Why is the slope a good predictor of excess returns?

An overreaction explanation

Samuel Watts
sam@samueldwatts.com

A thesis submitted for the degree of
MSc in Mathematical Finance

April 2017

Abstract

The slope of the yield curve is an important factor in predicting excess bond returns. Systematically borrowing at the short rate and investing long should not make a profit in the long run, according to the Strong Expectations Hypothesis. In reality, however, the unconditional long duration strategy in US Treasuries has yielded profits over the last few decades.

The classical explanation is that investors demand higher yields in compensation for the risk that their estimates of the future short rate are incorrect. However, we present statistical evidence that suggests that the excess returns are unlikely to be fully ascribable to 'simple' risk premia, as the observed sign of the excess returns does not correlate with the sign of the CAPM risk premium over the last four decades.

The yield curve has sloped upwards for the majority of recent history and this has been taken as prima facie evidence of risk premia. By running representative simulations of the yield curve we find, however, that in one-in-five cases we obtain an even steeper slope than the observed average 50-year slope just by chance alone. This demonstrates that it is possible for the average yield curve to be as steep as it has been without invoking any risk premia. If excess returns are just a 'quirk of fate' however, this does not explain why the slope is a good predictor.

If excess returns are not directly and completely due to risk premia, and since it is difficult to believe there are large institutional frictions in the US Treasury market, we turn to behavioural finance explanations instead. Following in the footsteps of well-established research, in particular Shiller (1980), we propose a simple investor overreaction explanation.

Our simulations of simple, but representative, models with risk-neutral investors overreacting in different ways to the central bank's interest rate policy all produce similar regression results to those observed in the real data. We demonstrate that small amounts of investor overreaction are sufficient to produce excess returns and cause the slope to become predictive.

We also contribute to the ongoing debate over the return predicting factor proposed by Cochrane and Piazzesi (2005). From our simple model we recover many of the different patterns for the return predicting factor found in the literature and show that minor changes in model parameters radically change the shape of the factor.

Though excess returns could be explained by chance and risk premia, and these are contributing factors, we propose that overreaction can convincingly explain why the slope is a good predictor of excess returns.

Contents

1 Introduction

The existence of excess returns from a long duration strategy in US Treasuries is a well-known fact. The nature, the cause and, most importantly for this dissertation, the predictability of these returns still remain open points of discussion, despite excess returns being a well-known phenomenon for many years.

The long duration strategy of investing in a T-maturity bond at time t, funded by borrowing at the one-year spot rate, and then unwinding the portfolio after one year yields linearised excess returns of[1]:

$$\chi^T_{t \to t+1} := p^{(T-1)}_{t+1} - p^{(T)}_t - y^{(1)}_t,$$

where $p^{(T)}_t$ denotes the log bond price and $y^{(1)}_t$ denotes the yield, of maturity T and 1 respectively. This is the testable quantity we will investigate in this dissertation.

One might naively think that simply investing at a higher rate than the current funding rate would guarantee excess returns. However, if "forwards come true", i.e. today's forward rate with expiry t is an unbiased estimate of the future spot rate r_t, then a simple calculation shows that the average profits from systematically following the long duration strategy are zero. This holds *no matter how steep the yield curve is.* The Strong Expectations Hypothesis states that this is the case.

Table 1 displays the Sharpe ratio from following the long duration strategy over different periods for various Treasury maturities. We can see that the unconditional Sharpe ratio over the whole period (top row) is positive, demonstrating that excess returns have been observed, contradicting the Strong Expectations Hypothesis.

This can also be clearly seen in Figure 2, which displays the monthly excess returns from the long duration strategy of investing in the 10-year Treasury bond and unwinding the portfolio after one year[2]. The average excess return from 1971-2014 was 2.85%. If an investor had followed this strategy from 1971, reinvesting after each year, they would have more than doubled their money by 2014.

[1] We will derive this expression in Chapter 2

[2] Our bond data is from "ACM" Treasury data at the New York Fed [1]

Period	**2-year**	**5-year**	**10-year**
Full Sample	0.20	0.20	0.16
1955 - 1986	0.04	-0.01	-0.07
1987 - 2014	0.59	0.56	0.49
Recession	0.82	0.72	0.59
Expansion	0.01	0.06	0.05
1st half Expansion	0.52	0.50	0.45
2nd half Expansion	-0.61	-0.50	-0.48
Tightening Cycles			
1979:Q3 - 1981:Q2	-1.06	-1.13	-1.23
1993:Q3 - 1995:Q1	-0.79	-0.86	-0.86
2004:Q2 - 2006:Q2	-1.52	-0.90	-0.50

Table 1: Sharpe ratios for the long duration strategy applied to US Treasuries during the period of 1955-2014, subdivided i) into different chronological sub-periods, ii) into periods of recessions or expansions, and iii) during tightening cycles. Data taken from work by Vasant Naik and Mukundan Devarajan at PIMCO and reproduced in Rebonato[14].

Therefore since we observe the existence of excess returns there must be a reason why investors price forward rates higher on average than the resulting spot rate turns out to be. This is normally referred to as the term premium.

If, in an affine setting, the market price(s) of risk depend on a set of state variables that describe the yield curve then two questions arise. Firstly, for which state variables is the market price of risk non-zero? Then in turn, what is the dependence of the market price of risk on these 'compensatory' state variables?

It has been known since the 1990s that investors in Treasuries are compensated for the uncertainty in the level of the yield curve (i.e. the first principal component), but the magnitude of this compensation is heavily dependent on the slope of the yield curve (the second principal component)[14].

From 1971 to 2014 there has been positive correlation between the slope (proxied by the difference between the 1- and 10-year yield) and excess returns at all maturities. The Pearson correlation increases with maturity, ranging from 0.1788 for excess returns of the 2-year bond, to 0.4143 for the 10-year bond, by our calculations. In other words, the slope is a good ex-ante predictor of excess returns. The crux of this dissertation is concerned with proposing a simple explanation for why this may be the case.

Cochrane and Piazzesi[7] have proposed that more complex return predicting factors may have ex-ante explanatory power for excess returns. Their results suggest a single linear "tent-shaped" combination of forward rates predicts excess returns of *all* maturities. This is consequential since it implies that including higher principal components of the yield curve than the usual three (level, slope and curvature) significantly increases the forecasting power, and captures something vital and distinct. Their results have been corroborated by Hellerstein[11] and Adrian et al.[2].

Contentions have been raised against these conclusions however, (see [8], [15], [17] and [3]) in particular that the excess return predictions produced by the factor are highly correlated with predictions from the slope. Rebonato[14], Villegas[17], and Bauer and Hamilton[3] claim that only the slope is robust. This dissertation contributes to this ongoing debate.

Regardless of the extra forecasting power of the Cochrane and Piazzesi return predicting factor, the fact remains that the slope accounts for a significant proportion of excess returns. Therefore the question of why the slope is a strong ex-ante predictor of excess returns remains outstanding.

This dissertation is structured as follows. Before diving into the question above, we first expound the theoretical background and state of the existing literature in Chapter 2.

We receive our first hint at the cause of the slope's predictive power by observing in Table 1 that the Sharpe ratios are strongly positive during recessions (when the curve tends to slope steeply upwards), close to zero during expansions, and strongly negative when the curve is flat or even inverted (observed during the tightening cycles when "the punch bowl is taken away"). Therefore this data suggests a business-cycle explanation as both the slope of the yield curve and excess returns are highly correlated with recession and expansion.

The efficient market explanation for the continued existence of consequential excess returns, in what is possibly the most liquid market in the world, is that they are compensation for risk. This is arguably the most intuitive explanation. In Chapter 3 we will investigate this classical risk premium explanation by comparing excess returns with the observed risk premium, which following the CAPM story is proportional to the correlation between equity and Treasury returns.

In Chapter 4, we explore a different potential explanation, which is that the observed excess returns over the last 60 years were a quirk of fate and are not statistically significant; thus the Strong Expectations Hypothesis would not be rejected out of hand.

In this proposed explanation, the discrepancy between future rates and the eventual spot rate was not due to poor or biased predicting on the part of investors but a once-in-a-lifetime secular event. The average yield curve over this period has been upwards sloping despite rates falling but, according to this hypothesis, investors who were long duration were not compensated for risk but just 'got lucky'.

As we expound in these chapters, we have good reasons to be sceptical of these two explanations. Instead in Chapter 5 we offer an alternative explanation of why the slope is a good predictor of excess returns: that investors overreact to changes in central bank policy.

This explanation follows a well-trodden path of established research into overreaction, in securities markets in general, such as Shiller[16], and bonds in particular. We contribute to this literature by asking whether a simple overreaction mechanism could account for the observed persistence of excess returns, the predictiveness of the slope, and also the observed shape of the return predicting factor.

We start with a simple overreaction model to confirm our intuition before developing it to make it more representative of the real-world economy. Simulations are run using a comprehensive set of initial parameters to fully explore the behaviour of the model and validate the ensuing conclusions.

Finally, we conclude in Chapter 6.

2 Theoretical background and existing literature

The following exposition is based on Rebonato[14]; in particular chapters 15, 23, 24 and 26. We start by explaining what explicitly the long duration strategy is, and what excess returns are.

2.1 Long duration strategy and excess returns

At time $t = 0$ we start with nothing. We wish to invest in a T-maturity bond so buy $1/P_0^{(T)}$ units of the bond $P_0^{(T)}$ costing one dollar[3]. To generate this missing dollar we borrow at the current one-year spot rate. This is equivalent to selling $1/P_0^{(1)}$ units of the 1-year bond $P_0^{(1)}$. Therefore our zero-cost initial portfolio is:

$$\Pi_0 = \left(\frac{1}{P_0^{(T)}}\right) P_0^{(T)} - \left(\frac{1}{P_0^{(1)}}\right) P_0^{(1)} = 0 \tag{1}$$

After one year we unwind the portfolio. The T-maturity bond is now worth $P_1^{(T-1)}$ and the 1-year bond is worth \$1. So our portfolio after one year is worth

$$\Pi_1 = \left(\frac{1}{P_0^{(T)}}\right) P_1^{(T-1)} - \left(\frac{1}{P_0^{(1)}}\right) \cdot 1 \tag{2}$$

We define the $t+1$ excess return of the 1-year, long duration strategy starting at t to be $\Pi_{t+1} - \Pi_t$. Since the initial portfolio is always zero cost, our general expression for the excess returns is as follows:

$$\chi^T_{t\to t+1} := \left(\frac{1}{P_t^{(T)}}\right) P_{t+1}^{(T-1)} - \left(\frac{1}{P_t^{(1)}}\right) \tag{3}$$

This expression can be simplified by approximating the bond price to the first order of the Taylor series:

$$P_t^{(T)} = e^{-y_t^{(T)}T} \simeq 1 - y_t^{(T)}T + ... \tag{4}$$

where $y_t^{(T)}$ is the yield at time t of the T maturity bond. This gives us

$$\chi^T_{t\to t+1} \simeq (1 + y_t^{(T)}T)(1 - y_{t+1}^{(T-1)}[T-1]) - (1 + y_t^{(1)}) \tag{5}$$

$$= y_t^{(T)}T - y_{t+1}^{(T-1)}[T-1] - y_t^{(1)} \tag{6}$$

If we define the log bond price as

$$p_t^{(T)} := \log P_t^{(T)}, \tag{7}$$

[3]We will use parentheses to make it clear that the superscript denotes maturity not exponent.

we get the expression most commonly used in the literature to express excess returns:

$$\chi^{T}_{t\to t+1} \simeq p^{(T-1)}_{t+1} - p^{(T)}_{t} - y^{(1)}_{t} \tag{8}$$

We use this approximation of the excess returns rather than the exact returns as the exact expression includes a contribution from convexity. By linearising the returns we can remove this convexity contribution and focus on the source of the remaining excess returns. In the actual data the difference is immaterial for maturities up to 10 years[14], which is a longer period than analysed in this dissertation, thus we will ignore convexity[4].

It can also be informative to express excess returns in terms of forward rates. Prices, yields and forward rates are all linear functions of each other so we can express excess returns as a linear function of maturity-spanning forward rates:

$$\chi^{T}_{t\to t+1} = \beta^{(T)}_{0} + \beta^{(T)}_{1} y^{(1)}_{t} + \beta^{(T)}_{2} f^{(2)}_{t} + ... + \beta^{(T)}_{i} f^{(i)}_{t}, \tag{9}$$

where $\beta^{(T)}_{j}$, $j = 0, ..., i$, are the parameters for maturity T, and i depends on the furthest available forward rate. Cochrane and Piazzesi[7], whose work we will discuss throughout this dissertation, used data of one- through five-year zero coupon bond prices, i.e. $i = 5$. To aid comparison we take the same approach and use these five forward rates.

2.2 Relevance of excess returns

Why are we looking at excess returns? Because doing so gives us direct information about risk premia. This can be seen as follows.

If the economy, and therefore the yield curve, is driven by n factors x_i, $i = 1, ..., n$ then the expected return from holding the T-maturity bond is

$$\mathbb{E}^{\mathbb{P}}[ret^{(T)}_{t}] = \mathbb{E}^{\mathbb{P}}\left[\frac{dP^{(T)}_{t}}{P^{(T)}_{t}}\right] = \left[r_t + \sum_{i=1}^{n} \mathcal{D}_i(T,t)\lambda_i\sigma_i\right], \tag{10}$$

where r_t is the short rate and λ_i and σ_i are respectively the market price of risk and volatility of the ith factor. $\mathcal{D}_i(T,t)$ is the duration of the bond associated with the ith term:

$$\mathcal{D}_i(T,t) := \frac{1}{P^{(T)}_{t}} \frac{\partial P^{(T)}_{t}}{\partial x_i} \tag{11}$$

[4]We have used $\simeq$ in equation (8) to make this difference clear but since we ignore convexity throughout, for the remainder of the dissertation we will use $=$ to lighten notation.

Thus the log price of the bond, neglecting convexity as before, can be expressed as

$$\log P_t^{(T)} = p_t^{(T)} = -\mathbb{E}^{\mathbb{P}}\left[\int_t^T \left(r_s + \sum_{i=1}^{n} \mathcal{D}_i(T-s,s)\lambda_i(s)\sigma_i(s)\right) ds\right] \tag{12}$$

Inserting equation (12) into (8) and using the tower law, $\mathbb{E}_t[\mathbb{E}_{t+1}[x]] = \mathbb{E}_t[x]$, we get

$$\mathbb{E}_t^{\mathbb{P}}[\chi_{t\to t+1}^T] = \mathbb{E}_t^{\mathbb{P}}\left[\int_t^{t+1} \left(r_s + \sum_{i=1}^{n} \mathcal{D}_i(T-s,s)\lambda_i(s)\sigma_i(s)\right) ds\right] - y_t^{(1)} \tag{13}$$

If we now discretise the integral, i.e. $\int_t^T f(s)ds \simeq \sum f_i \Delta t$ with $\Delta t = 1$ year, and note that then $\int_t^{t+1} r(s)ds \simeq r(t)\Delta t = y_t^{(1)}$, we obtain:

$$\mathbb{E}_t^{\mathbb{P}}[\chi_{t\to t+1}^T] = \sum_{i=1}^{n} \mathcal{D}_i(T-t,t)\lambda_i(t)\sigma_i(t) \tag{14}$$

That is, the expected excess return is equal to the sum over factors of the products of the market price of risk, times the volatility of the factor, times the duration of the T-maturity bond at the start of the investment period. In other words, there is a direct relationship between risk (or term) premia and the expected excess returns.

So to test the Strong Expectations Hypothesis we can investigate whether the realised expected excess returns are statistically different from zero by utilising equation (8):

$$\mathbb{E}_t^{\mathbb{P}}[\chi_{t\to t+1}^T] \equiv \mathbb{E}_t^{\mathbb{P}}[p_{t+1}^{(T-1)} - p_t^{(T)} - y_t^{(1)}] \stackrel{?}{=} 0 \tag{15}$$

Thus from above, with this test we are also simultaneously testing whether 'forwards come true' and whether term premia exist, both equivalent statements of the Strong Expectations Hypothesis.

As discussed in the introduction, Table 1 and Figure 2 suggest that the expected excess returns are non-zero, contradicting the Strong Expectations Hypothesis. However there is a corresponding 'weak' version of the hypothesis that states that: the expected excess returns from systematically following the long duration strategy are constant or, equivalently, forward term premia are constant. In other words, investors are compensated for systematically going long duration but this compensation is constant and there are no particularly better or worse times to invest in the strategy.

Equation (14) is useful here because we can use it to empirically test the Weak Expectations Hypothesis. We can take real bond data and compare the ex-ante expectations of the excess return using this equation with the actual ex-post profits from performing the long duration strategy. More explicitly, we are looking at the difference, $\Delta\chi^T_{t\to t+1}$, between the expected excess return from equation (14) and the realised excess return from equation (8):

$$\Delta\chi^T_{t\to t+1} := \sum_{i=1}^{n} \mathcal{D}_i(T-t,t)\lambda_i(t)\sigma_i(t) - \left(p^{(T-1)}_{t+1} - p^{(T)}_t - y^{(1)}_t\right) \tag{16}$$

If the Weak Expectations Hypothesis holds then the expectation of this difference should be zero:

$$\mathbb{E}^{\mathbb{P}}_t[\Delta\chi^T_{t\to t+1}] = \mathbb{E}^{\mathbb{P}}_t\left[\sum_{i=1}^{n} \mathcal{D}_i(T-t,t)\lambda_i(t)\sigma_i(t) - \left(p^{(T-1)}_{t+1} - p^{(T)}_t - y^{(1)}_t\right)\right] = 0 \tag{17}$$

From equation (14) this seems trivially true as

$$\mathbb{E}^{\mathbb{P}}_t[\mathbb{E}^{\mathbb{P}}_t[\chi^T_{t\to t+1}]] = \mathbb{E}^{\mathbb{P}}_t[\chi^T_{t\to t+1}] \tag{18}$$

However, we relied on some assumptions to reach equation (17). These are namely: that the market price of risk has an affine dependence on the state variables and that market expectations are unbiased. Therefore if we find that empirically (17) does not hold we can conclude that one or both of the above assumptions do not hold.

The former of these could be false due to faulty model specification: either inaccurate parameter calibration, e.g. wrong market price of risk, or the selection of an affine model was a specious choice. This would mean that equation (17) does not hold due to

$$\mathbb{E}^{\mathbb{P}}_t[\chi^T_{t\to t+1}] \neq \sum_{i=1}^{n} \mathcal{D}_i(T-t,t)\lambda_i(t)\sigma_i(t) \tag{19}$$

The second assumption is equivalent to saying that, instead of investors' expectations of the excess returns being formed using the real-world measure $\mathbb{P}$, they rather use a subjective measure which we denote $\mathbb{S}$:

$$\mathbb{E}^{\mathbb{S}}_t[\chi^T_{t\to t+1}] = \sum_{i=1}^{n} \mathcal{D}_i(T-t,t)\lambda_i(t)\sigma_i(t) \tag{20}$$

This way equation (17) would not hold as the tower law no longer applies:

$$\mathbb{E}^{\mathbb{P}}_t[\mathbb{E}^{\mathbb{S}}_t[\chi^T_{t\to t+1}]] \neq \mathbb{E}^{\mathbb{P}}_t[\chi^T_{t\to t+1}] \tag{21}$$

There are multiple reasons why investors may form their expectations using a 'wrong' measure. It may be that the historical information, on which they based their decisions, is not representative of the future. Fifty years ago investors did not have the 'lakes' of data we now swim in, or even simply data covering as long a period of time as now, so their estimates could have been statistically poor. This is the explanation offered by Piazzesi, Salomao and Schneider[13].

Investors could also use a subjective measure as their expectations are biased. In other words they are 'irrational' and the market is 'inefficient'. From this it would entail that excess returns are not purely due to risk premia but at least partially, if not wholly, due to repeatedly inaccurate future rate expectations. Whether this is part of the story is an ongoing debate. One of the seminal papers is Shiller[16], where he argues that investors overreact to new information when pricing equities. We will investigate this explanation but for bonds rather than equities in Chapter 5.

2.3 Predicting excess returns

If the expected difference between 'expected' excess returns and realised excess returns is non-zero then this is another way of saying that the excess returns are predictable to some degree.

Before the work of Cochrane and Piazzesi[7] the accepted wisdom was that the best yield-curve based predictor of excess returns was the slope of the yield curve[10][4].

Cochrane and Piazzesi[7] proposed a single factor model:

$$\chi^T_{t\to t+1} = b_T(\gamma_0 + \gamma_1 y^{(1)}_t + \gamma_2 f^{(2)}_t + ... + \gamma_5 f^{(5)}_t) + \epsilon^{(T)}_{t+1} \tag{22}$$

In order to make b_T and γ_n unique, b_T is normalised so that the average value of b_T across maturities is one, i.e. $\sum_{T=2}^{5} b_T/4 = 1$.

A single linear combination of forward rates therefore is the state variable for time-varying expected excess returns of *all* maturities. This contrasts with previous studies (Fama and Bliss[10], Campbell and Shiller[4]) that regress each bond's expected excess returns against the corresponding forward or yield spread. Cochrane and Piazzesi's[7] regressions using the single factor had much smaller p-values and more than double the forecast R^2 of these previous studies. Thus they argue it provides even stronger evidence against the Strong Expectations Hypothesis.

Their five component factor also contrasts with the received wisdom that three components are sufficient to capture yield curve dynamics[12][14]. Whilst by definition more components must be at least as accurate, Cochrane and Piazzesi argue that their factor captures something significant and vital not captured by the traditional three principal components.

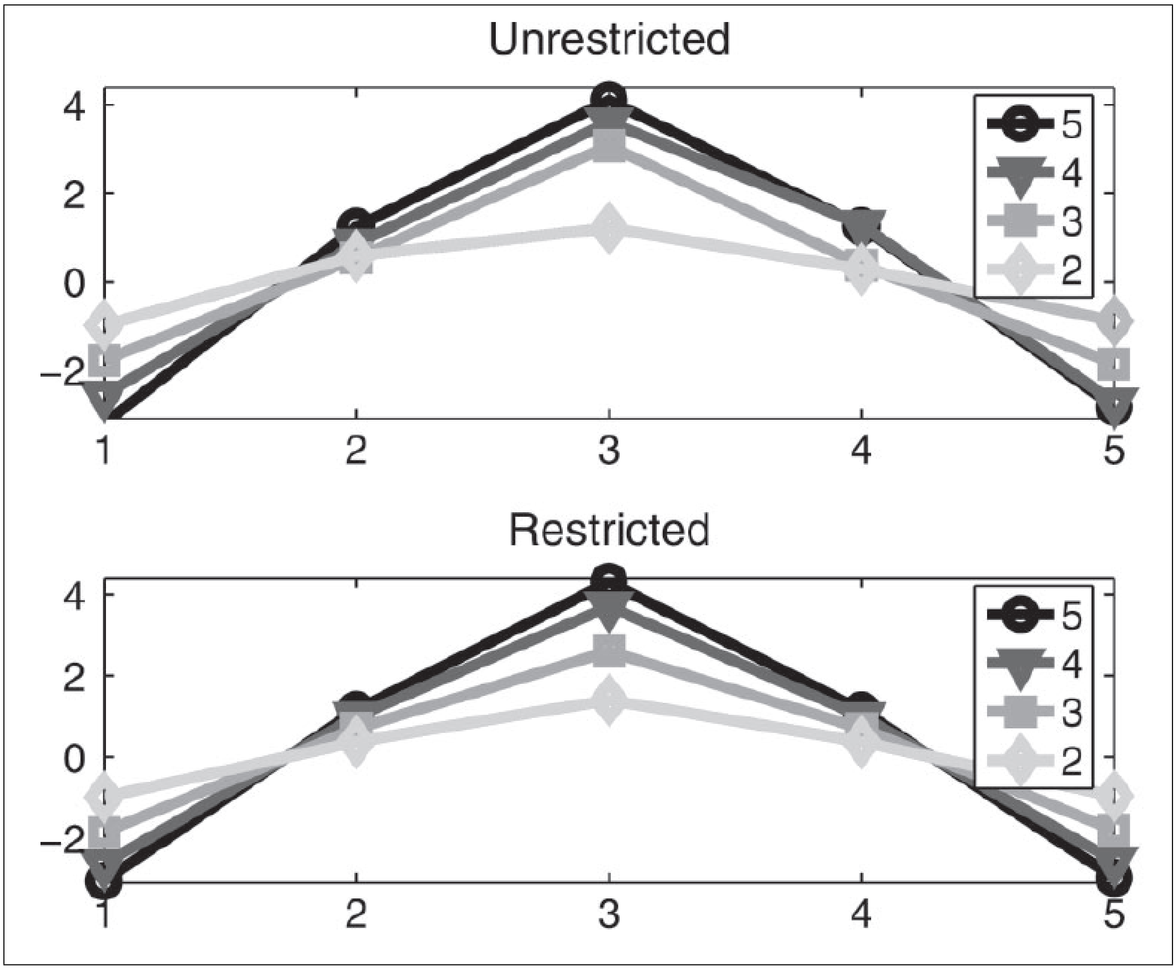

Figure 1: Figure reproduced from Cochrane and Piazzesi[7]. The top panel presents their estimates β from the unrestricted regressions (equation (9)) of bond excess returns on all forward rates. The bottom panel presents their restricted estimates $b\gamma$ from the single-factor model (equation (22)). The legend (5, 4, 3, 2) gives the maturity of the bond whose excess return is forecast. The x-axis gives the maturity of the forward rate on the right-hand side of the regression.

Since they claim it is "the *same* linear combination of forward rates that predicts bond returns at all maturities"[7] the shape of the factor and its robustness is of interest. Figure 1 shows their loadings $\beta_i^{(T)}$ from regressing the monthly overlapping annual excess returns from maturities of 2 to 5 years against the five forward rates at the start of the investment period[5]. This matches equation (9) from earlier. The lower panel shows the estimates for $b_T\gamma_i$ from the restricted single factor model in equation (22).

In both cases the loadings form a so-called "tent". They point out that restricting the coefficients preserves the shape and "does little damage" to the economic and statistical significance of the predictive regression.

Doubts have been raised about the robustness of this tent shape by Dai, Singleton and Yang[8]. They argue that the shape of the single factor is highly sensitive to the choice of spline used to construct the zero-coupon bond yields from data. Rather than a "tent", Dai et al. found the single factor formed a "bat" shape, which to me looks like a squashed letter 'M'[6].

Dai et al. consider their results of concern as "even small measurement errors, in the presence of highly correlated forward rates, can lead to large differences in fitted project coefficients".

Rebonato[15] discusses the two patterns, "tents" and "bats". He shows that the patterns produce economically indistinguishable predictions of excess returns. He then argues that a "tent"-shaped factor arises naturally if the yield curve slope is the most important explanatory factor for excess returns. Whereas on the other hand a "bat" shape cannot be easily produced by transformation from a slope factor in yields to loadings for forward rates.

The topic is expounded further by Villegas[17]. She adds a small amount of uninformative noise to the original data and shows that this can transform the pattern from a "bat" to a "distorted tent-shape", as she calls it. Her results suggest that the slope "is the most robust regressor for excess returns modelling as it remains robust on the prediction of the overall trend of the data, even when tested on extended out-of-sample periods".

Bauer and Hamilton[3] raised further contentions against this recent consensus (that additional variables to the first three principal components help predict excess

[5]The first forward rate is also the spot rate but following the literature we shall refer to all of the rates collectively as the 'forward rates'.

[6]Much like star constellations, these shape names are subjective and require some imagination plus some squinting.

bond returns). They claim that the statistical tests used by Cochrane and Piazzesi have been vitiated by "very large size distortions from a previously unrecognised problem arising from highly persistent regressors and correlation between the true predictors and lags of the dependent variable". Using tests they claim are more robust, they re-examine the data and conclude that only yield-curve variables (particularly the yield curve level and slope) can robustly predict excess returns.

Cochrane[6] has produced a rebuttal to Bauer and Hamilton. His response mainly revolves around econometric issues that are beyond the scope of this current dissertation. He does point out, however, that though the economic variables' growth and inflation "on their own do absolutely nothing to forecast returns", when regressed alongside the three principal components they become significant, as well as increasing the significance of the first two principal components.

Additionally he adds a linear trend amongst the regressors and finds that the R^2 value rises dramatically to 64%. Then he removes the third principal component as well as growth and inflation variables, leaving in the trend regressor, and finds the R^2 value still remains at 62%. Cochrane proposes that the reason for this is that "the inflation was just proxying (and poorly) for detrending of the level and slope factors" so business cycle patterns have a major effect on the data and must be handled carefully.

We contribute to this ongoing debate in Chapter 5.

3 Are the excess returns due to risk premia?

In Chapter 2 we drew a theoretical connection between excess returns and term premia. The classical explanation is that if the market is efficient these excess returns must be compensation for risk. We start from this viewpoint and infer the ensuing conclusions. The following exposition is based on Cochrane[5].

A positive risk premium is compensation for a negative payoff in bad states of the world and a positive payoff in good states of the world. Correspondingly, a negative risk premium is a cost for a good payoff in bad states of the world, i.e. hedging risk so we get a payoff when we really need it. This is why risky investments are cheaper and you have to pay for insurance.

Consider an investor who has an investment opportunity, say a stock S_t, and current wealth of W_t. The investor can consume at times t and $t+1$ so must decide how much of her wealth, a, to invest at time t to maximise her total utility, $u(C_t, C_{t+1})$, taking into consideration her coefficient of impatience β. This can be expressed by:

$$\max_{0 \le a \le 1} [u(C_t, C_{t+1})] = u(C_t) + \beta \mathbb{E}_t[u(C_{t+1})] \tag{23}$$

If each unit of stock S_t will pay x_{t+1} at time $t+1$ this gives us the expressions for the investor's consumption:

$$C_t = W_t - aS_t, \quad C_{t+1} = ax_{t+1} \tag{24}$$

To find the maximum we take the derivative with respect to a and equate it to zero:

$$\begin{aligned}\frac{\partial u(C_t, C_{t+1})}{\partial a} &= \frac{\partial u}{\partial C_t}\frac{\partial C_t}{\partial a} + \beta \mathbb{E}_t\left[\frac{\partial u}{\partial C_{t+1}}\frac{\partial C_{t+1}}{\partial a}\right] \\ &= u'(C_t)(-S_t) + \beta \mathbb{E}_t\left[u'(C_{t+1})x_{t+1}\right] = 0\end{aligned}$$

We can rearrange this to ascertain the fair price of S_t for the investor:

$$S_t = \beta \mathbb{E}_t\left[\frac{u'(C_{t+1})}{u'(C_t)}x_{t+1}\right] = \mathbb{E}_t[m_{t+1}x_{t+1}], \quad m_{t+1} := \beta\frac{u'(C_{t+1})}{u'(C_t)} \tag{25}$$

This derived quantity, m_{t+1}, is called the *stochastic discount factor* (SDF). Through equation (25) we capture the relationship between the asset price today on one side and, on the other, predictions of what it will pay in the future, the investor's impatience, and the shape of the investor's utility curve.

Applying the expectation of products rule we get

$$S_t = \mathbb{E}_t[m_{t+1}]\mathbb{E}_t[x_{t+1}] + \mathrm{Cov}_t[m_{t+1}, x_{t+1}] \tag{26}$$

We can calculate the value of the expectation of the SDF by considering the case of a one-period risk-free security. Investing \$1 at time t gives a risk-free gross return of $\$R_f$ at time $t+1$. Inserting this into equation (25) gives:

$$1 = \mathbb{E}_t[m_{t+1}R_f] = \mathbb{E}_t[m_{t+1}]R_f \tag{27}$$

$$\implies \mathbb{E}_t[m_{t+1}] = \frac{1}{R_f} \tag{28}$$

Putting this back into equation (26) gives for a generic asset:

$$S_t = \frac{1}{R_f}\mathbb{E}_t[x_{t+1}] + \mathrm{Cov}_t[m_{t+1}, x_{t+1}] \tag{29}$$

Looking at this equation we can see that the price of an asset can be split into two parts: the discounted expectation of the payoff, plus a term dependent on the covariance between the payoff and the stochastic discount factor. The latter part must be the risk premium since the first part is simply the price if the investor was completely risk-neutral.

This expression for the risk premium, $\mathrm{Cov}_t[m_{t+1}, x_{t+1}]$, matches our intuitive explanation earlier. Reasonable utility curves are concave, so the more we expect future utility to be larger than current utility the smaller the SDF is, as $u'(C_{t+1}) < u'(C_t)$. So for assets that have positive payoff when utility is high (i.e. good states of the world) we will get a negative covariance term in (29). This decreases the price of the asset S_t and increases the yield from holding the asset. Vice versa, if positive payoff is correlated with lower utility the covariance term will be positive increasing the price of the asset.

This risk premium expression is also very useful because we can use it to investigate risk premia by ascertaining the magnitude and sign of the covariance of m_{t+1} and x_{t+1}. In our case we select the asset return x_{t+1} as the 10-year bond excess returns. Rather than trying to exactly calculate the market's stochastic discount factor we use a proxy for the change in consumption, $u'(C_{t+1})/u'(C_t)$, which should be directly related to the SDF.

Our proxy is the inverse of the S&P 500 log returns[7]. Following the standard CAPM framework, we assume investors only consume from their investment returns and dividends. Taking a logarithmic utility function, $u(c) := \log(c)$, we get for the SDF:

$$m_{t+1} = \beta\frac{C_t}{C_{t+1}} \approx -\log\left(\frac{SP_{t+1}}{SP_t}\right), \tag{30}$$

[7]Data taken from the public, official S&P 500 data.

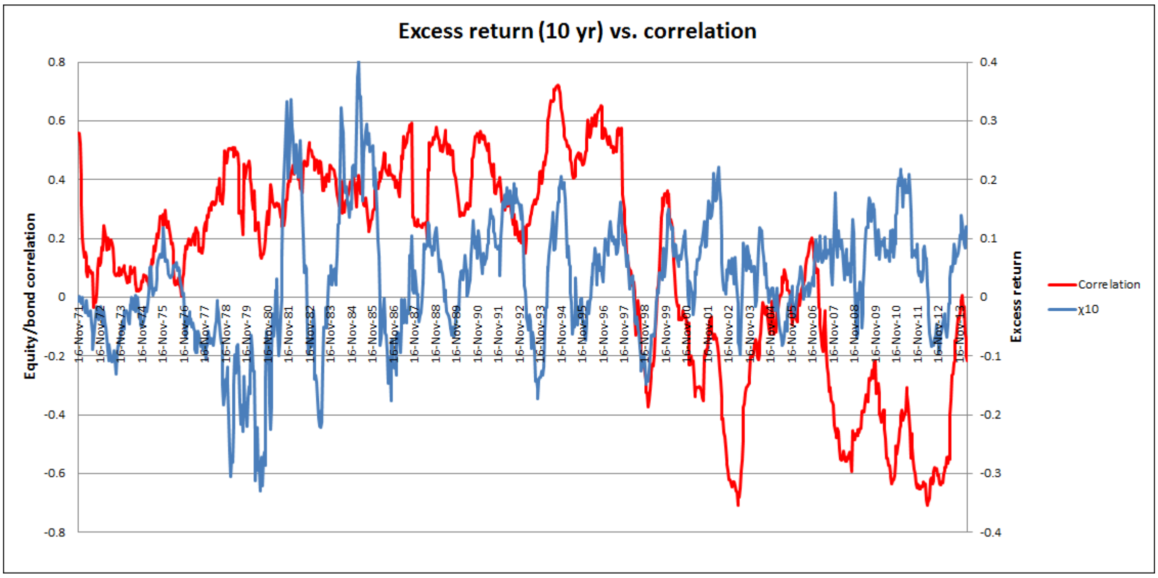

Figure 2: Correlation between returns from 10-year yield and the S&P 500, as well as the monthly overlapping annual excess returns from a long duration strategy from 1971-2013. Visually there is no strong consistent correlation between the two, in particular they appear positively correlated before the 1990s and negatively correlated afterwards.

where SP_t is the value of the S&P 500 at time t, taken as a representation of the investors' collective investments.

In the CAPM framework, bonds provide a hedge for stocks, so the covariance of bond returns and the above SDF proxies the CAPM risk premium. We want to test whether the claim that excess returns are due to risk premia is consistent with the asset pricing story. To do this we compare the excess returns with a rolling-window correlation of the 10-year bond returns and S&P returns, to check whether the sign of the excess returns matches that of our CAPM risk premium. Figure 2 displays the two series over the period from 1971 to 2014. Engle[9] gets similar results using the more complex dynamic conditional correlation (DCC), which stems from GARCH models.

If the excess returns are due to a risk premium then we would expect them to have the same sign and scale with the CAPM risk premium. Visually one can see that the signs of each are not consistent with each other and statistically the signs of the two only match 43% of the time. Before the 1990s the two appear positively correlated and then around 1992 they appear to diverge and afterwards become negatively correlated. The bond and S&P returns correlation becomes mainly negative at the turn of the millennium while the excess returns on the other hand remain mainly positive.

To check this quantitatively we calculate the covariance of the two from 1971 to 2014:

$$\mathrm{Cov}\left(xret^{(10)}_{t\to t+1}, \rho_t \left[-\log\left(\frac{SP_{t+1}}{SP_t}\right), y^{(10)}_{t+1} - y^{(10)}_t\right]\right), \tag{31}$$

where ρ_t here denotes the rolling correlation at time t.

We investigated different frequencies of the rolling correlation of bond and S&P returns and found the results detailed in Table 2. From these we can see that regardless of rolling correlation window length we do not find a significant dependence of the excess returns on the bond/equity correlation, and in fact the covariance is slightly negative.

Rolling correlation window	Covariance
6 months	-0.001958436
1 year	-0.002155091
18 months	-0.002687536
2 years	-0.002611631

Table 2: Covariance of 10-year excess returns and theoretical CAPM risk premium for a range of rolling correlation windows for the risk premium.

We have assumed here that consumption can be proxied by equity returns, which does not include other sources of consumption risk such as employment or inflation. However, even if not conclusive, this observation still suggests that the profits reaped from the long duration strategy are unlikely to be fully ascribable to 'simple' risk premia. By 'simple' we mean that it does not directly fulfil the role of providing a positive payoff in good states of the world. If they were, then we would expect that when the correlation changes sign the excess returns would change sign as well, but they do not. This leaves us with something that requires an explanation.

4 Are the realised excess returns a quirk of fate?

If the Strong Expectations Hypothesis was accurate then there would be no risk premia and risk-neutral expectations would equal the real-world expectations. This would mean the yield curve was the risk-neutral expectation of the short rate path and we would expect it to be downwards sloping as often as it was upwards sloping. In fact, the time average yield curve slope from 1961 to 2016 was 1.22%[8], i.e. upwards sloping. As can be seen in Figure 3, the slope has been positive for the majority of the last 55 years; 83% of the time to be more precise.

It has been argued that since it is implausible that investors should, over such a long period, consistently expect rates to 'go up' despite the evidence to the contrary, this is one of the strongest pieces of evidence that risk premia do exist. As we established in Chapter 2, if there are positive term premia this causes a more positive yield curve slope, as investors require increasing compensation for taking longer duration risk. How do we square this observation with the contention against the simple risk premium explanation in Chapter 3?

It is possible that the Strong Expectations Hypothesis holds but the half century or so of data we have is not a large enough data set. It could be that the observation of a mainly upwards slope is a so-called quirk of fate and that if time was replayed over and over again we would see average excess returns approaching zero.

In order to investigate this possibility, we ran representative simulations of 50 years of the yield curve. The results from this can tell us how likely the realised slope history was. If it is really uncommon to have such an upwards slope then this is an indication that risk premia, or some other phenomena, really are 'pushing up' longer-term yields.

To recreate the yield curve, its first three principal components were estimated and then simulated. Principal Component Analysis is a technique to reduce the dimensionality of multi-dimension problems and is an established method for term-structure modelling.

New York Fed economists Tobias Adrian, Richard Crump, and Emanuel Moench (or "ACM") present Treasury yields for maturities from one to ten years from 1961 to the present[1]. From this data the first three principal components were obtained by orthogonalising the covariance matrix of the yields.

[8]Again our bond data is from "ACM"[1]

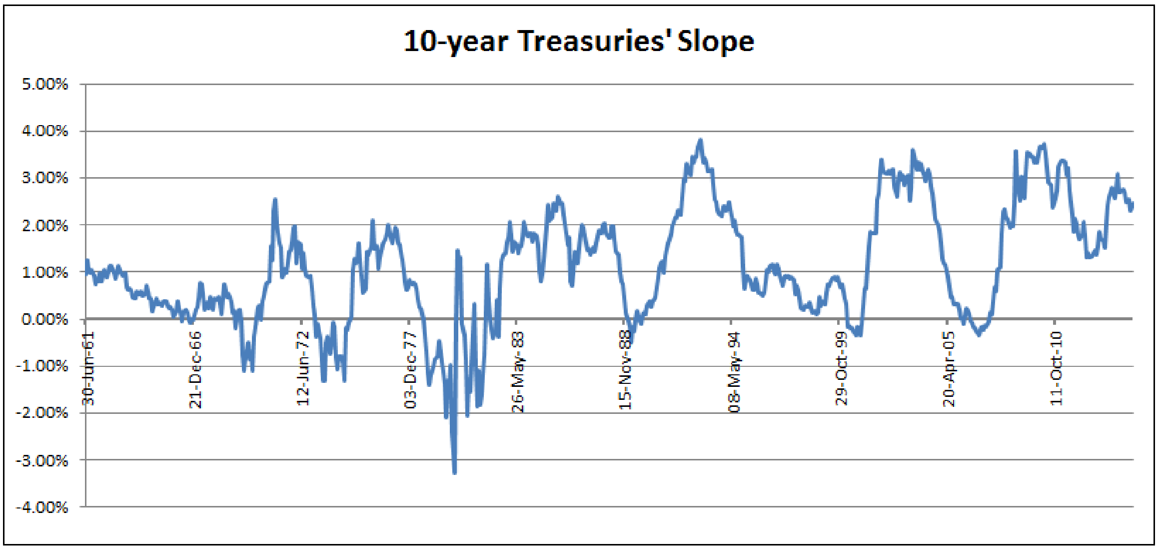

Figure 3: The yield curve slope, defined as the difference between the 10-year and 1-year yields, from 1961-2014. Note the slope is positive, i.e. the yield curve is upwards sloping, for the majority of the period.

These components are orthonormal, linear combinations of the yields:

$$x_i = \sum_{j=1,10} c_{ij} y^{(j)}, \quad i = 1, 2, 3 \tag{32}$$

They have been chosen so that each component explains the largest possible proportion of the overall variability. As such, they can be thought of as the three yield vectors that most capture the variation of the ten yields.

The conventional wisdom is that utilising three components is sufficient to capture 98+% of yield behaviour[12][14]. The first three components are also attractive as they have the intuitive interpretations of ‘level’, ‘slope’ and ‘curvature’ of the yield curve.

Once these components were estimated from the data, it was also possible to work out their mean reverting behaviour, assuming each yield $y_t^{(j)}$ follows an Ornstein-Uhlenbeck process:

$$dy_t^{(j)} = \kappa(\theta - y_t^{(j)})dt + \sigma dz_t \tag{33}$$

By regressing the principal components, their reversion speed and reversion level can be estimated. Let x_t denote a principal component at time t as in equation (32). Then we have the expected process of

$$x_{t+1} - x_t = \kappa\theta\Delta t - \kappa x_t \Delta t \tag{34}$$

$$\implies x_{t+1} = x_t(1 - \kappa\Delta t) + \kappa\theta\Delta t \tag{35}$$

Regressing x_{t+1} against x_t gives the slope β and intercept α:

$$\beta = 1 - \kappa\Delta t, \quad \alpha = \kappa\theta\Delta t.$$

Thus for a given Δt, in our case $\frac{1}{12}$ as ACM data is monthly, we can find the reversion speed:

$$\kappa = \frac{1-\beta}{\Delta t},$$

and then using this κ find the principal component's reversion level:

$$\theta = \frac{\alpha}{\kappa\Delta t}$$

This regression is performed for the first three principal components and the resultant empirical parameters then used to simulate the components, and hence the yield curve, in the $\mathbb{P}$ measure.

The simulations start from a flat zero yield curve, $\hat{y}_0 = \mathbf{0}$. Then we simulate each time step by iteratively calculating the Ornstein-Uhlenbeck process of the principal component:

$$\hat{y}_{t+1} = \hat{y}_t e^{-\hat{\kappa}\Delta t} + \hat{\theta}(\mathbf{I}_3 - e^{-\hat{\kappa}\Delta t}) + \hat{\sigma}\Delta t\mathbf{N}_3, \tag{36}$$

where $\hat{y}_t$ is the yield curve at time t, $\hat{\kappa}$ and $\hat{\theta}$ are respectively the reversion speed and level matrices derived as above, $\hat{\sigma}$ is the volatility matrix derived from the eigenvalue matrix of the principal component differences, and $\mathbf{N}_3$ is a vector of three random normally distributed numbers. Here $\hat{\kappa}$ is diagonal, since if we do a joint regression we get an almost diagonal $\hat{\kappa}$ matrix so this a justifiably close simplification.

1000 realisations of 50 years of the yield curve were simulated. Of these 19.4% had an average slope (defined as the 10-year yield minus the 1-year) steeper than the average realised slope of 1.22%. The standard deviation of the average slope from the simulations was 1.3 percentage points, which is considerable. Figure 4 displays the average yield curve, calculated as the average yield at each maturity, of the 1000 simulations as well as the average observed yield curve from 1961-2014.

Our simulations tell us that the yield curve is flat on average, contrasting with the 122 basis points slope observed in the actual yield curve history. However, we also found that there is approximately a one-in-five probability that the average 50-year slope would be as steep as observed.

Thus we can conclude that observing a consistently upwards-sloping yield curve is not conclusive evidence of the existence of term premia since our simulations reasonably often produce even steeper yield curve slopes purely by chance alone.

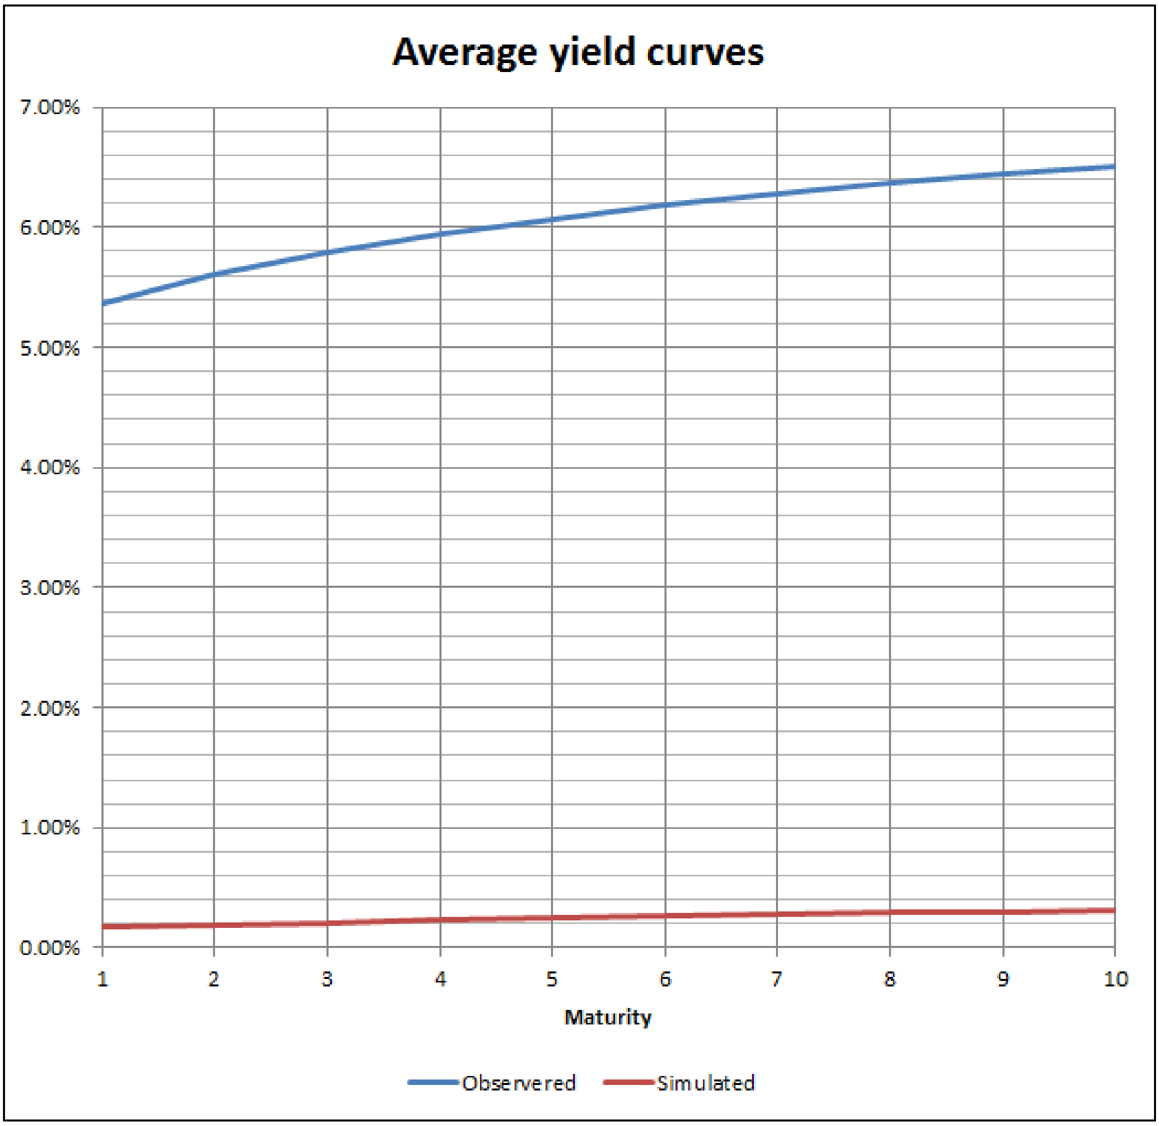

Figure 4: Average yield curve observed over 1961-2014 as well as of 1000 realisations of 50 years of yield curve simulations. The average curve of the simulations is nearly flat with a slope of less than 20 basis points, which contrasts considerably with the 122 basis points observed in the actual data.

So pure chance can explain the existence of excess returns and the upward-sloping average yield curve, without requiring any risk premia, but there remains a fundamental flaw: it does not explain why the slope is a good ex-ante predictor of excess returns. If the Strong Expectations Hypothesis holds and there is no phenomenon, such as risk premia, 'pushing up' the yield curve then we lose the direct relationship between the slope and excess returns. As our simulations in Chapter 5 show, randomness in the yield curve does not by itself cause the slope to have forecasting power.

This leaves a hole in the 'quirk of fate' explanation, one that we attempt to fill in the next chapter.

5 A proposed behavioural finance explanation

If excess returns are not due to risk premia then how do we explain their predictability? If they are not a fluke or compensation for risk then that leaves us with market inefficiency explanations. It would take an adventurous leap of faith to believe there are large enough institutional frictions in the US Treasury market that could create inefficiencies sufficient enough to generate the observed excess returns, since it is probably the most liquid market in the world. Therefore we focus on the second assumption of equation (17), that the investors' expectations are unbiased.

If investors do not form their expectations of the excess returns using the real-world measure, $\mathbb{P}$, but instead using a subjective measure, $\mathbb{S}$, then the expected difference between the expected and realised excess returns could be non-zero. We follow the work of Shiller[16] in proposing cognitive biases, namely overreaction. This suggests that the market is not 'rational' and gets overexcited upon receiving new information, overreacting to changes and thus having biased expectations.

To investigate whether investor overreaction can explain excess returns behaviour we study a model capturing the pertinent elements of the real-world economy.

5.1 Our model economy

Our model economy consists of two actors, the central bank and investors (a.k.a. the market). The central bank sets the current interest rate and the investors price interest rate products, in our case zero-coupon bonds. The investors are risk-neutral and, as described in Chapter 2, do not include the contribution from convexity in their pricing. Thus bond yields are the risk-neutral averages of the expected future path of the short rate. The model progresses in discrete time steps, that we set in our simulations to length of one year.

The central bankers would like to 'guide' the short rate from its current level to their desired reversion level at a preferred speed or 'aggressiveness'. However, they must respond to random shocks to the economy by annually adjusting their model parameters. At the beginning of each time period they publicly set the current short rate according to their model, updated to account for these parameter adjustments.

The investors meanwhile price bonds, and therefore the yield curve, according to their expectations of central bank behaviour. The problem the investors face is they only know the current short rate and that the central bank follows a particular mean-reverting affine model. They do not know any of the model parameters used by the central bank.

In our model the investors correctly guess the direction of the central bank's parameter adjustments but overreact by overestimating the size of the adjustments. In other words, if the central bank adjusts a parameter at the end of a time step, the investors adjust their corresponding parameter in the same direction but by a greater magnitude.

Bonds are therefore priced according to the risk-neutral investors' overreacted affine model: the yields at each time t are just computed as the expected path of the short rate from t to desired maturity T with overreacted parameters. The affine model is simply used to compute this average taking convexity considerations into account.

Convexity is small for bonds of maturity under 10 years[14] so to avoid its effects we equate the convexity-adjusted market yield with the average of the short rate in the investor measure. This is equivalent to pricing bonds as:

$$P_t^T = \exp\left(\mathbb{E}^{\mathbb{S}}\left[-\int_t^T r_s ds\right]\right), \tag{37}$$

where $\mathbb{S}$ is the investors' subjective (overreacted) measure.

We stress that in our economy investors are risk-neutral; therefore $\mathbb{P} = \mathbb{Q}$. This way any excess returns observed in the model cannot be due to risk premia as investors are not risk-averse.

5.2 Model simulation algorithm

In this dissertation we will present simulations run using two affine models. First the economy is simulated with the central bank using the Vasicek model, the simplest yield curve model. This is to bolster our intuition before using a more realistic model in the subsequent sections.

The Vasicek model consists of one state variable and three parameters:

- r_t - the short rate; the current central bank rate.
- σ_r - the volatility of the short rate.
- θ_t - the target rate; the rate that the central bank is currently aiming for.
- k_t^r - the reversion speed of the short rate to the target rate.

In this model the short rate follows the stochastic differential equation:

$$dr_t = k_t^r(\theta_t - r_t)dt + \sigma_r dz_t \tag{38}$$

For each simulation the short rate and parameters start at arbitrarily chosen initial values at time $t = 0$. We model each year as a discrete time step, τ. Over each time step the following steps occur.

One year passes ($\tau = 1$) and the central bank intends for the short rate to progress following the deterministic part of the Vasicek model:

$$r_t e^{-k_t^r \tau} + \theta(1 - e^{-k_t^r \tau}) \tag{39}$$

However, the central bank has to update its policy to respond to economic events. This policy revision at time $t+\tau$ involves adjusting the value of the short rate, target rate and reversion speed to levels different to those that were anticipated at time t for time $t+\tau$. This is modelled by drawing deltas, Δ_t^j, from log-normal distributions with zero drift and a corresponding constant volatility, σ_j, for each parameter[9]:

$$\Delta_t^j := \exp\left(-\frac{\sigma_j^2 \tau}{2} + \sigma_j \sqrt{\tau} \epsilon_t^j\right), \tag{40}$$

where $j = r, \theta, k^r$ and time step $\tau = 1$. Note that the expectation of Δ_t^j is 1.

Each delta is cross-sectionally independent of the others. However, since business cycles appear to be an important factor in excess returns, as discussed earlier, the deltas are auto-correlated in order to mimic economic cycles. The Brownian shocks, ϵ^j, are independent for each j and are iteratively drawn with the same auto-correlation of ρ.

$$\epsilon_t^j = \rho \epsilon_{t-1}^j + \sqrt{1-\rho^2}\eta_t^j, \quad \eta_t^j \text{ independently drawn from } N(0,1), \tag{41}$$

where again $j = r, \theta, k^r$.

The intended short rate and parameters are multiplied at each time step by their corresponding delta Δ_t^j giving us the actual (rather than intended) central bank model parameters at the next time step:

$$r_{t+1} = \Delta_t^r \left(r_t e^{-k_t^r \tau} + \theta(1 - e^{-k_t^r \tau})\right) \tag{42}$$

$$\theta_{t+1} = \Delta_t^\theta \theta_t \tag{43}$$

$$k_{t+1}^r = \Delta_t^{k^r} k_t^r \tag{44}$$

The central bank knows perfectly the parameters of their model at each time step. As explained earlier, our investors only ever know (for certain) at each t the short

[9] Note that for $j = r$ this is technically a different σ_r from the one in equation (38). However, since we only use the deterministic part of the Vasicek model, the short rate volatility in our simulation actually only comes from the deltas.

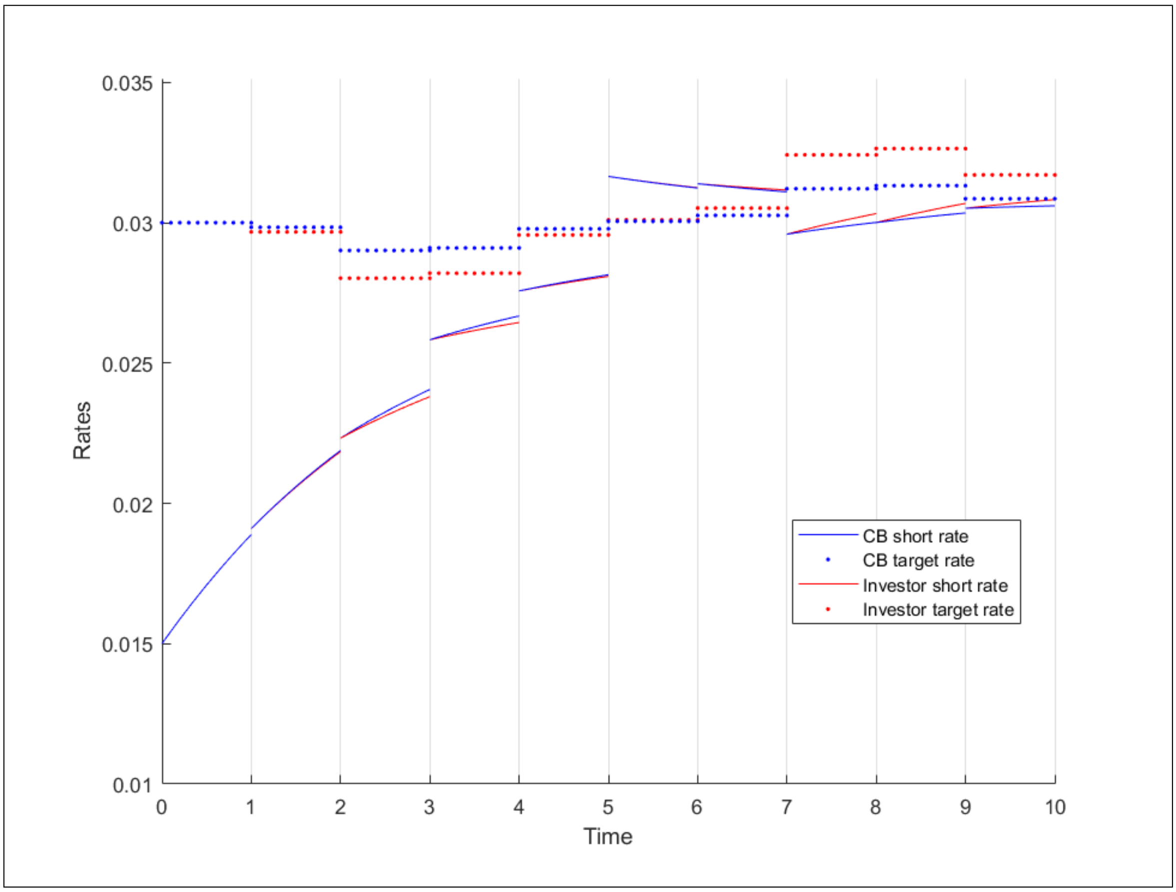

Figure 5: Example simulation of the short rate and target rate of the central bank and investors over 10 years.

rate, r_t, since this is the only value the central bank announces. Investors guess the parameters but overestimate the change to the central bank's target rate by a factor of $K > 1$.

In other words if the target rate has increased, the investors believe it has increased even more and vice versa. It is important to stress that the investors produce unconditionally unbiased estimates of the central bank's true process parameters.

The investors' target rate at each time step is

$$\theta^I_{t+1} := \theta^I_t + K(\theta_{t+1} - \theta_t) \tag{45}$$

Figure 5 displays an example simulation for ten time steps, i.e. ten years. Note how the rates change at each time step due to the deltas as well as the visible autocorrelation of the deltas.

Using the investors' overreacted model we calculate what the investors price the yields to be (neglecting convexity effects). As explained above, this is the expected path of the short rate using the investors' subjective measure. The yield for maturity T at time t is as follows:

$$y_t^T = \theta_t^I + (r_t - \theta_t^I)\frac{1 - e^{-k_t^r T}}{k_t^r T} \tag{46}$$

These yields are the continuation of the red lines in Figure 5.

From the yields we can calculate the excess returns from investing at time t in a T-maturity bond, funding the purchase at the 1-year bond, and unwinding the portfolio after one year, as detailed in Chapter 2.

$$\chi_{t\to t+1}^T = p_{t+1}^{T-1} - p_t^T - y_t^1, \tag{47}$$

where p_t^T again denotes the log-bond price. This gives in terms of the yields

$$\chi_{t\to t+1}^T = -(T-1)y_{t+1}^{T-1} + Ty_t^T - y_t^1 \tag{48}$$

This process is iterated over the desired number of years/time-steps. Once the simulation has been run we can analyse the behaviour of the excess returns.

We define the slope to be the difference between the 5-year maturity yield and the short rate (i.e. zero-year maturity yield):

$$sl_t := y_t^5 - r_t \tag{49}$$

Maturity of 5 years has been chosen to align with Cochrane and Piazzesi[7].

We regress the average excess return across maturities against the yield curve slope, giving the linear regression model:

$$\bar{\chi}_{t\to t+1} := \sum_{T=2}^{5} x_{t\to t+1}^T = \alpha + m_{sl} sl_t + \gamma_t, \tag{50}$$

where α is the intercept and γ_t are the errors.

We also regress the average excess return against the one-year forward rates of maturity T, which are calculated as per normal:

$$f_t^T = Ty_t^T - (T-1)y_t^{T-1} \tag{51}$$

This gives the model matching Cochrane and Piazzesi's in equation (22):

$$\bar{\chi}_{t\to t+1} = a + \sum_{T=1}^{5} m_T f_t^T + \delta_t, \tag{52}$$

where a is again the intercept and δ_t are the errors.

From this we can assess the coefficient of determination, R^2, and the regression coefficient(s) and intercept as well as p-values. Additionally we can calculate the Sharpe ratio of the annual returns from the unconditional long duration strategy (i.e. mean of $x^T_{t\to t+1}$ over time divided by standard deviation of $x^T_{t\to t+1}$).

For each simulation we also calculated the conditional Sharpe ratios. These are the Sharpe ratios of an investor who instead of blindly going long duration each period, instead invests conditional on the direction of the yield curve slope. With the 'up' strategy the investor enters the long duration position at the beginning of each period where the yield curve is upwards sloping and stays out of the market when it slopes downwards. The 'down' strategy is, naturally, the opposite position with the investor shorting duration when the yield curve slopes downwards and not investing when it is upwards-sloping.

The simulations were run for 1000 time steps, i.e. 1000 years. The short rate and target rate were initially set to the same value of 2%, in order to avoid any effects from an initial slope. We found that the simulations had variable results so in order to aid comparison between parameter configurations each set-up was run 100 times. So note that the Sharpe ratio and regression results in Table 3 are the means of 100 simulations.

The first thing to check is that the model behaves as expected when it is deterministic. We expect that when 'forwards come true' the Strong Expectations Hypothesis holds and there are zero excess returns. This is shown in the top row of Table 3. When there is no volatility and the investors guess the central bank's parameters perfectly there are no excess returns, which is as expected. This de facto means the slope and forwards are not predictive as the forwards always come true.

The next thing to note is that the unconditional Sharpe ratio of the duration strategy is tiny for all the parameter configurations. This is because our model is not directionally biased, so the slope is downwards sloping as often as it is upward sloping. In the long run any profits are offset by losses. Therefore unconditional Sharpe ratios *should* be close to zero in our model. Our simulations in Chapter 4 tell us that this is not inconsistent with the observed upwards-sloping average yield curve, as we could be in one of the one-in-five cases where just by chance the slope averages as steeply upwards.

We now look at what the effects of introducing volatility, autocorrelation and investor overreaction are on the predictiveness of the slope and the forwards. Table 3 displays data for five different parameter configurations. The values of the σs are based on an approximate calibration of the model to the observed short rate path in

	Model Parameters					Sharpe Ratio			Slope			Forwards		
Description	σ_r	σ_θ	σ_{k^r}	ρ	K	Uncond.	Up	Down	R^2	95%	99%	R^2	95%	99%
Deterministic	0	0	0	0	1	NaN	NaN	NaN	0	NaN	NaN	0	NaN	NaN
Only volatility	0.004	0.0012	0.05	0	1	-0.00573	0.00377	0.00129	8.07E-4	3	1	4.49E-3	4	1
Autocorrelation	0.004	0.0012	0.05	0.64	1	-0.00324	0.125	0.135	0.0590	97	94	0.0760	95	94
Overreaction	0.004	0.0012	0.05	0	1.6	-0.0141	0.0858	0.103	0.0348	52	44	0.0535	64	54
Complete	0.004	0.0012	0.05	0.64	1.6	0.0221	0.284	0.262	0.245	97	97	0.278	97	95

Table 3: Average statistics of regressions against the slope and against the one-year forward rates up to 5-year maturity from simulations using the Vasicek model. 100 simulations were run for each selected, representative parameter configuration, each with 1000 time steps and a flat initial curve: $r_0 = \theta_0 = 0.02$, $k_0^r = 0.1$. For each configuration we display the mean unconditional, up-conditional and down-conditional Sharpe ratios, as well as the mean R^2 of regressions against the slope and the forwards, and the percentage of simulations where each R^2 was significant to 95% and 99% confidence.

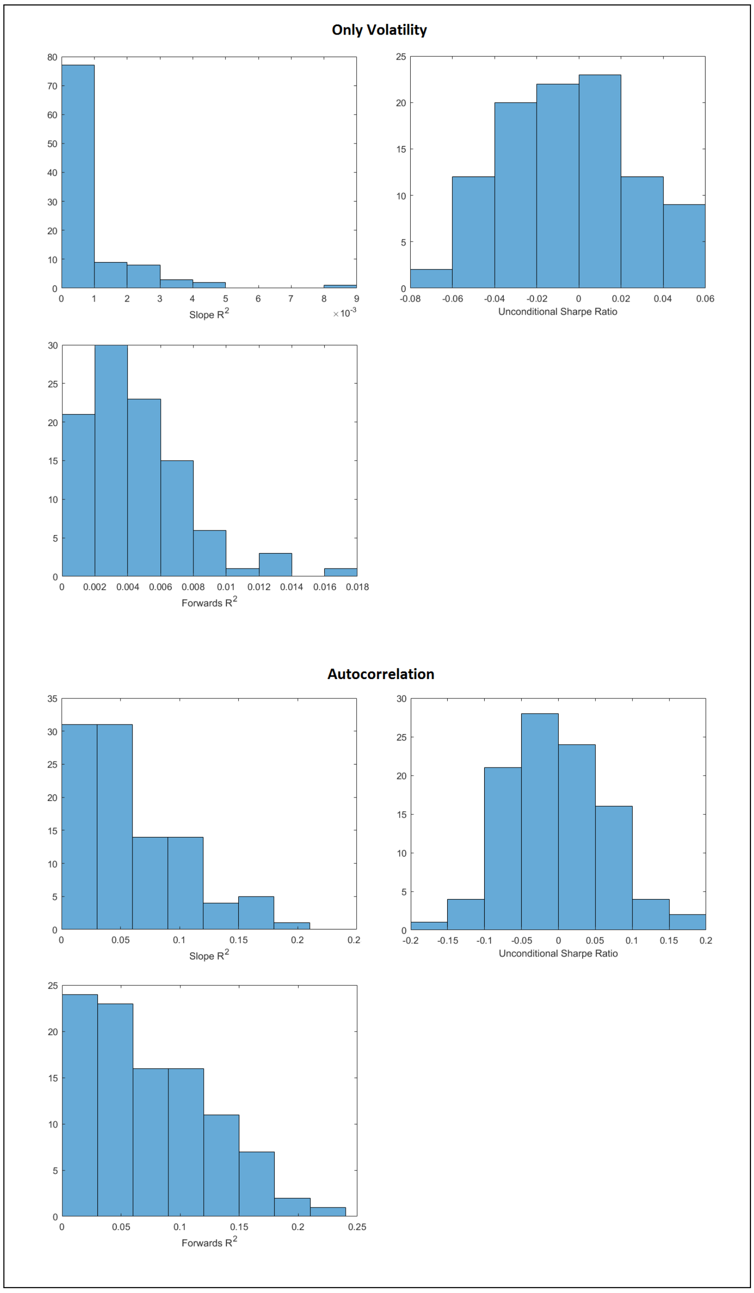

Figure 6: Histograms of R^2 values from slope and forwards regressions and unconditional Sharpe ratio from two sets of 100 simulations using the Vasicek model. The two parameter configurations "Only Volatility" and "Autocorrelation" correspond to those in Table 3.

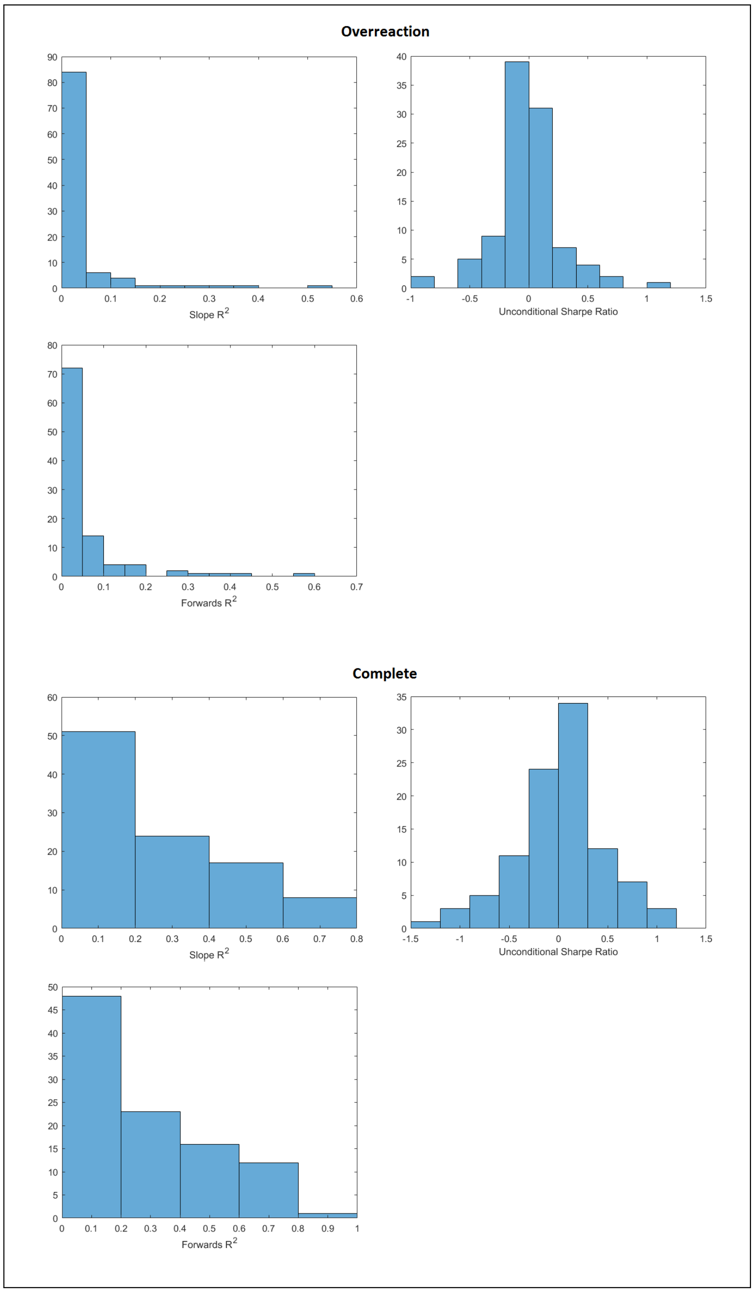

Figure 7: Histograms of R^2 values from slope and forwards regressions and unconditional Sharpe ratio from two sets of 100 simulations using the Vasicek model. The two parameter configurations "Overreaction" and "Complete" correspond to those in Table 3.

order to be reasonably representative.

It was found that only 'turning on' the constituent parts, so to speak, fundamentally changed the behaviour of the model; varying the magnitude of the parameters did not. So similar parameter values are used for all our results in this dissertation in order to focus attention on the pertinent behaviours of the model as well as aid comparison.

We discovered that when we introduce volatility by itself some excess returns are naturally generated but they remain unpredictable. The R^2 values of both the slope and the forwards regressions are tiny, and not significant the majority of the time. This can be seen in the "Only Volatility" parameter configuration: second row of Table 3 and top of Figure 6. Thus, in our model, randomness may cause excess returns but leaves the slope without forecasting power, substantiating the scepticism for the pure chance explanation in Chapter 4.

Next we looked at the effect of introducing business cycles into our economy, since Table 1 suggested that these are important. The results are displayed as the parameter configuration "Autocorrelation". When the changes in parameters, Δ_t^j, are autocorrelated ($\rho = 0.64$) we see in Figure 6 that both the slope and forwards R^2 values are now distributed between 0%-0.25% and focused around 5%-10%. The forwards are slightly more predictive than the slope, as they must be, although not considerably so. One can also see in the third row of Table 3 that the majority of slope and forwards R^2 values are significant now to 95% and 99% confidence. In addition to this, the conditional Sharpe ratios have increased, which is consistent as the slope has some forecasting power now. This all supports the claim that business cycles play a role in excess returns.

We now wish to investigate the effect of overreaction. In order to separate it from the effect of business cycles, we remove autocorrelation for now, and add a small amount of investor overreaction ($K = 1.6$). This is shown in the parameter configuration "Overreaction" in Table 3 and Figure 7. We can see that overreaction leads to much more varied R^2 values, although they are centred around a similar but slightly lower level to that of the "Autocorrelation" configuration. As you can see in the fourth row of Table 3, a lower percentage of the R^2 values are significant than with autocorrelation. Also the conditional Sharpe ratios are lower, although we see a wider range of unconditional Sharpe ratios. This would suggest that overreaction causes more variation of excess returns but business cycles are the more influential component.

What is interesting is when both serial autocorrelation and overreaction are combined, so that all the elements of our model are in play; we see a significant increase in the conditional Sharpe ratios and the R^2 values, as well as an increase in the number of coefficients of determination being significant. However, there remains no marked difference between the R^2 values from the slope and the forwards. Furthermore, the unconditional Sharpe ratio is more widely distributed, as can be seen in Figure 7. We additionally note that the distribution of the unconditional Sharpe ratios now encompasses the observed Treasury Sharpe ratio values in Table 1. We can therefore infer that business cycles and overreaction may have a compounding effect on the magnitude and predictiveness of excess returns.

We can conclude that if investors overreact to autocorrelated changes by the central bank then:

1. term premia become state-dependent;
2. excess returns from the difference between forward rates and the realised future short rate have some predictability;
3. the yield curve slope becomes a significant predictor of excess returns.

Furthermore, the size of the excess returns and the Sharpe ratios, as well as the forecasting power of the slope and forwards, are of similar magnitude to those observed in reality.

In order to further investigate Cochrane and Piazzesi's[7] "tent" factor we looked at the shape of the forward rate regressors. Our simulations produced a stark range of factor shapes, even for identical configurations, as can be seen in Figure 8.

This implies that the return predicting factor's shape is not robust to the fundamentals driving excess returns and is sensitive to minor changes in the data. Despite the five component factor being inevitably more predictive than just the slope, it is not remarkably so within this model.

One interesting behaviour of the model is the factor shape's sensitivity to σ_{k^r}. Figure 9 displays the shape of the forward rate regressors for two simulations with the "Complete" parameter configuration in row five of Table 3, but with σ_{k^r} equalling 0 and 0.0001 respectively. What can be seen is that a tiny addition of volatility to the reversion speed causes a massive change in the scale and shape of the factor. This 'knife-edge solution' behaviour is further evidence against the robustness of the Cochrane and Piazzesi return predicting factor shape.

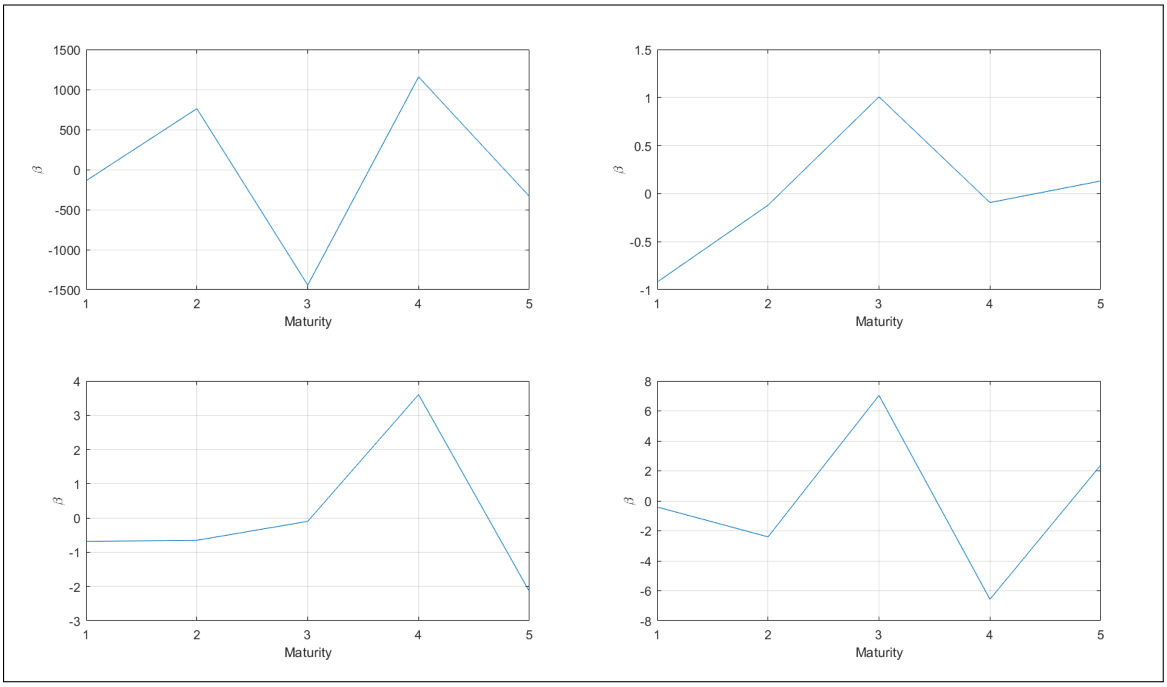

Figure 8: Shapes of Cochrane and Piazzesi's return predicting factor from four identically configured simulations using the Vasicek model and the "Complete" set of parameters from Table 3. The x-axis gives the maturity of the forward rate on the right-hand side of the regression. The y-axis is β from the unrestricted regressions (equation 9) of bond excess returns on all forward rates. The shapes as well as the magnitude of the βs are quite varied despite the identical simulation set-up.

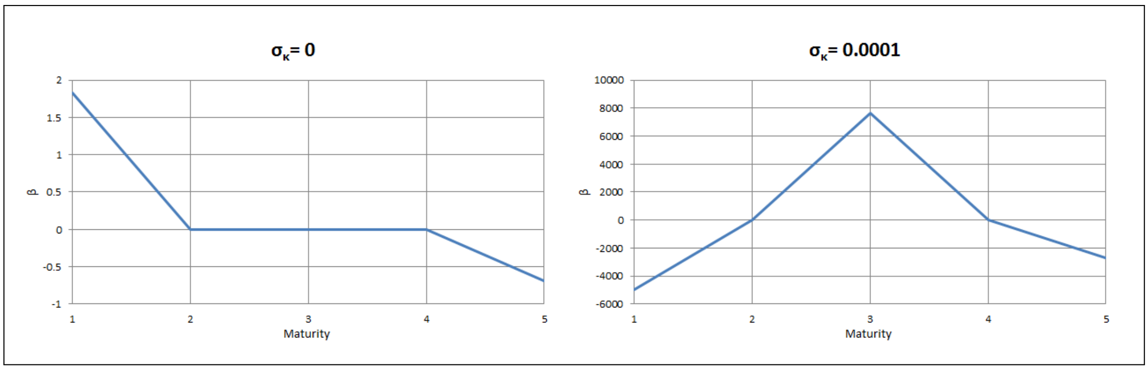

Figure 9: Shapes of Cochrane and Piazzesi's return predicting factor from two simulations using the "Complete" set of parameters from Table 3, but with σ_{k^r} of 0 and 0.0001 respectively. The x-axis gives the maturity of the forward rate on the right-hand side of the regression. The y-axis is β from the unrestricted regressions (equation (9)) of bond excess returns on all forward rates. The shape and magnitude of the βs are vastly different for only a tiny change in reversion speed volatility.

5.3 A more realistic model - DMRV

Our simple model has given us an intuitive feel for how business cycles and overreaction can affect excess returns and their predictiveness. However, the problem of using the Vasicek model for yield pricing is the monotonically decaying nature of the term structure of volatility; investors see the long-term level changing but are pricing with a constant reversion level. This is inconsistent and requires significantly more 'irrationality' from the investors than simply overreacting.

These problems can be resolved by instead using the Doubly-Mean-Reverting Vasicek (DMRV) model. In this model the target rate is no longer constant but itself reverts to a long-term reversion level. This allows the yield curve to form a 'hump' as the short rate quickly tends towards the target rate, which is itself reverting more slowly to the long-term reversion level. Thus, the DMRV model avoids inconsistent investor behaviour as well as allowing much more realistic yield curves to be produced, and thus is a closer model of the real-world economy.

The DMRV model has the same parameters as the Vasicek model, r_t, σ_r, θ_t and k_t^r, although θ_t becomes a state variable. In addition to these, it also has four extra parameters:

- σ_θ - the volatility of the target rate. To remain consistent with the financial interpretation this should be less than σ_r.
- r_t^∞ - the long-term reversion level; the long term rate that the central bank hopes to eventually reach. This represents the 'longer run' target of the FOMC as in Figure 10.
- k_t^θ - the reversion speed of the target rate to the long term reversion level. To remain consistent with the financial interpretation this should be less than k_t^r.
- ρ^z - the correlation between the Wiener processes of the short rate and the target rate, z_t^r and z_t^θ. This does not appear in the yield calculation so is not relevant for our investigation.

In the DMRV model the state variables, the short rate and the target rate, follow the SDEs:

$$dr_t = k_t^r(\theta_t - r_t)dt + \sigma_r dz_t^r \tag{53}$$

$$d\theta_t = k_t^\theta(r_t^\infty - \theta_t)dt + \sigma_\theta dz_t^\theta \tag{54}$$

$$\mathbb{E}[dz^r dz^\theta] = \rho^z dt \tag{55}$$

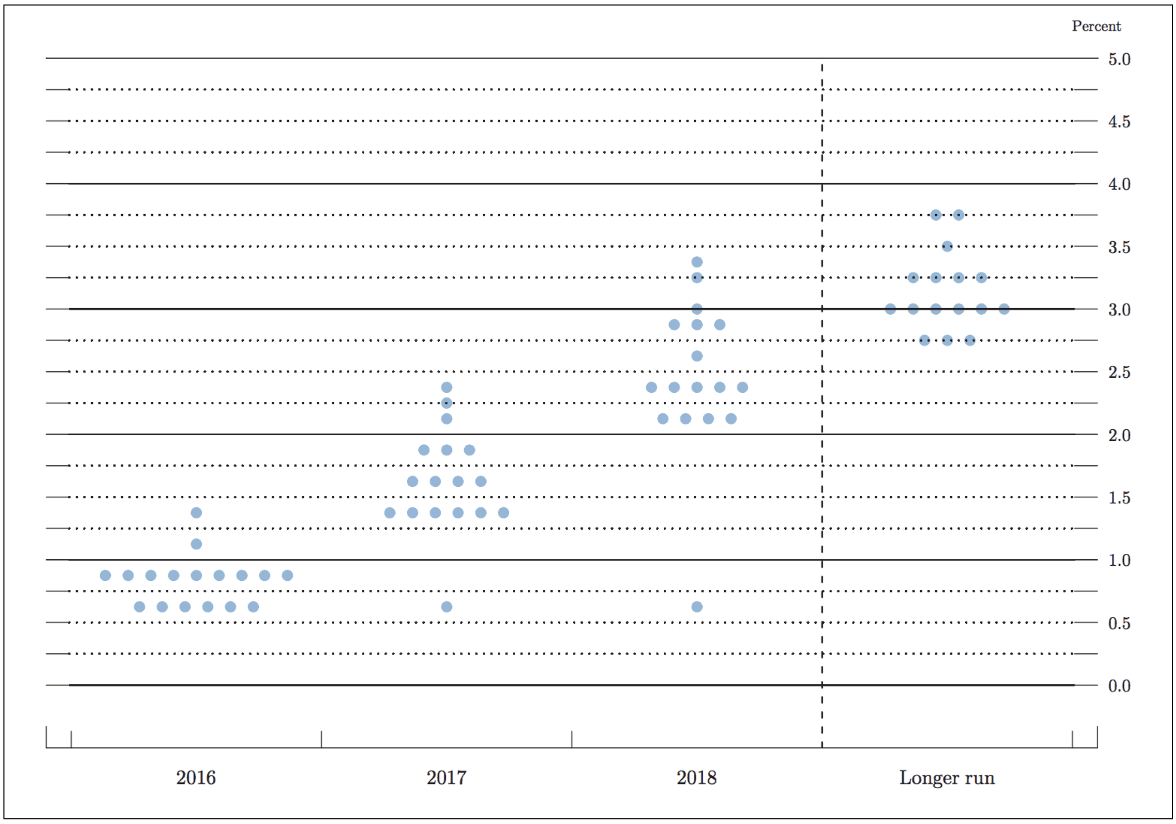

Figure 10: Federal Open Market Committee's (FOMC) "blue dots" chart. Each dot represents one FOMC participant's judgement of the midpoint of the appropriate target range for the federal funds rate at the end of the specified calendar year or long run. Values are rounded to the nearest 1/8 percentage point. This chart is used to give market participants an indication of the Fed's future expectations of the short rate. This particular chart was taken from the FOMC meeting notes for 15/6/2016.

Whilst the long-term reversion level is now the strongest factor in yield curve behaviour, all of the parameters affect its shape. High volatilities and low reversion speeds also ‘pull down’ the long end of the yield curve via convexity. This means that the reversion speeds play a greater role in the DMRV model than in the Vasicek model.

The simulation algorithm is the same as in section 5.2 when using the Vasicek model, but with the following modifications to adapt for using the DMRV model instead.

At the beginning of each time period, the central bank now sets an intended path for both the short rate and the target rate as well, following the deterministic part of the model. The model's parameters are once again updated as the central bank adjusts its policy at the end of the time period, before setting the path again for the following time period. This means we have Δ_t^j as in equation (40) but for $j = r, \theta, r^\infty, k^r, k^\theta$, which are autocorrelated as before with constant volatilities σ_j. So our actual central bank parameters are calculated iteratively as:

$$r_{t+1} = \Delta_t^r \left(r_t e^{-k_t^r \tau} + \frac{k_t^r(\theta_t - r_t^\infty)}{k_t^r - k_t^\theta}(e^{-k_t^\theta \tau} - e^{-k_t^r \tau}) + r_t^\infty (1 - e^{-k_t^r \tau}) \right) \tag{56}$$

$$\theta_{t+1} = \Delta_t^\theta \left(\theta_t e^{-k_t^\theta \tau} + r_t^\infty (1 - e^{-k_t^\theta \tau}) \right) \tag{57}$$

$$r_{t+1}^\infty = \Delta_t^{r^\infty} r_t^\infty \tag{58}$$

$$k_{t+1}^r = \Delta_t^{k^r} k_t^r \tag{59}$$

$$k_{t+1}^\theta = \Delta_t^{k^\theta} k_t^\theta \tag{60}$$

Now when the investors overreact they do not just overestimate changes to the target rate but also changes to the long-term reversion level. When using DMRV, the investors' target rate, θ^I, and long-term reversion level, $r^{I\infty}$, are:

$$\theta_{t+1}^I = \theta_t^I + K(\theta_{t+1} - \theta_t) \tag{61}$$

$$r_{t+1}^{I\infty} = r_t^{I\infty} + K(r_{t+1}^\infty - r_t^\infty) \tag{62}$$

Again the yields are calculated as the expected path of the short rate using the investors' overreacted model (i.e. using their subjective measure) and ignoring convexity effects. The yield for maturity T at time t is as follows:

$$y_t^T = r_t^{I\infty} + \frac{\theta_t^I - r_t^{I\infty}}{k_t^\theta T} + \frac{r_t - r_t^{I\infty}}{k_t^r T}(1 - e^{-k_t^r T}) + \frac{k_t^r(\theta_t^I - r_t^\infty)}{T(k_t^r - k_t^\theta)} \left(\frac{e^{-k_t^r T}}{k_t^r} - \frac{e^{-k_t^\theta T}}{k_t^\theta} \right) \tag{63}$$

Excess returns, one-year forward rates and the slope are then calculated using these yields, and the analysis performed on them accordingly. Regressions of excess returns against the slope and forward rates were performed for the same five parameter configurations as in the previous section, with the addition of the new DMRV parameters. The results are displayed in Table 4 as well as Figures 11 and 12.

We can see that the model still behaves as expected when deterministic. We find that the addition of volatility, autocorrelation and overreaction have the same effects with the Doubly-Mean-Reverting Vasicek model as with the plain Vasicek model that we saw in section 5.2. However with the DMRV model, overreaction by itself produces lower, as well as less significant, slope and forwards R^2 values. The conditional Sharpe ratios as well as the predictiveness of the slope and forwards are also less in the "Complete" configuration, compared to using the Vasicek model, but still remain sizeable and significant. We believe this is due to the ability of the DMRV model to form 'humps', which cancel out some of the excess returns.

The results all still remain at a similar magnitude to the actual observed values. Again we see no large difference in the excess returns forecasting power of the one-year forward rates compared to the slope.

Figure 13 displays four different return predicting factor shapes all generated from identically configured simulations using the "Complete" parameter configuration from Table 4. As in Figure 8 when using the Vasicek model, we have a variety of factor shapes as well as large changes in the magnitude of the β values from equation 9. It is interesting to note that in these simulations we never observed the "tent" shape found by Cochrane and Piazzesi but instead only saw "bat" shapes as observed by Dai et al.[8] and "distorted tent" shapes additionally observed by Villegas[17].

Thus we can deduce that the conclusions drawn in section 5.2 when using the simple Vasicek model still hold when changing to a more representative model, the Doubly-Mean-Reverting Vasicek model.

Description	Model Parameters							Sharpe Ratio			Slope			Forwards		
	σ_r	σ_θ	σ_{r^∞}	σ_{k^r}	σ_{k^θ}	ρ	K	Uncond.	Up	Down	R^2	95%	99%	R^2	95%	99%
Deterministic	0	0	0	0	0	0	1	NaN	NaN	NaN	0	NaN	NaN	0	NaN	NaN
Only volatility	0.004	0.0012	0.0005	0.05	0.05	0	1	-6.05E-4	-9.87E-3	-1.27E-3	1.15E-3	7	2	5.37E-3	11	2
Autocorrelation	0.004	0.0012	0.0005	0.05	0.05	0.64	1	-4.867E-6	0.139	0.154	0.0685	99	96	0.0816	94	92
Overreaction	0.004	0.0012	0.0005	0.05	0.05	0	1.6	0.0116	0.0479	0.0343	9.64E-3	36	30	0.0231	44	34
Complete	0.004	0.0012	0.0005	0.05	0.05	0.64	1.6	-0.0201	0.191	0.211	0.129	97	95	0.156	96	94

Table 4: Average statistics of regressions against the slope and against the one-year forward rates up to 5-year maturity from simulations using the DMRV model. 100 simulations were run for each selected, representative parameter configuration, each with 1000 time steps and a flat initial curve: $r_0 = \theta_0 = r_0^\infty = 0.02$, $k_0^r = 0.1$, $k_0^\theta = 0.05$. For each configuration we display the mean unconditional, up-conditional and down-conditional Sharpe ratios, as well as the mean R^2 of regressions against the slope and the forwards, and the percentage of simulations where each R^2 was significant to 95% and 99% confidence.

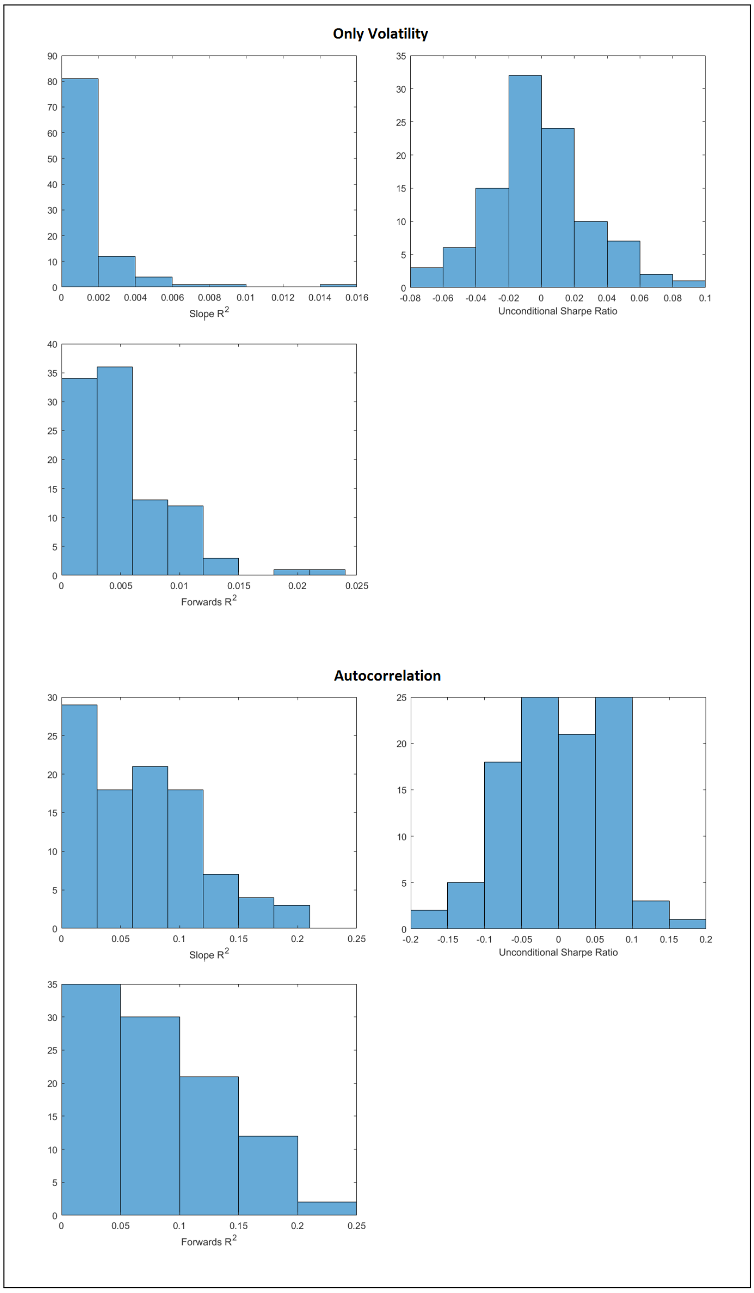

Figure 11: Histograms of R^2 values from slope and forwards regressions and unconditional Sharpe ratio from two sets of 100 simulations using the DMRV model. The two parameter configurations "Only Volatility" and "Autocorrelation" correspond to those in Table 4.

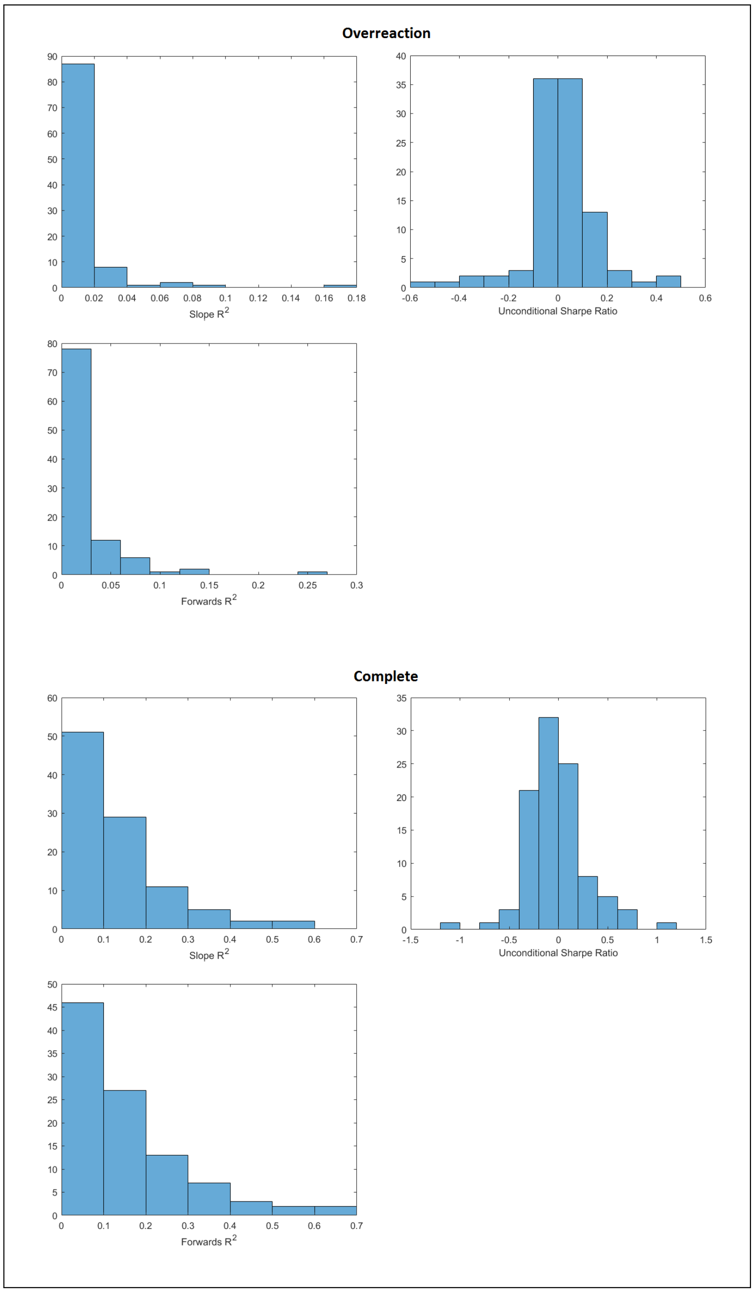

Figure 12: Histograms of R^2 values from slope and forwards regressions and unconditional Sharpe ratio from two sets of 100 simulations using the DMRV model. The two parameter configurations “Overreaction” and “Complete” correspond to those in Table 4.

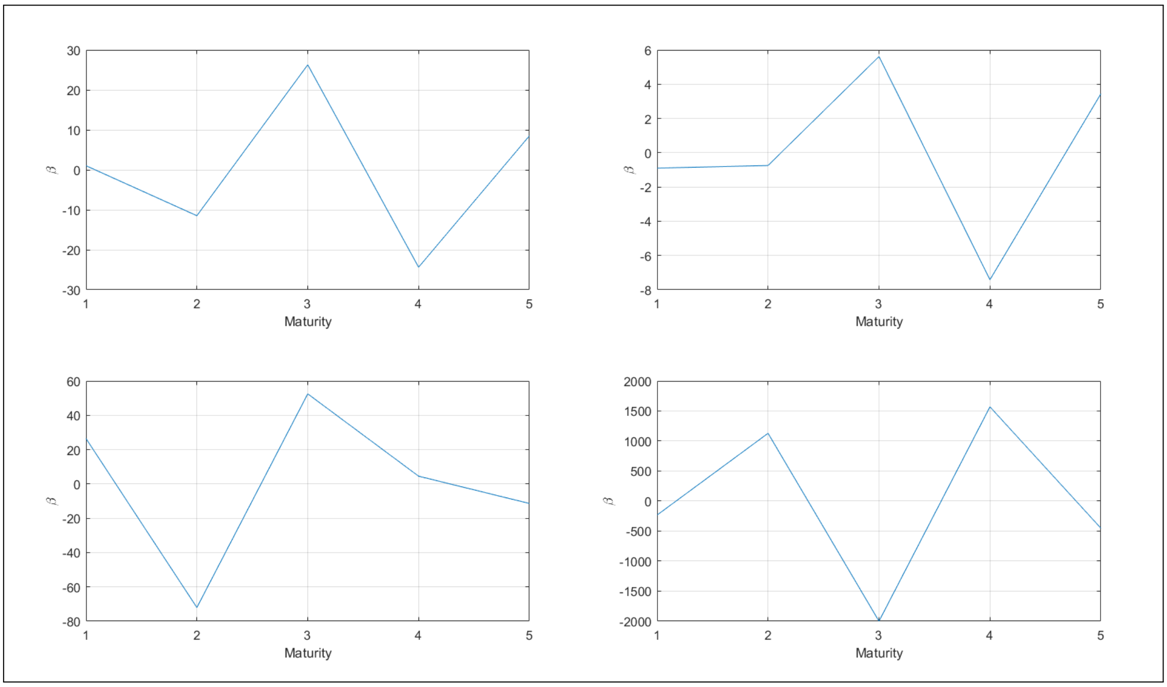

Figure 13: Shapes of Cochrane and Piazzesi's return predicting factor from four identically configured simulations using the DMRV model and the "Complete" set of parameters from Table 4. The x-axis gives the maturity of the forward rate on the right-hand side of the regression. The y-axis is β from the unrestricted regressions (equation 9) of bond excess returns on all forward rates. The shapes as well as the magnitude of the βs are quite varied despite the identical simulation set-up.

5.4 Overestimation of reversion speeds

So far we have only modelled investors overreacting to changes in the target rate and long-term level. However, as discussed earlier, the reversion speeds play an important role in the shape of the yield curve. Thus it is natural we also look at the effects of investors overreacting to changes in the reversion speeds.

We first look at simulations where investors only overreact to changes in the reversion speeds but not the central bank's target rates or long-term reversion level. This means the investors' parameters at each time step are:

$$\theta^I_{t+1} = \theta_{t+1} \tag{64}$$

$$r^{I\infty}_{t+1} = r^{\infty}_{t+1} \tag{65}$$

$$k^{Ir}_{t+1} = k^{Ir}_t + K(k^r_{t+1} - k^r_t) \tag{66}$$

$$k^{I\theta}_{t+1} = k^{I\theta}_t + K(k^{\theta}_{t+1} - k^{\theta}_t) \tag{67}$$

The yields are now calculated using the investors' reversion speeds:

$$y^T_t = r^{I\infty}_t + \frac{\theta^I_t - r^{I\infty}_t}{k^{I\theta}_t T} + \frac{r_t - r^{I\infty}_t}{k^{Ir}_t T}(1 - e^{-k^{Ir}_t T}) + \frac{k^{Ir}_t(\theta^I_t - r^{\infty}_t)}{T(k^{Ir}_t - k^{I\theta}_t)}\left(\frac{e^{-k^{Ir}_t T}}{k^{Ir}_t} - \frac{e^{-k^{I\theta}_t T}}{k^{I\theta}_t}\right) \tag{68}$$

The top row of Table 5 as well as Figure 14 display the results of simulations run with investors overreacting to reversion speed changes by the central bank. To aid comparison the simulation parameters are the same as the previous "Complete" parameters from Table 4.

One can see that the R^2 values are higher than when investors overreact to changes to the target rate and long term reversion level. We also see in Figure 14 that the peak of the distribution of the forwards R^2 is not at the far left for the first time. Additionally, a higher percentage of the simulations have significant R^2 values. The conditional Sharpe ratios are higher as well. This means that investors overreacting to reversion speed changes has a greater effect on the predictiveness of excess returns than overreacting to changes to the central bank's target rates.

What happens when investors overreact to all deltas? In other words, when investors overreact to changes to the target rate, long-term reversion level, and the reversion speeds of both. More precisely, when the investors' parameters are:

$$\theta^I_{t+1} = \theta^I_t + K(\theta_{t+1} - \theta_t) \tag{69}$$

$$r^{I\infty}_{t+1} = r^{I\infty}_t + K(r^{\infty}_{t+1} - r^{\infty}_t) \tag{70}$$

$$k^{Ir}_{t+1} = k^{Ir}_t + K(k^r_{t+1} - k^r_t) \tag{71}$$

$$k^{I\theta}_{t+1} = k^{I\theta}_t + K(k^{\theta}_{t+1} - k^{\theta}_t) \tag{72}$$

The results from these simulations are in the second row of Table 5 and the bottom of Figure 14. The R^2 values and conditional Sharpe ratios are even higher than with just one form of overreaction. The distributions of the slope and forwards R^2 values are now centred around 0.2-0.4 showing the forecasting power has significantly increased. Therefore the different forms of overreaction have additive effects on the excess returns predictability but investors misjudging central bank reversion speeds is more influential.

One can also see that the forwards return predicting factor is more predictive than the slope but, as with before, not markedly so.

5.5 Alternative forms of overreaction

We have assumed thus far that investor overreaction takes the form of overestimating the central bank's periodic adjustments. This is not the only way investors could potentially overreact to central bank policy. Thus, we look at an alternative form of overreaction to see whether it has a similar effect.

Instead of the annual changes, it is possible that investors instead overestimate the actual values of the central bank targets. In other words when the central bank is cutting rates, investors think they intend to cut them further than they are actually planning to and vice versa when hiking.

We consider this by modifying the calculation of the investors' target rate, θ^I, so that its distance from the short rate is K times the difference between the central bank's current target rate and short rate. So if the central bank's target rate is below the current short rate investors believe the target rate is K times lower. The investors' long-term reversion level, $r^{I\infty}$, is calculated similarly:

$$\theta^I_{t+1} = r_{t+1} + K(\theta_{t+1} - r_{t+1}) \tag{73}$$

$$r^{I\infty}_{t+1} = r_{t+1} + K(r^{\infty}_{t+1} - r_{t+1}) \tag{74}$$

These rates are then used in the calculation of the yields using equation (63).

The results from these simulations are displayed in the bottom row of Table 5 and in Figure 15. These show us that investors overestimating the values of the central bank's target rate and long-term reversion level also makes the slope and forwards predictive of excess returns. The effect is of a similar magnitude to when investors overreact to changes to the reversion speeds. Here, however, the difference between the forecasting power of the slope and the forwards is less.

	Model Parameters							**Sharpe Ratio**			**Slope**			**Forwards**		
Description	σ_r	σ_θ	σ_{r^∞}	σ_{k^r}	σ_{k^θ}	ρ	K	Uncond.	Up	Down	R^2	95%	99%	R^2	95%	99%
Overreact to rev. speeds	0.004	0.0012	0.0005	0.05	0.05	0.64	1.6	-0.0432	0.229	0.280	0.247	99	99	0.275	99	99
Overreact to all deltas	0.004	0.0012	0.0005	0.05	0.05	0.64	1.6	0.0528	0.355	0.313	0.313	99	98	0.375	99	99
Overreact to distance	0.004	0.0012	0.0005	0.05	0.05	0.64	1.6	-0.0232	0.259	0.273	0.241	100	100	0.251	100	100

Table 5: Average statistics of regressions against the slope and against the one-year forward rates up to 5-year maturity from simulations using the DMRV model and investors overreacting to: changes to the reversion speeds, changes to reversion speeds and rate levels, and distance of levels from short rate rather than changes. 100 simulations were run for each selected, representative parameter configuration, each with 1000 time steps and a flat initial curve: $r_0 = \theta_0 = r_0^\infty = 0.02$, $k_0^r = 0.1$, $k_0^\theta = 0.05$. For each configuration we display the mean unconditional, up-conditional and down-conditional Sharpe ratios, as well as the mean R^2 of regressions against the slope and the forwards, and the percentage of simulations where each R^2 was significant to 95% and 99% confidence.

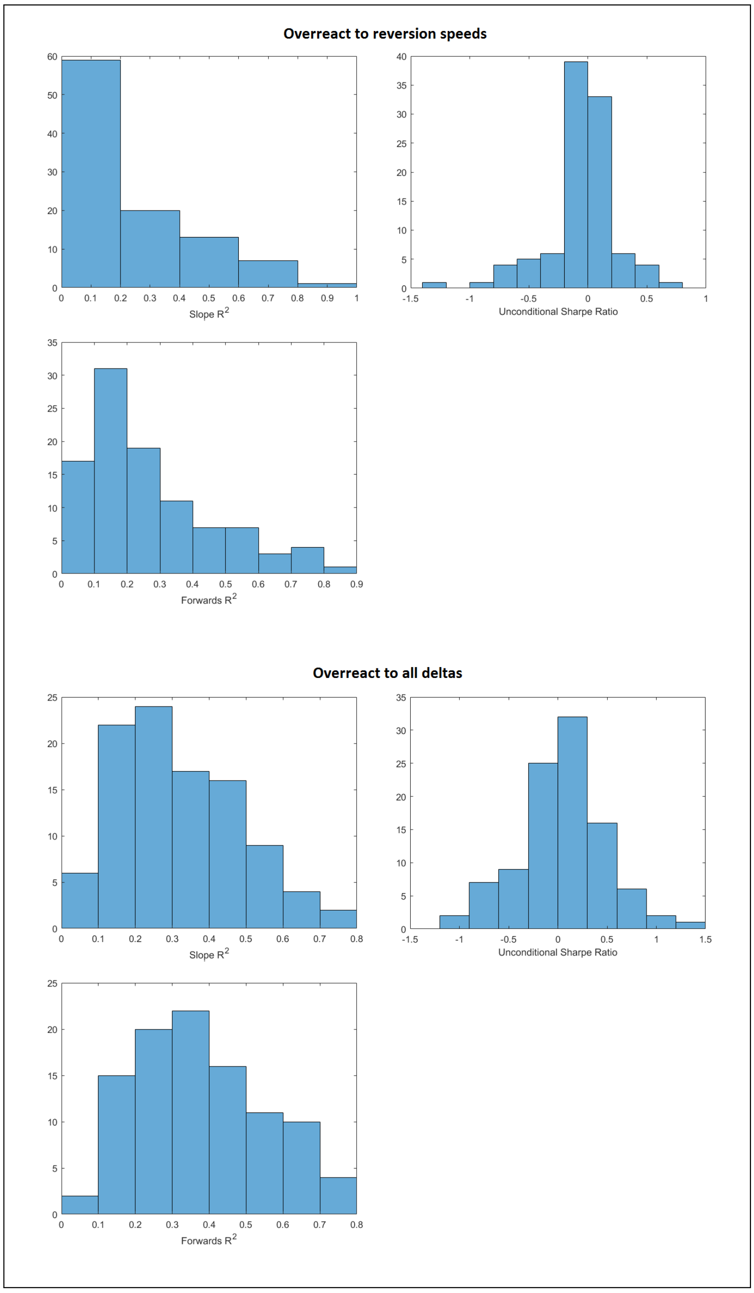

Figure 14: Histograms of R^2 values from slope and forwards regressions and unconditional Sharpe ratio from two sets of 100 simulations using the DMRV model. The two parameter configurations "Overreact to reversion speeds" and "Overreact to all deltas" correspond to those in Table 5

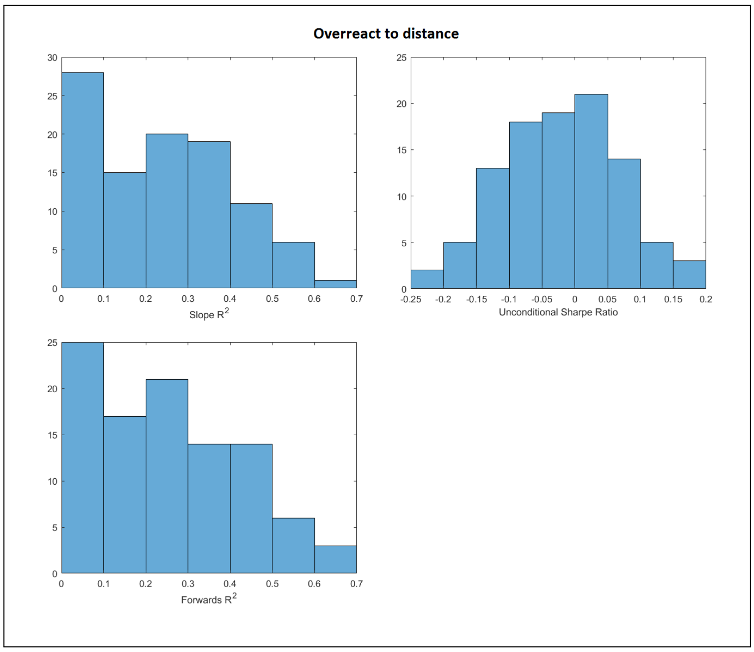

Figure 15: Histogram of R^2 values from slope and forwards regressions and unconditional Sharpe ratio from one set of 100 simulations using the DMRV model. The parameter configuration "Overreact to distance" corresponds to that in Table 5

6 Conclusion

In this dissertation we have investigated the established explanations for why the slope is a good predictor of excess bond returns.

If the Strong Expectations Hypothesis does not hold and excess returns will on average be reaped (as the evidence does suggest), we have provided evidence that seems inconsistent with the traditional asset pricing explanation, according to which a risk premium comes from the covariance between the payoff of a bond portfolio and consumption. Placing ourselves in a CAPM-like setting and assuming consumption is linked to equity returns, we observe a change in the sign of the correlation in the 1990s, separating the relationship history into two regimes. We fail to detect any systematic relationship between excess returns and the CAPM risk premium, proxied by the correlation between Treasury and equity returns. This finding suggests that if excess returns are due to risk premia, the relationship is not direct or simple.

The observation that the average yield curve slopes upwards has been taken as prima facie evidence of the existence of term premia. Our $\mathbb{P}$-measure simulations replaying the last half-century of yield curve history indicate this is not conclusive evidence for the existence of risk premia, as we found there was still a one-in-five likelihood of getting an even steeper yield curve slope than observed in reality purely by chance alone.

If the observed excess returns were purely a quirk of fate, however, then this does not explain why the slope is a good predictor. Our simulations have shown that randomness by itself only produces unpredictable excess returns. Therefore it is unlikely, although not impossible, that investors who were long duration were just 'once-in-a-lifetime' lucky.

We are left with the possibility that state-dependent excess returns may be due to factors other than risk premia. We have proposed that investor overreaction, in the vein of Shiller[16], may have a role to play.

Our simulations of a simple, but representative, economy with investors overreacting to the central bank's interest rate policy have clearly demonstrated that an overreaction mechanism, occurring over business cycles, generates state-dependent excess returns and causes the slope to be predictive of these returns. Only a small amount of overreaction is sufficient to create excess returns of a similar magnitude to those observed in the actual market, as well as achieve similar Sharpe ratios and slope predictiveness. This is achieved without any risk premia in our model.

We have shown that these effects also hold for a variety of ways investors might potentially overreact, establishing that our results are robust and that our conclusions are not contingent on the specific nature of the overreaction.

This dissertation contributes to the ongoing debate over Cochrane and Piazzesi's return predicting factor[7]. In our simple model we did not find the linear combination of forward rates to have significantly more forecasting power than the slope. The factor shapes we observed varied greatly and only rarely matched the "tent" shape proposed by Cochrane and Piazzesi, although they did correspond to the patterns observed elsewhere in the literature[15][17][3].

We have observed that the shape of the return predicting factor is also very sensitive to input changes and thus forms a 'knife-edge solution'; adding a tiny amount of volatility to the reversion speed causes the scale and shape of the factor to change drastically within our overreaction model. We believe this may be related to the sensitivity of linear estimation and the high correlation between the different forward rates.

Our proposed overreaction mechanism is not directionally biased so does not inherently lead to an upwards-sloping average yield curve. However, as our simulations recreating the yield curve history indicate, there is a reasonable probability of obtaining the observed average slope just by chance. Thus to fully explain all of the observations, we propose that overreaction combined with business cycles can explain the observed excess returns, Sharpe ratios and the predictiveness of the slope, and a 'quirk of fate' can explain the persistent upwards slope. We do not need to invoke risk premia.

The existence of excess returns and the effectiveness of the yield curve slope in predicting them could still be explained by a combination of chance and risk premia. However, investor overreaction is an alternative factor capable of explaining convincingly why the slope is a good predictor of excess returns.

References

[1] T. Adrian, R. K. Crump, B. Mills, and E. Moench. Treasury term premia: 1961-present. *Liberty Street Economics*, 2014. Data available at http://www.newyorkfed.org/research/data_indicators/term_premia.html, last accessed 27 January, 2017.

[2] Tobias Adrian, Richard K. Crump, and Emanuel Moench. Pricing the term structure with linear regressions. *Journal of Financial Economics*, 110(1):110–138, October 2013.

[3] Michael D. Bauer and James D. Hamilton. Robust bond risk premia. April 2015. revised January 2017.

[4] John Y. Campbell and Robert J. Shiller. Yield spreads and interest rate movements: A bird's eye view. *The Review of Economic Studies*, 58(3):495–514, May 1991.

[5] John H. Cochrane. *Asset Pricing:(Revised Edition)*. Princeton university press, 2009.

[6] John H. Cochrane. Comments on "Robust Bond Risk Premia" by Michael Bauer and Jim Hamilton. *Unpublished working paper. University of Chicago*, 2015.

[7] John H. Cochrane and Monika Piazzesi. Bond risk premia. *The American Economic Review*, 95(1):138–160, March 2005.

[8] Qiang Dai, Kennet Singleton, and Wei Yang. Predictability of bond risk premia and affine term structure models. *Standford University*, 2004.

[9] Robert Engle. *Anticipating correlations: a new paradigm for risk management*. Princeton University Press, 2009.

[10] Eugene F. Fama and Robert R. Bliss. The information in Long-Maturity forward rates. *The American Economic Review*, 77(4):680–692, 1987.

[11] Rebecca Hellerstein. Global bond risk premiums. 2011.

[12] Robert B. Litterman and Jose Scheinkman. Common factors affecting bond returns. *The Journal of Fixed Income*, 1(1):54–61, 1991.

[13] Monika Piazzesi, Juliana Salomao, and Martin Schneider. Trend and cycle in bond premia, March 2015. Available at http://web.stanford.edu/~piazzesi/trendcycle.pdf, last accessed 27 January 2017.

[14] Riccardo Rebonato. *Bond Pricing and Yield Curve Modelling: A Structural Approach*. Unpublished.

[15] Riccardo Rebonato. Return-Predicting factors for US treasuries: On the similarity of 'tents' and 'bats'. *Int. J. Theor. Appl. Finan.*, 18(04):1550028+, June 2015.

[16] Robert J. Shiller. Do stock prices move too much to be justified by subsequent changes in dividends? *National Bureau of Economic Research Working Paper Series*, pages 456+, February 1980.

[17] Daniela Villegas. Bond excess returns predictability and the expectation hypothesis. Master's thesis, 2015. Submitted in partial fulfilment of the requirements for the MSc in Mathematical Finance, University of Oxford.