Ex Ante Skewness and Expected Stock Returns*

Jennifer Conrad[†] **Robert F. Dittmar**[‡] **Eric Ghysels**[§]

First Draft: March 2007
This Draft: October 2008

Abstract

We use a sample of option prices, and the method of Bakshi, Kapadia and Madan (2003), to estimate the *ex ante* higher moments of the underlying individual securities' risk-neutral returns distribution. We find that individual securities' volatility, skewness and kurtosis are strongly related to subsequent returns. Specifically, we find a negative relation between volatility and returns in the cross-section. We also find a significant relation between skewness and returns, with more negatively (positively) skewed returns associated with subsequent higher (lower) returns, while kurtosis is positively related to subsequent returns. To analyze the extent to which these returns relations represent compensation for risk, we use data on index options and the underlying index to estimate the stochastic discount factor over the 1996-2005 sample period, and allow the stochastic discount factor to include higher moments. We find evidence that, even after controlling for differences in co-moments, individual securities' skewness matters. However, when we combine information in the risk-neutral distribution and the stochastic discount factor to estimate the implied physical distribution of industry returns, we find little evidence that the distribution of technology stocks was positively skewed during the bubble period–in fact, these stocks have the lowest skew, and the highest estimated Sharpe ratio, of all stocks in our sample.

*All errors are the responsibility of the authors. We thank Robert Battalio, Patrick Dennis, and Stewart Mayhew for providing data and computational code. We thank Andrew Ang, Leonce Bargeron, and Paul Pfleiderer for helpful comments and suggestions, as well as seminar participants at Babson College, Berkeley, Boston College, Cornell University, the Federal Reserve Bank of New York, National University of Singapore, New York University, Northern Illinois University, Queen's University Belfast, Stanford University and the Universities of Arizona, Michigan, Texas, and Virginia. A previous version of this paper was circulated under the title "Skewness and the Bubble."

[†]Department of Finance, Kenan-Flagler Business School, University of North Carolina at Chapel Hill

[‡]Department of Finance, Stephen Ross School of Business, University of Michigan, Ann Arbor, MI 48109

[§]Department of Finance, Kenan-Flagler Business School, and Department of Economics, University of North Carolina at Chapel Hill

1 Introduction

What role do higher moments play in investors' decisions about the choice of portfolios and the pricing of assets? Arditti (1967) shows that investors with decreasing risk aversion will display preference for greater skewness in asset payoffs, and Rubinstein (1973) and Kraus and Litzenberger (1976, 1983) formalize this preference in the context of a pricing model. More recently, Harvey and Siddique (2000) document empirical evidence supporting the role of skewness risk in explaining cross-sectional differences in returns, and Dittmar (2000) shows that skewness (and kurtosis) appear to play a significant role in pricing.

The common theme in these papers is that investors discount *aggregate* skewness. That is, investors are willing to pay more for a security with greater co-skew with some stochastic discount factor. A more recent literature has suggested that *total* rather than *co*-skewness plays a role in informing portfolio decisions and asset prices. Barberis and Huang (2007) suggest that, under cumulative prospect theory, agents will display a preference for stocks with more skewed returns. As a result, an asset with high total skewness will appear overpriced relative to a model with standard expected utility. Similar results are obtained with a different preference structure in Brunnermeier, Gollier, and Parker (2007). The models in these papers are consistent with the evidence in Mitton and Vorkink (2007) that suggests that individual investors with undiversified portfolios hold assets and portfolios that exhibit greater idiosyncratic skewness.

In this paper we examine the effect of total skewness on the pricing of equity securities. An important feature of the approach taken in our paper is that we focus on the *ex ante* distribution of returns by using information contained in option prices. Under the assumption of a no-arbitrage link between options and underlying markets, we retrieve risk-neutral measures of distributional moments following the procedure in Bakhsi, Kapadia, and Madan (2001). We suggest a number of advantages to this approach, compared to alternatives that measure distributional moments from the time series of underlying market asset returns. First, as noted by Bates (1991), Rubinstein (1985, 1994), and Jackwerth and Rubinstein (1996), option prices efficiently capture a market-based estimate of investors' beliefs. Second, the use of option prices eliminates the need for a long time series to reliably estimate higher moments of the distribution. This consideration is of particular importance in gauging beliefs about relatively new firms (i.e. Internet companies), or during periods in which beliefs may change relatively quickly. Third, options provide an ex ante measure of beliefs; they do not give us, as Battalio and Schultz (2006) note, the "unfair advantage of hindsight." As Jackwerth and Rubinstein (1996) state, "not only can the nonparametric method reflect the possibly complex logic used by market participants to consider the significance of extreme events, but it also implicitly brings a much larger set of information ... to bear on the formulation of probability distributions."

We first examine whether dispersion in skewness generates differences in expected returns across assets. We find that, indeed, assets with high *ex ante* skewness earn lower average returns than assets with low ex ante skewness. We then investigate the primary source of tension in the two streams of research discussed above; is the importance of skewness in pricing due to co-movement with some aggregate stochastic discount factor, or is is residual idiosyncratic skewness that matters in determining prices? We exploit no-arbitrage conditions in the options and cash markets to find evidence suggestive of a residual idiosyncratic skewness risk premium after accounting for systematic skewness. Finally, we ask whether differences in views of *ex ante* skewness can help explain why certain types of stocks, particularly tech stocks, had such high valuations in the late 1990s and early 2000s. We find that skewness had little to do with these valuations; rather, investors appear to have viewed these assets as good *ex ante* Sharpe ratio bets.

Two other recent papers also investigate measures of skewness and their relation to stock prices. Xing, Zhang, and Zhao (2007) find that portfolios sorted on differences in the slope of the volatility smirk generate differences in average returns. Since the slope of the smirk has been related to the probability of negative jumps in price levels, as suggested in Bates (1991) and Pan (2002), one may infer that the slope of the smirk is related to negative skewness. There are several differences between our paper and theirs. First, our measure of skewness includes information about both left-skewed and right-skewed behavior, since it uses information in both out-of-the-money puts and calls. Second, the focus in our paper differs: we are interested not only in the information that the risk-neutral skew may have for future stock returns, but also in the implications for the pricing of systematic and idiosyncratic risk.

A second study, Boyer, Mitton, and Vorkink (2008), examines the role of a measure of idiosyncratic skewness in explaining differences in returns across securities. The authors use a long-horizon cross-sectional model of forecasting the skew in individual security returns, and find a negative relation between idiosyncratic skewness and returns, as suggested by the theories discussed above. They also show that idiosyncratic skewness can help explain the role of idiosyncratic variance in generating cross-sectional dispersion in returns. Their measure of skewness is substantially different from ours, involving the use of a fairly long time-series (60 months) of *ex post* data; in addition, they do not explore the difference between systematic and idiosyncratic skewness.

The remainder of the paper is organized as follows. In section 2, we detail the method we employ for recovering measures of volatility, skewness, and kurtosis, following Bakshi, Kapadia, and Madan (2003). In Section 3, we discuss the data used in our analysis and present results of empirical tests performed on portfolios formed on the basis of the volatility, skewness, and kurtosis measures. In Section 4, we use data on the market portfolio, and its options, to estimate a stochastic discount factor which includes the information in higher moments, and use this stochastic discount factor to risk-adjust the raw returns related to higher moments. In Section 5, we discuss the estimation of implied physical distributions for individual securities, and present these estimates for industry portfolios. We conclude in Section 6.

2 Risk-Neutral Moments and Asset Prices

Throughout our discussion, we are assuming that securities are priced to eliminate risk-free arbitrage opportunities. As discussed in Harrison and Kreps (1979), the lack of arbitrage opportunities in the market implies the existence of a probability measure that prices payoffs by discounting at the risk free rate. Formally, this *risk-neutral* probability measure, Q, satisfies

$$P_t = e^{-r\tau} E_t^Q \left[P_t \left(1 + R_{t+\tau} \right) \right]. \tag{1}$$

where P_t represents the asset's price, r is the risk free rate, τ is the holding period, and $R_{t+\tau}$ represents the return on the asset. Equivalently, a stochastic discount factor, $M_{t+\tau}$, exists that discounts payoffs to current prices under the physical probability measure, P.

As noted in the introduction, there is a large body of theory and evidence that suggests that moments (variance, skewness, and kurtosis) of the physical distribution are important in determining investors' portfolio choice and the pricing of assets. Equation (1) similarly suggests that moments of the risk-neutral distribution will affect investors' pricing of assets.

We recover the risk-neutral moments above using the prices of options. Recovering risk-neutral distributions from option prices has a long history in the literature (see Figlewski (2007) for a review). One of the advantages of this approach is that it recovers moments from asset prices, rather than realized returns. Thus, the estimates are representative of the *ex ante* moments relevant for asset pricing, allaying the criticism leveled in Battalio and Schulz (2006) of the "unfair advantage of hindsight." Our specific approach follows Bakshi and Madan (2000) and Bakshi, Kapadia, and Madan (2003).

2.1 Computing Risk Neutral Moments

Bakshi and Madan (2000) show that any payoff to a security can be constructed and priced using a set of option prices with different strike prices on that security. Bakshi, Kapadia, and Madan (2003) demonstrate how to express the risk-neutral density moments in terms of quadratic, cubic, and quartic payoffs. In particular, Bakshi, Kapadia, and Madan (2003) show that one can express the τ-maturity price of a security that pays the quadratic, cubic, and quartic return on the base security as

$$\begin{aligned} V(t,\tau) &= \int_{S(t)}^{\infty} \frac{2(1 - ln(K/S(t)))}{K^2} C(t,\tau;K)dK \\ &+ \int_{0}^{S(t)} \frac{2(1 + ln(K/S(t)))}{K^2} P(t,\tau;K)dK \end{aligned} \tag{2}$$

$$\begin{aligned} W(t,\tau) &= \int_{S(t)}^{\infty} \frac{6ln(K/S(t)) - 3(ln(K/S(t)))^2)}{K^2} C(t,\tau;K)dK \\ &+ \int_{0}^{S(t)} \frac{6ln(K/S(t)) + 3(ln(K/S(t)))^2}{K^2} P(t,\tau;K)dK \end{aligned} \tag{3}$$

$$\begin{aligned} X(t,\tau) &= \int_{S(t)}^{\infty} \frac{12(ln(K/S(t)))^2 - 4(ln(K/S(t)))^3)}{K^2} C(t,\tau;K)dK \\ &+ \int_{0}^{S(t)} \frac{12(ln(K/S(t)))^2 + 4(ln(K/S(t)))^3}{K^2} P(t,\tau;K)dK \end{aligned} \tag{4}$$

where $V(t,\tau)$, $W(t,\tau)$, and $X(t,\tau)$ are the quadratic, cubic, and quartic contracts, respectively, and $C(t,\tau;K)$ and $P(t,\tau;K)$ are the prices of European calls and puts written on the underlying stock with strike price K and expiration τ periods from time t. As equations (2), (3) and (4) show, the procedure involves using a weighted sum of (out-of-the-money) options across varying strike prices to construct the prices of payoffs related to the second, third and fourth moments of returns.

Using the prices of these contracts, standard moment definitions suggest that the risk-neutral moments can be calculated as

$$\sigma^Q(t,\tau) = \sqrt{e^{r\tau}V(t,\tau) - \mu(t,\tau)^2} \tag{5}$$

$$\gamma^Q(t,\tau) = \frac{e^{r\tau}W(t,\tau) - 3\mu(t,\tau)e^{r\tau}V(t,\tau) + 2\mu(t,\tau)^3}{[e^{r\tau}V(t,\tau) - \mu(t,\tau)^2]^{3/2}} \tag{6}$$

$$\kappa^Q(t,\tau) = \frac{e^{r\tau}X(t,\tau) - 4\mu(t,\tau)W(t,\tau) + 6e^{r\tau}\mu(t,\tau)^2V(t,\tau) - \mu(t,\tau)^4}{[e^{r\tau}V(t,\tau) - \mu(t,\tau)^2]^2} \tag{7}$$

where

$$\mu(t,\tau) = e^{r\tau} - 1 - e^{r\tau}V(t,\tau)/2 - e^{r\tau}W(t,\tau)/6 - e^{r\tau}X(t,\tau)/24 \tag{8}$$

and r represents the risk-free rate. We follow Dennis and Mayhew (2002), and use a trapezoidal approximation to estimate the integrals in expressions (2)-(4) above using discrete data.[1]

2.2 Data

Our data on option prices are from Optionmetrics (provided through Wharton Research Data Services) . We begin with daily option price data for all out-of-the-money calls and puts for all stocks from 1996-2005.[2] Closing prices are constructed as midpoint averages of the closing bid and ask prices.

Some researchers have argued that option prices and equity prices diverged during our sample period. For example, Ofek and Richardson (2003) propose that the Internet bubble is related to the 'limits to arbitrage' argument of Shleifer and Vishny (1997). This argument requires that investors could not, or did not, use the options market to profit from mis-pricing in the underlying market, and, in fact, Ofek and Richardson (2003) also provide empirical evidence that option prices diverged from the (presumably misvalued) prices of the underlying equity during this period. However, Battalio and Schultz (2006) use a different dataset of option prices than Ofek and Richardson (2003), and conclude that shorting synthetically using the options market was relatively easy and cheap, and that short-sale restrictions are not the cause of persistently high Internet stock prices. A corollary to their results is that option prices and the prices of underlying stocks did not diverge during the 'bubble' period and they

[1] We are grateful to Patrick Dennis for providing us with his code to perform the estimation.

[2] We do not adjust for early exercise premia in our option prices. As Bakshi,Kapadia and Madan (2003) note, the magnitude of such premia in OTM calls and puts is very small, and the implicit weight that options receive in the estimation declines as they get closer to at-the-money. In their empirical work, BKM show that, for their sample of OTM options, the implied volatilities from the Black-Scholes model and a model of American option pricing have negligible differences.

argue that Ofek and Richardson's results may be a consequence of misleading or stale option prices in their data set. Note that if option and equity prices do not contain similar information, then our tests should be biased against finding a systematic relation between estimates of higher moments obtained from option prices and subsequent returns in the underlying market.[3] However, motivated by the Battalio and Schultz results, we employ a number of filters to try to ensure that our results are not driven by stale or misleading prices. We eliminate option prices below 50 cents, as well as options with less than one week to maturity. At the outset, we require that an option has a minimum of ten days of quotes during any month; in later robustness checks, we impose additional constraints on the liquidity in the option. We also eliminate days in which closing quotes on put-call pairs violate no-arbitrage restrictions.

In estimating equations (5) - (7), we use equal numbers of out-of-the-money (OTM) calls and puts for each stock for each day. Thus, if there are n OTM puts with closing prices available on day t we require n OTM call prices. If there are $N > n$ OTM call prices available on day t, we use the n OTM calls which have the most similar distance from stock to strike as the OTM puts for which we have data. We require a minimum n of 2; we perform robustness checks on our results when this minimum data constraint is increased.[4] The resulting set of data consists of 3,722,700 daily observations across firms and maturities over the 1996-2005 sample period.

In Table 1, we present descriptive statistics for the sample estimates of volatility, skewness and kurtosis. We report medians, 5th and 95th percentiles across time and securities for each year during the sample period. There are clear patterns in the time series of these moments through the sample period, as well as evidence of interactions between them. Volatility peaks in 2000, during the height of the bubble period, then declines through 2005. The median risk-neutral skewness is negative, indicating that the distribution is left-skewed; the median value stays relatively flat through 2000 after which it declines sharply, while the median kurtosis estimate increases during that same period, more than doubling from 2000 through 2005.

[3] Robert Battalio graciously provided us with the OPRA data used in their analysis; unfortunately, these data, provided by a single dealer, do not have a sufficient cross-section of data across calls and puts to allow us to estimate the moments of the risk-neutral density function in which we are interested.

[4] Dennis and Mayhew (2006) examine and estimate the magnitude of the bias induced in Bakshi-Kapadia-Madan estimates of skewness which is due to discreteness in strike prices. For $ 5 ($2.50) differences in strike prices, they estimate the bias in skewness is approximately -0.07 (-0.05). Since most stocks have the same differences across strike prices, however, the relative bias should be approximately the same across securities, and should not affect either the ranking of securities into portfolios based on skewness, or the nature of the cross-sectional relation between skewness and returns which we examine.

3 Risk-Neutral Moments and the Cross-Section of Returns

In this section, we examine whether portfolios formed on the basis of risk-neutral moments are associated with cross-sectional dispersion in subsequent returns. Data on stock returns are obtained from the Center for Research in Security Prices (again provided through Wharton Research Data Services). The basis for our analysis is the intersection of the the options data discussed above and monthly data on all individual securities with common shares outstanding.

3.1 Raw and Characteristic-Adjusted Returns

We begin by selecting daily observations of prices of out-of-the-money calls and puts on individual securities, which have maturities closest to 1 month, 3 months, 6 months and 12 months, and group these options into separate maturity bins. In each maturity bin, we estimate the moments of the risk-neutral density function for each individual security on a daily basis. Following Bakshi, Kapadia and Madan (2003), we average the daily estimates for each stock over time (in our case, the calendar quarter.) For each maturity bin, we further sort options into tercile portfolios based on the moment estimates (volatility, skewness or kurtosis); the 'extreme' portfolios contain 30% of the sample, while portfolio 2 contains 40% of the sample. We re-form portfolios each month, holding moment ranks constant over the calendar quarter. In each quarter, we also remove firms that are in the top 1% of the cross-sectional distribution of volatility, skewness or kurtosis to mitigate the effect of outliers.

In Table 2, we report results for portfolios sorted on the basis of estimated volatility (Panel A), estimated skewness (Panel B) and estimated kurtosis (Panel C). Specifically, we report the average of the subsequent raw returns of the equally-weighted moment-ranked portfolios over the next month in the first column of data. In the next column, we report the average characteristic-adjusted return over that same month. To calculate the characteristic-adjusted return, we perform a calculation similar to that in Daniel et al.$\tilde{(}$1997). For each individual firm, we assess to which of the 25 Fama-French size- and book-to-market ranked portfolios the security belongs. We subtract the return of that Fama-French portfolio from the individual security return and then average the resulting excess or characteristic-adjusted 'abnormal' return across firms in the moment-ranked portfolio. In the next three columns, we report the average risk-neutral volatility, skewness and kurtosis estimates for each of the ranked portfolios. Finally, we report average betas, average (log) market value and average book-to-market equity ratios of the securities in the portfolio.

Summary statistics in Panel A of Table 2 suggest a strong negative relation between volatility and subsequent raw returns; for example, in the shortest maturity options (maturity bin 1), the returns differential between high volatility (Portfolio 3) and low volatility (Portfolio 1) securities is -32 basis points per month; longer maturities have differentials between 50 and 56 basis points per month. The columns of data which report the average characteristics of securities in the portfolio show sharp differences in beta, size and book-to-market equity ratios across these volatility-ranked portfolios. Low (high) volatility portfolios tend to contain low (high) beta firms and larger (smaller) firms, while differences in book-to-market equity ratios across portfolios are relatively small and differ across maturity bins. We adjust for these differences in size and book-to-market equity ratio in the characteristic-adjusted return column. After adjusting for the differences in size and book-to-market observed across the volatility portfolios, the return differentials are somewhat attenuated in all four maturities. However, although the differential is reduced, it remains significant, with lowest volatility portfolios earning between 10 and 23 basis points per month more than the highest volatility portfolios in all four maturity bins.

Panel A also indicates that there is a weak negative relation between volatility and skewness; in all maturity bins, skewness has a tendency to decline as volatility increases, although the effect is not monotonic. The relation between volatility and kurtosis in Panel A is much stronger: as average volatility increases in the portfolio, kurtosis declines in all four maturity bins. Thus, the relation between volatility and returns may be confounded by the effect, if any, of other moments on returns; we examine this possibility in later sections of the paper. Finally, the average number of securities in each portfolio indicates that the portfolios should be relatively well-diversified. The top and bottom tercile portfolios average 273 firms, whereas the middle tercile portfolio averages 365 firms. Combined with the fact that we are sampling securities which have publicly traded options, this breadth should reduce the effect of outlier firms on our results.

Panel B of Table 2 sorts securities into portfolios on the basis of estimated skewness. Interestingly, we see significant differences in returns across skewness-ranked portfolios. The raw returns differential is negative for all four maturities, at 26, 43, 47 and 44 basis points per month, respectively. That is, on average, in each maturity bin the securities with lower skewness earn higher returns in the next month, while securities with less negative, or positive, skewness earn lower returns. The differentials in raw returns are of the same order of magnitude or larger than that observed in the volatility-ranked portfolios in Panel A. Compared to the volatility-ranked portfolios, the skewness-ranked portfolios show relatively little difference in their average market capitalization and betas, although differences in book-to-market equity ratios remain. When we adjust for the size- and book-to-market characteristics of securities, the characteristic-adjusted returns are reduced only slightly, and average 28, 43, 39 and 40 basis points per month, respectively, across the maturity bins.[5]

In addition to the differences in returns, the table indicates that there is a negative relation between skewness and both volatility and kurtosis. That is, both volatility and kurtosis decline as we move across skewness-ranked portfolios. As in Panel A, interactions between other moments and returns could be masking or exacerbating the relation between skewness and returns. Consequently, in later tests, we control for the relation of other higher moments to returns in estimating their effect.

Finally, Panel C of Table 2 reports the results when securities are sorted on the basis of estimated kurtosis. Generally, we see a positive relation between kurtosis and subsequent raw returns; the return differential is economically significant, at 12, 31, 35 and 37 basis points per month across the four maturities. As with the other moment-ranked portfolios, the effect is reduced after adjusting for book-to-market and market capitalization differences, but the differences are very slight and the effect remains highly significant, at 14, 30, 35 and 36 basis points per month across maturity bins. We also observe patterns in the other estimated moments, with both volatility and skewness decreasing as kurtosis increases. Again, this emphasizes the need to control for the relation of all higher moments to returns.

The results in Table 2, Panels A-C, suggest that, on average, higher moments in the distribution of securities' payoffs are related to subsequent returns. Consistent with the evidence in Ang, Hodrick, Xing and Zhang (2006a), we see that securities with higher volatility have lower subsequent returns. We also find that securities with higher skewness have lower subsequent returns, while higher kurtosis is related to higher subsequent returns. As a robustness check, in the next section we use a factor-adjustment method which controls for other characteristics of the firms.

3.2 Factor-Adjusted Returns

In Table 2 above, we adjust for the differences in characteristics across portfolios, following Daniel et al. (1997), by subtracting the return of the specific Fama-French portfolio to which an individual firm is assigned. However, Fama and French (1993) interpret the relation between characteristics and returns as evidence of risk factors. Consequently, we also adjust for differences in characteristics across our moment-sorted portfolios by estimating a time series regression of the 'factor-mimicking' portfolio returns on the three factors proposed in Fama

[5] In a different application, Xing, Zhang and Zhao (2007) find a positive relation between a skewness metric taken from option prices and the next month's returns. Their measure of skewness is the absolute value of the difference in implied volatilities in out-of-the-money call option contracts, where the strike price is constrained to be within the range of 0.8S to S, and preferably in the range of 0.95S to S. Thus, their skewness measure is related to the slope of the volatility smile over a smaller range of strike prices.

and French (1993). The dependent variable in these regressions is the monthly return from portfolios re-formed each month (as in Table 2), where the portfolios consist of a long position in the portfolio of securities with the highest estimated moments, and a short position in the portfolio of securities with the lowest estimated moments. The three factors used as independent variables in the regressions are the return on the value-weighted market portfolio in excess of the risk-free rate ($r_{MRP,t}$), the return on a portfolio of small capitalization stocks in excess of the return on a portfolio of large capitalization stocks ($r_{SMB,t}$), and and the return on a portfolio of firms with high book-to-market equity in excess of the return on a portfolio of firms with low book-to-market equity ($r_{HML,t}$). As in Table 2, firms are grouped by maturity and sorted into portfolios on the basis of estimated moments (volatility, skewness and kurtosis). We report intercepts, slope coefficients for the three factors, and adjusted R-squareds. Standard errors for the coefficients are presented in parentheses, and are adjusted for serial correlation and heteroskedasticity using the Newey and West (1987) procedure.

Panels A-D of Table 3 present results for options closest to one, three, six, and twelve months to maturity, respectively. The first row of each panel contains the results for the long-short portfolio constructed from volatility-sorted portfolios. Consistent with the results in Panel A of Table 2 for characteristic-adjusted returns, we observe negative alphas in our "high-low" portfolio in all four maturity bins. The absolute magnitude of the alphas declines from 57 to 41 basis points per month across maturity bins, with t-statistics of -1.77, -1.65, -1.54 and -1.16, respectively. These results are consistent with those of Ang, Hodrick, Xing and Zhang (2006), who show that firms with high idiosyncratic volatility relative to the Fama-French model earn "abysmally low" returns.

The patterns in the intercepts for skewness-sorted portfolios (row 2 of Panels A-D of Table 3) are also consistent with that observed in Panel B of Table 2. Alphas are negative in all four maturities, significant at the 10% level for the one month maturity and at the 5% level or better in the other three maturities. The alphas remain roughly constant in magnitude as we move from short-maturity options to long-maturity options, at 58, 67, 64 and 62 basis points per month, respectively. The negative alphas still suggest a 'low skewness' premium; that is, securities with more negative skewness earn, on average, higher returns in the subsequent months, while securities with less negative, or positive skewness, earn lower returns in subsequent months.

The evidence that skewness in individual securities is negatively related to subsequent returns is consistent with the models of Barberis and Huang (2004), and Brunnermeier, Gollier and Parker (2005). In their papers, they note that investors who prefer positively skewed distributions may hold concentrated positions in (positively skewed) securities–that is, investors may trade off skewness against diversification, since adding securities to a portfolio will increase diversification, but at the cost of reducing skewness. The preference for skewness will increase the demand for, and consequently the price of, securities with higher skewness and consequently reduce their expected returns. This evidence is also consistent with the empirical results in Boyer, Mitton and Vorkink (2008), who generate a cross-sectional model of expected skewness for individual securities and find that portfolios sorted on expected skew generate a return differential of approximately 67 basis points per month.

In the third rows of Panels A-D of Table 3, we report the results for kurtosis-sorted portfolios. Consistent with the results in Table 2, we see positive intercepts in portfolios that are long kurtosis. Alphas are positive and both economically and statistically significant, at 55, 62, 56 and 62 basis points per month, respectively, across the four maturities. Similar to the characteristic-adjusted returns in Table 3, there is no discernible trend in them across maturity bins. The magnitude of the alphas with respect to kurtosis is comparable to that observed in the skewness and volatility sorted portfolios.

There is one other noteworthy feature of Table 3. The explanatory power of the Fama-French three factors is, on average, lower for the kurtosis-sorted High-Low portfolios, and much lower for the skewness-sorted portfolios, than the volatility-sorted portfolios. Some of this difference is likely due to the fact that, as Table 2 shows, skewness and kurtosis-sorted portfolios exhibit much lower differences in size and beta than do the volatility-sorted portfolios. However, it is also possible that there are features of the returns on moment-sorted portfolios that are not captured well by the usual firm characteristics.

3.3 Additional robustness checks

We perform several additional robustness checks on our results to examine the possibility that return differentials are driven by liquidity issues, either in the underlying equity returns or by stale or illiquid option prices. To examine the latter possibility, we add an additional filter to our sample, and eliminate the observation if there is no trading in any of the out-of-the-money options on a particular day. These results are presented in Appendix Table A1. The principal impact of this restriction is to substantially reduce our sample. As discussed above, on average there are 911 firms per month in our original sample (273/365/273 by tercile). Imposing the trading restriction reduces the average number of firms to 307. However, as shown in the table, with the exception of short-maturity kurtosis-sorted portfolios, the magnitude of return differentials across portfolios remains stable, or actually increases. Thus, we continue to find that returns are negatively related to volatility and skewness, and positively related to kurtosis.[6]

Second, we add the liquidity factor of Pastor and Stambaugh (2003) to our time series regression and re-estimate the factor-adjusted returns. These results are presented in Appendix Table A2. The basic results change very little. The intercepts retain negative signs for volatility and skewness and positive signs for kurtosis across all three maturity bins. Statistical significance declines slightly; the alpha for the volatility portfolio loses its statistical significance for all maturities, as does the alpha for the skewness portfolio only in the shortest maturity options. However, the overall conclusions are similar: high volatility and high skewness stocks earn negative excess returns, and high kurtosis stocks earn positive excess returns.

Overall, both the characteristic-adjusted returns in Table 2 and the regression results in Table 3 provide evidence that higher moments in the returns distribution are associated with differences in subsequent returns, and that not all of the return differential observed can be explained by differences in the size, book-to-market, beta or liquidity differentials of the moment-sorted portfolios. That is, on average, we see some relation between the higher moments of risk-neutral returns distributions of individual securities and subsequent returns on these stocks in the underlying market. In the next section, we allow the risk adjustment for subsequent returns to incorporate higher co-moments as well.

4 Higher Moment Premia, Systematic, and Idiosyncratic Risk

In the previous section, we document a negative premium for ex ante volatility and skewness in stock returns, and a positive premium for kurtosis. As discussed in the introduction, an open question is whether these premia are related to systematic or idiosyncratic risk. In this section, we address that question. Specifically, we ask whether observed premia are related to measures of *ex ante* co-moment risk, *ex ante* idiosyncratic risk, or both.

Conceptually, we consider idiosyncratic risk as that portion of a security's return that is orthogonal to the stochastic discount factor, $M(s, t, t+\tau)$. That is, a security's payoff can be decomposed into two components:

$$R_{i,t+1} = R^s_{i,t+1} + e_{i,t+1}$$
$$E^P_t \left[R_{i,t+1} M_{t+1} \right] = E^P_t \left[R^s_{i,t+1} M_{t+1} \right] = 1$$

[6]For brevity, we report only the average and characteristic-adjusted average returns to these portfolios. The remaining characteristics exhibit similar patterns to those depicted in Table 2. These results are available from the authors upon request.

where R^s is the systematic component of gross returns and $e_{i,t+1}$ is the idiosyncratic component. In order to test for the presence of systematic risk, we consider the Euler equation specification

$$E_t^P \left[R_{i,t+1} M_{t+1} \right] - 1 = u_{i,t+1} \tag{9}$$

and test the restriction that

$$E \left[u_{i,t+1} \right] = \alpha = 0 \tag{10}$$

As discussed in Chen and Knez (1996), this α is analogous to Jensen's α.

Depending on one's null, the Euler equation restriction may be viewed as a test of the presence of idiosyncratic components of returns that generate mean returns or a test of model specification. In order to take the former view, one must assume that the stochastic discount factor, M_{t+1} is the correct stochastic discount factor for pricing the assets under consideration. As made clear in Hansen and Jagannathan (1991) and Hansen and Jagannathan (1997), in an incomplete market, the presence of multiple admissible stochastic discount factors makes this claim difficult to verify.

Nonetheless, we proceed by estimating a stochastic discount factor that is implied by a measure of the market portfolio. Coskewness and cokurtosis of returns with the market portfolio have been investigated in Harvey and Siddique (2000) and Dittmar (2002), and the authors find that these measures improve upon pricing of assets relative to the Fama and French (1993) three factors. Moreover, the notion of residual skewness and kurtosis rests on the idea of measurement relative to some diversified portfolio, presumably the tangency portfolio of aggregate wealth. While pricing models that are alternative to an extended CAPM as investigated in Harvey and Siddique and Dittmar implicitly propose such a portfolio, we do not have options traded on these portfolios, rendering retrieval of risk-neutral probability measures difficult. However, given the presence of index options, we have a measure of this portfolio in the context of a nonlinear CAPM. Thus, we proceed by using a market implied stochastic discount factor, with the caveat that our tests in this section may represent a test of model specification rather than a test of the presence of idiosyncratic skewness and kurtosis premia.

4.1 Estimating an implied stochastic discount factor

We begin by extracting an estimate of a stochastic discount factor from a benchmark market portfolio, the S&P 500 Index. If we assume that the market portfolio and its options are priced correctly, then the relation between the risk-neutral and physical density functions for the market for each state, s, can be expressed as:

$$M(s, t, t+\tau) * P(s, t, t+\tau) = e^{-r(t,t+\tau)} Q(s, t, t+\tau) \tag{11}$$

where $M(s, t, t+\tau)$ is the stochastic discount factor from time t to $t+\tau$, $P(s)$ is the physical density function for the market portfolio over the same period, and $Q(s)$ is the *ex ante* risk-neutral density function for the market portfolio implied by the options market. Thus, given estimates of the densities P and Q, we can construct a market stochastic discount factor.

To calculate an estimate of $M(s)$, we first compute the first four moments of the market's risk-neutral and physical density. The risk-neutral moments are calculated using the same method as individual securities, using S&P 500 index option prices in place of individual security option prices. That is, we first calculate equations (5) - (7) for OTM S&P 500 index options, using options closest to $\tau = 1$ month, 3 months, 6 months, and 12 months to maturity.

Physical moments are calculated by using historical data to generate sample analogues of the physical variance (VAR_P), skewness ($SKEW_P$), and kurtosis ($KURT_P$) of the underlying market return distribution. A number of issues arise in using historical sample data to measure conditional moments. First, Foster and Nelson (1996) address the question of optimal sample estimators for time-varying volatility, and suggest that reasonable estimates can, under most circumstances, be obtained with a calendar year of past daily returns data. There is less guidance on the appropriate window to use in calculating conditional higher moments, as much of the literature on sample moment estimators has focused on volatility. In their empirical work, Bakshi, Kapadia and Madan (2003) also note that skew and kurtosis may be underestimated using short windows. We therefore use a four-year period to estimate our moments, consistent with the length of historical returns data used in Jackwerth (2000) and Brown and Jackwerth (2001).

A second issue that arises is whether the sample moments can be viewed as conditional. In our application, we opt to use a four-year sample to estimate the conditional moments as of 1/31/96, and hold this physical distribution constant over our analysis period. We do so to provide a conservative view of the degree of time variation in conditional moments. In Section 5.1, we examine the sensitivity of our analysis to these assumptions. Specifically, we allow the moments to roll through time, so that in each period, we recalculate the physical moments, and thus the physical distribution. Additionally, we consider a parametric assumption for the moments, allowing them to follow an autoregressive process. We discuss these results further in Section 5.1, but note here that the qualitative implications of our analysis are unchanged.

Finally, estimation of the physical distribution requires a specification of the conditional mean of the S&P 500 return. Jackwerth (2000) suggests adding the historical risk premium of 8% to the risk-free rate observed at time t. In our analysis, we follow his suggestion and use the annualized yield on a 90-day Treasury bill, obtained from the Federal Reserve H.15 report, as our measure of the risk free rate. We experiment with alternative values of the risk premium and obtain similar results.

Once moments for both the risk-neutral and physical distribution are generated, the second step of the procedure involves estimating the density functions of both distributions using the method described in Eriksson, Forsberg and Ghysels (2004). This procedure uses the Normal Inverse Gaussian (NIG) family to estimate an unknown distribution of random variables. As they note, the appeal of the NIG family of distributions is that they can be completely characterized by the first four moments. As a consequence, given the first four moments, one can "fill in the blanks" to obtain the entire distribution and, as they show, the method is particularly well-suited when the distribution exhibits skewness and fat tails, as it does in the returns distributions which we examine in this application. Having the market risk-neutral and physical distributions approximated with an NIG distribution, we use two methods to estimate the stochastic discount factor.

In the first method, we simply use equation (11) to solve for $M(s)$ as the discounted ratio of the risk-neutral probability density function to the physical density function over a range of implied relative wealth (return) levels. We call the resulting stochastic discount factor M^*. In the second method, we begin with M^* and employ an additional step. We parameterize the stochastic discount factor from the first step by projecting it onto a polynomial in relative wealth levels. By controlling the form of the polynomial, we can force the stochastic discount factor to include (or exclude) sequential higher moments, allowing us to examine their incremental effect on the calculation of risk-adjusted returns. For example, the stochastic discount factor M^{VAR} includes only linear returns, while M^{SKEW} includes linear and squared terms (similar to that used in Harvey and Siddique (2000)) and M^{KURT} includes linear, squared and cubic terms (as in Dittmar (2002)). These stochastic discount factors more clearly indicate the role that co-moments with the aggregate market play in determining pricing.

Using each of the four stochastic discount factors, we calculate alphas following Chen and Knez (1996), who characterize pricing errors as:

$$\hat{\alpha} = \frac{1}{T}\sum_{t=1}^{T}\hat{M}(t,t+\tau)r(t,t+\tau). \tag{12}$$

The variable $r(t, t+\tau)$ represents overlapping τ-period returns on the (zero-cost) High-Low, or factor mimicking portfolios, for volatility, skewness, and kurtosis. As noted in Chen and Knez, under the null of zero pricing errors, $\alpha = 0$. As a consequence, we perform univariate tests for the null hypothesis using Newey-West (1987) standard errors. While the NIG class is versatile (e.g., as Eriksson, Forsberg and Ghysels (2004) note, its domain is much wider than Gram-Charlier or Edgeworth expansions), there are some restrictions on its use. In particular, the parameters of the NIG approximation may become imaginary and so the distribution cannot be computed. This constraint does not arise in the case of 3- and 12-month to maturity options, and arises in only one month for the 6-month maturity options. However, this condition is frequently violated in the case of 1-month to maturity options. As a result, we compute stochastic discount factors using only 3-, 6-, and 12-month maturity options.

Our data cover the period June, 1996 through December, 2005; consequently, we have 115 monthly observations for the three-month stochastic discount factor, 112 observations for the 6-month stochastic discount factor and 106 observations for the 12-month stochastic discount factor. The number of Newey-West lags used to compute standard errors reflects the number of overlapping months in each sample; for example, 12 lags are used in computing standard errors for the 12-month stochastic discount factor.

4.2 Comparing stochastic discount factors

The time series average of the four stochastic discount factors which we estimate are presented, over the range of possible market returns and the entire sample period, in Figure 1. For brevity, we focus on options closest to twelve months to maturity; results are qualitatively similar for 3- and 6-month maturities. In Part A of Figure 1, we present the four pricing kernels over the full support; in Part B, we present the three polynomial approximations M^{VAR}, M^{SKEW} and M^{KURT} over a partial support to better illustrate the differences over this range.

The linear stochastic discount factor M^{VAR} is downward sloping throughout its range, as is M^{SKEW}. The cubic stochastic discount factor, M^{KURT}, declines through most of its support, deviating only at extremely high and low values for the return on the market portfolio. These results are generally consistent with the behavior of investors who have declining relative risk-aversion. In contrast, note that the non-parametric stochastic discount factor M^* presented in the top graph has a segment in the mid-range of the graph which is increasing. Although an upward sloping segment of the stochastic discount factor implied from option prices is consistent with the evidence in Jackwerth (2000) and Brown and Jackwerth (2001), it is, as these papers point out, a puzzle–it suggests that the representative investor may be risk-seeking over the upward sloping range. Brown and Jackwerth (2001) examine several possibilities for this behavior. Although we do not focus specifically on this puzzle in the current paper, it is worth noting that we obtain a similar result despite the fact that our sample period does not overlap with the sample used in the Brown and Jackwerth (2001) paper, and the estimation methods used to estimate both the risk-neutral distribution and physical distribution are different. In addition, we observe the upward sloping segment over all three maturities (3-, 6- and 12-month maturity options) we examine. Thus, the empirical evidence suggests that the observation of an upward sloped segment in the non-parametric stochastic discount factor implied by option prices is robust to both sample and method. Moreover, the range over which Brown and Jackwerth (2001) observe their upward-sloping segment of the stochastic discount factor, at approximately 0.97 to 1.03, is associated with an upward-sloped segment in our estimation as well. [7]

Although the behavior of the polynomial approximations of M exhibit clear differences from the non-parametric discount factor M^*, the 'fit' of the polynomial approximations is reasonable; the average R^2's of M^{VAR}, M^{SKEW} and M^{KURT} are 5.5%, 13.2% and 16.3%, respectively. In the next section, we examine the implications of the estimated empirical stochastic discount factors for investors' expectations of the payoffs to individual securities, and consequently to the moment-sorted portfolios in Table 2.

4.3 Risk-adjusted returns

In Table 4, we report estimates of alphas calculated from each of the stochastic discount factors estimated above using options closest to 3, 6, and 12 months to maturity.[8] The alphas are calculated for each of the Hi-Lo moment-sorted portfolios (volatility, skewness and kurtosis) using equation (12). The results suggest that idiosyncratic skewness is important, even after allowing for the effects of higher moments on the stochastic discount factor. Specifically, the alphas for the skewness sorted portfolios have p-values of approximately 12% for 3-month options, and at the 5% level or better for 6- and 12-month options. The alphas related to skewness are economically significant as well, ranging from 54 to 64 basis points per month. In contrast, the alphas related to volatility are not statistically significant in any maturity bin, for any specification of the stochastic discount factor. The alphas for the kurtosis-sorted portfolio are marginally significant in the shortest maturity bin, but are not significant in the samples of either 6- and 12-month options. As with volatility, these results are not sensitive to the stochastic discount factor used to calculate alphas. Thus, although we observed some differences in the previous section between the stochastic discount factors M^*, M^{VAR}, M^{SKEW} and M^{KURT}, the inferences on residual returns are unaffected by this choice.

[7] Golubev et al. (2008) report a similar pattern of the pricing kernel using German DAX index data, and propose a statistical test for monotonicity. Using their test they find statistically significant against monotonicity; hence, their results also provide support for the presence of upward sloping segments.

[8] In the case of 6-month options, the NIG approximation assumptions were violated in only one month. This month is excluded from the calculations.

The residual importance of idiosyncratic skewness is consistent with models, such as Barberis and Huang (2004), and Brunnermeier, Gollier and Parker (2007), which suggest that investors have a preference for skewness in individual securities above and beyond their contribution to the co-skewness of the portfolio. It is also consistent with the empirical evidence in Mitton and Vorkink (2007), who find a relation between the skewness in individual securities in individuals' brokerage accounts and subsequent returns.

Our results do not necessarily imply that the alpha, or residual return, is an arbitrage profit. The estimates of the stochastic discount factor used to construct α control only for non-diversifiable risk (including the risk of higher co-moments) in the context of a well-diversified portfolio. If investors have a preference for individual securities' skewness, they may, as in Brunnermeier et al., hold concentrated portfolios and push up the price of securities which are perceived to have a higher probability of an extremely good outcome. As a consequence, the lower subsequent returns of high-skew stocks may represent an equilibrium result.

5 Implied Physical Probability Distributions

To this point, we have focused on the estimation of risk-neutral moments, and the relation of these moments to returns. However, the models that consider the effects on expected returns of skewness and fat tails in individual securities' distributions deal with investors' estimates of the physical distribution. Given an estimate of the stochastic discount factor, and risk-neutral distributions of individual securities, we can construct a market-based estimate of individual securities' physical distributions that does not rely on historical data. That is, we can directly estimate investors' expectations regarding the returns distributions of underlying equities. To our knowledge, this is the first time that market data have been used to construct an *ex ante* estimate of investors' subjective probability estimates. Since papers such as Brunnermeier, Gollier and Parker (2007) and Barberis and Huang (2004) are models in which investors have biased beliefs, we also compare this *ex ante* estimate of subjective probabilities to more traditional *ex post* estimates of distributions constructed from historical returns.

Specifically, we take the stochastic discount factors M^*, M^{VAR}, M^{SKEW} and M^{KURT} constructed from the market portfolio and its options, and, using equation (11), and individual firm options, reverse engineer an estimate of the underlying security's physical probability distribution. That is, for each security, i, we compute

$$P_i\left(s,t,t+\tau\right) = e^{-r(t,t+\tau)}\frac{Q_i\left(s,t,t+\tau\right)}{M(s,t,t+\tau)}. \tag{13}$$

For every firm i, we compute risk-neutral moments using daily option prices and equations (5) through (7). For each horizon τ, we use the risk-neutral moments and the NIG approximation to compute the risk-neutral density Q_i; using equation (11) and each of the stochastic discount factors that we have computed, we calculate implied physical distributions for firm i, for quarter t and horizon τ. We examine the differences in this measure of investors' *ex ante* distributional beliefs, and implications for moments of returns, across securities and across time.

In order to consider differences across firms, we aggregate securities into industry portfolios, using the ten industry groupings available on Kenneth French's website. We assign every individual security for which we can estimate risk-neutral moments into one of these industry groupings. The Utilities portfolio had very few firms in our sample, and some months were missing observations; consequently, we eliminate that industry portfolio as well as "Other," and report results for eight out of ten of the industry portfolios. As results across the three polynomial approximations and the NIG approximation to the stochastic discount factor exhibit little difference, we present results only using the NIG approximation, or M^*. For brevity, we also report results only for τ equal to 12 months.

For each industry, we present the equal-weighted imputed physical distribution in Figure 2 at intervals in our sample period. That is, at the end of each calendar quarter, we compute the industry physical distribution by averaging over the densities of the firms in the industry.[9] We then average the industry density over the sample depicted in each figure. The intervals presented are selected to accord roughly with interesting economic events (the Asian crisis, the 'bubble' period, the recession of 2000-2001 and the recovery); although the intervals we present are chosen with perfect hindsight, recall that the risk-neutral moments at any time t are *ex ante* in nature. For comparison, at each interval we present an estimate of the distribution taken from four years of historical data ending at time t for that industry portfolio.

The implied physical distributions constructed from options market data, and the estimated pricing kernel, appear much more stable than the distributions estimated from rolling historical data. For example, using historical data generates negative mean returns for three out of eight industries (Telecom, Tech, and Durables) in the fourth subperiod (Q103-Q405). This is clearly an artifact of the inclusion of the market downturn in the historical time series. In addition, skewness estimates based on historical data are much more variable, and the distributions tend to be left-skewed; approximately 70% (23/32) of the skewness estimates across subperiods and industries are negative. Although we do not report Sharpe ratios calculated

[9]Ideally, we would recover risk neutral moments for each industry using options on the industry indices. Unfortunately, since these contracts are not available, we employ averaging as a compromise procedure.

from historical distributions, they are extremely variable, ranging from -0.36 to 1.32. Overall, using historical data to generate estimates of investors' subjective probabilities generates distributions which are highly sensitive to prior events, and have very different implications for investors' opportunity sets.

In Table 5, we present estimates of the first four moments of the implied physical distribution for each industry for the full period and for the same intervals presented in Figure 2; these estimates are constructed by integrating over market states. We also present estimates of the Sharpe ratio for each industry portfolio. There are several striking results in Table 5. First, the Sharpe ratios are comparatively stable, ranging from 0.07 to 0.26. We do, however, observe a sharp increase in the *ex ante* estimates of the Sharpe ratio through our sample period. In the last interval (03Q1-05Q4), Sharpe ratios in every industry grouping are at least double what they are in the earliest interval (96Q2-98Q2). Thus, Sharpe ratios in the pre-crash period are significantly lower than Sharpe ratios in the recovery. Second, the skewness calculated from *ex ante* data is significantly higher than that observed from historical returns; in fact, it is positive for every industry in every interval, varying from 0.30 to 0.51. This may be evidence of investors' biased beliefs, such as the optimistic bias that Brunnermeier et al. (2007) discuss. In contrast, the differences in historical and implied kurtosis measures is smaller – although the implied kurtoses tends to be lower across industries than the historical estimates, the historical estimates are lower in 2 (out of 8) cases, and the magnitudes of the implied and historical kurtoses are quite similar. In fact, the average difference in the two measures is only 1.05, or approximately 22% of the average implied kurtosis across industries.

Consistent with our results in Table 2, we observe a strong negative correlation between skewness and the Sharpe ratio across industries–that is, industries which are perceived as more likely to have extreme positive outcomes also have lower *ex ante* Sharpe ratios. This correlation is very strongly negative throughout all four intervals examined, varying between -0.75 and -0.91, and averages -0.83 over the entire sample period. This result suggests that investors are trading off traditional risk-reward ratios for the likelihood of extreme 'good news', and is consistent with a preference for idiosyncratic skewness. However, the surprising result in Table 5 is that firms in the Technology portfolio have the *lowest* skew (and the highest Sharpe ratio); this relation also holds throughout the sample period. In contrast, firms in the Durables portfolio have the highest skew and the lowest Sharpe ratio. And, the difference in skewness across these portfolios is large–for example, the percentage increase in skewness from the Tech portfolio to the other industry groupings varies from 15% to 72% across the subperiods examined.

Thus, we find little evidence in implied physical probability distributions that the high prices of technology firms during the Internet bubble period is related to investors' expectations that these firms had a relatively high chance of extremely good outcomes. In fact, the comparatively high *ex ante* Sharpe ratios of these firms compared to firms in other industries suggests that investors believed that these firms were good 'mean variance' bets.

Two cautions are worth emphasizing. First, it may be that our requirement that firms have options traded on them prevents us from sampling the youngest, most highly skewed firms, particularly in those industries whose composition is changing the most rapidly. While the selection bias associated with option trading, and the use of option data to measure skewness, does not seem significantly larger than the selection bias associated with data requirements related to the use of historical returns to measure skewness, it is difficult to assess the extent of this selection bias for either method. We note simply that while the technology portfolio we analyze includes many 'established' firms such as IBM, Cisco and Microsoft, it also includes relative newcomers such as Amazon, Iomega, JDS Uniphase, Real Networks, Xilinx, etc., whose valuations during the 'bubble period' were quite high.

More importantly, the use of options data to infer higher moments limits us not just in the cross-section, but in the horizon over which we can estimate investors' expectations. If investors' expectations of moments over horizons longer than one year (the longest horizon for which we have data) are both relevant for prices, and significantly different from the shorter-horizon moments that we have estimated, then our results may be incomplete and/or our inferences may be incorrect. For example, if investors believed that technology stocks' extreme payoffs would occur over, say, five year intervals, then differences in five-year skewness may potentially explain high valuations.[10] While we cannot rule this out, it is worth pointing out that the relation between *shorter* horizon skewness and Sharpe ratios in Table 5 is remarkably stable across all four intervals we examine. Since these intervals include the 'pre-bubble' and 'post-bubble' periods, any separate effect of longer horizon skewness would suggest a very marked term premium in the skew which differs dramatically across these intervals.

5.1 Robustness checks

The implied physical distributions in the previous section are constructed with a combination of forward-looking data from option prices, and an estimate of the market's physical distribution estimated from historical market returns data. Since this use of historical market returns is the only instance where *ex post* data are used, we explore the sensitivity of our results to different choices of the historical record to estimate the market's underlying returns distribution. First, we maintain the length of the four-year window, but allow the window to

[10]We are grateful to Paul Pfleiderer for an analysis of a setting in which this situation could arise.

roll forward with the corresponding risk-neutral distribution obtained from option prices; we still require that the four-year window end before the risk-neutral moments are calculated, by ending the window in quarter t-1 relative to the option data used to estimate moments. Second, to ensure that the window includes relatively rare 'extreme' events, we lengthen the window to 46 years, and use the period from 1950 to 1995 to estimate the market's physical returns distribution. Third, we use a longer window of data and estimate a time series forecast of each of the three higher moments (volatility, skewness and kurtosis), using these forecasts as estimates of investors' expectations regarding the corresponding physical moments. More specifically, for each quarter t, we use quarterly data beginning in March 1953 and ending in quarter $t-1$ to estimate a separate AR(1)process for volatility, skewness and kurtosis, where the (quarterly) realized moments for this estimation are constructed from the previous 250, 500 and 750 days of daily market returns, respectively.

5.1.1 Implied Risk Aversion Measures Across Market Estimates

To compare the alternative estimates of the market's physical distribution, we combine these different estimates with the options-based estimate of the market's risk-neutral distribution to estimate an aggregate risk-aversion measure. In addition to highlighting the differences, if any, in inferences related to the market's physical distribution, these estimates of risk-aversion can serve as a diagnostic on the plausibility of our estimates of both the physical and risk-neutral distributions.

Following Jackwerth (2000) and Leland (1980), we estimate absolute risk aversion as:

$$RA = P'(s)/P(s) - Q'(s)/Q(s) \tag{14}$$

where $P(s)$ is the investors' subjective, or physical, distribution and $Q(s)$ is the risk-neutral distribution for the market; P' and Q' represent first derivatives. As in the previous sections, we use the Bakshi et al. algorithm, and the NIG approximation, to estimate the risk-neutral distribution, while we use the four historical market return samples, including the AR(1) estimation of moments, and the NIG approximation, to estimate the market's physical distribution. In addition to choosing different historical periods and methods to estimate variance, skewness and kurtosis, we explore the effect of setting the market risk premium at different levels. The results are presented in Table 6 and Figure 3; for brevity, we report levels of risk-aversion only for an 8% risk premium, and a risk-neutral distribution estimated from 12-month maturity options.

With an 8% risk premium, the estimated level of risk aversion is fairly reasonable, varying between 6.3 and 19.9, depending on the interval and the market moments used to generate the physical distribution. Although the level of the risk-aversion changes somewhat depending on the historical returns used to estimate moments, and the level of risk premium assumed, the pattern of the changes in estimated risk aversion through the sample period remains remarkably constant. Specifically, it declines in the middle of the sample, reaching its lowest levels in late 2002 and early 2003; subsequently, estimated risk-aversion increases sharply through the end of the sample period in 2005.

Overall, we find relatively little difference in estimates of aggregate risk-aversion among the different estimates of the market physical distribution which we use. However, there does appear to be evidence in the aggregate market that investors' attitudes towards risk change sharply, with risk-aversion at relatively low levels at the height of the bubble, and increasing thereafter. The evidence that market events can change investors' attitudes towards risk is consistent with the evidence in Jackwerth (2000) that the market crash of 1987 dramatically changed estimates of risk-aversion. The evidence that investors' attitudes toward risk change over time also increases the advantage of using *ex ante*, rather than *ex post*, data to generate estimates of abnormal returns, as well as moments.

5.1.2 Implied Physical Distributions Across Market Estimates

The evidence above indicates that different estimates of the market's physical distribution generate fairly similar levels and patterns in estimates of aggregate risk-aversion. When we use these alternative estimates of the market's physical distribution to generate implied physical distributions across industries, the results, not surprisingly, are also fairly similar. However, there are some differences worth mentioning. While the longer estimation window of 1950-1995 generate similar estimates of volatility to those of the original 1992-1996 sample period, it generates higher estimates of skewness and kurtosis across all industries. The AR(1)estimation, in contrast, generates lower estimates of skewness and kurtosis.

There is one other interesting feature of the physical distributions imputed when the AR(1) estimation is employed on moments which is noteworthy. In the second subperiod, which extends from 9/98 through 3/2000, the implied skewness of all the industries in our sample increases sharply–in fact, on average, the skewness in each industry increases by a factor of 4 from the prior interval. Thus, there is some evidence in the data that investors perceived the distribution of payoffs as being relatively right-skewed. However, this increase in skewness is observed across all industries, and not just in the industries associated with 'bubble' pricing.

Overall, however, while the implied physical distributions for industries which we estimate using these alternate estimates of the physical market distribution change, the cross-sectional inferences we draw remain the same: technology stocks have lower skewness and kurtosis than firms in other industries. Thus, while skewness affects prices of securities, we find no evidence that the prices of technology stocks in particular were higher in our sample period because of investors' perception that they were more likely to be associated with extreme positive payoffs.[11]

6 Conclusion

We explore the possibility that higher moments of the returns distribution are important in explaining security returns. Using a sample of option prices from 1996-2005, we estimate the moments of the risk-neutral density function for individual securities using the methodology of Bakshi, Kapadia and Madan (2003). We analyze the relation between volatility, skewness and kurtosis and subsequent returns.

We find a strong relation between these moments and returns. Specifically, we find that high (low) volatility firms are associated with lower (higher) returns over the next month. This result is consistent with the results of Ang, Hodrick, Xing and Zhang (2006). We also find that skewness has a strong negative relation with subsequent returns; firms with lower negative skewness, or positive skewness, earn lower returns. That is, investors seem to prefer positive skewness, and the returns differential associated with skewness is both economically and statistically significant. We also find a positive relation between kurtosis and returns. These relations are robust to controls for differences in firm characteristics, such as firm size, book-to-market ratios and betas, as well as liquidity and momentum.

We use index returns and index options to estimate an empirical stochastic discount factor, as well as polynomial approximations of the stochastic discount factor, to control for differences in higher co-moments, and their related compensation for risk. We use these stochastic discount factors to calculate risk-adjusted returns to portfolios sorted on the basis of volatility, skewness and kurtosis, where the risk-adjustment explicitly takes higher co-moments into account. After controlling for higher co-moments, we find weak evidence that idiosyncratic kurtosis matters for short maturities, and strong evidence that idiosyncratic skewness has significant residual predictive power for subsequent returns across maturities. This suggests that investors have a preference for skewness in individual securities, which is consistent with the models of Barberis and Huang (2004) and Brunnermeier, Gollier and Parker (2007).

Finally, we use the estimated stochastic discount factors, and the risk-neutral distribu-

[11] These results are available on request from the authors.

tions calculated for individual securities, to estimate implied physical distributions for securities. We find several interesting results. First, our results suggest that implied physical distributions are much more stable than those constructed using historical data. Second, in implied physical distributions, we find evidence of a trade-off between skewness in industry portfolios and *ex ante* estimates of the Sharpe ratios for the industry. That is, our results suggest a trade-off between expected returns and higher moments, with higher (lower) traditional risk-reward measures associated with lower (higher) skewness. However, we also find that the portfolio containing technology firms has low *ex ante* physical skew and kurtosis, and a high Sharpe ratio. Consequently, while we find *both* that higher moments matter, and that investors' expectations of higher moments change through time, our results do not appear to be an explanation of bubble pricing in the Internet period.

References

[1] Abreau, D. and M. Brunnermeier, "Synchronization risk and delayed arbitrage", 2002, *Journal of Financial Economics* 66, pp. 341-360.

[2] Ait-Sahalia, Y. and M. Brandt, "Consumption and Portfolio Choice with Option-Implied State Prices", 2007, Working Paper, Princeton University.

[3] Ang, A., R. Hodrick, Y. Xing and X. Zhang, "The cross-section of volatility and expected returns," 2006, *Journal of Finance* 51, pp. 259-299.

[4] Ang, A., R. Hodrick, Y. Xing and X. Zhang, "High idiosyncratic volatility and low returns: International and further U.S. evidence," 2006, forthcoming, *Journal of Financial Economics*.

[5] Bakshi, G. and D. Madan, "Spanning and derivative-security valuation", 2000, *Journal of Financial Economics* 55, pp. 205-238.

[6] Bakshi, G., N. Kapadia and D. Madan, "Stock return characteristics, skew laws and the differential pricing of individual equity options", 2003, *Review of Financial Studies* 16, pp. 101-143.

[7] Barberis, N. and M. Huang, "Stocks as lotteries: The implications of probability weighting for security prices", 2004, Unpublished manuscript, Yale University.

[8] Bates, D., "The Crash of '87: Was it expected? The evidence from options markets", 1991, *Journal of Finance* 46, pp. 1009-1044.

[9] Battalio, R. and P. Schultz, "Options and the bubble", 2006, *Journal of Finance* 61, pp. 2071-2102.

[10] Bekaert, G. and J. Liu, "Conditioning Information and Variance Bounds on Pricing Kernels", 2004, *Review of Financial Studies* 17, 339-378.

[11] Bekaert, G. and M. Urias, "Diversification, Integration, and Emerging Market Closed-End Funds", 1996, *Journal of Finance* 51, pp. 835-869.

[12] Boyer, B., T. Mitton and K. Vorkink, "Expected Idiosyncratic Skewness", 2008, Working Paper, Brigham Young University.

[13] Brunnermeier, M.K., C. Gollier and J. Parker, "Optimal Beliefs, Asset Prices, and the Preference for Skewed Return", 2007, Working Paper.

[14] Campbell, J.Y., M. Lettau, B.G. Malkiel and Y. Xu,“Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk”, 2001, *Journal of Finance* 56, pp. 1-43.

[15] Carhart, M., 1997, “On persistence in mutual fund performance,” *Journal of Finance* 52, pp. 57-82.

[16] Chabi-Yo, F., E. Ghysels and E. Renault, “Disentangling the Effects of Heterogeneous Beliefs and Preferences on Asset Prices”, 2006, Discussion Paper UNC.

[17] Chen, Z. and P. Knez, 1996, “Portfolio Performance Measurement: Theory and Applications,” *Review of Financial Studies* 9, 511-555.

[18] Chernov, M. and E. Ghysels, “A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation”, 2000, *Journal of Financial Economics* 56, pp. 407-458.

[19] Daniel, K., M. Grinblatt, S. Titman, and R. Wermers, 1997, “Measuring mutual fund performance with characteristic based benchmarks,” *Journal of Finance* 52, 1035-1058.

[20] Dennis, P. and S. Mayhew, “Risk-neutral skewness: Evidence from stock options”, 2002, *Journal of Financial and Quantitative Analysis* 37, pp. 471-493.

[21] DeSantis, G., 1995, “Volitality bounds for stochastic discount factors: Tests and implications from international stock returns,” working paper, University of Southern California.

[22] Dittmar, R.F., 2002, “Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns,” *Journal of Finance* 57, pp. 369-403.

[23] Eriksson, A., L. Forsberg and E. Ghysels, (2004) Approximating the probability distribution of functions of random variables: A new approach, Discussion Paper, UNC.

[24] Fama, E. and K. French, 1993, “Common risk factors in the returns on stocks and bonds,” *Journal of Financial Economics* 33, pp. 3-56.

[25] Fama, E. and K. French, 1997, “Industry costs of equity,” *Jouranl of Finance* 52.

[26] Ferson,W. and A.F. Siegel, “The Efficient Use of Conditioning Information in Portfolios”, 2001 *Journal of Finance* 56, pp. 967-982.

[27] Ferson,W. and A.F. Siegel, “Stochastic Discount Factor Bounds with Conditioning Information”, 2003, *Review of Financial Studies* 16, 567-595.

[28] Foster, D.F. and D.B. Nelson,"Continuous Record Asympttics for Rolling Sample Variance Estimators", 1996, *Econometrica* 64, 139-174.

[29] Gallant, R., L.P. Hansen, and G. Tauchen, "Using Conditional Moments of Asset Payoffs to Infer the Volatility of Intertemporal Marginal Rates of Substitution", 1990, *Journal of Econometrics* 45, 141-179.

[30] Geczy, C., D. Musto and A. Reed, "Stocks are special too: An analysis of the equity lending market," 2002, *Journal of Financial Economics* 66, pp. 241-269.

[31] Gomes, J., L. Kogan, and M. Yogo, 2007, "Durability of output and expected stock returns," Working paper, Wharton School, University of Pennsylvania.

[32] Golubev, Y., W. Härdle, and R. Timonfeev, 2008, "Testing Monotonicity of Pricing Kernels," SFB 649 Discussion Papers, Humboldt University.

[33] Hansen, L.P. and R. Jagannathan, "Implications of Security Market Data for Models of Dynamic Economies," 1991, *Journal of Political Economy* 91, 249-265.

[34] Harvey, C. and A. Siddique, "Conditional skewness in asset pricing tests," 2000, *Journal of Finance* 55, pp. 1263-1296.

[35] Huberman, G. and E. Kandel, "Mean-Variance Spanning", 1987, *Journal of Finance* 42, pp. 873-888.

[36] Jackwerth, J.,"Recovering Risk Aversion from Option Prices and Realized Returns", 2000, *Review of Financial Studies* 13, pp. 433-451.

[37] Jackwerth, J. and M. Rubinstein, "Recovering probability distributions from option prices", 1996, *Journal of Finance* 51, pp. 1611-1631.

[38] Mitton, T. and K. Vorkink, "Equilibrium Underdiversification and the Preference for Skewness", 2007, *Review of Financial Studies* 20, pp. 1255-1288.

[39] Ofek, E. and M. Richardson, 2003, "DotCom mania: The rise and fall and fall of internet stocks," forthcoming, *Journal of Finance*.

[40] Pastor, L. and R. Stambaugh, 2003, "Liquidity risk and expected stock returns," *Journal of Political Economy* 111, pp. 642-685.

[41] Rubinstein, M., "Nonparametric tests of alternative option pricing models using all reported trades and quotes on the 30 most active CBOE option classes from August 23, 1976 through August 31, 1978," 1985, *Journal of Finance* 40, pp. 455-480.

[42] Rubinstein, M., "Implied binomial trees," 1994, *Journal of Finance* 49, pp. 771-818.

[43] Shleifer, A. and R. Vishny, "The limits of arbitrage", 1997, *Journal of Finance* 52, pp. 35-55.

[44] Siegel, J., "Irrational exuberance, reconsidered," December 6, 2006, *Wall Street Journal*.

[45] Spiegel, M. and X. Wang, "Cross-sectional variation in stock returns: Liquidity and Idiosyncratic Risk," 2006, Unpublished manuscript, Yale University.

[46] Tversky, A., and D. Kahneman, "Advances in prospect theory: Cumulative representation of uncertainty," 1992, *Journal of Risk and Uncertainty* 5, pp. 297-323.

[47] Xing, Y., X. Zhang and R. Zhao, "What Does Individual Option Volatility Smirk Tell Us about Future Equity Returns?", 2007, Working Paper, Rice University.

Table 1: Descriptive Statistics: Risk Neutral Moments

Entries to the table are the 5th percentile, median, and 95th percentiles of risk neutral volatility, skewness, and kurtosis across securities by year. We calculate the risk neutral moments following the procedure in Bakshi, Kapadia, and Madan (2003) using data on out of the money (OTM) puts and calls. We require at least two OTM puts and two OTM calls to calculate the moments. Further, we restrict attention to options with prices in excess of $0.50 for which we have at least 10 quotes per month and are not expiring within one week. Finally, we eliminate any options that violate put-call parity restrictions and lie in the extreme 1% of the distribution of the risk neutral moments. The sample consists of 3,722,700 option-day combinations over the time period January 1996 through December 2005.

	Volatility			Skewness			Kurtosis		
Year	P5	P50	P95	P5	P50	P95	P5	P50	P95
1996	11.404	24.283	42.289	-3.495	-0.449	0.601	1.386	4.713	20.592
1997	11.311	23.591	41.568	-3.834	-0.539	0.624	1.390	4.868	24.632
1998	12.283	24.533	45.381	-3.486	-0.464	0.695	1.444	5.012	22.684
1999	13.543	26.837	51.576	-3.727	-0.601	0.564	1.313	4.940	24.514
2000	16.140	30.942	57.531	-3.083	-0.562	0.511	1.344	4.682	20.318
2001	15.000	30.594	67.485	-2.959	-0.648	0.456	1.549	4.756	18.596
2002	14.119	27.659	67.315	-3.353	-0.742	0.539	1.658	5.515	22.356
2003	12.093	25.549	75.391	-4.315	-1.297	0.309	1.820	6.836	28.889
2004	10.276	24.021	68.945	-4.652	-1.399	0.398	2.040	8.239	34.943
2005	8.710	22.365	53.033	-5.164	-1.609	0.337	2.119	9.584	39.102

Table 2: Descriptive Statistics

Panels A-C present summary statistics for portfolios sorted on measures of firms' risk-neutral moments. Firms are sorted on average risk-neutral volatility, skewness, and kurtosis within each calendar quarter into terciles based on 30^{th} and 70^{th} percentiles. We then form equally-weighted portfolios of these firms, holding the moment ranking constant for the subsequent calendar quarter. Risk-neutral moments are calculated using the procedure in Bakshi, Kapadia, and Madan (2003); the options used are those closest to one, three, six, and twelve months to maturity. The first column of each panel presents mean monthly returns. The second column presents characteristic-adjusted returns, calculated by determining, for each firm, the Fama-French 5X5 size- and book-to-market portfolio to which it belongs and subtracting that return. The next three columns present the average beta, log market value and book-to-market equity ratio of the portfolio, while the next three columns present the average volatility, skewness and kurtosis of the portfolio. Monthly return data cover the period 4/96 through 12/05, for a total of 117 monthly observations.

Panel A: Volatility-Sorted Portfolios

1 Month to Maturity

Rank	Mean	Char Adj	Vol	Skew	Kurt	Beta	ln MV	BM
1	1.213	0.268	16.104	-1.363	12.888	0.890	15.703	0.368
2	0.963	0.128	24.994	-0.968	8.842	1.281	14.304	0.393
3	0.893	0.172	44.033	-1.171	6.041	1.772	13.619	0.417
3-1	-0.320	-0.096	27.929	0.192	-6.847	0.883	-2.084	0.049

3 Months to Maturity

Rank	Mean	Char Adj	Vol	Skew	Kurt	Beta	ln MV	BM
1	1.237	0.273	17.139	-1.180	10.812	0.837	15.675	0.386
2	1.061	0.190	26.458	-0.934	8.166	1.290	14.299	0.391
3	0.738	0.062	45.890	-1.203	5.993	1.828	13.648	0.402
3-1	-0.499	-0.211	28.751	-0.023	-4.819	0.990	-2.028	0.016

6 Months to Maturity

Rank	Mean	Char Adj	Vol	Skew	Kurt	Beta	ln MV	BM
1	1.215	0.227	18.770	-0.713	6.621	0.816	15.617	0.397
2	1.137	0.266	28.656	-0.576	5.480	1.287	14.336	0.393
3	0.659	0.002	47.734	-0.749	4.148	1.861	13.658	0.386
3-1	-0.556	-0.225	28.964	-0.036	-2.473	1.045	-1.959	-0.012

12 Months to Maturity

Rank	Mean	Char Adj	Vol	Skew	Kurt	Beta	ln MV	BM
1	1.237	0.227	19.353	-0.692	6.416	0.824	15.453	0.402
2	1.060	0.195	29.541	-0.615	5.544	1.291	14.350	0.391
3	0.739	0.098	49.892	-0.826	4.259	1.844	13.807	0.384
3-1	-0.498	-0.129	30.539	-0.134	-2.156	1.020	-1.646	-0.018

Table continued on next page...

Panel B: Skewness-Sorted Portfolios

1 Month to Maturity

Rank	Mean	Char Adj	Vol	Skew	Kurt	Beta	ln MV	BM
1	1.233	0.391	26.626	-2.642	16.231	1.274	15.318	0.343
2	0.886	0.088	30.095	-0.975	6.847	1.365	14.351	0.398
3	0.975	0.110	26.699	0.116	5.365	1.228	13.961	0.436
3-1	-0.257	-0.281	0.074	2.758	-10.866	-0.046	-1.357	0.093

3 Months to Maturity

Rank	Mean	Char Adj	Vol	Skew	Kurt	Beta	ln MV	BM
1	1.281	0.421	29.444	-2.530	14.303	1.262	15.283	0.354
2	0.944	0.146	31.132	-0.923	6.303	1.366	14.377	0.395
3	0.849	-0.009	27.345	0.131	4.992	1.242	13.961	0.429
3-1	-0.432	-0.430	-2.099	2.661	-9.311	-0.020	-1.322	0.076

6 Months to Maturity

Rank	Mean	Char Adj	Vol	Skew	Kurt	Beta	ln MV	BM
1	1.341	0.476	32.812	-1.717	8.891	1.260	15.276	0.378
2	0.886	0.042	31.157	-0.527	4.181	1.318	14.444	0.390
3	0.867	0.085	30.353	0.190	3.614	1.311	13.874	0.412
3-1	-0.474	-0.391	-2.459	1.907	-5.277	0.051	-1.402	0.034

12 Months to Maturity

Rank	Mean	Char Adj	Vol	Skew	Kurt	Beta	ln MV	BM
1	1.325	0.487	35.105	-1.835	9.206	1.242	15.364	0.374
2	0.888	0.027	32.013	-0.524	4.008	1.326	14.398	0.392
3	0.881	0.090	30.842	0.195	3.522	1.319	13.847	0.413
3-1	-0.444	-0.397	-4.263	2.030	-5.684	0.077	-1.518	0.039

Table continued on next page...

Panel C: Kurtosis-Sorted Portfolios

1 Month to Maturity

Rank	Mean	Char Adj	Vol	Skew	Kurt	Beta	ln MV	BM
1	1.020	0.172	35.337	-0.358	2.864	1.387	13.665	0.461
2	0.925	0.110	27.185	-0.909	6.976	1.293	14.388	0.394
3	1.137	0.309	21.876	-2.254	18.556	1.217	15.556	0.325
3-1	0.117	0.138	-13.461	-1.896	15.691	-0.170	1.891	-0.136

3 Months to Maturity

Rank	Mean	Char Adj	Vol	Skew	Kurt	Beta	ln MV	BM
1	0.938	0.109	35.972	-0.346	2.775	1.418	13.650	0.452
2	0.901	0.075	28.714	-0.865	6.482	1.310	14.401	0.391
3	1.251	0.410	24.049	-2.129	16.280	1.169	15.550	0.338
3-1	0.313	0.301	-11.923	-1.782	13.505	-0.249	1.899	-0.115

6 Months to Maturity

Rank	Mean	Char Adj	Vol	Skew	Kurt	Beta	ln MV	BM
1	0.860	0.029	37.226	-0.189	2.155	1.449	13.661	0.427
2	0.987	0.149	30.722	-0.499	4.249	1.324	14.426	0.386
3	1.213	0.383	26.522	-1.375	10.259	1.124	15.507	0.370
3-1	0.353	0.354	-10.704	-1.186	8.104	-0.325	1.847	-0.057

12 Months to Maturity

Rank	Mean	Char Adj	Vol	Skew	Kurt	Beta	ln MV	BM
1	0.857	0.019	37.729	-0.179	2.097	1.465	13.669	0.424
2	0.980	0.158	32.052	-0.497	4.103	1.330	14.394	0.388
3	1.226	0.383	28.165	-1.496	10.502	1.102	15.542	0.369
3-1	0.369	0.364	-9.564	-1.316	8.404	-0.363	1.873	-0.055

Table 3: Time Series Regressions

The table presents the results of time series regressions of excess return differentials (Hi-Lo) between portfolios ranked on risk neutral volatility, skewness, and kurtosis on the three Fama and French (1993) factors MRP (the return on the value-weighted market portfolio in excess of a one-month T-Bill), SMB (the difference in returns on a portfolio of small capitalization and large capitalization stocks), and HML (the difference in returns on a portfolio of high and low book equity to market equity stocks). The moment-sorted portfolios are equally-weighted, formed on the basis of terciles and re-formed each quarter. The table presents point estimates of the coefficients and standard errors in parentheses. Data cover the period April 1996 through December 2005 for 117 monthly observations.

Panel A: 1 Month to Maturity

	α	β_{MRP}	β_{SMB}	β_{HML}	R^2
Vol	-0.57	0.51	0.83	-0.55	0.76
	-1.77	5.34	9.23	-5.29	
Skew	-0.58	0.14	-0.05	0.55	0.27
	-1.71	1.66	-0.37	4.23	
Kurt	0.55	-0.20	-0.30	-0.51	0.28
	2.49	-3.49	-3.17	-6.32	

Panel B: 3 Months to Maturity

	α	β_{MRP}	β_{SMB}	β_{HML}	R^2
Vol	-0.56	0.56	0.89	-1.03	0.84
	-1.65	5.00	10.21	-8.92	
Skew	-0.67	0.19	-0.13	0.37	0.17
	-1.99	2.34	-0.83	2.63	
Kurt	0.62	-0.30	-0.23	-0.16	0.21
	2.45	-4.42	-1.97	-1.37	

Panel C: 6 Months to Maturity

	α	β_{MRP}	β_{SMB}	β_{HML}	R^2
Vol	-0.55	0.59	0.90	-1.22	0.85
	-1.54	5.06	10.37	-9.73	
Skew	-0.64	0.18	0.00	0.14	0.05
	-2.43	2.48	0.00	1.10	
Kurt	0.56	-0.35	-0.26	0.10	0.44
	2.38	-4.25	-2.15	0.95	

Panel D: 12 Months to Maturity

	α	β_{MRP}	β_{SMB}	β_{HML}	R^2
Vol	-0.41	0.54	0.83	-1.29	0.85
	-1.16	4.70	10.20	-10.49	
Skew	-0.62	0.20	0.06	0.11	0.07
	-2.42	2.83	0.42	0.88	
Kurt	0.62	-0.38	-0.31	0.10	0.50
	2.65	-4.65	-2.51	0.96	

Table 4: Stochastic Discount Factor Risk Adjustments

The table presents risk adjustments for the volatility, skewness, and kurtosis factor mimicking portfolios using stochastic discount factors implied by the S&P 500 risk neutral and physical densities. The table presents returns in excess of those implied by discounting using the stochastic discount factor, calculated as

$$\hat{\alpha} = \frac{1}{T}\sum_{t=1}^{T} r_t m_t$$

where r_t is the return on the factor-mimicking portfolio at time t, and m_t is the stochastic discount factor. Columns "Linear," "Quad," and "Cubic" represent discount factors obtained by projecting the density-implied discount factor onto a linear, quadratic, and cubic polynomial, respectively. Panel A presents results using overlapping quarterly returns and the discount factor implied by 3 month maturity options; Panels B and C present similar results using 6 month and 12 month horizons. Point estimates are scaled to the monthly frequency. Newey-West standard errors are presented in parentheses below the point estimates. Data in Panel A extend from June, 1996 through December, 2005 for 115 monthly observations. Data in Panel B cover the period September, 1996 through December, 2005 for 112 monthly observations. Data in Panel C cover the period March, 1996 through December, 2005 for 106 monthly observations.

Panel A: 3 Months

	NIG	Linear	Quad	Cubic
Vol	-0.405	-0.358	-0.437	-0.192
SE	(0.844)	(0.893)	(0.895)	(0.797)
Skew	-0.594	-0.620	-0.619	-0.597
SE	(0.381)	(0.401)	(0.402)	(0.377)
Kurt	0.479	0.456	0.512	0.450
SE	(0.254)	(0.254)	(0.263)	(0.265)

Panel B: 6 Months

	NIG	Linear	Quad	Cubic
Vol	-0.081	-0.285	-0.089	0.645
SE	(1.180)	(1.182)	(1.265)	(1.205)
Skew	-0.631	-0.626	-0.636	-0.631
SE	(0.283)	(0.281)	(0.293)	(0.303)
Kurt	0.345	0.403	0.367	0.118
SE	(0.321)	(0.327)	(0.335)	(0.294)

Panel C: 12 Months

	NIG	Linear	Quad	Cubic
Vol	-0.130	-0.283	-0.228	0.371
SE	(0.557)	(0.561)	(0.570)	(0.528)
Skew	-0.555	-0.544	-0.543	-0.619
SE	(0.278)	(0.265)	(0.277)	(0.321)
Kurt	0.411	0.477	0.473	0.160
SE	(0.304)	(0.302)	(0.315)	(0.290)

Table 5: Imputed Physical Moments

The table presents moments of imputed physical distributions of eight industry portfolios. Distributions are imputed by letting the physical distribution, $f^P(x,s,\tau)$ be related to the risk neutral distribution, $f^Q(x,s,\tau)$ by

$$f^P(x,s,\tau) = e^{-r_f\tau}\frac{f^Q(x,s,\tau)}{m(x,s,\tau)}$$

where $m(x,s,\tau)$ is the stochastic discount factor implied by the S&P 500 index. The risk neutral distribution is the equally-weighted risk neutral distribution across firms implied by risk neutral moments retrieved from option prices and the NIG probability density. We calculate the moments for four subperiods: 1996 Q2 - 1998 Q2, 1998 Q3 - 2000 Q1, 2000 Q2 - 2002 Q4, and 2003 Q1 - 2005 Q4.

Panel A: Mean

Subperiod	NonDur	Dur	Mfg	Energy	Tech	Telecom	Wh/Ret	Health
I	0.0739	0.0753	0.0786	0.0770	0.0917	0.0859	0.0845	0.0882
II	0.0702	0.0720	0.0792	0.0859	0.0982	0.0940	0.0818	0.0893
III	0.0825	0.0860	0.0882	0.0886	0.1093	0.0947	0.0897	0.0988
IV	0.0834	0.0747	0.0884	0.0918	0.1004	0.0858	0.0863	0.0982

Panel B: Volatility

Subperiod	NonDur	Dur	Mfg	Energy	Tech	Telecom	Wh/Ret	Health
I	0.2724	0.2787	0.2856	0.2852	0.3422	0.3119	0.3222	0.3348
II	0.2730	0.2742	0.2801	0.2959	0.3294	0.3109	0.3079	0.3221
III	0.2811	0.2925	0.2936	0.2925	0.3625	0.3282	0.3063	0.3345
IV	0.2684	0.2601	0.2799	0.2729	0.3266	0.3019	0.2905	0.3205

Panel C: Skewness

Subperiod	NonDur	Dur	Mfg	Energy	Tech	Telecom	Wh/Ret	Health
I	0.4785	0.4918	0.4696	0.4861	0.3589	0.4339	0.3871	0.3544
II	0.6171	0.5555	0.4997	0.3823	0.3612	0.3905	0.4705	0.4238
III	0.4007	0.4507	0.4061	0.3690	0.2869	0.4062	0.3901	0.3613
IV	0.4328	0.5507	0.2784	0.2019	0.2135	0.3667	0.2824	0.2511

Panel D: Kurtosis

	NonDur	Dur	Mfg	Energy	Tech	Telecom	Wh/Ret	Health
I	6.0171	5.6818	5.3311	5.2650	2.9991	4.3173	3.6724	3.1631
II	6.8163	6.3679	6.0763	4.9900	3.6533	4.4124	4.6471	4.0329
III	5.6132	5.3446	5.1082	5.0510	2.5129	3.7166	4.4912	3.4917
IV	6.7235	7.3218	5.1764	5.7343	3.2445	4.7178	4.7419	3.5187

Sharpe Ratio

	NonDur	Dur	Mfg	Energy	Tech	Telecom	Wh/Ret	Health
I	0.0770	0.0807	0.0899	0.0859	0.1136	0.1065	0.0975	0.1056
II	0.0707	0.0784	0.1008	0.1186	0.1444	0.1400	0.1010	0.1193
III	0.1721	0.1777	0.1850	0.1848	0.2054	0.1830	0.1811	0.1926
IV	0.2353	0.2059	0.2392	0.2605	0.2414	0.2109	0.2227	0.2388

Table 6: Imputed Risk Aversion

The table presents subperiod estimates of imputed risk aversion. Risk aversion is calculated using estimates of the physical and risk neutral estimates of the probability density function following Leland (1980) and Jackwerth (2000):

$$RA(s) = P'(s)/P(s) - Q'(s)/Q(s)$$

where P is the physical probability measure and Q is the risk neutral measure. We choose the state, s to correspond to an 8% market risk premium. Both probability measures are calculated using the NIG approximation in Eriksson, Forsberg, and Ghysels (2004), which takes as its arguments the mean, variance, skewness, and kurtosis of the density. Risk neutral moments are retrieved from option prices on the S&P 500 closest to 12 months to maturity. Physical moments are calculated in one of four ways. "Fixed" indicates that the moments are sample moments computed over daily returns on the S&P 500 index over the period 1/92 - 12/95. "Roll" indicates that the moments are computed over a rolling lagged four year sample period. "Roll AR(1)" moments are computed by estimating an AR(1) on the S&P 500 four year sample moments over the period 1950-1995 and forecasting next period's moment on the basis of the current four year sample moment. "Unconditional" uses the unconditional moments estimated over the 1950-1995 time period. We average moments over four subperiods: I. 1996 Q2 - 1998 Q2, II. 1998 Q3 - 2000 Q1, III. 2000 Q2 - 2002 Q4, and IV. 2003 Q1 - 2005 Q4.

Subperiod	Fixed	Roll	Roll AR(1)	Uncond.
I	16.27	15.65	15.42	17.43
II	7.99	6.32	6.55	9.14
III	9.84	6.44	6.91	10.99
IV	18.73	14.91	15.04	19.88

Figure 1: Stochastic Discount Factors

The plots depict stochastic discount factors formed using risk neutral moments of S&P 500 index options at the 12-month maturity. The plot labeled 'NIG' represents stochastic discount factors, $m(x, s, tau)$, formed as

$$m(x,s,\tau) = e^{-r_f \tau} \frac{f^Q\,(x,s,\tau)}{f^P\,(x,s,\tau)}$$

where $f(\cdot)$ is the NIG probability density function, Q denotes the risk-neutral probability measure, and P denotes the physical measure. The risk neutral measure is calculated using risk neutral moments retrieved from option prices and the physical measure using the historical moments of the S&P 500 index from 1992 through 1995. 'Linear,' 'Quadratic,' and 'Cubic' represent linear, quadratic, and cubic polynomial fits to the NIG kernel. Subfigure A depicts plots of the average stochastic discount factor for all four kernels; Subfigure B depicts the polynomial kernels over a smaller range.

Figure A

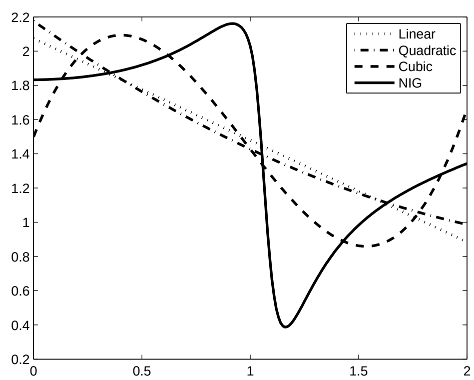

..

Figure B

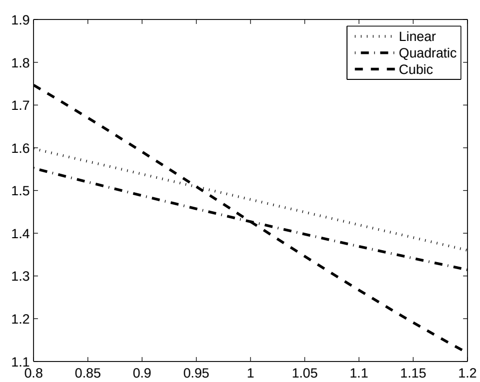

..

Figure 2: Imputed and Historical Probability Densities

The plots depict the probability densities for eight industry portfolios implied by historical and imputed moments. Historical moments are calculated from equally-weighted daily returns on each industry portfolio over the past four years, updated quarterly. Imputed moments are obtained by imputing the physical probability density for the industry portfolio using its risk neutral probability measure and the stochastic discount factor obtained from the S&P 500 index. Averages of moments over the relevant time periods are then used to calculate the NIG density function, evaluated at these moments. For each industry, we examine densities over four subperiods: 1996 Q2 - 1998 Q2, 1998 Q3 - 2000 Q1, 2000 Q2 - 2002 Q4, and 2003 Q1 - 2005 Q4.

Nondurables

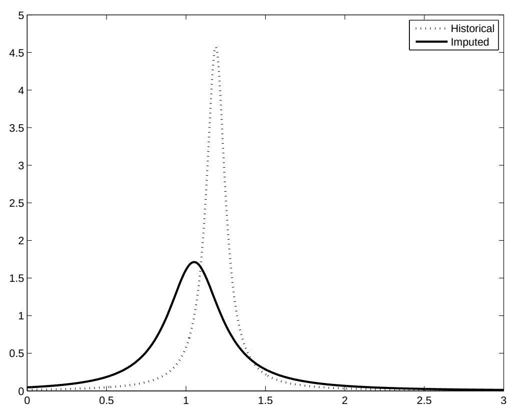

(a) 1996 Q2 - 1998 Q2

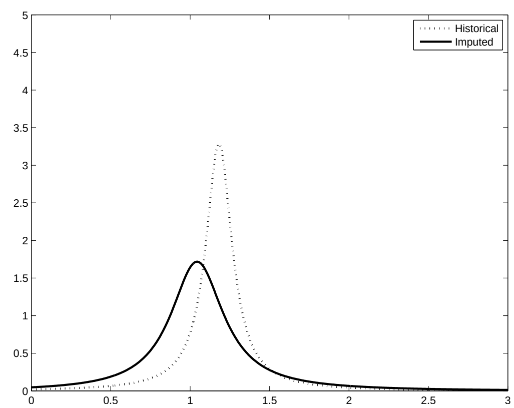

(b) 1998 Q3 - 2000 Q1

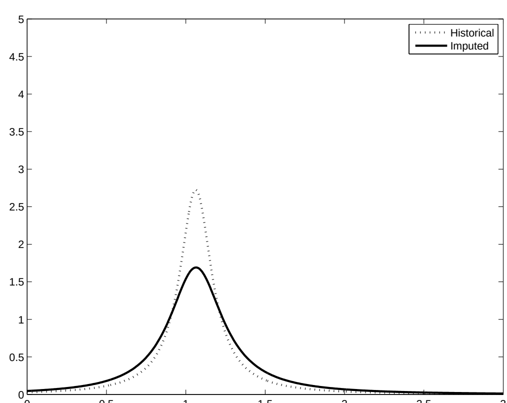

(c) 2000 Q2 - 2002 Q3

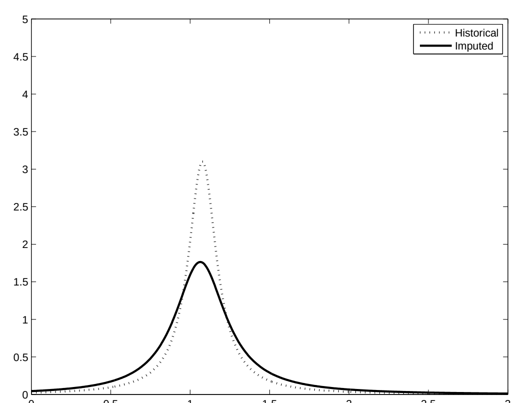

(d) 2003 Q1 - 2005 Q4

Figure continued on next page...

Durables

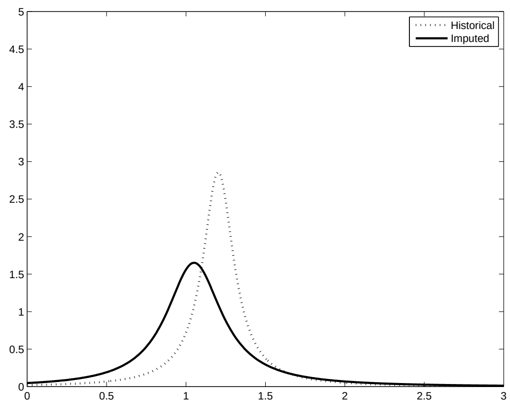

(a) 1996 Q2 - 1998 Q2

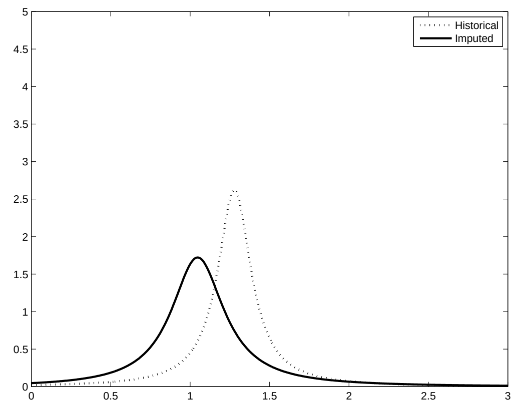

(b) 1998 Q3 - 2000 Q1

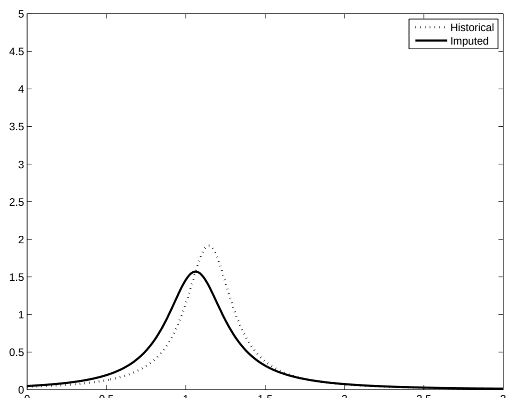

(c) 2000 Q2 - 2002 Q3

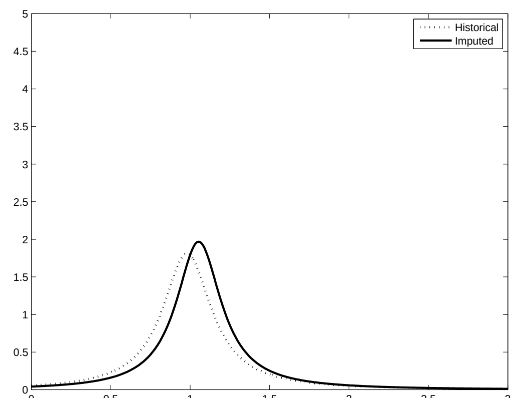

(d) 2003 Q1 - 2005 Q4

Figure continued on next page...

Manufacturing

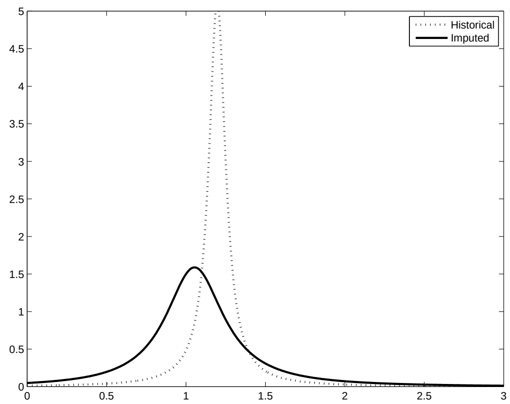

(a) 1996 Q2 - 1998 Q2

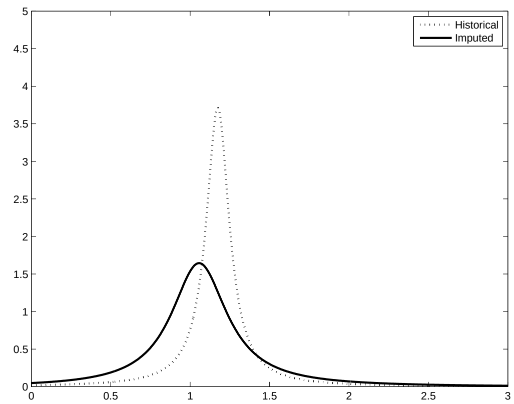

(b) 1998 Q3 - 2000 Q1

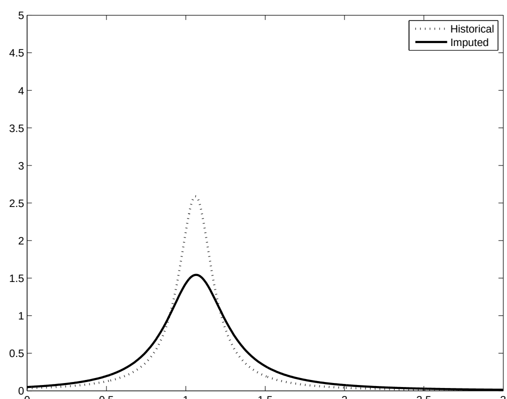

(c) 2000 Q2 - 2002 Q3

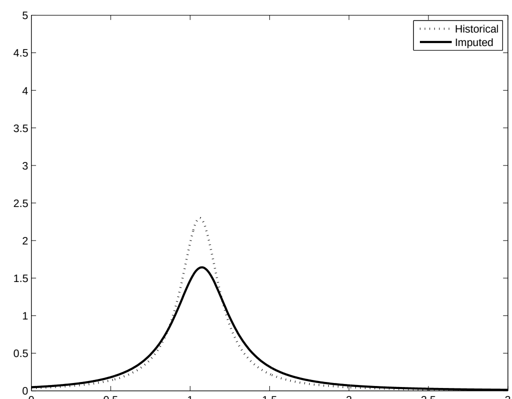

(d) 2003 Q1 - 2005 Q4

Figure continued on next page...

Energy

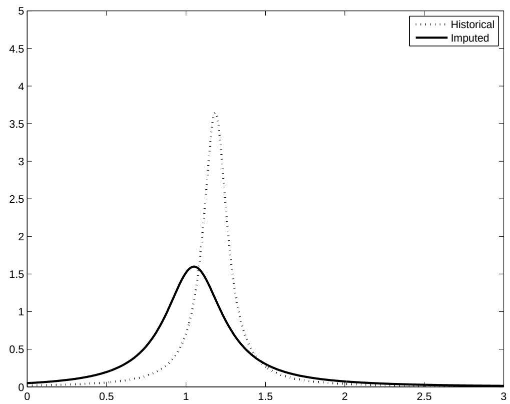

(a) 1996 Q2 - 1998 Q2

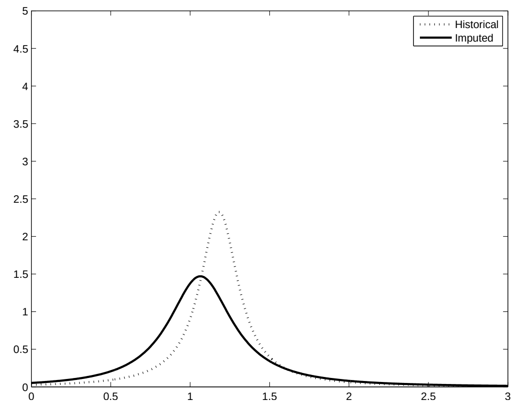

(b) 1998 Q3 - 2000 Q1

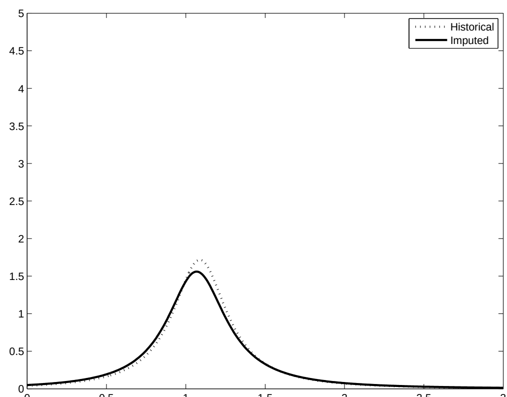

(c) 2000 Q2 - 2002 Q3

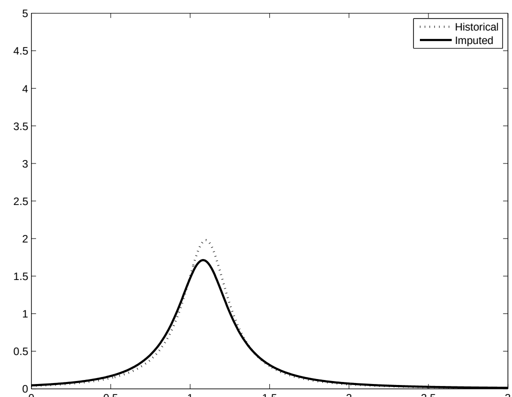

(d) 2003 Q1 - 2005 Q4

Figure continued on next page...

Tech

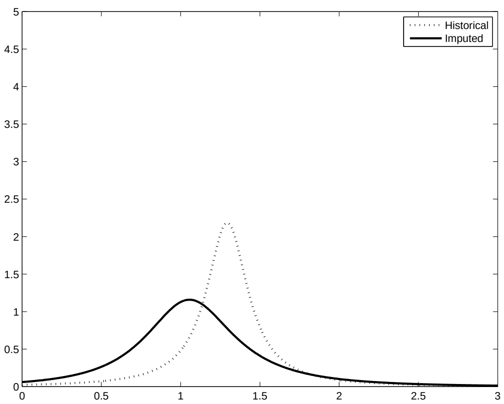

(a) 1996 Q2 - 1998 Q2

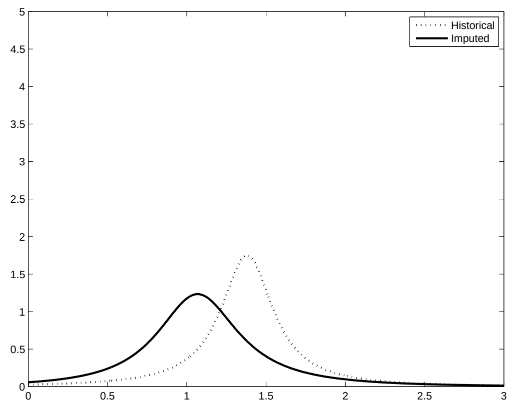

(b) 1998 Q3 - 2000 Q1

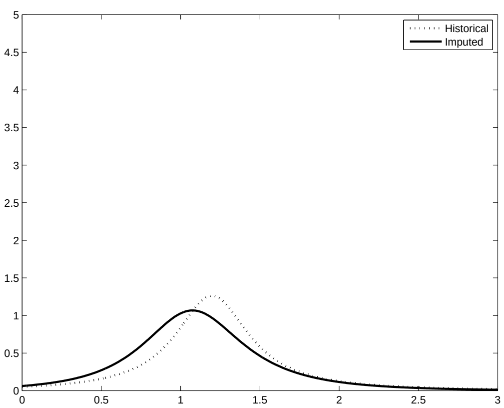

(c) 2000 Q2 - 2002 Q3

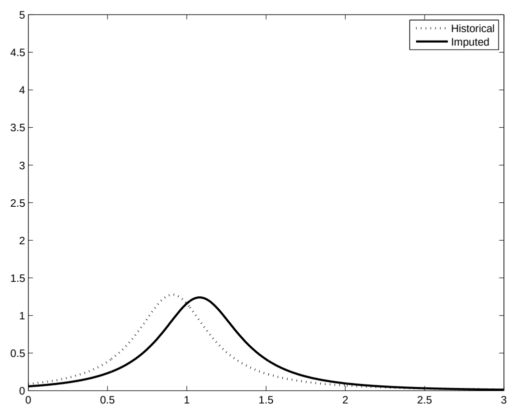

(d) 2003 Q1 - 2005 Q4

Figure continued on next page...

Telecom

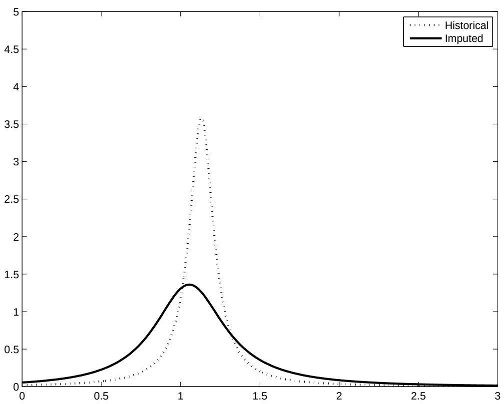

(a) 1996 Q2 - 1998 Q2

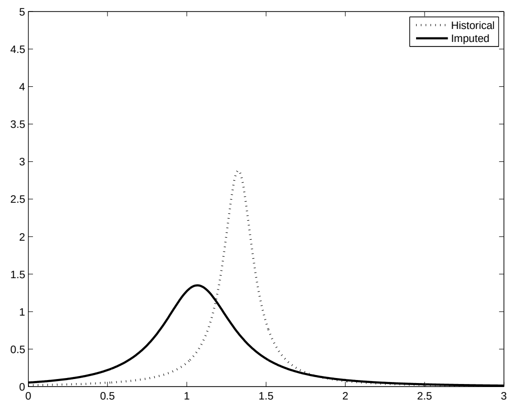

(b) 1998 Q3 - 2000 Q1

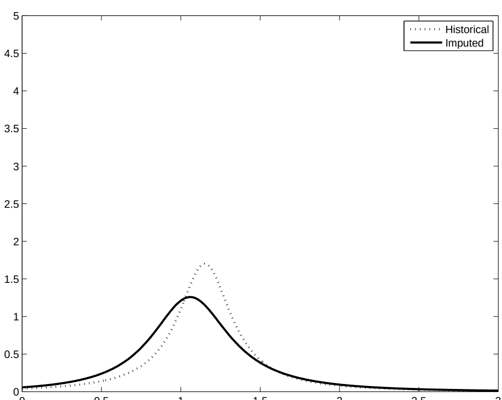

(c) 2000 Q2 - 2002 Q3

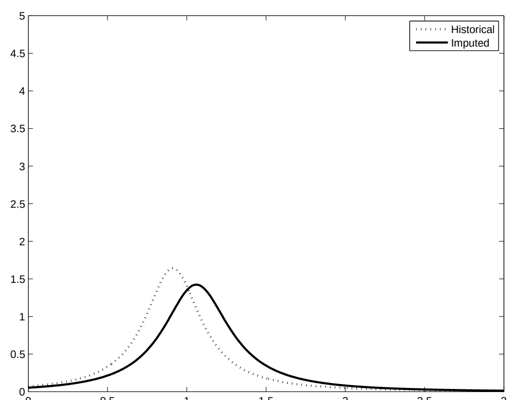

(d) 2003 Q1 - 2005 Q4

Figure continued on next page...

Wholesale/Retail

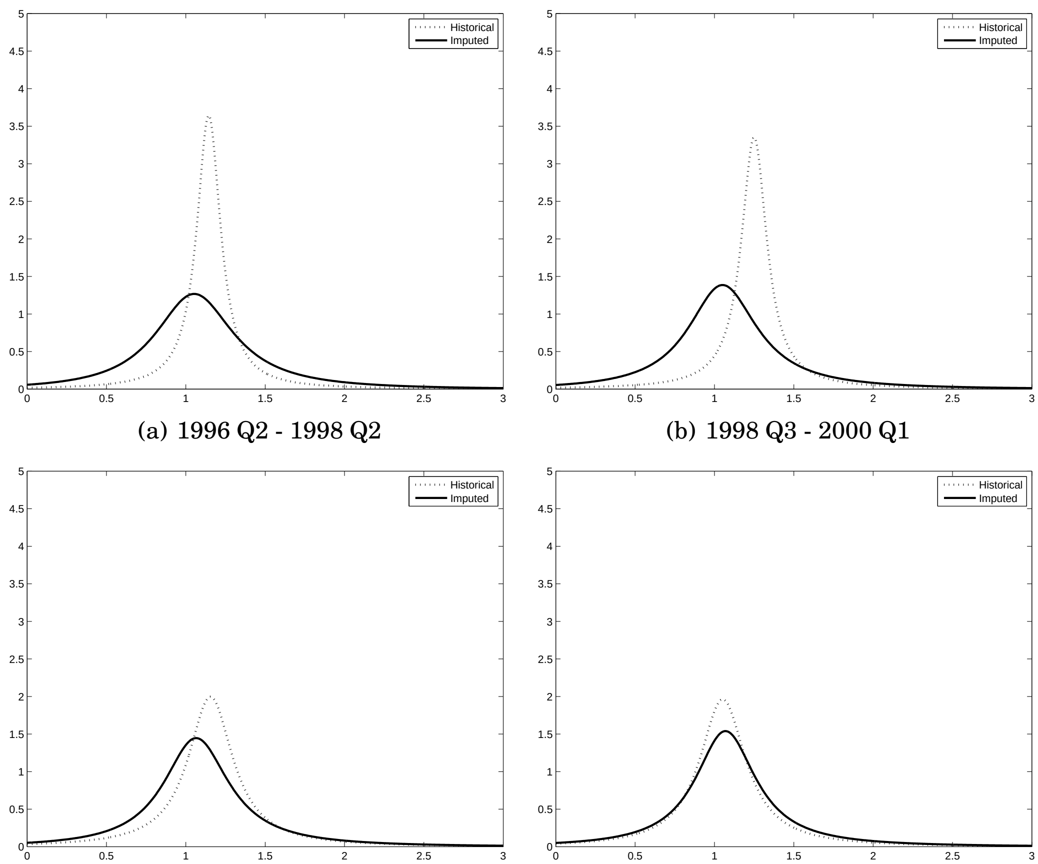

(a) 1996 Q2 - 1998 Q2

(b) 1998 Q3 - 2000 Q1

(c) 2000 Q2 - 2002 Q3

(d) 2003 Q1 - 2005 Q4

Figure continued on next page...

Healthcare

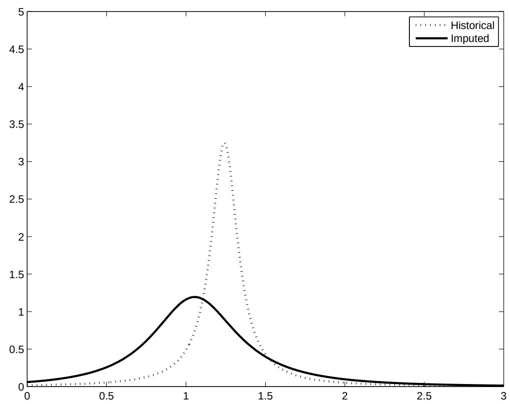

(a) 1996 Q2 - 1998 Q2

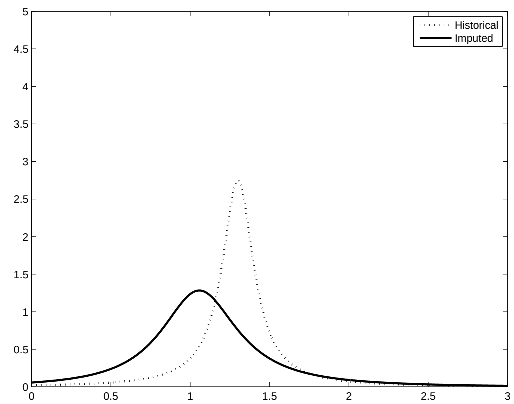

(b) 1998 Q3 - 2000 Q1

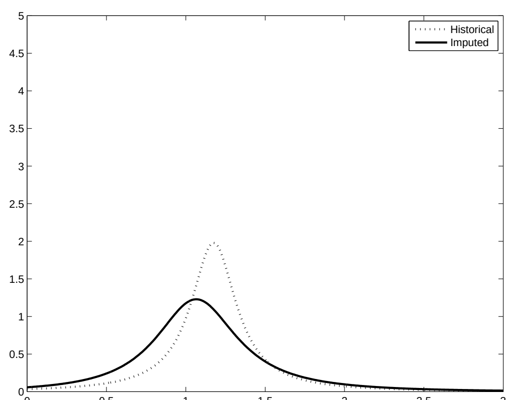

(c) 2000 Q2 - 2002 Q3

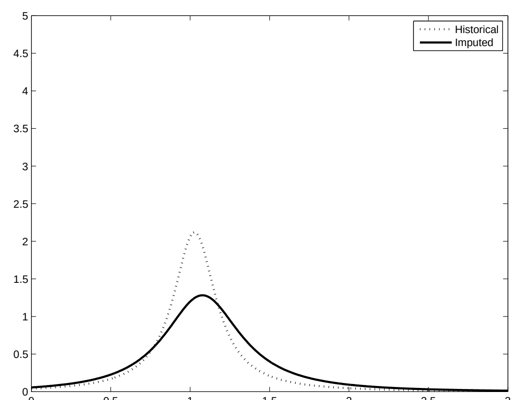

(d) 2003 Q1 - 2005 Q4

Figure 3: Risk Aversion

The plots depict the probability densities for eight industry portfolios implied by historical and imputed moments. Historical moments are calculated from equally-weighted daily returns on each industry portfolio over the past four years, updated quarterly. Imputed moments are obtained by imputing the physical probability density for the industry portfolio using its risk neutral probability measure and the stochastic discount factor obtained from the S&P 500 index. Averages of moments over the relevant time periods are then used to calculate the NIG density function, evaluated at these moments. For each industry, we examine densities over four subperiods: 1996 Q2 - 1998 Q2, 1998 Q3 - 2000 Q1, 2000 Q2 - 2002 Q4, and 2003 Q1 - 2005 Q4.

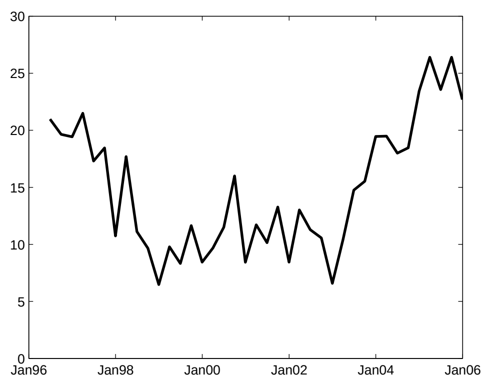

(a) Unconditional Moments

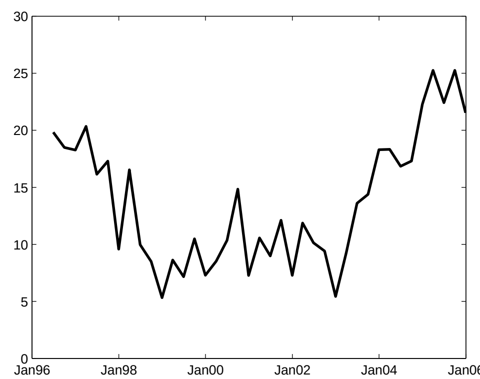

(b) Fixed Physical Moments

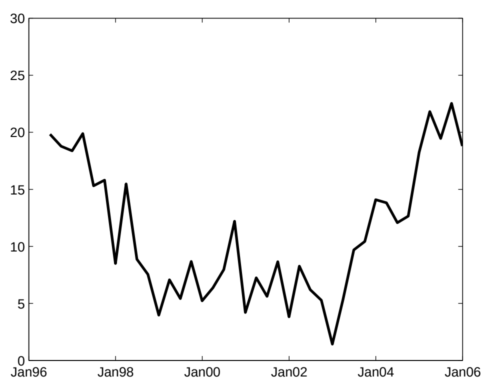

(c) Rolling Physical Moments

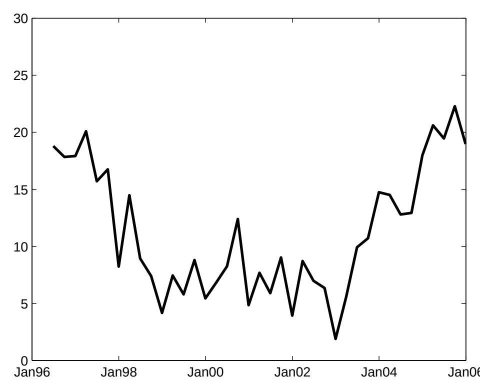

(d) AR(1) Physical Moments

Appendix

Table A1: Summary Statistics with Volume Screens

Panels A-C present summary statistics for portfolios sorted on measures of firms' risk-neutral moments. Firms are sorted on average risk-neutral volatility, skewness, and kurtosis within each calendar quarter into terciles based on 30^{th} and 70^{th} percentiles. We then form equally-weighted portfolios of these firms, holding the moment ranking constant for the subsequent calendar quarter. Risk-neutral moments are calculated using the procedure in Bakshi, Kapadia, and Madan (2003); the options used are those closest to one, three, six, and twelve months to maturity. We eliminate firms that do not have trading volume in at least one OTM put and OTM call in a calendar month. The first column of each panel presents mean monthly returns. The second column presents characteristic-adjusted returns, calculated by determining, for each firm, the Fama-French 5X5 size- and book-to-market portfolio to which it belongs and subtracting that return. Monthly return data cover the period 4/96 through 12/05, for a total of 117 monthly observations.

Panel A: Volatility-Ranked Portfolios

1 Month Maturity			3 Month Maturity			6 Month Maturity			12 Month Maturity		
Rank	Mean	Char Adj	Rank	Mean	Char Adj	Rank	Mean	Char Adj	Rank	Mean	Char Adj
1	1.331	0.395	1	1.289	0.375	1	1.297	0.405	1	1.339	0.432
2	0.861	0.012	2	1.007	0.107	2	1.026	0.134	2	0.995	0.088
3	0.972	0.226	3	0.819	0.101	3	0.782	0.049	3	0.781	0.090
3-1	-0.360	-0.169	3-1	-0.470	-0.274	3-1	-0.515	-0.355	3-1	-0.558	-0.342

Panel B: Skewness-Ranked Portfolios

1 Month Maturity			3 Month Maturity			6 Month Maturity			12 Month Maturity		
Rank	Mean	Char Adj	Rank	Mean	Char Adj	Rank	Mean	Char Adj	Rank	Mean	Char Adj
1	1.314	0.391	1	1.392	0.494	1	1.393	0.496	1	1.407	0.535
2	0.945	0.138	2	0.958	0.173	2	0.978	0.182	2	1.022	0.200
3	0.877	0.069	3	0.782	-0.099	3	0.751	-0.115	3	0.678	-0.183
3-1	-0.437	-0.322	3-1	-0.611	-0.593	3-1	-0.642	-0.611	3-1	-0.729	-0.718

Panel C: Kurtosis-Ranked Portfolios

1 Month Maturity			3 Month Maturity			6 Month Maturity			12 Month Maturity		
Rank	Mean	Char Adj	Rank	Mean	Char Adj	Rank	Mean	Char Adj	Rank	Mean	Char Adj
1	1.119	0.329	1	0.944	0.128	1	0.699	-0.159	1	0.691	-0.177
2	0.905	0.103	2	0.994	0.172	2	1.142	0.330	2	1.158	0.353
3	1.124	0.196	3	1.180	0.282	3	1.226	0.345	3	1.212	0.324
3-1	0.005	-0.133	3-1	0.236	0.154	3-1	0.527	0.505	3-1	0.521	0.501

Table A2: Time Series Regressions

The table presents the results of time series regressions of excess return differentials (Hi-Lo) between portfolios ranked on risk neutral volatility, skewness, and kurtosis on the three Fama and French (1993) factors MRP (the return on the value-weighted market portfolio in excess of a one-month T-Bill), SMB (the difference in returns on a portfolio of small capitalization and large capitalization stocks) HML (the difference in returns on a portfolio of high and low book equity to market equity stocks), and LIQ, the liquidity factor from Pastor and Stambaugh (2001). The moment-sorted portfolios are equally-weighted, formed on the basis of terciles and re-formed each month. The table presents point estimates of the coefficients and standard errors in parentheses. Data cover the period January 1996 through December 2004 for 107 monthly observations.

Panel A: 1 Month to Maturity

	α	β_{MRP}	β_{SMB}	β_{HML}	β_{LIQ}	R^2
Vol	-0.57	0.52	1.01	-0.61	-0.26	0.81
	-1.61	4.66	10.01	-5.15	-4.35	
Skew	-0.50	0.13	0.16	0.47	-0.35	0.54
	-1.64	2.01	1.67	4.00	-5.39	
Kurt	0.49	-0.19	-0.45	-0.45	0.23	0.49
	2.21	-3.78	-6.74	-5.91	6.08	

Panel B: 3 Months to Maturity

	α	β_{MRP}	β_{SMB}	β_{HML}	β_{LIQ}	R^2
Vol	-0.63	0.57	1.01	-1.06	-0.16	0.86
	-1.62	4.46	9.31	-8.18	-2.31	
Skew	-0.61	0.18	0.06	0.29	-0.33	0.42
	-1.89	2.71	0.56	2.21	-4.17	
Kurt	0.57	-0.29	-0.36	-0.11	0.20	0.33
	2.04	-4.00	-3.65	-0.96	3.13	

Panel C: 6 Months to Maturity

	α	β_{MRP}	β_{SMB}	β_{HML}	β_{LIQ}	R^2
Vol	-0.61	0.60	1.00	-1.26	-0.14	0.87
	-1.52	4.55	8.88	-8.96	-1.86	
Skew	-0.64	0.17	0.15	0.07	-0.26	0.27
	-2.47	2.57	1.57	0.57	-3.72	
Kurt	0.56	-0.33	-0.36	0.17	0.18	0.52
	2.10	-3.61	-3.85	1.64	3.07	

Panel D: 12 Months to Maturity

	α	β_{MRP}	β_{SMB}	β_{HML}	β_{LIQ}	R^2
Vol	-0.45	0.55	0.92	-1.32	-0.12	0.86
	-1.13	4.23	8.55	-9.67	-1.71	
Skew	-0.62	0.19	0.21	0.04	-0.26	0.28
	-2.39	3.00	2.25	0.36	-3.78	
Kurt	0.61	-0.37	-0.42	0.16	0.20	0.58
	2.32	-3.97	-4.64	1.72	3.38	