

The price of variance risk

Ian Dew-Becker, Stefano Giglio, Anh Le, and Marius Rodriguez*

November 12, 2014

Abstract

The average investor in the variance swap market is indifferent to news about future variance at horizons ranging from 1 month to 14 years. It is only purely transitory and unexpected realized variance that is priced. These results present a challenge to most structural models of the variance risk premium, such as the intertemporal CAPM, recent models with Epstein–Zin preferences and long-run risks, and models where institutional investors have value-at-risk constraints. The results also have strong implications for macro models where volatility affects investment decisions, suggesting that investors are not willing to pay to hedge shocks in expected economic uncertainty.

1 Introduction

The recent explosion of research on the effects of volatility in macroeconomics and finance shows that economists care about macroeconomic volatility. Investors, on the other hand, do not. We show in this paper that it is costless on average to hedge news about future volatility in aggregate stock returns; in other words investors are not willing to pay for insurance against volatility news. In recent macroeconomic models, uncertainty about the future, or expectations of high future volatility, can induce large fluctuations in the economy. But if increases in economic uncertainty can drive the economy into a recession, as in, e.g., Bloom (2009) and Gourio (2012, 2014), we would expect that investors would want to hedge

*Northwestern Kellogg, Chicago Booth, UNC Kenan–Flagler, and the Federal Reserve Board. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System. We appreciate helpful comments from and conversations with Dante Amengual, Jean-Francois Bacmann, Brandon Bates, Nick Bloom, John Campbell and Cam Harvey, and seminar participants at the Booth Econometrics Workshop, University of New South Wales, University of Sydney, University of Technology, Sydney, University of Adelaide, Northwestern Kellogg, and the 2014 Macro Finance Society Meeting. We are particularly grateful to Gerald Lim and Markit Group for providing variance swap data.

those shocks. The pricing of volatility shocks thus presents a challenge to the recent macro literature on the effects of volatility shocks.

As a concrete example, consider the legislative battles over the borrowing limit of the United States in the summers of 2010 and 2011. Those periods were associated with increases in both financial measures of uncertainty, e.g. the VIX, and also the measure of policy uncertainty from Bloom, Baker, and Davis (2014). Between June and July, 2011, the 1-month variance swap rate rose from 16.26 to 25.96 percent (in terms of annualized volatility). However, those shocks also had small effects on *realized* volatility in financial markets: annualized realized volatility in June and July, 2011, was 14.59 and 15.23 percent, respectively. The debt ceiling debate caused uncertainty about the future to be high, but did not correspond to high contemporaneous volatility.¹

Those facts make the debt-ceiling shocks the exact type of shock that is studied in the recent literature. It is precisely changes in expectations of *future* uncertainty that can have strong macroeconomic effects, because they affect all forward-looking decisions. In this paper, we directly measure how much people are willing to pay to hedge shocks to expectations of future volatility. We find that those news shocks are unpriced: any investor can buy insurance against volatility shocks for free, and therefore any investor could have freely hedged the type of uncertainty shock that materialized during the debt ceiling debate.

We measure the price of variance risk using novel data on a wide range of volatility-linked assets both in the US and around the world, focusing primarily on variance swaps with maturities between 1 month and 14 years. Variance swaps are assets that pay to their owner the sum of daily squared stock market returns from their inception to maturity. They thus give direct exposure to future stock market volatility at different horizons and seem to be the most natural and direct hedge for the risks associated with increases in aggregate economic uncertainty.

The analysis of the pricing of variance swaps yields two simple results. First, news about future volatility is unpriced – exposure to volatility news does not earn a risk premium. Second, though, exposure to *realized variance* is strongly priced, with an annualized Sharpe

¹The table below reports realized volatility and the 1-month variance swap rate (nearly identical to the VIX) for June to October of 2011:

	June	July	Aug.	Sept.	Oct.
1-month variance swap	16.26	25.96	31.68	42.32	28.53
Realized volatility	14.59	15.23	47.18	28.80	29.80

Except for August, when both the variance swap rate and realized volatility rose in tandem, for all other months the changes in the two series are essentially unrelated.

Some of the volatility in the fall of 2011 was also due to uncertainty about the state of the European economy.

ratio of -1.7 – five times larger than the Sharpe ratio on equities.² And Bollerslev and Todorov (2011) show that realized variance is priced due to its correlation with large negative jumps. In other words, investors are willing to pay a large amount of money for protection from an extreme negative shock to the economy that mechanically generates a spike in realized volatility. The evidence argues that variance swaps, options, and VIX forwards are used to hedge the realization of a disaster, but not news that uncertainty or the probability of a disaster has changed.

The results present a challenge to a wide range of models. In macroeconomics, there is now a large literature following Bloom (2009), who also studies the variance of the S&P 500, arguing that shocks to uncertainty can have important effects on the aggregate economy. If increases in future uncertainty have sufficiently important effects on the economy that they affect investor utility, though, we would expect them to carry a risk premium. The fact that they do not implies that volatility shocks may not be a major driver of welfare.

From a finance perspective, Merton’s (1973) intertemporal capital asset pricing model says that assets that have high returns in periods with good news about future investment opportunities are viewed as hedges and thus earn low average returns. Since expected future volatility is a natural state variable for the investment opportunity set, the covariance of an asset with shocks to future volatility should affect its expected return, but it does not.³

Consumption-based models with Epstein–Zin (1991) preferences have similar predictions. Under Epstein–Zin preferences, marginal utility depends on lifetime utility, so that assets that covary positively with innovations to lifetime utility earn high average returns because they have low payouts in bad states of the world.⁴ If high expected volatility is bad for lifetime utility (either because volatility affects the path of consumption or because volatility reduces utility simply due to risk aversion), then volatility news should be priced.⁵

²The Sharpe ratio is equal to the average excess return divided by the standard deviation. It thus gives a scale-free measure of the risk premium on an asset (though it obviously does not account for higher moments such as skewness and kurtosis).

³Recently, Campbell et al. (2014) and Bansal et al. (2013) estimate an ICAPM model with stochastic volatility and find that shocks to expected volatility (and especially long-run volatility) are priced in the cross-section of returns of equities and other asset classes. Although the focus on their paper is not the variance swap market, Campbell et al. (2014) test their specification of the ICAPM model also on straddle returns and synthetic volatility claims, and find that the model manages to explain only part of the returns on these securities. This suggests that the model is missing some high-frequency features of the volatility market.

⁴This is true in the most common calibrations with a preference for early resolution of uncertainty. When investors prefer a late resolution of uncertainty the risk prices are reversed.

⁵Also see Branger and Völkert (2010) and Zhou and Zhu (2012) for discussions. Barras and Malkhozov (2014) study the determinants of changes in the variance risk premium over time and find that an important component of the movements is explained by changes in the risk-bearing capacity of financial intermediaries, suggesting that a type of liquidity is relevant in these markets.

As a specific parameterized example, we study variance swap prices in Drechsler and Yaron’s (2011) calibrated long-run risk model with Epstein–Zin preferences. Drechsler and Yaron (2011) is a key benchmark because it is a quantitative model that can match a wide range of features of the dynamics of consumption growth, stock returns, and specifically volatility. While the model represents a major innovation in being able to both generate a large variance risk premium (the risk premium of short-term realized volatility shocks) and match results about the predictability of market returns, we find that its implications for the term structure of variance swap prices and returns are strongly at odds with the data: as one would expect, it predicts that shocks to future expected volatility should be strongly priced, counter to what we observe empirically.

We obtain similar results in Wachter’s (2013) model of time-varying disaster risk with Epstein–Zin preferences: the combination of predictability in the long-run probability of disaster and Epstein–Zin preferences results in a counterfactually high price for insurance against shocks to expected future volatility relative to current volatility. In both Wachter (2013) and Drechsler and Yaron (2011), Sharpe ratios (compensation per unit of risk exposure) earned by claims on future variance from 3 months to 14 years ahead are similar to those earned by claims to realized variance over the next month, whereas in the data the Sharpe ratios are all near zero (or positive) for claims to variance more than two months in the future. This fact means that the models can either generate high or low Sharpe ratios for claims to volatility at all horizons. They are unable to match our stylized fact that the Sharpe ratios are extremely large at the one-month horizon (hedging realized volatility) and then rapidly fall to zero for higher maturities.

Similar problems with matching term structures of Sharpe ratios in structural models have been studied in the context of claims to aggregate market dividends by Kojien, Brandt and van Binsbergen (2012). Our results thus support and complement theirs in a novel context.⁶

More positively, we show that Gabaix’s (2012) model of rare disasters can match the stylized fact that Sharpe ratios on variance claims fall to zero rapidly with maturity. In his model, the probability of a disaster is constant, but the exposure of the stock market – i.e. its expected decline if a disaster occurs – varies over time. The realization of a disaster is inevitably a state with high realized volatility (since if returns are highly negative, squared daily returns will mechanically be high), so variance swaps provide a direct hedge against the occurrence of a disaster, meaning they earn a large negative Sharpe ratio. But since changes in the exposure of the stock market to consumption disasters, which drive expected

⁶Our paper also relates to a large literature that looks at derivative markets to learn about general equilibrium asset pricing models, for example Backus, Chernov and Martin (2011) and Martin (2013, 2014).

future return variance, are uncorrelated with the *current level* of consumption, they are not priced shocks. The model is thus able to simultaneously generate a large negative premium on realized stock return variance and zero premium on news about future variance, just like in the data. Of course, that result comes at the cost of completely divorcing news about future volatility from the current level of consumption – that is, changes in expectations about future volatility must have no real effects in order for the model to fit the data.⁷

There is a large extant literature studying the pricing of volatility in financial markets.⁸ Most closely related to us is a small number of recent papers with data on variance swaps with maturities from two to 24 months, including Egloff, Leippold, and Wu (2010) and Aït-Sahalia, Karaman, and Mancini (2014), who study no-arbitrage term structure models. Aït-Sahalia, Karaman, and Mancini (2014) are notable in particular for analyzing a model with both jumps and diffusive shocks, and for linking the dynamics of variance swap prices with prices of other assets. They represent the current frontier of the literature that studies no-arbitrage models of the term structure of variance claims.

The pricing models we estimate are far less technically sophisticated than that of Aït-Sahalia, Karaman, and Mancini (2014), but we build on their work in three ways. First, we examine a vast and novel range of data sources. For S&P 500 variance swaps, our panel includes data at both shorter and longer maturities than in previous studies – from one month to 14 years. The one-month maturity is important for giving a claim to shorter-term realized variance, which is what we find is actually priced. Having data at very long horizons is important for testing models, like Epstein–Zin preferences, in which expectations at very long horizons are the main drivers of asset prices. In addition, to the best of our knowledge, we are the first to examine the term structure of variance swaps for major international indexes, as well as for the term structure of the VIX obtained from options on those indexes. We are thus able to confirm that our results hold across a wide range of markets, maturities, and time periods.

Our second contribution to the previous term structure literature is that rather than

⁷An alternative possibility is that the variance market is segmented from other markets, as in, e.g., Gabaix, Krishnamurthy, and Vigneron (2007). In that case, the pricing of risks might not be integrated between the variance market and other markets. We show, however, that our results hold not only with variance swaps, but also in VIX futures and in the options market, which is large, liquid, and integrated with equity markets, making it less likely that our results are idiosyncratic.

⁸A number of papers study the pricing of volatility in options markets, e.g. Jackwerth and Rubinstein (1996), Coval and Shumway (2001), Bakshi and Kapadia (2003), Broadie, Chernov and Johannes (2009), Christofferson, Jacobs, Ornathanalai, and Wang (2008), and Kelly, Pastor, and Veronesi (2014). Lu and Zhu (2010) and Mencia and Sentana (2013) study VIX futures markets, while Bakshi, Panayotov, and Skoulakis (2011) show how to construct forward claims on variance with portfolios of options. In the Treasury bond market, Cieslak and Povala (2014) find, similar to us, that short-run volatility is more strongly priced than long-run volatility. See also Amengual and Xiu (2014) for an important recent study of jumps in volatility.

working exclusively within the context of a particular no-arbitrage pricing model for the term structure of variance claims, we derive from the data more general and model-independent facts about pricing in this market. Our pricing results can be directly compared against the pricing implications of different structural economic models, which would be more difficult if the pricing results were only derived within a specific no-arbitrage model. Our key result, that purely transitory realized variance is priced while innovations to expectations are not, can be obtained from a simple reduced-form analysis and in data both for the United States and other countries. Nevertheless, we also confirm our results in a more formal no-arbitrage setting, whose main advantage is to yield much more precise estimates of risk prices.

Our third contribution is that, starting from our model-free empirical results, we then focus on what variance swaps can tell us about different structural economic models. This theoretical analysis leads us to the conclusion that the empirical facts in the variance swap market are most consistent with a model in which variance swaps are used to hedge the realization of market crashes and in which variation in expected future stock market volatility is not priced by investors, counter to the predictions of most standard asset pricing and macroeconomic models.

The remainder of the paper is organized as follows. Section 2 describes the novel datasets we obtain for variance swap prices. Section 3 reports unconditional means for variance swap prices and returns, which demonstrate our results in their simplest form. Section 4 analyzes the cross-sectional and time-series behavior of variance swap prices and returns more formally in standard asset pricing frameworks. In section 5, we discuss what structural general-equilibrium models can fit the data. We calibrate three leading models from the literature, comparing them to our data, showing that only one matches the key stylized facts. Section 6 concludes.

2 The data

Variance swaps are contracts in which one party pays a fixed amount at maturity, which we refer to as price of the variance swap, in exchange for a payment equal to the sum of squared daily log returns of the underlying occurring until maturity. In this paper, the underlying is the S&P 500 index unless otherwise specified. The payment at expiration of a variance swap initiated at time τ and with maturity m is

$$Payoff_{\tau}^m = \sum_{j=\tau+1}^{\tau+m} r_j^2 - VS_{\tau}^m \quad (1)$$

where time here is indicated in days, r_j is the log return on the underlying on date j , and VS_t^m is the price on date τ of an m -day variance swap.

Our main analysis focuses on two proprietary datasets of quoted prices for S&P 500 variance swaps for different periods and different maturities.⁹ Dataset 1 contains monthly variance swap prices for contracts expiring in 1, 2, 3, 6, 12, and 24 months, and includes data from December, 1995, to October, 2013. Dataset 2 contains data on variance swaps with expirations that are fixed in calendar time, instead of fixed maturities. Common maturities are clustered around 1, 3, and 6 months, and 1, 2, 3, 5, 10, and 14 years. Dataset 2 contains prices of contracts with up to 5 years of maturity starting in September, 2006, and up to 14 years starting in August, 2007, and runs up to February, 2014. We apply spline interpolation to each dataset to obtain the prices of variance swaps with standardized maturities covering all months between 1 and 12 months for Dataset 1 and between 1 and 120 months for Dataset 2 (though in the no-arbitrage model below we use the original price data without interpolation).¹⁰

To confirm the accuracy of the quoted variance swap prices, we compare them to prices reported by the Depository Trust & Clearing Corporation (DTCC), which has collected data on trades of variance swaps since 2013.¹¹

In addition to the prices of S&P 500 variance swaps, we also obtained prices for variance swaps in 2013 and 2014 for the FTSE 100 (UK), Euro Stoxx 50 (Europe), and DAX (Germany) indexes. And last, we obtained prices for options on those international indexes and on the S&P 500 and the CAC 40 (France) for various samples starting as early as 1996 that we use to construct a term structure of synthetic volatility swap prices (analogous to the VIX). Our main focus is the S&P 500 variance swap data, but we confirm that our key results hold on this wide range of other data sources. Finally, we also obtained data from Bloomberg on the prices of VIX futures with maturities of up to 6 months and confirm that the variance swaps and VIX futures have nearly identical prices.

From a macroeconomic perspective, our sample is somewhat short as it includes only two recessions. From a financial perspective, though, it gives a view of volatility behavior in a wide range of market states. In particular, there are two large declines in the market in

⁹Both datasets were obtained from industry sources. Dataset 2 is obtained from Markit Totem, and consists in an average of quotes obtained from dealers in the variance swap OTC market. Since the prices we observe are a composite of quotes from many different dealers (on average 11), the quality of this dataset is very high, and comparable to that of the widely used CDS dataset from Markit.

¹⁰For the times and maturities for which we have both datasets, the prices are effectively identical: the correlations between the two datasets are never below 0.996. We will also show below that the prices are well explained by only two principal components, suggesting that interpolation should accurately recover prices.

¹¹DTCC was the only swap data repository registered under the Dodd–Frank act to collect data on variance swaps in 2013. The Dodd–Frank act requires that all swaps be reported to a registered data repository.

our sample, including the biggest crash of the post-war period. It is thus unlikely that the average returns on variance swaps that we observed are downward biased – if anything, the number of market crashes in our 18-year sample is higher than the historical average for a sample of that length.

Both S&P 500 datasets are novel to the literature. Variance swap data with maturities up to 24 months as in Dataset 1 has been used before (Egloff, Leippold, and Wu, 2010, Ait-Sahalia, Karaman, and Mancini, 2014, and Amengual and Xiu, 2014), but the shortest maturity previous studies observed was two months, not one month. We show that the one-month variance swap is special in this market because it is the exclusive claim to next month’s realized variance, which is the *only* risk priced in this market. Observing the one-month variance swap price is critical for precisely measuring the price of realized-variance risk.

Furthermore, to the best of our knowledge, this is also the first paper to observe and use variance swap data with maturity longer than two years. While our time series for the longer maturities is shorter, starting in 2006, it does cover the period before and during the financial crisis. Since Epstein–Zin preferences imply that it is the very low-frequency components of volatility that should be priced (Branger and Volkert, 2010; Dew-Becker and Giglio, 2014), having claims with very long maturities is important for effectively testing the central predictions of Epstein–Zin preferences.

2.1 Characteristics of the market for volatility claims

There are a number of ways to trade claims on future volatility. In this paper we focus on variance swaps because their payoffs depend only on realized variance itself. However, there are three other related securities that are also exposed to volatility: VIX futures, equity index options, and exchange-traded funds and notes indexed to VIX futures. We discuss each in turn.

The variance swap market is large: the notional value of outstanding variance swaps at the end of 2013 was \$4 billion of notional vega.¹² This means that an increase in the annualized realized volatility by one point would induce total payments of \$4bn. This market is thus small relative to the aggregate stock market, but it is non-trivial economically. Current bid/ask spreads in the variance swap market average 1 to 3 percent, depending on maturity and trade size. Trading costs are thus much larger than for equities, but not large enough

¹²See the Commodity Futures Trading Commission’s (CFTC) weekly swap report. The values reported by the CFTC are consistent with data obtained from the Depository Trust & Clearing Corporation that we discuss below.

to have a major impact on our findings.

Table 1 shows the total volume (in notional vega) for all transactions reported to DTCC under the Dodd–Frank act between March 2013 and June 2014. The table shows that in little more than a year, the variance swap market saw \$7.2 billion of notional vega traded. Only 11 percent of the volume was traded in short maturity contracts (1-3 months); the bulk of the transactions occurred for maturities between 6 months and 5 years, and the median maturity was 12 months.

To check the accuracy of the quoted prices that we obtained, we compare them to those reported for actual trades by DTCC in the same period. Appendix Figure A.1 shows the distribution of the percentage difference between our quotes and transaction prices for different maturity baskets. Quotes and transaction prices are in most cases very close, with the median percentage difference across all maturities approximately 1 percent.

Futures have been traded on the VIX since 2004. The VIX futures market is significantly smaller than the variance swap market, with current outstanding notional vega of roughly \$500 million.¹³ Bid/ask spreads are smaller than what we observe in the variance swap market, at roughly 0.1 percent, but as the market is smaller, we would expect relatively more price impact (and market participants claim that there is). While the VIX futures market is smaller and potentially less liquid than the variance swap market, it has the advantage of being exchange traded, making it simpler to exit open positions and eliminating counterparty risk.

More recently, a market has developed in exchange-traded notes and funds available to retail investors that are linked to VIX futures prices. These funds currently have an aggregate notional exposure to the VIX of roughly \$5 billion, making them comparable in size to the value reported by the CFTC for the variance swap market.

The final major way to trade volatility is through options. Most well known is the straddle, which has a payoff that depends on the absolute value of the change in the underlying over the life of the option.¹⁴ Options are traded in numerous venues, have notional values outstanding of trillions of dollars, and have been thoroughly studied. Even in 1990, Vihj (1990) noted that the CBOE was highly liquid and displayed little evidence of price impact for large trades. In order to confirm our main results, we show below that returns obtained from investments in options are similar to those obtained from variance swaps.¹⁵

¹³According to the CBOE futures exchange market statistics. See: <http://cfe.cboe.com/Data/HistoricalData.aspx>

¹⁴A straddle is a portfolio consisting of a long position in an at-the-money put option and an at-the-money call option.

¹⁵Recently, Boguth et al. (2012a, b) argue that returns measured on options portfolios can be substantially biased by noise, one potential source of which is the bid/ask spread. The majority of our results pertain

Since variance swaps are traded over the counter, it is possible that counterparty risk could influence their prices (though this concern would not apply to our results on options and VIX futures). Given that variance swaps are standardized contracts covered by the ISDA Master Agreement, the margining of those contracts follows standard procedures for derivatives: an initial margin is posted by both parties, and variational margin is exchanged regularly depending on the value of the position. The residual counterparty risk in these contracts depends on the possibility of jumps in the value of the contracts between exchanges of collateral, and is therefore only a material issue when returns have high skewness and kurtosis at short horizons. As we will discuss later, only short-term variance swap contracts have payoffs that are far from Gaussian, and are therefore exposed to counterparty risk, and we argue below that for these contracts counterparty risk would push *against* the results we observe. We conclude that if counterparty risk was indeed priced by market participants, accounting for it would in fact make our results stronger.

3 The term structure of variance claims

In this section we study average prices and returns of variance swaps. The key result that emerges is that only very short-duration variance claims earn a risk premium.

3.1 Variance Swap Prices

The shortest maturity variance swap we consistently observe has a maturity of one month, so we treat a month as the fundamental period of observation. We define RV_t to be realized variance ($\sum r_j^2$) during month t .

Given a risk-neutral measure Q , the price of an n -month variance swap at the end of month t is

$$VS_t^n = E_t^Q \left[\sum_{j=1}^n RV_{t+j} \right] \quad (2)$$

where VS_t^n is the price of an n -month variance swap at the end of month t and the subscript from here forward always indexes months, rather than days.¹⁶ E_t^Q denotes the mathematical expectation under the risk-neutral measure conditional on information available at the end of month t .¹⁷

directly to prices of volatility claims, as opposed to their returns, meaning that the issues noted by Boguth et al. are unlikely to affect our analysis. Furthermore, when we analyze returns, the portfolios are not levered to the degree that Boguth et al. argue causes biases in results.

¹⁶Note again that the variance swap price is paid at the maturity rather than the inception of the contract.

¹⁷In the absence of arbitrage, there exists a probability measure Q such that the price of an asset with

Since an n -month variance swap is a claim to the sum of realized variance over months $t + 1$ to $t + n$, it is straightforward to compute prices of zero-coupon claims on realized variance. Specifically, we define an n -month zero-coupon variance claim as an asset with a payoff equal to realized volatility on date $t + n$. The absence of arbitrage implies

$$Z_t^n \equiv E_t^Q [RV_{t+n}] \quad (3)$$

$$= VS_t^n - VS_t^{n-1} \quad (4)$$

Z_t^n is the price of a claim to variance realized *only* during month $t + n$, so it represents the market's risk-neutral expectation of realized variance n months in the future. We use the natural convention that

$$Z_t^0 = RV_t \quad (5)$$

so that Z_t^0 is the variance realized *during* the current month t . Note that a one-month zero-coupon variance claim is exactly equivalent to a one-month variance swap, $Z_t^1 = VS_t^1$.

Figure 1 plots the time series of zero-coupon variance claim prices for maturities between one month and ten years. For readability, the figure shows all series in annualized percentage volatility units, rather than variance units: $100 \times \sqrt{12 \times Z_t^n}$ instead of Z_t^n . It also plots annualized realized volatility, $100 \times \sqrt{12 \times Z_t^0}$, in each panel. The top panel plots zero-coupon variance claim prices for maturities below one year, while the bottom panel focuses on maturities longer than one year.

The term structure of variance claim prices is usually weakly upward sloping. In times of distress, though, such as during the financial crisis of 2008, the short end of the curve spikes, temporarily inverting the term structure. That result is natural: volatility obviously was not going to continue at crisis levels, so markets priced variance swaps with the expectation that volatility would fall in the future. Interestingly, the bottom panel shows that long-term volatility forwards saw a sharp increase during the summer of 2010, indeed a larger movement than was recorded in 2008. This implies that the risk-neutral expectation of long-run volatility rose substantially in that period, which was when the US government was approaching the debt ceiling and some European governments were seen as close to default. For example, the 5-year price rose from 31 to 38 volatility points in three months.

Figure 2 reports the average term structure of zero-coupon variance claims Z_t^n for two

payoff X_{t+1} is $\frac{1}{R_{f,t+1}} E_t^Q [X_{t+1}]$, where $R_{f,t+1}$ is the risk-free interest rate. Under power utility, for example, we have $E_t^Q [X_{t+1}] = E^P \left[\frac{(C_{t+1}/C_t)^{-\rho}}{E_t^P [(C_{t+1}/C_t)^{-\rho}]} X_{t+1} \right]$, where ρ is the coefficient of relative risk aversion, C is consumption, and P is the physical probability measure. The price of a variance swap does not involve the interest rate because money only changes hands at the maturity of the contract.

different subperiods – the period 2008–2014, when we have data for longer maturities, is in the top panel, while the full sample, 1996–2013, is in the bottom panel. The term structure of zero-coupon variance claim prices is upward sloping on average. However, the figure also shows that it flattens out quickly as the maturity increases.

The average zero-coupon term structures in Figure 2 provide the first indication that the compensation for bearing risk associated with news about future volatility is small in this market. The return on holding a zero-coupon variance claim for a single month is $\frac{Z_{t+1}^{n-1} - Z_t^n}{Z_t^n}$. The average return is therefore closely related to the slope of the variance term structure. If the variance term structure is upward sloping then zero-coupon claims will have negative average returns, implying that it is costly to buy insurance against increases in future expected volatility. The fact that it is steep at short horizons and flat at long horizons is a simple way to see that it is only the claims to variance in the very near future that earn significant negative returns.

3.2 Returns on zero-coupon variance claims

We now study the monthly returns on zero-coupon variance claims. The return on an n -month claim corresponds to a strategy that buys the n -month claim and sells it one month later as an $(n - 1)$ -month claim, reinvesting then again in new n -month zero-coupon variance claims. We define the excess return of an n -period variance claim following Gorton, Hayashi, and Rouwenhorst (2013)¹⁸

$$R_{t+1}^n = \frac{Z_{t+1}^{n-1} - Z_t^n}{Z_t^n} \quad (6)$$

Given the definition that $Z_t^0 = RV_t$, the return on a one-month claim, R_{t+1}^1 is simply the percentage return on a one-month variance swap. We focus here on the returns for maturities of one to 12 months, for which we have data since 1995. All the results extend to higher maturities in the shorter sample.

Table 2 reports descriptive statistics for our panel of monthly returns. The table shows important differences between the very short end of the curve, especially the first return, for which the payoff is realized variance, and the other maturities. Only the average returns for the first and the second zero-coupon claims are negative, while all the others are zero or slightly positive. Similarly, the median return is strongly negative at the short end, and

¹⁸Note that $Z_{t+1}^{n-1} - Z_t^n$ is also an excess return on a portfolio since no money changes hands at the inception of a variance swap contract. Following Gorton, Hayashi, and Rouwenhorst (2013), we scale the return by the price of the variance claim bought. This is the natural scaling if the amount of risk scales proportionally with the price. We have reproduced all of our analysis using the unscaled excess return $Z_{t+1}^{n-1} - Z_t^n$ as well and confirmed that all the results hold in that case.

close to zero for maturities of three to 12 months.¹⁹ Return volatilities are also much higher at short maturities, though we note that the long end still displays significant variability. For example, returns on the 12-month zero-coupon claim have an annual standard deviation of 17 percent, as high as that of the aggregate stock market return. This indicates that markets' expectations of 12-month volatility fluctuates significantly over time.

Finally, we note that only very short-term returns have high skewness and kurtosis. A buyer of short-term variance swaps is therefore exposed to counterparty risk if realized variance spikes and the counterparty defaults. This should induce her to pay less for the insurance, i.e. we should expect the average return to be *less negative*. Therefore, the presence of counterparty risk on the short end of the term structure would bias our estimate towards not finding the large negative expected returns that we instead find. On the other hand, returns on longer-maturity zero-coupon claims have much lower skewness and kurtosis, which indicates that counterparty risk is substantially less relevant for longer maturities.

Given the different volatilities of the returns at different ends of the term structure, it is perhaps more informative to examine Sharpe ratios (average excess returns scaled by standard deviations). Sharpe ratios are useful because they measure compensation earned per unit of risk. Figure 3 shows the annual Sharpe ratios of the 12 zero-coupon claims. The Sharpe ratios are negative for the one-month zero-coupon claim (-1.4 annualized) and, to a lesser extent, the two-month zero-coupon claim (-0.5 annualized). All other Sharpe ratios are close to zero.

The results at the short end of the curve indicate that investors are willing to pay a large premium to hedge realized volatility.²⁰ What is new and surprising in this picture is the fact that agents are *not* willing to pay to hedge any innovations in *expected* volatility, even two or three months ahead. A claim to volatility at a horizon n beyond one month is purely exposed

¹⁹We note that there is a direct link between the average term structure of zero-coupon prices and the risk premia for shocks to volatility at different horizons. To see this, it is easier to work with unscaled returns, defined as:

$$\hat{R}_{t+1}^n = Z_{t+1}^{n-1} - Z_t^n$$

Taking unconditional expectations of this equation we obtain:

$$E[\hat{R}_{t+1}^n] = E[Z_{t+1}^{n-1}] - E[Z_t^n]$$

Since the term structure is stationary, $E[Z_{t+1}^{n-1}] = E[Z_t^{n-1}]$, and therefore:

$$E[\hat{R}_{t+1}^n] = E[Z_t^{n-1}] - E[Z_t^n]$$

The average slope of the term structure between maturities n and $n - 1$ corresponds to the risk premium on the n -maturity zero-coupon claim. A similar relation holds – after a loglinear approximation – for scaled returns as well: the average scaled returns at a certain maturity corresponds to the slope of the curve at that maturity relative to the level at the same maturity.

²⁰See Coval and Shumway (2001)

to news about future volatility: its return corresponds exactly to the change in expectations between periods t and $t + 1$ about volatility n months in the future. Pure news about future expected volatility will therefore affect its return, whereas purely transitory shocks to volatility that disappear before month n will not affect it at all. These zero-coupon variance claims, therefore, represent direct hedges against shocks to expectations of future volatility. Our results show that such shocks command a small to zero risk premium.²¹ The Sharpe ratios reinforce the result from Figure 2 that only the immediate realized volatility is priced in this market.

3.3 Alternative claims on variance

The results for variance swaps can also be confirmed in the options market. We exploit the well-known fact that if the S&P 500 follows a diffusion, a variance swap can be replicated by a portfolio of options with the same maturity.²² The term structure of synthetic zero-coupon variance swap prices constructed from options should then align well with the term structure of actual variance swap prices. The appendix reports details of the construction of the synthetic variance swap prices.

The top panel of Figure 4 shows the term structure of zero-coupon variance claims obtained from the variance swap data compared to the synthetic term structure for maturities up to 1 year. While the curve obtained using options data seems noisier, the two curves deliver the same message: the volatility term structure is extremely steep at the very short end but quickly flattens out for maturities above two months.

The bottom panel of Figure 4 shows the Sharpe ratios of the returns on the synthetic zero-coupon claims. Again, the estimated Sharpe ratios appear noisier than the corresponding ones computed from variance swaps (Figure 3), but substantiate the main finding that the only maturities where we see a large and negative Sharpe ratio are the very short ones.

We also compared the prices of claims to future variance in the variance swap market to

²¹The declining term structure of Sharpe ratios on short positions in volatility is consistent with the finding of van Binsbergen, Brandt, and Koijen (2012) that Sharpe ratios on claims to dividends decline with maturity, and that of Duffee (2011) that Sharpe ratios on Treasury bonds decline with maturity.

²²The most famous use of that result is in the construction of the VIX index, which uses 1-month options, and corresponds to the price of a 1-month variance swap. If returns are not a diffusion, the corresponding portfolio of options will not perfectly replicate the payoff of the variance swap. Indeed, the difference between the variance swap price and the portfolio of options contains some information about jumps that has been explored for example in Bollerslev and Todorov (2011) Ait-Sahalia et al. (2014). We note here that while this difference between variance swaps and option-based synthetic contracts is economically informative, the magnitude of this difference is much smaller than the magnitude of the differences in volatility across the term structure that are the focus of this paper. Despite the small difference between the prices of the two contracts, our main pricing result is easily visible in both markets.

those of VIX futures: if returns follow a diffusion, the price of an n month variance swap should correspond to the price of an $n - 1$ month VIX future. Appendix Figure A.3 shows that this relation holds closely both at the short end (1 month) and at the longer end (6 months) of the curve, confirming the integration across volatility markets.

Figure 5 replicates the option-based analysis using options for the Euro Stoxx 50, the FTSE 100, the CAC 40 and the DAX indexes. The top panel of Figure 5 plots the average term structure of synthetic zero-coupon volatility claims, together with the S&P 500 curve computed from U.S. variance swaps.²³ To ease the comparison across markets, in this figure we plot the term structures relative to the prices of the respective 2-month claims, so that all the curves are equal to 1 at 2 months maturity.

The bottom panel of figure 5 repeats the exercise using the prices of international variance swaps for the Stoxx 50, FTSE 100 and DAX indexes (for which only one year of data is available). In the appendix (Figure A.2) we confirm that for the indexes for which we have both variance swap prices and synthetic prices obtained from options, the two curves align well.

Both panels of the figure show that the international term structures have an average shape that closely resembles the one observed for the US (the solid line in both panels), suggesting that our results using US variance swaps extend to the international markets.

4 Asset pricing

We now formally examine the pricing of risks in the variance market.

4.1 Reduced-form estimates

We begin by exhibiting our main pricing result in a simple reduced-form setting: investors pay to hedge the immediate realized volatility but not shocks to expected volatility. To test that claim, we need to disentangle shocks to realized variance from shocks to expectations of future volatility. This subsection focuses just on the returns of the variance claims with maturity of 12 months or less since they require less interpolation; all the economically interesting results are clearly visible in this maturity range.

²³We do not plot Sharpe ratios for these markets because the data is of relatively poor quality, and the series of returns are very noisy.

4.1.1 Extracting innovations

As usual in the term structure literature, we begin by extracting principal components from the term structure of zero-coupon variance claims. We find that the first factor explains 97.1 percent of the variation in the term structure and the second explains an additional 2.7 percent. The loadings of the variance swaps on the factors are plotted in the top panel of Figure 6, while the time series of the factors are shown in the bottom panel of the figure. Similar to the term structure of Treasury bonds, the first factor captures the level of the term structure, while the second measures the slope. As we would expect, during times of crisis, the slope turns negative. The level factor captures the longer-term trend in volatility and clearly reverts to its mean more slowly.

The two factors explain 99.9 percent of the variation in variance swap prices and thus encode essentially all the information contained in variance swap prices. If risk premia are constant, then the factors represent market expectations of future variance. If risk premia vary over time, then the factors depend on both expectations of future variance and also risk premia. So if we find that the shocks to both factors are unpriced, then that means that no forward-looking information in the term structure, whether it is about risk premia or variances, is priced.

To extract shocks to variance and expectations, we estimate a first-order vector autoregression (VAR) with the two principal components and realized variance (RV). Including RV in the VAR allows us to separately identify shocks to the term structure of variance swaps and transitory shocks to realized volatility. The three estimated innovations are positively correlated: the correlation between RV and level shocks is 0.7, and that between RV and slope shocks is 0.6.

We rotate the three shocks using a Cholesky factorization where the first shock affects all three variables, the second affects only the slope and RV , and the third shock affects only RV . We will therefore refer to the third shock as the pure RV shock. The pure RV shock allows us to measure the price of risk for a shock that has only a transitory effect on realized variance and no effect on the term structure of variance swap prices, while the other two rotated shocks affect both current realized variance and also expectations of future variance.²⁴

The impulse response functions (reported in Appendix Figure A.4) confirm that the pure RV shock raises RV on impact and has no immediate effect on the level and slope. After

²⁴Note however that we cannot impose that the pure RV shock does not affect level or slope in the future; therefore, we cannot rule out with this decomposition that RV is priced *because* of its predictive power for future level and slope. This however seems implausible, since the first two shocks have much higher predictive power for future level and slope, and their price is not statistically different from zero.

the initial impact, the response of RV declines by over half within one month and falls to nearly zero after 10 months. The pure RV shock is thus a highly transitory shock to realized variance. As we would expect, the effect of the pure RV shock on RV is far larger than the effects of the other two shocks.

4.1.2 Risk prices

We estimate risk prices for the three shocks along with their standard errors using the Fama–MacBeth (1973) procedure.²⁵ We use the returns on the 1- to 12-month zero-coupon variance claims, and the pricing factors are the three orthogonalized shocks from the previous section. The top panel of Table 3 reports the betas of each variance swap return with respect to the three orthogonalized shocks. Short-maturity variance swaps are exposed to all three shocks with the expected signs. The higher maturities are mostly exposed to the level and slope shocks, with essentially no exposure to the pure RV shock.

The bottom panel of Table 3 reports the estimated annualized risk prices. Of the three shocks, only the pure RV shock has a statistically significant risk price. The risk price of that shock is also economically highly significant: it implies that an asset that was exposed only to the pure RV shock would earn an annualized Sharpe ratio of -2.72. Since the three shocks all have the same standard deviation, the magnitudes of the risk prices are directly comparable. Those for shocks 1 and 2 are five to eight times smaller than that for the pure RV shock, and thus economically far less important.

Overall, the analysis shows that investors do not price shocks to the level and slope, but they accept large negative returns to hedge transitory RV shocks.

4.1.3 Controlling for the market return

The analysis so far has shown that the only priced shock in the variance swap market is the immediate realized variance shock. The next natural question is why investors would be so eager to hedge realized variance.

One possibility is that realized volatility provides a good hedge for aggregate market shocks. To test that possibility, we add the market return as an additional factor in the estimation.²⁶ The first column of Table 4 shows that indeed the zero-coupon volatility

²⁵The results are robust to estimating the risk prices using one- and two-step GMM.

²⁶We add the market return as a test asset to impose discipline on its risk premium. For readability and to ensure that the risk premium on the market is matched relatively closely, we increase the weight on the market return by a factor of 12 as a test asset in our cross-sectional tests. That way, the market return carries as much weight in the pricing tests as do all the variance claims combined. The market factor, though, is still the monthly market return, as are all our zero-coupon variance returns.

claims are heavily exposed to the market return. The beta of the one-month forward is -7. But the betas do not rapidly decay with maturities in the way that expected returns do: even the higher maturity swaps have large negative betas – always below -2 – despite an average return of essentially zero. Adding the pure RV shock obtained from our VAR analysis above to the CAPM greatly improves the fit, as shown by column 2 of Table 4. It is the exposure to the transitory component of RV , not the exposure to the market, that explains the cross-section of variance returns.

4.2 The predictability of volatility

Since the key result of the paper concerns the pricing of volatility shocks at different horizons, a natural question is how much investors actually learn about volatility in the future. Perhaps the reason that we do not estimate significant risk prices for volatility news is that there simply is not much news (which would increase standard errors, though the estimates would still be unbiased). We would first note that if that is true, then the macro literature showing that volatility news can drive the business cycle would seem rather less relevant. Second, though, we now show that there is in fact substantial predictability of future volatility.

First, a large literature has shown that realized volatility is strongly predictable at horizons of a few months using high-frequency data, and that univariate and multivariate predictability extends to longer horizons as well.²⁷ Building on that literature, we report in Table 5 R^2 s from predictive regressions of realized volatility at different frequencies and horizons. The first pair of columns focuses on forecasts of monthly realized variance, while the second pair repeats the exercise at the annual frequency. The R^2 s for monthly volatility range from 45 percent at the 1-month horizon to 20 percent at the 12-month horizon. In predicting annual volatility, R^2 range between 56 and 21 percent for horizons of 1 to 10 years.

The third pair of columns in Table 5 reports, as a comparison, the results of forecasts of dividend growth.²⁸ R^2 s for dividend growth are never higher than 9 percent. So in the context of financial markets, there is an economically large amount of predictability of volatility. The appendix takes an extra step beyond Table 5 and shows, using Fama and Bliss (1987) and Campbell and Shiller (1991) regressions, that nearly all the variation in variance swap prices is actually due to variations in expected volatility, rather than risk premia. We

²⁷See for example Andersen et al. (2003), Ait-Sahalia and Mancini (2006), Bandi, Russell and Yang (2008), and Brownlees, Engle and Kelly (2011). Recently, Campbell et al. (2014) focus explicitly on longer horizons (up to 10 years) and show evidence of predictability of realized volatility in a multivariate setting: in particular, they show that both the aggregate price-earnings ratio and the BAA-AAA default spread are useful predictors of long-run volatility.

²⁸We compare predictability of volatility to that of dividends since realized variance in each month is the stochastic payment of the variance swap contract in that month.

thus conclude that investors' expectations of volatility in fact vary substantially over time.

4.3 A no-arbitrage model

In this section, we extend the pricing results reported above by considering a more formal estimation. We analyze a standard no-arbitrage term structure model for variance swaps. The model delivers implications strongly supportive of our reduced-form results. Because the no-arbitrage model uses the prices of the variance swaps, rather than just their returns, and because it uses a full no-arbitrage structure, it is able to obtain much more precise estimates of risk prices. We show that not only are the risk prices on the level and slope factors statistically insignificant, but they are also economically small.

The no-arbitrage model has three additional advantages over the reduced-form analysis: it explicitly allows for time-variation in the volatility of shocks to the economy and risk prices, the standard errors for the risk prices take into account uncertainty about the dynamics of the economy (through the VAR), and it links us more directly to the previous literature. Furthermore, because the inputs to the estimation of the no-arbitrage model are the observed variance swap prices rather than monthly returns, the results in this section do not rely on any interpolation and we can simultaneously use the full time series from 1996 to 2013 and every maturity from one month to 14 years.

4.3.1 Risk-neutral dynamics

As above, we assume that the term structure of variance swaps is governed by a bivariate state vector $(s_t^2, l_t^2)'$. Rather than state the factors as a level and slope, we now treat them as a short- and a long-term component, which will aid in the estimation process. s_t^2 is the one-month variance swap price: $s_t^2 = E_t^Q [RV_{t+1}]$. The other state variable, l_t^2 , governs the central tendency of s_t^2 . Under the risk-neutral measure, Q , $(s_t^2, l_t^2)'$ follows a first order VAR at the monthly frequency.

The factors $(s_t^2, l_t^2)'$ contain the information about the future contained in the term structure of variance swap prices. However, as discussed above, realized variance in month t is not fully explained by those factors: a regression of RV_t onto a panel of variance swap prices for six different maturities (ranging from one month to two years) over our sample period, 1996–2013, returns an adjusted R-squared statistic of only 83%. We therefore specify the joint dynamics of three variables: $(s_t^2, l_t^2, RV_t)'$.

We begin by specifying the conditional risk-neutral mean of the states,

$$E_t^Q \left[\begin{pmatrix} s_{t+1}^2 \\ l_{t+1}^2 \\ RV_{t+1} \end{pmatrix} \right] = \begin{pmatrix} \rho_s^Q & 1 - \rho_s^Q & 0 \\ 0 & \rho_l^Q & 0 \\ 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} s_t^2 \\ l_t^2 \\ RV_t \end{pmatrix} + \begin{pmatrix} 0 \\ v_l^Q \\ 0 \end{pmatrix} \quad (7)$$

where v_l^Q is a constant to be estimated which captures the unconditional mean of realized variance. l_t^2 can be viewed as the risk-neutral trend of s_t^2 . The first two rows of (7) are the discrete-time counterpart to the standard continuous-time setup in the literature, e.g. Egloff, Leippold, and Wu (2010) and Ait-Sahalia, Karaman, and Mancini (2014).²⁹ We diverge from Egloff, Leippold, and Wu (2010) and Ait-Sahalia, Karaman, and Mancini (2014) in explicitly specifying a separate process for realized variance, noting that it is not spanned by the other shocks. The specification of a separate shock to RV_{t+1} allows us to ask how shocks to both realized variance and the term structure factors are priced.³⁰

Given the assumption that $s_t^2 = E_t^Q [RV_{t+1}]$, the price of an n -period variance swap VS_t^n is

$$VS_t^n = E_t^Q \left[\sum_{i=1}^n RV_{t+i} \right] = E_t^Q \left[\sum_{i=1}^n s_{t+i-1}^2 \right] \quad (8)$$

which can be computed by applying (7) repeatedly, and which implies that VS_t^n is affine in s_t^2 and l_t^2 for any maturity.

²⁹For admissibility, we require that $0 < \rho_s^Q < 1$, $\rho_l^Q > 0$, and $v_l^Q > 0$. These restrictions ensure that risk-neutral forecasts of s_t^2 and l_t^2 , hence variance swap prices at various maturities, are strictly positive.

³⁰From a continuous-time perspective, it is not completely obvious how to think about a "shock" to realized variance that is completely transitory. There are two standard interpretations. One is that the innovation in RV_{t+1} represents the occurrence of jumps in the S&P 500 price. Alternatively, there could be a component of the volatility of the diffusive component of the index that has shocks that last less than one month. At some point, the practical difference between a jump and an extremely short-lived change in diffusive volatility is not obvious. In the end, the key feature of the specification is simply that there are shocks to the payout of variance swaps that are orthogonal to both past and future information contained in the term structure.

4.3.2 Physical dynamics and risk prices

Define $X_t \equiv (s_t^2, l_t^2, RV_t)'$. We assume that X follows a VAR(1) under the physical measure:³¹

$$\begin{pmatrix} s_{t+1}^2 \\ l_{t+1}^2 \\ RV_{t+1} \end{pmatrix} = \begin{pmatrix} 0 \\ v_l^Q \\ 0 \end{pmatrix} + \begin{pmatrix} \rho_s & 1 - \rho_s^Q & 0 \\ 0 & \rho_l & 0 \\ \rho_{s,RV} & 0 & 0 \end{pmatrix} \begin{pmatrix} s_t^2 \\ l_t^2 \\ RV_t \end{pmatrix} + \varepsilon_{t+1} \quad (9)$$

$$\varepsilon_{t+1} \sim N \left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, V_t(X_{t+1}) \right) \quad (10)$$

In our main results, we follow Egloff, Leippold, and Wu (2010) and Ait-Sahalia, Karaman, and Mancini (2014) and assume that the market prices of risks are proportional to the states,

$$\Lambda_t = \begin{pmatrix} \lambda_s s_t \\ \lambda_l l_t \\ \lambda_{RV} s_t \end{pmatrix} \quad (11)$$

By letting the market price of RV -risk be a scaled version of s_t , we are able to introduce some non-zero conditional covariance between RV_t and the term structure of variance swap prices, as observed empirically.³²

Given the risk prices, the innovation to the log SDF, m_{t+1} , is

$$m_{t+1} - E_t[m_{t+1}] = \Lambda_t' V_t(X_{t+1})^{-1/2} \varepsilon_{t+1} \quad (12)$$

where the superscript $1/2$ indicates a lower triangular Cholesky decomposition. The term $V_t(X_{t+1})^{-1/2}$ standardizes and orthogonalizes the shocks ε_{t+1} . Λ_t thus represents the price of exposure to a unit standard deviation shock to each component of X_{t+1} .

To maintain the affine structure of the model, we need the product $V_t(X_{t+1})^{1/2} \Lambda_t$ to be affine in X_t . The specification for Λ_t in (11) is therefore typically accompanied by a structure

³¹Admissibility requires that v_l^Q and the feedback matrix in (9) be non-negative, which ensures that the forecasts of X_t , and hence future volatility, be strictly positive.

³²Had we let the third entry of Λ_t be dependent on RV_t , then $\sigma_{s,RV}$ must be set to zero to ensure that the product $V_t(X_{t+1})^{1/2} \Lambda_t$ is affine. Such a parameterization would imply that RV_t is conditionally uncorrelated to both s_t^2 and l_t^2 , thus RV_t would be uncorrelated to the entire term structure of variance swap prices.

for the conditional variance similar to that of Cox, Ingersoll, and Ross (1985),

$$V_t(X_{t+1}) = \begin{pmatrix} \sigma_s^2 s_t^2 & 0 & \sigma_{s,RV} s_t^2 \\ 0 & \sigma_l^2 l_t^2 & 0 \\ \sigma_{s,RV} s_t^2 & 0 & \sigma_{RV}^2 s_t^2 \end{pmatrix} \quad (13)$$

which guarantees that $V_t(X_{t+1})^{1/2}\Lambda_t$ is affine in X_t .

It is important to note that the specifications of Λ_t in (11) and $V_t(X_{t+1})$ in (13) introduce tight restrictions on the difference $E_t(X_{t+1}) - E_t^Q(X_{t+1})$. In the appendix, we therefore consider two alternative specifications for the variance process $V_t(X_{t+1})$ and the risk prices Λ_t that are more flexible in certain dimensions. The results, both in terms of point estimates and standard errors, are essentially identical across the various specifications that we consider, so we report results for this specification here since it is most common in the literature.

4.3.3 Empirical results

The estimation uses standard likelihood-based methods. The appendix describes the details. We use both Dataset 1 and Dataset 2, meaning that the number of variance swap prices used in the estimation varies over time depending on availability.

Model fit Table 6 reports the means and standard deviations of the variance swap prices observed and fitted by our model together with the corresponding root mean squared errors (RMSE). The average RMSE across maturities up to 24 months is 0.73 annualized volatility points (i.e. the units in Figure 1).³³ For maturities longer than 24 months, since we do not have time series of variance swap prices with fixed maturities for the entire sample, we cannot report the sample and fitted moments for any fixed maturity. Instead, we stack all contracts with more than 24 months to maturity into one single series and compute the RMSE from the observed and fitted values of this series. The corresponding RMSE is reported in the last row of Table 6. At 0.87 percentage points, it compares favorably with the RMSE for the shorter maturities. Table 6 suggests that our models with two term structure factors plus RV are capable of pricing the cross-section of variance swap prices for an extended range of maturities. Even when maturities as long as 14 years are included in estimation, the data does not seem to call for extra pricing factors.

³³When we exclude the financial crisis, using a sample similar to that of Egloff, Leippold, and Wu (2010), we obtain an RMSE of 0.33 percentage points, which is nearly identical to their reported value. The increase in fitting error in the full sample is, not surprisingly, brought about by the large volatility spikes that occurred during the crisis.

Risk prices The steady-state risk prices in the model are reported in Table 7 along with their standard errors. As in the previous analysis, we find clearly that it is the purely transitory shock to realized variance that is priced (RV-risk). The Sharpe ratio associated with an investment exposed purely to the transitory RV shock – analogous to the pure RV shock above – is -1.70.

In the VAR analysis in the previous section, the pure RV shock had no immediate impact on the level and slope factors, but it could potentially indirectly affect future expected variance through the VAR feedback. In the no-arbitrage model, that effect is shut off through the specification of the dynamics. That is, the RV shock here is completely transitory – it has no impact on expectations of volatility on any future date. The other two shocks are forced to account for all variation in expectations. The fact that the results are consistent between the no-arbitrage model and the reduced-form analysis in the previous section helps underscore the robustness of our findings to different modeling assumptions.

The short- and long-term factors earn risk premia of only -0.11 and -0.18, respectively, neither of which is significantly different from zero. The lack of statistical significance is not due to particularly large standard errors; the standard errors for the risk prices for the s_t^2 and l_t^2 shocks are in fact substantially smaller than that for the RV shock. Moreover, Sharpe ratios of -0.11 and -0.18 are also economically small. For comparison, the Sharpe ratio of the aggregate stock market in the 1996–2013 period is 0.43. So the risk premia on the short- and long-term components of volatility are between 25 and 42 percent of the magnitude of the Sharpe ratio on the aggregate stock market. On the other hand, the Sharpe ratio for the RV shock is nearly four times larger than that for the aggregate stock market and 10 to 15 times larger than the risk prices on the other two shocks. Our no-arbitrage model thus clearly confirms the results from the previous sections.

Time-series dynamics The estimated parameters determining the dynamics of the state variables under the physical measure are (equation 9):

$$\begin{pmatrix} s_{t+1}^2 \\ l_{t+1}^2 \\ RV_{t+1} \end{pmatrix} = \begin{pmatrix} 0 \\ 0.99^{***} \\ 0 \end{pmatrix} + \begin{pmatrix} 0.82^{***} & 0.16^{**} & 0 \\ 0 & 0.98^{***} & 0 \\ 0.75^{***} & 0 & 0 \end{pmatrix} \begin{pmatrix} s_t^2 \\ l_t^2 \\ RV_t \end{pmatrix} + \varepsilon_{t+1}$$

The key parameter to focus on is the persistence of l_t^2 . The point estimate is 0.9814, with a standard error of 0.0013. At the point estimate, long-term shocks to variance have a half-life of 37 months. That level of persistence is actually *higher* than the persistence of consumption growth shocks in Bansal and Yaron’s (2004) long-run risk model, and only

slightly smaller than the persistence they calibrate for volatility, 0.987. Our empirical model thus allows us to estimate risk prices on exactly the type of long-run shocks that have been considered in calibrations. As we discuss further below, the fact that we find that the long-term shock to volatility is unpriced is strongly at odds with Epstein–Zin preferences.

5 Economic interpretation

The key message of our empirical analysis is that the average investor in the variance swap market is not willing to pay for protection against news about high future volatility. In other words, they do not hedge volatility intertemporally. That fact immediately suggests that models based on Epstein–Zin (1991) preferences, where intertemporal hedging effects are central, will struggle to match the data. To confirm that intuition, we simulate two models with Epstein–Zin preferences. The first is the long-run risk model proposed by Drechsler and Yaron (2011), and the second is a discrete-time version of the model with time-varying disaster risk proposed by Wachter (2013). In both cases, we show that the models generate substantially negative Sharpe ratios for zero-coupon volatility claims. Furthermore, both models imply that the Sharpe ratios earned from rolling over long-term zero-coupon variance claims are nearly as negative as those earned from holding just the one-month variance swap, counter to what we observe empirically in Figure 3.

The evidence we provide that there is no intertemporal hedging runs counter to many models beyond Epstein–Zin. Merton’s ICAPM, for example, implies that shocks to expected volatility should be priced since volatility affects the investment opportunity set.³⁴ The variance swap market thus is not well explained by the ICAPM. Similarly, in models with value-at-risk or leverage constraints, the constraint on financial intermediaries depends on expected volatility, rather than realized volatility.³⁵ In general, then, it is forward-looking volatility that is relevant in most asset pricing models.

The lack of intertemporal hedging in the variance swap market suggests a myopic model of investors. We therefore consider a simple model in which investors have power utility. While it is well known that the power utility model fails to match many asset pricing facts when consumption follows a process with low volatility, Rietz (1988), Barro (2006), Martin (2013), and others show that allowing for a small probability of a large decline in consumption can allow the power utility model to match many standard asset pricing moments. Gabaix

³⁴This is true even if volatility does not predict returns. If volatility rises but expected returns remain constant, then the investment opportunity set has deteriorated.

³⁵For example, financial intermediaries might be limited in the total amount of risk they may take. When expected volatility is higher, their demand for risky securities will fall.

(2012) extends the disaster model to allow for a time-varying exposure of firms to disasters, essentially a time-varying recovery rate. We find that Gabaix’s model is able to match both the qualitative and quantitative features of the variance swap market. This suggests that investors in the variance swap market are mostly worried about large negative shocks to the economy in which returns collapse and variance spikes, and are purchasing variance swaps to hedge these, and only these, shocks.

5.1 Structural models of the variance premium

5.1.1 A long-run risk model

The first model we examine is Drechsler and Yaron’s (DY; 2011) long-run risk model. They extend the model proposed by Bansal and Yaron (2004) to allow for jumps in both the consumption growth rate and volatility. DY show that the model can match the mean, volatility, skewness, and kurtosis of consumption growth and stock market returns, and generates a large variance risk premium (1-month variance swap price) that forecasts market returns, as in the data. DY is thus a key quantitative benchmark in the literature.

The structure of the endowment process is

$$\Delta c_t = \mu_{\Delta c} + x_{t-1} + \varepsilon_{c,t} \quad (14)$$

$$x_t = \mu_x + \rho_x x_{t-1} + \varepsilon_{x,t} + J_{x,t} \quad (15)$$

$$\bar{\sigma}_t^2 = \mu_{\bar{\sigma}} + \rho_{\bar{\sigma}} \bar{\sigma}_{t-1}^2 + \varepsilon_{\bar{\sigma},t} \quad (16)$$

$$\sigma_t^2 = \mu_{\sigma} + (1 - \rho_{\sigma}) \bar{\sigma}_{t-1}^2 + \rho_{\sigma} \sigma_{t-1}^2 + \varepsilon_{\sigma,t} + J_{\sigma,t} \quad (17)$$

where Δc_t is log consumption growth, the shocks ε are mean-zero and normally distributed, and the shocks J are jump shocks. σ_t^2 controls both the variance of the normally distributed shocks and also the intensity of the jump shocks. We follow their calibration for the endowment process exactly. While the calibration includes jumps that could potentially be interpreted as disasters, they are smaller and more frequent than the types of events studied in the disaster literature following Rietz (1988) and Barro (2006).

Aggregate dividends are modeled as

$$\Delta d_t = \mu_d + \phi x_{t-1} + \varepsilon_{d,t} \quad (18)$$

Equity is a claim on the dividend stream, and we treat variance claims as paying the realized variance of the return on equities.

DY combine that endowment process with Epstein–Zin preferences, and we follow their

calibration. Because there are many parameters to calibrate, we refer the reader to DY for the full details. However, the parameters determining the volatility dynamics are obviously critical to our analysis. Note that the structure of equations (16) and (17) is the same as the VAR in our no-arbitrage model in equation (9). The parameters governing volatility in DY’s calibration and the corresponding values from our estimation are:

	DY	Estimates
ρ_σ	0.87	0.82
$\rho_{\bar{\sigma}}$	0.987	0.9814
$\text{stdev}(\varepsilon_{\bar{\sigma},t})$	0.10	0.05
$\text{stdev}(\varepsilon_{\sigma,t} + J_{\sigma,t})$	1.10	1.48

The two feedback coefficients, ρ_σ and $\rho_{\bar{\sigma}}$, are nearly identical to our estimated values. Their long-term component, $\bar{\sigma}^2$, has a persistence of 0.987, which compares favorably with our estimate of 0.9814. Similarly, their calibration of $\rho_\sigma = 0.87$ is comparable to our estimate of 0.82. The calibration deviates somewhat more in the standard deviations of the innovations.

Overall, though, DY’s calibration implies volatility dynamics highly similar to what we observe empirically. The close match is not surprising as DY’s model was calibrated to fit the behavior of the (one-month) VIX and realized variance. As a robustness check, though, in the appendix we also simulate DY setting the standard deviations of the innovations to match our empirical estimates and obtain implications for variance swap prices that are essentially unchanged.

Given the high quality of DY’s calibration, if the long-run risk model fails to match the term structure of variance swap prices, it is not because it has an unreasonable description of the dynamics of volatility. Rather, we would conclude that the failure is due to the specification of the preferences, namely Epstein–Zin.

5.1.2 Time-varying disaster risk

The second model we study is a discrete-time version of Wachter’s (2013) model of time-varying disaster risk. In this case, consumption growth follows the process,

$$\Delta c_t = \mu_{\Delta c} + \sigma_{\Delta c} \varepsilon_{\Delta c,t} + J_{\Delta c,t} \tag{19}$$

where $\varepsilon_{\Delta c,t}$ is a mean-zero normally distributed shock and J_t is a disaster shock. The probability of a disaster in any period is F_t , which follows the process

$$F_t = (1 - \rho_F) \mu_F + \rho_F F_{t-1} + \sigma_F \sqrt{F_{t-1}} \varepsilon_{F,t} \quad (20)$$

The CIR process ensures that the probability of a disaster is always positive in the continuous-time limit, though it can generate negative values in discrete time. We calibrate the model similarly to Wachter (2013) and Barro (2006). Details of the calibration are reported in the appendix. The model is calibrated at the monthly frequency. In the calibration, the steady-state annual disaster probability is 1.7 percent as in Wachter (2013). σ_F is set to 0.0075 (ε_F is a standard normal), and $\rho_F = 0.87^{1/12}$, which helps generate realistically volatile stock returns and a persistence for the price/dividend ratio that matches the data. If there is no disaster in period t , $J_t = 0$. Conditional on a disaster occurring, $J_t \sim N(-0.30, 0.15^2)$, as in Barro (2006). Finally, dividends are a levered consumption claim, with λ representing leverage (calibrated to 2.8).³⁶

Wachter (2013) combines this specification of disasters with Epstein–Zin preferences. One of her key results is that a model with time-varying disaster risk and power utility has strongly counterfactual predictions for the behavior of interest rates and other asset prices. She thus argues that time-varying disaster risk should be studied in the context of Epstein–Zin preferences. We follow her in assuming the elasticity of intertemporal substitution is 1, and we set risk aversion to 3.6.³⁷

5.1.3 Time-varying recovery rates

The final model we study is a version of Gabaix’s (2012) model of disasters with time-varying recovery rates. Because the probability of a disaster is constant, power utility and Epstein–Zin are equivalent in terms of their implications for risk premia. We use power utility in our calibration, which eliminates the intertemporal hedging motives present in the two previous models. In this model, the expected value of firms following a disaster is variable. Specifically, we model the consumption process identically to equation (19) above, but with

³⁶The occurrence of a disaster shock implies that firm value declines instantaneously. To calculate realized variance for periods in which a disaster occurs, we assume that the shock occurs over several days with maximum daily return of -5 percent. For example, a jump of 20% would occur over 4 consecutive days, with a 5% decline per day. This allows for a slightly delayed diffusion of information and also potentially realistic factors such as exchange circuitbreakers. The small shocks $\varepsilon_{\Delta c,t}$ are treated as though they occur diffusively over the month, as in Drechsler and Yaron (2011).

³⁷Given the calibration of the endowment, if risk aversion is raised any higher the model does not have a solution. The upper bound on risk aversion is a common feature of models in which the riskiness of the economy varies over time.

the probability of a disaster, F_t , fixed at 1 percent per year (Gabaix’s calibration). Following Gabaix, dividend growth is

$$\Delta d_t = \mu_{\Delta d} + \lambda \varepsilon_{\Delta c,t} - L_t \times 1 \{J_{\Delta c,t} \neq 0\} \quad (21)$$

λ here represents leverage. $1 \{\cdot\}$ is the indicator function. Dividends are thus modeled as permanently declining by an amount L_t on the occurrence of a disaster. The value of L is allowed to change over time and follows the process

$$L_t = (1 - \rho_L) \bar{L} + \rho_L L_{t-1} + \varepsilon_{L,t} \quad (22)$$

We calibrate $\bar{L} = 0.5$ and $\rho_L = 0.87^{1/12}$ as in the previous model, and $\varepsilon_{L,t} \sim N(0, 0.16)$. We set the coefficient of relative risk aversion to 7 to match the Sharpe ratio on one-month variance swaps. Other than the change in risk aversion, our calibration of the model is nearly identical to Gabaix’s (2012), which implies that we will retain the ability to explain the same ten puzzles that he examines. He did not examine the ability of his model to match the term structure of variance claims, so this paper provides a new test of the theory.

5.2 Results

We now examine the implications of the three models for the zero-coupon variance curve. Figure 7 plots population moments from the models against the values observed empirically. The top panel reports annualized Sharpe ratios for zero-coupon variance claims with maturities from 1 to 12 months. Our calibration of Gabaix’s model with time-varying recovery rates matches the data well: it generates a Sharpe ratio for the one-month claim of -1.3, while all the forward claims earn Sharpe ratios of zero, similarly to what we observe in the data.

The two models with Epstein–Zin preferences, on the other hand, both generate Sharpe ratios for claims on variance more than one month ahead that are counterfactually large, especially when compared to the Sharpe ratio of the 1-month variance swap. For both the long-run risk and the time-varying disaster model, the Sharpe ratio on the three-month variance claim is roughly three-fourths as large as that on the one-month claim, whereas the three-month claim actually earns a slightly positive return empirically.

The economic intuition for the result is straightforward. In both models, there is a predictable component in RV : in DY both the jump probability and the volatility of the diffusive component are subject to persistent shocks, and in the time-varying disaster model the disaster probability is persistent. Since agents have Epstein–Zin preferences, they are

willing to pay to hedge shocks to expected future RV . Variance forwards are a financial instrument that precisely allows them to achieve this goal; agents are willing to pay nearly as much to protect themselves against the 1-month, transitory RV shock as they are to hedge news about variance in any month in the future.

The expected returns on the variance claims are closely related to the average slope of the term structure. The bottom panel of Figure 7 reports the average term structure in the data and in the models. The figure shows, as we would expect, that both models with Epstein–Zin preferences do not generate a curve that is much steeper at the short end than it is at the long end. Instead, the DY model generates a curve that is too steep everywhere (including on the very long end), while the time-varying disaster model generates a curve that is too flat everywhere.³⁸ On the other hand, the average term structure in the model with time-varying recovery rates qualitatively matches what we observe in the data – it is perfectly flat after the first month. Note that removing the forward-looking component of the time-varying disaster model (by using power utility instead of Epstein–Zin, and/or removing the predictability in disaster probability) would allow us to match the shape of the term structure. In that case, the model essentially collapses to the time-varying recovery rate model.

The comparison between the calibrated models and the data reported in Figure 7 does not take into account the statistical uncertainty due to the fact that we only observe variance swap prices in a specific sample. To directly test the models against the data, we simulate the calibrated models and verify how likely we would be to see a period in which the variance swap curve looks like it does in our data. In particular, the left-hand column of Figure 8 plots results from 10,000 215-month simulations of the models. In each simulation, we calculate the average term structure of the variance curve, and normalize the value at the third month to 1. We then plot the median value from the simulations along with the 95-percent sampling interval from the simulations.

Both the long-run risk model and time-varying disaster risk model have a hard time in matching the empirical shape of the variance term structure in the simulations, and particularly producing a steep slope at the short end of the curve and a small slope for higher maturities. Gabaix’s model with time-varying recovery rates performs qualitatively better: the 1-month variance swap is priced significantly higher than average realized volatility, but the slope is zero for all the rest of the curve. The long-run risk and time-varying disaster models are statistically rejected at the short end of the curve, while the long-run risk model is rejected at the long end of the curve.

³⁸The models have similar Sharpe ratios but different slopes of the term structure because the latter depends on the expected return, not the Sharpe ratio.

The right-hand column of figure 8 simulates variance swap prices in the models out to maturities of 10 years. The sampling intervals are wider because our sample with 10-year maturities only runs for 70 months. The story is similar to that in the right-hand column though: all three models fail quantitatively, but the time-varying recovery model is the one that best matches the qualitative features of the term structure.

We conclude from Figures 7 and 8 that models involving Epstein–Zin preferences, or likely any other source of intertemporal hedging demand, are unlikely to match the term structure of variance swaps. DY’s long-run risk model is calibrated in such a way that it nearly exactly matches our estimated volatility dynamics. When we adjust the calibration to match our estimates perfectly, the implications for variance swap prices are essentially unaffected. We therefore conclude that models with a major intertemporal hedging motive, such as Epstein–Zin preferences, do not match the features of the variance swap market. On the other hand, a model in which investors have power utility, and hence make investment choices myopically, is able to better match our data.

The model with time-varying recovery clearly does not perfectly match the data. For example, it implies that the variance of stock returns in the absence of a disaster is constant, and it implies that a single factor governs the term structure of variance swap prices, rather than the two factors that we observe in the data. As discussed in Gabaix (2012), the model is highly stylized. Nevertheless, it demonstrates a clear economic rationale for the basic patterns that we observe in the data. We leave it to future work to develop a fully quantitatively accurate structural model of the term structure of variance swaps.

It is also important to note that neither Drechsler and Yaron’s (2011) nor Wachter’s (2013) model were originally designed to match the full term structure of variance swap rates. Notably, though, the features of the data that they were designed to match are also matched by Gabaix’s (2012) model. In particular, Drechsler and Yaron (2011) focus on the ability of the variance risk premium to forecast market returns, while Wachter (2013) emphasizes excess volatility in asset prices induced by changes in the probability of a disaster. Gabaix (2012) shows (and we have confirmed in our implementation) that his model can match both of those patterns, among many others.

We also note that we do not rule out every possible model with Epstein–Zin preferences. Since consumption growth is i.i.d. in Gabaix’s (2012) model, its implications are identical whether we use power utility or Epstein–Zin preferences. Rather, we simply rule out Epstein–Zin preferences in settings where there is a strong force for intertemporal hedging, for example when volatility varies over time (as we observe in the data).

To conclude, the main features of the models that affect their ability to match our data

can be summarized as follows. In models with Epstein–Zin preferences, agents will pay to hedge shocks to expected future volatility, especially at long horizons. Long-term zero-coupon variance claims should thus have large negative returns because they hedge the agents against these shocks. In the data, we do observe shocks to future expected volatility, but find that their price is close to zero. Models with power utility, or where the variation in expected stock market volatility is independent of consumption volatility, solve that problem since investors are myopic and shocks to future expected volatility are not priced. However, the models also need to explain the high risk price associated with the realized volatility shock. In a power utility framework, this can be achieved if states of the world with high volatility are associated with large drops in consumption, as in a disaster model.

5.3 The behavior of volatility during disasters

In order for variance swaps to be useful hedges in disasters, realized volatility must be high during large market declines. A number of large institutional asset managers sell products meant to protect against tail risk that use variance swaps, which suggests that they or their investors believe that realized volatility will be high in future market declines.³⁹

In the spirit of Barro (2006), we now explore the behavior of realized volatility during consumption disasters and financial crises using a panel data of 17 countries, covering 28 events. We find two main results. First, volatility is indeed significantly higher during disasters. Second, the increase in volatility is not uniform during the disaster period; rather, volatility spikes for one month only during the disaster and quickly reverts. It is those short-lived but extreme spikes in volatility that make variance swaps a good product to hedge tail risk.

We collect daily market return data from Datastream for a total of 37 countries. Datastream’s data series begin in 1973. We compute realized volatility in each month for each country. To identify disasters, we use both the years marked by Barro (2006) as consumption disasters and the years marked by Schularick and Taylor (2012), Reinhart and Rogoff (2009) and Bordo et al. (2001) as financial crises.⁴⁰ Given the short history of realized volatility available, our final sample contains 17 countries for which we observe realized volatility and that experienced a disaster during the available sample. Table 8 shows for each country the starting year of our RV sample (last year is 2013), together with the years we use to identify consumption and financial disasters.

³⁹In particular, see Man Group’s TailProtect product (Man Group (2014)), Deutsche Bank’s ELVIS product (Deutsche Bank, 2010) and the JP Morgan Macro Hedge index.

⁴⁰See Giglio et al. (2014) for a more detailed description of the data sources.

The first three columns of Table 8 compare the monthly annualized realized volatility during disaster years and in all non-disaster years. We consider the reference period for a disaster event the calendar year in which it occurs. Column 1 shows the maximum volatility observed in any month of the year identified as a disaster averaged across all disasters for each country. Column 2 shows the average volatility during the disaster years, and column 3 shows the average volatility in all other years.

Comparing columns 2 and 3, we can see that in almost all cases realized volatility is indeed higher during disasters. For example, in the US the average annualized realized volatility is 25 percent during disasters and 15 percent otherwise. Column 1 reports the average across crises of the highest observed volatility. We see that within disaster years there is large variation in realized volatility: the maximum volatility is always much higher than the average volatility, even during a disaster. This result indicates that disasters are associated with large spikes in realized volatility, rather than a generalized increase in volatility during the whole period.

To confirm this result, in Figure 9 we perform an event study around the peak of volatility during a disaster. To produce the figure, for each country and for each disaster episode we identify the month of the volatility peak during that crisis (month 0) and the three months preceding and following that. We then scale the volatility behavior by the value reached at the peak, so that the series for all events are normalized to 1 at the time of the event. We then average the rescaled series across our 28 events.

The figure shows that indeed, the movements in volatility that we observe during disasters are short-lived spikes, where volatility is high for essentially only a single month. In the single months immediately before and after the one with the highest volatility, volatility is 40 percent lower than its peak, and it is lower by half both three months before and after the worst month.

What Table 8 and Figure 9 show is that an asset that provides protection for news about high future volatility provides only weak protection against market crashes since volatility is not particularly high *on average* during crashes. The asset that provides the best tail protection is one that provides protection against high realized volatility, rather than high expected volatility. That said, though, it is important to note that not all periods of high volatility coincide with large declines in consumption. The 1987 stock market crash is the prime example of an episode with high volatility in financial markets that had little or no effects on real activity. So while variance swaps clearly provide a hedge against crashes, their returns are not perfectly correlated with crashes.⁴¹

⁴¹That fact also implies that there must be assets with even larger Sharpe ratios than one-month variance swaps, and thus the annual Hansen–Jagannathan (1991) bound must be greater than 1.5.

6 Conclusion

This paper shows that it is only the transitory part of realized variance that is priced. That fact is not consistent with a broad range of structural asset pricing models. It is qualitatively consistent with a model in which investors desire to hedge rare disasters, but even that model does not match all the quantitative features of the data. Interestingly, the data is not consistent with all disaster models. The key feature that we argue models need in order to match our results is that variation in expected stock market volatility is not priced by agents, whereas the transitory component of volatility is strongly priced.

The idea that variance claims are used to hedge crashes is consistent with the fact that many large asset managers, such as Deutsche Bank, JP Morgan, and Man Group sell products meant to hedge against crashes that use variance swaps and VIX futures. These assets have the benefit of giving tail protection, essentially the form of a long put, but also being delta neutral (in an option-pricing sense). They thus require little dynamic hedging and yield powerful protection against large declines.

In the end, we conclude that the variance swap market is well described as a set of assets that investors use to hedge crashes. They do not seem to desire to hedge changes in the probability of a crash (or any other sort of volatility). That fact is at odds with models that imply that investors hedge intertemporally. An investor who does have an intertemporal hedging motive and wants protection against increases in future volatility would be well served to purchase that protection, essentially for free, from financial markets.

References


- Ait-Sahalia, Yacine, Mustafa Karaman, and Lorian Mancini**, “The Term Structure of Variance Swaps, Risk Premia and the Expectations Hypothesis,” 2014. Working paper.
- Amengual, Dante and Dacheng Xiu**, “Resolution of policy uncertainty and sudden declines in volatility,” 2014. Working paper.
- Andersen, Torben, Tim Bollerslev, Francis Diebold, and Paul Labys**, “Modeling and Forecasting Realized Volatility,” *Econometrica*, 2003, *71(2)*, 579–625.
- Backus, David, Mikhail Chernov, and Ian Martin**, “Disasters Implied by Equity Index Options,” *Journal of Finance*, 2011, *66(6)*, 1969–2012.
- Bakshi, Gurdip and Nikunj Kapadia**, “Delta-Hedge Gains and the Negative Market Volatility Risk Premium,” *The Review of Financial Studies*, 2003, *16(2)*, 527–566.
- , **George Panayotov, and Georgios Skoulakis**, “Improving the predictability of real economic activity and asset returns with forward variances inferred from option portfolios,” *Journal of Financial Economics*, 2011, *100(3)*, 475–495.
- Bandi, Jeffrey Russell Federico and Chen Yang**, “Realized Volatility Forecasting and Option Pricing,” *Journal of Econometrics*, 2008, *147(1)*, 34–46.
- Bansal, Ravi and Amir Yaron**, “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *Journal of Finance*, 2004, *59(4)*, 1481–1509.
- , **Dana Kiku, Ivan Shaliastovich, and Amir Yaron**, “Volatility, the Macroeconomy and Asset Prices,” *Journal of Finance*, 2013.
- Barras, Laurent and Aytex Malkhozov**, “Variance Risk Premium Dynamics in Equity and Option Markets,” 2014. Working paper.
- Barro, Robert J.**, “Rare Disasters and Asset Markets in the Twentieth Century,” *Quarterly Journal of Economics*, 2006, *121(3)*, 823–866.
- Bloom, Nicholas, Scott Baker, and Steve Davis**, “Measuring Economic Policy Uncertainty,” *Working Paper*, 2009.
- Boguth, Oliver, Murray Carlson, Adlai Fisher, and Mikhail Simutin**, “Leverage and the Limits of Arbitrage Pricing: Implications for Dividend Strips and the Term Structure of Equity Risk Premia.” Working paper.
- , —, —, and —, “More on Levered Noise and Arbitrage Pricing.” Working paper.
- Bollerslev, Tim and Viktor Todorov**, “Tails, fears, and risk premia,” *Journal of Finance*, 2011, *66(6)*, 2165–2211.

- Bordo, Michael, Barry Eichengreen, Daniela Klingebiel, and Maria Soledad Martinez-Peria**, “Is the crisis problem growing more severe?,” *Economic policy*, 2001, 16 (32), 51–82.
- Branger, Nicole and Clemens Volkert**, “What is the Equilibrium Price of Variance Risk? A Long-Run Risk Model with Two Volatility Factors,” 2010. Working paper.
- Broadie, Mark, Mikhail Chernov, and Michael Johannes**, “Understanding Index Option Returns,” *The Review of Financial Studies*, 2009, 22(11), 4493–4529.
- Brownlees, Christian, Robert Engle, and Bryan Kelly**, “A Practical Guide to Volatility Forecasting Through Calm and Storm,” *Journal of Risk*, 2011, 14(2), 1–20.
- Campbell, John Y., Stefano Giglio, Christopher Polk, and Robert Turley**, “An Intertemporal CAPM with Stochastic Volatility,” 2013. Working paper.
- Carr, Peter and Liuren Wu**, “Variance Risk Premiums,” *Review of Financial Studies*, 2009, 22(3), 1311–1341.
- Christoffersen, Peter, Kris Jacobs, Chayawat Ornthanalai, and Yintian Wang**, “Option Valuation with Long-run and Short-run Volatility Components,” *Journal of Financial Economics*, 2008, 90(3), 272–297.
- Cieslak, Anna and Pavol Povala**, “Information in the Term Structure of Yield Curve Volatility,” 2014. Working paper.
- Coval, Joshua D. and Tyler Shumway**, “Expected Option Returns,” *The Journal of Finance*, 2001, 56(3), 983–1009.
- Cox, John C., Jr. Jonathan E. Ingersoll, and Stephen A. Ross**, “A Theory of the Term Structure of Interest Rates,” *Econometrica*, 1985, 53(2), 385–407.
- Deutsche Bank Pension Strategies & Solutions**, “Tail Risk Hedging: A Roadmap for Asset Owners,” May 2010.
- Dew-Becker, Ian and Stefano Giglio**, “Asset Pricing in the Frequency Domain: Theory and Empirics,” 2013. Working paper.
- Drechsler, Itamar and Amir Yaron**, “What’s Vol Got to Do with it?,” *The Review of Financial Studies*, 2011, 24(1), 1–45.
- Duffee, Gregory R.**, “Sharpe ratios in term structure models,” 2011. Working paper.
- Egloff, Daniel, Madrus Leippold, and Liuren Wu**, “The Term Structure of Variance Swap Rates and Optimal Variance Swap Investments,” *Journal of Financial and Quantitative Analysis*, 2010, 45(5), 1279–1310.
- Epstein, Larry G. and Stanley E. Zin**, “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis,” *The Journal of Political Economy*, 1991, 99(2), 263–286.

- Exchange, Chicago Board Options**, “The CBOE Volatility Index - VIX,” Technical Report, Chicago Board Options Exchange 2009.
- Fama, Eugene F. and James D. MacBeth**, “Risk, Return, and Equilibrium: Empirical Tests,” *Journal of Political Economy*, 1973, *81(3)*, 607–636.
- and **Robert R. Bliss**, “The Information in Long-Maturity Forward Rates,” *The American Economic Review*, 1987, *77(4)*, 680–692.
- Gabaix, Xavier**, “Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance,” *Quarterly Journal of Economics*, 2012, *127(2)*, 645–700.
- , **Arvind Krishnamurthy**, and **Olivier Vigneron**, “Limits of Arbitrage: Theory and Evidence from the Mortgage-Backed Securities Market,” *Journal of Finance*, 2007, *62(2)*, 557–595.
- Giglio, Stefano, Matteo Maggiori, and Johannes Stroebe**, “Very Long-Run Discount Rates,” *Unpublished Manuscript New York University, and University of Chicago*, 2014.
- Gorton, Gary B., Fumio Hayashi, and K. Geert Rouwenhorst**, “The Fundamentals of Commodity Futures Returns,” *Review of Finance*, 2013, *17(1)*, 35–105.
- Hansen, Lars Peter and Ravi Jagannathan**, “Implications of Security Market Data for Models of Dynamic Economies,” *Journal of Political Economy*, 1991, *99(2)*, 225–262.
- Inc., Man Investments**, “The case for TailProtect Strategy,” 2014.
- Jackwerth, Jens Carsten and Mark Rubinstein**, “Recovering Probability Distributions from Option Prices,” *The Journal of Finance*, 1996, *51(5)*, 1611–1631.
- Kelly, Bryan, Lubos Pastor, and Pietro Veronesi**, “The Price of Political Uncertainty: Theory and Evidence from the Option Market,” 2014. Working paper.
- Lu, Zhongjin and Yingzi Zhu**, “Volatility Components: The Term Structure Dynamics of VIX Futures,” *The Journal of Futures Markets*, 2010, *30(3)*, 230–256.
- Martin, Ian**, “Simple Variance Swaps,” *Working Paper*, 2013.
- , “What is the Expected Return on the Market?,” *Working Paper*, 2014.
- Mencia, Javier and Enrique Sentana**, “Valuation of VIX Derivatives,” *Journal of Financial Economics*, 2013, *108(2)*, 367–391.
- Merton, Robert C.**, “An Intertemporal Capital Asset Pricing Model,” *Econometrica*, 1973, *41(5)*, 867–887.
- Mixon, Scott**, “The implied volatility term structure of stock index options,” *Journal of Empirical Finance*, 2007, *14*, 333–354.


- Reinhart, Carmen M. and Kenneth S. Rogoff**, *This time is different: Eight centuries of financial folly*, Princeton University Press, 2009.
- Rietz, Thomas A.**, “The Equity Risk Premium: A Solution,” *Journal of Monetary Economics*, 1988, *22(1)*, 117–131.
- Schularick, Moritz and Alan M. Taylor**, “Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870–2008,” *American Economic Review*, 2012, *102(2)*, 1029–1061.
- van Binsbergen, Jules H., Michael W. Brandt, and Ralph Koijen**, “On the Timing and Pricing of Dividends,” *American Economic Review*, 2012, *102*, 1596–1618.
- van Binsbergen, Jules, Michael Brandt, and Ralph Koijen**, “On the Timing and Pricing of Dividends,” *The American Economic Review*, 2012, *102(4)*, 1596–1618.
- Vijh, Anand M.**, “Liquidity of the CBOE Equity Options,” *Journal of Finance*, 1990, *45(4)*, 1157–1179.
- Wachter, Jessica A.**, “Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?,” *Journal of Finance*, 2013, *68(3)*, 987–1035.
- Zhou, Guofu and Yingzi Zhu**, “Volatility Trading: What Is the Role of the Long-Run Volatility Component?,” *Journal of Financial and Quantitative Analysis*, 2012, *47(2)*, 273–307.

Figure 1: Time series of zero-coupon variance claim prices




Note: The figure shows the time series of zero-coupon variance claim prices of different maturities. For readability, each line plots the prices in annualized volatility terms, $100 \times \sqrt{12 \times Z_t^n}$, for a different n . The top panel plots zero-coupon variance claim prices for maturities of 1 month, 3 months, and one year. The bottom panel plots zero-coupon variance claim prices for maturities of 1 year, 5 years and 10 years. Both panels also plot annualized realized volatility, $100 \times \sqrt{12 \times Z_t^0}$.

Figure 2: Average zero-coupon variance claim prices




Note: The figure shows the average prices of zero-coupon variance claims of different maturity, across different periods. The top panel shows average prices between 2008 and 2013, when we observe maturities up to 10 years. The bottom panel shows averages between 1996 and 2013, for claims of up to 1 year maturity. In each panel, the "x" mark prices of maturities we directly observe in the data (for which no interpolation is necessary). All prices are reported in annualized volatility terms, $100 \times \sqrt{12 \times Z_t^n}$. Maturity zero corresponds to average realized volatility, $100 \times \sqrt{12 \times Z_t^0}$.

Figure 3: Annualized Sharpe ratios for zero-coupon variance claims




Note: The figure shows the annualized Sharpe ratio for the zero-coupon variance claims. The returns are calculated assuming that the investment in an n-month variance claim is rolled over each month. Dotted lines represent 95% confidence intervals. The sample used is 1996-2013.

Figure 4: Synthetic zero-coupon variance claims: prices and Sharpe ratios




Note: See Figure 2. The solid line in the top panel plots average prices of zero-coupon variance claims calculated using the formula for the VIX index and data on option prices from the CBOE. The dotted line is the set of average prices of zero-coupon variance claims constructed from variance swap prices. The bottom panel plots annualized Sharpe ratios for zero-coupon variance claims returns with prices calculated using the VIX formula and CBOE option data. Dotted lines in the bottom panel represent 95% confidence intervals. The sample covers the period 1996-2013.

Figure 5: Average zero-coupon variance claim prices for international markets



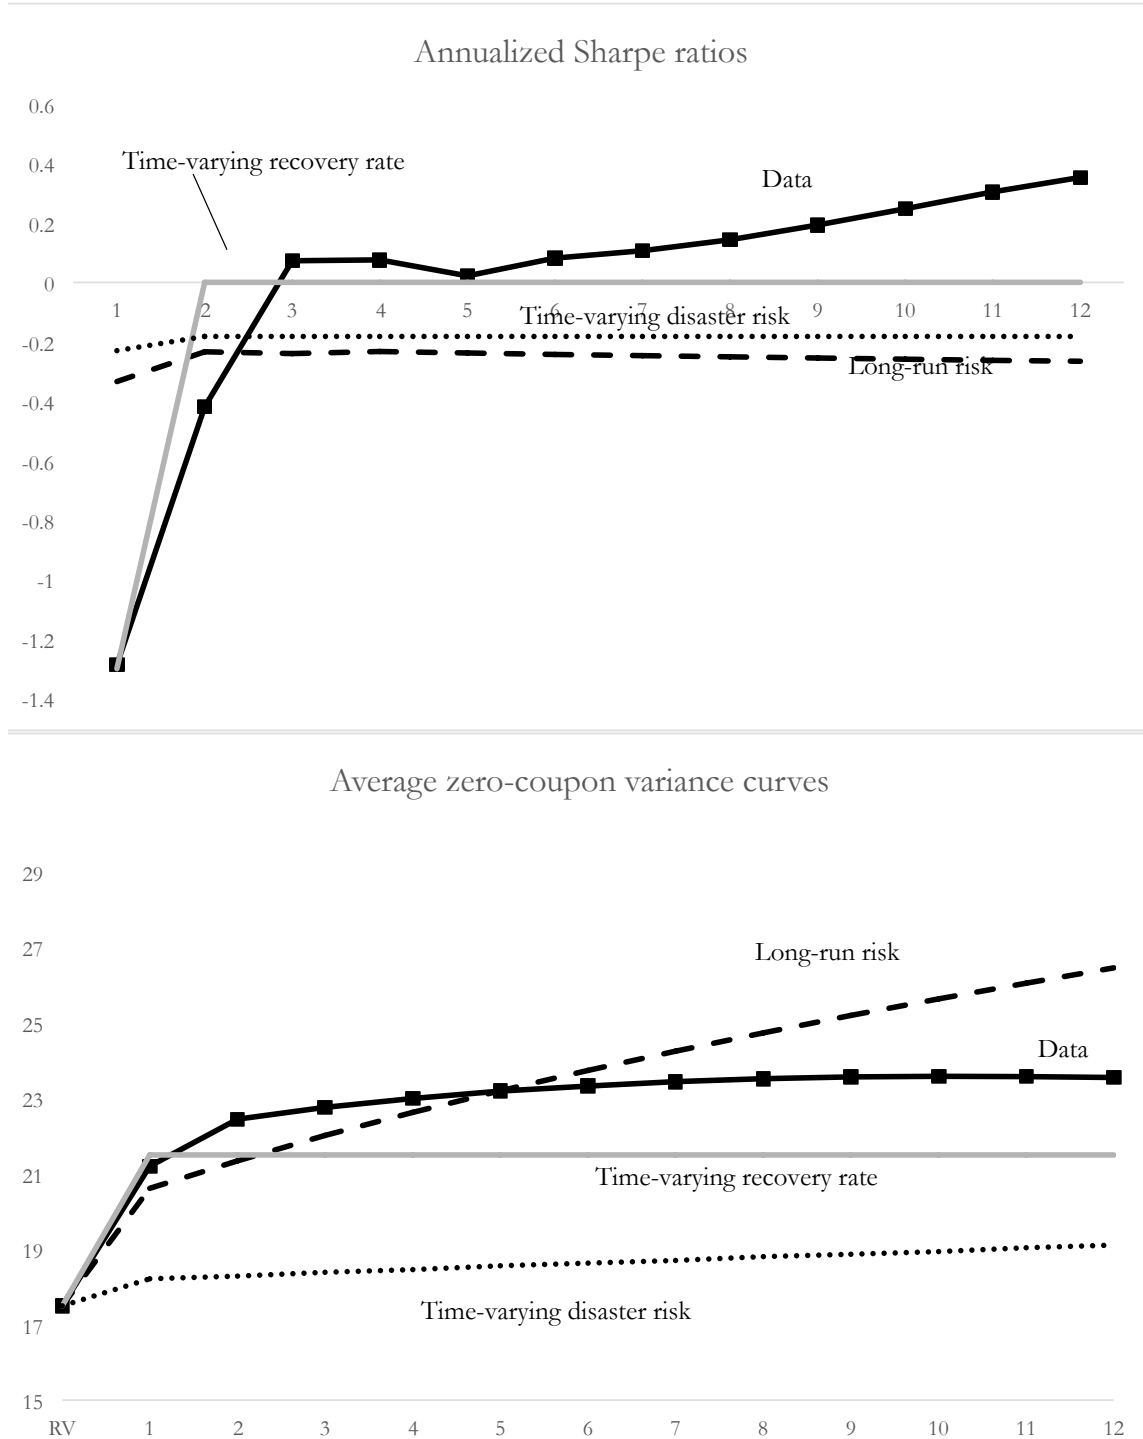
Note: The figure plots the average prices of zero-coupon variance claims as in Figure 2 for different international indices. The series for the S&P 500 (both in the top and bottom panel) is obtained from variance swaps (as in Figure 2). The top panel shows international curves obtained using option prices, using the same methodology used to construct the VIX for the S&P 500 (as in Figure 4). Options data is from OptionMetrics. The series cover FTSE 100, CAC 40, DAX, and STOXX 50, for the period 2006-2014. The bottom panel shows international curves obtained using variance swaps on the FTSE 100, DAX, and STOXX 50, for one year starting in April 2013. All series are rescaled relative to the price of the 2-month zero-coupon variance price (so they all cross 1 at maturity 2 months).

Figure 6: Principal components of variance swap prices



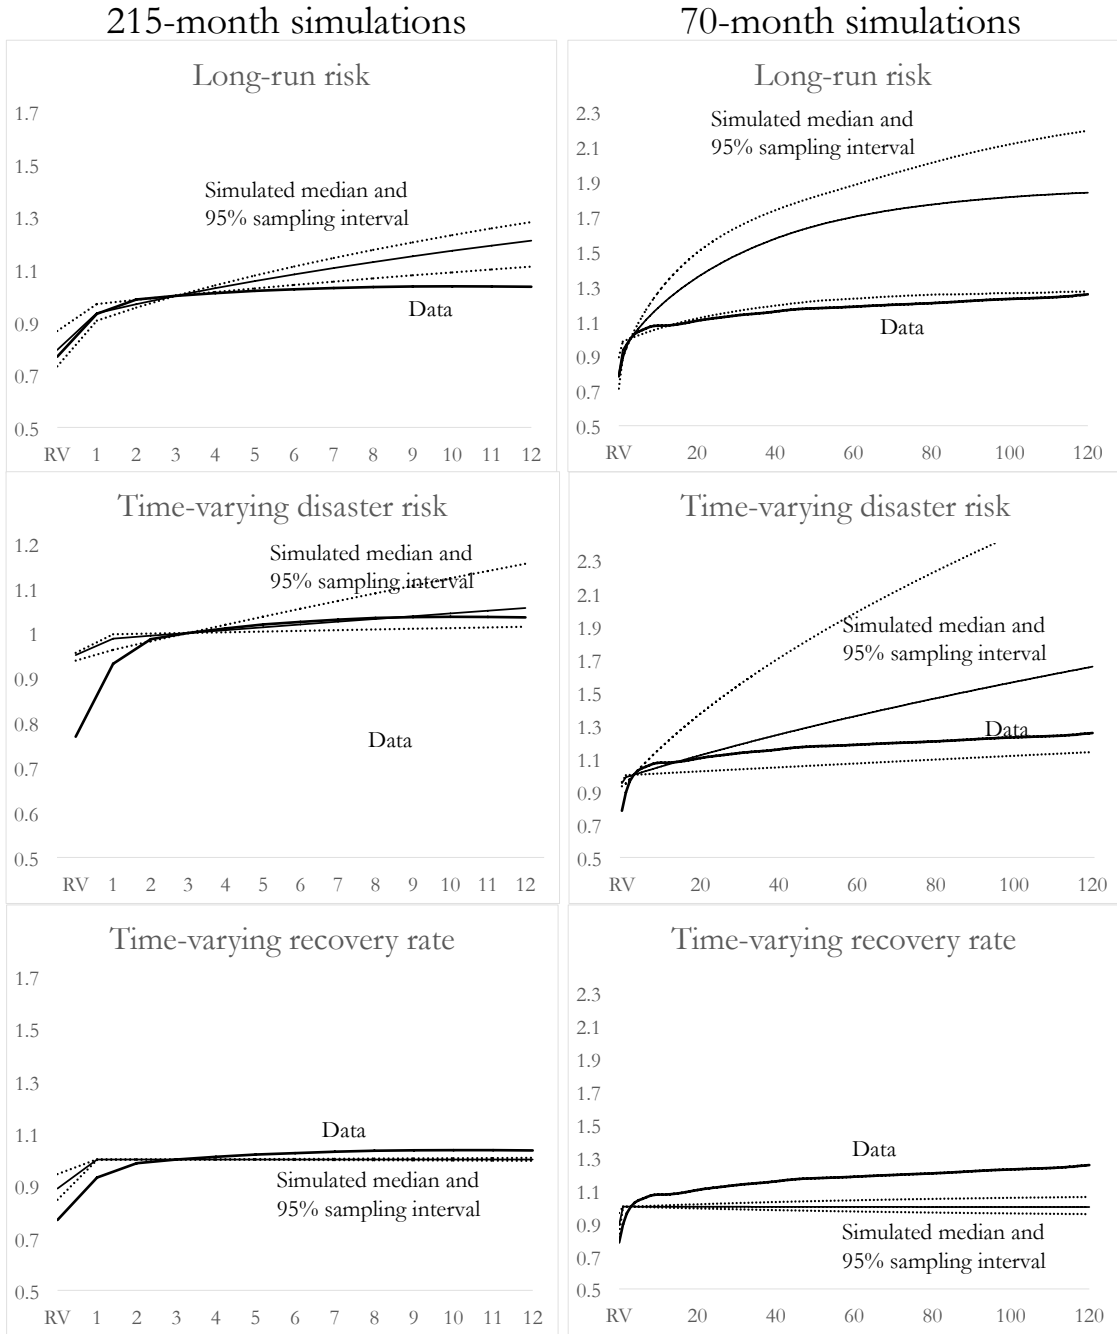
Note: The top panel plots the loadings of the variance swap prices on the level and slope factors (first two principal components). The bottom panel plots the time series of the level and slope factors. Both are normalized to have zero mean and unit standard deviation and are uncorrelated in the sample. The sample covers the period 1996-2013.

Figure 7: Sharpe ratios and average term structure in different models




Note: Notes: The top panel gives the population Sharpe ratios from the three models and the sample values from the data. The bottom panel plots population means of the prices of zero-coupon claims. All the curves are normalized to have the same value for the realized variance.

Figure 8: Slope of the term structure in different models



Note: Simulated zero-coupon variance claim prices from structural models. All curves are normalized to equal 1 at three months. The dotted lines are 95-percent sampling intervals from the simulations. The left-hand column simulates the models for 215 months to match our sample of 1- to 12-month variance claims. The right-hand column simulates 70 months to match the sample with up to 10-year claims.

Figure 9: Average behavior of RV during consumption disasters and financial crises



Note: We calculate realized variance in each month of a crisis and scale it by the maximum realized variance in each crisis. The figure plots the average of that scaled series for each country and crisis in terms of months relative to the one with the highest realized variance.

Table 1: Volume of variance swaps across maturities

Maturity (months)	Volume (million vega)	Volume (percentage)
1	402	6%
2	403	6%
3	78	1%
4-6	1037	14%
7-12	1591	22%
13-24	2371	33%
25-60	1315	18%
60+	48	1%
Total	7245	100%

Note: Total volume of variance swap transactions occurred between March 2013 and June 2014 and collected by the DTCC.

Table 2: Characteristics of returns

Maturity (months)	Mean	Std	Min	25th p.	Median	75th p.	Max	Skew	Exc.Kurt.
1	-25.6	69.0	-86.2	-59.2	-40.7	-16.5	687.1	6.1	54.3
2	-5.6	47.8	-59.3	-32.5	-18.2	9.5	376.0	3.9	23.3
3	0.8	34.0	-46.1	-21.4	-4.9	15.0	249.4	2.7	14.1
4	0.7	27.4	-42.2	-17.1	-5.4	11.4	170.4	1.9	7.5
5	0.2	22.5	-37.3	-14.1	-3.6	9.9	126.7	1.6	5.2
6	0.5	19.7	-31.0	-12.3	-3.7	12.9	100.6	1.3	3.3
7	0.6	18.7	-31.4	-12.4	-2.4	11.0	90.7	1.1	2.5
8	0.8	17.4	-29.8	-11.4	-2.9	11.7	81.6	1.0	2.0
9	0.9	16.2	-27.7	-10.3	-1.8	9.6	74.6	0.9	1.7
10	1.2	15.6	-30.0	-9.7	-1.9	10.0	70.8	0.9	1.5
11	1.5	16.0	-32.6	-9.7	-1.9	11.3	69.7	0.9	1.3
12	1.8	17.4	-35.0	-10.2	-2.1	12.2	70.4	1.0	1.4

Note: The table reports descriptive statistics of the monthly returns for zero-coupon variance claims (in percentage points). For each maturity n , returns are computed each month as $R_{t+1}^n = \frac{Z_{t+1}^{n-1} - Z_t^n}{Z_t^n}$. Given the definition that $Z_t^0 = RV_t$, the return on a one-month claim, R_{t+1}^1 is the percentage return on a one-month variance swap.

Table 3: Reduced-form pricing estimates

<i>Panel A: Betas</i>		Shock 1	Shock 2	Pure RV
	1	0.38	-0.19	0.32
	2	0.40	-0.11	0.05
	3	0.28	0.00	-0.01
	4	0.22	0.04	-0.01
	5	0.18	0.05	0.00
Maturities (months)	6	0.15	0.06	0.00
	7	0.13	0.06	0.01
	8	0.12	0.06	0.01
	9	0.11	0.06	0.00
	10	0.11	0.06	0.00
	11	0.11	0.06	0.00
	12	0.11	0.06	0.00
<i>Panel B: Risk Prices</i>				
	Risk prices	-0.51	-0.33	-2.72***
	Standard error	0.44	0.64	0.47
	R^2		0.998	

Note: Results of Fama–MacBeth regressions using the 12 zero-coupon claims as test assets and the three rotated VAR innovations as pricing factors. Shock 1 has effects on all three factors; shock 2 affects only the slope and RV, and pure RV only affects RV on impact. The top section reports betas on the three factors. The bottom section reports estimated risk prices and the Fama–MacBeth standard errors. *** denotes significance at the 1-percent level. Risk prices are annualized by multiplying by $\sqrt{12}$.

Table 4: Pricing results for the CAPM

		(1)	(2)	
		CAPM	CAPM + pure RV	
<i>Panel A: Betas</i>		Rm	Pure RV	Rm
	1	-7.11	0.36	-7.87
	2	-6.93	0.08	-7.11
	3	-5.04	0.01	-5.07
	4	-3.97	0.01	-3.98
	5	-3.14	0.01	-3.16
VS, maturities	6	-2.59	0.01	-2.62
	7	-2.32	0.02	-2.36
	8	-2.13	0.01	-2.17
	9	-2.03	0.01	-2.05
	10	-2.01	0.01	-2.03
	11	-2.06	0.01	-2.08
	12	-2.17	0.01	-2.18
	Market	1.00	0	1.00
<i>Panel B: Risk Prices</i>				
	Risk prices	0.0105***	-0.72***	0.0046
	Standard error	0.0034	0.12	0.0032
	R^2	0.377	0.997	

Note: See Table 3. Pricing results for the CAPM and for the CAPM with pure RV. The test assets are the 12 zero-coupon variance claims and the market portfolio. The market portfolio is given 12 times as much weight as the variance claims to ensure that it is priced correctly in the estimation.

Table 5: Forecasting volatility at different horizons: R2

Predictor: with PE_t, DEF_t	monthly RV_{t+n}		yearly RV_{t+n}		yearly Δd_{t+n}			
	RV_t	RV_t	RV_t	RV_t	Δd_t	Δd_t		
Months	\checkmark		Years	\checkmark		Years		
1	0.39	0.45	1	0.41	0.56	1	0.00	0.09
2	0.21	0.34	2	0.10	0.25	2	0.00	0.02
3	0.18	0.32	3	0.05	0.09	3	0.06	0.07
6	0.15	0.26	5	0.02	0.04	5	0.05	0.07
12	0.10	0.18	10	0.00	0.21	10	0.02	0.03

Note: The first column of the table reports R2 of predictive regressions of monthly volatility n months ahead at the monthly frequency. The second column reports R2 of predictive regressions of yearly volatility n years ahead at the yearly frequency. The second column reports R2 of predictive regressions of yearly log dividend growth n years ahead at the yearly frequency. The left side of each column reports univariate regressions using the lagged value of the target, while the right side of each column adds the market price-earnings ratio and the default spread as predictors. The sample is 1926-2014.

Table 6: Average prices and pricing errors for the no-arbitrage model

Maturity (months)	Sample		Fitted		RMSE
	Mean	Std	Mean	Std	
1	21.24	8.00	21.70	7.86	1.00
2	21.89	7.55	21.88	7.55	0.33
3	22.22	7.25	22.04	7.30	0.40
6	22.75	6.64	22.42	6.74	0.63
12	23.20	6.06	22.99	6.11	0.53
24	23.65	5.58	23.87	5.45	0.61
>24					0.87

Note: Prices are reported in annualized volatility terms. The RMSE is calculated using the deviation of the fitted price from the sample price in annualized volatility terms.

Table 7: Steady-state risk prices

Sources of risk	Estimates	Standard Error
s_t^2 -risk	-0.11	0.31
l_t^2 -risk	-0.18	0.12
RV-risk	-1.70***	0.48

Note: Estimates of steady-state risk prices from the no-arbitrage model. Risk prices are annualized by multiplying by $\sqrt{12}$. *** denotes significance at the 1-percent level.

Table 8: Realized volatility during disasters

Country	Peak Vol. during disaster	Mean Vol. during disaster	Mean Vol. outside disaster	Sample start year	Consumption disasters	Financial crises
US	47.5	25.2	14.9	1926	1933	1929, 1984, 2007
UK	24.6	16.4	15.1	1973		1974, 1984, 1991, 2007
France	72.1	31.4	16.6	1973		2008
Japan	40.9	21.4	15.1	1973		1992
Australia	33.7	13.8	15.1	1973		1989
Germany	83.1	28.1	14.3	1973		2008
Italy	55.1	23.0	19.2	1973		1990, 2008
Sweden	52.3	27.7	19.5	1982		1991, 2008
Switzerland	67.1	27.4	12.1	1973		2008
Belgium	66.1	32.0	12.4	1973		2008
Finland	29.3	18.9	25.0	1988	1993	1991
South Korea	80.0	43.6	24.6	1987	1998	1997
Netherlands	77.7	33.2	14.7	1973		2008
Spain	69.4	30.5	17.1	1987		2008
Denmark	37.2	14.7	14.4	1973		1987
Norway	44.2	20.2	20.7	1980		1988
South Africa	36.9	17.8	18.5	1973		1977, 1989

Note: Characteristics of annualized monthly realized volatility during and outside disasters across countries. Returns data used to construct realized volatility for the US is from CRSP, for all other countries from Datastream. Consumption disaster dates are from Barro (2006). Financial crisis dates are from Schularick and Taylor (2012), Reinhart and Rogoff (2009) and Bordo et al. (2001).

A.1 Synthetic variance swap prices

We construct option-based synthetic variance claims for maturity n , $VIX_{n,t}^2$ using the methods described by the CBOE (2009) in their construction of the VIX index, using data from Optionmetrics covering the period 1996 to 2012.

In particular, we construct $VIX_{n,t}^2$ for maturity n on date t as

$$VIX_{n,t}^2 \equiv \frac{2}{n} \sum_i \frac{\Delta K_i}{K_i^2} \exp(-nR_{n,t}) P(K_i) \quad (\text{A.1})$$

where i indexes options; K_i is the strike price of option i ; $\Delta K_i = \frac{K_{i+1} - K_{i-1}}{2}$ unless i is the first or last option used, in which case ΔK_i is just the difference in strikes between K_i and its neighbor; $R_{n,t}$ is the n -day zero-coupon yield (from Fama and Bliss (1987)); and $P(K_i)$ is the midpoint of the bid-ask spread for the out-of-the-money option with strike K_i . The summation uses all options available with a maturity of n days. We deviate slightly here from the CBOE, which drops certain options with strikes very far from the current spot. For each t and n , we require the presence of at least 4 out-of-the-money calls and puts. We create $VIX_{n,t}^2$ for all monthly maturities by interpolating between available option maturities, using the same techniques as the CBOE.

Under the assumption that the price of the underlying follows a diffusion (i.e. does not jump), it is the case that

$$VIX_{n,t}^2 = E_t^Q \left[\int_t^{t+n} \sigma_j^2 dj \right] \quad (\text{A.2})$$

$$\approx E_t^Q \left[\sum_{j=1}^n RV_{t+j} \right] \quad (\text{A.3})$$

where σ_t^2 is the instantaneous volatility at time t .¹ The second line simply notes that the integrated volatility is approximately equal to the sum of squared daily returns (where the quality of the approximation improves as the sampling interval becomes shorter). In other words, when the underlying follows a diffusion, $VIX_{n,t}^2$ corresponds to the price of an idealized variance swap where the squared returns are calculated at arbitrarily high frequency.

Given $VIX_{n,t}^2$ prices at the monthly intervals, we also construct zero-coupon variance claims, $VIXZ_{n,t}$ as

$$VIXZ_{n,t} = VIX_{n,t}^2 - VIX_{n-1,t}^2 \quad (\text{A.4})$$

¹See Carr and Wu (2009).

The zero-coupon prices obtained from options are very close to those obtained from variance swaps. Figure 4 in the text compares the two curves for the S&P 500. Figure A.2 compares the curves for the STOXX 50, FTSE 100 and DAX, and shows that the curves are similar for international markets as well (though in that case there is more noise).

A.2 Decomposing the source of variation of variance swap prices

In this section we investigate whether the variation in variance swap prices is primarily driven by changes in expected future volatility or changes in risk premia.

Using the definition of zero-coupon variance claims, the following identity holds, for each maturity n :

$$Z_t^n - Z_t^0 = [E_t RV_{t+n} - RV_t] + \left[-E_t \sum_{j=0}^{n-1} r_{t+1+j}^{n-j} \right] \quad (\text{A.5})$$

where r_t^n is the one-period return of the n -period zero-coupon claim. An increase in the n -period zero-coupon variance price must predict either an increase in future realized variance or lower future variance risk premia. Following Fama and Bliss (1987) and Campbell and Shiller (1991), we can then decompose the total variance of $Z_t^n - Z_t^0$ into the component that predicts future RV , and the component that captures movements in risk premia.²

The right side of Table A.1 shows this decomposition for different maturities between 1 month and 1 year. We see that most of the variation in variance swap prices can be attributed to movements in the expectation of future realized variance, not risk premia. In particular, at horizons of 3 to 12 months, essentially all the variation in prices is due to variation in expected volatility rather than variation in risk premia.

At the same time, we know from Table 2 and Figure 1 that prices of both short-term and long-term claims vary substantially: this indicates that the expectation of future realized variance changes dramatically over time. For example, the standard deviation of innovations in the 12-month zero-coupon claim is 17% per year. Given the finding in this section that the variance of Z_t^{12} is driven entirely by changes in volatility expectations, we see that investors' expectations of future volatility in fact vary substantially over time.

²A similar exercise was conducted by Mixon (2007), using S&P 500 options to predict future implied volatility.

A.3 Estimation of the no-arbitrage model

In this appendix, we provide more details on the estimation of the no-arbitrage model. We also report the estimated parameters and more detailed pricing results.

A.3.1 Estimation Strategy

For estimation purposes, a standard and convenient practice of the term structure literature is to assume that some fixed-weights “portfolios” of VS prices are priced perfectly. These portfolios in turn allow one to invert for the latent states which are needed in the computation of the likelihood scores of the data.

A challenge in implementing this practice in our context is that for parts of our sample the set of available maturities may change from one observation to the next. In the later part of our sample, we use VS prices with maturities up to 14 years, whereas the longest maturity for the earlier sample is only two years.

To tackle this issue, we maintain the assumption from the term structure literature that the current term structure of the VS prices perfectly reveals the current values of states. Nonetheless, we depart from the standard term structure practice by using some time-varying-weights “portfolios” of VS prices in identifying the states at each point in time. The portfolios weights are determined in a way to optimally accommodate different sets of maturities at each point in the sample.

To begin, let D_t denote the vector of observed data obtained by stacking up the vector of VS prices on top of the realized variance RV_t . Because VS prices are affine in states, we can write:

$$D_t = A + BX_t. \tag{A.6}$$

All entries of the last column of B , except for the last row, are zeros because VS prices are only dependent on s_t^2 and l_t^2 . In addition, since the last entry of D_t corresponds to RV_t , the last row of A is 0 and the last row of B is (0,0,1). Keep in mind that the length of D_t can vary from time to time due to different maturity sets for the observed VS data.

We assume that D_t is observed with iid errors:

$$D_t^o = D_t + e_t \tag{A.7}$$

where D_t^o denotes the observed counterpart of D_t . Adopting a standard practice in the term structure literature, we assume that the observational errors for the VS prices are uncorrelated and have one common variance σ_e^2 . Because observed RV_t is used in practice

to determine payoffs to VS contracts, it is natural to assume that RV_t is observed without errors. Combined, if a number of J VS prices are observed at time t , the covariance matrix of e_t , Σ_e , takes the following form:

$$\Sigma_e = \begin{pmatrix} I_J \sigma_e^2 & 0_{J \times 1} \\ 0_{1 \times J} & 0 \end{pmatrix}, \quad (\text{A.8})$$

where I_J denotes the identity matrix of size J .

We now explain how we can recursively identify the states. Assume that we already know X_t . Now imagine projecting X_{t+1} on D_{t+1}^o conditioning on all information up to time t . Our assumption (borrowed from the term structure literature) that the current term structure of the VS prices perfectly reveals the current values of the states implies that the fit of this regression is perfect. In other the words, the predicted component of this regression, upon observing D_{t+1}^o ,

$$E_t(X_{t+1}) + cov_t(X_{t+1}, D_{t+1}^o) var_t(D_{t+1}^o)^{-1} (D_{t+1}^o - E_t(D_{t+1}^o)), \quad (\text{A.9})$$

must give us the values for the states at time $t + 1$: X_{t+1} . All the quantities needed to implement (A.9) are known given X_t . Specifically, $E_t(X_{t+1}) = K_0 + K_1 X_t$ and $var_t(D_{t+1}^o) = V_t(X_{t+1}) + \Sigma_e$ where $V_t(X_{t+1})$ is computed according to each of our three specifications for the covariance. $E_t(D_{t+1}^o)$ is given by $A + B E_t(X_{t+1})$. And $cov_t(X_{t+1}, D_{t+1}^o)$ is given by $V_t(X_{t+1}) B'$. Clearly, the calculation in (A.9) can be carried out recursively to determine the values of X_t for the entire sample.

Our approach is very similar to a Kalman filtering procedure apart from the simplifying assumption that D_{t+1}^o fully reveals X_{t+1} . That is, $V_t(X_{t+1} | D_{t+1}^o) \equiv 0$. In a term structure context, Joslin, Le, and Singleton (2013) show that this assumption allows for convenient estimation, yet delivers typically highly accurate estimates.

We can view (A.9) as some “portfolios” of the observed data D_{t+1}^o with the weights given by: $cov_t(X_{t+1}, D_{t+1}^o) var_t(D_{t+1}^o)^{-1}$. As a comparison, whereas the term structure literature typically choose, prior to estimation, a fixed weight matrix corresponding to the lower principal components of the observed data, we do not have to specify *ex ante* any loading matrix. Our approach determines a loading matrix that optimally extracts information from the observed data and, furthermore, can accommodate data with varying lengths.

As a byproduct of the above calculations, we have available the conditional means and variances of the observed data: $E_t(D_{t+1}^o)$ and $V_t(D_{t+1}^o)$. These quantities allow us to compute

the log QML likelihood score of the observed data as (ignoring constants):

$$\mathcal{L} = \sum_t -\frac{1}{2} \|\text{var}_t(D_{t+1}^o)^{-1/2}(D_{t+1}^o - E_t(D_{t+1}^o))\|_2^2 - \frac{1}{2} \log |\text{var}_t(D_{t+1}^o)|. \quad (\text{A.10})$$

Estimates of parameters are obtained by maximizing \mathcal{L} . Once the estimates are obtained, we convert the above QML problem into a GMM setup and compute robust standard errors using a Newey West matrix of covariance.

A.3.2 Alternative variance specifications

Constant variance structure

In the first alternative specification, we let $V_t(X_{t+1})$ be a constant matrix Σ_0 . Since both $E_t(X_{t+1})$ and $E_t^Q(X_{t+1})$ are linear in X_t , Λ_t is also linear in X_t . We refer to this as the CV (for constant variance) specification.

Flexible structure

It is important to note that the specifications of Λ_t in (11) and $V_t(X_{t+1})$ in (13) introduce very tight restrictions on the difference: $E_t(X_{t+1}) - E_t^Q(X_{t+1})$. Simple algebra shows that the first entry of the product $V_t(X_{t+1})^{1/2}\Lambda_t$ is simply a scaled version of s_t^2 . This means that the dependence of $E_t(X_{t+1})$ and $E_t^Q(X_{t+1})$ on l_t^2 and RV_t and a constant must be exactly canceled out across measures. Similar arguments lead to the following restrictions on the condition mean equation:

$$E_t \begin{pmatrix} s_{t+1}^2 \\ l_{t+1}^2 \\ RV_{t+1} \end{pmatrix} = \underbrace{\begin{pmatrix} 0 \\ v_t^Q \\ 0 \end{pmatrix}}_{K_0} + \underbrace{\begin{pmatrix} \rho_s & 1 - \rho_s^Q & 0 \\ 0 & \rho_l & 0 \\ \rho_{s,RV} & 0 & 0 \end{pmatrix}}_{K_1} \begin{pmatrix} s_t^2 \\ l_t^2 \\ RV_t \end{pmatrix}. \quad (\text{A.11})$$

So in the CIR specification, the conditional mean equation only requires three extra degrees of freedom in: ρ_s , ρ_l , and $\rho_{s,RV}$. The remaining entries to K_0 and K_1 are tied to their risk-neutral counter parts. By contrast, all entries of K_0 and K_1 are free parameters in the CV specification. Whereas CV offers more flexibility in matching the time series dynamics of X_t , the parsimony of CIR, if well specified, can potentially lead to stronger identification. However, this parsimony of the CIR specification, if mis-specified, can be restrictive. For example, in the CIR specification, RV_t is not allowed to play any role in forecasting X .

In the flexible specification, we would like to combine the advantages of the CV specification in matching the time series dynamics of X and the advantages of the CIR specification in modeling time-varying volatilities.

For time-varying volatility, we adopt the following parsimonious structure:

$$V_t(X_{t+1}) = \Sigma_1 s_t^2, \quad (\text{A.12})$$

where Σ_1 is a fully flexible positive definite matrix. This choice allows for non-zero covariances among all elements of X .

The market prices of risks are given by:

$$\Lambda_t = V_t(X_{t+1})^{-1/2}(E_t(X_{t+1}) - E_t^Q(X_{t+1})), \quad (\text{A.13})$$

where the superscript $^{1/2}$ indicates a lower triangular Cholesky decomposition. Importantly, we do not require the market prices of risks to be linear, or of any particular form. This is in stark contrast to the CIR specification which requires the market prices of risks to be scaled versions of states. As a result of this relaxation, no restrictions on the conditional means dynamics are necessary. Aside from the non-negativity constraints, the parameters K_0, K_1 that govern $E_t(X_{t+1}) = K_0 + K_1 X_t$ are completely free, just as in the CV specification. In particular, RV_t is allowed to forecast X and thus can be important in determining risk premiums. We label this specification as the FLEX specification (for its flexible structure).

A.3.3 Additional estimation results

Table A.2 reports the risk neutral parameters of our no arbitrage models: ρ_s^Q , ρ_l^Q , and v_l^Q . As expected, these parameters are very strongly identified thanks to the rich cross-section of VS prices used in the estimation. Recall that our risk-neutral construction is identical for all three of our model specifications. As a result, estimates of risk neutral parameters are nearly invariant across different model specifications.

The effect of including the crisis is that the estimates for ρ_s^Q and v_l^Q are higher, whereas the estimate of ρ_l^Q is the same. A higher v_l^Q is necessary to fit a higher average VS curve. A higher estimate for ρ_s^Q implies that risk-neutral investors perceive the short-run factor s_t^2 as more (risk-neutrally) persistent. In other words, movements of the one-month VS price (s_t^2) will affect prices of VS contracts of much longer maturities.

We report the time series parameters – K_0, K_1 , and other parameters that govern the conditional variance $V_t(X_{t+1})$ – for the CV, CIR, and FLEX specifications in Tables A.3, A.4, and A.5, respectively.

The values of annualized steady-state risk prices implied by each model specification are reported in A.6. Regardless of how the quantity of risk ($V_t(X_{t+1})$) is modeled, RV-risk is always very significantly negatively priced. The point estimates and standard errors are similar across the various specifications for the variance process (and hence the physical dynamics), emphasizing the results of our findings. The table also shows that the results are similar regardless of whether the financial crisis is included in the estimation sample. Since the financial crisis was a period when the returns on variance swaps were extraordinarily high, excluding it from the data causes to estimate risk prices that are even more negative than in the full sample.

A.4 Calibration and simulation of the models

This section gives the details of the three models analyzed in the main text.

A.4.1 Long-run Risk (Drechsler and Yaron)

Our calibration is identical to that of Drechsler and Yaron (DY; 2011), so we refer the reader to the paper for a full description of the model. We have confirmed that our simulation matches the moments reported in DY (tables 6, 7, and 8, in particular). Here we report an extra set of results that modifies DY’s calibration to match our empirical estimates. Specifically, we change the standard deviation of the innovations to $\bar{\sigma}^2$ from 0.10 to 0.05. We also multiply the volatility of both the diffusive and jump innovations in σ^2 by 1.35. Those two changes make the processes for σ^2 and $\bar{\sigma}^2$ match the one we estimate in the no-arbitrage model.

The following table reports results from three simulations of the model – the original DY calibration, a calibration with risk aversion raised to 16, and the calibration with the change in the volatility parameters:

Moment	Original	High RRA	Empirical volatility
E[var premium]	9.06	11.30	11.79
stdev[var premium]	12.28	18.78	22.42

The variance premium here measures the difference between the 1-month variance swap price and realized variance (it is the “level difference” from DY). As we would expect, the variance premium rises when we raise risk aversion. The “original” column corresponds to the calibration used in the main text. Shifting to the calibration that matches the estimated

volatility process in the final column we can see reduces the mean and standard deviation of the variance premium, but only to a level similar to DY's original calibration. The reason for this is that Epstein–Zin preferences emphasize long-run shocks, and the estimated volatility process puts more weight on the short-run than the long-run shock to volatility compared to DY's original calibration.

Figures A.5 and A.6 replicate figures 7 and 8 but comparing now DY's original calibration, our version with higher risk aversion, and the version with the calibration of volatility to the empirical estimates. In all three cases, the Sharpe ratios at all maturities are substantially negative, roughly the value of the Sharpe ratio of the aggregate stock market. The term structures are also far too steeply upward sloping.

A.4.2 Time-varying disasters (Wachter)

The key equations driving the economy are

$$\Delta c_t = \mu_{\Delta c} + \sigma_{\Delta c} \varepsilon_{\Delta c,t} + J_{\Delta c,t} \quad (\text{A.14})$$

$$F_t = (1 - \rho_F) \mu_F + \rho_F F_{t-1} + \sigma_F \sqrt{F_{t-1}} \varepsilon_{F,t} \quad (\text{A.15})$$

$$\Delta d_t = \lambda \Delta c_t \quad (\text{A.16})$$

where $\varepsilon_{\Delta c,t}$ and $\varepsilon_{F,t}$ are standard normal innovations. This is a discrete-time version of Wachter's setup, and it converges to her model as the length of a time period approaches zero. The model is calibrated at the monthly frequency. Conditional on a disaster occurring, $J_{\Delta c,t} \sim N(-0.3, 0.15^2)$. The number of disasters that occurs in each period is a Poisson variable with intensity F_t . The other parameters are calibrated as:

Parameter	Value	Parameter	Value
$\mu_{\Delta c}$	0.02/12	$\sigma_{\Delta c}$	0.02/sqrt(12)
σ_F	0.0075	λ	2.8
ρ_F	$0.87^{1/12}$	μ_F	0.017/12

In the analytic solution to the model, the price/dividend ratio for a levered consumption claim takes the form $pd_t = z_0 + z_1 F_t$. The Campbell–Shiller approximation to the return (which becomes arbitrarily accurate as the length of a time interval shrinks) is

$$r_{t+1} = \theta pd_{t+1} + \Delta d_{t+1} - pd_t \quad (\text{A.17})$$

$$= \theta pd_{t+1} + \lambda \Delta c_{t+1} - pd_t \quad (\text{A.18})$$

It is straightforward to show analytically (the derivation is available on request) that

$$pd_t = z_0 + z_1 F_t \tag{A.19}$$

In the absence of a disaster, we treat the shocks to consumption and the disaster probability as though they come from a diffusion, so that the realized variance is $\theta^2 z_0^2 \sigma_F^2 F_{t-1} + \lambda \sigma_{\Delta c}^2$. When a disaster occurs, we assume that the largest daily decline in the value of the stock market is 5 percent. So, for example, a 30-percent decline would be spread over 6 days. The results are largely unaffected by the particular value assumed. The realized variance when a disaster occurs is then $\theta^2 z_0^2 \sigma_F^2 F_{t-1} + \lambda \sigma_{\Delta c}^2 - (0.05) J_{\Delta c, t}$ (assuming $J_{\Delta c, t} \leq 0$).

The model is solved analytically using methods similar to those in DY. Specifically, household utility, v_t , is

$$v_t = (1 - \beta) c_t + \frac{\beta}{1 - \alpha} \log E_t [\exp((1 - \alpha) v_{t+1})] \tag{A.20}$$

α is set to 3.6 and $\beta = 0.98^{1/12}$. The recursion can be solved because the cumulant-generating function for a poisson mixture of normals (the distribution of $J_{\Delta c}$) is known analytically (again, see DY).

The pricing kernel is

$$M_{t+1} = \beta \exp(-\Delta c_{t+1}) \frac{\exp((1 - \alpha) v_{t+1})}{E_t [\exp((1 - \alpha) v_{t+1})]} \tag{A.21}$$

Asset prices, including those for claims on realized variance in the future, then follow immediately from the solution of the lifetime utility function and cumulant-generating function for $J_{\Delta c}$. The full derivation and replication code is available upon request.

We attempted to keep the calibration as close as possible to Wachter's. The two differences are that we increase risk aversion somewhat in order to try to generate a larger variance risk premium and that we use a normal distribution for the disasters rather than the empirical distribution used by Wachter (to allow us to obtain analytic results). It is important to note that risk aversion cannot be increased past 3.7 because the model no longer has a solution.

As with the DY model, we checked that the moments generated by our solution to the model match those reported by Wachter. We confirm that our results are highly similar, in particular for her tables 2 and 3, though not identical since we use slightly different risk aversion and a different disaster distribution.

A.4.3 Disasters with time-varying recovery (Gabaix)

$$\Delta c_t = \mu_{\Delta c} + \varepsilon_{\Delta c,t} + J_{\Delta c,t} \quad (\text{A.22})$$

$$L_t = (1 - \rho_L) \bar{L} + \rho_L L_{t-1} + \varepsilon_{L,t} \quad (\text{A.23})$$


$$\Delta d_t = \lambda \varepsilon_{\Delta c,t} - L \times 1 \{J_{\Delta c,t} \neq 0\} \quad (\text{A.24})$$

The model is calibrated at the monthly frequency. Conditional on a disaster occurring, $J_{\Delta c,t} \sim N(-0.3, 0.15^2)$. The probability of a disaster in any period is 0.01/12. The other parameters are calibrated as:

Parameter	Value
$\mu_{\Delta c}$	0.1/12
$\text{stdev}(\varepsilon_{\Delta c})$	$0.02/\sqrt{12}$
ρ_L	$0.87^{1/12}$
\bar{L}	0.5
$\text{stdev}(\varepsilon_L)$	0.035
λ	5


Agents have power utility with a coefficient of relative risk aversion of 7 and a time discount factor of $0.96^{1/12}$.

Figure A.1: Quotes vs. transaction prices



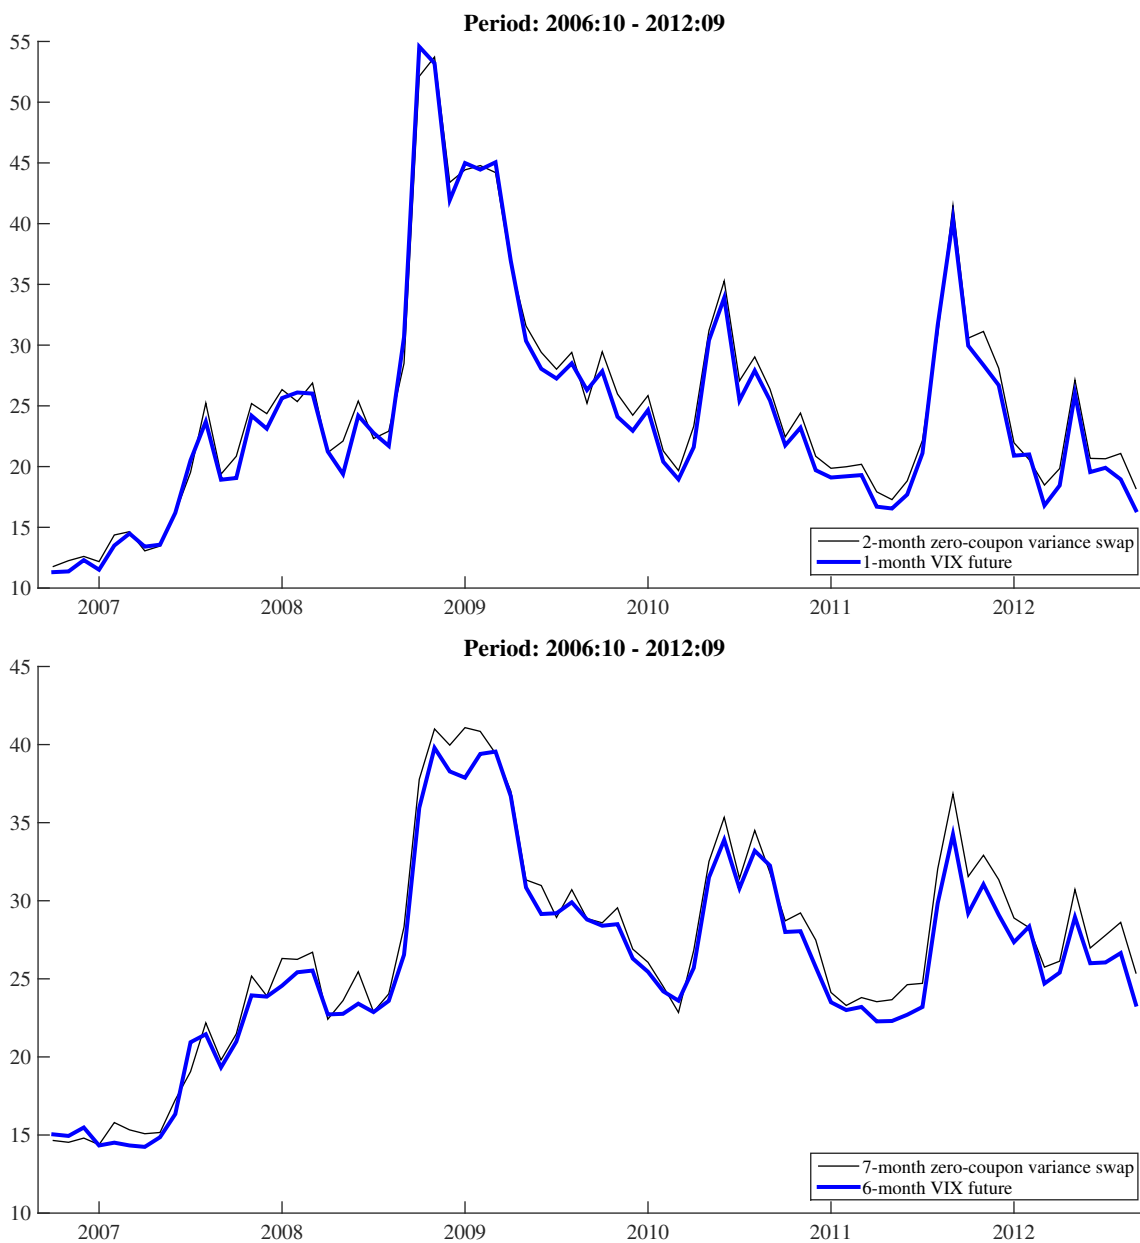
Note: The figure shows the distribution of percentage difference between variance swap price quotes and actual transaction prices, computed as $(\text{transaction price} - \text{quote})/\text{quote}$. The quotes are our main sample, while transaction prices are obtained from the DTCC and begin in 2013. Each panel shows the histogram for a different bucket of maturity of the variance swap contracts.

Figure A.2: International zero-coupons variance claim prices from options and variance swaps




Note: In each panel (corresponding to STOXX 50, FTSE 100 and DAX), the solid line plots average prices of zero-coupon variance claims calculated using the formula for the VIX index and data on international option prices from Optionmetrics. The dotted line plots the average prices of the same claims computed from international variance swaps. The sample covers the period 2013:4-2014:2.

Figure A.3: VIX futures vs. zero-coupon variance swap prices




Note: the top panel plots the time series of the 2-month zero-coupon variance swap and the 1-month VIX future price from the CME, in annualized volatility terms. The bottom panel plots the time series of the 7-month zero-coupon variance swap and the 6-month VIX future price from the CME, in annualized volatility terms. The sample covers the period 2006:10-2012:9.

Figure A.4: Impulse response functions of level, slope and RV



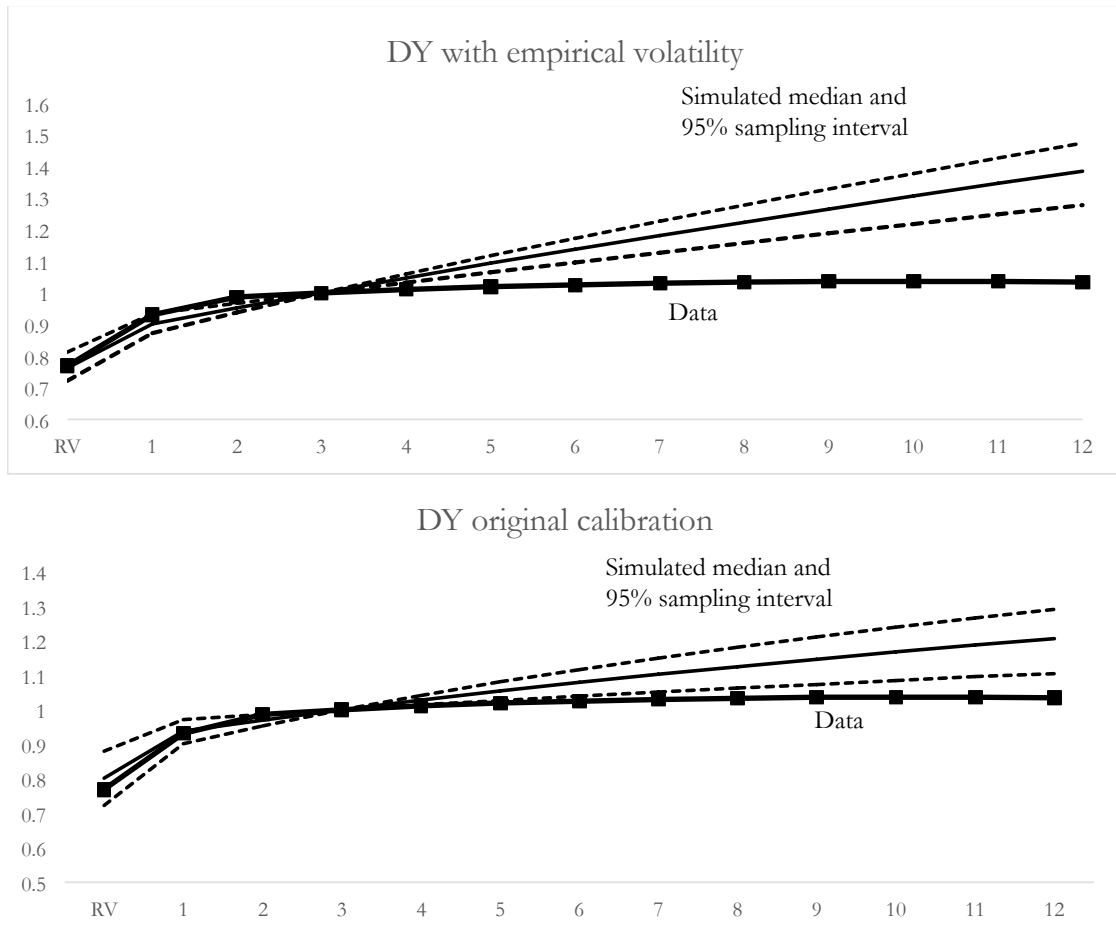
Note: Each panel plots the response of one of the variables in the VAR (level, slope, and RV) to one of the three orthogonalized shocks. The shocks are orthogonalized with a Cholesky factorization with the ordering level-slope-RV.

Figure A.5: Sharpe ratios and average term structure in the long-run risk model



Note: The top panel gives the population Sharpe ratios from the two models and the sample values from the data. The bottom panel plots population means of the prices of zero-coupon claims. All the curves are normalized to have the same value for the realized variance.

Figure A.6: Slope of the term structure in the long-run risk model



Note: Simulated zero-coupon variance claim prices from structural models. All curves are normalized to equal 1 at three months. The dotted lines are 95-percent sampling intervals from the simulations.

Table A.1: Predictive regressions

Horizon (months)	Predictor: Slope and Curvature		Predictor: $Z_t^n - Z_t^0$					
	Dep var: $RV_{t+n} - RV_t$	R^2	Dep var: $RV_{t+n} - RV_t$	R^2	coeff.	Dep var: $-\sum_{j=0}^{n-1} r_{t+1+j}^{n-j}$	R^2	coeff.
1	0.08		0.17	0.63***	(0.22)	0.07	0.37*	(0.22)
2	0.22		0.18	0.71***	(0.17)	0.04	0.29*	(0.17)
3	0.31		0.24	0.83***	(0.14)	0.01	0.17	(0.14)
6	0.36		0.34	0.95***	(0.09)	0.001	0.05	(0.09)
12	0.38		0.41	1.10***	(0.09)	0.001	-0.10	(0.09)

Note: Results of regressions forecasting changes in realized variance. The left side reports the R^2 of a regression of changes in realized volatility between month t and month $t+n$ on the level and the slope at time t . The right side reports the coefficients of univariate regressions of changes in realized volatility (left column) and returns to volatility claims from t to $t+n$ (right column) on the difference between the zero-coupon prices of maturity n (Z_t^n) and realized volatility (Z_t^0) at time t . * indicates significance at the 10-percent level, ** the 5-percent level, and *** the 1-percent level.

Table A.2: Risk neutral parameters

Specifications		1996:2007			1996:2013		
		ρ_s^Q	ρ_l^Q	v_l^Q	ρ_s^Q	ρ_l^Q	v_l^Q
CV	Est.	0.66***	0.99***	0.65***	0.83***	0.99***	1.01**
	Stderr.	0.06	0.00	0.14	0.07	0.01	0.48
CIR	Est.	0.68***	0.99***	0.65***	0.84***	0.99***	0.99**
	Stderr.	0.05	0.00	0.14	0.06	0.01	0.44
FLEX	Est.	0.66***	0.99***	0.64***	0.83***	0.99***	1.01**
	Stderr.	0.06	0.00	0.14	0.07	0.01	0.47

Note: The table reports the risk-neutral estimated dynamics of the term structure model, for the three specifications CV, CIR and FLEX, and separately for the full sample (1996-2013) and the pre-crisis sample (1996-2007).

Table A.3: Time series parameter estimates for the CV specification

Sample		K_0	K_1			Σ^*		
1996:2007	Est.	6.49***	0.65***	0.14**	0.03	15.42***	0	0
		3.36***	0.08	0.83***	0.02	3.87***	5.25***	0
		1.98	0.56***	0	0.12	10.57***	-2.62	14.08***
	Std. err.	2.42	0.08	0.07	0.12	3.23	-	-
		1.12	0.10	0.08	0.07	0.73	0.73	-
		2.11	0.07	-	0.10	2.36	2.63	2.07
1996:2013	Est.	8.99***	0.56***	0.08	0.19**	20.59***	0	0
		2.43***	0	0.91***	0.06***	3.26	7.04***	0
		5.67*	0.20	0	0.57***	30.00**	-2.14	25.49***
	Std. err.	2.48	0.13	0.05	0.09	5.07	-	-
		0.90	-	0.03	0.01	2.09	0.97	-
		2.92	0.13	-	0.15	13.38	2.94	4.23

Note: The table reports the time-series parameter estimates for the CV specification. Σ^* is the lower triangular Cholesky decomposition of Σ_0 . For admissibility, K_0 and K_1 are constrained to be non-negative. Those entries for which the non-negativity constraint is binding are set to zero and thus standard errors are not provided. The table reports the estimates separately for the full sample (1996-2013) and the pre-crisis sample (1996-2007).

Table A.4: Time series parameter estimates for the CIR specification

Sample		K_0	K_1			Σ^*		
1996:2007	Est.	0	0.66***	0.32***	0	5.16***	0	3.53***
		0.65***	0	0.99***	0	0	0.98***	0
		0	0.69***	0	0	3.53***	0	6.74***
	Std. err.	-	0.07	0.05	-	1.82	-	1.18
		0.14	-	0.00	-	-	0.23	-
		-	0.01	-	-	1.18	-	1.69
1996:2013	Est.	0	0.82***	0.16**	0	7.17**	0	9.42*
		0.99**	0	0.98***	0	0	1.42**	0
		0	0.75***	0	0	9.42*	0	23.56*
	Std. err.	-	0.04	0.06	-	2.84	-	5.42
		0.44	-	0.00	-	-	0.58	-
		-	0.09	-	-	5.42	-	12.88

Note: The table reports the time-series parameter estimates for the CIR specification. The diagonal elements of Σ^* correspond to the variance parameters σ_s^2 , σ_l^2 , and σ_{RV}^2 . The (3,1) entry of Σ^* correspond to the covariance parameter $\sigma_{s,RV}$. The table reports the estimates separately for the full sample (1996-2013) and the pre-crisis sample (1996-2007).

Table A.5: Time series parameter estimates for the FLEX specification

Sample		K_0	K_1			Σ^*		
1996:2007	Est.	2.31	0.74***	0.19**	0.01	2.26***	0	0
		2.18***	0.09*	0.87***	0	0.58***	0.78***	0
		0.11	0.60***	0	0.13	1.56***	-0.24	2.04***
	Std. err.	1.49	0.08	0.08	0.08	0.42	-	-
		0.81	0.05	0.06	-	0.08	0.08	-
		0.75	0.08	-	0.09	0.29	0.31	0.27
1996:2013	Est.	4.54***	0.68***	0.08**	0.19**	2.59***	0	0
		1.37**	0	0.93***	0.07***	0.44*	1.09***	0
		0.67***	0.36***	0	0.51**	3.32***	-0.34	3.20***
	Std. err.	1.30	0.08	0.04	0.08	0.48	-	-
		0.55	-	0.02	0.02	0.23	0.12	-
		0.05	0.11	-	0.20	1.24	0.43	0.48

Note: The table reports the time-series parameter estimates for the FLEX specification. Σ^* is the lower triangular Cholesky decomposition of Σ_0 . For admissibility, K_0 and K_1 are constrained to be non-negative. Those entries for which the non-negativity constraint is binding are set to zero and thus standard errors are not provided. The table reports the estimates separately for the full sample (1996-2013) and the pre-crisis sample (1996-2007).

Table A.6: Annualized steady state risk prices, all specifications

Specifications	Sources of risks	1996:2007		1996:2013	
		Estimates	Standard Error	Estimates	Standard Error
CV	s_t^2 -risk	-0.23	0.22	-0.08	0.14
	l_t^2 -risk	0.05	0.18	-0.21	0.14
	RV-risk	-2.78***	0.65	-1.44***	0.43
CIR	s_t^2 -risk	-0.18	0.36	-0.11	0.31
	l_t^2 -risk	0.04	0.14	-0.18	0.12
	RV-risk	-3.92***	0.91	-1.70***	0.48
FLEX	s_t^2 -risk	-0.23	0.30	-0.14	0.16
	l_t^2 -risk	0.06	0.22	-0.17	0.19
	RV-risk	-3.17***	0.78	-1.69**	0.73

Note: Estimates of steady-state risk prices from the no-arbitrage model. Risk prices are annualized by multiplying by $\sqrt{12}$. *** denotes significance at the 1-percent level. Results are reported for the full sample (1996-2013), and restricted to 1996-2007.