

NBER WORKING PAPER SERIES

INFORMATION, ANIMAL SPIRITS, AND THE MEANING OF INNOVATIONS
IN CONSUMER CONFIDENCE

Robert B. Barsky
Eric R. Sims

Working Paper 15049
<http://www.nber.org/papers/w15049>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2009

The authors gratefully acknowledge helpful comments and suggestions from Susanto Basu, Alan Blinder, John Cochrane, Olivier Coibion, Daniel Cooper, Mike Elsby, Yuriy Gorodnichenko, Christopher House, Peter Ireland, Lutz Kilian, Miles Kimball, Guido Lorenzoni, Serena Ng, Matthew Shapiro, Robert Solow, Frank Vella, two anonymous referees, and seminar participants at Boston College, Johns Hopkins, Michigan, the Wharton School, the Russell Sage Foundation, and the NBER Summer Institute in Monetary Economics. We also thank John Fernald for providing his quarterly utilization-corrected TFP measure. Robert Barsky acknowledges generous support as a visiting scholar at the Russell Sage Foundation and the Bank of Israel. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2009 by Robert B. Barsky and Eric R. Sims. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence
Robert B. Barsky and Eric R. Sims
NBER Working Paper No. 15049
June 2009
JEL No. E2,E3,E32

ABSTRACT

Innovations to measures of consumer confidence convey incremental information about economic activity far into the future. Comparing the shapes of impulse responses to confidence innovations in the data with the predictions of a calibrated New Keynesian model, we find little evidence of a strong causal channel from autonomous movements in sentiment to economic outcomes (the “animal spirits” interpretation). Rather, these impulse responses support an alternative hypothesis that the surprise movements in confidence reflect information about future economic prospects (the “information” view). Confidence innovations are best characterized as noisy measures of changes in expected productivity growth over a relatively long horizon.

Robert B. Barsky
Department of Economics
University of Michigan
Ann Arbor, MI 48109-1220
and NBER
barsky@umich.edu

Eric R. Sims
University of Notre Dame
eric Sims@umich.edu

I. Introduction

In the popular press and much of the business community it continues to be an article of faith that “consumer confidence” has an important role – both prognostic and causal – in macroeconomics. On the other hand, the stance of the rather limited academic literature on confidence is far more ambiguous. The judgments range from the conclusion that confidence measures have an important role both in prediction and understanding the cause of business cycles, to the view that they contain important information but have little role in the assignment of causality, to the verdict that they have no value even in forecasting.

There are, broadly speaking, two contrasting approaches to the role of confidence in macroeconomics. The first, which we will refer to as the “animal spirits” view, posits autonomous fluctuations in beliefs and consumption that in turn have causal effects on economic activity. In the proceedings of a symposium on the causes of the 1990-1991 recession, both Hall (1993) and Blanchard (1993) regard exogenous movements in consumption as a cause of business cycles.¹ Indeed, Blanchard proposes that the cause of the recession was a powerful, long-lasting negative consumption shock associated with an exogenous shift in pessimism that had a causal effect on consumption and overall aggregate demand. While not fully pursuing the idea in his brief paper, Blanchard proposes that one might be able to test this hypothesis on the basis of the observation that such an exogenous shift in pessimism ought to have only temporary effects on consumption.²

The second view of confidence – what we will call the “information view” – suggests that a relationship between innovations in measures of consumer confidence and subsequent macroeconomic activity arises because confidence measures contain fundamental information about the current and future states of the economy. For example, Cochrane (1994b) proposes that consumption surprises proxy for news that consumers receive about future productivity that does not otherwise show up in econometricians’ information sets. His attempt to reconcile VAR

¹ In an interesting but almost forgotten early contribution, Hall (1986) – partially repudiating Hall (1978) – argues that an important fraction of the random walk in consumption comes not from the expectational surprise in the Euler equation but from a second disturbance that he has more recently referred to as “spontaneous consumption”. In Hall (1993), this is interpreted as a shock to the taste for consumption relative to leisure.

² In some ways, a limiting case of animal spirits appears in the “sunspot” literature. Though pinned down only by extrinsic coordinating variables, expectations in the equilibria of these models are self-fulfilling, and thus not irrational (see Farmer (1999)). The existence of sunspot equilibria depend on strong increasing returns, supply externalities, or other mechanisms that are typically not accepted as empirically plausible. The notion of animal spirits in this paper does not encompass sunspots.

evidence with theory closely anticipates the “news approach to business cycles” of Beaudry and Portier (2004, 2006). They analyze models where agents become aware of changes in future productivity orthogonal to current productivity, and argue that stock price innovations proxy for future technological improvement not reflected in current technology. The “information view” of confidence supposes that confidence innovations might contain similar information.

In Section II of the paper, we first show that unexplained innovations in several variables representing survey responses to forward-looking questions from the Michigan Survey of Consumers have powerful predictive implications for the future paths of macroeconomic variables. In particular, within the context of augmented consumption-income VARs, we show that unexplained innovations in the responses to several consumer confidence questions have significant, slowly building, and apparently permanent implications for output and consumption. Confidence is not highly Granger-caused by income or consumption, nor are its innovations highly correlated with innovations in those variables. Responses to little-used survey questions on “news heard” do help to somewhat explain confidence innovations, but with only a very modest incremental R^2 . These observations point to the conclusion that these measures of consumer confidence are not merely noise, nor are they simply reflections of macroeconomic news reports or innovations in other variables with which they are correlated.

In Section III we attempt to distinguish the hypothesis that these impulse responses indicate a causal channel from sentiment to economic outcomes (the “animal spirits” view) from the alternative interpretation that the surprise confidence movements summarize information about economic prospects known to consumers (the “information” view). To provide a framework for distinguishing these alternative views of confidence, we present a highly stylized New Keynesian model with three kinds of shocks. The first shock is an immediate and unexpected improvement in productivity (a “level shock”). The second is a reflection of genuine news that productivity will grow more rapidly for a substantial period of time into the future (a “growth shock”, also to be referred to as an “information shock” because it conveys information about future productivity that cannot be fully inferred from current productivity).³ We only permit households to observe a noise-ridden signal of the information shock to technology. We interpret the noise innovation in the signal as an “animal spirits shock” as it is associated with

³ We employ the term “information” in the same way Cochrane (1994b) and Beaudry and Portier (2006) use the word “news”. Lorenzoni (2008), somewhat confusingly, uses the term “news” to refer to noise in a public signal, which functions much like our animal spirits shock.

erroneous consumer optimism or pessimism. This shock can be given alternative less structural interpretations, and in equilibrium its implications are similar to those of an exogenous innovation to the Euler equation. Regardless of the particular interpretation, a series of positive animal spirits shocks might capture the putative “irrational exuberance” of the 1920s or 1990s, while a predominance of negative shocks would usher in a period of excessive pessimism.

The model has clear implications for the response of the endogenous variables to each of the three shocks. “Animal spirits” shocks behave as aggregate demand shocks – they are associated with transitory increases in output that attenuate over time, and they produce both inflation and increases in real interest rates. “Information shocks” regarding future productivity and shocks to current productivity are followed by gradual movements in the macroeconomic variables that are not subsequently reversed. Both of these fundamental shocks are also associated with rising real interest rates. Thus, the model yields two primary criteria by which to distinguish animal spirits from fundamental shocks: positive animal spirits shocks are followed by transitory movements in real activity and increases in inflation, while favorable fundamental shocks may result in permanent movements in activity and may be either inflationary or deflationary.

In Section IV, we estimate an expanded VAR with the variables implied by the model augmented with a measure of confidence. As in the three variable systems of Section II, the results show that confidence innovations are associated with little immediate response of real activity but prolonged growth in consumption, income, and measured productivity. There is no evidence of reversion in these variables – in particular, the point estimates suggest that income and consumption are higher by more than two-thirds of a percent in the long future in response to a confidence innovation, with the confidence bands associated with these impulse responses lying above zero at horizons in excess of ten years. Confidence innovations are associated with transitory increases in real interest rates and hours of work, and also lead to a large and persistent reduction in inflation. These empirical responses are not at all similar to the implications of animal spirits shocks in our model, nor are they particularly consistent with the theoretical responses to level shocks.

We next postulate a structural equation in which surprise movements in confidence are attributable to the signal agents receive about the growth rate and to the innovation in the current state of productivity. We estimate a subset of the structural parameters of the model via a

modified version of simulated method of moments. We are able to resoundingly reject the hypothesis that animal spirits shocks (as specified in this paper) are an important source of the observed relationships between confidence innovations and macroeconomic variables. On the other hand, we do find convincing evidence in favor of the information interpretation of consumer confidence. The implications of confidence innovations for output and spending at short horizons are far too small for confidence to be primarily a reflection of changes in current fundamentals, yet the longer horizon implications are far too large and significant for confidence innovations to not be conveying information about fundamentals. Our results suggest that there are information shocks about future productivity not wholly reflected in current productivity, and that these shocks account for a significant fraction of the innovation in measured confidence.

II. Income, Consumption, and Confidence

(a) Cochrane's Bivariate VAR

We begin with the dynamics of income and consumption as implied by the bivariate vector autoregression discussed by Cochrane (1994a). In particular, we estimate a two variable system consisting of the log of real GDP and the log of real consumption of services plus non-durables, both in per capita terms after dividing by the civilian non-institutionalized population aged sixteen and over. The data are seasonally adjusted measures at a quarterly frequency from the first quarter of 1960 to the third quarter of 2007. The data strongly suggest that the variables are cointegrated, and the estimated cointegrating vector is sufficiently close to $[1,-1]$ that we follow Cochrane and others in imposing it. While popular information criteria generally favor a small number of lags (one or two), we take a conservative stance and estimate the VAR with four lags.

Cochrane orthogonalizes the innovations so that consumption is ordered first. This ordering is implied by a simple permanent income model in which all information is immediately reflected in consumption.⁴ However, the line of inquiry in his subsequent paper (Cochrane (1994b)) suggests a focus on the alternative ordering with income first; there the focus is on the information about future income embodied in consumption but not in current income. Figure 1

⁴ The consumption \rightarrow income ordering also splits income fairly neatly into permanent and transitory components, as any innovation to income not reflected in consumption ought to be transitory under a partial equilibrium view of the PIH.

presents impulse responses under both orderings, with the solid line referring to the ordering with consumption first and the dashed line to the orthogonalization with income ordered first. The key feature of these impulse responses is that innovations in consumption – whether or not they are orthogonalized with respect to income – are associated with powerful and prolonged subsequent increases in income. At the shorter horizons, most of the movement in income is explained by its own innovation, but the “effects” of a consumption innovation build over time so that much or all of the permanent component of GDP appears to be captured by innovations in consumption. In short, results from this two variable VAR suggest that “consumption shocks” convey news about income many periods into the future.

As Cochrane (1994b) stresses, a natural explanation for the finding that consumption innovations predict much of future output is that agents have some advance knowledge about future income that they use when making consumption decisions. Forward-looking questions on surveys of consumer expectations and attitudes might potentially provide a direct measure of such information, and thus a direct test of Cochrane’s hypothesis. Is much or most of the information embodied in consumption picked up by survey expectations of future output? Do the survey data indicate, on the other hand, that consumers receive a great deal of news that is not reflected in current consumption? We turn to these questions now, introducing some expectational measures from the Michigan Survey of Consumers and augmenting the bivariate consumption-income VARs with these variables.

(b) Confidence Data

The survey measure that we will make the most use of in this paper, which we call E5Y, summarizes responses to the following question: “Turning to economic conditions in the country as a whole, do you expect that over the next five years we will have mostly good times, or periods of widespread unemployment and depression, or what?” The variable is constructed as the percentage giving a favorable answer minus the percentage giving an unfavorable answer plus one hundred.⁵ Our particular affinity for this question arises from the fact that it is aimed at gauging expectations over a relatively long horizon, and because of its specificity as to the

⁵ Thus a value of 100 is a “neutral” position, while a value of 140 means that the fraction of responses reflecting optimism about the future exceeds the fraction reflecting pessimism by forty percentage points.

relevant time frame.⁶ However, its correlation with the response to a similar question specifying a horizon of only twelve months (a variable we call E12M) is 91 percent, and its correlation with another concerning expected changes in personal financial situation over the next twelve months is 85 percent. The correlation of E5Y with the overall expectations component of the Michigan index exceeds 95 percent. Our results in this section are essentially unchanged by the substitution of either of these alternative expectations variables. The alternative questions are described in more detail in the Appendix.

Figure 2 plots E5Y and E12M against time. Both series undergo repeated dramatic swings though (as we would expect) the twelve-month-ahead expectations are more volatile than the expectations over a five year horizon. Both variables are quite stationary. The cross-correlogram between E5Y and the conventional Hodrick-Prescott detrended GDP (not shown) indicates that the expectations are by no means a reflection of current output; the contemporaneous correlation between detrended GDP and E5Y is essentially zero. E5Y is negatively correlated with the output gap lagged several periods, and positively correlated with the gap several quarters ahead.

(c) Augmenting the Bivariate VAR with Confidence Measures

We begin by augmenting Cochrane's income-consumption VAR with E5Y. As before, the system is estimated allowing cointegration between consumption and income, with four lags of each variable. Because confidence measures are clearly stationary, E5Y cannot enter into the long run equilibrium relationship, and we once again impose that the cointegrating vector between consumption and income is $[1, -1]$.⁷ It is necessary to make some choices as to how to orthogonalize the innovations. It is important to understand that alternative orthogonalizations in this context are not to be thought of as minimum delay restrictions that delineate alternative structural models; in almost any sensible model, innovations in the underlying structural shocks

⁶ Some might argue as well that this question gives the animal spirits hypothesis its "best shot". One argument is that individuals are likely to be more sober-minded in assessing family resources than in forming expectations about the national economy. Another is based on animal spirits models that focus on strategic complementarity; in those models beliefs about the economic activities of other agents are central.

⁷ Formally, the system features E5Y in levels, consumption and income in first differences, and the lagged (log) difference between consumption and income as an exogenous variable. Our results are virtually identical when estimating the full system in levels.

should affect all three variables instantaneously. Attempts to think about ordering should instead focus on “assigning” the common component of the information in innovations to one or another variable so as to provide upper and lower bounds for the amount of information content in each of the series.

To begin to assess the extent to which the “consumption shocks” in the bivariate VAR are in fact “information shocks” that are well captured by innovations in the survey expectations, we compute impulse responses with E5Y ordered first. As in Cochrane (1994a), income is ordered last, though our results from the augmented consumption-income VARs are largely invariant to the placement of income in the ordering. Figure 3(a) presents the impulse responses to E5Y and consumption innovations under this orthogonalization. The dashed lines represent 90 percent bias-corrected bootstrap confidence bands.⁸ As in Cochrane’s two variable system, consumption behaves roughly like a random walk in response to its own innovation. In response to a consumption innovation output jumps up on impact, follows a slight hump-shape, and levels off at roughly 0.4 percent higher than its pre-shock value. Though not shown, output displays a large and significant response to its own innovation that dissipates rather quickly. The part of the output innovation that is orthogonal to consumption predicts no significant movement in consumption at any horizon.

An innovation to E5Y has very small (though statistically significant) implications for both consumption and output on impact. The small impact effects are followed by slowly-building, statistically and economically significant, and apparently permanent responses of both consumption and output. In particular, a one standard deviation innovation to E5Y predicts levels of output and consumption that are roughly 0.7 percent higher forty quarters hence; further, the long run responses of both consumption and GDP to an E5Y innovation are both statistically significant at better than the 90 percent level. E5Y responds significantly neither to income nor consumption innovations; its own innovation accounts for more than 95 percent of its forecast error variance at all horizons under this ordering.

E5Y innovations thus clearly convey important information about the future time paths of real variables, with “effects” that show no tendency to attenuate even at long horizons. However, to what extent are innovations in E5Y simply reflective of information contained in

⁸ In particular, we generate the confidence bands from the empirical distribution of impulse responses based on 2000 bootstrap draws using bias-corrected OLS slope coefficients as proposed by Kilian (1998).

consumption? To address this possibility, we re-order the variables in the system such that E5Y is orthogonalized with respect to consumption. As before, output is ordered last in the system. Figure 3(b) presents impulse responses with this particular ordering.

The qualitative features of the impulse responses are unaffected by the alternative orthogonalization. In particular, E5Y innovations orthogonal to consumption still predict slowly-building and permanent responses of both output and consumption. The point estimates are slightly smaller than in the case with E5Y ordered first, with a one standard deviation innovation to E5Y prognostic of long run increases in both consumption and output of slightly more than 0.5 percent (as opposed to 0.7 percent with E5Y ordered first). E5Y also responds significantly (in the statistical sense) to a consumption innovation, but the point estimate is small and the response is statistically significant only for a few quarters.

Figure 4 graphically depicts the variance decompositions of consumption, income, and E5Y under both orthogonalizations. Regardless of ordering, own innovations account for the bulk of the forecast error variance of output at short horizons and virtually none at longer horizons. Ordered first, E5Y innovations account for more than 60 percent of the forecast error variance of income and consumption at long horizons. Even after orthogonalization with respect to consumption, innovations to E5Y still account for more than 30 percent of the long horizon forecast error variance of both income and consumption. We can thus fairly easily reject the hypothesis that E5Y simply reflects information available in consumption. Rather, innovations in E5Y and in consumption each convey news about future output that is not subsumed in the other.

We now examine several variations on the three variable VAR using alternative measures of consumer confidence. First, we substitute the relative score from the question on the Michigan Survey concerning expected personal financial situation (PFE) in place of E5Y. This question gauges expectations analogously to E5Y and E12M, although it specifically asks for expectations concerning personal situations as opposed to aggregate expectations.⁹ The second

⁹ Dominitz and Manki (2004) express doubt that consumers can give meaningful responses to survey questions concerning aggregate as opposed to individual expectations, and they point to the higher volatility of responses to questions like E5Y versus questions like PFE as support. Given the structure of the questions, however, we would in fact expect aggregate questions to have greater volatility even if individuals are equally capable of answering both kinds of questions accurately. For example, even in severe recessions most people do not personally experience layoffs. The typical respondent who says that the national economy will exhibit “periods of widespread unemployment or depression” is predicting that a significant minority of *others* will experience layoffs while his or her own income is stable by comparison.

modification is to use the Index of Consumer Sentiment (ICS) in place of the purely forward-looking survey questions. While the ICS is the most reported measure of consumer confidence (both by the press and in the academic literature), it is an average of survey responses to both forward-looking and retrospective questions, and thus its interpretation is unclear. For a more detailed description of these alternative confidence measures and their statistical relationships with E5Y, the interested reader is referred to the Appendix.

Figure 5 presents impulse responses to confidence innovations in our three variable system with three alternative measures of confidence: E5Y, PFE, and ICS. We order the confidence measure first in the system, impose cointegration between consumption and output, and employ a lag structure of four.¹⁰ There is very little qualitative or quantitative difference between the results using E5Y or any of the other broad confidence measures. The seeming disparity between some of our results and others in the academic literature thus does not appear to be attributable to different measures of confidence.¹¹ Use of other alternative confidence measures – such as E12M or the expectations index of the Michigan Survey – and alternative measures of consumption and output (for example, durable goods consumption or private sector GDP) also produce very similar impulse responses.

In summary, innovations in expectational variables from the Michigan Survey of Consumers are powerful predictors of changes in output and future spending that last for the foreseeable future. This finding obtains regardless of whether or not the confidence innovations are orthogonalized with respect to current spending. In Section III we will argue, based on model with both shocks to information and animal spirits shocks, that the apparent permanence of the impulse responses of consumption and output to confidence shocks is more consistent with an information view of confidence than it is with an animal spirits interpretation.

Our finding that unexpected increases in confidence imply predictably higher subsequent consumption is somewhat related to the results of Carroll, Fuhrer, and Wilcox (1994), who focus

¹⁰ Alternative orderings with the confidence measure after consumption also produce quite similar results.

¹¹ Among papers in this literature that find a small role for consumer confidence measures in predicting the future time path of economic variables are Mishkin (1978), Leeper (1992), Mehra and Martin (2003), and Croushore (2005). Matsusaka and Sbordone (1995) and Howrey (2001) report a much stronger prognostic role for confidence, while Ludvigson (2004) takes something of a middle ground. Souleles (2004) analyzes the micro data underlying aggregate confidence data used in the present paper. However, the most important difference between our results and the results in these papers is that by looking at impulse responses to confidence innovations many periods into the future, we are able to recover the longer run implications of confidence innovations that are in fact more powerful than are the short run business cycle “effects”.

on one-period-ahead consumption growth. These authors regard Granger causality from confidence to consumption growth partly as a failure of the PIH along the lines of short-term stickiness of consumption.¹² This focuses excessively on the short run and reflects a decidedly partial equilibrium approach. Our finding that consumption tracks predictable income increases over periods of several years suggests that the predictability of consumption growth is better thought of in terms of an endowment economy along the line of Lucas (1978), in which consumers may believe that income will be higher in the future, but can in the aggregate do little to increase current consumption in anticipation. One implication of this interpretation is that positive confidence innovations should be associated with increases in expected real rates of return. This implication will be explored in more detail in the next section, and we will see that, in addition to being an implication of a simple general equilibrium model, it also holds in the data.

(d) What is the News?

In the augmented consumption-income VAR, E5Y and other overall confidence measures are roughly exogenous. With E5Y ordered first, more than 95 percent of the forecast error variance of confidence is explained by its own innovation at every horizon. Even when confidence is allowed to respond contemporaneously to consumption innovations, the fraction of the forecast error variance of confidence attributable to its own innovation always exceeds 85 percent.

What kinds of news might explain these surprise movements in consumer confidence? The Michigan Survey of Consumers, in addition to the questions already discussed, also asks respondents to report any recent “news heard” concerning the economy. It seems natural to include a brief investigation of the relationship between this reported economic news and responses to the survey questions concerning overall expectations of aggregate and individual economic conditions. For a complete description of the news heard questions, see the Appendix.

Respondents give answers to a question asking them to report favorable or unfavorable economic news, and their answers are tabulated into arbitrary, but generally well-defined, categories. Figure 6 presents spike plots for several of the more popular response categories

¹² The proposed solution is that increases in confidence measures summarize information possessed by “rule of thumb consumers” whose consumption is excessively tied to current income. The authors do reject that this hypothesis is a complete explanation of the Granger causality from confidence to consumption.

across time. Most categories (such as trade deficit, government budget deficit, etc.) record very few responses in a typical quarter. Rather clearly, the most consistently popular concern news about prices and news about employment. Other responses stand out in particular time periods. Examples are a high incidence of mention of “energy crisis” during periods of the 1970s and early 1990s as well as news heard concerning the stock market sporadically across the sample period, but most frequently during the 1990s.¹³

In Table 1 we present coefficient estimates from regressions of the E5Y innovations from the three variable VAR on selected categories of news. Most of the news heard categories have coefficients of the expected signs – an increase in the percentage of respondents reporting favorable news is positively correlated with the confidence innovation and vice versa. Favorable or unfavorable news about general prices and favorable news about the stock market are significant covariates with the E5Y innovation at the 10 percent level or better. News about employment and favorable news about the stock market have no significant correlation with the E5Y innovation. Unfavorable news about government policies also has a statistically significant coefficient at the 10 percent level. The adjusted R^2 from these regressions ranges from 0.10 to 0.15, suggesting that the bulk of E5Y innovations remain inexplicable from particular categories of news heard. Use of other more obscure categories of news heard produce insignificant coefficient estimates that frequently reduce the adjusted R^2 in the regressions. We also ran a specification that included the news heard variables in the income-consumption VARs directly. This produced impulse responses of consumption and income which were much weaker than when using the broader confidence measures.

Innovations to measures of consumer confidence evidently convey information about income many periods into the future, much of which is not reflected in current consumption or income innovations, and the surprise movements in the confidence measures are not attributable to tangible news. Some might find it surprising that the answers of largely naïve respondents to rather crude questions could be so informative. As emphasized in Cochrane (1994b), however, such expressions of surprise fail to recognize the role of information aggregation. As Cochrane puts it (see p. 350), “Ask a consumer about next year’s GDP, and he will say ‘I don’t know.’ But

¹³ The data summarizing responses to the “news heard” questions do not have the statistical properties of “news” in the rational expectations sense. Rather, the data on news reports are highly serially correlated. This may be due to gradual diffusion of news reports along the lines of Carroll’s (2003) epidemiological model, or it may reflect merely the wording of the question, which refers to news heard in the “last several months”.

he may know that his factory is closing, and hence he is consuming less. This idiosyncratic shock is correlated with future GDP.” Just as consumption data aggregate idiosyncratic information, consumer confidence data aggregate information from many sources and many individuals.¹⁴

III. Information and Animal Spirits in a New Keynesian Model

The results of the previous section suggest that survey measures of consumer confidence ought to be taken seriously. The observation that unexpected movements in confidence appear to have permanent implications for output and consumption seems inconsistent with an interpretation in which confidence innovations represent autonomous fluctuations in sentiment (i.e. animal spirits), but perhaps consistent with the notion that confidence reflects households’ information about current and/or future fundamentals. To subject these statements to further scrutiny requires reference to a theoretical model that contains both fundamental and animal spirits shocks.

In this section, we develop a simple New Keynesian general equilibrium model with three structural disturbances. The two fundamental shocks are a level shock and what we call an information shock. The level shock is an immediate and permanent innovation to the level of technology, while the information shock is a persistent but transitory innovation to the growth rate of technology. We call it an information shock because it portends of a permanent change in technology orthogonal to the present. We only allow households to observe a noise-ridden signal of the growth rate of technology, and interpret a pure noise innovation as an animal spirits shock, as it is associated with erroneous consumer optimism or pessimism.

We then develop the implications of each of the structural shocks for the endogenous variables of the model. The level and information shocks are associated with permanent movements in measures of real activity, while the animal spirits shock is associated with transitory increases in spending. Guided by the theoretical impulse responses of the model, we take up a more rigorous analysis of the meaning of consumer confidence innovations in Section IV.

¹⁴ Some who accept the notion that intangible news could be responsible for large movements in confidence might nevertheless be surprised at the volatility of responses to questions like E5Y. Our claim is not that all of the movements in measured confidence reflect genuine information, but rather that whatever relationship obtains between confidence and subsequent income or consumption is likely to reflect information. Our methodology does not unveil the meaning of innovations in measured confidence associated with words alone and not actions.

(a) Model

(i) Households

Households have standard preferences over consumption and leisure, live forever, and are identical. They consume a final consumption good, c , and supply labor, n , to intermediate goods producers. Each period, they choose consumption, labor supply, and holdings of a riskless one period bond so as to maximize expected discounted lifetime utility subject to a nominal budget constraint:

$$\max_{c_t, b_t, n_t} \sum_{t=0}^{\infty} \beta^t E_0 \left\{ \ln(c_t - \alpha c_{t-1}) - \frac{n_t^{1+1/\eta}}{1+1/\eta} \right\}$$

s.t.

$$p_t c_t + b_t \leq w_t n_t + (1 + i_{t-1}) b_{t-1} + \Pi_t$$

β is a subjective discount factor; α is the degree of internal habit persistence; η is the Frisch labor supply elasticity; p is the price of the final consumption good; w is the nominal wage; b is a riskless one period bond paying nominal interest i ; and Π denotes any lump sum profits or transfers households might receive.

The first order conditions characterizing the solution to the household's optimization problem are:

$$MU(c_t) = E_t \left(MU(c_{t+1}) (1 + i_t) \frac{p_t}{p_{t+1}} \right) \quad (1)$$

$$n_t^{1/\eta} = MU(c_t) \frac{w_t}{p_t} \quad (2)$$

Equation (1) is the intertemporal consumption Euler equation and equation (2) is the labor supply condition. The marginal utility of consumption depends positively on lagged and led consumption and negatively on current consumption:

$$MU(c_t) = \frac{1}{c_t - \alpha c_{t-1}} - \alpha \beta E_t \left(\frac{1}{c_{t+1} - \alpha c_t} \right)$$

(ii) *Final Goods*

The final good is a CES aggregate of a continuum of intermediate goods, indexed by j along the unit interval:

$$y_t = \left[\int_0^1 (y_{j,t})^{\frac{\xi-1}{\xi}} dj \right]^{\frac{\xi}{\xi-1}}$$

The parameter ξ has the interpretation as the price elasticity of demand for intermediate goods, and is assumed to be greater than unity. Similarly, the price index for final goods is given by:

$$p_t = \left[\int_0^1 p_{j,t}^{1-\xi} dj \right]^{\frac{1}{1-\xi}}$$

The model has neither capital nor a storage technology, so all final output must be consumed each period:¹⁵

$$y_t = c_t \tag{3}$$

(iii) *Intermediate Goods*

Intermediate goods are produced according to a linear production function:

$$y_{j,t} = A_t n_{j,t} \tag{4}$$

A denotes technology, which is common and freely available to all intermediate goods firms. It can be shown that profit maximization in the final goods sector implies a downward-sloping demand curve for each intermediate good:

$$y_{j,t} = \left(\frac{p_{j,t}}{p_t} \right)^{-\xi} y_t \tag{5}$$

We assume that intermediate goods firms are not freely able to adjust prices each period. In particular, following Calvo (1983), firms face a constant hazard of being able to adjust their price in any period equal to $1 - \theta$. Whenever a firm gets an opportunity to adjust its price, it solves the following maximization problem:

¹⁵ As is commonplace in the sticky price literature, we abstract from the presence of capital in the model. While the central lessons we draw are unaffected by this simplifying assumption, the presence of capital does matter in an essential way for certain aspects of the model. We address some of these issues in the next section.

$$p_t^* = \arg \max \sum_{t=0}^{\infty} \theta^t E_0 \frac{\beta^t MU(c_t)}{p_t} (p_{j,t} - mc_{j,t}) \left(\left(\frac{p_{j,t}}{p_t} \right)^{-\xi} y_t \right)$$

Above, p^* is the firm's optimal reset price and mc denotes nominal marginal cost. Marginal cost can be found in the firm's static labor demand condition:

$$w_t = mc_{j,t} A_t \quad (6)$$

The optimal reset price will be a present discounted value of expected nominal marginal costs:

$$p_t^* = \frac{\xi}{\xi - 1} \frac{\sum_{t=0}^{\infty} (\theta\beta)^t MU(c_t) p_t^{\xi-1} y_t mc_{j,t}}{\sum_{t=0}^{\infty} (\theta\beta)^t MU(c_t) p_t^{\xi-1} y_t} \quad (7)$$

Because all firms face the same wage and technology, expected marginal costs will be the same across firms, implying that all firms with the ability to update their price will choose the same reset price. The aggregate price level will thus evolve according to:

$$p_t = (\theta p_{t-1}^{1-\xi} + (1 - \theta) p_t^{*1-\xi})^{(1/(1-\xi))} \quad (8)$$

Log-linearizing these conditions about the zero inflation steady state and simplifying gives rise to the familiar New Keynesian Phillips Curve relating current inflation to real marginal cost and expected future inflation. See Galí and Gertler (1999) or Woodford (2003) for a more thorough discussion.

(iv) *Technology*

We assume that log technology ($a = \ln(A)$) obeys a random walk with drift:

$$a_t = g_{t-1} + a_{t-1} + u_t \quad (9)$$

The random variable u represents a *level shock* – a permanent and immediate innovation to the level of technology, while g is a drift term that is itself stochastic. We assume that g obeys a stationary autoregressive process:

$$g_t = (1 - \kappa) \bar{g} + \kappa g_{t-1} + e_t \quad (10)$$

Where $\kappa < 1$, \bar{g} denotes the steady state growth rate and e (which is assumed orthogonal to u) is a growth shock – a persistent but stationary innovation to the growth rate of technology, heralding periods of above or below average growth. We call e an *information shock* because it portends changes in future levels of technology. It is simply a smooth version of the “news

shocks” studied by Beaudry and Portier (2004) and Jaimovich and Rebelo (2006). Because of the assumed nominal rigidities in the model, there is an avenue here for output to expand upon the arrival of good news about the future and our model is not subject to the “bust” feature of neoclassical models in which agents receive advance signals about future technology.¹⁶

(v) *Perceptions and Animal Spirits*

While households observe the level of technology at each point in time, we assume that they never explicitly observe level shocks to technology, u , and observe only a noisy signal of growth rate shocks, e . The signal they receive is equal to:

$$s_t = e_t + v_t \quad (11)$$

v is a mean zero white noise disturbance uncorrelated with both growth and level shocks.

The setup described above implies that households imperfectly observe the drift term. We posit that they update their perceptions according to a simple linear filter:

$$g_t^p = \kappa(1 - \Omega_1)g_{t-1}^p + \kappa\Omega_1(a_t - a_{t-1}) + \Omega_2s_t \quad (12)$$

κ is the autoregressive coefficient from equation (10), and the coefficients Ω_1 and Ω_2 are functions of the variances of the shocks in the economy. In particular,

$$\Omega_1 = \frac{\sigma_v^2}{\sigma_v^2 + \sigma_u^2} \quad \text{and} \quad \Omega_2 = \frac{\sigma_e^2}{\sigma_e^2 + \sigma_v^2}$$

To see why these coefficients look the way they do, it is helpful to consider a couple of extreme cases. If $\sigma_v^2 = 0$ (i.e. there is no noise in the signal concerning the growth rate shock) then $\Omega_2 = 1$, $\Omega_1 = 0$, and the perceived drift term is equal to the truth at all times. If $\sigma_u^2 = 0$ (i.e. there are no shocks to the current level of technology), then $\Omega_1 = 1$. In this case agents will be uncertain about the growth rate between today and tomorrow due to the noise in the signal, but the realization of technology tomorrow will reveal perfectly to them today’s actual growth rate shock, so that there will be no endogenous persistence of a false signal for more than one period.

¹⁶ By “bust” feature we are referring to the tendency of neoclassical models to yield output *declines* in response to good news of the future (see Beaudry and Portier (2004) or Jaimovich and Rebelo (2006)). This is because in these models output is completely supply determined, and, in the standard framework, there is no explicitly dynamic dimension to the firm’s problem. As such, the wealth effect on the household side of the model usually induces a decrease in labor supply. Coupled with no change in labor demand, this results in a reduction of output in equilibrium. We do not have this problem because the price stickiness allows the “demand” effect of news about the future to work in the right direction. In particular, the increased desire to consume induces an (undesired) reduction in firm markups, which leads to an increase in labor demand, thus allowing employment and output to expand in anticipation of the realization of good news.

Intermediate cases are more interesting. As the variance of the noise term in the signal grows, Ω_2 becomes smaller and Ω_1 gets bigger – people will place little weight on a very noisy signal but will place a lot of weight on the realization of actual technology growth relative to their previous period’s perception in updating their current belief. As σ_u^2 gets bigger, Ω_1 becomes smaller, so household perceptions about the technology drift term will be more persistent. Intuitively, a very high variance of level technology shocks means that a realization of technology growth different from what was expected is less likely to mean that the original perception of the persistent growth rate was wrong, and more likely that there was simply an offsetting level shock.

While we assume that households observe the drift term in technology with imprecision, we allow firms to view both level and information shocks without noise. Although it seems both intuitive and realistic that firms have superior information relative to individuals, this setup is incompatible with the usual structure in which firms are owned by households. To avoid this complication, we can simply assume that management is separated from ownership, with managers risk neutral agents with the sole objective to maximize profits.

The disparate information to which households and firms are privy presents an additional complication. Even though households are unable to immediately differentiate between legitimate news about the drift term and pure noise, the equilibrium effects of noise and genuine information shocks on the endogenous variables of the model will be different, owing to the fact there is a shock to the supply side of the model in the case of a true growth rate shock, whereas there is not in response to a noise shock to the households’ signal. Therefore, the linear filter given by equation (11) is not the optimal filter when firms have better information than households. In particular, the optimal filter would include information revealed to households through wages, interest rates, and prices, whose equilibrium behavior would reveal to them the underlying nature of the signal. We simply assume this complication away. The filtering specification in (11), though not fully optimal, is both intuitive and simple, and household perceptions converge to the truth in the long run.¹⁷

¹⁷ Under a fully optimal filter, there would be no endogenous persistence of the noise shock on households’ perceptions of the drift term. This is because the general equilibrium behavior of the endogenous variables would immediately reveal to households the true nature of the signal. A more complicated version of this model which would preserve the endogenous persistence of these noise shocks under a fully optimal filter would introduce an additional shock, unobservable to households, into the model (e.g. a markup shock or a monetary policy shock).

We will interpret a noisy innovation to the households' signal of the drift term as an *animal spirits shock*. A positive v means that households erroneously believe that the future will be better. Given this belief, they will desire to consume more immediately. Because firms do not share this belief, there is no shock on the supply side of the model. In this way, our animal spirits shock is a pure demand shock, and is similar to the kind of shock studied in Lorenzoni (2008).¹⁸ The animal spirits shock will play a role in equilibrium nearly identical to a preference shock manifesting itself as an exogenous innovation in the consumption Euler equation. As such, one could give this disturbance an alternative, less structural interpretation as a taste or rate of time preference shock.¹⁹

(vi) *Monetary Policy Rule*

We close the model with a nominal interest rate rule. In particular, we postulate that the central bank sets nominal interest rates according to a partial adjustment mechanism where the interest rate in any period is equal to a convex combination of the lagged interest rate and the central bank's target rate. The target rate is adjusted in response to deviations of output growth and inflation from constant targets.

$$\dot{i}_t = \rho \dot{i}_{t-1} + (1 - \rho) \left(\phi_y (y_t - y_{t-1} - \Delta y^*) + \phi_\pi (\pi_t - \pi^*) \right) \quad (13)$$

The parameter ρ captures the degree of interest rate something. We abstract from the presence of purely monetary disturbances, so that we show no error term.

Our specification of the policy rule differs slightly from the ubiquitous Taylor rule (1993) in which the nominal rate is adjusted in response to inflation and the output gap. We prefer our specification for two reasons. First, the informational requirements imposed on the central bank are much more reasonable when assuming that it responds to output growth relative to its long

Even though the general equilibrium effects of a pure noise disturbance would differ from those of a true growth rate shock, households could attribute the different equilibrium movements of the endogenous variables to an unobserved shock elsewhere in the model, which would in turn allow their erroneous belief about the drift term to persist. Our specification here thus provides a sort of upper bound on the degree of persistence of erroneous news on household perceptions of the underlying growth rate.

¹⁸ Conceptually, the only fundamental difference between Lorenzoni's specification and ours is that in his paper agents receive a noisy signal about the *current* level of technology, whereas in our framework the noisy signal concerns *future* levels of technology. His specification gives animal spirits a better shot at inducing significant fluctuations, and we address this possibility in further detail below.

¹⁹ Some might prefer to think of an animal spirits shock this way in the first place. Irrational exuberance, for instance, could be interpreted as an emphasis on the enjoyment of current consumption at the implicit expense of future consumption.

term trend as opposed to its “potential”, which is itself time-varying. Secondly, a rule such as this is capable of matching certain features of the data which a standard Taylor rule is not. In particular, we know from ongoing work that information shocks about future productivity appear to be strongly deflationary (Sims (2008)). In general equilibrium, predictable increases in output and consumption must be associated with rising real interest rates. It is extremely difficult to simultaneously generate a large increase in real interest rates and a large disinflation under a standard Taylor rule.²⁰ A rule in which the bank reacts to output growth as opposed to the gap is capable of matching the data along this dimension, and is thus the one which we adopt here.

A policy rule similar to (13) is also not without precedent in the literature. In particular, a number of recent papers make use of very similar rules – for example, Coibion and Gorodnichenko (2007), Fernandez-Villaverde and Rubio-Ramirez (2007), and Ireland (2004). In particular, our exclusion of a theoretical output gap from the policy rule is consistent with Ireland’s (2004) finding that the coefficient on the gap in an estimated rule does not differ significantly from zero. Orphanides (2003) argues that a rule responding to output growth provides at least as good a description of actual US monetary policy over the last thirty years as does the more standard formulation in which the central bank responds to an output gap. We will discuss the implications of alternative monetary policy rules for our results below.

(b) Theoretical Impulse Responses to the Structural Shocks

We solve the model by log-linearizing the equations above about the non-stochastic balanced growth path. Solving the model requires picking values of the structural parameters. We assume the following: $\beta = 0.995$ (with the interpretation of the unit of time as one quarter, this corresponds to an annual discount rate of roughly two percent), $\alpha = 0.5$, $\eta = 1.0$, $\theta = 0.66$ (meaning that firms get to update their prices on average once every three quarters), $\rho = 0.75$, $\phi_y = 2.5$, $\phi_\pi = 4.5$, $\kappa = 0.85$, $\sigma_u = 1$, $\sigma_e = 0.125$, and $\sigma_v = 0.125$.²¹ Because of the assumed high degree of persistence to information shocks ($\kappa = 0.85$), it is necessary that the standard

²⁰ One can see the intuition with a simple approximation. The standard Taylor rule is: $i_t = \theta_1(y_t - y_t^*) + \theta_2\pi_t$. Assume that $y_t \approx y_t^*$ and $\pi_t \approx E_t\pi_{t+1}$ (both of these approximations will be good for a standard calibration of the parameters of the model). Then $r_t = i_t - E_t\pi_{t+1} \approx (\theta_2 - 1)\pi_t$. With $\theta_2 > 1$, an increase in the real rate must be associated with an increase in inflation. This approximation also works in a partial adjustment specification of the rule in which it is assumed that $i_t \approx i_{t-1}$.

²¹ This calibration of the policy rule parameters is consistent with the determinacy of a rational expectations equilibrium. See Woodford (2003).

deviation of shocks to the drift term be small relative to that of level shocks in order to generate data in which actual productivity growth is approximately white noise, which appears to be the case in the US. We calibrate the variance of the animal spirits shock so that agents place a fairly high weight (in this case $\Omega_2 = 1/2$) on the signal in updating their perceptions of the drift term. The choice of the habit persistence term is similar to the estimates in Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2003). The calibration of the labor supply elasticity is in the middle of the range of estimates from micro studies (which are typically small) and those in the business cycle literature (which are usually much higher than unity), and is equal to the central point estimate in Kimball and Shapiro (2003). Our calibration of the parameters of the monetary policy reaction function is in line with empirical estimates from similarly specified rules (Coibion and Gorodnichenko (2007), Ireland (2004), and Paciello (2008)).

Figure 7 shows the theoretical responses of output, technology, hours, inflation, and real interest rates to the level and information shocks. The sizes of the shocks are chosen so that each leads to an ultimate increase in technology of one percent. For output, technology, and hours, the figures show the percentage response relative to the initial non-stochastic steady state. For inflation and the real interest rate, the figures show the annualized percentage point response (for example, a response of inflation of -0.2 means that inflation falls from, say, 4.0 percent to 3.8 percent at an annualized rate, and similarly for the real interest rate).

By construction, the level shock leads to an immediate jump in technology that is expected to remain forever at the new higher level, whereas the information shock is orthogonal to current technology but portends a sustained period of smooth growth. In response to both the level and information shocks, output jumps on impact and is expected to rise towards its new steady state value. Quite naturally, the impact jump in output is smaller for the information shock than for the level shock. Relative to the perfect information version of the model (where the variance of the animal spirits shock is zero), output overshoots in response to a level technology shock and undershoots in response to an information shock. The intuition for these effects is clear. When agents receive a signal that productivity growth will be higher, they place some weight on the possibility that the signal is purely noise, and so they react less than if they knew the shock with certainty. Likewise, when the level of technology jumps up unexpectedly, households place some weight on the possibility that there was an unseen growth rate shock at some point in the past that was buried in noise. They therefore place some weight on the

possibility that productivity growth will be higher in the near future, and so react more than if they knew that it were only a one time level shock.

Employment rises on impact in response to the growth rate shock, while it falls on impact following the level shock. The fall in hours in response to the level shock is the well-known “contractionary technology shock” due to the undesired increase in firm markups following immediate technological improvement.²² The response of hours to any shock in the model is constrained to be transitory because household preferences are in the class of preferences described by King, Plosser, and Rebelo (1988) consistent with balanced growth. Both kinds of shocks are associated with a rising real interest rate, which is to be expected, as both shocks make the future plentiful relative to the present.

Both the level and information shocks are disinflationary in the model, with the magnitude of the disinflation smaller but more persistent for the information shock than the level shock. A different behavior of prices could be rationalized under a different policy rule. In a standard Taylor rule where rates are adjusted in response to inflation and the deviation of output from the theoretical gap, for example, inflation would be almost completely stabilized in response to both shocks and would in fact rise slightly in response to the information shock. An exogenous time path for the money supply with a quantity type money demand equation, on the other hand, would produce a behavior of prices quite similar to what is shown here. As such, there is no robust implication of the model for prices in response to the two fundamental shocks.

Figure 8 shows the responses to the animal spirits shock in the model. The size of shock is chosen so that it is the same as the information shock (i.e. both shocks raise the signal by an amount prognostic of an ultimate increase in the level of technology of one percent). By construction, the shock never has any effect on the actual level of technology. The animal spirits shock is differentiated from the level or information shocks in that it is associated with a transitory response of output and rising prices. All three kinds of shocks raise the real interest rate.

For the given calibration of the parameters of the model, the impact of animal spirits on output is small (the maximal output effect is roughly 0.06 percent), though the effects are fairly persistent. The reason for the weak response of output is straightforward – this response is essentially the “aggregate demand” effect of an information shock (i.e. the increase in output

²² See, for example, Basu, Fernald, and Kimball (2006).

coming from the consumer side of the model following a good signal). The response to the true information shock combines this demand effect with the “aggregate supply” effect, which is positive given the forward-looking nature of the Phillips Curve. The impact effect of animal spirits is thus bounded from above by the impact effect of a true information shock, which is itself modest.

The animal spirits shock also leads to a fairly significant increase in inflation. In the New Keynesian model, inflation is equal to a present discounted value of real marginal costs. Real marginal costs should be weakly higher at all horizons following an animal spirits shock, and inflation should thus rise. The intuition for this effect is straightforward. Real marginal cost in the model is equal to the log difference between real wages and technology. Following a positive animal spirits disturbance, households feel wealthier and thus demand a higher real wage for a given level of employment. Since there is never any effect on actual or perceived technology on the firm side of the model, the wealth effect on the household side dictates that real marginal costs are always weakly higher following a noise innovation to the household signal. As such, the model has the implication that animal spirits shocks (as specified) are inflationary.

There are alternative calibrations of the model’s parameters which yield more significant responses to the animal spirits shock. In particular, the impact effects of animal spirits are larger for very low values of κ (the parameter governing the persistence of information shocks). The intuition for this is straightforward. When κ is small, most of the expected improvement in productivity occurs sooner as opposed to later, resulting in a larger innovation to perceived permanent income, which in turn leads to a larger effect on overall aggregate demand. We will allow the data to inform us on the value of this parameter in the next section.

IV. Expanded Empirical Analysis and a Structural Model of Consumer Confidence

(a) Reduced Form Empirical VAR

To begin to assess the relative importance of the structural shocks of the theoretical model in the determination of consumer confidence innovations, we first estimate a VAR with E5Y, annualized CPI inflation, the three month Treasury Bill rate, the BLS measure of aggregate per capita hours in the non-farm business sector, real non-durables plus services consumption per

capita, and real GDP per capita.²³ Aside from the fact that we include separate measures of output and consumption, the variables in this empirical VAR coincide with those in the theoretical model of Section III.²⁴ The real interest rate is implicitly defined as the nominal three month Treasury Bill rate less the VAR forecast of one quarter ahead inflation. As in the empirical VARs of Section II, the data are quarterly from 1960:01 – 2007:03, we choose a lag order of four, and we impose cointegration between output and consumption.²⁵ Labor hours, E5Y, the interest rate, and inflation enter in the VAR in levels.²⁶

Figure 9 presents the impulse responses of the variables in the empirical VAR to an E5Y innovation (ordered first in a block recursive system). As before, the dashed lines represent 90 percent confidence bands from a bias-corrected bootstrap procedure. Consumption and output both jump slightly on impact in response to an E5Y innovation, but thereafter continually rise, with no tendency to attenuate. The point estimates suggest that a one standard deviation to E5Y is prognostic of consumption and output that are higher in the long run by roughly 0.67 percent. Even at a horizon of forty quarters, these responses are statistically different from zero at better than the 90 percent level. The E5Y innovation is associated with transitory and significant increases in both real interest rates and hours of work, with both responses following hump-shaped patterns. Inflation falls by roughly one quarter of a percentage point on impact and is persistently below its initial value for a number of quarters. Though there is no significant impact effect of E5Y on measured labor productivity (imputed within the VAR as the output

²³ Earlier versions of this paper reported results with the civilian unemployment rate in place of hours, and expressed some concern about reverting impulse responses of consumption and income to E5Y innovations. This reversion, which apparently depends on a marginally significant coefficient implying that higher unemployment is associated with higher consumer confidence, largely disappears as one increases the lag length. Nevertheless, that reversion is nonetheless statistically indistinguishable from a permanent response, and, at any rate, the responses of output and consumption to E5Y innovations with unemployment in the VAR (even with very low lag lengths) are far too persistent to be taken as positive evidence in favor of an important animal spirits component.

²⁴ The theoretical model without capital does not distinguish between consumption and output. We include consumption as well as GDP in the empirical VARs of this section so as to facilitate comparison with the results of Section II. The empirical results are unaffected by using either consumption or output in isolation.

²⁵ As before, we impose that the cointegrating vector between consumption and output is [1,-1]. The imposed lag order of four is somewhat higher than the choices of a variety of widely accepted information criteria (which, on average, favor two lags). Alternative lag structures produce nearly identical results. A VAR with all variables entering in levels yields nearly identical impulse responses.

²⁶ There is a large debate over whether labor hours are $I(1)$ or $I(0)$ (see, for example, Christiano, Eichenbaum, and Vigfusson (2004)). Our results are qualitatively similar whether hours enter the VAR in levels, first differences, or as deviations from a trend.

response less the hours response), an E5Y innovation is prognostic of a permanent increase in productivity of more than two thirds of a percent, with the long run response statistically different from zero at horizons in excess of forty quarters.

A cursory comparison of the responses in Figures 7 and 9 reveals that the empirical responses to a confidence innovation look similar to the theoretical responses to what we have deemed an information shock in our model. In particular, a positive innovation to E5Y is associated with a prolonged and permanent increase in real activity, a transitory rise in both real rates and hours of work, and a strong and persistent disinflation. These are roughly the qualitative predictions of the model in response to a favorable information shock. It therefore seems natural to associate innovations in consumer confidence with information shocks – in particular, persistent shocks to productivity growth. Does this observation mean that there is no role for animal spirits, and no noise in measured confidence? In the next subsection we specify and estimate a variance components model of confidence innovations that allows us to address these questions.

(b) A Model of Consumer Confidence

In the context of the theoretical model of the previous section, we assume that a measure of consumer confidence follows a stationary autoregressive process, with its innovation a linear combination of the unexpected change in the current state of the economy and the signal concerning the persistent growth term:

$$CC_t = \delta CC_{t-1} + \lambda_1(a_t - E_{t-1}a_t) + \lambda_2 s_t \quad (14)$$

The conditional relationships between confidence innovations and the other variables of the model will depend both on the λ s and the variances of the structural shocks. For instance, $\lambda_1 > 0$ and $\lambda_2 = 0$ would mean that that the confidence innovation is purely a reflection of the change in the current state of the economy, while $\lambda_2 > 0$ would mean that innovations to confidence at least partially reflect signals households receive about the future. If λ_2 is relatively large and the signal, s , is not very noisy, the relationships between confidence innovations and macroeconomic variables in the model will look similar to the theoretical responses to an information shock. On the other hand, if λ_2 is large and the signal is quite noisy, then confidence innovations may generate patterns similar to the theoretical responses to an animal spirits shock.

As written, consumer confidence responds only to structural shocks in the economy. One should not take this description too literally. No one would wish to maintain that all of the variation in the observed confidence data reflects genuine information (or even changes in beliefs, however formed) – there is always sampling error, misunderstandings on the part of respondents, etc. As such, any realistic specification of confidence should include some kind of measurement error. We abstract from the presence of pure measurement noise here simply because, using our empirical method detailed below, it is not possible to separately identify its empirical properties.

We estimate the parameters of the confidence equation using a modified version of the simulated method of moments. In particular, we would like to know what parameter configuration is most likely to generate data yielding impulse responses similar to what we see in the actual data. As such, our SMM estimator tries to match impulse responses to confidence innovations from simulated data from the model to the impulse responses from the actual data. This approach is similar to that in Christiano, Eichenbaum, and Evans (2005).²⁷

As our main focus is only on the parameters directly influencing consumer confidence, we first calibrate many of the other parameters of the model. In particular, we set the preference parameters, the parameter governing price stickiness, and the policy coefficients as in the calibration of the previous section, and we set the autoregressive coefficient in the confidence equation at 0.8.²⁸ We normalize the variance of level technology shocks to be unity. The remaining parameters to be estimated are given by the vector $\Theta = [\lambda_1 \ \lambda_2 \ \kappa \ \sigma_e \ \sigma_v]'$.

For a given guess of the parameter vector Θ , we simulate a data set of length τT , where T is the length of the actual data set and $\tau = 10$ (the shocks are drawn from normal distributions). After discarding the first 100 observations from the simulated data set (so as to limit the influence of arbitrary starting values), we estimate a five variable VAR similar to the one in subsection (a) with simulated confidence, output growth, inflation, hours of work, and the interest rate and compute impulse responses to the confidence innovation (ordered first in the

²⁷ Our approach is similar to that of CEE in that we are choosing parameter values so as to match impulse responses as opposed to unconditional moments. CEE differ slightly from us in that their objective function is to match theoretical impulse responses from their model to those in the data, whereas we match impulse responses from VARs estimated on simulated data from our model.

²⁸ We experimented with many different values for the other parameters of the model. Alternative calibrations of these parameters have little noticeable impact on our results. $\delta = 0.8$ is approximately the estimated autoregressive coefficient for ESY in the data.

VAR).²⁹ Letting $\mathbf{m}(\Theta)$ denote the stacked vector of impulse responses for the given guess of Θ and \mathbf{m}^* the corresponding stacked vector of impulse responses from the data, we iterate on our guess of Θ so as to minimize the following:³⁰

$$\Theta^* = \operatorname{argmin} (\mathbf{m}(\Theta) - \mathbf{m}^*)' \mathbf{W} (\mathbf{m}(\Theta) - \mathbf{m}^*)$$

As in Christiano, Eichenbaum, and Evans (2005), we set $\mathbf{W} = \mathbf{V}^{-1}$, where \mathbf{V} is a diagonal matrix with elements equal to the variances of the impulse responses from the data. The optimal parameter vector is then that which minimizes the weighted sum of squared deviations between the estimated impulse responses on model simulated data and the corresponding responses from the actual data. This choice of weighting matrix places more weight on those responses which are most precisely estimated in the data.

The estimates of the parameters of interest and corresponding confidence bands are in Table 2. These coefficients are of the expected signs, and with the exception of the coefficient on the unexpected change in the current state of the economy, are all different from zero at conventional levels of significance. In order to provide some interpretation to these quantitative estimates, the total variance of the structural confidence innovation is seen to be:

$$\operatorname{var}(e_{cc}) = \lambda_1^2 \operatorname{var}(a_t - E_{t-1} a_t) + \lambda_2^2 (\sigma_e^2 + \sigma_v^2)$$

Given the low estimate of λ_1 , we see that the unexpected change in the current state of the economy evidently accounts for less than one half of one percent of the innovation variance of confidence in the model. Information and animal spirits shocks each account for roughly one half of the innovation variance in measured consumer confidence, with the noise disturbance mattering slightly more.

Figure 10 presents the average impulse responses of output, inflation, hours, the real interest rate, and confidence to a confidence innovation from simulated data using these

²⁹ As laid out above, our model only has three structural shocks, meaning that any combination of three or more simulated series from the model would be perfectly collinear – i.e. the model suffers from “stochastic singularity”. So as to be able to estimate a VAR with more than three variables, we introduce three additional shocks into the model for the purposes of estimation. In particular, we introduce a shock to the monetary policy rule, a “cost-push” shock in the Phillips Curve, and a preference shock in the Euler equation. The model E5Y innovation is unrelated to these disturbances, so they should not (in large enough samples) impact the conditional correlations between confidence and the other variables. Rather, the only role of these shocks is to ensure that five variables in the model are not perfectly collinear. The estimated VAR here is identical to the empirical VAR of subsection (a), except that the model does not differentiate between income and consumption.

³⁰ The impulse responses making up our objective function include the impact effect on confidence itself, the impulse response of the level of output over ten years, and the responses of hours, interest rates, and inflation over five years. We also include the autocorrelation of productivity growth in the objective function. The observed first order correlation of measured labor productivity growth in our sample is roughly 0.045.

parameter estimates. In particular, we simulated 2000 sets of data with 200 observations each based on these parameters, with the structural shocks drawn from a normal distribution. For each simulated data set we estimated a five variable VAR with confidence, output growth, hours, interest rates, and inflation, with four lags of each variable. For output, interest rates, and inflation, these responses look similar to those from the empirical VAR depicted in Figure 9. In the simulations, the average E5Y innovation is associated with a small impact effect on output followed by a sustained period of growth, a significant and persistent disinflation, and persistently high real interest rates.

The dimension along which the model is least successful in matching the empirical impulse responses is in the response of hours. In the data confidence innovations are associated with a small but reasonably persistent increase in hours. While the model produces data roughly matching the impact response of hours to an E5Y innovation, it fails to match the persistence. After the small positive impact effect and a few quarters of being above trend, the response of hours is slightly negative for a number of quarters. The intuition for this response is reasonably straightforward. A significant portion of the E5Y innovation is accounted for by true information shocks, which begin to exert a contractionary effect on employment in the model once the technological improvement starts to take hold.

Two seemingly contradictory conclusions emerge from our structural estimation results. On the one hand, animal spirits disturbances seem to account for an important portion of the structural confidence innovation. On the other hand, the model responses to a confidence innovation look very much like the theoretical responses to an information shock, and nothing at all like the responses to an animal spirits disturbance. The resolution of this apparent contradiction is that information shocks have implications for the other variables of the model which simply dwarf those of animal spirits shocks. As such, the conditional correlations between confidence innovations and the other variables of the model are dominated by the information shock, even though animal spirits shocks account for a significant component of the structural confidence innovation.

Nevertheless, our results do allow us to reject the hypothesis that animal spirits account for a significant portion of the observed relationship between consumer confidence and macroeconomic variables. Given the estimated persistence of information shocks, the implications of an animal spirits shock for the variables of the model are very small (see the

discussion on p. 21 or the theoretical impulse responses in Figure 8). As such, the animal spirits shock is difficult to differentiate from pure measurement noise in data generated from the model. The wide confidence bands on the estimate of the standard deviation of animal spirits shocks shown above confirm this and seem to suggest that this parameter is probably poorly identified. Forcing the variance of animal spirits shocks to zero and re-estimating the model leads to little noticeable difference in the estimates of the other parameters or in the overall fit of the model. Eliminating information shocks from the model and re-estimating, however, leads to a much poorer overall fit.

While confidence innovations evidently reflect both information and animal spirits (which are in practice difficult to differentiate from pure measurement error), our results suggest that the relationships between confidence and macroeconomic variables are largely driven by information about future fundamentals. The impulse responses in Figure 11 make this point perhaps even more clear. These responses are from a bivariate VAR featuring the growth rate of a utilization-corrected measure of aggregate total factor productivity (TFP) and E5Y.³¹ This exercise is similar to the stock price-TFP VARs in Beaudry and Portier (2006). Confidence is ordered second, so that the structuralized innovation in E5Y is contemporaneously orthogonal to TFP. Two observations stand out. First, the confidence innovation orthogonal to TFP predicts a permanent increase in TFP of roughly 0.7 percent, with this effect highly significant even at very long horizons. In quantitative terms, this long run response of TFP is about the same magnitude as TFP's response to its own innovation, which looks very much like a pure random walk. The TFP response to E5Y is both smooth and prolonged, and looks similar to the theoretical response to a growth rate shock discussed in the previous section. Secondly, consumer confidence does not respond significantly (either statistically or economically) to the TFP innovation. Both of these findings corroborate our estimation results above, which did not make explicit use of any TFP measure. In particular, there appear to be information shocks about the future which may account for significant component of long run productivity. These information shocks appear to be reflected in consumer confidence innovations, which are evidently unrelated to contemporaneous productivity shocks.

³¹ We are grateful to John Fernald for providing us with this measure. The responses show the level response of TFP, which is simply the cumulated growth rate response.

(c) Information Shocks and Business Cycle Fluctuations

We close the body of the paper with a brief comment on the differences and similarities between our model and those of Beaudry and Portier (2004) and Jaimovich and Rebelo (2006). One difference is that their models are laid out in neoclassical settings, whereas we assume the presence of nominal rigidities. The nominal rigidities turn out to be important, for they introduce an explicitly forward-looking dimension into the firm's profit maximization problem which provides an avenue by which output can expand in response to good news about future productivity. The other primary difference is that we abstract from the presence of capital. As both of the above papers make clear, it is difficult to generate a simultaneous increase in consumption, output, and investment in response to favorable news about future productivity. In a pure neoclassical model, there is no explicitly dynamic dimension to the firm's problem. As such, the wealth effect of higher future productivity usually induces a reduction in labor supply and an increase in consumption, the combined implications of which are reduced output and investment.

As our goal has been to study the meaning of surprise movements in consumer confidence – and not the requisite model structures needed to deliver broad-based co-movement following information shocks – we have made the deliberate decision to sidestep the issue of co-movement by abstracting from endogenous capital accumulation altogether. While not without loss of generality, this decision has enabled us to elucidate the apparent necessity of shocks to future fundamentals orthogonal to the present in generating confidence data consistent with what we see in the world. That we have found an apparent empirical counterpart to the kinds of information shocks studied by other authors suggests that further study of more sophisticated models with these kinds of shocks – as well as the model features which will produce positive co-movement – is likely to be a fruitful avenue for future research.

V. Conclusion

While many in the popular press and business community regard measures of consumer confidence as essential in understanding the evolution of the aggregate economy, economists have devoted little attention to the economic interpretation of variation in measured confidence. Most of the scant academic literature focuses on the extent to which confidence measures help to improve forecasts of spending and output over relatively short horizons. While related to that

line of research, this paper goes further in attempting to ascertain the underlying meaning of surprise movements in confidence.

We began our inquiry with an analysis of simple consumption-income VARs augmented with forward-looking measures of confidence. As noted by Cochrane (1994a), innovations to consumption are powerful predictors of subsequent movements in income. We demonstrated that measures of consumer confidence play a role similar to that of consumption innovations in that they foretell important movements in future output. Even after orthogonalization with respect to consumption, confidence innovations remain prognostic of significant movements in output and spending, especially at longer horizons.

We then turned more formally to the question of what economic concept underlies surprise movements in confidence, beginning with two polar hypotheses. The first – which we deemed “animal spirits” – posits that surprise movements in measured confidence proxy for exogenous changes in sentiment, which in turn have causal effects on aggregate demand. Such an interpretation of confidence was given by Blanchard (1993) in a paper on the causes of the 1990-1991 recession. The second hypothesis – the “information view” – supposes that there exists no causal relationship from confidence to economic activity, but rather that measured confidence reflects aggregated information individuals possess regarding present and future economic fundamentals.

We developed a New Keynesian model incorporating both animal spirits and fundamental shocks. The animal spirits shock is a manifestation of overly optimistic or pessimistic perceptions on the part of households, and leads them to desire more or less consumption than is optimal under perfect information. The two fundamental shocks in the model are a current level shock and an anticipated growth rate shock to productivity. In general equilibrium, the animal spirits disturbance plays the role of an aggregate demand shock – it is associated with transitory movements in spending and with higher inflation. Both fundamental shocks, on the other hand, are likely to be disinflationary and are associated with movements in spending that are not subsequently reversed. The information shock is distinguished by a small initial response of output followed by a prolonged period of growth.

We estimated an empirical VAR including the variables in the model as well as a measure of consumer confidence. Income and consumption appear to respond permanently to a confidence innovation, with a positive innovation to confidence associated with income,

consumption, and labor productivity that are appreciably higher in the long run. The implications of a confidence innovation for output, spending, and productivity are much larger at longer horizons than at shorter ones, and positive confidence innovations are associated with a strong and persistent disinflation. In light of the theoretical model, the impulse responses from the empirical VAR provide essentially no support for animal spirits and point strongly to the information interpretation of confidence.

After positing a structural equation for consumer confidence, we then turned to a more formal estimation of the parameters of the model. We can resoundingly reject the hypothesis that animal spirits shocks (as specified in this paper) can account for the bulk of the relationships between consumer confidence and macroeconomic variables. If ever one hoped to find empirical support for animal spirits like shocks, surely it would be found in survey responses of seemingly naïve consumers. That we are unable to find compelling evidence in support of the animal spirits hypothesis thus casts expectations-driven theories of demand shocks such as Lorenzoni (2008) into serious doubt. On the other hand, we do find convincing evidence in favor of the information interpretation of consumer confidence. The implications of confidence innovations for output and spending at short horizons are far too small for confidence to be primarily a reflection of changes in current fundamentals, yet the longer horizon implications are far too large and significant for confidence innovations to not be conveying information about fundamentals. Putting the two together, it would appear as though confidence innovations are likely conveying information about future fundamentals, and in particular long run productivity. A bivariate TFP-confidence VAR seems to lend credence to this conclusion.

A recent line of research studies the extent to which news about future fundamentals can drive the business cycle (Beaudry and Portier (2004, 2006) and Jaimovich and Rebelo (2006)). Our results provide empirical support for the notion that agents do receive advance signals about future fundamentals, but they do not yet indicate that such information shocks play a pivotal role in short run fluctuations. Our ongoing research builds on the results of this paper and further addresses the business cycle implications of information shocks.

References

- Basu, Susanto, John Fernald, and Miles Kimball. "Are Technology Improvements Contractionary?" *American Economic Review*, December 2006, 1418-1448.
- Beaudry, Paul and Franck Portier. "An Exploration into Pigou's Theory of Cycles." *Journal of Monetary Economics*, September 2004, 1183-1216.
- Beaudry, Paul and Franck Portier. "Stock Prices, News, and Economic Fluctuations." *American Economic Review*, September 2006, 1293-1307.
- Blanchard, Olivier J. "Consumption and the Recession of 1990-1991." *American Economic Review*, May 1993, 270-274.
- Calvo, Guillermo. "Staggered Prices in a Utility Maximizing Framework." *Journal of Monetary Economics*, September 1983, 383-398.
- Carroll, Christopher D. "Macroeconomic Expectations of Households and Professional Forecasters." *Quarterly Journal of Economics*, February 2003, 269-298.
- Carroll, Christopher D., Jeffrey C. Fuhrer and David W. Wilcox. "Does Consumer Sentiment Forecast Household Spending? If so, Why?" *American Economic Review*, December 1994, 1397-1408.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, February 2005, 1-45.
- Christiano, Lawrence J., Martin Eichenbaum, and Robert Vigfussion. "The Responses of Hours to a Technology Shock: Evidence Based on Direct Measures of Technology." *Journal of the European Economic Association*, May 2004, 381-395.
- Cochrane, John H. "Permanent and Transitory Components of GNP and Stock Prices." *Quarterly Journal of Economics*, February 1994a, 241-266.
- Cochrane, John H. "Shocks." *Carnegie-Rochester Conference Series on Public Policy*, December 1994b, 295-364.
- Coibion, Olivier and Yuri Gorodnichenko. "Strategic Interaction Among Price-Setters in an Estimated DSGE Model." Working paper, 2007.
- Croushore, Dean. "Do Consumer Confidence Indexes Help Forecast Consumer Spending in Real Time?" *North American Journal of Economics and Finance*, December 2005, 435-450.

- Dominitz, Jeff and Charles Mankiw. "How Should We Measure Consumer Confidence?" *Journal of Economic Perspectives*, Spring 2004, 51-66.
- Farmer, Roger. *The Macroeconomics of Self-Fulfilling Prophecies*. Cambridge, MA: MIT Press. 1999.
- Fernandez-Villaverde, Jesus, and Juan F. Rubio-Ramirez. "How Structural are Structural Parameters?" Forthcoming, *NBER Macroeconomics Annual 2007*.
- Gali, Jordi. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, March 1999, 249-271.
- Gali, Jordi and Mark Gertler. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics*, October 1999, 195-222.
- Gorodnichenko, Yuri and Matthew D. Shapiro. "Monetary Policy When Potential Output is Uncertain: Understanding the Growth Gamble of the 1990s." *Journal of Monetary Economics*, May 2007, 1132-1162.
- Hall, Robert E. "Macro Theory and the Recession of 1990-1991." *American Economic Review*, May 1993, 275-279.
- Hall, Robert E. "Stochastic Implications of the Life Cycle Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy*, December 1978, 971-987.
- Hall, Robert E. "The Role of Consumption in Economic Fluctuations." *The American Business Cycle: Continuity and Change*, edited by Robert J. Gordon. Chicago: UCP, 1986, pp. 237-255 and 265-266.
- Howrey, E. Phillip. "The Predictive Power of the Index of Consumer Sentiment." *Brookings Papers on Economic Activity*, January 2001, 175-207.
- Ireland, Peter N. "Technology Shocks and the New Keynesian Model." *Review of Economics and Statistics*, April 2004, 923-936.
- Jaimovich, Nir and Sergio Rebelo. "Can News About the Future Drive the Business Cycle?" Stanford University Working Paper, August 2006.
- Kilian, Lutz. "Small Sample Confidence Intervals for Impulse Response Functions." *Review of Economics and Statistics*, May 1998, 218-230.
- Kimball, Miles S. and Matthew D. Shapiro. "Labor Supply: Are the Income and Substitution Effects Both Large or Both Small?" University of Michigan Working Paper, May 2003.

- King, Robert G., Charles I. Plosser, and Sergio Rebelo. "Production, Growth, and Business Cycles I: The Basic Neoclassical Model." *Journal of Monetary Economics*, 21 (2-3), 195-232.
- Leeper, Eric M. "Consumer Attitudes: King for a Day." *Federal Reserve Bank of Atlanta Economic Review*, July-August 1992, 1-15.
- Lorenzoni, Guido. "A Theory of Demand Shocks." MIT Working Paper, January 2008.
- Lucas, Robert E. "Asset Prices in an Exchange Economy." *Econometrica*, November 1978, 1429-1445.
- Ludvigson, Sidney. "Consumer Confidence and Consumer Spending." *Journal of Economic Perspectives*, Spring 2004, 29-50.
- Matuszaka, John G. and Argia M. Sbordonne. "Consumer Confidence and Economic Fluctuations." *Economic Inquiry*, April 1995, 296-318.
- Mehra, Yash P and Elliot W. Martin. "Why Does Consumer Sentiment Predict Household Spending?" *Federal Reserve Bank of Richmond Economic Quarterly*, Fall 2003, 51-67.
- Mishkin, Frederic S. "Consumer Sentiment and Spending on Durable Goods." *Brookings Papers on Economic Activity*, 1, 1978, 217-231.
- Orphanides, Athanasios. "Historical Monetary Policy Analysis and the Taylor Rule." *Journal of Monetary Economics*, July 2003, 983-1022.
- Paciello, Luigi. "The Response of Prices to Technology and Monetary Policy Shocks Under Rational Inattention." Northwestern University working paper, February 2008.
- Ruge-Murcia, Francisco J. "Methods to Estimate Dynamic Stochastic General Equilibrium Models." *Journal of Economic Dynamics and Control*, August 2007, 2599-2636.
- Sims, Eric R. "Expectations Driven Business Cycles: An Empirical Evaluation." Working paper, University of Michigan, 2008.
- Smets, Rafael and Frank Wouters. "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association*, September 2003, 1123-1175.
- Souleles, Nicholas S. "Expectations, Heterogeneous Forecast Errors, and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys." *Journal of Money, Credit, and Banking*, February 2004, 39-72.
- Woodford, Michael. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press, 2003.

Consumer Confidence Data Appendix

Questions:

E5Y: Looking ahead, which would you say is more likely – that in the country as a whole we'll have continuous good times during the next five years, or that we'll have periods of widespread unemployment or depression, or what?

E12M: Now turning to business conditions in the country as a whole – do you think that during the next twelve months we'll have good times financially or bad times or what?

PFE: Now looking ahead – do you think that a year from now you (and your family living there) will be better off financially, worse off, or just about the same as now?

News Heard: During the last few months, have you heard of any favorable or unfavorable changes in business conditions?

Follow up news question: What did you hear?

Answer Choices and Variable Construction: For most questions (including E5Y, E12M, and PFE), individuals are given three answer choices that amount to “favorable”, “neutral” or “don’t know”, and “unfavorable”. The “relative score” – the variable we use in this paper – is then constructed as the percentage giving a favorable response less the percentage giving an unfavorable response plus one hundred.

- Thus, a relative score of 100 indicates that an equal number of people gave a favorable response as an unfavorable response. If 30 percent of respondents give a favorable response and 20 percent given an unfavorable response, with the remaining 50 percent either “neutral” or “don’t know”, then the relative score will be 110 (i.e. $30 - 20 + 100$).
- If, out of 100 people, 1 person switches from an unfavorable response to a neutral response, the index score will go up by 1. If that person switches from unfavorable to favorable, the index score goes up by 2. If someone leaves the state of “neutral” to either “favorable” or “unfavorable” the index score moves up or down by 1.

The Index of Consumer Expectations (ICE) is constructed based on the relative scores for PFE, E12M, and E5Y as follows:

$$ICE = \frac{PFE + E12M + E5Y}{4.1134} + 2.0$$

The Index of Consumer Sentiment (ICS) is constructed based on the relative scores for the PFE, E12M, and E5Y, plus two other questions. The first we’ll call PFP and is similar to PFE, except that it asks respondents to make a comparison of their current financial situation relative to one year ago. The second we’ll call DUR and it asks respondents whether or not it is currently a good time to buy “large household items” (i.e. durable goods). The ICS is constructed as:

$$ICS = \frac{E12M + E5Y + PFE + DUR + PFP}{6.7558} + 2.0$$

For more, see: <http://www.sca.isr.umich.edu/documents.php?c=i>

Descriptive Statistics:

	E5Y	E12M	PFE	ICE	ICS
Mean	94.54545	114.7219	123.7861	82.69626	88.11551
Median	96.00000	118.0000	125.0000	84.50000	91.30000
Maximum	136.0000	168.0000	141.0000	106.0000	110.1000
Minimum	46.00000	35.00000	92.00000	47.60000	54.40000
Std. Dev.	20.04972	30.89033	9.645726	13.69135	11.55991

Correlations Among Confidence Variables:

(a) Levels

Correlation Matrix

	E5Y	E12M	PFE	ICE	ICS
E5Y	1.000000	0.916905	0.849950	0.970639	0.926224
E12M	0.916905	1.000000	0.783679	0.974327	0.909826
PFE	0.849950	0.783679	1.000000	0.875248	0.877126
ICE	0.970639	0.974327	0.875248	1.000000	0.955998
ICS	0.926224	0.909826	0.877126	0.955998	1.000000

(b) First differences

Correlation Matrix

	D(E5Y)	D(E12M)	D(PFE)	D(ICE)	D(ICS)
D(E5Y)	1.000000	0.807549	0.487918	0.890646	0.809159
D(E12M)	0.807549	1.000000	0.582231	0.960898	0.914735
D(PFE)	0.487918	0.582231	1.000000	0.693076	0.710045
D(ICE)	0.890646	0.960898	0.693076	1.000000	0.938775
D(ICS)	0.809159	0.914735	0.710045	0.938775	1.000000

Table 1
Regressions of Confidence Innovations on News Heard Categories

News Heard Category	Coefficient		
Favorable Employment	0.248 ^{**} (0.13)	0.113 (0.13)	0.140 (0.13)
Favorable Prices	1.001 ^{**} (0.51)	0.889 [*] (0.51)	1.005 [*] (0.58)
Unfavorable Employment	-0.064 (0.05)	-0.071 (0.05)	0.035 (0.06)
Unfavorable Prices	-0.363 ^{***} (0.13)	-0.342 ^{***} (0.13)	-0.312 ^{***} (0.15)
Favorable Stocks		0.915 ^{**} (0.38)	0.845 ^{**} (0.38)
Unfavorable Stocks		-0.235 (0.16)	-0.259 (0.17)
Favorable Government			0.342 (0.53)
Unfavorable Government			-0.604 ^{**} (0.24)
Favorable Credit			-0.342 (0.27)
Unfavorable Credit			0.124 (0.19)
Energy Crisis			-0.393 [*] (0.22)
Adj. R^2	0.10	0.12	0.15

The above are coefficient estimates from a regression of the reduced form innovation in E5Y obtained from the three variable system described in Section II on the percentage of respondents reporting having heard either favorable or unfavorable news concerning employment, prices, or stock prices. The sample period is 1961:1 – 2007:3. OLS standard errors are in parentheses.

Table 2
Structural Parameter Estimates

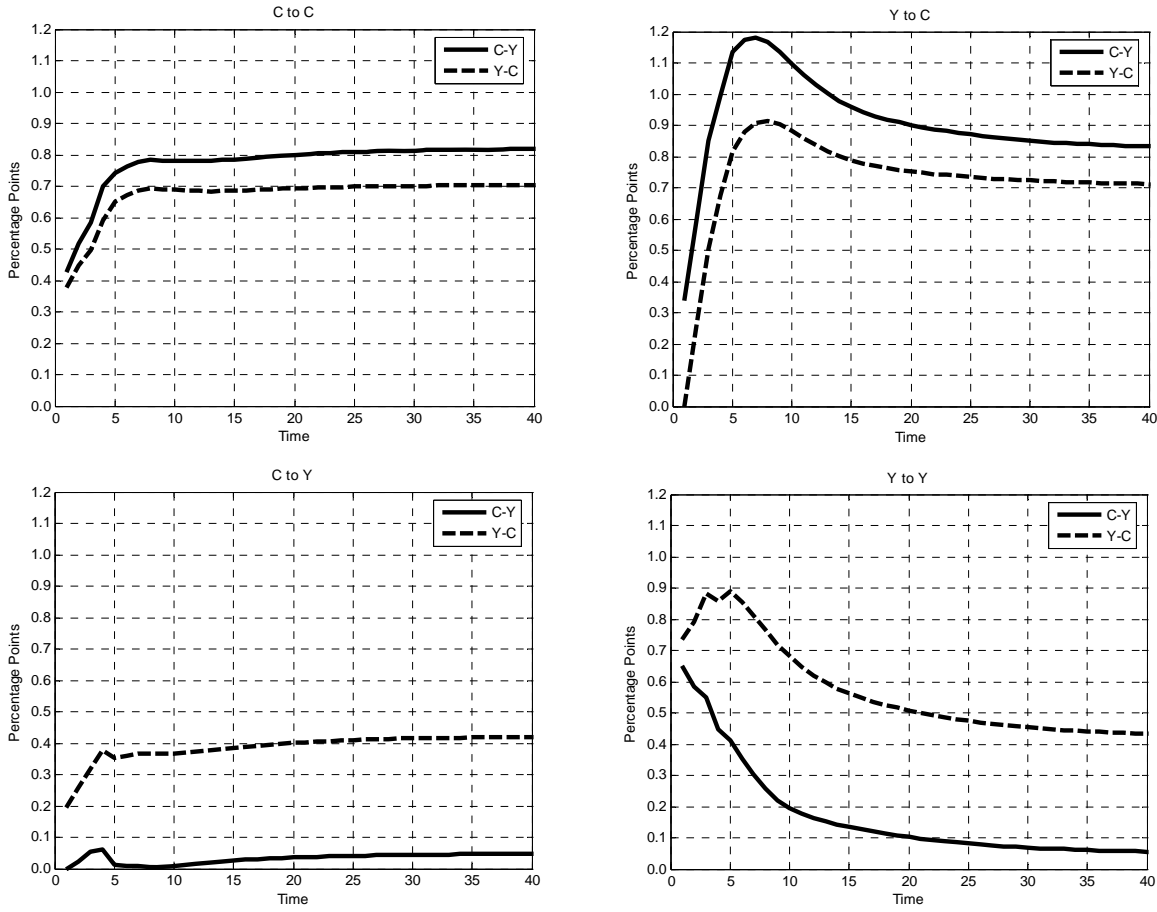
Parameter	Estimate	90 percent confidence interval
λ_1	0.15	[-1.13, 1.38]
λ_2	29.11	[7.44, 57.91]
κ	0.76	[0.48, 0.91]
σ_e	0.17	[0.01, 0.09]
σ_v	0.21	[0.001, 1.00]

Innovation variance in consumer confidence:

Due to unexpected change in current state:	0 percent
Due to information shocks:	41 percent
Due to animal spirits shocks:	59 percent

The above are estimates of the parameters of the model augmented with a structural specification of consumer confidence, as described in Section IV. The confidence intervals are computed as the 5th and 95th percentiles of a Monte Carlo simulated distribution.

Figure 1
Impulse Responses in Cochrane's Bivariate VAR



These are impulse responses from the bivariate consumption-income VAR discussed in Cochrane (1994a). The solid line shows responses from the orthogonalization with consumption ordered first (Cochrane's principle interpretation), while the dashed line refers to an orthogonalization with income ordered first.

Figure 2
E5Y and E12M






Figure 3
Impulse Responses to $E5Y$ and C Innovations in Augmented Consumption-Income VAR
(a) Orthogonalization $E5Y \rightarrow C \rightarrow Y$



These are impulse responses from the three variable VARs described in Section II. We omit the responses to output innovations. Dashed lines represent 90 percent confidence bands based on the bootstrap after bootstrap of Kilian (1998).

Figure 3
Impulse Responses to $E5Y$ and C Innovations in Augmented Consumption-Income VAR
(b) Orthogonalization $C \rightarrow E5Y \rightarrow Y$



These are impulse responses from the three variable VARs described in Section II. We omit the responses to output innovations. Dashed lines represent 90 percent confidence bands based on the bootstrap after bootstrap of Kilian (1998).

Figure 4
Forecast Error Variance Decompositions from Augmented Consumption-Income VAR





Figure 5
Impulse Responses to Confidence Innovations in Augment Consumption-Income VARs
Alternative Measures of Confidence: E5Y, PFE, and ICS



These are impulse responses from the three variable VAR described in Section II. The solid line depicts the responses when E5Y is the confidence measure, the dashed line when PFE is used, and the dotted line shows responses when the ICS is the confidence measure.

Figure 6
Spike Plots of Responses in News Heard Categories



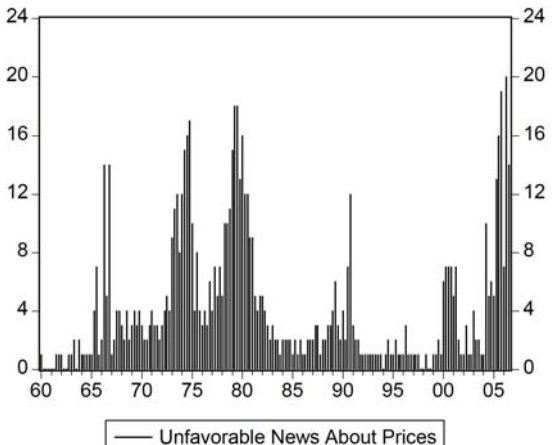


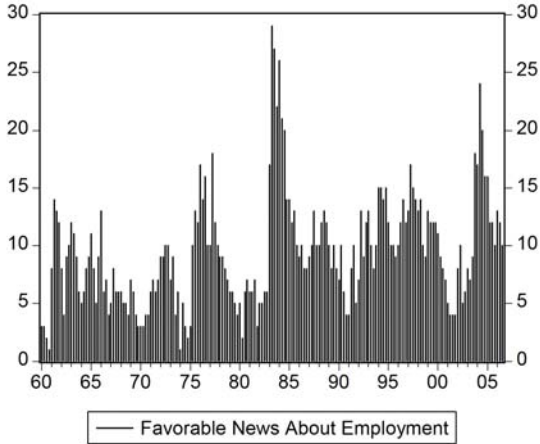
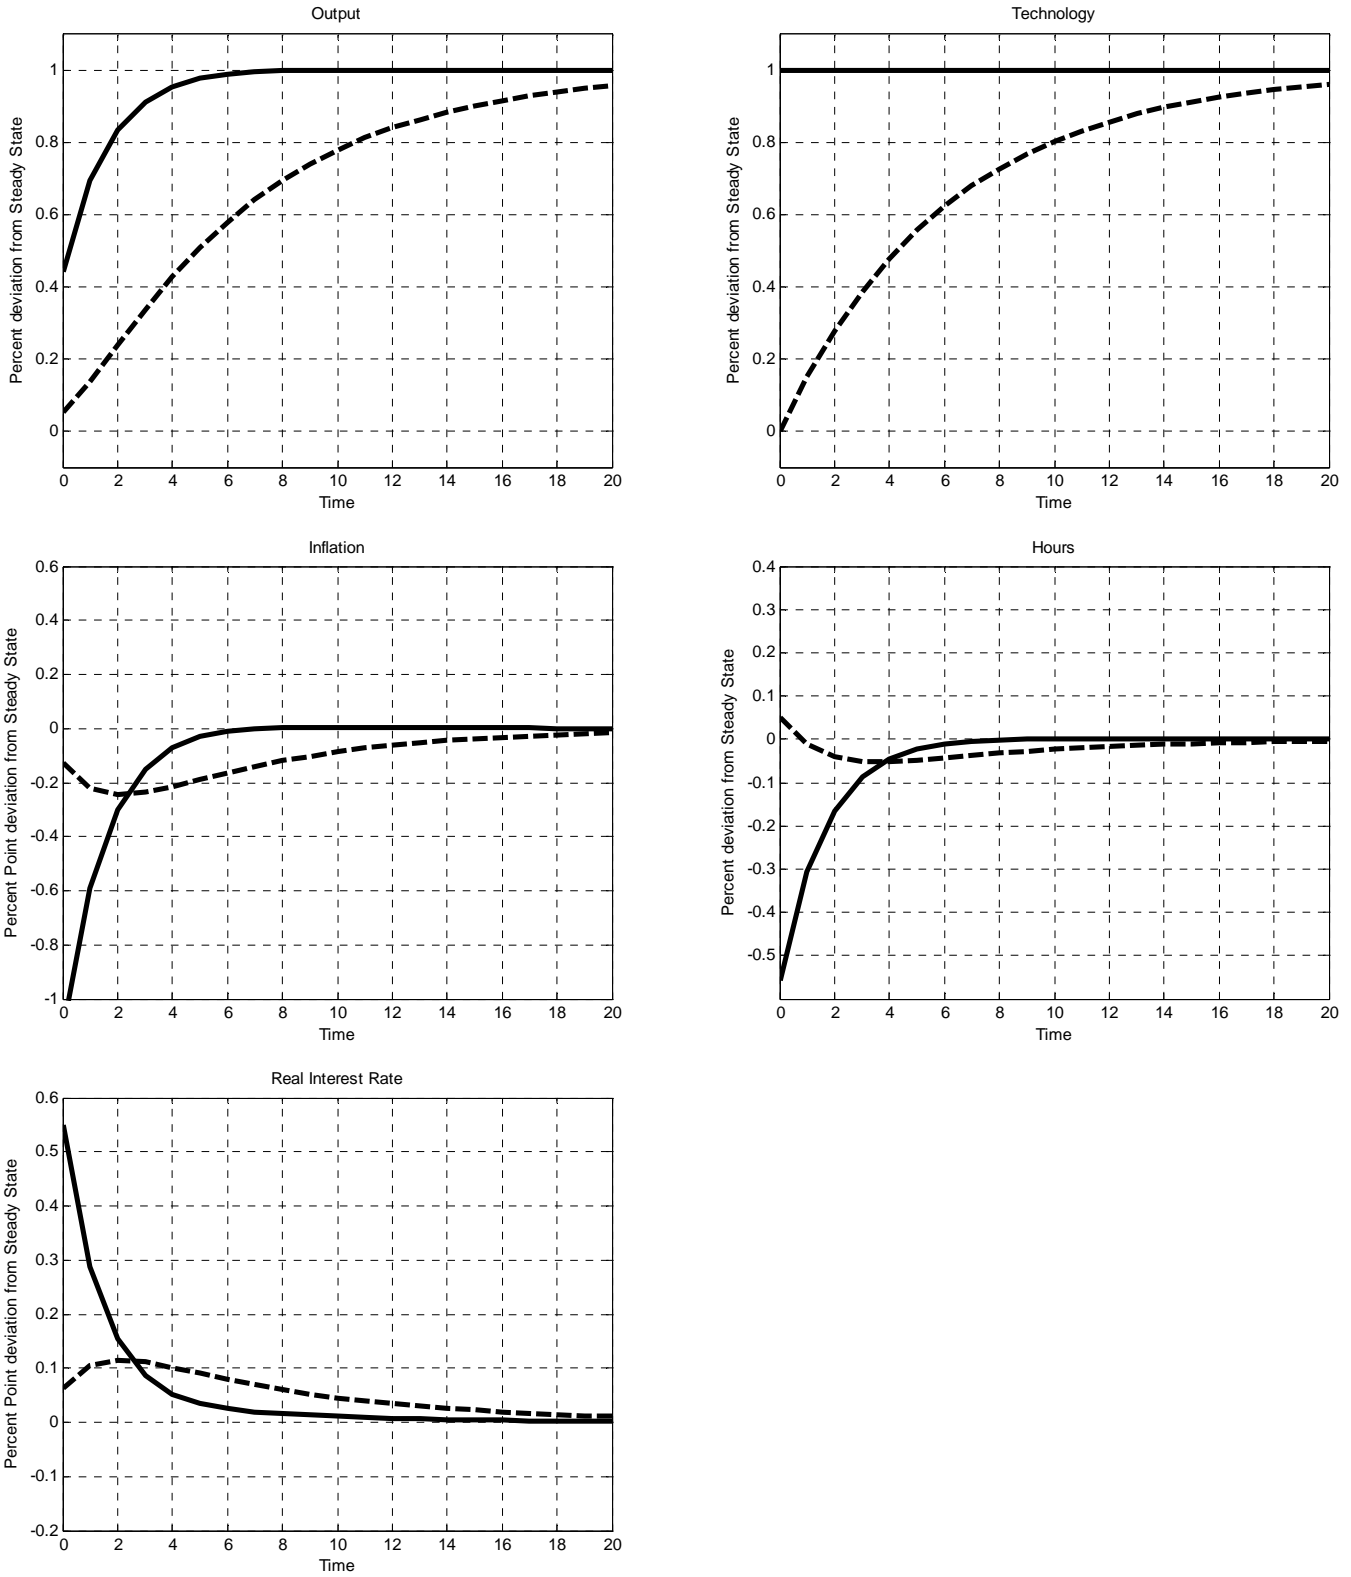



Figure 7
Theoretical Responses to Level and Information Shocks in Model




The above are theoretical impulse responses from the model using the calibration noted in Section III. The solid line depicts the responses to a level shock to technology of one percent, while the dashed line shows the responses to an information shock that technology will eventually be one percent higher.

Figure 8
Theoretical Response to Animal Spirits Shock in Model




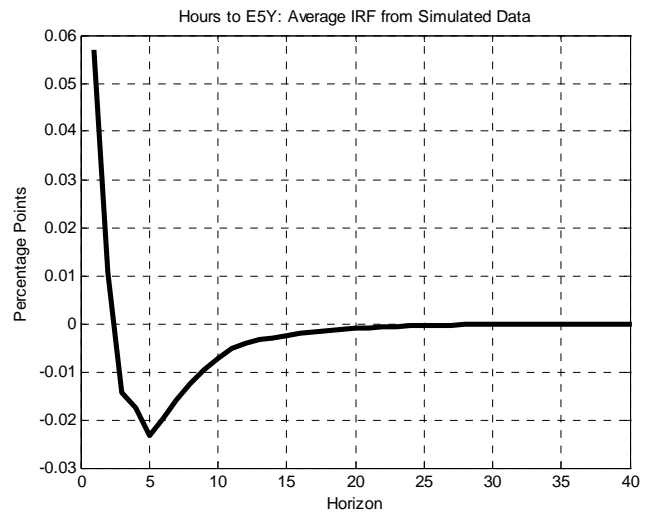
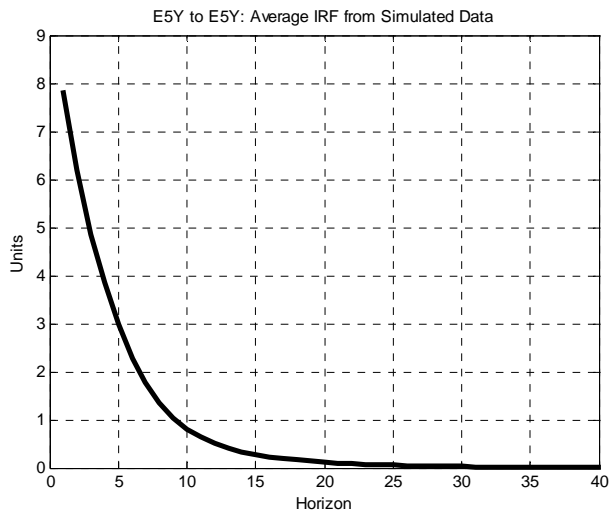
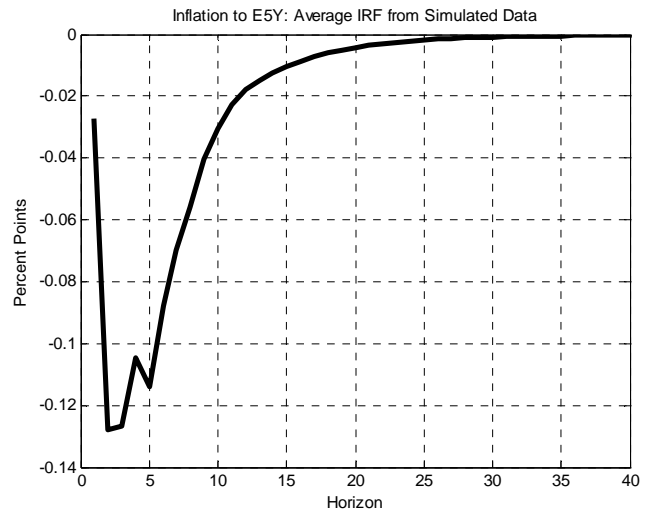

The above are theoretical impulse responses to an animal spirits shock from the model using the calibration noted in Section III.

Figure 9
Responses to E5Y Innovation in Expanded Empirical VAR




These are empirical responses from the six variable VAR described in Section IV. The dashed lines represent 90 percent bootstrap after bootstrap confidence bands.

Figure 10
Average Responses from Simulated Data Using Parameter Estimates



These figures show average responses from 2000 simulations with 200 observations each using the parameter estimates.

Figure 11
Impulse Responses from TFP-E5Y VAR



These are impulse responses from a bivariate VAR featuring a utilization corrected measure of total factor productivity and E5Y. The innovations are orthogonalized so that the structuralized E5Y innovation has no contemporaneous effect on the level of TFP. The dashed lines are 90 percent confidence bands from a bias-corrected bootstrap.