



Universität St.Gallen

C-CAPM without Ex Post Data

Paul Söderlind

September 2006 Discussion Paper no. 2006-22

Editor:

Prof. Jörg Baumberger
University of St. Gallen
Department of Economics
Bodanstr. 1
CH-9000 St. Gallen
Phone +41 71 224 22 41
Fax +41 71 224 28 85
Email joerg.baumberger@unisg.ch

Publisher:

Department of Economics
University of St. Gallen
Bodanstrasse 8
CH-9000 St. Gallen
Phone +41 71 224 23 25
Fax +41 71 224 22 98

Electronic Publication:

<http://www.vwa.unisg.ch>

C-CAPM without Ex Post Data

Paul Söderlind

Author's address:

Professor Paul Söderlind, PhD
Swiss Institute of Banking and Finance
Rosenbergstrasse 52
9000 St. Gallen
Tel. +41 71 224 7080
Fax +41 71 224 7088
Email Paul.Soderlind@unisg.ch
Website www.sbf.unisg.ch

Abstract

Survey and option data are used to take a fresh look at the equity premium puzzle. Survey data on equity returns (Livingston survey) shows much lower expected excess returns than ex post data. At the same time, option data suggests that investors perhaps overestimate the volatility of equity returns. Both facts reduce the puzzle. However, data on beliefs about output volatility (Survey of Professional Forecasters) shows marked overconfidence. On balance, the equity premium is somewhat less of a puzzle than in ex post data.

Keywords

equity premium puzzle, Livingston survey, S&P 500 options, Survey of Professional Forecasters

JEL Classification

G12, E130, E320

1 Introduction

This paper studies if the consumption based asset pricing model is compatible with survey data on subjective beliefs.

Asset pricing models are statements about how expected returns are related to the perceived risk exposure of different assets. This is seldom tested directly. Instead, the models are typically assessed by comparing theoretical pricing implications with the properties of historical (ex post) data. The results can therefore only be given a clear interpretation under the maintained hypothesis that the historical sample is a good representation of subjective beliefs. For instance, the finding of an “equity premium puzzle” relies on the assumption that investors have really expected an excess return of 6%–8% on U.S. equity.

There are several reasons to believe that the moments of historical data can be quite poor approximations of investors’ expectations. Early empirical evidence from survey data suggests that expected returns may deviate from historical averages (for instance, Lakonishok, 1980) and more recent evidence suggests that investors may underestimate the uncertainty of risk (for instance, Thaler, 2000, and Giordani and Söderlind, 2003). These findings could be driven by either small sample problems or some sort of distorted expectations.

The small sample problems are quite likely, since equity returns are highly volatile and sometimes exposed to unusually large shocks. It can easily happen that a fairly long sample has sample moments that deviate substantially from the true values (and subjective, ex ante, beliefs).

Even in the absence of small sample problems, historical data may be poorly suited for testing asset pricing models. Recent findings in behavioural economics (for instance, Hirshleifer (2001)) often point to overconfidence among economic agents. Research on learning (for instance, Lewellen and Shanken, 2002) and on robust decision making (for instance, Tornell, 2000, and Anderson, Hansen, and Sargent, 2003) emphasise that there may be good theoretical reasons for why ex ante beliefs deviate systematically from ex post data.

The approach in this paper is to evaluate the consumption based asset pricing model by using survey data on expected returns and the volatility of the risk factors. This circumvents the problems with expectation errors—and is therefore a way to test the model more directly.

To get a reasonably long sample with high-quality data, I combine several data sources. The Livingston survey and the Survey of Professional Forecasters are used to measure subjective beliefs of expected equity returns and consumption volatility. Both surveys are focused on the beliefs of economists close to the financial markets and are administered by the Federal Reserve Bank of Philadelphia. Implied volatilities from options on the S&P 500 index are used as a measure of equity return uncertainty.

The plan of the paper is as follows. Section 2 summarises the standard consumption based asset pricing model. Sections 3 and 3.4 present the survey data and the empirical results. Section 5 sums up. The appendix gives details on data.

2 The Standard Consumption Based Asset Pricing Model

This section gives a brief summary of the standard CRRA model and the equity premium puzzle.

The standard consumption based asset pricing model assumes a utility function with constant relative risk aversion, $C_t^{1-\gamma}/(1-\gamma)$, where C_t is consumption and γ the risk aversion coefficient. The Euler equation for optimal portfolio choice can be written as $E[R_t^e (C_t/C_{t-1})^{-\gamma}] = 0$, where R_t^e is the excess return on an asset—and $E()$ denotes the expectations of the (representative) investor.

To simplify the analysis, assume that the investor thinks that the excess return and consumption growth have a bivariate normal distribution. Use Stein's lemma¹ to rearrange the Euler equation as

$$E(R_t^e) = \text{Cov}(R_t^e, \Delta c_t) \gamma \quad (1)$$

$$= \text{Corr}(R_t^e, \Delta c_t) \sigma(R_t^e) \sigma(\Delta c_t) \gamma, \quad (2)$$

where Δc_t is the growth rate of consumption, $\ln(C_t/C_{t-1})$. It is worth emphasising that this equation is part of the investor's decision process—and therefore holds only for his/her beliefs. It really does not matter where those beliefs come from, if they are rational, or if they are the result of a learning process (as in, for instance, Brennan (1998)).

We can relax the assumption that the excess return is normally distributed: (2) holds also if R_t^e and Δc_t have a bivariate mixture normal distribution—provided Δc_t has the

¹Stein's lemma says that if x and y have a bivariate normal distribution and $h(y)$ is a differentiable function such that $E[|h'(y)|] < \infty$, then $\text{Cov}[x, h(y)] = \text{Cov}(x, y) E[h'(y)]$. See Cochrane (2001).

same mean and variance in all the mixture components (see Söderlind (2006)). This restricts consumption growth to have a normal distribution, but allows the excess return to have a distribution with fat tails and skewness.

The gain from assuming a normal distribution of consumption growth (using Stein's lemma) is that the unknown coefficient of relative risk aversion, γ , enters multiplicatively in (1)–(2). This allows us to work with correlations and standard deviations of returns and consumption growth—for which it might be possible to find survey data. In contrast, if we did not assume normally distributed consumption growth rates, then we would need information about, for instance, the standard deviation of $(C_t/C_{t-1})^{-\gamma}$. There is no survey data on such things.²

The equity premium puzzle (see Mehra and Prescott (1985)) is that (2) does not hold for ex post data unless γ is very high. To illustrate this, I use quarterly US data for 1952Q3–2005Q4. Consumption growth is measured as the growth rate of per capita consumption of nondurables and services (see Appendix A) and the return is the real return on the S&P 500 in excess of the real return on a 3-month T-bill rate.

In annualised terms (quarterly growth rates are multiplied by 4 and quarterly standard deviations by 2) the ex post data gives

$$\underbrace{E(R_t^e)}_{0.06} = \underbrace{\text{Corr}(R_t^e, \Delta c_t)}_{0.17 \text{ or } 0.33} \underbrace{\sigma(R_t^e)}_{0.14} \underbrace{\sigma(\Delta c_t)}_{0.01} \gamma. \quad (3)$$

Two numbers are given for the correlation: the lower is when consumption growth (a flow variable) is measured as the growth between the current and lagged quarter; the higher when it is measured as the growth between the next and current quarters. In any case, even if the correlation was perfect (one), then we need a relative risk aversion, γ , of 43 to make (3) hold. This is the puzzle.

3 Evidence from Survey Data

The survey and option data used in this paper come from different sources, cover different periods, and have different sampling frequencies. There really are not many alternatives, at least not if we want to work with reasonably long samples.

²The general expression is $E(R_t^e) = -\text{Cov}[R_t^e, (C_t/C_{t-1})^{-\gamma}] / E[(C_t/C_{t-1})^{-\gamma}]$; the covariance can be expanded in terms of the correlation and the two standard deviations.

3.1 Relation Between Survey Data and the Asset Pricing Equation

Survey data gives information on conditional moments of subjective beliefs of heterogeneous agents. In contrast, the asset pricing equation (3) is expressed in terms of unconditional moments of a representative investor. This section discusses how the survey data may still be useful to understand the properties of the asset pricing equation.

The evidence from surveys and options is on conditional moments—both data sets essentially contain answers to a question like “based on your information today, what do you believe about x in the near future?” We therefore need a rule for transforming back to unconditional moments. This is straightforward for the expected return, $E(R_t^e)$, since the unconditional expectation is best estimated by the average conditional expectation (by the law of iterated expectations). For the standard deviations in (3), matters are slightly more complicated. But, if the excess returns and the consumption growth rates are unpredictable (which is a good approximation), then the unconditional variance is the average conditional variance. The unconditional standard deviation is then the square root of the average conditional variance—which is the approach used below.

The Euler equation (2) is valid for each investor—provided we are careful enough to use the moments of his/her beliefs about the return and *his/her own consumption*. In (3) it is further assumed that all investors are identical: they share the same beliefs and it is aggregate consumption that matters.

The survey data has a different structure: it asks for his/her beliefs about the return and *aggregate output*. To bridge the gap between data and the asset pricing equation, we need strong assumptions. First, that results on over/underestimation of the volatility of output (which we have data on) carries over to consumption (which enters the asset pricing equation). Second, that equity premia are not affected by non-insurable idiosyncratic risk (since we have no information on this risk). Third, that aggregation of beliefs (since the respondents in the survey disagree) is unimportant.

None of these assumptions is likely to hold exactly, but they may well be reasonable approximations. First, the time series properties of output and consumption are very similar in many respects: they are strongly correlated over all horizons and so are the forecast errors. In addition, most macroeconomic theory would suggest a very strong link. Second, the evidence on (the importance for asset pricing of) idiosyncratic risk is mixed (see, for instance, Lettau (2002) and Cogley (1998)). Third, aggregation of heterogeneous beliefs may well suggest that an average belief is a good approximation of the represen-

tative investor's belief—in particular when the relative risk aversion is high. For instance, Giordani and Söderlind (2005) show (by extending a model by Varian (1985)) that a representative investor should be assigned a mean equal to the cross-sectional mean, and a variance equal to a combination of the (cross sectional average of) individual uncertainty and the cross-sectional disagreement. It turns out that the relative weight on disagreement is $1/(\gamma + 1)$: with log utility disagreement is as important as individual uncertainty (this is the case in Rubinstein (1974) and Detemple and Murthy (1994)), but a more realistic value of the relative risk aversion gives a very small role to disagreement. This suggests that it might be reasonable to disregard the aggregation issue.

3.2 Expected Stock Return

The *Livingston survey* summarizes the forecasts of economists in industry, government, banking, and academia. It is the oldest continuous survey of economists' forecasts, started in 1946 by the financial columnist Joseph Livingston. It was taken over by the Federal Reserve Bank of Philadelphia in 1990. It contains questions about the expected Standard & Poor (S&P) index level 6 and 12 months ahead in time.

For the period June 1952 to June 1990, the Livingston survey asked for the S&P Industrials index (labelled S&P's 425 Industrials for most of period, but is different from the more recent S&P 400 Midcap).³ From December 1990 and onwards the survey has asked for the S&P 500 Composite. For further details and references, see Lakonishok (1980), Federal Reserve Bank of Philadelphia (2005), Federal Reserve Bank of Philadelphia (2004) and Appendix A.

The June and December surveys are sent out in late May and November respectively. The June survey asks for the index levels at the end of December the same year and June next year. Similarly, the December survey asks for the index levels at the end of June and December next year.

It is, however, somewhat unclear what base value to use in order to calculate expected capital gains (growth rates). There are two main alternatives in the literature. First, the survey contained base values—typically from the middle of the month before the survey. This gives investment horizons of 7.5 and 13.5 months respectively. Second, Pearce

³The June 1952 survey asked about the S&P 365 Industrials, the December 1952 to December 1956 asked about the S&P 420 Industrials. These indices had 1935–39=100 as base. For June 1957 to June 1990 the survey asked about the S&P (425) Industrials, with 1941–43=10 as base.

(1984) use base value from the last day of month before the survey: which gives investment horizons of 7 and 13 months. There are issues with both: some of the base values seem unreliable and it is not entirely clear that the respondents had access to the end-of-month values (of the right index) when filling out the forms. I will mostly focus on a third alternative where I combine the 6- and 12-month forecasts to calculate an implied expected growth rate over a 6-month horizon starting 6 months from now. This avoids the problems with the base level and the investment horizon is undoubtedly 6 months. (None of the conclusions are affected by this choice.)

Table 1 shows the (annualised) results for the 6-month horizon starting 6 months from now for the sample June 1952 to December 2005. It shows the average capital gains in excess of a riskfree rate (approximated by a 3-month T-bill rate) as forecasted by the survey and in ex post data. The ex post returns are for precisely the same periods as the survey contains forecasts of, so the results are comparable.

	1952–1990	1990–2005	1952–2005
Survey expectations (capital gain)	−1.1	1.5	−0.4
S&P Industrials, ex post (capital gain)	3.2		
S&P 500, ex post (capital gain)		5.3	
S&P combined, ex post (capital gain)			3.8
Dividend yield S&P 500	4.0	2.1	3.5

Table 1: Average capital gains in excess of a riskfree rate, dividend yield, annualised %. The table shows results for the 6-month horizon starting 6 months ahead in time, annualised by multiplying by 2. The survey data are from the Livingston survey and show the expected change of the stock index (S&P Industrials for 1952:6–1990:6 and S&P 500 Composite for 1990:12–2005:12) — in excess of a riskfree rate. The survey data is based on the cross-sectional median forecast in each period. The ex post results are for the corresponding actual changes in the stock index. The S&P combined uses S&P Industrials for 1952:1–1990:6 and S&P 500 Composite for 1990:12–2005:12.

For the early sample (1952–1990) the Livingston survey has an expected capital gain (in excess of a riskfree rate) of −1.1% (annualised), while the ex post capital gain on the S&P Industrials was 3.2%. For the late sample (1990–2005), the survey has an expected capital gain of 1.5% while the ex post capital gain on the S&P 500 was 5.3% (both in excess of the riskfree rate).

Clearly, the capital gain on the indices is not the total return from investing on the stock market. Table 1 therefore also gives the ex post dividend yield on the S&P 500 index, which is around 4.0% and 2.1% for the early and late samples respectively. The survey

contains no information about expected dividends, but it seems reasonable to assume that there is no bias in the 12-month dividend forecasts—especially since dividends are fairly smooth over time. It also seems reasonable to use the dividend yield for the S&P 500 index for both sub samples (the S&P Industrials and 500 indices are very similar). In that case, the expected excess returns were 2.9% and 3.6% for the two samples.⁴ In contrast, the ex post excess returns were around 7.2% – 7.4%. This suggests that, over the last half century, ex post returns on U.S. equities have been substantially higher than expected (in the Livingston survey). This squares well with the notion of the U.S. market as one of the “lucky cases” (see, for instance, Brown, Goetzmann, and Ross (1995)), or with the idea of too pessimistic investors (see Abel (2002)).




Figure 1: Expected capital gains on the S&P index in excess of a riskfree rate, annualised %. This figure shows the expected capital gains (in excess of a riskfree rate) for the 6-month horizon starting 6 months ahead in time (with and without adding the ex post dividend yield) and the NBER recessions (shaded areas). The data is for the S&P Industrials for the period up to June 1990, and the S&P 500 after that.

This evidence comes from a small sample of forecasters, so we cannot really tell if the results are representative of the market expectations or not. Studying the properties of the forecasts is one (informal) way of assessing the evidence.

⁴The Survey of Professional Forecasters also asks for the expected 10-year S&P 500 return as well as the average 3-month T-bill rate. For the 1992–2006 sample the implied expected excess return is 3.9%.




Figure 2: Expected and ex post capital gains on the S&P index in excess of a riskfree rate, annualised %. This figure shows the 6-month ex post capital gains (in excess of a riskfree rate) and the forecast of it made 1 year earlier. The data is for the S&P Industrials for the period up to June 1990, and the S&P 500 after that.

The implied average expectation error is around 3% (annualised) for the combined sample—which is large. However, the standard deviation of the expectation errors is around 16% and the sample has just a bit more than 100 observations, so a standard test of the null hypothesis of no (average) expectation errors is borderline significant (a t-stat of $3/(16/\sqrt{100}) = 1.9$). The volatile stock market returns makes learning hard, even in a sample of more than 50 years.

To assess the time pattern of the expectations and the expectation errors, consider *Figure 1*. The upper panel shows the implied expected capital gains in excess of a riskfree rate for months 6 to 12—with and without adding the ex post dividend yield. The NBER recessions are marked by shaded areas. The Livingston forecasters clearly did not believe in the random walk hypothesis: there are distinct movements in the expected capital gains. In particular, the expectations have (local) maxima in almost all recessions. In contrast, the expectations for the next 6 months (not shown), have local minima in most recessions. This suggests a belief in short term momentum, but a medium term mean reversion. Indeed, the correlation between current ex post capital gains (over the last 6 months) and the expected gains over the next 6 months is positive (0.13 for the combined

sample), while the correlation with the expected gains over months 6 to 12 is negative (-0.17). In any case, the forecasting performance is poor. Apart the (already documented) average underprediction, the forecasters have virtually zero R-squares in trying to predict the movements around the mean.

Figure 1 also shows the expected values plus the ex post dividend yield—in an attempt to illustrate the time profile of expected returns. This curve is, more or less, a vertical shift of the expected capital gains—although the shift is larger in the early part of the sample. There are still a number of periods with negative expected returns, which are somewhat unappealing. Clearly, if this represents the market returns, and we assume a text book mean-variance setting, then a risk averse representative investor cannot have negative expected returns. On the other hand, in a consumption-based model nothing rules out negative risk premia. For instance, Duffee (2005) estimates a time-varying covariance of consumption growth and the CRSP stock market returns and finds negative values of many periods, for instance, in the early 1960s and 1980s (compare with Figure 1). If we instead take the position of Campbell and Thompson (2005) that negative expected excess returns should be replaced by zeros, then the ex post returns would still be 1.5% higher than the survey data. This half of the previous difference, but still a sizeable figure.

Figure 2 compares the expectations with the ex post capital gains. The figure suggests that the differences in average capital gains may well be driven by a few exceptional periods. For instance, there were *prolonged* periods of consistent under estimation in 1954–1955 (when the market value grew by almost 70%) and 1995–1999 (when the market value trebled). Although there are a few periods of consistent over prediction (mainly 2001–2002), they tend to be less dramatic. All other periods appear more like random noise, where short-lived under and over predictions more or less cancel.

While the post-Korean-war and dot-com booms are defining periods in U.S. stock market history, it does not seem unreasonable to conclude that they may well have been mostly unexpected. If so, the survey data on stock market expectations tells a very different story about average returns than ex post data.

In the end, this evidence suggests that the *expected excess return might be around 3.0%–3.5%. This is about half of the historical average.* This will, in itself, reduce the required risk aversion (γ) in asset pricing equation (2) by half.

3.3 Consumption Growth Volatility

Most surveys ask the respondents to supply their point forecasts—but very few ask for measures of uncertainty. The *Survey of Professional Forecasters* (SPF) is an exception, since it asks for histograms of the respondents' subjective probability distribution of real GDP growth.⁵ I will use this information to compare the subjective uncertainty to actual uncertainty (the volatility of forecast errors). Although this provides evidence only on output growth volatility, it is plausible that the results (in relative terms) carry over to consumption growth.

The idea of the SPF is to ask professional forecasters about their views on key macroeconomic and financial variables. The survey dates back to the late 1960s, but the question about the subjective probability distribution was introduced only in 1981Q3. The survey is now administered by the Federal Reserve Bank of Philadelphia (before 1990, by Victor Zarnowitz and others of the American Statistical Association and the National Bureau of Economic Research).

To elicit subjective probability distribution of real GDP growth the respondent is asked to assign probabilities of different growth intervals (bins) defined by the survey manager. For instance, for growth rates between 0% and 0.9%, between 1% and 1.9%, and so on.⁶ The growth rates are defined as year-on-year (the value in calendar year t divided by the value in calendar year $t - 1$, minus one). I use forecasts of this growth rate at four forecasting horizons: one to four quarters ahead. The four-quarter forecast is made in Q1 of year t , the three-quarter forecast is made in Q2 and so forth.

The survey is sent out towards the end of the first month of the quarter and the replies are due around two weeks later. This makes the definition of the forecasting horizon a moot point. Unlike financial data, GDP data is published with a considerable lag (and later revised). For instance, it can be argued that the four-quarter survey (with replies due around 15 February) is anything between a 12-month forecast to a 10.5-month forecast. I will use the former convention, but that will not affect the results on the relative under-/overestimation of volatility since I combine the information in the subjective distribution and the point forecasts reported at the same time (see below).

⁵To be precise, real GNP during 1981Q3–1991Q4 and real GDP since 1992Q1.

⁶Before 1992, the bins were as follows: less than -2, -2 to -0.1, 0 to 1.9, 2 to 3.9, 4 to 5.9, and 6 and above. Since 1992, they are: less than -2, -2 to -1.1, 1 to -0.1, 0 to 0.9, 1 to 1.9, 2 to 2.9, 3 to 3.9, 4 to 4.9, 5 to 5.9, and 6 and above.

Giordani and Söderlind (2003) and Giordani and Söderlind (2005) discuss how to estimate the underlying distribution from a histogram reported by a forecaster. They find that fitting a normal distribution to each histogram typically works well. Therefore, for each forecaster/forecasting horizon/survey period, I fit a normal distribution to his/her histogram by minimizing the sum of squared errors (implied probability of a bin minus the actual probability according to the histogram). The fitted normal distribution is then an estimate of the forecaster's subjective distribution. This produces a subjective variance for every forecaster in every period (that he/she participated in the survey). I then calculate a cross-sectional "consensus" variance for each period by taking the cross-sectional median, and then form a time-average. (Using the cross-sectional means produce a similar time-average.)

For the "ex post volatility" we need a conditional variance, that is, the variance of the forecast errors. I estimate this by using the consensus (cross-sectional median) forecasts as the point forecasts.

	4 quarters	3 quarters	2 quarters	1 quarter
Survey	0.8	0.7	0.5	0.4
Ex post	1.1	0.9	0.8	0.8

Table 2: Standard deviation of year-on-year output growth, %. The table shows results for four different forecasting horizons. The survey data are from the Survey of Professional Forecasters, and covers the period 1982Q1 to 2005Q4. The survey standard deviation is calculated as the square root of the time average of the consensus variance. The consensus variance for a given period is calculated as the cross-sectional (different forecasters) median variance. The ex post standard deviation is for the forecast errors, using the consensus point forecast.

Table 2 reports results for all four horizons.⁷ For instance, the ex post standard deviation for the 2-quarter horizon is 0.8%, while the standard deviation from the subjective distributions is only 0.5%. This difference is significant at the 5% level (based on a delta method approach). The other horizons give similar results. This suggests that the forecasters have underestimated uncertainty considerably: the ratio of the ex post and subjective standard deviations varies between 1/2 (at the 1-quarter horizon) to 3/4 (at the 3-quarter horizon). Giordani and Söderlind (2005) reach a similar conclusion by using somewhat different methods (the coverage ratio of subjective confidence bands).

⁷Some data points (1985Q1 and 1986Q1) for the 4-quarter horizon are likely to be wrong (see Federal Reserve Bank of Philadelphia (2000)). They are therefore excluded from the estimations.

Clearly, this result refers to GDP growth—not consumption. However, it is likely to carry over to consumption since output and consumption are so strongly related (both theoretically and empirically).

To sum up, this evidence suggests that the *consumption growth volatility might be around 0.6 times the historical volatility*. This will, in itself, increase the required risk aversion (γ) in asset pricing equation (2) by two thirds (that is, scale it up by 5/3).

3.4 Evidence from Options: Stock Return Volatility

There is only very little survey evidence on the return volatility, but implied volatilities from option prices provide an (imperfect) alternative.

In the Black-Scholes model, an option price depends directly on the market’s belief about the return volatility until expiration of the option. It is straightforward to calculate this “implied volatility” from the Black-Scholes formula—and to compare with the ex post (“realized”) return volatility over the same horizon. This is basically what I am doing in this section.

The option data (daily 1985 and onwards) contain information on call and put options on the S&P 500 index, for various strike prices and maturities. I use only those observations for which there was trade, to avoid using non-informative price quotes. For each option, I calculate the implied volatility, and then form an average of the two call options and the two put options with strike prices closest to the futures price. This is a well established approach (similar to the Bloomberg approach and also the old VIX), and one that Ederington and Guan (2002) found to be close to optimal in a forecasting sense.

The reason for averaging calls and puts is that their implied volatilities are likely to have negatively correlated measurement errors (due to, for instance, nonsynchronous data)—and the reason for using options that are close to “at the money” is that they are heavily traded so the prices are informative. To get results for four different fixed horizons (1 to 4 quarters), I interpolate between maturities.

Table 3 reports results for four horizons (1 to 4 quarters). This evidence suggests that investors have overestimated volatility by around 15%. Tests of differences (correcting for overlapping forecasting periods and using a delta method approach) show that the overestimation is clearly significant at a 5% level.

This results can be criticised for being based on a faulty option pricing model which assumes that volatility is non-random (and asset returns are normally distributed). I will

	4 quarters	3 quarters	2 quarters	1 quarter
Options	19.8	19.7	19.3	19.2
Ex post	16.8	16.7	16.7	16.6

Table 3: Standard deviation of S&P 500 returns, annualised %. The table shows results for four different forecasting horizons. The results for options are implied volatilities calculated from options on the S&P 500 index (daily data 1985–2005). The ex post volatility is estimated from daily S&P 500 returns. The standard deviations are calculated as the square root of the time averages of the variances.

discuss two extensions.

First, in a stochastic volatility model like that of Hull and White (1987) where the stochastic volatility does not imply an additional risk premium (because the asset price and volatility are uncorrelated), the distribution becomes non-normal so the Black-Scholes formula does not apply. A Ball and Roma (1994) approximation then shows that we should add around 0.5 to the implied volatilities in Table 3, so the overestimation becomes even worse.

Second, in a model like that of Heston (1993), the stochastic volatility gives an additional risk premium. Even if we could approximate the variance as before, this would only be the variance of the riskneutral distribution. In this case, it is possible that the “overestimation” in Table 3 is just risk premia. This is essentially an open issue (see, for instance, Jackwerth (2000) for a discussion about “overpriced” options). Without survey data we rally cannot tell. However, *Figure 3* shows an interesting pattern. For clarity, it plots only the non-overlapping data points for the 2-quarter horizon. The “overestimation” is fairly persistent, but it seems to be related to the level of uncertainty (measured by either of the series): more overestimation at low uncertainty. This makes the interpretation in terms of risk premia a bit counter-intuitive, since the periods of low uncertainty also have fairly stable uncertainty. Instead, this could be systematic expectations errors.

Together this evidence suggest that the *return volatility might perhaps be around 1.15 times the historical volatility*, but only if the risk premium for stochastic volatility is not too important.




Figure 3: Standard deviation of the S&P 500 index, annualised %. This figure implied volatilities (from options) and the ex post volatility of the S&P 500 index, for the 6-month horizon. The ex post volatility is shifted 6 months back in time to allow for a direct comparison between the forecast (implied volatility from the options) and the outcome.

4 Implications for the Equity Premium

The results from survey and option data suggest that the historical mean excess returns should be scaled by 0.5, the standard deviation of excess returns by 1.15 (if we choose to disregard risk premia for stochastic volatility), and the standard deviation of consumption growth by 0.6. This changes the asset pricing equation (3) to

$$\underbrace{E(R_t^e)}_{0.5 \times 0.06} = \underbrace{\text{Corr}(R_t^e, \Delta c_t)}_{?} \underbrace{\sigma(R_t^e)}_{1.15 \times 0.14} \underbrace{\sigma(\Delta c_t)}_{0.6 \times 0.01} \gamma. \quad (4)$$

This means that the left hand side is scaled by 0.5 and the right hand side by 0.7—so the equation would hold for a lower value (factor 0.7) of the relative risk aversion, γ . For instance, with a perfect correlation we need $\gamma = 30$ (instead of 43 in ex post data) to make this equation hold. This is a step in the right direction—although not a very large one.

5 Summary

This paper performs a traditional analysis of the equity premium puzzle—but without ex post data. Instead, data from surveys and options are used.

It is found that survey data on equity returns (from the Livingston survey) show much lower expected excess returns than in ex post data for the same period: around 3% instead of around 6%. At the same time, option data indicate that investors might overestimate the standard deviation of equity returns by a factor 1.15. Both these factors work in favour of the standard consumption-based asset pricing model.

Unfortunately (for the model), data on subjective beliefs about output growth volatility (Survey of Professional Forecasters) shows a marked overconfidence—the standard deviation implied by the survey data is just 0.6 of the historical standard deviation. Although no direct evidence on consumption volatility exists, it is likely to look similar.

Putting these results together leads to some improvement of the consumption based model. Using survey and option data, the coefficient of relative risk aversion need only be 0.7 of what ex post data requires.

A Data Appendix

The data of the Livingston Survey and the Survey of Professional Forecasters are from the Federal Reserve Bank of Philadelphia (<http://www.phil.frb.org/>). For details, see Federal Reserve Bank of Philadelphia (2005) and Croushore (1993).

The daily S&P 500 and S&P 100 price indices are from Datastream. The option data is from CBOE.

The monthly S&P 500 and S&P Industrials price indices are from S&P's Trade and Securities Statistics (1976) and Compustat. The dividends on the S&P 500 is from Shiller (<http://www.econ.yale.edu/%7Eshiller/data.htm>), see Shiller (2000). For a discussion of the historical S&P series, see Jones and Wilson (2002).

The 3-month T-bill rate is from FRED II (<http://research.stlouisfed.org/fred2/>).

Quarterly growth of real consumption per capita of nondurables and services is calculated from the seasonally adjusted number in the NIPA tables (available at <http://www.bea.doc.gov/>). The growth rate is calculated as a weighted average of the growth rate of nondurables and the growth rate of services (chained 2000 dollars), where the (time-varying) weight is the

relative (current dollar) size of nondurables in relation to services.

Real returns are calculated by dividing the nominal gross return by the gross inflation rate over the same period. Inflation is calculated from the seasonally adjusted CPI for all urban consumers (available at <http://research.stlouisfed.org/fred2/>).

References

- Abel, A. B., 2002, "An Exploration of the Effects of Pessimism and Doubt on Asset Returns," *Journal of Economic Dynamics and Control*, 26, 1075–1092.
- Anderson, E., L. P. Hansen, and T. J. Sargent, 2003, "Robustness, Detection, and the Price of Risk," Stanford University Working Paper.
- Ball, C. A., and A. Roma, 1994, "Stochastic Volatility Option Pricing," *Journal of Financial and Quantitative Analysis*, 29, 589–607.
- Brennan, M. J., 1998, "The Role of Learning in Dynamic Portfolio Decisions," *European Finance Review*, 1, 295–306.
- Brown, S. J., W. N. Goetzmann, and S. A. Ross, 1995, "Survival," *Journal of Finance*, 50, 853–873.
- Campbell, J. Y., and S. B. Thompson, 2005, "Predicting the Equity Premium Out of Sample: Can Anything Beat the Historical Average," Discussion paper, NBER, WP 11468.
- Cochrane, J. H., 2001, *Asset Pricing*, Princeton University Press, Princeton, New Jersey.
- Cogley, T., 1998, "Idiosyncratic Risk and Equity Premium: Evidence from the Consumer Expenditure Survey," mimeo, Federal Reserve Bank of San Francisco.
- Croushore, D., 1993, "Introducing: The Survey of Professional Forecasters," *Federal Reserve Bank of Philadelphia Business Review*, pp. 3–13.
- Detemple, J., and S. Murthy, 1994, "Intertemporal Asset Pricing with Heterogeneous Beliefs," *Journal of Economic Theory*, 62, 294–320.

- Duffee, G. R., 2005, "Time Variation in the Covariance between Stock Returns and Consumption Growth," *Journal of Finance*, 60, 1673–1712.
- Ederington, L. H., and W. Guan, 2002, "Measuring Implied Volatility: Is an Average Better? Which Average?," *Journal of Futures Markets*, 22, 811–837.
- Federal Reserve Bank of Philadelphia, 2000, "Survey of Professional Forecasters: Documentation for Median Forecast Data," <http://www.phil.frb.org/econ/spf/spfmeddoc.html>.
- Federal Reserve Bank of Philadelphia, 2004, "Documentation for the Discontinued Stock Price Forecasts in the Federal Reserve Bank of Philadelphia's Livingston Survey," <http://www.philadelphiafed.org/econ/liv/S&P.html>.
- Federal Reserve Bank of Philadelphia, 2005, "Livingston Survey Documentation," <http://www.philadelphiafed.org/econ/liv/index.html>.
- Giordani, P., and P. Söderlind, 2003, "Inflation Forecast Uncertainty," *European Economic Review*, 47, 1037–1059.
- Giordani, P., and P. Söderlind, 2005, "Is there Evidence of Pessimism and Doubt in Subjective Distributions? Implications for the Equity Premium Puzzle," *Journal of Economic Dynamics and Control*, 30, 1027–1043.
- Heston, S. L., 1993, "A closed-form solution for options with stochastic volatility with applications to bond and currency options," *Review of Financial Studies*, 6, 327–343.
- Hirshleifer, D., 2001, "Investor Psychology and Asset Pricing," *Journal of Finance*, 56, 1533–1597.
- Hull, J., and A. White, 1987, "The Pricing of options on assets with stochastic volatilities," *Journal of Finance*, 42, 281–300.
- Jackwerth, J. C., 2000, "Recovering risk aversion from option prices and realized returns," *Review of Financial Studies*, 13, 433–451.
- Jones, C. P., and J. Wilson, 2002, "An Analysis of the S&P 500 Index and Cowles's Extensions: Price Indexes and Stock Returns 1870–1999," *Journal of Business*, 75, 505–533.

- Lakonishok, J., 1980, "Stock Market Return Expectations: Some General Properties," *Journal of Finance*, 35, 921–931.
- Lettau, M., 2002, "Idiosyncratic Risk and Volatility Bounds, or, Can Models with Idiosyncratic Risk Solve the Equity Premium Puzzle?," *Review of Economics and Statistics*, 84, 376–380.
- Lewellen, J., and J. Shanken, 2002, "Learning, Asset-Pricing Tests, and Market Efficiency," *Journal of Finance*, 57, 1113–1145.
- Mehra, R., and E. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145–161.
- Pearce, D. K., 1984, "An Empirical Analysis of Expected Stock Price Movements," *Journal of Money, Credit and Banking*, 16, 317–327.
- Rubinstein, M., 1974, "An Aggregation Theorem for Securities Markets," *Journal of Financial Economics*, 1, 225–244.
- Shiller, R. J., 2000, *Irrational Exuberance*, Princeton University Press, data available at <http://www.econ.yale.edu/shiller/data.htm>.
- Söderlind, P., 2006, "C-CAPM Refinements and the Cross-Section of Returns," *Financial Markets and Portfolio Management*, 20, 49–73.
- Thaler, R. H., 2000, "From Homo Economicus to Homo Sapiens," *Journal of Economic Perspectives*, 14, 133–141.
- Tornell, A., 2000, "Robust- H_∞ Forecasting and Asset Pricing Anomalies," NBER Working Paper No. 7753.
- Varian, H. R., 1985, "Divergence of Opinion in Complete Markets: A Note," *Journal of Finance*, 40, 309–317.