

Contractionary volatility or volatile contractions?

David Berger, Ian Dew-Becker, and Stefano Giglio*

February 7, 2016

Abstract

There is substantial evidence that the volatility of the economy is countercyclical. This paper provides new empirical evidence on the relationship between aggregate volatility and the macroeconomy. We aim to test whether (1) increases in uncertainty about the future cause recessions or (2) recessions are periods of high volatility. We measure volatility expectations using market-implied forecasts of future stock return volatility. According to both simple cross-correlations and a standard VAR, shocks to *realized* volatility are contractionary, while shocks to *expected* volatility in the future are, if anything, expansionary. Furthermore, investors have historically paid large premia to hedge shocks to realized volatility, but the premia associated with shocks to volatility expectations are not statistically different from zero. We argue that these facts are inconsistent with models in which increases in expected future volatility cause contractions, but they are in line with the predictions of a simple model in which aggregate technology shocks are negatively skewed. The volatile contractions hypothesis is also consistent with evidence that equity returns and real activity are negatively skewed.

1 Introduction

Volatility in financial markets and the real economy appears to be countercyclical, and a large body of recent macroeconomic research explores the effects of shocks to volatility and uncertainty on the economy.¹ The theoretical literature has focused on mechanisms through which shocks to volatility drive business-cycle fluctuations. However, it is also possible that there is no causal relationship, and instead, volatility is simply high when other shocks to the economy are negative. The simplest example of that scenario is when shocks are negatively skewed: then large realized values tend to be negative ones, so realized volatility – driven by extreme observations – tends to be high when shocks are negative.

*Northwestern University, Northwestern University and the University of Chicago. We appreciate helpful comments from Larry Christiano and seminar participants at the CITE conference.

¹Gilchrist, Sim, and Zakrajsek (2014) use the same fact as a starting point for an analysis of volatility, irreversible investment, and financial frictions. See Campbell et al. (2001) (equity volatility at the index, industry, and firm level is countercyclical); Storesletten, Telmer, and Yaron (2004) and Guvenen, Ozkan, and Song (2014) (household income risk is countercyclical); Eisfeldt and Rampini (2006) (dispersion in industry TFP growth rates is countercyclical); Alexopoulos and Cohen (2009) and Bloom, Baker, and Davis (2015) (news sources use uncertainty-related language countercyclically); among many others, some of which are discussed below.

This paper provides novel evidence on the question of whether expectations of future volatility cause downturns, or whether, instead, volatility is simply high in bad times. That is, does volatility lead to contractions, or are contractions simply volatile? The key distinction that we draw is between *realized* and *expected* volatility. When we refer to volatility in this paper, what we mean is dispersion in the continuous flow of shocks that hits the economy. Models are typically clear: in some, increases in the variance of agents’ subjective distributions over future outcomes induce recessions; in others, recessions are periods when volatility is high. In the data, though, realized and expected volatility are not easy to disentangle. The central empirical exercise in Bloom (2009) highlights that difficulty. That paper estimates a vector autoregression (VAR) where “uncertainty” is measured as a single time series that splices together option-implied variances since 1986 with the history of actual squared returns on the S&P 100 in the period prior to the availability of the option-implied series. The option-implied variance measures investor expectations of future volatility, while squared returns measure realized volatility.² The question we ask is whether those two factors have the same effect on the economy, using novel data to distinguish between realized and expected future volatility.

We examine a sample of S&P 500 index options that has been little studied in the literature, but that allows us to measure volatility expectations since 1983 and horizons of up to 6 months. We thus have a sample spanning 31 years, three recessions, and a wide range of financial and real conditions in the economy. We measure realized volatility as the sum of daily squared returns on the market. This concept differs from that used by some other recent papers, but has a number of advantages, including directly measuring the concept of interest (the realized dispersion in aggregate shocks) and being linked in a well-defined way to our volatility forecasting variable. While there are many different potential measures of volatility and uncertainty in the economy, if one wants to cleanly distinguish between expectations and realizations, equity indexes are the only source that is available.

Using both cross-correlations and VARs, we find that increases in realized volatility are associated with declines in output, consumption, investment, and employment, consistent with findings in Bloom (2009) and Basu and Bundick (2015). More surprisingly, though, we find that increases in six-month *expected* volatility are, if anything, associated with *expansions* in activity. In other words, increases in uncertainty about future stock returns do not appear to reduce output, even though realized volatility tends to be high in bad times. The results are consistent across estimation methods and they are robust to changes in the ordering of variables in the VAR, changing the data sample and state variables, choices about detrending, and estimating the model at the quarterly or monthly frequency.

The simplest theoretical explanation for our results is that fluctuations in economic activity are negatively skewed. That could either be because the fundamental shocks are skewed, or because symmetrical shocks are transmitted to the economy asymmetrically (e.g. because constraints, such as financial frictions, bind more tightly in bad times). Skewness literally says that the squared

²Basu and Bundick (2015) make a similar observation.

value of a variable is correlated with the variable itself, which is essentially what we find when realized volatility is associated with contractions. There are two important pieces of evidence in favor of this hypothesis. First, changes in a wide variety of measures of real activity are negatively skewed, as are stock returns. Second, the observed asset prices imply that investors have paid large premia for insurance against high realized volatility and extreme negative stock returns (known as the variance risk premium and the option skew or put premium, respectively) in the last 30 years, whereas the premium paid for protection against increases in *expected* volatility has been near zero or even positive.

Two important caveats apply to our results. First, our analysis is of the effects of fluctuations in *aggregate* uncertainty. We do not measure variation in cross-sectional uncertainty. There are obviously many dimensions along which uncertainty can vary, and we try to understand one here. Second, there are other measures of aggregate uncertainty that may still cause contractions even after controlling for realized volatility. The key contribution of this paper is to show that when a single concept of volatility can be split into components coming from realizations and expectations, it is the realization component that appears to drive the results. That does not imply, though, that no other measures of uncertainty (which do not distinguish between expectations and realizations) can affect the economy.

The remainder of the paper is organized as follows. Section 1.1 reviews the extensive literature on the interaction of volatility and the real economy. Sections 2 and 3 describe our data and examines its basic characteristics. Section 4 describes our main analysis of the relationship between realized volatility, expected volatility, and the real economy. Finally, section 5 examines implications of the contractionary volatility and volatile contractions hypothesis and argues that our evidence is consistent with the view that economic shocks are negatively skewed and that contractions are inherently risky. Section 6 concludes.

1.1 Literature review

There is now a broad array of models that describe possible mechanisms through which volatility could affect the economy. The recent literature tends to cite Bloom (2009), who examined the effects of volatility empirically in a VAR and theoretically in an Ss model with fixed costs of adjustment. Bloom et al. (2014) extend Bloom (2009) to a general-equilibrium setting with shocks to both aggregate and idiosyncratic volatility. Notably, their shocks to volatility are correlated with shocks to the level of productivity itself, which can induce the skewness that we observe.

Gourio (2012, 2013) examines the effects of time-varying disaster risk. Basu and Bundick (2012, 2015) and Liu and Leduc (2015) examine the effects of volatility shocks in New Keynesian DSGE models.

There is also an empirical literature in macroeconomics providing evidence on the effects of volatility shocks. Bloom (2009) and Bloom, Baker, and Davis (2015) examine vector autoregressions using the VIX and a measure of policy uncertainty. Alexopoulos and Cohen (2009) study a similar

measure. Gilchrist, Sim, and Zakrajsek (2014) provide evidence from a VAR and a calibrated DSGE model. Stock and Watson (2012) examine a dynamic factor model and conclude that the recession of 2008–2009 was driven largely by shocks to uncertainty (measured as in Bloom (2009) and Bloom, Baker, and Davis (2015)).

Most recently, Jurado, Ludvigson, and Ng (2015) measure uncertainty based on forecast errors for a wide range of macroeconomic and financial factors and find that uncertainty shocks are related to recessions. Ludvigson, Ma, and Ng (2015), like us, distinguish between different types of uncertainty. They show that variation in uncertainty about macro variables is largely an endogenous response to business cycles, whereas shocks to financial uncertainty cause recessions. The key distinction between our work and theirs is that we focus on the distinction between uncertainty expectations and realizations, while they distinguish between uncertainty in different sectors of the economy.

Our empirical analysis uses option prices to infer investor expectations of future volatility in the economy. We therefore draw from and build on a large literature in finance examining the pricing and dynamics of volatility. There is a long history of studies of fluctuations in aggregate (i.e. stock index) volatility, including, among many others, Adrian and Rosenberg (2008), Bollerslev et al. (2009), Heston (1993), Ang et al. (2006), Carr and Wu (2009), Bakshi and Kapadia (2003), Egloff, Leippold, and Wu (2010), and Ait-Sahalia, Karaman, and Mancini (2013) (see Dew-Becker et al. (2014) for a review).

2 Data

The past literature has examined a number of different measures of risk and uncertainty in the economy. There is work that examines volatility in productivity and household income, text-based measures of uncertainty, and many other concepts of risk. The concept of volatility studied in this paper is aggregate equity return volatility. We focus on stock market volatility for a number of reasons.

The feature of the data that we desire to measure is the variance of the flow of aggregate shocks that hit the economy. We thus do not aim to measure cross-sectional dispersion in shocks or even forecast uncertainty. Jurado, Ludvigson, and Ng (2015), for example, construct a monthly measure of forecast uncertainty for a wide range of macroeconomic variables. Our goal, on the other hand, is to measure the variance of the *common* shocks to measures of activity, rather than the total dispersion of each measure.

In principle, realized volatility can be measured with data sampled at many different frequencies, and much past work has focused on monthly or quarterly data. We use data on daily stock returns for two reasons. First having more observations allows one to measure the dispersion of shocks more accurately (a standard deviation can be measured with only one observation (assuming zero mean), but with very little accuracy). Second, by using high-frequency data, we eliminate the need to estimate the mean of the distribution. With sufficiently high-frequency data, the observed

squared values reveal the variance, with the conditional mean having no effects.³

So stock market volatility both measures what we are interested in – the variance of the flow of shocks – and also does so accurately. We are not the only ones to study stock market volatility, and for good reason. Equity prices summarize information about the future path of the economy, so volatility in the economy should be expected to be related to volatility in the stock market. One would expect that almost any factor that affects risk in the economy would affect the riskiness of firms, since the revenue and profitability of firms ultimately depend on all the features of the economy.

In almost any conceivable model where the volatility of aggregate shocks to the economy fluctuates over time, that variation in volatility would also pass through to firm equity returns. For example, in standard investment theories, stock prices are closely related to the discounted present value of the marginal product of capital (in q theory, that link is exact). So volatility in stock prices measures volatility in the stream of future marginal products that determines incentives to invest (as opposed to, for example, uncertainty about just a single month of profits or income).

That said, the stock market is obviously not the only way to measure volatility, and there are certainly aspects of the economy that it will not capture. But while alternative volatility concepts, like survey and text-based measures (like the Michigan survey or the measure of Baker, Bloom, and Davis (2015)), are important for providing insights into how consumers feel about the future, they are sometimes difficult to interpret quantitatively. Stock market volatility, on the other hand, is clearly defined, and has a direct link to economic activity.

Finally, by measuring stock market volatility, we are able to draw a clear distinction between expectations and realizations of volatility. None of the other measures of uncertainty that we are aware of has such a feature. Some sources clearly measure volatility realizations, some clearly measure expectations, and some do both. None, though, yield both expectations and realizations of volatility for the *same* underlying variable. And not only can we construct expectations, but those expectations are available at a range of different maturities. News shocks – i.e. shocks to expectations of future rather than current values of state variables – play an important role in many recent macroeconomic studies (e.g. Beaudry and Portier (2006); Alexopoulos (2011); Christiano, Motto, and Rostagno (2014)). Since we have measures of expectations at multiple horizons, we are able to measure these news shocks for volatility directly.

The most important drawback of our data is that it does not measure variation in purely idiosyncratic risk, which plays an important role in some recent models. This paper is not able to shed light on that class of models, but in future work the methods we study here may be able to be extended to examine fluctuations in idiosyncratic risk.

³Specifically, suppose that there is some fundamental unit of time, dt , indexed by t , and shocks are drawn in each period. In each period, shocks are independently drawn from a distribution with mean μ_{dt} and standard deviation σ_{dt} . The expected squared value of a variable is equal to the mean squared plus the variance. The mean squared is proportional to dt^2 , while the variance is proportional to dt . So if we measure the realized variance of a series as the sum of squared realizations using data sampled at a sufficiently high rate – with dt small – then the conditional mean of the series eventually has no impact on the realized variance.

2.1 Model-free implied volatility (MFIV)

Equity return volatility is measured in this paper as the volatility of returns on the S&P 500 index. In particular, realized volatility in each month is measured as the sum of squared daily returns. We denote realized volatility in month t as RV_t . We obtain data on daily stock returns from the CRSP database, which has coverage since 1926.

We construct measures of expected future volatility using option prices. The prices of S&P 500 options are obtained from the Chicago Mercantile Exchange (CME), with maturities from 1 to 6 months since 1983.⁴

The most widely used measure of expected volatility is the VIX. The VIX is an index, based on option prices, of implied volatility 30 days ahead. There are two main issues with using the VIX as a measure of expected future volatility in our empirical analysis. First, since realized volatility is persistent, one-month expectations are very highly correlated with realizations. But at longer horizons, we observe larger differences between realized and expected volatility, allowing for their effects to be separately identified. Moreover, the very short maturity of the VIX makes it less than ideal to capture the kind of uncertainty about the future that might matter to firms and consumers, whose horizons are typically much longer (it is possible to calculate a one-day implied volatility, but it is hard to believe it would be relevant for typical economic decisions).

Those problems with the VIX are solved by examining implied volatilities at longer maturities. From a theoretical point of view, the expectation of future volatility 6 or 12 months ahead lines up more closely with the horizons that matter for firms' decisions in macroeconomic models. In addition, as we will argue below, the correlation of shocks to expectations of future volatility with shocks to realized volatility are low enough that they allow separate identification of the effects of the two shocks.

In this section, we describe how we construct a version of the VIX (the price of a claim to future realized volatility), extending it beyond the usual one-month horizon. We will refer to the extension of the VIX to arbitrary maturities as *model-free implied volatility*, or $MFIV_n$, where n refers to the horizon in months (so that the VIX is $MFIV_1$).

The VIX and MFIV are claims to future realized volatility. The fundamental time period that we analyze in this paper is a single month, which is indexed by t . Within any month, the realized volatility of equity returns is calculated as

$$RV_t \equiv \left(\sum_{days \in t} r_i^2 \right)^{1/2} \quad (1)$$

where r_i is the log return on the S&P 500 on day i in month t .

There is a large literature in finance that examines model-free implied variances for various

⁴Maturities above 6 months are also available at least for some time periods, but the scarcity of long-maturity options with different strike prices makes it impossible to obtain a continuous and reliable time series for such maturities. Therefore, we focus here on maturities up to 6 months.

assets. Under very general conditions (Jiang and Tian (2005) and Carr and Wu (2009)), the price of a claim to realized volatility between dates $t + 1$ and $t + n$ can be written as a function of time- t option prices:

$$MFIV_{t,n} = \left(2 \int_0^\infty \frac{O_t(n, K)}{B_t(n) K^2} dK \right)^{1/2} \approx \left(E_t \left[\sum_{m=1}^n \frac{M_{t,t+m} RV_{t+m}^2}{E_t M_{t,t+m}} \right] \right)^{1/2} \quad (2)$$

where $M_{t,t+m}$ represents the stochastic discount factor between dates t and $t+m$. The $1/2$ exponent in all these formulas implies that all the quantities we look at are in volatility (not variance) terms, the same units in which the VIX is expressed.

Model-free implied volatility (including the VIX) is calculated as an integral over option prices, where K denotes strikes, $O_t(n, K)$ is the price of an out-of-the-money option with strike K and maturity n , and $B_t(n)$ is the price at time t of a bond paying one dollar at time $t + n$. $MFIV_{t,n}$ is approximately equal to total expected volatility over the next n months (to be most precise, MFIV is the square root of expected variance; RV is squared in the calculation since variances are additive over time, and RV^2 is realized variance). The approximate equality is due to the discretization in the calculation of realized variance.

The MFIV makes extremely minimal assumptions about the dynamics of stock returns, which is why $MFIV_{t,1}$ is used by the CBOE as its definition of the VIX. Crucially, unlike the Black–Scholes (1973) implied volatility, it does not require that volatilities be constant over time.

$MFIV_{t,n}$ represents the expectation of realized volatility between times t and $t + n$. To isolate the expectation of future volatility at different horizons, we can also look at the prices of *volatility forwards*, claims to future volatility in one specific month in the future. Forwards allow us to measure news shocks, in the sense that they isolate expectations about volatility some number of months in the future, as opposed to including volatility both in the next month and at longer horizons.

Because of the additivity of variances, the price of volatility forwards can be constructed directly using the set of MFIVs of different maturity. Specifically, the price of a forward claim to variance n periods from t is

$$F_{t,n} \equiv (MFIV_{t,n}^2 - MFIV_{t,n-1}^2)^{1/2} \approx E_t \left[\frac{M_{t,t+n} RV_{t+n}^2}{E_t M_{t,t+n}} \right]^{1/2} \quad (3)$$

Expectations involving the term $\frac{M_{t,t+n}}{E_t M_{t,t+n}}$ are often referred to as *risk-neutral*, or *risk-adjusted expectations*,

$$E_t^Q [X_{t+1}] \equiv E_t \left[\frac{M_{t,t+1}}{E_t M_{t,t+1}} X_{t+1} \right] \quad (4)$$

The MFIV (and hence the VIX), as well as the volatility forward prices $F_{t,n}$, are risk-neutral expectations of future realized volatility. A risk-neutral expectation depends on both the physical expectation of future volatility and also any risk adjustment due to covariation with the pricing

kernel (marginal utility). While we generally expect risk-neutral expectations to be biased on average (for volatility, risk premia will tend to push the price above the physical expectation), to the extent that risk premia are stable, changes over time in risk-neutral expectations are equivalent to changes in physical expectations. That is why the past literature on the effects of uncertainty shocks has used the VIX (a risk-neutral expectation) as a measure of expected future volatility.

We construct the MFIV for the S&P 500 at maturities between 1 and 6 months. We calculate it using data on option prices from the Chicago Mercantile Exchange (CME), which allows us to construct a time series dating back to 1983. Throughout the paper and appendix, we examine a range of robustness tests, including constructing volatility expectations using alternative sources of option prices and using variance swaps, whose payoffs are directly linked to realized volatility (but for which the sample is only half as long). Computing the MFIV with real-world data requires several steps; the appendix provides an extensive description of our calculation methods and analyzes the data to confirm its accuracy.

Finally, in the remainder of the paper we often look at the logs of realized variance (RV) and forwards (F). Given the skewed nature of realized variance, looking at relations in logs makes the results less dependent on the occasional volatility spikes and ensures that these extreme observations do not drive our results. We have confirmed all of the results both in logs and in levels. We refer to variables in logs everywhere with lower-case letters, e.g. $rv_t = \log RV_t$.

2.2 The time series of realized volatility and its expectations

Figure 1 plots the history of (annualized) realized volatility along with market expectations, $F_{t,n}$, at horizons of 1, 3, and 6 months. Both realized volatility and volatility expectations vary considerably over the sample. The two most notable jumps in volatility are the financial crisis and the 1987 market crash, which both involved realized volatility above 60 annualized percentage points and rises of the 6-month forward volatility above 40 percent. At lower frequencies, the periods 1997–2003 and 2008–2012 are associated with persistently high volatility expectations, while expectations are lower in other periods, especially the early 1980’s, early 1990’s, and mid-2000’s. There are also distinct spikes in expected volatility in the summers of 2010 and 2011 due to concerns about the stability of the Euro and the willingness of the United States to continue to pay its debts.

Panel A of Table 1 reports descriptive statistics for the series in figure 1. The means increase with the horizon, which is due to the risk-adjustment mentioned above. Specifically, there is a negative risk premium on volatility (Coval and Shumway (2001) and Dew-Becker et al. (2016)), which causes the prices of variance forwards to be upward biased estimates of future volatility. As we would expect, the standard deviations of the expectations decline with horizon, implying that investors have more information about volatility over the next 30 days than they do about volatility 6 months in the future.

The various series are highly correlated; Panel B of Table 1 reports all pairwise correlations. While correlations are all high, they fall substantially as the horizon lengthens. In particular, the

correlation between RV_t and expected volatility 6 months in the future, F_6 , is 0.68, suggesting that it should be possible to separately identify the effects of variation in RV_t and expected future volatility (in logs the results are similar: the correlation between rv_t and $f_{t,6}$ is 0.71).

The innovations in the various series are even less correlated than the levels. The residuals from a VAR(1) in aggregate realized volatility RV and the 6-month expectation F_6 are only 55 percent correlated (54 percent in logs), implying that shocks to realized volatility explain only one third of the variance of the shocks to volatility expectations. The other two thirds is entirely independent (and similarly, expected volatility shocks explain only 1/3 of the variance of the shocks to realized volatility).

Finally, table 1 (panel C) reports raw correlations of realized and 6-month expected log volatility with measures of real economic activity – capacity utilization, the unemployment rate, and returns on the S&P 500 (correlations are similar in levels). Volatility does appear somewhat correlated with these macroeconomic variables; what the rest of this paper will explore is whether this association is due to simple correlation with realized volatility, or whether expected future volatility is causally linked to macroeconomic activity.

3 The dynamics of volatility expectations

We now examine the dynamics of realized volatility and its market expectations more formally. We first run simple forecasting regressions to measure how well realized volatility is predicted by lagged market expectations. It is obviously critical that we show that MFIV actually forecasts volatility, otherwise we cannot claim to identify volatility news. Next, we estimate a variance decomposition measuring the fraction of the variation in realized volatility that is anticipated. Finally, we compare market-based forecasts with those of econometric models. In all of this section, we will run our regressions using logs of realized and expected volatility; all of the results are essentially unchanged if we perform our analysis in levels.

3.1 Forecasting regressions for realized volatility

Under the assumption that risk premia are constant, the 6-month volatility forward, $f_{6,t}$, should correspond to the expectation of realized volatility 6 months ahead, rv_{t+6} , plus a constant term. If risk premia are time-varying, $f_{6,t}$ may not be entirely driven by expectations of future volatility, but we should still expect it to have forecasting power for future rv even after controlling for other predictors like lagged rv .

The top panel of Figure 2 plots the coefficient β_h in the univariate regression

$$rv_t = \alpha + \beta_h f_{t-h,6} + \varepsilon_t \quad (5)$$

for different lags h (in the x axis). The figure shows that market-implied expected volatility indeed has highly significant forecasting power for future realized volatility, with a statistically and eco-

nomically significant coefficient even 12 months ahead. The coefficients are well below 1, though, which implies that $f_{t,6}$ is not a statistical expectation of future realized volatility.

Since rv is a persistent process, a natural question is whether f_6 contains any information about future rv after controlling for rv itself. We therefore estimate the multivariate specification,

$$rv_t = \alpha + \beta_h f_{t-h,6} + \gamma_h rv_{t-h} + \varepsilon_t \quad (6)$$

The second row of Figure 2 shows the two sets of coefficients, β_h and γ_h , for different lags h . The left panel shows the coefficient on lagged rv , γ_h . As expected, lagged rv forecasts future rv , with a coefficient declining with the horizon. More interestingly, though, the right panel (that reports the coefficients on lagged f_6) shows that f_6 has significant predictive power for future volatility at all horizons, even after controlling for lagged rv . In fact, this coefficient does not decline substantially with maturity.

Overall, then, at short horizons, current rv is a better predictor of future rv than the 6-month market expectation. However, as the horizon increases, the forecasting power of rv diminishes quickly, whereas the 6-month forward maintains a stable coefficient. f_6 thus seems to predict a more persistent component of realized volatility that persists even after short-run variation dissipates.

3.2 Decomposing the variance of realized volatility

The results above show that model-free implied volatility can forecast future realized volatility. The next natural question is how much news there actually is about future volatility, and to what extent fluctuations in realized volatility are surprises.

Exploiting properties of expectations, the variance of rv_t can be decomposed into three components,

$$V(rv_t) = V(rv_t - E_{t-1}[rv_t]) + V(E_{t-1}[rv_t] - E_{t-6}[rv_t]) + V(E_{t-6}[rv_t]) \quad (7)$$

where $V(X)$ denotes the unconditional variance of a variable X . The first term is the variance of the surprise in rv conditional on information available in the previous period. The second component is the news about rv_t that occurs between months $t-6$ and $t-1$. That is, it is the variance of the innovations in the expectation of rv_t over those months. It thus measures how much investors learn about rv_t on average in the five months before it is realized. Finally, $V(E_{t-6}[rv_t])$ is the variance of expectations of rv six months ahead.

To implement this decomposition, we must construct expectations of future volatility at 1- and 6-month horizons. We examine a range of methods for forming expectations. First, we construct expectations using the corresponding forward volatility claims. In particular, to account for the possibility that risk premia vary, to measure $E_t[rv_{t+n}]$, we project rv_{t+n} on $f_{t,n}$ (as we did in the previous section).

Second, we model rv_t and the volatility forwards as being driven by a VAR, following Dew-Becker et al. (2014), Egloff et al. (2010), and Ait-Sahalia, Karaman, and Mancini (2013),

$$\begin{bmatrix} rv_t \\ f_{t,1} \\ f_{t,6} \end{bmatrix} = \begin{bmatrix} 0 & b_{rv,\sigma} & 0 \\ 0 & b_{\sigma,\sigma} & b_{\sigma,f} \\ 0 & b_{f,\sigma} & b_{f,f} \end{bmatrix} \begin{bmatrix} rv_{t-1} \\ f_{t-1,1} \\ f_{t-1,6} \end{bmatrix} + \begin{bmatrix} \varepsilon_{rv,t} \\ \varepsilon_{\sigma,t} \\ \varepsilon_{f,t} \end{bmatrix} \quad (8)$$

This model is typical in modeling volatility dynamics as having two persistent factors – measured by $f_{t,1}$ and $f_{t,6}$ – and also allowing for transitory shocks through $\varepsilon_{rv,t}$. The VAR imposes some structure on the dynamics by forcing $f_{t-1,1}$ to determine the conditional expectation of rv_t . The two predictive factors follow an unconstrained VAR(1).

Finally, The last two methods that we examine for forming expectations involve estimating simple univariate ARMA models for rv_t , thus ignoring the volatility forward prices entirely. We estimate an ARMA(1,1) and an ARMA(2,2).

Panel D of Table 1 reports results for the variance decomposition under the various methods. The table shows that, independently of the method used to construct expectations, approximately 45% of the variance of rv is due to the purely unexpected component, $rv_t - E_{t-1}rv_t$. That is, almost half of the entire variance of rv_t is a surprise, whether we form expectations using univariate models or taking into account information from asset prices that directly reflect investor expectations of future equity market variance.

Of the remaining 55 percent of the variance of rv , the vast majority – around 40 percentage points – is due to the news that investors gain between months $t - 6$ and $t - 1$. The variance of $E_{t-6}[rv_t]$ accounts for only 10 to 15 percent of the total variance of rv_t .

These results do not mean that variance is unforecastable; indeed, there is significant evidence (see Campbell et al. 2015) that variance is forecastable even at long horizons. However, the results in table 1 show that when modeling realized variance, it is important to take explicitly into account the fact that almost half of the variation in rv is unpredictable, and only 15% of it is predictable at horizons 6 months or longer. That is, there appears to be a high-frequency term, like a low-order moving average, that must be accounted for. Furthermore, to the extent that there is predictability in volatility, it comes mainly from expectations over horizons of six months or less.

3.3 Physical vs. market-based expectations

Finally, we can compare the six-month forecasts we obtain using f_6 to the six-month forecasts produced by a model for the dynamics of volatility. In particular, we compare the forecast of rv_{t+6} based on $f_{t,6}$ to the forecast from an AR(1) model for rv_t , as well as to the forecast using a set of six variables used in Campbell et al. (2015) to forecast volatility: P/E ratio of the S&P 500, value spread, default spread, term spread, lagged rv and S&P return.

Panel E of Table 1 shows that the correlations between these three forecasts are high, at least 0.7. More interestingly, the model-free implied volatility forecast aligns more closely (correlation of 0.84) with the richer forecast of the six-variable model. Relative to that model, however, it has the advantage of not requiring the estimation of many parameters, since it only uses the forward

as a predictor.

We conclude that the variance forwards we construct aligns well with econometric models of expected volatility based on observables. In the remainder of the paper, we will use $f_{t,6}$ as our main variable to capture the expectations of future volatility six months ahead.

4 Variance shocks and the real economy

We now examine the relationship between shocks to volatility and the real economy.

4.1 Data

Since our data sample is only 32 years long, our main analysis uses monthly data to help improve statistical power. Moreover, because fluctuations in both expected and realized volatility are rather short-lived, using higher-frequency data helps estimate the dynamics of volatility well. We show also that our results hold in quarterly data.

We measure real activity using the Federal Reserve’s measure of industrial production for the manufacturing sector. Employment and hours worked are measured as that of the total private non-farm economy. Inflation is measured using the CPI.

All of the variables except volatility are non-stationary, so we detrend them with a one-sided HP filter with a smoothing parameter of 1.296×10^7 . The appendix shows that the results are similar without the detrending included, or if we also detrend the volatility series.

4.2 Cross-correlations

We begin by examining the raw correlations between measures of real activity and leads and lags of changes in realized and expected volatility. Rather than taking a stand on a causal interpretation of the dynamics relationship between volatility shocks and macroeconomic outcomes, in this section we simply report reduced-form empirical patterns of these variables.

In particular, we estimate a bivariate VAR with rv and f_6 as state variables. We then regress four macroeconomic variables – industrial production, hours worked, employment, and inflation – on 12 leads and lags of the two residuals from the VAR. The residuals are not orthogonalized – the regression coefficients simply represent the conditional covariance of the macroeconomic variables with leads and lags of innovations to rv and f_6 . Specifically, denoting the residuals from the VAR as $\varepsilon_{t,rv}$ and $\varepsilon_{t,f}$, we estimate regressions of the form

$$Y_t = \sum_{j=-12}^{12} (b_{rv,j} \varepsilon_{rv,t-j} + b_{f,j} \varepsilon_{f,t-j}) + \mu_t \quad (9)$$

where $b_{rv,j}$ and $b_{f,j}$ are coefficients and μ_t is a residual. Y_t here denotes any of the four macroeconomic series used. We rescale both the residuals and the macroeconomic variables to have unit standard deviations. Figure 3 reports the coefficients, $b_{rv,j}$ and $b_{f,j}$.

The figure shows a very consistent pattern. For all macroeconomic variables except the CPI, an increase in $\varepsilon_{rv,t}$ is followed by a significant decline in the economy within the next year. Furthermore, $\varepsilon_{rv,t}$ is weakly positively predicted by positive lagged economic conditions. That is, a strong economy is associated with low realized volatility in the past and higher volatility in the future. On the other hand, $\varepsilon_{f,t}$ does not predict declines in activity, and if anything predicts expansions. That is, increases in expected volatility appear to be associated with high future output and employment. The relationship between the innovations and the CPI can be viewed as something of a placebo test. There is little reason to expect that volatility shocks would have substantial effects on inflation, so it is good to see that the estimated coefficients for inflation are all well-estimated zeros.

While these graphs only reveal cross-correlations between different types of volatility shocks and the macroeconomy, and have no direct causal interpretation, they indicate a substantive difference between increases in realized and expected volatility.

4.3 Vector autoregressions

We now examine more standard vector autoregressions (VARs) to measure the impact of shocks to realized volatility (rv) and expected volatility (f_6) on the economy. For all the VARs that we run, we include four lags, as suggested by the Akaike information criterion for our main specification.

Following Bloom (2009), Basu and Bundick (2015), and Leduc and Liu (2015), we identify impulse response functions by ordering volatility first in a Cholesky factorization. Unlike those papers, though, we must also choose the ordering of expected and realized volatility. For our main results, we assume that realized volatility moves first. In other words, what we refer to as an identified “realized volatility shock” is simply the residual from a regression of realized volatility on the lagged variables in the VAR. To the extent that realized volatility is persistent, we would expect it to be associated with an increase in expected future volatility. The “expected volatility shock”, since it is ordered second, is orthogonalized to the realized volatility shock, so it may be interpreted as a *pure news* shock – it is the change in volatility expectations that is unrelated to current volatility.

We do not interpret the Cholesky ordering here as a statement about the timing of the ordering of shocks. Rather, we interpret the volatility expectations shock as what it literally is in an econometric sense: the change in six-month volatility expectations that cannot be explained by news about realized volatility in the same period.⁵ We discuss below the implications for the

⁵Christiano, Motto, and Rostagno (2014) study a model with a data-generating process that gives a formal justification for the ordering we use. Suppose realized volatility, rv_t follows the process

$$rv_t = \phi rv_{t-1} + \varepsilon_{0,t} + \varepsilon_{1,t-1} \quad (10)$$

where $\varepsilon_{1,t}$ is observable to agents on date t . Investors’ expectation of volatility at date $t+1$ is thus $E_t rv_{t+1} = \phi rv_t + \varepsilon_{1,t}$. In period t , the innovation in rv_t is $\varepsilon_{0,t}$, while the innovation in $E_t rv_{t+1}$ is $\phi \varepsilon_{0,t} + \varepsilon_{1,t}$. If those innovations are rotated with a Cholesky factorization in which rv is ordered first, the identified shock to rv is $\varepsilon_{0,t}$ and the identified shock to $E_t rv_{t+1}$ is $\varepsilon_{1,t+1}$. That is how we intuitively understand the method we use.

results of changing the ordering of the VAR.

In figure 4, we plot impulse responses from the VAR that has been run in the previous literature that just includes the *one-month* option-implied volatility.⁶ The shock to volatility has a half-life of approximately 8 months and reduces employment and industrial production by statistically and economically significant amounts. The peak reduction in both following a volatility shock is approximately 0.35 percent. The magnitude of these responses is in line with those obtained by Basu and Bundick (2015), for example, though slightly larger. As noted above, by only looking at one volatility indicator (in this case VXO), it is impossible to distinguish whether the macroeconomic effects are caused by, or related to, the component of realized volatility to which VXO is correlated or the expectations component.

Figure 5 presents our main VAR results. The figure has eight panels. We measure responses to unit standard deviation shocks to both realized and expected volatility (with realized volatility ordered first and expected volatility second). Expected volatility here is measured as the log 6-month volatility forward, f_6 . The responses (in the columns) are for expected and realized volatility, log employment, and log industrial production.

The effect of realized volatility on itself appears to be less persistent than that of the VXO – the IRF falls by half within two months, and by three fourths within 5 months, compared to the eight-month half-life of shocks to the VXO. So, consistent with the results in section 3, realized volatility appears to have a highly transitory component. Naturally, the shock to realized volatility also affects volatility expectations – the dynamics of 6-month volatility expectations appear to line up reasonably well with what is implied by the IRF for realized volatility itself.

As to the real economy, a unit standard deviation increase in realized volatility has strong negative effects on both employment and industrial production, reducing them by 0.4 to 0.5 percent. The effect of realized volatility on the economy actually seems to be more strongly negative than that of the VXO (though this difference is not statistically significant).

The effects of a shock to volatility expectations are much different. First, as we would expect, a shock to volatility expectations forecasts high realized volatility in the future, though not statistically significantly (we attribute the difference between the results here and those in the previous section to the generally smaller statistical power in a VAR with many coefficients to estimate). It also forecasts high volatility expectations for a number of months, and the increase is of a similar magnitude to the increase in expected volatility following a realized volatility shock.

Much more surprisingly, though, increases in volatility expectations are associated with, if anything, *increases* in employment and industrial production. Moreover, the effects are of similar order of magnitude as the contractionary effects of realized volatility, peaking at 0.20 percent. While the responses to the expected volatility shock are never statistically significantly positive, the difference between the responses to expected and realized volatility are in fact significant at the

⁶We use the VXO for this analysis, which is the equivalent of the VIX for the S&P 100 rather than the S&P 500, since it has a longer time series than the VIX; this is why previous literature has focused on the VXO rather than the VIX. In practice, the same results hold with VIX and VXO.

5-percent level.

In order to examine the effects of our two shocks on a wider range of variables, we now examine the results of a quarterly VAR that includes, in addition to the two volatility series, GDP, consumption, investment, hours, the GDP deflator, the M2 money supply, and the Fed Funds rate (using the Wu and Xia (2014) shadow rate when the zero lower bound binds). Figure 6 shows that following an increase in realized volatility, we obtain the same comovement emphasized by Basu and Bundick (2015): output, consumption, investment, and hours worked all decline, all statistically significantly. Similarly, following an increase in expected volatility, those same four variables all appear to increase, though, as in the monthly VAR, those increases are, at best, marginally significant.

To summarize, then, we confirm the usual result that increases in one-month stock market volatility expectations are contractionary when they are included alone. But when the VAR includes both realized and expected volatility, we find that it is the increase in realized volatility that is associated with contractions, while increases in expected volatility are weakly associated with expansions.

4.4 Robustness tests

We examine a range of perturbations of our main specification from figure 5. First, we consider alternative orderings of the variables in the VAR. The effects of the ordering depend ultimately on the correlation matrix of the innovations, which we report below:

	<i>rv</i>	Exp. vol.	Fed Funds	Empl.	IP
<i>rv</i>	1				
Exp. vol.	0.53	1			
Fed Funds	-0.04	-0.06	1		
Empl.	0.05	0.11	0.01	1	
IP	-0.02	0.11	0.02	0.55	1

The shocks to realized and expected volatility are strongly correlated (though far from collinear), so obviously their ordering in the VAR is relevant for the results. However, their innovations are only weakly contemporaneously correlated with those to the other variables, implying that the relative ordering of the financial and macro variables is unlikely to affect the results. We confirm that intuition in the appendix.

Figure 7 reports results from a monthly VAR analogous to that of figure 5 where we reverse the ordering of realized and expected volatility. In figure 7, the shock to expected volatility is the entire innovation to f_6 , while the orthogonalized shock to realized volatility is the residual component uncorrelated with expected volatility. The responses of employment and industrial production to the full expected volatility shock are now slightly negative, though far from statistically significantly so, and the magnitudes are much smaller than those for a realized volatility shocks. This shock

essentially combines the positive shock to expected volatility in figure 5 (orthogonalized to realized volatility) and adds some of the realized volatility shock, thus mixing positive and negative effects.

Importantly, though, figure 7 shows that even after we orthogonalize the realized volatility shock to expected volatility – i.e. if we just look at a change in realized volatility that has no direct effect on expectations, therefore a purely transitory shock to RV – we continue to find a negative effect on the real economy. Note also that the point estimate of the effect is of almost the exact same magnitude as the one we obtained in the main specification in figure 5. The negative relationship between the real economy and realized volatility thus appears to be a robust feature of the data, while the effects of expected future volatility depend on how the innovations are rotated, but are never statistically or economically significantly negative.

Figure A.5 and A.6 in the appendix report a range of additional robustness tests. Figure A.5 shows the response of log employment to the two volatility shocks, and figure A.6 reports the response of log industrial production. In each figure, each row corresponds to a different specification of the model. The left panels report responses to a 1-standard deviation shock to rv , while the right panels report the responses to a 1-standard deviation shock to volatility expectations, f_6 .

The first row in the two figures shows the results obtained without detrending the macroeconomic series, which appears to have little effect on the results from a qualitative and quantitative standpoint. The second row reports a version of the main VAR using quarterly data instead of monthly. The results are again very similar.

The third row re-estimates the VAR using only the period 1988 to 2006, i.e. excluding the two large volatility spikes in 1987 and 2008. The results are again similar to our baseline estimates. The fourth row of the figures orders rv and f_6 last in the VAR, and obtains similar results.

4.5 Alternative measures of volatility

While stock market volatility has been widely studied and is our preferred measure, there are a number of other measures that have been used recently. We focus two important alternatives: the index constructed by Baker, Bloom, and Davis (BBD; 2015) and the Michigan consumer survey’s measure of the fraction of respondents who say they plan on delaying car purchases due to economic uncertainty. Those two measures both potentially measure information that economic agents have that is independent of realized volatility. We examine VARs with those variables in them, both with and without including realized volatility.

Figures 8 to 11 report the results of the VARs. Figure 8 reports the response of the economy to a BBD shock in a VAR with only BBD to capture uncertainty. We find that innovations in the BBD index are associated with declines in employment and IP. Figure 9, though, shows that when we control for rv (ordering it first in the VAR), the BBD index becomes only trivially contractionary, and the point estimates imply that it is in fact expansionary after a year. The estimated effects of an rv shock on the economy appear unchanged whether we include the BBD index or f_6 in the

VAR.

We obtain highly similar results for the Michigan survey in figures 10 and 11. Again, the Michigan survey is contractionary when it is included alone (as found by Leduc and Liu (2015)), but it is driven out by rv in figure 11. In appendix figures A.7–A.8, we show that similar results are obtained if we reverse the ordering in the VARs, placing the BBD and Michigan indexes before rv .

Jurado, Ludvigson, and Ng (2015) provide an important measure of uncertainty that depends on forecast uncertainty for a wide range of macro and financial variables. In VARs, their measure drives out both the VIX and realized stock market volatility as a predictor of future output growth. Their time series for uncertainty is measured purely based on the history of realized volatility (squared forecast errors). We therefore view the results of Jurado, Ludvigson, and Ng (2015) as consistent with ours, in the sense that they both say that realized volatility (or some transformation of it) is an important driver of fluctuations in the real economy. Our results are also consistent with more recent work by Ludvigson, Ma, and Ng (2015), who find that a large fraction of the variation in uncertainty in the economy is actually an endogenous response to shocks, rather than an exogenous driving force.

4.6 Confounding effects of time-varying risk premia

Our main results measure volatility expectations with f_6 , which is a market-based measure of volatility expectations. f_6 therefore depends not only on the statistical expectation of volatility, but also on a risk premium. A natural question, then, is whether variation in the risk premium embedded in f_6 would be expected to affect our results.

Risk premia are generally viewed as varying countercyclically (e.g. Campbell and Cochrane (1999), among many others). Since volatility is high in bad times, it earns a negative risk premium (Coval and Shumway (2001)). A countercyclical risk premium would thus imply that when output is low, f_6 should be *higher* than average (i.e. investors are particularly risk averse, so they are willing to pay more for the protection that a volatility forward like f_6 provides against high-volatility states).

So countercyclical variation in risk premia should induce a *negative* relationship between volatility expectations and the state of the economy, exactly the opposite of what we see in the data. In order for our results on the effects of volatility expectations to be due to fluctuations in the risk premium, it would have to be the case that investors are relatively *less* risk averse when output is low or falling.

As to realized volatility, risk premia should not have any direct effects. Realized volatility is measured simply as the sum of squared daily returns on the S&P 500. A day is sufficiently small that any realistic variation in the conditional mean return on the market would have a trivial effect on realized volatility.

5 Contractionary volatility or volatile contractions?

The analysis above shows that fluctuations in expected future stock market volatility are associated with, if anything, expansions in output, while shocks to realized volatility are robustly associated with contractions. Those results immediately suggest that contractions are simply volatile, rather than news about volatility being contractionary. In this section, though, we discuss general theoretical implications of the two hypotheses about the relationship between volatility and the real economy. Those implications then motivate us to enrich our results above.

While there is broad interest in the effects of shocks to volatility, there is no canonical model of volatility fluctuations, either causing or caused by business cycles. We therefore attempt to draw testable implications from a more basic description of contractionary volatility and volatile contractions. However, we also show that the testable implications that we draw arise in a pair of models that we view as representative.

5.1 Contractionary volatility

When we refer to contractionary volatility, we mean the hypothesis that an exogenous increase in volatility in the economy induces a recession (reducing output, employment, consumption, and investment). In general, such a hypothesis assumes that the uncertainty is about the *future*, implying that the impact of the shock occurs before the realization of the volatility. That is, the assumption is that on date t , we receive information that the variance of shocks on date $t + 1$ or later will be high. So the news arrives before the volatility itself does.

Fernandez-Villaverde et al. (2013), Basu and Bundick (2015), Leduc and Liu (2015) are leading recent examples of such models. We claim that there are at least three basic implications that can be drawn from the contractionary volatility hypothesis. As a specific example of this type of model, we have calibrated a model similar to that of Basu and Bundick (2015) and show that the implications claimed below hold in that case.

5.1.1 Implications:

1. *Unexpected increases in expected volatility, i.e. news about future volatility, cause low or falling economic activity.* This is really the basic assumption of the contractionary volatility hypothesis. In the most basic form of the hypothesis, then, there should be a negative relationship between economic activity and shocks to expected volatility.
2. *News about high future volatility should earn a negative risk premium.* The hypothesis is that news about high future volatility is associated with reductions in output and consumption. If periods of low output are also periods with high marginal utility, then people will be willing to pay to hedge shocks to volatility expectations, in the sense that assets whose returns covary positively with expected volatility will earn negative returns. For example, under power utility, if consumption is low when expected volatility is high, then the Arrow–Debreu price

of such states is relatively high, and assets that insure investors against high volatility should earn negative risk premia.

3. *Stock returns and innovations to output will not be skewed.* The key assumption of the contractionary volatility hypothesis is that the key shock is one to expected *future* uncertainty. That is, in period t there is news that risk on dates $t + 1$ or later is higher. Stock returns and output growth will thus be unexpectedly negative on date t and volatile on date $t + 1$. But skewness is a measure of the correlation between the level and volatility (i.e. $E[\varepsilon^3] = E[\varepsilon \cdot \varepsilon^2]$). Since negative shocks are not contemporaneously correlated with volatile shocks in this model, there is no force generating skewness.

5.2 Volatile contractions

The volatile contractions hypothesis says that recessions are periods of high volatility. It is fundamentally a story about skewness: negative shocks are volatile shocks, so denoting shocks as ε , $E[\varepsilon^3] < 0$. There is no claim about causation or timing. Rather the view is that bad news is volatile news. This type of model may be illustrated with a simple real business cycle model that features negatively skewed innovations to the level of technology. There is a long literature studying mechanisms that can generate skewness endogenously. Models with frictions that bind mainly in recessions (e.g. financial frictions) tend to generate skewness (see, for example, Kocherlakota (2000) and Ordóñez (2013)). Other mechanisms, such as learning (van Nieuwerburgh and Veldkamp (2006)) can also generate asymmetries, or the fundamental shocks might simply be skewed to the left.

5.2.1 Implications:

The subjects of the implications here parallel those for contractionary volatility, but the direction is different.

1. *Increases in realized volatility are correlated with low or falling economic activity.* Under the volatile contractions hypothesis, there is no causal relationship between expected volatility and activity. There will be, however, a clear correlation between realized volatility and low output. Again, this is really just the basic assumption of the volatile contractions hypothesis.
2. *Surprises in realized volatility should earn a negative risk premium.* The key implication of the volatile contractions hypothesis is that high realized volatility occurs in bad states of the world. We would thus expect investors to be willing to pay for insurance against realized volatility.
3. *Stock returns and innovations to output will be negatively skewed.* Again, the key assumption in this case is that negative shocks are volatile shocks. This is fundamentally a statement that $E[\varepsilon^3] < 0$.

5.3 Distinguishing the hypotheses

The results we obtain above showing that realized volatility is contractionary while expected volatility is not provide direct evidence on the first implication of the two models. However, there are two additional implications of the contractionary volatility and volatile contractions models that allow them to be distinguished. We now show that the data on risk premia and skewness also support the view that contractions are volatile, rather than that volatility is contractionary.

5.4 Risk premia

In the main analysis above, we used forward claims on stock return variance as state variables in a VAR. But since those forwards are financial assets, they themselves have returns. Assets that are risky, in the sense that their performance is positively correlated with the state of the economy (more formally, negatively correlated with marginal utility), should earn positive risk premia, while assets that are hedges, in the sense that their performance is *negatively* correlated with the state of the economy, should earn *negative* risk premia.

Claims to future variance earn high returns when variance is high. Specifically, the return on a 6-month variance forward is high over the next month if we receive information that volatility 6 months in the future will be high. On the other hand, the return on the 1-month variance forward (a claim to realized variance over the next month) is exposed to *realized* variance – the return is high when realized variance is high.

So under the contractionary volatility hypothesis, we would expect that the 6-month variance forward earns a negative risk premium, while there is no particular prediction about the premium for realized volatility. Conversely, the volatile contractions hypothesis says that news about future volatility need not earn any premium. But volatile contractions imply that the 1-month forward, since it is exposed to realized volatility, which is high in downturns, should earn a negative risk premium.

The question, then, is whether risk premia are higher for short-term (1-month) or longer-term (6-month) variance forwards. Dew-Becker et al. (2016) study that question in detail. We report here a simple summary of the risk premia and also extend their sample back to the 1980's.

Figure 12 plots the average Sharpe ratios earned by forward variance claims between 1996 and 2015. We focus on this sample because it allows us to construct the returns from data on variance swaps, which have the highest quality price information (they are direct claims on variance, whereas the MFIVs we study require numerically integrating across many option prices). The solid line plots the sample mean Sharpe ratio, while the dotted lines bound the 95-percent confidence interval. The results in the figure are stark: only the one- and two-month variance forwards earn returns that are statistically significantly negative – all other point estimates are positive, and some are even statistically significantly so.

As an economic matter, the Sharpe ratios that we observe on the one-month variance swap are extremely large, at -1.3. The return earned by an investor who sells one-month variance claims is

three to five times larger than that earned by an investor in the aggregate stock market. In other words, investors are willing to pay enormous premia for protection against periods of high realized volatility, but the premia associated with shocks to expected future volatility are not distinguishable from zero (and are statistically significantly smaller than that on realized volatility).

The Sharpe ratios directly reveal what states of the world investors have paid for protection against. Specifically, under power utility, one may show that for any asset return x , the Sharpe ratio is

$$\frac{ER[x]}{SD[x]} = -RRA \times std(\Delta c) \times corr(x, \Delta c) \quad (11)$$

where RRA is the coefficient of relative risk aversion. Assets with larger Sharpe ratios therefore have larger correlations with consumption growth. More generally, assets that have a higher correlation with marginal utility should earn more negative Sharpe ratios. The fact that the largest Sharpe ratios are earned by one-month variance claims tells us that investors view realized volatility as being more correlated with their marginal utility than news about future volatility.

In order to be able to use our longer sample running back to the 1980's, figure 13 reports the average shape of the term structure of variance forwards constructed using the CME options data (which is estimated more precisely than the month-to-month returns). The term structure reported here is still extremely informative about risk premia. The average return on an n -month variance claim is:⁷

$$\begin{aligned} E \left[\frac{P_{n-1,t} - P_{n,t-1}}{P_{n,t}} \right] &\approx \frac{E[P_{n-1}] - E[P_n]}{E[P_n]} & (12) \\ &= E[R_{n,t}] - 1 & (13) \end{aligned}$$

The slope of the average term structure thus indicates the average risk premia. If the term structure is upward sloping, then the prices of the variance claims fall on average as their maturities approach, indicating that they have negative average returns. If it slopes down, then average returns are positive.

Figure 13 plots the average term structure of variance forward prices for the period 1983–2013. The term structure is strongly upward sloping for the first two months, again indicating that investors are willing to pay large premia for assets that are exposed to realized variance and expected variance one month in the future. But the curve quickly flattens, indicating that the risk premia for exposure to fluctuations in expected variance farther in the future are much smaller.

The asset return data says that investors are highly averse to news about high realized volatility, but shocks to expected volatility in the future earn relatively small returns. The confidence intervals that we obtain are sufficiently wide that we cannot claim that shocks to expected future volatility do not earn an economically meaningfully negative risk premium. What we can say, though, is that investors seem to care *much more* about surprises in realized volatility than in expected volatility.

⁷Whereas $F_{n,t}$ is in volatility terms, we work here in variance term since variance swaps returns are expressed in these units. The two are closely related: $P_{n,t} = F_{n,t}^2$.

5.5 Skewness

The volatile contractions model is fundamentally about skewness in shocks. There are large literatures studying skewness in both aggregate stock returns and economic growth. We therefore provide just a brief overview of the literature and the basic evidence.

Table 2 reports the skewness of monthly and quarterly changes in a range of measures of economic activity. Nearly all the variables that we examine are negatively skewed, at both the monthly and quarterly levels. One major exception is monthly growth in industrial production, but that result appears to be due to some large fluctuations in the 1950's. When the sample is cut off at 1960, the results for industrial production are consistent with those for other variables.

In addition to real variables, table 2 also reports realized and option-implied skewness for S&P 500 returns.⁸ The implied and realized skewness of monthly stock returns is substantially negative, and in fact surprisingly similar to the skewness of capacity utilization. The realized skewness of stock returns is less negative than option-implied skewness, which is consistent with investors demanding a risk premium on assets that have negative returns in periods when realized skewness is especially negative (i.e. that covary positively with skewness).

In addition to the basic evidence reported here, there is a large literature providing much more sophisticated analyses of asymmetries in the distributions of output and stock returns. Morley and Piger (2012) provide an extensive analysis of asymmetries in the business cycle and review the large literature. They estimate a wide range of models, including symmetrical ARMA specifications, regime-switching models, and frameworks that allow nonlinearity. The models that fit aggregate output best have explicit non-linearity and negative skewness. Even after averaging across models using a measure of posterior probability, which puts substantial weight on purely symmetrical models, Morley and Piger find that their measure of the business cycle is substantially skewed to the left, consistent with the results reported in table 2.

Finally, the finance literature has long recognized that there is skewness in aggregate equity returns and in option-implied return distributions (see Campbell and Hentschel (1992), Ait-Sahalia and Lo (1998), and Bakshi, Kapadia, and Madan (2003), for recent analyses and reviews). The skewness that we measure here appears to be pervasive and has existed in returns reaching back even to the 19th century (Campbell and Hentschel (1992)).

Taken as a whole, then, across a range of data sources and estimation methods, there is a substantial body of evidence that fluctuations in the economy are negatively skewed. In a world of negative skewness, it is not surprising that measures of realized volatility are correlated with declines in activity, simply because skewness is related to the third moment: $E[\varepsilon^3] = E[\varepsilon \cdot \varepsilon^2]$.

⁸We obtain option-implied skewness from the CBOE's time series of its SKEW index, which is defined as $SKEW = 100 - 10 \times Skew(R)$. We thus report $10 - SKEW/10$.

6 Conclusion

The goal of this paper is to understand whether shocks to uncertainty have negative effects on the economy. Our contribution is to estimate a range of models that include measures of both volatility expectations and also realized volatility. We find that shocks to expected volatility in the future, after controlling for current realized volatility, do not have negative effects on the economy – if anything, they are expansionary. The evidence we present favors the view that bad times are volatile times, not that volatility causes bad times. A leading hypothesized explanation for the slow recovery from the 2008 financial crisis has been that uncertainty since then has been high. Our evidence suggests that uncertainty may not have been the driving force, and that economists should search elsewhere for an explanation to the slow recovery puzzle.

References


- Adrian, Tobias and Joshua Rosenberg**, “Stock Returns and Volatility: Pricing the Short-Run and Long-run Components of Market Risk,” *The Journal of Finance*, 2008, *63*(6), 2997–3030.
- Aït-Sahalia, Yacine and Andrew W Lo**, “Nonparametric estimation of state-price densities implicit in financial asset prices,” *The Journal of Finance*, 1998, *53* (2), 499–547.
- Ait-Sahalia, Yacine, Mustafa Karaman, and Lorian Mancini**, “The Term Structure of Variance Swaps, Risk Premia and the Expectations Hypothesis,” 2014. Working paper.
- Alexopoulos, Michelle**, “Read All about It!! What Happens Following a Technology Shock?,” *American Economic Review*, 2011, *101* (4), 1144–1179.
- and **Jon Cohen**, “Uncertain Times, Uncertain Measures,” 2009. Working paper.
- Andersen, Torben G. and Oleg Bondarenko**, “Construction and Interpretation of Model-Free Implied Volatility,” 2007. Working paper.
- Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang**, “The Cross-Section of Volatility and Expected Returns,” *The Journal of Finance*, 2006, *61*(1), 259–299.
- 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.
- , —, and **Dilip Madan**, “Stock return characteristics, skew laws, and the differential pricing of individual equity options,” *Review of Financial Studies*, 2003, *16* (1), 101–143.
- Basu, Susanto and Brent Bundick**, “Uncertainty Shocks in a Model of Effective Demand,” Working Paper 18420, National Bureau of Economic Research 2015.
- and —, “Uncertainty Shocks in a Model of Effective Demand,” 2015. Working paper.

- Beaudry, Paul and Franck Portier**, “Stock Prices, News, and Economic Fluctuations,” *The American Economic Review*, 2006, *96* (4), 1293–1307.
- Black, Fischer and Myron Scholes**, “The Pricing of Options and Corporate Liabilities,” *The Journal of Political Economy*, 1973, *81* (3), 637–654.
- Bloom, Nicholas**, “The Impact of Uncertainty Shocks,” *Econometrica*, 2009, *77*(3), 623–685.
- , **Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry**, “Really Uncertain Business Cycles,” 2014. Working paper.
- , **Scott Baker, and Steven J. Davis**, “Measuring Economic Policy Uncertainty,” 2015. Working paper.
- Bollerslev, Tim, George Tauchen, and Hao Zhou**, “Expected Stock Returns and Variance Risk Premia,” *Review of Financial Studies*, 2009, *22*(11), 4463–4492.
- Broadie, Mark and Jerome Detemple**, “American option valuation: new bounds, approximations, and a comparison of existing methods,” *Review of Financial Studies*, 1996, *9* (4), 1211–1250.
- Campbell, John Y. and John H. Cochrane**, “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 1999, *107* (2), 205–251.
- and **Ludger Hentschel**, “No news is good news: An asymmetric model of changing volatility in stock returns,” *Journal of Financial Economics*, 1992, *31* (3), 281–318.
- , **Martin Lettau, Burton G. Malkiel, and Yexiao Xu**, “Have Individual Stocks Become More Volatile? an Empirical Exploration of Idiosyncratic Risk,” *Journal of Finance*, 2001, *56* (1), 1–43.
- Carr, Peter and Liuren Wu**, “Variance Risk Premiums,” *Review of Financial Studies*, 2009, *22*(3), 1311–1341.
- Coval, Joshua D. and Tyler Shumway**, “Expected Option Returns,” *The Journal of Finance*, 2001, *56*(3), 983–1009.
- Dew-Becker, Ian, Stefano Giglio, Anh Le, and Marius Rodriguez**, “The Price of Variance Risk,” 2016. Working paper.
- , —, —, and —, “The Price of Variance Risk,” 2016. Working paper.
- Egloff, Daniel, Markus 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.

- Eisfeldt, Andrea L. and Adriano A. Rampini**, “Capital Reallocation and Liquidity,” *Journal of Monetary Economics*, 2006, 53 (3), 369–399.
- Fernandez-Villaverde, Jesus, Pablo Guerron-Quintana, Keith Kuester, , and Juan Rubio-Ramirez**, “Fiscal Volatility Shocks and Economic Activity,” 2013. Working paper.
- Gatheral, Jim and Antoine Jacquier**, “Convergence of Heston to SVI,” *Quantitative Finance*, 2011, 11 (8), 1129–1132.
- and — , “Arbitrage-free SVI volatility surfaces,” *Quantitative Finance*, 2014, 14 (1), 59–71.
- Gilchrist, Simon, Jae W. Sim, and Egon Zakrajek**, “Uncertainty, Financial Frictions, and Investment Dynamics,” Working Paper 20038, National Bureau of Economic Research April 2014.
- Gourio, Francois**, “Disaster Risk and Business Cycles,” *American Economic Review*, 2012, 102(6), 2734–2766.
- , “Credit Risk and Disaster Risk,” *American Economic Journal: Macroeconomics*, 2013, 5(3), 1–34. Working paper.
- Güvenen, Fatih, Serdar Ozkan, and Jae Song**, “The Nature of Countercyclical Income Risk,” *Journal of Political Economy*, 2014, 122 (3), 621–660.
- Homescu, Cristian**, “Implied volatility surface: Construction methodologies and characteristics,” *Available at SSRN 1882567*, 2011.
- Jiang, George J. and Yisong S. Tian**, “The Model-Free Implied Volatility and its Information Content,” *Review of Financial Studies*, 2005, 18 (4), 1305–1342.
- Jiang, George J and Yisong S Tian**, “Extracting model-free volatility from option prices: An examination of the VIX index,” *Journal of Derivatives*, 2007, 14 (3).
- Jurado, Kyle, Sydney Ludvigson, and Serena Ng**, “Measuring Uncertainty,” *American Economic Review*, 2015, 105 (3), 1177–1216.
- Kocherlakota, Narayana**, “Creating business cycles through credit constraints,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 2000, 24 (3), 2–10.
- Leduc, Sylvain and Zheng Liu**, “Uncertainty Shocks are Aggregate Demand Shocks,” 2015. Working paper.
- Lee, Roger W**, “The moment formula for implied volatility at extreme strikes,” *Mathematical Finance*, 2004, 14 (3), 469–480.


- Morley, James and Jeremy Piger**, “The asymmetric business cycle,” *Review of Economics and Statistics*, 2012, *94* (1), 208–221.
- Nieuwerburgh, Stijn Van and Laura Veldkamp**, “Learning asymmetries in real business cycles,” *Journal of Monetary Economics*, 2006, *53* (4), 753–772.
- Ordonez, Guillermo**, “The Asymmetric Effects of Financial Frictions,” *Journal of Political Economy*, 2013, *121* (5), 844–895.
- Stock, James H. and Mark W. Watson**, “Disentangling the Channels of the 2007–2009 Recession,” *Brookings Papers on Economic Activity*, 2012, pp. 81–135.
- Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron**, “Cyclical Dynamics in Idiosyncratic Labor Market Risk,” *Journal of Political Economy*, 2004, *112* (3), 695–717.
- Taylor, Stephen J, Pradeep K Yadav, and Yuanyuan Zhang**, “The information content of implied volatilities and model-free volatility expectations: Evidence from options written on individual stocks,” *Journal of Banking & Finance*, 2010, *34* (4), 871–881.
- Wu, Jing Cynthia and Fan Dora Xia**, “Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound,” 2014. Working paper.
- Zeliade**, “Quasi-explicit calibration of Gatheral’s SVI model,” 2009.

Figure 1: Time series of realized variance and expectations




Note: Time series of realized variance, 1-month, and 6-month expectations (F_n), in annualized units. Grey bars indicate NBER recessions.

Figure 2: Predictive regressions




Note: Coefficients of regressions of rv on lagged rv and 6-month forwards, f_6 , at different lags (X axis). Top row reports coefficient of univariate regressions of rv on lagged f_6 , for different lags. Bottom row reports coefficients of multivariate regressions of rv onto lagged rv and lagged f_6 , for different lags (respectively, left and right panel).

Figure 3: Correlation of macro variables to rv and expected volatility shocks




Note: The figure shows the response of macroeconomic variables to shocks to rv (left) and Expected Volatility measured by f_6 (right), before and after the shock (lag 0 on the X axis). In particular, the solid line reports the coefficient of regressions of macroeconomic variables (Industrial Production, first row, Hours Worked, second row, Employment, third row, and Inflation, fourth row) on lags and lags of shocks to rv (left column) and Expected Volatility (right column). Dashed lines indicate 95% confidence intervals. The sample covers 1983 to 2013.

Figure 4: Impulse response functions from VAR (with only VXO)




Note: The figure shows impulse response functions (with 90% and 95% CI) of volatility (measured by VXO), employment and industrial production to a shock to VXO, in a VAR with VXO, federal funds rate, log employment and log industrial production. Impulse response functions reported correspond to structural shocks with Choleski decomposition (in the order reported above). Sample covers 1986-2014. All macroeconomic series were detrended with a one-sided HP filter.

Figure 5: Impulse response functions from VAR (ordering rv first and f_6 second)




Note: The figure shows impulse response functions (with 90% and 95% CI) of rv , volatility expectations (measured by f_6), employment and industrial production to shocks to rv and f_6 , in a VAR with rv , f_6 , federal funds rate, log employment and log industrial production. Impulse response functions reported correspond to structural shocks with Choleski decomposition (in the order reported above). Sample covers 1983-2014. All macroeconomic series were detrended with a one-sided HP filter.

Figure 6: Quarterly VAR specification




Note: The figure shows impulse response functions (with 90% and 95% CI) from a VAR with rv , volatility expectations (measured by f_6), and the macroeconomic series from Basu and Bundick (2015): GDP (Y), consumption (C), investment (I), hours (H), the GDP deflator (DEF), M2 and the FFR. The figure reports IRF of Y, C, I, H to shocks to rv and f_6 . Impulse response functions reported correspond to structural shocks with Choleski decomposition (in the order reported above). Sample covers 1983-2014.

Figure 7: Impulse response functions from VAR (ordering f_6 first and rv second)




Note: The figure shows impulse response functions (with 90% and 95% CI) of rv , volatility expectations (measured by f_6), employment and industrial production to shocks to f_6 and rv , in a VAR with rv , f_6 , federal funds rate, log employment and log industrial production. Impulse response functions reported correspond to structural shocks with Choleski decomposition (in the order reported above). Sample covers 1983-2014. All macroeconomic series were detrended with a one-sided HP filter.

Figure 8: VAR with Bloom, Baker and Davis measure (BBD; 2015)




Note: The figure shows impulse response functions (with 90% and 95% CI) of uncertainty (measured as in Bloom, Baker and Davis (BBD, 2015)), employment and industrial production to a shock to the BBD measure, in a VAR with BBD, federal funds rate, log employment and log industrial production. Impulse response functions reported correspond to structural shocks with Choleski decomposition (in the order reported above). Sample covers 1985-2014. All macroeconomic series were detrended with a one-sided HP filter.

Figure 9: VAR with both rv and Bloom, Baker and Davis measure (BBD; 2015)




Note: The figure shows impulse response functions (with 90% and 95% CI) of rv , uncertainty (measured as in Bloom, Baker and Davis (BBD, 2015)), employment and industrial production to shocks to rv and BBD, in a VAR with rv , BBD, federal funds rate, log employment and log industrial production. Impulse response functions reported correspond to structural shocks with Choleski decomposition (in the order reported above). Sample covers 1985-2014. All macroeconomic series were detrended with a one-sided HP filter.

Figure 10: VAR with uncertainty measure from Michigan survey




Note: The figure shows impulse response functions (with 90% and 95% CI) of uncertainty (from the Michigan survey), employment and industrial production to a shock to the Michigan uncertainty measure, in a VAR with the Michigan measure, federal funds rate, log employment and log industrial production. Impulse response functions reported correspond to structural shocks with Choleski decomposition (in the order reported above). Sample covers 1983-2014. All macroeconomic series were detrended with a one-sided HP filter.

Figure 11: VAR with both rv and uncertainty from the Michigan survey




Note: The figure shows impulse response functions (with 90% and 95% CI) of rv , uncertainty (from the Michigan survey), employment and industrial production to shocks to rv and the Michigan uncertainty measure, in a VAR with rv , the Michigan measure, federal funds rate, log employment and log industrial production. Impulse response functions reported correspond to structural shocks with Choleski decomposition (in the order reported above). Sample covers 1983-2014. All macroeconomic series were detrended with a one-sided HP filter.

Figure 12: Annualized Sharpe ratios for forward variance claims



Note: The figure shows the annualized Sharpe ratio for the forward variance claims, constructed using Variance Swaps. 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. All tests for the difference in Sharpe ratio between the 1-month variance swap and any other maturity confirm that they are statistically different with a p-value of 0.03 (for the second month) and < 0.01 (for all other maturities). The sample used is 1996-2013. For more information on the data sources, see Dew-Becker et al. (2015).

Figure 13: Average variance forward prices, 1983–2013



Note: The figure shows the average prices of forward variance claims of different maturity, for the period 1983–2013. All prices are reported in annualized volatility terms. Maturity zero corresponds to average realized volatility.

Table 1: Descriptive statistics

Panel A: Descriptive statistics				
	Mean	Std.	Skewness	
RV (annualized)	15.22	8.64	3.05	
F_1 (annualized)	19.22	7.75	2.07	
F_3 (annualized)	20.01	6.51	1.38	
F_6 (annualized)	21.05	6.65	1.08	

Panel B: Correlations				
	RV	F_1	F_3	F_6
RV	1.00	0.89	0.77	0.68
F_1	0.89	1.00	0.94	0.88
F_3	0.77	0.94	1.00	0.97
F_6	0.68	0.88	0.97	1.00

Panel C: Correlations with economic activity					
	rv_t	f_6	Unempl.	Cap. Ut.	S&P 500 ret.
rv_t	1.00	0.71	-0.03	-0.29	-0.29
f_6	0.71	1.00	0.12	-0.40	-0.10
Unemployment	-0.03	0.12	1.00	-0.70	0.10
Capacity Utilization	-0.29	-0.40	-0.70	1.00	-0.06
S&P 500 return	-0.29	-0.10	0.10	-0.06	1.00

Panel D: Decomposition of $V(rv)$			
	$V(rv_{t+1} - E_t^1)$	$V(E_{t+5}^1 - E_t^6)$	$V(E_t^6)$
1. Using f_6	42%	42%	16%
2. VAR	46%	36%	18%
3. ARMA(1,1)	46%	43%	11%
4. ARMA(2,2)	44%	38%	17%

Panel E: Correlation of 6-month forecasts, based on:			
	Lagged f_6	Lagged rv	6 predictors
Lagged f_6	1.00	0.71	0.84
Lagged rv	0.71	1.00	0.72
6 predictors from Campbell et al. (2015)	0.84	0.72	1.00

Note: The table reports various statistics on realized volatility, forwards and their relationship. Panel A reports the mean, standard deviation and skewness of realized volatility and the 1, 3, and 6-month forwards (in annualized volatility units). Panel B reports the correlations between those variables. Panel C shows the correlation of log RV and log F_6 (respectively, denoted rv and f) with unemployment, capacity utilization, and the S&P 500 return. Panel D computes a variance decomposition for $V(rv)$ into the surprise component $V(rv_{t+1} - E_t^1)$, the volatility news 1 to 5 months ahead, $V(E_{t+5}^1 - E_t^6)$, and the volatility news 6 months ahead, $V(E_t^6)$. Each row constructs the volatility expectations using a different model. The first row uses f_6 as a univariate predictor. The second row uses a VAR with variables rv , f_1 and f_6 . The third and fourth row compute expectations using and ARMA(1,1) and an ARMA(2,2) model for rv . Panel E reports correlations of 6-month volatility forecasts from different models. The first row uses lagged f_6 as a predictor. The second row uses lagged rv as a predictor. The last row uses as predictors the 6 state variables from Campbell et al. (2015): P/E ratio of the S&P 500, value spread, default spread, term spread, rv and S&P return. Sample period is 1983-2013.

Table 2: Skewness

<i>Panel A: real economic activity</i>	Monthly	Quarterly	Start of sample (year)
Employment	-0.41	-0.41	1948
Capacity Utilization	-1.02	-1.30	1967
IP	0.17	-0.16	1948
IP, starting 1960	-0.93	-1.28	1960
Y		-0.11	1947
C		-0.28	1947
I		-0.03	1947

<i>Panel B: skewness of S&P 500 monthly returns</i>	
Implied (since 1990)	-1.81
Realized (since 1926)	0.36
Realized (since 1948)	-0.42
Realized (since 1990)	-0.61

Note: Panel A reports the skewness of changes of employment, capacity utilization, industrial production (beginning both in 1948 and in 1960), GDP, consumption and investments. The first column reports the skewness of monthly changes, the second column the skewness of quarterly changes. Panel B reports the realized skewness of S&P 500 monthly returns in different periods, as well as the implied skewness computed by the CBOE using option prices.

A.1 MFIV Construction

In this section we describe the details of the procedure we use to construct MFIV at different horizons, starting from our dataset of end-of-day prices for American options on S&P 500 futures from the CME.

A.1.1 Main steps of MFIV construction

A first step in constructing the MFIV is to obtain implied volatilities corresponding to the observed option prices. We do so using a binomial model.¹ For the most recent years, CME itself provides the implied volatility together with the option price. For this part of the sample, the IV we estimate with the binomial model and the CME’s IV have a correlation of 99%, which provides an external validation on our implementation of the binomial model.

Once we have estimated these implied volatilities, we could in theory simply invert them to yield implied prices of European options on forwards. These can then be used to compute the MFIV directly as described in equation (2).

In practice, however, an extra step is required before inverting for the European option prices and integrating to obtain the MFIV. The MFIV defined in equation (2) depends on the integral of option prices over *all strikes*, but option prices are only observed at discrete strikes. We are therefore forced to interpolate option prices between available strikes and also extrapolate beyond the bounds of observed strikes.² Following the literature, we fit a parametric model to the Black–Scholes implied volatilities of the options and use the model to then interpolate and extrapolate across all strikes (see, for example, Jiang and Tian (2007), Carr and Wu (2009), Taylor, Yadav, and Zhang (2010), and references therein). Only after this extra interpolation-extrapolation step, the fitted implied volatilities are then inverted to yield option prices and compute MFIV according to equation (2). To interpolate and extrapolate the implied volatility curve, we use the SVI (stochastic volatility inspired) model of Gatheral and Jacquier (2014).

In the next sections, we describe in more detail the interpolation-extrapolation step of the procedure (SVI fitting) as well as our construction of MFIV after fitting the SVI curve. Finally, we report a description of the data we use and some examples and diagnostics on the SVI fitting method.

A.1.2 SVI interpolation: theory

There are numerous methods for fitting implied volatilities across strikes. Homescu (2011) provides a thorough review. We obtained the most success using Gatheral’s SVI model (see Gatheral and Jacquier 2014). SVI is widely used in financial institutions because it is parsimonious but also known to approximate well the behavior of implied volatility in fully specified option pricing models (e.g.

¹See for example Broadie and Detemple (1996) and Bakshi, Kapadia, and Madan (2003), among others.

²See Jiang and Tian (2007) for a discussion of biases arising from the failure to interpolate and extrapolate.

Gatheral and Jacquier (2011)); SVI also satisfies the limiting results for implied volatilities at very high and low strikes in Lee (2004), and, importantly, ensures that no-arbitrage conditions are not violated.

The SVI model simply assumes a hyperbolic relationship between implied variance (the square of the Black–Scholes implied volatility) and the log moneyness of the option, k (log strike/forward price).

$$\sigma_{BS}^2(k) = a + b \left(\rho(k - m) + \sqrt{(k - m)^2 + \sigma^2} \right)$$

where $\sigma_{BS}^2(k)$ is the implied variance under the Black–Scholes model at log moneyness k . SVI has five parameters: a , b , ρ , m , and σ . The parameter ρ controls asymmetry in the variances across strikes. Because the behavior of options at high strikes has minimal impact on the calculation of model-free implied volatilities, and because we generally observe few strikes far above the spot, we set $\rho = 0$ (in simulations with calculating the VIX for the S&P 500 – for which we observe a wide range of options – we have found that including or excluding ρ has minimal impact on the result).

We fit the parameters of SVI by minimizing the sum of squared fitting errors for the observed implied volatilities. Because the fitted values are non-linear in the parameters, the optimization must be performed numerically. We follow the methodology in Zeliade (2009) to analytically concentrate a and b out of the optimization. We then only need to optimize numerically over σ and m (as mentioned above, we set $\rho = 0$). We optimize with a grid search over $\sigma \times m = [0.001, 10] \times [-1, 1]$ followed by the simplex algorithm.

For many date/firm/maturity triplets, we do not have a sufficient number of contract observations to fit the implied volatility curve (i.e. sometimes fewer than four). We therefore include strike/implied volatility data from the two neighboring maturities and dates in the estimation. The parameters of SVI are obtained by minimizing squared fitting errors. We reweight the observations from the neighboring dates and maturities so that they carry the same amount of weight as the observations from the date and maturity of interest. Adding data in this way encourages smoothness in the estimates over time and across maturities but it does not induce a systematic upward or downward bias. We drop all date/firm/maturity triplets for which we have fewer than four total options with $k < 0$ or fewer than two options at the actual date/firm/maturity (i.e. ignoring the data from the neighboring dates and maturities).

When we estimate the parameters of the SVI model, we impose conditions that guarantee the absence of arbitrage. In particular, we assume that $b \leq \frac{4}{(1+|\rho|)T}$, which when we assume $\rho = 0$, simplifies to $b \leq \frac{4}{T}$. We also assume that $\sigma > 0.0001$ in order to ensure that the estimation is well defined. Those conditions do not necessarily guarantee, though, that the integral determining the MFIV is convergent (the absence of arbitrage implies that a risk-neutral probability density exists – it does not guarantee that it has a finite variance). We therefore eliminate observations where the integral determining MFIV fails to converge numerically. Specifically, we eliminate observations where the argument of the integral does not approach zero as the log strike rises above two standard deviations from the spot or falls more than five standard deviations below the strike (measured

based on the at-the-money implied volatility).

A.1.3 Construction of MFIV from the SVI fitted curve

After fitting the SVI curve for each date and maturity, we compute the integral in equation (2) numerically, over a range of strikes from -5 to +2 standard deviations away from the spot price.³ We then have an MFIV for every firm/date/maturity observation. The MFIVs are then interpolated (but not extrapolated) to give MFIVs at maturities from 1–12 months for each firm/date pair.

A.1.4 Data description and diagnostics of SVI fitting

Our dataset consists of 2.3 million end-of-day prices for all American options on S&P 500 futures from the CME.

When more than one option (e.g. a call and a put) is available at any strike, we compute IV at that strike as the average of the observed IVs. We keep only IVs greater than zero, at maturities higher than 9 days and lower than 2 years, for a total of 1.9 million IVs. The number of available options has increased over time, as demonstrated by Figure A.2 (top panel), which plots the number of options available for MFIV estimation in each year.

The maturity structure of observed options has also expanded over time, with options being introduced at higher maturities and for more intermediate maturities. Figure A.1 (top panel) reports the cross-sectional distribution of available maturities in each year to estimate the term structure of MFIV. The average maturity of available options over our sample was 4 months, and was relatively stable. The maximum maturity observed ranged from 9 to 24 months and varied substantially over time.

Crucial to compute the MFIV is the availability of IVs at low strikes, since options with low strikes receive a high weight in the construction of MFIV. The bottom panel of Figure A.1 reports the minimum observed strike year by year, in standard deviations below the spot price. In particular, for each day we computed the minimum available strike price, and the figure plots the average of these minimum strike price across all days in each year; this ensures that the number reported does not simply reflect outlier strikes that only appear for small parts of each year.

Figure A.1 shows that in the early part of our sample, we can typically observe options with strikes around 2 standard deviations below the spot price; this number increases to around 2.5 towards the end of the sample.

These figures show that while the number of options was significantly smaller at the beginning of the sample (1983), the maturities observed and the strikes observed did not change dramatically over time.

Figure A.3 shows an example of the SVI fitting procedure for a specific day in the early part of our sample (November 7th 1985). Each panel in the figure corresponds to a different maturity. On

³In general this range of strikes is sufficient to calculate the MFIV. However, the MFIV technically involves an integral over the entire positive real line. Our calculation is thus literally a calculation of Andersen and Bondarenko's (2007) corridor implied volatility. We use this fact also when calculating realized volatility.


that day, we observe options at three different maturities, of approximately 1, 4, and 8 months. In each panel, the x's represent observed IVs at different values of log moneyness k . The line is the fitted SVI curve, that shows both the interpolation and the extrapolation obtained from the model.

Figure A.4 repeats the exercise in the later part of our sample (Nov 1st 2006), where many more maturities and strikes are available.

Both figures show that the SVI model fits the observed variances extremely well. The bottom panel of Figure A.2 shows the average relative pricing error for the SVI model in absolute value. The graph shows that the typical pricing error for most of the sample is around 0.02, meaning that the SVI deviates from the observed IV by around 2% on average. Only in the very first years (up to 1985) pricing errors are larger, but still only around 10% of the observed IV.


Overall, the evidence in this section shows that our observed option sample since 1983 has been relatively stable along the main dimensions that matter for our analysis – maturity structure, strikes observed, and goodness of fit of the SVI model.

Figure A.1: Maturities and strikes in the CME dataset




Note: Top panel reports the distribution of maturities of options used to compute the VIX in each year, in months. Bottom panel reports the average minimum strike in each year, in standard deviations below the forward price. The number is obtained by computing the minimum observed strike in each date and at each maturity (in standard deviations below the forward price), and then averaging it within each year to minimize the effect of outliers.

Figure A.2: Number of options to construct the VIX and pricing errors




Note: Top panel reports the number of options used to compute the VIX in each year, in thousands. Bottom panel reports the average absolute value of the pricing error of the SVI fitted line relative to the observed implied variances, in proportional terms (i.e. 0.02 means absolute value of the pricing error is 2% of the observed implied variance).

Figure A.3: SVI fit: 11/7/1985




Note: Fitted implied variance curve on 11/7/1987, for the three available maturities. X axis is the difference in log strike and log forward price. x's correspond to the observed implied variances, and the line is the fitted SVI curve.

Figure A.4: SVI fit: 11/1/2006




Note: Fitted implied variance curve on 11/1/2006, for the three available maturities. X axis is the difference in log strike and log forward price. x's correspond to the observed implied variances, and the line is the fitted SVI curve. On 11/1/2006 also a maturity of 5 months was available (not plotted for reasons of space).


Figure A.5: Robustness (I): response of Employment to RV and expectations shocks across specifications




(a) Without detrending the macroeconomic time series



(b) VAR estimated at the quarterly frequency




(c) Subperiod 1988-2006 (excluding 1987 crash and financial crisis)




(d) Ordering RV and expectations last in the VAR

Note: The Figure reports the response of employment to RV shocks (left panels) and volatility expectations (right panels) in different specification. Each row of the figure corresponds to a different model specification. Row (a) does not detrend the macroeconomic time series. Row (b) repeats the exercise using quarterly, rather than monthly, data. Row (c) estimates the VAR in the subsample 1988-2006, which excludes both RV peaks (1987 crash and financial crisis). Row (d) orders the RV shock second to last and the expectation shock last in the VAR.


Figure A.6: Robustness (II): response of IP to RV and expectations shocks across specifications




(a) Without detrending the macroeconomic time series



(b) VAR estimated at the quarterly frequency




(c) Subperiod 1988-2006 (excluding 1987 crash and financial crisis)



(d) Ordering RV and expectations last in the VAR


Note: The Figure reports the response of IP to RV shocks (left panels) and volatility expectations (right panels) in different specification. Each row of the figure corresponds to a different model specification. Row (a) does not detrend the macroeconomic time series. Row (b) repeats the exercise using quarterly, rather than monthly, data. Row (c) estimates the VAR in the subsample 1988-2006, which excludes both RV peaks (1987 crash and financial crisis). Row (d) orders the RV shock second to last and the expectation shock last in the VAR.

Figure A.7: VAR with both rv and Bloom, Baker and Davis measure (BBD; 2015), BBD ordered first



Note: The figure shows impulse response functions (with 90% and 95% CI) of rv , uncertainty (measured as in Bloom, Baker and Davis (BBD, 2015)), employment and industrial production to shocks to rv and BBD, in a VAR with BBD, rv , federal funds rate, log employment and log industrial production. Impulse response functions reported correspond to structural shocks with Choleski decomposition (in the order reported above). Sample covers 1985-2014. All macroeconomic series were detrended with a one-sided HP filter.

Figure A.8: VAR with both rv and uncertainty from the Michigan survey, survey ordered first



Note: The figure shows impulse response functions (with 90% and 95% CI) of rv , uncertainty (from the Michigan survey), employment and industrial production to shocks to rv and the Michigan uncertainty measure, in a VAR with the Michigan measure, rv , federal funds rate, log employment and log industrial production. Impulse response functions reported correspond to structural shocks with Choleski decomposition (in the order reported above). Sample covers 1983-2014. All macroeconomic series were detrended with a one-sided HP filter.