

Taxonomy of Global Risk, Uncertainty, and Volatility Measures

Datta, Deepa, Juan M. Londono, Bo Sun, Daniel Beltran, Thiago Ferreira, Matteo Iacoviello, Mohammad R. Jahan-Parvar, Canlin Li, Marius Rodriguez, and John Rogers

Please cite paper as:

Datta, Deepa, Juan M. Londono, Bo Sun, Daniel Beltran, Thiago Ferreira, Matteo Iacoviello, Mohammad R. Jahan-Parvar, Canlin Li, Marius Rodriguez, and John Rogers (2017). Taxonomy of Global Risk, Uncertainty, and Volatility Measures. International Finance Discussion Papers 1216.

<https://doi.org/10.17016/IFDP.2017.1216>



International Finance Discussion Papers

Board of Governors of the Federal Reserve System

Number 1216

November 2017

Board of Governors of the Federal Reserve System

International Finance Discussion Papers

Number 1216

November 2017

Taxonomy of Global Risk, Uncertainty, and Volatility Measures

Deepa Datta, Juan M. Londono, Bo Sun, Daniel Beltran, Thiago Ferreira, Matteo Iacoviello, Mohammad R. Jahan-Parvar, Canlin Li, Marius Rodriguez, and John Rogers

NOTE: International Finance Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. References in publications to International Finance Discussion Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors. Recent IFDPs are available on the Web at <https://www.federalreserve.gov/econres/ifdp/>. This paper can be downloaded without charge from Social Science Research Network electronic library at <http://www.ssrn.com>.

Taxonomy of Global Risk, Uncertainty, and Volatility Measures

Deepa Datta, Juan M. Londono, Bo Sun,
Daniel Beltran, Thiago Ferreira, Matteo Iacoviello,
Mohammad R. Jahan-Parvar, Canlin Li, Marius Rodriguez, and John Rogers¹

Abstract:

A large number of measures for monitoring risk and uncertainty surrounding macroeconomic and financial outcomes have been proposed in the literature, and these measures are frequently used by market participants, policy makers, and researchers in their analyses. However, risk and uncertainty measures differ across multiple dimensions, including the method of calculation, the underlying outcome (that is, the asset price or macroeconomic variable), and the horizon at which they are calculated. Therefore, in this paper, we review the literature on global risk, uncertainty, and volatility measures drawing on internal and external academic research as well as ongoing monitoring conducted by the Federal Reserve Board's economics divisions to catalog measures by method of data collection, computation, and subject. We first explore a set of non-asset-market-based measures of risk and uncertainty, including news-based and survey-based uncertainty measures of monetary policy and macroeconomic outcomes. We then turn to asset-market-based measures of risk uncertainty for equity prices, interest rates, currencies, oil prices, and inflation.

Keywords: risk, uncertainty, volatility, monetary policy, geopolitical risk, equities, interest rates, exchange rates, commodities, inflation, variance risk premium

JEL classifications: E6, G1, G15

¹ The authors are economists in the Division of International Finance, Board of Governors of the Federal Reserve System, Washington, D.C. 20551 USA. This paper draws from academic research internal and external to the Board of Governors of the Federal Reserve System as well as from ongoing monitoring efforts in the divisions of International Finance, Monetary Affairs, and Research and Statistics. We thank Daniel Covitz, George Eckerd, Eric Engstrom, Gustavo Suarez, Charles Thomas, Min Wei, Beth Anne Wilson, and Jason Wu for their valuable comments and suggestions. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System. Corresponding author: Juan M. Londono. Telephone: +1 202-973-7478. E-mail: juan-miguel.londono-yarce@frb.gov.

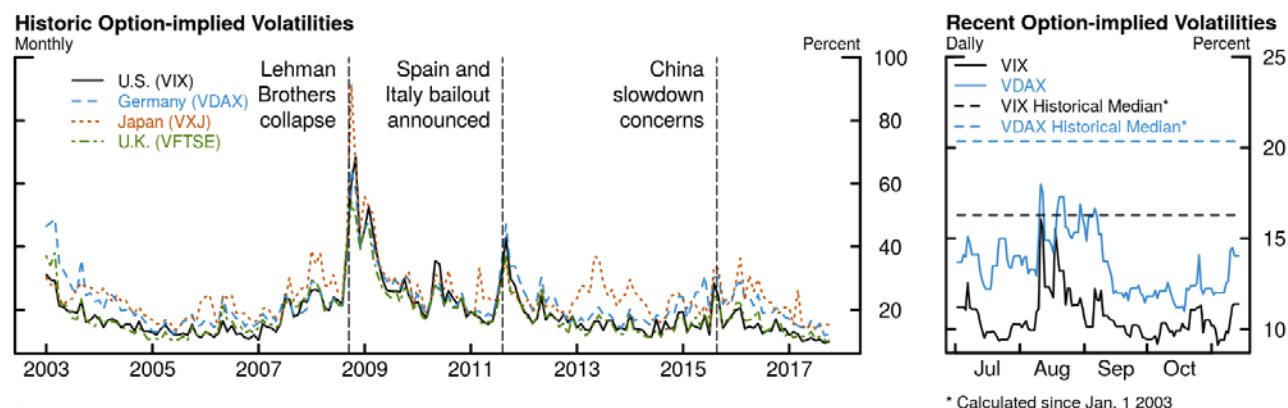
Table of contents

Introduction	2
Chapter 1. Non-Asset-Market Indicators	4
1.1. Economic Policy Uncertainty.....	5
1.2. Monetary Policy Uncertainty	8
1.3. Geopolitical Risk.....	11
1.4. Survey-Based Macroeconomic Uncertainty.....	12
1.5. Econometric Measures of Macroeconomic Uncertainty	14
Chapter 2. Asset-Market Indicators	16
2.1. Realized Volatility.....	17
2.2. Cross-Sectional Distribution of Stock Market Returns.....	18
2.3. Global Risk-On/Risk-Off Index.....	20
2.4. Derivative-Implied Risk and Uncertainty Measures.....	24
2.4.1. Option-Implied Volatility for Equity Indexes.....	25
2.4.2. Variance Risk Premium and Its Components	28
2.4.3. Option-Implied Probability Distribution for Equity Indexes.....	30
2.4.4. Option-Implied Probability Distribution for Individual Stocks.....	33
2.4.5. Interest Rates.....	33
2.4.6. Currency Option-Implied Distributions and Risk Premiums.....	35
2.4.7. Oil Price Option-Implied Probability Distributions and Risk Measures	37
2.4.8. Inflation.....	40
References	43

Introduction

Despite the heightened political risks, equity option-implied volatility remains near historic lows not just in the United States but across most international markets. As shown in Figure 1, global volatility has fallen significantly since January 2016 and in recent months, option-implied volatilities for headline equity indexes in both the U.S. and in Germany have remained below their respective historical medians. Despite the relatively low levels of stock market volatility, many have been struck by the recent global disparity between these measures and the relatively high level of geopolitical and economic policy uncertainty. Motivated by these observations, this taxonomy reviews a broad set of measures that have been proposed in the literature and that are frequently used by market participants, policy makers, and researchers to monitor risk and uncertainty surrounding macroeconomic and financial outcomes.

Figure 1. Option-implied Volatilities for Headline Equity Indexes



Notes: Figure 1 plots model-free (VIX methodology) option-implied volatility, calculated as the weighted average of the implied volatility of options at different degrees of moneyness on each country's representative index. Source: Bloomberg Finance LP.

Risk and uncertainty measures differ across multiple dimensions: method of calculation, the underlying outcome (i.e., asset price or macroeconomic variable), and the horizon at which they are calculated, among others. In this taxonomy, we review the literature on global risk and uncertainty measures. We catalog measures by methods of data collection, computation, and subject, drawing on internal and external academic research as well as ongoing monitoring conducted by the Board's economics divisions. We first explore a set of non-asset market based measures of risk and uncertainty, including news-based and survey-based uncertainty measures for monetary policy and macroeconomic outcomes. We then turn to asset market-based measures

of risk uncertainty for equity prices, interest rates, currencies, oil prices, and inflation. No single measure can be the silver bullet to understanding uncertainty across all markets or outcomes. However, taken together, we present a variety of measures covering a broad spectrum of the sources of risk, uncertainty, and volatility.

CHAPTER 1
NON-ASSET-MARKET INDICATORS

1.1. Economic Policy Uncertainty

One of the most widely used non-asset-market indicators of uncertainty is the economic policy uncertainty (EPU) index developed by Baker, Bloom, and Davis (2016). For the United States, the EPU index is constructed from three components: The first component quantifies policy-related uncertainty by searching the archives of 10 major U.S. newspapers for articles that contain terms related to EPU; the second component gauges uncertainty regarding the federal tax code by counting the number of federal tax code provisions set to expire in future years; and the third component measures disagreement among economic forecasters as an indication of uncertainty. EPU indexes are constructed for almost 20 other countries or country aggregates but are based on only the first component—newspaper articles regarding policy uncertainty.²

For the United States, the news component of the EPU index is constructed by counting the number of articles in 10 leading U.S. newspapers that contain the words “economic” or “economy”; “uncertain” or “uncertainty”; and one or more of “Congress,” “deficit,” “Federal Reserve,” “legislation,” “regulation,” or “White House.” Analogous searches are performed for the other countries.

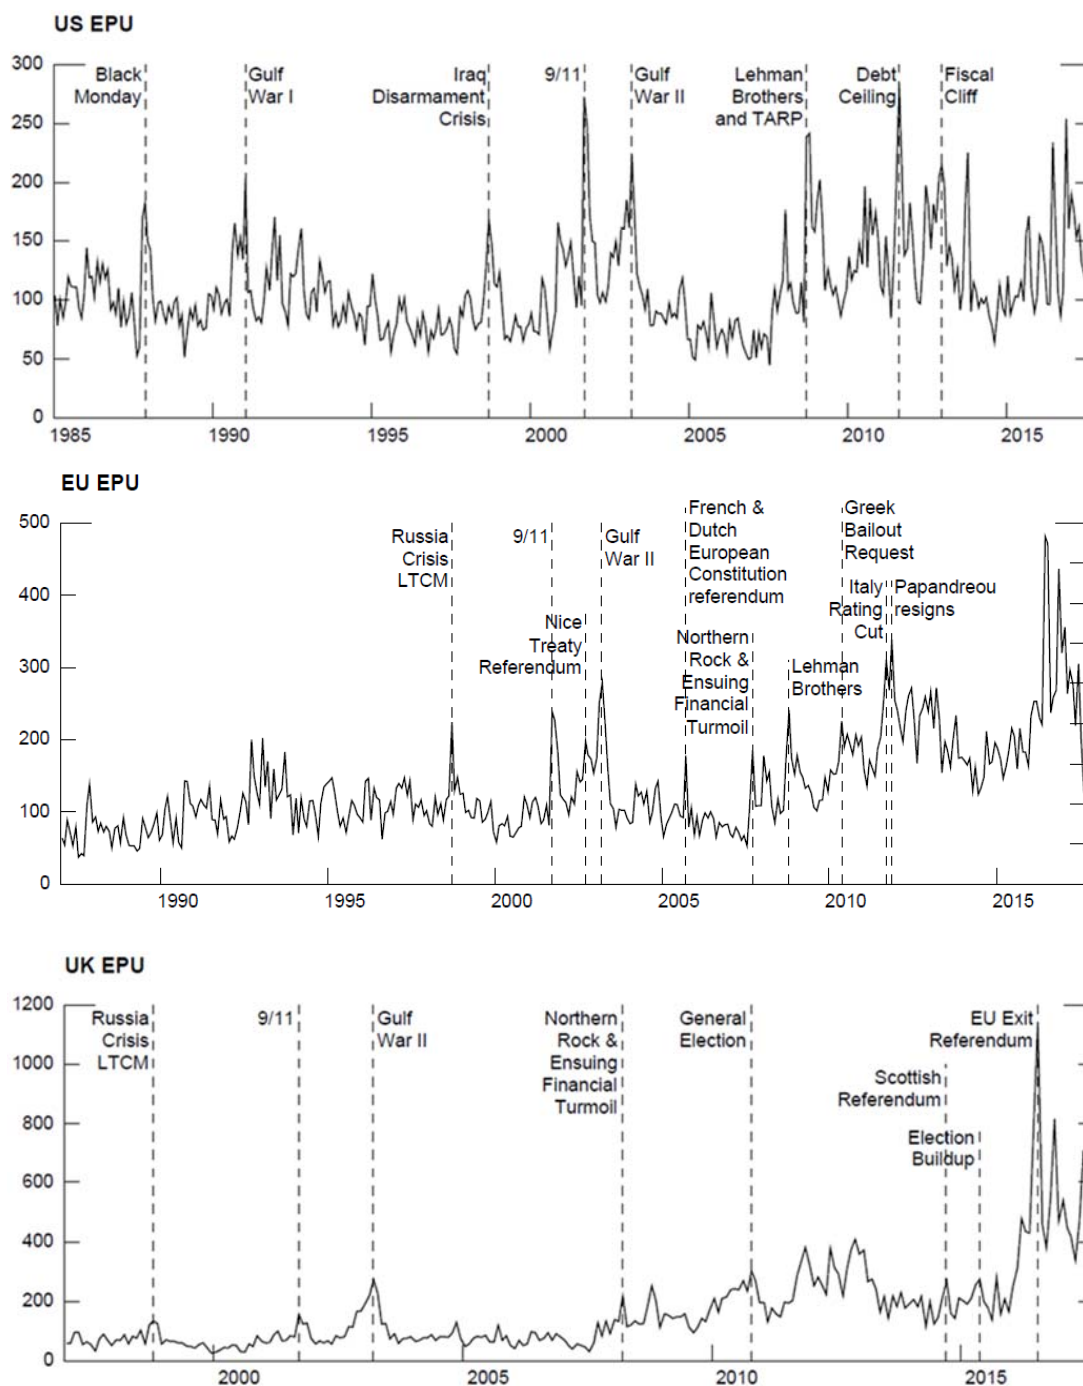
An obvious difficulty with these raw counts is that the overall volume of articles varies across newspapers and time. Thus, Baker, Bloom, and Davis (2016) scale the raw monthly counts for each newspaper by the total number of articles in that newspaper and in that month to produce a monthly EPU series for each newspaper. They scale each newspaper-level series to ensure that each has a unit standard deviation for the 1985 to 2009 period, and then take the average of these 10 monthly series. Finally, they normalize the 10-paper average series to a mean of 100 from 1985 to 2009. The series for all countries are standardized similarly.

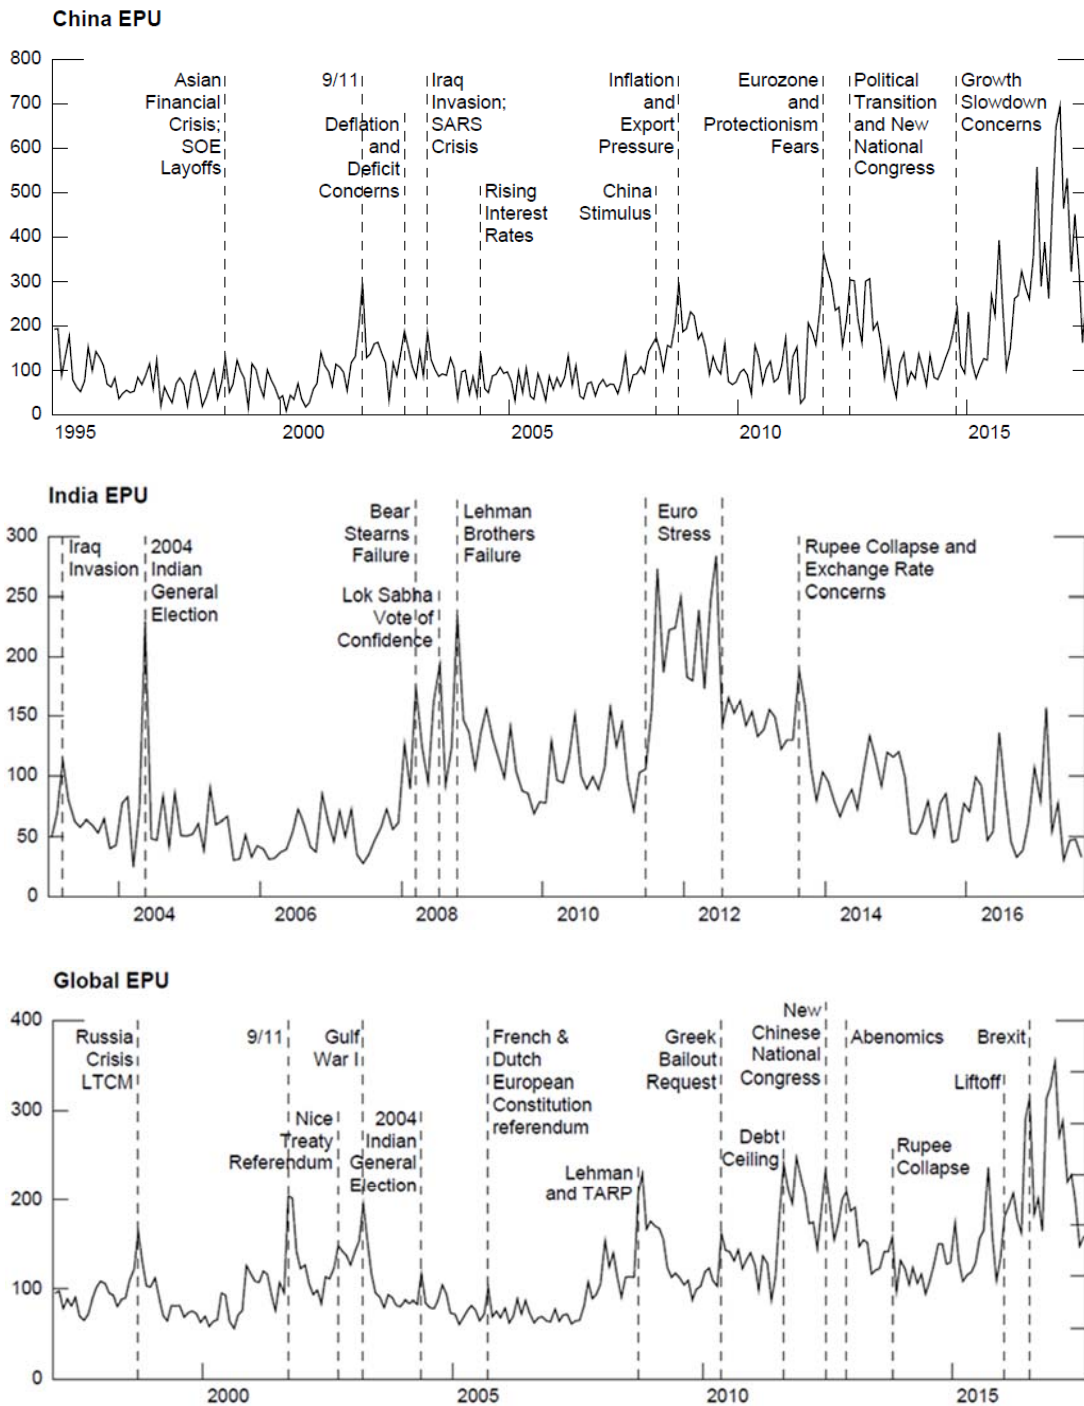
The first panel in figure 2 shows the monthly time series for the U.S. EPU index constructed based on newspaper coverage since 1985. The EPU index shows clear spikes around events and developments that may affect uncertainty, such as the Gulf wars, presidential elections, the terrorist attacks on September 11, 2001, the stimulus debate in early 2008, the Lehman Brothers bankruptcy and the subsequent Troubled Asset Relief Program, or TARP, legislation in late 2008, the summer

² EPU indexes are available for the following countries and country aggregates: the global aggregate, Australia, Brazil, Canada, Chile, China, Europe, France, Germany, India, Ireland, Italy, Japan, Korea, the Netherlands, Russia, Singapore, Spain, Sweden, and the United Kingdom.

2011 debt ceiling dispute, and the battle over the “fiscal cliff” in late 2012. The EPU indexes for the euro area, the United Kingdom, China, and India as well as the global EPU index are also exhibited in figure 2.

Figure 2. Economic Policy Uncertainty (EPU)





Notes: Panels 1 through 6 depict the EPU index for the following areas: the United States, Europe, the United Kingdom, China, India, and the world. Each panel shows some of the significant historical events that could have increased EPU.

Source: www.policyuncertainty.com.

Using the Access World News database of over 2,000 U.S. newspapers, Baker, Bloom, and Davis (2016) also developed subindexes for policy categories by counting the number of articles that not only meet the criteria for inclusion in the EPU index but also contain terms relevant to the specific category in question. For example, they constructed a health care EPU index by searching for articles that discuss rising EPU as well as terms such as “health care,” “Medicaid,” “Medicare,” “health insurance,” “affordable care act,” and “medical insurance reform.” Similarly, they created a trade policy EPU index by searching for words including “import tariffs,” “world trade organization,” and “trade policy” in addition to their baseline trio of search terms. They have other category-specific indexes covering monetary policy, taxes, fiscal policy and government spending, national security, entitlement programs, regulation, financial regulation, and sovereign debt and currency crises.

Researchers have also used EPU indexes to show that policy uncertainty can affect the economy and asset prices. For example, using firm-level data, researchers have shown that policy uncertainty seems to reduce investment and employment, especially in firms that are more dependent on government spending (Gulen and Ion, 2015; Baker, Bloom, and Davis, 2016). At the macro level, researchers have shown that higher policy uncertainty can lead to lower investment, output, and employment (Bachmann, Elstner, and Sims, 2013; Baker, Bloom, and Davis, 2016) as well as reduced monetary policy effectiveness (Aastveit, Natvik, and Sola, 2013). Financial research has also shown that policy uncertainty can increase stock volatility, stock co-movement, and equity premiums (Pastor and Veronesi, 2012, 2013; Brogaard and Detzel, 2015) as well as default risk and credit spreads (Manzo, 2013) and financial intermediation costs (Francis, Hasan, and Zhu, 2014).

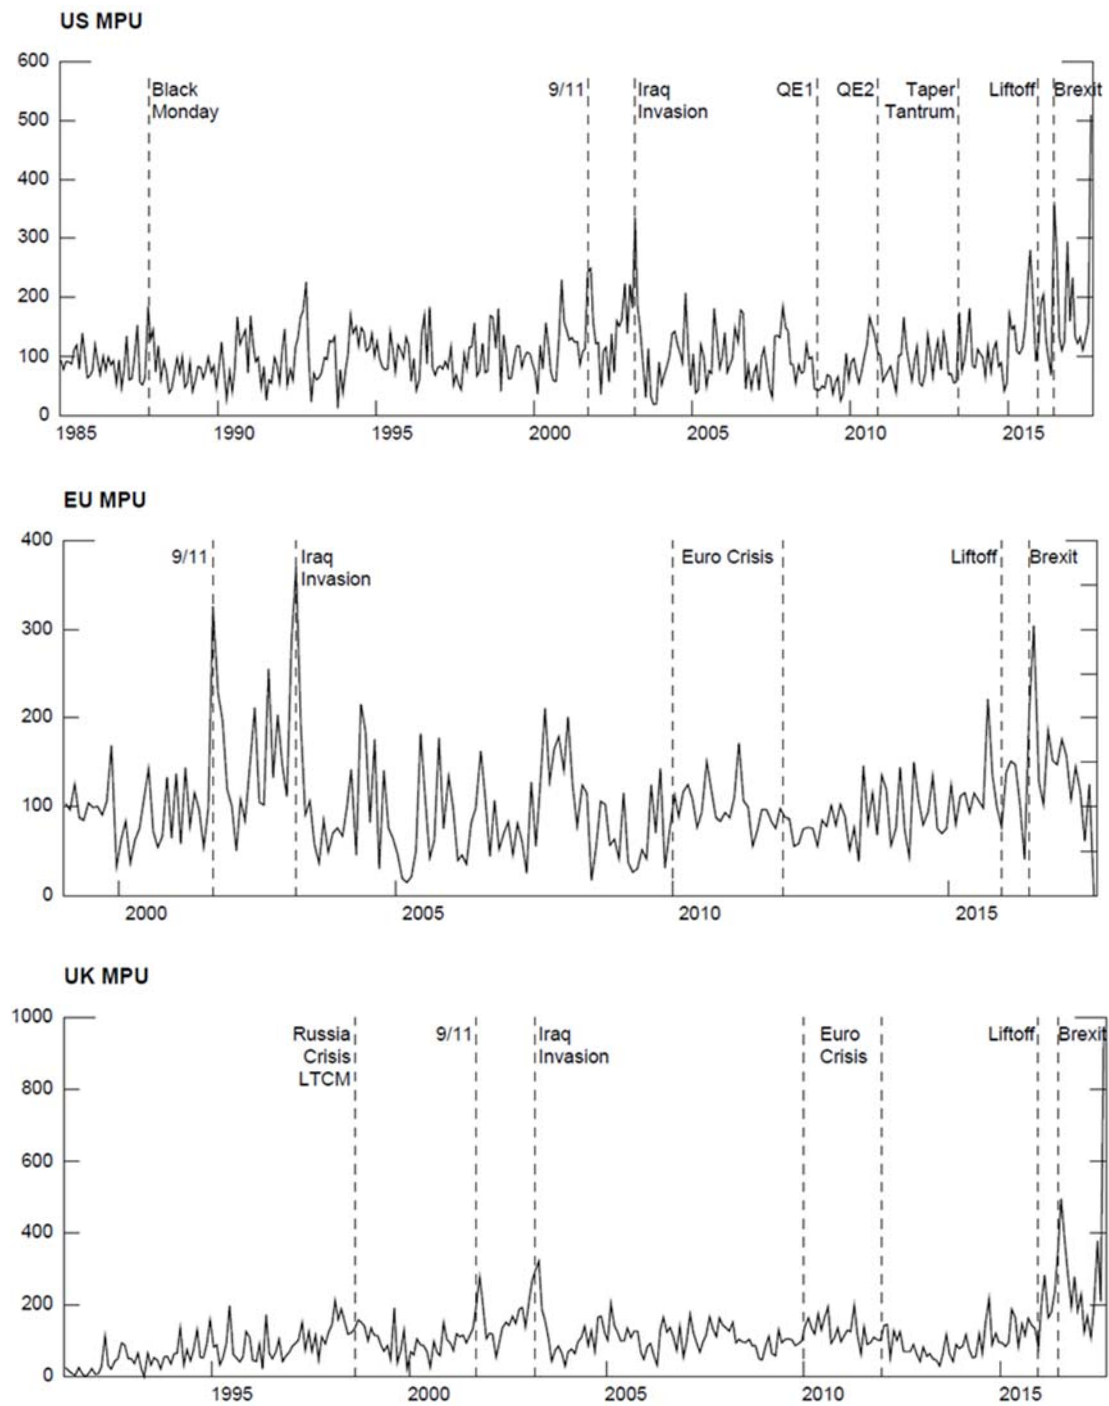
1.2. Monetary Policy Uncertainty

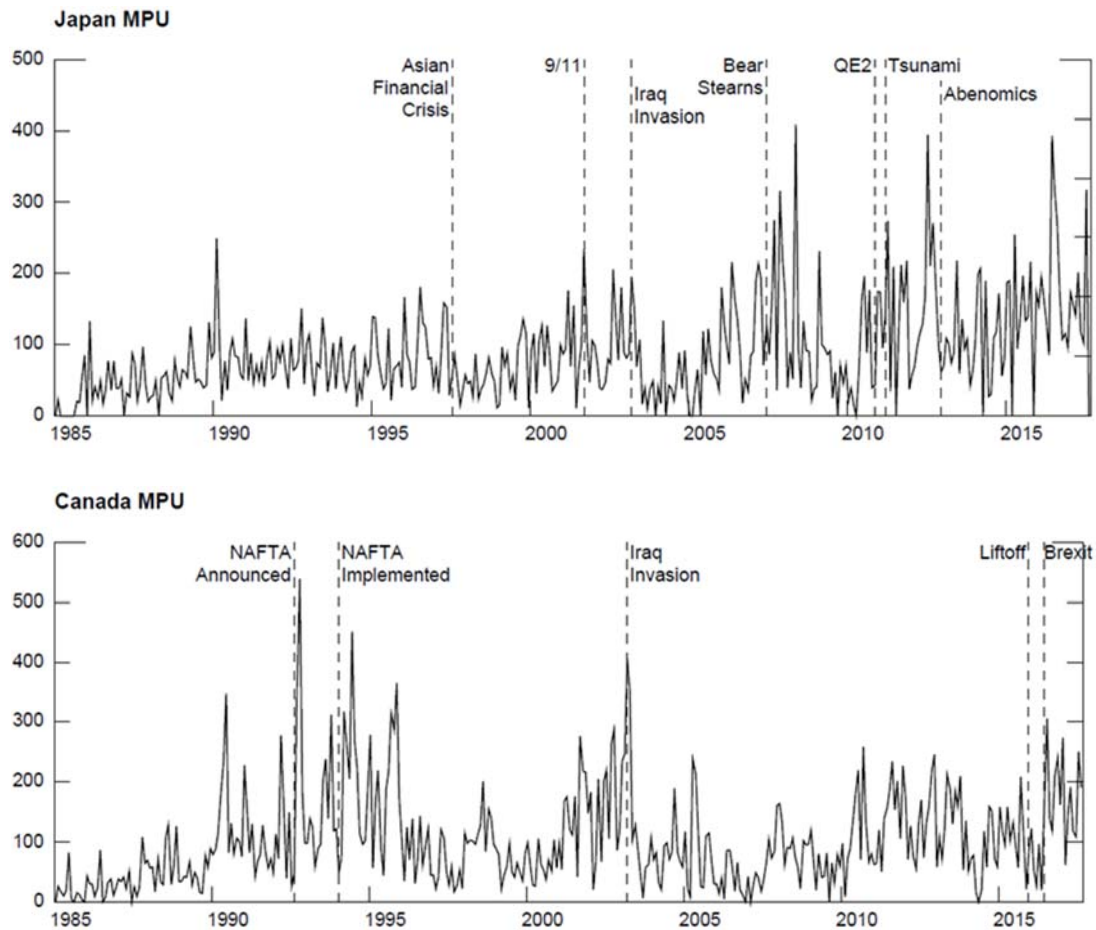
To capture uncertainty related to policies of central banks, Husted, Rogers, and Sun (2017) apply the same text-based methodology as Baker, Bloom, and Davis (2016) to construct an index of monetary policy uncertainty (MPU) by tracking the frequency of newspaper articles related to MPU. For the United States, the MPU index measures the perceived uncertainty surrounding the Federal Reserve Board’s policy decisions and their consequences.

Using ProQuest Newsstand and historical archives as a primary source, Husted, Rogers, and Sun (2017) construct their index by searching for keywords related to monetary policy from the following three sets: (1) “uncertainty” or “uncertain”; (2) “monetary policy(ies),” “interest rate(s),” “federal fund(s) rate,” or “fed fund(s) rate”; and (3) “Federal Reserve,” “the Fed,” “Federal Open Market Committee,” or “FOMC.” The index is normalized following the methodology of Baker, Bloom, and Davis (2016). The narrow word search used for this MPU index gives rise to an index that isolates MPU, relative to the relatively broader word search used for the Baker, Bloom, and Davis (2016) EPU index that results in a broader measure of uncertainty. Husted, Rogers, and Sun (2016) also show that U.S. output and inflation fall and credit costs become tighter following positive shocks to the MPU index.

The first panel in figure 3 shows the time series for the U.S. MPU index since 1985. Notably, the MPU spikes during the taper tantrum in 2013 and right before the beginning of the increase of the Federal Reserve Board’s benchmark rate (liftoff) in 2015. The timing of these spikes relative to policy decisions shows the ability of the index to capture the ex post and ex ante uncertainty of different Federal Open Market Committee decisions. Major macroeconomic events with the capacity to affect monetary policy, like the invasion of Iraq in 2003, also move the index. The MPU indexes for the euro area, the United Kingdom, Japan, and Canada are also provided in figure 3.

Figure 3. Monetary Policy Uncertainty (MPU)





Notes: Panels 1 through 5 depict the MPU index for the following respective areas: the United States, the European Union, the United Kingdom, Japan, and Canada. Each panel shows some of the significant historical events that could have increased MPU.

Source: Husted, L., Rogers, J., and Sun., B. (2016).


1.3. Geopolitical Risk

Caldara and Iacoviello (2016) construct an index that measures geopolitical risk (GPR) based on a tally of newspaper stories that contain a fairly broad set of terms related to geopolitical tensions. The GPR index measures the risk associated with events, such as wars, political tensions, and terrorist acts that affect the normal course of domestic politics and international relations. Caldara and Iacoviello (2016) also propose two indexes that distinguish between geopolitical acts and geopolitical threats. Country-specific GPR indexes are also available for the following countries: the United States, Argentina, Brazil, China, Colombia, India, Indonesia, Israel, Korea, Malaysia, Mexico, the Philippines, Russia, Saudi Arabia, South Africa, Thailand, Turkey, and Venezuela.

The GPR index is constructed by counting the occurrence of words related to geopolitical tensions in leading international newspapers. In particular, the GPR index reflects automated text searches in the electronic archives of 11 national and international newspapers for articles that contain several keywords, including “risk of war,” “terrorist threats,” and “geopolitical tensions.”

Figure 4 shows the time series for the GPR index since 1990. The GPR index spikes during the Gulf War, the 2003 invasion of Iraq, and on September 11, 2001, as well as in periods of increased bilateral tensions. Moreover, the index has remained heightened since the beginning of 2017, a sign of mounting tensions between the new U.S. administration and their global partners as well as growing instability in the Levant. Caldara and Iacoviello (2016) show that increased geopolitical risk leads to declines in real activity and is associated with increases in the VIX, lower oil prices, and higher corporate credit spreads.

Figure 4. Geopolitical Risk Index (GPR)



Notes: Figure 4 depicts the GPR index plotted against the VIX. The figure shows some of the significant historical events that could have increased GPR.

Source: <https://www2.bc.edu/matteo-iacoviello/gpr.htm>.

1.4. Survey-Based Macroeconomic Uncertainty

Economic surveys sometimes provide useful information about survey participants’ probabilistic assessments of future events. Economic surveys typically ask about point predictions (that is, mean or mode expectation) of future events of each individual respondent. Aggregating these individual responses allows us to calculate a measure of the dispersion across the respondents regarding the point prediction. However, this aggregation does not provide meaningful information regarding the uncertainty that each individual may attach to his or her point forecast.


Therefore, some surveys also ask about the uncertainty surrounding point forecasts. Such surveys allow us to calculate both the dispersion across individuals and the uncertainty of each individual regarding his or her own forecast.³ For example, the Survey of Professional Forecasters, a quarterly publication, asks respondents to provide both point estimates and probabilistic assessments of the outlook for inflation and 10-year interest rates, from which we can construct aggregate proxies for both dispersion and uncertainty.

Scotti (2016) uses macroeconomic news and survey forecasts to construct an ex post, realized measure of uncertainty about the state of the economy. The author's macroeconomic uncertainty index measures uncertainty based on weighted averages of economic data surprises, which are measured by examining deviations of recent economic data releases from consensus expectations. The index is calculated using the weighted average of the squared surprises from a sample of macroeconomic variables where the surprises are measured as the differences between the actual data and the median Bloomberg forecasts an hour before the data release. A dynamic factor model is employed to estimate monthly business condition indexes and compute the weights representing the contribution of the economic indicators to these business condition indexes. Those weights are then used to average the squared surprises to construct the uncertainty index.

Figure 5 shows Scotti's uncertainty index since 1990. The index closely follows the VIX leading up to and after the global financial crisis of 2008 and Brexit. Scotti (2016) shows that a higher uncertainty index is associated with lower real activity.

³ Both the dispersion across individuals and the uncertainty that each individual feels are relevant for studying market surprises and the effects of realized future events. For a detailed discussion of this issue, see the speech by Vice Chairman Stanley Fischer on April 17, 2017, "Monetary Policy Expectations and Surprises," delivered at the Columbia University School of International and Public Affairs, <https://www.federalreserve.gov/newsevents/speech/fischer20170417a.htm>.

Figure 5. Scotti's Uncertainty Index



Notes: Figure 5 depicts Scotti's uncertainty index plotted against the VIX. The figure shows some of the significant historical events that could have increased economic uncertainty.

Source: <https://sites.google.com/site/chiarascottifrb/research/surprise-and-uncertainty-indexes>.

1.5. Econometric Measures of Macroeconomic Uncertainty

Jurado, Ludvigson, and Ng (2015) construct indexes of macroeconomic uncertainty using the uncertainty around objective statistical forecasts for hundreds of economic series.⁴ They use a monthly dataset comprising the information from hundreds of macroeconomic indicators to construct direct econometric estimates of time-varying macroeconomic uncertainty. Their key insight is that macroeconomic uncertainty can be constructed as an appropriately weighted average of the forecast error variance of all the included macroeconomic indicators.


Formally, Jurado, Ludvigson, and Ng define the h -period ahead uncertainty in a single variable as the conditional volatility of the unforecastable component of the future value of the variable (that is, the difference between the future value of the variable and its expectation based on the information available at time t). The aggregate uncertainty at the macro level is the average of the uncertainty measures across all macro variables. This measure differentiates uncertainty from traditionally used measures of volatility, such as conditional volatility. Conditional volatility—as embedded in common measures of stock market volatility—does not necessarily remove the forecastable component of a time series, while the Jurado, Ludvigson, and Ng index (henceforth

⁴ Jurado, Ludvigson, and Ng (2015) use a similar methodology to construct a measure of financial uncertainty. The financial uncertainty measure is explored further in Ludvigson, Ma, and Ng (2016).

referred to as the JLN index) does so by incorporating a large number of indicators into the forecasting model for each individual time series.

Relative to the Baker, Bloom and Davis (2016) EPU measures, the JLN index (shown in figure 6) is more persistent and exhibits fewer spikes. Moreover, this measure of uncertainty doubles during recessions. Additionally, when embedded in a vector autoregression, increases in the JLN index are associated with large declines in real activity.

Figure 6. Jurado, Ludvigson, and Ng (JLN) Macroeconomic Uncertainty Index



Notes: Figure 6 depicts the JLN macroeconomic uncertainty index. The figure shows some significant historical events that could have increased economic uncertainty.

Source: <https://www.sydneyludvigson.com/data-and-appendixes/>.

CHAPTER 2

ASSET-MARKET INDICATORS

2.1. Realized Volatility

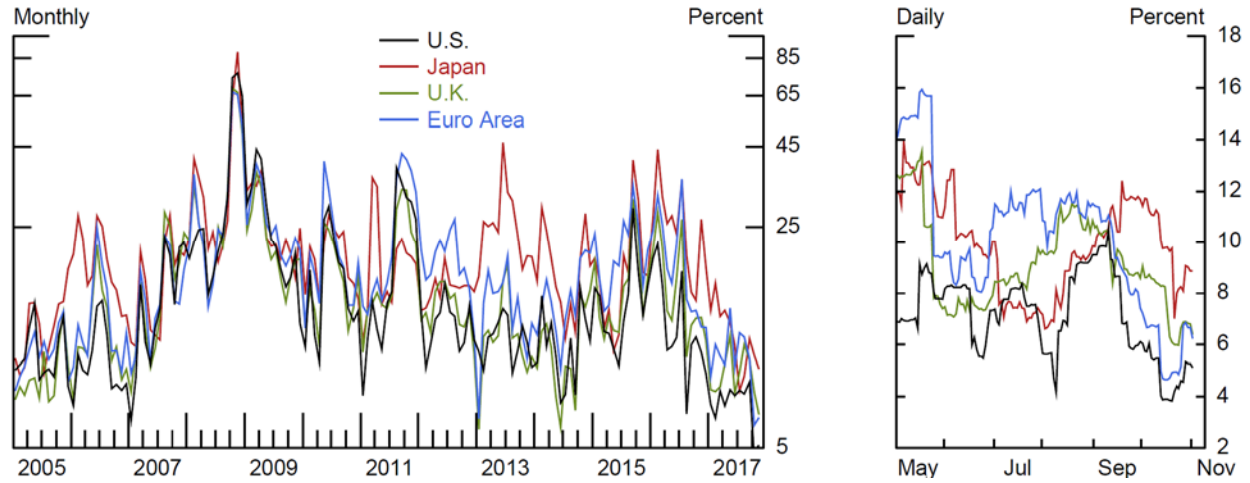
Before 2000, volatility was largely modeled using parametric methods, such as generalized autoregressive conditional heteroskedasticity (GARCH) or stochastic volatility (SV) methods. Although these methods are extensively used to model financial and macroeconomic time series for daily to annual frequencies, they have several limitations. For example, GARCH-based models generally fail to capture the magnitude of sudden increases in the level of volatility, such as the increase that occurred on October 19, 1987. These models are also relatively hard to extend to and implement in multivariate settings.

Pioneered by Andersen, Bollerslev, Diebold, and Labys (2001, 2003), realized volatility (RV)—defined as the scaled sum of squared daily returns—offers a nonparametric alternative to traditional parametric volatility measures. RV estimators are feasible in multivariate applications and can separate the volatility contributions of jumps from continuous changes in asset prices. In addition, they are flexible and easy to implement.

The properties of RV-style estimators are well documented in the literature, and they are routinely used for forecasting volatility (Alizadeh, Brandt, and Diebold, 2002; Corsi, 2009; Patton and Shephard, 2016; among many others) and for predicting returns (Bollerslev and Zhou, 2006; among many others). However, while RV-style measures have proved successful in predicting future volatility, their ability to predict financial returns is somewhat limited.

Figure 7 shows the time series of the Standard & Poor's (S&P) 500 index RV (the black line) between January 2005 and April 2017. The figure also shows time series of RV-style indexes for headline equity indexes for Japan, the United Kingdom, and the euro area. These RV measures are based on daily returns (as opposed to intradaily data), and seem to be particularly high around episodes of market uncertainty, which are usually associated with unexpected news or events. Not surprisingly, the most notable volatility spikes in this sample occurred around the collapse of Lehman Brothers in 2008 when the RV reached a maximum of around 80 percent. These measures were also particularly high in July and August of 2011 at the peak of the euro-area crisis. The later episodes of relatively high volatility spikes occurred between the second quarter of 2015 and the third quarter of 2016. Interestingly, the U.S. RV has remained at near-historical lows since the 2016 presidential election.

Figure 7. Realized Volatility (RV)



Notes: This figure plots the RV for equity indexes in the United States (black), the United Kingdom (green), Japan (red), and the euro-area (blue). RV is calculated as the square root of the sum of daily squared returns over the last month (22 days).

Source: Bloomberg Finance LP.

2.2. Cross-Sectional Distribution of Stock Market Returns

The RV measures in section 2.1 are calculated using time series of data on the aggregate stock market. Additional measures of volatility can be computed by exploiting the distribution of stock returns across firms at each point in time (for example, all stocks in the S&P 500 index). For instance, Bloom (2009) and Christiano, Motto, and Rostagno (2014) use the variance across individual stock returns at each point in time as a measure of cross-sectional uncertainty and show that exogenous shocks to these measures are important sources of business cycle fluctuations.

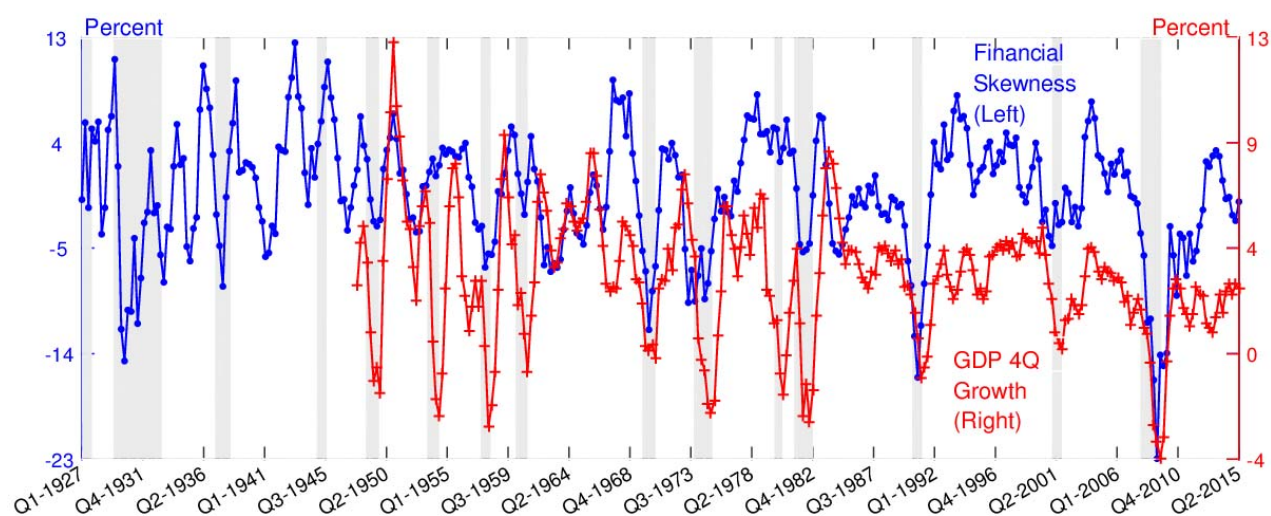
Higher-order moments of the cross-sectional distribution of stock returns can also provide useful information about the economic cycle. In particular, Ferreira (2017) focuses on the skewness of the distribution of log returns across firms and assesses the balance between upside and downside risks.⁵ The author shows that financial skewness—the skewness of the cross-sectional distribution of stock returns of financial firms—not only closely tracks the business cycle, but predicts

⁵ Skewness is calculated as the difference between the magnitudes of two equally probable tails of the distribution of returns: the right tail, measured by the difference between the 95th and 50th percentiles, and the left tail, measured by the difference between the 50th and 5th percentiles.

economic activity better than many well-known bond spreads (for example, the term spread and the spread in Gilchrist and Zakrajšek (2012)) and other cross-sectional moments.

Figure 8 sheds light on the relationship between financial skewness and the business cycle since 1926. The figure shows that sharp decreases in financial skewness—that is, when the left tail of the distribution becomes larger than the right tail—coincide with slowing GDP growth and recessions, especially over the past two decades.⁶ Financial skewness was at its lowest during the recent global financial crisis, the recession following the savings and loan crisis of the 1990s, and the Great Depression. During the global financial crisis, the measure of financial skewness started to plummet well before the drop in economic activity, consistent with the predictive analysis provided by Ferreira (2017).

Figure 8. Financial Skewness and Economic Activity



Notes: Shading indicates National Bureau of Economic Research recessions. The financial skewness shown as a four-quarter moving average.

Source: Ferreira (2017).

Ferreira (2017) then shows evidence supporting the interpretation of financial skewness as a measure of the balance of risks across economy-wide investment projects as well as vulnerabilities of the financial sector. Moreover, the author estimates that shocks to financial skewness have sizable effects on economic activity, credit growth, and corporate credit spreads.

⁶ The National Bureau of Economic Research classification is used to define recession periods.

2.3. Global Risk-On/Risk-Off Index

The global risk-on/risk-off index (ROI) is a proxy for risk appetite in global financial markets. It is constructed from changes in a variety of asset prices, including those of bonds, equities, exchange rates, and gold. During “risk-on” periods, investors typically reallocate their portfolios away from safe-haven assets (for example, Treasury securities, bunds, gold) toward riskier assets (for example, equities, high-yield bonds, emerging market bonds). The ROI measures the extent to which these portfolio shifts lead to increases in the prices of risky assets relative to those of safe assets. Positive values indicate risk-on days, while negative values indicate “risk-off” days. The ROI is global in scope, covers the major asset classes, captures high-frequency changes in risk appetite, and is intuitive and easy to compute.

The ROI is an equally weighted sum of changes in 15 assets or indexes, with their respective signs aligned such that positive changes are consistent with risk-on behavior and negative changes with risk-off behavior. Changes are scaled by their respective historical standard deviations. An index value of 2 means that the components of the index increased in the risk-on direction by 2 standard deviations, on average. The magnitude of the ROI is larger when the changes in the underlying assets are directionally consistent with each other.⁷

The ROI includes the following components:

Equities:

Percent change in the MSCI (Morgan Stanley Capital International) local currency advanced economies equity index (↑ equity prices = risk on)

Percent change in MSCI local currency emerging market economies (EME) equity index (↑ equity prices = risk on)

Percent change in U.S. S&P 500 index (↑ equity prices = risk on)

(-1) * Basis point change in S&P 500 VIX (↓ implied volatility = risk on)

(-1) * Basis point change in Dow Jones Euro Stoxx implied volatility (VSTOXX) index (↓ implied volatility = risk on)

⁷ With 15 components, idiosyncratic noise largely cancels out.

Bonds:

Basis point change in 10-year Treasury yield (↑ yield = risk on)

Basis point change in 10-year bund yields (↑ yield = risk on)

(-1) * Basis point change in U.S. high-yield bond spread (↓ spread = risk on)

(-1) * Basis point change in euro-area high-yield bond spread (↓ spread = risk on)

(-1) * Basis point change in J.P. Morgan's EMBI+ (Emerging Markets Bond Index Plus) spread (↓ spread = risk on)

Exchange rates and gold:

Percent change U.S. dollar (USD)–yen exchange rate (yen depreciation = risk on)

Percent change Swiss Franc (CHF)–Euro exchange rate (CHF depreciation = risk on)

(-1) * Percent change in emerging market dollar index (emerging market currency appreciation = risk on)⁸

(-1) * Basis point change in implied volatility of EME currencies⁹ (↓ implied volatility = risk on)

(-1) * Percent change in gold spot price (↓ gold price = risk on)

The standard ROI is based on daily changes and, thus, captures high-frequency changes in risk appetite. It is also mean reverting with a mean of zero and a standard deviation of 0.6. The average duration of risk-on (positive ROI) and risk-off (negative ROI) episodes are about 2 days. The longest risk-on and risk-off episodes each lasted 13 days. To capture longer-term trends in risk sentiment, another version of the index is calculated based on changes over a rolling window of 120 business days (ROI-120).

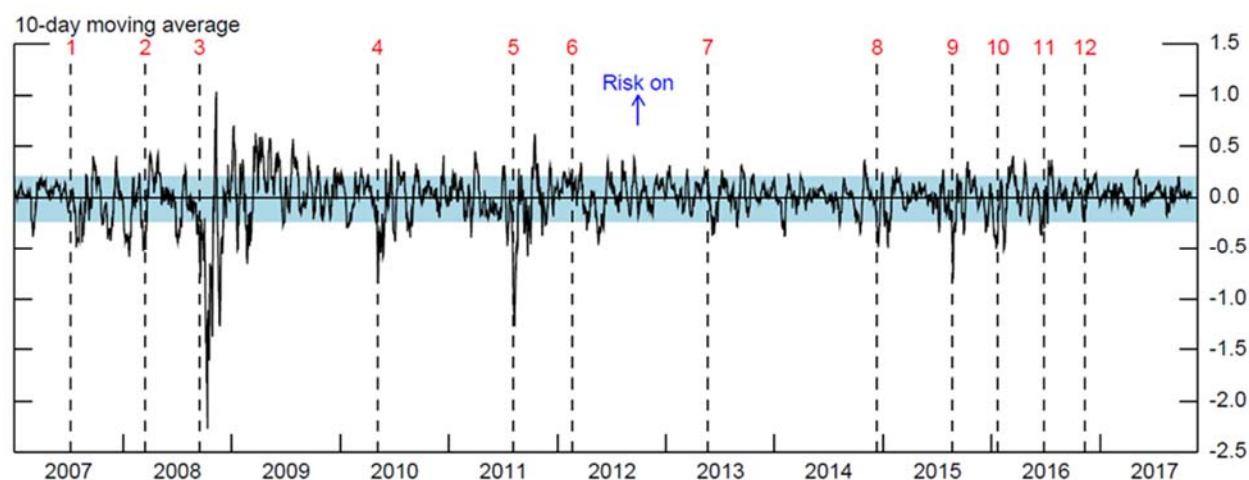
Figure 9 shows the 10-day moving average of the daily ROI between January 2007 and October 2017. Some of the worst risk-off days included June 24, 2016, the day after the Brexit referendum, August 8, 2011, the day S&P downgraded U.S. debt, May 6, 2010, the U.S. “flash crash,” and the days following the Lehman Brothers bankruptcy in the fall of 2008. During all of these episodes, all 15 components of the ROI moved in the risk-off direction. Some of the best risk-on days

⁸ Specifically, the ROI uses the Nominal OITP (Other Important Trading Partners) Dollar Index available at https://www.federalreserve.gov/releases/h10/summary/indexo_b.htm.

⁹ The ROI uses the J.P. Morgan Emerging Market Currency Volatility Index, computed from three-month options.

followed the worst risk-off days. The best risk-on days include October 27, 2011, when a European Union summit reached agreement for Greek debt relief, May 10, 2010, when the European Central Bank announced measures to support peripherals, September 19, 2008, when the Federal Reserve Board and the U.S. Treasury announced measures to improve market functioning), and October 13, 2008, when a number of central banks took coordinated action to boost liquidity and funding to financial institutions.

Figure 9. Daily Global Risk-On/Risk-Off Index (ROI)



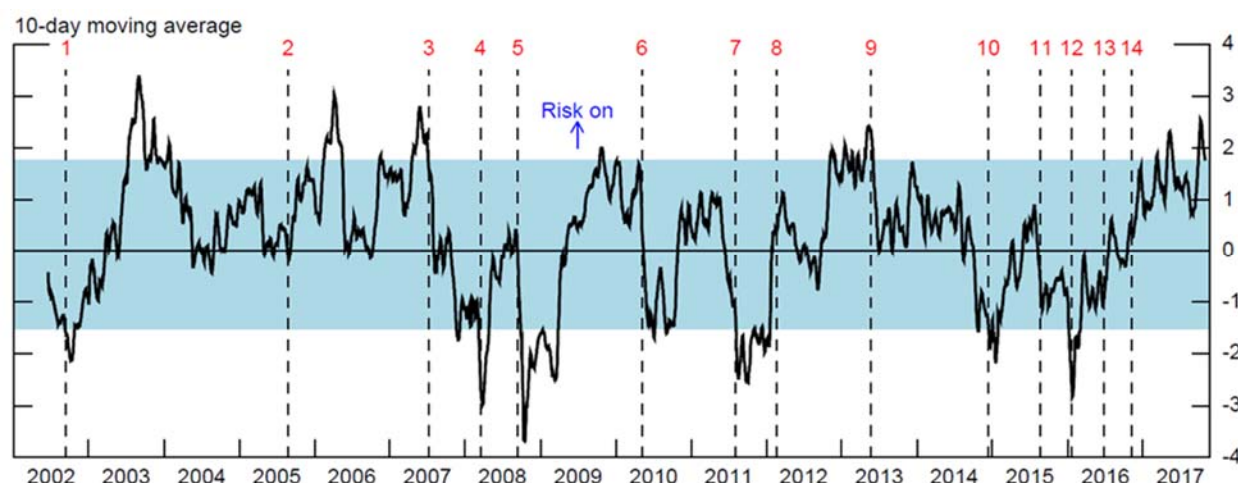
Notes: Blue shading denotes the range spanned by the 10th and 90th percentiles of the distribution. Vertical lines denote the following dates:

1. July 10, 2007: Moody's downgrades U.S. residential mortgage-backed securities
2. March 16, 2008: J.P. Morgan acquires Bear Stearns
3. September 15, 2008: Lehman Brothers bankruptcy
4. May 6, 2010: U.S. flash crash
5. August 5, 2011: Standard & Poor's downgrades U.S. long-term debt
6. February 21, 2012: Eurozone deal on second Greek rescue package
7. May 22, 2013: Federal Reserve Chairman Bernanke suggests eventual tapering of asset purchases
8. December 12, 2014: International Energy Agency cuts global oil demand forecast, sharp declines in crude oil prices
9. August 24, 2015: Shanghai Composite Index declines 8.5 percent, global equities decline on global growth concerns
10. January 20, 2016: Brent crude falls to \$26 per barrel amid global growth concerns
11. June 24, 2016: Brexit referendum
12. November 8, 2016: U.S. presidential election

Source: Federal Reserve Board staff calculations using data from Bloomberg Finance LP.

Figure 10 shows the 10-day moving average of the 120-day version of the ROI between July 2002 and October 2017. This measure is useful for gauging gradual improvement or deterioration in risk appetite over longer stretches of time. By this measure, risk appetite has been improving since early 2016.

Figure 10. 120-Day Global Risk-On/Risk-Off Index (ROI-120)



Notes: Blue shading denotes the range spanned by the 10th and 90th percentiles of the distribution. Vertical lines denote the following dates:

1. September 12, 2002: U.S. President George Bush signals possible military action against Iraq in speech to U.N. general assembly
2. August 25, 2005: Hurricane Katrina hits U.S. Gulf Coast
3. July 10, 2007: Moody's downgrades U.S. residential mortgage-backed securities
4. March 16, 2008: J.P. Morgan acquires Bear Stearns
5. September 15, 2008: Lehman bankruptcy
6. May 6, 2010: U.S. flash crash
7. August 5, 2011: Standard & Poor's downgrades U.S. long-term debt
8. February 21, 2012: Eurozone deal on second Greek rescue package
9. May 22, 2013: Federal Reserve Chairman Bernanke suggests eventual tapering of asset purchases
10. December 12, 2014: International Energy Agency cuts global oil demand forecast, sharp declines in crude oil prices
11. August 24, 2015: Shanghai Composite Index declines 8.5 percent, global equities decline on global growth concerns
12. January 20, 2016: Brent crude falls to \$26 per barrel amid global growth concerns
13. June 24, 2016: Brexit referendum
14. November 8, 2016: U.S. presidential election

Source: Federal Reserve Board staff calculations using data from Bloomberg Finance LP.

Previous literature has found that a lower risk appetite can result in higher cost of capital (and lower investment), as investors demand higher excess returns for holding risky assets, such as equities and bonds.¹⁰ Investment banks have also developed and relied on risk appetite measures to inform trading strategies.¹¹ Extremely bullish risk sentiment is often associated with risky asset prices overshooting their long-run trends. Sudden deteriorations in risk appetite have been associated with contagion in global financial markets (Kumar and Persaud, 2002). The literature on risk appetite indexes has also shown that oscillations in risk sentiment are related to global growth and liquidity cycles (Rey, 2015).

2.4. Derivative-Implied Risk and Uncertainty Measures

Derivative prices reflect investor preferences as well as investor beliefs about the likelihood of future realizations of the underlying asset's price. For instance, buying a put option is a profitable strategy only if the price of the underlying asset falls below a certain threshold—the strike price of the option. The price of a put option will increase with the probability an investor assigns to an outcome in which the price of the asset will drop below its strike price, and will also increase with the value an investor places on a positive return in the event of a price drop. Thus, at any point in time, the prices of derivatives at different strikes contain commingled information about the probabilities assigned to each possible market outcome as well as investor preferences.

Derivative-implied distributions allow us to calculate derivative-implied moments, such as the derivative-implied volatility or skewness, as well as the cost of insurance against any potential outcome (for example, a price drop of a certain magnitude).

The derivative-implied distribution used to generate these moments is often referred to as the risk-neutral distribution because, by construction, this is the probability measure that makes the expected return on a risky investment equal to the risk-free rate. It is not called a risk-neutral measure because we assume that agents are risk neutral, but rather because, under this measure, probabilities are calculated as though agents only cared about the mean return. Because investors

¹⁰ As investors become more risk averse, equity risk premiums rise, translating into a higher cost of capital. Barro (1990) also finds that lagged changes in real stock market prices predict future investment.

¹¹ See, for example, Credit Suisse (2004), Misina (2003), and Illing and Aaron (2005).

are not risk neutral in most cases, derivative-implied distributions contain information about risk premiums.

Risk-neutral or derivative-implied probabilities are different from actual (usually referred to as physical) probabilities. Typically, the actual or physical distribution cannot be known. However, if we assume that the process is stationary (that is, tomorrow's draw will come from the same distribution as historical draws), then the physical distribution can be estimated from historical realizations. Comparing the estimated physical distribution with the derivative-implied distribution can provide some information about investors' risk preferences—that is, about investors' outcome-specific preferences, such as their preference for having positive returns in one state of the economy (for example, a large drop in asset prices) versus another. For example, if the risk-neutral distribution systematically has wider tails than the physical distribution (that is, more probability assigned to extreme market outcomes), we can infer that either investors systematically overestimate the probability of tail events or that their estimations are correct but they particularly value positive returns in those tail events.

2.4.1. Option-Implied Volatilities for Equity Indexes

Option-implied volatilities for headline equity indexes are calculated using the methodology in Carr and Madan (1998) and Britten-Jones and Neuberger (2000). For each underlying index, implied volatility is calculated as a weighted average of the price of put and call options that expire in more than 23 days but less than 37 days. The weight assigned to each option depends on its strike and its maturity, and is intended to generate a portfolio of options that isolates the expected volatility of the underlying equity index at the 30-day horizon (see Carr and Madan, 1998; and Britten-Jones and Neuberger, 2000).¹²

Option-implied volatilities are available for headline equity indexes of the following countries: the United States (S&P 500), Germany (DAX 30), Japan (Nikkei 225), the United Kingdom (FTSE 100), Switzerland (SMI), the Netherlands (AEX 25), and France (CAC 100). Implied-volatility is also available for the euro area (Euro Stoxx 50). The U.S. option-implied

¹² Option-implied volatilities are also available at horizons other than 30 days. For example, the S&P 500 VIX Short-Term Futures Exchange Traded Notes (ticker: VXX) is similar to the VIX, the S&P 500 option-implied volatility, but with a 9-day horizon.

volatility, the VIX, is perhaps the most popular derivative-implied risk measure and is frequently used by researchers and market participants to gauge fear or uncertainty with respect to the U.S. equity market and even with respect to global equity markets. Although there has been extensive research on the usefulness of the VIX as a tool to monitor equity and other financial asset markets, its informational content is often misunderstood and abused. The VIX and equivalent measures for foreign equity markets are formally defined as the risk-neutral expectation of the volatility of the equity index over the next 30 days.

While analogous measures for longer horizons are also available, the 30-day measure is the most widely used because of the relatively high liquidity for the options around this horizon. This relatively short horizon implies that this index likely does not capture expected volatility beyond the 30-day horizon, and this short horizon could be one possible driver of the discrepancy between the low readings of the VIX observed in early 2017 and the higher perceived policy uncertainty at that time.

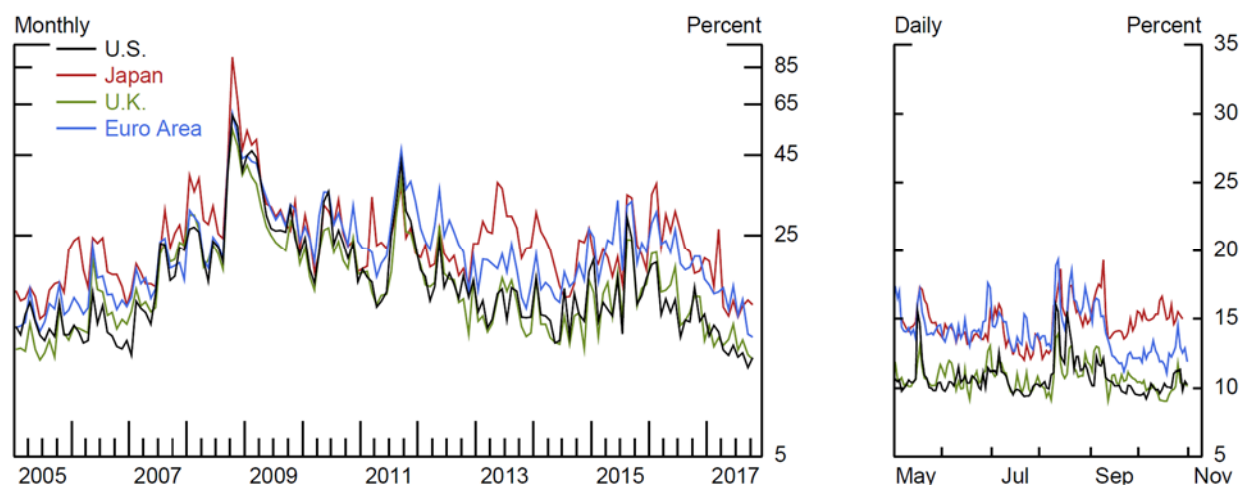
Next, we briefly discuss longer-horizon financial volatility measures. Variance swap contracts with maturities ranging from one month to two years are traded as over-the-counter assets.¹³ Variance swap contracts allow us to examine the expectations of changes in market volatility beyond the 30 days captured by the VIX. In comparison with the VIX, variance swaps rose less during the global financial crisis; however, they did not return to their pre-crisis levels as rapidly. Variance swap levels seem to follow a pattern more similar to a macroeconomic uncertainty index, such as the EPU index. For examples of use of these contracts in research see Dew-Becker, Giglio, Le, and Rodriguez (2017), Amengual and Xiu (2017), and Ait-Sahalia, Karaman, and Mancini (2014), among others.

Figure 11 shows the time series for the VIX (the black line) between January 2005 and May 2017. Also shown are the equivalent option-implied volatility measures for the headline equity indexes of Japan, the United Kingdom, and the euro area. The VIX seems to be particularly high around episodes of high market uncertainty, usually associated with unexpected news or events, which is why the VIX is commonly known as the “investor fear gauge” (Whaley, 2000). Not surprisingly,

¹³ Variance swaps are contracts in which one party pays a fixed amount at maturity, which we refer to as the price of the variance swap, in exchange for a payment equal to the sum of squared daily log returns of the underlying asset (in this case, S&P 500 returns) occurring until maturity.

the most notorious VIX spikes in this sample occurred around the collapse of Lehman Brothers in 2008 when the VIX reached a maximum of around 80 percent. The VIX was also particularly high in July and August 2011, at the peak of the euro-area crisis. The most recent episodes of relatively high VIX realizations occurred in August 2015 and February 2016, two times when the Chinese equity market experienced substantial losses, and June 2016, with the unexpected results of the U.K. referendum campaign. Interestingly, the VIX has remained at near-historical lows since the 2016 presidential election and has only ticked up recently (right panel) following geopolitical tensions involving Syria and North Korea. The figure also suggests that option-implied indexes are highly correlated across countries and tend to spike simultaneously (see Londono, 2016).

Figure 11. Option-Implied Volatilities for Headline Equity Indexes



Notes: Figure 11 plots model-free (VIX methodology) option-implied volatility, calculated as the weighted average of the implied volatility of options at different degrees of moneyness on each country's representative index.

Source: Bloomberg Finance LP.

The informational content of option-implied volatilities, especially the VIX, has been extensively explored in the literature. In particular, the VIX has been shown to be a useful predictor of future RV (see, for example, Whaley, 2009; and Jiang and Tian, 2005). Additionally, as mentioned previously, the price of derivatives at different strikes contains information about the probability assigned to each possible market outcome as well as investor preferences (for example, having positive returns in one state of the economy). It can be difficult to disentangle how movements in the VIX reflect changes in expected volatility or changes in attitudes or preferences. However, as mentioned previously, under some assumptions (for example, the stationarity assumption that

tomorrow's draw will come from the same distribution as historical draws), we can extract the component of the VIX related to investor attitudes or preferences toward volatility. This concept of the variance risk premium, which will be explained further in section 2.4.2, is usually linked to risk aversion (see, for example, Bollerslev, Tauchen, and Zhou, 2009; and Bekaert, Hoerova, and Lo Duca, 2013) and has been documented to have predictive power for international equity index returns (Londono, 2016; Bollerslev, Marrone, Xu, and Zhou, 2014).

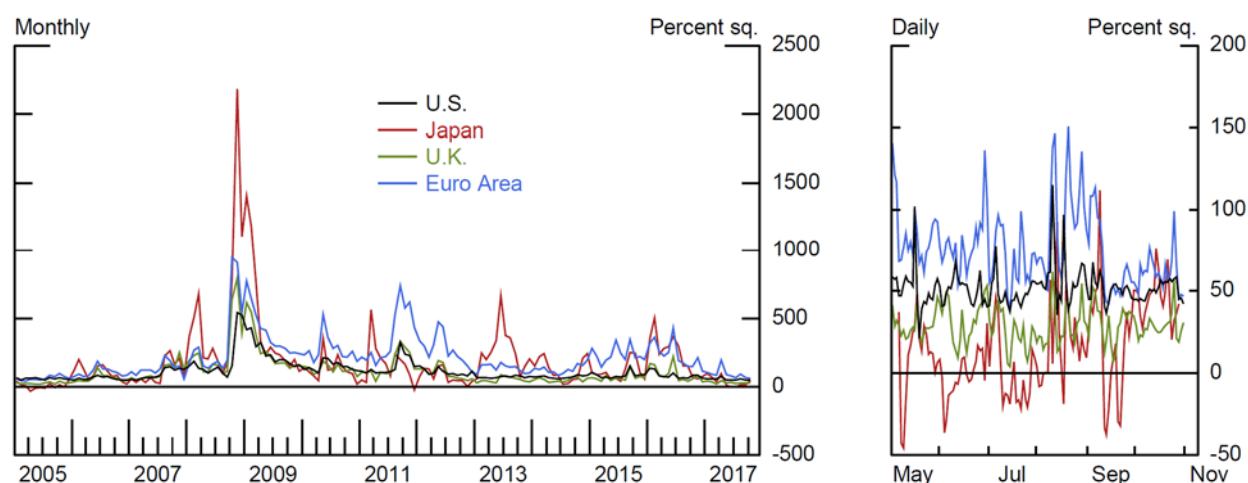
2.4.2. Variance Risk Premium and Its Components

The variance risk premium is a measure of the compensation investors demand for bearing volatility risk, or, in other words, a measure of investor preference for volatility. Formally, it is defined as the difference between a risk-neutral measure of expected variance (for example, the squared value of VIX) and a physical measure of expected realized variance (Bollerslev, Tauchen, and Zhou, 2009). The variance risk premium is often used as a time-varying and state-dependent measure of risk aversion (Rosenberg and Engle, 2002). This measure is also used as a gauge of macroeconomic risk compensation (Bollerslev, Tauchen, and Zhou, 2009; Drechsler and Yaron, 2011). Empirically, it has been shown that the variance risk premium is one of the most successful short-term (between one month and one quarter ahead) predictors of returns across a broad range of financial assets (Bollerslev, Tauchen, and Zhou, 2009; Bollerslev, Gibson, and Zhou, 2011; Dew-Becker, Giglio, Le, and Rodriguez, 2017; Feunou, Jahan-Parvar, and Okou, 2017; and Londono and Zhou, 2017).

Figure 12 shows the time series of the variance risk premium for the United States (the black line). As can be seen in the figure, the magnitude of this premium is positively correlated with macroeconomic uncertainty: The variance risk premium rose significantly after the collapse of Lehman Brothers, during the European debt crisis, and during the debt ceiling negotiations in the United States. However, similar to both the VIX and RV, the variance risk premium is low and has fallen since the 2016 U.S. presidential election. Variance risk premiums are also available, although for a shorter sample, for Canada, Japan, the United Kingdom, France, Germany, and the Netherlands. Interestingly, the variance risk is highly correlated across countries, which suggests that there is a common or global component in variance risk premiums (see Londono, 2016; and Bollerslev, Marrone, Xu, and Zhou, 2014).

A number of studies have investigated the recent historically low levels of volatility and the variance risk premium since the Great Recession. In particular, Feunou, Jahan-Parvar, and Okou (forthcoming) attribute the low levels of variance risk premium to the additive nature of the variance risk premium and specifically to the fact that volatility-based measures treat market movements due to negative and positive realizations the same. However, market participants view positive and negative market movements differently. Feunou, Jahan-Parvar, and Okou (forthcoming) decompose the variance risk premium into upside and downside variance risk premiums. In practice, upside and downside variance risk premiums have different properties. The former represents market participants' interest in being exposed to upside risk and the higher gains it generates, while the latter represents the premium that market participants demand as compensation to bear downside risk and the possible losses it may generate. Upside and downside variance risk premiums may even have different signs. Because the variance risk premium is, by construction, the sum of these two risk premiums, a low variance risk premium may mean calm markets or uncertainty in both positive and negative directions of similar magnitudes.

Figure 12. Variance Risk Premiums for Headline Equity Indexes



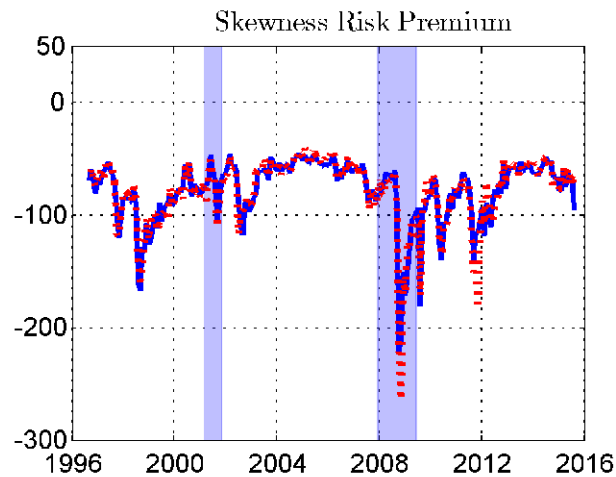
Notes: In Figure 12, each country's variance risk premium is calculated as the difference between the square of the option-implied volatility and the expected realized variance. The expected realized variance is calculated as an in-sample forecast of realized variance using the one-month-lagged realized variance, the square of the VIX, the square of each country's option-implied volatility, and two measures of RV that rely heavily on recent stock returns.

Source: Federal Reserve Board staff calculations based on data from Bloomberg Finance LP.

In addition, using earlier work by Kim and White (2004) as well as more recent studies by Feunou, Jahan-Parvar, and Tedongap (2013, 2016) and Patton and Sheppard (2016), Feunou, Jahan-Parvar, and Okou (2017) show that the difference between upside and downside variance risk premiums,

also known as the signed-jump premium, is a measure of the skewness risk premium. This measure (figure 13), which shares many similarities with the Bollerslev and Todorov (2011) “fear index,” is a better reflection of the direction of uncertainty and market participant concerns about tail risks.

Figure 13. The Skewness Risk Premium for S&P 500 Returns



Notes: Figure 13 plots the skewness risk premium as the difference between the upside and downside variance risk premiums, also known as signed jumps. The blue lines are based on random walk forecasts of RV and the red dotted lines are based on heteroskedastic autoregressive forecasts of RV as in Corsi, 2009.

Source: Feunou, Jahan-Parvar, and Okou (forthcoming).

2.4.3. Option-Implied Probability Distribution for Equity Indexes

Options on equity indexes, unlike those on individual stocks, are fairly liquid and available for a wide range of strikes and time horizons, which facilitates the computation of option-implied probability distributions. This section explains a semiparametric method used to calculate option-implied probability distributions for headline equity indexes.

Breeden and Litzenberger (1978) provide a direct mapping between observed option prices and the option-implied probability distribution, which can be applied to calculate option-implied equity index distributions. In particular, Breeden and Litzenberger (1978) show that the second difference of the price of a European call with respect to the strike price is equivalent to the risk-free rate discounted probability distribution function. However, because option prices are observed for a discrete set of strike prices and the strike prices do not extend to all possible values

of the underlying asset, implementing this method requires that we both interpolate across observed option prices and make some assumptions about the distribution above and below the range of observed strikes. Also, as with the calculation of the VIX discussed in section 2.4.1, to make option-based distributions comparable across dates, we can fix a reference horizon (for example, 30 days or 1 year). Because most options have a fixed expiration date and not a fixed horizon, we need to interpolate across expiration dates to obtain a distribution at a given fixed horizon.

This semiparametric method usually yields smooth option-implied distributions that are easy to interpret and are, therefore, suitable as a policy tool to monitor equity markets. However, this method does involve several choices that can influence the shape of the estimated distribution. First, instead of interpolating option prices directly, we interpolate in the space of Black-Scholes option-implied volatilities to reduce the sensitivity of the results to options with strike prices that are very distant from the current price of the equity index (away-from-the-money options). Second, because strike prices are often far from each other, we use a smoothing spline to interpolate volatilities.¹⁴ This method tends to produce smooth probability distributions with moderately low pricing errors (Datta, Londono, and Ross, 2017). Finally, we assume that option-implied volatilities remain constant beyond the region of strikes for which no options are traded, and we estimate the tails of the distribution under this assumption.

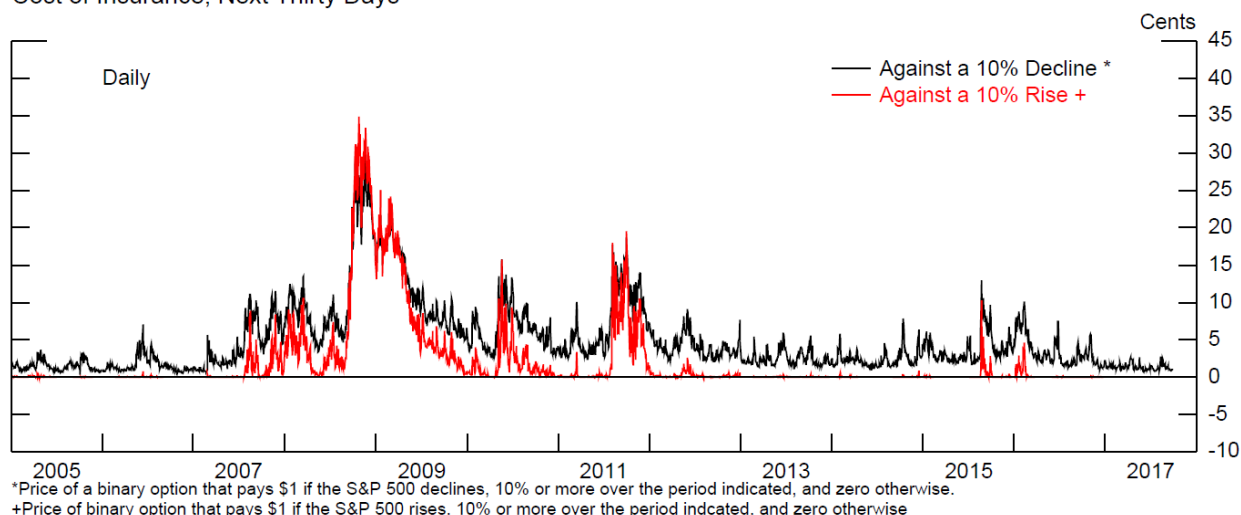
Figure 14 shows the cost of insurance against 10 percent changes in the S&P 500 index for the next 30 days (upper panel) and for the next 3 months (lower panel) as calculated from the option-implied probability distribution of the S&P 500 index. This measure is calculated as the interpolated price of a binary option that pays \$1 if the price of the index declines or rises beyond 10 percent within the next 30 days or 3 months, and otherwise pays zero. The cost of insurance against 10 percent changes tends to increase in episodes of high uncertainty, in line with the dynamics of the VIX (see figure 11). As can be seen in figure 14, in the short term (30 days), investors are usually willing to pay more to hedge a drop in prices than to hedge a potential increase. Interestingly, the cost of insurance against a 10 percent drop has risen following

¹⁴ The method used to interpolate across implied volatilities presents important challenges especially because, in most cases, option prices obtained from interpolated option-implied volatilities ignore the law of one price.

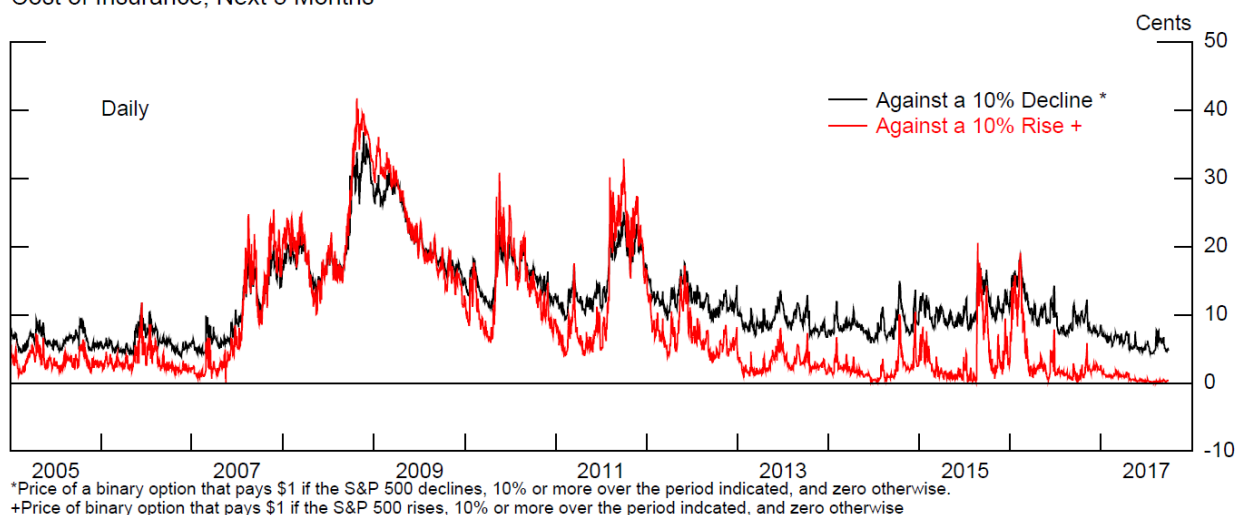
increasing geopolitical tensions earlier this year (upper panel). In contrast, investors appear to be more sanguine about longer-term (3-month) prospects, as the cost of insurance against a 10 percent increase within the next 3 months has ticked up recently, while the cost of insurance against a 10 percent drop for this horizon has remained relatively low since the November 2016 elections.

Figure 14. Cost of Insurance against a 10 Percent Change in the S&P 500 Index

Cost of Insurance, Next Thirty Days



Cost of Insurance, Next 3 Months



Notes: The cost of insurance against a 10 percent decline (rise) is calculated as the price of a binary option that pays \$1 if the S&P 500 index declines (rises) 10 percent or more over the next 30 days (upper panel) or 3 months (lower panel); otherwise the pay is zero.

Source: Federal Reserve Board staff calculations using OptionMetrics data.

2.4.4. Option-Implied Probability Distribution for Individual Stocks

There is increasing interest in measuring the probability of large negative movements in individual companies' equity prices, especially as a large drop in prices may induce default or even bankruptcy. For large, systemically important firms, such events are materially important for both investors and policy makers. There is an active market for options written on stock prices of publicly traded, and typically large, companies. These options are not as frequently traded as index options, most of them are American options (they can be exercised before the expiration date of the option), and they tend to have strike prices close to the underlying stock price. In other words, in sharp contrast to index options, individual company option prices do not span potentially large and negative drops in the underlying stock price. As a result, these options do not provide the level of comprehensive protection against large negative underlying price moves or convey extensive information about market participants' expected probability of such moves.

Since many single-firm options are relatively thinly traded and are concentrated around strikes close to the underlying stock price, we do not apply the methods that are used for index options to extract risk-neutral distributions for individual company option prices. To overcome this issue, Aramonte, Jahan-Parvar, Rosen, and Schindler (2017) exploit the link between credit default swaps (CDS) and options documented by Carr and Wu (2010) to provide a method to extract a derivative-implied distribution for individual stocks by blending option prices and CDS spreads (with the former providing information about the central part of the distribution and the latter determining the left tail) and interpolating the remaining part of the risk-neutral distribution.

2.4.5. Interest Rates

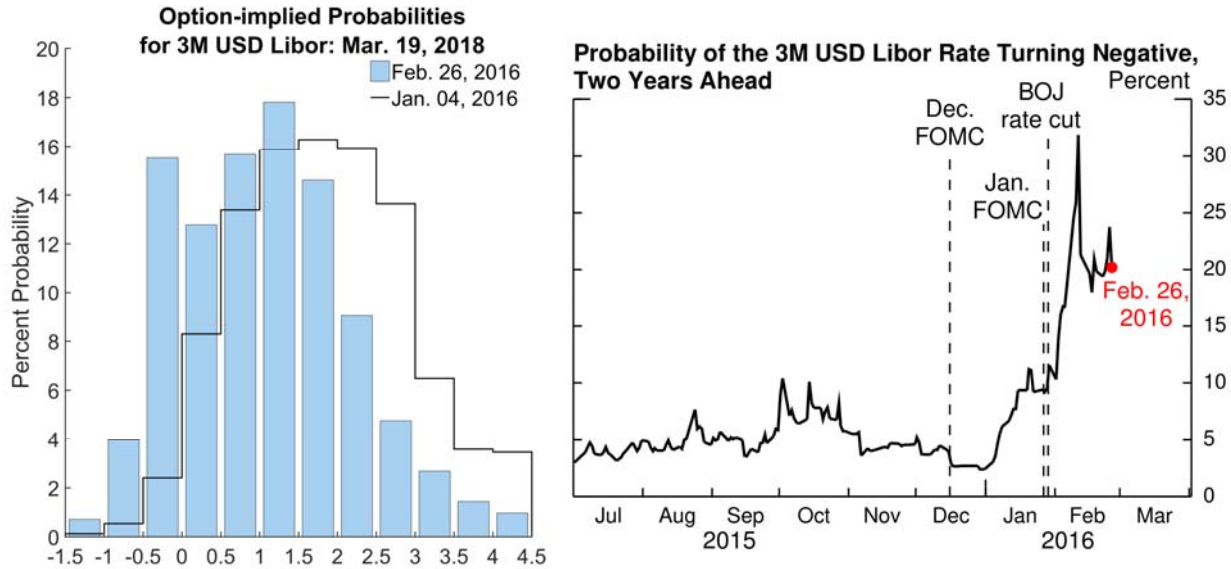
Risk-neutral distributions for interest rates can be either continuous or discrete and are typically calculated from interest rate options, such as federal funds futures options, Eurodollar futures options, Treasury futures options, interest rate caps and floors, or swaptions. Risk-neutral moments can be then calculated either by using these risk-neutral distributions or directly by using other distribution-free methods as outlined by Bakshi, Kapadia, and Madan (2003).

Risk-neutral distributions can be calculated using either a parametric or nonparametric method. In a parametric method, interest rate distributions are assumed to follow a particular specified

distribution. Typically, a mixture of normal distributions is used because interest rate distributions tend to be skewed and have fat tails. Alternatively, a mixture of lognormal distributions can be used if one wants to ensure that interest rates do not fall below zero. The unknown parameters of the mixture distribution can be solved by minimizing the differences between the observed option prices and the model-implied option prices at various strikes. In a nonparametric method, the probability is calculated at discrete strike values using the prices of “butterfly” portfolios of traded options, which have positive payoff at a particular strike and zero payoffs at other strikes. This method can be applied to continuous strikes if interpolation is used to make option prices a continuous function of strikes.

Figure 15 shows the risk-neutral distribution of the London interbank offered rate (LIBOR). The distribution plots the probabilities associated with various outcomes for the three-month LIBOR, two years ahead, for a horizon of three months. The figure compares the distribution on January 4, 2016, with the distribution on February 26, 2016, and shows the market expectation for negative interest rates in the United States in early 2016 following global market turmoil. Increased concerns about the global economic outlook since the beginning of that year shifted the interest rate distribution notably to the left, and option prices suggested that there was a substantial risk-neutral probability that the three-month LIBOR would be negative in two years. The probability of negative rates was low and relatively stable throughout the second half of 2015 (right panel) even during the market turmoil in the summer of 2015 when China unexpectedly devalued its currency. However, at the beginning of 2016, this probability rose substantially amid renewed market stress and following the Bank of Japan’s unexpected move to lower its benchmark policy rate into negative territory. An important caveat, however, is that these elevated probability point estimates may have simply reflected increased risk premiums instead of changed investor expectations.

Figure 15. Interest Rates Risk-Neutral Distributions



Notes: This figure plots the probabilities as implied by the options on the three-month Eurodollar futures expiring on March 19, 2018. In the left panel, the bars located just below -1.00 and above 4.00 percent include all tail probability as measured on January 4, 2016 and February 26, 2016.

Source: Federal Reserve Board staff calculations using CME (Chicago Mercantile Exchange) data.

In a study of the economics of USD and euro-dollar interest rate swaptions-implied moments, Trolle and Schwartz (2014) show that swaption-implied volatility and skewness contain predictive information about future RV and skewness. By comparing these option-implied moments to realized moments, the authors also showed that interest rate variance and, to a lesser extent, skewness risk premiums are time varying and, on average, positive (that is, the risk-neutral moment is higher than the physical expectation of the moment), which suggests that investors dislike higher interest rate variance and skewness.

2.4.6. Currency Option-Implied Distributions and Risk Premiums

The foreign exchange (FX) derivative market is one of the largest and most liquid in the world. However, while a drop in prices is usually considered an unfavorable event for stocks, for currencies, whether the appreciation or depreciation of a currency with respect to another is an unfavorable outcome depends on several factors, especially investor location. Therefore, unlike

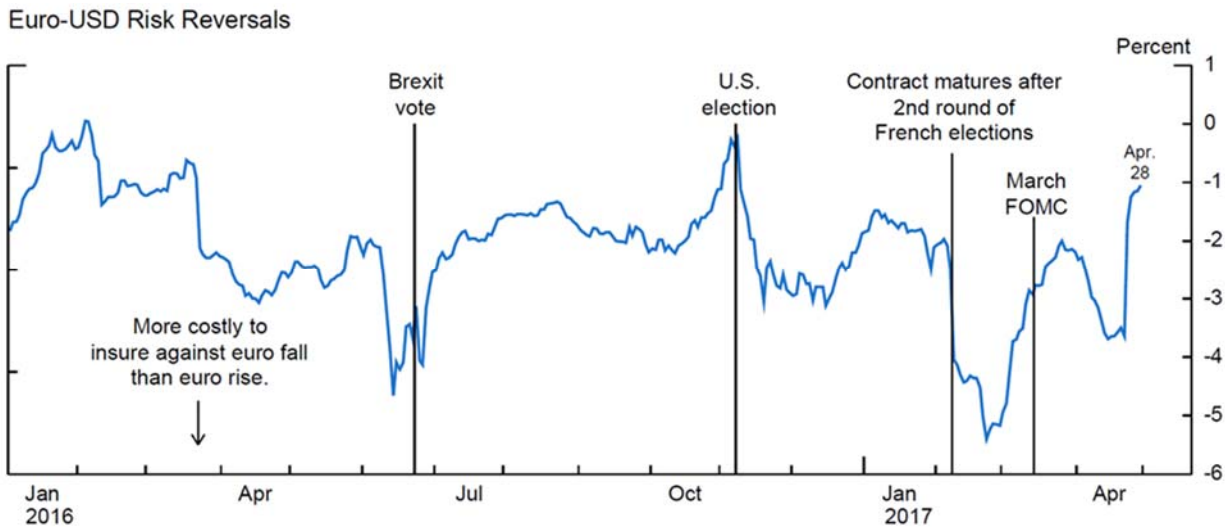
the options written on equity markets described in section 2.4.3, most exchange rate derivatives are written as a combination of put and call options with the same deltas (the sensitivity of the option price to changes in the price of the underlying asset). Therefore, we can still use these combination derivatives, or “strategies,” to derive the risk-neutral distribution of currencies.

The most common strategies are risk reversals and strangles. Risk reversals provide information about the cost of insurance against the depreciation of a currency relative to the cost of insurance against the appreciation of such currency. Specifically, a long position in a risk reversal is equivalent to purchasing a call option and selling a put option on a single bilateral exchange rate. Thus, this strategy protects the investor against an unfavorable drop in the exchange rate (for example, a drop in the dollar with respect to another currency for an exporter located in the United States) but limits investor gains if there is a favorable increase in the exchange rate. In a strangle, the investor buys out-of-the-money calls and puts that have the same maturities and deltas. With this strategy, the investor can profit when a currency appreciates or depreciates significantly.

The prices of call and put options at different strikes can be extracted from the prices of the different FX strategies. Consequently, all the methods to calculate option-implied distributions for equities described previously can be used for currencies once the options prices are extracted. In addition, the strategies also give us direct readings of the cost of insurance against a currency depreciation.

Figure 16 shows an example relating the cost of insurance derived from risk reversals to the uncertainty of some economic events. The figure shows the time series for euro-dollar risk reversals for the three-month horizon between January 2016 and April 2017. As is clear from the figure, concerns about the outcome of the Brexit vote and the French presidential elections contributed a notable increase in the cost of insuring against a euro depreciation relative to a euro appreciation.

Figure 16. Euro-USD Risk Reversals



Notes: This figure plots a time series derived from the prices of risk reversals for the cost of insurance against a euro depreciation. These are over-the-counter derivatives with a 10-delta strike and 3-month maturity.

Source: Thomson Reuters.

As with equities (see section 2.4.2), we can also calculate the currency variance risk premium and its components for all currency pairs with option data available. Londono and Zhou (2017) show that the global variance risk premium, which is an average of currency-specific variance risk premiums, is a useful predictor for future appreciation rates. In particular, an increase in the global currency variance risk premium is followed by an appreciation of the USD with respect to other currencies. Interestingly, the currency variance risk premium contains additional information relative to the equity variance risk premium, which has also been documented to be a useful predictor of currency appreciation. The additional informational content of the global variance risk premium can be empirically and theoretically related to global inflation uncertainty.

2.4.7. Oil Price Option-Implied Probability Distributions and Risk Measures

Options on oil futures contracts for West Texas Intermediate (WTI) crude oil are available on the New York Mercantile Exchange (NYMEX). As with other assets, the prices of options with different strikes and different horizons to maturity contain information about the probability assigned to each possible market outcome for crude oil prices as well as investor preferences. These option prices can be used to generate option-implied distributions, which, in turn, can be

used to calculate option-implied moments, such as implied volatility or the cost of insurance against particular market outcomes. Figure 17 (panel 5) depicts the implied volatility calculated from option-implied distributions as well as the Chicago Board Options Exchange (CBOE) Oil VIX, which is an alternative summary measure of implied volatility for the WTI price of crude oil that is analogous to the S&P 500 VIX.


Datta, Londono, and Ross (2017) apply the semiparametric method developed by Breeden and Litzenberger (1978) to map between observed option prices and the option-implied probability distribution, analogous to the method used for equity distributions discussed previously. This method generally produces smooth probability distributions with low pricing errors. To reduce the sensitivity of the results to options with strike prices that are very distant from the current oil price, the authors use a smoothing spline and interpolate across the volatilities implied by the observed strike prices. Then, to generate the distribution above and below the range of observed strikes, the authors assume that option-implied volatilities remain constant beyond the region of strikes for which no options are traded.

Lastly, to make option-implied distributions comparable across dates, the authors fix a reference horizon (for example, 30 days or 1 year). As with options on equity indexes, these options are available for fixed expiration dates that coincide with futures contract availability, which have a monthly frequency. Datta, Londono, and Ross (2017) find that a fixed horizon of 90 days generally has enough liquidity to support estimation of informative option-implied distributions. Distributions with fixed December expiration dates (for example, December 2017) are also informative.

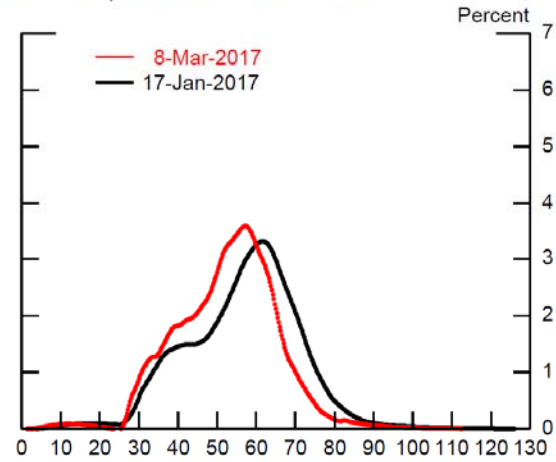
Panels 1 and 2 of figure 17 illustrates the additional information provided by the entire WTI option-implied probability density function (PDF) compared to summary moments. In the December 2017 PDF (shown in panel 2), the significant implied probability of oil prices falling to the \$40-to-\$50 per barrel range likely indicates market participant beliefs that the November 2016 deal among OPEC and some non-OPEC countries to constrain production in support of prices may unravel by December 2017, prompting a significant decline in prices relative to the modal outcome.

Figure 17. Option-Implied Measures for WTI Crude Oil

1. PDFs, 3-Month Futures



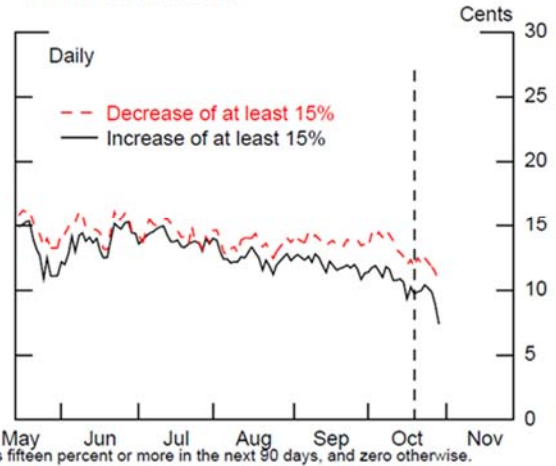
2. PDFs, December 2017 Futures



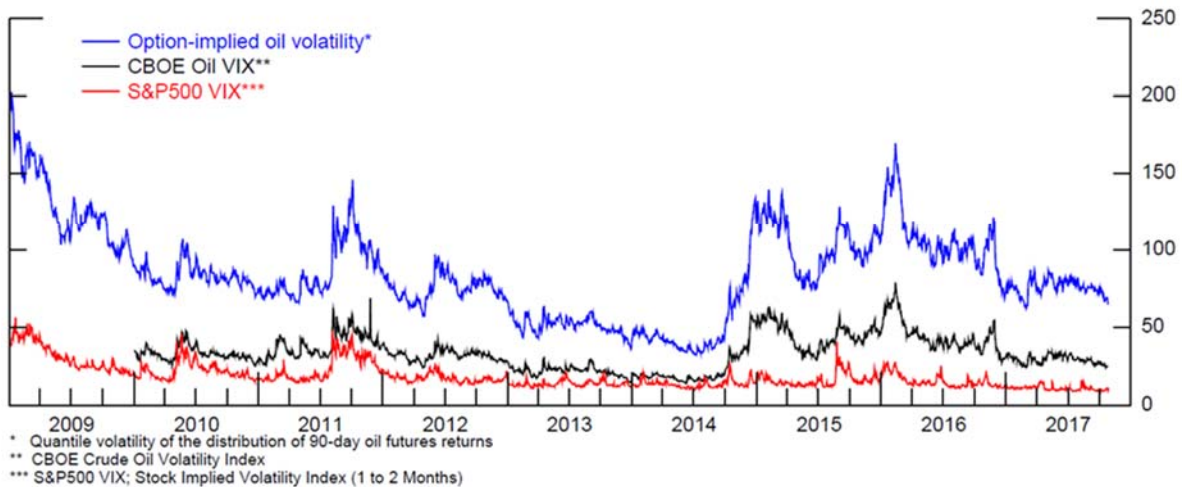
3. Cost of Insurance*



4. Cost of Insurance*



5. Options-Implied Volatility



* Quantile volatility of the distribution of 90-day oil futures returns
 ** CBOE Crude Oil Volatility Index
 *** S&P500 VIX; Stock Implied Volatility Index (1 to 2 Months)

Source: Federal Reserve Board staff calculations using NYMEX data.

While we show here the distribution of the expected price level, the distribution can be expressed as a distribution of expected returns (that is, changes in future oil prices with respect to the current price). Using the distribution of returns, we calculate summary measures, which can be plotted over time. In particular, panels 3 and 4 show the monthly and daily frequency series for the cost of insurance against 15 percent changes in the price of oil in the next 90 days. Panel 5 shows the CBOE Oil VIX along with option-implied quantile volatility.¹⁵

Panel 3 of figure 17 shows that after the OPEC deal was announced in late November 2016, the cost of insurance against a 15 percent move in oil prices fell to levels not seen since before the oil price drop in 2014. More recently, the cost of insurance has ticked up as U.S. crude oil production gains strength, prompting concerns that the OPEC deal could unravel or not be extended past its current end date of June 2017. Notably, unlike the implied volatility in equity markets, which has fallen to historic lows in recent months (as indicated by the red VIX line in panel 5), implied volatility for oil has continued to respond to major oil supply and demand developments.

2.4.8. Inflation

Inflation can significantly erode the real values of nominal investments over time, whereas deflation can significantly increase the real burdens of nominal liabilities. Inflation-linked derivatives can help investors manage both types of inflation risks by offering a way to capture incomes or hedge liabilities pegged to future inflation growth, and, as such, their prices contain useful information about investors' perceived likelihood and preferences with respect to various future inflation outcomes.¹⁶ The most common types of inflation derivatives are inflation swaps

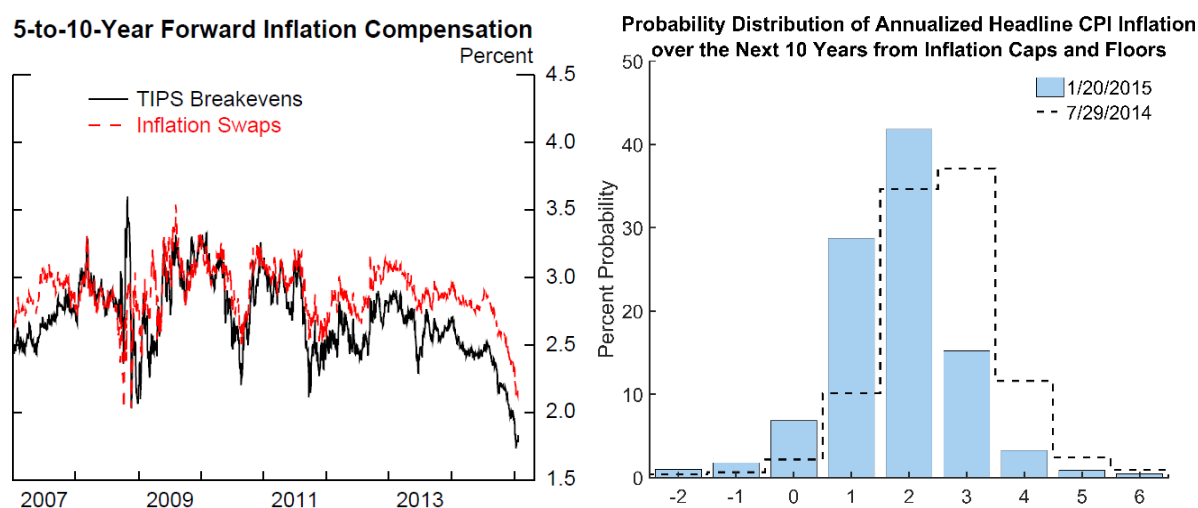
¹⁵ Datta, Londono, and Ross (2017) show that relative to the usual calculations for volatility, skewness, and kurtosis, quantile moments are more likely to reflect meaningful underlying daily market movements than data anomalies. In particular, because quantile moments do not use the most extreme tails of the distribution, they are less reliant on the assumptions made to generate the distribution above and below the range of observed strike prices. Quantile volatility is calculated as the difference between the returns at the 75th and 25th percentiles of the option-implied probability density functions. Quantile skewness is calculated as the difference between the returns at the 75th and 50th percentiles minus the difference between the returns at the 50th and 25th percentiles, with this difference then divided by the difference between the returns at the 75th and 25th percentiles. Finally, quantile kurtosis is calculated as the difference between the returns at the 95th and 5th percentiles divided by the difference between the returns at the 75th and 25th percentiles minus a coefficient adjustment.

¹⁶ In contrast to other financial derivatives, the market for inflation derivatives is still relatively small. Currently, the European inflation derivatives market is by far the largest, but the other markets have reached a reasonable liquidity as well.

and inflation caps and floors. In an inflation swap, one party pays the other party a predetermined fixed rate and, in return, receives a floating payment that is linked to future realized inflation.¹⁷ In an inflation cap (floor) contract, in exchange for a premium, the holder of the contract pays the other party the excess (shortfall) of realized inflation over some specified strike if the realized inflation exceeds (falls below) that level over the life of the contract.

Because inflation caps and floors have options-like payoffs, we can calculate their implied probability distribution using methods similar to those described in the previous sections. Figure 18 shows an example of how the probability distributions implied by inflation caps can be useful for policy analysis. In the fall of 2014, following the sharp decline in oil prices that started in the summer, measures of longer-term inflation compensation, such as 5-to-10-year-forward inflation compensation, had declined notably to levels comparable to those last seen during the 2008 financial crisis. The decline in inflation compensation measures prompted market commentary and policy discussion about the risks of deflation.

Figure 18. Inflation Risk-Neutral Distributions



Note: This figure plots the 5-to-10-year forward inflation compensation in the left panel and the probability distribution of annualized headline Consumer Price Index inflation as implied by inflation caps and floors in the right panel. The distribution is derived under the assumption that average inflation takes discrete values (for example, the bar for 3 percent covers roughly the area between 2.5 and 3.5 percent).

Source: Federal Reserve Board staff calculations using Bloomberg data.

¹⁷ The indexes used to calculate realized inflation are the CPI (Consumer Price Index) in the United States and France, the RPI (Retail Prices Index) in the United Kingdom, and the HICP (Harmonised Index of Consumer Prices) excluding tobacco in the euro area. All are non-seasonally adjusted. Therefore, the headline and non-core inflation are the focus.

At the time, the relative changes in the distributions of future inflation derived from inflation caps and floors suggested that investors had become more concerned about lower inflation outcomes and less concerned about higher inflation outcomes. These distributions, as well as measures of inflation compensation, reflect not only expected inflation but also an inflation risk premium.¹⁸ Consequently, the shift in the distribution in the fall of 2014 could have reflected an increase in the perceived likelihood of low inflation outcomes or an increase in the willingness to pay higher premiums to insure against such outcomes, perhaps because investors at that time increasingly associated them with poor economic performance. Staff term structure models can help us provide estimates of each component separately to better understand the contribution of expectations and risk premiums to movements in inflation compensation. The model results suggest that in the fall of 2014, inflation expectations remained relatively stable and pointed to lower inflation risk premiums. This result was consistent with survey-based distributions of 5-to-10-year forward inflation, which had generally remained stable over the period.

Kitsul and Wright (2013) study how options-implied distributions of inflation respond to news announcements in the United States and find that the implied probabilities of deflation are sensitive to certain macroeconomic news releases. They also find that the option-implied distributions point to higher probabilities for extreme outcomes (deflation or high inflation) relative to the physical distributions, implying that investors have high marginal utility in these states of the world.

¹⁸ Inflation compensation also reflects, to a lesser extent, other premiums due to liquidity differences and shifts in the relative supply and demand of nominal versus inflation-indexed securities.

References

- Aastveit, K. A., Natvik, G. J., and Sola, S. (2013), "Economic Uncertainty and the Effectiveness of Monetary Policy," Norges Bank Working Paper 17.
- Ait-Sahalia, Y., Karaman, M., and Mancini, L. (2014), "The Term Structure of Variance Swaps, Risk Premia and the Expectations Hypothesis," Working Paper, Princeton University.
- Alizadeh, S., Brandt, M. W., and Diebold, F. X. (2002), "Range-Based Estimation of Stochastic Volatility Models," *Journal of Finance* 57(3): 1047-1091.
- Amengual, D., and Xiu, D. (2017), "Resolution of Policy Uncertainty and Sudden Declines in Volatility," Working Paper, Chicago Booth and CEMFI.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2001), "The Distribution of Realized Exchange Rate Volatility," *Journal of the American Statistical Association* 96(453): 42-55.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2003), "Modeling and Forecasting Realized Volatility," *Econometrica* 71(2): 579-625.
- Aramonte, S., Jahan-Parvar, M. R., Rosen, S., and Schindler, J. (2017), "Firm-Specific Risk-Neutral Distributions: The Role of CDS Spreads," Working Paper, Federal Reserve Board.
- Bachmann, R., Elstener, S., and Sims, E. (2013), "Uncertainty and Economic Activity: Evidence from Business Survey Data," *American Economic Journal: Macroeconomics* 5(2): 217-249.
- Baker, S. R., Bloom, N., and Davis, S. J. (2016), "Measuring Economic Policy Uncertainty," *The Quarterly Journal of Economics* 131(4): 1593-1636.
- Bakshi, G., Kapadia, N., and Madan, D. (2003), "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options," *Review of Financial Studies* 16(1): 101-143.
- Barro, R. (1990), "The Stock Market and Investment," *Review of Financial Studies* 3(1): 115-131.
- Bekaert, G., Hoerova, M., Lo Duca, M. (2013), "Risk, Uncertainty and Monetary Policy," *Journal of Monetary Economics* 60(7): 771-788.
- Bloom, N. (2009), "The Impact of Uncertainty Shocks," *Econometrica* 77(3): 623-685.
- Bloomberg Finance LP. Bloomberg Terminals (Open, Anywhere, and Disaster Recovery Licenses).
- Bollerslev, T. and Todorov, V. (2011), "Tails, Fears, and Risk Premia," *Journal of Finance* 66(6): 2165-2221.
- Bollerslev, T., Gibson, M., and Zhou, H. (2011), "Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities," *Journal of Econometrics* 160(1): 235-245.

- Bollerslev, T., Marrone, J., Xu, L., and Zhou, H. (2014), "Stock Return Predictability and Variance Risk Premia: Statistical Inference and International Evidence," *Journal of Financial and Quantitative Analysis* 49(3): 633-661.
- Bollerslev, T., Tauchen, G., and Zhou, H. (2009), "Expected Stock Returns and Variance Risk Premia," *Review of Financial Studies* 22(11): 4463-4492.
- Bollerslev, T. and Zhou, H. (2006), "Volatility Puzzles: A Simple Framework for Gauging Return-Volatility Regressions," *Journal of Econometrics* 131(1): 123-150.
- Breeden, D. and Litzenberger, R. (1978), "Prices of State-Contingent Claims Implicit in Option Prices," *Journal of Business* 51(4): 621-651.
- Britten-Jones, M. and Neuberger, A. (2000), "Option Prices, Implied Price Processes, and Stochastic Volatility," *Journal of Finance* 55(2): 839-866.
- Brogaard, J. and Detzel, A. (2015), "The Asset Pricing Implications of Government Economic Policy Uncertainty," *Management Science* 61(1): 3-18.
- Caldara, D. and Iacoviello, M. (2016), "Measuring Geopolitical Risk," Working Paper, Federal Reserve Board.
- Carr, P. and Madan, D. (1998), "Towards a Theory of Volatility Trading," In R. Jarrow (Ed.), *Volatility: New Estimation Techniques for Pricing Derivatives*: 417-427. Risk Books.
- Carr, P. and Wu, L. (2010), "Stock Options and Credit Default Swaps: A Joint Framework for Valuation and Estimation," *Journal of Financial Econometrics* 8(4): 409-449.
- Christiano, L. J., Motto, R., and Rostagno, M. (2014), "Risk Shocks," *The American Economic Review* 104(1): 27-65.
- Credit Suisse (2004), "Global Risk Appetite Index," *Market Focus*, February 20.
- Corsi, F. (2009), "A Simple Approximate Long-Memory Model of Realized Volatility," *Journal of Financial Econometrics* 7(2): 174-196.
- Datta, D. D., Londono, J. M., and Ross, L. J. (2017), "Generating Options-Implied Probability Densities to Understand Oil Market Events," *Energy Economics*, 64: 440-457.
- Dew-Becker, I., Giglio, S., Le, A., and Rodriguez, M. (2017), "The Price of Variance Risk," *Journal of Financial Economics* 123(2): 225-250.
- Drechsler, I. and Yaron, A. (2011), "What's Vol Got to do with It?," *Review of Financial Studies* 24(1):1-45.
- Ferreira, T. (2017), "Stock-Market Cross-Section Skewness and Business Cycle Fluctuations," Working paper, Federal Reserve Board.

- Feunou, B., Jahan-Parvar, M. R., and Okou, C. (forthcoming), "Downside Variance Risk Premium," *Journal of Financial Econometrics*.
- Feunou, B., Jahan-Parvar, M. R., and Tedongap, R. (2013), "Modeling Market Downside Volatility," *Review of Finance* 17(1): 443-481.
- Feunou, B., Jahan-Parvar, M. R., and Tedongap, R. (2016), "Which Parametric Model for Conditional Skewness?," *European Journal of Finance* 22(13): 1237-1271.
- Francis, B. B., Hasan, I., and Zhu, Y. (2014), "Political Uncertainty and Bank Loan Contracting," *Journal of Empirical Finance* 29: 281-286.
- Gilchrist, S. and Zakrajšek, E. (2012), "Credit Spreads and Business Cycle Fluctuations," *The American Economic Review* 102(4): 1692-1720.
- Gulen, H. and Ion, M. (2015), "Policy Uncertainty and Corporate Investment," Working Paper, Purdue University.
- Husted, L., Rogers, J., and Sun., B. (2016), "Measuring Cross Country Monetary Policy Uncertainty," IFDP Notes, <https://www.federalreserve.gov/econresdata/notes/ifdp-notes/2016/measuring-cross-country-monetary-policy-uncertainty-20161123.html>.
- Husted, L., Rogers, J., and Sun., B. (2017), "Monetary Policy Uncertainty," Working Paper, Federal Reserve Board.
- Illing, M. and Aaron, M. (2005), "A Brief Survey of Risk-Appetite Indexes," *Bank of Canada Financial System Review*, June: 37-43.
- Jiang, G. J. and Tian, Y. S. (2005), "The Model-Free Implied Volatility and Its Information Content," *The Review of Financial Studies* 18(4): 1305-1342.
- Jurado, K., Ludvigson, S. D., and Ng, S. (2015), "Measuring Uncertainty," *American Economic Review* 105(3): 1177-1216.
- Kim, T. H. and White, H. (2004), "On More Robust Estimation of Skewness and Kurtosis," *Finance Research Letters* 1(1): 56-73.
- Kitsul, Y. and Wright, J. H. (2013), "The Economics of Options-Implied Inflation Probability Density Functions," *Journal of Financial Economics* 110(3): 696-711.
- Kumar, M. and Persaud, A. (2002), "Pure Contagion and Investors' Shifting Risk Appetite: Analytical Issues and Empirical Evidence," *International Finance* 5(3): 401-36.

- Londono, J. M. (2016), "The Variance Risk Premium around the World," Working Paper, Federal Reserve Board.
- Londono, J. M. and Zhou, H. (2017), "Variance Risk Premiums and the Forward Premium Puzzle," *Journal of Financial Economics* 124(2): 415-440.
- Ludvigson, S., Ma, S., and Ng, S. (2016), "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?," Working Paper, New York University.
- Manzo, G. (2013), "Political Uncertainty, Credit Risk Premium and Default Risk," Working Paper, University of Chicago Booth School of Business.
- Misina, M. (2003), "What does Risk-Appetite Index Measure," Working Paper No. 2003-23, Bank of Canada.
- OptionMetrics. Ivy DB US.
- Pastor, L. and Veronesi, P. (2012), "Uncertainty about Government Policy and Stock Prices," *Journal of Finance* 67(4): 1219-1264.
- Pastor, L. and Veronesi, P. (2013), "Political Uncertainty and Risk Premia," *Journal of Financial Economics* 110(3): 520-545.
- Patton, A. and Sheppard, K. (2016), "Good Volatility, Bad Volatility: Signed Jumps and the Persistence of Volatility," *Review of Economics and Statistics* 97(3): 683-697.
- Rey, H. (2015), "Dilemma not Trilemma: The Global Financial Cycle and Monetary Policy Independence," Working Paper No. 21162 National Bureau of Economic Research.
- Rosenberg, J. V., and Engle, R. F. (2002), "Empirical Pricing Kernels," *Journal of Financial Economics* 64(3): 341-372.
- Scotti, C. (2016), "Surprise and Uncertainty Indexes: Real-Time Aggregation of Real-Activity Macro Surprises," *Journal of Monetary Economics* 82: 1-19.
- Thomson Reuters. Thomson Reuters Worldscope, <http://financial.thomsonreuters.com/en/products/data-analytics/company-data/company-financials.html>.
- Trolle, A. B. and Schwartz, E. S. (2014), "The Swaption Cube," *Review of Finance Studies* 27(8): 2307-2353.
- Whaley, R. E. (2000), "The Investor Fear Gauge," *Journal of Portfolio Management* 26(3): 12-17.