Jun Pan
junpan@mit.edu

253-3083
E62-624

Class 22: Fixed Income, Term Structure Models

This Version: November 23, 2016[1]

1 Term Structure Models

- **The Challenge from the Data:** In the fixed income market, term structure models are used to model interest rates. The challenge from the data has two dimensions. First, it should take into account of how the interest rates move over time. Second, for a given time, it should be able to model the yield curve, also called the term structure of interest rates. Figure 1 is a good summary of these two challenges from the data: a good term structure model should be able to capture the dynamic variations of the level of interest rates and the shape of the yield curve.

  These two demands from the data are very similar to those in the equity market. A good model for stock market returns should be able to take into account of how stock returns vary over time, as well as how, for a give time, the cross-section of stocks are priced in relation to one another. In the equity market, an i.i.d. model for stock returns is a reasonable approximation. As such, the dynamics for stock returns are really simple: constant expected return μ, constant volatility σ, and unpredictable random shocks ϵ_{t+1}. Cross-sectionally, the expected stock returns are linked to one another through their exposures (i.e., betas) to risk factors in a model such as the CAPM. As such, the CAPM model is a static model with constant expected returns and constant beta.

  The need for a dynamic model shows up when we investigated the time-varying volatility in our volatility class and stochastic volatility in our options class. Here in this class, we have a chance to take a closer look at these dynamic models.

- **Term Structure Models, Historical Development:** Term structure models were developed in the mid-1970s by Cox, Ingersoll and Ross (1985) and Vasicek (1997). You

[1] A small correction of Figure 3, which was missing a portion of the Fed target rate.

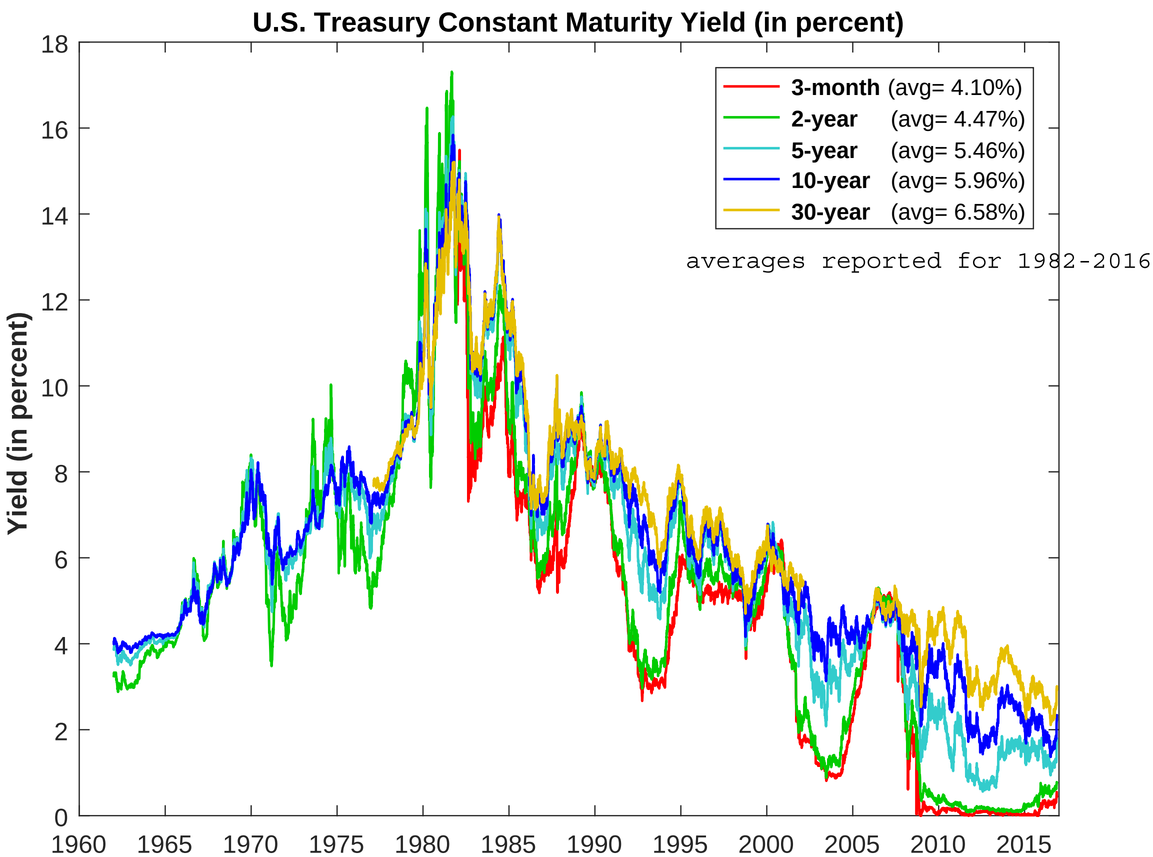

Figure 1: Time-Series of Treasury Constant Maturity Yields.

might notice that the CIR paper was published in 1985, but it was really a product of the mid-1970s. These term structure models were a continuation of the work done by Black, Merton, and Scholes, who popularized the application of continuous-time models in Finance. Like the Black-Scholes model before them, these term-structure models use the stochastic processes studied by mathematicians and physicists. For example, the CIR model builds on the Feller process and the Vasicek model builds on the Ornstein-Uhlenbeck process. In both cases, the starting point is the instantaneous short-rate r_t, which is modeled by a stochastic process (OU or Feller). The entire yield curve is priced using the dynamics of this one short rate. As such, the CIR and Vasicek models are one-factor short-rate models.

The second wave of term structure models came in the 1990s. When I entered the Stanford PhD program in 1995, I was just in time to catch the excitement surrounding term structure models. Relative to the original models of CIR and Vasicek, the effort of the new generation of term structure models is to be empirically relevant. From the work of Litterman and Scheinkman (1991), it became clear that a one-factor model will not be able to capture the entire shape of the yield curve. Unlike the stock market, where you can dismiss the risk uncaptured by the model as idiosyncratic risk, we do not have the luxury of dismissing common risk factors in the fixed income market (e.g., the slope factor).

These multifactor models quickly found their way into the "real" world. It is my understanding that each investment bank has its own proprietary term-structure model. And I was told by some practitioners that the industry has the best and most sophisticated term structure models. And they use these models to manage and hedge interest-rate risk (level, slope, convexity, volatility, etc) as well as to price interest-rate derivatives and other rate-sensitive instruments such as MBS. Looking back, I can now understand why during the mid-1990s, the Wall Street hired so many physicists and mathematicians. Most of my classmates in Physics ended up on Wall Street. I can also understand the sudden demand for more sophisticated term structure models in the 1990s. The fixed income desks were very profitable and the range and trading volume of their fixed income products were also expanding very rapidly during that time.

By now, the excitement surrounding term structure models has all but fizzled out. As a PhD student at Stanford, I spent much more time learning and working on term-structure models than anything else I did there. Since coming to Sloan in 2000, I have not made much use of that part of my training. Nevertheless, I am very grateful to my advisers at Stanford for having trained me in this area. As I wrote earlier in my lecture notes, not everything we do in life is of practical use. Still, they are useful and meaningful in our growth process.

For our class, however, I don't want to emphasize too much on the modeling part, because it takes quite a bit of mathematical skills. Instead, I would like to use the term structure models as a way for us to understand conceptually how the various parts of the yield curve are connected through a pricing model and the role of the risk factors in generating the pricing results.

- **Bond Pricing in Continuous-Time:** Let r_t be the time-t instantaneous short rate. Let today be time 0, and let P_0 be the present value of a dollar to be paid in T years. Discounting this future dollar all the way from T to today using the short rate, we have:

$$P_0 = E\left(e^{-\int_0^T r_t\, dt}\right) \tag{1}$$

  Let me explain this expression in sequence:

  - The reason why we need to do $\int_0^T r_t\, dt$ is because we have to add up all of the future short rates along the path from 0 to T. Take the extreme example of a constant short rate r. We have $\int_0^T r_t\, dt = r\, T$ and $P_0 = e^{-r\,T}$.
  - We put $\int_0^T r_t\, dt$ onto $e^{-\int_0^T r_t\, dt}$ because the rates are continuously compounded. (You will find that working with e^x and $\ln(x)$ typically gives us a lot of tractability in Finance.)
  - Later on, we will see how r_t is going to be driven by a random risk factor. Because of this, there could be many paths of r_t, depending on the random outcomes of the risk factor. And the present value of a future dollar to be paid in year T is an expectation, $E(\cdot)$, taken over all potential random paths of r_t with t running from 0 to T.

- **Relating back to Option Pricing:** The calculation in Equation (1) is similar to the calculation of $E^Q\left(e^{-rT}\,(K - S_T)\,\mathbf{1}_{S_T<K}\right)$ in option pricing. The difference is that we do not have to deal with the random variation in S_T. But we have to deal with the random variation in the riskfree r, which turns out to be more difficult to deal with.

  Instead of fixing a maturity date for this interest rate r (as in yields to maturity), we choose to work with the "short rate" so that this one rate can be used to discount future cashflows over any horizon. We just need to add them up via $\int_0^T r_t\, dt$.

  A by-product of this modeling choice is that we now have to keep track of the entire path of r_t from 0 to T in order to calculate $\int_0^T r_t\, dt$. Remember that when you performed option pricing via simulation in your Assignment 3, you didn't have to keep track the path of S_T from 0 to T. You only needed to know the values of S_T. So in order to have one million scenarios of S_T, you needed to simulate one million random variables.

To price bonds, however, you need to simulate the entire path of r_t from 0 to T. Suppose we decide to discretize the time interval from 0 to T into monthly intervals, then pricing a one-year bond with one million scenarios would involve simulating 12 $\times$ one million random variables; pricing a 10-year bond would involve simulation 120 $\times$ one million random variables. In short, pricing bond is generally more involving than pricing equity options and pricing bond derivatives would be even more challenging. That is why models with closed-form solutions are very useful. Otherwise, we will have to resort to either simulations or solving partial differential equations.

Also notice that to be precise, I should take the expectation in Equation (1) under the risk-neutral measure. For this class, however, let me not make a distinction between the two, just to keep things simple.

- **The Vasicek Model:** In the Vasicek model, the short rate r_t follows

$$dr_t = \kappa\,(\bar{r} - r_t)\;dt + \sigma\,dB_t\,, \tag{2}$$

where, as in the Black-Scholes model, $\sigma\,dB_t$ is the diffusion component with B as a Brownian motion. This model has three parameters:

  – $\bar{r}$: The long-run mean of the interest rate, $\bar{r} = E(r_t)$.

  – κ: The rate of mean reversion. When r_t is above its long-run mean $\bar{r}$, $\bar{r} - r_t$ is negative, exerting a negative pull on r_t to make it closer to $\bar{r}$. A larger κ amplifies this pull of mean reversion and a smaller κ dampens it. Conversely, when r_t is below its long-run mean $\bar{r}$, $\bar{r} - r_t$ is positive, exerting a positive pull on r_t, again to make it closer to its long-run mean $\bar{r}$.

  – σ: controls the volatility of the interest rate.

- **Bond Pricing under Vasicek:** Bond pricing under the Vasicek model turns out of be very simple. Let today be time t and let r_t be today's short rate, then the time-t value of a dollar to be paid T years later at time $t+T$ is

$$P_t = e^{A+B\,r_t}\,,$$

where

$$B = \frac{e^{-\kappa T} - 1}{\kappa}$$
$$A = \bar{r}\left(\frac{1 - e^{-\kappa T}}{\kappa} - T\right) + \frac{\sigma^2}{2\kappa^2}\left(\frac{1 - e^{-2\kappa T}}{2\kappa} - 2\frac{1 - e^{-\kappa T}}{\kappa} + T\right)$$

2 Calibrating the Model to the Data

- **The Vasicek Model:** As usual, we work with models in order to understand, at a conceptual level, the key drivers in the pricing of a security. Applying the model to the data, we further understand quantitatively how well the model works and what's missing in the model.

  For a one-factor model such as the Vasicek model, we know its limitation even before applying it to the data. In the fixed income market, the level of the interest rates is the number one risk factor in terms of its importance, but it is not the only risk factor.

  In Assignment 4, I ask you to work with a discrete-time version of the Vasicek model by first estimating the model parameters, $\bar{r}$, κ, and σ, using the time-series data of 3-month Tbill rates. Basically, I am asking you to calibrate the model only to the time-series information of the short-end of the yield curve, without allowing you to take into account of the information contained in the other parts of the yield curve. Then I ask you to price the entire yield curve. Not surprising, you will find that the calibrated model does not work very well to accommodate the different shapes of the yield curve.

  An alternative approach is to calibrate the model using the yield curve. For example, on any given day, we estimate the model parameters, $\bar{r}$, κ, and σ, so that the pricing errors between the model yields and the market yields are minimized. By doing so, the model will do a much better job in matching the market observed yield curve, but it will miss the time-series information. Moreover, you will have one set of parameters per day, which is inconsistent with the assumption that these parameters are constant.

  The better solution is to introduce more factors to the model. For example, instead of forcing the long-run mean $\bar{r}$ to be a constant, we can allow it to vary over time by modeling it as a stochastic process. Instead of forcing the volatility coefficient σ to be a constant, we can allow it to vary as another stochastic process. There, you have a three-factor model. The pricing will be more complicated and so will be the estimation. Working with these multi-factor models requires some patience, perseverance, and the love for the subject matter. Indeed, it is not for everybody.

- **Curve Fitting:** On a topic related to model calibration is yield curve fitting. In this approach, there is no consideration along the time-series dimension. The zero rate $r(\tau)$ of maturity τ is modeled as a parametric function, which is then used to price all market traded coupon-bearing bonds. On any given day, the parameters in that parametric function will be chosen so that the pricing errors between the model yields and the market yields are minimized. This exercise of yield curve fitting is repeated daily and the model parameters are updated daily as well.

  Figure 2 plots the yield curve during the depth of the 2008 crisis. It uses the Svensson model for curve fitting. The parameters in the Svensson model are first optimized so that the model can price all of the market-traded bonds on December 11, 2008 with minimum pricing errors. Using these parameters, the black line is the corresponding par coupon curve. The blue or purple dots are the market yields for the market-traded bonds. For each dot, there is a companion red "+", which is the model yield for the corresponding bond. In a fast decreasing interest rate environment such as December 2008, most of the existing bonds are premium bonds. As we discussed earlier, with an upward sloping term structure, the yields of these bonds are lower than the corresponding par-coupon yields of the same maturity. That is why most of the red "+"s are below the par coupon curve. If there are many discount bonds being traded at the time, then you will see some red "+"s above the par coupon curve.

  This curve fitting exercise is useful in connecting the yields of different maturities through a parametric function of zero rates. For example, there is quite a bit of overlap in discount rates between a ten-year yield and a ten-year minus one-month yield. The presence of a parametric function of zero rates acknowledges the overlap (ten years minus one month) and the pricing difference between these two yields will be sensitive only to the one-month gap. But the usefulness of a curve fitting exercise stops at this level. If you would like to use a model to help you with derivatives pricing on the yield curve (e.g., Bond options, swaptions, caps/floors, etc), a curve-fitting model will not be helpful at all because it does not take into consideration of how yields vary over time. For derivatives pricing on the yield curve, you need to use dynamic models. The usual approach is to use multi-factor versions of CIR or Vasicek models. Affine models are examples of these multi-factor versions of CIR and Vasicek.

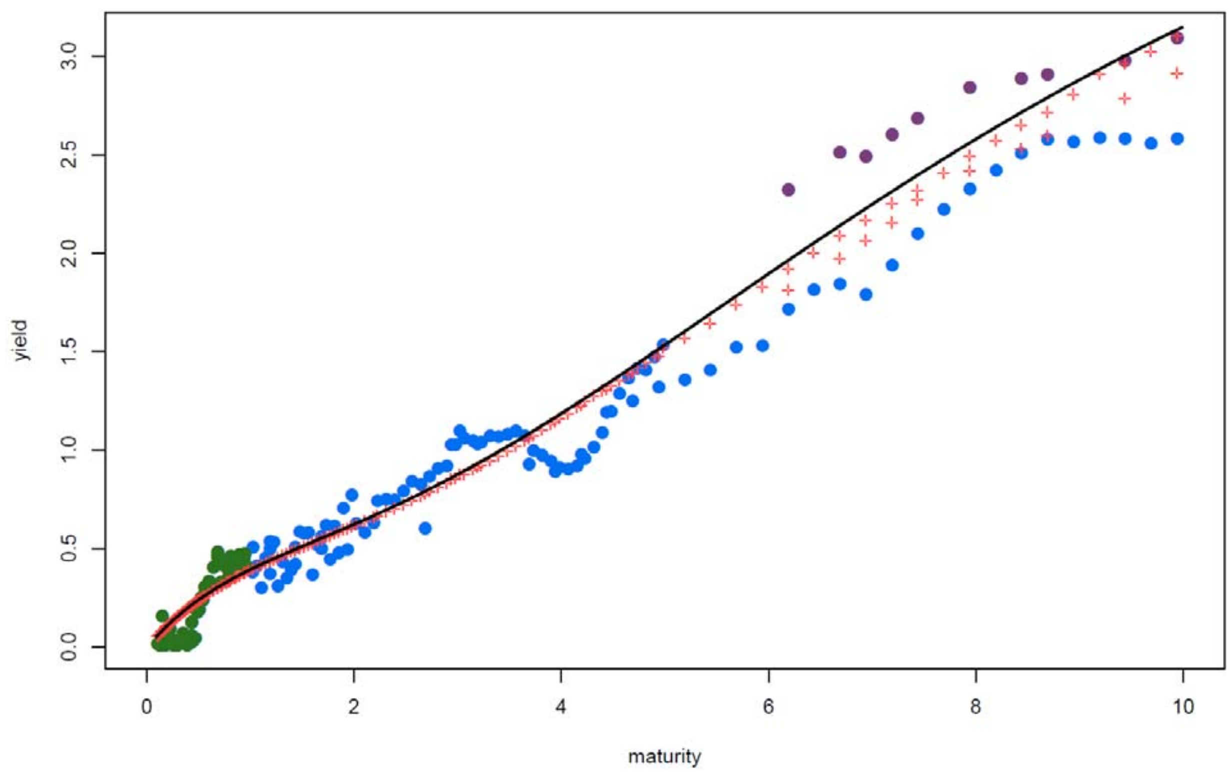

Figure 2: Treasury Yield Curve on December 11, 2008.

3 Relative Value Trading with a Term Structure Model

In March 2011, Chifu Huang (a former MIT Sloan Finance professor) came to Prof. Merton's class to give a guest lecture. I found his talk to be very informative and the following is based on one portion of his talk.

- **How to Use a Term-Structure Model to Identify Trading Opportunity:** Relative value trading in the fixed income market does not make a judgment on the level of interest rates or the slope of the curve. It assumes that a few points on the yield curve are always fair. For example, the time-series data on the 10yr, 2yr, and 1-month rates can be used to estimate a three-factor term structure model.

  Recall that in the Vasicek model, the short rate is the only risk factor (i.e., state variable). That is why in your Assignment 4, I ask you to estimate the model using only the 3M Tbill rates. With a three-factor model, we have three risk factors (i.e., state variables) and we need three points on the yield curve to help us estimate the model. Intuitively, the 10yr gives us information about the level of long-term interest rates; the 2yr together with the 10yr informs us about the slope of the curve; and the 1-month Tbill rate captures the short-term interest rate (including expectations on monetary policy in the near term).

  Once you have the model estimated by the time-series data (which is a non-trivial task if you would like to do it properly), this model is going to give you predictions about the level of interest rates across the entire yield curve. You can then compare the model price with the market price to judge for yourself whether or not a market price is cheap or expensive. Once you convince yourself that your model helps you pick up a trading opportunity, you would structure a trade around it. You can buy cheap maturities and sell expensive maturities, and, at the same time, hedge your portfolio so that it is insensitive to the changes of the level or the slope of the yield curve.

  The main judgment call is to understand why your model identifies some maturities as cheap or expensive. If it is due to institutional reasons (which does not show up in your model but does show up in the data), then you can make judgment as to whether or not such institutional reasons will dissipate over time (and how fast).

- **An Example:**

  One example was given by Chifu. In August 1998, Russian defaulted on its local currency debt, and the effect lingered well into September and was later known as the LTCM crisis. As shown in Figure 3, in September 1998, bond markets rallied in

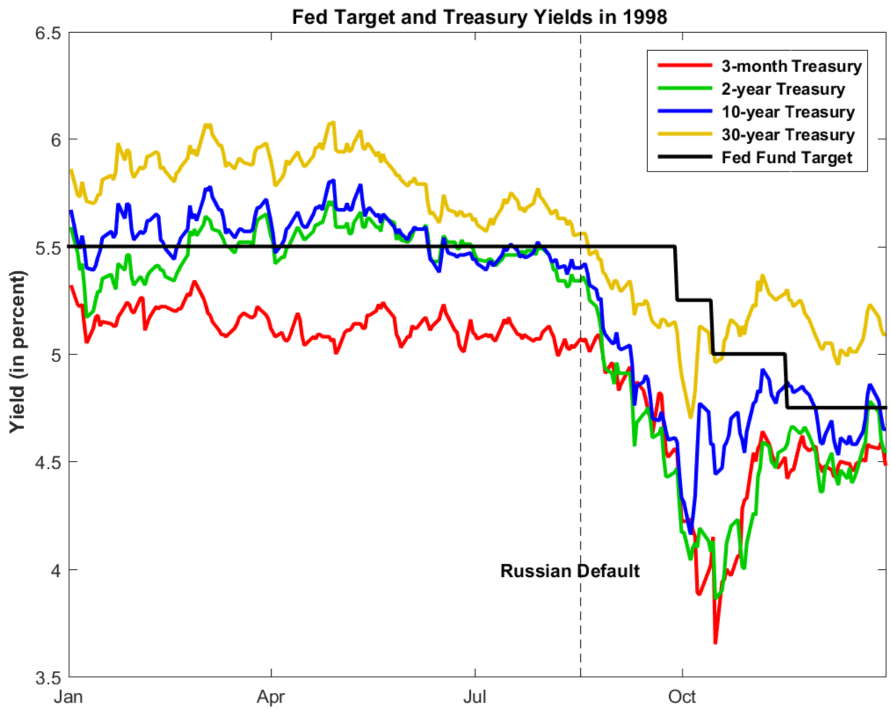

Figure 3: Fed Target and Treasury Yields in 1998.

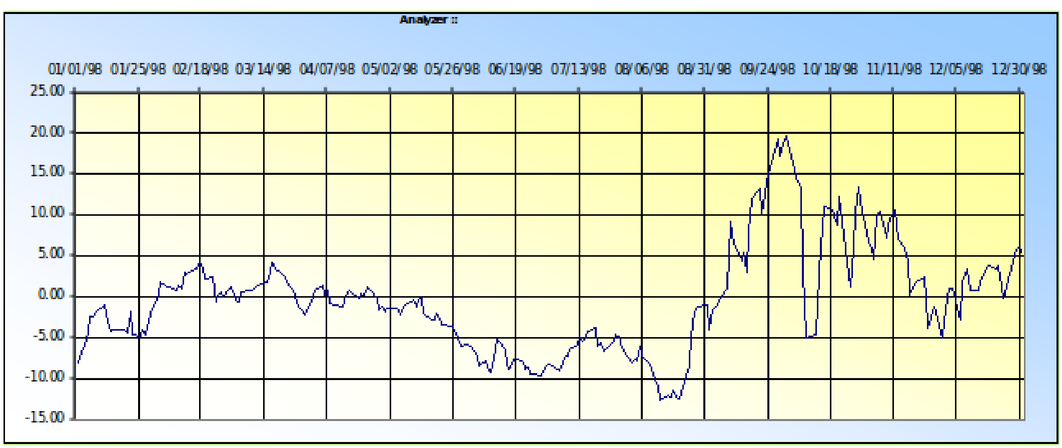

Figure 4: Cheapness and Richness of US 30-Year Swap Rate Based on a Two-Factor Model.

anticipation of a rate cut. On September 29, the Fed cut the fed funds target rate by 25 bps.

Figure 4 is a slide presented by Chifu in his talk. In September 1998, his two-factor model picks up a trading opportunity regarding the 30yr bond. According to the model, the market price for the 30yr bond is cheap relative to the model price. The deviation between the data and the model was at the range of 10 to 20 bps. The 30-year rate was around 5.5% at that time, implying a modified duration of about 15 years. So a 10 bps price deviation in 30yr would translate to 10 bps $\times$ 15 = 150 bps in bond return. And a 20 bps deviation will translate to 3% in bond return.

So what are the reasons for this cheapening of 30yr? It is because residing over the 30yr region are pension funds and life insurance companies who are either inactive "portfolio rebalancers" or rate-targeted buyers. As a result, the rally that happened in the rest of the yield curve didn't find its way to the 30yr region. There is a lag in how information (regarding an impending rate cut) gets transmitted to this region. As you can see from Figure 3 and 4, it was only after the Fed's rate cut on September 28 when the 30yr yield was brought back in alignment with the rest of the yield curve.