The term-structure of systematic risk

Nuno Clara

Abstract

Using information embedded in option prices, we uncover the existence of a non-trivial term-structure of betas for individual stocks and portfolios. The slope of this term-structure is a priced factor in the cross-section of returns and spikes following relevant macroeconomic and firm-specific events. The slope of the term-structure of systematic risk is mainly driven by the slope of the term-structure of variance swaps. An investment model with uncertainty shocks in the spirit of Bloom (2009) can quantitatively explain the relation between the time-series and cross-sectional dynamics of the term-structure.

1 Introduction

Betas are a central concept in finance. Academics and practitioners tend to think of risk-return tradeoff in terms of beta - stocks with higher exposure to a common risk factor (systematic risk) should earn higher average returns. The majority of the risk-return tradeoff literature has focused on the cross-sectional properties of betas. However, in addition to studying the cross-section of systematic risk, exploring fluctuations of, and the information contained in, the term-structure of systematic risk can provide many insights into the riskiness of a particular stock or portfolio and shape the way we think about traditional macro-finance models.

While there are many studies in the literature focusing on the term structures of nominal and real interest rates and, more recently, the term-structure of equity, there has been almost no research regarding the term-structure of systematic risk, despite the evidence that betas are not equal across different horizons.[1] In this paper, we make use of options data to estimate forward looking capital asset pricing model (CAPM) betas for several portfolios across different time horizons and jointly study the cross-section and the term-structure of systematic risk. The betas we estimate are "spot" betas, meaning that for any security at any given point in time we have forward looking estimates of systematic risk with different maturities.

The question of whether systematic risk has a term-structure only makes sense at the portfolio or individual stock level. By definition the term-structure of the market is flat - the market beta is always one regardless of the horizon - and therefore it is fairly uninteresting. We uncover the existence of a time-varying, non-flat term-structure of systematic risk for both individual stocks and portfolios. We find that what really matters for pricing the cross-section of returns is the slope of the term-structure, i.e., the difference between the long-term beta and the short-term beta. After controlling for the level of the term-structure (either the short-term or the long-term beta), portfolios with a negative slope earn higher excess returns than portfolios with a positive slope. These portfolio returns cannot be explained by the standard Fama and French (1993) three factor model. This implies that investors should care about how systematic risk evolves over time and therefore becomes relevant to economically document the properties of this term-structure.

First, the slope of the term-structure is a priced factor in the cross-section, significantly improving the fit of the Fama and French (1993) model when pricing the standard 25 book-to-market and size sorted portfolios. Economically, this pricing power can be understood by

[1] Ferson and Harvey (1991) provided some evidence on betas not being constant across different horizons. van Binsbergen, Hueskes, Koijen, and Vrugt (2013), Van Binsbergen, Brandt, and Koijen (2012) study the term-structure of equity.

studying the time-series properties of the slope of the term-structure. It spikes on relevant macroeconomic and firm-specific events. For example, during bad times such as the Great Recession, the slope of the term-structure of systematic risk of value (growth) portfolios spikes down (up). This means that the short-term beta for value stocks spikes up vis-a-vis its long-term beta. Value firms have many assets in place; thus, in recessions, they are much riskier than growth firms. Market participants price this riskiness: the portfolio's term-structure becomes very negatively sloped as investors recognize the higher risk of these stocks in the short-term but expect it to diminish as the economy recovers.

Second, we decompose the slope of the term-structure of betas into its main components. Most of the variation in the term-structure comes from the short-end of the curve - the short-term beta. This effect can be decomposed further into a variance swap and a correlation effect. Most of the heterogeneity in the term-structure across portfolios comes from the variance swap. This finding has important theoretical implications. The literatures on the cross-section of returns and on the term-structure of variance have evolved somewhat separately. Important references in the literature on the cross-section of returns (e.g., Zhang (2005), Gomes, Kogan, and Zhang (2003), Kiku (2006), Lettau and Wachter (2007), among others), do address the term-structure of variance and, analogously, literature on the term-structure of variance swaps (e.g., Dew-Becker, Giglio, Le, and Rodriguez (2015), Aït-Sahalia, Karaman, and Mancini (2014), Egloff, Leippold, and Wu (2010)) is silent regarding the cross-section of returns. Our paper makes a contribution to these two strands of the literature by jointly modelling the term-structure of variance swaps and the cross-section of returns.

We build a investment model similar to that of Lin and Zhang (2013) but augment it with a time varying second moment in the spirit of Bloom (2009). Our model has heterogeneous firms and time-varying uncertainty which allows us to link the cross-section of returns with the term-structure of systematic risk. Firms face both economy-wide and individual specific shocks and decide whether to invest or to distribute dividends. The model allows us to quantitatively match the observed empirical dynamics: the superiority of the option implied CAPM and the dynamics of the term-structure of systematic risk. When volatility spikes are accompanied by productivity decreases, value stocks are burdened with more unproductive capital. Firms want to cut down capital, and the presence of capital adjustment costs makes them riskier. Therefore, value stocks have on average a higher variance swap than growth stocks and consequently a higher option implied beta. The option implied beta therefore lines up extremely well with book-to-market sorted portfolios. Furthermore, at the inception of a volatility shock, the implied variance of value stocks spikes more than the implied variance of growth stocks and this heterogeneous effect accounts for the dynamics of the term-structure of option implied betas. This paper is the first to estimate and analyze the behavior of the term-structure of systematic risk both in the cross-section and the time-series. Our results have several important implications for empirical and theoretical asset pricing.

The remainder of the paper is organized as follows. Section 2 describes the data and methodology. Section 3 documents the main empirical results of the paper - uncovering the existence of a term-structure of systematic risk and studying its economic content. In section 4, we build an investment model provides an economic explanation for the results in the empirical section. Section 5 concludes.

2 Data and Methodology

2.1 Data

Our study is based on data on the S&P 500 Index and its constituents between January 4, 1996 and August 29, 2014, or a total of 3,978 trading days. We obtain daily price data from the Center for Research in Security Prices (CRSP) and option data from OptionMetrics. To obtain the S&P 500 constituents on any given day, we use the CRSP S&P 500 constituents file. This file has information on the addition and deletion date of each stock by PERMNO. Our CRSP sample includes 976 unique stocks which we match to the OptionMetrics SECIDs. We are able to match a total of 891 stocks which form our base sample.

The data on equity and index options are taken from the IvyDB OptionMetrics Surface file. It provides a smoothed volatility surface for a range of maturities and strikes. The use of the surface file has been a standard in this literature (e.g. Buss and Vilkov (2012), An, Ang, Bali, and Cakici (2014), among others) and it has the advantage of making this study more easily replicable by other researchers. For each security and at each point in time the volatility surface file stores implied volatilities and strike prices of calls and puts for several maturities (30, 60, 91, 122, 152, 182, 273, 365, 547 and 730 days). We only keep out-of-the money calls and puts since these are more liquid instruments, and option dates that match the underlying CRSP trading days.

2.2 Historical Betas, Option Implied Betas and Realized Betas

To estimate option-implied betas, we follow closely to the methodology proposed by Buss and Vilkov (2012). Their proposed forward-looking beta follows the standard definition of beta, i.e. it is the ratio of stock-to-market covariances to the market variance, but replacing the historical moments with their risk-neutral counterparts. Thus, at time t the option implied market beta of stock i is given by:

$$\beta^Q_{iM,t,\tau} = \frac{\sigma^Q_{i,t,\tau} \sum_{j=1}^{N} w_j \sigma^Q_{j,t,\tau} \rho^Q_{ij,t}}{(\sigma^Q_{M,t,\tau})^2} \tag{1}$$

where $\sigma^Q_{i,t,\tau}$ denotes the volatility of stock or portfolio i at time t under the risk-neutral measure Q, $\rho^Q_{ij,t}$ denotes the pairwise stock correlations, and w_j are the weights of security j on the market portfolio. Similarly, $\sigma^Q_{M,t,\tau}$ denotes the volatility of the overall market under the same Q measure. Note that the we add an extra subscript τ to our betas to denote the maturity of the options with which the beta was estimated.

Before outlining the means of obtaining the elements of equation (1), two precautionary notes are, however, appropriate. First, the risk-neutral betas are estimated using risk-neutral variances and correlations and these measures are usually different from their objective counterparts due to the existence of a variance and correlation risk premium. This has been thoroughly documented by Carr and Wu (2009), Han and Zhou (2012) and Driessen, Maenhout, and Vilkov (2009). Second, there are no traded options on a basket on any combination of two stocks; therefore, one needs to make an appropriate parametric choice for modelling option implied correlations. This means that $\beta^Q_{iM,t}$ might be a biased estimator of $\beta^P_{iM,t}$. Buss and Vilkov (2012) extensively discussed these issues and specified a simple parametric form for implied correlations that is consistent with several empirical facts despite the caveats outlined above.[2] We make use of the same parametric form, which we describe below.

To study whether option implied betas across several maturities successfully capture the risk-return relationship and to assess their realized beta predictive power vis-a-vis historical betas, we compute: (i) historical betas in the standard way by regressing the asset (excess) return on the market (excess) return and (ii) realized betas using the same methodology as Andersen, Bollerslev, Diebold, and Wu (2006), i.e., we use daily log-returns to estimate the realized beta between $t^\star$ and $T^\star$:

$$\beta^R_{iM,t} = \frac{\sum_{t^\star=1}^{T^\star} r_{i,t^\star} r_{M,t^\star}}{\sum_{t^\star=1}^{T^\star} r^2_{M,t^\star}} \tag{2}$$

where $r_{i,t^\star}$ and $r_{M,t^\star}$ are the log excess returns of stock i and the market at time $t^\star$, respectively and $T^\star$ is the number of days on the period under analysis.

2.3 Option Implied Moments

In order to approximate model-free implied volatility we closely follow the methodology of Demeterfi, Derman, Kamal, and Zou (1999) and Bakshi, Kapadia, and Madan (2003), who

[2]We refer the reader to their paper regarding the details on the assumed parametric form.

showed that if one owns a portfolio of options across all strikes inversely weighted by the squared strike, then one obtains a variance exposure that does not depend on the price. The moment free implied variance from period t to period τ (or the variance swap rate) can be approximated by:

$$(\sigma^Q_{i,t,\tau})^2 = \int_{S_i(t)}^{\infty} \frac{2\Big(1 - \log[\frac{K}{S_i(t)}]\Big)}{K^2} C_i(t,\tau,K)dK + \int_0^{S_i(t)} \frac{2\Big(1 - \log[\frac{K}{S_i(t)}]\Big)}{K^2} P_i(t,\tau,K)dK \quad (3)$$

where $P_i(t,\tau,K)$ and $C_i(t,\tau,K)$ are the prices of out-of-the-money puts and calls options with maturity τ and strike K, respectively. In practice, a continuum of option strikes does not exist, so that one has to interpolate and extrapolate strikes and implied volatilities for the remaining moneyness levels. We define a moneyness grid between 1/3 and 3 with 1,001 points. We then find the maximum and minimum points of the moneyness grid for which implied volatility is available. We interpolate implied volatility inside this interval by fitting smooth cubic splines. For moneyness levels above (below) the highest (lowest) available strike we use the implied volatility of the highest (lowest) strike.

The last element from equation (1) that we need to estimate is the risk-neutral correlation between each pair of stocks: $(\rho^Q_{ij,t})$. The presence of the correlation premia led Buss and Vilkov (2012) to estimate this risk-neutral correlation by letting:

$$\rho^Q_{ij,t} = \rho^P_{ij,t} - \alpha_t(1 - \rho^P_{ij,t}) \quad (4)$$

Combining equation (4) with the identifying restriction that equates the observed implied variance of the market index $(\sigma^Q_{M,t,\tau})^2$ with the calculated implied variance of a portfolio of all market index constituents $i = 1, ..., N$:

$$(\sigma^Q_{M,t,\tau})^2 = \sum_{i=1}^{N}\sum_{j=1}^{N} w_i w_j \sigma^Q_{i,t,\tau}\sigma^Q_{j,t,\tau}\rho^Q_{ij,t} \quad (5)$$

one can solve for α and consequently estimate $\rho^Q_{ij,t}$.

3 The term-structure of systematic risk

In this section, we document the key empirical facts on the term-structure of betas of individual stocks and portfolios. We start with a motivating example. Figure 1 plots the term-structure of betas for Apple Corporation at two different points in time.

[Figure 1 here]

It is worth mentioning that the term-structure is not stable - sometimes it slopes upward and occasionally, it slopes downward and even its level varies widely over time. For example, in September 1999, Apple had an upward sloping term-structure of systematic risk, implying that agents on the market were expecting the systematic risk of Apple to increase over time. The opposite pattern occurred in June of 2003, at which time Apple had a downward sloping term-structure of systematic risk.

This figure allow us to illustrate several important points about the term-structure of systematic risk. First, this figure illustrates that, due to the nature of the options data, we cannot estimate betas longer than a two-year horizon. Although we cannot take a stance on very long-term betas (such as a 10-year beta), empirically, the term-structure of systematic risk seems to flatten out after one year. In this paper, we will be focusing on the slope of the term-structure of betas, i.e., the difference between the long-term and the short-term beta, and the results are virtually unchanged if we use the 2-year or the 1-year option implied beta as the long-term beta.

Second, practitioners often shrink their long-term beta estimates. For example, Bloomberg has argued that beta estimates have a tendency to revert towards one, so to better forecast future betas, many analysts use: $\bar{\beta} = \frac{2}{3}\hat{\beta} + \frac{1}{3}$, where $\hat{\beta}$ is usually the CAPM beta, estimated using a standard Ordinary Least Squares (OLS) regression on historical data. Below we will show that option implied betas are better than the standard OLS estimates of beta in terms of explaining the cross-section of returns and predicting realized beta. What is, however, more striking is that these shrinkage methods might even worsen the estimation error, as they always imply a downward (upward) sloped term-structure of systematic risk if the beta is higher (lower) than one. By taking another look at the left-hand panel of Figure 1, it becomes clear that there is no way a shrinkage method would capture an upward sloped term-structure of systematic risk, for a company with a beta above one, as it is the case for Apple in September 1999.

Finally, recall that, mechanically, the slope of term-structure of systematic risk of the market is zero. Therefore, if, at any given point in time, a term-structure of a stock slopes upward, another term-structure must slope downward.

In the next two subsections, we study the cross-sectional and time-series properties of the term-structure of systematic risk. Before doing so, it is important to establish whether these option implied betas have good properties - if we have any hope of extracting any economic content from the term-structure of betas, it must be the case that our option implied betas correctly capture the risk-return tradeoff and have a good predictive power in respect to realized betas. Buss and Vilkov (2012) and Hollstein and Prokopczuk (2014) have shown that 1-year option implied betas line up with expected returns better than the standard CAPM betas.

We extend their analysis for option implied betas with different maturities. We examine the relationship between stock historical or option implied market betas and expected returns, as well as the ability of historical betas and implied betas to predict the realized beta.

To establish the risk-return relationship, we follow the standard procedure in the literature. At the end of each month t, we sort stocks into five portfolios according to their estimated historical and option implied beta and compute the value-weighted return of each portfolio over the next τ months, where τ is the option maturity underlying the estimation.[3] We then calculate a time-series average of the portfolios' betas and value-weighted returns to infer the beta-return relationship. We do this exercise for both historical betas and for option-implied betas across several maturities. Table 1 provides the results.

[Table 1 here]

Regardless of the maturity, the option implied betas provide a better spread across average realized returns in comparison with the historical betas. In fact, portfolios sorted on historical betas yield a risk-return relationship which is not linear and a spread in average realized returns of 2.77%. Conversely, portfolios sorted on option betas yield a relation which is more linear and a wider spread on realized returns (3.45% to 5.75%). This is in line with our conjecture that option implied betas successfully capture the risk-return relation.

Estimated betas should also have good predictive power in respect to realized betas. At the end of each month we compute historical betas and option implied betas for all maturities. We then sort all S&P 500 stocks into five portfolios and compute their value-weighted historical and option implied expected beta and compare them to the realized beta over the subsequent periods. Figure 2 plots the estimated beta against the realized beta for each of the five portfolios at each point in time.[4] At the top of each panel in figure 2 we also report the out-of-sample estimation accuracy of the beta estimates using the root mean squared error (RMSE) criterion.[5]

[Figure 2 here]

If the model fits well, we expect the points on the scatter plot to cluster around the 45-degree line as well as a low RMSE. It is clear that our option implied methodology outperforms the

[3] For historical betas, we compute value-weighted returns over the next month.

[4] We report the results for realized betas over the next 6 and 12 months, but the results hold for other maturities as well. The results are available from the author.

[5] The RMSE criterion is defined as the sum of the squared difference between the predicted beta and the realized beta for every portfolio, at any given month: RMSE $= \sqrt{\frac{1}{T-t}\sum_{t=1}^{T}(\beta_{t,T}^{R} - \hat{\beta})^2}$ where $\beta_{t,T}^{R}$ is the realized beta between period t and T and $\hat{\beta}$ is the ex-ante beta estimate (either historical or option implied).

historical beta methodology given that, for option implied betas, the points do indeed cluster around the diagonal of the plot, implying a good fit. Also, the RMSE is around 0.16, which is half the RMSE of the historical betas. From figure 2 becomes clear that the superiority of 1-year maturity option implied betas denoted by Buss and Vilkov (2012) also generalizes to other beta maturities.[6]

Having established the economic superiority of option implied betas in explaining the cross-section of returns, we now study its term-structure. To analyze the shapes of the term-structures of implied betas, we adopt a concept similar to the one used within the interest rate term-structure literature. Several authors have shown that the shape of the term-structure of interest rates can be captured via three simple factors: level, slope and curvature. For example, Diebold and Li (2006) fit, period-by-period, a Nelson-Siegel (NS) exponential components model to the entire term-structure of interest rates and show that the three parameters are enough to describe the term-structure of interest rates. In Appendix A we conduct a similar analysis and show that the term-structure of betas can also be successfully captured by those three factors.

Empirically, we define the level factor, γ_0, as the long-term beta ($\beta^Q_{iM,t,t+12m}$); the slope, γ_1, as the difference between the twelve-month implied beta and the one-month implied beta ($\beta^Q_{iM,t,t+12m} - \beta^Q_{iM,t,t+1m}$); and the curvature, γ_2, as the difference between twice the three-month implied beta and the sum of the one-month implied beta with the twelve-month implied beta ($2 \times \beta^Q_{iM,t,t+3m} - (\beta^Q_{iM,t,t+1m} + \beta^Q_{iM,t,t+12m})$). These definitions of level, slope and curvature are standard in the term-structure literature (Diebold and Li (2006)).

Given the Apple example above, it is interesting to ask whether betas on average have a flat term-structure. To look into this question, we compute period-by-period the term-structure of betas for all stocks in our sample and classify a term-structure to be upward (downward) sloping if $\gamma_1 > 0.1$ ($\gamma_1 < 0.1$) and classify it as positively (negatively) humped if the curvature factor γ_2 is above (below) 0.05 (-0.05). All shapes in between are classified as having no slope, no hump or both. The distribution of shapes of the term-structure is reported in Table 2.[7]

[Table 2 here]

[6] Figure 2 only shows the pattern for 6-month and 1-year option implied betas, but the result holds true for all other maturities. Results are available from the author.

[7] The cutoffs chosen are somewhat arbitrary. Different cutoffs would obviously imply a different distribution of shapes. For our purpose, it is sufficient to note that, on average there is a non-flat term structure. Given the large cross-section and time-series of our data, it is difficult to summarize the shapes of the term-structure in a single number as there is considerable heterogeneity in the data. Take, as an example, Panel B from figure 11 in Appendix A. There significant variation in the slope (curvature) of the term-structure of Apple Inc., which ranges from -0.5 to 0.5 (-0.75 to 0.6).

During approximately half of our panel sample term-structures are sloped (top and bottom rows from table 2). This means that on a relevant proportion of periods, average betas for short and long time-horizons are different. This difference may be very significant for capital budgeting and investment decisions. If systematic risk is different at different horizons, then a firm's cost of capital and expected return are different for different holding periods.[8]

3.1 Cross-section

In this subsection, we begin to address the question of whether there is any economic content in the slope of the term-structure. We begin by carefully exploring whether the slope of the term-structure matters for the cross-section of returns. We take the canonical asset-pricing approach of double-sorting portfolios on the slope and level of the term-structure of betas. We find that stocks with a negative term-structure slope outperform stocks with a positive slope and that the slope of the term-structure is a priced factor in the cross-section of returns, even after controlling for level of the term-structure and for the Fama and French (1993) three-factor model.

Suppose there is an investor with a one-year investment horizon whose main concern is his/her exposure to systematic risk.[9] This investor would choose stocks or portfolios based on their one-year option implied beta. However, given that the term-structure of systematic risk is usually not flat, this investor might also be concerned with how the systematic risk will evolve throughout the one-year holding period. In particular, if this investor rebalances his portfolio monthly, then investing in two stocks with the same one-year beta but with different short-term betas should yield different systematic risk exposures throughout the holding period.

Our hypothesis is that, despite the differences in short-term beta, these portfolios should yield the same expected return over the one-year holding period. We find that they do not.[10] To address this hypothesis, we build double sorted portfolios based on the level and slope of their term-structure. At the end of each month, we rank stocks according to their one-year option implied beta and the slope of their term-structure (the difference between the 1-year

[8]In terms of humps, the distribution is slightly more even. A third of the sample shows no hump and the remaining two-thirds show either a positive hump or a negative hump (with a 50% chance conditional on having hump). In this paper, we will focus on the economics of the slope of the term-structure and leave a more careful analysis of the complete shape of the term-structure as a suggestion for future research.

[9]This simple setup could be motivated through a static model wherein agents have mean-variance preferences and the same beliefs about the underlying return distribution, but different levels of risk-aversion. All agents would pick portfolios on the security market line, but more risk-loving agents would choose portfolios with higher beta.

[10]This is akin to the failure of the expectations hypothesis of interest rates.

beta and the 1-month beta). We follow a procedure similar to the one developed by Jegadeesh and Titman (1993), and hold the portfolios for one year. Basically, we select a portfolio based on its end-of-month 1-year beta and slope and hold it for 12 months. Therefore, each month, we close the position initiated 12 months prior and open a new position, implying that on any given month we revise the weights on 1/12 of the securities in the entire portfolio. Table 3 reports the excess returns of this strategy.

[Table 3 here]

Table 3 shows that portfolios with higher beta earn higher expected returns (as the CAPM would imply), and stocks with negative slope, i.e., for which the short-term beta is higher than the long-term beta, also earn higher expected returns after controlling for the level of beta.[11]

This implies that an investor with a fixed-horizon should not only be concerned about the expected level of systematic risk over his investment horizon but also about how systematic risk is expected to change. Suppose our investor has a high degree of risk aversion and therefore picks stocks with low-beta and holds the stocks from the first row of Table 3. He would earn higher expected returns on the portfolios in which systematic risk is expected to decrease. This is a very robust result. Below, as we look at portfolios sorted on characteristics (industry, size and book-to-market) and we will see the same result. As an example, value stocks term-structure is very negatively sloped during recessions, when marginal utility is high, which is exactly the time at which this portfolio has very high expected returns going forward.

The second panel of Table 3 reports the time-series α's of these portfolios in relation to the Fama and French (1993) three-factor model. The vast majority are positive and economically significant. At the bottom of the table we report the Gibbons, Ross, and Shanken (1989) F-statistic (GRS) for the joint significance of the α's being different from zero and its p-value. This statistic has the simple economic interpretation - a large GRS F-statistic implies that the distance between the sample Sharpe ratio of the factor portfolio and the maximal sample Sharpe ratio attainable, using all the test assets, is "too big". This indicates that the factor portfolio is "too far" from being mean-variance efficient. Given the large F-statistics, we reject the hypothesis that the returns on our portfolios are explained by the Fama-French 3-factor model.

To establish whether the slope of the term-structure is a priced factor in the cross-section of returns, we augment the Fama-French 3-factor model to include a slope factor. We test whether

[11] The only exception to the monotonic positive relation between term-structure slope and average returns is among the highest beta portfolios, although the difference between the smallest and largest quintiles is still a healthy four percentage points per year.

the slope of the term-structure of betas is a priced factor in the cross-section by running standard Fama-Macbeth regressions that include the slope factor in the first-stage regression. That is, for each portfolio i, we estimate the following time-series regression:

$$R^e_{i,t+1} = \beta_{0,i} + \beta_{Rm,i}R_{m,t+1} + \beta_{HML,i}R_{HML,t+1} + \beta_{SMB,i}R_{SMB,t+1} + \beta_{Slope,i}Slope_{t+1,i} + \epsilon_{t+1,i} \quad (6)$$

Panel A from Table 4 reports the Fama-Macbeth prices of risk of each factor.

[Table 4 here]

The prices of risk for all the Fama-French factors (Market, Size, and Book-to-Market) are insignificant. All explanatory power seems to come from the slope factor (with a t-stat of 2.3.). As a robustness check, we also test the performance of the slope factor using the Fama-French 25 double sorted portfolios on size and book-to-market. It is important to show that the slope factor is also priced beyond the 25 portfolios built based on the level and slope of the term-structure. This addresses some of the criticisms of traditional asset pricing tests raised by Lewellen, Nagel, and Shanken (2010).

Panel B of Table 4 reports the results of Fama-Macbeth regressions on the 25 size and book-to-market portfolios including the slope as a factor. In this specification, as expected, the high-minus-low factor is significant and substantially explains the cross-sectional variation. However, once we augment the model to include the slope as a factor, the cross-sectional R^2 of the model doubles (from 0.21 to 0.41). The slope of the term-structure of betas seems to have pricing power for the cross-section of returns. We also run the same tests but include the level of the term-structure as a control. We find that after controlling for the slope of the term-structure, its level has no explanatory power (second row of Panel B of Table 4). Thus, it is indeed the difference between the long-term beta and the short term-beta that improves the performance of the model. These results are robust to the exclusion of the financial crisis period.[12]

Economically, the reason why the slope of the term-structure is priced in the cross-section is its strong relation with the risk of the underlying portfolios. The relative magnitudes between the short-term beta and the long-term beta successfully capture the risk-return tradeoff. This will become more evident once we study the time-series properties of the slope of the term-structure of betas.

[12]See appendix B

3.2 Time-series

In the previous section, we conjectured that an upward sloping term-structure of option estimated market betas implies that investors expect an increase in risk in the underlying portfolio and therefore should earn higher returns on that portfolio. By sorting stocks on the ex-ante slope of their term-structure we established that portfolios with a negative (positive) slope earn higher (lower) expected returns, even after controlling for the level of the term-structure. We went a step further and found that the slope of the term-structure has pricing power for the 25 portfolios sorted on size and book-to-market.

We now examine the time-series dynamics of the slope. We start by plotting its time-series for the 25 size and book to market portfolios and then analyze what drives its dynamics.[13] Panel A of figure 3 plots the slope of the term-structure for value and growth stocks.

[Figure 3 here]

During unsettled periods such as the 2007-2009 financial crisis, value stocks have a negative term-structure slope, implying that market participants view these stocks as having a high systematic risk at the moment, but they expect it to decrease. This is in accordance with Lettau and Wachter (2007) argument: value firms with a low duration of cash-flows are deemed to be riskier during a recession in relation to growth stocks that have a higher duration of cash-flows. Therefore, value stocks during the crisis had a high short-term implied market beta and negatively sloped term-structure.

The difference between several maturities betas might shed some light on how long investors expect the recession to last. The reason that the slope of the term-structure is a priced factor in the cross-section is precisely this: when a particular stock becomes riskier, its short-term beta spikes vis-a-vis its long-term beta. The great benefit of estimating forward-looking CAPM betas is exactly the fact that they are akin to conditional CAPM betas. They take into account the current state of the nature and move around daily, capturing underlying unobserved factors.

In appendix D, we plot the dynamics of the slope of the term-structure for a portfolio of technology stocks and a portfolio of financial stocks. In line with the previous argument, during the turn of the millennium tech-bubble, technology stocks had a high systematic risk and thus a negatively sloped term-structure. The same pattern was evident for financial stocks during the recent financial crisis. This reinforces the fact that these conditional forward-looking betas are very powerful at capturing adverse states of nature that matter for the underlying stock or portfolio.

[13] In appendix C, we look into other interesting implications of the slope of the term-structure: how it behaves in aggregate and for individual stocks.

The important question now is what drives the slope of the term-structure of these portfolios. Changes in the slope of the term-structure can come from several sources. To start with, the slope can move around due to movements in the short-term beta β^Q_{iM,t,τ_s} or movements in the long-term beta β^Q_{iM,t,τ_l}.[14]

$$\beta^Q_{iM,t,\tau} = \frac{\sigma^Q_{i,t,\tau} \sum_{j=1}^{N} w_j \sigma^Q_{j,t,\tau} \rho^Q_{ij,t,\tau}}{(\sigma^Q_{M,t,\tau})^2} \tag{7}$$

Second, conditional on the change in the slope being driven by either the long-end or the short-end, that change might be due to either the first term on the numerator of equation (7) (the variance swap rate) or the second term (the correlation effect). In summary, the slope of the term-structure can change due to:

1. A short-end or long-end effect (or both):

$$\gamma_{1,t,i} = \beta^Q_{iM,t,\tau_l} - \beta^Q_{iM,t,\tau_s}$$

2. A variance swap effect:

$$\sigma^Q_{value,t,\tau} > \sigma^Q_{growth,t,\tau}$$

3. A correlation effect (in the spirit of Driessen, Maenhout, and Vilkov (2009)):

$$\sum_{j=1}^{N} w_j \sigma^Q_{j,t,\tau} \rho^Q_{value,j,t,\tau} > \sum_{j=1}^{N} w_j \sigma^Q_{j,t,\tau} \rho^Q_{growth,j,t,\tau}$$

Figure 4 shows the decomposition of the slope of the term-structure into its three components. It is clear that most of the effect comes from the short-end of the curve (Panel A of figure 4). This occurs because term-structures slope either upward or downward is mainly due to a change at the short-end of the curve.[15] In addition, most of the effect seems to be driven by the variance swap rate and not by correlation. In fact, throughout our sample, the correlation effect is essentially the same for both value and growth portfolios (Panel B of figure 4). It is important to emphasize that although most of the variation in the slope of the term-structure comes from the short-end of the curve, it is still the relation between the short-term and the long-term that matters for cross-sectional asset pricing.

[14] As we will see below, most of the effect is on the short-end of the curve.

[15] Actually, changes in both ends of the curve are strongly positively correlated, i.e. an increase in the short-term beta is likely to be associated with an increase in the long-run beta, but the effect is more pronounced on the short-term beta.

[Figure 4 here]

This new empirical evidence has strong implications for asset pricing models. On one hand standard asset pricing models such as the Consumption CAPM of Lucas (1976), the habit model of Campbell and Cochrane (1999), the long-run risk model of Bansal and Yaron (2004) or the duration model of Lettau and Wachter (2007) imply a negligible term-structure of variance. Furthermore, these models do not focus on the cross-section of firms. On the other hand, most studies that have investigated the term-structure of variance relied on no-arbitrage affine term-structure models (e.g. Aït-Sahalia, Karaman, and Mancini (2014), Egloff, Leippold, and Wu (2010) among others). Most structural models do not address on the term-structure of variance and the few that do fail to take into account the cross-section of stocks (e.g. Dew-Becker, Giglio, Le, and Rodriguez (2015)). In the next section, we therefore take a step further into the theoretical asset pricing literature and build a model that concurrently examines the cross-section of returns and the term-structure of variance swaps.

4 An investment model with uncertainty shocks

In this section we build a dynamic investment-based model, with heterogeneous firms and time-varying uncertainty, to link the cross-section of returns and the slope of the term-structure of systematic risk. We model time-varying uncertainty in the same way as Bloom (2009), by assuming that volatility follows a two-state markov chain. Uncertainty shocks drive both the value-premium and the term-structure of variance swaps. During normal times, the variance swap term-structure is almost flat (see panel D of figure 4). This is standard: during normal times, short-term expected variance is low and as uncertainty shocks are unlikely, long-term expected variance is also low. On the other hand, at the inception of a uncertainty shock, short-term variance spikes inducing a negatively sloped term-structure of variance. The persistence of the uncertainty shock keeps the term-structure of variance negatively sloped for a long-time. These uncertainty shocks have a heterogeneous effect on the cross-section of stocks, depending on firms' holdings of growth options. Growth firms have low assets in place so that, when uncertainty increases, they perform better than value stocks and thus have lower expected variance. Given that growth stocks variance spikes less than value stocks, this accounts for the heterogeneous slope of the option implied beta term-structure observed in the data.

Our choice for the volatility process relies on the empirical observation made by Bloom (2009) that uncertainty appears to dramatically increase after major economic and political shocks. In figure 5 we replicate figures 1 and 2 from Bloom (2009). In the left panel of figure

5 we plot the VIX index.[16] It is clear that market volatility - a common proxy for uncertainty - displays large spikes after major shocks. At the same time, these second-moment shocks generate large falls, of around 1% on impact, in output (right panel of figure 5). These two empirical observations will be important for our modelling choices.[17]

[Figure 5 here]

4.1 The economic environment

The economy is populated by N firms, indexed by j, that use physical capital ($K_{j,t}$) to produce a homogeneous good ($Y_{j,t}$). Firms have a standard neoclassical production function given by:

$$Y_{j,t} = X_t Z_{j,t} K_{j,t}^{\alpha} \tag{8}$$

where X_t is the aggregate productivity and $Z_{j,t}$ is firm-specific productivity, and $0 < \alpha < 1$ is the curvature parameter. The production technology exhibits decreasing-returns-to-scale.

Let x_t and $z_{j,t}$ denote the log of X_t and $Z_{j,t}$ respectively. Both follow a first order auto-regressive process:

$$x_{t+1} = \bar{x} J_t (1 - \rho_x) + \rho_x x_t + \sigma_t^x \epsilon_{t+1}^x \tag{9}$$

$$z_{j,t+1} = \rho_x z_{j,t} + \sigma_t^z \epsilon_{j,t+1}^z \tag{10}$$

in which ϵ_{t+1}^x and $\epsilon_{j,t+1}^z$ are independent and identically distributed (i.i.d.) uncorrelated normal shocks, $\bar{x}$ is the long-term mean of aggregate productivity; ρ_x and ρ_z are the persistence of aggregate and firm-level productivity, respectively; σ_t^x and σ_t^z are the conditional volatility of innovations to aggregate and firm-specific productivity, respectively. J_t is a very small productivity jump. We will thoroughly discuss how we calibrate the jump and why it is important in our setup in the next sections.

Firms can either pay dividends or invest in capital. Physical capital accumulation is given by:

$$K_{j,t+1} = (1 - \delta) K_{j,t} + I_{j,t} \tag{11}$$

[16] We only have VIX data starting from mid-80s, so we plot actual realized volatility before that.

[17] In appendix E we discuss how the average level of uncertainty in the market (VIX) relates with the average absolute slope of the term-structure of betas.

where δ is the depreciation rate and $I_{j,t}$ is investment at time t. Firms incur convex asymmetric adjustment costs of investment ($G_{j,t}$) as in Zhang (2005), Lin and Zhang (2013) and Bloom (2009):

$$G(I_{j,t}, K_{j,t}) = \begin{cases} a^+ K_{j,t} + \frac{c^+}{2}\left(\frac{I_{j,t}}{K_{j,t}}\right)^2 K_{j,t}, & \text{for } I_{j,t} > 0 \\ 0 & \text{for } I_{j,t} = 0 \\ a^- K_{j,t} + \frac{c^-}{2}\left(\frac{I_{j,t}}{K_{j,t}}\right)^2 K_{j,t}, & \text{for } I_{j,t} < 0 \end{cases} \tag{12}$$

where $a^- > a^+ > 0$, and $c^- > c^+ > 0$ capture nonconvex adjustment costs.

4.2 Time-varying uncertainty

Bloom (2009) documents that uncertainty appears to dramatically increase following major economic and political shocks (just like the slope of the term-structure of systematic risk). We follow Bloom (2009) and extend the standard Lin and Zhang (2013) model with a time varying second moment and use the model to simulate the impact of a large temporary uncertainty shock. The stochastic volatility process for both aggregate (σ_t^x) and firm-specific (σ_t^z) productivity is assumed to follow a two-point Markov Chain, i.e.:

$$\sigma_t^k = \{\sigma_L^k \sigma_H^k\}, \text{ where } Pr(\sigma_{t+1}^k = \sigma_j^k | \sigma_t^k = \sigma_i^k) = \pi_{k,j}^{\sigma^k}, \text{with } k \in \{x, z\} \tag{13}$$

Time-varying uncertainty generates periods of low and high volatility both for the aggregate productivity process and for the firm specific productivity process. We will assume these two to be perfectly correlated, meaning that periods with high aggregate uncertainty also correspond to periods with hight firm-specific uncertainty. In this model, due to the presence of nonconvex adjustment costs, when uncertainty spikes, firms delay investment decisions as the real option value of waiting increases. As we will see below, growth stocks have low assets in place and therefore higher growth options, which provide a hedge against uncertainty shocks. The opposite is true for value stocks, making them riskier.

4.3 Stochastic Discount factor

Following Zhang (2005) we directly parameterize the stochastic discount factor which is a function of the mean reverting state variable x_t as follows:

$$\log(m_{t+1}) = \log(\eta) + \gamma_t(x_t - x_{t+1}) \tag{14}$$

The risk aversion parameter γ_t is given by:

$$\gamma_t = \gamma_0 + \gamma_1 (x_t - \bar{x}) \tag{15}$$

with $\gamma_0 > 1$ and $\gamma_1 < 0$. This allows the model to generate time-varying risk aversion and a counter-cyclical price of risk *à la* Campbell and Cochrane (1999).

4.4 Equity value, returns, and variance swap rates

The profit function for an individual firm j, with capital stock $K_{j,t}$, idiosyncratic productivity $Z_{j,t}$ and an aggregate state of X_t is given by:

$$\pi_{j,t} = X_t Z_{j,t} K_{j,t}^{\alpha} - f \tag{16}$$

where f denotes fixed costs of production. Firm's maximize the present value of profits less adjustment costs. Denote $V_{j,t}$ the cum-div the market value of firm j. The Bellman equation of the optimization is given by:

$$\begin{aligned} V_{j,t} &= V(K_{j,t}, X_t, Z_{j,t}, \sigma_t^x, \sigma_t^z) = \\ &= \max[\pi_{j,t} - I_{j,t} - G(I_{j,t}, K_{j,t}) + E_t[M_{t+1} V(K_{j,t+1}, X_{t+1}, Z_{j,t+1}, \sigma_{t+1}^x, \sigma_{t+1}^z))]] \end{aligned} \tag{17}$$

subject to equation (11). Notice that the cum-dividend value of equity of each firm, depends on the five state variables: $K_{j,t}$ (the current level of capital stock of each firm), $Z_{j,t}$ (the current productivity level specific to each firm), X_t (the aggregate level of technology), σ_t^x and σ_t^z (the level of aggregate and firm-specific uncertainty).

The above setup implies that at the optimum, $V_{j,t} = D_{j,t} + E_t[M_{t+1} V_{j,t+1}]$ with $D_{j,t} = \pi_{j,t} - I_{j,t} - G(I_{j,t}, K_{j,t})$. It follows that the stock return of firm j is simply given by: $r_{j,t+1} = V_{j,t+1}/(V_{j,t+1} - D_{j,t})$ and that the ex-div stock price S at time t is given by: $S_{j,t} = V_{j,t} - D_{j,t}$.

Given a risk-neutral (pricing) measure Q, the price of an τ-month variance swap at the end of month t for stock j, $(\sigma_{j,t,\tau}^Q)$ is given by:

$$\sigma_{j,t,\tau}^Q = E_t^Q \Big[\sum_{l=1}^{\tau} RV_{j,t+l} \Big] \tag{18}$$

where E_t^Q denotes the mathematical expectation under the risk-neutral measure conditional on the state of the world at time t and $RV_{j,t}$ is the realized variance of stock j during month t.

To keep as close as possible to the empirics in the paper, we backout variance swap rates for each single stock using again the methodology of Demeterfi, Derman, Kamal, and Zou (1999) and Bakshi, Kapadia, and Madan (2003).[18] In order to do so, we need to value call and put options within our model. Given that the stochastic discount factor is known the value of a call option C, on stock j, at time t with maturity τ and strike K is given by:

$$C_j(t,\tau,K) = E_t[M_{t+\tau}\max\{S_{j,t+\tau}(.) - K, 0\}] \tag{19}$$

where we suppress the state variables of the policy function $S(K_{j,t+\tau}, X_{t+\tau}, Z_{j,t+\tau}, \sigma^x_{t+\tau}, \sigma^z_{t+\tau})$ for convenience. To evaluate the expectation above we employ numerical methods and make extensive use of Gauss-Hermite Quadrature techniques and splines for points that do not lie on the grid for the state variables. It is straightforward to backout call option values for any individual stock j. However, given the large number of state variables, it is impossible to integrate equation (19) for the overall market portfolio. In order to do so we rely on an aggregation technique similar to the one proposed by Krusell and Smith (1998) and Den Haan et al. (1997) and approximate the equity market-value by means of a regression.[19] Therefore we conjecture that any given point in time t the aggregate market value of equity, $S_{M,t}$ is a function of the aggregate state variables and the lagged market value of equity:

$$S_{M,t} = \Gamma(X_t, \sigma^x_t, S_{M,t-1}) \tag{20}$$

We assume that the market value of equity is log-linear in the state space. The above specification yields an R^2 of 90% if one does not include the lagged market value of equity on the specification and an R^2 of 98% once we include it. So it is an overall good approximation, and allows us to reduce the dimensionality of the problem and compute the option implied variance of the market.

The model is solved using value function iteration on a discrete state space. Tauchen-Hermite quadrature methods and numerical interpolation are extensively used to compute expectations.

4.5 Quantitative results

In subsection 4.5.1 we calibrate the model, then subsection 4.5.2 presents the main quantitative results and subsection 4.5.3 investigates the model key mechanism to generate the term-structure of variance and systematic risk.

[18]Remember that we estimate variance swap rates using the following approximation: $(\sigma^Q_{j,t,\tau})^2 = \int_{S_j(t)}^{\infty} \frac{2\left(1-\log[\frac{K}{S_j(t)}]\right)}{K^2} C_i(t,\tau,K)dK + \int_0^{S_j(t)} \frac{2\left(1-\log[\frac{K}{S_i(t)}]\right)}{K^2} P_i(t,\tau,K)dK$

[19]The main difference between our setup and theirs is that the equity market value is not a state variable in the problem, so we do not need to iterate for a fixed point.

4.5.1 Calibration

We calibrate all model parameters at monthly frequency to be consistent with the empirical section of this paper. Table 5 reports the parameters values used to solve and simulate the model. Most parameters are taken from the literature - namely from Lin and Zhang (2013) and Bloom (2009) - and allow us to match selected moments in the data.

[Table 5 here]

The parameters $\beta = 0.99$, $\gamma_0 = 6$ and $\gamma_1 = -1000$ are set to match the average real interest rate (1.80%), its volatility (3.00%) and the average market Sharpe Ratio (0.32). The curvature of the production function, α, is set to be 0.7 as in Hennessy and Whited (2007). The monthly rate of depreciation, δ, is set to be 0.01, which implies an annual rate of 12%. The persistence of aggregate productivity ρ_x is set to be $0.95^{\frac{1}{3}}$ consistent with Cooley and Prescott (1995). The persistence ρ_z of firm-specific productivity, is set to 0.96. The uncertainty process parameters are calibrated using Bloom (2009), i.e. an uncertainty shock is expected every 3 years and have a 2-month half-life. Uncertainty shocks double the baseline uncertainty. We set aggregate unconditional volatility and firm-specific volatility to the same value as in Lin and Zhang (2013). We calibrate the jump, J_t, to 1.003 which allows us to match the 1% drop in output at the inception of a volatility shock (as documented by Bloom (2009)). Finally, the parameter $\bar{x}$ is a scaling variable and has no implication on the target moments. We set it to -3.65 which implies a steady state value of capital equal to one.

In total 2,500 samples of artificial data are simulated at monthly frequency, with 15 years and 1,500 firms each, with all values averaged across these samples. Similar to Bloom (2009) in each simulation we hit the economy with an uncertainty shock in the first month of the eleventh year, defined as $\sigma_t = \sigma_H$. In any given sample, some economies will already be in the high uncertainty state whereas others will forcefully move to that state. This allows us to understand the impact of an uncertainty shock in returns, the term-structure of variance and the term-structure of betas. All other shocks are randomly drawn.

The comparison between target moments from data and those from model simulations is given in table 6. Table 6 shows that the model does a reasonable job of matching the key return and quantity moments.[20]

[20] The average market return, volatility of market return and Sharpe ratio in the data are taken from Goyal and Welch (2008). The data moments of the real interest rate are from Campbell, Lo, MacKinlay, et al. (1997). The average output fall conditional on a volatility shock is from Bloom (2009). The remainder of moments are from Pontiff and Schall (1998).

[Table 6 here]

Importantly, the fit seems reasonable not only for the moments that serve as immediate targets of calibration, but also for other moments. The mean and volatility of the market return are comparable to those computed using the data from Goyal and Welch (2008). The median of aggregate book-to-market ratio is 1.49, close to that of 1.53 reported by Pontiff and Schall (1998). The average rate of investment is 0.135 in the model, close to 0.15 in the data reported by Abel and Eberly (2001).

4.5.2 Stock returns, the term-structure of variance and the term-structure of systematic risk

In this section we investigate the empirical predictions of the model for the cross-section of returns, the term-structure of variance and systematic risk.

We start by studying the performance of the option implied CAPM, within the model, in pricing the ten book-to-market decile portfolios. For each of the 2,500 artificial simulated samples we use the exact timing of Fama and French (1993) to sort stocks into deciles based on their book-to-market. We then run the standard CAPM regressions to estimate betas and use the same methodology from our empirical section to estimate option implied betas. Table 7 reports the results. The first thing to note is that the model generates a value premium of around 5% per annum which is close to the spread in the data of 4.88% as reported by Lettau and Wachter (2007). Second, in the model, standard CAPM betas (fourth row of table 7) do not line up that well with expected returns. The spread in betas between the top and the bottom market-to-book decile is a dismal 0.2. On the contrary, betas from the option implied CAPM (fifth row of table 7) line up extremely well with expected returns and have a spread of 1.5 between the bottom and top market-to-book decile. This is similar to the pattern we uncovered in our empirical section where the option implied CAPM betas line up better with expected returns than the standard CAPM. The next subsection will look thoroughly as to why this is the case.

[Table 7 here]

Before dwelling into the term-structure of variance and betas, it is useful to show the precise impulse that will drive the term-structure results. Figure 7 plots the average value of σ_t^x normalized to unity before the shock (the results are virtually unchanged for the firm-specific volatility σ_t^z). It is clear that the uncertainty shock generates a spike in volatility, and that the shock dies out fairly quickly; within two months the volatility is already half-way through its long-term value (remember that we calibrated the shock to be consistent with the estimates by Bloom (2009)). The shock almost doubles the average σ_t^x. The rise is less than 100% since at the inception of the shock some of our 2,500 economies are already in state of the nature with high volatility.

[Figure 7 here]

These volatility shocks have a heterogenous impact on the term-structure of variance of individual securities. Figure 8 plots the average slope of the term-structure of variance for value and growth stocks. During normal times, i.e. during the twelve months that precede the volatility shock the term-structure of variance is nearly flat, with value stocks having a slightly higher slope than growth stocks. The difference is due to value stocks being riskier than growth stocks as a result of the convex adjustment costs, which make assets in place very risky when aggregate productivity decreases. Further, at the inception of an uncertainty shock the term-structure of variance of value stocks spikes much more than the one of growth stocks. This is in accordance with the empirical evidence outlined in Panel D of figure 4, which plots the slope of the term-structure of variance swap rates for value and growth stocks. As uncertainty increases, growth stocks are effectively hedged against this shock as the increased uncertainty expands the upside of future outcomes. Growth stocks have more growth opportunities than value stocks, and therefore perform better. To the best of my knowledge this is the first paper to look into the term-structure of individual stocks variance swap rates within a structural model. Most models focus only on the aggregate market implied variance (see for example Dew-Becker, Giglio, Le, and Rodriguez (2015) and Bollerslev, Tauchen, and Zhou (2009)).

[Figure 8 here]

This heterogeneity in the dynamics of the slope of the term-structure of variance swaps is the key driver of the term-structure of systematic risk in the model. Figure 9 plots the slope of the term-structure of betas in a twenty-four month window around a volatility shock.

[Figure 9 here]

There are several interesting things to note. First, value stocks have on average a negative slope of their term-structure of betas (i.e., the short-term beta is higher than their long-term beta). Value stocks suffer much more from a given volatility shock. Therefore, ex-ante, they have a higher short-term beta. At the inception of a volatility shock the systematic risk of these stocks soar, their option implied beta spikes, and the persistence of the volatility shock makes their short-term beta slowly decay to its long-run mean.

Second, just as in the data, the absolute change in the slope of the term-structure of value stocks betas is slightly higher than the one for growth stocks.

Third, it is worth mentioning that although the changes in the slope of the term-structure of variance swaps are somewhat in line with the data (see the shaded are in panel D from figure 4), the slopes of the term-structure of betas seem to high. In the model we estimate the long-term market implied variance using an approximate aggregation; this approximate aggregation only depends on aggregate state variables, thus leaving out the firm-specific components; although on the aggregate these elements might not matter to describe the expected path of $S_{t,Market}$ they are likely to matter for the value of the call option. Not including them is likely to decrease its time-value leading to an underestimation of the true long-term variance swap rate, and increasing the estimates of the slope. Despite this issue, the bottom line of this paper still goes through: systematic risk does have a term-structure which is driven by the combination of volatility shocks and heterogeneous firm-specific growth opportunities. In the next section, we take a closer look at the mechanism that drives this result.

4.5.3 Inspecting the mechanism

In this section we take a closer look at the mechanism that underlies the relative success of the option implied CAPM vis-a-vis the standard CAPM. We do so by using the investment model we outlined above. There are two things that account for the relative success of the option implied CAPM. First, unlike the standard CAPM that relies on regressions on historical data, which implicitly assume that the past is sufficiently close to the future, the option implied betas take into account the current state of the nature and agents' expectations regarding the future. This is also true in the model.

Second, the presence of volatility shocks and small productivity jumps, as empirically shown by Bloom (2009), change the relative risk of different portfolios according to their growth options and this again is well captured, in the model, by our option implied betas. Figure 10 breaks down the mechanism that generates the heterogeneity of the term-structure of variance swaps (and consequently of the term-structure of betas). Panel A of figure 10 plots the same figure we have seen before: the term-structure of variance swaps spikes following a volatility shock, with a higher impact on firms that have a higher share of assets in place. This higher impact on value firms is mainly driven by the presence of the small productivity jump (Panel B of figure 10). If there was no productivity jump, the impact would be very similar among value and growth stocks.

It is this empirical correlation between volatility and productivity that drives the heterogeneity. Finally, in Panel C from figure 4.5.3 we shut down both the volatility shocks and the productivity jumps. This is exactly the model of Zhang (2005) and Lin and Zhang (2013). This model implies a flat term-structure of variance swaps for portfolios and the market. This is at odds with the empirical evidence that there a term-structure of variance swaps for the market (see, for example, Dew-Becker, Giglio, Le, and Rodriguez (2015)). Therefore, it seems important for asset pricing models to take into account volatility jumps and their correlation with productivity. Taking into account these features might also help to explain the term-structure of equity (see Van Binsbergen and Koijen (2015)).

5 Conclusion

In this paper we have uncovered the existence of a term-structure of systematic risk. The slope of the term-structure is priced in the cross-section of returns. In the time-series it spikes following relevant macroeconomic and firm-specific events. We also took a step into the asset pricing theory literature by showing that volatility jumps are very important to match the underlying cross-section and time-series dynamics of the term-structure of systematic risk. We conclude this paper by briefly discussing two potential avenues for future research in this area. It would be interesting to extend the large literature on real and nominal term structure models to match the term-structure of systematic risk. We have given a first step in this direction, but further work is needed. Also the term-structure literature on equity should evolve hand-in-hand with the literature on the term-structure of systematic risk, as these two things should be inherently connected.

Figure 1: Term-structure of betas for Apple Corporation

The blue points on this figure are the option implied betas for Apple Corporation for the 10 maturities under analysis (1, 2, 3, 4, 5, 6, 9, 12, 18, 24 months) on two different dates. The red line is the Nelson-Siegel model fit (equation 21). We fix theta to maximize the loading of the medium-term factor at six months which is usually where the hump occurs and estimate the remaining parameters by Ordinary Least Squares.

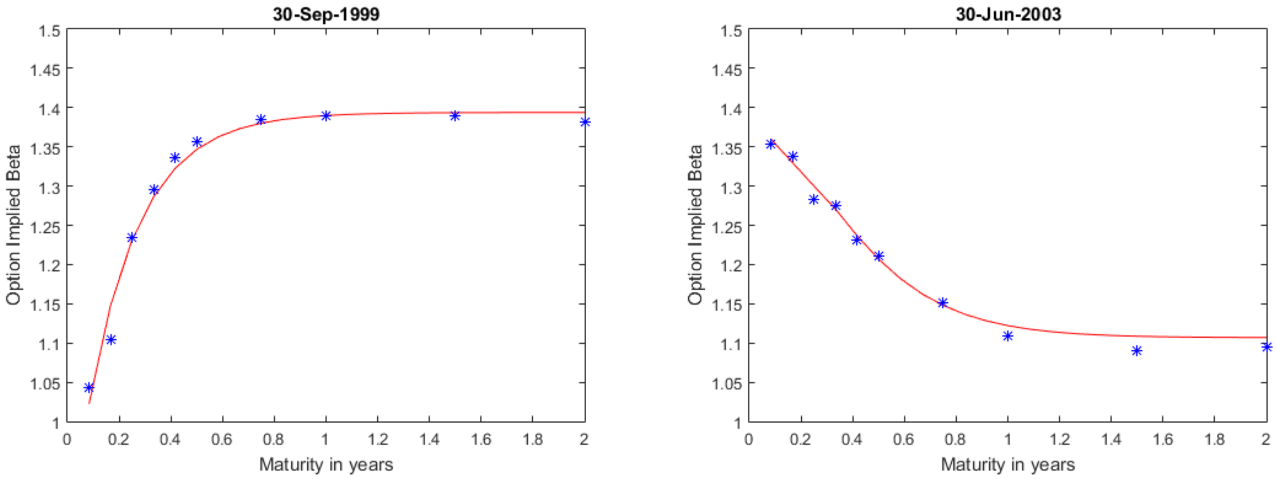

Figure 2: Portfolio Beta Predictions

The figure shows the scatter plots of expected and realized quintile portfolio betas over the sample period from January 1996 to August 2014. At the end of each month we sort stocks in five portfolios based on their implied or historical beta. The first portfolio thereby contains the stocks with the lowest expected market betas, and the last portfolio contains the stocks with the highest expected market betas. For each portfolio, month and methodology we then compute the value-weighted realized beta over the next six months (Panel A and B) and the next twelve months (Panel C and D).The figure plots the realized quintile portfolio betas against the expected quintile portfolio betas for all five quintiles, separately for each methodology.

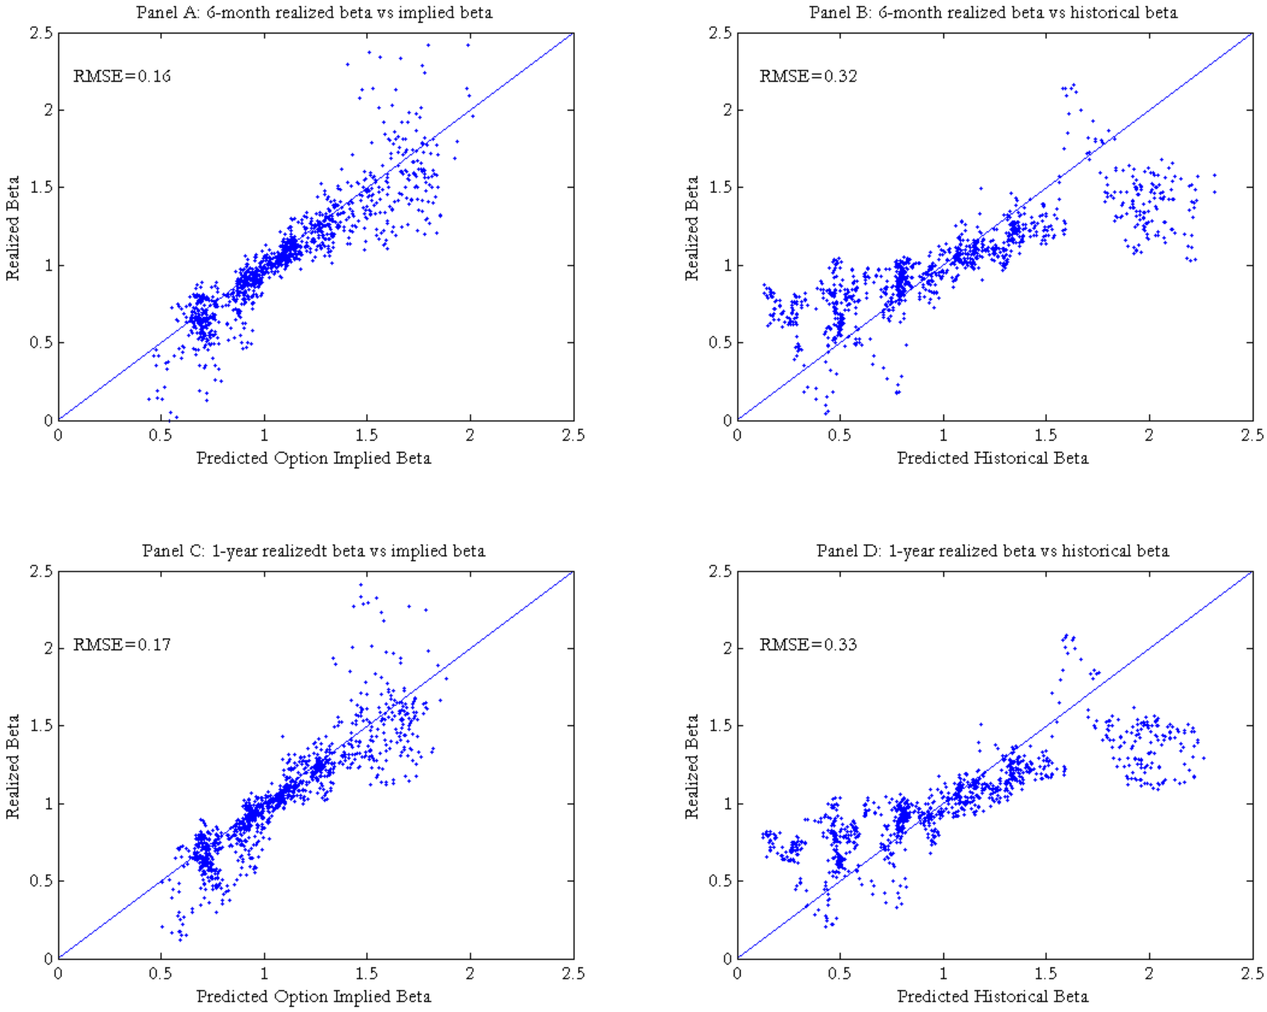

Figure 3: Beta term-structure slope

This figure plots a monthly time-series of the slope of the term-structure between January 1996 to August 2014. Panel A reports the slope for value (high book-to-market) and growth (low book-to-market) stocks and Panel B reports the slope for small and large firms in terms of market capitalization. Following the same methodology of Fama and French (1993) we sort stocks on five portfolios based on size and book-to-market.

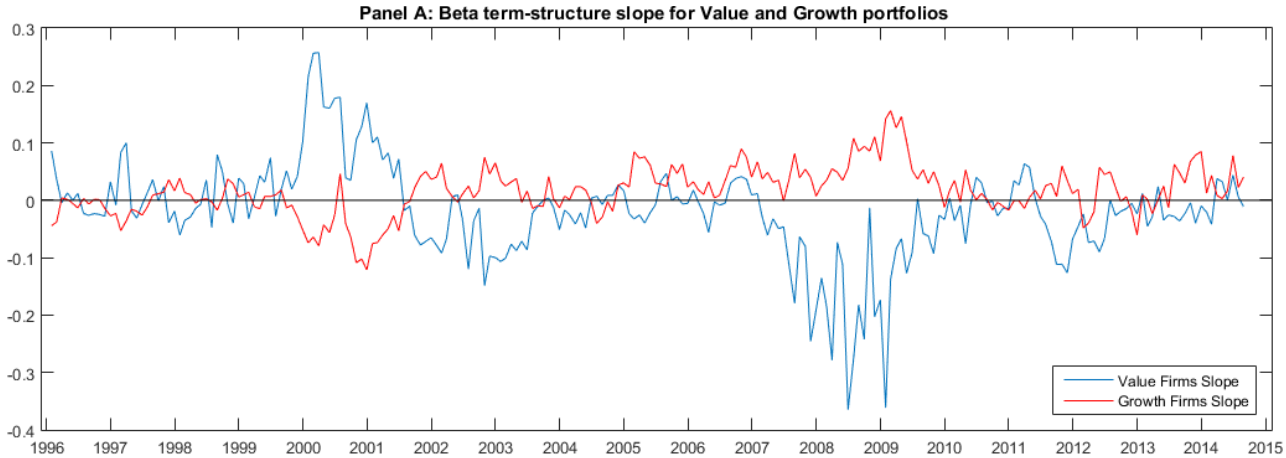

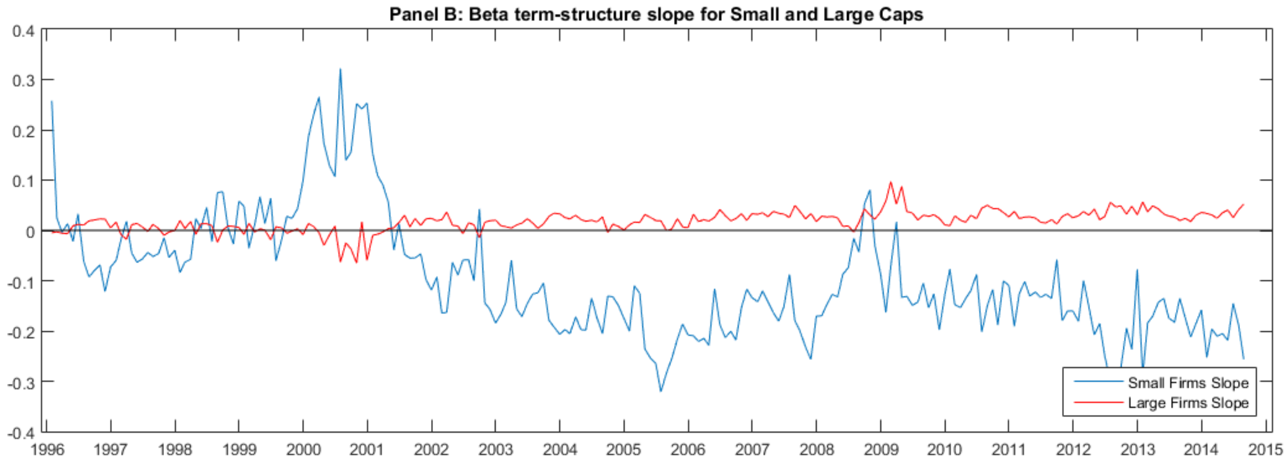

Figure 4: Decomposition of the slope of value and growth stocks into its three components

This figure decomposes the slope of the term-structure of betas of value and growth portfolios into its three components: (i) short-term vs long-term effect (Panel A), (ii) correlation effect (Panel B) and (iii) variance swap effect (Panel C).

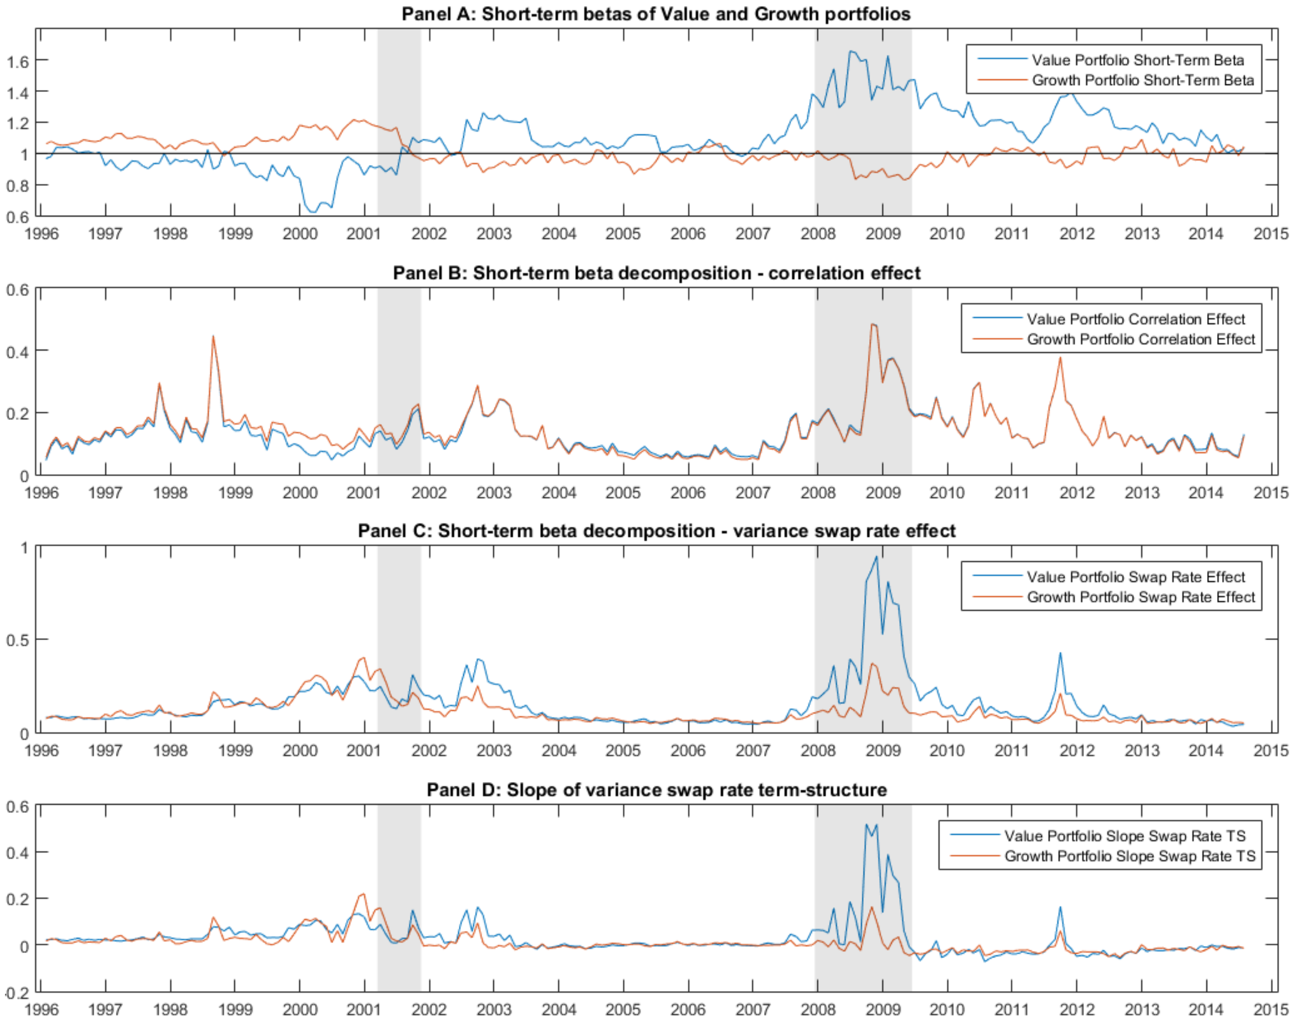

Figure 5: Monthly U.S. stock market volatility and the impact of a volatility shock on production

This figure replicates figures 1 and 2 from Bloom (2009). Panel A plots the VIX index between 1986 onwards. Pre-1986 the VIX index is unavailable, so actual monthly returns volatilities are calculated as the monthly standard deviation of the daily S&P500 index normalized to the same mean and variance as the VIX index when they overlap from 1986 onward. Panel B plots a VAR estimation of the impact of a volatility shock on industrial production. Dashed lines are 1 standard-error bands around the response to a volatility shock. A more detailed description on the data and methodology underlying this figure can be found in Bloom (2009).

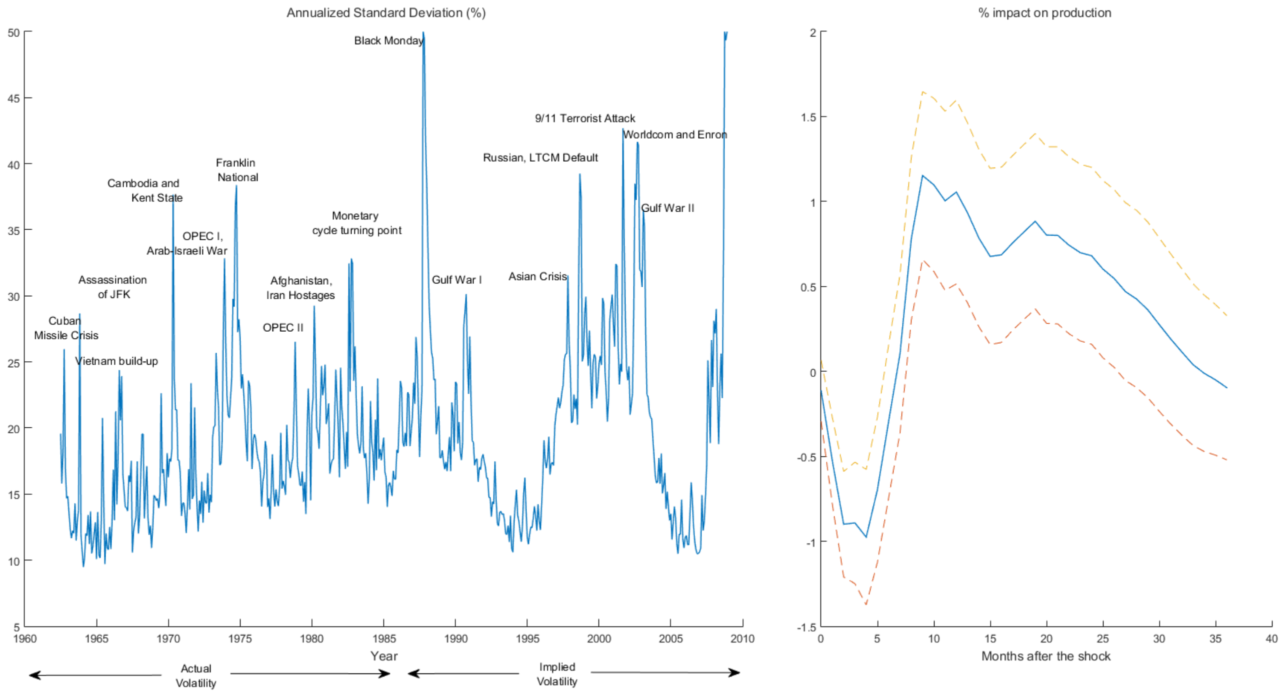

Figure 6: Average beta term-structure

This figure plots the VIX index and the value-weighted average beta term-structure slope (equation 9) between January-1996 and August-2014.

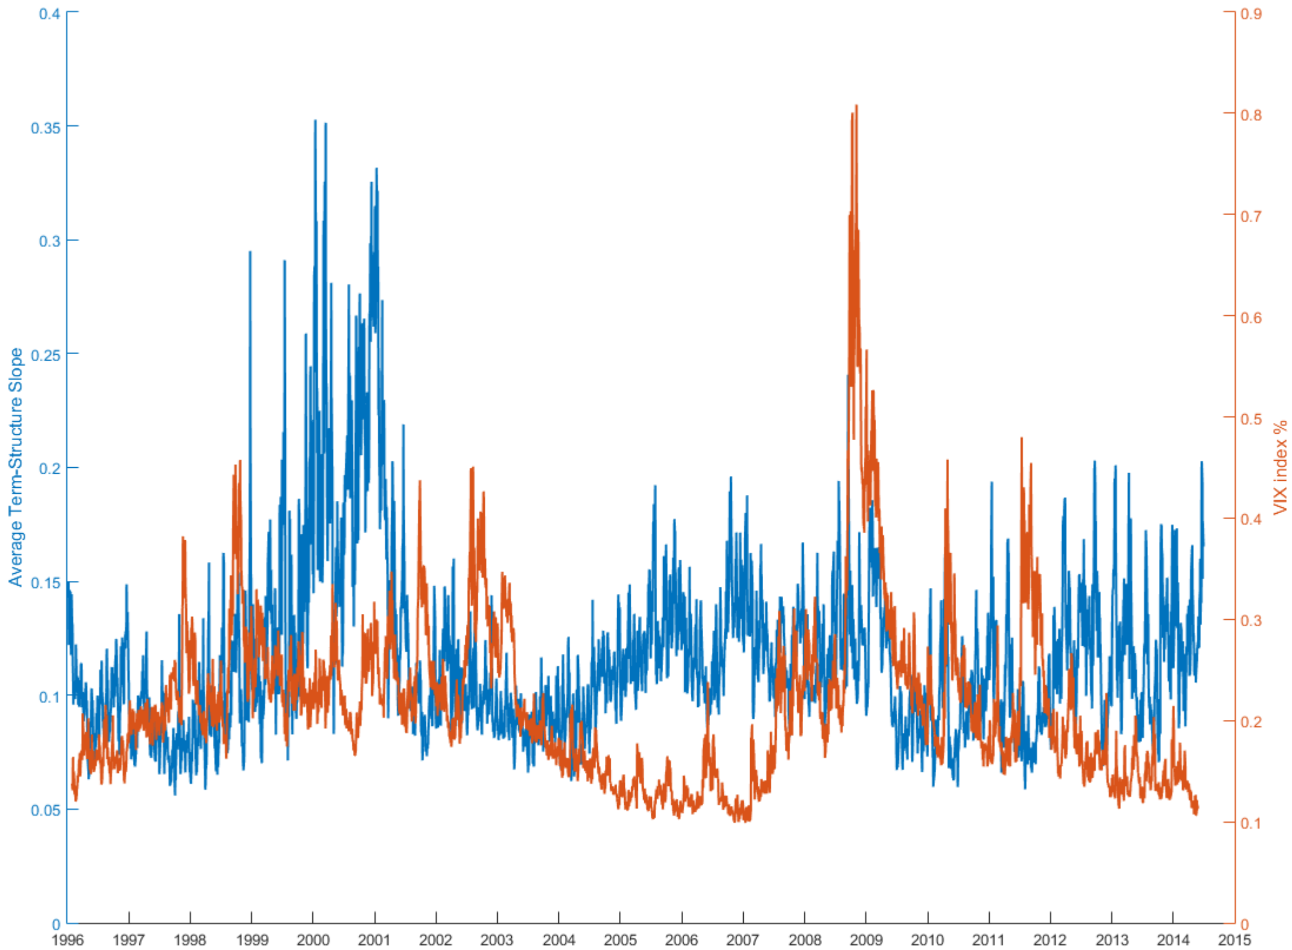

Figure 7: The simulation has a large second-moment shock

We simulate 2,500 artificial data samples at monthly frequency, with 15 years each In each simulation we hit the economy with an uncertainty shock in the first month of the eleventh year, defined as $\sigma_t^x = \sigma_t^H$. This figure plots the average value of σ_t^x around a 24-month window of the shock. The shock is normalized to unity before the shock date. It is plotted on a monthly basis, with the month normalized to zero on the date of the shock.

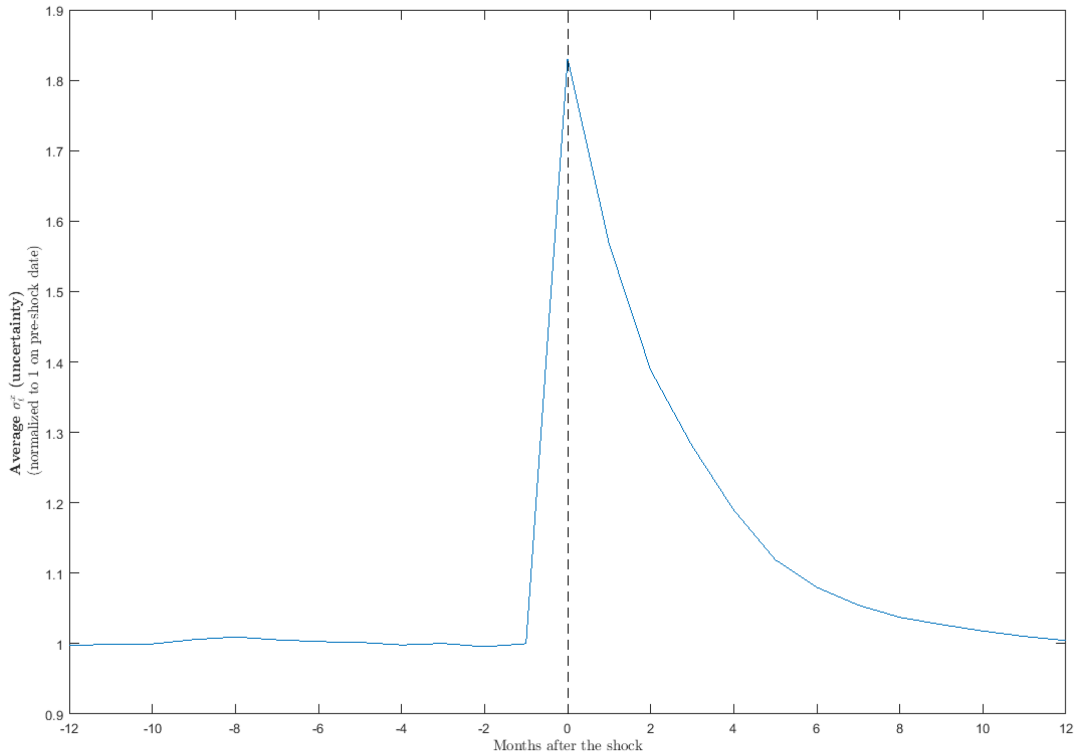

Figure 8: Slope of term-structure of option implied volatility for value and growth stocks

This figure plots $\sigma^Q_{j,t,\tau_l} - \sigma^Q_{j,t,\tau_s}$ for value and growth stocks on a window of 12 months before and after the uncertainty shock. We simulate 2,500 artificial data samples at monthly frequency, with 15 years each In each simulation we hit the economy with an uncertainty shock in the first month of the eleventh year, defined as $\sigma^x_t = \sigma^H_t$. We then average across all simulations to understand the impact of a volatility shock on the term-structure of variance swaps.

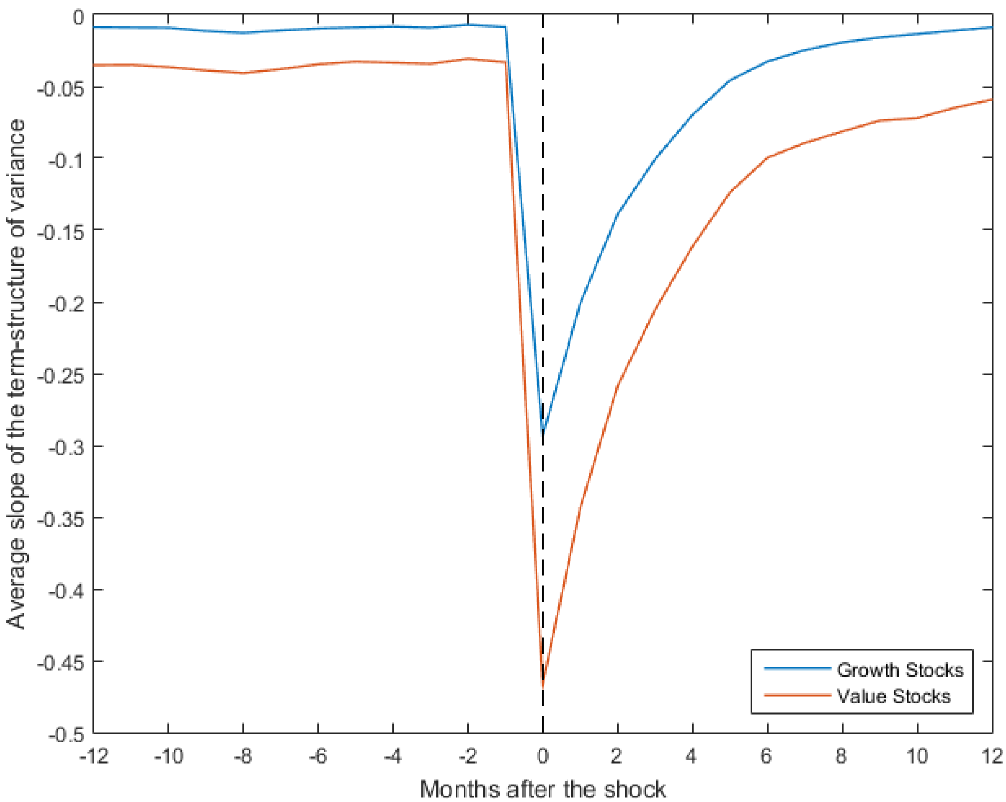

Figure 9: Slope of the term-structure of betas for value and growth stocks

This figure plots the slope of the term-structure of betas ($\beta^Q_{iM,t,\tau_l} - \beta^Q_{iM,t,\tau_s}$) for value and growth stocks, on a window of 12 months before and after the uncertainty shock. We simulate 2,500 artificial data samples at monthly frequency, with 15 years each In each simulation we hit the economy with an uncertainty shock in the first month of the eleventh year, defined as $\sigma^x_t = \sigma^H_t$. We then average across all simulations to understand the impact of a volatility shock on the term-structure of betas.

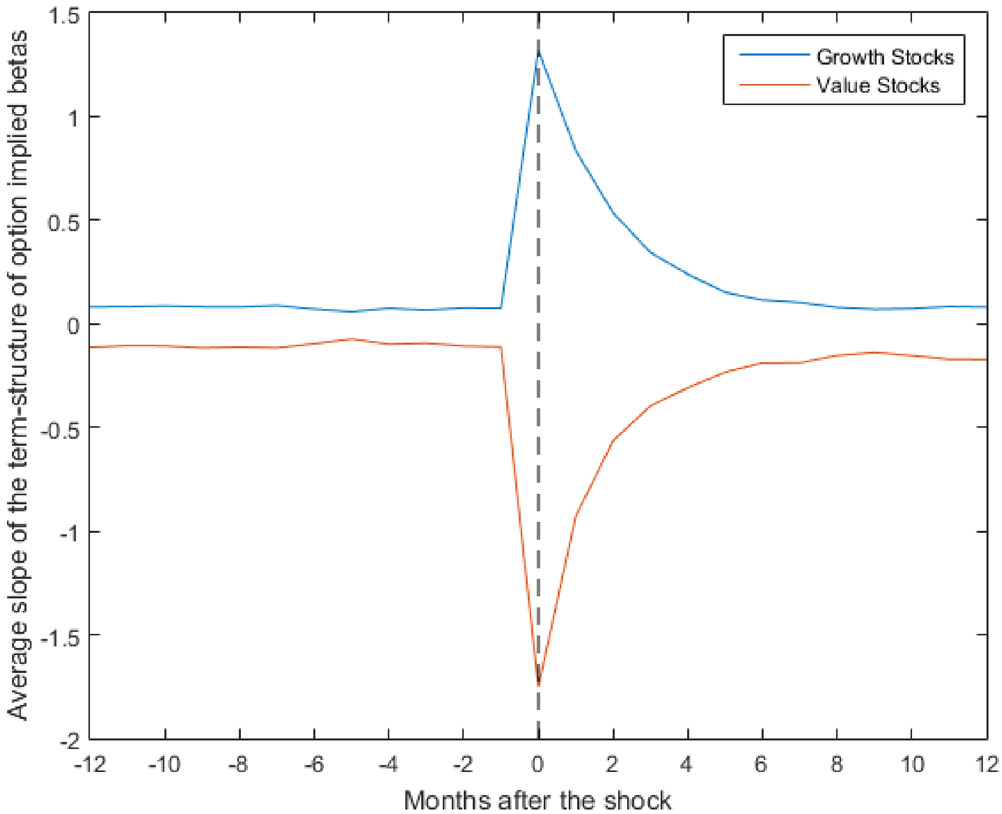

Figure 10: Term-structure of variance swaps - mechanism breakdown

In this figure we breakdown the model mechanism to generate a term-structure of variance swaps. Panel A plots the term-structure of variance swaps before and after a volatility shock for the benchmark model. In Panel B we plot the same thing but shutting down the small productivity jump. Finally, Panel C plots the standard term-structure of variance swaps for the standard Lin and Zhang (2013) model, i.e., a model with no volatility shocks nor productivity jumps.

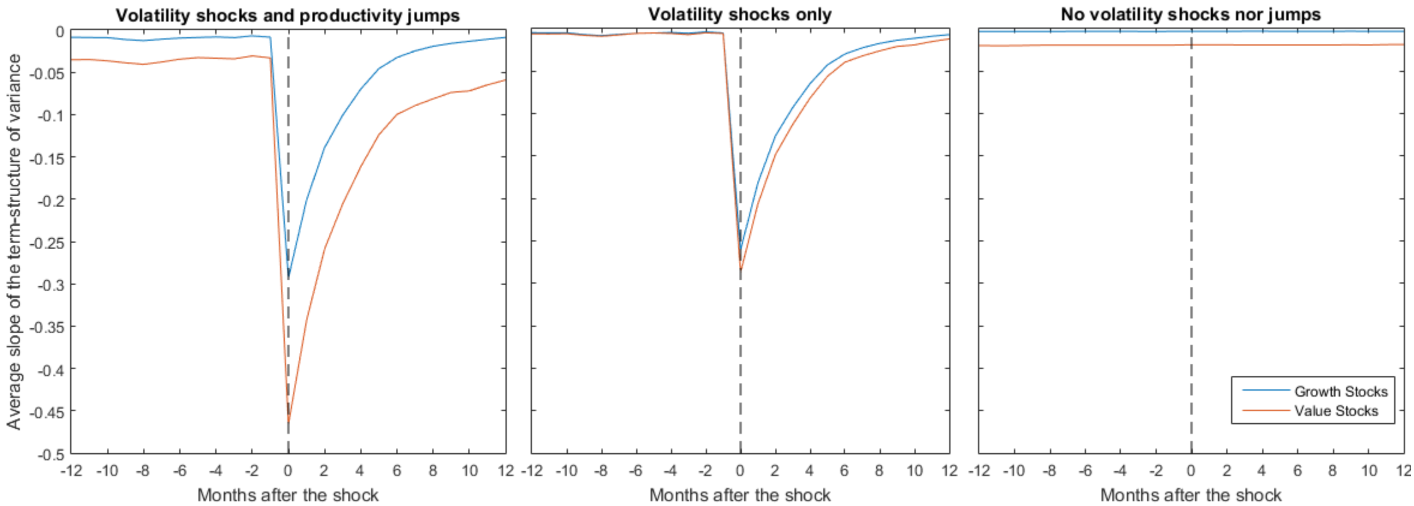

Table 1: Risk-Return Option Implied Betas and Historical Betas

This table reports mean expected beta and mean realized return for five portfolios sorted on ex-ante beta. At the end of each month for we sort portfolios in accordance to their historical or option implied beta. Then we compute the value-weighted monthly return of the portfolio over the next τ-months where τ is the maturity of the options underlying the estimation. For historical betas, the returns are over the following month. The numbers in the table are the time-series means of these values. The first (last) portfolio contains stocks with lowest (highest) expected beta.

	1	2	3	4	5	5-1
Historical						
Expected Beta	0.39	0.69	0.95	1.25	1.85	-
Realized Return	8.14	9.97	9.19	11.89	10.91	2.77
60 day options						
Expected Beta	0.65	0.88	1.06	1.27	1.68	-
Realized Return	8.13	10.83	10.91	11.14	11.58	3.45
91 day options						
Expected Beta	0.66	0.88	1.05	1.26	1.66	-
Realized Return	7.73	10.06	11.04	11.16	12.08	4.36
122 day options						
Expected Beta	0.67	0.89	1.05	1.25	1.65	-
Realized Return	7.72	9.70	10.75	11.12	13.13	5.41
152 day options						
Expected Beta	0.67	0.89	1.05	1.25	1.63	-
Realized Return	7.80	9.87	10.46	10.99	13.55	5.75
182 day options						
Expected Beta	0.68	0.89	1.05	1.24	1.63	-
Realized Return	7.84	10.43	11.19	11.43	13.09	5.24
365 day options						
Expected Beta	0.70	0.90	1.05	1.23	1.58	-
Realized Return	7.60	9.98	11.09	10.37	12.40	4.80

Table 2: Shapes of term-structure of betas: distribution

This table reports the distribution of shapes of term-structure of betas. For every day in our sample and every stock on the S&P 500 we compute a term-structure of betas using the methodology described in section 2. We classify a term-structure to be upward (downward) sloping if $\gamma_1 > 0.1$ ($gamma_1 < 0.1$) and classify it as positively (negatively) humped if the curvature factor γ_2 is above (below) 0.05 (-0.05).

	Positive Hump	No Hump	Negative Hump	Total
Upward Sloping	5.83%	7.97%	10.73%	24.54%
No Slope	14.05%	21.70%	14.68%	50.44%
Downward Sloping	10.66%	6.09%	8.28%	25.03%
Total	30.54%	35.76%	33.70%	100.00%

Table 3: Expected Returns on double sorted portfolios on beta and slope of term-structure

This table reports the expected returns on double sorted portfolios on the one-year option implied beta and the slope of the term-structure of beta. At the end of each month, stocks are ranked according to their one-year option implied beta and the slope of their term-structure and assigned to a portfolio. The portfolios are held for one year, such that each month, the weights are revised on 1/12 of the securities in the portfolio as in Jegadeesh and Titman (1993). Excess returns are reported in annual terms. The sample period ranges from January-1996 to August-2014.

Panel A: Average Excess returns					
	Neg. Slope	2	3	4	Pos. Slope
Low Beta	2.10%	1.64%	0.67%	-0.33%	-0.66%
2	4.83%	3.19%	2.11%	1.39%	1.96%
3	7.42%	4.35%	4.45%	4.03%	3.37%
4	6.24%	4.60%	4.17%	4.85%	4.35%
High Beta	8.57%	5.24%	6.03%	7.16%	12.36%

Panel B: t-statistics					
	Neg. Slope	2	3	4	Pos. Slope
Low Beta	1.38	0.78	0.24	1.71	0.66
2	1.58	1.84	1.91	2.38	1.58
3	2.00	2.29	1.68	2.47	1.42
4	2.21	2.49	1.96	1.94	1.18
High Beta	2.25	2.16	1.99	1.59	-0.49
GRS F-Test		1.58		P-Value	0.03

Table 4: Fama-MacBeth cross-sectional test

This table summarizes Fama-MacBeth cross-sectional regressions (average slopes, R^2s) when monthly returns are regressed on the Fama and French (1993) 3-factor model and the slope of the term-structure of the portfolio. T-statistics for the slopes are Newey-West corrected. Panel A reports the results of cross-sectional pricing of the double sorted portfolios on option-implied beta the slope of the term-structure (see table 3 for details). Panel B reports the same results for the 25 Fama-French size and book-to-market double sorted portfolios.

Panel A: Fama-MacBeth regressions using 25 portfolios sorted on Option implied beta and term-structure slope

		FF-Factors				
	Constant	Market	HML	SMB	Slope	R2
coef	-0.005	0.007	0.003	0.002	0.043	0.660
t-stat	-1.371	1.609	0.445	0.536	2.393	

Panel B: Fama-MacBeth regressions using 25 portfolios sorted size and book-to market

		FF-Factors					
	Constant	Market	HML	SMB	Slope	Level	R2
coef	0.010	-0.004	0.022	-0.002	0.033		0.41
t-stat	1.211	-0.441	4.406	-0.585	2.175		
coef	0.010	-0.003	0.022	-0.002	0.030	-0.039	0.42
t-stat	1.157	-0.390	4.034	-0.705	2.046	-1.284	
coef	0.016	-0.010	0.025	-0.001			0.22
t-stat	1.309	-1.324	5.179	-0.200			

Table 5: Calibration

This table lists the benchmark parameter values used to solve and simulate our model

Parameter		Value
Stochastic Discount Factor		
Subjective discount rate	β	0.99
Aggregate price of risk	γ_0	6
	γ_1	-1000
Technology		
Depreciation Rate	δ	0.01
Curvature of production function	α	0.7
Fixed costs of production	f	0.0032
Productivity		
Average aggregate productivity	$\bar{x}$	-3.55
Persistence of aggregate productivity	ρ_x	$0.95^{\frac{1}{3}}$
Persistence of idiosyncratic productivity	ρ_z	0.97
Productivity Jump	J_t	1.003
Adjustment costs		
Linear adjustment costs	a_0	0.01
	a_1	0.03
Convex Adjustment Costs	c^+	20
	c^-	200
Uncertainty Process		
Low volatility of aggregate productivity	σ_L^x	0.003
High volatility of aggregate productivity	σ_H^x	0.006
Low volatility of idiosyncratic productivity	σ_L^z	0.108
High volatility of idiosyncratic productivity	σ_H^z	0.217
Probability of low vol state given low vol	$\pi_{L,L}$	0.97
Probability of low high state given high vol	$\pi_{H,H}$	0.71

Table 6: Target Moments

This table reports unconditional moments from the simulated data and the real data. We simulate 2,500 artificial economics, each with 1.500 firms and 150 months. The average market return, volatility of market return and Sharpe ratio in the data are taken from Goyal and Welch (2008). The data moments of the real interest rate are from Campbell, Lo, MacKinlay, et al. (1997). The average output fall conditional on a volatility shock is from Bloom (2009). The remainder of moments are from Pontiff and Schall (1998).

	Data	Model
Average risk-free rate (%)	1.80	2.68
Volatility of risk-free rate (%)	3.00	2.01
Average market return (%)	6.33	6.49
Volatility of market return (%)	19.41	16.80
Sharpe Ratio	0.33	0.39
Annual average rate of investment	0.15	0.13
Volatility of firm-level investment rate (%)	22.30	12.01
Median market-to-book ratio	1.49	1.53

Table 7: Properties of market-to-book deciles in the Model

In total 2,500 artificial panels are simulated from the model in Section 4. Each panel contains 1,500 firms and 150 months. The Ten book-to-market deciles are constructed on each of the artificial panels, the CAPM regressions are performed (third row from the table), and the cross-sample averaged results are reported. We also investigate the properties of the option implied CAPM (fourth row). We compute option implied betas using the same methodology from the empirical section and then report cross-sectional averages.

	Growth	2	3	4	5	6	7	8	9	Value
$\mathrm{E}[R^j - R^f]$	0.048	0.05	0.053	0.057	0.064	0.072	0.079	0.087	0.091	0.096
$\sigma[R^j - R^f]$	0.18	0.19	0.21	0.22	0.23	0.24	0.24	0.25	0.25	0.26
Sharpe Ratio	0.32	0.34	0.37	0.41	0.44	0.47	0.49	0.5	0.52	0.54
Standard CAPM Beta	0.88	0.95	1.02	0.98	1.02	0.99	1.01	1.02	1.03	1.08
Option Implied Beta	0.24	0.60	0.77	0.91	1.02	1.13	1.24	1.36	1.50	1.75

References

AÏT-SAHALIA, Y., M. KARAMAN, AND L. MANCINI (2014): "The term structure of variance swaps and risk premia," *Working Paper*.

AN, B.-J., A. ANG, T. G. BALI, AND N. CAKICI (2014): "The joint cross section of stocks and options," *The Journal of Finance*, 69(5), 2279–2337.

ANDERSEN, T. G., T. BOLLERSLEV, F. X. DIEBOLD, AND J. WU (2006): "Realized beta: Persistence and predictability," *Advances in Econometrics*, 20(1), 39.

BAKSHI, G., N. KAPADIA, AND D. MADAN (2003): "Stock return characteristics, skew laws, and the differential pricing of individual equity options," *Review of Financial Studies*, 16(1), 101–143.

BANSAL, R., AND A. YARON (2004): "Risks for the long run: A potential resolution of asset pricing puzzles," *The Journal of Finance*, 59(4), 1481–1509.

BLOOM, N. (2009): "The impact of uncertainty shocks," *econometrica*, 77(3), 623–685.

BOLLERSLEV, T., G. TAUCHEN, AND H. ZHOU (2009): "Expected stock returns and variance risk premia," *Review of Financial Studies*, 22(11), 4463–4492.

BUSS, A., AND G. VILKOV (2012): "Measuring equity risk with option-implied correlations," *Review of Financial Studies*, 25(10), 3113–3140.

CAMPBELL, J. Y., AND J. H. COCHRANE (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107(2).

CAMPBELL, J. Y., A. W.-C. LO, A. C. MACKINLAY, ET AL. (1997): *The econometrics of financial markets*, vol. 2. princeton University press Princeton, NJ.

CARR, P., AND L. WU (2009): "Variance risk premiums," *Review of Financial Studies*, 22(3), 1311–1341.

COOLEY, T. F., AND E. C. PRESCOTT (1995): "Economic growth and business cycles," *Frontiers of business cycle research*, 1.

DEMETERFI, K., E. DERMAN, M. KAMAL, AND J. ZOU (1999): "A guide to volatility and variance swaps," *The Journal of Derivatives*, 6(4), 9–32.

Den Haan, W. J., et al. (1997): "Solving dynamic models with aggregate shocks and heterogeneous agents," *Macroeconomic dynamics*, 1, 355–386.

Dew-Becker, I., S. Giglio, A. Le, and M. Rodriguez (2015): "The price of variance risk," Discussion paper, National Bureau of Economic Research.

Diebold, F. X., and C. Li (2006): "Forecasting the term structure of government bond yields," *Journal of Econometrics*, 130(2), 337–364.

Driessen, J., P. J. Maenhout, and G. Vilkov (2009): "The price of correlation risk: Evidence from equity options," *The Journal of Finance*, 64(3), 1377–1406.

Egloff, D., M. Leippold, and L. Wu (2010): "The Term Structure of Variance Swap Rates and Optimal Variance Swap Investments," *Journal of Financial and Quantitative Analysis*, 45(5), 1279–1310.

Fama, E. F., and K. R. French (1993): "Common risk factors in the returns on stocks and bonds," *Journal of Financial Economics*, 33(1), 3–56.

Ferson, W. E., and C. R. Harvey (1991): "The variation of economic risk premiums," *Journal of Political Economy*, pp. 385–415.

Gibbons, M. R., S. A. Ross, and J. Shanken (1989): "A test of the efficiency of a given portfolio," *Econometrica*, pp. 1121–1152.

Gomes, J., L. Kogan, and L. Zhang (2003): "Equilibrium Cross Section of Returns," *Journal of Political Economy*, 111(4).

Goyal, A., and I. Welch (2008): "A comprehensive look at the empirical performance of equity premium prediction," *Review of Financial Studies*, 21(4), 1455–1508.

Han, B., and Y. Zhou (2012): "Variance risk premium and cross-section of stock returns," *unpublished paper, University of Texas at Austin.*

Hennessy, C. A., and T. M. Whited (2007): "How costly is external financing? Evidence from a structural estimation," *The Journal of Finance*, 62(4), 1705–1745.

Hollstein, F., and M. Prokopczuk (2014): "Estimating Beta," *Journal of Financial and Quantitative Analysis, Forthcoming.*

Jegadeesh, N., and S. Titman (1993): "Returns to buying winners and selling losers: Implications for stock market efficiency," *The Journal of Finance*, 48(1), 65–91.

Kiku, D. (2006): "Is the value premium a puzzle?," *Manuscript, Wharton Business School, University of Pennslyvania, Philadelphia.*

Krusell, P., and A. A. Smith, Jr (1998): "Income and wealth heterogeneity in the macroeconomy," *Journal of political Economy*, 106(5), 867–896.

Lettau, M., and J. A. Wachter (2007): "Why is long-horizon equity less risky? a duration-based explanation of the value premium," *The Journal of Finance*, 62(1), 55–92.

Lewellen, J., S. Nagel, and J. Shanken (2010): "A skeptical appraisal of asset pricing tests," *Journal of Financial Economics*, 96(2), 175–194.

Lin, X., and L. Zhang (2013): "The investment manifesto," *Journal of Monetary Economics*, 60(3), 351–366.

Pontiff, J., and L. D. Schall (1998): "Book-to-market ratios as predictors of market returns," *Journal of Financial Economics*, 49(2), 141–160.

Van Binsbergen, J., M. Brandt, and R. Koijen (2012): "On the Timing and Pricing of Dividends," *The American Economic Review*, pp. 1596–1618.

van Binsbergen, J., W. Hueskes, R. Koijen, and E. Vrugt (2013): "Equity yields," *Journal of Financial Economics*, 110(3), 503–519.

Van Binsbergen, J. H., and R. S. Koijen (2015): "The Term Structure of Returns: Facts and Theory," *Working Paper.*

Zhang, L. (2005): "The value premium," *The Journal of Finance*, 60(1), 67–103.

Appendix A Term-structure of betas: level, slope and curvature

On this appendix we show that three factors are enough to capture the shapes of the term-structure of betas. For every stock in our sample we model its option implied beta term-structure using Nelson-Siegel exponential components model which imposes a structure on factor loadings thus reducing the estimation error. At the end of each month t we fit the following curve to the estimated option implied betas:

$$\beta^{Q}_{iM,t,\tau} = \gamma_{0,t} + \gamma_{1,t} e^{-\tau/\theta} + \gamma_{2,t} \frac{T}{\theta} e^{-\tau/\theta} \tag{21}$$

We conjecture that in line with Diebold and Li (2006) we may interpret the γ coefficients as three latent dynamic factors and interpret them as the level, slope and curvature of the term structure.

For long-term maturities the beta estimates approach asymptotically γ_0; then γ_1 represents the deviation from the asymptote; and γ_2 determines the hump that happens at time τ. The parameter θ governs the decay, so a high (low) value of θ allows for a better fit for short (long) maturities. Following Diebold and Li (2006) we fix θ to maximize the loading of the medium-term factor at six months which is when the hump occurs on average. This allows to compute the values of γ using ordinary least squares (OLS). The results are robust to the value of θ chosen as long as it lies in a reasonable range (i.e. allowing for the hump to occur at around 6 months).

The Nelson-Siegel model is flexible enough to capture different shapes of the term-structure (upward sloping, downward sloping and hump-shaped). Figure **??**, on the main text illustrates this by plotting the model fit at two different points in time. For all stocks in our sample the model works fairly well. In table 8 we report the residual statistics from the in-sample estimation of equation 21 for all the stocks. We compute each statistic for each stock and then average them across stocks. The average error is fairly low ranging from -0.004 to 0.003, implying than on average the Nelson-Siegel model is properly capturing the shape of the term-structure. Further, the one-month auto-correlation of the errors increase with maturity meaning that if we misprice the long-term beta on one period we are likely to misprice it on the next period. This error persistence on the highest maturities might be due to the lack of liquidity of the long-term options. More interestingly, the maximum and minimum pricing errors (third and fourth column of table 8) are very low - given that we estimate a term-structure for each point in time for all S&P 500 stocks one would expect higher maximum (minimum) pricing errors. Furthermore, RMSE are low and do not have a pattern across maturities implying that the moment when the hump on the term-structure occurs might change but the model, on average, is able to accurately capture.

The three estimated γ can be interpreted in terms of level, slope and curvature of the betas term-structure. Define the the level factor, γ_0, as the long-term beta ($\beta^Q_{iM,t,t+12m}$); the slope, γ_1, as the difference between the twelve-month implied beta and the one-month implied beta ($\beta^Q_{iM,t,t+12m} - \beta^Q_{iM,t,t+1m}$); and the curvature, γ_2, as the difference between twice the three-month implied beta and the sum of the one-month implied beta with the twelve-month implied beta ($2 \times \beta^Q_{iM,t,t+3m} - (\beta^Q_{iM,t,t+1m} + \beta^Q_{iM,t,t+12m})$). Figure 11 plots the estimated level, slope and curvature of the Apple Corporation stock across time against its empirical counterpart. The figure confirms our conjecture that the three factors in our model can indeed be interpreted as the slope, level and curvature of the curve. This result is robust for all other stocks in our sample.

Figure 11: Level, Slope and Curvature of Term-Structure

This figure plots the model based level, slope and curvature (γ_0, γ_1 and γ_2) against their empirical counterparts for the Apple Corporation. We define the empirical level as the long-term beta ($\beta^Q_{iM,t,t+24m}$); the slope, γ_1, as the difference between the twenty-four-month implied beta and the one-month implied beta ($\beta^Q_{iM,t,t+24m} - \beta^Q_{iM,t,t+1m}$); and the curvature, γ_2, as the difference between twice the three-month implied beta and the sum of the one-month implied beta with the twenty-four-month implied beta ($2 \times \beta^Q_{iM,t,t+3m} - (\beta^Q_{iM,t,t+1m} + \beta^Q_{iM,t,t+24m})$).

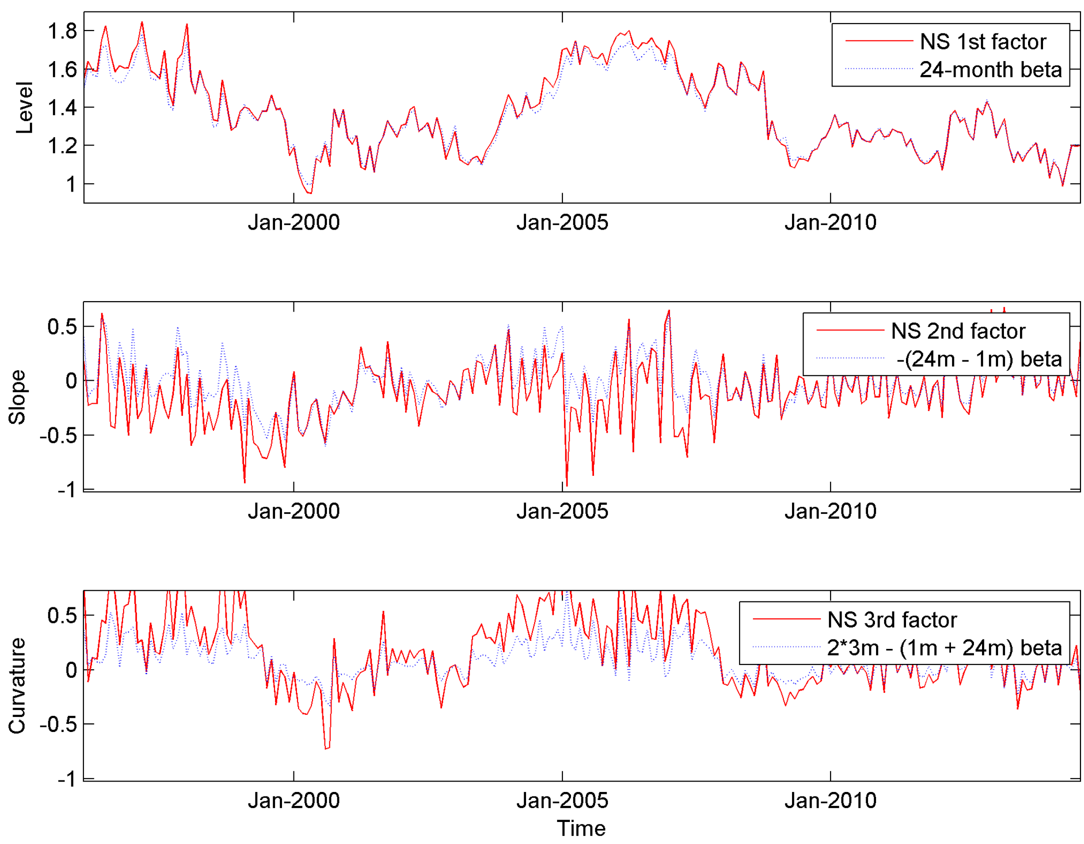

Table 8: Nelson-Siegel In-Sample Performance

This table presents the in-sample performance of the Nelson-Siegel model in fitting the option implied betas term-structure. At the end of each month we fit equation 21 to the estimated option implied betas. The table reports the residual statistics for all the S&P 500 stocks as an average across stocks. The residual at time t for the estimated beta with maturity τ is defined as: $\hat{e}_t = \beta^Q_{iM,t,\tau} - \hat{\beta}^Q_{iM,t,\tau}$, where $\hat{\beta}^Q_{iM,t,\tau}$ is the Nelson-Siegel fit. The first two columns present the average residuals for each maturity and their standard deviation. The third and fourth columns present the maximum and minimum pricing errors. The RMSE is a performance measure defined as: RMSE $= \sqrt{\frac{1}{n}\sum_{t=1}^{n}(\beta^Q_{iM,t,\tau} - \hat{\beta}^Q_{iM,t,\tau})^2}$. ρ_1 and ρ_{12} are the order one and twelve auto-correlation coefficients respectively.

	Average	St. Dev	Max	Min	RMSE	ρ_1	ρ_{12}
1M	-0.002	0.018	0.052	-0.068	0.017	0.105	0.157
2M	0.002	0.037	0.127	-0.124	0.034	-0.038	0.196
3M	0.003	0.028	0.101	-0.078	0.026	0.015	0.239
4M	0.001	0.016	0.059	-0.052	0.015	-0.044	0.221
5M	-0.002	0.017	0.053	-0.061	0.015	0.040	0.180
6M	-0.003	0.021	0.056	-0.080	0.020	0.278	0.130
9M	-0.004	0.025	0.060	-0.101	0.027	0.529	0.142
12M	0.000	0.018	0.050	-0.070	0.019	0.531	0.207
18M	0.003	0.015	0.064	-0.040	0.016	0.437	0.146
24M	0.002	0.025	0.096	-0.062	0.028	0.567	0.142

Appendix B

Table 9: Fama-MacBeth cross-sectional test - excluding financial crisis period

This table summarizes Fama-MacBeth cross-sectional regressions (average slopes, R^2s) when monthly returns are regressed on the Fama and French (1993) 3-factor model and the slope of the term-structure of the portfolio. T-statistics for the slopes are Newey-West corrected. Panel A reports the results of cross-sectional pricing of the double sorted portfolios on option-implied beta the slope of the term-structure (see table 3 for details). Panel B reports the same results for the 25 Fama-French size and book-to-market double sorted portfolios. We cut the sample in January 2007 to exclude the financial crisis period.

Panel A: Fama-MacBeth regressions using 25 portfolios sorted on Option implied beta and term-structure slope

		FF-Factors				
	Constant	Market	HML	SMB	Slope	R2
coef	-0.008	0.007	0.008	0.004	0.023	0.680
t-stat	-1.403	1.224	0.794	0.750	1.790	

Panel B: Fama-MacBeth regressions using 25 portfolios sorted size and book-to market

		FF-Factors					
	Constant	Market	HML	SMB	Slope	Level	R2
coef	0.007	-0.003	0.023	-0.002			0.23
t-stat	1.196	-0.343	3.811	-0.563			
coef	0.002	0.002	0.019	-0.003	0.036		0.38
t-stat	0.278	0.307	3.033	-0.721	2.404		
coef	0.001	-0.003	0.016	-0.003	0.047	-0.040	0.38
t-stat	0.127	0.427	2.433	-0.721	2.944	-1.452	

Appendix C Other time-series facts regarding the term-structure of betas

In this appendix we show some other empirical facts regarding the term-structure of betas, namely we document the behavior of the slope of the term-structure at the aggregate level and at the individual firm-level. If the slope of the term-structure indeed captures changes in the underlying risk of stocks and portfolios, we should expect higher changes around crisis periods, major economic events, and expect these changes to provide information about which sectors of the economy are more risky at a given point in time. For individual stocks changes in the slope of the term-structure should also spike following corporate relevant events. In this appendix we show that the change in the slope of the term-structure is rather informative about the riskiness of the underlying stock/portfolio. For the aggregate market, absolute value-weighted changes of the slope of the term-structure are higher during systemic events such as FOMC meetings, the Lehman Brothers bankruptcy, the Asian mini-crash, DJIA biggest one-day crash, the Boston marathon disaster, among others. For individual stocks, the changes in slope are more significant in company specific relevant events such as mergers and acquisitions, ratings cut, bankruptcy announcements, company bailouts, major market disruptions, etc.

To assess whether the how the slope of the term-structure behaves at the aggregate level we define a market risk measure (RM_t) as:

$$RM_t = \sum_{i=1}^{n} w_i |\Delta\gamma_{1,t+1,i}| \tag{22}$$

where $\Delta\gamma_{1,t+1,i}$ is the change in slope of the term-structure of stock i between date t and $t{+}1$ and w_i is the market weight of stock i.[21] Given that the value-weighted changes on the slope of all the market constituents must add up to zero, the absolute value in equation 22, ensures that we can capture the points in time where slope changes are higher. The slope of the term-structure can change either due to: (i) changes in the short-term beta, (ii) changes in long-term beta, (iii) changes in both. Usually, the changes in short-term beta are more pronounced and account for most of the changes in the slope. However, the changes in long-term beta act frequently as a shrinking parameter. An increase in both short and long-term beta means that the underlying portfolio significantly increased its expected risk whereas an increase in short-term together with a decrease in long-term beta implies that the underlying increased its short-term expected risk but decreased its long-term expected risk.

[21] On this section we use the empirical slope instead of the Nelson-Siegel estimated slope on our computations. The results are not sensitive to the use of one or the other.

Figure 12 plots the daily time-series of equation 22 between 1996-2014. During normal days, the risk reallocation measure fluctuates around its time-series average of 0.032. However, on days with relevant events - either macroeconomic shocks or stock market crashes - it spikes implying a high degree of risk-reallocation. It is clear that RM_t was clearly above its time series average during turbulent periods such as the technology bubble or the financial crisis and there are clear spikes on the asian crash in 1997, the dot-com bubble in 2000, the Lehman bankruptcy in 2008 and two large DJIA (Dow Jones Industrial Average) crashes. Also RM_t clearly spikes on important macroeconomic events such as FOMC meetings, the congress debt ceiling debate and the Greek crisis. At the individual stock level, the measure spikes, as well, on relevant firm specific events. Table 10 reports the 20 largest changes in the term-structure slope of all S&P 500 stocks on a given month (we average the daily absolute changes on the slope of individual stocks to get a monthly figure and check which individual stock had the highest average change on the slope). Two patterns emerge (i) the slope of the term-structure significantly changes following a merger, acquisition or spin-off and (ii) economic and financial distress or major market disruptions also lead to significant changes in the slope.

Figure 12: Risk-reallocation: Value-weighted absolute changes in beta term-structure

This figure plots the daily time-series of equation $RM_t = \sum_{i=1}^{n} w_i |\Delta\gamma_{1,t+1,i}|$ where $\Delta\gamma_{1,t+1,i}$ is the change in slope of the term-structure of stock i between date t and $t+1$ and w_i is the market weight of stock i. The sample ranges from January-1996 until August-2014.

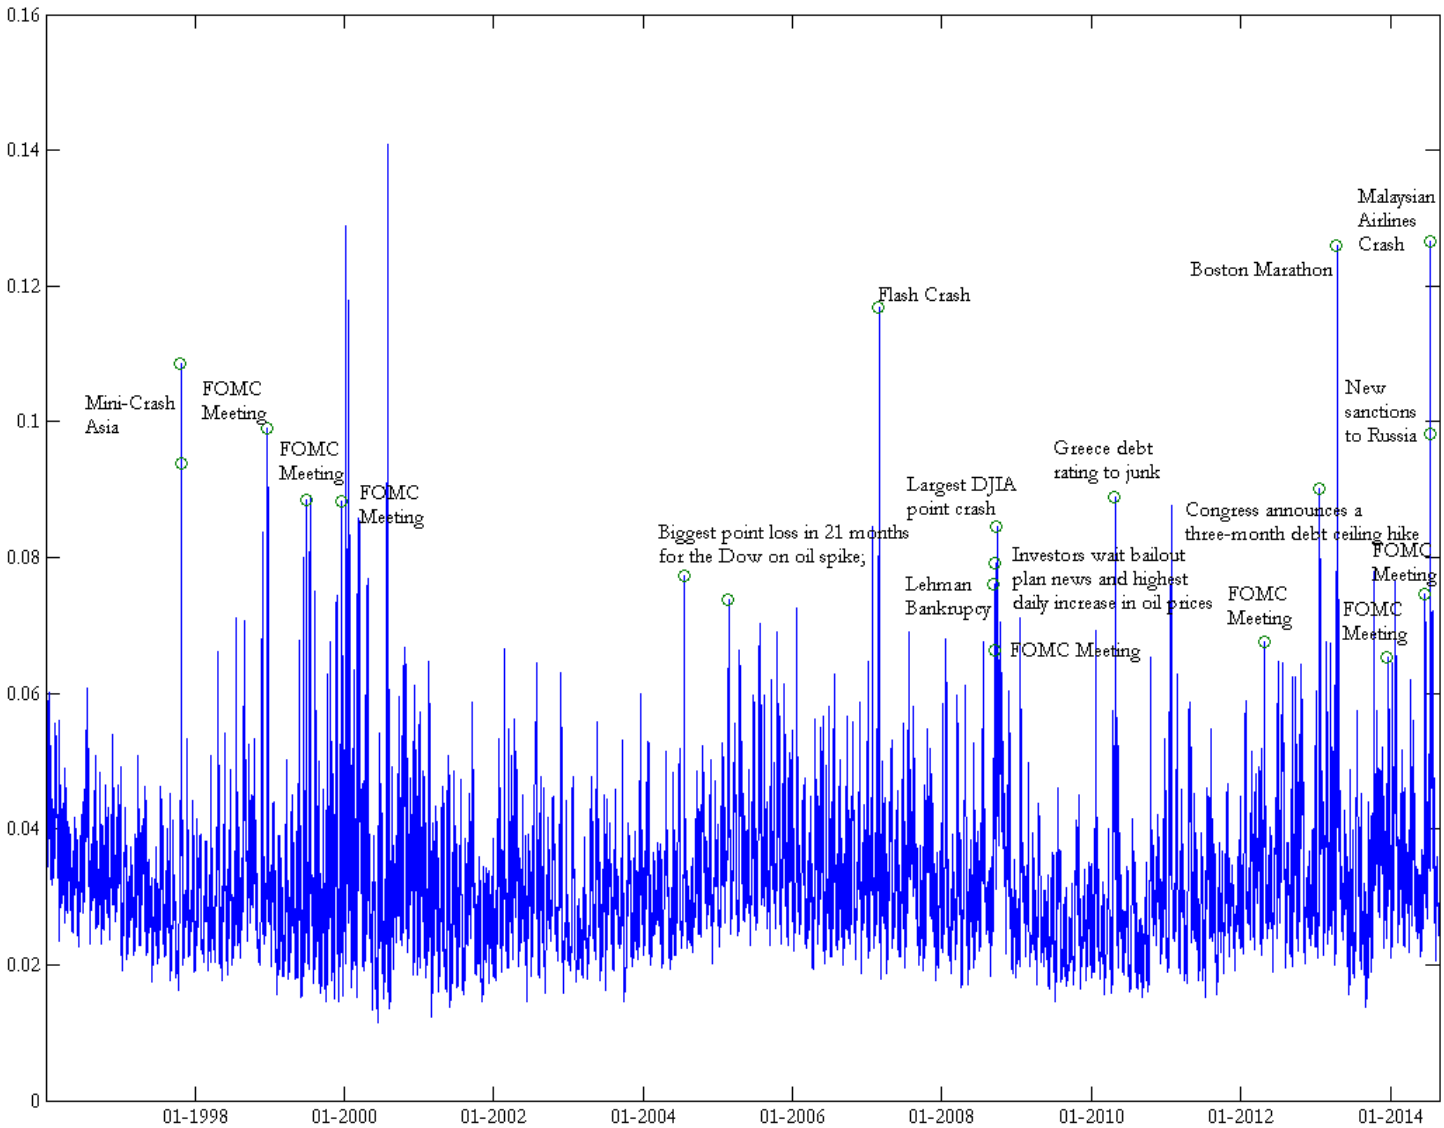

Table 10: Individual stocks risk-reallocation

This table reports 20 of the largest absolute changes in the term-structure slope of individual S&P 500 stocks on a given month. Each month we average the daily absolute changes in the term-structure slope of each S&P 500 stock and check which stock had the largest change. The first column reports the month under analysis, the second column reports the company with the largest absolute change in slope on that month, and the last column reports key events related with the company on that month.

Date	Company	Event
Jul-1997	General Instrument	Spin off of two major divisions
Aug-1997	Txu Gas Co	Merger with ENSERCH Corp
Sep-1997	Sunamerica Inc	Discloses intention to acquire a major competitor
Aug-2001	Quaker Oats Co	Acquired by Pepsi Co
Jul-2002	Qwest Communication Intl	Admits it has used improper accounting methods (26% of market value wiped off)
Sep-2002	William Cos	Almost bankrupt. Accepts an emergency high interest loan from Warren Buffet
Jan-2006	Mercyry Interactive	Fell to the Pink Sheets
Nov-2006	Amsouth Bancorporation	Acquired by Regions Financial Corporation
Dec-2006	Freescale Semiconductor Inc	Largest buyout of a technology firm
Feb-2007	Integrys Energy Group Inc	Moody downgrades its rating
Mar-2008	Bear Stearns Companies Inc	Federal Reserve Bank of New York agreed to provide a $25 billion liquidity loan
Jun-2008	Mbia Inc	Several lawsuits were filed against MBIA Inc
Jul-2008	Federal Home Loan Mortg corp	Government attempted to ease market fears by reiterating their view that company plays a central role in the US housing finance system
Sep-2008	Lehman Brothers holdings inc	Filled for Bankruptcy
Apr-2009	Rohm and Haas	Acquired by The Dow Chemical Company
Aug-2012	First Solar	Halted panel deliveries to the worlds largest photovoltaic power plant
Oct-2013	Dell inc	Privatization deal completed
Apr-2014	Cliffs Natural Resources inc	Proxy fight
Jun-2014	Vertex Pharmaceuticals inc	FDA drug approval stock jumps 40%

Appendix D Term-structure slope of Value/Growth portfolios and Financial/Technology stocks

Figure 13: Beta term-structure slope

This figure plots a monthly time-series of the slope of the term-structure between January 1996 to August 2014. Panel A reports the results for value (high book-to-market) and growth (low book-to-market) stocks. We sort stocks on five portfolios based on their book-to-market following the same methodology of Fama and French (1993). Panel B reports the results for financial firms and technology stocks. As industry definition we use the 12 industry portfolios available at Kenneth French's website.

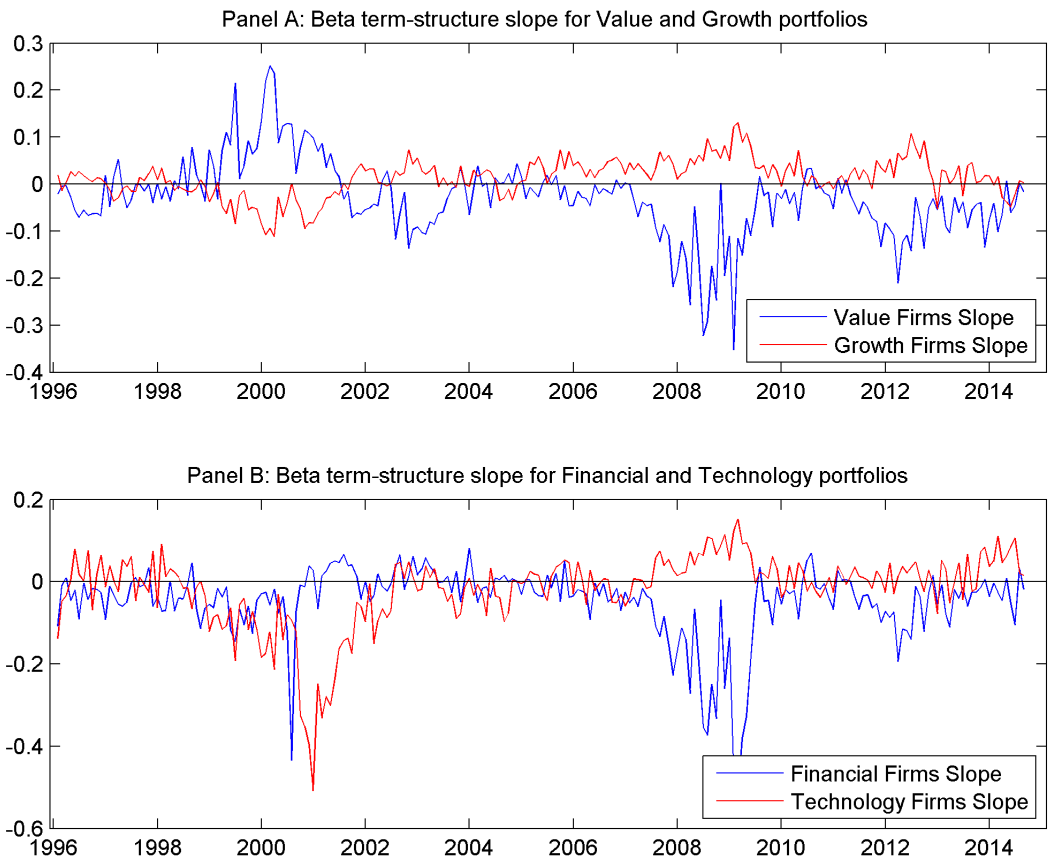

Appendix E VIX and term-structure of betas

In this appendix we discuss how does the average cross-sectional slope of these term-structures relate with the market level of uncertainty or VIX. This is an important question but with no straightforward answer. To begin with, the market slope of systematic risk is a flat line, whereas the VIX fluctuates over time. Bearing this in mind, we define the following measure to capture the average slope of the term-structure:

$$AverageSlope_t = \sum_{i=1}^{n} w_i |\gamma_{1,t+1,i}| \tag{23}$$

The modulus in equation (23) ensures that we can capture how much on average do term-structures slope. Figure 6 plots the VIX index against the time-series of equation (23). From the figure it is clear that the levels of the VIX and average term-structure slope are not related.

[Figure 6 here]

However, changes in VIX and changes in the average slope have a daily correlation of -0.48 (with a t-stat of -38.0) which implies that on average when uncertainty rises term-structures slope increases, whereas when uncertainty decreases the term-structures get less sloped. This correlation is almost mechanical as the VIX index shows up on the denominator of the short-term option implied beta. Also, VIX is an aggregate measure of uncertainty or risk-aversion of the market. The advantage of looking at the slope of term-structures is that we can do it for any individual security at any given point in time. So despite the fact overall market uncertainty plays and important role in the term-structure of systematic risk, it proxies a different thing.