

# DISCUSSION PAPER SERIES

No. 5407

**C-CAPM WITHOUT *EX POST* DATA**

Paul Söderlind

***INTERNATIONAL MACROECONOMICS***



**Centre for Economic Policy Research**

**[www.cepr.org](http://www.cepr.org)**

Available online at:

**[www.cepr.org/pubs/dps/DP5407.asp](http://www.cepr.org/pubs/dps/DP5407.asp)**

## **C-CAPM WITHOUT *EX POST* DATA**

**Paul Söderlind**, Swiss Institute of Banking and Finance and CEPR

Discussion Paper No. 5407  
December 2005

Centre for Economic Policy Research  
90–98 Goswell Rd, London EC1V 7RR, UK  
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999  
Email: [cepr@cepr.org](mailto:cepr@cepr.org), Website: [www.cepr.org](http://www.cepr.org)

This Discussion Paper is issued under the auspices of the Centre's research programme in **INTERNATIONAL MACROECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as a private educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Paul Söderlind

CEPR Discussion Paper No. 5407

December 2005

## **ABSTRACT**

### **C-CAPM Without Ex Post Data\***

Survey and option data are used to take a new 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 (CBOE's VIX) indicates that investors 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.

JEL Classification: E130, E320 and G12

Keywords: CBOE VIX, equity premium puzzle, Livingston survey and Survey of Professional Forecasters

Paul Söderlind  
Swiss Institute for Banking & Finance  
University of St. Gallen  
Rosenbergstr. 52  
CH-9000 St. Gallen  
Switzerland

Tel: (41 71) 224 7080

Fax: (41 71) 224 7088

Email: paul.soderlind@unisg.ch

For further Discussion Papers by this author see:

[www.cepr.org/pubs/new-dps/dplist.asp?authorid=114406](http://www.cepr.org/pubs/new-dps/dplist.asp?authorid=114406)

\* I thank Michael Fischer for excellent research assistance and Anna Cieslak for comments.

Submitted 17 November 2005

# 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 tested 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 expectations errors—and is therefore a way to test the model more directly.

Under some simplifying assumptions, the key implication of the consumption based asset pricing model is that expected excess equity returns are a product of a risk aversion coefficient, the standard deviations of consumption growth and equity returns, and their correlation.

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. The CBEO volatility index (VIX) is 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. Section 3 presents the survey data and the empirical results. Section 4 sums up. There are several appendices giving details on data and derivations of some analytical results.

## 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  $\gamma$  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 lemma<sup>1</sup> to rearrange the Euler equation as

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

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

where  $\Delta c_t$  is the growth rate of consumption,  $\ln(C_t/C_{t-1})$ . It is worth emphasising

---

<sup>1</sup>Stein'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).

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  $\Delta c_t$  has a bivariate mixture normal distribution—provided  $\Delta c_t$  has the same mean and variance in all the mixture components (see Appendix B). 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 is that the unknown relative risk aversion,  $\gamma$ , 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.<sup>2</sup>

The equity premium puzzle (see Mehra and Prescott (1985)) is that (2) does not hold for ex post data unless  $\gamma$  is very high. To illustrate this, I use quarterly US data for 1957Q1–2003Q4. 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

$$\underbrace{E(R_t^e)}_{0.06} = \underbrace{\text{Corr}(R_t^e, \Delta c_t)}_{0.14 \text{ or } 0.31} \underbrace{\sigma(R_t^e)}_{0.16} \underbrace{\sigma(\Delta c_t)}_{0.01} \gamma. \quad (3)$$

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,  $\gamma$ , of 38 to make (3) hold.

---

<sup>2</sup>The 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 Evidence from Survey Data and Options

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.

#### 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,  $E[E_t(R_t^e)] = E(R_t^e)$ ). 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 deviations are then the square roots of the average conditional variances—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 representative 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 from 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 Federal Reserve Bank of Philadelphia in 1990. It has questions about the expected Standard & Poor (S&P) index level 6 and 12 months ahead in time.

The documentation of the Livingston survey before 1990 (when the Federal Reserve Bank of Philadelphia took over) is somewhat sketchy. To make things worse, the Standard&Poor made major changes to their indices on several occasions. However, it is clear that for the period June 1957 (and perhaps even earlier) to June 1990, the Livingston survey asked for the S&P 400 Industrial index (which is different from the more recent S&P 400 Midcap). 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 (2004) and Appendix A. The median forecasts (as a robust proxy of the mean) and the actual indices are displayed in *Figure 1* in form of the logarithm of the index. The survey forecasts follow the general movements of the index, but are mostly





**Figure 1: S&P indices and Livingston survey.** The figure shows the S&P index and the Livingston forecast made in the same period.

above the current index level.

The survey is sent out in late May and November, respectively and asks for the index level six and twelve months ahead. This makes it somewhat unclear what base value to use in order to calculate expected capital gains (growth rates). I follow the approach in Pearce (1984). For the June survey, this means this following: the base value for the June survey is from the last day of May, and it is assumed that the 6-month (12-month) forecast refers to the index level at the end of December (next June). The implied expected growth rate therefore refers to a 7-month (13-month) horizon. A similar logic applies to the December survey. A way to circumvent both the problem with the base level and the investment horizon is to combine the 6- and 12-month forecasts to calculate an implied expected growth rate over a 6-month horizon starting 6 months from now. I will use this

to check the other results.

*Table 1* shows the (annualised) results for the 7-month horizon in the form of capital gains in excess of a riskfree rate (approximated by a 3-month T-bill rate). The results for the 13-month horizon and the 6-month horizon starting 6 months from now are very similar.

For the early sample the Livingston survey has an expected capital gain (in excess of a riskfree rate) of  $-1.3\%$ , while the ex post capital gain on the S&P 400 (Industrials) was  $1.3\%$ . For the late sample, the survey has an expected capital gain of  $0.7\%$ , while the ex post capital gain on the S&P 500 (Composite) was  $5.6\%$ .

|                               | 1957–1990 | 1990–2003 | 1957–2003 |
|-------------------------------|-----------|-----------|-----------|
| Survey expectations           | $-1.3$    | $0.7$     | $-0.7$    |
| S&P 400 Industrial, ex post   | $1.3$     |           |           |
| S&P 500 Composite, ex post    | $1.6$     | $5.6$     | $2.7$     |
| S&P combined 400/500, ex post |           |           | $2.5$     |
| Dividend yield S&P 500        | $4.0$     | $2.2$     | $3.5$     |

**Table 1: Capital gains in excess of riskfree rate, dividend yield, annualised %.** The table shows results for the 7-month horizon, annualised by multiplying by  $12/7$ . The survey data are from the Livingston survey and show the expected change of the stock index (S&P 400 Industrial for 1957:6–1990:6 and S&P 500 Composite for 1990:12–2003:6) — in excess of a riskfree rate. The ex post results are the corresponding actual change in the stock index. The S&P combined 400/500 uses S&P 400 Industrial for 1957:6–1990:6 and S&P 500 Composite for 1990:12–2003:6.

To get a result for the full sample, we have to think of an index which consists of S&P 400 during the early period and S&P 500 during the late sample. This is probably not too far-fetched since the S&P 400 and 500 are so similar during the overlapping period (early sample): the mean capital gains were  $1.3\%$  and  $1.6\%$  respectively, and their correlation was  $0.96$ . For this spliced sample, the survey has an expected capital gain of  $-0.7\%$  while the ex post has a capital gain of  $2.7\%$ .

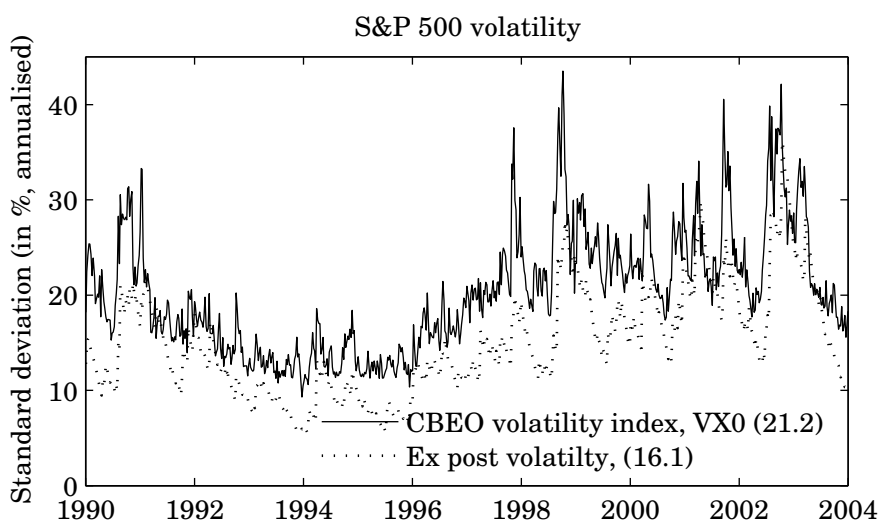
Clearly, the capital gain on the indices is not the total return from investing on the stock market. Table 1 therefore also gives the ex post dividend yield on the S&P 500 index, which is around  $3.5\%$  for the spliced sample. The survey contains no information about expected dividends, but it seems reasonable to assume that there is no bias in the dividend forecasts—especially since dividends are fairly smooth over time. In that case, the results for the whole sample suggest that the expected total return was around  $2.8\%$

(−0.7% in capital gain and 3.5% in dividend yield).<sup>3</sup>

In short, this evidence suggest that the *expected excess return*  $E(R_t^e)$  in (2) *might be around 3%. This is half of the historical average.*

### 3.3 Stock Return Volatility

We now study the return volatility,  $\sigma(R_t^e)$ , in (2). There is only very little survey evidence on this, but implied volatilities from option prices provide a good alternative. CBOE compiles an index of S&P 500 volatility (denoted VIX) by averaging the implied volatilities of eight near-the-money options (with 30 days to expiration) on the S&P 500 index. However, data is only available from 1990 onwards.



**Figure 2: Volatility of S&P 500 changes.** The volatility index is the CBOE VIX index and the ex port volatility is estimated as exponential weighted moving average of de-meanded price changes (weekly) of the S&P 500 changes. The numbers in parentheses are the implied unconditional standard deviations.

*Figure 2* compares the volatility index to volatility of ex post data for the S&P 500 price changes. The ex post volatility is estimated by an exponential moving average similar to the RiskMetrics approach ( $\sigma_t^2 = 0.1u_{t-1}^2 + 0.9\sigma_{t-1}^2$ , where  $u_t$  are the demeaned price changes). It can be shown that a GARCH-model produces very similar results. The two measures of volatility are strongly correlated, but the VIX index is generally higher.

<sup>3</sup>The Survey of Professional Forecasters ask for the expected 10-year S&P 500 return as well as the average 3-month T-bill rate. For the 1992–2004 sample the implied expected excess return is 4%.

If returns are unpredictable, then the unconditional standard deviation is the square root of the average conditional variance. This is 21% from the VIX (annualised) but only 16% for the ex post returns: the subjective standard deviations seems to be around 1.3 times the historical volatility. Results for the VXO index from 1986 onwards (the “old VIX”) on S&P 100 volatility are very similar.

However, there are several potential problems with these results. First, the horizon is only 30 days, which is much shorter than the survey data for the expected returns and consumption volatility. Second, the implied volatility of options near-the-money may not be a good estimator of the standard deviation of the underlying distribution unless option prices are well approximated by the Black-Scholes model.

It is therefore interesting to compare with the results of Lynch and Panigirtzoglou (2003). They study (among other things) 6-months S&P500 options over the period 1983–2002 and find that the implied volatility (at the money) is on average 17% and that the standard deviation of the riskneutral distribution (estimated using options at a wide range of strike prices) is on average 18.5%. Since these figures refer to the average standard deviation rather than the unconditional standard deviation discussed before (the square root of the average variance), they are likely to be downward “biased” compared with the previous results. The data for the VXO and the ex post volatility suggests that this bias is around 1 percentage point. This means that the unconditional standard deviation derived from the riskneutral distributions is around 19.5%, which should be compared with an ex post volatility of 16% (for the same sample). Provided the risk premium does not change the shape of the distribution much, this suggests a ratio of 1.2 between the subjective standard deviation and the ex post volatility.

Together this evidence suggest that the *return volatility*  $\sigma(R_t^e)$  in (2) *might be around 1.25 times the historical volatility.*

### **3.4 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.<sup>4</sup> I will use this information to compare the subjective uncertainty to actual

---

<sup>4</sup>Real GNP 1981Q3–1991Q4 and real GDP 1992Q1 and onwards,

uncertainty (the volatility of forecast errors). Although this provides evidence only on output growth volatility, output growth and consumption growth are so strongly related that it is plausible that the results carry over to consumption growth.

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

The histogram bins are set by the survey manager, and forecasters assign probabilities to each bin.<sup>5</sup> The growth rates are defined as the value in (calendar) year  $t$  divided by the value in year  $t - 1$ , minus one. I use forecasts of this growth rate at four forecasting horizons: one to four quarters ahead. In practice, the four quarters ahead forecast is made in Q1 of year  $t$ , the three quarters ahead forecast is made in Q2 and so forth.

Giordani and Söderlind (2005) and Giordani and Söderlind (2003) discuss how to estimate the underlying distribution from such a histogram. They find that fitting a normal distribution to each histogram typically works well. They also find that some forecasters supply occasionally odd answers, so cross-sectional moments (central location and dispersion) need to be estimated by “robust” methods—for instance, the median instead of the mean.

Therefore, for each forecaster and quarter, 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 time series (for each forecaster) of his/her subjective variance. To get an aggregate number, the cross-sectional median (of the individual subjective variances) is calculated for every time period—and a time average is then formed. This procedure ensures that the aggregate number gives equal weight to all time periods—which circumvents the problem with that the number of participating forecasters has changed over time.

For the “ex post volatility” I calculate the forecast error of every forecast and then the forecast error variance for each forecaster. To get an aggregate number that is comparable to the aggregate subjective variance, I first calculate the cross-sectional median variance

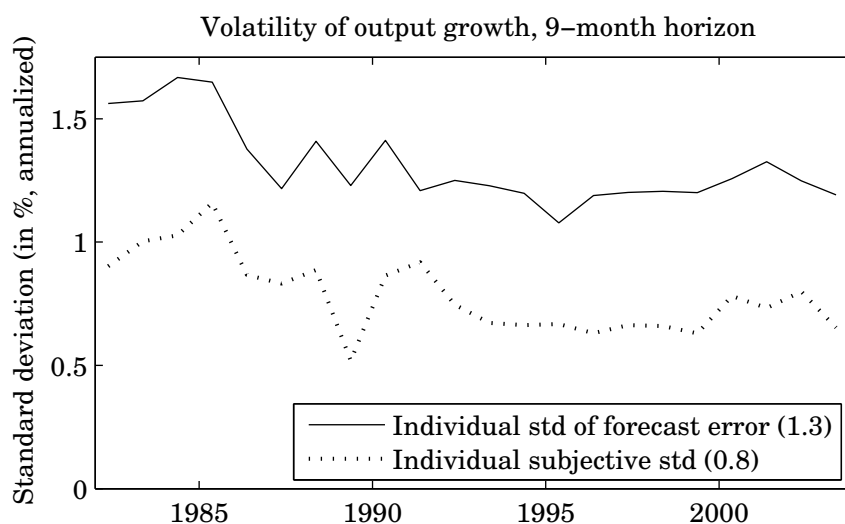
---

<sup>5</sup>Before 1992, the bins boundaries were -2, 0, 2, 4, 6. Since 1992, they are -2, -1, 0, 1, 2, 3, 4, 5, 6.

of those forecasters that participate in the survey at that point in time, then I form a time average.

Figure 3 shows results for the 9-months horizon. The other horizons look very similar.<sup>6</sup> It is clear that forecasters have underestimated the uncertainty: the subjective standard deviation is only 60% of the standard deviation of the forecasters' own forecast errors. Giordani and Söderlind (2005) reach a similar conclusion by using somewhat different methods (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).

To sum up, this evidence suggest that the *consumption growth volatility*  $\sigma(\Delta c_t)$  in (2) might be around 0.6 times the historical volatility.



**Figure 3: Survey data on output growth volatility on the 9-month horizon.** The numbers in parentheses are the implied unconditional standard deviations.

### 3.5 Implications

The results from survey and option data suggests that the historical mean excess returns should be scaled by 0.5, the standard deviation of excess returns by 1.25, and the standard

<sup>6</sup>I show results for the 9-month horizon since some data points (1985Q1 and 1986Q1) for the 12-month horizon are likely to be wrong (see Federal Reserve Bank of Philadelphia (2000)).

deviation of consumption growth by 0.6. This changes the asset pricing equation (3) to

$$\underbrace{E(R_t^e)}_{0.5 \times 0.06} = \underbrace{\text{Corr}(R_t^e, \Delta c_t)}_{?} \underbrace{\sigma(R_t^e)}_{1.25 \times 0.16} \underbrace{\sigma(\Delta c_t)}_{0.6 \times 0.01} \gamma. \quad (4)$$

This means that the left hand side is scaled by 0.5 and the right hand side by 0.75—so the equation would hold for a lower value (factor 2/3) of the relative risk aversion,  $\gamma$ . For instance, with a perfect correlation we need  $\gamma = 25$  (instead of 38 in ex post data) to make this equation hold. This is a step in the right direction—although not a very large one.

## 4 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 (CBOE’s VIX) indicate that investors might overestimate the standard deviation of equity returns by a factor 1.25. Both these factors work in favour of the standard consumption-based asset pricing model.

Unfortunately (for the model), data on subject beliefs about output growth volatility (Survey of Professional Forecasters) show 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 2/3 of what ex post data requires.

## A Data Appendix

The data from the Livingston Survey is from the Federal Reserve Bank of Philadelphia (<http://www.phil.frb.org/>). For details, see Federal Reserve Bank of Philadelphia (2004).

The S&P 500 and S&P 100 price indices, as well as the CBOE VIX and VXO volatility indices are from Datastream. Details on the volatility indices are found at

<http://www.cboe.com/micro/vix/introduction.asp>.

The S&P 400 (industrial) is from the Federal Reserve Bulletin (various issues). 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 (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 NIPA Table 8.7 (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/>).

## B Derivation of Stein's Lemma for a Special Case of Mixture Normals

This section proves that Stein's lemma continues to hold if  $x$  and  $y$  has a bivariate mixture normal distribution, but that the marginal distribution of  $y$  is normal.

Let the pdf of  $(x, y)$  be a mixture of  $n$  bivariate normal distributions

$$pdf \left( \begin{bmatrix} x \\ y \end{bmatrix} \right) = \sum_{i=1}^n \alpha_i \phi \left( \begin{bmatrix} x \\ y \end{bmatrix}; \begin{bmatrix} E_i(x) \\ E_i(y) \end{bmatrix}, \begin{bmatrix} \text{Var}_i(x) & \text{Cov}_i(x, y) \\ \text{Cov}_i(x, y) & \text{Var}_i(y) \end{bmatrix} \right), \text{ with } \sum_{i=1}^n \alpha_i = 1,$$

and where  $\phi(z; \mu, \Sigma)$  is the normal pdf with mean vector  $\mu$  and covariance matrix  $\Sigma$ .

Direct calculations give

$$\text{Cov}[x, h(y)] = \sum_{i=1}^n \alpha_i \{ \text{Cov}_i[x, h(y)] + E_i(x) E_i[h(y)] \} - E(x) E[h(y)].$$

If  $E_i[h(y)]$  is a constant  $E[h(y)]$ , then this simplifies to

$$\text{Cov}[x, h(y)] = \sum_{i=1}^n \alpha_i \text{Cov}_i[x, h(y)].$$

Since  $\text{Cov}_i[x, h(y)] = \text{Cov}_i(x, y) E_i[h'(y)]$  for all states, we get

$$\text{Cov}[x, h(y)] = E[h'(y)] \sum_{i=1}^n \alpha_i \text{Cov}_i(x, y).$$

If  $E_i(y) = E(y)$ , then the sum equals  $\text{Cov}(x, y)$ . Note, however, that the combination of  $E_i(y) = E(y)$  and  $E_i[h(y)] = E[h(y)]$  requires that  $\text{Var}_i(y) = \text{Var}(y)$ , which means that the marginal distribution of  $y$  is a normal distribution.



## References

- 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," Stanford University Working Paper.
- Brennan, M. J., 1998, "The Role of Learning in Dynamic Portfolio Decisions," *European Finance Review*, 1, 295–306.
- 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.
- Detemple, J., and S. Murthy, 1994, "Intertemporal Asset Pricing with Heterogeneous Beliefs," *Journal of Economic Theory*, 62, 294–320.
- 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, "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*, forthcoming.
- Granger, C. J., and R. Ramanathan, 1984, "Improved Methods of Combining Forecasts," *Journal of Forecasting*, 3, 197–204.
- Hirshleifer, D., 2001, "Investor Psychology and Asset Pricing," *Journal of Finance*, 56, 1533–1597.
- Jones, C. P., 2002, "An Analysis of the S&P 500 Index and Cowles's Extensions: Price Indexes and Stock Returns, 1870–19999," *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*, forthcoming.
- Lewellen, J., and J. Shanken, 2002, "Learning, Asset-Pricing Tests, and Market Efficiency," *Journal of Finance*, 57, 1113–1145.
- Lynch, D., and N. Panigirtzoglou, 2003, "Summary Statistics of Option-Implied Probability Density Functions and their Properties," mimeo, Bank of England.
- 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.
- Thaler, R. H., 2000, "From Homo Economicus to Homo Sapiens," *Journal of Economic Perspectives*, 14, 133–141.
- Tornell, A., 2000, "Robust- $H_\infty$  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.