An Interpretation of the Cieslak-Povala Return-Predicting Factor

Riccardo Rebonato
Oxford University

July 3, 2015

Abstract

This paper presents a simple reformulation of the restricted Cieslak and Povala (2010) return-predicting factor which retains by construction exactly the same (impressive) explanatory power as the original one, but affords an alternative and attractive interpretation. What determines the future returns, the new factor shows, is a function of the distance of the yield-curve level and the slope not from a fixed reference level, but from a *conditional* one, determined by a function of the long-term inflation.

The decomposition also allows a clear attribution of the predictive power of the Cieslak and Povala factor between the conditional level and slope deviations.

Finally, the new reformulation shows that once the conditionality is taken into account, level deviations are important predictors of excess returns. (Hardly any predictability was found in earlier studies for the *unconditional* level.)

1 Background[1]

Before the work by Cochrane and Piazzesi (2004, 2008), the conventional wisdom about bond excess returns was that the time-varying risk premium should be related to the slope of the yield curve, or some of its proxies. See, eg, Fama and Bliss (1987), Campbell and Shiller (1991).

In their well-known 2004 paper, Cochrane and Piazzesi (2004) showed that five forward rates, arranged in the shape of a 'tent', provide a more powerful return-predicting factor. It has been shown that the precise tent shape is not really essential (see Rebonato (2014)), but the finding that as many yield-curve-based factors as five have a greater predictive power than the slope (however defined), has been confirmed by several studies (see, eg, Hellerstein (2011), Adrian Moench and Crump (2013)).

[1] It is a pleasure to acknowledge the contribution provided by Dr Pottinton with discussions and calculations.

In a parallel strand of work, Ludvigson and Ng (2009) show that yield-curve variables do not fully span the uncertainty in the risk premium, and that adding some skillfully constructed macroeconomic variables significantly increases the predictive power of yield-curve-only-based return-predicting factors. Their conclusions are that 'by rendering most popular predictors insignificant, our forecasting factor aggregates a variety of macro-finance risks into a single quantity'.

Cieslak and Povala (2010) then presented an approach (briefly described in Section 2) that attempts to combine the yield-curve-based return-predicting-factors (such as the tent or the slope), and the macroeconomic predictors (such as the linear combinations of economic indicators in Ludvigson and Ng (2009). They describe the dynamics of the yield curve in terms of two cycles of very different frequencies, and produce an intuitively appealing and financially motivated return-predicting-factor that mixes yield-curve-based and macroeconomic factors and that performs significantly better than the Cochrane tent, and even better than the combination of the Cochrane tent and the Ludvigson and Ng factors.

In this paper we present a very simple reformulation of the restricted version of the Cieslak and Povala return-predicting factor. This formulation affords a very natural – and, we find, illuminating – alternative interpretation, and qualifies some well-known findings, such as the supposedly little explanatory power of the level of yields.

The simple insight our interpretation affords it that it is not the absolute (unconditional) slope or the level that matter in order to assess the magnitude of the expected excess return, but where the slope and the level of the yield curve stand with respect to where they 'should be' – where 'should be' refers to the state-dependent values determined by the macroeconomic quantity (a long-term proxy for inflation) that controls the low-frequency behaviour of the yield curve. So, the decomposition shows clearly that it is the *conditional* values of level and slope that matter, not their absolute values.

Since the conditioning depends on a macroeconomic variable (in the case of the Cieslak and Povala model on a proxy for the long-term inflation), the interpretation is consistent with the Ludvigson and Ng observation that yield-curve variables by themselves do not contain all the available information about excess returns, and helps achieving the Cieslak and Povala's goal of combining the two predictive strands.

Our simple decomposition also allows us to answer the question: is it the deviation of the slope or of the level from where they 'should be' that better explains excess returns? We find that both the 'conditional distance' of level and slope from their long-term-inflation prediction are important, with the level deviation explaining excess returns more at the short end, and the slope deviation at the long end. This should be contrasted with traditional findings that the unconditional level has very little (if any) explanatory power for excess returns.

The paper is organized as follows.

To help the reader following the argument and to establish the notation, we first give in Section 2 a very brief summary of those results in Cieslak and Povala necessary for our analysis. We give special emphasis to the restricted version of their return-predicting factor in Section 3. Then in Section 4 we present our equivalent decomposition and provide its interpretation. Finally, we carry out the attribution analysis mentioned above in Section 5. We present our conclusions in Section 6.

2 The Approach by Cieslak and Povala

Cieslak and Povala (2010) start from the definition of yields as the $\mathbb{P}$-measure (real-world) expectation of the future path of the short rate plus a risk premium component. In discrete time this means:

$$y_t^{(n)} = \underbrace{\frac{1}{n}E^{\mathbb{P}}\left[\sum_{i=0}^{n-1} r_{t+i}\right]}_{\text{expectation}} + \underbrace{rpy_t^{(n)}}_{\text{risk premium}} \tag{1}$$

In keeping with their two-frequency description of the economy, the Authors then posit that the short rate should evolve as the sum of a highly-persistent (actually, unit-root) component, τ_t, and a quickly-mean-reverting $AR(1)$ component, x_t, with autoregression coefficient ϕ_x:

$$r_t = \rho_0 + \overbrace{\underbrace{\rho_\tau \tau_t}_{\text{unit-root part}} \quad + \quad \underbrace{\rho_x x_t}}^{\text{two factors differing in persistence}} \tag{2}$$

Given the value of the short rate at time t, one can easily write the expression both for the value of the short rate i periods ahead:

$$r_{t+i} = \rho_0 + \rho_\tau \tau_{t+i} + \rho_x x_{t+i} \tag{3}$$

and for the quickly-mean-reverting part

$$x_{t+1} = \mu_x + \phi_x x_t + \sigma_x \epsilon_{t+1}^x. \tag{4}$$

By iterating and by direct substitution into Equation (1) one gets

$$y_t^{(n)} = \underbrace{b_0^{(n)} + b_\tau^{(n)} \tau_t + b_x^{(n)} x_t}_{\text{expectation}} + \underbrace{rpy_t^{(n)}}_{\text{risk premium}} \tag{5}$$

$$y_t^{(n)} = \underbrace{b_0^{(n)} + b_\tau^{(n)} \tau_t}_{\text{persistent component}} + \underbrace{b_x^{(n)} x_t + rpy_t^{(n)}}_{\text{transitory component}} \tag{6}$$

with

$$b_\tau^{(n)} = \rho_\tau \tag{7}$$

and

$$b_x^{(n)} = \frac{1}{n}\rho_x \frac{\phi_x^n - 1}{\phi_x - 1}. \tag{8}$$

Cieslak and Povala therefore define as the maturity-dependent cycle the quantity $\widetilde{c}_t(n)$, given by

$$\widetilde{c}_t^{(n)} \equiv b_x^{(n)} x_t + rpy_t^{(n)}. \tag{9}$$

It is important to note that the contributions to the cycle from the risk premium and the expectation components vary with maturity. So, for instance, the risk premium for $n = 1$ is, of course, 0, and the cycles purely captures changes in short-rate expectations. However, as the maturity increase, because of its fast mean-reversion, the contribution from the $AR(1)$ process, ie, the coefficient $b_x^{(n)}$, becomes smaller and smaller, expectations become less and less important and risk premia contribute more and more.

As for the persistent component, it is linked by Cieslak and Povala to the time series of inflation, CPI_t, which is assumed to follow

$$CPI_t = \tau_t^{CPI} + CPI_t^c, \tag{10}$$

where CPI_t^c denotes a cyclical contribution to inflation, and τ_t^{CPI} a persistent (no-mean-reversion) inflation time-varying endpoint:

$$\tau_t^{CPI} = \tau_{t-1}^{CPI} + \epsilon_t^{\tau}. \tag{11}$$

A proxy for this persistent endpoint, τ_t^{CPI}, can be created as weighted moving average of past inflation data:

$$\tau_t^{CPI} = \frac{\sum_{i=0}^{t-1} v^i CPI_{t-i}}{\sum_{i=0}^{t-1} v^i}. \tag{12}$$

In Equation (12) the weight parameter, v, is set to 0.987[2], and the sum runs over 120 monthly observations.[3]

Operationally, the cycles are therefore obtained as follows. First yields are regressed against the *contemporaneous* persistent component calculated as per Equation (12):

$$y_t^{(n)} = b_0^{(n)} + b_\tau^{(n)} \tau_t + \epsilon_t^{(n)}. \tag{13}$$

Looking at Equations (6) and (9), it is clear that what has been 'left over' in the regression (13) is just the cycle, $c_t^{(n)}$. Therefore the cycle is calculated as the 'residual'[4]

$$c_t^{(n)} = y_t^{(n)} - \left[b_0^{(n)} + b_\tau^{(n)} \tau_t\right] =$$

[2] It can be shown that the estimator (12) is optimal and maximally robust when one wants to estimate a parameter, but is uncertain about the true data-generating process, and wants to make sure that the estimator is robust across different models. (See, in this respect, Evans, Honkapohja and Williams (2010) for a precise discussion, or Evans and Honkapohja (2009) for a more general account.)

[3] We appended a superscipt 'CPI' to indicate the the persistent component was estimated using inflation data. A parallel treatment can be carried out using as macroeconomic variable the savings rate. Cieslak and Povala show that the final results are little changed by the precise choice of the persistent variable.

[4] To be precise, one should distinguish between the 'true' (and unknown) coefficients a and b in the regression $y = a + bx + \epsilon$, and their estimated counterparts, $\widehat{a}$ and $\widehat{b}$. To keep the notation light we will not make this distinction.

$$y_t^{(n)} - \left[b_0^{(n)} + b_\tau^{(n)} \left(\frac{\sum_{i=0}^{t-1} v^i CPI_{t-i}}{\sum_{i=0}^{t-1} v^i} \right) \right]. \tag{14}$$

Cycles are important because, in the Authors' model, they are highly significant predictors of yield changes. To show that this is indeed the case, the Authors carry another regression, this time using as left-hand variable the time-t change in the n-maturity yield, $\Delta y_t^{(n)}$, against the *previous-time* values of the cycle, $c_{t-\Delta t}^{(n)}$, of the persistent component, $\tau_{t-\Delta t}$, and of the the yield change, $\Delta y_{t-\Delta t}^{(n)}$,

$$\Delta y_t^{(n)} = a_0 + a_c c_{t-\Delta t}^{(n)} + a_y \Delta y_{t-\Delta t}^{(n)} + a_\tau \tau_{t-\Delta t} + \epsilon_{t+1}. \tag{15}$$

When they do so, they find that the (negative) coefficient a_c associated with the cycles is highly significant.[5] It is the statistical significance of the negative coefficient that suggests the interpretation offered by Cieslak and Povala that a *higher* value of the cycle at time $t - \Delta t$ is associated with a *lower* value for the yield at time t, and hence to a *higher* excess return.

This intuition is made precise with the following analysis of excess returns. After estimating the cycles using Equation (14), these are now 'known quantities'. As we have seen, cycles are maturity-dependent. Six of these maturity-dependent cycles (ie, those associated with maturities of 1, 2, 5, 7, 10 and 20 years) are then used as regressors to explain excess returns:

$$rx_{t+1}^{(n)} = \delta_0 + \sum_{i=1,6} \delta_i c_t^{(i)} + \epsilon_{t+1}^{(n)}. \tag{16}$$

To establish a fair horse race against the Cochrane and Piazzesi return-predicting factor, Cieslak and Povala also use the six forward rates of maturities of 1, 2, 5, 7, 10 and 20 years as alternative regressors. This is therefore an extended Cochrane-Piazzesi model.

They find that in all samples 'and across all maturities, cycles give much stronger evidence of predictability than do forward rates',[6] a conclusion that is not changed by choosing different maturities, or increasing the number of forward rates. More precisely, the R^2 obtained using cycles varies from 0.43 to almost 0.60 for the whole-sample studies, to be contrasted with an R^2 ranging from (approximately) 0.20 to 0.30 when forward rates are used.

To illustrate the approach, Fig (1) shows a time series of inflation, of the ten yields of maturities from 1 to 10 years, of the proxy for the secular inflation (the variable τ) and of the results of the regression of the yields on the secular inflation proxy (the curves labelled 'predyield$n$').

The cycles (not shown) are then the differences between the jagged and the smooth lines (ie, between the yields and their predictions by the secular inflation).

[5] The regression in Equation (14) is carried out against the *same-time* values of the presitent component, τ_t. The predictive regression in Equation (15) links changes in yields at time t to cyles (and changes in yields, and changes in the persitstent varibale, τ) at time $t - \Delta t$.

[6] page 13.

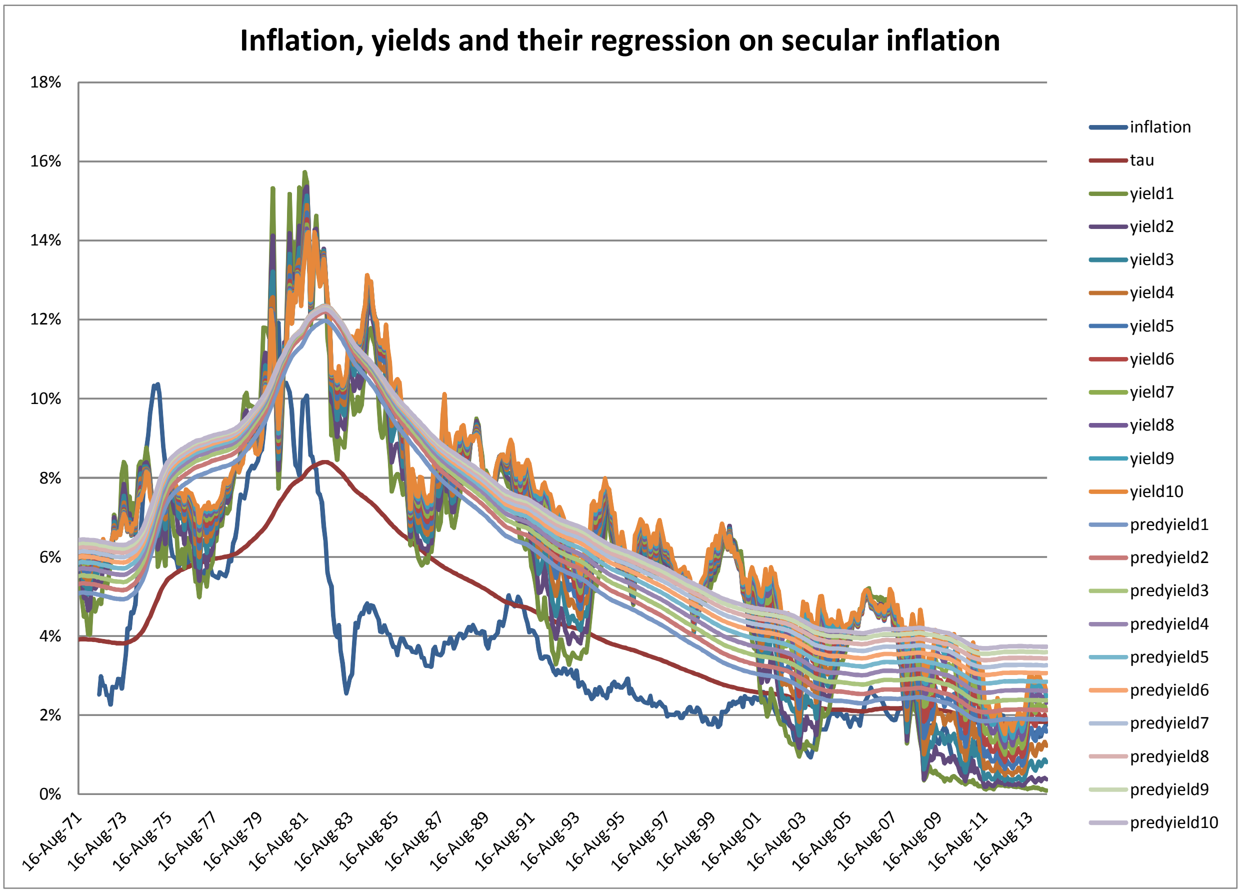

Figure 1: A time series of inflation, of the ten yields of maturities from 1 to 10 years, of the proxy for the secular inflation (the variable τ) and of the results of the regression of the yields on the secular inflation proxy (the curves labelled 'predyield$n$').

The reported R^2 statistics (closely reproduced in our analysis using a different data source) are so impressive that they prompt an obvious question: why does the Cieslak and Povala return-predicting factor perform *so* much better than the best forward-rate-base one (taking the Cochrane-and-Piazzesi's extended tent as the yield-curve based benchmark to beat)?

3 The Restricted RPF

In order to answer the question as clearly as possible by focussing on the essence of the approach, it pays to look at the Cieslak and Povala restricted cycle-based return-predicting factor, and compare it with the Cochrane-Piazzesi restricted factor.

The Cieslak and Povala restricted return-predicting forecasting factor, $\widehat{cf}$, is built as follows. They start from the definition of the maturity-dependent cycle in Equation (9):

$$\widetilde{c}_t^{(n)} \equiv b_x^{(n)} x_t + rpy_t^{(n)}. \tag{17}$$

Recall that the 'composition of' (the information conveyed by) the cycle depends on its maturity, with the cycle for $n = 1$ containing no information about risk premia (because $rpy_t^{(1)} = 0$), and a lot of information about expectations. As a consequence $\widetilde{c}_t^{(1)}$ "captures variation[s] in short rate expectations ($b_x^{(1)} x_t$), but not in premia.[7] [...] Therefore a natural way to decompose the transitory variation in the yield curve into the expectations part and the premium part is by estimating[8]:

$$rx_{t+1}^{(n)} = \alpha_0 + \alpha_1^{(n)} c_t^{(1)} + \alpha_2^{(n)} c_t^{(n)} + \epsilon_{t+1}^{(n)} \qquad n \geq 2." \tag{18}$$

Next Cieslak and Povala look at the average excess returns, $\overline{rx}_{t+1}$, (averaged over maturities), and use as regressors to predict the average excess returns i) the expectation-only-related factor, $c_t^{(1)}$, and ii) the average, $\overline{c}_t$, of the maturity-dependent cycle factors, $c_t^{(i)}$, $i = 2, 3, ..., n$:

$$\overline{rx}_{t+1} = \gamma_0 + \gamma_1 c_t^{(1)} + \gamma_2 \overline{c}_t + \overline{\epsilon}_{t+1}. \tag{19}$$

Their single (restricted) forecasting factor, $\widehat{cf}_t$ is therefore given by:

$$\widehat{cf}_t \equiv \gamma_0 + \gamma_1 c_t^{(1)} + \gamma_2 \overline{c}_t. \tag{20}$$

Cieslak-and-Povala then pit their single forecasting factor against the restricted single-tent factor built by Cochrane and Piazzesi. When they do so, once again, they find that the cycle-based factor, $\widehat{cf}$, explains a lot more than the restricted tent factor. (The R^2 coefficients they find range from a minimum of 0.41 (2-year maturity) to a maximum of 0.56 (15-year maturity).)

[7] page 8.
[8] page 14.

4 Interpretation of the Cieslak-Povala RPF

Let's look in some detail at the restricted Cieslak and Povala return-predicting factor:

$$\overline{rx}_{t+1} = \gamma_0 + \gamma_1 c_t^{(1)} + \gamma_2 \overline{c}_t + \overline{\epsilon}_{t+1}^{(n)} \tag{21}$$

It is important for the future discussion to note that, when the regression is carried out on US\$ data, the coefficient γ_1 (ie, the loading onto the short cycle) turns out to be significantly negative, the coefficient γ_2 (ie, the loading onto the average cycle) is significantly positive, and $|\gamma_2| > |\gamma_1|$. For reasons that will become apparent in a few lines, we prefer to work with positive coefficients and we introduce γ_3:

$$\gamma_3 = -\gamma_1 \tag{22}$$

With some foresight we also choose to write

$$\gamma_2 = \gamma_3 + \Delta\gamma \tag{23}$$

with $\Delta\gamma > 0$ because $|\gamma_2| > |\gamma_1|$.

With these two definitions we therefore have

$$\overline{rx}_{t+1} = \gamma_0 - \gamma_3 c_t^{(1)} + (\gamma_3 + \Delta\gamma)\,\overline{c}_t + \overline{\epsilon}_{t+1}^{(n)}. \tag{24}$$

Next, recall that we defined for the two cycles to be given by the residuals of the regression of the 1-year yield and average yield against the slowly moving inflation proxy, τ_t:

$$c_t^{(1)} = y_t^{(1)} - \left[b_0^{(1)} + b_\tau^{(1)}\tau_t\right] \tag{25}$$

and

$$\overline{c}_t = \overline{y}_t - \left[\widehat{b}_0 + \widehat{b}_\tau \tau_t\right] \quad , \tag{26}$$

respectively.

Substituting these expressions into Equation (24) gives

$$\begin{gathered}\overline{rx}_{t+1} = \\ \gamma_3 \left\{ \underbrace{\left(\overline{y}_t - y_t^{(1)}\right)}_{\text{actual slope}} - \left[\underbrace{\left(\widehat{b}_0 + \widehat{b}_\tau \tau_t\right) - \left(b_0^{(1)} + b_\tau^{(1)}\tau_t\right)}_{\text{regression-predicted slope}} \right] \right\} + \\ \Delta\gamma \left[\underbrace{\overline{y}_t}_{\text{actual level}} - \underbrace{\left(\widehat{b}_0 + \widehat{b}_\tau \tau_t\right)}_{\text{regression-predicted level}} \right] + \\ \overline{\epsilon}_{t+1}^{(n)}. \end{gathered} \tag{27}$$

Let's pause on the various terms for a moment, beginning from the quantities in curly brackets.

The first, $\left(\overline{y}_t - y_t^{(1)}\right)$, is just a reasonable proxy for the observed yield-curve slope at time t. Then the terms $\left(\widehat{b}_0 + \widehat{b}_\tau \tau_t\right)$ and $\left(b_0^{(1)} + b_\tau^{(1)} \tau_t\right)$ are the CPI-regression-based predictions of the average yield and the one-year yield, respectively. Therefore the difference $\left(\widehat{b}_0 + \widehat{b}_\tau \tau_t\right) - \left(b_0^{(1)} + b_\tau^{(1)} \tau_t\right)$ is just the regression-predicted slope. This means that the quantity in curly brackets is the difference between the actual slope and the regression-predicted slope.

Moving to the term in square brackets that multiplies $\Delta\gamma$, $\overline{y}_t$ is an obvious proxy for the level of the yield curve, and the quantity $\left(\widehat{b}_0 + \widehat{b}_\tau \tau_t\right)$ is, by definition, the regression-based prediction of the same quantity. The square bracket in the third line of Equation (27) therefore contains the difference between the actual yield-curve level and the level prediction by the regression. We call in the following the distances of level and slope from their local regression-predicted values the 'conditional' distances.

This simple rearrangement therefore shows that in the Cieslak and Povala restricted factor there is no single 'typical' slope or typical 'level' of the yield curve: there are instead conditional 'typical' levels and slopes of the yield curve associated with different values of the slow-moving inflation proxy, τ_t. Indeed, from Equation (27) we see that, if the long-term expected inflation, τ_t, is high, then the level of the yield curve should be high – as shown by the term $\left(\widehat{b}_0 + \widehat{b}_\tau \tau_t\right)$. Then, if we re-write $\left(\widehat{b}_0 + \widehat{b}_\tau \tau_t\right) - \left(b_0^{(1)} + b_\tau^{(1)} \tau_t\right)$ as $\left(\widehat{b}_0 - b_0^{(1)}\right) + \left(\widehat{b}_\tau - b_\tau^{(1)}\right) \tau_t$, and we remember that $\widehat{b}_\tau - b_\tau^{(1)}$ is negative, we see just as easily that the regression predicts that the slope should be low when the long-term inflation is high.

This allows us to understand very clearly the Cieslak and Povala return-predicting factor: we expected a high excess return (the return predicting factor is high)

- when the actual slope of the yield curve is higher than what it 'should be' (given the regression prediction);
- when the actual level of the yield curve is higher than what it 'should be' (again, given the regression prediction).

This interpretation gives an interesting twist to the slope (and level) interpretation and significance that has been the traditional wisdom since the study by Fama and Bliss (1987). Taken together, the success of the Cieslak and Povala predictions, and the interpretation just given, show that whether one can expect to make money by engaging in the 'carry' trade does not depend on whether the slope of the yield curve is higher or lower than a fixed reference level. Rather it depends on whether it is higher or lower than the long-term-inflation-*predicted slope*. And the same applies to the level – which by itself, as all statistical studies to date had shown, is either statistically insignificant or, when marginally significant, a poor contributor to the explanation of excess returns.

This simple decomposition explain very clearly why the Cieslak and Povala return-predicting factor can perform so much better than the slope or the Cochrane-Piazzesi factor: because it contains information about the *conditional* slope and level that the yield-curve-based approaches do not have. It also gives us a simple way to answer the question: are the level or the slope 'deviations' more important in explaining excess returns? We answer this question in the next section.

5 Attribution of the Predictability

Equation (27) is an exact restatement of the Cieslak-Povala return-predictive factor. As such, it affords exactly the same R^2 as the original factor. The decomposition below however allows for a very simple attribution study.

We first perform a bivariate regression of excess returns for investment horizons from 2 to 15 years as in Equation (27).

We then perform two univariate regressions of the same excess returns, one against the factor that multiplies γ_3, the other against the factor that multiplies $\Delta\gamma$. (After correcting for overlapping returns, all the coefficients were significant.)

We report in Fig (2) the R^2 statistics obtained for the bivariate regression (curve labelled 'R^2 Full') and for the two univariate regressions (curve labelled 'R^2 Slope' and 'R^2 Level'). The curve labelled 'Sum Single R^2' shows the sum of the R^2 statistics from the individual regressions.

We now see the importance of looking at the *conditional* level and slope. For instance, it has been well known since the late 1980s that the absolute (unconditional) slope has a significant explanatory power for excess returns. In this respect, therefore, the present decomposition merely qualifies and makes this well-known result more precise. However, when it comes to the level, conventional wisdom held that the level factor explains very little of excess returns. This is no longer true if we look at the conditional level, whose stand-alone explanatory power is similar to the stand-alone explanatory power of the conditional slope, and actually dominates at the short end.

In general, we find that both the deviations of the instantaneous levels and slopes from their respective conditional levels play an important explanatory role. This makes intuitive sense. The actual and predicted slope are strongly co-integrated, and so are the actual and predicted yield-curve levels. So, when 'rates are higher than they should be' they will, on average, come down, and the 'carry' strategy will make money; and when the yield curve is 'steeper than it should be', it will on average flatten, and this will also make money for the fund-short, invest-long carry strategy.

Finally, we observe that the sum of the R^2 statistics from the two separate return predictive factors falls well short of adding up to the full explanatory power of the bivariate regression. This suggests that the *relative* values of the two conditional distances is important. Our decomposition still does not make it obvious why this should be the case.

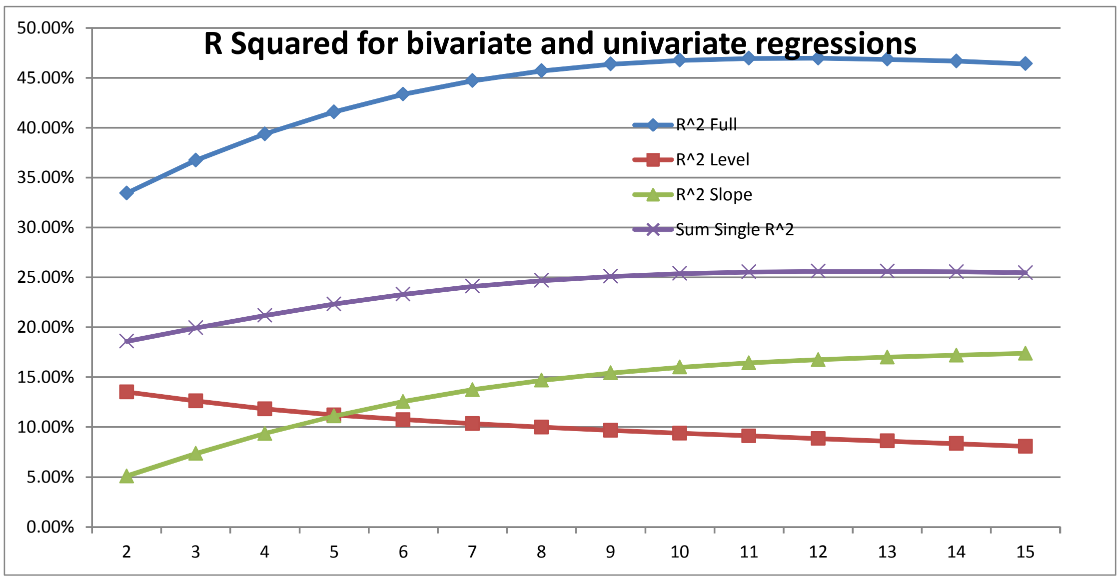

Figure 2: The R^2 statistics obtained for the bivariate regression (curve labelled 'R^2 Full') and for the two univariate regressions (curve labelled 'R^2 Slope' and 'R^2 Level') described in the text. The curve labelled 'Sum Single R^2' shows the sum of the R^2 statistics from the individual regressions.

6 Conclusions

We have presented a simple reformulation and decomposition of the restricted Cieslak and Povala (2010) return-predicting factor. The new factor retains by construction exactly the same explanatory power as the original one, but affords an alternative and attractive interpretation: what matters for excess return, it shows, it not whether the level or the slope of the yield curve are high or low in absolute, but whether they are higher or lower than where the should be with respect to where a long-term inflation proxy suggests they should be.

The decomposition also allows a clear attribution, which shows that both the conditional deviations are important predictors of excess returns.

7 References

Adrian T, Crump R K, Moench E, (2013), *Pricing the Term Structure with Linear Regressions*, Federal Reseve Bank of New York, Staff Report No 340, August 2008, Revised 2013, available at

http://www.newyorkfed.org/research/staff_reports/sr340.pdf, last accessed 15th October, 2014

Cieslak A, Povala P, (2010), *Understanding Bond Risk Premia*, working paper, Kellog School of Management, Wisconsin University and University of Lugano, available at

https://www.gsb.stanford.edu/sites/default/files/documents/fin_01_11_CieslakAnna.pdf, last accessed 5th May 2015

Campbell J, Shiller, R, (1991), *Yield Spreads and Interest Rate Movements: A Bird's Eye View*, Review of Economic Studies, 58, 495-514

Fama E, Bliss R, (1987), *The Information in Long-Maturity Forward Rates,* American Economic Review, 77, 680-692

Hellerstein R, (2011), *Global Bond Risk Premiums*, Federal Reserve Bank of New York Staff Report no 499, June 2011

Ludvigson S C and NG S, (2009), *Macro Factors in Bond Risk Premia*, Review of Financial Studies, 22, 5027-5067, available at http://rfs.oxfordjournals.org/content/22/12/5027.abstract last accessed 6th May, 2015