0.38$							
Six-month European put price	1.438	3.297	6.419	10.994	17.011	24.282	32.525
Black-Scholes implied vol.	42.50%	41.16%	40.03%	39.10%	38.35%	37.77%	37.34%
Percent increase in put price: Firm 2 vs. Firm1	20.14%	8.30%	2.09%	−0.80%	−1.81%	−1.88%	−1.57%

Panel B: Options on the market portfolio

The benchmark parameters are as follows: stock price $P=100$, interest rate $r=0$, annualized long run variance $V_0=0.04$, current variance $V=0.04$, mean reversion in variance $\kappa=1$, and volatility of variance $\eta=0.4$. Regime 1 is assumed to have a value of $\rho=0$. Regime 2 is assumed to have a value of $\rho=-0.33$. These values of ρ imply that the difference in daily skewness between Regimes 1 and 2 is equivalent to that created by a two-standard-deviation move in the market DTURNOVER variable, using our time-series estimates from Table 7, col. 1.

	85	90	95	100	105	110	115
Regime 1: $\rho=0$							
Six-month European put price	0.859	1.693	3.121	5.330	8.367	12.093	16.298
Black-Scholes implied vol.	20.36%	19.61%	19.09%	18.91%	19.07%	19.49%	20.04%
Regime 2: $\rho=-0.33$							
Six-month European put price	1.070	1.912	3.258	5.289	8.134	11.755	15.955
Black-Scholes implied vol.	21.84%	20.68%	19.63%	18.76%	18.21%	18.01%	18.10%
Percent increase in put price: Regime 2 vs. Regime 1	24.66%	12.91%	4.39%	−0.77%	−2.79%	−2.80%	−2.10%

statistically much weaker than those from the cross-section, they are no less indicative of important economic effects.

Although they are not shown in Table 8, we have also done similar calculations to measure the economic significance of our results for past returns. As it turns out, the quantitative influence of past returns on skewness is stronger than that of trading volume. With individual stocks, a shock of 40% to the RET variable in the most recent six-month period (note from Table 1 that 40% is approximately a two-standard-deviation shock for a firm in the next-to-largest quintile) translates into a movement of 0.087 in the NCSKEW variable. This in turn is equivalent to ρ going from zero to -0.920, which causes the put with the 70 strike to rise in value from 1.20 to 1.73, a 44.86% increase.

In the case of the aggregate market, the coefficients on past returns suggest effects on skewness that are so large that they cannot generally be captured in the context of a pure diffusion model like that of Das and Sundaram (1999). For example, even if the RET variable has moved by only 7% in the last six months, one has to adjust ρ from zero to -0.978 to reflect the corresponding predicted change in NCSKEW. Given that the standard deviation of RET for the market is about 11%, this means that we cannot even capture a one-standard-deviation shock to six-month returns without violating the constraint that the absolute value of ρ not exceed one. Rather, we are left to conclude that, for the market as a whole, large movements in past returns give rise to conditional negative skewness so substantial that it cannot be adequately represented in terms of a pure diffusion process – one would instead need some type of mixed jump-diffusion model.

7. Conclusions

Three robust findings about conditional skewness emerge from our analysis of individual stocks. In the cross-section, negative skewness is greater in stocks that (1) have experienced an increase in trading volume relative to trend over the prior six months, (2) have had positive returns over the prior 36 months, and (3) are larger in terms of market capitalization. The first two results also have direct analogs in the time-series behavior of the aggregate market, though the statistical power of our tests in this case (especially with respect to trading volume) is quite limited.

Let us try to put each of these findings into some perspective. The first, regarding trading volume, is the most novel, and is the one we were looking for based on a specific theoretical prediction from the model of Hong and Stein (1999). Clearly, our results here are supportive of the theory. At the same time, this does not mean that there are not other plausible interpretations. While we have attempted to control for some of the most obvious alternative stories, no doubt others can be thought up. This caveat would seem to be particularly relevant given that there has been so little research to date on conditional skewness at the individual stock level.

The second and third findings, having to do with the effects of past returns and size on skewness, do not speak directly to predictions made by the Hong-Stein model. Rather, these variables are included in our regressions because prior work (Harvey and Siddique, 2000) suggests that they might enter significantly, and we want to be careful to isolate the effects of trading volume from other factors.

Having verified the importance of past returns, we have found it helpful to think about it in terms of models of stochastic bubbles, such as that developed by Blanchard and Watson (1982). However, we would stop well short of claiming to have strong evidence in favor of the existence of bubbles. Indeed, there is a large body of research from the 1980s (see, e.g., reviews by West, 1988; Flood and Hodrick, 1990) that focuses on a very different set of implications of bubble models – having to do with the relation between prices and measures of fundamentals such as dividends – and tends to reach mostly skeptical conclusions on this question. Rather, the more modest statement to be made is that previous research has not examined the implications of bubble models for conditional skewness, and that on this one score, the bubble models look pretty good.

With respect to the third finding – that small-cap stocks are more positively skewed than large-cap stocks – we are not even aware of an existing theory that provides a simple explanation. Instead, we have developed an informal hypothesis after the fact, based on the ideas that (1) managers prefer to disclose good news right away, while dribbling bad news out slowly, and (2) managers of small companies have more scope for hiding bad news from the market in this way. This discretionary-disclosure hypothesis in turn yields the further prediction that, controlling for size, positive skewness ought to be more pronounced in stocks with fewer analysts – a prediction which is clearly supported in the data.

A fair criticism of this whole line of discussion is that we have three main empirical results, and a different theoretical story for each: the Hong-Stein (1999) model to explain the effect of turnover on skewness, stochastic bubbles to explain the effect of past returns, and discretionary disclosure to explain the effect of size. This lack of unity is unsatisfying, and it serves to further underscore our previous caveat about the extent to which one should at this point consider any single theory to be strongly supported by the data. A natural challenge for future work in this area is to come up with a parsimonious model that captures these three patterns in a more integrated fashion.

References

Almazan, A., Brown, B.C., Carlson, M., Chapman, D.A., 1999. Why constrain your mutual fund manager?. Working paper, University of Texas, Austin.

Bakshi, G., Cao, C., Chen, Z., 1997. Empirical performance of alternative option pricing models. Journal of Finance 52, 2003–2049.

Bates, D., 1991. The crash of '87: was it expected? The evidence from options markets. Journal of Finance 46, 1009–1044.

Bates, D., 1997. Post-'87 crash fears in S&P 500 futures options. NBER working paper 5894.

Bekaert, G., Wu, G., 2000. Asymmetric volatility and risk in equity markets. Review of Financial Studies 13, 1–42.

Black, F., 1976. Studies of stock price volatility changes. Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economical Statistics Section, 177–181.

Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. Journal of Political Economy 81, 637–659.

Blanchard, O.J., Watson, M.W., 1982. Bubbles, rational expectations, and financial markets. In: Wachtel, P. (Ed.), Crises in Economic and Financial Structure. Lexington Books, Lexington, MA, pp. 295–315.

Braun, P.A., Nelson, D.B., Sunier, A.M., 1995. Good news, bad news, volatility, and betas. Journal of Finance 50, 1575–1603.

Campbell, J.Y., Hentschel, L., 1992. No news is good news: an asymmetric model of changing volatility in stock returns. Journal of Financial Economics 31, 281–318.

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

Coval, J.D., Hirshleifer, D., 1998. Trading-generated news, sidelined investors, and conditional patterns in security returns. Working paper, University of Michigan Business School, Ann Arbor.

Damodaran, A., 1987. Information bias: measures and implications. Working paper, New York University.

Das, S.R., Sundaram, R.K., 1999. Of smiles and smirks: a term structure perspective. Journal of Financial and Quantitative Analysis 34, 211–239.

DeLong, J.B., Shleifer, A., Summers, L.H., Waldmann, R.J., 1990. Noise trader risk in financial markets. Journal of Political Economy 98, 703–738.

Duffee, G.R., 1995. Stock returns and volatility: a firm-level analysis. Journal of Financial Economics 37, 399–420.

Dumas, B., Fleming, J., Whaley, R.E., 1998. Implied volatility functions: empirical tests. Journal of Finance 53, 2059–2106.

Engle, R.F., Ng, V.K., 1993. Measuring and testing the impact of news on volatility. Journal of Finance 48, 1749–1778.

Fama, E.F., 1998. Market efficiency, long-term returns, and behavioral finance. Journal of Financial Economics 49, 283–306.

Fama, E.F., MacBeth, J.D., 1973. Risk, return and equilibrium: empirical tests. Journal of Political Economy 81, 607–636.

Flood, R.P., Hodrick, R.J., 1990. On testing for speculative bubbles. Journal of Economic Perspectives 4, 85–101.

French, K.R., Schwert, G.W., Stambaugh, R.F., 1987. Expected stock returns and volatility. Journal of Financial Economics 19, 3–29.

Gallant, A.R., Rossi, P.E., Tauchen, G., 1992. Stock prices and volume. Review of Financial Studies 5, 199–242.

Genotte, G., Leland, H., 1990. Market liquidity, hedging and crashes. American Economic Review 80, 999–1021.

Glosten, L., Jagannathan, R., Runkle, D.E., 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. Journal of Finance 48, 1779–1801.

Greene, W.H., 1993. Econometric Analysis. Macmillan, New York.
Griffin, J.M., Lemmon, M.L., 1999. Does book-to-market equity proxy for distress risk or overreaction? Working paper, Arizona State University, Tempe, AZ.
Grossman, S.J., 1988. An analysis of the implications for stock and futures price volatility of program trading and dynamic hedging strategies. Journal of Business 61, 275–298.
Harris, M., Raviv, A., 1993. Differences of opinion make a horse race. Review of Financial Studies, 473–506.
Harvey, C.R., Siddique, A., 1999. Autoregressive conditional skewness. Journal of Financial and Quantitative Analysis 34, 465–487.
Harvey, C.R., Siddique, A., 2000. Conditional skewness in asset pricing tests. Journal of Finance 55, 1263–1295.
Hong, H., Lim, T., Stein, J.C., 2000. Bad news travels slowly: size, analyst coverage and the profitability of momentum strategies. Journal of Finance 55, 65–295.
Hong, H., Stein, J. C., 1999. Differences of opinion, rational arbitrage and market crashes. NBER Working paper.
Jacklin, C.J., Kleidon, A.W., Pfleiderer, Paul., 1992. Underestimation of portfolio insurance and the crash of October 1987. Review of Financial Studies 5, 35–63.
Kandel, E., Pearson, N.D., 1995. Differential interpretation of public signals and trade in speculative markets. Journal of Political Economy 103, 831–872.
Koski, J.L., Pontiff, J., 1999. How are derivatives used? Evidence from the mutual fund industry. Journal of Finance 54, 791–816.
Lakonishok, J., Smidt, S., 1986. Volume for winners and losers: taxation and other motives for stock trading. Journal of Finance 41, 951–974.
Nelson, D., 1991. Conditional heteroskedasticity in asset returns: a new approach. Econometrica 59, 347–370.
Odean, T., 1998a. Volume, volatility, price and profit when all traders are above average. Journal of Finance 53, 1887–1934.
Odean, T., 1998b. Are investors reluctant to realize their losses? Journal of Finance 53, 1775–1798.
Pindyck, R.S., 1984. Risk, inflation, and the stock market. American Economic Review 74, 334–351.
Poterba, J.M., Summers, L.H., 1986. The persistence of volatility and stock market fluctuations. American Economic Review 76, 1142–1151.
Romer, D., 1993. Rational asset-price movements without news. American Economic Review 83, 1112–1130.
Schwert, G.W., 1989. Why does stock market volatility change over time? Journal of Finance 44, 1115–1153.
Shefrin, H., Statman, M., 1985. The disposition to sell winners too early and ride losers too long: theory and evidence. Journal of Finance 40, 777–790.
Varian, H.R., 1989. Differences of opinion in financial markets. In: Stone, C. (Ed.), Financial Risk: Theory, Evidence and Implications: Proceedings of the 11th Annual Economic Policy Conference of the Federal Reserve Bank of St. Louis. Kluwer Academic Publishers, Boston, pp. 3–37.
West, K.D., 1988. Bubbles, fads and stock price volatility tests: a partial evaluation. Journal of Finance 43, 639–656.
Wu, G., 2000. The determinants of asymmetric volatility. Review of Financial Studies, in preparation.