Bond Variance Risk Premiums*

Hoyong Choi
Rotterdam School of Management, Erasmus University†

Philippe Mueller
London School of Economics‡

Andrea Vedolin
London School of Economics§

Abstract

This paper studies variance risk premiums in the Treasury market. We first develop a theory to price variance swaps and show that the realized variance can be perfectly replicated by a static position in Treasury futures options and a dynamic position in the underlying. Pricing and hedging is robust even if the underlying jumps. Using a large options panel data set on Treasury futures with different tenors, we report the following findings: First, the term structure of implied variances is downward sloping across maturities and increases in tenors. Moreover, the slope of the term structure is strongly linked to economic activity. Second, returns to the Treasury variance swap are negative and economically large. Shorting a variance swap produces an annualized Sharpe ratio of almost two and the associated returns cannot be explained by standard risk factors. Finally, the returns remain highly statistically significant even when accounting for transaction costs and margin requirements.

First Version: November 2010
This Version: November 2016

JEL code: G12.

*We would like to thank Dante Amengual, Torben Andersen, Ryan Carlson, Mike Chernov, Andrea Eisfeldt, Jean-Sébastien Fontaine, Anh Le, Ian Martin, André Siegrist, Alireza Tahbaz-Salehi, Roméo Tédongap, Anders Trolle, Grigory Vilkov, Hao Zhou and seminar participants at BlackRock, University of Bern, Bank of Canada, Bank of England, Copenhagen Business School, European Central Bank, Econometric Society Meeting Chicago, EFA Stockholm, ESEM Oslo, 2011 Hedge Fund Conference at Imperial College, IE Madrid, Inquire Europe, London School of Economics, Lund University, New York Fed, Oxford-Man Institute, San Francisco Fed, SFS Cavalcade Virginia Darden, Structured Portfolio Management, and the Workshop on Advances in Fixed Income Modeling at the Bank of Spain for useful comments.

†Department of Finance, Email: choi@rsm.nl

‡Department of Finance, Email: p.mueller@lse.ac.uk

§Department of Finance, Email: a.vedolin@lse.ac.uk

Markets for volatility derivatives have grown a lot over the last decade. Nowadays, investors have both exchange-traded and over-the-counter instruments available to hedge and trade volatility in equity markets. While a plethora of research has focused on understanding and pricing equity variance risk, the same risk is much less well understood in fixed income markets.

In this paper, we propose the generalized Treasury variance swap (GTVS) which offers a pure exposure to fixed income variance. We theoretically derive the fair value of the contract and empirically document significant returns to variance trading in Treasury markets that are comparable to those earned in the equity variance market. We quantitatively compare our GTVS strategy to alternative volatility trading strategies based on fixed income options that may not leave the investor with a clean volatility exposure. Calculating the variance swap rate for various horizons, we obtain a term structure of Treasury implied variances and ex post bond variance risk premiums.[1] Finally, we show that the term structure of Treasury implied variances is significantly related to economic activity and stress indicators for financial markets.

The main contribution of this paper is twofold. First, we study the theoretical properties of variance swaps in Treasury markets. Different from standard variance contracts, our strategy is model-free and allows for stochastic interest rates. Variance swaps consist of two legs: (i) a realized leg and (ii) the fair strike. The latter is defined as the strike that makes the net present value of the swap equal to zero at initiation. Building on the pioneering work of Neuberger (1994) and Demeterfi, Derman, Kamal, and Zhou (1999), we show that the fair strike can easily be approximated using a portfolio of puts and calls. The way we define the realized leg is crucial. For the generalized Treasury variance swap we use a generalized measure of realized variance which allows for perfect replication of the contract even in the presence of jumps (see Neuberger (2012) for an application in the equity market). This property is particularly useful in face of the recent extreme events as standard variance swaps (the log Treasury variance swap or LTVS) rely on squared log returns and are therefore exposed to cubed returns, resulting in extremely inefficient hedges for investors.

[1]Treasury Implied Volatility (TIV) data is available from the authors' webpages.

Second, using a large panel data set of daily option prices on Treasury futures with different tenors, we study the payoffs of the GTVS and the associated ex post variance risk premiums, defined as the difference between the realized variance and the fair strike or variance swap rate. Our main findings can be summarized as follows: Consistent with the literature in the equity market we find that the variance risk premiums are negative and economically large. The average excess returns for a strategy that shorts the variance swap with a one-month maturity, independent of the tenor of the underlying, is around 20% per month (21.2% for the 30 year Treasury futures (t-statistic of 8.98), 27.6% for the 10 year Treasury futures (t-statistic of 12.71) and 18.7% for the 5 year Treasury futures (t-statistic of 6.61), respectively) and the associated annual Sharpe ratio is just below two.[2] Since our theory is based on (European) options on forwards, we need to convert the (American) options on futures to European options on forwards. We show that the adjustment is overall very small, especially for options with maturities of less than one year.

Traditionally, the most common strategy to exploit variance is to invest in a delta-hedged at-the-money straddle. A drawback of straddles compared to variance swaps is that the sensitivity of the trading strategy with respect to volatility or variance is non-linear. Moreover, the sensitivity to other factors is typically non-zero (volatility of volatility effect). We find that trading a one-month delta-hedged at-the-money straddle generates significantly smaller rewards compared to a position in a one-month variance swap over our sample period: The average return on the 10y Treasury futures straddle is around -7% (with an associated t-statistic of 4.91) and the annualized Sharpe ratio is around 1.05, which is about half in size compared to the GTVS.

We then study whether the alpha of these strategies can be explained by common factors found to explain variance trading strategies. Not surprisingly, we find that the market return explains very little of the excess returns on variance swaps. However, even if we include a battery of additional factors such as size, book-to-market, or bond market liquidity, the time-series variation of these risk premiums remains largely unexplained.

[2] The excess return is calculated as a simple return, meaning it is the variance risk premium normalized by the fair strike price or the implied variance.

Moreover, the alpha of the variance swap strategy continues to be highly significant. For example, for the 30y Treasury futures, we find that the alpha is 26% per month, even when including all controls.

As a by-product of our analysis, we construct the unconditional term structure of implied variances and variance risk premiums. The unconditional variance risk premiums are obtained as the unconditional average of the difference between the realized variance under the physical measure and the expected variance under the risk-neutral measure. We find that, on average and for all tenors, the term structure of implied variances is downward sloping. The slope, defined as the difference between the implied variance of a one-year and a one-month option is strongly pro-cyclical. During crisis periods, the slope becomes extremely negative. Moreover, the slope has strong predictive power for future economic growth and proxies of economic stress, especially at short horizons: A steeper slope predicts higher growth or lower stress up to eight months ahead. When we add the slope of the term structure of Treasury yields, which itself is considered a good predictor of future growth, we find that the predictive power of the implied variance slope remains unchanged. Moreover, at short horizons, the slope of the yield curve has no power at all.

The term structure of variance risk premiums is also downward sloping in absolute terms. Expressed in monthly squared percent, the variance risk premiums for the 30y Treasury futures are ranging from -4.9 for a ten-day horizon to essentially zero for the one-year horizon. For 10y Treasury futures, the short- and longer-term variance risk premiums are -3.1 and -0.4, while for the 5y Treasury futures they are -1.1 and -0.4, respectively. Note that unlike the excess returns on the generalized Treasury variance swap, the variance risk premiums are not normalized by the implied variance and they are, thus, distinctly increasing in the variance level (which is increasing in the underlying tenors).

We run different robustness checks to challenge our findings. First, one might be worried about statistical significance as option returns are known to be non Gaussian. We therefore use a studentized bootstrap to obtain confidence intervals on the means, alphas, and Sharpe ratios of our trading strategies. We find all performance measures to be statistically different from zero.

Second, we also study the impact of bid and ask spreads on the profitability of the variance swap strategy and find that while average returns decrease by around 3%, shorting variance is still a very attractive strategy: Annualized Sharpe ratios remain above one and the alpha is highly significant.

Third, margin requirements limit the notional amount of capital that can be invested in trading strategies. Moreover, large losses in a position can force investors out of a trade, potentially at the worst possible time. Thus, in order to realistically assess the profitability of our proposed variance swap trading strategy we take into account the impact of margins on the realized returns. Earlier literature suggests that margin requirements can have an economically significant effect when investors do not have access to unlimited capital when the market is in a downturn. We confirm that there is a difference between average returns for an unrestricted and a margined variance swap strategy but we find that average returns and Sharpe ratios remain significant and economically large.

Fourth, we ask to what extent it matters how the variance swap is constructed. As mentioned, a perfect replication of the variance swap that is robust to jumps in the underlying and the choice of the re-balancing frequency requires a particular definition to measure the realized variance. This definition differs from the standard way to calculate realized variance, namely to use squared log returns. On average, we find that the payoffs of the GTVS and those from a contract based on realized variance measured using log returns (LTVS) are very similar. This stems from the fact that positive and negative jumps cancel each other out. The devil is in the details, however. We show that whenever there are jumps in the underlying, regardless of whether they are positive or negative, the returns to the two trading strategies are distinctly different. Negative (positive) jumps render the payoff to the LTVS larger (smaller) compared to the GTVS. More importantly, while the payoff to the GTVS can still be perfectly hedged, this is no longer the case for the LTVS. This is in line with the findings of Broadie and Jain (2008) who study variance swaps in the presence of jumps in the equity index market and document that negative jumps have the most severe impact on replication errors.

Finally, one might wonder whether it matters for the profitability of our trading strategies whether we include crisis periods or not. Hence, a natural additional robustness check is to study variance swap returns for different time periods. We confirm that trading strategies that go short variance during periods that include the October 1987 crash or the height of the credit crisis in August 2008 are still very profitable over time. Overall, we find almost no quantitatively relevant differences across various sub-periods.

Related Literature: Our paper draws from the large literature on variance trading in equity markets starting with the work of Dupire (1994) and Neuberger (1994). For example, Carr and Wu (2009) use portfolios of options to approximate the value of the variance swap rate for different stock indices and individual stocks. They then compare the synthetic variance swap rates to the ex post realized variance to determine the size of the variance risk premium. Wu (2010) estimates variance risk dynamics by combining the information in realized variance estimators from high frequency returns and the VIX. Egloff, Leippold, and Wu (2010) directly use variance swap quotes and study the term structure of variance swap rates. Motivated by the recent financial crisis, much attention has been paid to the pricing and hedging of equity variance swaps in the presence of jumps. Both Schneider (2015) and Schneider and Trojani (2015) study tradeable properties of volatility risk, where the latter focus on higher-order risk premiums attached to time-varying disaster risk.

Other papers investigate the term structure of variance risk premiums and prices in the equity index market. Dew-Becker, Giglio, Le, and Rodriguez (2016) estimate an affine term structure model using variance swap data and find that realized volatility is the only priced risk factor which implies a term structure that is steeply negative at the short-end but flat beyond a one-month maturity. They conclude that these stylized facts are hard to reconcile within standard asset pricing models. Similarly, Andries, Eisenbach, Schmalz, and Wang (2015) study the term structure of variance risk premiums and find that a model where investors feature horizon-dependent risk aversion matches the data well. Our paper is different from this strand of the literature, as our approach is completely model-free and we do not take a stand on the microfoundations.

We are not the first to explore variance contracts in fixed income markets. Trolle (2009) estimates variance risk premiums in two ways: First, he estimates a dynamic term structure model that allows for unspanned stochastic volatility and, second, he corroborates his findings using a model-independent approach similar to ours. Both approaches lead him to the conclusion that the market price of variance risk is highly negative. While his focus is on a dynamic portfolio choice problem which includes interest rate derivatives, his approach also differs from ours as he derives the model-free variance risk premiums under different assumptions. Furthermore, he does not study the term structure of variance risk premiums. Trolle and Schwartz (2014) study variance and skewness across different swap maturities and option tenors in the swaptions market. They find that both variance and skewness risk premiums are negative and highly time-varying. The authors then propose a dynamic term structure model that fits the dynamics of these risk premiums. We see our work complimentary to theirs as our focus is on documenting empirical facts about a variance trading strategy rather than asking what model is best suited to capture the dynamic behavior of conditional swap rate moments. Recently, Cieslak and Povala (2016) study yield volatility risk and suggest that investors willingness to pay large premiums to hedge volatility can be linked to uncertainty about the future path of monetary policy.

In contemporaneous theoretical work, Mele and Obayashi (2013) explore variance contracts on Treasury futures similar to ours. There are important differences, however. While their fair strike is constructed in the same way than ours, their definition of the realized leg is different as it is constructed using log returns. In our paper, we show that using realized variance based on log returns leaves an exposure to cubed returns. This has important consequences during periods where the underlying exhibits jumps. Moreover, we also conduct an empirical analysis of the proposed contract while they are mainly interested in the theoretical derivation of the option-implied leg. Finally, Merener (2012) constructs a variance strategy on forward swap rates and studies dynamic hedging strategies. He assumes that the forward curve is flat which implies that the no-arbitrage condition is violated. In our setting, we study a variance contract where the underlying is a traded asset and derive explicit solutions for the hedge positions. Our assumptions are very general and, in particular, we only assume that no-arbitrage holds.

Our paper is also related to Aït-Sahalia, Karaman, and Mancini (2015) who estimate a two-factor affine stochastic volatility model to study the term structure of variance swaps in the equity index market. They show that the risk premiums contain a large jump risk component, especially at short horizons. Filipović, Gourier, and Mancini (2015) propose a quadratic term structure model for equity variance swaps. Using data on over-the-counter variance swaps, they also find a downward sloping term structure of variance swap payoffs.

The findings in this paper are also related to Duarte, Longstaff, and Yu (2007) who study risk and return for different fixed income arbitrage strategies featuring, among others, a volatility trading strategy through delta-hedged caps. Depending on the cap maturity, the (annualized) Sharpe ratios can be quite attractive, reaching 0.82 for a four-year maturity cap. The difference between their study and ours is that their results depend on a particular model. Hedge ratios to calculate the delta of caps are based on Black (1976). Our results are model-independent, moreover, shorting delta-hedged caps leaves the investor with Gamma exposure, similar to the straddle strategies that we consider as an alternative to the variance swap.

The rest of the paper is organized as follows. Section 1 provides the expressions to price variance swaps in Treasury futures markets and introduces the generalized Treasury variance swap. Section 2 describes our data set, explains the calculation of the variance risk premiums, and documents the term structures of implied volatilities and variance risk premiums in the fixed income market. Section 3 outlines the construction of the various trading strategies and presents the results of our empirical study, and Section 4 concludes. The Appendix contains some proofs and derivations; a detailed description of data filters and additional robustness checks are deferred to an Online Appendix.

1 Theory

This section theoretically derives the payoff of a variance swap in the Treasury bond market. While the contract we propose is robust to jumps in the underlying, we start with a standard contract that assumes a continuous process for the underlying before relaxing this assumption to present our main result. The model-free implementation requires that we use forward contracts instead of the futures that we have available in the data. Thus, in the empirical implementation, we also show how to convert American options on futures (as observed empirically) to European options on forwards (that are used in the theoretical derivation). Quantitatively, we find the differences between the two types of options to be negligible, especially for options with short maturities.

Variance swaps consist of two different legs: The floating leg (realized variance) and the fixed leg (expected variance). The difference between the two is then the variance risk premium (see, e.g., Carr and Wu (2009)). We fix the current time $t = 0$ and study contracts which pay at some future date T. We denote by $F_{t,T}$ the price of a forward contracted at time t with maturity T on the underlying X_T.

1.1 Log Treasury Variance Swap

To start, we assume that $F_{t,T}$ follows a continuous process. By no-arbitrage, this implies that the dynamics of $F_{t,T}$ are

$$\frac{dF_{t,T}}{F_{t,T}} = \sigma_t dW^{\mathbb{Q}_T},$$

where $W^{\mathbb{Q}_T}$ is a standard Brownian motion under the T-forward measure $\mathbb{Q}_T$ and σ_t is the instantaneous volatility.

A standard variance swap exchanges the realized variance defined as

$$\mathrm{RV}_{t,T}^{\log} := \left(\log \frac{F_{t_1,T}}{F_{t_0,T}}\right)^2 + \left(\log \frac{F_{t_2,T}}{F_{t_1,T}}\right)^2 + \ldots + \left(\log \frac{F_{t_n,T}}{F_{t_{n-1},T}}\right)^2$$

with some fair strike $\tilde{F}_{t,T}$ at maturity T. We assume that the realized variance is constructed using some sampling partition $\mathcal{T} = [t_0, t_1, \ldots, t_n]$ with trading dates $t = t_0 < t_1 < \ldots < t_n = T$. For a variance swap, the fair strike $\tilde{F}_{t,T}$ is chosen such that at initiation of the contract at time t no money is exchanged. Dupire (1994) and Neuberger (1994) define the variance swap rate $\tilde{F}_{t,T}$ in terms of the expected payoff under the risk-neutral measure $\mathbb{Q}$ from a so-called Log contract. To account for stochastic interest rates, we make use of the $\mathbb{Q}_T$ measure

$$\tilde{F}_{t,T} := -2\mathbb{E}_t^{\mathbb{Q}_T}\left[\log \frac{F_{T,T}}{F_{t,T}}\right]. \tag{1}$$

We now want to derive an expression for Equation (1). Let us start with the fundamental theorem of asset pricing, which implies that for any traded asset X_t

$$\frac{X_t}{p(t,T)} = \mathbb{E}_t^{\mathbb{Q}_T}\left[X_T\right], \tag{2}$$

where $p(t,T)$ is the price of a zero-coupon bond. Relation (2) holds in general and in particular under stochastic interest rates. Consider now the payoff $\log F_{T,T}$. Using the results in Carr and Madan (1998), we can write

$$\log F_{T,T} = \log F_{t,T} + \frac{F_{T,T} - F_{t,T}}{F_{t,T}} - \int_0^{F_{t,T}} \frac{(K - F_{T,T})^+}{K^2} dK - \int_{F_{t,T}}^{\infty} \frac{(F_{T,T} - K)^+}{K^2} dK.$$

Re-arranging yields that

$$-2(\log F_{T,T} - \log F_{t,T}) = 2\left(-\frac{F_{T,T} - F_{t,T}}{F_{t,T}} + \int_0^{F_{t,T}} \frac{(K - F_{T,T})^+}{K^2} dK + \int_{F_{t,T}}^{\infty} \frac{(F_{T,T} - K)^+}{K^2} dK\right).$$

The forward $F_{t,T}$ is a $\mathbb{Q}_T$-martingale.[3] By taking $\mathbb{Q}_T$ expectations on both sides, we get

$$\begin{aligned} -2\mathbb{E}_t^{\mathbb{Q}_T}\left[\log F_{T,T} - \log F_{t,T}\right] &= \mathbb{E}_t^{\mathbb{Q}_T}\left[\int_0^{F_{t,T}} \frac{(K - F_{T,T})^+}{K^2} dK + \int_{F_{t,T}}^{\infty} \frac{(F_{T,T} - K)^+}{K^2} dK\right] \\ &= \frac{2}{p_{t,T}}\left(\int_0^{F_{t,T}} \frac{P_{t,T}(K)}{K^2} dK + \int_{F_{t,T}}^{\infty} \frac{C_{t,T}(K)}{K^2} dK\right), \end{aligned} \tag{3}$$

where $P_{t,T}(K)$ and $C_{t,T}(K)$ are puts and calls, since by Equation (2)

$$\begin{aligned} \frac{P_{t,T}(K)}{p(t,T)} &= \mathbb{E}_t^{\mathbb{Q}_T}\left[(K - F_{T,T})^+\right] \\ \frac{C_{t,T}(K)}{p(t,T)} &= \mathbb{E}_t^{\mathbb{Q}_T}\left[(F_{T,T} - K)^+\right]. \end{aligned}$$

[3] It is easy to see that

$$\mathbb{E}_t^{\mathbb{Q}_T}\left[F_{T,T}\right] = \mathbb{E}_t^{\mathbb{Q}_T}\left[S_T\right] = \mathbb{E}_t^{\mathbb{Q}}\left[\frac{\exp\left(-\int_t^T r_s ds\right)}{p_{t,T}} S_T\right] = F_{t,T}.$$

Hence, it follows that the Log contract can be written as a portfolio of puts and calls with the same strike K and the same maturity T. To see that Equation (3) indeed represents expected variance, note that by applying Itô's Lemma, we get that

$$\log F_{T,T} - \log F_{t,T} = -\frac{1}{2}\int_t^T \sigma_u^2 du + \int_t^T \sigma_u dW_u^{\mathbb{Q}_T}. \tag{4}$$

Note that our objective is to price a contract on forward volatility which is the same as the volatility on futures. Forwards are martingales under the $\mathbb{Q}_T$-measure while futures are martingales under the risk-neutral measure. By Girsanov's theorem, the volatilities are the same while their drifts are not. In addition, Equation (4) also implies that $\tilde{F}_{t,T}$ is a fair strike for exchanging the quadratic variation of $F_{t,T}$ rather than $\mathrm{RV}_{t,T}^{\log}$, although the latter is considered a good approximation to the former under a fine enough time partition $\mathcal{T} = [t_0, t_1, \ldots, t_n]$. Indeed, $\mathrm{RV}_{t,T}^{\log}$ converges to $\int_t^T \sigma_u^2 du$ as the partition goes to zero, and then the quadratic variation can be perfectly replicated by a static position in options and a dynamic position in the underlying. We summarize our findings in a first Proposition.

Proposition 1. *Assume that $F_{t,T}$ is continuous. Then, the payoff $\int_t^T \sigma_u^2 du$ can be perfectly replicated by a static position in*

$$\tilde{F}_{t,T} = \frac{2}{p_{t,T}}\left(\int_0^{F_{t,T}} \frac{P_{t,T}(K)}{K^2}dK + \int_{F_{t,T}}^{\infty} \frac{C_{t,T}(K)}{K^2}dK\right),$$

and a dynamic position in the underlying which at any time $s \in [t, T]$ holds $2\left(\frac{1}{F_{s,T}} - \frac{1}{F_{t,T}}\right)$. Hence,

$$\tilde{F}_{t,T} = \mathbb{E}_t^{\mathbb{Q}_T}\left[\int_t^T \sigma_u^2 du\right].$$

Proof: See Appendix A.

The square root of Equation (3) resembles the definition of the VIX but instead of deriving everything under the risk-neutral measure, we derive it under the T-forward measure. This is important as it allows for stochastic interest rates, which is needed since we want to study contracts in the fixed income market. The VIX is in general referred to as being model-free, because we have not made any assumption on $F_{t,T}$ other than it being an Itô process. However, once we deviate from this assumption, i.e., once we allow for jumps, VIX^2 is no longer the fair strike of a variance swap. For example, Carr and Wu (2009) note that

$$\mathbb{E}_t^{\mathbb{Q}_T}\left[\int_t^T \sigma_u^2 du\right] = \text{VIX}_t^2 + \text{error from jumps}.$$

Negative (positive) jumps induce an upward (downward) bias in VIX^2. The sensitivity of the standard variance swap presented in Proposition 1 has been extensively documented in the literature, see for example Broadie and Jain (2008).

1.2 Generalized Treasury Variance Swap

We now proceed to relax the assumption that the underlying follows a continuous process. Martin (2013) studies so called simple variance swaps which are robust to jumps. He does this by altering the fair strike, $\tilde{F}_{t,T}$. In the following, we are deviating from this approach by changing the realized variance leg of the contract instead. The idea is that rather than focussing on the unobservable quantity $\tilde{F}_{t,T}$, we concentrate on the observed realized variance. This closely follows Neuberger (2012) and Bondarenko (2014) who study generalized variance swaps in the equity market. We extend their approach by allowing for stochastic interest rates.

We define the generalized Treasury variance swap (GTVS) as an agreement to exchange

$$\widetilde{RV}_{t,T} = 2\left[\left(\frac{F_{t_1,T}}{F_{t_0,T}} - 1 - \log\frac{F_{t_1,T}}{F_{t_0,T}}\right) + \ldots + \left(\frac{F_{t_n,T}}{F_{t_{n-1},T}} - 1 - \log\frac{F_{t_n,T}}{F_{t_{n-1},T}}\right)\right], \tag{5}$$

with the fair strike $\tilde{F}_{t,T}$ where $t = t_0 < t_1 < \ldots < t_n = T$. At first sight, $\widetilde{RV}_{t,T}$ may not look like a variance measure but it turns out to be the same as realized variance computed from log returns ($\text{RV}_{t,T}^{\log}$) plus cubed returns.[4] This new measure of realized variance has the useful property that it allows for a perfect replication of the variance contract for every price path and every partition. We summarize our findings in the following Proposition.

[4] A formal proof can be found in Appendix A.

Proposition 2. *For any process $F_{t,T}$, the payoff $\widetilde{RV}_{t,T}$ can be perfectly replicated by a static position in*

$$\tilde{F}_{t,T} = \frac{2}{p_{t,T}} \left(\int_0^{F_{t,T}} \frac{P_{t,T}(K)}{K^2} dK + \int_{F_{t,T}}^{\infty} \frac{C_{t,T}(K)}{K^2} dK \right), \tag{6}$$

and a dynamic position in the underlying, which at any time $s \in \mathcal{T} = [t_0, t_1, \ldots, t_n]$ holds $2\left(\frac{1}{F_{s,T}} - \frac{1}{F_{t,T}}\right)$. Hence,

$$\tilde{F}_{t,T} = \mathbb{E}_t^{\mathbb{Q}_T}\left[\widetilde{RV}_{t,T}\right].$$

Proof: See Appendix A.

Hence, it follows that the replicating strategy for the realized variance consists of two parts: First, a path-independent payoff from options and, second, a dynamic strategy in the underlying. Note that we have not made any assumption about $F_{t,T}$ or about the frequency with which we re-balance the portfolio. Proposition 2 looks almost identical to Proposition 1, and indeed, if $F_{t,T}$ is continuous, both $\widetilde{RV}_{t,T}$ and $\mathrm{RV}_{t,T}^{\log}$ converge to $\int_t^T \sigma_u^2 du$ as the partition goes to zero.

Note, however, that the results of Proposition 2 are valid under any partition $\mathcal{T}$ as long as re-balancing in the underlying takes place on the same dates, $t_0, t_1, \ldots, t_n$. It, hence, does not matter whether we sample $\widetilde{RV}_{t,T}$ from daily, weekly, or monthly data as long as we re-balance at the daily, weekly, or monthly frequency. The reason for this lies in the aggregation property of the realized variance estimator given in Equation (5), see Neuberger (2012). The aggregation property essentially tells us that for any real-valued function g, for any martingale process X_t, for any measure $\mathbb{M}$, and for any times $0 \leq s \leq t \leq u \leq T$,

$$E_t^{\mathbb{M}}\left[g(X_u - X_s) - g(X_u - X_t) - g(X_t - X_s)\right] = 0.$$

In other words, for any function g which satisfies this restriction and if X_t is the forward price $F_{t,T}$—which we know is a martingale under the $\mathbb{Q}_T$-measure—then the discretely-sampled payoff $\sum_{i=1}^{n} g(F_{t_i,T} - F_{t_{i-1},T})$ has the same market price as the path-independent time T payoff of $g(F_{T,T} - F_{t,T})$. Thus, the realized variance which we calculate from higher frequency future returns (daily) is an unbiased estimate of the lower frequency counterpart (monthly). A priori, we do not expect the effect of discrete sampling to be large for pricing (see Broadie and Jain (2008)), however, the aggregation property tells us that it is exactly zero for the measure of realized variance that we propose.

Next, we implement the generalized Treasury variance swap introduced in Proposition 2 using a large panel of Treasury options data. We first outline how to convert American option prices on futures into European option prices on forwards that are needed to calculate the model-free implied variance measure. We will also compare the payoffs of the GTVS (Proposition 2) with those of the LTVS (Proposition 1).

2 Data and Measurement of Variance Swaps

In this section, we briefly introduce the data used in our analysis. We use futures and options data to construct the bond variance swap payoffs. To put our results for the Treasury market in perspective, we also calculate returns to variance trading strategies using S&P500 futures and options. Before we can start, however, we need to account for two features of the data that mainly affect the calculation of the fair strike. First, the observed options are written on futures rather than on forwards as implied by our theoretical contracts in Propositions 1 and 2. Second, the futures options are American rather than European. Based on existing evidence for the Eurodollar futures market, we expect both effects to be small. For example, Flesaker (1993) and Cakici and Zhu (2001) show in the Eurodollar futures market that the effect of having futures as the underlying as opposed to forward contracts is very small especially for options with shorter maturities.

The data is available from October 1982, May 1985, and June 1990 to May 2012 for the 30 year, 10 year, and 5 year Treasury bond futures and options, respectively. Using a monthly frequency throughout the paper, we have at most 355, 325, and 264 observations available, respectively. To make our trading strategies comparable, we present our baseline results using a sample starting in June 1990.

2.1 Futures and Options Data

Treasury Futures and Options: To calculate implied and realized variance measures for Treasury bonds, we use futures and options data from the Chicago Mercantile Exchange (CME). For our benchmark results we use end-of-day price data for the 30-year Treasury bond futures, the 10- and 5-year Treasury notes futures, and end-of-day prices of options written on the underlying futures, respectively.[5]

Treasury futures are traded electronically as well as by open outcry. While the quality of electronic trading data is higher, the data only becomes available in August 2000. To maximize our time span, we use data from electronic as well as pit trading sessions.

The contract months for the Treasury futures are the first three (30y Treasury futures) or five (10y and 5y Treasury futures) consecutive contracts in the March, June, September, and December quarterly cycle. This means that at any given point in time, up to five contracts on the same underlying are traded. To get one time series, we roll the futures on the 28^{th} of the month preceding the contract month.

For options, the contract months are the first three consecutive months (two serial expirations and one quarterly expiration) plus the next two (30y futures) or four (10y and 5y futures) months in the March, June, September, and December quarterly cycle. Serials exercise into the first nearby quarterly futures contract, quarterlies exercise into futures contracts of the same delivery period. We roll our options data consistent with the procedure applied to the futures.[6]

S&P500 Index Futures and Options: In line with our approach for Treasuries, we calculate the implied and realized variance measures for the stock market using futures and options on the S&P500 index from CME. The sample period is from January 1983 to May 2012.[7]

[5] We use settlement prices for both options and futures which do not suffer from stale trading or bid-ask spreads. CME calculates settlement prices simultaneously for all options based on their last bid and ask.

[6] Detailed information about the contract specifications of Treasury futures and options can be found on the CME website, www.cmegroup.com.

[7] We compare our results to the VIX and VXO measures that are calculated using options on the S&P500 cash index instead of S&P500 index futures. The VIX is the implied volatility calculated using a

2.2 Differences between Futures and Forwards and the Effect on Option Prices

Options on Treasury and S&P500 index futures are American. We need to adjust for these two features since Equation (6) is derived under the assumption that the options are European and written on forwards. In order to obtain implied volatilities from the available option prices, we hence need a suitable model for the futures and American feature.

The approach we use closely follows Cakici and Zhu (2001) and is outlined in the Online Appendix. In line with the results in Cakici and Zhu (2001), we find the adjustment from American futures options to European forward options to be very small and never exceeding two percent of the implied volatility level for all tenors and option maturities we consider. For the 30y Treasury options the difference ranges between 0.60% and 0.91% of the price for short-term and long-term options, respectively. For the 10y Treasury options, the average differences range between 0.65% and 0.90%, and finally for 5y Treasury options they range between 0.56% and 0.77%, respectively.[8]

2.3 Construction of Implied Volatility

In this section we show how to empirically calculate the fair strike of the variance swap at time t with maturity T. In line with the equity literature, we report our benchmark results for a horizon $\tau = T - t$ of one-month. However, we also consider a term structure using horizons ranging from ten days to one year.

Before including an option in the calculation, we apply a set of filters to clean the data: (i) We eliminate all data where either the futures or option price, the strike, the maturity, or the open interest are equal to zero. (ii) We also delete data when option prices fail to pass the no arbitrage boundary conditions.[9] (iii) We eliminate deep in-the-money

model-free approach, whereas the VXO is calculated using the Black and Scholes (1973) implied volatility. The VIX is available starting in January 1990 and the VXO is available since January 1986. Over the common sample period, the VIX and our implied volatility measure from index futures options using the same methodology have a correlation of over 99.4% and the root mean squared error is below 1%.

[8]To save space, we defer more detailed summary statistics to the Online Appendix.

[9]The boundary condition for the call options is $C_{t,T}(K) - \max{(F_{t,T} - K)^+} \geq 0$ and for the put options it is $P_{t,T}(K) - \max{(K - F_{t,T})^+} \geq 0$.

options, i.e., we eliminate calls if the strike is less than $0.94 \times F_{t,T}$ and puts if the strike is greater than $1.06 \times F_{t,T}$, where $F_{t,T}$ is the underlying futures price. Thus, note that we do not restrict ourselves to out-of-the-money options only. Using strike prices greater than $0.94 \times F_{t,T}$ for calls or less than $1.06 \times F_{t,T}$ for puts means we still include some in-the-money and near-the-money options. At each maturity, we then fit a spline for the available implied volatilities against their corresponding strike prices and from the fitted spline, we obtain a fine grid of implied volatilities. We then convert the grid of implied volatilities into European option prices and numerically evaluate Equation (6) with daily options data on 30y, 10y, and 5y Treasury futures, and the S&P500 futures adjusted for the forward/futures feature.

2.4 Construction of Realized Volatility

We now describe the construction of the realized variance measures used to calculate the payoffs at maturity T to a generic variance swap that pays the difference between realized variance and the strike $\tilde{F}_{t,T}$.

Proposition 2 implies that the fair strike of the GTVS is given by $\tilde{F}_{t,T}$ irrespective of the sampling partition and the underlying price process if we use the definition given in Equation (5) for the realized variance leg, $\widetilde{RV}_{t,T}$. As variance swaps generally use a daily sampling frequency, the benchmark results we report are based on daily data. In addition, we also calculate the standard measure of realized variance, $RV^{\log}_{t,T}$, using daily log returns.

Note that the realized variance measures $\widetilde{RV}_{t,T}$ and $RV^{\log}_{t,T}$ are only observed ex post at maturity T of the variance swap. For our benchmark horizon of one month and using a daily sampling frequency, $\widetilde{RV}_{t,T}$ and $RV^{\log}_{t,T}$ are therefore based on the futures prices of the previous 21 trading days. We use end-of-day prices measured at 14:00 CT in line with the end of pit trading hours at the CME. In addition to the one-month horizon, we calculate a term structure of realized variances ranging from ten days to one year (252 trading days) to match the term structure of implied variances.

2.5 Summary Statistics and Variance Risk Premiums

Table I reports summary statistics of implied volatilities as described in Equation (6) and Figure 1 provides a plot of the average implied volatilities. As it is custom in practice to report volatilities, we present the square root of the annualized variances expressed in percent.

Insert Table I and Figure 1 here.

We note that all term structures are downward sloping in the maturity dimension. For 30y (10y, 5y) Treasury options, average implied volatilities range from 11.8% (8.3%, 5.5%) for the ten-day options to 9.9% (6.7%, 4.5%) for options with one-year to maturity. Hence, along the tenor dimension, the term structure is upward sloping. The slope of the term structure of implied variances, which we define as the difference between the implied variance of a one-year option and a one-month option, is negative on average. The equity index option literature finds both upward and downward-sloping implied volatility term structures depending on the method used (see, e.g., Aït-Sahalia, Karaman, and Mancini (2015)). Using options on S&P500 futures, we find a slightly downward sloping term structure ranging between 20.1% and 19.4%.

In analogy to the calculation of the equity risk premium, we call the difference between the (ex post) realized and the implied variance in the variance swap the (unconditional) variance risk premium. In line with this definition, we use the terms ex post or realized variance risk premiums to describe the time-varying differences between the realized and the implied variances.[10] The realized variance measure is calculated by summing up daily data.[11] We report average variance risk premiums in Table II together with the corresponding Sharpe ratios for different subsamples. Panel A reports variance risk premiums

[10]Note that we explicitly take an *ex post* view by focusing on realized variance which is in line with the payoff traders face in a variance swap. Taking an *ex ante* view to study *conditional* variance risk premiums requires to form expectations about variance using, for example, an econometric model. See, e.g., Mueller, Sabtchevsky, Vedolin, and Whelan (2016) for an exploration of conditional variance risk premiums in the stock and bond markets.

[11]For different maturities we sum up daily data over various horizons. The sample averages are then taken over the whole sample. Regardless of the maturity, this is essentially the same as taking an unconditional average of all daily values.

for the period June 1990 to May 2012 which is the time span for which we have available data for all tenors. There are several noteworthy observations. First, variance risk premiums are negative in line with expectations (investors are willing to pay a premium to be protected against volatility spikes) and existing research for the equity market (with the sole exception of 30y Treasury variance risk premiums for the one year horizon). Second, the downward sloping implied volatility term structures imply an upward sloping term structure of the ex post variance risk premiums. Given that the variance risk premiums are essentially all negative on average, this of course means that in absolute terms, the term structure of variance risk premiums is downward sloping. The same is true for the Sharpe ratios. These declining Sharpe ratios are consistent with findings of van Binsbergen and Koijen (2016), who document that Sharpe ratios in a range of markets (stocks, bonds, corporate bonds, index straddles, and housing) decline with maturity.[12] Third, the variance risk premium term structures look very similar across tenors. In absolute terms, the variance risk premium term structures are all downward sloping while the Sharpe ratios decline almost monotonically with the horizon. At the same time, average variance risk premiums increase with the tenor (again in absolute terms) and, for example, average 30y variance risk premiums are always larger than the corresponding variance risk premiums for the 10y and 5y tenors.

Insert Table II and Figure 2 here.

To check robustness of this pattern about the average shape of the term structure of variance risk premiums, we examine term structures for different subsamples. Figure 2 plots the average variance risk premiums and Sharpe ratios for a sample that excludes the financial crisis and ends in December 2007 (left panels) and a sample that starts in 2008 (right panels). We note that the variance risk premium surface looks very similar for both subsamples. However, the slope across tenors becomes much steeper during the post crisis period. For example, the difference between the ten-day variance risk premium

[12] Applying a hedging based method for swaptions with different maturities, Duyvesteyn and de Zwart (2015) also find a downward sloping term structure of variance risk premiums in the swaptions market. Similarly, Aït-Sahalia, Karaman, and Mancini (2015) and Dew-Becker, Giglio, Le, and Rodriguez (2016) also find that the term structure of variance risk premiums is downward sloping for the S&P500.

on the 30y and 5y Treasury futures is -3.17 during the pre-crisis period, but increases to -6.45 after 2007. At the other end of the maturity spectrum, for one year options, we find that the difference between the variance risk premiums for 30y and 5y Treasury futures turns from being negative during the pre-crisis period to a positive value after the crisis, implying that while the 30y variance risk premium was more negative than the 5y variance risk premium in the early sample, this relationship has been reversed since 2008. A similar pattern emerges for Sharpe ratios which increase in absolute terms during the post-crisis periods, especially for shorter horizons while the Sharpe ratios for longer horizons are now slightly smaller.

Finally, we examine to what extent the implied variances across tenors and option horizons are driven by a common set of factors. To this end, we perform a principal components analysis (PCA) on monthly changes in the implied variances for the three tenors and horizons ranging from ten days up to one year. We find that the first principal component (PC) explains around 70% of the overall variation while the second PC accounts for an additional 15%. Since the third PC captures 5% of the overall variation, the first three PCs are sufficient to explain almost 90% of the variation across tenors and for all horizons.[13] Figure 3 plots the factor loadings on the implied variance surface for the first two PCs. The first PC acts as a level factor, having a roughly uniform impact on implied variances for all tenors and option maturities. On the other hand, the second PC acts as a slope factor, having a more prominent impact on long-horizon variances and monotonically decreasing loadings across tenors. To summarize, only a small number of factors is needed to capture the variation in implied variances across tenors and option maturities, and the factors present the usual level and slope effects as in other asset classes (see, for example, Litterman and Scheinkman (1991) and Lustig, Roussanov, and Verdelhan (2011)).

Insert Figure 3 here.

[13]This is in line with results for the swaptions implied volatility surface reported in Trolle and Schwartz (2014).

3 Empirics

In this section, we evaluate different variance trading strategies using fixed income options. We first present summary statistics for the variance swap as well as straddles, which also provide exposure to variance risk. Our findings reveal that trading variance in fixed income markets is attractive even if we condition on other risk factors. We further benchmark our results for the fixed income markets against those realized by trading variance in the equity index market using options on S&P500 futures. As it is well established that variance trading in the equity market has been popular and profitable over the last decade, comparing the returns from the Treasury and equity market helps to put our novel results into perspective.

3.1 Trading Strategies

First, we study returns to a strategy based on our generalized Treasury variance swap which are calculated as follows

$$r_{t,T}^{GTVS} = \frac{\widetilde{RV}_{t,T}}{\mathbb{E}_t^{\mathbb{Q}_T}\left[\widetilde{RV}_{t,T}\right]} - 1, \tag{7}$$

where T is the maturity of the contract, $\widetilde{RV}_{t,T}$ is the realized variance as defined in Equation (5) and the denominator is given by Equation (6). This means that the return is the ex post variance risk premium scaled by the fair strike price. At the same time, this is the excess returns to a fully collateralized long position in the variance swap that posts $\mathbb{E}_t^{\mathbb{Q}_T}\left[\widetilde{RV}_{t,T}\right]$ dollars of collateral and receives $\widetilde{RV}_{t,T}$ at expiration plus interest on the collateral. Alternatively, one can label what we define as the realized variance risk premium the realized excess return on a unit position in the variance swap, whereas Equation (7) represents the scaled version thereof. However, we prefer the term "ex post variance risk premium" to highlight the fact that our main object of interest is the realized quantity of the ex ante variance risk premium that the literature generally focuses on.

We then compare these returns to those of two standard volatility trading strategies using straddles. First, we consider an at-the-money straddle, a classical position for getting exposure to volatility. Unfortunately, such a position loses sensitivity to volatility as the underlying moves away from the strike. To this end, in addition to an unhedged straddle, we also consider a delta-hedged position. Coval and Shumway (2001) and Santa-Clara and Saretto (2009) show that trading in straddles yields very attractive (annualized) Sharpe ratios above one for options on the S&P500 index futures.

To calculate the delta-hedged returns, we proceed as follows. Each month, we simultaneously purchase an at-the-money call and put option with 30 days to expiration (or the closest to 30 days). We track the path of the straddle until its expiration date and on a daily basis we go long or short the corresponding underlying future such that the whole position becomes delta-neutral at the end of each day. Denoting the delta of a straddle on a given date t as $\Delta_{S,t}$, the accumulated profit and loss from this hedging activity can be written as:

$$\sum_{i=t+1}^{T} -\Delta_{S,i-1}\left(F_{i,T} - F_{i-1,T}\right).$$

Hence, hold to expiration returns of the delta-hedged straddle strategy are defined as:

$$r_{t,T}^{\Delta S} = \frac{\left(K - F_{T,T}\right)^{+} + \left(F_{T,T} - K\right)^{+} - \sum_{i=t+1}^{T} \Delta_{S,i-1}\left(F_{i,T} - F_{i-1,T}\right)}{P_{t,T}\left(K\right) + C_{t,T}\left(K\right)} - 1. \tag{8}$$

3.1.1 Summary Statistics

Table III reports summary statistics of the trading strategies together with different performance measures, the Sharpe ratio and Jensen's alpha. Note that while Equation (7) is defined for an arbitrary maturity $T - t$ we focus on the one-month horizon in line with the sampling frequency for our benchmark results. Hence, in what follows we study one-month excess returns on one-month variance swaps only. To calculate the alpha, we employ two different market indices depending on whether we use Treasury or equity options, respectively. The market return for the options on the S&P500 futures is the value weighted excess return on all stocks in CRSP and for the bond options, we use the total return on the Barclays US Treasury bond index available from Datastream.

Insert Table III here.

Panel A summarizes the annualized returns on the generalized Treasury variance swap, our main object of interest. We note that shorting a variance swap produces a monthly average return of around 20% for options on 30y, 10y, and 5y Treasury futures. All average returns are highly significantly different from zero as indicated by the t-statistics which range between 6.61 for the 5y and 12.71 for the 10y futures. The volatilities of the variance swap trading strategies are relatively small, leading to annualized Sharpe ratios ranging between 1.4 for the 5y and 2.7 for the 10y Treasury futures, respectively.[14] The associated alphas are only marginally smaller than the average returns ranging between between 18% and 27% and, hence, the market return does not explain the variance swap returns at all. Trolle and Schwartz (2014) study variance swaps in the swaptions market and report even higher average returns ranging between 44% and up to 66% per month. The associated volatilities are considerably higher than in the Treasury options market, so the Sharpe ratios they find are in line with ours. The returns to shorting a variance swap are negatively skewed and exhibit excess kurtosis, however, the values are comparable to the strategies based on straddles. Trading a variance swap on the S&P500 index futures is similarly attractive: The return is 31% per month, with an annualized Sharpe ratio of two. However, the reward comes with some additional risk as the strategy has much fatter tails with a kurtosis as high as 22.

Panel B reports the summary statistics for at-the-money (ATM) straddles without taking a position in the underlying futures. The average returns are much smaller than for the variance swap strategy. In fact, shorting an ATM straddle with options written on Treasury futures produces an average return close to, and not significantly different from, zero. Similarly, we also find the strategies' alpha to be insignificant. Moreover, the associated risk as proxied by the volatility is more than 50% higher compared to the variance swap strategies.

Panel C in Table III presents the results for the delta-hedged straddle strategies. By delta-hedging the straddle, the volatility of the strategy becomes considerably smaller. Average returns are highly significant and range between −3.4% (30y) and −7% (10y) per month. The corresponding annualized Sharpe ratios are 0.5 for the 30y and 1.1 for the 10y Treasury futures, respectively. The associated alphas are significant and also range between −3.4% (30y) and −7% (10y). Overall, we conclude that delta-hedged straddle

[14] Note that the tables report monthly Sharpe ratios.

strategies yield attractive returns but variance swaps still produce average returns and alphas that are on average larger by a factor of roughly four, both for Treasury as well as for S&P500 futures strategies.

Insert Figures 4 and 5 here.

To further investigate the variance swap strategy, we plot the time series of the GTVS returns for the 30y Treasury futures in Figure 4 (upper panel) together with the associated realized and implied volatility measures (lower panel). We note that most of the time, implied volatility exceeds the ex post realized one. Hence, on average, a strategy that is long realized and short implied variance produces a negative return. At the same time, there are some very distinct positive spikes which, interestingly, correspond to general periods of distress as can be gauged from the annotations in the upper panel. However, there is one single spike which is very specific to the bond market and it coincides with the large bond-market sell-off in July 2003 due to mortgage hedging activity (see Malkhozov, Mueller, Vedolin, and Venter (2016b)). To compare, we also plot the returns to the variance swap on the S&P500 futures in Figure 5. Different from the bond market, there is for example no spike in July 2003. On the other hand, there is one large positive spike in July 2002, when the S&P500 index lost 8% between June and July.

3.1.2 Statistical Significance

One might worry about statistical significance as it is well known that option trading strategies are non-Gaussian and the extant literature shows that performance measures such as the Sharpe ratio can change dramatically if returns do not follow a Normal distribution (see Lo (2002)).

Thus, in the following, we use a studentized bootstrap to obtain confidence intervals on the mean, Sharpe ratio, and alpha of the trading strategies discussed earlier. Using a sample of 10,000 bootstrapped repetitions, we report 95% confidence intervals in Table IV..

Insert Table IV here.

In line with the previous results, average returns for the variance swap and delta-hedged straddle strategies are all significantly different from zero, as none of the confidence intervals includes the zero itself. The same applies to the Sharpe ratios as well as for the alphas.

3.1.3 Risk-Adjusted Returns

In this section, we explore how and whether the returns of these trading strategies are related to market risk. To this end, we regress the strategy returns onto the returns of different equity and bond portfolios. In particular, we control for the market, size (SMB), book-to-market (HML), and momentum (MOM) portfolios. We also include two liquidity factors for bond and equity markets: The liquidity factor extracted from bonds used in Malkhozov, Mueller, Vedolin, and Venter (2016a) and the Pástor and Stambaugh (2003) equity liquidity factor. In order to make the two measures comparable, we multiply the latter by minus one to get an illiquidity measure. Table V presents the regression results for each strategy. We report the alpha together with its associated t-statistic while for the other regressors we only report (Newey and West (1987) adjusted) t-statistics.

Insert Table V here.

Panel A reports the results for variance swap returns. We first note the high significance of the momentum factor (t-statistics all above two for the Treasury variance swap returns) and the borderline significance of the equity momentum factor for variance swaps on the S&P500. The alpha of the strategy is still negative and highly significant, which indicates that while these factors explain some of the variation (adjusted R^2 range between 4% to 10%), the majority is left unexplained.

The alphas for the un-hedged straddles (Panel B) are not statistically significant except for the S&P500. On the other hand, once the straddles are delta-hedged (Panel C) we find alphas to be highly significant. Naturally, due to the delta-hedging, the market return is insignificant. Moreover, we find the momentum portfolio to have no correlation with the strategy returns.

Overall, we reconfirm that trading variance in the fixed income market produces high average returns and attractive Sharpe ratios. The associated alphas are large and statistically significant even when controlling for standard risk factors.

3.1.4 Transaction Costs

It is well known that transaction costs lower the returns of option strategies and that the impact is increasing with decreasing moneyness, i.e., it is worst for deep out-of-the-money options (see, for example, George and Longstaff (1993) and Santa-Clara and Saretto (2009)). In the following, we explore the impact of bid-ask spreads onto the profitability of the trading strategies. To this end, we recalculate the summary statistics in Table III by assuming that we always buy at the ask price and sell at the bid price. The results are summarized in Table VI.

Insert Table VI here.

Not surprisingly, we find that average returns drop but shorting variance is still attractive: Annualized Sharpe ratios vary between 1.1 and 2.4 for the variance swaps. Moreover, average returns and alphas remain statistically significant both for the variance swaps and the delta-hedged straddles. We conclude that while bid-ask spreads lower average returns by up to one third, trading variance swaps is still very profitable.

3.1.5 Margins

The strategies we consider are implemented using options traded on the CME. In practice, this means that investors are subject to margin constraints which are likely going to influence a strategy's profitability. For example, if a strategy of shorting variance leads to large losses, an investor could be forced to close down the position if she does not have unlimited funds. In the following, we follow a similar procedure as in Santa-Clara and Saretto (2009) and study the impact of margin requirements on the trading strategies discussed earlier.

In practice, variance swap strikes are quoted in terms of volatility (expressed in percent), not variance (see, for example, Allen, Einchomb, and Granger (2006)). The payoff to a variance swap is the difference between the ex post realized variance and the squared strike price multiplied by the so-called variance notional, which represents the profit or loss per point difference between realized and implied variance. Since market participants often think in terms of volatility, the vega notional is often used instead of the variance notional. The vega notional shows the profit or loss from a 1% change in volatility and it is calculated as $N_{vega} = N_{var} \times 2K$, where K is the strike of the variance swap expressed in terms of volatility.[15]

Margin requirements in general depend on the type of strategy employed. Variance swaps are margined in a similar way to options. In the following, we assume that the required margin is nine times the vega notional. For our sample period, the assumed margin is sufficient to withstand a daily adverse move in volatility in excess of three standard deviations based on monthly data. This means a trading portfolio that can be re-balanced daily will never be wiped out in a single day.[16] We further assume that during the life of the swap, the variation margin as well as the initial margin are set in the same fashion. For example, the variation margin for a short position with notional value of one dollar at time t is set as

$$PV_t(T) \times \left[\left\{ \frac{t}{T} \times \widetilde{RV}_{0,t} + \frac{T-t}{T} \times (\text{Implied Vol}(t,T) + 9)^2 \right\} - K^2 \right], \tag{9}$$

where K is the fair strike at initiation and $PV_t(T)$ is the t-present value of one dollar at time T. The minimum required margin is then the maximum of the initial margin and the variation margin at each point in time t.[17]

[15] Thus, a vega notional of one dollar means a trader with a short position in a variance swap with a strike of 10% will lose one dollar if the volatility increases to 11%.

[16] Allen, Einchomb, and Granger (2006) provide an example of a term sheet with a collateral requirement equal to three times the vega notional, which is not sufficient to prevent a complete loss in a single trading day. In practice, payoffs to variance swaps are often capped. As a result, a three notional vega margin may be enough. However, this further complicates hedging and pricing (see also Martin (2013)). To abstract from the pricing issues, we thus impose a higher margin.

[17] As the month progresses, Equation (9) requires the calculation of the variance swap rate with a decreasing time to maturity. To ensure that our results are not biased by interpolation, we enter into a position when the options have exactly one month to maturity instead of sampling at the end of the month.

Margins influence our strategies along two dimensions: i) they limit the number of contracts that an investor can write (execution) and ii) they may force the investor to close down positions and take losses (profitability). To evaluate these effects, we assume in line with Santa-Clara and Saretto (2009) a zero-cost strategy. In particular, at the beginning of each month the investor borrows one dollar and allocates that amount to a risk-free rate account which she can use to cover the margin requirements. Then, the investor takes a short position in a variance swap contract for a notional amount which is equivalent to a fraction of that one dollar. This quantity is referred to as the target notional, and the corresponding vega amount of the contract is referred to as the target vega notional. The initial margin requirement is then approximately equal to nine times the target vega notional.

During the month, we assume that the investor cannot borrow additional capital. In other words, margin calls are met by liquidating the investment in the risk-free account. When the risk-free account is no longer sufficient to meet the margin call, then the position is liquidated at the swap value. The investor is then allowed to open a new position such that the new margin does not exceed 90% of the available wealth. The maturity date of the new contract remains the same as before, but a new strike is defined for the value of the swap to be zero. Moreover, the (vega) notional of the new contract is adjusted accordingly. Hence, unless no re-balancing occurs during the month, the effective (vega) notional will differ from the target (vega) notional. At the end of the month, the variance swap position is closed and the proceeds are added to the risk-free account. Together with the interest earned on the risk-free account (which in general is negligible) this allows to calculate both the P/L of the strategy as well as the return in month t compared to the initial position of one dollar.

Given the relatively high margin requirements, the maximum target vega notional possible is 10 cents on the dollar or 10%. Table VII reports results using target vega notionals ranging between 1% and 10%, meaning that the initial margin ranges from just under 10% to 90% of the one dollar borrowed at the beginning of each month. For comparison, we also calculate un-margined returns based on the same zero-cost strategy described above but without applying the variational margin. This means that the initial position is never forcibly closed out and remains the same until the end of the month. Note that the returns in Table VII are calculated for a short position in a variance swap and they are, thus, positive on average.[18] We report bootstrapped confidence intervals for the quantities of interest.

Insert Table VII and Figure 6 here.

First, the results in Table VII show that margined returns are increasing in the target vega notional. This is not surprising as shorting variance is profitable and a higher target vega notional is tantamount to increasing the exposure to variance for a given amount of capital. At the same time, a higher vega notional also leads to a higher probability of forced rescaling if the volatility increases too much and a margin call cannot be met with the available capital. For example a 10% vega notional leads to a forced rescaling in every month of our sample and for all underlying tenors, while a 1% target vega notional never leads to a forced rescaling.

The relationship between the target vega notional and the volatility of our strategy is less straightforward. Absent rescaling, the volatility of the strategy is increasing in the vega notional. However, once the margin requirements take effect, forced rescaling often leads to a reduction in the volatility of the trading strategy as positions are often scaled down during periods of turmoil. This has an overall positive effect on Sharpe ratios that is at times, however, offset by a reduction in the mean return. Hence, Sharpe ratios are not necessarily monotonic in the target vega notional. However, they remain on average very sizable and range between 1 and 1.8 (annualized) depending on the underlying and the target vega notional.

As mentioned above, the assumed margin requirements are rather restrictive. This ensures that the investor cannot loose the total available capital in a daily move and before rescaling is allowed, which makes our strategy viable in the long-run. Applying for example the same setup to variance swaps on S&P500 futures results in the investor losing all the capital before being able to rescale a total of six times during the sample

[18] Moreover, note that the returns in Table VII are defined slightly differently than those presented in Table III and, hence, they cannot be directly compared.

period. Hence, the strategy with variance swaps on equity is much riskier compared to using Treasury variance swaps and thus, in practice, margins would have to be higher in order to ensure that the portfolio can withstand even periods of severe turmoil in financial markets.[19]

Figure 6 plots in log scale the accumulation of wealth for different target vega notional short positions in variance swaps on 30y Treasury futures. The starting wealth is one dollar and it is assumed that the full amount is reinvested each month. As benchmark, we plot the total return index from investing one dollar in the S&P500 stock index. The figure confirms the results from Table VII, namely that that shorting variance in the Treasury market remains a profitable strategy when taking margins into account. Measured over 22 years, the Treasury variance swap strategies with vega notionals ranging between 5% and 10% deliver on average around 40% return per year, compared with just above 6% for the S&P500 index. Moreover, the results also highlight that with proper margin requirements, the strategy remains very attractive even in a setting where investors do not have unlimited access to capital and may need to close out or scale down their positions during market turmoil.

3.2 The Impact of Realized Variance

In this section we gauge to what extent it makes a difference whether we use realized variance as defined in Equation (5) or whether we follow common practice and use daily squared log returns. The results are reported in Table VIII. The summary statistics reveal that *on average* it does not matter whether one uses one or the other measure of realized variance. This is intuitive, as for the data and sample period we study, the number and size of positive and negative jumps is roughly the same. The distinction between the two approaches becomes more evident, however, once we consider the time series of the differences between a variance swap that is defined using daily squared log returns and of our GTVS that uses realized variance as defined in Equation (5). The time series of the differences is depicted in Figure 7.

[19] As mentioned earlier, one can alternatively impose a cap on the payoff as is regularly done in practice.

We note that there are both positive and negative differences in line with the notion that there are positive and negative jumps. The largest positive spikes correspond to the same distinct spikes that we observe in Figure 4. The three largest spikes are: July 2003 which corresponds to the month with the largest mortgage refinancing activity, August 1990 which was when Iraq invaded Kuwait, and September 2008 right after the Lehman default. Positive spikes correspond to negative jumps. In Appendix A we derive the difference between the realized variance measure that we use for our variance swap contract and realized variance calculated using daily squared log returns. It is straightforward to show that $\widetilde{RV}_{t,T} = RV_{t,T}^{\log} +$ cubed returns. Hence, large negative jumps render the log realized variance measure larger than $\widetilde{RV}_{t,T}$. Since the payoff is the realized leg minus the fair strike, the payoff becomes larger (smaller) compared to using $\widetilde{RV}_{t,T}$ in the presence of negative (positive) jumps.

Insert Table VIII and Figure 7 here.

3.3 Treasury Implied Volatility and Economic Activity

The VIX is often referred to as a fear gauge and, hence, it seems natural to ask whether the VIX has any predictive power for future economic activity.[20] In the following, we study whether Treasury implied variance and the slope of the implied variance term structure has any predictive power for economic activity as captured by the Chicago Fed National Activity Index (CFNAI).[21] We also check how implied volatility and the slope of implied variances is related to a measure of economic stress, namely the St. Louis Fed Stress Index (STLFSI).[22]

[20]For example Bekaert and Hoerova (2014) present empirical evidence that the VIX^2 is an excellent predictor of economic activity at the monthly, quarterly, and annual frequency.

[21]The CFNAI is available on the Chicago Fed web page as a weighted average of 85 existing monthly indicators of national economic activity including (i) production and income, (ii) employment, unemployment, and hours, (iii) personal consumption and housing, and (iv) sales, orders, and inventories. A positive index corresponds to growth above trend and a negative index corresponds to growth below trend.

[22]The STLFSI is available from FRED and is a principal component from 18 different time-series including (i) interest rates, (ii) yield spreads, and (iii) other indicators which include variables like the VIX. A positive (negative) index corresponds to a higher (lower) degree of stress compared to the trend.

Even though we have implied variance available for different tenors, we focus on the longest maturity, i.e., the 30y Treasury futures options. To simplify notation, we call implied volatility from options on 30y Treasury futures, TIV (Treasury Implied Volatility, the implied variance is then TIV^2). We denote $\mathrm{slope}^{\mathrm{TIV}}$ the slope of the term structure of implied variances which is defined as the difference between the implied variance from a one-year and one-month option.

Figure 8 plots $\mathrm{slope}^{\mathrm{TIV}}$ together with the CFNAI, and the STLFSI multiplied by minus one. The co-movement between the three time series is strikingly high, especially during the recent crisis period. Indeed, the unconditional correlation between the slope and the CFNAI is 68% and it is 76% with the STLFSI. To test the relationship more formally, we run predictive regressions from the economic activity/stress index onto TIV^2, $\mathrm{slope}^{\mathrm{TIV}}$, VIX^2, the slope of the implied variance on the S&P500 ($\mathrm{slope}^{\mathrm{S\&P500}}$), and the slope of the term structure at different horizons, n, ranging from zero to twelve months:

$$\mathrm{CFNAI}_{t+n}/\mathrm{STLFSI}_{t+n} = \beta_n^{\mathrm{TIV}}\mathrm{TIV}_t^2 + \beta_n^{\mathrm{slope\ TIV}}\mathrm{slope}_t^{\mathrm{TIV}} + \ldots + \epsilon_{t+n}.$$

The results are reported in Tables IX and X.[23] TIV^2 is an excellent predictor of future economic activity up to eight months: For the contemporaneous regressions, we find that for any one standard deviation shock in the TIV^2, there is a 0.3 standard deviation shock to economic activity. t-statistics range from -6.45 to -3.01 for the eight month horizon. The R^2s drop fast with the horizon: For contemporaneous regressions, the R^2 is 29% but it drops by half, i.e., to 14% after four months. Similarly to the level, the slope of the TIV has strong predictive power up to eight months with R^2s monotonically decreasing from 20% for contemporaneous regressions to virtually zero for horizons longer than eight months. When we include both the level and the slope in the regression, the significance of both factors is not affected.

Adding VIX^2 or the slope of the VIX term structure does not alter the results much: The predictive power of the slope of the TIV is quantitatively the same, however, some of the predictive power of TIV^2 is subsumed by VIX^2. Adding the slope of the Treasury yield term structure itself does not change the results either. The slope of the term structure

[23] Variables are standardized, meaning, we de-mean and divide by the standard deviation.

is known to be a long-term predictor of the business cycle rather than a predictor of short-term fluctuations. Hence, at horizons up to eight months, we observe basically no predictive power from the slope of the term structure, whereas at longer horizons, where the other right-hand side variables lose any power, the slope of the term structure becomes significant.

Insert Figure 8 and Tables IX and X here.

We run the same exercise but now use the stress index as the left-hand side variable.[24] The results are reported in Table X. The results depict a similar picture as for the activity index but with opposite signs: An increase in TIV^2 implies an increase in the stress index. Estimated coefficients are significant up to four months. The estimated slope coefficients of $\text{slope}^{\text{TIV}}$ are negative and highly significant up to a horizon of eight months. The effects are not only statistically but also economically significant: For any one standard deviation change in the slope, there is almost a 0.8 standard deviation change in the stress index.

We conclude that both TIV^2 and the slope of the TIV are excellent predictors of future economic activity or stress. The predictive power is only at the short horizon and dies out after eight months. One might now obviously suspect that our results are mainly driven by the large co-movement between all series during the summer of 2008. We therefore run robustness checks and study the predictive power using a sample which ends in August 2008. Indeed, the TIV^2 has no predictive power for CFNAI if we end the sample in 2008, however, the slope of the TIV is still significant but only at short horizons. For the stress index, we still find that both TIV^2 and the slope are highly significant. To save space, these results are deferred to the Online Appendix.

4 Conclusions

This paper studies the returns of variance trading strategies in the Treasury options market. We first derive theoretically how to replicate variance swaps. The fair strike can

[24] We do not run a regression which includes the slope from yields, VIX^2, or the slope of the VIX as the stress index contains both the slope of the term structure and the VIX itself.

easily be constructed using a continuum of put and call options. The way we formalize the realized leg is critical: Instead of using squared log returns, we use simple returns which allows for a perfect replication of the variance swap payout even in the presence of jumps and regardless of the sampling partition.

Using a large panel data set of Treasury options with different tenors, we juxtapose the variance swap with delta-hedged ATM straddle strategies—an alternative strategy well known to be very profitable in the equity markets. Trading variance through straddles is attractive but the profitability of these strategies is dwarfed by the variance swaps: The average excess return on a variance swap is about 20% per month regardless of the tenor. Moreover, these returns come at reasonable risk as the annualized Sharpe ratio is above two. We then examine how and whether these returns can be explained by standard risk factors such as the market, book-to-market, size, momentum, or liquidity factors. We find that none of these factors have significant explanatory power for the returns of variance swaps. The alpha of the strategy is large and highly significant. Using realistic assumptions on margins, we verify that variance trading in fixed income markets yields substantial profits even when investors cannot rely on an unlimited supply of capital to implement their strategies.

We also study the term structure of implied variances and variance risk premiums. We document that variance risk premiums are negative and their term structure is downward sloping in absolute terms. At longer maturities, the premiums are smaller in magnitude. The term structure of implied variances is also downward sloping on average while over time, it is highly time-varying and strongly pro-cyclical. In particular, the slightly negative slope becomes extremely negative during crisis periods. Using this observation, we find that both the level of implied volatility and the slope are excellent predictors of both economic activity and stress, especially at short horizons.

References

Aït-Sahalia, Y., M. Karaman, and L. Mancini (2015): "The Term Structure of Variance Swaps and Risk Premia," Working Paper, Princeton University.

Allen, P., S. Einchomb, and N. Granger (2006): "Variance Swaps," Discussion paper, JPMorgan.

Andries, M., T. Eisenbach, M. Schmalz, and Y. Wang (2015): "Price vs. Quantity in the Term Structure of Variance Risk Premia," Working Paper, University of Toulouse.

Bekaert, G., and M. Hoerova (2014): "The VIX, the variance premium and stock market volatility," *Journal of Econometrics*, 183, 181–192.

Black, F. (1976): "The Pricing of Commodity Contracts," *Journal of Financial Economics*, 3, 167–179.

Black, F., and M. Scholes (1973): "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637–659.

Bondarenko, O. (2014): "Variance Trading and Market Price of Variance Risk," *Journal of Econometrics*, 180, 81–97.

Broadie, M., and A. Jain (2008): "The Effect of Jumps and Discrete Sampling on Volatility and Variance Swaps," *International Journal of Theoretical and Applied Finance*, 11, 761–797.

Cakici, N., and J. Zhu (2001): "Pricing Eurodollar Futures Options with the Heath-Jarrow-Morton Model," *Journal of Futures Markets*, 21, 655–680.

Carr, P., and R. Lee (2009): "Volatility Derivatives," *Annual Review of Financial Economics*, 1, 1–21.

Carr, P., and D. Madan (1998): "Towards a Theory of Volatility Trading," in *Volatility: New Estimation Techniques for Pricing Derivatives*, ed. by R. Jarrow, pp. 417–427. RISK Publications.

Carr, P., and L. Wu (2009): "Variance Risk Premiums," *Review of Financial Studies*, 22(3), 1311–1341.

Cieslak, A., and P. Povala (2016): "Information in the Term Structure of Yield Curve Volatility," *Journal of Finance*, 71, 1393–1434.

Coval, J., and T. Shumway (2001): "Expected Option Returns," *Journal of Finance*, 56, p. 983–1009.

Demeterfi, K., E. Derman, M. Kamal, and J. Zhou (1999): "A Guide to Volatility and Variance Swaps," *Journal of Derivatives*, 6, 9–32.

Dew-Becker, I., S. Giglio, A. Le, and M. Rodriguez (2016): "The Price of Variance Risk," *forthcoming, Journal of Financial Economics.*

Duarte, J., F. Longstaff, and F. Yu (2007): "Risk and Return in Fixed-Income Arbitrage: Nickels in Front of a Steamroller?," *Review of Financial Studies*, 20, 769–811.

Dupire, B. (1994): "Pricing with a Smile," *Risk*, 7, 18–20.

Duyvesteyn, J., and G. de Zwart (2015): "Riding the swaption curve," *Journal of Banking & Finance*, 59, 57–75.

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, 1279–1310.

Filipović, D., E. Gourier, and L. Mancini (2015): "Quadratic Variance Swap Models," *forthcoming, Journal of Financial Economics*.

Flesaker, B. (1993): "Arbitrage Free Pricing of Interest Rate Futures and Forward Contracts," *Journal of Futures Markets*, 13, 77–91.

George, T. J., and F. A. Longstaff (1993): "Bid-Ask Spreads and Trading Activity in the S&P100 Index Option Market," *Journal of Financial and Quantitative Analysis*, 28, 381–397.

Litterman, R., and J. Scheinkman (1991): "Common Factors Affecting Bond Returns," *Journal of Fixed Income*, 1, 54–61.

Lo, A. W. (2002): "The Statistics of Sharpe Ratios," *Financial Analyst Journal*, 58, 36–52.

Lustig, H., N. Roussanov, and A. Verdelhan (2011): "Common Risk Factors in Currency Markets," *Review of Financial Studies*, 24(11), 3731–3777.

Malkhozov, A., P. Mueller, A. Vedolin, and G. Venter (2016a): "International Illiquidity," Working Paper, London School of Economics.

——— (2016b): "Mortgage Risk and the Yield Curve," *Review of Financial Studies*, 29(5), 1220–1253.

Martin, I. (2013): "Simple Variance Swaps," Working Paper, London School of Economics.

Mele, A., and Y. Obayashi (2013): "Volatility Indexes and Contracts for Government Bonds and Time Deposits," Working Paper, Applied Academics LLC.

Merener, N. (2012): "Swap Rate Variance Swaps," *Quantitative Finance*, 12, 249–261.

Mueller, P., P. Sabtchevsky, A. Vedolin, and P. Whelan (2016): "Variance Risk Premia on Stocks and Bonds," Working Paper, London School of Economics.

Neuberger, A. (1994): "Hedging Volatility - the Case for a New Contract," *Journal of Portfolio Management*, 20(2), 74–80.

——— (2012): "Realized Skewness," *Review of Financial Studies*, 25, 3423–3455.

Newey, W. K., and K. D. West (1987): "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55(3), 703–708.

Pástor, L., and R. F. Stambaugh (2003): "Liquidity Risk and Expected Stock Returns," *Journal of Political Economy*, 111, 642–685.

Santa-Clara, P., and A. Saretto (2009): "Option Strategies: Good Deals and Margin Calls," *Journal of Financial Markets*, 12, 391–417.

Schneider, P. (2015): "Generalized Risk Premia," *Journal of Financial Economics*, 116, 487–504.

Schneider, P., and F. Trojani (2015): "Fear Trading," Working Paper, University of Lugano.

Trolle, A. B. (2009): "The Price of Interest Rate Variance Risk And Optimal Investments in Interest Rate Derivatives," Working Paper, EPFL.

Trolle, A. B., and E. S. Schwartz (2014): "The Swaption Cube," *Review of Financial Studies*, 27, 2307–2353.

van Binsbergen, J. H., and R. S. Koijen (2016): "The Term Structure of Returns: Facts and Theory," *forthcoming, Journal of Financial Economics.*

Wu, L. (2010): "Variance Dynamics: Joint Evidence from Options and High-Frequency Returns," *Journal of Econometrics*, 160, 280–287.

Appendix A Proofs and Derivations

Proof of Proposition 1. Following Carr and Madan (1998), we assume that there exist a function $g(F_{T,T})$ and that it is twice differentiable. It then follows that

$$\begin{aligned} g(F_{T,T}) &= g(x) + g'(x)(F_{T,T} - x) + \int_0^x g''(K)(K - F_{T,T})^+ dK \\ &+ \int_x^\infty g''(K)(F_{T,T} - K)^+ dK \end{aligned}$$

for any $x \geq 0$. If $x = F_{t,T}$, then

$$\begin{aligned} g(F_{T,T}) &= g(F_{t,T}) + g'(F_{t,T})(F_{T,T} - F_{t,T}) + \int_0^{F_{t,T}} g''(K)(K - F_{T,T})^+ dK \\ &+ \int_{F_{t,T}}^\infty g''(K)(F_{T,T} - K)^+ dK \end{aligned}$$

Now assumr $g(F) = \log F$. Then, we get

$$\begin{aligned} \log F_{T,T} &= \log F_{t,T} + \frac{1}{F_{t,T}}(F_{T,T} - F_{t,T}) \\ &- \left(\int_0^{F_{t,T}} \frac{(K - F_{T,T})^+}{K^2} dK + \int_{F_{t,T}}^\infty \frac{(F_{T,T} - K)^+}{K^2} dK \right). \end{aligned} \tag{A-1}$$

Given that $d(\log F_{t,T}) = \frac{dF_{t,T}}{F_{t,T}} - \frac{1}{2}\sigma_t^2 dt$ due to Itô's lemma, the quadratic variation of $F_{t,T}$ can be written as

$$\begin{aligned} \int_t^T \sigma_u^2 du &= -2 \log \frac{F_{T,T}}{F_{t,T}} + \int_t^T \frac{dF_{u,T}}{F_{u,T}} \\ &= 2\left(\frac{F_{T,T} - F_{t,T}}{F_{t,T}} - \log \frac{F_{T,T}}{F_{t,T}} \right) + 2 \int_t^T \left(\frac{1}{F_{u,T}} - \frac{1}{F_{t,T}} \right) dF_{u,T} \\ &= 2 \left(\int_0^{F_{t,T}} \frac{(K - F_{T,T})^+}{K^2} dK + \int_{F_{t,T}}^\infty \frac{(F_{T,T} - K)^+}{K^2} dK \right) \\ &+2 \int_t^T \left(\frac{1}{F_{u,T}} - \frac{1}{F_{t,T}} \right) dF_{u,T}. \end{aligned}$$

where the last equality follows from Equation (A-1). Since $F_{t,T}$ is a martingale under the $\mathbb{Q}_T$ measure, the $\mathbb{Q}_T$ expectation of the dynamic strategy is zero and hence this implies that the market price of the realized variance is equal to $\tilde{F}_{t,T}$, i.e.

$$\tilde{F}_{t,T} = \mathbb{E}_t^{\mathbb{Q}_T} \left[\int_t^T \sigma_u^2 du \right].$$

In words, replicating the expression for the variance $\mathrm{RV}_{t,T}^{\log}$ involves

1. a path-independent payoff in out-of-the-money options

2. a dynamic trading strategy which is re-balanced to hold $2\left(\frac{1}{F_{s,T}}-\frac{1}{F_{t,T}}\right)$ of the underlying at any time $s\in[t,T]$.

$\square$

Proof of Proposition 2.

$$\begin{aligned}
\widetilde{RV}_{t,T} &= -2\log\frac{F_{T,T}}{F_{t,T}}+2\sum_{i=1}^{T-t}\frac{F_{t+i,T}-F_{t+i-1,T}}{F_{t+i-1,T}} \\
&= 2\left(\frac{F_{T,T}-F_{t,T}}{F_{t,T}}-\log\frac{F_{T,T}}{F_{t,T}}\right) \\
&\quad +2\sum_{i=1}^{T-t}\left(\frac{1}{F_{t+i-1,T}}-\frac{1}{F_{t,T}}\right)(F_{t+i,T}-F_{t+i-1,T}) \\
&= 2\left(\int_0^{F_{t,T}}\frac{(K-F_{T,T})^+}{K^2}dK+\int_{F_{t,T}}^{\infty}\frac{(F_{T,T}-K)^+}{K^2}dK\right) \\
&\quad +2\sum_{i=1}^{T-t}\left(\frac{1}{F_{t+i-1,T}}-\frac{1}{F_{t,T}}\right)(F_{t+i,T}-F_{t+i-1,T})
\end{aligned}$$

where the last equality follows from Equation (A-1). $\square$

Relationship between squared log returns and new variance measure. The calculations follow Carr and Lee (2009). We start with a Taylor expansion of $2\log F_{t+i,T}$ around $F_{t+i-1,T}$,

$$\begin{aligned}
2\log F_{t+i,T} &= 2\log F_{t+i-1,T}+\frac{2}{F_{t+i-1,T}}(F_{t+i,T}-F_{t+i-1,T})-\left(\frac{F_{t+i,T}-F_{t+i-1,T}}{F_{t+i-1,T}}\right)^2+ \\
&\quad \frac{2}{3}\left(\frac{F_{t+i,T}-F_{t+i-1,T}}{F_{t+i-1,T}}\right)^3+O\left(r_{t+i,T}^4\right) \\
&= 2\log F_{t+i-1,T}+2r_{t+i,T}-r_{t+i,T}^2+\frac{2}{3}r_{t+i,T}^3+O\left(r_{t+i,T}^4\right), \qquad \text{(A-2)}
\end{aligned}$$

where $r_{t+i,T}:=F_{t+i,T}/F_{t+i-1,T}-1$. Rearranging the right and left hand sides of the above equation yields

$$\log\frac{F_{t+i,T}}{F_{t+i-1,T}}=r_{t+i,T}-\frac{1}{2}r_{t+i,T}^2+\frac{1}{3}r_{t+i,T}^3+O\left(r_{t+i,T}^4\right).$$

Squaring both sides,

$$\left(\log\frac{F_{t+i,T}}{F_{t+i-1,T}}\right)^2=r_{t+i,T}^2-r_{t+i,T}^3+O\left(r_{t+i,T}^4\right).$$

Solving the above for $r_{t+i,T}^2$ and substituting it into Equation (A-2), we have

$$2\log F_{t+i,T}=2\log F_{t+i-1,T}+2r_{t+i,T}-\left(\log\frac{F_{t+i,T}}{F_{t+i-1,T}}\right)^2-\frac{1}{3}r_{t+i,T}^3+O\left(r_{t+i,T}^4\right)$$

which implies

$$\left(\log \frac{F_{t+i,T}}{F_{t+i-1,T}}\right)^2 = 2\left(r_{t+i,T} - \log \frac{F_{t+i,T}}{F_{t+i-1,T}}\right) - \frac{1}{3} r^3_{t+i,T} + O\left(r^4_{t+i,T}\right)$$

Finally, summing over different i's, we get

$$RV^{\log}_{t,T} = \widetilde{RV}_{t,T} - \frac{1}{3}\sum_{i=1}^{T-t} r^3_{t+i,T} + \sum_{i=1}^{T-t} O\left(r^4_{t+i,T}\right)$$

where $RV^{\log}_{t,T} = \sum_{i=1}^{T-t}\left[\log\left(1+r_{t+i,T}\right)\right]^2$ and $\widetilde{RV}_{t,T} = 2\sum_{i=1}^{T-t}\left[r_{t+i,T} - \log\left(1+r_{t+i,T}\right)\right]$. $\square$

Table I
Summary Statistics Treasury Implied Volatilities

The table reports means, standard deviations, minima and maxima of annualized implied volatilities for the three different tenors (30y, 10y, and 5y) and different maturities ranging from ten days to one year. Volatilities are the square root of the implied variances extracted from daily option prices using equation (6). All numbers are annualized and expressed in percent. Data is monthly and runs from June 1990 to May 2012.

maturity	10d	20d	1m	3m	6m	9m	12m
	Panel A: 30y Treasury						
mean	11.83	11.21	10.98	10.47	10.35	10.06	9.88
std	3.57	2.85	2.65	2.41	2.11	2.09	2.13
min	4.21	6.01	6.60	5.78	6.07	5.82	5.41
max	31.18	26.20	24.15	21.12	20.30	20.02	19.88
	Panel B: 10y Treasury						
mean	8.27	7.65	7.43	6.90	6.87	6.77	6.72
std	2.56	1.77	1.63	1.54	1.66	1.86	1.98
min	0.31	4.32	4.26	2.99	3.11	2.91	2.66
max	17.08	14.76	14.33	13.08	12.52	13.69	14.24
	Panel C: 5y Treasury						
mean	5.54	4.99	4.79	4.49	4.49	4.48	4.47
std	1.94	1.29	1.16	1.22	1.36	1.45	1.50
min	0.74	2.34	2.59	2.06	1.77	1.46	1.28
max	11.88	9.90	9.69	9.40	9.33	9.30	9.29

Table II
Term Structure of Variance Risk Premiums: Means and Sharpe Ratios

The table reports average variance risk premiums (mean) and average Sharpe ratios (SR) for the three different tenors (30y, 10y, and 5y) and different maturities ranging from ten days to one year. Variance risk premiums are computed by subtracting the implied variance as in equation (6) from the ex-post realized variance as in equation (5). They are monthly and expressed in squared percent. Sharpe ratios are calculated as the average variance risk premiums divided by the corresponding standard deviation of the monthly variance risk premiums. Panel A reports the results for the benchmarks sample period from June 1990 until May 2012 while Panel B shows the corresponding numbers for the maximally available data set that goes back to October 1982, May 1985, and June 1990 for the respective tenors, 30y, 10y, and 5y. Data is sampled monthly.

maturity		10d	20d	1m	3m	6m	9m	12m
Panel A: June 1990 - May 2012								
mean	30y	-4.90	-3.22	-2.41	-1.35	-0.88	-0.32	0.08
	10y	-3.15	-1.81	-1.33	-0.61	-0.49	-0.46	-0.40
	5y	-1.08	-0.61	-0.44	-0.16	-0.25	-0.27	-0.37
SR	30y	-0.81	-0.67	-0.51	-0.34	-0.23	-0.08	0.02
	10y	-0.94	-0.76	-0.69	-0.36	-0.25	-0.22	-0.18
	5y	-0.62	-0.48	-0.45	-0.19	-0.24	-0.25	-0.24
Panel B: all available data								
mean	30y	-6.82	-3.75	-2.84	-1.31	-0.56	0.18	0.66
	10y	-4.68	-2.46	-1.62	-0.52	-0.34	-0.26	-0.17
	5y	-1.08	-0.61	-0.44	-0.16	-0.25	-0.27	-0.37
SR	30y	-0.64	-0.50	-0.45	-0.26	-0.12	0.04	0.15
	10y	-0.67	-0.54	-0.60	-0.23	-0.14	-0.10	-0.07
	5y	-0.62	-0.48	-0.45	-0.19	-0.24	-0.25	-0.24

Table III
Summary Statistics Option Trading Strategies

This table reports monthly summary statistics for one-month returns on three different trading strategies described in Section 3.1: mean, standard deviation, maximum, skewness, kurtosis, Sharpe ratio (SR), and alpha. Panel A presents the summary statistics of monthly returns on the generalized Treasury variance swap defined in equation (7) across three tenors (5y, 10y and 30y) and with a maturity $T-t=1$ month. Panel B and Panel C report the summary statistics of monthly returns on un-hedged or delta-hedged at-the-money straddles with a maturity of one month, respectively. The delta-hedged straddle return is defined in equation (8). The un-hedged return is the same but without the accumulated profit and loss from the hedging activity, i.e., $\sum_{i=t+1}^{T} -\Delta_{S,i-1}\left(F_{i,T}-F_{i-1,T}\right)$ in equation (8). In addition, we also report the corresponding summary statistics for the trading strategies using options on S&P500 index futures. t-Statistics reported in parentheses are adjusted according to Newey and West (1987). Data is sampled monthly and runs from June 1990 to May 2012.

Panel A: Variance Swap

	mean	t-stat	std	max	skew	kurt	SR	alpha	t-stat
30y	-0.212	(-8.98)	0.382	2.167	2.128	7.805	-0.554	-0.202	(-8.54)
10y	-0.276	(-12.71)	0.353	2.123	1.893	8.104	-0.784	-0.270	(-11.14)
5y	-0.187	(-6.61)	0.460	2.313	2.191	7.857	-0.407	-0.181	(-5.39)
S&P500	-0.314	(-9.40)	0.542	4.182	3.711	21.634	-0.580	-0.279	(-7.33)

Panel B: Straddles not hedged

	mean	t-stat	std	max	skew	kurt	SR	alpha	t-stat
30y	0.007	(0.17)	0.710	2.409	0.912	3.723	0.011	0.005	(0.10)
10y	-0.019	(-0.42)	0.724	4.097	1.227	6.461	-0.026	-0.018	(-0.36)
5y	0.033	(0.70)	0.758	4.340	1.365	6.876	0.044	0.023	(0.45)
S&P500	-0.125	(-2.95)	0.686	2.895	1.379	5.607	-0.182	-0.150	(-3.22)

Panel C: Straddles hedged

	mean	t-stat	std	max	skew	kurt	SR	alpha	t-stat
30y	-0.034	(-2.47)	0.219	1.167	0.831	6.614	-0.153	-0.034	(-2.19)
10y	-0.070	(-4.91)	0.231	0.711	0.146	3.287	-0.304	-0.069	(-4.31)
5y	-0.045	(-2.99)	0.244	0.815	0.366	3.611	-0.185	-0.042	(-2.53)
S&P500	-0.075	(-4.11)	0.296	1.421	1.599	7.537	-0.254	-0.081	(-4.04)

Table IV
Bootstrapped Confidence Intervals

This table reports 95% bootstrapped confidence intervals for mean, Sharpe ratio (SR), and alpha for three trading strategies presented in Table III. The empirical distribution of returns is obtained from 10,000 studentized bootstrap repetitions of our sample. Data is monthly and runs from June 1990 to May 2012.

Panel A: Variance Swap						
	mean		SR		alpha	
30y	-0.248	-0.177	-0.741	-0.415	-0.205	-0.199
10y	-0.311	-0.250	-1.050	-0.643	-0.272	-0.267
5y	-0.242	-0.138	-0.581	-0.254	-0.184	-0.177
S&P500	-0.376	-0.252	-0.851	-0.373	-0.282	-0.276

Panel B: Straddles not hedged						
	mean		SR		alpha	
30y	-0.097	0.078	-0.128	0.091	-0.056	0.042
10y	-0.099	0.060	-0.163	0.047	-0.041	0.048
5y	-0.051	0.116	-0.066	0.121	-0.032	0.280
S&P500	-0.167	-0.048	-0.316	-0.078	-0.200	-0.104

Panel C: Straddles hedged						
	mean		SR		alpha	
30y	-0.065	-0.014	-0.287	-0.045	-0.042	-0.021
10y	-0.091	-0.045	-0.432	-0.178	-0.088	-0.070
5y	-0.063	-0.014	-0.276	-0.056	-0.068	-0.027
S&P500	-0.099	-0.037	-0.415	-0.117	-0.095	-0.072

Table V
Risk-Adjusted Returns Trading Strategies

This table reports Newey and West (1987) adjusted t-statistics of OLS coefficients from regressing the returns of three different trading strategies (the one-month generalized Treasury variance swap, un-hedged straddle, and delta-hedged straddle) on the market excess return (MRKT), size (SMB), book-to-market (HML), momentum factors (MOM), and an illiquidity factor.

$$r_t^i = \alpha + \beta_1 r_t^{\text{MRKT}} + \beta_2 r_t^{\text{SMB}} + \beta_3 r_t^{\text{HML}} + \beta_4 r_t^{\text{MOM}} + \beta_5 r_t^{\text{ILLIQ}} + \epsilon_t$$

For the bond option trading strategies we use the Barclays US Treasury bond index as a proxy for the market return. For the S&P500 futures strategies we use the excess return on the value-weighted return of all CRSP firms. The illiquidity measure for the bond market is taken from Malkhozov, Mueller, Vedolin, and Venter (2016a) and for the equity market we use the negative of the Pástor and Stambaugh (2003) liquidity factor such that a high value again measures illiquidity. The first column reports Jensen's alpha together with its t-statistic in parentheses. Data is monthly and runs from June 1990 to May 2012.

Panel A: Variance Swap

	alpha	t-stat	MRKT	SMB	HML	MOM	ILLIQ	Adj. R^2
30y	-0.261	(-4.44)	-2.317	-3.061	1.924	-2.572	1.562	10.37%
10y	-0.418	(-8.21)	-1.997	-2.195	1.839	-2.483	3.211	10.48%
5y	-0.336	(-3.80)	-0.370	-2.106	2.466	-2.082	1.697	4.79%
S&P500	-0.321	(-9.65)	0.146	-1.133	-1.376	-1.653	4.249	31.72%

Panel B: Straddles not hedged

	alpha	t-stat	MRKT	SMB	HML	MOM	ILLIQ	Adj. R^2
30y	-0.031	(-0.36)	0.264	-0.660	0.329	-0.985	0.591	0.72%
10y	-0.082	(-0.93)	0.261	-0.907	1.277	-1.076	0.903	1.61%
5y	-0.025	(-0.27)	0.895	-0.522	1.554	-1.018	0.564	1.56%
S&P500	-0.230	(-4.07)	0.111	-0.313	1.304	-1.126	0.440	2.76%

Panel C: Straddles hedged

	alpha	t-stat	MRKT	SMB	HML	MOM	ILLIQ	Adj. R^2
30y	-0.071	(-2.58)	0.180	-0.238	0.888	-0.185	1.623	1.58%
10y	-0.125	(-4.48)	0.581	-1.191	1.922	-1.040	2.482	5.14%
5y	-0.083	(-2.83)	0.566	-0.245	2.410	-1.067	1.513	3.82%
S&P500	-0.176	(-4.67)	1.091	-1.095	0.410	-0.334	2.628	5.28%

Table VI
Summary Statistics Option Trading Strategies Bid-Ask Spread Adjusted

This table reports monthly summary statistics for one-month returns on three different trading strategies described in Section 3.1, taking into account bid and ask spreads: mean, standard deviation, Sharpe ratio (SR), and alpha. First, we report returns on the generalized Treasury variance swap defined in equation (7) across three tenors (5y, 10y and 30y) and with a maturity $T - t = 1$ month. Second, we report the summary statistics of monthly returns on un-hedged or delta-hedged at-the-money straddles with a maturity of one month, respectively. The delta-hedged straddle return is defined in equation (8). The un-hedged return is the same but without the accumulated profit and loss from the hedging activity, i.e., $\sum_{i=t}^{T} -\Delta_{S,i-1}\left(F_{i,T} - F_{i-1,T}\right)$ in equation (8). t-Statistics reported in parentheses are adjusted according to Newey and West (1987). Data is sampled monthly and runs from June 1990 to May 2012.

	Variance Swap			Straddles not hedged			Straddles hedged		
	30y	10y	5y	30y	10y	5y	30y	10y	5y
mean	-0.138	-0.254	-0.173	-0.001	-0.004	0.020	-0.021	-0.062	-0.035
t-stat	(-5.32)	(-11.42)	(-5.95)	(-0.01)	(-0.21)	(0.41)	(-2.11)	(-4.24)	(-2.57)
std	0.421	0.362	0.474	1.625	0.309	0.793	0.162	0.238	0.221
SR	-0.327	-0.703	-0.364	-0.001	-0.013	0.025	-0.130	-0.261	-0.158
alpha	-0.127	-0.245	-0.166	-0.002	-0.005	0.014	-0.027	-0.046	-0.038
t-stat	(-5.00)	(-10.69)	(-5.17)	(-0.00)	(-0.01)	(0.21)	(-1.96)	(-3.15)	(-3.26)

Table VII
Treasury Variance Swaps and Margin Requirements

This table analyzes the impact of margin requirements on the returns to writing one-month Treasury variance swaps. The table reports the effective vega notional, monthly average and monthly Sharpe ratio (SR) of margined returns at different levels of target vega notional in addition to the corresponding un-margined returns. The proportion of months with forced rescaling (out of 264) is reported in the third column of the table. 95% confidence intervals (in brackets) are obtained using a bootstrap with 10,000 draws. Data is monthly and runs from June 1990 to May 2012.

	margined				un-margined	
target vega	effective vega notional	rescaled month (fraction)	mean	SR	mean	SR
Panel A: 30y Treasury						
0.010	0.010 [0.010,0.010]	0.00 [0.00,0.00]	0.01 [0.00,0.01]	0.33 [0.18,0.50]	0.01 [0.00,0.01]	0.33 [0.19,0.50]
0.050	0.050 [0.050,0.050]	0.01 [0.00,0.02]	0.03 [0.02,0.05]	0.32 [0.18,0.50]	0.03 [0.02,0.05]	0.33 [0.19,0.50]
0.075	0.068 [0.067,0.069]	0.52 [0.46,0.58]	0.05 [0.03,0.07]	0.29 [0.15,0.47]	0.05 [0.03,0.07]	0.33 [0.19,0.50]
0.100	0.067 [0.066,0.068]	1.00 [1.00,1.00]	0.05 [0.03,0.06]	0.30 [0.15,0.47]	0.07 [0.04,0.09]	0.33 [0.19,0.50]
Panel B: 10y Treasury						
0.010	0.010 [0.010,0.010]	0.00 [0.00,0.00]	0.01 [0.01,0.01]	0.53 [0.37,0.72]	0.01 [0.01,0.01]	0.53 [0.37,0.72]
0.050	0.050 [0.050,0.050]	0.00 [0.00,0.03]	0.04 [0.03,0.04]	0.52 [0.37,0.71]	0.04 [0.03,0.04]	0.53 [0.37,0.72]
0.075	0.060 [0.059,0.061]	0.93 [0.90,0.96]	0.04 [0.03,0.05]	0.49 [0.34,0.68]	0.05 [0.04,0.07]	0.53 [0.37,0.72]
0.100	0.060 [0.059,0.060]	1.00 [1.00,1.00]	0.04 [0.03,0.05]	0.50 [0.34,0.68]	0.07 [0.06,0.09]	0.53 [0.37,0.72]
Panel C: 5y Treasury						
0.010	0.010 [0.010,0.010]	0.00 [0.00,0.00]	0.00 [0.00,0.01]	0.43 [0.28,0.62]	0.00 [0.00,0.01]	0.43 [0.28,0.62]
0.050	0.048 [0.047,0.048]	0.25 [0.20,0.31]	0.02 [0.01,0.03]	0.44 [0.29,0.63]	0.02 [0.01,0.03]	0.43 [0.28,0.62]
0.075	0.050 [0.049,0.051]	1.00 [1.00,1.00]	0.02 [0.02,0.03]	0.43 [0.28,0.62]	0.03 [0.02,0.04]	0.43 [0.28,0.62]
0.100	0.050 [0.049,0.051]	1.00 [1.00,1.00]	0.02 [0.02,0.03]	0.43 [0.28,0.62]	0.04 [0.03,0.05]	0.43 [0.28,0.62]

Table VIII
Summary Statistics Variance Swap Different RV

This table reports monthly summary statistics, mean, standard deviation, kurtosis, skewness, and the Sharpe ratio (SR) for one-month variance swap returns that are based on different measures of realized variance. GTVS corresponds to the variance swap payoff using realized variance as defined in equation (5) that allows for perfect replication. LTVS corresponds to a variance swap that uses realized variance calculated using daily squared log returns. Data is monthly and runs from June 1990 to May 2012.

	30y Treasury		10y Treasury		5y Treasury	
	GTVS	LTVS	GTVS	LTVS	GTVS	LTVS
mean	-0.212	-0.212	-0.276	-0.276	-0.187	-0.187
t-stat	(-8.98)	(-8.98)	(-12.71)	(-12.71)	(-6.61)	(-6.61)
std	0.383	0.383	0.353	0.353	0.460	0.460
kurtosis	10.635	10.824	10.929	11.163	10.686	10.896
skewness	2.116	2.131	1.882	1.901	2.178	2.197
SR	-0.553	-0.552	-0.782	-0.782	-0.407	-0.406

Table IX
Predictive Regressions Economic Activity

This table reports estimated coefficients from predictive regressions from the Chicago Fed National Activity Index (CFNAI) onto the Treasury implied volatility index squared (TIV^2), the slope of the TIV, and other factors:

$$\mathrm{CFNAI}_{t+n} = \beta_n^{\mathrm{TIV}} \mathrm{TIV}_t^2 + \beta_n^{\mathrm{slope\ TIV}} \mathrm{slope}_t^{\mathrm{TIV}} + \ldots + \epsilon_{t+n}, \quad n = 0, \ldots, 12.$$

Variables are standardized, meaning we de-mean and divide by the standard deviation. t-Statistics presented in parentheses are calculated using Newey and West (1987). Regressions are run contemporaneously and for forecast horizons up to twelve months. Data is monthly and runs from June 1990 to May 2012.

horizon	contemp.	3m	6m	9m	12m
TIV^2	-0.302	-0.226	-0.121	-0.051	-0.038
	(-6.45)	(-5.24)	(-3.30)	(-2.51)	(-1.34)
Adj. R^2	28.89%	16.20%	4.64%	0.81%	0.47%
slope TIV	0.622	0.552	0.327	0.152	0.121
	(3.11)	(4.49)	(3.61)	(1.72)	(0.99)
Adj. R^2	19.71%	15.46%	5.32%	1.16%	0.74%
TIV^2	-0.237	-0.154	-0.074	-0.026	-0.018
	(-6.05)	(-4.21)	(-1.80)	(-0.69)	(-0.38)
slope TIV	0.330	0.360	0.232	0.118	0.098
	(3.74)	(3.47)	(2.23)	(0.91)	(0.59)
Adj. R^2	32.84%	20.83%	6.28%	0.94%	0.43%
TIV^2	-0.172	-0.022	0.019	0.004	0.010
	(-3.43)	(-0.38)	(0.31)	(0.07)	(0.13)
slope TIV	0.339	0.375	0.261	0.183	0.179
	(3.01)	(3.06)	(2.01)	(1.12)	(0.84)
VIX^2	-0.064	-0.131	-0.097	-0.046	-0.046
	(-1.22)	(-2.08)	(-1.58)	(-0.79)	(-0.72)
slope VIX	-0.037	-0.073	-0.101	-0.173	-0.210
	(-0.32)	(-0.67)	(-0.71)	(-0.95)	(-1.09)
Adj. R^2	33.52%	25.20%	7.66%	1.23%	1.37%
TIV^2	-0.188	-0.015	0.039	0.038	0.054
	(-3.82)	(-0.26)	(0.69)	(0.63)	(0.71)
slope TIV	0.309	0.388	0.298	0.243	0.261
	(2.90)	(3.11)	(2.24)	(1.41)	(1.19)
VIX^2	-0.053	-0.136	-0.112	-0.072	-0.081
	(-1.03)	(-2.16)	(-1.81)	(-1.17)	(-1.14)
slope VIX	-0.014	-0.084	-0.135	-0.232	-0.288
	(-0.13)	(-0.78)	(-0.94)	(-1.23)	(-1.39)
slope yields	-0.099	0.044	0.124	0.204	0.268
	(-1.73)	(0.72)	(1.65)	(2.18)	(2.38)
Adj. R^2	34.34%	25.12%	8.96%	5.40%	9.14%

Table X
Predictive Regressions Stress Index

This table reports estimated coefficients from predictive regressions from the St. Louis Fed Stress Index (STLFSI) onto the Treasury implied volatility index squared (TIV^2) and the slope of the TIV:

$$\text{STLFSI}_{t+n} = \beta_n^{\text{TIV}} \text{TIV}_t^2 + \beta_n^{\text{slope TIV}} \text{slope}_t^{\text{TIV}} + \epsilon_{t+n}, \quad n = 0, \ldots, 12.$$

Variables are standardized, meaning we de-mean and divide by the standard deviation. t-Statistics presented in parentheses are calculated using Newey and West (1987). Regressions are run contemporaneously and for forecast horizons up to twelve months. Data is monthly and runs from January 1994 to May 2012.

horizon	contemp.	3m	6m	9m	12m
TIV^2	0.276	0.192	0.088	0.025	0.006
	(2.98)	(2.43)	(1.53)	(0.71)	(0.27)
Adj. R^2	28.74%	13.99%	2.93%	0.24%	0.02%
slope TIV	-0.791	-0.600	-0.356	-0.285	-0.220
	(-4.14)	(-4.12)	(-3.69)	(-2.21)	(-1.47)
Adj. R^2	35.37%	20.14%	6.95%	4.43%	2.55%
TIV^2	0.160	0.098	0.022	-0.048	-0.055
	(2.03)	(1.42)	(0.41)	(-0.98)	(-1.06)
slope TIV	-0.576	-0.466	-0.325	-0.354	-0.299
	(-3.89)	(-3.22)	(-2.47)	(-1.98)	(-1.65)
Adj. R^2	42.18%	22.44%	6.65%	4.60%	2.86%

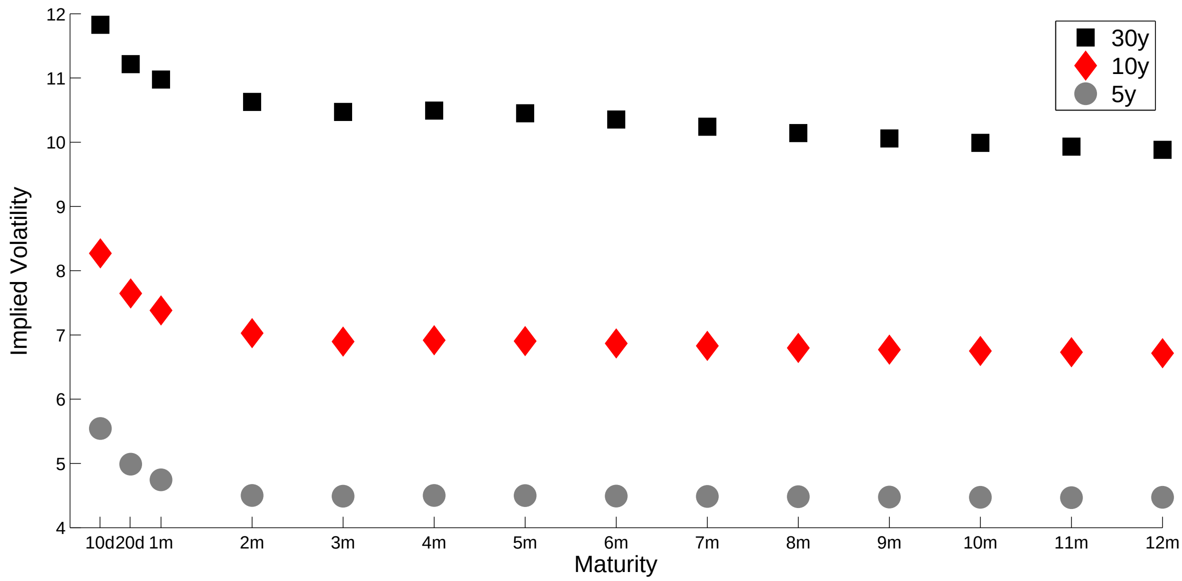

Figure 1. Term Structure of Implied Volatilities

The figure plots average implied volatilities for the three different tenors (30y, 10y, and 5y) and different maturities ranging from ten days to one year. Volatilities are the square root of the implied variances extracted from daily option prices using equation (6). All numbers are annualized and expressed in percent. Data is monthly and runs from June 1990 to May 2012.

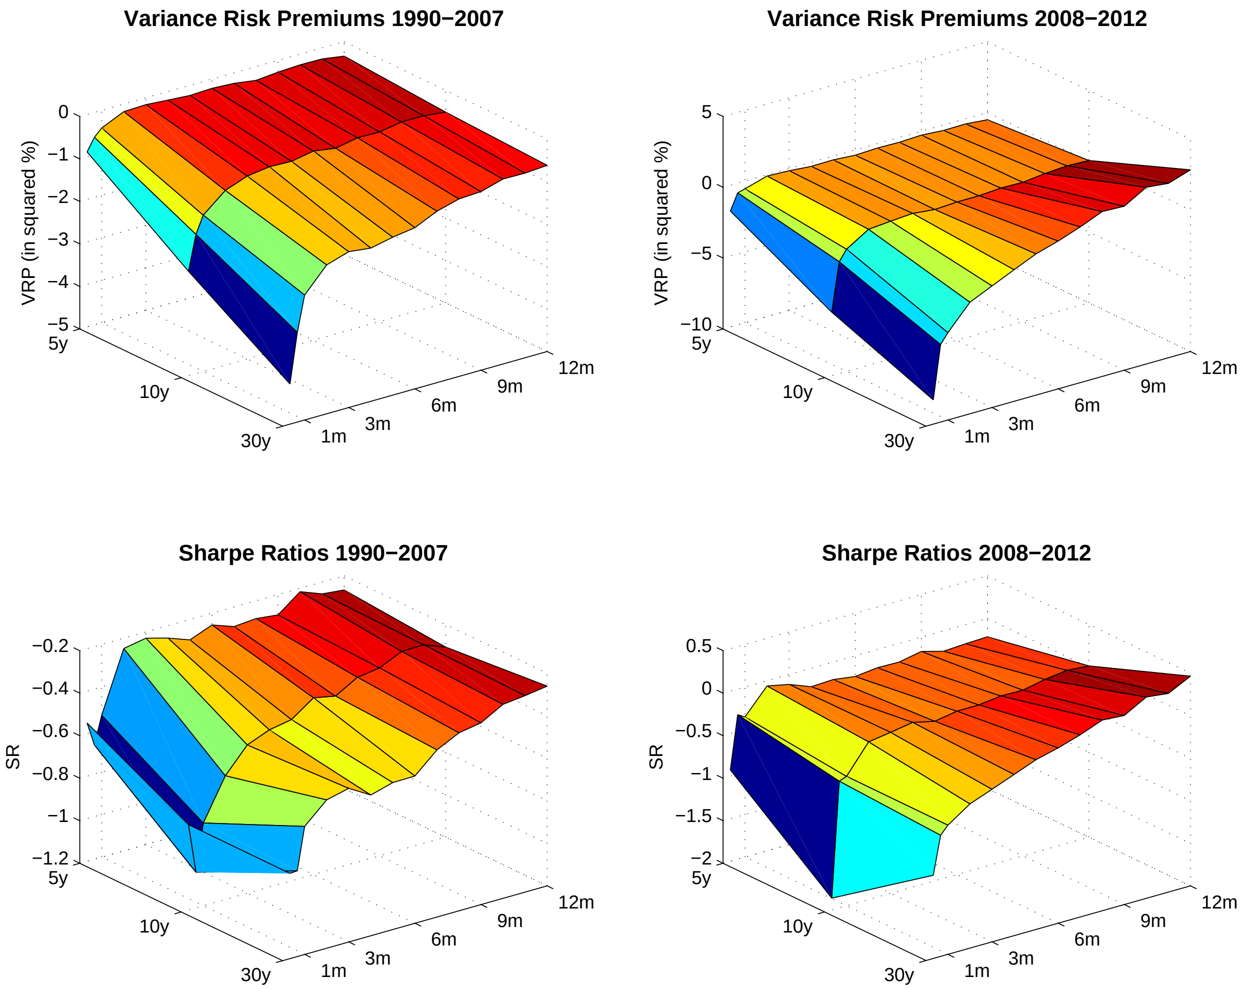

Figure 2. Term Structure of Variance Risk Premiums and Sharpe Ratios

The figure plots average variance risk premiums (upper panels) and average Sharpe ratios (lower panels) for the three different tenors (30y, 10y, and 5y) and different maturities ranging from ten days to one year. The left side panels display the results for the pre-crisis period from June 1990 until December 2007 while the right side panels show the corresponding numbers for the crisis and post-crisis period from January 2008 to May 2012. Variance risk premiums are computed by subtracting the implied variance as in equation (6) from the ex-post realized variance as in equation (5), and then expressed in squared percent after scaling them to monthly measures. Sharpe ratios are calculated as the average variance risk premiums divided by the corresponding standard deviation of the variance risk premiums. Data is monthly and runs from June 1990 to May 2012.

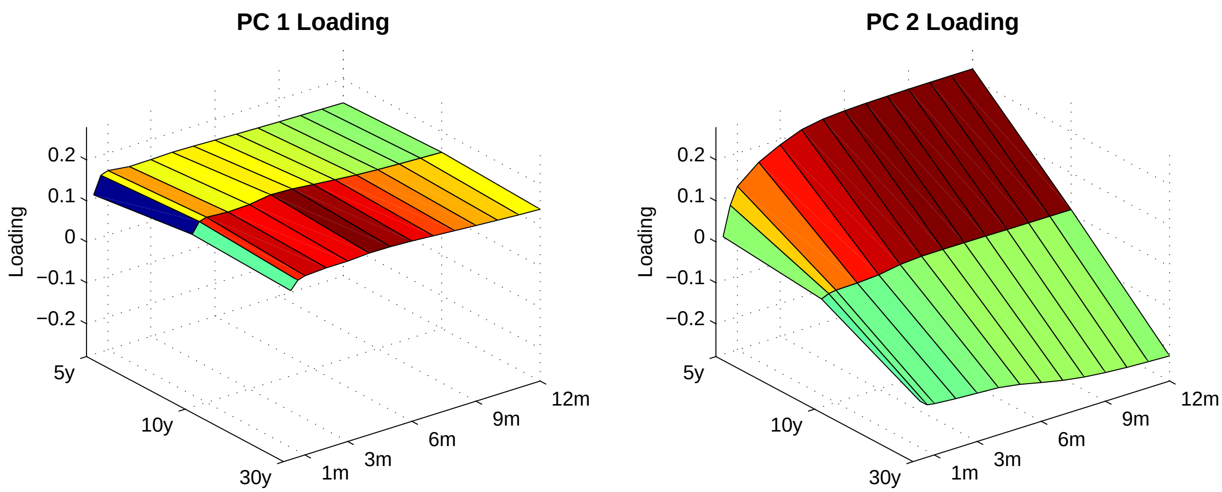

Figure 3. Factor Loadings on (Standardized) Implied Variances

The figure plots the factor loading of the first (left panel) and second (right panel) principal component of implied variances as in equation (6) across the three tenors (5y, 10y and 30y) and for maturities ranging between ten days and one year. The loadings are constructed from the eigenvectors corresponding to the two largest eigenvalues of the correlation matrix of monthly changes in implied variances. Data is monthly and runs from June 1990 to May 2012.

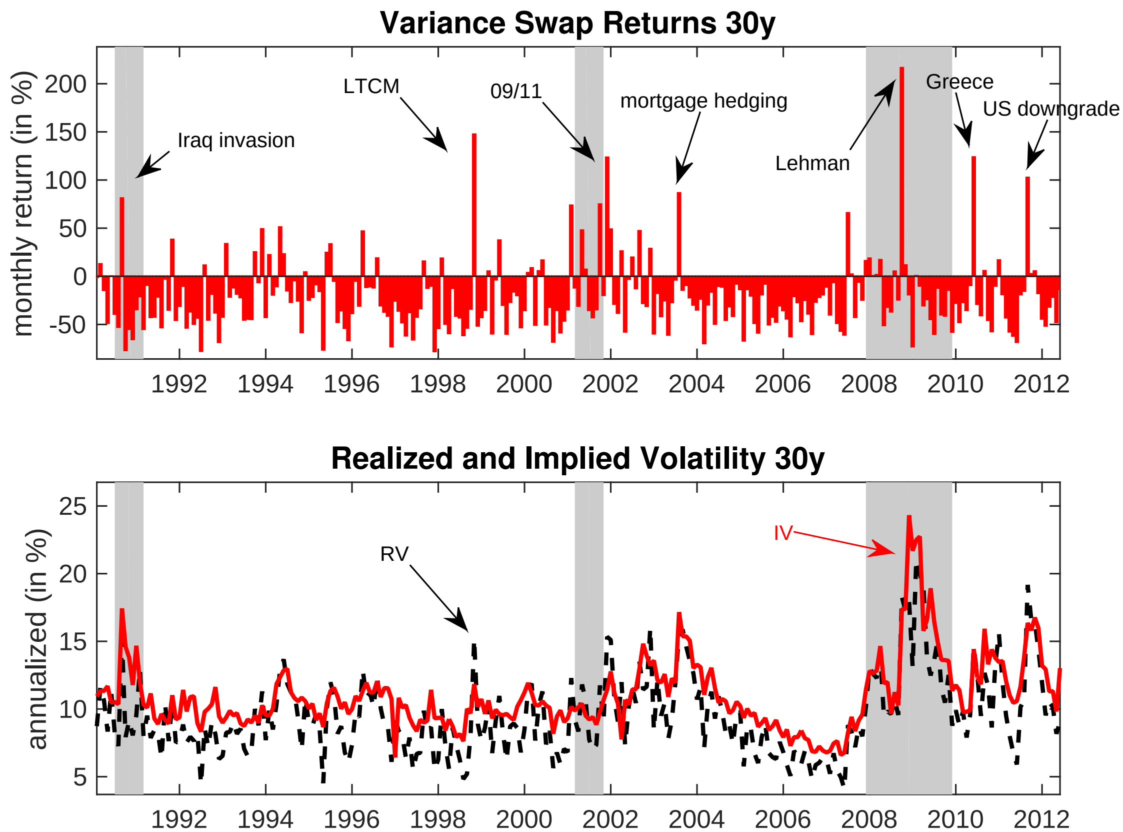

Figure 4. Generalized Treasury Variance Swap Returns 30y Treasury

The upper panel plots the monthly returns of the generalized Treasury variance swap (GTVS) with a 30y tenor and one month to maturity. The return is computed as the payoff of the one-month variance swap (implied variance minus ex-post realized variance) scaled by the fair strike of the variance swap (implied variance). The lower panel plots (annualized) realized and implied volatilities for the 30y Treasury futures. Gray bars indicate NBER recessions. Data is monthly and runs from January 1990 to May 2012.

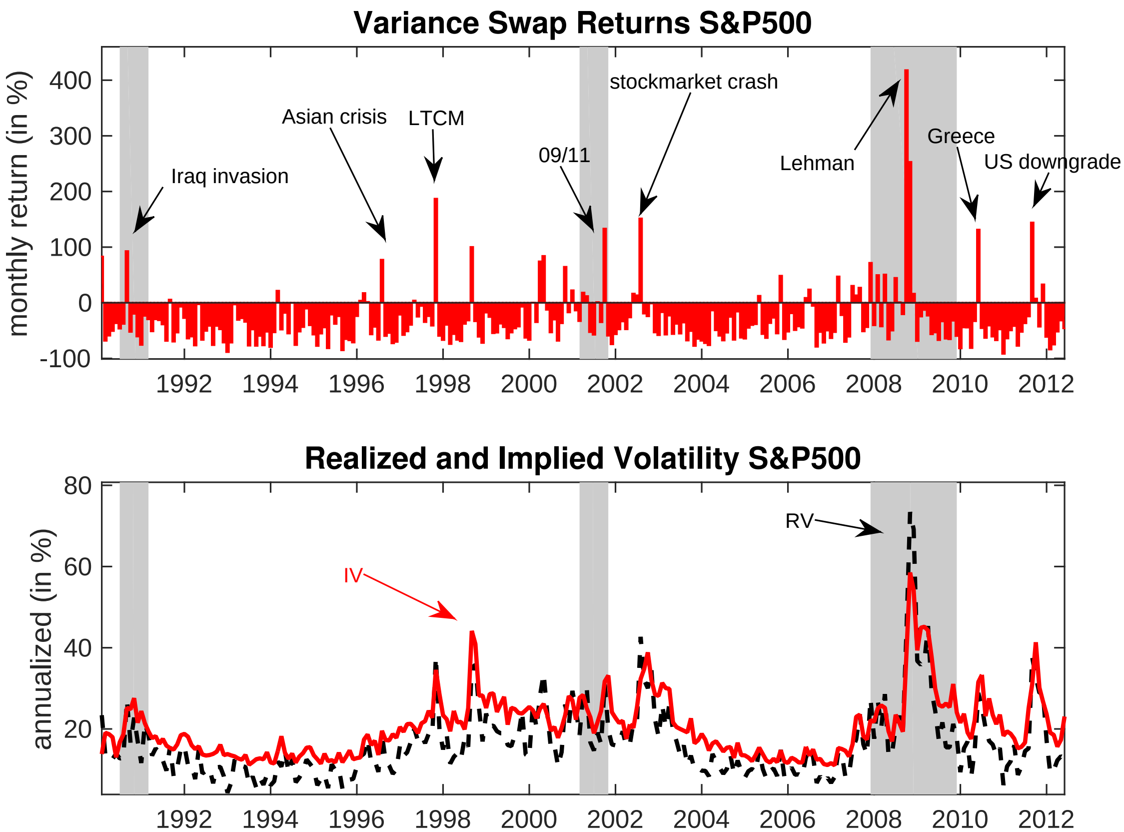

Figure 5. Variance Swap Returns S&P500

The upper panel plots the monthly returns of the variance swap on the S&P500 with a maturity of one month. The return is computed as the payoff of the one-month variance swap (implied variance minus ex-post realized variance) scaled by the fair strike of the variance swap (implied variance). The lower panel plots (annualized) realized and implied volatilities for the S&P500. Gray bars indicate NBER recessions. Data is monthly and runs from January 1990 to May 2012.

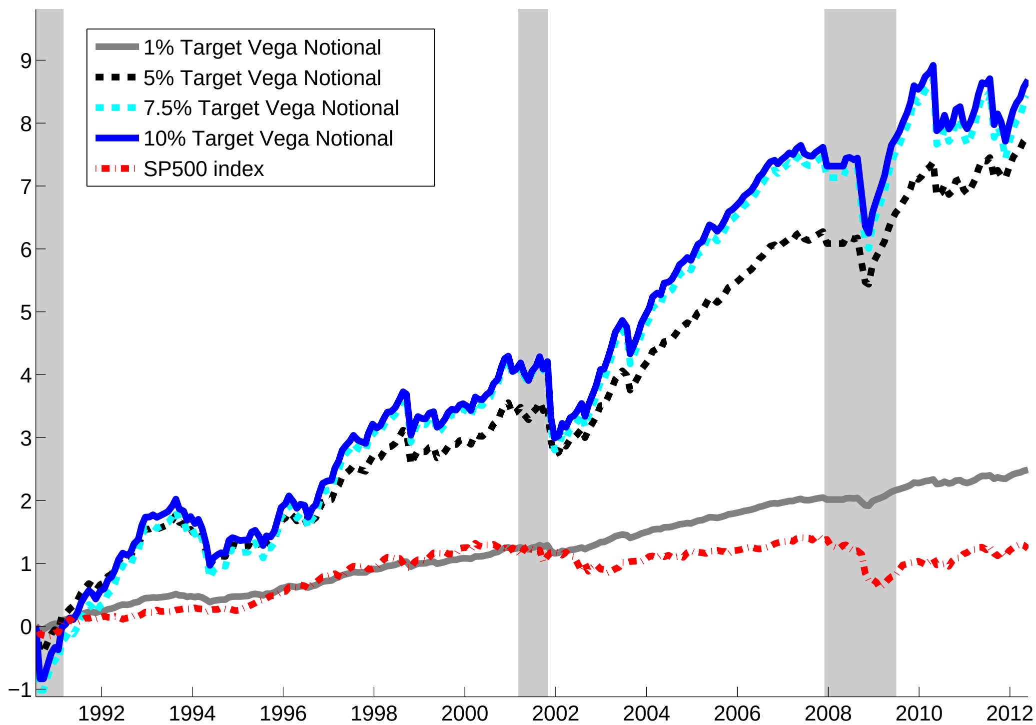

Figure 6. Margined Variance Swap Returns for 30y Treasury

This figure plots the wealth evolution in log scale for portfolios with an initial value of one USD. Each month the investor takes a short position in a 30y generalized Treasury variance swaps with a given target vega notional (ranging from 1% to 10%) and a maturity of one month. The required margin is nine times the vega notional. If losses within the month exceed the margin, the investor has to rescale the position. For comparison purposes, we also plot the evolution (also in log scale) of an investment in the S&P500 market index. Gray bars indicate NBER recessions. Data is monthly and runs from June 1990 to May 2012.

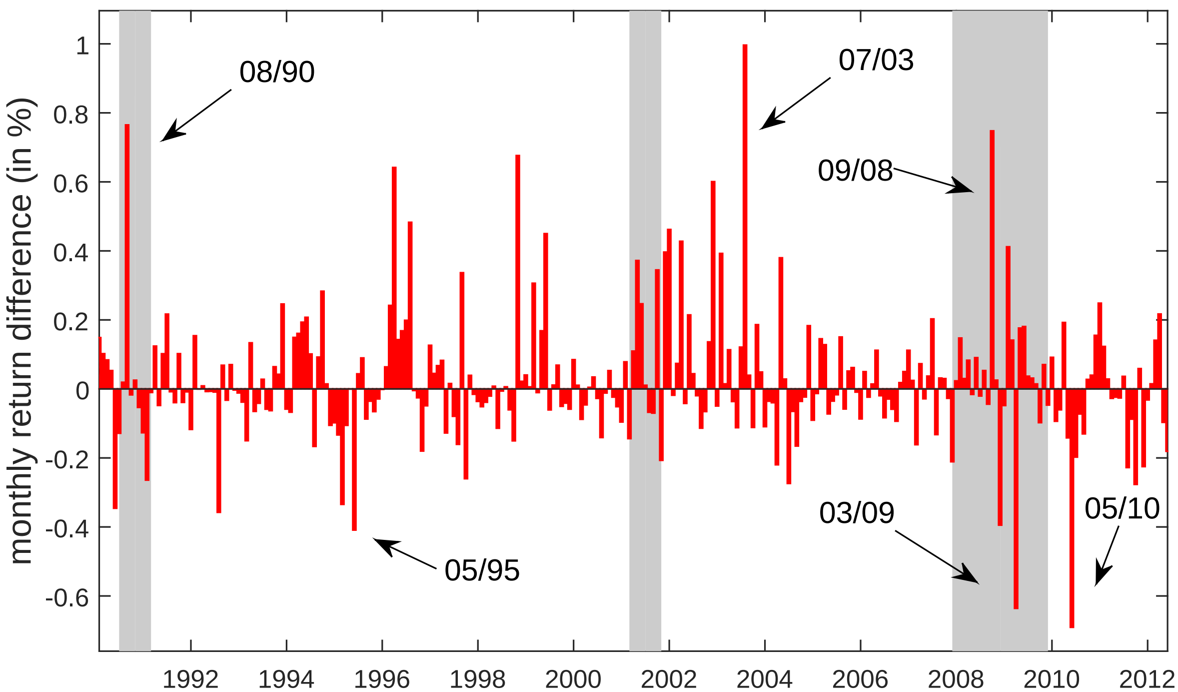

Figure 7. Difference in 30y Variance Swap Returns between $RV^{\log}$ and $\widetilde{RV}$

This figure plots the differences in monthly returns on 30y variance swap with a maturity of one month, but from two different measures of realized variance, log squared returns ($RV_{t,T}^{\log}$) and realized variance $\widetilde{RV}_{t,T}$ defined in equation (5). Each return is computed as the payoff of the one-month variance swap (implied variance minus $RV_{t,T}^{\log}$ or $\widetilde{RV}_{t,T}$ respectively) scaled by the fair strike of the variance swap (implied variance). Gray bars indicate NBER recessions. Data is monthly and runs from January 1990 to May 2012.

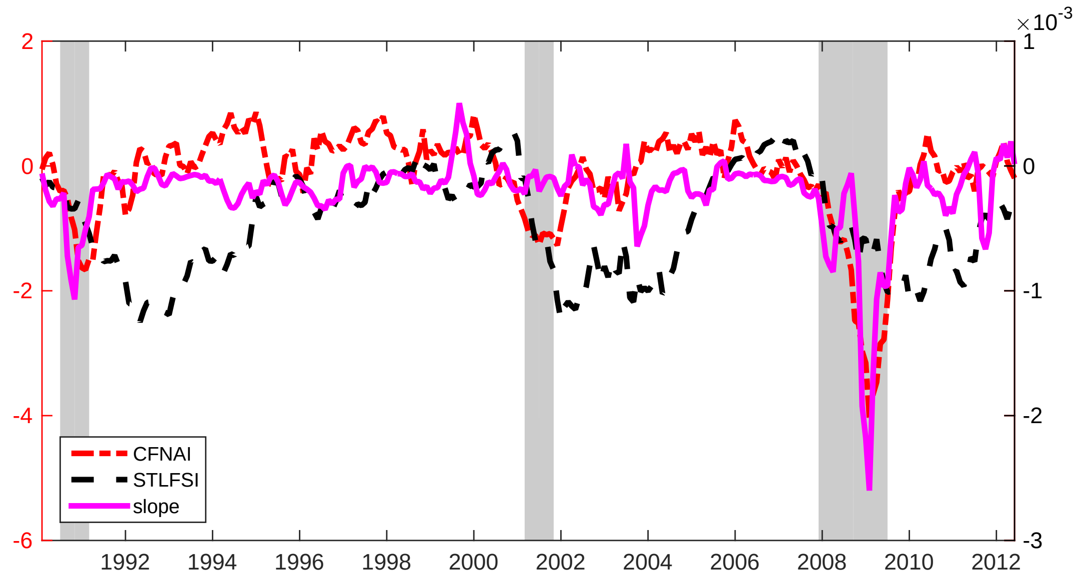

Figure 8. Slope of 30y Implied Variance, Chicago Fed National Activity Index (CFNAI), & St. Louis Fed Stress Index (SLFSI)

This figure plots the slope of the implied variance term structure which is defined as the difference between the one-year and one-month implied variance together with the Chicago Fed National Activity Index and the (negative of the) St. Louis Fed Stress Index. The variables are de-meaned and standardized. Gray bars indicate NBER recessions. Data is monthly and runs from January 1990 to May 2012.