

Quantitative Structuring vs the Equity Premium Puzzle

Andrei N. Soklakov*

25 July 2015

Quantitative Structuring is a rational framework for manufacturing financial products. It shares many of its components with mainstream economics. The Equity Premium Puzzle is a well known quantitative challenge which has been defying mainstream economics for the last 30 years. Does Quantitative Structuring face a similar challenge? We find Quantitative Structuring to be in remarkable harmony with the observed equity premium. Observed values for the equity premium (both expected and realized) appear to be a real and transparent phenomenon which should persist for as long as equities continue to make sense as an investment asset. Encouraged by this finding, we suggest a certain modification of mainstream economics.

1 Quantitative Structuring

Each and every financial product is completely defined by its payoff function F which states how the benefits (usually cash flows) depend on the underlying variables. In order to price a product, defined by its payoff F , we compute a quantity of the form

$$\text{Price}(F) \propto \sum_x F(x)Q(x), \quad (1)$$

where the summation is taken over all possible values of the underlying variables and where Q is given by a mathematical model for the variables. Equation (1) is probably the most famous formula in the whole of mathematical finance. It shows, among other things, that the value of a product is determined by its payoff structure F and the model Q in a nearly symmetric way.

*Head of Equities Model Risk and Analytics, Deutsche Bank.

The views expressed herein should not be considered as investment advice or promotion. They represent personal research of the author and do not necessarily reflect the view of his employers, or their associates or affiliates. Andrei.Soklakov@(db.com, gmail.com).

Product design clearly deserves as much technical attention and respect as modeling. In fact, one can argue that products are much more important than modeling for they define the very nature of a business. Quantitative Structuring recognizes the importance of financial products and provides a technical framework for their design [1].

Within Quantitative Structuring all investments begin with research. Ahead of any proposals, a minimum of two learning steps must happen. The investor needs to form an opinion on the market and to learn their own preferences (risk aversion). Mathematically these two steps are described by two equations:

$$b = f m \quad (2)$$

$$\frac{d \ln F}{d \ln f} = \frac{1}{R}. \quad (3)$$

These equations can be introduced by making just a couple of observations. Firstly, we observe that each and every investment is an exercise in optimization. Secondly, we note that the above equations are obeyed by a payoff function $F(x)$ which solves the following optimization problem [2]

$$\max_F \int b(x) U(F(x)) dx \quad \text{subject to budget constraint} \quad \int F(x) m(x) dx = 1. \quad (4)$$

The risk aversion coefficient R is connected to the utility U through the standard Arrow-Pratt formula: $R = -FU''_{FF}/U'_F$. The economic meaning of the market-implied and investor-believed distributions $m(x)$ and $b(x)$ follows from the above optimization.

For further explanations of these equations, including motivation, derivations, intuitive illustrations as well as concrete numerical examples, we refer the reader to [1], [2], [3], [4] and [5].

2 Confronting the Equity Premium Puzzle

In 1985 Mehra and Prescott investigated historical data on the excess returns achieved by equities over government bonds [6]. These excess returns, known as the equity premium, appeared to be surprisingly high. Mehra and Prescott concluded that the equity premium was an order of magnitude greater than could be rationalized within the standard utility-based theories of asset prices.

Given the importance of the challenge, proposals to resolve the puzzle quickly snowballed. More than two decades later Mehra and Prescott revisited the progress on the problem only to reinforce their original conclusions [7]. They estimated the equity premium to be 2-8% in arithmetic terms or up to 6% in terms of geometric (compound) returns and reiterated the Equity Premium Puzzle as a standing challenge to explain these values.

The work on understanding the equity premium continues. Many insightful observations have been made. The scope of proposals has widened enormously. It now ranges from plausible denials of the puzzle to behavioral explanations. The complexity of individual proposals also increased. With some proposals still awaiting adequate independent analysis, it would be fair to say that no single explanation of the puzzle has yet received general acceptance and the search for a clear dominant explanation continues.

A balanced review of the 30 year history of the puzzle is a major task in its own right which would lead us away from the main focus of this paper. For our purposes we need to know only one historical fact. We need to note that the puzzle has posed a major challenge to utility-based economic models. This makes the Equity Premium Puzzle a perfect challenge to Quantitative Structuring which, as we can see from the optimization (4), heavily relies on the expected utility theory.

How would we know if Quantitative Structuring survived the challenge? Of course, it would have to explain the numerical premium of 6% annualized compounded returns. Mehra and Prescott set additional guidelines in their most recent review [7]. They urge clear differentiation between expected and realized returns. They emphasize long-time historical horizons. Furthermore, they set an expectation that any theory which takes on the puzzle must be able to say something about the future of the puzzle. In other words, are the equity returns real and likely to persist or were they a statistical fluke with no material probability of re-occurring?

We accept the challenge with all of the above conditions. We investigate separately the expected and the realized returns. We use long-time horizons when talking about realized returns. Within Quantitative Structuring the observed numerical values of the equity premium appear to be absolutely real and natural. In fact, if these numerical values were somehow not known, Quantitative Structuring would have predicted them.

3 Expected premiums

Using the notation of (4), we can write the investor-expected continuously-compounded rate of return as

$$\text{ER} = \int b(x) \ln F(x) dx. \quad (5)$$

This quantity is determined by two things – the structure of the investment $F(x)$, and the investor-believed distribution $b(x)$.

As we focus on equity investments, we describe the investment structure as:

$$F(x) = x, \quad (6)$$

where x is a total return on one unit of wealth invested in the equity.

To get the believed distribution we need to know the investor's risk aversion. For example, in the case of a growth-optimizing investor $R = 1$, equation (3) becomes redundant, i.e. $F(x) = f(x)$, and Eq. (2) gives us the believed distribution

$$b_{\text{GO}}(x) = F(x) m(x) = x m(x). \quad (7)$$

The corresponding expected return becomes

$$\text{ER} \rightarrow \text{ER}_{\text{GO}} = \int (x \ln x) m(x) dx. \quad (8)$$

As an example, consider a log-normal market-implied distribution

$$\frac{m(x)}{\text{DF}} = \frac{1}{x\sigma\sqrt{2\pi}} \exp \left\{ -\frac{(\ln x - \mu)^2}{2\sigma^2} \right\}, \quad \mu = r - \sigma^2/2, \quad (9)$$

where DF is the discount factor, r is the risk free return and σ is the volatility. In this case the integral in Eq. (8) can be computed analytically with the result:

$$ER_{GO} \rightarrow ER_{GO}^{LN} = r + \sigma^2/2. \quad (10)$$

Mehra and Prescott considered an investor with arbitrary constant relative risk aversion. Generalization of the above calculations to this case is very easy. All we have to do is to bring into play Eq. (3) with a constant value of R . Equation (10) is then replaced by a slightly more general quantity (see Eq. (33) in the Appendix):

$$ER_R^{LN} = r + (R - 1/2)\sigma^2. \quad (11)$$

This gives us the expected premium of

$$EP_R^{LN} \stackrel{\text{def}}{=} ER_R^{LN} - r = (R - 1/2)\sigma^2. \quad (12)$$

In their pioneering paper [6], Mehra and Prescott argue that the acceptable values for R must be below 10. In fact, all of the actual estimates of R which they cite to support their argument were below 3. Even staying within this tight range below 3 and making the standard assumption of 20% for typical equity volatility we can easily explain premia as high as 10% in terms of continuously compounded annual returns. This ball-park range is in remarkable agreement with the values observed by Mehra and Prescott.

In the remainder of this section we are going to examine independent quotes for the expected risk premia and see what values of R they imply. Before we do that, let us restore the generality of our arguments by removing the above made assumption of log-normality. In the case of arbitrary market-implied distributions, Eq. (12) is replaced by the expression (see Eq. (30) in the Appendix):

$$EP_R \stackrel{\text{def}}{=} ER_R - r = \frac{1}{\text{Price}(x^R)} \frac{\partial \text{Price}(x^R)}{\partial R} - r. \quad (13)$$

Implying the value of R from this expression is considerably less convenient than using Eq. (12). Nevertheless, it is a simple root-finding problem which can be solved. In terms of technology, we just need the ability to price power payoffs, x^R , which can be done by replication with vanillas.

In terms of independent quotes for the equity premium we reach out to the field of equity valuations where the expected premium is a very important factor. On Fig. 1 we display expected equity premia as reported by Damodaran [8] using SPX data. It is important to note that these values are just as large as noted by Mehra and Prescott – at least an order of magnitude above 0.35%.




Figure 1: Implied Equity Premia as reported by Damodaran [8]. The records are updated on a monthly basis starting from September 2008. The quoted values refer to the beginning of each month. In our calculations we interpreted this as the first business day of each month.

There are always limits to how far in the future one can look using available market data. According to Damodaran [8], his quotes for the premia accurately reflect detailed market information (such as market-implied dividends) of up to five years into the future.

At five year horizons, equity skew is quite flat. This makes Eq. (12) useful as a test calculation which requires very little access to market data. On Fig. 2 we compute relative risk aversion from the quoted premia using both the exact Eq. (13) and the test Eq. (12).

In the former case we made no simplifying assumptions and used complete historical records of 5-year volatility curves. In the latter case we used 5-year at-the-money-forward implied volatilities (displayed for convenience on Fig. 3). The graphs for the two cases show good agreement.

All computed values of risk aversion are comfortably within the realistic range. We conclude that, in terms of investors' expectations, Quantitative Structuring is consistent with the observed equity premia.




Figure 2: Implied risk aversion. Solid and dashed lines correspond to Eqs. (13) and (12) respectively. In both cases the timing of investments is chosen consistently with the quoted values of implied risk premia, i.e. they are assumed to mature in five years starting on the first business day of each month.




Figure 3: SPXT 5-year at-the-money-forward values of implied volatility.

4 Realized premiums

In the above section we managed to reconcile rational expectations of equity premiums. In terms of numerical values, these expectations were just as high as reported by Mehra and Prescott [6]. In this section we would like to understand how such expectations materialize, with investors doing no more than just keeping their money in the equity.

Let S_t be the value of the total return version of some equity index at time t . The return on the equity investment can be partitioned arbitrarily into N imaginary reinvestment steps:

$$S_N = S_0 \cdot \frac{S_1}{S_0} \cdot \frac{S_2}{S_1} \cdots \frac{S_N}{S_{N-1}}. \quad (14)$$

Defining $x_i = S_i/S_{i-1}$ we compute

$$S_N = S_0 \prod_{i=1}^N x_i = S_0 e^{\sum_{i=1}^N \ln x_i} = S_0 e^{N \cdot \text{Rate}}, \quad (15)$$

where

$$\text{Rate} = \frac{1}{N} \sum_{i=1}^N \ln x_i. \quad (16)$$

Let us now look at the time series x_1, \dots, x_N using the standard statistical approach. In this approach the individual elements $\{x_i\}$ are viewed as realizations of a random variable X with some (possibly unknown) distribution $P(X)$. For the basic statistical concepts, like the average, to make practical sense, the law of large numbers is assumed to hold.¹ In this framework, as N increases, the average (16) converges almost surely to the expectation

$$\text{Rate} \xrightarrow{\text{a.s.}} \int P(x) \ln x \, dx. \quad (17)$$

Let us compare this equation with Eq. (5) (remember $F(x) = x$ for equity investments). We see that the investor-expected returns can be achieved provided that the time series is long enough (i.e. N is sufficiently large) and, crucially, that the investor correctly determines the probabilities, i.e. $b(x) \approx P(x)$. This gives us some information about equity investors. Our task now is to understand enough detail to see if it is realistic.

Mehra and Prescott describe the Equity Premium Puzzle as a long-term phenomenon. This discourages us from considering very short reinvestment periods. Ideally, we want to consider the case of smallest possible N that is large enough to ensure noticeable convergence (17). The standard deviation of the sum (16) from its mean (17) scales as $N^{-1/2}$. For the first significant digit of the sum (16) to emerge with some reasonable probability, the convergence must reduce the standard deviation by an order of magnitude ($N^{-1/2} \sim 0.1$). This means that we must choose N which is not much lower than 100.

We managed to find full market data, including volatility surfaces, for SPXT (total return version of SPX) going back to 17 May 2000. At the time of writing, this was about 15 years worth of data (daily records). Some researchers might argue the need for longer historical records. However, 15-year investments are already at the limit of what many

¹This can be ensured if the individual values are sufficiently independent.

people would consider practical, so we choose to accept it. Viewing 15 years of the entire investment history (14) as if it was a sequence of bi-monthly reinvestments we get $N = 90$ reinvestment periods.

We need access to the distribution $P(x)$. One way of defining a probability distribution is to imagine a source of numbers distributed according to this distribution. Given such a source one can estimate expectations using the Monte-Carlo method. In terms of such a definition for the distribution of the actual realized returns, $P(x)$, all we have is a set of $N = 90$ values $\{x_i\}_{i=1}^N$. As discussed above, this is just enough to talk about expectations like (16).

Consider an investor whose original belief happened to coincide with the actual realized distribution, $b(x) = P(x)$. For this investor, the expected return is given by equation (16) which, by construction, evaluates to the actual realized returns exactly. The analysis of the realized equity premium boils down to the analysis of whether such an investor is realistic. Following Mehra and Prescott, this means computing and examining the investor's risk aversion.

Using Eqs. (2 - 3) and recalling that for the simple equity investment $F(x) = x$ we compute

$$R = \frac{d \ln f}{d \ln F} = \frac{d \ln(b/m)}{d \ln x} = \frac{m}{b} \left(\frac{b}{m} \right)'_x x. \quad (18)$$

Theoretically, this gives us the complete risk-aversion profile for the investor in question. Right now, however, we have a bare minimum of statistical information regarding b . So, as many other researchers before us have done, we choose to focus on the overall level of risk aversion and defer the very interesting topic of the shape of risk-aversion profiles to further research. As a measure of the overall risk aversion we consider the investor's own expectation of it

$$\langle R \rangle_b \stackrel{\text{def}}{=} \int R(x) b(x) dx. \quad (19)$$

Put together, the above two equations give

$$\langle R \rangle_b = \int m \left(\frac{b}{m} \right)'_x x dx = \int x m d \left(\frac{b}{m} \right). \quad (20)$$

Integrating by parts and noticing that $xb \Big|_0^\infty = 0$, we obtain

$$\langle R \rangle_b = - \int \frac{b}{m} d(xm) = - \int \frac{b}{m} (m dx + x dm) = -1 - \int b x \frac{dm}{m}. \quad (21)$$

Finally, using the notation of (19) we derive

$$\langle R \rangle_b = -1 - \langle x (\ln m)'_x \rangle_b. \quad (22)$$

This formula does not look very intuitive so, before using it, let us spend a few lines understanding it. To this end, let us see what it implies for a log-normal market-implied distribution. From Eq. (9) we derive

$$(\ln m)'_x \stackrel{\text{LN}}{=} \left(- \ln x - \frac{(\ln x - \mu)^2}{2\sigma^2} + \text{const} \right)'_x = -\frac{1}{x} - \frac{\ln x - \mu}{\sigma^2 x}. \quad (23)$$

Substitution into Eq. (22) gives

$$\langle R \rangle_b \stackrel{\text{LN}}{=} \frac{\langle \ln x \rangle_b - \mu}{\sigma^2} = \frac{1}{2} + \frac{\langle \ln x \rangle_b - r}{\sigma^2}. \quad (24)$$

Compare this to Eq. (12) which we studied above. We recognize Eq. (22) as a generalized analog of Eq. (12). The extent of generalization is very substantial: the market can have any implied distribution, and the investor can have an arbitrary profile of risk-aversion.

As discussed above, we now substitute $b(x) = P(x)$ into Eq. (22) and obtain the formula for the expected risk aversion for the equity investor who correctly expressed an accurate long-term view on the market

$$\langle R \rangle_P = -1 - \frac{1}{N} \sum_{i=1}^N x_i \left(\ln m(x_i) \right)'_{x_i}. \quad (25)$$

We are now in a position to compute $\langle R \rangle_P$ as of any day for which we have market information, m . We should remember, however, that our investor took a 15-year view and is completely ignoring all intermediate updates from the markets. The level of risk aversion for such an investor should be measured in a way that represents most of the actual investment period and is not sensitive to daily market fluctuations. Below we report two kinds of experiments which achieve this. In the first kind we look at the averaged value of $\langle R \rangle_P$ across the entire 15-year investment period. In the second type we get a glimpse of the term structure of risk aversion by looking at a 10-year moving average.

Above we explained our choice to partition historical investments into bi-monthly reinvestment periods. This choice has a useful side effect. A single experiment would skip most of the available market data using only what it needs at bi-monthly intervals. The skipped market data can be used to repeat the experiment (42 times in total) – we just need to start the bi-monthly sequence on a different business day within the first two months for which we have data.

The horizontal green lines on Fig. 4 report the levels of $\langle R \rangle_P$ averaged across the entire (~ 15 -year) investment period. Different lines correspond to the 42 different runs of the experiment. The red line on Fig. 4 is a bi-monthly report of the 10-year moving average of $\langle R \rangle_P$ for the investment which started on the 17th of May 2000 – the first day for which we have market data. The 42 runs of this experiment are plotted by faint hashed lines across the same graph.

As in the case of the expected equity premia considered in the previous section, we see completely normal levels of risk aversion. Even our attempt to glimpse the term structure, which misaligned investment horizon with the measurement of risk aversion, returned reasonable values.

Speaking about historical premia, we must mention that the performance of equities over the last 15 years has been rather patchy. This has reduced the magnitude of the relevant historical equity premia.² However, the reduction was not strong or persistent enough to remove large equity premia across the entire data set used in this paper. Out of the 42

²This might be partially responsible for the slight dip of risk aversion below zero on Fig. 4, although the confidently positive values for the averages (represented by the green lines) indicate that this is probably just noise.

investments represented by the green lines on Fig. 4, the worst and the best-performing ones delivered around 2% and 2.6% per annum in terms of the annualized equity premium. All of these values are well above the threshold of 0.35% reported by Mehra and Prescott [6].




Figure 4: Historical risk aversion. 10-year moving averages are computed on the bi-monthly grid as described in the main text. Within the 15-years of history this produces sequences of 30 (or 29) values (depending on the availability of data for the last period).

As a final remark, we would like to point the reader back to the discussion around Eqs. (22-24) which brings together the separate investigations of the expected and the realized premia. The two types of premia are different in terms of their precise interpretations. They also come with their own inherent challenges such as high levels of statistical noise in the case of realized premia. Yet, whether we talk about expected or realized equity premia, it is important to note that the underlying mathematics addressing the equity premium puzzle is basically identical.

5 Epilogue

Quantitative Structuring successfully survives the challenge from the Equity Premium Puzzle. In fact, it shows how the puzzle can be resolved. Indeed, given realistic values of risk aversion, Quantitative Structuring predicts the correct expected premia and shows how such expectations materialize over long time horizons. We expect the equity premia

to stay at the levels given by our formulae (Eq. (12), or more generally, Eq. (13)) for as long as investing in equities makes rational sense.

Our analysis is highly generalizable. In this paper we focused on equity investments, which happened to have a linear payoff function $F(x) = x$, but just as easily we could have examined any other investment strategy with a very different payoff function.

This is interesting because economic environments emerge from the successes and failures of individual strategies. It is not unreasonable to think that we might understand an economy by understanding the performance of its key strategies. Due to the potential importance of this line of thinking, let us conclude this paper with a few paragraphs articulating what our approach can offer to the wider subject of economics.

Detailed economics

Investments thrive on information. The information content of an investment is compressed into its economic structure – the payoff function. In the field of economics it has been a popular custom to replace the detailed payoff structure of an investment by simpler ad-hoc representations such as a point on a mean-variance diagram. The resulting loss of information is hard to quantify and even harder to compensate for, even with the most reasonable of assumptions.

Ideally, economic theories should mirror the reality and consider investors as individuals: each one with their own views and goals. Every attempt to get closer to this ideal inevitably faces the formidable challenge of practicality. More detailed models need more detailed information. Quantitative Structuring fulfills this need by providing access to the deep information content of payoff functions.

This is how we escaped the Equity Premium Puzzle. We consider investors as individuals which are allowed to hold any views they want. At the same time we leave no room for speculation about what these views actually are. It is crucial that the views are not assumed, they are derived using the knowledge of payoff functions (see Eqs. (7) and (28)).

Equity investors express strong directional views. Investment premia of over 6% per annum are not unusual in such circumstances. Similar premia can be seen in much more subtle investment strategies [5]. The expected premia are achieved in the long term, provided, of course, that the views are correct.

6 Appendix

Equation (3) can be rewritten as

$$d \ln f = R d \ln F. \tag{26}$$

For the case of constant but otherwise arbitrary R the above equation is immediately integrated to obtain

$$f(x) \propto e^{R \ln F(x)} = F^R(x). \tag{27}$$

This result together with Eq. (2) give us the investor-believed distribution

$$\begin{aligned} b(x) &= f(x) m(x) \\ &= \frac{e^{R \ln F(x)} m(x)}{\int e^{R \ln F(y)} m(y) dy}, \end{aligned} \quad (28)$$

where we used the fact that $b(x)$ is normalized. For the expected logarithmic return we compute

$$\text{ER}_R = \int b(x) \ln F(x) dx \quad (29)$$

$$= \frac{1}{Z} \frac{\partial Z}{\partial R}, \quad (30)$$

where

$$Z = \int F^R(x) m(x) dx. \quad (31)$$

In this paper we focus on the straightforward equity investment. In this case $F(x) = x$, and Z becomes essentially the R th moment of m . In the special case of log-normal market-implied distribution, this can be computed analytically (see Eq. (9) for notation)

$$Z = \int x^R m(x) dx = \text{DF} \cdot \exp \left\{ R\mu + \frac{1}{2} R^2 \sigma^2 \right\}, \quad (32)$$

and therefore

$$\text{ER}_R \rightarrow \text{ER}_R^{\text{LN}} = \mu + R\sigma^2. \quad (33)$$

References

- [1] Soklakov, A., “Why quantitative structuring?”, July (2015). arXiv:1507.07219.
- [2] Soklakov, A., “Elasticity theory of structuring”, April (2013). arXiv:1304.7535.
- [3] Soklakov, A., “Bayesian lessons for payout structuring”, RISK, Sept. (2011), 115-119. arXiv:1106.2882.
- [4] Soklakov, A., “Deriving Derivatives”, April (2013). arXiv:1304.7533.
- [5] Soklakov, A., “Model Risk Analysis via Investment Structuring”, July (2015). arXiv:1507.07216.
- [6] Mehra, R., and E. Prescott, “The equity premium: a puzzle”, Journal of Monetary Economics **15**, 145-161 (1985).
- [7] Mehra, R., and E. Prescott, “The equity premium in retrospect”, Handbook of the Economics of Finance, 887-936 (2003).
- [8] Damodaran, A., “Equity Risk Premiums (ERP): Determinants, Estimation and Implications – The 2014 Edition” ssrn-id2409198, updated annually since 2008. The numerical data used in this presentation are available at <http://people.stern.nyu.edu/adamodar/>