Information Shocks, Animal Spirits, and the Meaning of Innovations in Consumer Confidence*

ME
7/13/06
1:55 PM

Robert B. Barsky†
University of Michigan and NBER

Eric R. Sims‡
University of Michigan

June 19, 2006

Abstract

We show that unexplained innovations in several variables representing survey responses to forward-looking questions on the Michigan Survey of Consumers have powerful prognostic implications for the future paths of macroeconomic variables. We attempt to distinguish the hypothesis that these impulse responses indicate a causal channel from autonomous movements in sentiment to economic outcomes (the "animal spirits" view) from the alternative interpretation that the surprise confidence movements summarize information about future economic prospects (the "information" view). In natural rate models, "animal spirits" shocks are associated with "overshooting" of (among other variables) consumption that attenuates when agents come to grips with their overreaction, while "information shocks" about the long future are followed by gradual movements in macroeconomic variables that are not subsequently reversed. In a baseline vector autoregression involving consumption, income, and confidence, the data come down sharply in favor of the information view. The impulse responses of consumption and GDP show no tendency to attenuate even after a number of years. Further, confidence innovations have strong implications for labor productivity many quarters into the future. In somewhat larger VARs with an information block that includes inflation and/or stock prices, the impulse responses to confidence innovations continue to have the permanent shape that defines information shocks, but they are smaller in magnitude. We demonstrate that this is due to the fact that both inflation and stock price innovations have prognostic implications for future productivity that are very similar to the implications of innovations in confidence. Addition of unemployment to the system induces a transitory component that changes the shape of the impulse responses and somewhat weakens the previously airtight case against animal spirits, but it does not provide constructive evidence of such effects.

†: barsky@umich.edu
‡: ericsims@umich.edu

* Preliminary draft. Comments welcome. The authors gratefully acknowledge helpful comments and suggestions from Susanto Basu, Alan Blinder, Daniel Cooper, Christopher House, Peter Ireland, Serena Ng, Robert Solow, and seminar participants at Boston College, the Johns Hopkins University, the University of Michigan, and the Russell Sage Foundation. Robert Barsky acknowledges generous support as a visiting scholar at the Russell Sage Foundation.

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 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]

---

[1] In an interesting but almost forgotten early contribution, Hall (1986) – partially repudiating Hall (1978) – argued 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 (1992), 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. The existence of sunspot equilibra 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

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 information about the future evolution 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 evidence with theory closely anticipates the "news approach to business cycles" of Beaudry and Portier (2004, 2005).

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 small VARs, we show that unexplained innovations in consumer confidence have significant, slowly building, and apparently permanent effects on measures of real economic activity. The predictive implications of innovations in confidence for activity variables are far stronger when one uses a select few of the purely forward looking individual questions from the survey as opposed to the overall index number.

We next 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 future economic prospects known to consumers but unobservable to the econometrician (the "superior information" view). In a variety of models, "animal spirits" shocks are associated with "overshooting" of consumption that attenuates when agents come to grips with their erroneous expectations. (See Lorenzoni (2005) for a recent model with this flavor). On the other hand,

---

not cover sunspots. However, Susanto Basu pointed out to us that the most common sunspot models have stationary equilibria, and that these would presumably generate the same transitory impulse responses that are associated with animal spirits effects as defined here.

"information shocks" regarding the long future are followed by gradual movements in the macroeconomic variables that are not subsequently reversed.

We formalize an *information shock* as a structural innovation that is orthogonal to current output but has a permanent effect on future output. Likewise, a *transitory shock* or *potential animal spirits shock* is one that raises output over at least several quarters, but not over several years. Using agnostic identification procedures of the sort proposed by Faust (1998) and Uhlig (1999), we identify an entire range of shocks satisfying each of these criteria. We find that reduced form innovations in confidence are highly correlated with structural innovations satisfying the criteria of an information shock. Further, confidence innovations are essentially uncorrelated with structural innovations identified as having transitory consequences for output. The first result further reinforces the notion that movements in confidence are to an important extent proxies for information shocks. The second indicates that the trivariate VAR does not provide evidence of an important animal spirits channel from confidence to output.

The robustness of the "information shock" interpretation to the inclusion of information beyond output is studied in the context of slightly larger systems in Section III. Regressing observed confidence innovations on "news heard" categories from the Michigan Survey, we find that confidence innovations – which appear exogenous in the three variable system – are somewhat explicable by news about inflation, with news about unemployment and stock prices playing slightly less important, though non-negligible, roles. The inclusion of inflation and/or stock prices into the VAR generally reduces the magnitude of the impulse responses of output and consumption to orthogonalized innovations in confidence. Importantly, the responses still appear permanent – and thus consistent with the information approach and not the animal spirits view of confidence.

We show that – taken separately – innovations in confidence measures, innovations in inflation, and innovations in stock prices each predict large and permanent changes in future labor productivity, even after orthogonalization with respect to current productivity.[3] The effects of each are of approximately the same magnitude and have the same temporal pattern. Within our proposed information shock context, the finding of a strong commonality between the inflation and confidence innovations helps explain the reduced magnitude of the effects of shocks to confidence in the larger block VARs discussed in Section III. This connection between inflation, confidence, and real activity variables is an intriguing one that poses questions that we only begin to address in this paper.

Finally, we expand the VARs to include unemployment as an additional variable. In contrast to the previous cases, recursive identifications of the shocks in a VAR with unemployment included indicate that the responses of output and consumption to confidence innovations show more stationary behavior. Agnostic identifications of shocks having transitory effects on output produce a wide range of possibilities – from the earlier finding that such shocks are essentially uncorrelated with innovations in consumer confidence to the polar extreme that such shocks are essentially collinear with innovations in confidence. As such, we are unable to definitively rule out an interpretation of consumer confidence in accord with the "animal spirits" hypothesis. However, none of the findings constitutes an affirmative case for an animal spirits channel.

[3] The result for stock prices is similar to the finding of Beaudry and Portier (2004b).

II. Income, Consumption, and Confidence

We begin with the dynamics of income and consumption as implied by the bivariate VARs introduced by Cochrane (1994a).[4] Figure 1, Panel A shows the impulse responses when C is ordered before Y, and Panel B presents the case in which Y is ordered first. As Cochrane (1994a, 1994b) stressed, the key feature of these impulse responses is that innovations in consumption (whether or not they are orthogonalized with respect to Y) are associated with powerful and prolonged subsequent increases in output. If C is ordered first, almost the entire permanent component of Y is explained by consumption. If Y is ordered first, the finding is less dramatic; orthogonalized innovations in both C and Y lead to large and permanent movements in Y. Yet here too innovations in consumption generate more prolonged impulses than income innovations, and more than a third of the forecast error variance in output at long horizons is accounted for by shocks to consumption conditional on income. In short, the results from the two variable VAR suggest that "consumption shocks" convey news about income many periods in the future.[5]

The natural explanation of this phenomenon 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 provide an alternative measure of such information. How much of the information about the future embodied in consumption is picked up by survey expectations? Conversely, do the survey expectations indicate that

[4] The variables are real GDP and real consumption of nondurables, the data are quarterly from 1960 through the second quarter of 2005, and the VARs are run in log levels with two lags. We do not impose cointegration between C and Y in the VARs reported on here because, theory aside, it is rather inconsistent with the data. We have, however, run the VAR both in first differences and with various assumptions about cointegration, and the results are very similar.

[5] Cochrane (1994) emphasized the ordering with consumption first because it is suggested by the simple permanent income model in which all information is immediately reflected in consumption, and because it splits income fairly neatly into permanent and transitory components. However, the idea of an information shock, as defined below, leads to a focus on the alternative orthogonalization.

consumers receive news that is not reflected in current consumption? We turn to these questions now, introducing some expectational measures from the Michigan survey and augmenting the bivariate consumption-income VARs with these variables.

The survey measure that we will make the most use of in this paper, which we call E5Y, summarizes responses to the question "Now 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. 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 relevant time frame. However, its correlations both with the response to a similar question specifying a horizon of only twelve months (a variable we call E12M) and with the overall expectations component of the Michigan index (based on a weighted average of these two questions and another concerning expected changes in personal financial situation over the next twelve months) exceed ninety percent. The results in this section are essentially unchanged by the substitution of either of these alternative expectations variables.

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. Table 1 shows the cross-correlogram between E5Y and the conventional Hodrick-Prescott detrended GDP. 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.

Table 2 presents estimated coefficients and associated statistics for a trivariate VAR featuring real GDP and real nondurable consumption, along with the variable E5Y. The VAR is once again run in log levels, with two lags of each variable (as suggested by the Schwartz Information Criterion).[6] To address the relative informational content of E5Y and consumption, we focus on two alternative orthogonalizations: C→E5Y→Y and E5Y→C→Y. If all relevant news about future income is immediately reflected in consumption, an E5Y innovation orthogonalized with respect to consumption will have no predictive implications for the path of output. The other polar hypothesis (tested by the E5Y→C→Y orthogonalization) is that consumption contains no information about future Y beyond that conveyed by our expectations variable.

Figure 3 presents impulse responses and Figure 4 variance decompositions for both orderings. Regardless of the ordering, orthogonalized innovations in E5Y are followed by slowly-building, large, and permanent responses of both C and Y; the long-run effects are in the neighborhood of one percentage point higher consumption and output for a one standard deviation positive innovation in E5Y. When E5Y is ordered first, the long-horizon implications of E5Y shocks actually exceed those of consumption shocks, and E5Y accounts for about half of the variance of output at twenty-four quarters. However, the response to an E5Y shock builds so slowly that consumption innovations have much larger effects than E5Y shocks at short horizons. When E5Y is orthogonalized with respect to consumption, the expectational variable never quite catches up. Still, the variance decompositions show that one third of the variance of output at long horizons is accounted for by the survey measure. Thus neither of the above polar

[6] Both consumption and GDP have unit roots, while E5Y is clearly stationary. Estimation of the VARs in this section with consumption and GDP in first differences produces nearly identical impulse responses. While our data are not in favor of a cointegrating relationship between consumption and GDP (because of the sharp upward trend in consumption relative to income in the last ten years or so), imposing such a relationship also produces very similar results.

cases applies. Innovations in E5Y and in consumption each convey relevant news about future output that is not subsumed in the other.

Regardless of ordering, E5Y responds almost entirely to its own innovations, a fact which is evident from both the impulse responses and variance decomposition. At business cycle frequencies, more than one-third, and up to more than one half, of the forecast error variances of both output and consumption are attributable to innovations in E5Y when E5Y is ordered first. When consumption is ordered first, the fraction of the forecast error variance of consumption attributable to E5Y innovations is somewhat smaller at most horizons, but still tops one third at business cycle frequencies.

For our principal interpretation of the VAR, we focus on an orthogonalization in which income is ordered first and consumption last, with E5Y in the middle position. This is motivated by the idea of an "information shock" – news about future output (or productivity) that is not reflected in current output. The impulse responses and variance decompositions for the trivariate VAR with the GDP→E5Y→C orthogonalization are shown in Figures 5 (A) and 5 (B), respectively. The key observation is the gradual but powerful and prolonged response of output and consumption to the E5Y shock. An "unexplained confidence shock" raises consumption by a modest (though non-negligible) 0.1 percent on impact (as compared with a 0.27 percent response of consumption to an income innovation and a 0.56 percent response to its own innovation). However, after twelve quarters a one standard deviation shock to E5Y has raised consumption by 0.63 percent, compared to the 0.59 percent that it would have been raised by a one standard deviation shock to income. Likewise, after twelve quarters a one standard deviation shock to E5Y has raised income by 0.67 percent. The variance decomposition confirms that negligible shares of consumption and GDP are explained by E5Y at short horizons,

but a much larger percent are explained at longer horizons. Strikingly, the responses of both income and consumption show no tendency to attenuate at long horizons.

Although it shows a statistically significant positive response to both income and consumption shocks, E5Y responds almost exclusively to its own innovation. Surprises in E5Y must be reflective either of information not contained elsewhere in the VAR or of truly exogenous animal spirits. Later we will see that a nontrivial fraction, though not the lion's share, of E5Y innovations are explained by included variables in a larger system, especially prices.

We now examine several variations on our baseline three variable VAR. First, we use durable consumption expenditures in place of non-durables. Second, we substitute the relative score from the question on the Michigan survey concerning expected personal financial situation (PFE) in place of E5Y, which has an interpretation analogous to E5Y,[7] and the third is to use the commonly employed ICS, which is an equally weighted composite of five backward-looking and forward-looking questions from the Michigan Survey. The impulse responses under the first two modifications are presented in Figures 6 (A) and 6 (B), respectively. The conclusions are virtually identical to those from the baseline specification with E5Y and non-durable consumption. Both income and consumption continue to have large and persistent responses to innovations in consumer confidence as measured by forward-looking questions.

Figure 6 (C) presents the impulse responses of our baseline VAR using the oft reported Index of Consumer Sentiment in place of E5Y. The structuralized innovation to the ICS has far weaker implications for the evolution of both total output and non-durable consumption expenditures than does either of the specific forward looking questions from the survey. In

[7] This question, which we call PFE, is as follows: "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?"

particular, the point estimates of the impulse responses of both GDP and consumption to ICS innovations are roughly half in magnitude compared with use of E5Y or PFE, and the ninety-five percent confidence bonds lie outside of zero only for the first few quarters following the innovation. This is one important explanation for the apparent incompatibility between our result that confidence measures have important predictive implications for output and spending and the results in a number of recent papers that argue that confidence surveys have little predictive value.[8,9]

In summary, innovations in expectational variables from the Michigan survey not explained by innovations in current output are powerful predictors of output and spending for the foreseeable future. There is a great deal of initial undershooting; consumer confidence shocks explain relatively little of the variation in income and spending at short horizons. The effects of a confidence shock gradually build and show no tendency to attenuate as time passes. This lack of attenuation in the responses of either GDP or consumption is evidence in support of an information view of consumer confidence. Were the confidence innovation primarily reflective of exogenous changes in "animal spirits" we would expect such innovations to have predictably short-lived effects on output, at least within the context of natural rate models with a unique equilibrium. Our empirical results might be consistent with a large role for animal spirits if the economy exhibits multiple equilibria of a sort that allows the economy to shift permanently from

[8] Among papers in this literature that find different predictive roles for measures of consumer sentiment are Mishkin (1978), Carroll, Fuhrer, and Wilcox (1994), Matsusaka and Sbordonne (1995), Mehra and Martin (2003), Ludvigson (2004), and Croushore (2005).

[9] There are several other differences as well. One is that we focus on impulse responses, whereas other papers have focused on incremental R-squared statistics associated with Granger causality tests. Another issue is what happens to the importance of confidence as additional variables are added to the conditioning information set, an issue to which we will return in greater depth shortly. For some purposes, it is most interesting to consider the *incrementa*l role of confidence after accounting for a substantial list of macroeconomics variables. For other purposes – e.g. the above exercise studying the extent to which consumption shocks should be regarded as information shocks – it is not.

one level of productivity to another via supply externalities or other mechanisms producing very large increasing returns. The conditions under which an economy may exhibit such behavior are stringent, and the empirical literature has found little evidence that such conditions might be satisfied in the data.

The observation that responses of output and consumption to confidence shocks are permanent suggests that the innovation in confidence reveals information about future productivity in much the same fashion as the information shocks in Lorenzoni (2005) or Beaudry and Portier (2004c). As a test of this conjecture, we first run a simple bivariate VAR with the log of labor productivity and E5Y, with labor productivity ordered first. The impulse responses are presented in Figure 7. Shocks to E5Y do indeed have a quantitatively and statistically significant and approximately permanent effect on labor productivity. As we show later in the paper in the context of a larger VAR, the fact that prices fall in response to confidence innovations is further evidence in support of a productivity knowledge interpretation of innovations in confidence.

It might be desirable to have an alternative test, not dependent on vector autoregressions, of the implication of the animal spirits hypothesis that when confidence is unusually high, consumption eventually reverts. Table 3 reports the results of regressing consumption growth over various horizons (one quarter, one year, two years, five years, and seven years) on E5Y. The regressor in each equation is E5Y at date t, while the regressand is the growth rate in consumption from period t to $t+k$. Newey-West standard errors are reported because of the high degree of serial correlation associated with overlapping observations. The null hypothesis is that consumption growth at each horizon is invariant to E5Y. The animal spirits alternative implies that the coefficient of E5Y is negative at long horizons. If increases in E5Y induce unsustainable

increases in spending at short horizons, consumers must eventually face the reality of their lifetime budget constraints and reverse the initial increases in consumption.

Table 3 reveals that at each horizon the coefficient estimate is significantly positive at better than the one percent level. Though the magnitudes may appear small, the standard deviation of E5Y exceeds twenty points, and the implied swings in consumption growth are rather large; predicted five-year consumption growth, for example, ranges from less than ten percent (two percent per annum) to more than sixteen percent (roughly three percent per annum). These long-horizon regressions corroborate the findings from the VAR that increases in E5Y lead to permanent and statistically significant increases in consumption.[10] Clearly, they do not provide support for the animal spirits hypothesis.

To what extent are our results from the recursive identifications that support the information shock interpretation of confidence robust to a range of alternative, potentially non-recursive, identifications? To address this question, we employ "agnostic" identification procedures first expounded in Faust (1998) and Uhlig (1999). The idea is to consider (separately) the universe of shocks that have the appearance of information shocks and the universe of those that behave like animal spirits shocks.

First, we address the information shock. How well does a shock identified as being contemporaneously uncorrelated with a measure of output but having a "large" influence on output at a moderately long horizon explain E5Y innovations? To answer this question, we employ the method of maximum forecast error variance identification developed in Faust (1998)

[10] The finding that one-period-ahead consumption growth is predictably higher when confidence is high is the essential result of Carroll, Fuhrer, and Wilcox (1994). One might be tempted to ask, in the spirit of Carroll, et al, why our results seem so inconsistent with the permanent income hypothesis; i.e. why consumption predictably tracks productivity increases rather than jumping sharply in anticipation of them. The long-horizon results make it clear that the answer must be a general equilibrium one, and it will be taken up in a subsequent paper on the information approach to business cycles.

and Francis, Owyang, and Roush (2005). (For a more detailed sketch of the procedure, see Appendix A). The procedure involves computing an entire family of shocks identified to be orthogonal to current output by creating orthonormal rotation vectors of an arbitrary orthogonalization of the reduced form. For each shock, we then compute the fraction of the forecast error variance of output at a horizon of eight years attributable to this shock. The shock explaining the maximum of the forecast error variance of output at this horizon is then deemed a "best case" for a shock satisfying the criteria of an "information shock".

Among the family of shocks orthogonal to current output, we can find a shock that accounts for more than forty-five percent of the forecast error variance of output eight years subsequent to the shock. Strikingly, such a shock accounts for seventy-nine percent of the contemporaneous innovation in E5Y and upwards of eighty-five percent at longer horizons, substantially reinforcing the notion that innovations in E5Y are manifestations of news consumers receive concerning future fundamentals. Figure 8 (A) presents historical simulations of the included variables in the VARs in response to the "best case" for the identified information shock. Though it does not account for much of the business cycle variation in output, such a shock does a rather good job at explaining historical movements in consumption. It does an even better job at explaining historical movements in E5Y, lending further support to the idea that measured confidence is, in significant measure, a reflection of genuine information about the future.

How sure can we be that there is not an alternative identification that suggests that confidence includes an important animal spirits component? We proceed in a way similar to the previous exercise, but instead find orthonormal rotations of the reduced form trivariate VAR that

lead to predictably temporary responses of measures of real activity. We then ask how well such shocks can account for the contemporaneous innovation in E5Y.

We are able to find shocks that have large transitory effects on both income and consumption. In particular, the "best" case for a transitory shock accounts for more than two-thirds of the forecast error variance of output at high frequencies and leads to an impulse response in which output rises by more than 0.6 percent after a year before reverting to its pre-shock level. However, such a shock accounts for less than fifteen percent of the innovation in E5Y and – worse – moves E5Y and output in opposite directions. Of all linear rotations of the VAR, only three percent lead to simultaneous increases in E5Y and income, and among these three percent the increase in E5Y is always less than one point and the associated change in output is small. It is thus impossible, within the three variable context, to find a shock which simultaneously leads to a large positive increase in E5Y and a large, temporary response of output.

Historical simulations further reinforce the case against an animal spirits interpretation of consumer confidence. As shown in Figure 8 (B), the best case for the role of the transitory shock does a reasonable job at explaining historical movements in output and a poorer job for consumption. It does a much worse job at explaining the historical movements in observed confidence. The historical simulation for E5Y from the max forecast error variance identification and that from the impulse response shape identification portray a sharp contrast.[11] The identified information shock does quite a good job at explaining the historical movements in E5Y, the transitory shocks do not. In the next section of the paper, we examine the extent to which expanding the conditioning set of the VAR might alter these findings.

[11] Note that the respective "best-case" shocks are from separate identifications and their contributions are not additive. Neither method identifies a full set of structural shocks, so that these are not complete historical decompositions.

III. The Role of Consumer Confidence with Richer Conditioning Information

Up to this point we have focused on a trivariate setting with only GDP and consumption in addition to our forward-looking measure of consumer confidence. This small system was motivated naturally as an extension of Cochrane's bivariate system inspired by the permanent income hypothesis. It made sense to ask about the extent to which "consumption shocks" could be interpreted as "news shocks" that reveal information about future fundamentals, and the exercise generated results that were sensible and easily interpretable. Clearly innovations in E5Y convey important news about the relatively distant future, and we were able to conclude that in this context it is this information content and not a self-fulfilling aggregate demand effect that accounts for the lion's share of the relationship between confidence and subsequent real activity.

On the other hand, it is also important to investigate the role of confidence variables after conditioning on a more expansive set of macroeconomic indicators, for several reasons. First, it is interesting in and of itself to understand what – if anything – helps explain movements in confidence, as we found it to be more or less exogenous in the three variable setting. Second, previous authors have found that the predictive content of consumer confidence is weakened once financial variables in particular are taken into account.[12] Third, we need to know whether or not orthogonalizing confidence with respect to other macroeconomic indicators changes the dynamics of the VAR. To the extent that it does, we may have to reassess whether the animal spirits hypothesis has greater plausibility than the impulse responses from the trivariate system led us to believe. Finally, we will see that confidence variables serve as a lens through which

[12] See, for example, Mishkin (1978) and Leeper (1992).

we discover an intriguing and potentially important relationship between future productivity and current inflation.

III (a): Determinants of Confidence Innovations

In the simple trivariate setting, E5Y and other overall confidence measures are roughly exogenous. With E5Y ordered in between GDP and non-durable consumption in a recursive VAR, close to ninety-five 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 eighty-five percent.

Are there variables other than GDP and consumption that might help explain a sizable share of the 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.[13] It seems natural to begin with an investigation of the relationship between this reported economic news and subsequent responses to survey questions concerning macroeconomic expectations in attempting to answer the question of what – if anything – helps explain consumer confidence.

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 9 presents spike plots for several of the more popular response categories 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

[13] The specific question is phrased as follows: "During the last few months, have you heard of any favorable or unfavorable changes in business conditions. What did you hear?"

about prices and news about unemployment. Other responses stand out in particular time frames. 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.

In Table 4 we present coefficient estimates from regressions of the E5Y innovations from the three variable VAR on selected categories of news. The news heard categories have coefficients of the expected signs; an increase in the percentage of respondents reporting favorable news concerning employment or prices is positively correlated with confidence innovations. Unfavorable news about employment is evidently an insignificant covariate, while favorable news about unemployment and inflation are both significant at the five percent level and unfavorable news about prices is significant at nearly the one percent level. In short, news about prices has statistically and substantively significant explanatory power for the innovation in consumer confidence, while news about employment also has marginally significant effects.

III (b): Confidence, Inflation, and Stock Prices

Figure 10 plots inflation (right scale) and E5Y (left scale) across time.[14] The general negative association between the two series across time is quite marked, though it is most evident during the 1970s and early 1980s when inflation was most volatile. Since the 1990s there have been important movements in confidence unrelated to inflation.

We now add inflation to our baseline VAR alongside consumption, income, and E5Y. As with prior VARs, the system is run in log levels on quarterly data from 1960-2005 with a lag length of two. Figure 11 presents the impulse responses of all variables to orthogonalized inflation and confidence innovations when inflation is ordered prior to confidence, while Figure

[14] Inflation is measured as the first difference of the natural logarithm of the Consumer Price Index.

12 covers the case in which E5Y is ordered prior to inflation. When E5Y is ordered first, the responses of consumption and income to inflation are – qualitatively and quantitatively – the mirror images of the responses to E5Y. Because the correlation between the reduced form innovations in inflation and confidence in this system is nearly -0.3, the orthogonalization matters; attributing the common component of the innovations in inflation and E5Y to inflation significantly reduces the responses of consumption and income to the structuralized E5Y innovation. Indeed, when inflation is ordered first, the impulse responses to E5Y at all horizons are in the neighborhood of one half of the magnitude obtained in the system without inflation. Importantly, the responses of the activity variables to confidence innovations have the same non-reverting shape that they had in the trivariate case. Once again, the recursive identifications do not produce anything resembling an animal spirits shock.

As in the three variable case, we employ agnostic identification procedures as robustness checks. We can find a shock explaining more than forty percent of the forecast error variance of output among those shocks orthogonal to current output. Such a shock explains slightly more than half of the contemporaneous innovation in E5Y. At somewhat longer horizons, the contribution of this shock to the E5Y forecast error variance is greater than fifty percent and reaches as high as seventy percent. While the correlation between this information shock and the reduced form innovation in confidence is somewhat smaller than that found in the three variable system, it is nonetheless impressive.

We also impose a shape restriction to identify the family of shocks having transitory effects on output. We define the "best case" as that shock satisfying the shape restriction that has the maximal effect on output at a horizon of one year. Such a shock accounts for an increase in output of slightly more than 0.6 percent on impact. In contrast to the three variable case –

where it was essentially impossible to find a transitory shock that moved confidence and output in the same direction – E5Y now rises on impact, but by a very modest one point. This "best case" for the transitory shock explains only fifteen percent of the innovation in E5Y. However, we need to investigate further other transitory shocks that produce more modest increases in output but explain much more of the variance of E5Y.

Figures 13 (A) and (B) present historical simulations of output, inflation, confidence, and consumption for the "best cases" for the information and transitory shocks, respectively. The best case for the information shock does a good job at explaining the time series variation in E5Y and inflation, and, to a lesser extent, output and consumption. In contrast, the best case transitory shock does a poor job at matching the historical movements in E5Y and inflation, though it does pick up several important movements in real activity.

Figure 14 plots the S & P 500 index (deflated by the CPI) along with E5Y. The positive correlation between the two series is not difficult to detect and is prevalent throughout the entire sample period. On any reading – information, animal spirits, or a third alternative – stock prices and confidence measures ought to reflect similar forces. Thus the finding of previous authors (Leeper (1992) and Mishkin (1978)) that the inclusion of asset prices reduces the predictive content of measures of consumer confidence for subsequent real activity is neither unexpected nor an adverse judgment on the value of confidence data.

In Figure 15, we present impulse responses when stock prices are included in the VAR and ordered prior to E5Y. As has been the case in all VARs examined thus far, there is no tendency for either income or consumption to attenuate in response to orthogonalized innovations in E5Y. Stock price innovations appear to have qualitatively and quantitatively much the same implications for income, consumption, and E5Y as do inflation innovations

(though of opposite sign). Ordering E5Y prior to inflation attributes more of the common component of the stock price and E5Y innovations to E5Y, and leads to larger and more statistically significant impulse responses of real activity to the E5Y innovation (not shown). Orthogonalizing stock prices with respect to E5Y rather significantly weakens the quantitative and statistical significance of the responses of income and consumption to stock price innovations, something which we did not see in the case of inflation when it was included in the VAR and ordered after E5Y.

Figure 16 presents the impulse responses when both inflation and real stock prices are included in the VAR and are ordered prior to E5Y. For the presented plots, inflation is ordered prior to stock prices. Innovations to E5Y continue to have the same qualitative effects on both consumption and income even after orthogonalization with respect to stock prices and inflation. In particular, the point estimates indicate no tendency for either income or consumption to revert following a positive innovation in E5Y. As in the cases in which either inflation or stock prices alone were included in the VAR, the responses of real activity to E5Y are smaller and less statistically significant than in the three variable case.

In isolation, innovations confidence, inflation, and stock prices each lead to slowly-building, quantitatively significant, and permanent responses of both income and consumption. This suggests they may possess much of the same informational content. Figure 17 plots the first principal component of inflation and stock prices (right scale) along with our forward-looking measure of consumer confidence (left scale) across time. The high degree of co-movement among the series is highly apparent. With few exceptions, the peaks and valleys of each appear to be nearly identical images of one another, with essentially every large movement in E5Y associated with a similar movement in the principal component of inflation and stock

prices.[15] Gordon and Dew-Becker (2005) highlight the strong negative relationship between inflation and productivity growth, while Piazessi and Schneider (2006) stress the negative relationship between inflation and consumption growth. Modigliani and Cohn (1979) and Campbell and Vuolteenho (2004) take note of the striking correlation between inflation and real stock prices.

We showed in Figure 7 that innovations in E5Y orthogonalized with respect to current labor productivity have slowly-building, large, and permanent implications for a measure of labor productivity. Figures 18 (A) and (B) present impulses response from a bivariate system analogous to that presented in Figure 7 but with either inflation or stock prices used in place of confidence. The results from this set of bivariate VARs, in conjunction with those presented in Figure 7, are striking. As already shown by Beaudry and Portier (2004b), stock price innovations are potent predictors of future productivity. We now see that inflation has equally powerful prognostic implications for long run productivity.[16]

The very similar response of productivity to innovations in either confidence, inflation, or stock prices, in conjunction with other evidence concerning the close connection between these three series, suggests that confidence, inflation, and stock prices convey much the same information about future productivity. In particular, all three anticipate the productivity

[15] One might well conclude on the basis of this diagram that the confidence series are essentially redundant – a synthetic confidence series could be created from inflation and stock market series without incurring any data collection costs. But this relationship cannot possibly be structural. As Howard Rosenstein pointed out, in the Great Depression period inflation would have been welcomed as a harbinger of economic recovery. It just so happens that in our sample period, inflation is an excellent indicator of slow long-term productivity growth.

[16] The underlying reason that inflation innovations predict productivity slowdowns is the subject of a follow-up paper in preparation. The version of the New Keynesian Phillips Curve spelled out by Gali and Gertler (2000) in the presence of *real* wage rigidity provides a possible answer. Inflation today reflects expectations of future real unit labor costs. Expectations of low productivity in the future combined with a refusal of real wages to fall accordingly means high expected unit labor costs and thus high inflation today. However, the nexus between inflation and these real variables may be too low frequency in nature to be merely a consequence of sticky prices and wages.

slowdown beginning around 1970 and its reversal in the 1980s and 1990s. The fact that these series are so interchangeable in this respect helps to explain why conditioning on either inflation or stock prices (or both) tends to reduce the predictive content of E5Y innovations. The fact that the responses of consumption and income to E5Y innovations are weakened upon the inclusion of inflation and stock prices undermines the incremental value of consumer confidence as a forecasting tool but does not in any way reverse the finding that E5Y innovations are, to an important extent, manifestations of information that consumers have about the relatively far-off future.

III (c): Confidence, Unemployment and Transitory Components

As discussed in the beginning of this section, it is consistently the case that a relatively large fraction of the survey respondents report having heard either favorable or unfavorable news concerning employment. Indeed, we might reasonably expect there to be a significant correlation between unemployment and forward-looking measures of consumer confidence. Figure 19 plots the civilian unemployment rate with E5Y across time from 1960-2005. While there are fewer high frequency fluctuations in the unemployment rate than in E5Y, large movements in the two series seem to track one another reasonably well, with the period of the early 1980s being one notable exception, where confidence generally rose from its forty year low while the unemployment rate generally also trended upward. Because confidence is related to both inflation and unemployment, it is not surprising that the strongest co-movements between confidence and unemployment appear during stagflationary periods.

Figure 20 presents the impulse responses from a recursive identification of our baseline three variable VAR augmented with the unemployment rate, which is ordered after income and

before E5Y. While most of the differences are not statistically significant, the presence of unemployment in the VAR does alter the general dynamics of the system. In sharp contrast to all of our previous results, the impulse responses to E5Y innovations now display a tendency towards reversion after a couple of years.

Not only does the presence of unemployment alter the system's dynamics in response to E5Y innovations, it also does so for both income and consumption innovations. Income's response to its own innovation is now very clearly transitory, and its response to the consumption innovation is essentially zero at all horizons. Like income, consumption has a hump-shaped impulse response to its own innovation as well as to the innovation in income. Rather surprisingly, it now appears that the primary source of permanent movements in both income and consumption is innovations to unemployment. These dynamic responses turn out not to be a result of the particular ordering of unemployment in the VAR, but rather its presence. For all recursive identifications, income and consumption have transitory responses to E5Y innovations, and unemployment innovations want to be the primary source of the unit roots in income and consumption.

Figure 21 presents impulse responses from our basic system with both unemployment and inflation included – both ordered after income but prior to E5Y. If anything, the result that unemployment is the primary source of the non-stationarity of income and consumption is only strengthened by the inclusion of inflation in the VAR. Responses of real activity variables to orthogonalized E5Y innovations continue to appear stationary. Whereas innovations in inflation led to permanent negative responses of consumption and output in systems without the unemployment rate, they now appear to have transitory effects on real activity – much like E5Y

innovations. The overall dynamics of the VARs with unemployment are unchanged by the inclusion of stock prices in replacement of, or in addition to, inflation.

It is not possible to rule out an animal spirits channel via shape restrictions when unemployment is included in the VAR. Agnostic identifications of shocks having transitory effects on output produce a wide range of possibilities – from the earlier finding that such shocks are essentially uncorrelated with innovations in consumer confidence to the polar extreme that such shocks are essentially collinear with innovations in confidence. Importantly, these impulse responses cannot be used to construct an affirmative case for animal spirits. A reverting impulse response is a necessary, but certainly not sufficient, condition for an animal spirits shock. The impulse responses to many other shocks – monetary shocks and temporary supply disturbances, for example – are also hump-shaped and reverting. Future research might conceivably lead to a more complete specification under which affirmative identification of an animal spirits shock may be possible.

IV. Concluding Remarks

Taken as a whole, our results support an interpretation of consumer confidence based neither on autonomous movements in sentiment nor on precautionary responses to uncertainty about future economic conditions, but rather on the aggregation of knowledge concerning future fundamentals, and, in particular, productivity. In this sense, we have found an empirical analogue of "information shocks". Does this mean that we have found support for the "news approachto business cycles"? Beaudry and Portier (2004, 2005) attempt to ascertain both the information structure and real features of the economy necessary for a viable theory in which news about future productivity generates broad-based movements in measures of economic

activity well in advance of the realization of the productivity change.[17] The fact that Beaudry and Portier's (2004c) solution requires multiple sectors and a large amount of complementarity has somewhat dampened enthusiasm for this brand of business cycle theory. More recently, Jaimovich and Rebelo (2005) have developed a one-sector model which delivers the desired positive co-movement, but the assumptions of the model seem rather tenuous.

Our paper points to even more serious obstacles for this approach on the empirical front. Our results do support the notion that agents receive advance signals concerning future productivity, and provide data in addition to Beaudry and Portier's (2004b) stock prices that serve as an empirical counterpart of these information shocks. Further, the estimated impulse responses to such shocks display co-movement of the proper signs and not the troublesome contrarian movements predicted by the simplest neoclassical models. However, the movements in consumption, output, and other activity variables in the quarters directly following an information shock are rather negligible. Only after a long delay do consumption and output show substantial responses to a confidence shock. The behavior of consumption and output appear to track rather than anticipate movements in productivity. Consequently, information shocks alone are unlikely to produce major economic fluctuations.

[17] Earlier, Cochrane (1994) had foreshadowed the difficulties in producing positive co-movement of consumption, investment, and labor input in response to such information shocks.

References

Beaudry, Paul and Franck Portier. "An Exploration into Pigou's Theory of Cycles." *Journal of Monetary Economics*, September 2004a, 1183-1216.

_____"Stock Prices, News and Economic Fluctuations." *National Bureau of Economic Research Working Paper #10548*, 2004b.

_____. "When Can Changes in Expectations Cause Business Cycle Fluctuation in Neo-Classical Settings?" *National Bureau of Economic Research Working Paper #10776*, 2004c.

Blanchard, Olivier J. "Consumption and the Recession of 1990-1991." *American Economic Review*, May 1993, 270-274.

Campbell, John Y. and Tuomo Voulteenho. "Inflation Illusion and Stock Prices." *American Economic Review*, May 2004, 19-23.

Carroll, Christopher D., Fuhrer, Jeffrey C., and David W. Wilcox. "Does Consumer Sentiment Forecast Household Spending? If so, why?" *American Economic Review*, December 1994, 1397-1408.

Cochrane, John H. "Permanent and Transitory Components of GNP and Stock Prices" *Quarterly Journal of Economics*, February 1994a, 241-266.

_____. "Shocks." *Carnegie-Rochester Conference Series on Public Policy*, December 1994b, 295-364.

Croushore, Dean. "Do Consumer Confidence Indexes Help Forecast Consumer Spending in Real Time?" Mimeo, University of Richmond, January 2005.

Faust, Jon. "The Robustness of Identified VAR Conclusions About Money." *Carnegie-Rochester Conference Series on Public Policy*, December 1998, 207-244.

Francis, Neville, Owyang, Michael T., and Jennifer E. Roush. "A Flexible Finite-Horizon Identification of Technology Shocks." *Federal Reserve Bank of St. Louis Working Paper Series*, September 2005.

Gali, Jordi and Mark Gertler. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics*, 44, 2, 1999, 195-222.

Gordon, Robert J. and Ian Dew-Becker. "Where Did the Productivity Growth Go? Inflation Dynamics and the Distribution of Income." *National Bureau of Economic Research Working Paper #11842*, December 2005.

Hall, Robert E. "Macro Theory and the Recession of 1990-1991." *American Economic Review*, May 1993, 275-279.

_____. "Stochastic Implications of the Life Cycle Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy*, December 1978, 971-987.

_____. "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*, 1, 2001, 175-207.

Jaimovich, Nir and Sergio Rebelo. "Can News About the Future Drive Business Cycles?" Mimeo, Northwestern University, 2005.

Leeper, Eric M. "Consumer Attitudes: King for a Day." *Federal Reserve Bank of Atlanta Economic Review*, July-August 1992, 1-15.

Matsusaka, 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.

Modigliani, Franco and Richard Cohn. "Inflation, Rational Valuation, and the Market." *Financial Analysts Journal* 35(3), 1979, 24-44.

Piazzesi, Monika and Martin Schneider. "Equilibrium Yield Curves". Forthcoming in *NBER Macroeconomics Annual 2006.*

Uhlig, Harald. "What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure." *CEPR Discussion Paper #2137*, May 1999.

Appendix A: Sketch of Agnostic Identification Procedures

(A) Max Forecast Error Variance

Letting x be the vector of endogenous variables, suppose that an arbitrary structural moving average representation can be written as:

$$x_t = C(L)GG^{-1}e_t$$

Where G is any matrix satisfying $GG' = \Sigma$, where Σ is the variance-covariance matrix of reduced form innovations. For any such orthogonal decomposition, the fraction of the forecast error variance of the i^{th} variable to the j^{th} structural shock at horizon h is given by:

$$\omega_{ij}(h) = \frac{e_i' \left[\sum_{t=0}^{h-1} C_t G e_j e_j' G' C_t' \right] e_i}{e_i' \left[\sum_{t=0}^{h-1} C_t \Sigma C_t' \right] e_i}$$

Here e_i and e_j are selection vectors with one in the i^{th} or j^{th} places and zeros elsewhere.

For any orthonormal matrix D, the structural moving average representation can equivalently be written as:

$$x_t = C(L)GDD'G^{-1}e_t$$

The above gives us the entire space of orthogonal identifications. After applying this linear rotation, the expression for the forecast error variance is now given by:

$$\omega_{ij}(h) = \frac{e_i' \left[\sum_{t=0}^{h-1} C_t G \alpha \alpha' G' C_t' \right] e_i}{e_i' \left[\sum_{t=0}^{h-1} C_t \Sigma C_t' \right] e_i}$$

Where α is the j^{th} column of the orthonormal matrix D. The maximum forecast error variance identification is accomplished by choosing α to maximize the contribution of the j^{th} shock to the forecast error variance of the i^{th} variable at some arbitrarily chosen horizon, subject to the restriction that $\alpha'\alpha = 1$. The j^{th} structural shock is then identified as $\alpha' G^{-1} e_t$ and the structural moving average coefficients to the j^{th} shock are $C(L)G\alpha$. That the vector α is restricted to have unit length maintains the scaling of the variance of the structural shock to unity.

To identify the "best case" for an information shock, we adhere to the following procedure:

(1) Estimate the reduced form

(2) Perform a Choleski decomposition to attain an arbitrary orthogonalization of the reduced form, G
(3) Create several thousand vectors α satisfying the unit length assumption with a zero in the i^{th} position, where i corresponds to the position of output in the VAR
(4) Compute the forecast error variance decomposition of output for each rotation vector α
(5) The "best case" for the information shock is then defined as that rotation which explains the maximal portion of the forecast error variance of output at a horizon of eight years

(B) Impulse Response Shape Restriction

Similar to the max forecast error variance identification, suppose that an arbitrary structural moving average representation can be written as:

$$x_t = C(L)GG^{-1}e_t$$

Where G is any arbitrary orthogonalization of the reduced form. For any orthonormal matrix D, the moving average representation can alternatively be written:

$$x_t = C(L)GDD'G^{-1}e_t$$

Taking the j^{th} shock as the shock of interest, the moving average representation can then be written:

$$x_t = C_tG\alpha\alpha'G'C_t'e_{tj}$$

Where α is the j^{th} column of the matrix D. $C_tG\alpha$ then gives the moving average coefficient of all variables in the VAR to shock j while $\alpha'G'C_t'e_{tj}$ defines structural shock j.

To identify the "best case" for a transitory shock, we adhere to the following procedure:

(1) Estimate the reduced form
(2) Perform a Choleski decomposition to attain an arbitrary orthogonalization of the reduced form, G
(3) Create several thousand vectors α satisfying the unit length assumption (as opposed the max forecast error variance case, we no longer require that i^{th} element of α be zero)
(4) Keep all rotations which satisfy the following conditions
    a. The response of output on impact is non-negative
    b. The response of output at one year is positive
    c. The response of output at one year is greater than the response at four years
    d. The response of output at eight years is within 0.2 percent of zero

(5) The “best case” for a transitory shock is then defined as that rotation yielding the maximal effect on output at one year of those rotations satisfying the requirements in (4).

Appendix B: Tables

Table 1
Cross-Correlogram of E5Y and HP Detrended GDP

Cross Correlogram of E5Y and LRGDP-HPTREND01

Date: 11/24/05 Time: 22:51
Sample: 1960:1 2005:4
Included observations: 181
Correlations are asymptotically consistent approximations

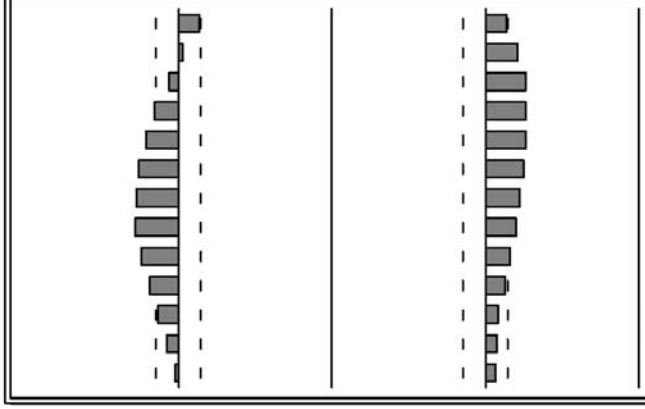

E5Y,LRGDP-HPTREND01	E5Y,LRGDP-HPTREND01	i	lag	lead
		0	0.1447	0.1447
		1	0.0285	0.2210
		2	-0.0647	0.2707
		3	-0.1592	0.2777
		4	-0.2211	0.2726
		5	-0.2647	0.2547
		6	-0.2793	0.2340
		7	-0.2876	0.2080
		8	-0.2530	0.1642
		9	-0.1896	0.1314
		10	-0.1336	0.0887
		11	-0.0808	0.0789
		12	-0.0228	0.0721

Table 2
Estimates from VAR with GDP, E5Y, and Non-Durable Consumption

	Dependent Variable		
	GDP	**E5Y**	**C**
GDP(-1)	1.18590*** (0.08118)	-65.29044 (84.6742)	0.175881*** (0.06669)
GDP(-2)	-0.154147* (0.08329)	75.97706 (86.8752)	-0.167825** (0.06842)
E5Y(-1)	0.000154** (0.000075)	0.811427*** (0.07775)	0.00148** (0.000061)
E5Y(-2)	-0.0000841 (0.000072)	0.097227 (0.07539)	-0.0000765 (0.000059)
C(-1)	0.231661** (0.10211)	189.8963* (106.502)	1.000133*** (0.08388)
C(-2)	-0.189288* (0.10453)	-203.3877* (109.028)	-0.009382 (0.08587)
Constant	-0.030584 (0.03493)	20.60121 (36.4294)	0.004470 (0.02869)
Adjusted R^2	0.999633	0.836197	0.999618

* Significant at 10 % level
** Significant at 5 % level
*** Significant at 1 % level

Chi-Square Statistics for Block Exogeneity:

	Chi-Square Statistic	*p* value
E5Y → GDP	7.092	0.028
GDP → E5Y	0.575	0.750
C → GDP	6.440	0.039
GDP → C	5.216	0.073
E5Y → C	11.014	0.004
C → E5Y	3.376	0.185

Table 3
Long Horizon Regressions

$$100 \times (\ln C_{t+h} - \ln C_t) = \alpha + \beta E5Y_t + \varepsilon_t$$

Horizon	E5Y	R^2
1 quarter	0.011	.09
	(0.003)	
1 year	0.032	.15
	(0.011)	
2 years	0.043	.13
	(0.017)	
5 years	0.078	.19
	(0.021)	
7 years	0.083	.16
	(0.026)	

Notes: Sample period is fixed at 1960.1 – 1998.2 (quarterly) to account for the seven years of leading observations. Consumption is non-durable consumption expenditures. Robust (Newey-West) standard errors are in parentheses and *p*-values are in square brackets

Table 4
Regressions of Confidence Innovations on Selected News Heard

News Heard	Coefficient
Fav. Employment	0.320
	(0.101)
Un. Employment	-0.049
	(0.042)
Fav. Prices	1.028
	(0.520)
Un. Prices	-0.393
	(0.164)

Adjusted $R^2 = 0.12$

Notes: Sample period is 1960.1 – 2005.1 (quarterly). The dependent variable is the innovation in E5Y from a three variable system including consumption and GDP with two lags. The independent variables are defined as the percentage of all respondents in a given quarter reporting having heard news in that particular category.

Appendix C: Figures

Figure 1
Bivariate VARs with Output and non-durable Consumption

(A) Ordering: C→Y

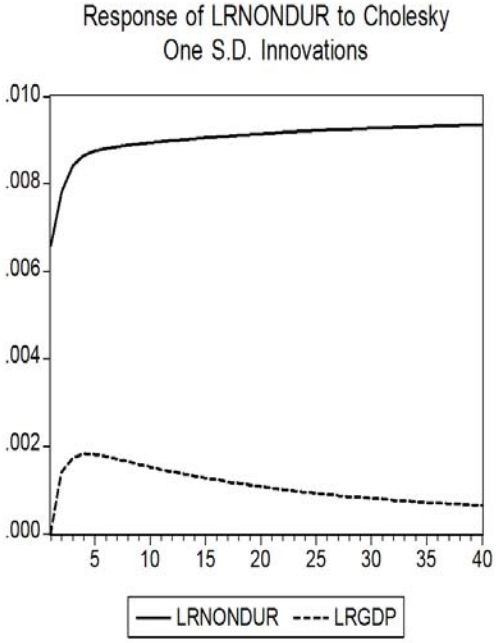

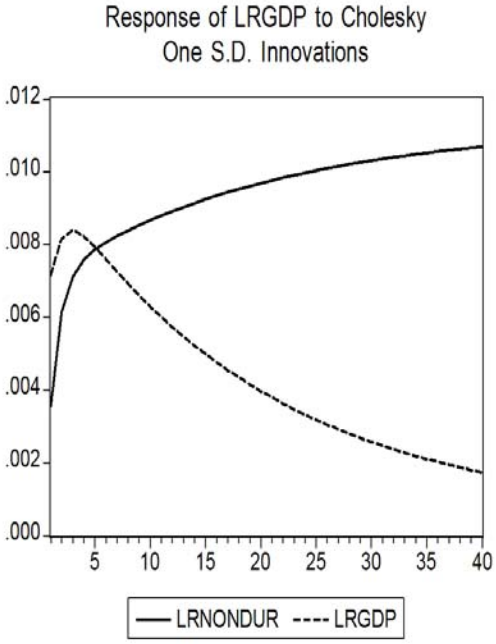

(B) Ordering: Y→C

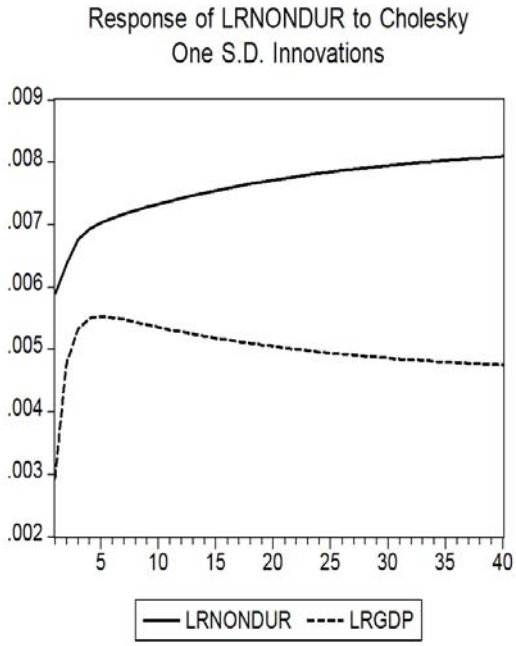

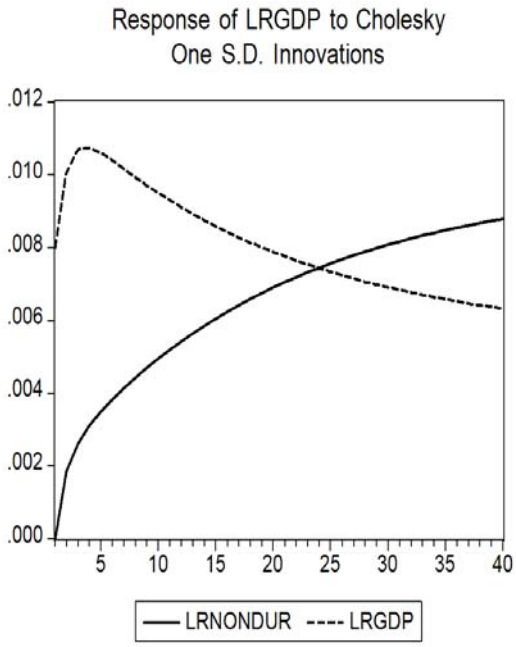

Figure 2
E5Y and E12M Across Time

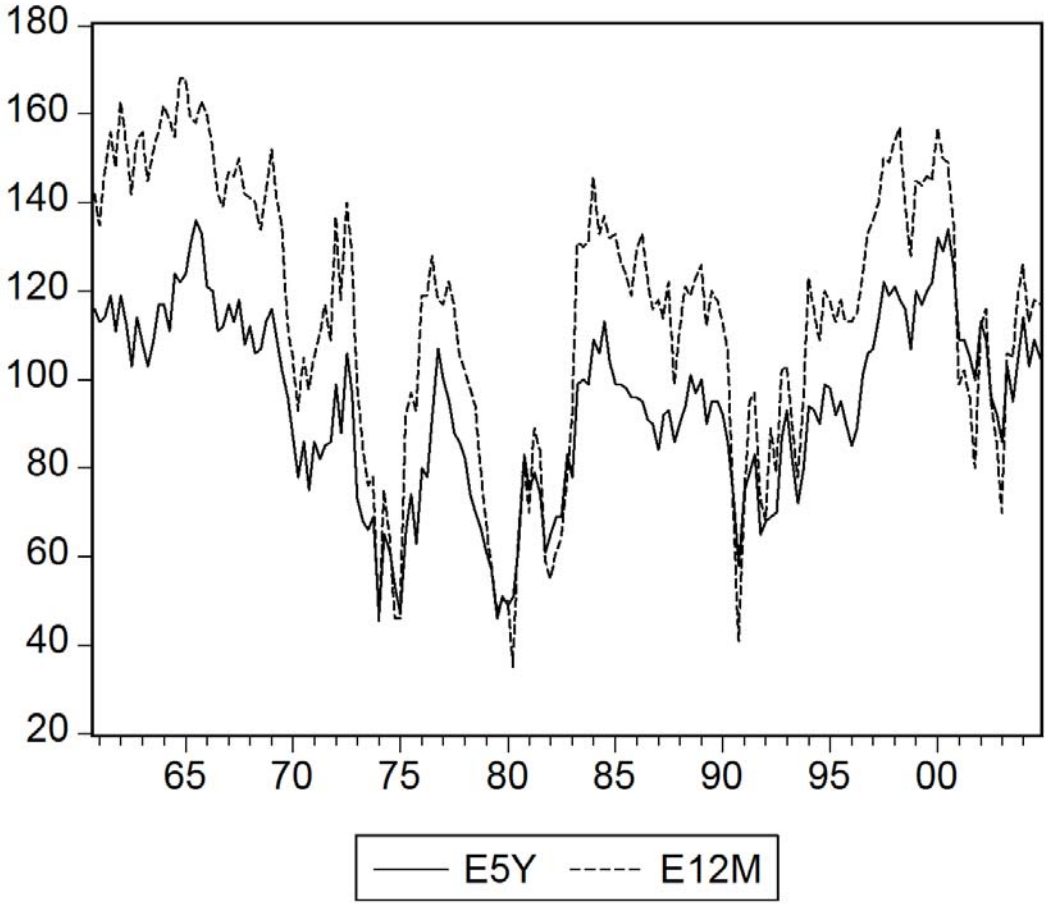

Figure 3
Trivariate VARs with C, E5Y, and Y
Impulse Responses

(A) Ordering: E5Y→C→Y

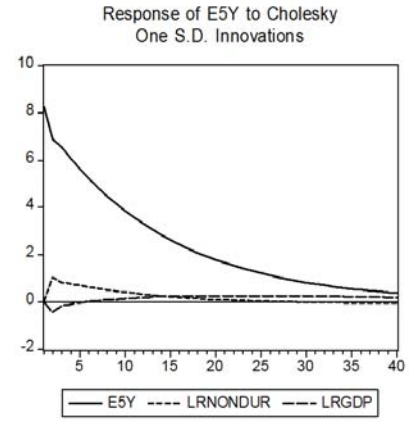

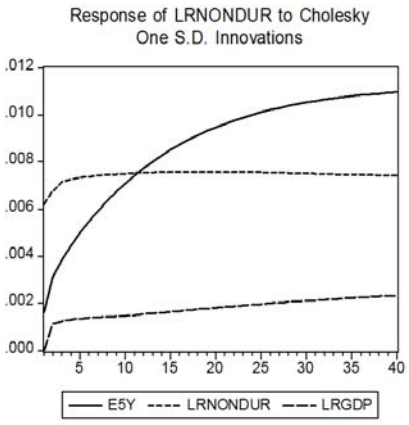

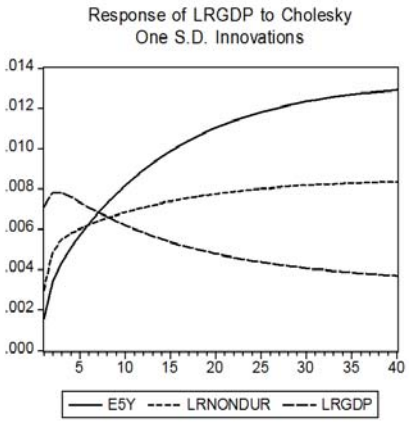

(B) Ordering: C→E5Y→Y

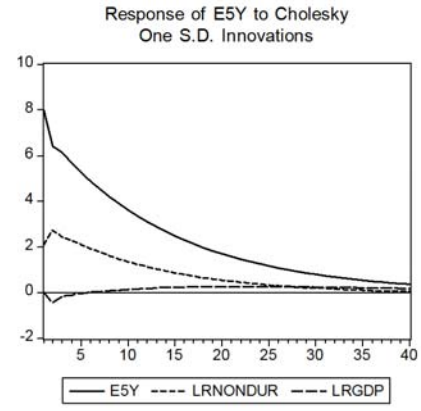

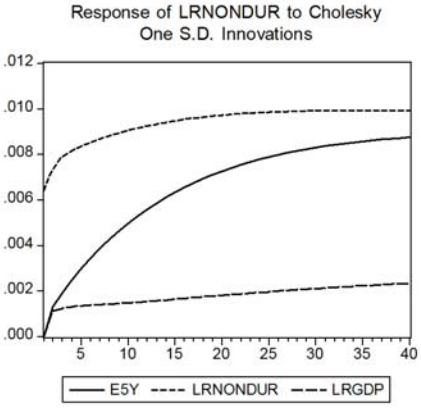

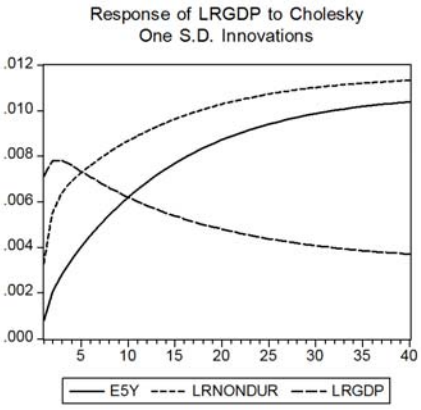

Figure 4
Trivariate VARs with C, E5Y, and Y
Variance Decompositions

(A) Ordering: E5Y→C→Y

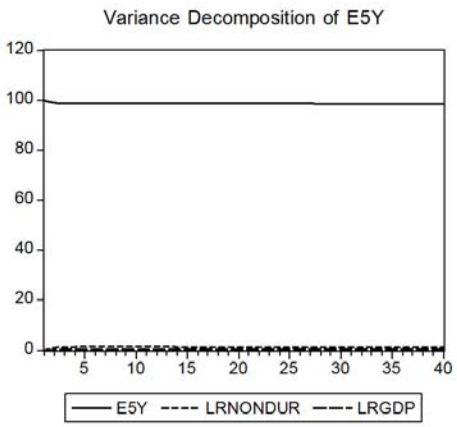

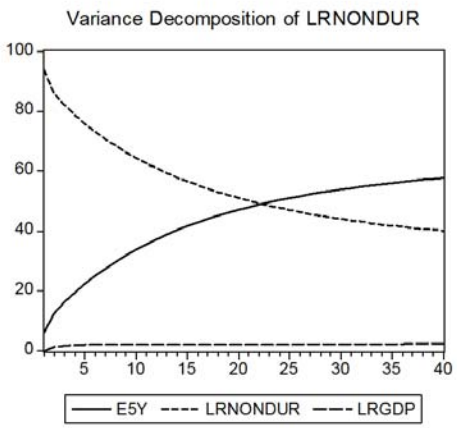

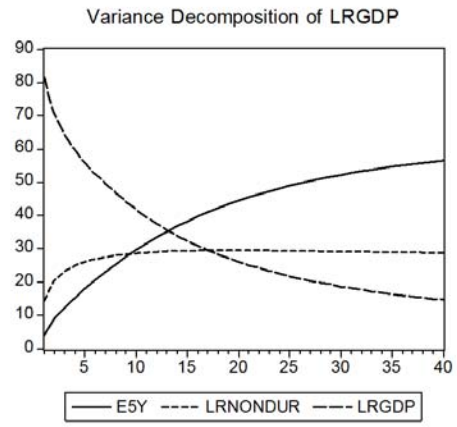

(B) Ordering: C→E5Y→Y

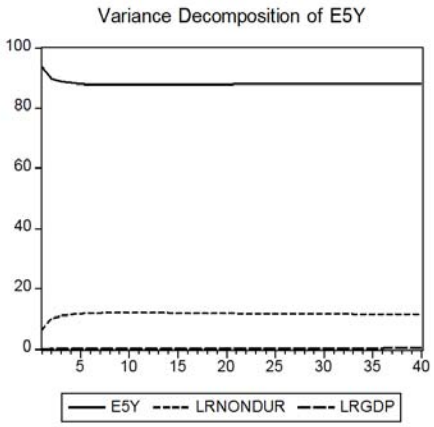

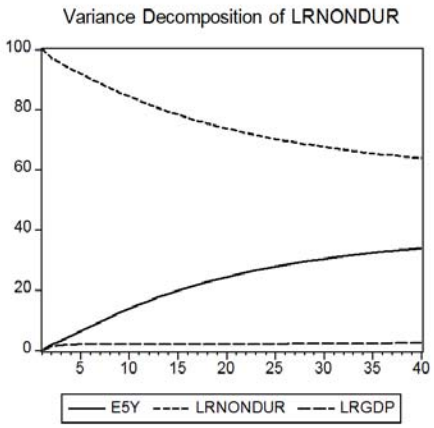

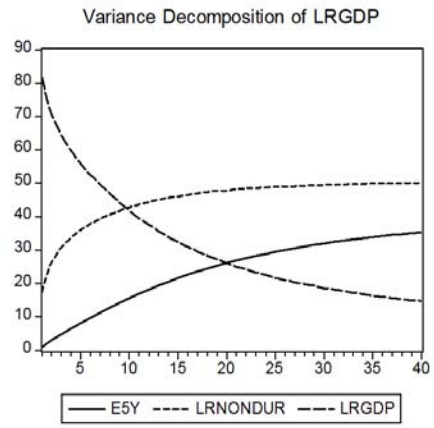

Figure 5 (A)
Impulse Responses from VAR with ordering GDP→E5Y→C

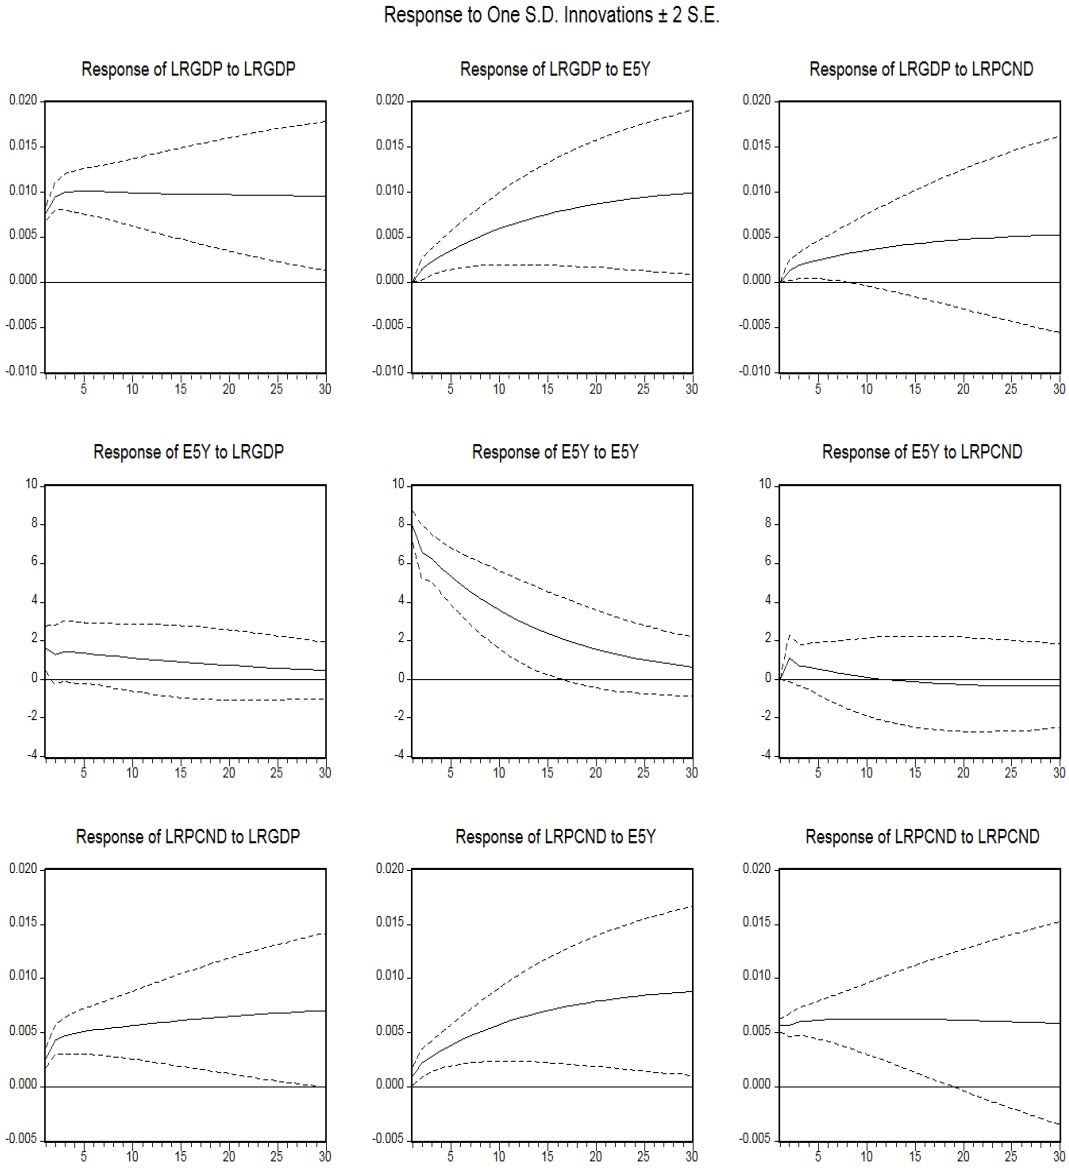

Figure 5 (B)
Variance Decomposition with ordering GDP, E5Y, Non-Durable C

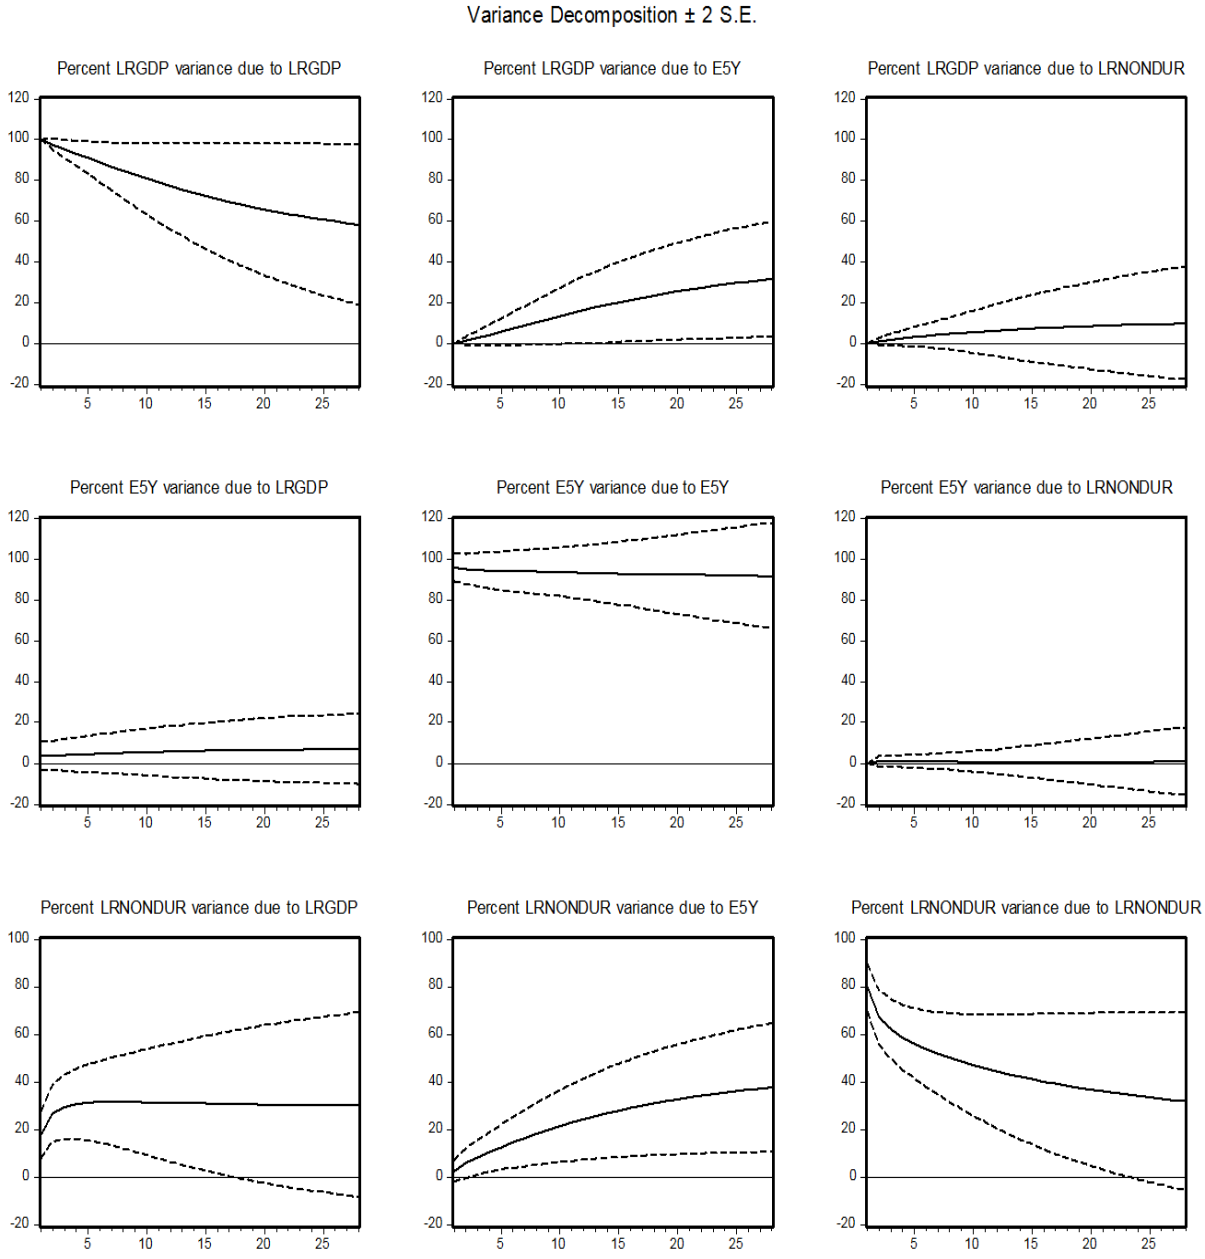

Figure 6 (A)
Impulse Responses from VAR with ordering GDP→E5Y→Durable C

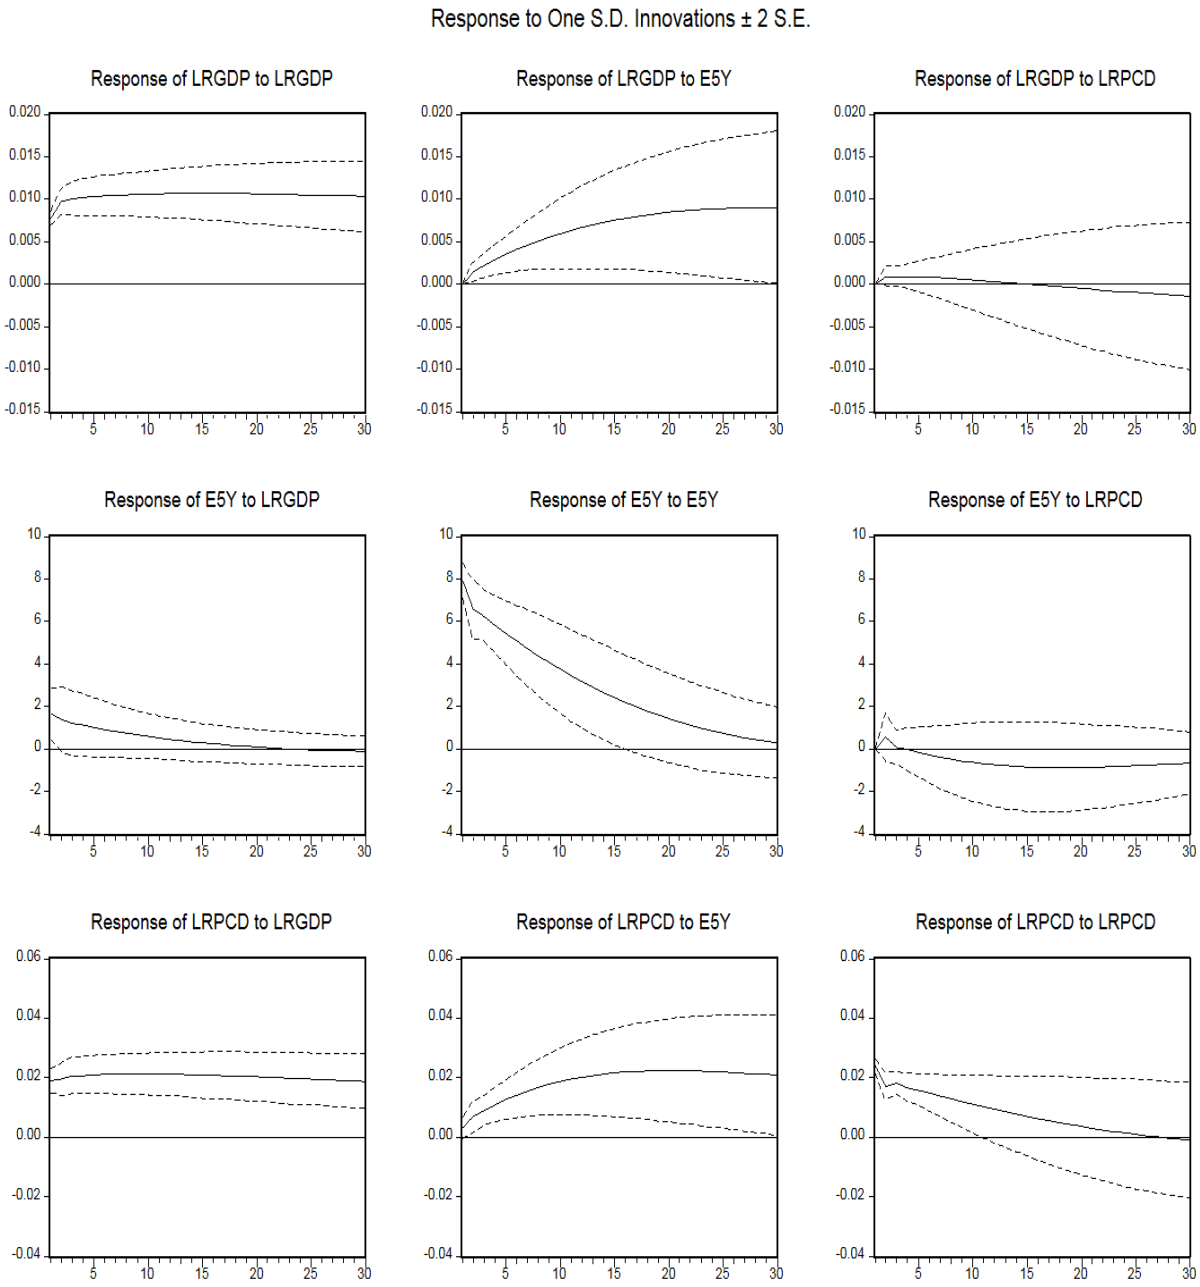

Figure 6 (B)
Impulse Responses from VAR with ordering GDP→PFE→C

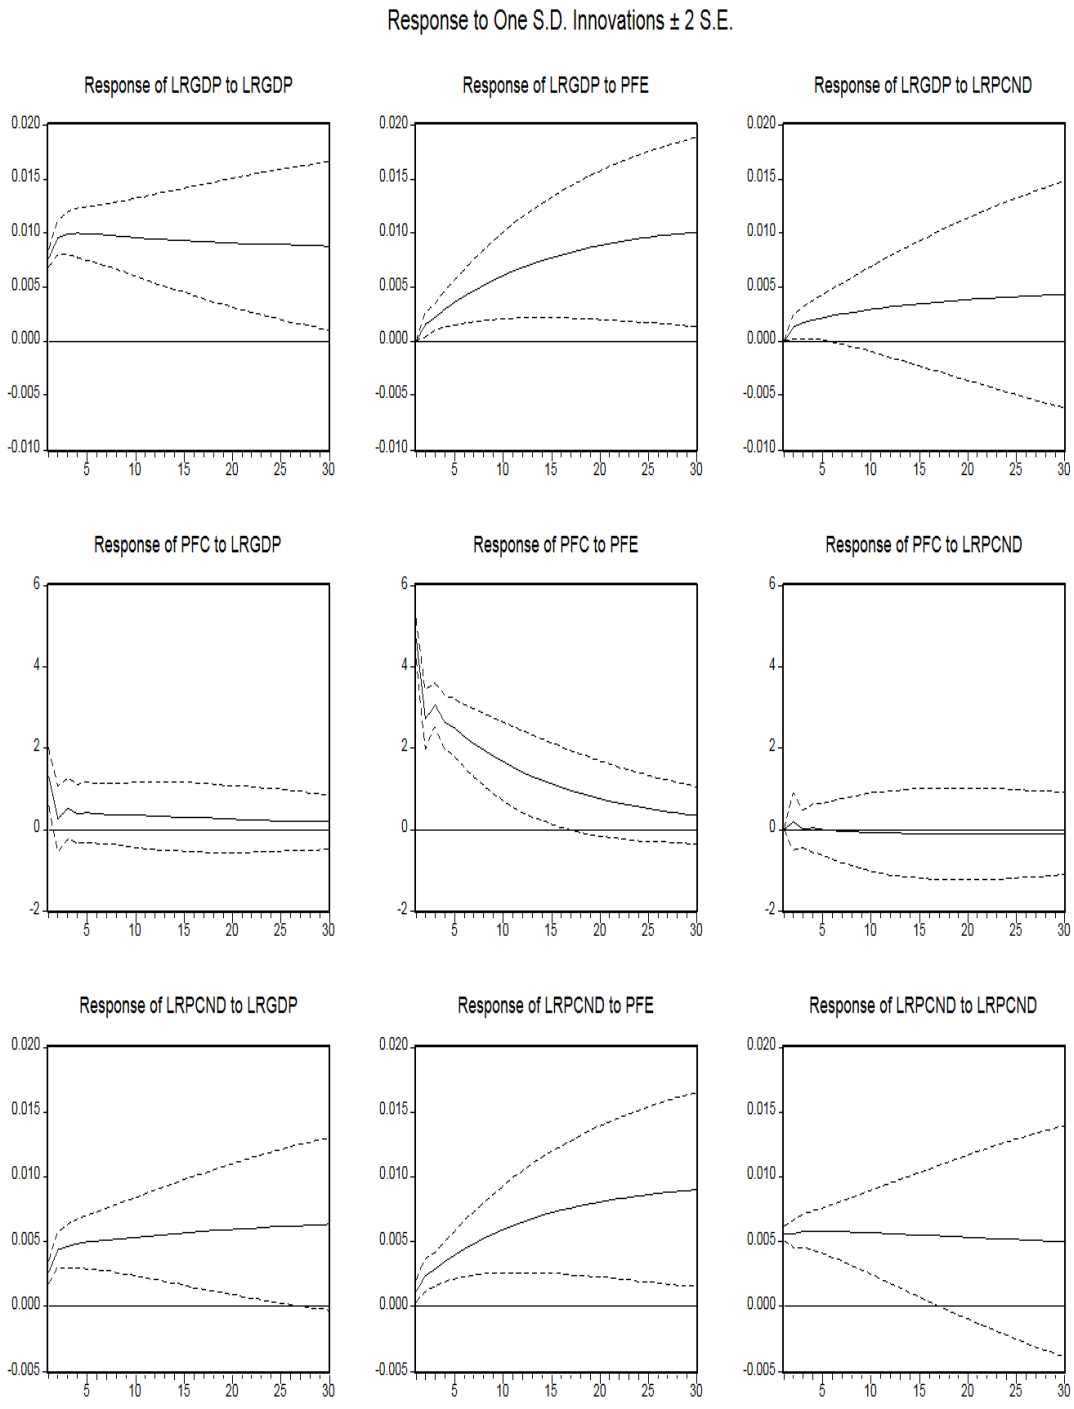

Figure 6 (C)
Impulse Responses from VAR with ordering GDP→ICS→C

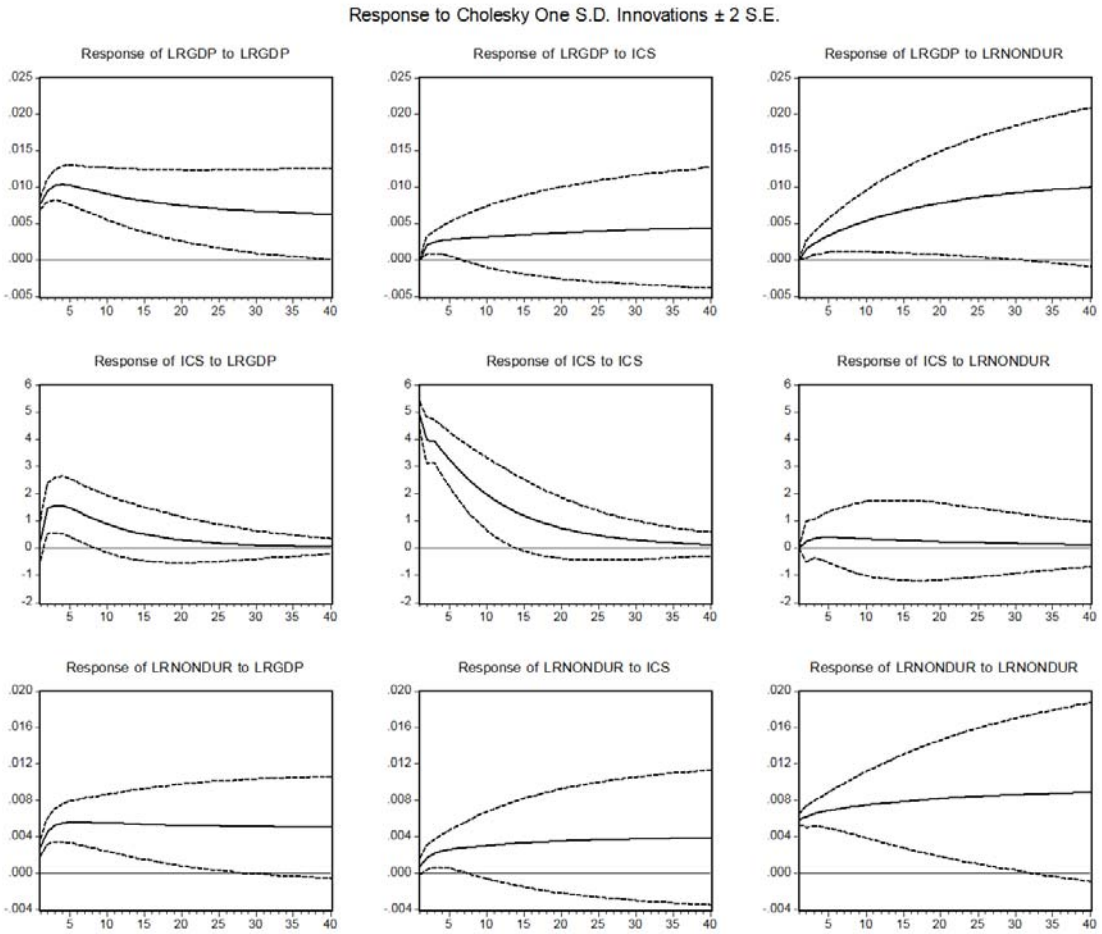

Figure 7
Bivariate VAR with labor productivity and E5Y
Ordering: productivity, confidence

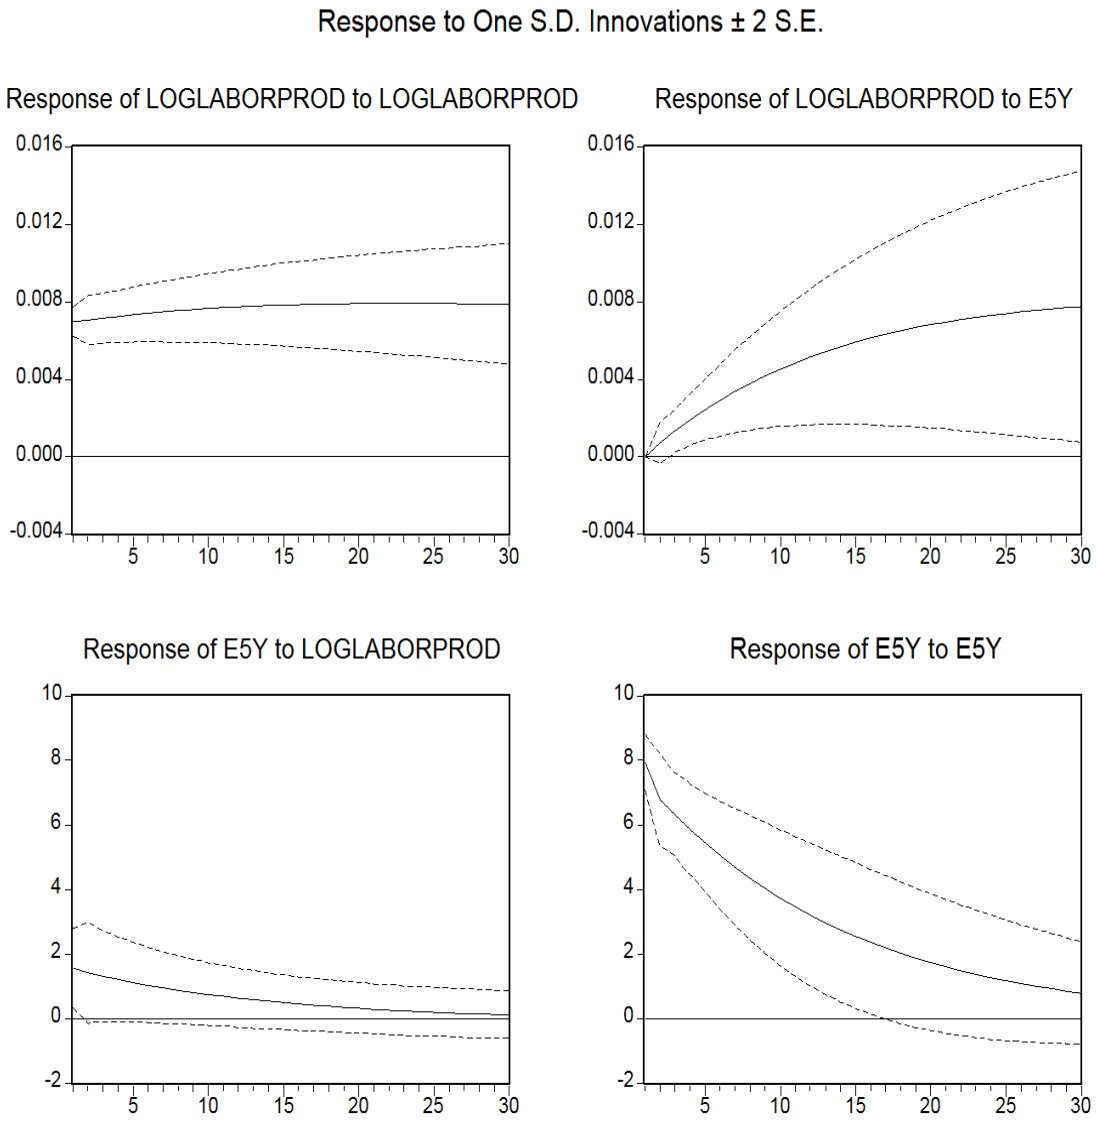

Notes: For these and future figures including productivity, productivity is defined as the output to civilian employment ratio. Systems using the BLS measure of output per hours or standard Solow residuals (at an annual frequency) produce similar results.

Figure 8
Historical Simulations

(A) Contribution of Information Shock: Historical Simulations

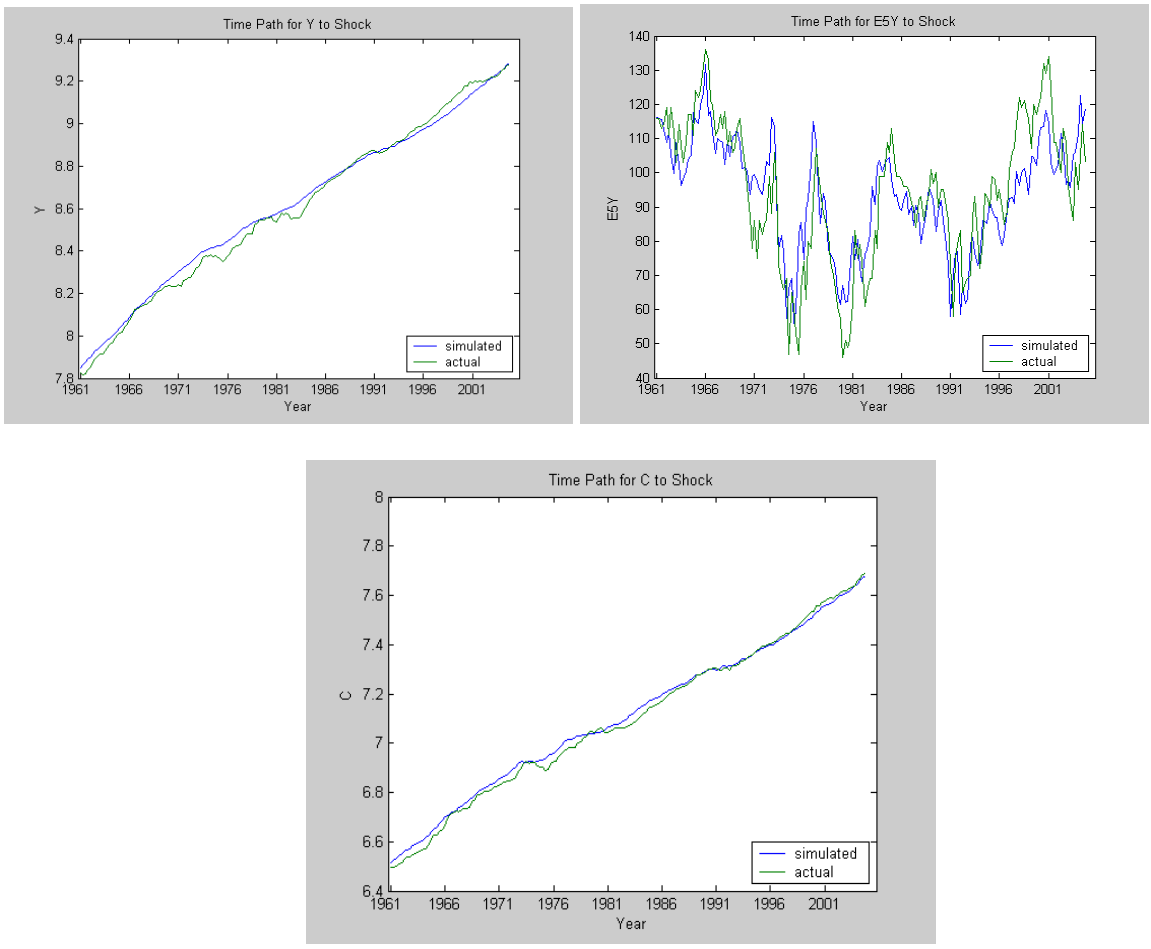

Note: The above simulation is from the identification of a shock satisfying the best case scenario for an information shock in a three variable system with output, E5Y, and non-durable consumption. The shock is identified as that shock explaining the most of the forecast error variance of output at a horizon of eight years among all shocks contemporaneously orthogonal to output. The actual series is given by the more lightly shaded line while the simulated is given by the darker line.

Figure 8 (continued)
Historical Simulations

(B) Contribution of Transitory Shock: Historical Simulations

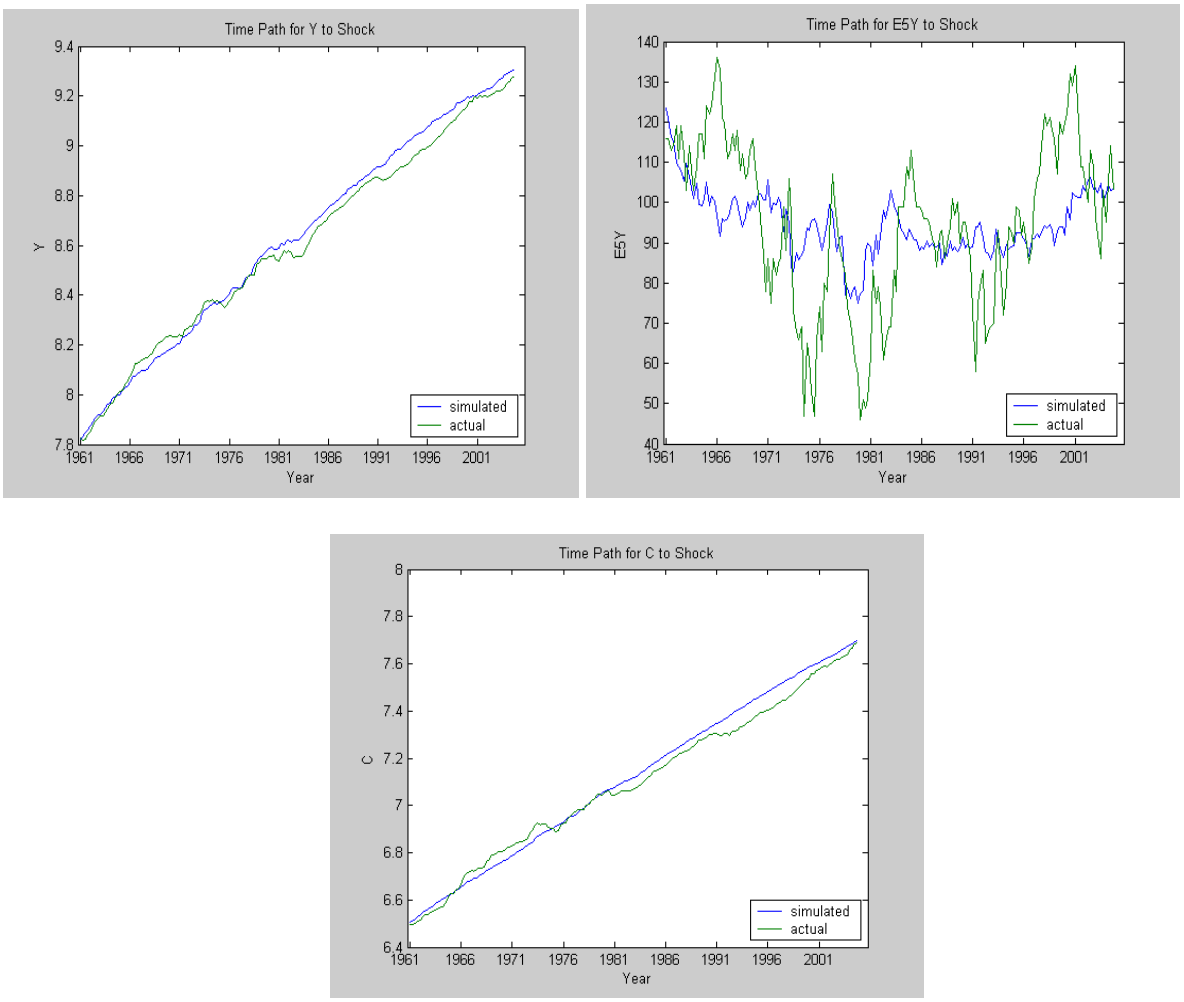

Note: The above simulation is from the identification of a shock satisfying the best case scenario for a transitory shock in a three variable system with output, E5Y, and non-durable consumption. The shock is identified as that shock explaining the most of the forecast error variance of output at a horizon of one year among all shocks which have a transitory effect on output. The actual series is given by the more lightly shaded line while the simulated is given by the darker line.

Figure 9
Spike Plots of Responses in News Heard Categories: Favorable Employment, Unfavorable Employment, Unfavorable Energy Crisis, Unfavorable Prices, and Favorable Stocks

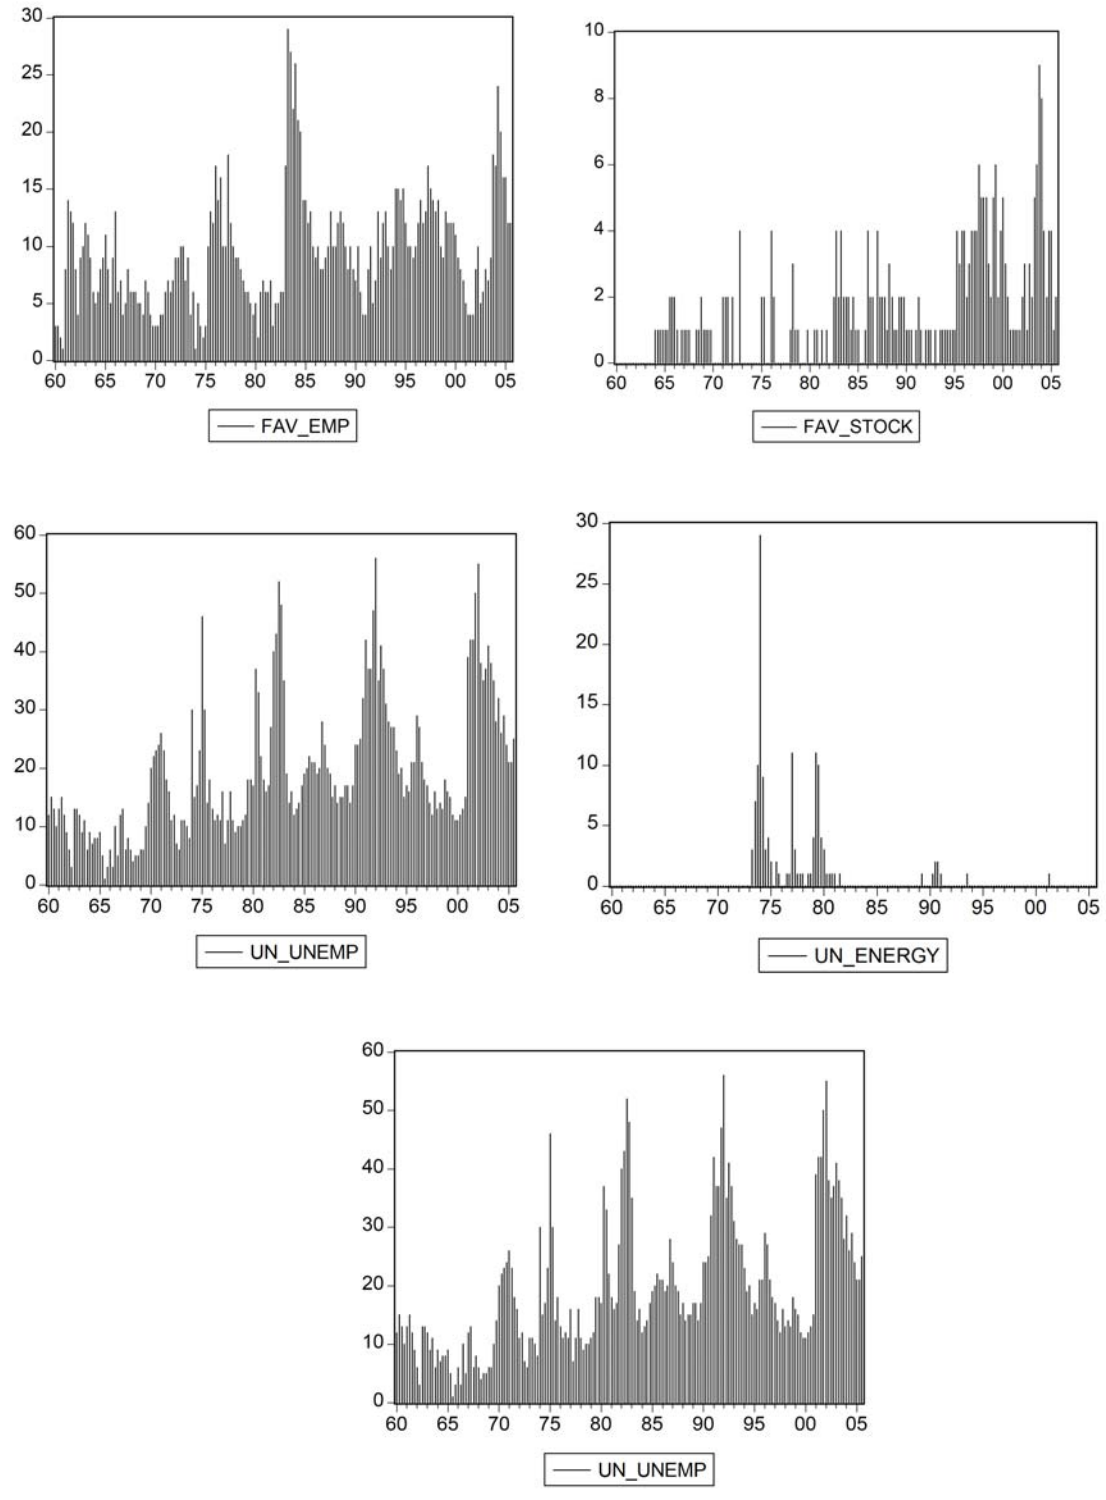

Figure 10
Inflation and E5Y Across Time

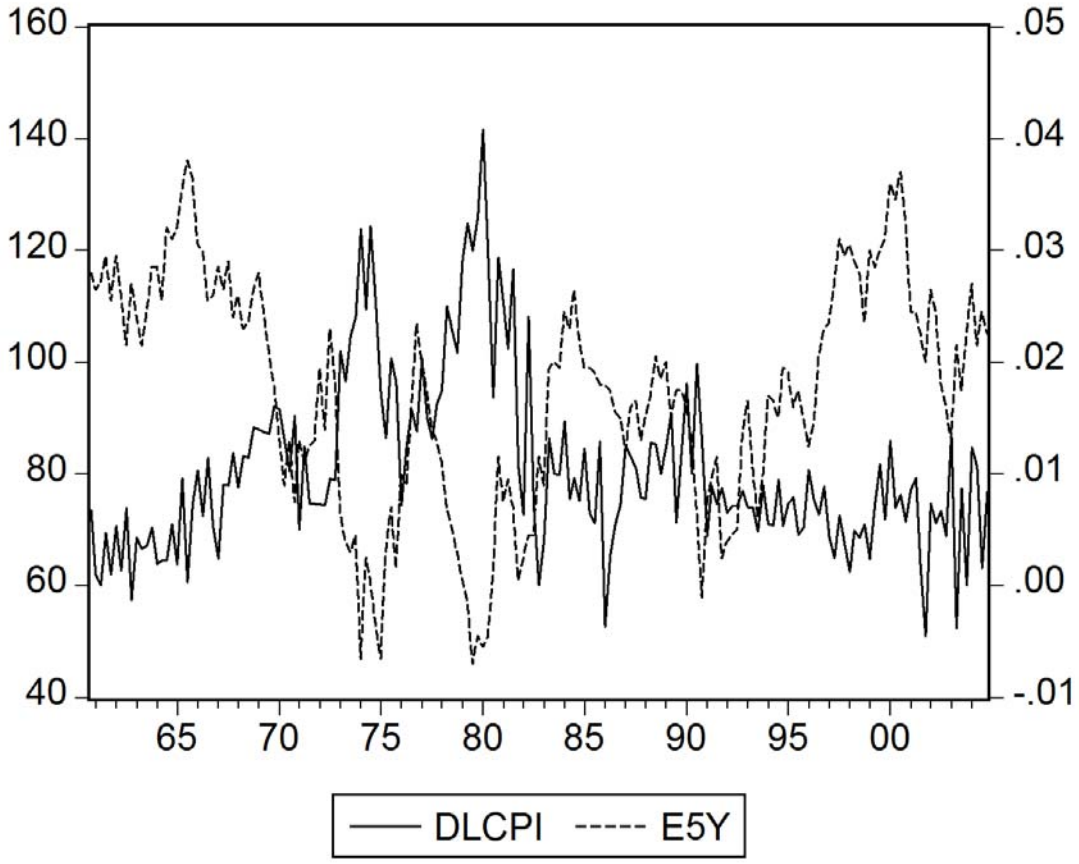

Figure 11
Impulse Responses with ordering: Y→Inflation→E5Y→C

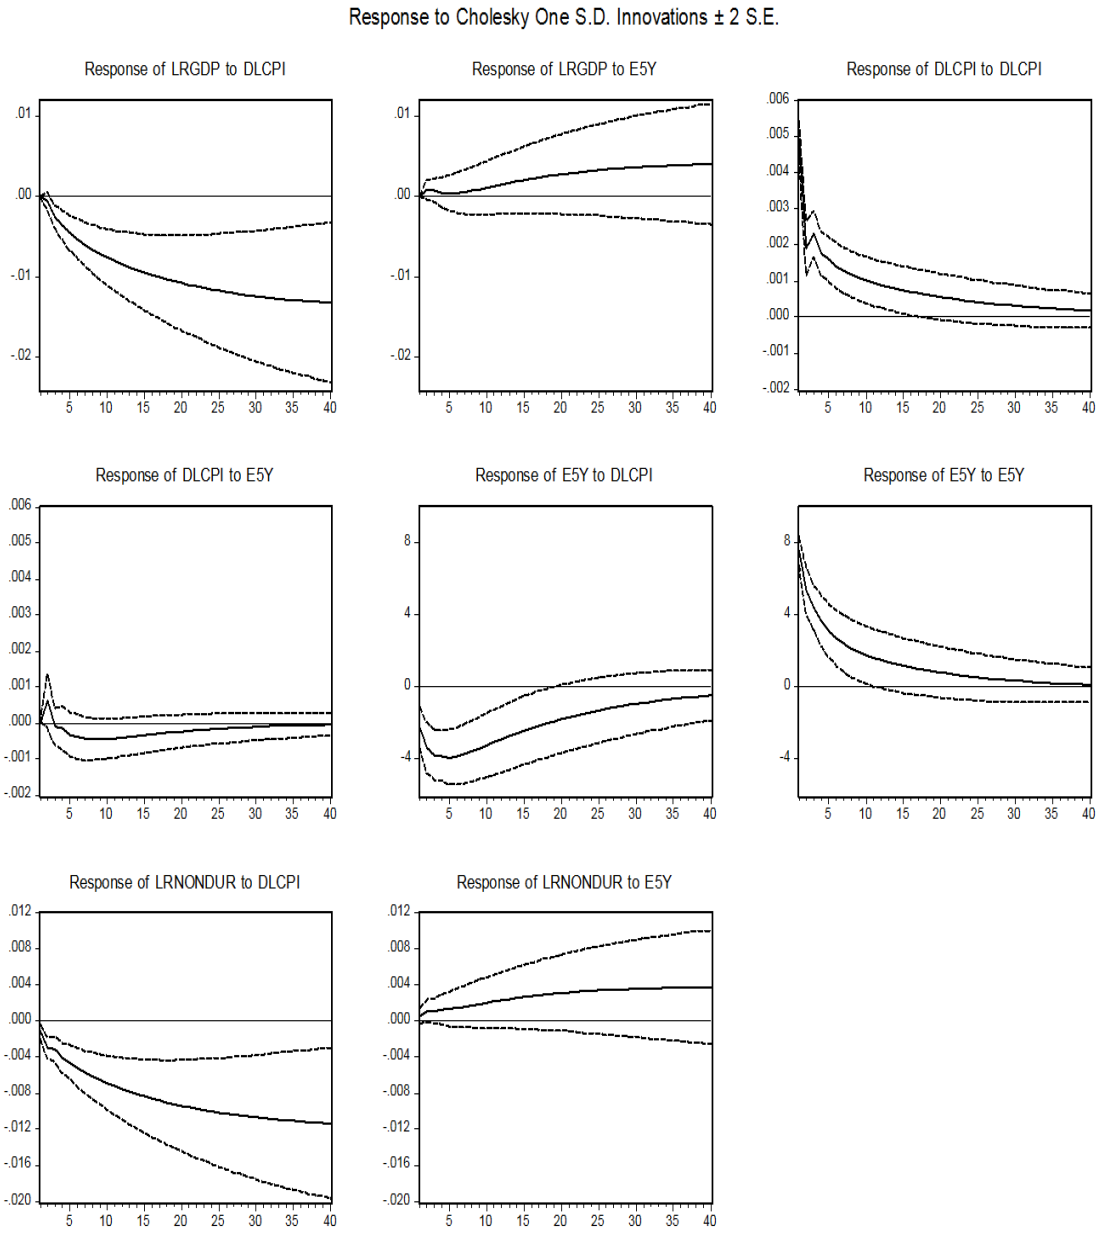

Figure 12
Impulse Responses with ordering: Y→E5Y→Inflation→C

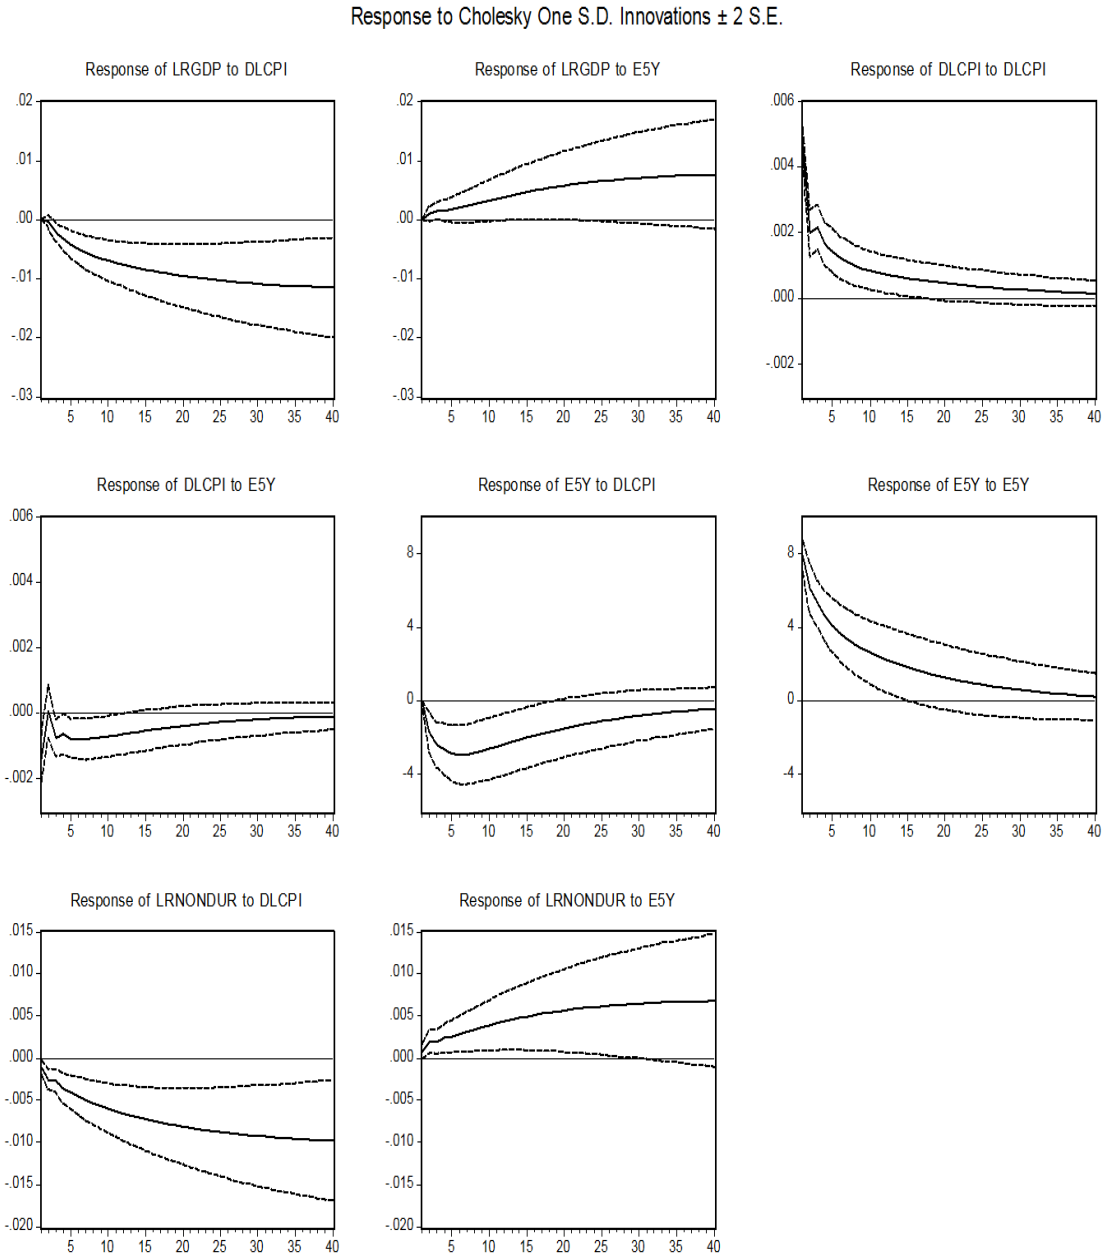

Figure 13
Historical Simulations in Four Variable Case
Included Variables: Output, inflation, E5Y, and non-durable consumption

(A) Contribution of Information Shock: Historical Simulations

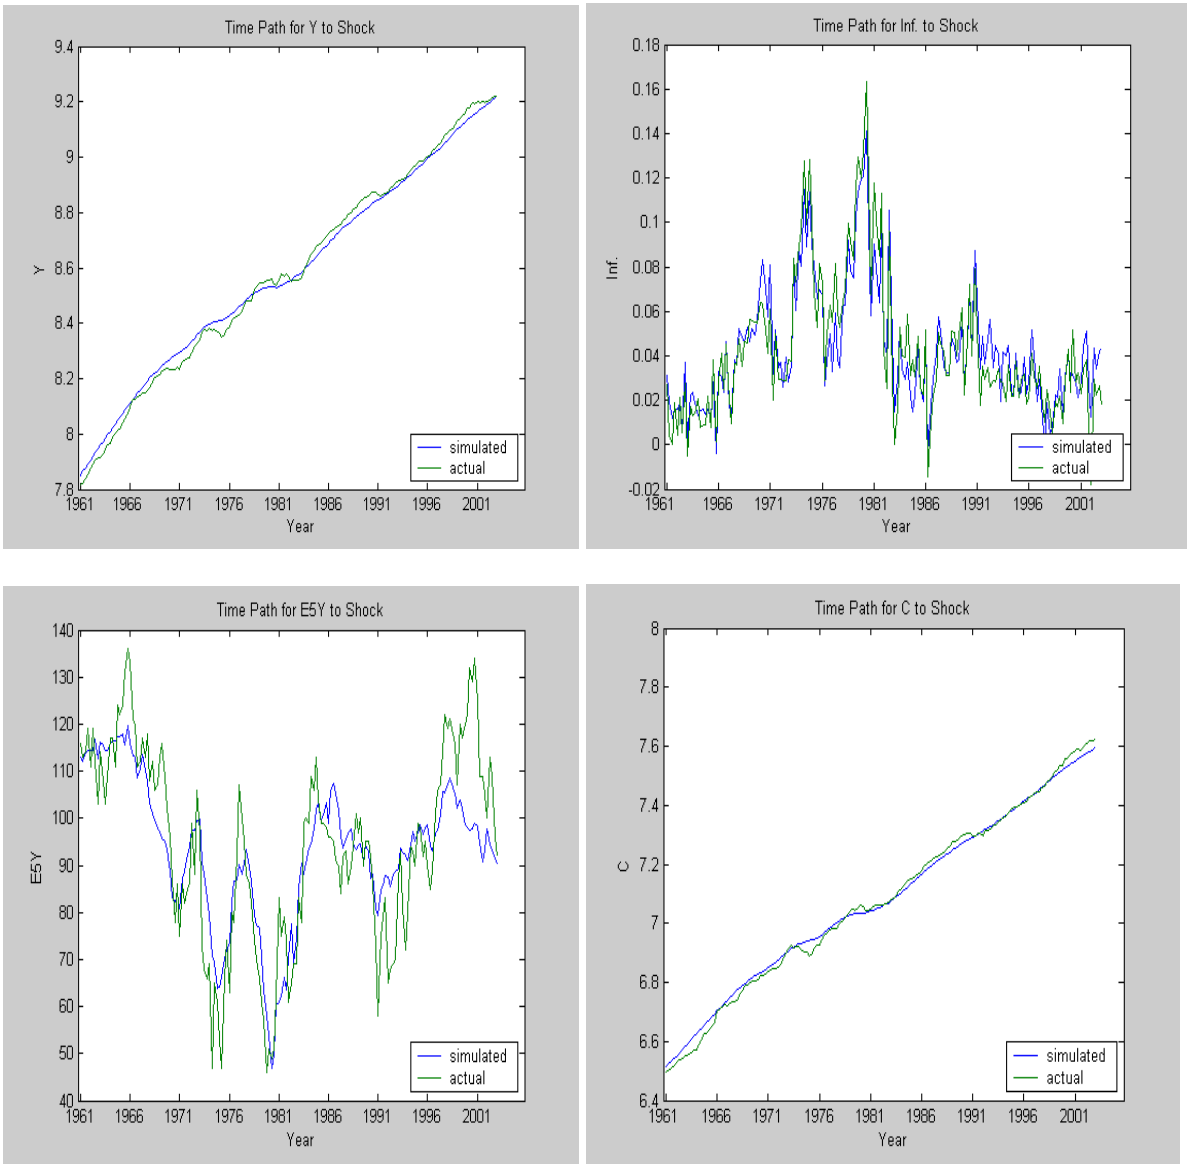

Note: The above simulation is from the identification of a shock satisfying the best case scenario for an information shock in a four variable system with output, consumer price inflation, E5Y, and non-durable consumption. The shock is identified as that shock explaining the most of the forecast error variance of output at a horizon of eight years among all shocks contemporaneously orthogonal to output. The actual series is given by the more lightly shaded line while the simulated is given by the darker line.

Figure 13 (continued)
Historical Simulations in Four Variable Case
Included Variables: Output, inflation, E5Y, and non-durable consumption

(B) Contribution of Transitory Shock: Historical Simulations

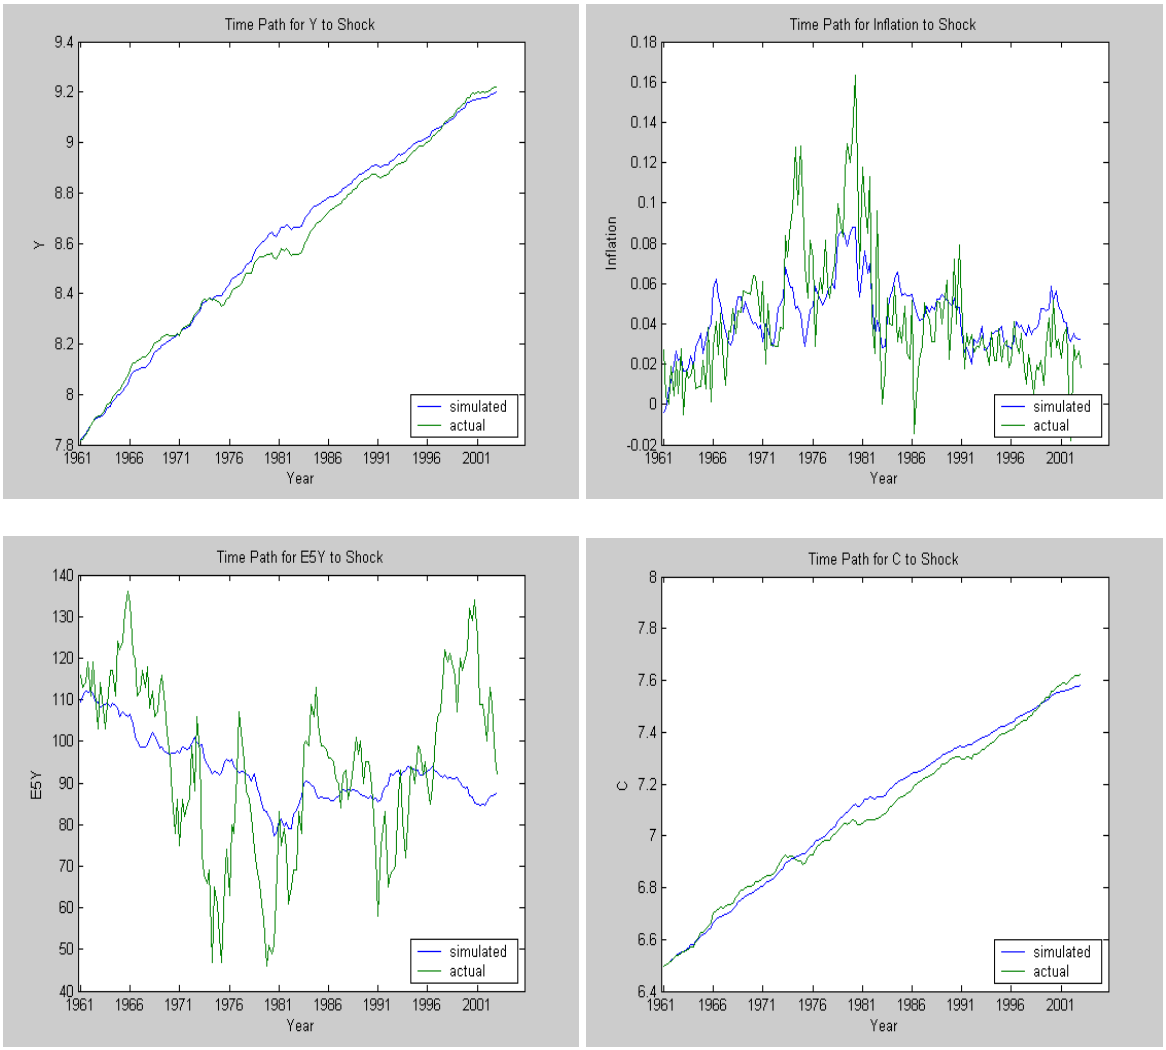

Note: The above simulation is from the identification of a shock satisfying the best case scenario for a transitory shock in a four variable system with output, consumer price inflation, E5Y, and non-durable consumption. The shock is identified as that shock explaining the most of the forecast error variance of output at a horizon of one year among all shocks which have a transitory effect on output. The actual series is given by the more lightly shaded line while the simulated is given by the darker line.

Figure 14
S & P 500 and E5Y Across Time

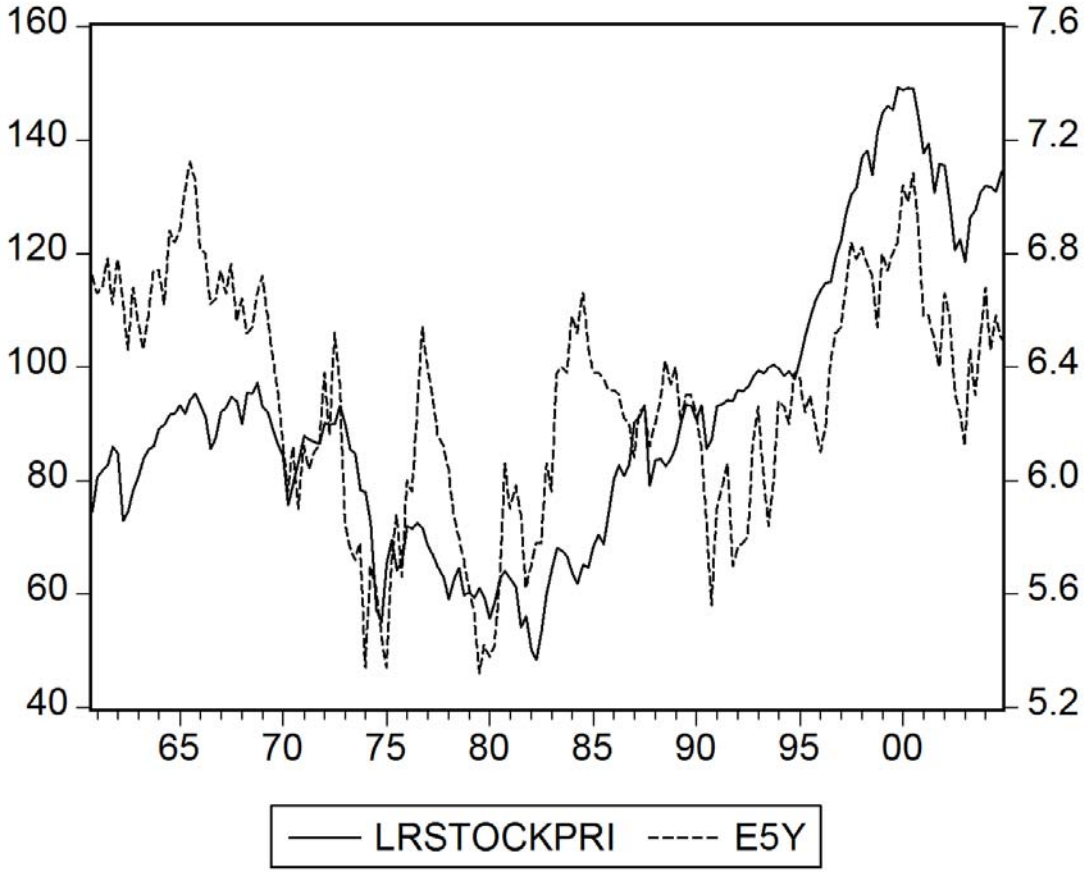

Figure 15
Impulse Responses with ordering: Y→S & P 500→E5Y→C

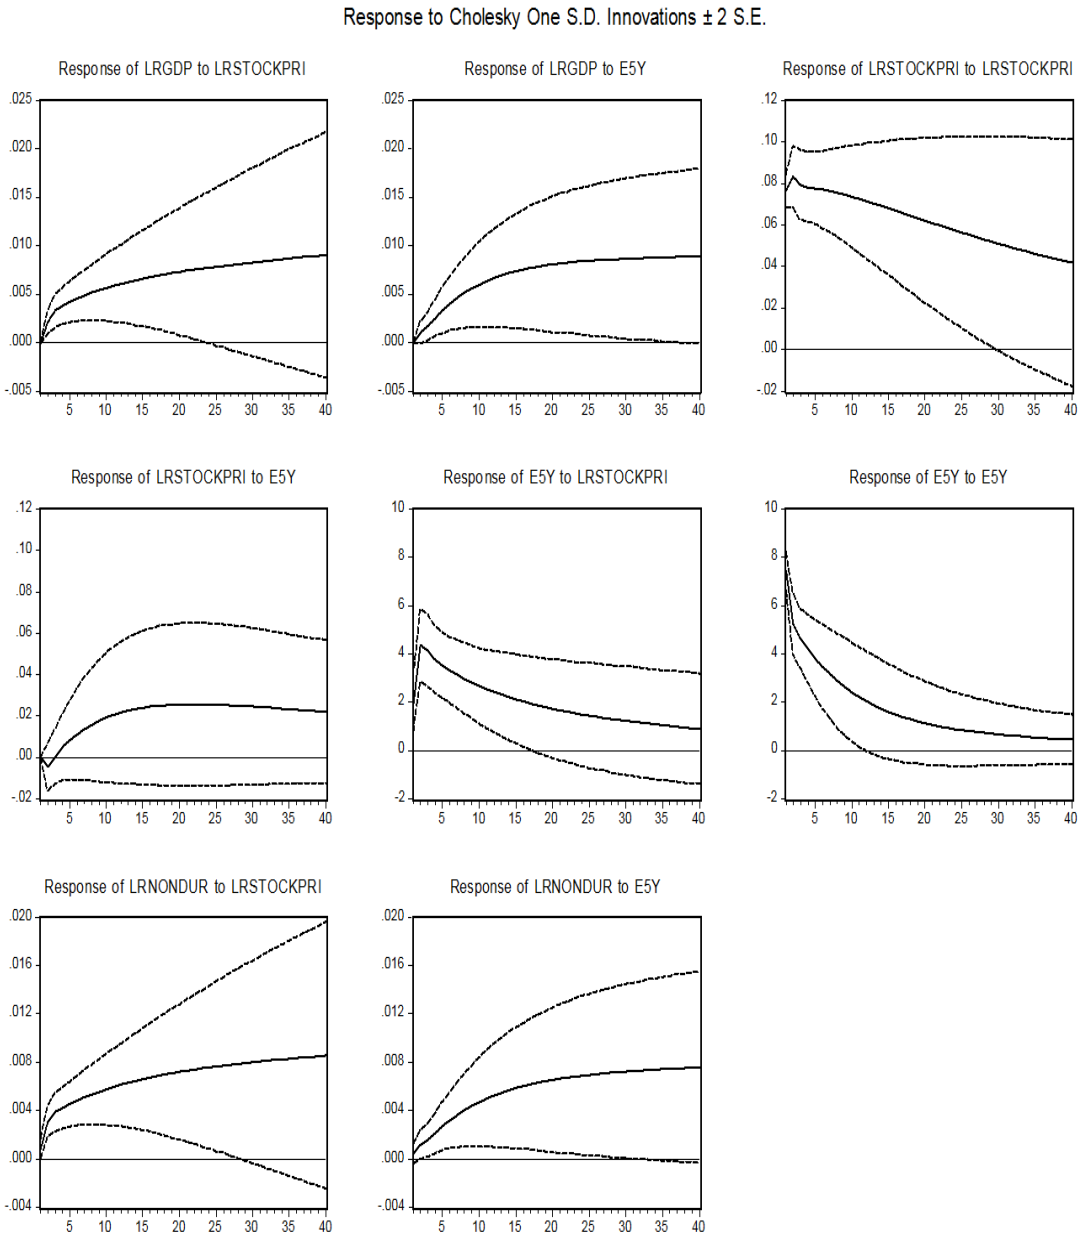

Figure 16
Impulse Responses with ordering: Y→Inflation→S & P 500→E5Y→C

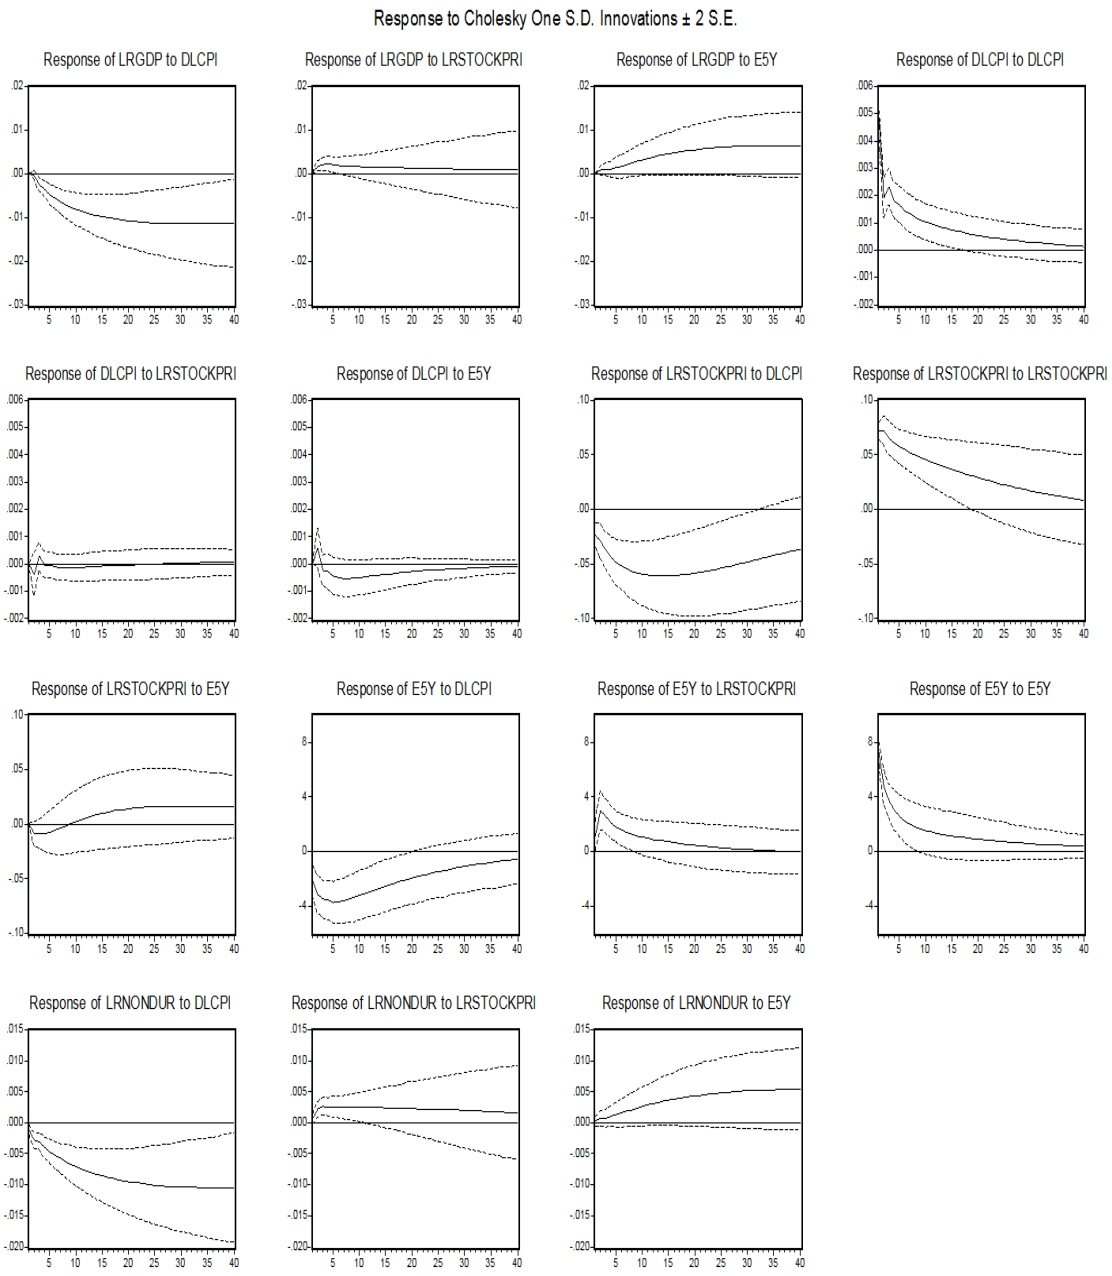

Figure 17
E5Y and the First Principal Component of Inflation and the Stock Market

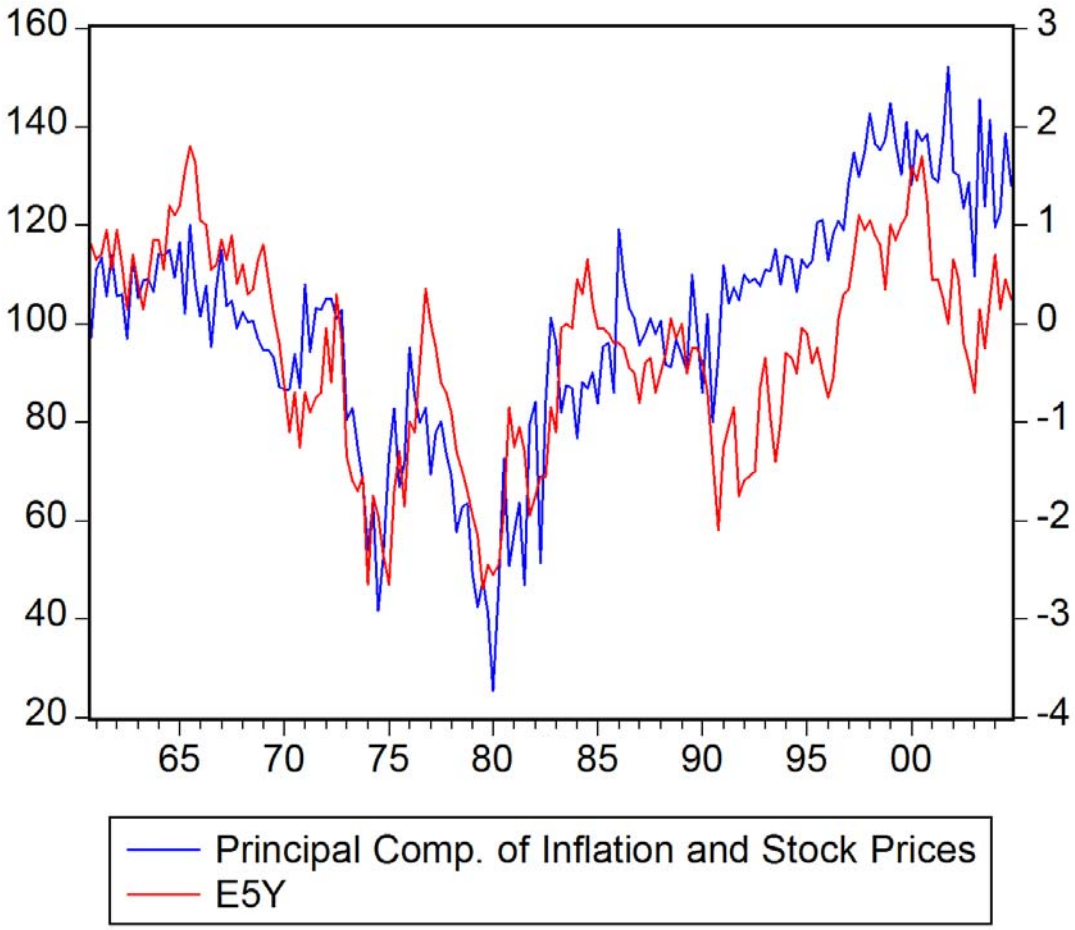

Figure 18 (A)
Bivariate VAR
Labor Productivity and Inflation

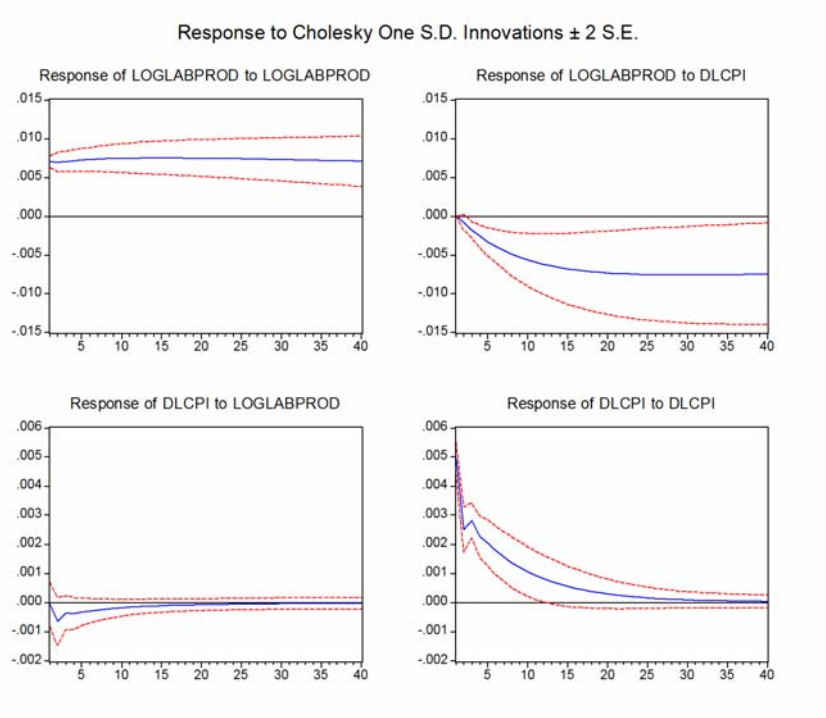

Figure 18 (B)
Bivariate VAR
Labor Productivity and S & P 500

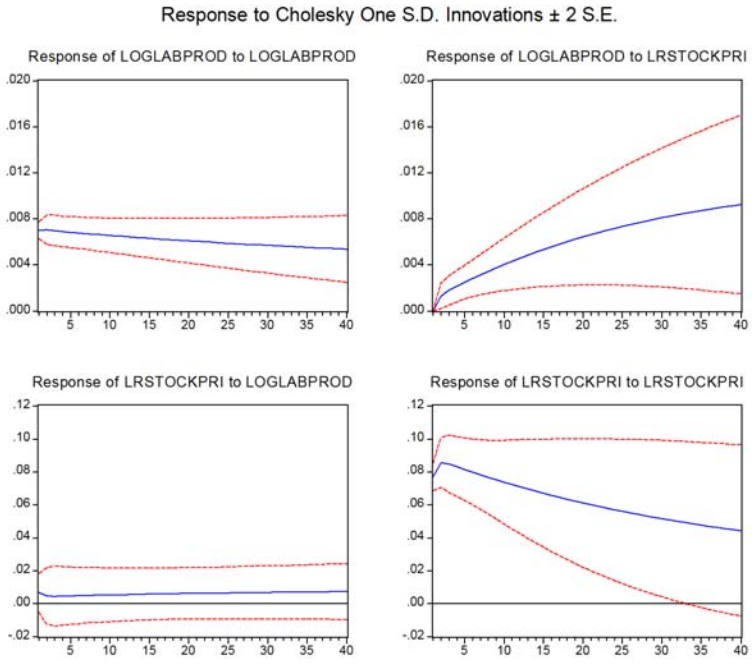

Figure 19
Unemployment and E5Y Across Time

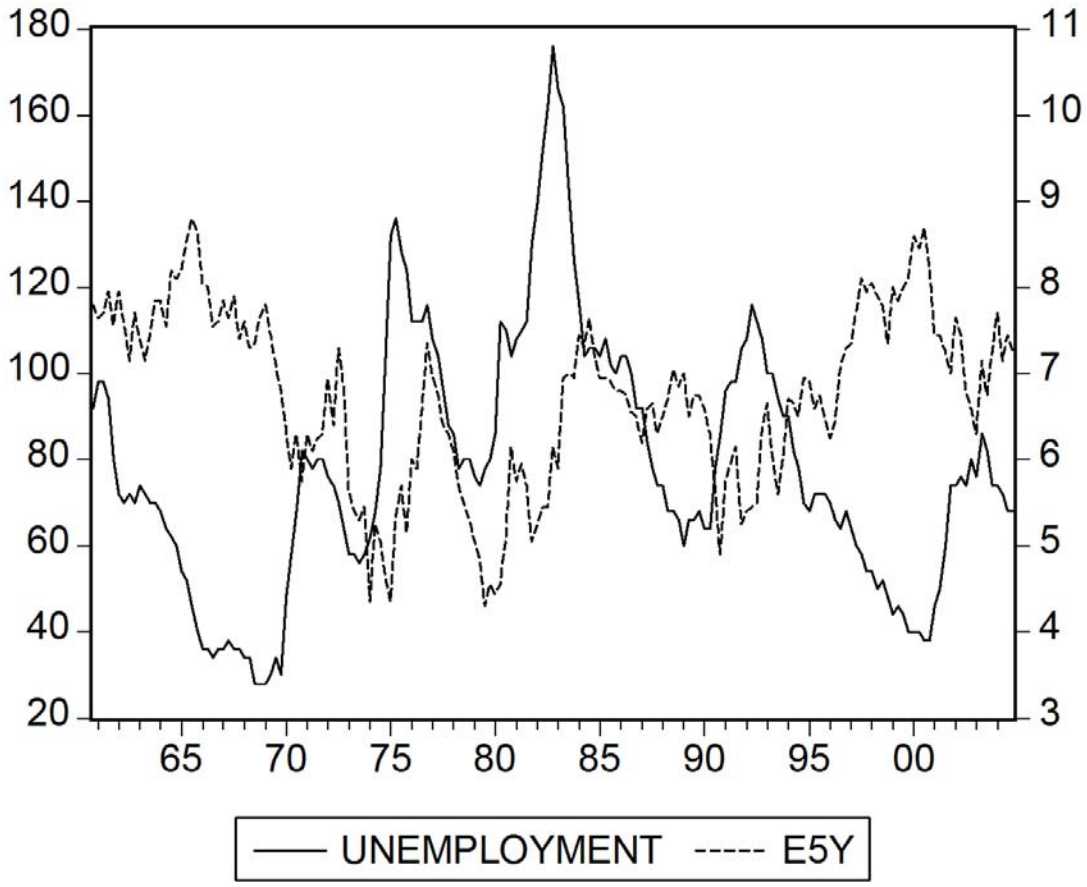

Figure 20
Impulse Responses with Ordering Y→Unemployment→E5Y→C

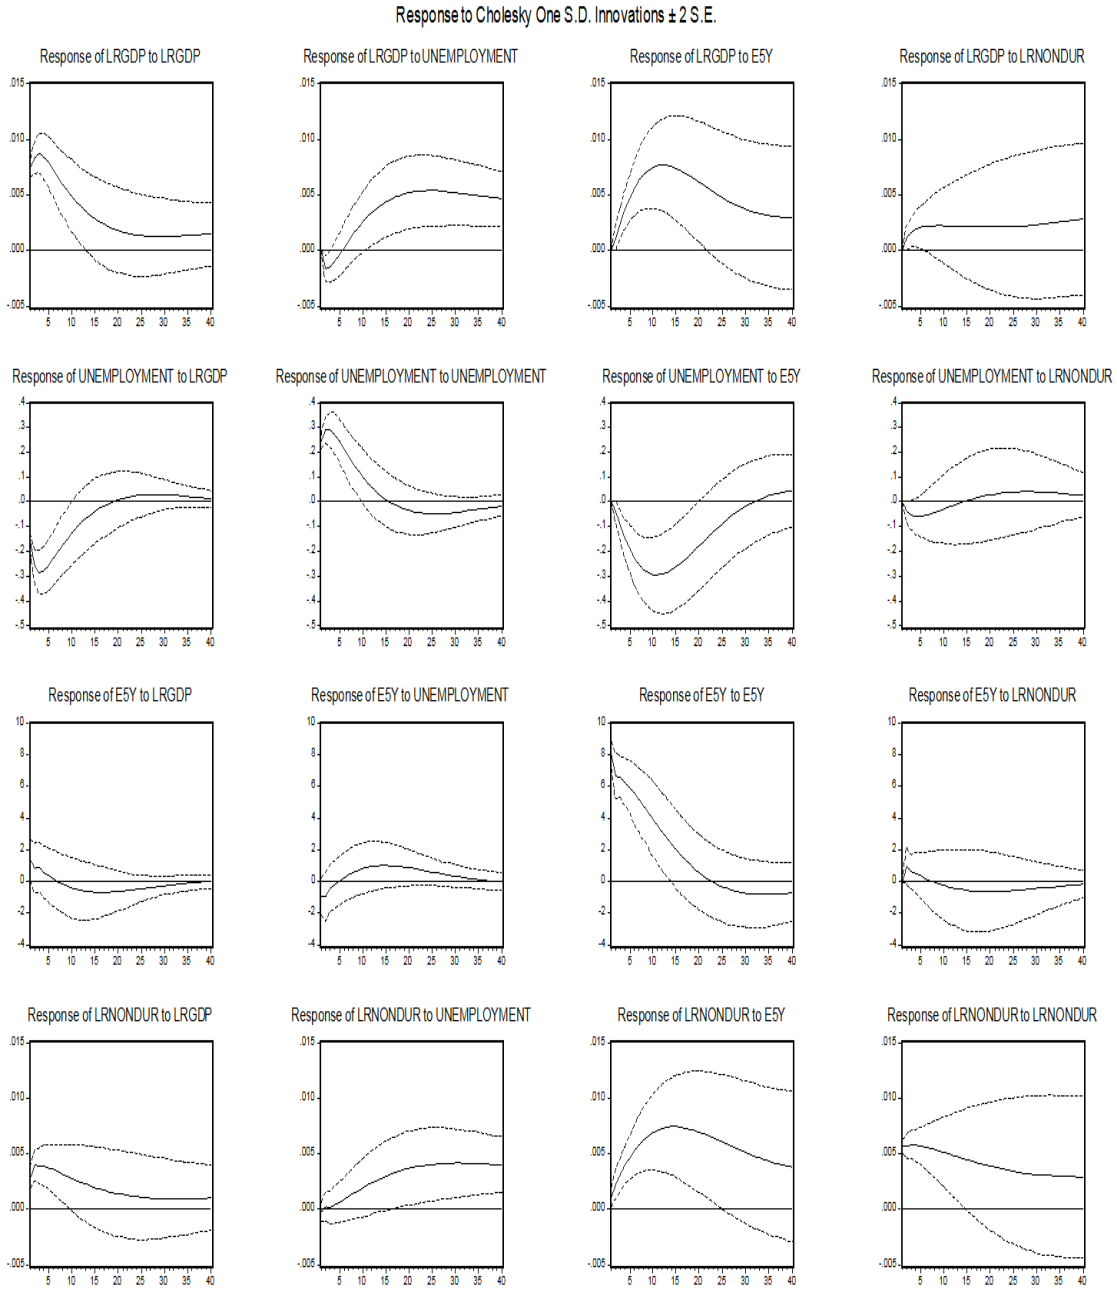

Figure 21
Impulse Responses with Unemployment and Inflation Ordered Prior to E5Y