The C-CAPM without Ex-Post Data

Paul Söderlind*

16 January 2009

Forthcoming in *Journal of Macroeconomics*

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 tend to overestimate the volatility of equity returns. Both facts contribute towards solving 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 Numbers: G12, E130, E320

*University of St. Gallen and CEPR. *Address:* s/bf-HSG, Rosenbergstrasse 52, CH-9000 St. Gallen, Switzerland. *E-mail:* Paul.Soderlind@unisg.ch. I thank Magnus Andersson, Jan Bernhard, Michael Fischer, and Felix Moldenhauer for help with data; and three anonymous referees for comments.

1 Introduction

This paper studies whether the equity premium puzzle can be solved through the use of survey and option data. The answer turns out to be yes, to a certain degree (two fifths, to be exact).

Several suggestions have been made on how to solve the equity premium puzzle: by using complicated utility functions (see Campbell and Cochrane (1999), among others), long-run consumption changes (Parker and Julliard (2005)) or data for shareholders only (Vissing-Jorgensen and Attanasio (2003)). This paper takes another route by focusing on measured market beliefs. In order to highlight this (potential) explanation for the equity premium puzzle, the rest of the setting is very standard: a traditional utility function with constant relative risk aversion and short run (2–4 quarters) aggregate data.

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, 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 (ex ante) 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 suggested that expected returns may deviate from historical averages (for instance, Lakonishok, 1980) and more recent evidence suggests that investors may underestimate the volatility of risk factors (for instance, Thaler, 2000, and Giordani and Söderlind, 2003). These findings could be driven either by small sample problems or some sort of distorted expectations.

Small sample problems are quite likely to occur, 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 beliefs). For instance, Merton (1980) has highlighted the problems of estimating expected returns—and suggested several ways of circumventing them by using the more precisely estimated market volatility.

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 taken in this paper is to evaluate the consumption based asset pricing model by using survey (and some option) data on expected returns and the volatility of the risk factors. This avoids the issue of whether *ex post* data is a good approximation of subjective beliefs or not—and is therefore a more direct way of testing the economic model (the Euler equation).

Several data sources are combined to get a reasonably long sample with high-quality data. The Livingston survey and the Survey of Professional Forecasters are used to measure subjective beliefs of expected equity returns and consumption uncertainty—and implied volatilities from options on the S&P 500 index are used to measure equity return uncertainty.

The paper is organised in the following way. Section 2 summarises the standard consumption based asset pricing model and discusses how it corresponds to the available data. Section 3 presents the data and the empirical results. Finally, Section 4 sums up the findings. The appendix provides further details on the data.

2 Asset Pricing and Survey Data

2.1 The Standard Consumption Based Asset Pricing Model

This section provides 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 is the risk aversion coefficient. The Euler equation for optimal portfolio choice can be expressed 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 expectation of the investor. (The time discount factor cancels when excess returns are used.)

To simplify the analysis, assume that the investor thinks that the excess return and consumption growth have a bivariate normal distribution. Then use Stein's lemma (see Cochrane (2001)) to rearrange the Euler equation as

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

where Δc_t is the growth rate of consumption, $\ln(C_t/C_{t-1})$, and $\sigma()$ denotes the standard deviation. It is worth emphasising that this equation is part of the investor's decision process—and therefore holds only for his/her beliefs. That is, the expectation, correlation and standard deviations are those from the subjective distribution of the investor. It really does not matter where those beliefs come from, if they are rational, or if they are the result of a learning process—at least as long as the learning does not depend on investing (see, for instance, Brennan (1998)).

We can relax the assumption that the excess return is normally distributed: (1) holds also if R_t^e and Δc_t have a bivariate mixture normal distribution—provided that Δc_t has the same mean and variance in all the mixture components (see Söderlind (2008a)). This restricts consumption growth to having a normal distribution, but allows the distribution of excess returns to be skewed and have fat tails.

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

The equity premium puzzle (see Mehra and Prescott (1985)) is that (1) does not hold when the beliefs are approximated by the moments of ex post data—unless γ is very high. To illustrate this, quarterly US data for 1952Q3–2007Q4 is used here. Consumption growth is measured as the growth rate of per capita consumption of nondurables and services (see Appendix) 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 expected returns are multiplied by 4 and quarterly stan-

dard deviations by 2) the ex post data gives

$$\underbrace{E(R_t^e)}_{0.07} = \underbrace{\text{Corr}(R_t^e, \Delta c_t)}_{0.33} \underbrace{\sigma(R_t^e)}_{0.14} \underbrace{\sigma(\Delta c_t)}_{0.01} \gamma. \quad (2)$$

The relative risk aversion, γ , that fits this equation is 150. Even if the correlation was perfect, a relative risk aversion of 50 would still be required. This is the puzzle.

The idea of this paper is to measure the components of (2) from surveys and options. This allows us to evaluate the asset pricing model separately from the issue of whether the ex post distribution coincides with subjective (ex ante) beliefs or not. To highlight the contribution from using survey and option data, the results are systematically compared with those from ex post data. In practice, this means that the Euler equation (2) is studied according to two sets of beliefs: (i) as measured by survey and option data; (ii) as proxied by the moments of ex post data.

2.2 The Relationship Between Survey Data and the Asset Pricing Equation

This section discusses some issues to do with the data used in the paper.

The survey and option data used in this paper come from different sources, cover different periods and have different sampling frequencies. There are few alternatives if we want a reasonably long sample.

I assume that the survey data is representative of the expectations that drive US equity pricing and consumption. This is not uncontroversial, especially since the survey participants are economic experts rather than (a properly weighted) cross-section of investors. However, these economic experts are often close to large investors and are likely to influence the public.¹ I therefore consider the survey data to be a reasonable approximation of market beliefs.

The evidence from surveys and options is on conditional moments. Both data sets essentially contain answers to questions like “based on your information today, what do you believe about x in the near future?” This corresponds to a conditional distribution. A rule is therefore required for transforming back to the unconditional moments used in the asset pricing equation (2). This is quite straightforward for the expected return, $E(R_t^e)$, since the unconditional expectation is best estimated by the average conditional expectation (by

¹For instance, the Livingston panel in December 2006 included the Chief US Economist of Morgan Stanley and the Chief Economist of Standard&Poor’s. Both are very active in the media.

the law of iterated expectations). Similarly, for the excess return the average conditional variance is a good approximation of the unconditional variance, since the returns are almost unpredictable. Matters are slightly more complicated for the standard deviation of the consumption growth rate, which is somewhat predictable. I therefore assume that the under-/overestimation of the conditional volatility carries over to the unconditional volatility.

The Euler equation is valid for each investor—provided that care is taken to use the moments of his/her beliefs about the return and his/her *own consumption*. The survey data is structured differently: it asks for his/her beliefs about the return and *aggregate output*. To bridge the gap between data and the asset pricing equation, strong assumptions are required. First, it is assumed that results on over/underestimation of the volatility of output (for which data is available) carry over to consumption (which enters the asset pricing equation). Second, it is assumed that equity premia are not affected by non-insurable idiosyncratic risk (since no information on this risk is available). Third, it is also assumed that an aggregation of beliefs can be done by averaging across respondents.

These assumptions are unlikely to hold exactly, but they may 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 that there is 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, theoretical results on aggregation of heterogeneous beliefs suggest that an average belief may be a good approximation of the representative investor's belief—in particular when the relative risk aversion is high. For instance, Söderlind (2008b) shows (through an extension of a model proposed by Varian (1985)) that a representative investor should be assigned a mean equal to the cross-sectional mean, and a variance equal to a weighted average of the (cross sectional average of) individual uncertainty and the cross-sectional disagreement. The relative weight on disagreement is $1/(\gamma + 1)$: with log utility disagreement is as important as individual uncertainty (this was shown to be the case in Rubinstein (1974) and Detemple and Murthy (1994)). A higher value of risk aversion is more relevant from the perspective of the C-CAPM—which gives a much smaller role to disagreement. This suggests that it is reasonable to abstract from the disagreement and aggregate by calculating the average individual uncertainty.

3 Evidence from Survey and Option Data

3.1 Expected Stock Return

The *Livingston survey* summarises 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 lists the expected Standard & Poor (S&P) index level 6 and 12 months ahead.

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

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

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 problems to do with the base level used to calculate returns and the investment horizon is undoubtedly 6 months (see Dokko and Edelstein (1989) for a discussion). However, none of the conclusions are affected by using alternative base levels or horizons.²

Table 1 shows annualised results for the sample from June 1952 to December 2007. 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 (median across respondents) and in ex post data. The survey and ex post samples are matched. For instance, the last data point compares the forecast made in December 2006 to the ex post return for July–December 2007.

In the early sample (1952–1990) the Livingston survey has an expected capital gain

²There are two main alternatives in the literature. First, the survey contains 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 (1984) use base value from the last day of the 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. Since 1992, the survey has also asked for a forecast of the current month, which could be 'used' as a base value (see Park (2006)).

	1952–1990	1990–2007	1952–2007
Survey expectations (capital gain)	−1.1	1.6	−0.3
S&P Industrials, ex post (capital gain)	3.2		
S&P 500, ex post (capital gain)		5.2	
S&P combined, ex post (capital gain)			3.8
Dividend yield S&P 500	4.0	2.0	3.4

Table 1: Average capital gains in excess of a riskfree rate, dividend yield, annualised %.

The table shows results for the 6-month horizon, with forecasts made 12 months earlier, 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 the surveys 1952:6–1990:6 and S&P 500 Composite for the surveys since 1990:12) — 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 1990:12–2007:12.

(in excess of a riskfree rate) of -1.1% (annualised), while the ex post capital gain on the S&P Industrials is 3.2% .³ In the late sample (1990–2007), the survey has an expected capital gain of 1.6% while the ex post capital gain on the S&P 500 is 5.2% (both in excess of the riskfree rate).

The capital gain does not include dividends. Therefore, Table 1 also reports the ex post dividend yield on the S&P 500 index: it is around 4.0% and 2.0% for the early and late samples, respectively. The survey contains no information on 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 S&P 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% for both samples.

To sum up, the result of the Livingston data suggest that, over the last half century, expected returns on U.S. equities have been substantially lower than in ex post data. This corroborates earlier findings of substantial underestimation of average excess returns. For instance, Lakonishok (1980) and Park (2006) both document an underestimation in the

³The surveys conducted in December 1989 and June 1990 were potentially faulty, but excluding these dates is of no consequence: the average survey expectation does not change and the average ex post capital gain changes from 3.2% to 3.4% .

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. During 1992–2007 the implied expected excess return is 3.8% .

Livingston survey (for 1947–1974 and 1992–2002, respectively) and Claus and Thomas (2001) do the same from analysts earnings forecasts for 1985–1998. The underestimation also squares well with the notion of the U.S. market as one of the “lucky cases” (see, for instance, Brown, Goetzmann, and Ross (1995)) and with the idea of investors who are too pessimistic (see Abel (2002)).

Figure 1: Expected capital gains on the S&P index in excess of a riskfree rate, %. This figure shows the expected capital gains (in excess of a riskfree rate) for the 6-month horizon starting 6 months ahead (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 it is not really obvious that the results are representative of the market expectations. Studying the properties of the forecasts is one (admittedly informal) way of assessing the evidence. The rest of this section is devoted to this.

The implied average expectation error is around 4% (annualised) for the combined sample—which is large. The t-stat for the null hypothesis of no (average) expectation error is 2.5, which would be borderline significant once the fat tails of the return distribution are taken into account. It therefore cannot be argued that the survey data is implausible.

To assess the time pattern of the survey expectations, consider *Figure 1*. It shows the implied expected capital gains in excess of a riskfree rate, with and without adding the ex post dividend yield. The NBER recessions are marked by the shaded areas. The

Figure 2: Expected and ex post capital gains on the S&P index in excess of a riskfree rate, %. This figure shows the 6-month ex post capital gains (in excess of a riskfree rate) and its forecast 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.

Livingston forecasters clearly did not believe in the random walk hypothesis: there are distinct movements in the expected capital gains. In particular, the expectations (which are for 6 to 12 months ahead) have local maxima in almost all recessions. In contrast, expectations for 1 to 6 months ahead (not shown), have local minima in most recessions. This suggests a belief in short term momentum, but a medium term mean reversion. In any case, the forecasting performance is poor: the R-square of the survey forecasts (in trying to predict the movements around the mean) is virtually zero.

The dotted curve in Figure 1 adds 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 a number of periods with negative expected returns, which are somewhat unappealing. Clearly, if this represents the market returns and we assume a mean-variance setting, then a risk averse representative investor cannot expect negative excess returns. On the other hand, in a consumption-based model nothing rules out negative risk premia. Duffee (2005) estimates a time-varying covariance of consumption growth and stock market returns and finds negative values in many periods, for instance, in the early 1960s and 1980s (which fits the negative expected returns in Figure 1 fairly well).

If we instead take the position of Campbell and Thompson (2008) that negative expected excess returns should be replaced by zeros, then the average expected capital gain (in the survey) would be 1.5% instead of -0.3% . This is still a lot lower than the average ex post capital gain (3.8%).

Figure 2 compares the expectations with the ex post capital gains. The figure suggests that the underestimation may well be driven by a few exceptional periods. For instance, there were *prolonged* periods of consistent underestimation 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 overprediction (mainly 2001–2002), they tend to be less dramatic. All other periods look like random noise, where short-lived under- and over-predictions are more or less cancelled out. While the post Korean war and dot-com booms are defining periods in U.S. stock market history, it seems reasonable to assume that they were mostly unexpected. If so, the survey data tells a very different story about expected returns than the ex post data.

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

3.2 Consumption Growth Volatility

Most surveys ask their respondents to supply their point forecasts—and very few ask for measures of uncertainty. The *Survey of Professional Forecasters* (SPF) is an exception, as it asks for histograms of the respondents' subjective probability distribution of real GDP growth. 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 (there are no histograms for consumption growth), it is plausible that the results (in relative terms) carry over to consumption growth volatility.

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 of real GDP was only introduced in 1981Q3. The Federal Reserve Bank of Philadelphia now administers the survey (before 1990 it was administered Victor Zarnowitz and others from the American Statistical Association and the National Bureau of Economic Research).

Survey respondents are asked to assign probabilities of different GDP growth intervals

defined by the survey manager, for instance, for GDP growth rates between 0% and 0.9%, between 1% and 1.9%, and so on. The growth rates are defined as year-on-year (the value in calendar year t compared to the value in calendar year $t - 1$). Therefore, the Q1 survey can be thought of as referring to approximately 4-quarter growth, the Q2 survey referring to 3-quarter growth, etc. See Croushore (1993) and the Appendix for more details.

Giordani and Söderlind (2003, 2006) discuss how to estimate the underlying distribution from a histogram reported by a SPF respondent. 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 non-linear least squares. 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 calculate a cross-sectional “consensus” variance for each period by taking the cross-sectional median—and then form a time-average.

To be comparable, the “ex post volatility” should be a conditional variance, that is, the variance of forecast errors. I estimate this by the variance of the actual GDP growth rate minus the consensus (cross-sectional median) forecast. As was done earlier, the survey and ex post samples are matched.

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

Table 2: Standard deviation of year-on-year output growth 1981–2007, %. The table shows results for three different forecasting horizons. The survey data are from the Survey of Professional Forecasters, and includes the surveys 1981Q3 to 2006Q4. 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 (for 1981–2007), using the consensus point forecast and actual outcomes.

Table 2 reports on the results for three horizons (2 to 4 quarters). The standard deviation from the subjective distributions for the 3-quarter horizon is 0.7%, while the ex post (forecast error) standard deviation is 0.9%: a ratio of 0.8. This difference is significant at the 5% level (based on a delta method approach). The 2- and 4-quarter horizons show similar results (ratios of 0.6 and 0.7).⁵ Giordani and Söderlind (2006) came to a

⁵Some data points in 1985Q1 and 1986Q1 are likely to be wrong (see Federal Reserve Bank of Philadelphia (2000)). However, the figures in the table remain unchanged after those dates are excluded.

similar conclusion using a somewhat different method (the coverage ratio of subjective confidence bands).

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), suggesting that the *subjective consumption growth volatility is only around 0.7 times the historical volatility*. This scales up the required risk aversion (γ) in asset pricing equation (2) by 10/7.

3.3 Evidence from Options: Stock Return Volatility

There is very little survey evidence on 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 the expiration of the option. It is straightforward to calculate this “implied volatility”—and to compare it with the ex post (realised) return volatility over the same horizon. This is (essentially) what this section does.

The option data (daily, from 1985 onwards) contains information on call and put options on the S&P 500 index—for various strike prices and maturities. I calculate the implied volatility for each option 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 to the old VIX), and one that Ederington and Guan (2002) found 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. In addition, only those observations for which there was trade are utilised—to avoid using non-informative price quotes.

To get results for four different fixed horizons (2 to 4 quarters), I interpolate between maturities. As was done earlier, the option and ex post data samples are matched. For instance, the last data point for the 2-quarter horizon compares the implied volatilities on 29 June 2007 (a Friday) with the realised volatility from July to December 2007.

Table 3 reports the results for three horizons (2 to 4 quarters). The evidence shows that the implied volatilities are 15%–20% higher than the volatility of ex post data. This difference is clearly significant at the 5% level (adjusting the test for overlapping forecasting

periods and using a delta method approach).

	4 quarters	3 quarters	2 quarters
Options	19.4	19.3	18.9
Ex post	16.3	16.3	16.3

Table 3: Standard deviation of S&P 500 returns 1985–2007, annualised %. The table shows results for three different horizons. The results for options are implied volatilities calculated from options on the S&P 500 index (daily data 1985–2007:6). The ex post volatility is estimated from daily S&P 500 returns, 1985–2007:12. The standard deviations are calculated as the square root of the time averages of the variances.

This result may be criticised for being based on a faulty option pricing model which assumes that volatility is non-random (and asset returns are normally distributed). For this reason, two extensions will be discussed.

First, in a stochastic volatility model 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 (Hull and White (1987)). A Ball and Roma (1994) approximation then shows that around 0.5 should be added to the implied volatilities in Table 3—strengthening the results.

Second, in a model like that of Heston (1993), the stochastic volatility generates an additional risk premium. The options are then informative only about the variance of the riskneutral distribution. In this case, it is possible that the “overestimation” in Table 3 is just a risk premium. This is essentially an open issue (see, for instance, Jackwerth (2000) for a discussion of “overpriced” options). Without survey data, it is really impossible to tell. However, *Figure 3* shows an interesting time pattern for the implied and ex post volatilities. For the sake of clarity, it plots only the non-overlapping data points for the 2-quarter horizon (while the table above is based on daily data). The “overestimation” is persistent, but it seems to be related to the level of uncertainty: more overestimation when there is less uncertainty. This makes the interpretation in terms of risk premia a bit counter-intuitive, since the periods of low uncertainty levels also have a stable level of uncertainty. This could be systematic expectations errors instead.

This evidence suggests that the *subjective return volatility is around 1.15 times the historical volatility*, but only if the risk premium for stochastic volatility is not too important. This reduces the required risk aversion (γ) in asset pricing equation (2) by the same fraction.

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 implied volatility is shifted forward 6 months to allow for a direct comparison between the forecast (implied volatility from the options) and the outcome.

3.4 Implications for the Equity Premium

The results from the 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 0.7. Unfortunately, there is no survey or option data on the correlation of consumption and returns. Therefore, the conclusions are based on the assumption that the (unobservable) correlation in subjective beliefs is not considerably smaller than in ex post data.

Together this changes the asset pricing equation (2) to

$$\underbrace{E(R_t^e)}_{0.07 \times 0.5} = \underbrace{\text{Corr}(R_t^e, \Delta c_t)}_{?} \underbrace{\sigma(R_t^e)}_{0.14 \times 1.15} \underbrace{\sigma(\Delta c_t)}_{0.01 \times 0.7} \gamma. \quad (3)$$

This means that the left hand side is scaled by 0.5 and the right hand side by 0.8—so the equation would hold for a lower value (factor 3/5) of the relative risk aversion, γ . For instance, with a perfect correlation $\gamma = 30$ (instead of 50 in ex post data) is required to make this equation hold. This is a step in the right direction.

4 Summary

This paper performs a traditional analysis of the equity premium puzzle—but using data from surveys and options.

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 3.5% instead of around 7%. At the same time, option data indicate that investors may overestimate the standard deviation of equity returns by a factor of 1.15. Both findings work in favour of the standard consumption-based asset pricing model.

Unfortunately (for C-CAPM), data on subjective beliefs about output growth volatility (Survey of Professional Forecasters) show marked overconfidence—the standard deviation implied by the survey data is just 0.7 of the ex post 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. Through the use of survey and option data, the coefficient of relative risk aversion need only be three fifths of what ex post data requires.

A Data Appendix

The data of the Livingston Survey and the Survey of Professional Forecasters are taken 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 June 1952 Livingston survey asked about the S&P 365 Industrials and the December 1952 to December 1956 about the S&P 420 Industrials. These indices had 1935–39=100 as a base. For June 1957 to June 1990 the survey asked about the S&P (425) Industrials, with 1941–43=10 as a base.

The Survey of Professional Forecasters asked about probabilities for real GNP during 1981Q3–1991Q4 and real GDP since 1992Q1. Before 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.

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 are 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, "A quartet of semi-groups for model specification, robustness, prices of risk, and model detection," *Journal of the European Economic Society*, 1, 68–123.
- 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 J. H. Cochrane, 1999, "By force of habit: a consumption-based explanation of aggregate stock market behavior," *Journal of Political Economy*, 107, 205–251.

- Campbell, J. Y., and S. B. Thompson, 2008, "Predicting the equity premium out of sample: can anything beat the historical average," *Review of Financial Studies*, 21, 1509–1531.
- Claus, J., and J. Thomas, 2001, "Equity premia as low as three percent? Evidence from analysts' earnings forecasts for domestic and international stock markets," *The Journal of Finance*, 56, 1629–1666.
- 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.
- Dokko, Y., and R. H. Edelstein, 1989, "How well do economists forecast stock market prices? A study of the Livingston surveys," *The American Economic Review*, 79, 865–871.
- 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, 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, 2006, "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.
- Merton, R. C., 1980, "On estimating the expected return on the market," *Journal of Financial Economics*, 8, 323–361.
- Park, C., 2006, "Rational beliefs or distorted beliefs: equity premium puzzle and micro survey data," *Southern Economic Journal*, 72, 677–689.

- Parker, J., and C. Julliard, 2005, "Consumption risk and the cross section of expected returns," *Journal of Political Economy*, 113, 185–222.
- 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.
- Söderlind, P., 2008a, "An extended Stein's lemma for asset pricing," *Applied Financial Economics Letters*, forthcoming.
- Söderlind, P., 2008b, "Why disagreement may not matter (much) for asset prices," *Finance Research Letters*, forthcoming.
- 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.
- Vissing-Jorgensen, A., and O. P. Attanasio, 2003, "Stock market participation, intertemporal substitution and risk aversion," *American Economic Review*, 93, 383–391.