W.E. UPJOHN INSTITUTE
FOR EMPLOYMENT RESEARCH

Upjohn Research

Upjohn Institute Working Papers

Upjohn Research home page

6-7-2013

Expenditure, Confidence, and Uncertainty: Identifying Shocks to Consumer Confidence Using Daily Data

Marta Lachowska
W.E. Upjohn Institute for Employment Research, marta@upjohn.org

Follow this and additional works at: https://research.upjohn.org/up_workingpapers

Part of the Behavioral Economics Commons

Citation

Expenditure, Confidence, and Uncertainty: Identifying Shocks to Consumer Confidence Using Daily Data

Marta Lachowska
W.E. Upjohn Institute for Employment Research
e-mail: lachowska@upjohn.org

June 7, 2013

ABSTRACT

The importance of consumer confidence in stimulating economic activity is a disputed issue in macroeconomics. Do changes in confidence represent autonomous fluctuations in optimism, independent of information on economic fundamentals, or are they a reflection of economic news? I study this question by using high-frequency microdata on spending and consumer confidence, and I find that consumer confidence contains information relevant to predicting spending, independent from other indicators. The exogenous movements in consumer confidence lead to very short fluctuations in consumer spending, consistent with the hypothesis that more consumer confidence reflects less uncertainty about the future.

1 Introduction

Consumer confidence indexes are regularly reported by the press as an indicator of economic prospects. Despite the ubiquitous presence of the word "confidence" in public discourse, it is not well understood whether confidence measures merely reflect information contained in other economic indicators or whether they contain additional independent information about the economy. Researchers have addressed this question by trying to estimate a causal relationship between indexes of consumer confidence and other macroeconomic variables, such as consumption and output. (See Ludvigson [2004] for a survey of the literature.)

The standard approach to inferring the meaning of consumer confidence is to observe its impact on some margin of economic activity, typically consumption, holding other factors constant. The theoretical framework, as well as the empirical strategy, used to study this issue has varied. For example, Mishkin (1978), Carroll, Fuhrer, and Wilcox (1994), Acemoglu and Scott (1994), and Souleles (2004) discuss their results in the context of variants of the life-cycle model and reach different conclusions regarding the role of consumer confidence.[1]

In a recent paper, Barsky and Sims (2012) study the joint dynamics of consumer confidence, income, and consumption in a structural vector autoregressive model. In their interpretation, the empirical response of consumption following a shock to confidence informs us about the meaning of consumer confidence. In their model, confidence may represent an autonomous change in beliefs that affects economic activity (the "animal spirits" view) or may incorporate future news about the economy (the "news" view). Barsky and Sims find that confidence, to a large degree, reflects news about future output. In contrast, Starr (2012) concludes that a substantial part of variation in consumer confidence is due to nonfundamentals. In sum, why consumer confidence predicts spending remains in dispute.[2]

---

[1]For instance, Carroll, Fuhrer, and Wilcox (1994) find that confidence helps to forecast changes in spending, which violates the rational-expectations version of the permanent income hypothesis. In addition, they find that consumer confidence and consumption growth are positively correlated. On the other hand, Souleles (2004), who imputes consumer confidence in the Consumer Expenditure Survey, finds that consumption growth and consumer confidence are negtively correlated.

[2]There also exists a literature studying consumer confidence in the context of sunspot theories; see for example Matsusaka and Sbordone (1995) and Chauvet and Guo (2003).

The aim of this paper is to understand whether consumer confidence contains information relevant to spending decisions that is not contained in other economic indicators. As Starr (2012) and Barsky and Sims (2012) have done, I study this issue in a vector autoregressive (VAR) model.

The primary source of data is the G1K, a survey that collects information on spending and consumer confidence at a daily frequency for a large number of respondents. The unique feature of the survey is that the number of daily interviews is large enough to construct a daily time series of expenditures and consumer confidence from. The available data cover the year 2008, which is a period of particular interest because it is the first full year of a deep U.S. recession. In particular, the last quarter of 2008 witnessed a contraction of the economy that was not anticipated, as is evident from the revised expectations in the Survey of Professional Forecasters.[3] With the available G1K data, it is possible to study the sensitivity of daily spending to a highly uncertain and changing economic environment.

The other economic indicators come from various sources. First, I collect a set of "Wall Street" indicators, consisting of indexes of stock market prices and stock market volatility. Second, borrowing from the taxonomy in Alexopoulos and Cohen (2009), I collect a "Main Street" indicator of economic prospects, namely a daily series of newspaper articles that mention the word "recession." Since 2008 saw a lot of news coverage about the unfolding economic turmoil, it is plausible to expect consumer confidence and spending to adjust to article shocks. I also include oil prices and a measure of unemployment risk.

In the empirical part of the paper, I estimate impulse response functions of spending following a shock to consumer confidence. First, I study the dynamic response of spending to a shock to consumer confidence in a VAR model with only two variables, spending and consumer confidence. Second, I examine the response of spending to confidence shocks in a richer VAR model that conditions the response of spending on the movements of additional economic indicators. The objective is to understand whether consumer confidence has additional predictive power for spending, controlling for the variation in these other economic variables.

---

[3]http://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/2008/spfq308.pdf

The identification strategy used is based on the daily frequency of the data. I assume that, within a day, spending and consumer confidence may be affected by economic indicators, but not the other way around. A similar identification strategy is used by Alexopoulos and Cohen (2009), Knotek II and Khan (2011), and Starr (2012), who assume that within a *month* consumption is affected by news, but not vice versa. A possible concern with this identification is that within a month a drop in consumption may itself become a news event affecting consumption. This ought to be less likely within a day.[4]

The paper has two main findings. First, consumer confidence contains relevant information for predicting spending, separate from information contained in news reports about the economy, but partly explained by fluctuations of the stock market. Second, spending reacts to innovations in consumer confidence, but this reaction fades within weeks. The pattern of the reaction of spending following a confidence shock is robust to various alternative estimation orders, which suggests that confidence carries information independent of other available economic information. The shape of the impulse responses suggests that the behavior of spending following an innovation in confidence is consistent with the interpretation that more consumer confidence reflects less uncertainty about the future.

The rest of the paper is organized as follows. Section Two describes the data. The primary source of data is the G1K, a survey conducted each day with about 1,000 individuals living in the United States. First, I discuss this data set and how I construct the variables. In addition to the G1K survey, which collects daily data on spending, consumer confidence, and job market conditions, I describe other daily measures of economic information, including stock market prices, stock market volatility, article counts of the word "recession," and oil prices. In Section Three, I discuss the identification strategy of the paper. I describe the VAR models and the estimation orders used to identify the response of spending following a confidence shock. I begin by fitting a simple two-variable VAR consisting of consumer spending and confidence. This simple model is the

---

[4]Since the paper uses news article counts as one of the economic indicators, the identification strategy is also related to the "narrative approach" studies; see, for example, Romer and Romer (1990), Ramey and Shapiro (1998), and Ramey (2011).

benchmark for further analysis. Later, I study a system where, in addition to consumer spending and confidence, other economic indicators are introduced—first alternatively, then jointly. Section Four describes the results, and Section Five interprets the findings in the light of economic theory and the nature of the daily expenditure data. The final section draws conclusions.

2 Data

2.1 The G1K survey

The data on spending and consumer confidence come from a Gallup Organization survey, G1K, which is conducted daily by telephone interviews with a random sample of about 1,000 individuals aged 18 and older living in the United States. Each day a new cross section is drawn, and the survey is conducted seven days a week, excluding major holidays.

Gallup collects the data using a dual-frame random-digit dialing of both landlines and cellular phones. The interviews are conducted with a respondent who is 18 years of age or older, living in the household, and who had his or her birthday most recently. In order to make the sample representative, Gallup provides survey sampling weights to correspond to the national distribution of age, gender, race, region, and educational level.[5] The data available in this paper collected information on about 359,000 individuals surveyed from January 2, 2008, to January 5, 2009.[6]

The G1K survey poses the following question about daily expenditures to a random half-sample of the respondents:

> Next, we'd like you to think about your spending yesterday, not counting the purchase
> of a home, motor vehicle, or your normal household bills. How much money did you
> spend or charge yesterday on all other types of purchases you may have made, such as
> at a store, restaurant, gas station, on-line, or elsewhere?

---

[5]Krueger and Kuziemko (2011) use the data to estimate the price elasticity of the demand for health insurance. The G1K data were also used by Deaton and Arora (2009) in a study on the benefits of height. Deaton (2012) uses the same data to track how the financial crisis has affected subjective well-being in the United States.

[6]The suvey was not conducted on the following days: January 21, February 21, March 23, May 26, May 30, June 18, June 27, July 4, November 27, December 23-25, December 29, December 31, and January 1.

The G1K question is a total expenditure question, which measures the dollar amount spent on goods and services while excluding some of the biggest durables, such as the purchase of a home and car. Browning, Crossley, and Weber (2003) discuss the usefulness of such "one-shot" total expenditure questions asked in general purpose surveys. They recommend that when collecting retrospective information on total expenditure, among other items, one ought to 1) ask about nondurables and services, 2) specify an exact time frame of the recall, 3) and provide a list of specific subitems as cues for the respondent.[7] The G1K question broadly follows this recommendation by specifying a time frame (yesterday) listing specific subitems (purchases at the store, restaurant, gas station, or on-line), and asking the respondent to exclude large durables.

At the same time, in the very short run, most goods are storable or, in effect, "durable." Browning and Crossley (2009) show that within a month about 20 percent of household expenses go to buying "small durables," such as home entertainment equipment, cosmetics, and clothes. Many such items have little resale value and are hence irreversible. Interestingly, Browning and Crossley show that following transitory shocks to income, households use the timing of replacement of small durables as a smoothing mechanism. Households tend to cut back relatively more on clothing expenditures and dining out than on necessities such as food at home. Because the available data do not allow one to distinguish particular categories of goods and services, the results of this paper ought to be considered in the light of Browning and Crossley's findings.

Survey data on spending often contain measurement error and, when compared to aggregate data equivalents, spending tends to get underreported. When comparing the level of the G1K spending to that from different subaggregates of the national accounts, I find that the survey measure is lower than its subaggregate counterparts. In order to gauge the ratio of true variation in the G1K data to total variation, I use a method of correlating two noisy measures in order to extract a signal-to-noise metric; see Ashenfelter and Krueger (1994). The estimate of the reliability ratio in the expenditure measure is about two-thirds, suggesting that the spending variable is not dom-

---

[7]Browning, Crossley, and Weber (2003) point to the tendency of such total expenditure questions to underreport, but conclude that such questions still contain informative variation.

inated by noise and contains a large fraction of information.[8] Ashenfelter and Krueger use this method to validate a measure of years of schooling by correlating two noisy measures of years of schooling; they obtain a reliability ratio of about 90 percent. This is perhaps not surprising given that the recall problem with regard to one's education might be smaller than recall problems with regard to how much one spent yesterday. Although one would like as high a reliability ratio as possible, given the nature of the data, a ratio of two-thirds may be considered as high.

The consumer confidence measure is collected from the same subsample as the spending measure. The respondent is asked to evaluate the present economic conditions in the United States. The question reads as follows:

> Right now, do you think that economic conditions in this country, as a whole, are getting better or getting worse?

The possible responses are measured as "getting worse," "the same," and "getting better."

The G1K confidence question puts an emphasis on the present ("right now"), but asks the respondent about the economic conditions using a present progressive tense ("getting better/worse"). This is different from the questions examining current and future conditions contained in the University of Michigan and Conference Board surveys, which ask about conditions as they are right now (good/bad) or give a clear time frame for the future conditions question (in the next 12 months; 5 years from now); see Ludvigson (2004). it is not clear how the respondents understand the question, but one may speculate that the G1K confidence measure is mix of perceptions about present economic conditions and, perhaps, perceptions of the very near future.

Most previous research has used the composite indexes of consumer confidence from either the University of Michigan or Conference Board surveys. The indexes are weighted averages on a variety of questions regarding past, current, and future economic conditions. There are also studies that focus on only one question. For example, Carroll and Dunn (1997) use the University of Michigan question on unemployment expectations as their measure of consumer confidence.

---

[8]To estimate the reliability ratio, I use a random number generator to split the G1K in two halves and construct two time series of spending. I then regress one series on the other. These results are available from the author.

Also, Barsky and Sims (2012) note that in a VAR analysis it is hard to interpret the meaning of an index mixing past, present, and future perceptions and focus on the forward-looking question about economic conditions.

2.2 Economic indicators

The data on other daily economic indicators come from several sources. First, I include information on stock market prices. I use the daily series of closing prices on the Standard & Poor's (S&P) 500 index. Poterba and Samwick (1995) contrast two not necessarily exclusive reasons for why stock prices may affect consumption. On one hand, there might be a dependency of consumption on stock market prices through a wealth effect. On the other hand, stock prices might act as a leading indicator of economic prospects. Poterba and Samwick conclude that "the effect of stock price fluctuations on consumption operates through channels other than a direct wealth effect, for example by altering consumer confidence" (p. 356). Therefore, the information stemming from stock market indexes could be relevant to individuals' spending decisions and consumer confidence, even if the individuals interviewed in the G1K do not own stock themselves.[9] Recently, Beaudry and Portier (2006) have highlighted the importance of stock prices as predictors of future productivity, and stock prices are commonly included in VARs as measures of "news" shocks.

Second, in addition to including information on stock market prices, I also include the index of stock market volatility, the VIX. The VIX is a measure of the market's expectations of the volatility of S&P 500 index options over the next 30 days and is commonly used as an indicator of uncertainty; see Bloom (2009). Hence, I allow shocks to stock market volatility to affect spending and confidence separately from shocks to the stock market level.

Third, the expenditure variable in the G1K specifically lists gas prices as an example of daily purchases. As gas expenses are a nonnegligible part of U.S. consumers' consumption baskets, oil price shocks could be a relevant predictor of spending. Edelstein and Killian (2007) find that shocks to energy prices reduce consumer spending mainly because of a spending-power effect;

---

[9]A similar conclusion was drawn by Romer (1990) in her study of the effects of the Great Crash on spending.

following an increase in energy prices, once consumers have paid their energy bills, there is less income with which to purchase goods and services.[10] As gasoline consumption was highly volatile during 2008 (see Petev, Pistaferri, and Saporta-Eksten [2012]) with dramatic variations in the retail price of gas (see Hamilton 2009), it seems appropriate to control for oil prices in the estimation. Because of its availability, I use daily data from the United States Oil Fund (USO), which tracks the movements of West Texas Intermediate (WTI) crude oil.

Fourth, following the approach in Doms and Morin (2004) and Alexopoulos and Cohen (2009), I collect information on how often the press reports about the recession. The idea is that in 2008 the volume of news reports about the recession was very high, especially in the last quarter, and this reporting could be more salient to consumer spending and confidence than gyrations of the stock market per se. To construct this measure, I calculate the number of articles in the *New York Times* on a given day that contain the term "recession" in the headline or the lead of the article.[11] As Alexopoulos and Cohen argue, the *New York Times* is considered the national newspaper of the United States and a leading news source both for the public and for other news outlets. I concentrate on the headline and lead sections of the articles, as presumably their purpose is to capture the attention of the reader and sum up the focus of the story. I focus on the word "recession," as it was arguably one of the most frequently reported economic words of 2008 and a word with an unambiguously negative connotation.[12]

---

[10]Among other possible effects considered by Edelstein and Killian (2007) is the uncertainty effect: a shock to oil prices increases uncertainty about the future and increases the option value of delaying durable expenditures. Increased uncertainty may also lead to a precautionary saving effect, which induces consumers to increase their buffer stock by cutting back on nondurables.

[11]I use the LexisNexis Academic PowerSearch to search for:

> Search Terms: HLEAD(recession) AND SUBJECT(recession); Index Terms Added: * United States * Economy & Economic Indicators; Select Source: New York Times.

In order to avoid including articles that have a different context, such as a story about the rap album titled "The Recession," released in September 2008, I restrict the search to only include articles indexed by LexisNexis as articles about the economic recession.

[12]I have also conducted a search for the article count using the word "recovery." Not surprisingly, this search generated very few hits and had little forecasting power for consumer confidence and spending.

Finally, unemployment risk has been shown be a predictor of spending and consumer confidence.[13] Since there is no available daily measure of the unemployment rate, I construct a proxy for daily unemployment risk from the G1K job market question:[14]

> Now thinking more generally about the company or business you work for, including all of its employees—based on what you know or have seen, would you say that, in general, your company or employer is: 1) hiring new people and expanding the size of its workforce, 2) not changing the size of its workforce, or 3) letting people go and reducing the size of its workforce?

In subsection 2.4, I show that this daily question can be used as a proxy for aggregate risk of unemployment.

2.3 Variable construction and aggregation

Following the literature, I compute a "diffusion measure" of consumer confidence; see Ludvigson (2004) and Barsky and Sims (2012). This measure is defined as the daily percentage of respondents who reply "getting better," minus the daily percentage of respondents who reply "getting worse," plus 100 (making 100 the neutral position).

With regard to expenditure, following Attanasio and Weber (1993), I first transform the expenditure variable into a logarithm and then average the data using sampling weights.[15] The raw expenditure series hence consists of a daily series grouping respondents at each day t, the day when the expenditure took place. As the G1K expenditure question is answered by about 500 people per

---

[13]Carroll and Dunn (1997) show that unemployment risk is an important determinant of consumer spending. Also, Doms and Morin (2004) show that risk of unemployment, constructed from news searches for the terms "layoff" or "downsizing," is a predictor of consumer confidence.

[14]This question is only asked of individuals who are employed, who make up some 60 percent of the respondents.

[15]When working with aggregate data on consumption, the data are only available in levels, which are often transformed into logarithms. However, the recommendation of Attanasio and Weber (1993) for the analysis of averaged microdata on expenditures is to first compute the nonlinear transformation, then average the data. This approach gets complicated when using daily data, as, on average, 30 percent of the sample report that they spent zero dollars. In practice, as the time-series behavior of the logarithm of average spending and the average of the logarithm are similar (see Figure A.1 in the appendix), the two approaches to transformation yield very similar results.

day, such a cell size allows for a daily time series aggregated across consumers that would not be possible with the daily diary of the Consumer Expenditure Survey (CEX).

To reduce the influence of outliers, I trim the sample by dropping individuals who report expenditures above the 98th percentile—this corresponds to spending more than 600 dollars on a daily basis. Following this, the average level of spending in the analysis sample equals about 50 dollars. When I use the total expenditure numbers from Bee, Meyer, and Sullivan (2012) and divide them by the civilian noninstitutional population and the number of days in the year, I find that the average daily spending in 2010 was about 43 dollars in the CEX's Interview Survey but only 27 dollars in the CEX's Diary Survey, which suggests that the G1K is not dramatically off from these other measures.

The S&P 500, the VIX, and oil prices are transformed to natural logarithms, while the article counts from the *New York Times* are in levels. I define my measure of unemployment risk as the difference between the percentage of respondents saying their workplace is downsizing and the percentage of respondents saying their workplace is hiring, plus 100. Hence, it is similar to Carroll and Dunn's (1997) measure of unemployment expectations, defined as the difference between the fraction of respondents who think that unemployment will rise and the fraction who think it will fall.

In all of the analyses the dates reflect when the expense was made, rather than the date of the interview. As the stock market measures are not available on the weekends and major holidays, the analysis sample consists of 238 days of weekday data.

2.4   Descriptive statistics

Figure 1 plots the time series of consumer confidence and the log of spending. Both series are smoothed over lagged 14 days to reduce noise. In the estimations, I use the raw data and control for time effects directly. Throughout the period, consumers are quite gloomy about the economic conditions—the average of the diffusion measure is about 30. The confidence series is quite volatile, with dramatic swings occurring in the last quarter of 2008. The spending series is station-

ary but displays seasonality, evident in the increase in spending closer to Christmas and during the middle of the year.

Figure 2 shows the same plot of the consumer confidence variable, with article counts of "recession" from the *New York Times* superimposed as spikes. The takeaway from this figure is that confidence tends to decrease with an increase in the volume of reporting about the recession. Also, Doms and Morin (2004), Barsky and Sims (2012), and Starr (2012) document that confidence covaries with the volume of economic news reporting.

Figure 3 shows the G1K measure of risk of unemployment (i.e., "the reducing minus hiring gap") along with the number of new weekly UI claims (in thousands). The unemployment risk series is normalized, so that 100 means that the fraction of respondents who say their workplace is downsizing is equal to the fraction of respondents who say their workplace is hiring. The two series are strikingly similar. There is a close comovement of the two series and, indeed, the coefficient of correlation between unemployment risk and new UI claims is 0.84, suggesting that this measure is a very strong predictor of UI claims, which in turn is known to be a leading indicator of the monthly unemployment rate. Given this close relationship with new UI claims, and given that the G1K data is available at a finer frequency than the weekly UI reports, I proceed with using the G1K measure as my proxy for the risk of unemployment.

3   Estimation and Identification

In order to study the dynamic effect of a shock to consumer confidence on spending, I estimate a range of VAR models. I interpret the results by comparing how the impulse response function following a shock to consumer confidence differs once I control for the other economic indicators. If the response of spending to a confidence shock is the same regardless of whether a larger conditioning set is used, then confidence must contain information independent of the other variables.

All of the VARs are estimated in levels and include time variables, $\mathbf{x}_t$, consisting of day-of-week dummies, a weekly linear and square time trend meant to flexibly control for seasonality across the year, and an indicator for the last week of the year.

First, I examine a benchmark case by studying the dynamics of a simple two-variable VAR model:

$$\mathbf{y}_t = \mathbf{x}_t + \sum_{j=1}^{J} \mathbf{A}_j \mathbf{y}_{t-j} + \mathbf{u}_t \ ,$$

where $\mathbf{y}$ is a vector consisting of daily consumer confidence (confidence) and the log of daily spending (spending): $\mathbf{y} = (\text{confidence, spending})'$. Each $\mathbf{A}$ is a matrix of coefficients on the lagged terms of $\mathbf{y}$ (up to a lag J) and $\mathbf{u}_t$ is a vector of normally distributed disturbance terms.

Following Barsky and Sims (2012), I order consumer confidence first, followed by spending, and use a Cholesky decomposition to orthogonalize the residuals. This amounts to assuming that consumer confidence may affect spending *within a day*, but that spending does not affect confidence within a day. As the contemporaneous correlation between the reduced-form residuals from this VAR is very low, a reverse ordering does not change the results much.

Second, I build a richer model by including other variables discussed in subsection 2.2. In addition to spending and consumer confidence, there are five other variables: 1) log of VIX, 2) log of S&P 500, 3) log of oil prices, 4) *NYT* article counts, and 5) G1K's measure of risk of unemployment. In order to get a sense of whether and how the results from the bivariate VAR are sensitive to conditioning on more variables, I estimate five sets of three-variable VARs, each consisting of one of the five economic indicators, followed by consumer confidence, followed by spending. The motivation for this ordering follows the assumption that *within a day*, many economic variables change more slowly than do confidence and spending; a drop in daily confidence or spending will most likely not set off a bank collapse within a day, but news of a bank collapse might convey information relevant to confidence and spending within a day.

Third, I estimate a joint multivariate VAR consisting of all of the available variables in the following estimation order: $\mathbf{y} = (\text{log of S\&P 500, log of VIX, } \textit{NYT} \text{ article counts, log of oil prices,}$ unemployment risk, confidence, spending$)'$.

As the identification is using a Cholesky decomposition, the order of the variables preceding confidence and spending might matter. It is not obvious whether news shocks, stock market shocks, volatility shocks, or other shocks should precede any of the others within a day. The ordering above is motivated by Bloom (2009), who places stock market prices first, followed by volatility measures, oil prices, and unemployment indicators. Again, I make confidence contemporaneously orthogonal to spending. In this specification, *within a day*, consumer confidence is allowed to be affected by stock market prices, stock market volatility, *NYT* article counts, oil prices, and unemployment risk, but not by spending. Spending is affected by all the variables *within a day*, but is only allowed to affect the remaining variables with a day's lag.

As there are seven variables in total, it is not feasible to check all the other orders. However, to check the robustness of the results, I present four alternative estimation orders. First, I change the order of *NYT* article counts and log of VIX. In a second specification, I place the price variables—log of oil prices, followed by log of S&P 500—last. This is in keeping with the literature on "news" shocks, which has shown stock market innovations to predict future productivity and which hence suggests that they should be placed last. Third, I change the order of spending and confidence, so that spending comes before confidence. In this specification, log of S&P 500, log of VIX, *NYT* article counts, log of oil prices, unemployment risk, and spending are all controlled for when I look at the effect of a confidence shock on spending. In a fourth specification, I place confidence followed by spending first in the estimation order. This way, all of the other economic indicators are allowed to affect confidence and spending, only with one day's lag.

4   Results

I select the number of lags with the help of standard information criteria. In the postestimation phase, I perform the Lagrange multiplier test for serial correlation in the residuals of the VARs. If there is an indication of residual autocorrelation, which there typically is, I add more lags to the

model. All the responses are plotted following standard deviation shocks. I bootstrap the standard errors 1,000 times and plot one standard error confidence bands.

4.1 Results from a two-variable VAR

The left panel of Figure 4, labeled (a), traces the impulse response function of spending following a shock to consumer confidence. A standard deviation shock to confidence results in a statistically insignificant increase in daily expenditures in day 2, followed by a drop in day 3. Spending then increases back to trend for about 30 days. Panel (b) shows the impulse response function following a shock to spending. The results from this bivariate VAR point to a sensitivity of spending to confidence shocks. The economic impact of confidence on spending is not large—this is, however, in line with the findings of previous research. Confidence is not Granger-caused by spending to the same extent that spending is forecast by consumer confidence.

The correlation between the reduced-form residuals of this bivariate system is close to zero, suggesting that an alternative ordering of this small VAR, with spending ordered before confidence, does not alter the results much; Figure A.2 in the appendix confirms this.

The overall effect of a shock to confidence has a negative effect on spending. It is as if the consumers, following an initial pause, overreact by cutting back on expenses. This delayed negative reaction of spending is the most robust feature of the impulse response functions reported throughout the paper.

4.2 Results from three-variable VARs

In Figure 5, I show how the impulse response functions of spending following a shock to confidence change once I condition the VAR on either one of the five economic indicator variables: log of S&P 500, log of VIX, log of oil prices, *NYT* article counts, or unemployment risk. The impulse response functions in all of the subplots are strikingly similar to the response function found in

14

Figure 4 (a): shortly after a shock to confidence, there is a decrease that lasts for about a month.[16] Hence, innovations to confidence contain information relevant for spending that is not reflected in either one of the other economic variables.

4.3 Results from multivariate VARs

In order to study whether the innovations in confidence contain information relevant for spending, but separate from the *joint* information in all of the other available economic indicators combined, I estimate a joint multivariate VAR model including all seven variables. The estimation order is log of S&P 500, log of VIX, *NYT* article counts, log of oil prices, unemployment risk, confidence, and lastly, spending.

Figure 6 plots the impulse response function of spending following a shock to confidence from this VAR. Again, following a shock to confidence, spending takes on a very similar pattern to the reaction found in Figure 4 (a). The most salient feature of the dynamic effect is the negative reaction of spending and a slow recovery that lasts for about a month.

Analysis of forecast-error variance decompositions shows that spending is overwhelmingly explained by its own innovation; see Figures A.3–A.5 in the appendix. In comparison, a shock to consumer confidence explains up to 12 percent of the forecast-error variance in spending. The other innovations that help to explain the forecast-error variance of spending are shocks to the stock market level and the volatility index.

Confidence reacts almost entirely to its own shock, and, at the end of the 30-day horizon, it accounts for some 60 percent of the forecast-error variance. In comparison, stock market innovations explain at most 12 percent of the forecast error variance of confidence. Variance decomposition of confidence after a shock to spending shows that spending shocks do not explain any variation. Noticeably, shocks to the *New York Times* article count variable do not explain any of the forecast-

---

[16]Note that the five models use different lag lengths. This is due to how the information criteria select the number of lags. Increasing the number of lags does not alter the qualitative shape of the impulse responses.

error variance of spending or confidence. In sum, the variation in consumer confidence appears to be independent from the variation in the other available economic variables.

In order to study how robust the pattern of spending is to alternative estimation orders of the VAR, Figure 7 shows four alternative orthogonalizations. The various estimation orders are explained below the figure. Each figure shows that following a shock to confidence, the reaction is qualitatively very similar regardless of the ordering. This robustness to alternative orderings suggests that innovations in consumer confidence convey information relevant for spending separately from the other shocks. Analysis of variance decomposition from the alternative estimation orders confirms that, regardless of the estimation order, confidence reacts mostly to its own innovations and, to a lesser extent, to stock market shocks.[17]

5 Discussion

The empirical results suggest that consumer confidence has a modest but statistically independent impact on consumer spending; that is, the informational content of consumer confidence reflects something other than news reports, stock market prices or volatility, oil prices, or unemployment risk. In sum, the measure of consumer confidence in the G1K appears to have independent power to predict consumer behavior, and this result is robust to various estimation orders.

In all of the results, the shape of the impulse response functions of spending differs from the pattern found in Starr (2012) and Barsky and Sims (2012), where the reported response is slow-building and permanent. The empirical response of spending to a shock to confidence, reported in Figures 4–7, looks nothing like this—it is fairly immediate and follows an initial zigzag pattern. After a few days spending drops, and it takes about a month for it to recover back to the initial level. Overall, the observed drop and recovery of spending lasts for about a month.

What can explain such a short-run fluctuation? On one hand, given the short duration and size of the drop in spending, it seems unlikely that the effects reported in this paper are driven

---

[17]These results are available from the author.

by a wealth effect. If it were so, the time required to accumulate a necessary balance should occur over a reasonably long period of time; one would expect a gradual increase in spending that lasts longer than a month. On the other hand, Barsky and Sims (2012) hypothesize that a reversing pattern of consumption following a confidence shock is consistent with an "animal spirits" shock—consumers receive a positive shock and increase their spending, but then they realize the shock's lack of content and adjust their consumption path downward. This "animal spirits" shock would, however, imply an initial increase in spending, which does not occur. Instead, spending takes an initial pause, followed by a drop and rebound.

Whereas the model in Barsky and Sims treats all consumption as instantaneous, in order to interpret the results one may require a theoretical framework that allows for partial irreversibility. This is because within a day, almost all expenditures become "durable," and these daily purchases typically have poor resale value. Hence, the interpretation of the results requires partial irreversibility for the goods that are likely to be durable over a short period of time. Furthermore, the reported impulse response resembles the reaction of durable spending following a decrease in uncertainty (that is, a negative uncertainty shock).[18]

In response to an uncertainty shock, durable spending is predicted to follow a "drop-rebound-overshoot" dynamic due to two opposing effects: an uncertainty effect and a volatility effect. An increase in uncertainty makes consumers postpone purchases of durable items; as a result, spending falls in the short run (the uncertainty effect). However, since the shock is only temporary, as uncertainty falls, pent-up demand rises and purchases increase. The effect of reduced uncertainty is that the economic conditions become more volatile as consumers buy more volume to adjust their stock of durables back to optimum. In the medium run, there is a "volatility overshoot" in spending (the volatility effect).[19] At longer horizons, spending converges back to trend. Note

---

[18]Other research has also shown that durable purchases are volatile and sensitive to uncertainty. See, for example, Romer (1990), Caballero (1993), Hassler (2001), Knotek II and Khan (2011), and Alexopoulos and Cohen (2009).

[19]The importance of time-varying uncertainty for economic activity has recently gained much attention; see Bloom (2009). Although Bloom does not estimate the response of spending, the response of output and investment spending to an uncertainty shock is characterized by a short-run fluctuation.

that in order for the drop-rebound-overshoot fluctuation to occur, there must be an option value to delaying spending.

If a confidence shock reflects decreased uncertainty, then following this shock, spending increases as consumers feel more confident and adjust their stock of durables. But as the shock is temporary, once confidence fades, uncertainty rises and volatility increases. Because of the increase in ex-post volatility, demand falls and takes time to converge back to trend.

The most robust feature of the results of this paper is that following a shock to confidence, there is an initial period of a few days when spending meanders around the trend, then within a week of the shock it falls sharply and recovers sometime after a month. This behavior is consistent with consumers experiencing an unanticipated decrease in uncertainty with a dominating volatility effect: following the shock there is a pause in spending, followed by a decline due to more realized volatility of the economic conditions. The initial increase in spending predicted by the uncertainty effect does not manifest itself in the data.

Does this interpretation make sense? First, consider the "zeitgeist" of this time period. In both the second and third quarters of 2008, the respondents to the Survey of Professional Forecasters revised their estimates of real GDP growth downward. For example, in the third quarter of 2008 the estimate was revised from 1.8 percent down to 0.7 percent.[20] This suggests that the contractions occurring later in 2008 were largely unexpected only three months before. By the end of 2008, another indicator of uncertainty, the stock market volatility index, had quadrupled compared to its level at the beginning of the year. Hence, if following a shock to confidence the volatility of economic conditions turned out to be higher than what the consumers initially expected, spending would be predicted to drop.

Second, Browning and Crossley (2009) show that households smooth utility through two channels: 1) by postponing the purchases of small durables, such as clothing and small electronic equipment, and 2) by cutting back on eating out and entertainment. The available data prevent me from directly answering whether following a confidence shock there is a change in the structure of

[20]http://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/2008/spfq308.pdf

demand. Still, one might justifiably speculate that, following a confidence shock with a dominating volatility effect, consumers 1) delay purchases on items such as clothing, gas, and certain grocery items and 2) cut back on other forms of spending, such as eating out. Then, over the longer run, as volatility subsides, spending converges back to trend.

6 Conclusion

This paper uses high-frequency identification to determine whether consumer confidence contains information relevant to spending that is independent of other predictors of economic activity. The available data cover the turbulent first full year of the "Great Recession," which generates much day-to-day variation in stock market prices as well as in the volume of reporting on economic news. This variation is needed to identify the effect of shocks to consumer confidence on spending. The daily frequency of data makes it possible to orthogonalize shocks to consumer confidence with respect to other economic shocks.

Like the authors of other recent studies, I find that confidence carries relevant information that is orthogonal to other measures of economic conditions. A novel result of the paper is that, following a positive shock to confidence, spending first pauses, then decreases. This very short fluctuation in consumer spending following a shock to confidence is robust to various alternative estimation orders, but has not been observed when studying data collected at monthly or quarterly frequencies.

The temporary fluctuation in spending suggests that consumer confidence is a measure of uncertainty. For example, in Bloom's (2009) framework of time-varying uncertainty, a second-moment shock induces a drop in economic activity, followed by a rebound, and finally an overshoot. The observed response of spending following a shock to confidence suggests that consumer confidence might capture decreased uncertainty about the future. This finding supports the conclusions drawn by Mishkin (1978), who, though applying a different approach to identification, concludes that consumer confidence is a measure of uncertainty.

Results like these might be of interest to policymakers, because as Bachmann and Sims (2012) show, in a recession, consumer confidence matters for an effective fiscal policy. In Bloom's (2009) model the economy is responsive to economic stimuli if uncertainty is low. Thus, if consumer confidence is a measure of uncertainty, keeping confidence high and uncertainty low should be a goal for policymakers.

References

Acemoglu, D. and A. Scott (1994). Consumer Confidence and Rational Expectations: Are Agents' Beliefs Consistent with the Theory? *Economic Journal 104*(422), 1–19.

Alexopoulos, M. and J. Cohen (2009). Uncertain Times, Uncertain Measures. University of Toronto Working Paper No. 352.

Ashenfelter, O. and A. B. Krueger (1994). Estimates of the Economic Return to Schooling from a New Sample of Twins. *American Economic Review 84*(5), 1157–1173.

Attanasio, O. P. and G. Weber (1993). Consumption Growth, the Interest Rate and Aggregation. *Review of Economic Studies 60*(3), 631–649.

Bachmann, R. and E. R. Sims (2012). Confidence and the Transmission of Government Spending Shocks. *Journal of Monetary Economics 59*(3), 235–249.

Barsky, R. B. and E. R. Sims (2012). Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence. *American Economic Review 102*(4), 1343–1377.

Beaudry, P. and F. Portier (2006). Stock Prices, News, and Economic Fluctuations. *American Economic Review 96*(4), 1293–1307.

Bee, A., B. D. Meyer, and J. X. Sullivan (2012). The Validity of Consumption Data: Are the Consumer Expenditure Interview and Diary Surveys Informative? NBER Working Paper 18308.

Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica 77*(3), 623–685.

Browning, M. and T. F. Crossley (2009). Shocks, Stocks, and Socks: Smoothing Consumption over a Temporary Income Loss. *Journal of the European Economic Association 7*(6), 1169–1192.

Browning, M., T. F. Crossley, and G. Weber (2003). Asking Consumption Questions in General Purpose Surveys. *Economic Journal 113*(491), 540–567.

Caballero, R. J. (1993). Durable Goods: An Explanation for Their Slow Adjustment. *Journal of Political Economy 101*(2), 351–384.

Carroll, C. D. and W. E. Dunn (1997). Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets. In B. S. Bernanke and J. Rotemberg (Eds.), *NBER Macroeconomic Annual*, Volume 12, pp. 165–230. MIT Press.

Carroll, C. D., J. C. Fuhrer, and D. W. Wilcox (1994). Does Consumer Sentiment Forecast Household Spending? If So, Why? *American Economic Review 84*(5), 1397–1408.

Chauvet, M. and J.-T. Guo (2003). Sunspots, Animal Spirits, and Economic Fluctuations. *Macroeconomic Dynamics 7*(1), 140–169.

Deaton, A. (2012). The Financial Crisis and the Well-Being of Americans. *Oxford Economic Papers 64*(1), 1–26.

Deaton, A. and R. Arora (2009). Life at the Top: The Benefits of Height. *Economics and Human Biology 7*(2), 133–136.

Doms, M. and N. Morin (2004). Consumer Sentiment, the Economy, and the News Media. Finance and Economics Discussion Series Working Paper 2004-51.

Edelstein, P. and L. Killian (2007). Retail Energy Prices and Consumer Expenditures. Unpublished.

Hamilton, J. D. (2009). Causes and Consequences of the Oil Shock of 2007–08. *Brookings Papers on Economic Activity 40*(1), 215–283.

Hassler, J. (2001). Uncertainty and the Timing of Automobile Purchases. *Scandinavian Journal of Economics 103*(2), 351–366.

Knotek II, E. S. and S. Khan (2011). How Do Households Respond to Uncertainty Shocks? *Federal Reserve Bank of Kansas City Economic Review* (Second Quarter), 63–92.

Krueger, A. B. and I. Kuziemko (2011). The Demand for Health Insurance among Uninsured Americans: Results of a Survey Experiment and Implications for Policy. NBER Working Paper 16978.

Ludvigson, S. C. (2004). Consumer Confidence and Consumer Spending. *Journal of Economic Perspectives 18*(2), 29–50.

Matsusaka, J. G. and A. M. Sbordone (1995). Consumer Confidence and Economic Fluctuations. *Economic Inquiry 33*(2), 296–318.

Mishkin, F. S. (1978). Consumer Sentiment and Consumer Durable Expenditure. *Brookings Papers on Economic Activity 1*, 217–231.

Petev, I., L. Pistaferri, and I. Saporta-Eksten (2012). Consumption and the Great Recession: An Analysis of Trends, Perceptions, and Distributional Effects. In D. Grusky, B. Western, and C. Wimer (Eds.), *Analyses of the Great Recession*. Russel Sage Foundation.

Poterba, J. M. and A. A. Samwick (1995). Stock Ownership Patterns, Stock Market Fluctuations, and Consumption. *Brookings Papers on Economic Activity 26*(2), 295–372.

Ramey, V. A. (2011). Identifying Government Spending Shocks: It's All in the Timing. *Quarterly Journal of Economics 126*, 1–50.

Ramey, V. A. and M. D. Shapiro (1998). Costly Capital Reallocation and the Effects of Government Spending. *Carnegie-Rochester Conference Series on Public Policy 48*(1), 145–194.

Romer, C. D. (1990). The Great Crash and the Onset of the Great Depression. *Quarterly Journal of Economics 105*(3), 597–624.

Romer, C. D. and D. H. Romer (1990). Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz. NBER Working Paper 2966.

Souleles, N. S. (2004). Expectations, Heterogenous Forecast Errors, and Consumption: Micro Evidence from the Michigan Sentiment Surveys. *Journal of Money, Credit, and Banking 36*(1), 39–72.

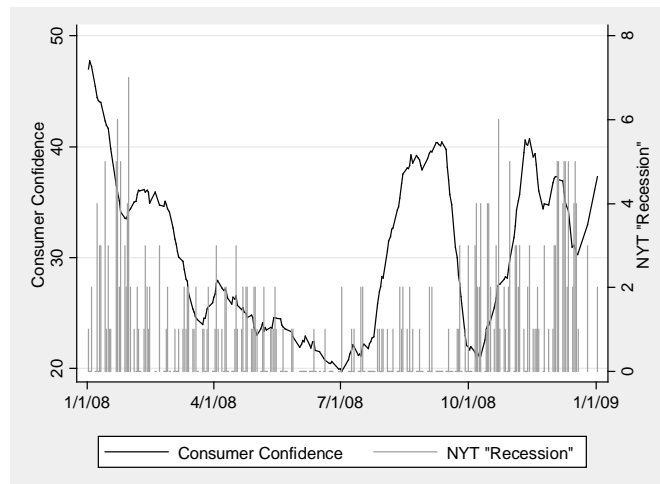Starr, M. (2012). Consumption, Sentiment, and Economic News. *Economic Inquiry 50*(4), 1097–1111.

Results

Figure 1: Consumer confidence and the log of expenditure.
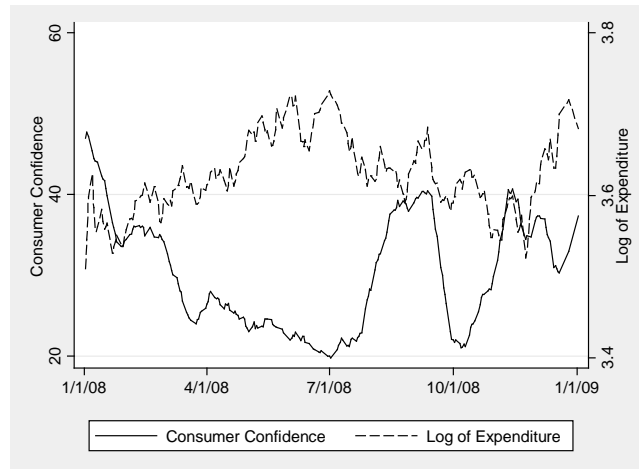
Notes: Both series are smoothed using a lagged 14-day moving average.

Figure 2: Consumer confidence and daily article counts for "recession" in the *New York Times.*

Notes: Consumer confidence is smoothed using a lagged 14-day moving average.

Figure 3: Unemployment risk (reducing/hiring gap) and initial weekly unemployment insurance (UI) claims.
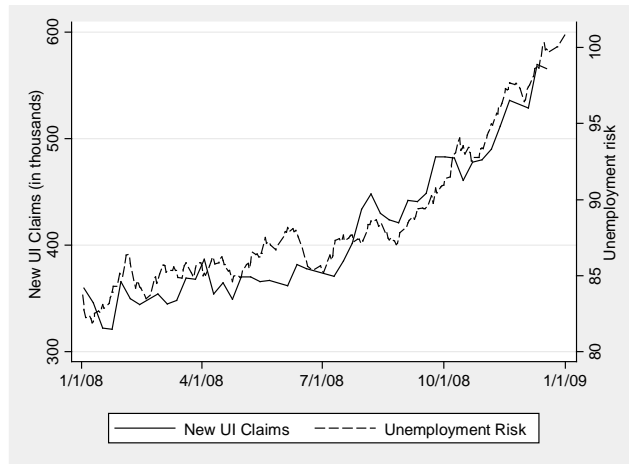
Notes: Unemployment risk is calculated as the daily gap between the percentage of respondents who say their workplace is reducing its workforce minus the percentage of respondents who say their workplace is hiring new people, plus 100. The plot shows the series smoothed using a lagged 14-day moving average. The solid line shows the initial seasonally adjusted weekly UI claims.

Figure 4: Impulse responses, two-variable VAR.
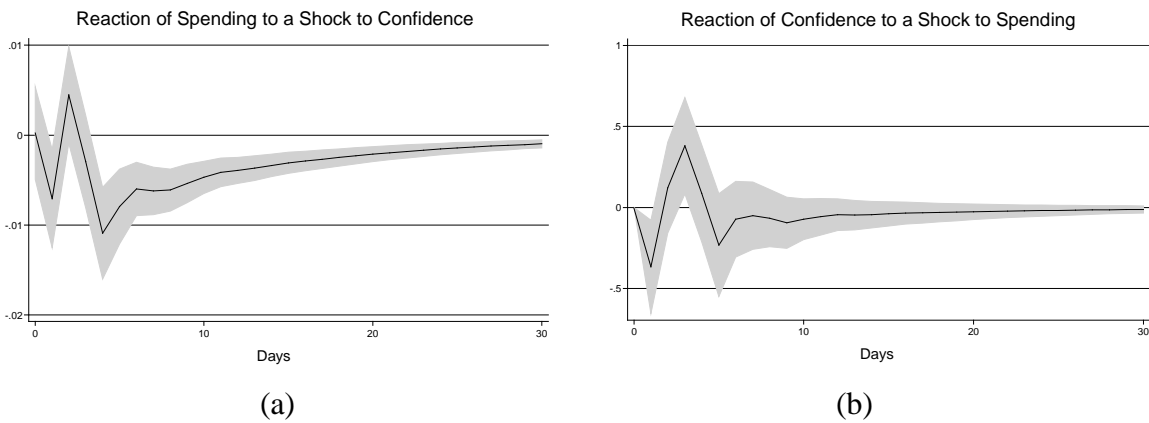
Notes: VAR with two variables: consumer confidence and log of expenditure. Other exogenous variables included: weekly linear and square trend, day-of-week dummies, and an indicator for the last week of the year. One-standard-error confidence intervals are computed using bootstrapping method with 1,000 draws. Lags used: 5.

Estimation order: confidence → spending.

Granger causality tests: *p*-values
H_0: Spending does not Granger-cause confidence: 0.214
H_0: Confidence does not Granger-cause spending: 0.029

25

Figure 5: Impulse response of spending following a shock to confidence—three-variable VARs.
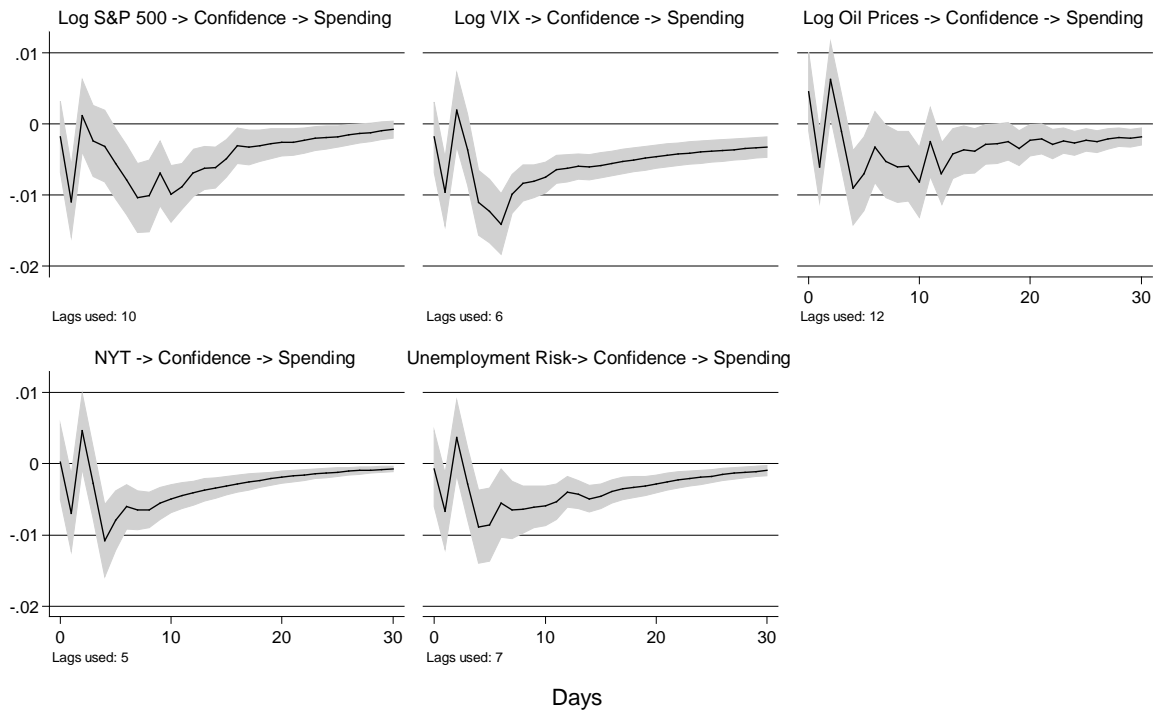
Notes: Each VAR is fitted separately and contains three variables: "economic indicator," consumer confidence, and log of expenditure. "Economic indicator" is either: the log of S&P 500, log of VIX, log of oil prices, *New York Times* article counts for "recession," or a proxy for unemployment risk (the reducing/hiring gap). Other exogenous variables included: weekly linear and square trend, day-of-week dummies, and an indicator for the last week of the year. One-standard-error confidence intervals are computed using bootstrapping method with 1,000 draws. Lags used: varies; see the plots for details.

Estimation order: "economic indicator" → confidence → spending.

Figure 6: Impulse response of spending following a shock to confidence—multivariate VAR.
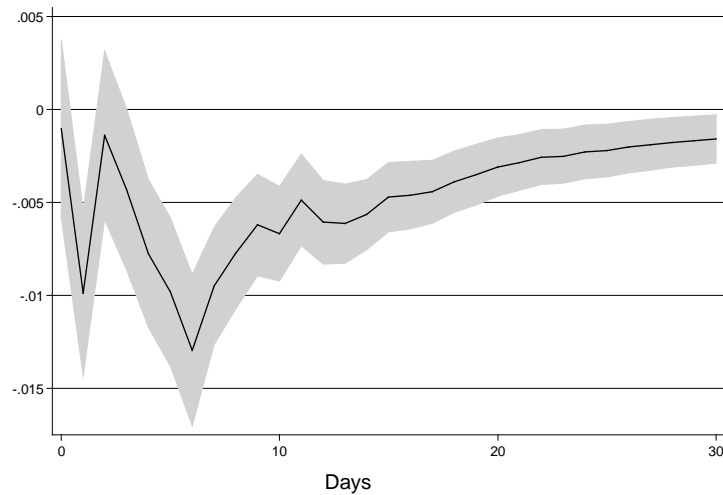
Notes: VAR contains seven variables: log of S&P 500, log of VIX, log of oil prices, *New York Times* article counts for "recession," a proxy for unemployment risk (the reducing/hiring gap), consumer confidence, and log of expenditure. Other exogenous variables included: weekly linear and square trend, day-of-week dummies, and an indicator for the last week of the year. One-standard-error confidence intervals are computed using bootstrapping method with 1,000 draws. Lags used: 6.

Estimation order: log of S&P 500 → log of VIX → NYT "Recession" → log oil prices→ unemployment risk → confidence → spending.

Figure 7: Sensitivity of the multivariate VAR to alternative estimation orders—impulse response of spending following a shock to confidence.
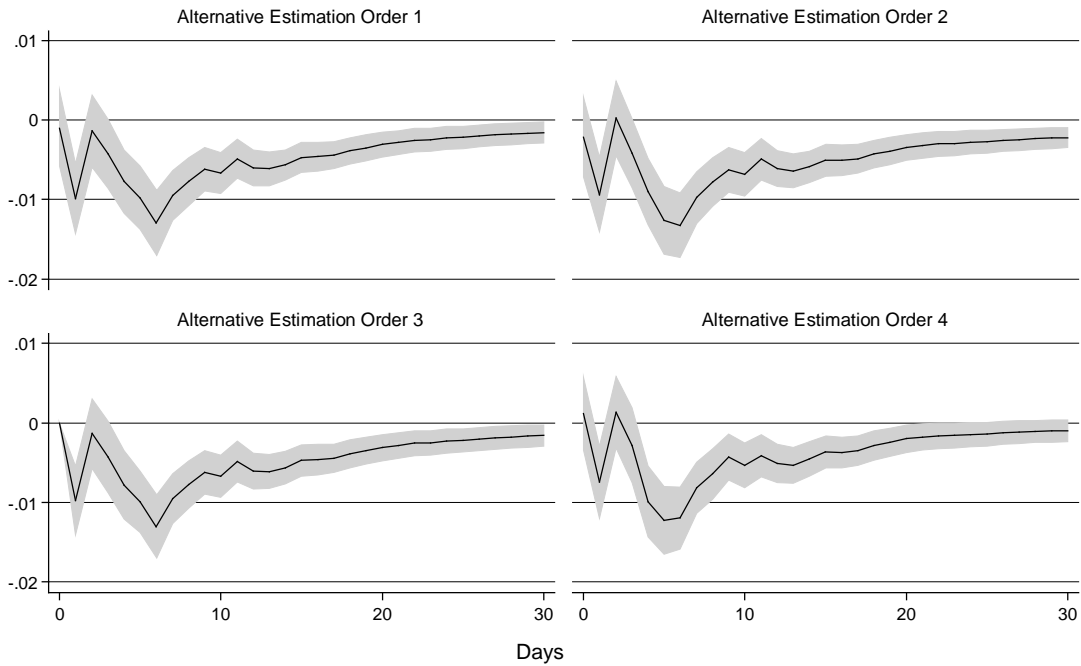
Notes: VAR contains seven variables: log of S&P 500, log of VIX, log of oil prices, *New York Times* article counts for "recession," a proxy for unemployment risk (the reducing/hiring gap), consumer confidence, and log of expenditure. Other exogenous variables included: weekly linear and square trend, day-of-week dummies, and an indicator for the last week of the year. One-standard-error confidence intervals are computed using bootstrapping method with 1,000 draws. Lags used: 6.

Alternative estimation order 1: log of S&P 500 → NYT "Recession" → log of VIX → log oil prices→ unemployment risk → confidence → spending.

Alternative estimation order 2: log of VIX → NYT "Recession" → unemployment risk → confidence → spending → log oil prices → log of S&P 500.

Alternative estimation order 3: log of S&P 500 → log of VIX → NYT "Recession" → log oil prices→ unemployment risk → spending → confidence.

Alternative estimation order 4: confidence → spending → log of S&P 500 → log of VIX → NYT "Recession" → log oil prices → unemployment risk.

Appendix A

Figure A.1: Average of log of spending and the log of average spending.
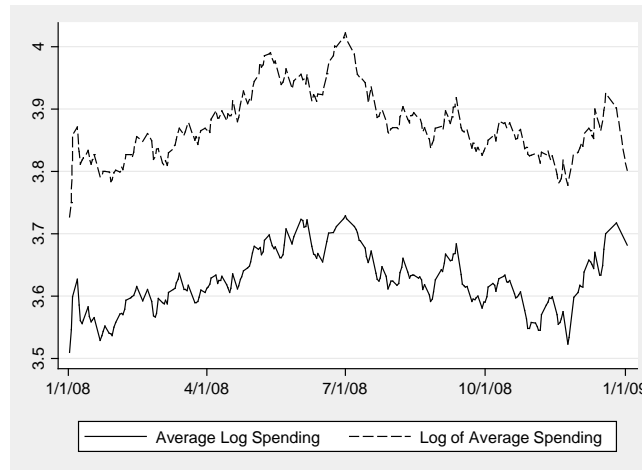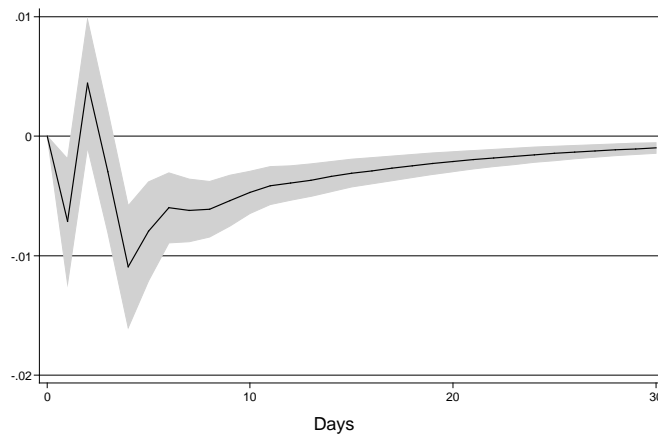

Figure A.2: Sensitivity of the two-variable VAR to estimation order—impulse response of spending following a shock to confidence.

Notes: VAR with two variables: consumer confidence and log of expenditure. Other exogenous variables included: weekly linear and square trend, day-of-week dummies, and an indicator for the last week of the year. One-standard-error confidence intervals are computed using bootstrapping method with 1,000 draws. Lags used: 5.

Estimation order: spending $\rightarrow$ confidence.

Figure A.3: Forecast-error variance decomposition (FEVD) of confidence and spending following a shock to itself.
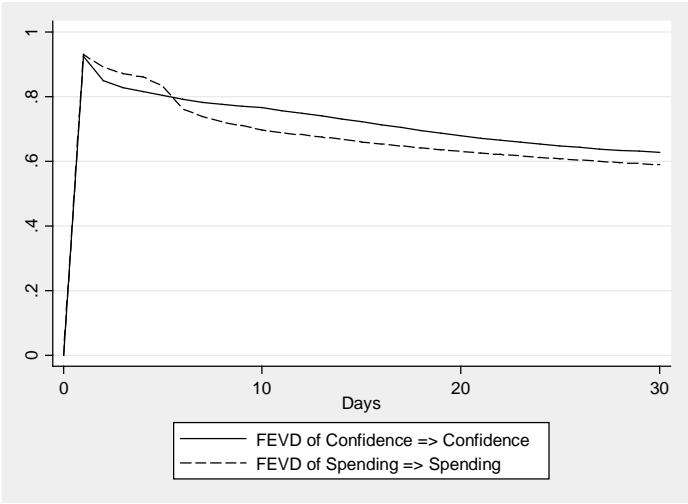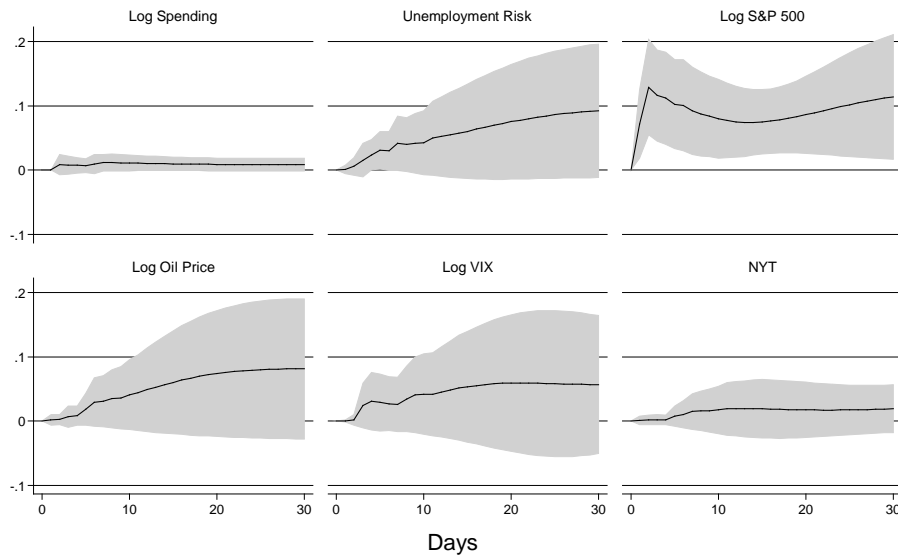

Notes: FEVD from the model in Figure 6.

Figure A.4: Fraction of forecast-error variance of <u>spending</u> that can be attributed to a shock to…
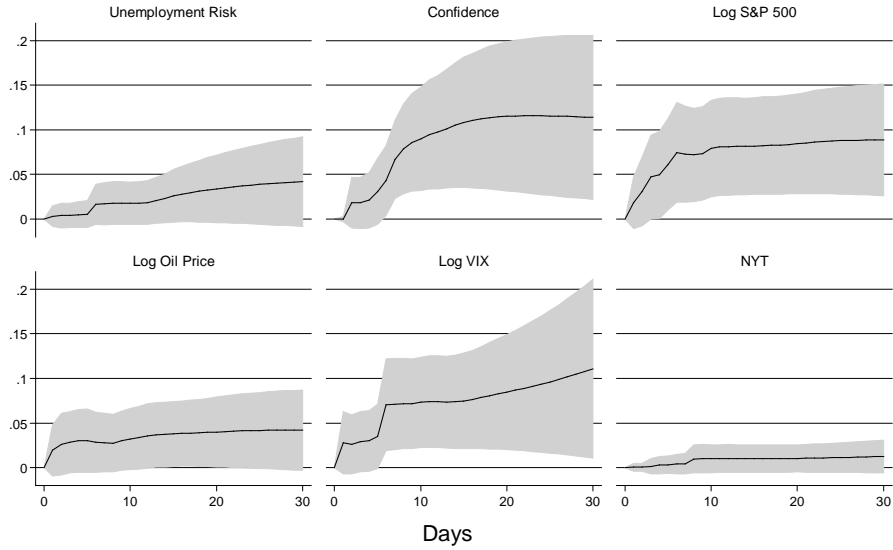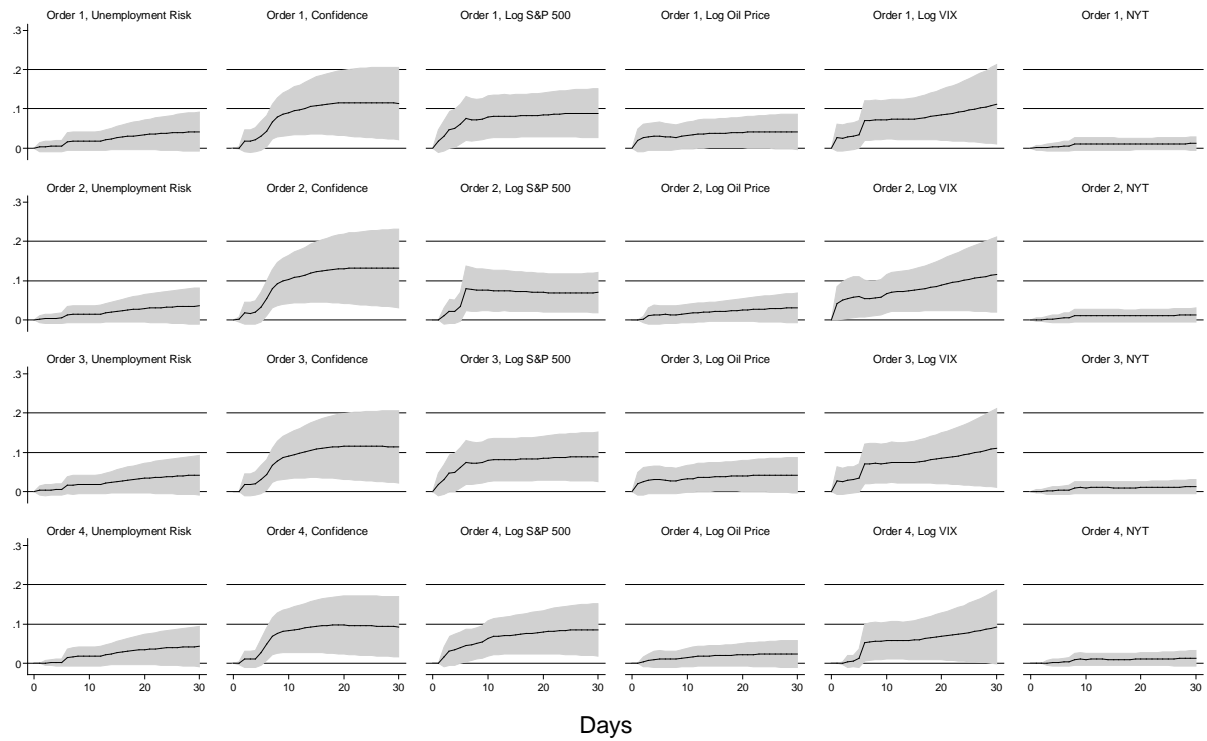

Notes: Forecast-error variance with 90 percent error confidence intervals based on the model in Figure 6.

Figure A.5: Fraction of forecast-error variance of <u>confidence</u> that can be attributed to a shock to…
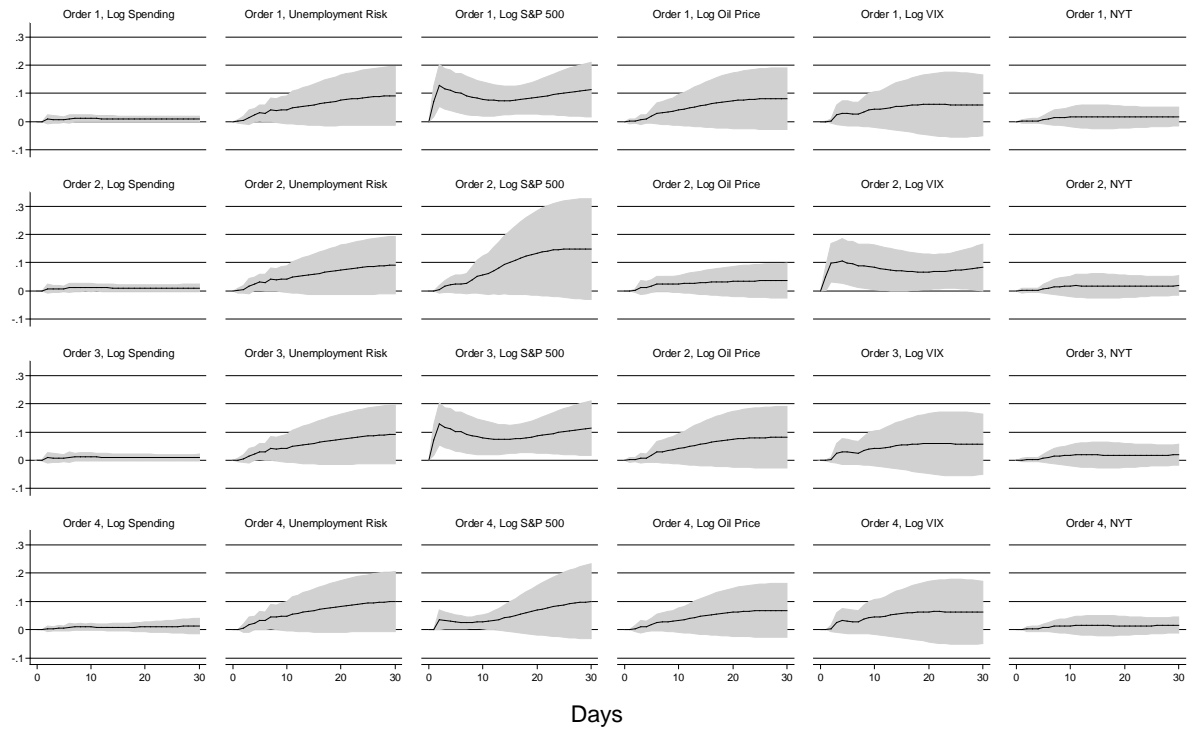

Notes: Forecast-error variance with 90 percent error confidence intervals based on the model in Figure 6.

Figure A.6: Fraction of forecast-error variance of <u>spending</u> that can be attributed to a shock to…

Days

Notes: Graphs by estimation order in Figure 7 and impulse variable. 90 percent confidence intervals based on the models in Figure 7.

Figure A.7: Fraction of forecast-error variance of <u>confidence</u> that can be attributed to a shock to…

Notes: Graphs by estimation order in Figure 7 and impulse variable. 90 percent confidence intervals based on the models in Figure 7.