How Do Households Respond to Uncertainty Shocks?

By Edward S. Knotek II and Shujaat Khan

Economic disruptions generally coincide with heightened uncertainty. In the United States, uncertainty increased sharply with the recent housing market crash, financial crisis, deep recession, and uneven recovery. In July 2010 congressional testimony, Federal Reserve Chairman Bernanke described conditions as "unusually uncertain." The uncertain landscape was also cited as a factor in the slow recovery from the 2001 recession, when the March 2003 Federal Open Market Committee statement highlighted the "unusually large uncertainties" at the time.

Uncertainty is a standard feature of most macroeconomic models, in which consumers and firms make decisions today based on expectations of an unknown (and hence uncertain) future. But in light of real-world events, economists have begun to think more critically about the role of uncertainty in the economy. Recent research has allowed the degree of uncertainty to vary over time and examined how these fluctuations affect business activity. The results have been mixed thus far, with some authors finding that fluctuations in uncertainty are a key factor in the business cycle, while others have found little such evidence.

*Edward S. Knotek II is a vice president and economist, and Shujaat Khan is a research associate, both at the Federal Reserve Bank of Kansas City. This article is on the bank's website at **www.KansasCityFed.org**.*

This article takes a similar approach in studying levels of uncertainty that can vary over time, but it focuses on household responses to changes in uncertainty. Because uncertainty can take many forms, the article considers two measures of uncertainty: one based on references to uncertainty in newspaper articles and another derived from the stock market.

While economic theory predicts sudden, sharp pullbacks of household purchases following increases in uncertainty, the empirical results suggest that household spending reductions are modest and may only appear after a considerable time has passed. In addition, movements in uncertainty account for only a small portion of the total fluctuations in household spending. These results suggest that variations in the amount of uncertainty—at least as they are commonly captured—do not appear to be a key factor driving household spending decisions and, in turn, economic weakness.

The first section of this article provides a framework for thinking about uncertainty and how it affects economic activity. The second section describes two separate measures that have been proposed by economists to quantify uncertainty. The third section assesses the importance of fluctuations in uncertainty on household purchases, first using a simple bivariate model and then using a model that incorporates more relevant information.

I.   A FRAMEWORK FOR UNCERTAINTY SHOCKS AND ECONOMIC ACTIVITY

Uncertainty is an important feature of the real world and plays a central role in modern macroeconomics. Yet uncertainty can be interpreted in a number of ways. To fix ideas, this section provides a framework for thinking about uncertainty and how it might change over time. With this framework in place, the section then considers how economic agents respond to fluctuations in uncertainty in theoretical models. The common theme of these theories is that an increase in uncertainty tends to dampen spending immediately, as businesses and consumers enter a "wait-and-see" mode. This response is especially characteristic of expensive purchases that are difficult to undo. Over time, as the increase in uncertainty wears off, economic activity tends to rebound sharply.

What is an uncertainty shock?

Economic models incorporate uncertainty about the future through random disturbances, or shocks. For example, a household may know its wealth today, but its wealth tomorrow may be uncertain because it is subject to shocks that are ultimately beyond the household's control.

To analyze the effects of such shocks, economists usually assume that the probability distribution of the shocks—that is, the likelihood of a shock of a given size occurring—is both known and constant over time. The blue bars in Chart 1 provide an example of a distribution of shocks that can affect future wealth. The possible outcomes range from negative shocks that reduce wealth by a certain amount to positive shocks that increase wealth. By assumption, the household knows both the range of possible shocks and their associated probabilities (the heights of the bars). These probabilities are typically assumed not to change over time. Not knowing exactly which of the five shocks will occur in the future is thus the only source of uncertainty in this example.

Recent research has considered the possibility that there may be an additional source of uncertainty. In particular, the shape of the probability distribution of shocks may change over time.[1] As an example, in Chart 1 the probabilities of the different shocks may change from the blue bars to the black bars and vice versa. This change in the probability distribution represents an uncertainty shock.

An uncertainty shock does not affect the level of wealth directly, as a wealth shock does. Rather, an uncertainty shock affects the distribution of shocks that can affect wealth by changing the variance of the distribution while leaving the mean constant.[2] An increase in uncertainty (a positive uncertainty shock) makes outcomes more uncertain in the sense that tail events, or realizations at the extremes of the distribution, have a higher probability of occurring. This is equivalent to moving from the blue distribution to the black distribution in the chart.[3]

From a household's perspective, uncertainty shocks can take a number of different forms. For example, an uncertainty shock may be associated with a more volatile stock market. Such volatility simultaneously raises the probabilities of much higher and much lower wealth in the future. But uncertainty shocks can also affect a household's income prospects. Greater uncertainty could simultaneously raise the likelihood

Chart 1
TWO DISTRIBUTIONS OF FUTURE OUTCOMES

Probability of each outcome

Shock Outcome	
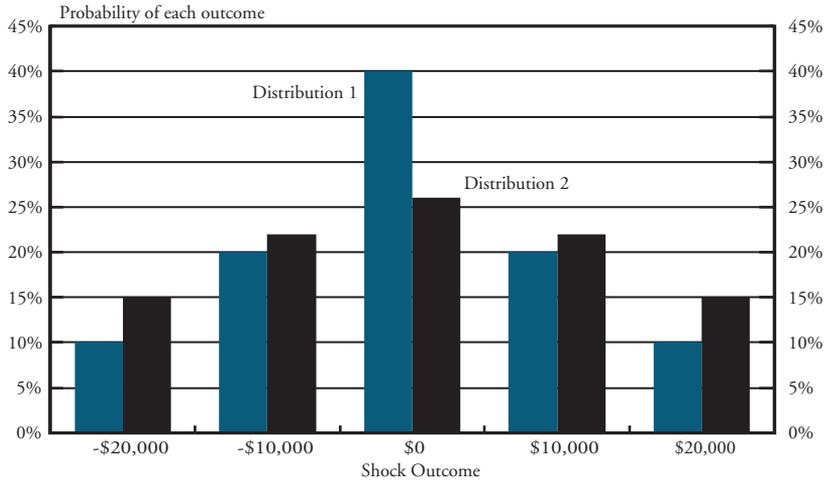
that an individual may lose a job, resulting in lower near-term income prospects, or find a better job, resulting in higher income prospects.

What does economic theory predict?

While theoretical models generally incorporate random shocks, they do not typically allow the probability distribution of shocks to change over time. Thus, they do not account for uncertainty shocks as defined above. To understand the effects of uncertainty shocks, economists have primarily relied on the real options framework.[4]

The real options framework draws on insights from financial markets. In the market for financial assets, the buyer of a call option acquires the right—but not the obligation—to purchase a specified financial asset at a given price by a particular time in the future. This call option gives the buyer additional time to acquire more information and determine whether to buy the financial asset. Economists use the phrase "real options" to denote investments in economic (or "real") assets, as opposed to investments in financial assets. At its essence, the real options framework captures the notion that under some circumstances there may be a benefit to waiting and acquiring more information before making a decision to invest in a real asset.

A key feature underlying the real options framework is irreversibility. For instance, a firm may purchase a new, customized piece of equipment that cannot be sold to another firm without incurring a loss. Alternatively, a firm may face costs associated with hiring new workers (such as posting a vacancy, screening applicants, and training new hires) and firing workers (such as severance packages). In either case—the purchase of a new piece of capital or the hiring of a new worker—the decision cannot be reversed costlessly.

By themselves, irreversibilities can cause firms to delay acting when faced with an uncertain future. In particular, firms typically want to make sure that times are "good enough" or "bad enough" to justify making an irreversible decision that they may later wish they could undo. As a result, they may not respond to very small shocks, or they may wait for enough positive or negative shocks to accumulate before acting.[5]

Irreversibilities combined with uncertainty shocks make firms even more cautious. An uncertainty shock raises the likelihood of extreme outcomes; for instance, after an uncertainty shock firms may see the potential for much higher or lower future profits than previously expected. Faced with this increased level of uncertainty, firms' first choice is to wait for more information, if possible, before acting. In other words, the value of the real option of waiting is greater after an uncertainty shock than in normal times. Firms that do act would want to ensure that times are "really good enough" or "really bad enough" before doing so to limit the number of occasions when the firm makes an irreversible decision it may later regret if it were to face an extreme outcome.[6]

This phenomenon has also been termed the "bad news principle" because the increased threat of negative outcomes swamps the simultaneous increased possibility of positive outcomes, thereby reducing activity (Bernanke). If enough firms follow the bad news principle, then this uncertainty shock can produce an economic downturn.[7]

But the effects of uncertainty shocks are not limited to firms. Uncertainty shocks are usually felt throughout the economy, affecting households as well. In addition, like firms, households must sometimes make irreversible decisions. For example, purchases of houses cannot be reversed quickly and costlessly due to real estate agent commissions, closing costs associated with obtaining a mortgage, and moving costs. Purchases of durable goods, notably cars, are also difficult to reverse.

The prices of barely used cars are considerably less than the prices of brand new cars—the well-known lemons problem (Akerlof).[8] Such irreversibilities, however, may be less pronounced or even absent when consumers purchase other nondurable goods and services.[9]

Even under normal circumstances, irreversibilities make households cautious about purchasing houses or durable goods because they would rather not make a purchase they must quickly undo. The arrival of an uncertainty shock may cause households to become even more cautious. While households know their current income and wealth, the uncertainty shock may lead them to believe that the distribution of future income and wealth prospects favors more extreme outcomes, as illustrated in Chart 1. As a result, the uncertainty shock amplifies households' caution, making them less likely to act. That is, the real option value of waiting temporarily increases. As a consequence, households optimally respond by reducing their purchases of houses and durable goods below normal levels. The decline in purchases of durable goods may even coincide with a rise in spending on nondurable goods, as households shift how they spend their current wealth (Romer).

Over time, the uncertainty shock subsides and households gain more information about their income and wealth prospects. As uncertainty returns to a more normal level, households realize they own fewer houses and durable goods than they consider optimal given their reduced uncertainty. Their pent-up demand thus causes a temporary surge in spending.

Chart 2 illustrates this bust-boom pattern for the durable goods that face irreversibilities. When the uncertainty shock occurs in period 0, households immediately postpone purchases and activity falls below trend. Within a short time, activity rebounds as the uncertainty is resolved and households act on their pent-up demand.

II. MEASURING UNCERTAINTY

Considering variations in the level of uncertainty in a macroeconomic model may seem straightforward. But capturing the intangible movements of uncertainty in the real world can be problematic. To this end, economists have considered many different proxies for the level of uncertainty, which in turn can be used to identify uncertainty shocks.

Chart 2

AN ILLUSTRATION OF THE EFFECTS OF AN UNCERTAINTY SHOCK
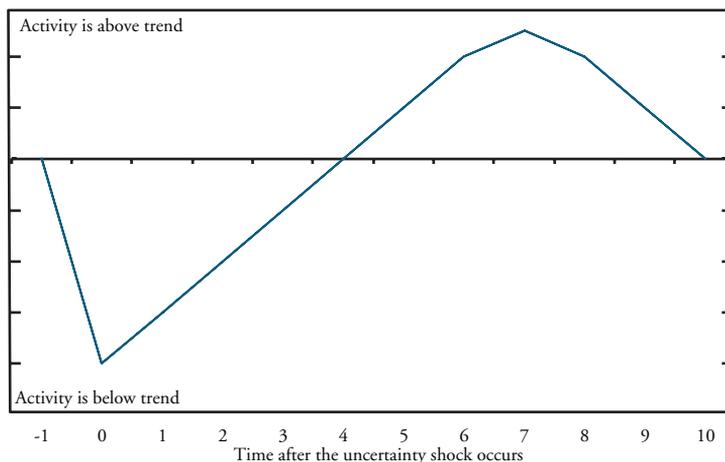
Note: The uncertainty shock occurs at time 0 on the horizontal axis.

To examine the impact of uncertainty shocks on household purchases, this article focuses on two measures of uncertainty.

The first measure is based on the monthly appearance of the words "uncertainty" or "uncertain" in articles in *The New York Times* (Alexopolous and Cohen). Highly uncertain times are likely to be reported as such by journalists. When uncertainty is not as important, it is unlikely to be newsworthy. Thus, the written record can provide one measure of intangible uncertainty. To ensure that the use of the word relates to the economy rather than to, say, fall fashion lines, the article must also contain the words "economy" or "economic."

The New York Times has a number of attributes that make it attractive for such a project. It is one of the nation's most widely read newspapers and a primary New York City metropolitan newspaper. Hence, its published articles capture both the nation's and Wall Street's sentiments on the economy. The *Times* website also has an online archive that allows for relatively easy searches of previously published articles dating back to its founding in 1851. This provides ample coverage of the time period considered in this article, which begins in the 1960s.

Over time, variation in the number of articles mentioning uncertainty could reflect either changes in the frequency of such articles in the *Times* or changes in the total number of articles published within a given month. To correct for this possibility, this article computes the frequency of uncertainty references by calculating the ratio of monthly articles on economic uncertainty to the total number of articles.[10]

The second measure of uncertainty is derived from monthly stock market volatility.[11] More uncertain times may be associated with highly volatile stock price movements, reflecting rapid shifts in investor sentiment between positive and negative outlooks. To protect themselves from this volatility, investors may also increase their use of options to hedge their positions. Several previous studies have focused on stock market volatility as a measure of uncertainty, including Bloom (2009).[12] The volatility-uncertainty measure in this analysis updates the series from Bloom's paper, which begins in 1962, through the end of the sample period in 2010.

In general, the two measures of uncertainty tend to move with each other (Chart 3). Their correlation is 0.51. Both uncertainty measures exhibit their highest readings during the 2007-09 recession, and their lowest readings came during the 1960s and the mid-2000s. During the peak of the crisis in 2008, almost 2 percent of all articles in the *Times* contained references to economic uncertainty, compared with only about 0.25 percent of all articles at the beginning of 2007. The volatility index also surged about seven-fold during this period.

Episodes of economic and financial disruption are marked by an increase in the level and volatility of the uncertainty measures. The 1970s, for instance, witnessed large oil price spikes, the breakdown of the Bretton Woods system of fixed exchange rates, multiple recessions, and stagflation. The *Times* uncertainty measure remained elevated and volatile during that period. Stock market volatility exhibited a number of spikes but returned to more normal levels for parts of the decade. Events like the Gulf Wars and the terrorist attacks of September 11, 2001, have also coincided with abrupt spikes in uncertainty.

Closely tying the uncertainty measures to the state of the business cycle is more challenging. Recessions (the gray bars in Chart 3) tend to coincide with spikes in uncertainty using both measures. However, not all spikes in uncertainty are associated with a recession, as evidenced by

Chart 3

TWO UNCERTAINTY MEASURES: *THE NEW YORK TIMES*
RATIO AND STOCK MARKET VOLATILITY
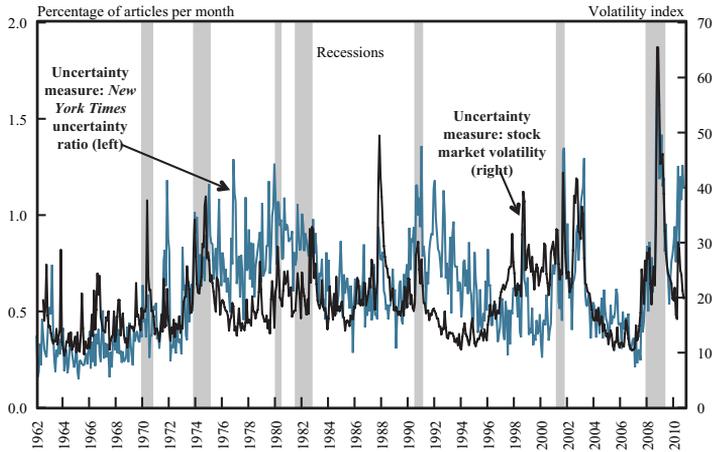
Notes: The *New York Times* uncertainty ratio and stock market volatility are defined in the text. The gray recession bars use the dates provided by the National Bureau of Economic Research.

Black Monday on October 19, 1987, when the S&P 500 index fell 20 percent. At the other extreme, long economic expansions with strong growth tend to coincide with low or declining uncertainty, as in the 1960s or 1990s. But even these expansions coincided with some spikes in uncertainty, such as around the collapse of the Long Term Capital Management hedge fund in 1998.

The two measures exhibit some notable differences. In particular, the *Times* identified the 1970s as a period of generally higher uncertainty than the stock market volatility measure. In addition, during the boom of the 1990s the volatility-uncertainty measure started low and then moved higher, while the *Times* measure started high and moved steadily lower. More generally, the *Times* uncertainty ratio exhibited more spikes during nonrecession periods than the volatility index, perhaps because the *Times* is a broader measure of uncertainty.

III.   REGRESSION ANALYSIS

Recessions tend to coincide with spikes in uncertainty. This raises the possibility that uncertainty shocks might be an important factor in generating recessions as households pull back on their spending. But

not all uncertainty spikes coincide with recessions. Moreover, it is possible that recessions cause uncertainty spikes instead of the reverse. Disentangling these factors in a rigorous way therefore requires empirical analysis to determine uncertainty's role in the business cycle.

This section provides evidence from vector autoregressions (VARs) on how households respond to uncertainty shocks. In simple bivariate VAR models, households appear to behave largely as the theoretical models would suggest: An uncertainty shock produces an immediate drop in spending followed by a rebound.

But when the VAR is augmented to include more variables, these results tend to disappear. In these cases, uncertainty shocks typically are associated with modest reductions in household purchases that can take a considerable time to appear. Moreover, uncertainty shocks account for only a small portion of fluctuations in household spending. These results suggest that uncertainty shocks—at least as they are commonly captured—do not appear to be a key factor driving household spending decisions.

Bivariate models of uncertainty and household spending

As a first pass, it is useful to consider how the uncertainty measures affect household spending in the simplest model: a bivariate VAR. In this model, one uncertainty measure—either the *Times* uncertainty ratio or the stock market volatility index—is paired with one measure of household spending. For household spending, this article considers four separate measures: new home sales, real durable goods consumption, real nondurable goods consumption, and real services consumption.[13] Economic theory predicts that uncertainty shocks should have the largest impact on irreversible spending decisions, such as housing and durables. In contrast, spending on nondurables and services should be little affected, or may actually increase if consumers shift purchases to these items in the face of an uncertainty shock. All data are monthly to capture high-frequency responses to uncertainty shocks and run from July 1962 to October 2010.[14]

Following a *Times* uncertainty shock, new home sales follow a bust-boom pattern similar to the pattern suggested by the real options framework (Chart 4, panel A).[15] In the month that the shock occurs, sales decline immediately and by a considerable amount. As the un-

Chart 4
HOUSEHOLD RESPONSES TO A *NEW YORK TIMES* UNCERTAINTY SHOCK, BIVARIATE MODEL
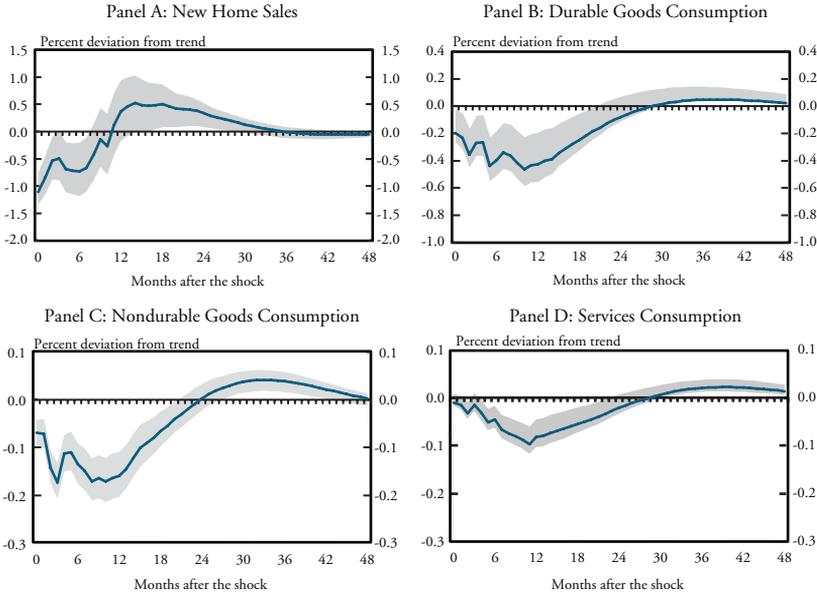
Panel A: New Home Sales

Panel B: Durable Goods Consumption

Panel C: Nondurable Goods Consumption

Panel D: Services Consumption

Notes: Uncertainty shocks are derived from *The New York Times* ratio measure. The uncertainty shock occurs in period 0. The gray shaded area is the one-standard-error band. Given the different sizes of the responses, the vertical axes differ between most panels.

certainty shock subsides, new home sales quickly rebound. Within 12 months of the shock, sales move above their trend level for a time and, within three to four years, return to their normal trend level.

The *Times* uncertainty shock also produces bust-boom dynamics for the three components of consumer spending, though the patterns differ slightly from the response of new home sales. Durable consumption (panel B) and nondurable consumption (panel C) both decline immediately, but consumption of services (panel D) does not. In all three cases, the largest declines are about one year after the shock occurs rather than on impact. Consumer spending moves above its normal trend levels for a time about two years after the shock. This is considerably longer than the bust-boom pattern documented by Bloom (2009) for the manufacturing sector, where activity drops sharply below trend following an uncertainty shock but remains there for only six to 12 months. (The Appendix replicates Bloom's empirical findings

for industrial production and employment in the manufacturing sector in response to large uncertainty shocks.) There is no evidence that an uncertainty shock produces an increase in spending on nondurables and services as consumers substitute away from irreversible purchases, which contrasts with the results of Romer.

The relative sizes of the responses conform to the predictions of the real options framework. The spending declines are largest for the most expensive, partly irreversible items—houses and durable goods. Household spending on nondurable goods and services is less affected.

The household responses are similar following an uncertainty shock based on stock market volatility (Chart 5). The ordering of the size of the responses is the same, although the magnitudes differ modestly. Additionally, new home sales go through a longer period of retrenchment in this case, remaining below trend for one to two years, and do not exhibit the bust-boom cycle seen earlier.

The results from the bivariate regressions generally suggest that uncertainty shocks substantively dampen household spending. The declines in spending occur quickly, often immediately. The largest reductions are on irreversible purchases of new homes and durables. These results are in line with the theoretical predictions of the real options framework developed earlier.[16]

Multivariate models with uncertainty shocks

While bivariate models can be useful, they have important shortcomings. Most notably, bivariate models may omit other relevant explanatory variables. For instance, wealth and income prospects may affect household spending. If a shock to one of these factors affects the uncertainty measures before it affects household spending, then the bivariate regression would mistakenly attribute the subsequent movements in spending to uncertainty shocks rather than the true source.

Capturing more completely the relevant determinants of household spending thus requires incorporating more variables into the VAR. The more complete empirical model contains the following variables: the level of the S&P 500 stock market index, the measure of uncertainty, the federal funds rate, wages, the consumer price index, hours worked, the number of employees on nonfarm payrolls, new single-family home sales, real durable goods consumption, real nondurable goods con-

Chart 5

HOUSEHOLD RESPONSES TO A VOLATILITY-UNCERTAINTY SHOCK, BIVARIATE MODEL
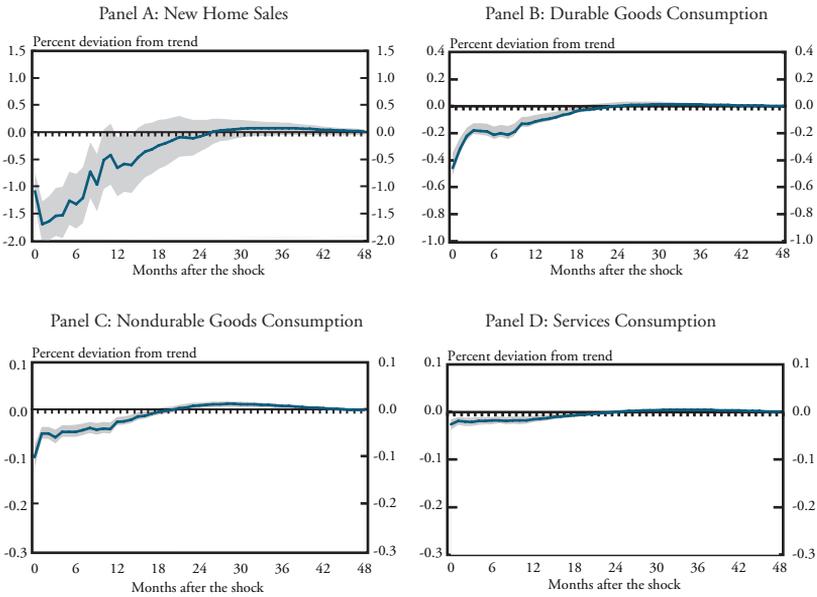
Panel A: New Home Sales

Panel B: Durable Goods Consumption

Panel C: Nondurable Goods Consumption
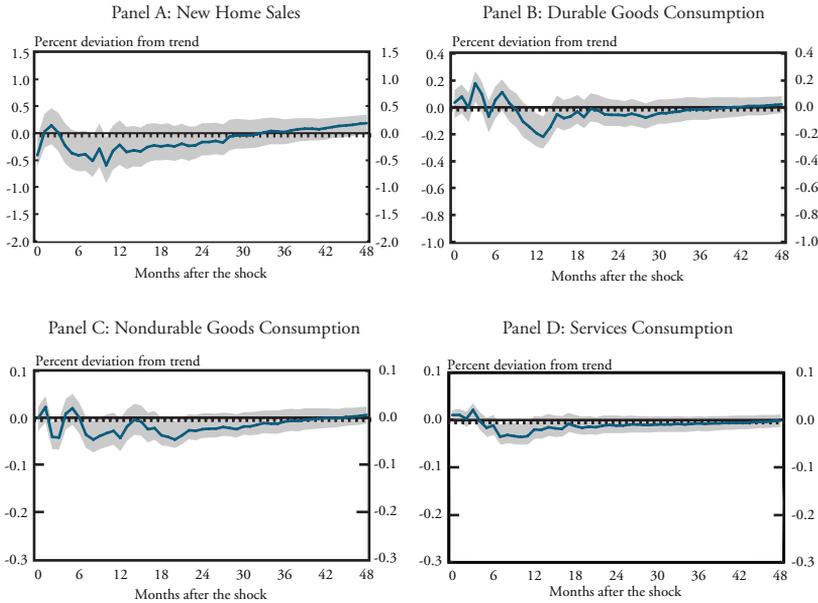
Panel D: Services Consumption

Notes: Uncertainty shocks are derived from stock market volatility. The uncertainty shock occurs in period 0. The gray shaded area is the one-standard-error band. Given the different sizes of the responses, the vertical axes differ between most panels.

sumption, and real services consumption.[17] Data limitations shorten the timeframe to January 1965 through October 2010.

Uncertainty and The New York Times. In the wake of a shock that increases the *Times* uncertainty measure, households behave quite differently than suggested by the bivariate results (Chart 6).[18] In the bivariate model, new home sales experienced a sharp, immediate decline that persisted about a year. But the multivariate results (panel A) show a different pattern. The point estimate (the blue line) implies an immediate drop, then a rebound and subsequent decline that persists almost three years. But the vast majority of this response is not statistically significant. The gray one-standard-error band includes zero for all but months zero, eight, and ten.

The picture is similarly murky for durable goods consumption (panel B), nondurable goods consumption (panel C), and services consumption (panel D). All the point estimates for these measures turn

Chart 6

HOUSEHOLD RESPONSES TO A *NEW YORK TIMES* UNCERTAINTY SHOCK, MULTIVARIATE MODEL

Panel A: New Home Sales

Panel B: Durable Goods Consumption

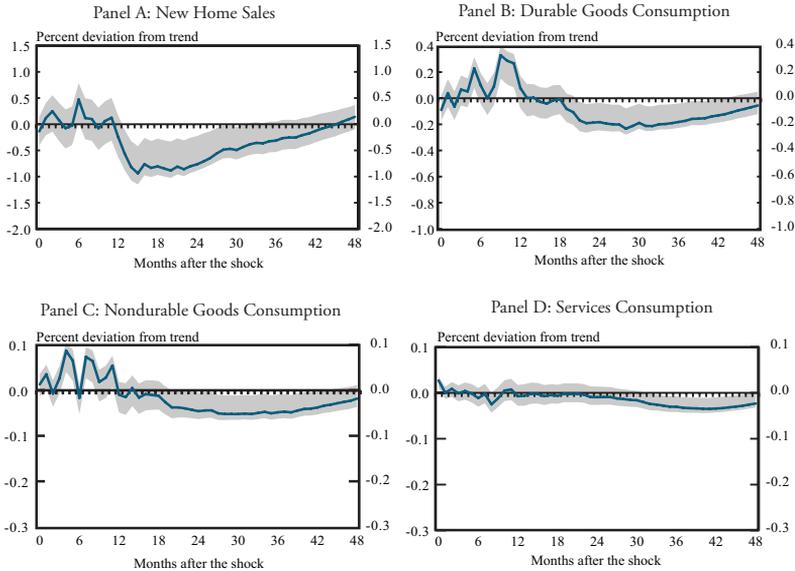Panel C: Nondurable Goods Consumption

Panel D: Services Consumption

Notes: Uncertainty shocks are derived from *The New York Times* ratio measure. The uncertainty shock occurs in period 0. The gray shaded area is the one-standard-error band. Given the different sizes of the responses, the vertical axes differ between most panels.

negative at some point, but not immediately. Notably, spending on consumer durables is initially above trend for a time before turning negative toward the end of the first year after the shock.

Overall, these results suggest that uncertainty shocks as measured by *The New York Times* index only modestly curtail household purchases, after controlling for other factors. There is no evidence of substantial immediate declines in spending, especially on irreversible purchases of items such as new homes and durable goods. These results contradict the intuition from the real options framework and the results from the manufacturing sector presented by Bloom (2009), where an uncertainty shock can have powerful, immediate effects.[19]

Uncertainty and stock market volatility. When uncertainty is measured with the stock market volatility measure, the multivariate analysis yields somewhat surprising results (Chart 7).[20] Immediately following an

Chart 7

HOUSEHOLD RESPONSES TO A VOLATILITY–UNCERTAINTY SHOCK, MULTIVARIATE MODEL

Notes: Uncertainty shocks are derived from stock market volatility. The uncertainty shock occurs in period 0. The gray shaded area is the one-standard-error band. Given the different sizes of the responses, the vertical axes differ between most panels.

uncertainty shock, little happens to new home sales (panel A) and services consumption (panel D). While their behavior is volatile, household spending on durable goods (panel B) and nondurable goods (panel C) actually both rise modestly above trend for a time.[21] Eventually, households respond by reducing all categories of spending relative to their normal levels, but only after a delay of one to three years following the shock.

The household responses for new home sales and durable goods consumption are at odds with the predictions of the real options framework. According to that framework, spending that is highly irreversible should decline immediately as households wait for uncertainty to be resolved.[22]

Overall, these results suggest that uncertainty shocks captured through stock market volatility can weigh on household purchases. However, the dynamics are not immediate and take a considerable time to appear.[23]

Variance decompositions. Evidence from the VARs on the response of household spending to uncertainty shocks suggests shocks can moderately reduce household purchases. Further evidence on the quantitative importance of uncertainty shocks comes from estimating the extent to which they explain movements in household spending relative to other factors. In particular, this exercise—a forecast error variance decomposition—estimates the percentage of movements in the four measures of household purchases that can be explained by the variables in the VAR (Table 1).

The exact quantitative estimates depend on the measure of uncertainty used in the VAR and the time horizon. The first column shows the contributions of uncertainty shocks using the two measures over two different time horizons. For instance, over a 12-month horizon, shocks to the *Times* uncertainty measure explain about 1 percent of the forecast variance in new home sales. By contrast, shocks to the level of the S&P 500 explain about 7 percent of the variance in new home sales. Over a 48-month horizon, the volatility uncertainty measure explains about 5 percent of the variance in new home sales, and the stock market explains about 14 percent.

The variance decomposition shows that uncertainty shocks explain only a small fraction of the movements in household purchases, whether for new home sales, durable goods, nondurable goods, or services. This is true at both short horizons and long horizons. Using stock market volatility as a proxy for household uncertainty explains a larger fraction of the variances, but the percentages remain below 10 percent at all horizons considered. In general, the variance decompositions suggest that uncertainty shocks are a small source of fluctuations in household purchases.

By contrast, shocks to the level of the S&P 500, the federal funds rate, and the consumer price index explain a much larger share of fluctuations.[24] Of these variables, shocks that depress the level of the S&P 500 generate household responses most similar to the predictions of the real options model: Spending on irreversible purchases falls immediately and by a sizable amount, only to rebound shortly thereafter to above-normal levels for a time. However, it is not clear that these responses are necessarily due to uncertainty because stock market declines operate through a number of channels. For example, a decline in the

Table 1

FORECAST ERROR VARIANCE DECOMPOSITION

Percentage of forecast variance explained by shocks to:							
	Uncertainty: NYT ratio	S&P 500	Federal funds rate	Wages	Consumer price index	Hours worked	Employment
Time horizon: 12 months							
New home sales	1	7	15	2	8	7	6
Consumption of durable goods	1	10	10	0	16	3	4
Consumption of nondurable goods	1	17	4	1	16	1	5
Consumption of services	2	10	13	1	4	3	1
Time horizon: 48 months							
New home sales	1	15	15	4	9	5	7
Consumption of durable goods	1	10	19	1	17	2	6
Consumption of nondurable goods	2	15	8	2	17	1	5
Consumption of services	2	14	18	3	9	3	1
	Uncertainty: stock market volatility	S&P 500	Federal funds rate	Wages	Consumer price index	Hours worked	Employment
Time horizon: 12 months							
New home sales	0	8	15	1	6	5	6
Consumption of durable goods	3	12	9	1	13	2	4
Consumption of nondurable goods	3	19	4	1	12	1	5
Consumption of services	0	13	12	1	4	2	0
Time horizon: 48 months							
New home sales	5	14	13	4	11	3	8
Consumption of durable goods	5	12	17	3	17	2	7
Consumption of nondurable goods	4	18	6	5	14	1	7
Consumption of services	3	19	15	5	9	2	3

Notes: NYT stands for *The New York Times*. The remaining terms in the VAR are omitted.

stock market can exert negative wealth effects on households or reflect expectations of future household spending.

Alternative explanations. Why do consumers not exhibit "wait-and-see" behavior following an uncertainty shock, as the real options theory would predict? One possible explanation is that the model generating real options behavior is misspecified. That is, it does not capture the relevant factors that actually determine household spending, so households are not affected by real options to the extent predicted by the theory. A second possibility is that consumers might ignore small uncertainty shocks but respond to large uncertainty shocks. Because the VARs are linear, they assume that household responses are proportional

to the size of the uncertainty shock. This precludes the possibility of radically different responses to small and large uncertainty shocks. But, as shown in the Appendix, focusing only on large uncertainty shocks generates effects on household spending similar to the other measures used in the text. A third possible explanation is that economists have not found the right proxy for uncertainty. Measuring uncertainty is inherently difficult, given its intangible nature.[25] Uncertainty may therefore be captured within the VAR by some other variable.[26]

IV.   CONCLUSION

Uncertainty surged during the financial crisis in 2008 and remained high through a considerable portion of the recovery into 2010. Since then, uncertainty has risen again due to the recent oil price spikes and the March 11, 2011, earthquake and tsunami in Japan. This heightened uncertainty raises the question: How does it affect economic activity?

This article focuses on how households respond to uncertainty shocks—sudden, unexpected events that raise the possibility of extreme future outcomes, either good or bad. Economic theory predicts that household purchases would decline immediately following an uncertainty shock because households would find a value in waiting to make big, irreversible purchases to see how the uncertain environment plays out.

The empirical results, however, suggest that uncertainty shocks tend to curtail household spending only modestly. In some cases, these responses manifest themselves only after a considerable period. In addition, uncertainty shocks account for only a small portion of the total fluctuations in household spending. These results suggest that commonly used measures of uncertainty shocks do not appear to be a key factor driving households' spending decisions and, in turn, economic weakness.

APPENDIX

THE EFFECTS OF LARGE UNCERTAINTY SHOCKS

The uncertainty measures considered in this article allow for a continuum of shock sizes. In addition, the empirical VAR models are all linear: Doubling the size of a given shock doubles the size of the response. It is possible that these assumptions are incorrect when considering uncertainty. In particular, households may ignore small uncertainty shocks, but they may react to large uncertainty shocks in a manner consistent with the real options framework. It may therefore be necessary to separate the small and large shocks when conducting the analysis.[27]

This Appendix considers a measure of large uncertainty shocks suggested by Bloom (2009): A large uncertainty shock occurs in a given month if stock market volatility, as explained in Section II, experiences a large jump in that month.[28] The Appendix first replicates some of Bloom's findings for the manufacturing sector. It then analyzes how households respond to large uncertainty shocks.

The results from large uncertainty shocks on household purchases are generally similar to those from the other measures of uncertainty shocks. While simple bivariate models tend to find that large uncertainty shocks immediately depress spending, the multivariate results are more complicated. In fact, some spending measures actually increase before they experience a substantial decline in the wake of a large uncertainty shock. Large uncertainty shocks based on jumps in stock market volatility therefore do not appear to be a key factor behind weakness in household spending.

Large uncertainty shocks and manufacturing firms

Much of the recent emphasis on uncertainty shocks has been stimulated by evidence from the manufacturing sector (Bloom 2009). A VAR similar to the multivariate VARs in the body of the article can capture how manufacturing firms respond to large uncertainty shocks.[29]

In the aftermath of a large uncertainty shock, manufacturing production and employment follow patterns broadly consistent with the real options framework. Manufacturing production exhibits a short,

sharp decline that lasts about six months (Chart A1, panel A).[30] After that point, the bust gives way to a boom, with production surging above its normal level for one to two years. About three years after the large uncertainty shock occurs, this bust-boom cycle has effectively played itself out and production returns to its normal trend level.

Employment in the manufacturing sector follows these bust-boom movements in production with a slight delay (panel B).[31] Immediately following a large uncertainty shock, manufacturers reduce payrolls below their normal levels; given that production is declining, they do not need the same amount of labor. As production overshoots its normal level, employment rebounds, rising above its normal level as well. About three years after the large uncertainty shock occurs, employment and production both return to normal.

Large uncertainty shocks and household purchases

With this same measure of large uncertainty shocks, it is also possible to use VARs to analyze household behavior. To match the endpoint of Bloom's measure of large uncertainty shocks, the sample ends in June 2008. The VAR specifications are the same as those used in Section III, but with the measure of large uncertainty shocks replacing the old measure of uncertainty.[32]

The bivariate regression results bear a number of similarities to those seen earlier (Chart A2). Quantitatively, the responses are larger, which is not surprising given the fact that the exercise now considers only large uncertainty shocks. But qualitatively, the patterns are generally similar to those from Charts 4 and 5. All measures of household spending decline immediately following a large uncertainty shock, with the greatest declines coming from new home sales and durable goods consumption. Similar to the *Times* uncertainty shock response, the decline in new home sales (panel A) is short-lived and gives way to a surge in sales above trend. This behavior seems to conform closely to the real options framework and the results from the manufacturing sector. However, the other series take a considerable time to return to trend—far longer than the declines in manufacturing.

While all four measures of household spending decline in the wake of the large uncertainty shock in bivariate regressions—as suggested by the real options framework—incorporating large uncertainty shocks

Chart A1

MANUFACTURING RESPONSES TO A LARGE UNCERTAINTY SHOCK
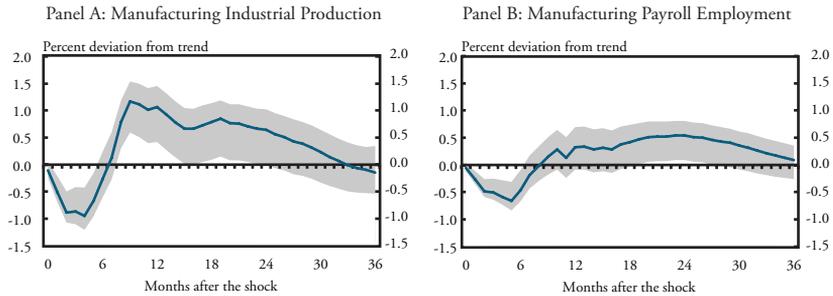
Panel A: Manufacturing Industrial Production
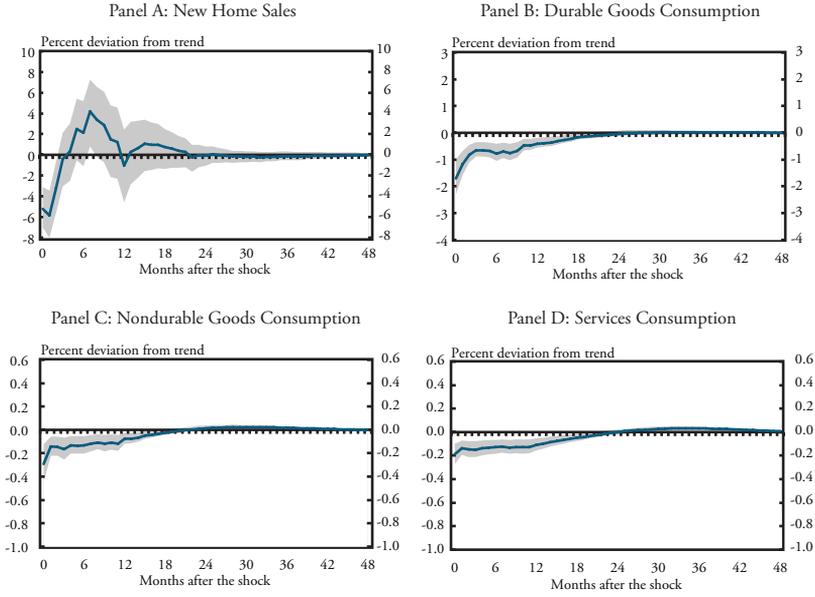
Panel B: Manufacturing Payroll Employment

Notes: The definition of a large uncertainty shock follows from Bloom (2009). The uncertainty shock occurs (takes on the value of 1) in period 0. The gray shaded area is the one-standard-error band.

into a multivariate model produces much different, counterintuitive responses (Chart A3). After declining slightly on impact, new home sales (panel A) rise abruptly above trend for a year before dropping below trend for the subsequent two years. This response is essentially the exact opposite of what the real options framework would predict. Durable goods consumption (panel B) falls on impact but then immediately rebounds, spending some time above trend, before eventually falling back below normal levels long after the large uncertainty shock occurred.

The responses of new home sales and durable goods consumption suggest that large uncertainty shocks based on jumps in stock market volatility are not a key factor behind weakness in household spending.
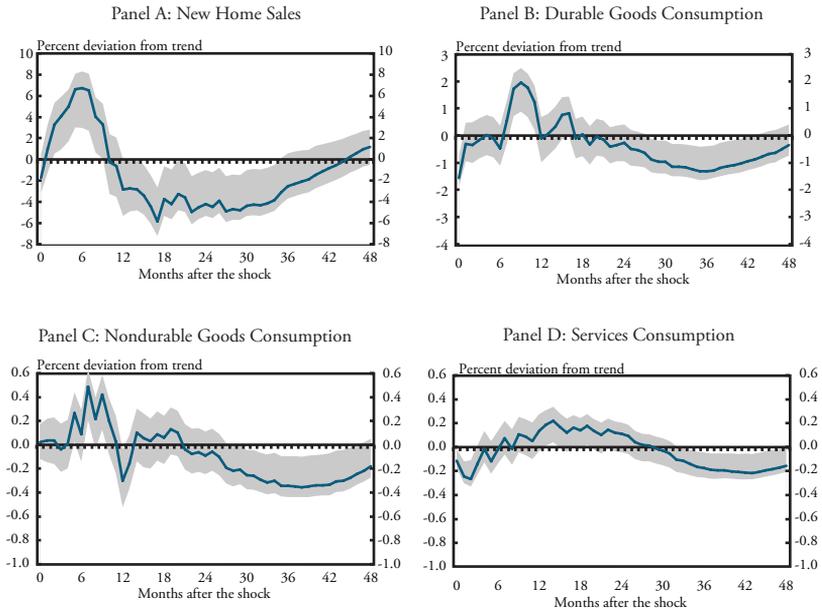
Chart A2

HOUSEHOLD RESPONSES TO A LARGE UNCERTAINTY SHOCK, BIVARIATE MODEL

Panel A: New Home Sales

Panel B: Durable Goods Consumption

Panel C: Nondurable Goods Consumption

Panel D: Services Consumption

Notes: The definition of a large uncertainty shock follows from Bloom (2009). The uncertainty shock occurs (takes on the value of 1) in period 0. The gray shaded area is the one-standard-error band. Given the different sizes of the responses, the vertical axes differ between most panels.

Chart A3

HOUSEHOLD RESPONSES TO A LARGE UNCERTAINTY SHOCK, MULTIVARIATE MODEL

Panel A: New Home Sales

Panel B: Durable Goods Consumption

Panel C: Nondurable Goods Consumption

Panel D: Services Consumption

Notes: The definition of a large uncertainty shock follows from Bloom (2009). The uncertainty shock occurs (takes on the value of 1) in period 0. The gray shaded area is the one-standard-error band. Given the different sizes of the responses, the vertical axes differ between most panels.

ENDNOTES

[1]That is, there are second-moment shocks that occur alongside the more traditional first-moment shocks that are a standard feature of modern dynamic stochastic general equilibrium (DSGE) models.

[2]Note that given the possible outcomes and the associated probabilities, the means of the two distributions in the chart are identical, but the variance of the distribution in the black bars is higher. Most DSGE models allow for normally distributed stochastic shocks. Consequently, an increase in the variance of the distribution does not imply a wider set of possible outcomes; with a normal distribution, these are infinite. However, an increase in uncertainty under a discrete approximation, such as in the chart, could witness a wider range of potential outcomes as well as higher probabilities for the tail events.

[3]While uncertainty shocks allow for changes to the probability distribution, all the possible outcomes are still known. This contrasts with the notion of "Knightian uncertainty," in which uncertainty and possible outcomes are ultimately unmeasurable and unquantifiable (Knight 1921). While this type of uncertainty is relevant for the real world, capturing it in even the most state-of-the-art macro models has proven challenging.

[4]Bernanke (1983) is one of the early works that introduced real options theory into firms' investment decisions over the business cycle. Dixit and Pindyck (1994) provide a thorough treatment of real options. Bloom (2009) has restimulated interest in uncertainty shocks and real options for firms' investment decisions. Davig and Hakkio (2010) consider the real options framework for their analysis of the effects of financial stress on economic activity, but ultimately they focus on a model with a financial accelerator.

[5]Economists usually use (S,s) models to explain firms' decisions in the presence of fixed costs and irreversibilities. These (S,s) models generate thresholds, or (S,s) bands, such that the firm will only take action if it finds itself pushed beyond those thresholds. Between the thresholds is a range of inaction, in which the firm will choose to wait before acting. When a firm does act, it reoptimizes and places itself back in the range of inaction. In this sense, the firm may decide not to respond to small shocks if it is still in the interior of the range of inaction, instead allowing shocks to accumulate before the firm is pushed to its (S,s) band and acts.

[6]In an (S,s) framework, the uncertainty shock expands the (S,s) bands, which can lead firms to delay taking actions.

[7]Bloom (2007, 2009); Bloom, et al. (2010); and Alexopoulos and Cohen (2009) provide results supporting a key role for uncertainty shocks in business cycle fluctuations. The results presented in Bachmann and Bayer (2011) and Bachmann, et al. (2010) question these findings.

[8]Romer (1990); Eberly (1994); Carroll and Dunn (1997); Foote, et al. (2000); Hassler (2001); and Bertola, et al. (2005) provide additional detail on

households' purchases of housing and durables under various aspects of uncertainty. While some of these studies focus on consumer spending at the level of the individual household, this article focuses on aggregate household spending.

[9]Some categories of nondurable goods, such as clothing, might be better referred to as "semidurable"; others, such as fresh fruit, are clearly "perishable." Romer's (1990) study using data from around the beginning of the Great Depression in 1929 and 1930 grouped goods into the categories of durable goods, semidurable goods, and perishable goods. The more recent convention, however, is to divide consumer spending into durable goods, nondurable goods, and services.

[10]This ratio was constructed in multiple steps. The first step required searching through *The New York Times'* website (*http://www.nytimes.com*) for articles that contained the words "uncertain" or "uncertainty," along with the words "economy" or "economic." Because the *Times'* web search function is relatively rudimentary, this required conducting nine different searches for each month. These nine searches were:

> uncertain economy -uncertainty -economic
> uncertain economic -uncertainty -economy
> uncertainty economy -uncertain -economic
> uncertainty economic -uncertain -economy
> uncertain economy uncertainty -economic
> uncertain economy economic -uncertainty
> uncertain economic uncertainty -economy
> uncertainty economy economic -uncertain
> uncertain economic uncertainty economy

A hyphen in front of an identifier indicated exclusion from the search criterion. Hence, the combination "uncertain economy -uncertainty -economic" searched all the articles that contained the words "uncertain" and "economy," and that did not contain the words "uncertainty" and "economic."

The second step required determining the total number of news articles published during a month. To approximate this number, the word "the" was searched, since this was assumed to be the most common and essential word in an article. *The Times* uncertainty measure is thus the ratio of the sum of the nine "uncertainty" searches to the number of articles containing "the."

During the sample period, several months provided erroneous search results. The search results for August 1978 produced only 23 percent of the average number of monthly articles published in previous months, and *The New York Times* search engine provided no results for September and October 1978. To smooth through these missing data points, linear interpolation was used between July and November 1978.

Search results for May and July 2010 produced repetitions for the "uncertainty" searches. (That is, some articles appeared in the search results more than once.) The sum of the nine "uncertainty" searches was corrected by subtracting the number of repeated articles from the total, where the number of repeated

articles was computed by hand. Although it was feasible to make this correction for the "uncertainty" articles, the same correction for the total number of articles containing "the" was infeasible. (The number of returned "the" articles was greater than 12,000 in these months; surrounding months had fewer than 9,000 articles.) To correct the denominator of the frequency measure, the ratio of non-repeated "uncertainty" searches to total "uncertainty" searches was multiplied by the number of "the" searches.

[11]Following Bloom (2009), the stock market volatility-uncertainty measure is built using the VXO index (*http://www.cboe.com/micro/vxo/*) of percentage implied volatility of the S&P 100 30-day option. This index dates back to 1986. Prior to that, the actual volatility of the S&P 500 index is used as a proxy for uncertainty. VXO data are available on the Chicago Board of Options Exchange website (*http://www.cboe.com/micro/vix/historical.aspx*) at a daily frequency and were converted to a monthly frequency by averaging. More details can be found at Nicholas Bloom's website, *http://www.stanford.edu/~nbloom/.*

[12]Among others, see also Romer (1990) and Hassler (2001). Bloom (2009) identified uncertainty shocks as discrete events. An uncertainty shock occurred within a given month (i.e., took on a value of one) if the detrended stock market volatility level rose more than 1.65 standard deviations above the mean; otherwise, an uncertainty shock did not occur (i.e., took on a value of zero). This resulted in a set of 17 uncertainty shocks between 1962 and 2008. See the Appendix of Bloom (2009) for more details. The Appendix to this article takes a closer look at Bloom's identified uncertainty shocks and the responses of household spending.

[13]The measure of new single-family home sales is based on signed contracts: To the extent that households would immediately pull back in the wake of an uncertainty shock, this measure would capture such behavior better than existing home sales, which are based on closings. While housing starts and permits could be another way to capture housing activity, some of this activity would reflect speculation on the part of builders rather than purchase activity on the part of households. The measures of monthly real consumer spending on durables, nondurables, and services come from the Bureau of Economic Analysis for 1995 onward. Prior to 1995, the monthly real measures were constructed by deflating nominal consumer spending by the respective price deflator for that type of spending.

[14]All variables enter the baseline VARs in detrended form after using the Hodrick-Prescott filter with weight 129,600, following Bloom (2009). The measures of household spending are detrended in natural log form. All VARs use 12 lags given the monthly data frequency and include a constant term. The uncertainty shock is one standard deviation. All standard error bands are constructed via 1,000 bootstraps.

[15]The bivariate VARs use a Cholesky decomposition to identify the orthogonalized uncertainty shocks, with the uncertainty measure ordered first and the household spending measure second.

[16]In many cases, these results are sensitive to the decision to detrend the data before running the VARs. One alternative is to run the bivariate VARs using one uncertainty measure and the natural log of one measure of household spending, neither of which is detrended. In this case, all eight combinations show declines in household spending. New home sales decline immediately and eventually return within one to two years to their starting values using either measure for the uncertainty shock. For the other six combinations, the point estimates for the spending responses do not return to their original starting points within four years. A second alternative is to consider a Blanchard-Quah decomposition, in the spirit of Bachmann, et al. (2010), where uncertainty shocks are identified using the long-run restriction that they cannot have permanent effects on real variables. In this case, new home sales using either uncertainty measure experience statistically significant declines following an uncertainty shock. But the results for the other six combinations are markedly different. Using *The New York Times* measure, point estimates for real spending on durables, nondurables, and services all rise following the shock, though the increases are rarely statistically significant. Using the stock market volatility measure, the responses are highly volatile but largely centered around zero.

[17]The choice of variables and the baseline ordering is based on Bloom (2009); alternative orderings are considered below. Ordering potentially matters because orthogonalized shocks are identified using a Cholesky decomposition. The measure of wages is total private average hourly earnings of production workers. The measure of hours is total private average weekly hours of production workers. All variables enter the VAR in detrended form after using the Hodrick-Prescott filter with weight 129,600. The measures of household spending, the S&P 500 index level, wages, the consumer price index, hours, and nonfarm payrolls are detrended in natural log form.

[18]The shock is one standard deviation, so the uncertainty ratio index immediately jumps one standard deviation above its normal, trend level in month zero. The jump in uncertainty is short-lived: it spikes in month 0, and within three to four months it has returned to its normal level, based on either *The New York Times* measure or the stock market volatility measure.

[19]These results also contrast with those of Alexopolous and Cohen (2009), who find a bust-boom pattern in consumer spending (though they do not consider housing) similar to Bloom (2009). There are several potential explanations for this. First, Alexopolous and Cohen use the total number of uncertainty articles in *The New York Times* in their regressions (i.e., the level), rather than the ratio of uncertainty articles to total articles. In our baseline VAR with the detrended level of uncertainty substituted for the ratio, the four measures of household spending all experience statistically significant declines following an uncertainty shock. However, it takes some time for these declines to appear—from about four months to one year after the shock. None of the measures experiences a statistically signifi-

cant boom thereafter. Second, Alexopolous and Cohen consider smaller VARs, which will tend to push the results closer to the bivariate models.

[20]Romer (1990) and Hassler (2001) both focus on stock market volatility as a proxy for uncertainty when studying household behavior.

[21]This is reminiscent of the "price puzzle," in which VAR studies often find that the price level initially rises following a positive shock to (i.e., a tightening of) the federal funds rate.

[22]The positive response of nondurables consumption and the immediate (though small) jump in services consumption are consistent with the predictions of Romer (1990).

[23]It is also possible to consider different orderings of the variables in the VAR, for both *The New York Times* and stock market volatility measures of uncertainty, though with 11 variables it is not possible to try all combinations.

For instance, Christiano, et al. (2005) suggest an ordering of real variables, then prices, then the federal funds rate. This suggested: employment; hours; real consumption of services; real consumption of nondurables; real consumption of durables; new single-family home sales; the consumer price index; wages; the federal funds rate; the S&P 500 stock index; and *The New York Times* uncertainty measure. This ordering produced generally similar impulse responses to Chart 6: The standard error bands continued to straddle zero for most of the horizon, suggesting modestly negative impacts of uncertainty shocks on household purchases. Putting the S&P 500 first, followed by *The New York Times* uncertainty measure, followed by the remaining variables in order did not change these results.

Beaudry and Portier (2006) put the S&P 500 last, suggesting: the New York Times uncertainty measure; the federal funds rate; wages; the consumer price index; hours worked; the number of employees on nonfarm payrolls; new single-family home sales; real consumer spending on durable goods; real consumer spending on nondurable goods; real consumer spending on services; and the S&P 500. In this case, new single-family home sales immediately fall in response to an uncertainty shock by a statistically significant amount, but within two months the response is once again statistically insignificant; the drop in new home sales lasts for four months using stock market volatility as the uncertainty measure. The responses of the other series are volatile and modestly negative on net, with some statistically significant readings and others that are not.

Another possibility is to run the VARs with the variables in levels (natural log levels for all variables other than the uncertainty measure and the federal funds rate) rather than transforming them to deviations from trend. Doing so did not greatly alter the impulse responses for *The New York Times* uncertainty measure. Once again, the point estimates for the responses all experienced declines. However, the declines were rarely statistically significant. In the VAR using (non-detrended) stock market volatility as the uncertainty measure, the point estimates for the four household spending measures actually increased rather than

decreased; the increases were statistically significant for the three measures based on real consumption.

[24]This result contrasts with the findings of Romer (1990), who finds that stock market volatility is a more significant factor in driving consumer spending than the level of the stock market. However, it is worth pointing out that Romer defines volatility as a backward-looking measure: It is the average squared monthly change in real stock prices over the previous 12 months. This article follows Bloom (2009) in defining volatility from stock market options for the post-1986 period and the monthly volatility of the S&P 500 index before that period.

[25]There are other possible uncertainty measures, including Carroll and Dunn's (1997) preferred measure of household uncertainty: the difference between the fraction of consumers in the University of Michigan's survey who thought unemployment would rise over the next 12 months minus the fraction who thought unemployment would fall. Using monthly data from January 1978 through October 2010 (this measure was not available on a monthly basis prior to 1978), incorporating this measure for uncertainty into the multivariate VAR did not produce responses consistent with the real options framework.

[26]The measures of uncertainty considered in this article focus on aggregate uncertainty. Instead, it is possible that individual households respond using the real options framework to their own uncertainty shocks, but these household-level uncertainty shocks are not necessarily correlated with aggregate uncertainty.

[27]Davig and Hakkio (2010) consider a two-regime world that captures the essence of this argument, though they do not focus on household purchases.

[28]Bloom identifies the uncertainty shocks as discrete events when the detrended stock market volatility level in a month rose more than 1.65 standard deviations above its mean. See the Appendix of Bloom for more details.

[29]As in Bloom (2009), the VAR includes the following variables, in their Cholesky ordering: S&P 500, uncertainty indicator, federal funds rate, wages in manufacturing, consumer price index, hours worked in manufacturing, employment in manufacturing, and industrial production in manufacturing. With the exception of the uncertainty indicator, all other variables enter as deviations from their Hodrick-Prescott filtered trends, with smoothing parameter 129,600 for monthly data. The data come from the website of Nicholas Bloom; they run from July 1962 through June 2008.

[30]All impulse responses in this Appendix are normalized so that the large uncertainty shock indicator takes on the value of 1 in month 0. This chart is a reproduction of Figure 2 in Bloom (2009).

[31]This chart is a reproduction of Figure 3 in Bloom (2009).

[32]The 0-1 large uncertainty shock indicator series is not detrended. In the baseline model results presented in the charts, all other variables enter the VAR in detrended form.

REFERENCES

Akerlof, George A. (1970). "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics,* 84(3): 488-500.

Alexopolous, Michelle and Jon Cohen (2009). "Uncertain Times, Uncertain Measures," University of Toronto Department of Economics Working Paper 352.

Bachmann, Rüdiger, and Christian Bayer (2011). "Uncertainty Business Cycles—Really?" unpublished manuscript.

Bachmann, Rüdiger, Steffen Elstener, and Eric Sims (2010). "Uncertainty and Economic Activity: Evidence from Business Survey Data," unpublished manuscript.

Beaudry, Paul and Franck Portier (2006). "Stock Prices, News, and Economic Fluctuations," *American Economic Review* 96(4): 1293-1307.

Bernanke, Ben S. (1983). "Irreversibility, Uncertainty, and Cyclical Investment," *Quarterly Journal of Economics,* 98(1): 85-106.

Bertola, Giuseppe, Luigi Guiso, and Luigi Pistaferri (2005). "Uncertainty and Consumer Durables Adjustment," *Review of Economic Studies,* 72: 973-1007.

Bloom, Nicholas (2009). "The Impact of Uncertainty Shocks," *Econometrica,* 77(3): 623-85.

Bloom, Nicholas (2007). "Uncertainty and the Dynamics of R&D," *American Economic Review Papers and Proceedings,* 97(2): 250-255.

Bloom, Nicholas, Max Floetotto, and Nir Jaimovich (2009). "Really Uncertain Business Cycles," unpublished manuscript.

Carroll, Christopher D. and Wendy E. Dunn (1997). "Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets," *NBER Macroeconomics Annual,* 12: 165-217.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (2005). "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy,* 113(1): 1-45.

Davig, Troy and Craig Hakkio (2010). "What Is the Effect of Financial Stress on Economic Activity?" Federal Reserve Bank of Kansas City, *Economic Review* 95(2), second quarter: 35-62.

Dixit, Avinash K. and Robert S. Pindyck (1994). *Investment under Uncertainty,* Princeton University Press: Princeton, New Jersey.

Eberly, Janice C. (1994). "Adjustment of Consumers' Durables Stocks: Evidence from Automobile Purchases," *Journal of Political Economy* 102(3): 403-436.

Foote, Christopher, Erik Hurst, and John Leahy (2000). "Testing the (S,s) Model," *American Economic Review,* Papers and Proceedings, 90(2): 116-119.

Hassler, John (2001). "Uncertainty and the Timing of Automobile Purchases" *Scandinavian Journal of Economics,* 103(2): 351-366.

Knight, Frank (1921). *Risk, Uncertainty, and Profit,* Houghton Mifflin Company: Boston, MA.

Romer, Christina D. (1990). "The Great Crash and the Onset of the Great Depression," *Quarterly Journal of Economics,* 105(3): 597-624.