Oxford Economic Papers, 68(4), 2016, 920–944
doi: 10.1093/oep/gpw009
Advance Access Publication Date: 31 March 2016

Expenditure and confidence: using daily data to identify shocks to consumer confidence

By Marta Lachowska

W. E. Upjohn Institute for Employment Research, 300 S. Westnedge Avenue, Kalamazoo, MI 49007, USA; e-mail: marta@upjohn.org.

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? This article uses novel daily data to understand what can be learned about the dynamics of consumer confidence and spending. In contrast to the existing literature that uses data collected at lower frequencies, I find that the estimated relationship between daily consumer confidence and daily spending is weak. I interpret this finding as an indication that on a day-to-day basis, consumers are rationally inattentive and do not react to small and temporary fluctuations in consumer confidence.

JEL classifications: E21, E32, C32

1. Introduction

Consumer confidence indices are regularly reported by the media as indicators of economic prospects. However, 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 indices 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 *et al.* (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 non-fundamentals. In sum, why consumer confidence predicts spending remains in dispute.[2]

This article uses daily data to understand what can be learned about the dynamics of consumer confidence and spending. As Starr (2012) and Barsky and Sims (2012) have done, I study the dynamics 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. The main data cover the year 2008, a period of particular interest because it is the first full year of a deep US recession. Using the 2008 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 indices 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 gas and oil prices and a measure of unemployment risk.

In the empirical part of the article, 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.

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

1 For instance, Carroll *et al.* (1994) find that confidence helps 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 negatively correlated.

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

and Cohen (2009), Knotek 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.[3]

The use of daily data has both strengths and weaknesses. On one hand, daily data may help uncover interesting dynamics between spending and consumer confidence that data collected at a monthly or a quarterly frequency average over. On the other hand, daily data may be too fine to expect a reaction of spending to occur. This may be the case if consumers face costs of processing economic information and therefore do not continuously adjust their behavior to small shocks; see, for example, Sims (2003) and Reis (2006). Because it is *ex ante* unclear if and how confidence affects daily spending, it is interesting to study whether the day-to-day relationship between consumer confidence and spending differs from the relationship estimated at more conventional data-collection frequencies such as a month or a quarter.

The article has the following main findings. First, the estimated impulse response of daily spending following an innovation in confidence is different from that observed when using monthly or quarterly data. At a monthly or quarterly frequency, spending has been shown to increase following a confidence shock. Using daily data, following a confidence shock, the reaction of spending is statistically and economically small. I show that this result holds when I use daily data for 2008 as well as for 2011, the other year of G1K data available to me. I interpret my results in the context of Reis's (2006) model of consumer inattention.

The rest of the article is organized as follows. Section 2 describes the data. In Section 3, I discuss the identification strategy. 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 benchmark two-variable VAR consisting of consumer spending and confidence. Next, I study a system where, in addition to consumer spending and confidence, other economic indicators are introduced—first alternatively, then jointly. Section 4 describes the results and conducts robustness checks using data from a post-recession year, 2011. Section 5 interprets the findings and tries to answer the question 'What is consumer confidence?' The final section concludes.

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 or older living in the USA. 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 had his or her birthday most recently. To make the sample representative, Gallup provides survey sampling weights to correspond to the national

3 Since the article 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).

distribution of age, gender, race, region, and educational level.[4] The main data used in this article collected information on about 359,000 individuals surveyed from 2 January 2008 to 5 January 2009.[5] As a robustness check, later I also use another wave of G1K data, which was collected between 2 January 2011 and 29 December 2011.

2.1.1 The expenditure question

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, online, or elsewhere’?

The G1K question is a total expenditure question, which measures the dollar amount spent on goods and services whilst excluding some of the biggest durables, such as the purchase of a home and car. 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% of household expenses go to buying small durables, such as home entertainment equipment, cosmetics, and clothes. I discuss the extent to which the G1K spending measure matches aggregate consumption data from the US Bureau of Economic Analysis (BEA) and measurement error in the Online Appendix, section A.1.

2.1.2 The forward-looking consumer confidence question

The forward-looking consumer confidence measure is collected from the same subsample as the spending measure. The respondent is asked to evaluate the 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 forward-looking 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).

Note that this forward-looking question is not the same as the Gallup Economic Confidence Index (ECI).[6] My reason for focusing on the forward-looking confidence

4 Krueger and Kuziemko (2013) 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 (2011) uses the same data to track how the financial crisis has affected subjective well-being in the USA.

5 The survey was not conducted on the following days: 21 January, 21 February, 23 March, 26 May, 30 May, 18 June, 27 June, 4 July, 27 November, 23–25 December, 29 December, 31 December, and 1 January.

6 The Gallup ECI is based on the combined response to two questions, one of them being the forward-looking consumer confidence question. According to Gallup: ‘the Gallup Economic Confidence Index is based on the combined responses to two questions asking Americans, first, to rate economic conditions in the country today, and second, whether they think economic conditions in the country as a whole are getting better or getting worse. The Gallup ECI is computed by

question is that I want to follow other papers that also have focused on forward-looking questions. For example, Barsky and Sims (2012) write 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. Dominitz and Manski (2004) express a similar concern regarding the use of composite indices. I discuss the differences between the forward-looking consumer confidence question, the Gallup ECI, the University of Michigan, and Conference Board indices in more detail in the Online Appendix. From here on, I refer to the forward-looking consumer confidence question as 'consumer confidence' or, simply 'confidence'.

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 of 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 (1995, p. 356). Therefore, the information stemming from stock market indices could be relevant to individuals' spending decisions and consumer confidence, even if the individuals interviewed in the G1K do not own stock themselves.[7] Also 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 include the index of stock market volatility, the VIX. The VIX is a measure of the implied 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 US 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; following an increase in energy prices, once consumers have paid their energy bills, there is less income with which to purchase goods and services.[8] As gasoline consumption was highly volatile during 2008 (see Petev *et al.*, 2012) with dramatic variations in the retail price of gas (see Hamilton, 2009), it seems appropriate to control for

adding the percentage of Americans rating current economic conditions (("excellent" + "good") minus "poor") to the percentage saying the economy is ("getting better" minus "getting worse"), and then dividing that sum by 2'; http://www.gallup.com/poll/123323/understanding-gallup-economic-measures.aspx (accessed 19 January 2016). As I show in the Online Appendix, Section A.2, my results remain qualitatively unchanged if I use the Gallup ECI as a measure of confidence.

7 Romer (1990) draws similar conclusions in her study of the effects of the Great Crash on spending.

8 Amongst 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 non-durables.

gas prices in the estimation. Because on a day-to-day basis oil and gasoline prices may have different impact on spending, I use daily data from the US Oil Fund (USO), which tracks the movements of West Texas Intermediate (WTI) crude oil, and daily data on gasoline prices from the Federal Reserve Economic Data base (FRED).

Fourth, following the approach in Doms and Morin (2004) and Alexopoulos and Cohen (2009), I collect information on how often and how much 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.[9] As Alexopoulos and Cohen argue, the *New York Times* is considered the national newspaper of the United States and a leading news source for the public and 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.[10]

Finally, unemployment risk has been shown to be a predictor of spending and consumer confidence.[11] Because there is no available daily measure of the unemployment rate, I construct a proxy for daily unemployment risk from the G1K job market question:[12] ‘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’?

Below, I argue that this question can be used as a proxy for 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

9 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. 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.

10 Also, *The Economist* tracks the number of newspaper articles mentioning the word *recession* as a predictor of whether the economy is facing a recession: http://www.economist.com/blogs/dailychart/2011/09/r-word-index. I also conducted a search for the word *recovery.* Not surprisingly, this search generated very few hits and had little forecasting power for consumer confidence and spending.

11 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.

12 This question is only asked of individuals who are employed, who make up some 60% of the respondents.

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 regards to expenditure, following Attanasio and Weber (1993), I first transform the expenditure variable into a logarithm and then average the data using sampling weights.[13] The raw expenditure series hence consists of a daily series grouping respondents at each day *t*, the day the expenditure took place. As the G1K expenditure question is answered by about 500 people per 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 on a daily basis. Trimming does not affect the main results but the time series of expenditure reported in this article is different from the one presented on Gallup’s website; see section A.1.3 in the Online Appendix. Following the trimming, the average level of spending in the analysis sample equals about $50.[14]

The S&P 500, the VIX, and gas and oil prices are transformed to natural logs, whilst the article counts from the *New York Times* are in levels.

I define unemployment risk as the difference between the percentage of respondents who say their workplace is downsizing and the percentage of respondents who say 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 of the 2008 sample

Figure 1 plots the time series of the log of expenditure (in panel a) and consumer confidence (in panel b). 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 stationary but displays seasonality, evident in the increase in spending closer to Christmas and during the middle of the year.

Panel (a) in Fig. 2 shows consumer confidence and the daily article counts of the word recession appearing in the headline or lead section of the *New York Times*. The takeaway

13 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 non-linear transformation, then average the data. This approach gets complicated when using daily data, as on average 30% of the sample report that they spent $0. In practice, as the time-series behavior of the log of average spending and the average of the log are similar (see Fig. A.10 in the Online Appendix), the two approaches to transformation yield very similar results.

14 When I use the total expenditure numbers from Bee *et al.* (2012) and divide them by the civilian non-institutional population and the number of days in the year, I find that the average daily spending in 2010 was about $43 in the CEX’s Interview Survey but only $27 in the CEX’s Diary Survey, which suggests that the G1K is not dramatically off from these other measures.

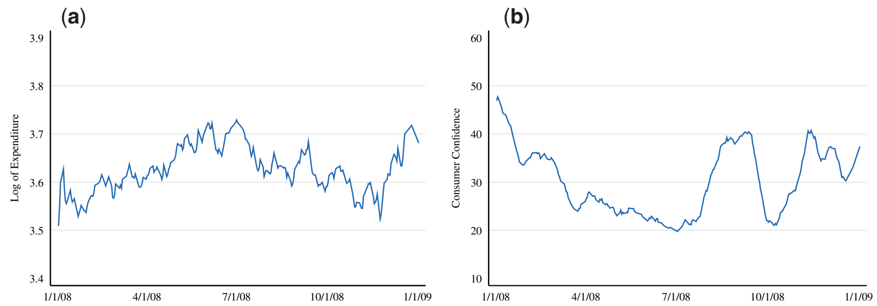

Fig. 1. The log of expenditure and consumer confidence
Source: 2008 G1K.
Notes: Consumer confidence is based on Gallup's 'economic outlook' question. It is calculated as the percentage saying the economy is 'getting better' minus the percentage saying the economy 'getting worse' plus 100. Log of expenditure is calculated as the average of the log of daily spending. Both time series are averaged using survey sampling weights. In the figure, both series are smoothed using a lagged 14-day moving average.

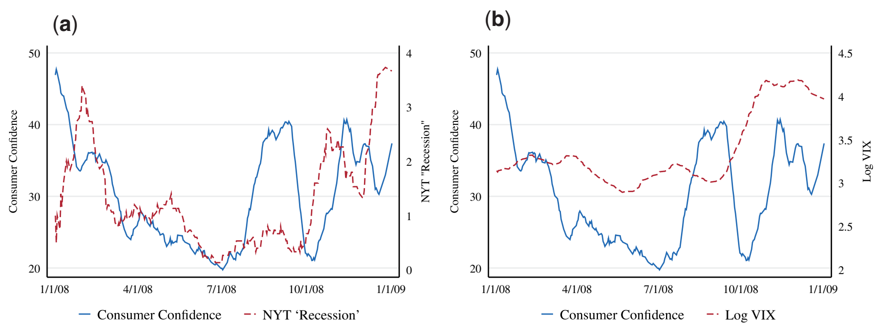

Fig. 2. Consumer confidence, the daily article counts for *recession* in the *New York Times*, and the log of the VIX
Source: 2008 G1K.
Notes: The series are smoothed using a lagged 14-day moving average.

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 negatively with the volume of economic news reporting. Panel (b) plots consumer confidence and the log of the VIX index. Confidence covaries negatively with high levels of the VIX in the later part of 2008.

Figure 3 shows the G1K measure of unemployment risk (i.e., the reducing/hiring gap) along with the number of new weekly unemployment insurance (UI) claims. 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. This suggests that the G1K measure is a strong predictor of new 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

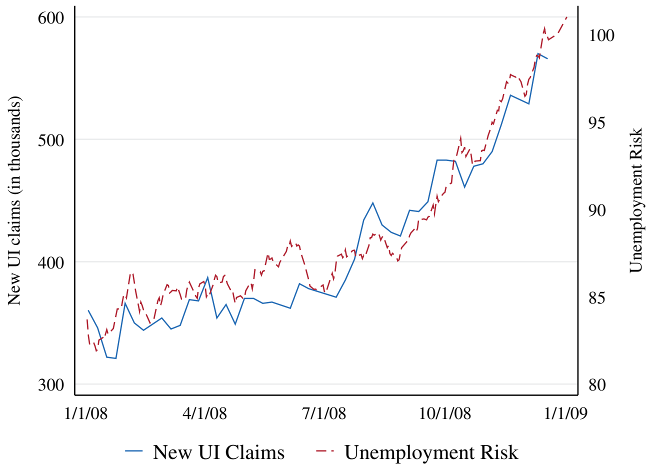

Fig. 3. Unemployment risk (reducing/hiring gap) and initial weekly unemployment insurance (UI) claims

Source: 2008 G1K and US Department of Labor.

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.

Table 1. Correlation matrix of economic indicators

Economic indicator	Log S&P 500	Log VIX	NYT	Log oil prices	Log gas prices	Unemployment risk
Log S&P 500	1.00					
Log VIX	−0.94	1.00				
NYT	−0.35	0.41	1.00			
Log oil prices	0.84	−0.83	−0.51	1.00		
Log gas prices	0.86	−0.85	−0.51	0.97	1.00	
Unemployment risk	0.81	0.72	0.30	−0.68	−0.68	1.00

Source: 2008 G1K and other sources; see text for details.

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.

G1K data are available at a finer frequency than the weekly UI reports, I use the G1K measure as my proxy for unemployment risk.

Table 1 presents the correlation coefficients between the six daily economic indicators discussed in section 2.2. The stock market volatility index and the stock market price index correlate very closely, and so do oil and gas prices. As expected, the *New York Times* article count for the word *recession* correlates negatively with the stock market price index, but positively with the volatility index and unemployment risk. Unemployment risk, on the other hand, correlates positively with both the stock market level and stock market volatility.

3. Estimation and identification

To study the dynamic effect of a shock to consumer confidence on spending, I estimate a range of VAR models using the 2008 data. 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, x_t, consisting of day-of-week dummies, a weekly linear and quadratic time trend meant to flexibly control for seasonality across the year, an indicator for the last week of the year, and, to account for changes in monetary policy, the dates of the Federal Open Market Committee meetings.

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

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

where $\mathbf{y}$ is a vector consisting of daily consumer confidence (confidence) and the log of daily spending (spending): $\mathbf{y}$ = (confidence, spending)'. Each A is a matrix of coefficients on the lagged terms of y (up to a lag J) and 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 order 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 section 2.2. In addition to spending and consumer confidence, there are six other variables: log of VIX, log of S&P 500, log of oil prices, log of gas prices, *NYT* article counts, and G1K's measure of unemployment risk. To get a sense of whether and how the results from the bivariate VAR are sensitive to conditioning on more variables, I estimate six sets of three-variable VARs, each consisting of one of the six 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: $\mathbf{y}$ = (log of S&P 500, log of VIX, *NYT* article counts, log of oil prices, log of gas prices, unemployment risk, confidence, spending)'.

Because 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 is motivated by Bloom (2009, p.630), 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, gas 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.

Because there are eight 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 and log of gas 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, log of gas 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.

Finally, I check external validity by estimating multivariate VARs using the five different estimation orders described above using the 2011 G1K data.

4. Results

I select the number of lags with the help of standard information criteria. In the post-estimation 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 95% confidence bands.

4.1 Results from a two-variable VAR

Figure 4 shows results from a bivariate VAR, which is estimated using five lags.[15] The left panel of Fig. 4, labeled (a), traces the impulse response function of spending following a shock to consumer confidence. A positive 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 of consumer confidence following a shock to spending. The Granger-causality p-values reported below the figure suggest that consumer confidence is not Granger-caused by spending, but spending is Granger-caused by consumer confidence at a 5% significance level. Panel (c) traces the response of confidence following a shock to itself, and panel (d) traces the response of spending following a shock to itself.

A positive 1 standard deviation shock to consumer confidence increases consumer confidence by about 5 points or, equivalently, by 17%.[16] The response of confidence to itself is temporary and lasts less than 30 days. This short-lived response is different from the reaction observed when using quarterly data. For example, in Barsky and Sims (2012), following a confidence shock, confidence increases and remains above trend for about 10 quarters.

15 The AIC picks four lags, whilst BIC picks two. At four lags the Lagrange multiplier (LM) test suggests residual autocorrelation. At five lags, the LM test-statistic's p-value is 0.12.

16 That is, a 5-point increase relative to the average level of confidence, which equals 30.

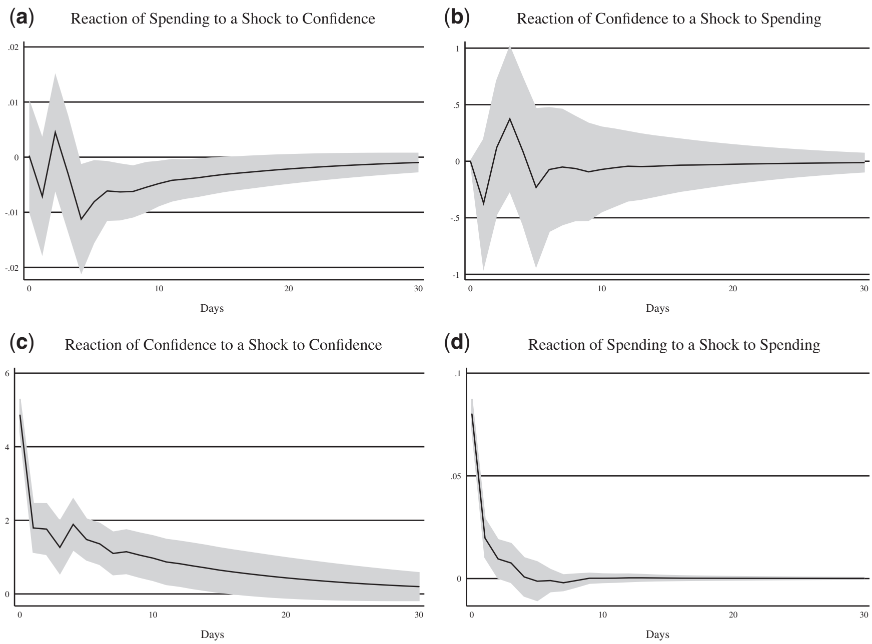

Fig. 4. Impulse responses, two-variable VAR

Source: 2008 G1K.

Notes: Impulse responses show the effect of a positive standard deviation shock. VAR with two variables: consumer confidence and log of expenditure. Other exogenous variables included weekly linear and quadratic trend, day-of-week dummies, Federal Open Market Committee meeting dates, and an indicator for the last week of the year. 95% 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: 0.216

H_0: confidence does not Granger-cause spending: 0.026

The results from this bivariate VAR point to some sensitivity of spending to confidence shocks. The economic impact of confidence on spending is not large—a 17% increase in consumer confidence increases spending by about 0.5% (in day 2), implying an elasticity of about 0.03. Barsky and Sims (2012) find that a positive 1 standard deviation shock to confidence predicts about 0.1–0.2% increase in consumption one quarter after the shock and about 0.5% increase in consumption several quarters after the shock.

What is strikingly different between the results in Fig. 4 and the existing literature is the shape of the impulse response of spending following a positive confidence shock and its lack of precision. Whereas in Barsky and Sims (2012) the response is positive, permanent, and statistically significant at a 5% significance level, at a daily frequency, the response is short-lived and mostly imprecisely estimated. Noticeably, the overall effect of a positive shock to confidence has a slight negative effect on spending.[17]

17 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.

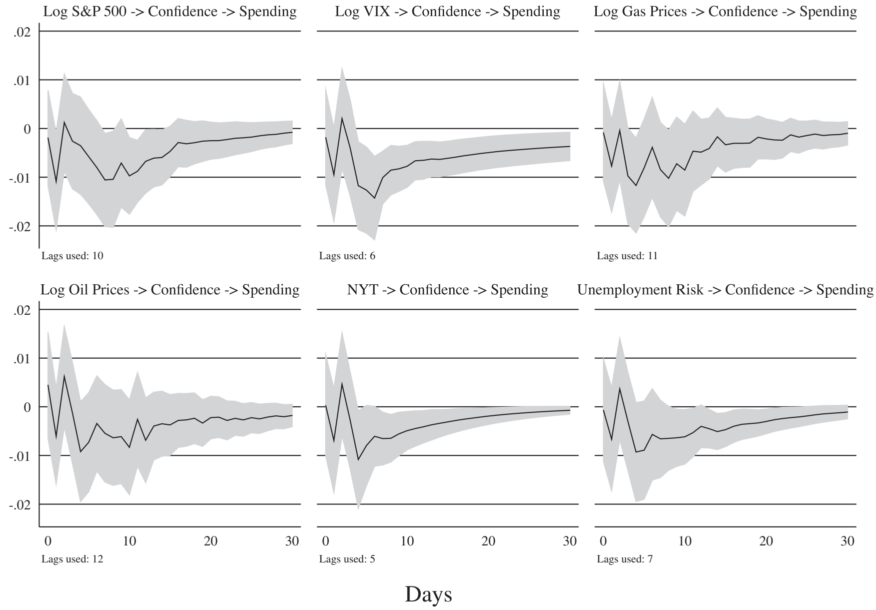

Fig. 5. Impulse response of spending following a positive shock to confidence—three-variable VARs
Source: 2008 G1K and other sources; see text for details.
Notes: Each VAR is fitted separately and contains three variables: 'economic indicator', consumer confidence, and log of expenditure. 'Economic indicator' is either the log of S&P 500, log of VIX, log of gas prices, 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 quadratic trend, day-of-week dummies, Federal Open Market Committee meeting dates, and an indicator for the last week of the year. 95% 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.

4.2 Results from three-variable VARs

In Fig. 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 six economic indicator variables: log of S&P 500, log of VIX, log of oil prices, log of gas prices, *New York Times* article counts, or unemployment risk. The impulse response functions in all of the sub-plots are strikingly similar to the response function found in Fig. 4(a).[18]

As is shown in the figure, at the 5% significance level, the reaction of spending in the first days following a confidence shock is mostly imprecise. Overall, the innovations to confidence seem to contain information that has a modest impact on spending and that is not reflected in either one of the other economic variables. However, the cumulative effect of confidence on spending is small.

18 Note that the six models use different lag lengths, each denoted in the figure. This is due to how the information criteria select the number of lags. AIC tends to pick four lags, whilst BIC tends to favor a smaller number, such as one or two. In general, increasing the number of lags does not alter the qualitative shape of the impulse responses, although with more lags, they become more zigzag; see Fig. A.14 in the Online Appendix.

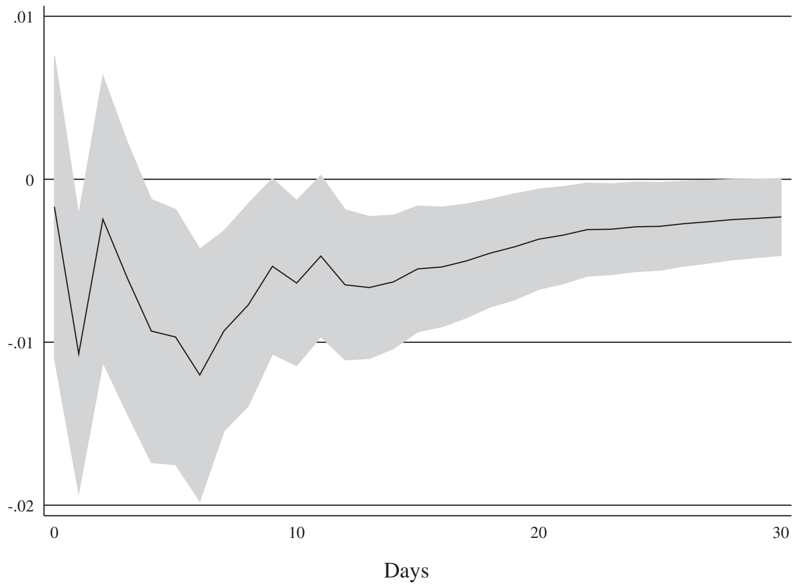

Fig. 6. Impulse response of spending following a positive shock to confidence—multivariate VAR
Source: 2008 G1K and other sources; see text for details.
Notes: VAR contains eight variables: log of S&P 500, log of VIX, *New York Times* article counts for *recession*, log of oil prices, log of gas prices, a proxy for unemployment risk (the reducing/hiring gap), consumer confidence, and log of expenditure. Other exogenous variables included weekly linear and quadratic trend, day-of-week dummies, Federal Open Market Committee meeting dates, and an indicator for the last week of the year. 95% 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 → log gas prices → unemployment risk → confidence → spending.

4.3 Results from the multivariate VAR

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, I estimate a multivariate VAR model including all eight variables. The estimation order is log of S&P 500, log of VIX, *NYT* article counts, log of oil prices, log of gas prices, unemployment risk, confidence, and spending. For this multivariate model, the AIC selects only two lags, whilst BIC picks one. At six lags, the p-value of the Lagrange multiplier test is 0.16. Figure 6 plots the impulse response function of spending following a positive shock to confidence from this VAR. Again, following a shock to confidence, spending takes on a very similar pattern to the reaction found in Fig. 4(a).

4.4 Robustness check: alternative estimation orders

To study how robust the pattern of spending is to alternative estimation orders of the VAR, Fig. 7 shows four alternative orthogonalizations. The various estimation orders are explained below the figure. Each figure shows that following a positive shock to confidence, the reaction is qualitatively very similar regardless of the ordering. Analysis of variance decomposition from the alternative estimation orders confirms that regardless of the

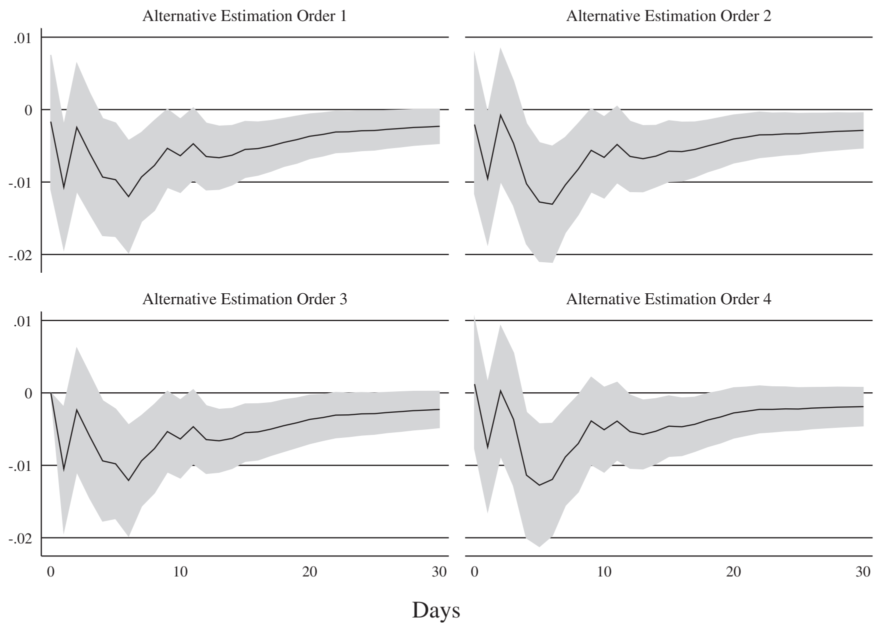

Fig. 7. Sensitivity of the multivariate VAR to alternative estimation orders—impulse response of spending following a positive shock to confidence

Source: 2008 G1K and other sources; see text for details.

Notes: VAR contains eight variables: log of S&P 500, log of VIX, *New York Times* article counts for *recession*, log of oil prices, log of gas prices, a proxy for unemployment risk (the reducing/hiring gap), consumer confidence, and log of expenditure. Other exogenous variables included: weekly linear and quadratic trend, day-of-week dummies, Federal Open Market Committee meeting dates, and an indicator for the last week of the year. 95% 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 → log gas prices → unemployment risk → confidence → spending.

Alternative estimation order 2: log of VIX → *NYT* 'recession' → unemployment risk → confidence → spending → log oil prices → log gas prices → log of S&P 500.

Alternative estimation order 3: log of S&P 500 → log of VIX → *NYT* 'recession' → log oil prices → log gas prices → unemployment risk → spending → confidence.

Alternative estimation order 4: confidence → spending → log of S&P 500 → log of VIX → *NYT* 'recession' → log oil prices → log gas prices → unemployment risk.

estimation order, confidence reacts mostly to its own innovations and, to a lesser extent, stock market shocks; these results are available from the author.

The most salient feature of the dynamic effect in Figs 6 and 7 is that following a positive shock to confidence, there is a small but puzzling decrease in spending. However, in the days right after the confidence shock, the estimated impulse responses have confidence intervals wide enough that it is difficult to rule out a range of effects. This difference is striking when comparing these impulse responses with the clearly positive and statistically significant reaction reported in Barsky and Sims (2012) and Starr (2012). Overall, when compared with lower frequency data, there simply appears not to be much of a reaction of spending following a shock to confidence.

4.5 Robustness check: external validity

One reason the estimated impulse responses in Figs 4–7 are economically small and largely imprecise could be that 2008 was simply a very different year. The unusual abundance of negative news during this year creates an interesting setting for studying consumer behavior, but a possible downside is that there are few positive events throughout this year; hence the estimated statistical link between confidence and spending might not be the same as in a normal year. As a robustness check, I estimate multivariate VARs using the five different estimation orders described in the previous subsection using data for another year, 2011. The benefit of using 2011 is that it is a post-recession year; hence one can hope that the statistical relationships in 2011 are more 'normal' than in 2008.

Using the 2011 G1K, I construct my measure of consumer confidence and the log of spending the same way as for the 2008 data. To estimate a joint multivariate VAR for 2011, I collected the same set of economic indicators as in the analysis for 2008: log of S&P 500, log of VIX, *NYT* article counts for the word recession, log of oil prices, log of gas prices, and the G1K measure of unemployment risk. Figures A.15 and A.16 and Table A.3 in the Online Appendix present descriptive statistics of the 2011 data.

Figure 8 plots the impulse response function of spending following a positive shock to confidence from a VAR using the estimation order from Fig. 6 and, as a further robustness check, the four alternative estimation orders from Fig. 7. In keeping with the models for 2008 (i.e., the models presented in Figs 6 and 7), in each model for 2011, I use six lags.[19] Figure 8 shows that following a shock to confidence, spending initially rises, and then zigzags back to trend. The estimated impulse response of spending is only statistically different from zero some 20 days after the confidence shock.

As with the results for 2008, the analysis for 2011 is best described as a 'null result' — following a positive shock to confidence, daily spending reacts little. Hence, repeating the analysis in a post-recession year such as 2011 and finding a null result as well suggests that the 2008 null result is more likely due to the daily frequency of the data rather than how different 2008 was.

5. Discussion

The estimates suggest that consumer confidence has at best a modest impact on consumer spending. In all the results, the shape of the impulse response function of spending differs from the pattern found in Starr (2012) and Barsky and Sims (2012), where the estimated response is slow-building and permanent. The empirical response of spending to a positive shock to confidence, reported in Figs 4–7, looks nothing like this—spending follows an initial zigzag pattern, drops after a few days, then takes about one month to recover to its initial level.[20]

Two explanations for this pattern can be ruled out. First, we can rule out the explanation of consumer confidence being a manifestation of animal spirits (Barsky and Sims, 2012). In the Barsky-Sims framework, an animal spirits shock implies an initial increase in spending, which we do not observe. Second, we can rule out the explanation that an

19 Figure A.17 in the Online Appendix presents a similar analysis where the models have been estimated using 12 lags instead.

20 The presence of durables, aggregation issues, and filtering further complicate direct comparisons between papers using time series collected at different frequencies.

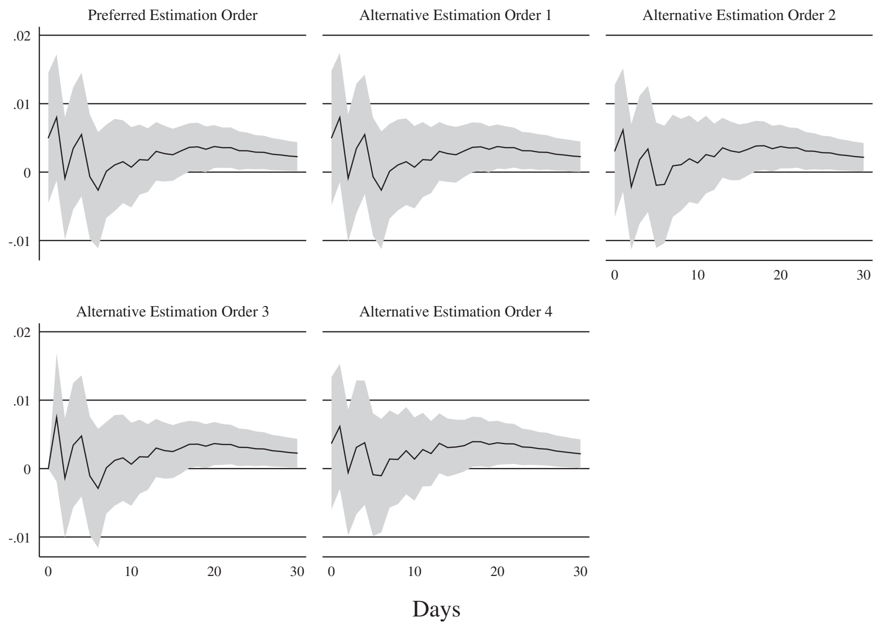

Fig. 8. Robustness check: impulse responses of spending following a positive shock to confidence—multivariate VAR estimated using daily data for 2011

Source: 2011 G1K and other sources; see the text for details.

Notes: VAR contains eight variables: log of S&P 500, log of VIX, *New York Times* article counts for *recession*, log of oil prices, log of gas prices, a proxy for unemployment risk (the reducing/hiring gap), consumer confidence, and log of expenditure. Other exogenous variables included weekly linear and quadratic trend, day-of-week dummies, Federal Open Market Committee meeting dates, and an indicator for the last week of the year. 95% confidence intervals are computed using bootstrapping method with 1,000 draws. Lags used: 6.

Preferred estimation order (same as in Fig. 6): log of S&P 500 → log of VIX → *NYT* 'recession' → log oil prices → log gas prices → unemployment risk → confidence → spending.

Alternative estimation order 1: log of S&P 500 → *NYT* 'recession' → log of VIX → log oil prices → log gas prices → unemployment risk → confidence → spending.

Alternative estimation order 2: log of VIX → *NYT* 'recession' → unemployment risk → confidence → spending → log oil prices → log gas prices → log of S&P 500.

Alternative estimation order 3: log of S&P 500 → log of VIX → *NYT* 'recession' → log oil prices → log gas prices → unemployment risk → spending → confidence.

Alternative estimation order 4: confidence → spending → log of S&P 500 → log of VIX → *NYT* 'recession' → log oil prices → log gas prices → unemployment risk.

increase in consumer confidence amounts to a decrease in uncertainty (Bloom, 2009). This is because if more confidence means less uncertainty, the predicted response of spending following a confidence shock would be a boom-and-bust dynamic. Instead, in the 2008 results, spending takes an initial pause, followed by a drop and rebound. Overall, in 2008 and in 2011, the point estimates of the impulse response of spending in the days following the confidence shock are often so imprecise that it is difficult to rule out a range of effects.

Is there an economic rationale for this imprecise, puzzling response? One explanation for is that consumers face a capacity constraint in their ability to continuously process economic information. For example, models of consumer inattention (Reis, 2006) assume that consumers are forward-looking and choose a level of consumption that reflects their present knowledge about current and future income. At the same time, it is costly to continuously adjust the optimal consumption path. Hence, following the arrival of new information, consumers choose to re-optimize their consumption plans only if the opportunity cost of not adjusting outweighs the fixed cost of adjustment. In other words, consumers are rationally inattentive to small and temporary news shocks.

If on a day-to-day basis consumers are inattentive, then the pattern of spending within a month could be different from the consumption smoothing behavior studied at a month-to-month frequency. To see this, suppose that consumers maximize utility by smoothing consumption across some period of time longer than a day, such as a month. If it is costly for consumers to update their behavior on a day-to-day basis, then following a confidence shock, spending may not display much of a reaction even if confidence contains relevant economic information. However, on a month-to-month or quarter-to-quarter basis, such adjustment costs may not be binding, leading spending to increase following a confidence shock.

Reis's (2006) model illustrates that consumer inattention has different implications for individual and aggregate variables. At the aggregate level, inattention manifests itself in a sluggish but permanent response of consumption following a news shock. How sluggish this response is depends on how many consumers are inattentive at a given point in time. If consumer confidence contains relevant information, but consumers are inattentive, then following a confidence shock, aggregate consumption would smoothly adjust to a permanent new level. Indeed, using quarterly data, Barsky and Sims (2012) find that the reaction of consumption to a confidence shock is gradual but persistent.

It is worthwhile to note that the consumer confidence shocks that lead to a gradual and persistent reaction in spending are themselves long-lived—Barsky and Sims (2012) show that a confidence shock leads to an increase in confidence that lasts over two years. In contrast, the reaction of daily consumer confidence following a shock to itself is short-lived—following an initial increase, confidence returns to trend within a month.[21] One may speculate that following a daily confidence shock, consumers receive information about a change in economic prospects. Since the increase in consumer confidence is only temporary, consumers might decide that the cost of adjusting their daily spending level to reflect this new information is greater than the utility loss of not re-optimizing their expenditure pattern and do not adjust spending.

An indirect test of this hypothesis is to aggregate the level of analysis to a lower frequency level, such as a week. The general pattern of a weekly impulse response can provide suggestive evidence if the rational inattention interpretation holds. Figure 9 shows the estimated impulse responses for spending following shock to confidence estimated using weekly aggregates for 2008, which amounts to 53 observations (52 weeks in 2008 plus the first days of 2009), for different lag selections. When one or two lags are picked, the approximate shape of the impulse responses is more similar to that observed using quarterly data: an increase in spending, following decay back to trend. However, at this level of aggregation, the impulse responses are not statistically significant. Oddly, once three (or more) lags are picked, the impulse response 'bends' in the middle.

21 Figure A.7 in the Online Appendix shows that by the end of 2008, consumer confidence is close to its January 2008 level, whereas at the same time, the Gallup ECI is below its January 2008 level. However, when re-estimating the VARs using the Gallup ECI, the shape of the impulse response of spending is qualitatively very similar to the VAR in Fig. 6; see Fig. A.8 in the Online Appendix.

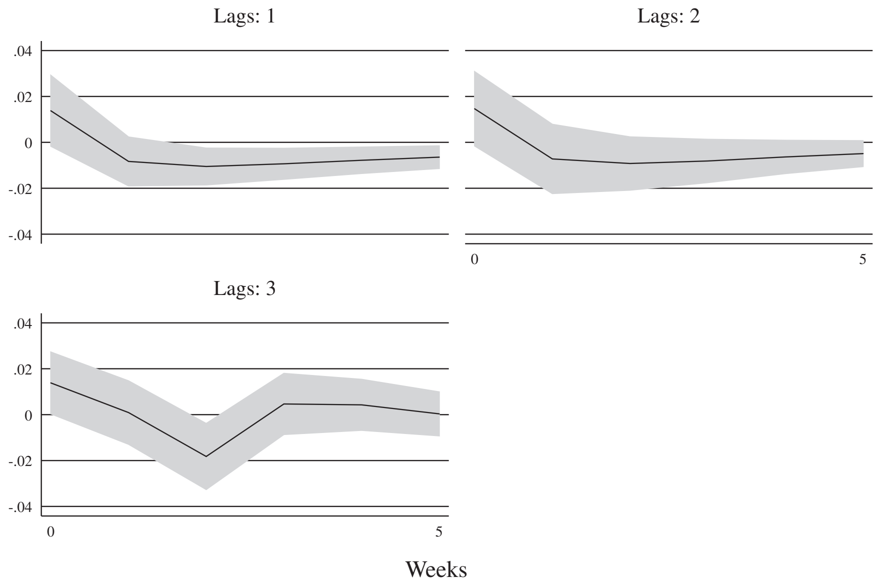

Fig. 9. Impulse response of spending following a positive shock to confidence using weekly aggregates

Source: 2008 G1K.

Notes: VAR with two variables: consumer confidence and log of expenditure. Other exogenous variables included weekly linear trend. 95% confidence intervals are computed using bootstrapping method with 1,000 draws. Lags used: varies. Estimation order: confidence → spending.

Although the available data prevent me from directly answering whether the observed response is due to consumer inattention, the contrast between the reaction of daily spending and weekly spending and the relative similarity between the weekly reaction and quarterly reaction found in previous literature provides suggestive support for this explanation.

5.1 What is consumer confidence?

Analysis of forecast-error variance decompositions (FEVD) from the models in Figs 6 and 7 show that confidence reacts almost entirely to its own shock, and at the end of the 30-day horizon, it accounts for some 60% of the forecast-error variance.[22] Hence, with the exception of stock market shocks and unemployment risk shocks (which each explain about 10% of the forecast-error variance of confidence), the variation in consumer confidence appears to be mostly explained by its own shocks and so prompts the question 'What accounts for day-to-day fluctuations in confidence?'

The G1K microdata underlying the consumer confidence question allow me to investigate if anything can be learned from regressing consumer confidence on the available data. Table 2, column (1), shows the estimated coefficients from ordinary least squares models of consumer confidence, converted to a dummy that equals 1 if the respondent says the economic conditions are 'getting better' and 0 if she says they are 'getting worse', regressed on

22 See Figs A.11–A.13 in the Online Appendix.

Table 2. Regressions of consumer confidence on various covariates

Covariates	(1)	(2)	(3)	(4)
Demographics				
Age/100	−0.321***	−0.277***	−0.332***	−0.331***
	(0.027)	(0.027)	(0.028)	(0.028)
Woman	0.001	0.001	0.003	0.002
	(0.006)	(0.006)	(0.006)	(0.006)
Race/ethnicity				
Other	0.012	0.012	0.011	0.010
	(0.015)	(0.014)	(0.014)	(0.015)
African American/black	−0.084***	−0.090***	−0.097***	−0.097***
	(0.011)	(0.011)	(0.011)	(0.011)
Hispanic	−0.024	−0.022	−0.027	−0.026
	(0.018)	(0.018)	(0.018)	(0.018)
Asian	0.009	0.011	0.008	0.008
	(0.023)	(0.023)	(0.022)	(0.022)
No children	0.012*	0.010	−0.004	−0.004
	(0.006)	(0.006)	(0.007)	(0.007)
Married	0.037***	0.037***	0.026***	0.026***
	(0.006)	(0.006)	(0.006)	(0.006)
Highest completed level of education				
High school degree or diploma	−0.068***	−0.068***	−0.070***	−0.069***
	(0.021)	(0.021)	(0.021)	(0.021)
Technical/vocational school	−0.061***	−0.060***	−0.061***	−0.061***
	(0.022)	(0.022)	(0.021)	(0.021)
Some college	−0.052**	−0.052**	−0.052**	−0.051**
	(0.022)	(0.022)	(0.021)	(0.021)
College graduate	−0.059***	−0.057***	−0.056***	−0.055***
	(0.022)	(0.022)	(0.021)	(0.021)
Post-graduate work or degree	−0.092***	−0.086***	−0.082***	−0.081***
	(0.021)	(0.021)	(0.021)	(0.021)
Gross monthly household income				
Under $60	0.114	0.101	0.077	0.079
	(0.107)	(0.105)	(0.104)	(0.104)
$60 to $499	0.194***	0.166**	0.133**	0.139**
	(0.066)	(0.065)	(0.063)	(0.063)
$500 to $999	0.117*	0.095	0.080	0.085
	(0.065)	(0.064)	(0.062)	(0.062)
$1,000 to $1,999	0.061	0.041	0.024	0.030
	(0.060)	(0.059)	(0.058)	(0.057)
$2,000 to $2,999	0.050	0.031	0.002	0.007
	(0.058)	(0.058)	(0.056)	(0.056)
$3,000 to $3,999	0.056	0.034	−0.004	0.002
	(0.059)	(0.058)	(0.057)	(0.056)
$4,000 to $4,999	0.057	0.035	−0.006	−0.000
	(0.059)	(0.058)	(0.056)	(0.056)
$5,000 to $7,499	0.088	0.064	0.018	0.024
	(0.059)	(0.058)	(0.056)	(0.056)

(continued)

Table 2. Continued

Covariates	(1)	(2)	(3)	(4)
$7,500 to $9,999	0.116*	0.091	0.042	0.047
	(0.059)	(0.058)	(0.056)	(0.056)
$10,000 and over	0.116**	0.088	0.036	0.041
	(0.059)	(0.058)	(0.056)	(0.056)
Yesterday's experiences				
Did you worry about money yesterday?			−0.093***	−0.092***
			(0.007)	(0.007)
Experienced enjoyment yesterday			0.060***	0.059***
			(0.008)	(0.008)
Experienced physical pain yesterday			−0.017**	−0.018**
			(0.008)	(0.008)
Experienced happiness yesterday			0.029***	0.029***
			(0.009)	(0.009)
Experienced worry yesterday			−0.043***	−0.042***
			(0.008)	(0.008)
Experienced sadness yesterday			−0.014*	−0.015*
			(0.008)	(0.008)
Experienced stress yesterday			−0.045***	−0.044***
			(0.007)	(0.007)
Experienced anger yesterday			−0.003	−0.003
			(0.009)	(0.009)
Economic indicators				
Is your company hiring or letting go?(Unemployment risk)				
Hiring		0.117***	0.107***	0.107***
		(0.007)	(0.007)	(0.007)
Letting go		−0.105***	−0.076***	−0.076***
		(0.006)	(0.006)	(0.006)
Log S&P500				−0.090
				(0.174)
Log oil prices				−0.116
				(0.093)
Log VIX				−0.272***
				(0.059)
Log gas				−0.070
				(0.068)
NYT recession				−0.002
				(0.003)
Adjusted R^2	0.015	0.029	0.043	0.045

Source: 2008 G1K and other sources; see the text for details.

Notes: The dependent variable is a dummy that equals 1 if economic conditions are 'getting better' and 0 if economic conditions are 'getting worse'. Robust standard errors are in parentheses (*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$) and are clustered by the date of the interview. The regressions are estimated using ordinary least squares. All of the regressions control for day-of-week dummies and month-of-year dummies. The omitted categories are male, white, has children, not married, has less than a high school degree, has no income, company is not hiring or reducing, did not worry about money yesterday, did not experience enjoyment, pain, happiness, worry, sadness, stress, or anger yesterday.

a set of demographic variables.[23] Although several of these demographics are individually significant, overall these characteristics explain very little of the variation in confidence.

Since the FEVD analysis showed that unemployment risk predicted consumer confidence, in column (2) of Table 2, I additionally control for whether the respondent says that his workplace is hiring or letting people go (I omit the neutral category—neither hiring or letting go). Although including these dummies doubles the adjusted R^2 from 0.015 to 0.029, most of the variation in confidence remains unexplained.

Interestingly, the G1K survey includes several yes or no questions asking the respondents about yesterday's experiences: 'Did you worry about money yesterday?' and 'Did you experience the following feelings a lot yesterday: enjoyment, pain, happiness, worry, sadness, stress, anger?' Although these variables have little time-series variation, they may be useful in explaining the cross-sectional variation in confidence. In column (3) of Table 2, I add this set of dummies to the regression. As it turns out, these largely idiosyncratic experiences correlate strongly with consumer confidence and help increase R^2 to 0.043. Hence, worries about money, pain, and stress decrease confidence whilst enjoyment and happiness improve confidence.

Last, in column (4), I include the daily 'economic indicator' variables used in the VAR models. Although a higher VIX decreases confidence, adding this set of variables only marginally increases the R^2 to 0.045.[24]

The conclusion from Table 2 is that most of the variation in daily consumer confidence remains unexplained, although unemployment risk and idiosyncratic events such as experiencing positive or negative affect the day before help explain some of the variation in confidence.

On a final note in their paper, Barsky and Sims (2012) report that the 'news heard' variables in the University of Michigan Survey explain up to 15% of the variation in the University of Michigan confidence measure, thus further strengthening their interpretation of confidence as 'news'. Bearing this in mind, it is interesting to note that in Table 2, the daily *NYT* article counts are not statistically significant. The discrepancy between how news affects monthly confidence and daily confidence may offer further evidence in favor of the rational inattention interpretation. It appears that on a day-to-day basis, fluctuations in confidence have less to do with news and more to do with recent idiosyncratic experiences of the respondent.

6. Conclusion

This article uses high-frequency identification to study the dynamics of consumer confidence and spending in a VAR model. 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. I use this variation to identify the effect of shocks to consumer confidence on spending. The daily frequency of data makes it

23 The regressions also control for day-of-week dummies and month-of-year dummies.

24 Some of the time-series variation in the 'economic indicator' variables is subsumed by the month-of-year dummies. If I do not control for any time effects, then the economic indicators become individually significant, but including these variables only increases the goodness of fit marginally, from 0.035 to 0.038 (not shown). *NYT* article counts are never statistically significant.

possible to orthogonalize shocks to consumer confidence with respect to other economic shocks.

I find that the estimated relationship between consumer confidence and spending is weak. Following a positive confidence shock spending first pauses, then decreases, but overall this reaction is imprecisely estimated. This unexpected pattern is robust to various alternative estimation orders but is in sharp contrast with the pattern observed when studying data collected at a lower frequency.

I speculate that this finding can be explained by consumer inattention. If it is costly for consumers to revise their optimal monthly consumption plans from one day to another, consumers will only alter their behavior if the cost of not acting on the arrival of new information exceeds a fixed cost of adjustment. As fluctuations in daily consumer confidence are only short-lived, a shock to daily confidence may not warrant a statistically discernible change in day-to-day consumer behavior.

Since the media devote considerable attention to changes in consumer confidence, results of this article might be of interest to a broad audience. Although unexpected improvements in quarterly consumer confidence have been shown to have a persisting positive effect on consumption, the day-to-day fluctuations in confidence are short-lived and have little impact on daily spending.

Supplementary material

Supplementary material is available online at the OUP website.

Acknowledgements

I thank Alan Krueger, Jim Harter, and Dennis Jacobe for the opportunity to work with the G1K data. I am grateful to the editor, Raoul Minetti, and two anonymous referees for helpful and constructive comments. I am also indebted to Henry Farber, John Hassler, Lars Lefgren, Matthew Lindquist, Andreas Mueller, Åsa Rosén, Sam Schulhofer-Wohl, Eric Sims, and Stephen Woodbury for their comments. I thank the seminar participants at Princeton University, Stockholm University, Uppsala University, the University of Cambridge, Copenhagen Business School, the University of Notre Dame, and the AEA and EEA conferences for their comments.

References

Acemoglu, D. and Scott, A. (1994) Consumer confidence and rational expectations: are agents' beliefs consistent with the theory?, *Economic Journal*, 104, 1–19.

Alexopoulos, M. and Cohen, J. (2009) Uncertain times, uncertain measures, Working Paper 352, Department of Economics, University of Toronto.

Attanasio, O. P. and Weber, G. (1993) Consumption growth, the interest rate and aggregation, *Review of Economic Studies*, 60, 631–49.

Barsky, R.B. and Sims, E.R. (2012) Information, animal spirits, and the meaning of innovations in consumer confidence, *American Economic Review*, 102, 1343–77.

Beaudry, P. and Portier, F. (2006) Stock prices, news, and economic fluctuations, *American Economic Review*, 96, 1293–307.

Bee, A., Meyer, B.D., and Sullivan, J.X. (2012) The validity of consumption data: are the consumer expenditure interview and diary surveys informative?, NBER Working Paper 18308, Cambridge, MA.

Bloom, N. (2009) The impact of uncertainty shocks, *Econometrica*, 77, 623–85.

Browning, M. and Crossley, T.F. (2009) Shocks, stocks, and socks: smoothing consumption over a temporary income loss, *Journal of the European Economic Association*, 7, 1169–92.

Carroll, C.D. and Dunn, W.E. (1997) Unemployment expectations, jumping (S,s) triggers, and household balance sheets, in B.S. Bernanke and J. Rotemberg (eds), *NBER Macro-economic Annual*, vol. 12, MIT Press, Cambridge, MA.

Carroll, C.D., Fuhrer, J.C., and Wilcox, D.W. (1994) Does consumer sentiment forecast household spending? If so, why?, *American Economic Review*, 84, 1397–408.

Chauvet, M. and Guo, J.-T. (2003) Sunspots, animal spirits, and economic fluctuations, *Macroeconomic Dynamics*, 7, 140–69.

Deaton, A. (2011) The financial crisis and the well-being of Americans, *Oxford Economic Papers*, 64, 1–26.

Deaton, A. and Arora, R. (2009) Life at the top: the benefits of height, *Economics and Human Biology*, 7, 133–6.

Dominitz, J. and Manski, C.F. (2004) How should we measure consumer confidence?, *Journal of Economic Perspectives*, 18, 51–66.

Doms, M. and Morin, N. (2004) Consumer sentiment, the economy, and the news media, Finance and Economics Discussion Series Working Paper 2004-51, Divisions of Research and Statistics and Monetary Affairs, Federal Reserve Board, Washington, DC.

Edelstein, P. and Killian, L. (2007) Retail energy prices and consumer expenditures, Centre for Economic Policy Research Discussion Paper 6255, Centre for Economic Policy Research, London.

Hamilton, J.D. (2009) Causes and consequences of the oil shock of 2007–08, *Brookings Papers on Economic Activity*, 40, 215–83.

Knotek, E .S. II. and Khan, S. (2011) How do households respond to uncertainty shocks?, *Federal Reserve Bank of Kansas City Economic Review* (Second Quarter), 63–92.

Krueger, A.B. and Kuziemko, I. (2013) The demand for health insurance among uninsured Americans: results of a survey experiment and implications for policy, *Journal of Health Economics*, 32, 780–93.

Ludvigson, S.C. (2004) Consumer confidence and consumer spending, *Journal of Economic Perspectives*, 18, 29–50.

Matsusaka, J.G. and Sbordone, A.M. (1995) Consumer confidence and economic fluctuations, *Economic Inquiry*, 33, 296–318.

Mishkin, F.S. (1978) Consumer sentiment and consumer durable expenditure, *Brookings Papers on Economic Activity*, 1, 217–31.

Petev, I., Pistaferri, L., and Saporta-Eksten, I. (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*, Russell Sage, New York, NY.

Poterba, J. and Samwick, A.A. (1995) Stock ownership patterns, stock market fluctuations, and consumption, *Brookings Papers on Economic Activity*, 26, 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 Shapiro, M.S. (1998, June) Costly capital reallocation and the effects of government spending, *Carnegie-Rochester Conference Series on Public Policy*, 48, 145–94.

Reis, R. (2006) Inattentive consumers, *Journal of Monetary Economics*, 53, 1761–800.

Romer, C.D. (1990) The great crash and the onset of the Great Depression, *Quarterly Journal of Economics*, 105, 597–624.

Romer, C.D. and Romer, D.H. (1990) Does monetary policy matter? A new test in the spirit of Friedman and Schwartz, NBER Working Paper 2966, Cambridge, MA.

Sims, C.A. (2003) Implications of rational inattention, *Journal of Monetary Economics*, 50, 665–90.

Souleles, N.S. (2004) Expectations, heterogeneous forecast errors, and consumption: micro evidence from the Michigan sentiment surveys, *Journal of Money, Credit, and Banking*, 36, 39–72.

Starr, M. (2012) Consumption, sentiment, and economic news, *Economic Inquiry*, 50, 1097–111.