Very Preliminary Draft Prepared for the 2014 NBER Macroeconomics Annual

Whither News Shocks?

Robert B. Barsky, Susanto Basu, and Keyoung Lee*

Does news about future productivity cause business-cycle fluctuations? What other effects might it have? We explore the answer to this question using semi-structural VARs, where "news" is defined as the innovation in the expectation of TFP at a fixed horizon in the future. We find that systems incorporating a number of forward-looking variables, including stock prices, consumption, consumer confidence and inflation, robustly predict two outcomes. First, following a news shock, TFP rises for several years. Second, inflation falls immediately and substantially, and stays low, often for 10 quarters or more. Consumption typically rises following good news about the future, but investment, consumer durables purchases and hours worked typically fall on impact. All the quantity variables subsequently rise, as does TFP. Depending on the specification of the reduced form VAR, the activity variables may lead TFP to some extent -possibly lending support to the hypothesis of news-driven business cycles – or they may move in lockstep with productivity. For the most part, the quantity and inflation responses are quite consistent with the predictions of a standard, flexible-price RBC model augmented with a Taylor rule for the nominal interest rate. In such models, news shocks typically play at most a small role in explaining business-cycle fluctuations.

* Federal Reserve Bank of Chicago and NBER, barsky@frbchi.org; Boston College and NBER, susanto.basu@bc.edu; and Federal Reserve Bank of Chicago, klee@frbchi.org. We thank Larry Christiano for many helpful discussions. The paper does not reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.

I. Introduction

In the last decade or so, the old "Pigouvian" hypothesis that good news about the future may cause business cycle expansions – which in turn are reversed by less favorable news - has again become the subject of an important discussion within macroeconomics, due in large part to the pioneering work of Beaudry and Portier (2004, 2006). Apart from interest in the question for its own merits, the investigation of news shocks has also been driven by the growing realization that it is difficult to identify structural shocks that might be responsible for the bulk of business cycle fluctuations. Identified shocks to technology do not appear to induce the right business-cycle comovements[1], nor can technology shocks yield anything close the degree of forecastable variation in hours and output seen in real data (Rotemberg and Woodford, 1996). Monetary shocks may induce the right comovements, but their variance is too small. The importance of investment-specific technology shocks appears to decline when the shock is asked to explain the relative price of investment.[2] Other shocks are often either implausible (for example, high-frequency labor supply shocks) or difficult to identify without conditioning on a full DSGE model (for example, markup shocks).

News shocks are appealing for several reasons. First, they are *a priori* plausible sources of fluctuations. Second, they fix some of the unappealing features of technology-driven business-cycle models. For example, recessions can occur without technological regress in models with news shocks; downturns can take place after even positive realizations of technology change, as long as the realizations were smaller than expected. Third, news-driven models of cycles do not require large and frequent changes in current fundamentals, since news

[1] See, for example, Gali (1999), Francis and Ramey (2005), and Basu, Fernald and Kimball (2006).
[2] Justiniano, Primiceri and Tambalotti (2010, 2011).

about the same realization of future technology can change many times. Thus, it is unsurprising that news-driven stories feature in casual explanations of the recession of 2001, and even the Great Recession of 2008. There is considerable appeal to the story that agents' demand for durable goods (including housing) preceding each recession was predicated on high growth expectations that—whether or not rational at the time—proved to be excessively optimistic, requiring a reduction in subsequent purchases and (at least in the case of the Great Recession) a concomitant deleveraging process.

On the other hand, news-driven models have a fundamental theoretical difficulty of their own. Notably, in standard neoclassical models, news shocks do not produce the correct business-cycle comovements. The reasoning follows directly from the insight of Barro and King (1984): without a current increase in the marginal product of labor, consumption and labor supply must move in opposite directions. So good news shocks may raise consumption, but they will then lower investment, labor input and output. Much theoretical work has been devoted to augmenting neoclassical models so that good news shocks will be expansionary in these models.[3]

This paper focuses primarily on the empirical identification of news shocks and their effects, with considerable attention also to understanding and rationalizing what we find in the data in terms of "DSGE reasoning," with and without formal DSGE models.

As far as the identification issue is concerned, we focus the lion's share of our effort in the area of structural (or semi-structural) VARs, for two reasons. First, the bulk of the substantive debate about the empirical effects of news shocks has used the methods of VAR analysis. Second, we believe that the substantially (though certainly not completely) nonparametric nature of the VAR approach (in contrast to fully specified DSGE models)

[3] For example, Beaudry and Portier (2007) and Jaimovich and Rebelo (2009).

maximizes the likelihood that the answers will be driven by the data rather than by model structure. Having identified the properties of news shocks and shown that they are robust to reasonable variations in specification, we turn to interpreting the results through the lens of fully specified DSGE models, using Monte Carlo simulation methods to evaluate the ability of structural VARs to get the "right" answer at least in a controlled context.

Before asking whether news shocks cause business cycles, one ought to ask whether news shocks do in fact exist to a nontrivial extent. Somewhat to our surprise, we find that they do. More precisely, we find that an innovation to the expectation of TFP at t+20 quarters that is orthogonal to current TFP explains between 20 and 40 percent of the variance of TFP at a horizon of five years and typically in excess of 50 percent of the variance at a horizon of 10 years. Since we take the information set to be orthogonal to current TFP, the innovation we identify is a news shock by a stringent definition. It is indisputable that this shock explains a large fraction of the variance of TFP. Thus, it appears that news shocks exist, and are quantitatively important.

Using a battery of forward-looking variables, including consumer confidence, is key to our ability to identify an important role for predictable fluctuations in the level of future TFP. Thus, even though the univariate representation of TFP is best characterized as a random walk, implying that no variables other than current TFP should help predict future TFP, the same is not true in a multivariate context. This point was made in the context of GDP by Cochrane (1994) and Rotemberg and Woodford (1996), but to our knowledge has not been understood previously for TFP. The historical decompositions of our VAR system show that to a large extent our identified news shock is a change in consumer confidence, suggesting that consumer confidence

may aggregate news about future fundamentals in a way consistent with discussions of "the wisdom of crowds" and the anecdote of Galton's ox.[4]

We go on ask whether news shocks are plausibly significant drivers of economic fluctuations. As noted earlier, few if any of the shocks that economists have identified so far (at least using data-driven procedures consistent with large classes of interesting models) can simultaneously match the comovement that is a defining characteristic of business cycles and also explain a large fraction of economic fluctuations at frequencies identified with economic fluctuations. Can our identified news shocks fill this gap?

We believe the answer is no, while admitting that the case is not completely clear-cut. The impact effects of news shocks clearly does not induce the kind of comovement that is characteristic of business cycles. In most of our specifications, we find that consumption rises when there is good news, but investment, consumer durables purchases and hours worked all fall. As discussed above, this is exactly what one would predict in neoclassical models where there is a representative consumer with time-separable preference, the case analyzed by Barro and King (1984). In such settings, news shocks have an income effect but no substitution effect, so consumption and work hours must move in opposite directions, as must consumption and investment.

In subsequent periods consumption, investment and durables purchases all rise strongly, and sometimes work hours do as well. However, unlike Beaudry and Portier (2006), we find that TFP begins rising strongly one or two periods after the news shock. Thus, it is difficult to tell whether the rise in business-cycle variables after a news shock is due to the effects of news per se, or simply to the change in fundamentals that is preceded by a news innovation.

[4] For example, Surowiecki (2005).

Unfortunately, our data-driven procedure cannot definitely answer the counterfactual question: What would happen if there were a news shock, but TFP subsequently did not change? We *can* show that the impact effect of a news shock, when by definition TFP does not change, does not resemble a business cycle. Furthermore, in many ways the impact effect and subsequent dynamics of the key business-cycle variables are quantitatively as well as qualitatively consistent with the predictions of the simple one-sector RBC model. In that model, we know that news cannot explain business-cycle comovements at any horizon, not just on impact. Thus, while we cannot establish that news does not cause business-cycle fluctuations, we can restrict the set of models consistent with a significant role for news. In particular, such models must be ones where positive news causes a number of key quantity variables to "overshoot"—decline on impact, before then rising strongly several quarters after the shock. We do not know of such models, but we admit that they might exist.

We show that news-driven models of business cycles must be consistent with another fact as well. Perhaps our most robust empirical finding is that inflation declines by 30-40 basis points after a news shock and stays below its pre-shock level for 10 quarters or more. This is unusual behavior if one thinks of news shocks as "demand shocks"—that is, shocks that raise output but are not technological in nature. One might surmise that such shocks would create the positive comovement between inflation and output that is also a characteristic of business cycles.[5] We show, however, that a standard RBC model augmented with a Taylor rule can match the behavior of inflation and nominal interest rates, as well as the responses of real variables following a news shock.

[5] As documented by, for example, Stock and Watson (1999). However, Christiano (2010) and Christiano et al. (2010) show how expectations of low future marginal cost from an improvement in technology could lead to declines in inflation on impact of a news shock in a medium-scale New Keynesian model.

Perhaps our most controversial conclusion is that our evidence on the effects of news shocks can be explained by very straightforward economic models with flexible prices and without adjustment costs for either the capital stock or investment. We leave to future research the task of reconciling the unadorned flexible price model that fits news shocks well with the sophisticated New Keynesian models apparently necessary to understand many other macroeconomic phenomena in a unified model.

II. Semi-Structural VAR Identification of News Shocks

A. Basic Method

Let's turn to VAR identification. This is of course intimately tied up with the question of the extent to which quantitatively important news shocks exist in the first place. If the data supports a finding that indeed news shocks do exist, we will be interested in the shape of the impulse responses of standard macroeconomic variables (including the usual real variables as well as inflation and asset prices) to the news shock, as well as to the surprise technology shock. We will also study variance decompositions, and – perhaps more interestingly, historical simulations. At this point, the reader should regard what we are doing as establishing facts regarding news shocks in a relatively nonparametric way, not constrained by a structural model. But of course, we ultimately are interested, at least in part, in what the data say about the effects of news shocks in a DSGE model. In a later section, of the paper, we will compare the VAR results to impulse responses and (in an appendix) Monte Carlo simulations of a simple DSGE model that turns out to fit a number of the facts rather well.

A structural news shock will be defined as an advance signal that agents receive about future productivity. One example of a stochastic process for log TFP, *a*, that contains a news shock would be:

$$a_t = \rho a_{t-1} + v_{t-1} + \varepsilon_t \quad 0 < \rho < 1$$

This is the baseline news specification of Christiano, et al (2010). A second example (Barsky and Sims, 2009, 2011; Jinai, 2013) is:

$$a_t = a_{t-1} + g_{t-1} + \varepsilon_{1t}$$
$$g_t = (1-\rho)\,\bar{g} + \rho g_{t-1} + \varepsilon_{2t}$$

In this specification productivity has a time-varying trend, and agents observe the innovation in that trend one period in advance. Hence the news shock is $\varepsilon_{2,t-1.}$ Since the impulse response to the news shocks we identify via vector autoregression typically displays a rather smooth increase in productivity over several quarters, the latter process is closer to what is nonparametrically estimated in the data.

How do we identify news shocks in a vector autoregression? For the purposes of this paper, an identified news shock will be a linear combination of reduced-form innovations that 1) predicts future productivity; and 2) is not "excessively" correlated with current productivity (more particularly, with a productivity measure stripped of the endogenous, utilization-driven component). Criterion (1) needs no discussion (except for issues about the forecast horizon) – it is certainly a *sine qua non* for a news shock. Nor is it difficult to see why criterion (2) is critical. Consider the extreme case where productivity is a univariate random walk that is also not Granger-caused by any known variables. In that case, the best forecast of future productivity is current productivity; the innovations in the two are equivalent at all horizons, and a shock to current technology and a shock to the forecast of future technology are one and the same. Such

an innovation is no doubt a kind of news (*all* innovations are), but hardly the kind of advance information about future technology with which the expectations-driven business cycle literature is concerned.

It is well-known that TFP growth is approximately white noise; more formally, it is a martingale difference with respect to its own past. Since the univariate TFP process cannot predict future TFP, the possibility of identifying news shocks from the VAR arises if and only if there are observable variables that Granger-cause TFP. These variables may or may not be (indeed in general they will not be) the actual signals seen by the agents in the model, but they serve as *indicators* of the signals agents observe. They may be asset prices, survey measures of expectations, or macroeconomic variables such as consumption, investment, and hours that reflect – and hopefully reveal – the information possessed by the agents.

A positive innovation in the stock price orthogonal to the current productivity innovation but Granger-causing future productivity was *the* operational definition of a news shock in the pioneering empirical paper of Beaudry and Portier (2006). While there are a number of reasons not to rely too heavily on the stock market as an indicator of technological news (some of them discussed in Barsky and Sims, 2011), we retain the stock price as one indicator of a news shock. Consumer confidence also turns out to be a surprisingly valuable indicator (Barsky and Sims, 2009, 20012).

As a warm-up that helps to motivate what we are doing, we show in Figure I impulse responses from three separate bivariate VARs, each involving TFP and one forward-looking variable. In each case, TFP is ordered first. In Panel A we essentially reproduce the key result in the Beaudry and Portier (2006) – an innovation in stock prices orthogonal to current productivity presages a rather long period of increased productivity growth without any apparent subsequent

reversion. Panel B shows that almost precisely the same pattern – with a very similar estimate of the long-run increase in the level of productivity – holds when the forward-looking variable is E5y, the measure of five year business expectations from the Michigan Survey of Consumers (though the error bands are regrettably wide). In Panel C, we consider the implications for productivity growth of the orthogonalized innovation in consumption of nondurables and services. Once again, we observe more or less the same pattern, with a somewhat more rapid step-up in productivity and a slightly smaller long-run effect (though the difference is clearly not statistically significant). Finally - though the full motivation will become clear only later in the paper - inspired by the canonical New Keynesian model in which inflation is a jump variable that presages future real marginal costs, we choose inflation for our forward-looking variable. The results are a mirror image of what we saw in the previous three panels.

Beaudry and Portier (2006), in their previously discussed bivariate VAR with TFP and stock prices, in fact consider two alternative "structuralizations" – a recursive one with stock price ordered after TFP and a "long-run" identification that assigns to a single shock the role of explaining all of the zero frequency variation in TFP. They report that the two alternative "structural" shocks are very highly correlated. In a second-stage regression procedure, they find that a positive realization of the shock is expansionary for consumption, investment, and hours.

Barsky and Sims (2012), on the other hand, use an agnostic VAR identification with a medium-sized VAR that includes consumption, GDP, hours, inflation, stock prices, consumer confidence and interest rates. Their identification strategy considers all shocks that are orthogonal to the innovation in current productivity, and among these – following Uhlig (2004) - chooses the shock that maximally explains a weighted average of future levels of productivity. They find that their identified news shocks raise consumption, but reduce hours, investment, and

GDP in the short run. This finding holds across different VAR specifications in their paper, and is consistent with the results of standard neoclassical models. Additionally, news shocks seem not to account for historical recessions.

The maximization-based identification in Barsky and Sims is not entirely transparent, and there is an arbitrariness (inherited from Uhlig (2004)) about the weights attached to the various horizons over which technology shocks are to be explained. Furthermore, as we discuss shortly, it is not absolutely clear that we want to impose *a priori* that the news shock be strictly uncorrelated with the contemporaneous technology innovation. The approach that we now describe renders it trivial to study both cases.

In this paper, we focus our attention on the object:

$$\left[E_t - E_{t-1}\right] TFP_{t+k} ,$$

where E is the expectations operator and the forecast variable is quarterly utilization-adjusted TFP (from Fernald [2012], who uses a subset of the corrections to standard TFP proposed by Basu, Fernald, Kimball [2006] to create a quarterly TFP series purged of its endogenous utilization component). This object is the innovation in the optimal VAR forecast of TFP at some *fixed* point in the future. We vary this forecast horizon k sequentially, though we focus on appears to be the reasonable benchmark of five years (20 quarters) . For the bulk of the paper we will - in keeping with the news shock literature - orthogonalize this forecast innovation with respect to the contemporaneous TFP innovation. Near the end of the paper we briefly consider the unorthoganalized case, and give it an interpretation that is an alternative to the pure news shock notion.

We see this forecast innovation approach as preferable to the scheme of Barsky and Sims because it is more closely connected with the fundamental definition of a news shock[6]. Note that the forecast is computed entirely from the reduced form VAR. In the case where the VAR is entirely in differences, and *k* goes to infinity it amounts to the Beveridge-Nelson procedure (which is not to say that either the differencing or the infinity are advisable choices.) We see this as a major virtue, because it allows us to go as far as we can with the data alone before imposing any *a priori* restrictions.

To see how we compute the above object, denote a VAR in companion form as

$$Y_t = AY_{t-1} + \varepsilon_t .$$

Note that in this case, $\left[E_t - E_{t-1}\right]Y_{t+k} = A^k \varepsilon_t$. We choose the vector of the matrix A^k associated with TFP to arrive at the linear combination of reduced form residuals. At a mechanical level, one can always regard the shock as one of a number of shocks in a structural or semi-structural VAR. We then compute the impulse responses to this shock after orthogonalizing with respect to current TFP (until the last section in the paper in which we included unorthoganalized forecast innovations). We focus on *k* equal to 20 quarters (5 years). We also consider horizons of 12 and 40 quarters. 40 quarters (10 years) is on the long side, and can be thought of as the finite-horizon analogue of the long-run identification advocated by Beaudry and Portier (2007) and Beaudry and Lucke (2011). Such finite horizon or "medium run" identification approaches have been shown to have considerably better small sample properties in some controlled setting (Francis, Owyang and Roush), though we are concerned that at the 10-year horizon the downward-biased nature of the estimated largest autoregressive root might be the cause of qualitatively misleading

[6] Our method is of course conceptually related to Barsky-Sims (after all, optimal forecasts are the solution to a maximization problem) and overall it gives rather similar – though certainly not identical - results.

results. This is on top of (and quite separate from) the fact that ten years might be an implausibly long lead time for agents to foresee information about forthcoming technological developments.

Is it clear that we ought to orthogonalize the forecast innovation with respect to current innovations in TFP to arrive at our "identified news shock?" On one hand, one might think that a "pure news shock," by definition, should not affect current TFP. The current line of DSGE modeling for news shocks reflects this idea that news shocks are shocks that affect productivity in the future while being unrelated to current (true) productivity. It is for this reason that orthogonalization works well in simulations of DSGE models (Barsky and Sims, 2009, 2011).

On the other hand, a "news shock" in the real world may not have this kind of feature. It is possible that news about future productivity arrives along with innovations in productivity today. Positive technological innovations today may be introduced today with the understanding that significant improvements to the technology will occur in the years to come. Or there may be gradual diffusion of a new general purpose technology across sectors. In these cases, a single structural shock raises TFP on impact and then raises it further over time. In such cases, orthogonalizing the news shock with respect to innovation in current TFP is simply the wrong identification. Therefore, it would be nice to examine both kinds of "news shocks," one where the innovation in the *k*-horizon-ahead forecast is orthogonalized with respect to the reduced form innovation in current TFP and one where it is not.

If we find that the unorthogonalized news shock looks quite similar to the orthogonalized one, so that (subject to small sample bias) the *data alone* tell us the effects of TFP news unrelated to current TFP , then we would be in an ideal world in terms of defining and identifying news shocks – no identifying restrictions would be needed. Indeed, this happy

coincidence obtains in the levels specification for the 40-quarter horizon. Unfortunately, the same is not the case for the shorter horizons of 12 and 20 quarters.[7] Indeed, three years ahead much of the forecast innovation is due to the contemporaneous innovation. We regard five years as a reasonable benchmark horizon for predictions of future productivity, and this is the horizon on which we focus. Here it is necessary to orthogonalize with respect to current TFP if the resultant impulse response is to long anything like what has come to be regarded in the literature as a news shock.

B. A Brief and Informal Digression on Invertibility

Much has been written about the "invertibility problem" (see, for example, Fernandez-Villaverde et al., 2005). A DSGE model (which typically has a state space representation in the form of a VARMA) is *invertible* if that VARMA can be reduced to a VAR – i.e. if the structural shocks in the DSGE model are linear combinations of reduced-form VAR residuals. At this level, invertibility is about the suitability of the *reduced-form* VAR for studying the underlying DSGE model, and prior to the problem of imposing identifying restrictions on the VAR. If it is know that the news shock cannot be represented as a linear combination of VAR residuals, there is no point in trying to "structuralize" the VAR. In our context the underlying economic reason for potential noninvertibility is that agents in the economy might well receive signals about the future that are not innovations with respect to variables observed by the econometrician.

Earlier writers on news shocks (e.g. Leeper and Walker, 20xx; Dupor ,; Chari, Kehoe, and McGrattan, 2008) often wrote as if the invertibility problem is a disaster for structural

[7] A finding that at the 40 quarter horizon the unorthoganalized forecast variance is more or less identical to the orthogonalized one is equivalent to saying that the effect on TFP of the contemporaneous innovation in productivity has largely disappeared after ten years. Once must be concerned that the 40 quarter result is due in part to downward bias in the largest autoregressive root in the productivity equation.

inference from VARs. We take the (still controversial but increasingly respectable) position that invertibility is not a devastating problem. This is one place in which we are consonant with Beaudry and Portier (2013). The reason has to do with recent refinements in understanding invertibility issues at both the theoretical and practical level. The new understanding relates invertibility to insufficient richness either in the variables included in the underlying DSGE model or in the data available as indicators of news shocks rather than to the existence of news itself.

The whole idea of a model of news-driven business cycles revolves around concrete economic decisions that agents make in response to news. Signals that are observed by economic actors but not reflected in their behavior are of limited interest. If news shocks are quantitatively important, and a rich enough set of observable indicators is included in the model and available to the econometrician, then news should be reflected in such activity variables as consumption, investment, and hours as well as in stock and bond prices. News may also be partially revealed in measures of survey confidence.

Roughly speaking, the richer the underlying DSGE model is in observable variables, the more likely it is that we will be able to uncover the news shock and their impulse responses, either exactly (in large samples) or approximately. At an analytical level, this point is reflected in the example given at the end of Fernandez-Villaverde et al. (2005). These authors exhibit a permanent-income model that is noninvertible because the econometrician does not have enough information to infer news that agents receive about permanent income. In a somewhat anticlimactic turn of event at the end of the paper, it is noted that the information set of the econometrician did not include consumption. Once consumption is observed, the system

becomes invertible and once again a VAR can be used to recover the underlying shocks of the model. Not every case will be this straightforward, but the general lesson is clear.

Moreover, at a practical level, it may not even be critical that the DSGE model be invertible in the analytical sense. In an important paper, Sims (2012) argues – and demonstrates via monte carlo simulations - that VARs with news shocks and rich information sets may give good estimates of theoretical impulse response from DSGE models that are technically not invertible.[i] In this sense, invertibility is not an "on/off" issue. Simulations of a New Keynesian model without capital in Barsky and Sims (2009) and a real business cycle model with capital in Barsky and Sims (2011) are favorable to the VAR identification method (conditional, of course, on the "true" model being of the same variety as the DSGE model studied).

III. Data

We use standard national income accounts data on gross investment, purchases of consumer durables, and consumption of non-durables and services (aggregated into a single Divisia index). We express each variable in per-capita terms by dividing by the civilian, non-institutional population. Hours worked are the BLS measure of aggregate nonfarm payrolls hours, again put on a per-capita basis. The stock price variable is Shiller's real S&P 500 index; the interest rate is the 3-month Treasury bill rate; inflation is measured by the CPI-U. The consumer confidence measure is from the Michigan Survey of Consumers. Data on quarterly, utilization-adjusted TFP are from Fernald (2012), who uses a subset of the procedures proposed by Basu, Fernald and Kimball (2006) to create a quarterly TFP series purged of the endogenous utilization component. When necessary, we convert the growth rates in Fernald's TFP series to

an index in log levels. We take logs of the quantity variables and the stock price. See the Data Appendix for details.

IV. Identified News Shocks

For this next long stretch, we adopt the now standard approach in the literature that treats news shocks and impact technology shocks as completely separable—i.e. we orthogonalize the innovation in the *k*-horizon TFP forecast with respect to current technology, and provisionally call this the *identified news shock*. In a later section we will interpret the unorthogonalized forecast innovation as a *composite shock* associated with technology diffusion. Thus, all the news shocks we discuss in this sub-section are orthogonalized with respect to the unanticipated TFP shock.

We have run a number of specifications in our VAR exercises with news shocks. The differences in specification concern the handling of nonstationarity (and the associated issue of cointegration), and the forecast horizon. The *levels* specification needs no explanation, except to point out that it is often thought to be desirable because of its robustness to specification error concerning the number of cointegrating vectors, and that it can be given a Bayesian justification even when some of the variables are neither stationary nor cointegrated with others. The *hybrid* specification puts nonstationary variables in differences (thus assuming an absence of cointegration - a feature which is in fact not particularly supported by statistical tests), and presumptively stationary variables (consumer confidence, inflation, and interest rates) in levels. Of course, if cointegration holds then there are potential significant benefits to imposing it. . However, evidence has been accumulating in recent decades that the US economy might be better understood through the lens of a two-sector growth model rather than a one-sector model.

In a two-sector model, aggregate TFP would no longer be cointegrated with either nondurables consumption or with investment and consumer durables purchases. The hybrid specification, by simply differencing the non-stationary variables, does not impose a model that would be misspecified in this sense. That said, we also present results of an error correction model. We impose cointegrating factors of unity relating TFP to nondurables and services, durables, and investment, but let the data determine the cointegrating relationship between TFP and the stock market.

Figures 2 and 3 display the impulse responses to news shocks and unanticipated technology shocks, respectively, for the levels specification. Figure 4 shows the impulse responses to news shocks for the hybrid specification, with Figure 5 showing the corresponding responses for unanticipated technology. Figures 6 and 7 display the impulse responses to news shocks and unanticipated technology shocks for the cointegration specification. As previously noted, our baseline model is run with a 5-year forecast horizon in identifying the news shock. We have also run the model with alternative forecast horizons, identifying news shocks as orthogonalized innovations in the expectation of TFP three and ten years in the future. To save space, we do not report the results here, but we discuss which of our findings are sensitive to the choice of horizon. The full results for the different horizons are available from the authors on request.

A. Discussion of Results for Identified News Shocks

Across all our specifications for the effects of news shocks, our most consistent finding of the effects of news on a standard business-cycle variable is that inflation falls immediately, substantially and persistently in response to a "good" news shock. The results in the level specification of Figure 3 are representative. A current news shock that by assumption has no

effect on TFP today but raises TFP by about 0.15 percent in three quarters and an additional 0.15 percent 20 quarters after the shock, lowers inflation on impact by more than 0.3 percentage points. Inflation continues to fall over the next year, reaching a maximum drop of 0.6 percentage points four quarters after the shock. Inflation then rises slowly from this trough, but is still 0.1 percentage point below its pre-shock value 10 quarters after the news shock.

The literature on news shocks has not focused much attention on this very robust finding, with Barsky and Sims (2009), Christiano et al. (2010) and Jinnai (2013) being notable exceptions. We feel this result merits substantially more discussion, in part because it is difficult to understand without a model where there is substantial interplay between real and nominal variables. We thus devote a separate section to understanding and explaining this finding in the context of popular DSGE models of business cycle.

The next most robust findings are the effects of news shocks on consumer confidence and the stock market. In every specification in our baseline VAR model, stock prices jump up when there is a positive news shock, and typically continue to rise for about two years. At its peak, the stock market reaches a level between 2 and 5 percent above its pre-shock value.

A positive innovation in the stock price orthogonal to the current productivity innovation but Granger-causing future productivity was *the* operational definition of a news shock in the pioneering empirical paper of Beaudry and Portier (2006). While there are a number of reasons not to rely too heavily on the stock market as an indicator of technological news (some of them discussed in Barsky and Sims, 2011), we retain the stock price as one indicator of a news shock, and the results are fully consistent with variants of our specifications that omit the stock market.

Unlike the stock market, which takes between five and 10 quarters to reach its peak, consumer confidence jumps on impact and either is at its peak immediately, or reaches that peak

within a quarter or two. Confidence then remains significantly above its pre-shock level for years, typically 15 quarters or more.

To demonstrate that the news shock explains much of the behavior of these three variables, particularly inflation and consumer confidence, in Figure 8 we plot the time series of our variables and the fitted values that can be attributed to news and unanticipated technology shocks. We use results from the hybrid specification to generate the figure. To isolate business-cycle comovements, we filter both the data and the fitted values using the bandpass filter, set to isolate cycles with frequencies between 6 and 32 quarters.

The historical decomposition for inflation shows that the news shock explains a large fraction of cyclical inflation movements. The comovement is particularly impressive in the 1970s and early 1980s, but unlike the case for other variables, the news shocks continue to play an important role in explaining inflation fluctuations in the 1990s and 2000s.

If the decomposition of inflation is striking, that of consumer confidence is even more so. It is clear that to a large extent a news shock *is* a consumer confidence shock. But whereas confidence is often discussed as a measure of animal spirits, an exogenous driver of aggregate demand, it is clear from our results that confidence is driven heavily by fluctuations in expectations of TFP. This is a welcome confirmation of Barsky and Sims (2012), who come to this conclusion using a quite different approach involving a New Keynesian structural model. The decomposition for the stock market, by contrast, shows that while both news and unanticipated TFP are positively correlated with changes in the real value of equities, the relationship is not a close one, and is especially weak in recent decades.

We now turn to exploring the effects of news shocks on quantity variables. First is the effect of the shock on the future path of TFP. Strikingly, news shocks predict future TFP growth

at high levels of significance. Furthermore, once realized, the path of TFP following a news innovation resembles a growth rate shock, rising predictably over time. As is well known, growth rate shocks create different incentives for intertemporal substitution than level TFP shocks. Thus, in addition to the fact that (orthogonalized) news shocks create a wealth effect with no substitution effect on impact, even when the news shock leads on average to an increase in future TFP, the wealth effects in the early periods after TFP begins to rise should be strong relative to the substitution effects. The historical decomposition shows that news shocks are definitely correlated with TFP changes, although the bulk of changes in TFP are due to the unanticipated shocks.

We now turn to the more controversial results regarding the impulse responses of quantities to news. This of course leads to the most controversial question about news shocks – are they an important cause of business cycles? In almost all of our specifications, we find that investment, hours and consumer durables purchases do not rise on impact of a news shock. Thus, news shocks do not on impact move these endogenous variables in the direction that would be expected by an advocate of news-driven business-cycles. In some cases, one or more of these variables declines in a statistically significant fashion. This is particularly true for hours worked. The failure of consumption to rise more strongly may appear to be at odds with the notion of a news shock as a positive wealth shock or with the idea that consumption should also, as a forward-looking variable, react strongly to news. Interestingly, the historical decomposition shows a pattern that will become familiar as we discuss the other quantity variables. It is clear that news does drive consumption to some extent, but it is also apparent that this pattern was much stronger in the early part of the sample, especially around the time of the dramatic oil price increases of the 1970s. However, consumption is expected to rise over time following a positive

news shock, which in turn requires that real interest rates rise at the time of the news. Indeed, this is what happens: the nominal interest rate falls, but expected inflation falls further, so real interest rates rise when there is good news.

As noted, in two out of our three specifications hours worked decline significantly when a news shock hits. In no case do hours worked rise significantly on impact of a news shock. This pattern of results is consistent with the idea that news shocks constitute a positive wealth shock for consumers. If consumption and leisure are both normal goods, then an increase in wealth with no offsetting substitution effect should cause hours worked to decline and consumption to rise, which is the pattern we generally see in the data. It then follows that investment and/or purchases of durable goods must fall. We see this qualitative pattern in Figures 2, 4 and 6, but the declines are rather small, except in Figure 6.

Turning to the historical decompositions, we see that for hours worked news shocks were important early in the sample, but cease to be so around the early 1980s. Similar results hold for investment and, to a lesser extent, for consumer durables.

Finally, we study the behavior of unanticipated TFP shocks. Here the results vary somewhat more depending on specification. With the level and cointegrating specifications (Figures 2 and 6), we see that the time path of TFP following the shock is mean-reverting, while in the hybrid specification (Figure 4) TFP follows basically a random walk. However, in general we see that consumption, investment, hours worked and durables purchases all rise following a positive TFP shock. Interestingly, it appears that inflation rises after an increase in TFP, which appears difficult to reconcile with sticky-price models, where a mean-reverting TFP shock typically lowers marginal cost for several periods. It is also inconsistent with the findings of Basu, Fernald and Kimball (2006), who found that inflation fell significantly for several years

after a surprise improvement in technology. In all the specifications, we see that stock prices either decline or fail to rise significantly, while consumer confidence typically rises by a small amount on impact and then falls. In general, the results following unanticipated TFP shocks are difficult to reconcile with the theoretical predictions for positive technology shocks, especially in sticky-price models with limited monetary accommodation (see, for example, Gali, 2008).

These results lead us to question whether our "purified" TFP series has really succeeded in purging the Solow residual of its non-technological component. As Fernald (2012) discusses carefully, the quarterly series that he constructed and we use cannot make all the corrections that Basu, Fernald and Kimball (2006) were able to make using annual data, because some of the source data are not available at a quarterly frequency. But if more non-technological components remain in Fernald's TFP series, then an unanticipated rise in his measure may be a composite of a change in technology and an endogenous change in utilization coming from a non-technological "demand" shock. This hypothesis would explain why we find, contrary to Basu, Fernald and Kimball (2006), that an innovation in TFP does not lower inflation and raises hours worked.

Since our identification of news shocks relies on predicting changes in Fernald's TFP measure, one might ask whether this problem calls our identification of news shocks into question. We believe it does not, because our identification rests on forecasting TFP k quarters in the future, where we take k to be a large value (such as 20, or even 40, quarters). Over such long periods of time, predictable changes in utilization should be essentially zero. This can be verified by checking the predictability of hours per worker at such horizons, which according to the model of Basu, Fernald and Kimball (2006) is an observable proxy for unobserved factor utilization.

Looking further at the impulse responses, we see that hours, though contracting a bit on impact (less than in Barsky and Sims), *expand* quite a bit at medium frequencies, as do the various components of output. We confirm this visual impression with the variance decompositions for the hybrid specification in Table 1, which show that at a 20-quarter horizon news shocks explain about 30 percent of the variance in consumption and purchases of consumer durables and 16 percent of investment, although only 5 percent of hours. Table 2 shows variance decompositions for the cointegrated specification, which appear more favorable for the news-driven business cycle hypothesis. At a 20-quarter horizon, news shocks account for 68 percent of the variance of consumption, about 30 percent of the variation in investment and consumer durables, and 40 percent of the variance of hours worked. This looks quite a bit more like a news-driven business cycles, albeit with a delay.

To what extent are the variance decompositions in Table 2 evidence for the news theory of business cycles? To answer this question, it is important to take into consideration the behavior of TFP following a news shock. Recall that *today's news shock is (on average) tomorrow's rise in actual productivity*. Higher productivity, whether anticipated or not, *will* raise output permanently and (probably) hours temporarily—although perhaps after a delay, as in Basu, Fernald, and Kimball (2006) and Gali (1999). That hours rise at medium frequencies may not be the effect of news *per se*, but rather the effect of the subsequent actual improvement in technology. Note from Table 1 that at a 20-quarter horizon, the news shock "explains" 40 percent of the variance of TFP (rising ultimately to 55 percent at 40 quarters) while only accounting for a relatively smaller fraction of hours (5 percent) and investment (16 percent). So on one hand, the rise of hours and investment at medium frequencies is an important result. But

on the other hand, if those variables are merely rising in conjunction with TFP, then it is difficult to say that news per se is causing the business cycle movements.

The critical issue is the extent to which hours, investment, and durables, rise ahead of TFP, or only contemporaneously with it. To the extent that activity variables rise sharply before cumulative TFP growth has been substantial, we have evidence of news-driven cycles. In the hybrid specification (Table 1) this is clearly not the case. In the levels and error correction specifications, the case for news-driven cycles is somewhat stronger. Some substantial movement in investment and hours leads productivity. Yet these activity variables continue to rise as productivity growth picks up. This is suggestive of some combination of a pure news effect and a direct productivity effect on activity, much as might be seen in a standard real business cycle model.

In the case of the hybrid specification, noting that activity variables pick up only as does productivity, we propose a heuristic exercise to decompose the effects of a news shock into the "pure news" effect and the effect of actual (predictable) technology change. Since this is a counterfactual experiment, it is difficult to do without a full structural model. However, we can get an idea of the importance of "pure news" versus the realization of technology by combining results from our estimated impulse responses to news shocks and to unanticipated technology shocks. The idea is to use the impulse response to unanticipated technology to uncover the "direct effect" of technology on the various outcome variables. Then the pure news effect is computed as the impulse response to news subtracting the imputed direct effect of the technology change.

This procedure would be rigorous if the path of TFP following a news shock (after TFP began to increase) followed exactly the same stochastic process as an unanticipated TFP shock.

To see why, consider an RBC model, where the state variables are just capital and TFP, and the current level of TFP summarizes all the expectations for the future path of TFP. Suppose that TFP always follows the same stochastic process whether it is anticipated or unanticipated. However, some TFP changes are anticipated *j* periods in advance. Then the response of economic variables after anticipated and unanticipated TFP changes would differ only because the starting values of the capital stock would be different: an anticipated news shock would create incentives to accumulate or de-cumulate capital over those *j* periods, while an unanticipated shock by definition would not. Assuming that the effects of capital and technology are additive (which would be true in a linearized model), the impulse responses to unanticipated technology would allow us to uncover the effects due to the change in the capital stock from the reaction of agents to the anticipated shock, which would be the effects of news *per se*. The same result would hold in more complex linearized models with more state variables (for example, nominal prices or wages).

Unfortunately, the impulse responses to anticipated and unanticipated TFP shocks are not similar in shape. The path of TFP following a news shock looks like the response to a stationary growth rate shock, with the univariate time series process for technology resembling an AR(1) in first differences. The response of TFP to its own unanticipated shock looks like a stationary AR(1) in levels or, in the hybrid specification, resembles a random walk. Assuming rational expectations on the part of economic agents, the differences in shape imply different incentives for intertemporal substitution in consumption and labor supply. Thus, our calculations can only be suggestive.

However, to mitigate the problems just outlined, we take as our baseline the effects of an unanticipated TFP shock 10 quarters after the shocks occurs. This is roughly when TFP is

halfway back to its pre-shock value in Figures 2 and 6, and where TFP and other variables are roughly at their long-run values in the hybrid specification shown in Figure 4.

Thus, we implement the following calculation:

$$\text{"Pure News Effect"}_k = \text{IRF}|\text{News}_k - \frac{\text{TFP}|\text{News}_k}{\text{TFP}_{10}}\text{IRF}|\text{TFP}_{10}\,.$$

We implement this imputation of the pure news effect using the hybrid specification in Figures 3 and 4. The results are shown in Figure 8. We will also show results for a specification based on a two-sector growth model in the next subsection.

We find that in the hybrid case this correction is significant. The news responses of consumption, investment and durables are cut roughly in half by correcting for the TFP change following a news shock. The hours response is reduced even further, and hours are now negative for the first 6 quarters following a news shock. A news shock that ultimately raises TFP by about 0.7 percent has a maximum positive effect on hours of about 0.1 percent. On the other hand, the correction actually increases the decline in inflation and interest rates, leaves consumer confidence basically unaffected, and has only a modest effect on the stock price response. These results cast some doubt on the claim that the data support the news-driven business cycle hypothesis. On the other hand, the corresponding corrections for the level and error correction models are – as expected – smaller. Once again, these latter two specifications are somewhat more supportive of the notion of purely news-driven cycles.

C. Restrictions from a 2-sector model

In our earlier discussion of the potential pitfalls of cointegration – and the concomitant case for the hybrid specification – we noted that evidence supports idea that US data are better modeled as a two-sector growth model. In particular, relative price data strongly indicate that

technology in the production of investment goods and consumer durables has been advancing at a rate faster than the technology for producing other forms of output. Thus, we check the robustness of our previous results by estimating a specification that allows for two different shocks to technology, and therefore news about these shocks separately. Thus, in addition to our basic specifications, we have run a model where we included consumption-specific and investment-specific TFP as separate variables, TFPC and TFPI. The TFP series used in our previous specifications is a weighted average of the two, but standard economic models imply that the effects of the two shocks should be quite different, so it makes sense to separate the two measures.[8]

The measure of investment-specific technology is Fernald's (2012) utilization-corrected TFP in equipment and consumer durables. Our TFPC is also from Fernald, and is intended to measure technology in the other industries. Thus, we can estimate impulse responses to investment- or consumption-specific TFP news shocks, which are the forecast revisions in TFPI or TFPC at a horizon of 20 quarters, orthogonalized with respect to the unanticipated TFPC and TFPI shocks. Figure 11 presents the impulse responses to a consumption news shock. Figure 12 presents the impulse responses to an investment news shock.

We find that news about consumption TFP has an interesting constellation of effects. It predicts a rise in consumption TFP of 0.2 percent starting about a year after the shock, but an even larger and more significant increase in investment TFP, almost 1 percent after 20 quarters. Consumption is unchanged on impact, but significantly higher starting about 3 quarters after the shock. Durables consumption jumps down significantly, but becomes significantly positive four

[8] For discussion of the business-cycle implications of investment-specific technical change, see Greenwood, Hercowitz and Krusell (2000) and Basu, Fernald, Fisher and Kimball (2011).

quarters after the shock. Hours worked are basically unchanged on impact, and also become significantly positive four quarters after the shock. Investment, by contrast, jumps up and generally keeps increasing. As is our earlier results, investment falls significantly on impact, and reaches its trough several quarters after the shock. The nominal interest rate also declines. Stock prices rise, but the increase is not significant until four quarters after the shock. Consumer confidence, by contrast, rises significantly on impact, and stays high for an extended period.

From 5 to 15 quarters after the shock, investment, consumption, hours and durables purchases are significantly higher, so over this horizon the fluctuations in these variables exhibit the comovement one expects over the business cycles. But of course this is also the period when consumption and investment TFP are noticeably higher following the news shock.

Figure 12 shows that in response to investment-specific news shocks, consumption- and investment-specific TFP both rise as before, but this time the rise in TFPC is not significantly different from zero. As before, TFPI increases by about 1 percent after 20 quarters. The responses of business-cycle variables are more clearly positive after an investment news shock: consumption, investment and hours all increase on impact, as do stock prices and consumer confidence. The only quantity variable that declines significantly after the shock is durables purchases, and it becomes significantly positive 3 quarters after the shock. CPI inflation declines as before, but not as much. This fact is unsurprising, since investment and consumer durables account for a small fraction of GDP, and for a particularly small part of the basket of consumption basket for the CPI.

We also experimented with a third approach to the empirical identification of news shocks, an approach based on publication counts of science and technology titles, based on the neering work of Alexopoulos (2011). While developed to discuss contemporaneous technology

shocks, her time series data on the number of science and technology titles published seemed to us to be potentially suited for interpretation in terms of the news shock concept, albeit at a relatively short horizon. Interestingly, Alexopoulos (2011, table 4) finds that her measure of computer publications Granger-causes the utilization-adjusted technology measure of Basu, Fernald and Kimball (2006) (p-value of 0.02). The correlation between the two series at one- and two-year horizons is very small (on the order of two percent). Alexopoulos also finds that an innovation in her measure is positively correlated with changes in GDP, investment and hours worked. After two years, the correlation between the measure of computing-related publications and utilization-adjusted TFP jumps to about 0.20 (her table 3B). This suggested to us that the effects on GDP and other business-cycle variables in the first two years following an innovation in Alexopoulos's publication series might be fruitfully interpreted as being due to the effects of news *per se*. If so, her measure, interpreted as a news shock, has potential for explaining business cycles as traditionally defined.

It thus seemed natural to run our reduced form VARs with the addition of the Alexopoulos' time series of computer titles as an additional variable, and interpret the innovation in that measure (or perhaps the innovation in the first principal component of that measure and stock prices) as a proxy for news shocks. Unfortunately, these data are available only at an annual frequency. The error bands in all of the specifications we tried were far too wide to make any inference that would add to what Alexopoulos already found, and we do not include the impulse responses in the paper, though they – along with the data – are available on request.

In recent work, Christiano, Motto and Rostagno (2010, 2013) [henceforth CMR] find that in an estimated DSGE model with financial frictions along with the now-standard features of sophisticated New Keynesian empirical models, "risk news shocks" drive out technology news

shocks. Risk shocks as viewed by CMR manifest themselves as increases in the cross-sectional standard deviation of TFP at the firm level. Since the underlying theoretical model is closely related to that of Bernanke, Gertler, and Gilchrist, the ultimate significance of the CMR measure is presumably that those firms that do relatively poorly have a high probability of defaulting on their obligations. CMR find that *news* about *future* risk is a dominant source of the business cycle, while tfp news is of relatively trivial importance.

It is important to note that CMR have a quite different aim than do we in the present paper. They are seeking a shock (or at most, a few shocks) that maximize(s) the ability to explain business cycle variation. We, on the other hand, are seeking to identify and understand the effects – large or small – of a particular pre-specified shock (namely, the tfp news shock) on real and nominal macroeconomic variables. Nonetheless, in light of the CMR results in the DSGE context, it seems appropriate to ask whether the inclusion of a proxy for risk news shocks in a VAR changes the magnitude of the identified tfp news shocks, or the behavior of their impulse responses.

To address this question we choose as our proxy for risk news shocks the interest rate spread between Moody's Aaa bonds of twenty or more years of maturity and the corresponding Baa bonds. Since the duration of these bonds (at least at recent interest rates) is quite long, this spread responds to information about future bankruptcy risk over a long horizon, and thus represents news about future risk.

We included this spread variable in addition to our other variables in both the levels and hybrid specifications, ordering it first so as to maximize its potential to dwarf the importance of TFP news shocks. Figures 13 and 15 show the responses of all of the variables to the "risk news shock" in the levels and hybrid specifications, respectively, while Figures 14 and 16 show the

responses to the TFP news shock orthogonalized with respect to both the risk spread and the unanticipated TFP shock. In both the levels and hybrid specifications, durables, investment, hours, inflation, treasury bill yields fall substantially on impact in response to a spread shock, as expected[9]. Most importantly, the risk news shock has essentially no effect on the impulse responses to the TFP news shock.

V. Why are News Shocks So Disinflationary? A Flexible Price DSGE Explanation

By far the most robust result that we obtain across different specifications is that good news shocks are highly disinflationary. The strong disinflationary response - on the order of half a percentage point or more for a one standard deviation news shock - appears immediately on impact and in most specifications lasts for a number of quarters. *Whether or not they are a cause of business cycles, news shocks are very important for the understanding of inflation.*

The disinflationary effects of news shocks appeared in previous literature in two guises. Barsky and Sims (2009, 2011) included inflation in their VAR systems – just as in the present paper - and observed that the identified news shock was associated with a sharp drop in inflation on impact. Christiano, Motto, and Rostagno (2010) took a more indirect approach, in which the empirical observation and its theoretical rationale were more closely intertwined. They looked at a number of stock market booms, and noted that they were consistently associated with low inflation.

They then constructed a model in which news shocks – in combination with a monetary policy in which the Fed did not properly account for changes in the natural rate of interest (e.g. by including it as an intercept in the Taylor Rule as in Woodford [2003]) – accounted for both

[9] Somewhat surprisingly, confidence and stock prices respond positively in the levels specification, though not in the hybrid specification.

the rise in the stock market and the drop in inflation. They went on to make the normative point that this model tends to call for inclusion of asset prices or leverage measures in the monetary policy rule – an important counterexample to the well-known result of Bernanke and Gertler (200x and 200y) that the Fed need pay attention to asset market considerations only to the extent that they have implications for the optimal forecast of inflation.

Both Barsky and Sims (2009, 2011) and Christiano et al (2011) regard this phenomenon as a manifestation of the forward-looking nature of inflation in the New Keynesian model. Consider the now-classic New Keynesian relationship between inflation and future expected real marginal costs (Gali and Gertler, 1999; Sbordone, 199x and 200y):

$$\pi_t = \frac{(1-\theta)-(1-\theta\beta)}{\theta\beta}\sum_{j=0}^{\infty}\beta^j E_t mc_{t+j}$$

where θ is the Calvo price adjustment parameter and β is the discount factor. In the benchmark model, the log deviation of real marginal cost is equal to the log difference between real wage and labor productivity: $mc_t = w_t - p_t - y_t - n_t$. Expected future productivity improvements, ceteris paribus, lower expected real marginal costs. If the real wage does not rise too sharply (either on impact or over time), inflation will jump down and stay down, potentially for a number of quarters – just as in the data. Whether this happens or not depends on the persistence of the news process, the monetary policy rule (particularly the extent to which the monetary authority observes a good proxy for the natural real interest rate), and the degree of price and especially wage stickiness. Barsky and Sims (2009) at least saw the disinflationary nature of news shocks as rather persuasive evidence for the forward-looking New Keynesian model. Jinnai (2013) offered significant modifications to Barsky and Sims (2009) to allow the disinflationary consequence of news shocks to continue to hold in a model with capital.

We now show that perhaps surprisingly the disinflationary nature of news shocks can be accounted for in a fully flexible price model supplemented with a Taylor rule to determine inflation and nominal interest rates. In an effort to see what model features are critical for matching the estimated impulse responses to news shocks, we began with a baseline flexible price DSGE model that includes news shocks. To this model, we add a Taylor rule to produce definite predictions for inflation following shocks. Our assumed law of motion for technology is once again

$$a_t = a_{t-1} + g_{t-1} + \varepsilon_{1,t}$$

$$g_t = (1-\kappa)\bar{g} + \kappa g_{t-1} + \varepsilon_{2,t}$$

where a_t is log of technology, and $\bar{g}$ is the steady-state technology growth rate. In this framework, $\varepsilon_{1,t}$ is a surprise technology shock and $\varepsilon_{2,t-1}$ is the news shock, which has no effect on actual technology in the period that the shock hits. The process is chosen so that the news shock leads to a sustained growth in technology starting from the following period, as in our VAR results.[10] Technology finally asymptotes to a permanently higher level following a news shock. Since the growth-rate shocks are stationary, the model can be solved by log-linearization around the non-stochastic balanced growth path.

The real block of the model is a standard, one-sector, flexible-price model, with capital accumulation, Cobb-Douglas production, log utility in consumption, and a unitary Frisch elasticity of labor supply. The log-linearized model equations and the parameter values are listed in the appendix.

[10] Some of the VAR impulse responses suggest that the main rise in TFP begins two quarters after the shock. It would obviously be easy to modify the stochastic process for technology to introduce a two-quarter delay, rather than our current choice of one quarter.

Since the model has flexible prices, the Classical Dichotomy holds. Thus, the monetary policy rule is important only for determining the behavior of inflation and nominal interest rates following a shock. In particular, we know that the real rate will always be at the natural rate. We assume a Taylor-type rule with nominal interest rate smoothing:

$$i_t = \rho i_{t-1} + (1-\rho)\left(\phi_y \left(y_t - y_{t-1}\right) + \phi_\pi \pi_t\right)$$

Notice that the monetary authority responds both to output growth and the inflation rate, which is crucial to understanding the dynamics of inflation in response to a news shock. With a parameter value of ϕ_y = 0.65 (or higher), we can generate a disinflationary response following a good news shock, while ensuring that the responses of real variables match our VAR impulse responses reasonably well.

The impulse responses to both a news shock and to a surprise technology shock are presented in Figure 17. For our parameter values, a positive news shock raises consumption and lowers hours on impact due to the wealth effect. With increased consumption, lower hours worked, and unchanged technology and capital (on impact), investment must fall. Even when TFP changes, the shock implies that technology will rise predictably over time. As is well known, growth-rate shocks have strong wealth effects on impact, which are not offset by substitution effects since the level of technology changes by little in the early periods. As a consequence, investment is still below its pre-shock level for another quarter. As TFP rises, hours, output and investment increase rapidly, and investment and output asymptote to permanently higher levels. These are standard responses to a growth rate shock.

As Figure 17 shows, inflation falls for several periods after a news shock, although the low inflation does not persist as long as it does in our empirical estimates. Why does inflation

fall in the model with these parameters? A good news shock raises the natural rate of interest, since consumption is expected to grow over time. Since prices are flexible, the real interest rate will always match the natural rate. Thus, the monetary authority's choice of a nominal rate simultaneously determines a rate of inflation such that the nominal rate minus expected inflation matches the natural interest rate. While the effect of a rise in nominal rates on inflation is ambiguous, the effect of a fall of nominal rates on inflation is unambiguously negative. Thus any monetary policy rule that lowers the nominal rate following a news shock will cause disinflation. But recall that the RBC model reduces output on impact. With a sufficiently large coefficient on output in the Taylor rule, the nominal rate will indeed fall as well.

As noted above, this simple model does not have a role for stock prices. However, we conjecture that we could introduce a non-trivial stock market into this model, even without adding investment adjustment costs. (Note that a simple *q* model of investment would not match the joint behavior of investment and stock prices, since investment jumps down on impact of a news shock, but stock prices jump up.) Neither is the model consistent with the empirically preferable "i-dot" adjustment costs, as investment jumps in this model, as in the data. Our hypothesized solution is to make the Hayashi (1982) theorem inapplicable by introducing diminishing returns to variable factors, and thus rents accruing to a fixed factor, into the model. As a modeling device, the easiest method would be to introduce a fixed factor of production into the production function ("land"), which would receive rents equal to its marginal product. More realistically, we think of land as a metaphor for slow-moving state variables, such as brand capital or "owned ideas" that do not require costly investment to acquire (Laitner and Stolyarov, 2005). If this "land" is owned by firms, the value of the firms is likely to rise along with the rents coming from higher TFP and output. Indeed, the stock market value of firms will rise as

long as the effect of higher future rents is stronger than the discounting coming from higher real interest rates. Thus, we suspect that it is quite possible to match the empirical rise in stock prices occasioned by a news shock in an extended version of this simple model.

VI. Concluding Remarks

So far, we have discussed news shocks using the framework pioneered by Beaudry and Portier (2004), where the actual fundamental change comes as manna from heaven, unrelated to current TFP, and "news" is a noisy signal of its future advent. Now we briefly discuss two significant departures from this framework. First, we consider the possibility that the growth in TFP is in fact an endogenous outcome of economic activity. Second, we consider an interpretation of the unorthogonalized news shock as a slowly diffusing technology process.

Perhaps predictable growth in productivity is not a realization of news of an exogenous process, but rather the outcome of purposeful economic activity. For instance, an increase in demand may lead to the production of knowledge via learning by doing. Although activity falls on impact when our statistical procedure tells us that productivity will soon rise, there is a subsequent rise in investment that could conceivably lead to productivity increases through learning by doing. We considered three possible sources of demand-driven expansions that might plausibly lead to endogenous productivity expansion: monetary policy shocks, increases in government expenditure associated with potential preparation for war (Ramey 2013), and sharp increase in oil prices occasioned by physical supply interruptions in the Middle East as discussed by Hamilton (2006). Only oil had any effect at all, and that effect was small and did not materially change the impulse response to news shocks as we identify them. Thus endogenous

growth does not appear to provide an important challenge to the view of news shocks laid out by Beaudry and Portier (2006) and reflected in this paper.

A second and less radical alternative to the pure news shock hypothesis is suggested by computing impulse responses to our forecast innovation $E_t - E_{t-1}[TFP_{t+k}]$ without orthogonalizing with respect to the current innovation in TFP. These impulse responses involve a modest increase in productivity on impact followed by further TFP growth. Such would be the case, for example, if technology diffuses slowly over time, or if technology is embodied in new capital goods and the new capital displaces the old over some length of time. Apparently, the distinction between a technology shock and a news shock is a clean one only in a limiting case.

Let us suppose a distinction between "science"—basic discoveries, such as the transistor—and "technology," practical innovations such as integrated circuits, made from millions or billions of transistors and used in devices of all sorts. Suppose that the scientific frontier advances exogenously, but the technology frontier, which is what we measure as the level of TFP, only advances because of purposeful investments in adapting scientific discoveries to commercial use.[11] Such a situation, which is a good description of many important advances, presents both practical and conceptual difficulties for our investigation of news shocks. At a practical level, the lags between the advent of news—the scientific discovery—and the subsequent effects on technology may be both long and variable, both because inventions are heterogeneous and for economic reasons, such as those discussed in Shleifer (1986). At a conceptual level, there may be several types of news associated with the discovery of a basic invention: the date the transistor was invented, the date that transistors could be miniaturized, the date it was obvious that computers, and hence integrated circuits, had many more applications

[11] A similar model is proposed by Comin, Gertler and Santacreu (2009).

and hence much greater demand than had previously been believed. It is not clear that one can represent this complex process by a single type of technology shock and a single type of news. And, in keeping with the multi-sector model of Long and Plosser (1983) and the critique of single-shock RBC models by Blanchard and Fischer (1989), it suggests that news, like technology, is sector-specific, and it is difficult to see how most sectoral news shocks (or sectoral technology shocks) can have large aggregate effects.

In this paper we explored the effect of news shocks using semi-structural VARs, where "news" is defined as the innovation in the expectation of TFP at a fixed horizon in the future. We found that systems incorporating a number of forward-looking variables, including stock prices, consumption, consumer confidence and inflation, robustly predicted two outcomes. First, following a news shock, TFP rises for several years. Second, inflation falls immediately and substantially, and stays low, often for 10 quarters or more. Consumption typically rises following good news about the future, but investment, consumer durables purchases and hours worked typically fall on impact. All the quantity variables subsequently rise, as does TFP. Depending on the specification of the reduced form VAR, the activity variables may lead TFP to some extent -possibly lending support to the hypothesis of news-driven business cycles – or they may move in lockstep with productivity. For the most part, the quantity and inflation responses are quite consistent with the predictions of a standard, flexible-price RBC model augmented with a Taylor rule for the nominal interest rate. In such models, news shocks typically play at most a small role in explaining business-cycle fluctuations.

Appendix: Log-Linearized DSGE Model with News Shocks

The law of motion for technology is

$$a_t = a_{t-1} + g_{t-1} + \varepsilon_{1,t}$$

$$g_t = (1-\kappa)\bar{g} + \kappa g_{t-1} + \varepsilon_{2,t}$$

The following equations summarize the behavior of consumption, labor supply, production and capital accumulation:

$$0 = E_t\left[c_t - c_{t+1} + i_t - \pi_{t+1}\right]$$

$$0 = E_t\left[c_t - c_{t+1} + \left(1-\beta e^{-\bar{g}}(1-\delta)\right)r_{t+1}\right]$$

$$y_t = \frac{C}{Y}c_t + \frac{X}{Y}x_t$$

$$y_t = \alpha k_t + (1-\alpha)(a_t + n_t)$$

$$w_t = y_t - n_t$$

$$r_t = y_t - k_t$$

$$\frac{1}{\eta}n_t = w_t - c_t$$

$$e^{\bar{g}}k_{t+1} = (1-\delta)k_t + \left(e^{\bar{g}} - 1 + \delta\right)x_t$$

$$i_t = \rho i_{t-1} + (1-\rho)\left(\phi_y\left(y_t - y_{t-1}\right) + \phi_\pi \pi_t\right)$$

The parameters are from Barsky and Sims (2009) with some additions from Jinnai (2013): β = 0.99, η = 1, ρ = 0.75, ϕ_π = 1.1, ϕ_y = 0.65, κ = 0.5, $\bar{g}$ = 0.0025, α = 0.33 and δ = 0.025.

Table 1: Variance Decompositions (Hybrid Specification)

H	TFP	Consumption	Dur Cons.	Investment	Hours	Stock Prices	Confidence	Inflation	3-mo T-bills
0	0.000	0.037	0.025	0.012	0.065	0.004	0.573	0.219	0.154
4	0.092	0.047	0.085	0.023	0.015	0.030	0.659	0.448	0.141
8	0.181	0.141	0.206	0.069	0.014	0.063	0.718	0.484	0.156
12	0.268	0.215	0.287	0.112	0.032	0.074	0.744	0.504	0.184
16	0.346	0.259	0.335	0.139	0.045	0.076	0.751	0.519	0.214
20	0.406	0.283	0.366	0.161	0.054	0.074	0.751	0.528	0.240
40	0.545	0.280	0.409	0.222	0.062	0.048	0.740	0.530	0.288

H is the forecast horizon. The numbers are fraction of forecast error variance of each variable from our identified news shock. The variance decomposition was done with our hybrid specification, where confidence, inflation, and t-bills are included in levels and all other variables in differences.

Table 2: Variance Decompositions (Error Correction Specification)

H	TFP	Consumption	Dur Cons.	Investment	Hours	Stock Prices	Confidence	Inflation	3-mo T-bills
0	0.000	0.073	0.110	0.009	0.003	0.273	0.328	0.281	0.172
4	0.025	0.380	0.117	0.068	0.085	0.446	0.495	0.474	0.086
8	0.042	0.546	0.198	0.185	0.262	0.574	0.545	0.510	0.087
12	0.084	0.618	0.242	0.234	0.334	0.640	0.563	0.462	0.100
16	0.139	0.662	0.281	0.274	0.372	0.668	0.551	0.410	0.109
20	0.200	0.685	0.315	0.304	0.393	0.669	0.532	0.390	0.118
40	0.467	0.645	0.412	0.316	0.328	0.562	0.476	0.382	0.182

H is the forecast horizon. The numbers are fraction of forecast error variance of each variable from our identified news shock. The variance decomposition was done with our error correction specification, where confidence, inflation, and t-bills are included in levels and all other variables in differences, and cointegration was imposed between TFP and Consumption, TFP and Durables Consumption, TFP and Investment, and TFP and Stock Prices.

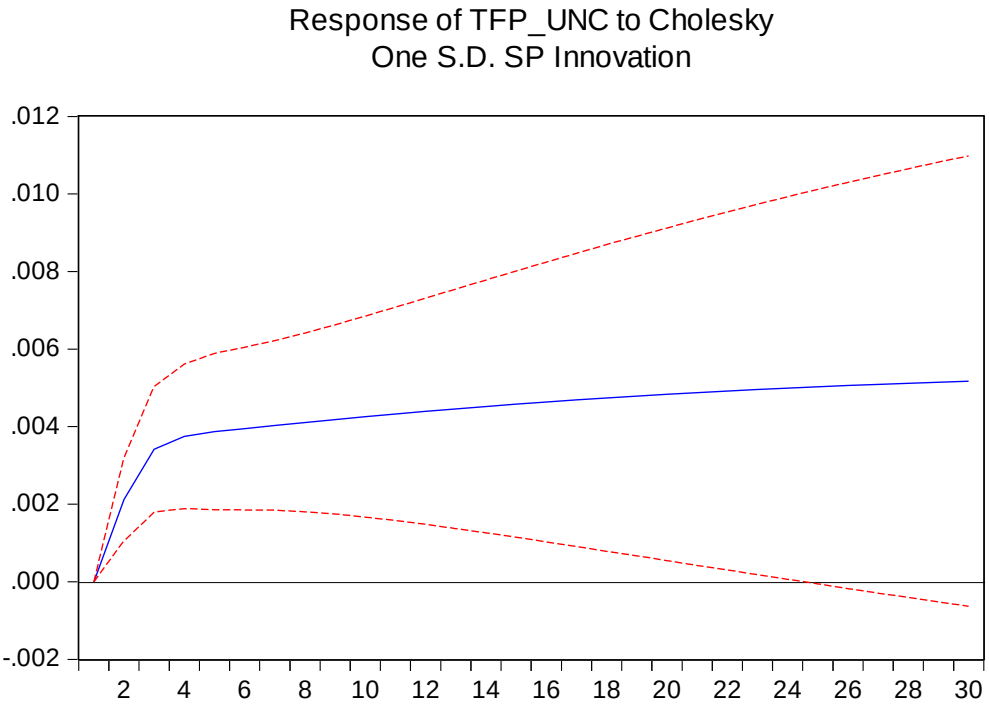
Table 3: Variance Decompositions (Levels Specification)

H	TFP	Consumption	Dur Cons.	Investment	Hours	Stock Prices	Confidence	Inflation	3-mo T-bills
0	0.000	0.085	0.032	0.001	0.185	0.043	0.366	0.167	0.230
4	0.067	0.432	0.281	0.131	0.054	0.141	0.512	0.368	0.217
8	0.147	0.648	0.539	0.344	0.206	0.228	0.559	0.402	0.183
12	0.232	0.725	0.643	0.452	0.334	0.279	0.557	0.350	0.158
16	0.318	0.739	0.683	0.486	0.364	0.291	0.532	0.302	0.135
20	0.392	0.722	0.687	0.487	0.347	0.282	0.506	0.278	0.121
40	0.534	0.569	0.631	0.436	0.222	0.213	0.446	0.249	0.107

H is the forecast horizon. The numbers are fraction of forecast error variance of each variable from our identified news shock. The variance decomposition was done with our levels, where all variables were included in levels.

Figure 1: IRFs from Bivariate VARs with Productivity and a Single Forward-looking Variable

Panel A



Panel B

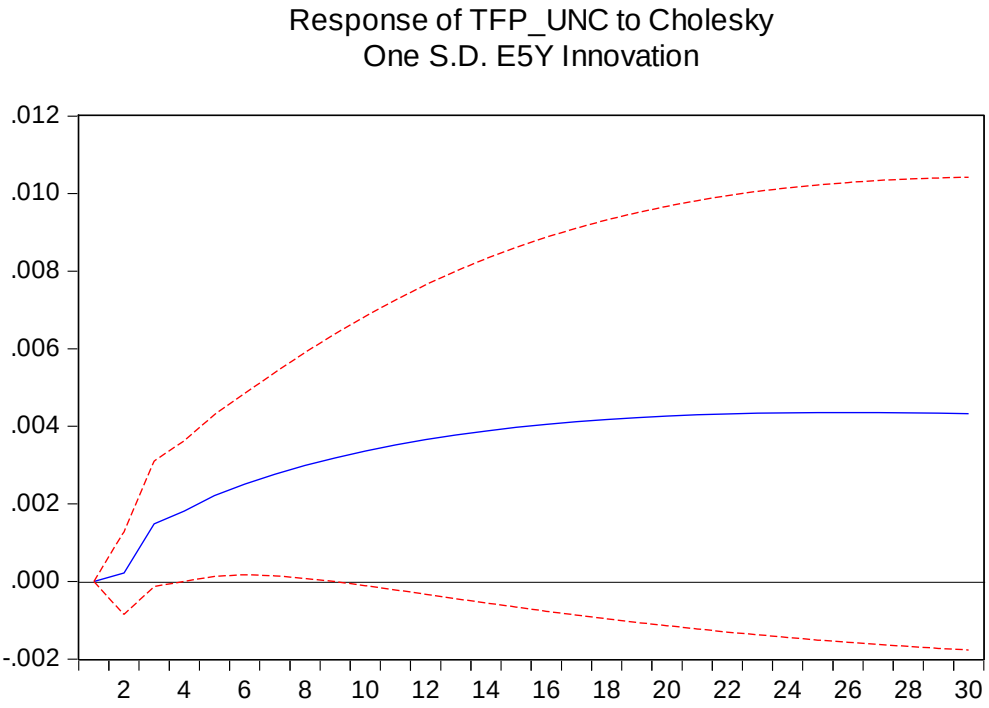
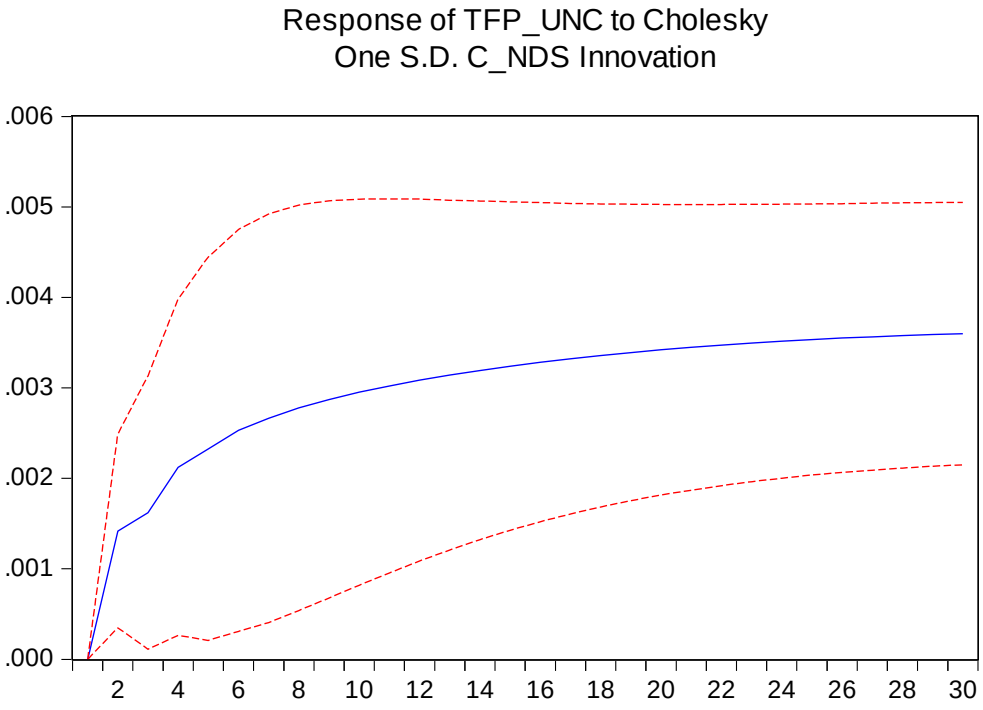


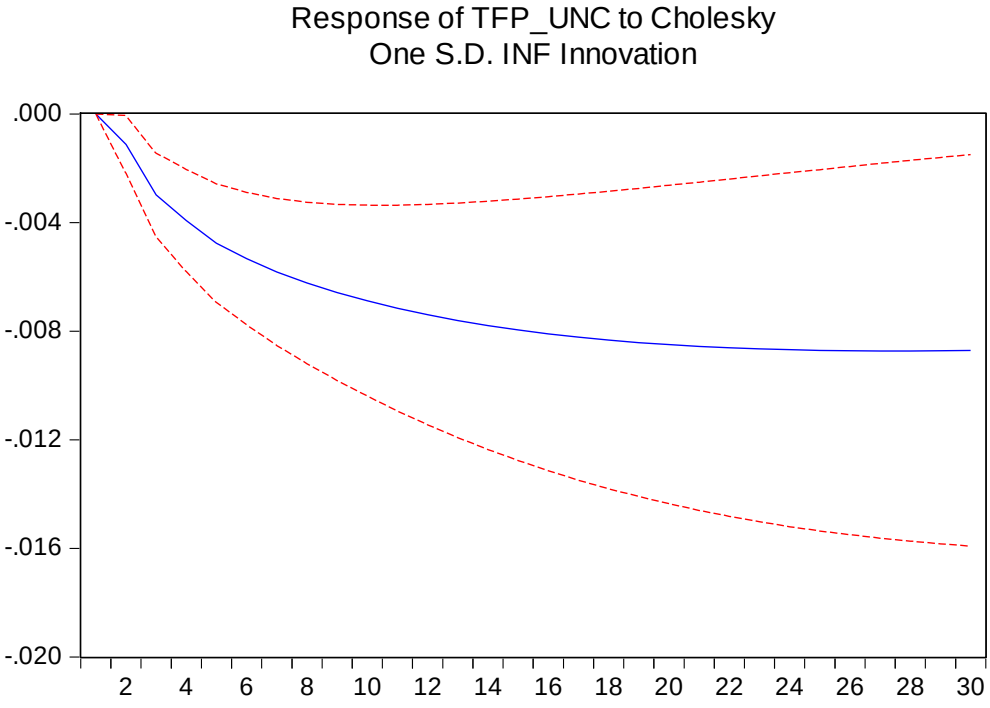
Figure 1: All estimation was done with sample period 1960Q1 to 2012Q2, in levels with constant and 3 lags.

Figure 1 (cont.)

Panel C



Panel D



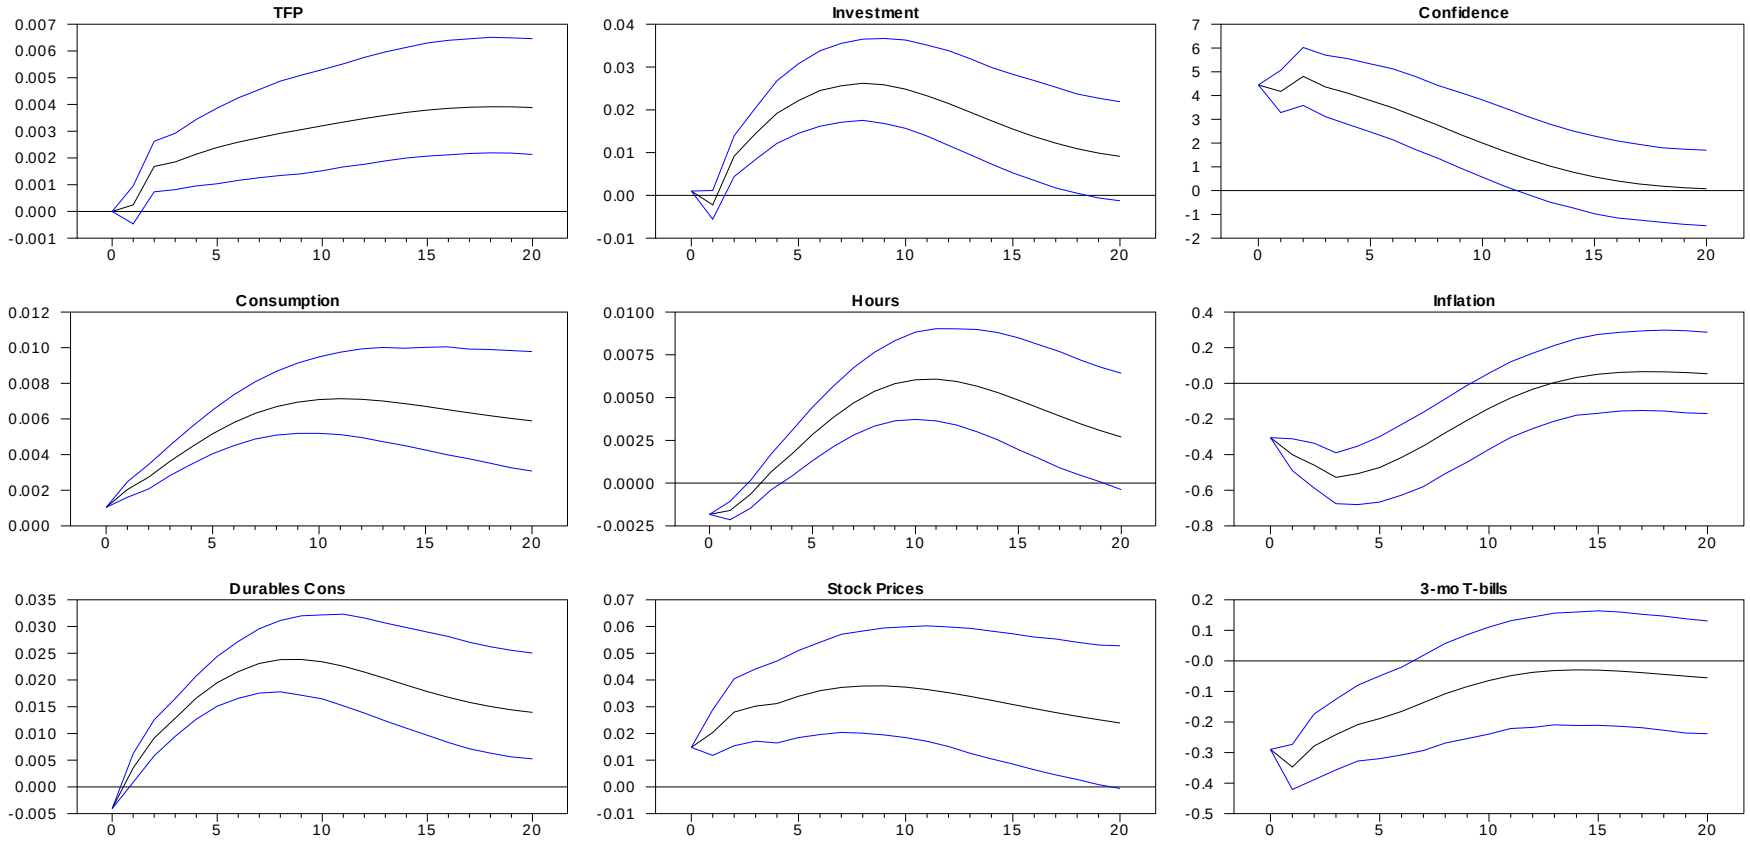


Figure 2: Impulse responses to identified news shock using levels specification. Black lines are point estimates of the impulse responses and blue lines are 5% and 95% percentile error bands obtained from Monte Carlo integration with 2000 replications.

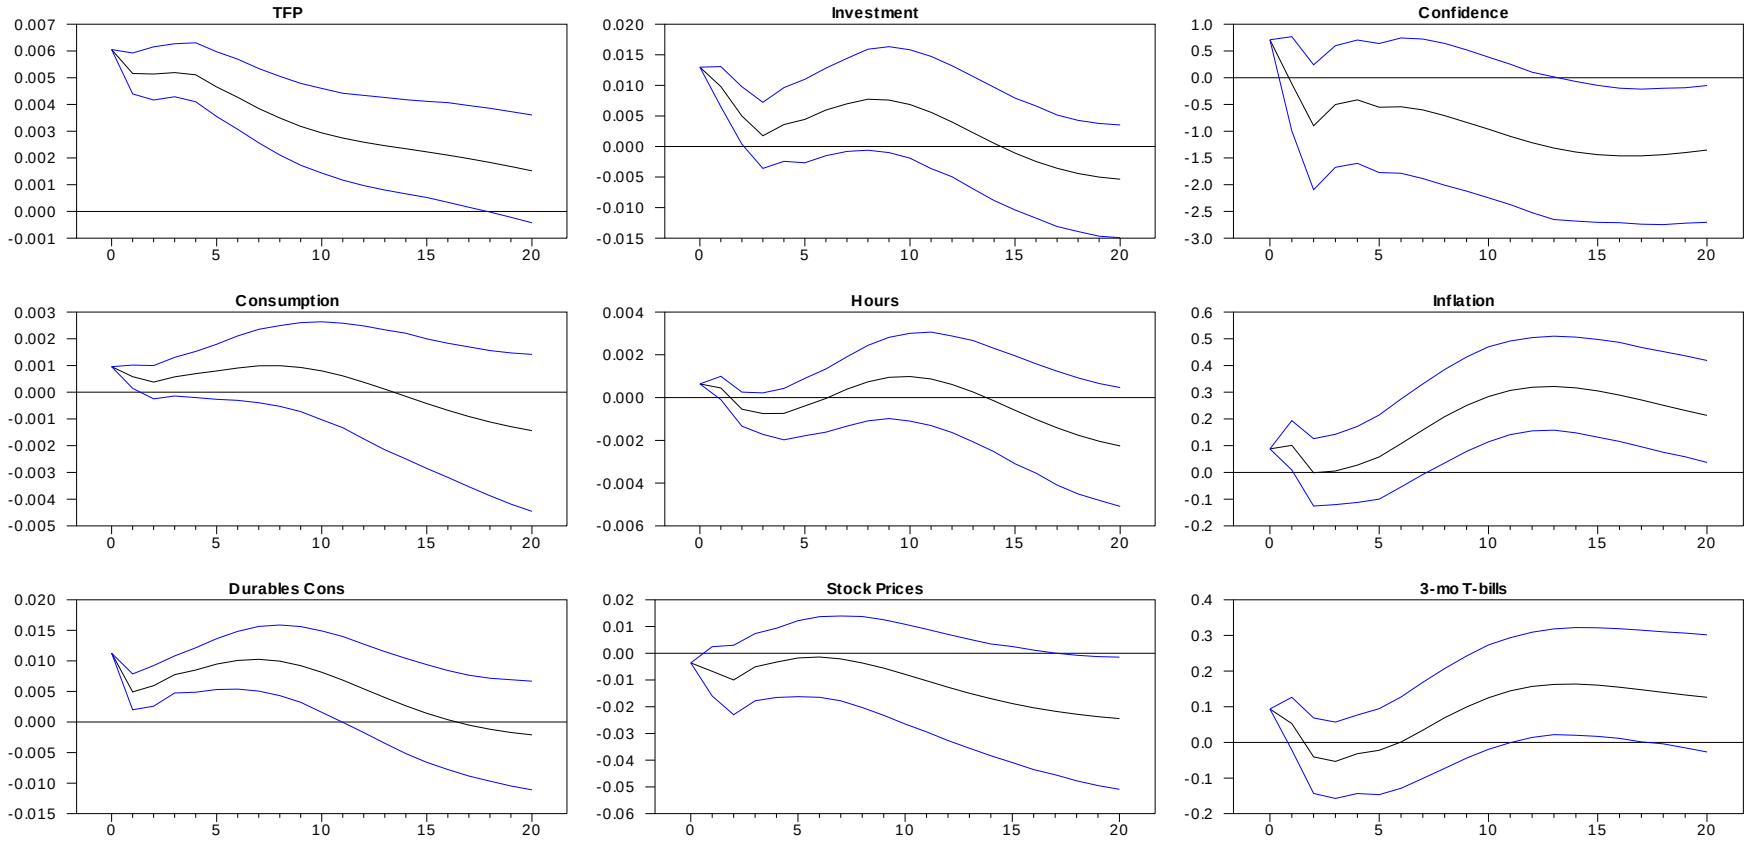


Figure 3: Impulse responses to unanticipated TFP shock, using levels specification. Black lines are point estimates of the impulse responses and blue lines are 5% and 95% percentile error bands obtained from Monte Carlo integration with 2000 replications.

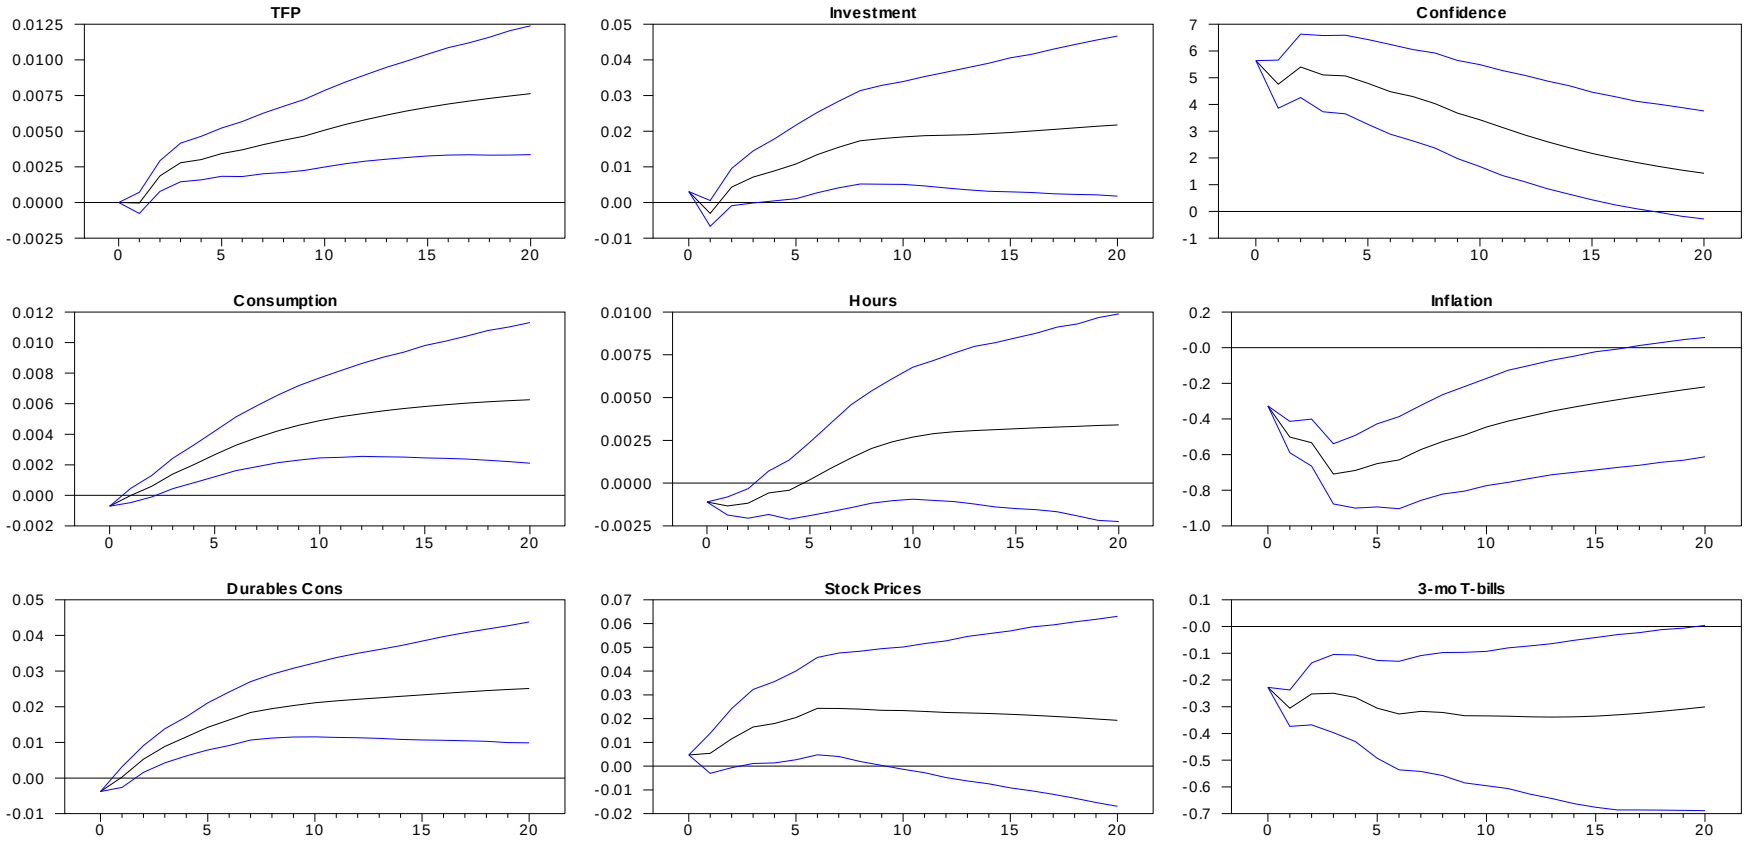


Figure 4: Impulse responses to identified news shock, using hybrid specification, where confidence, inflation, and t-bills are in levels and all else in differences. Black lines are point estimates of the impulse responses and blue lines are 5% and 95% percentile error bands obtained from Monte Carlo integration with 2000 replications.

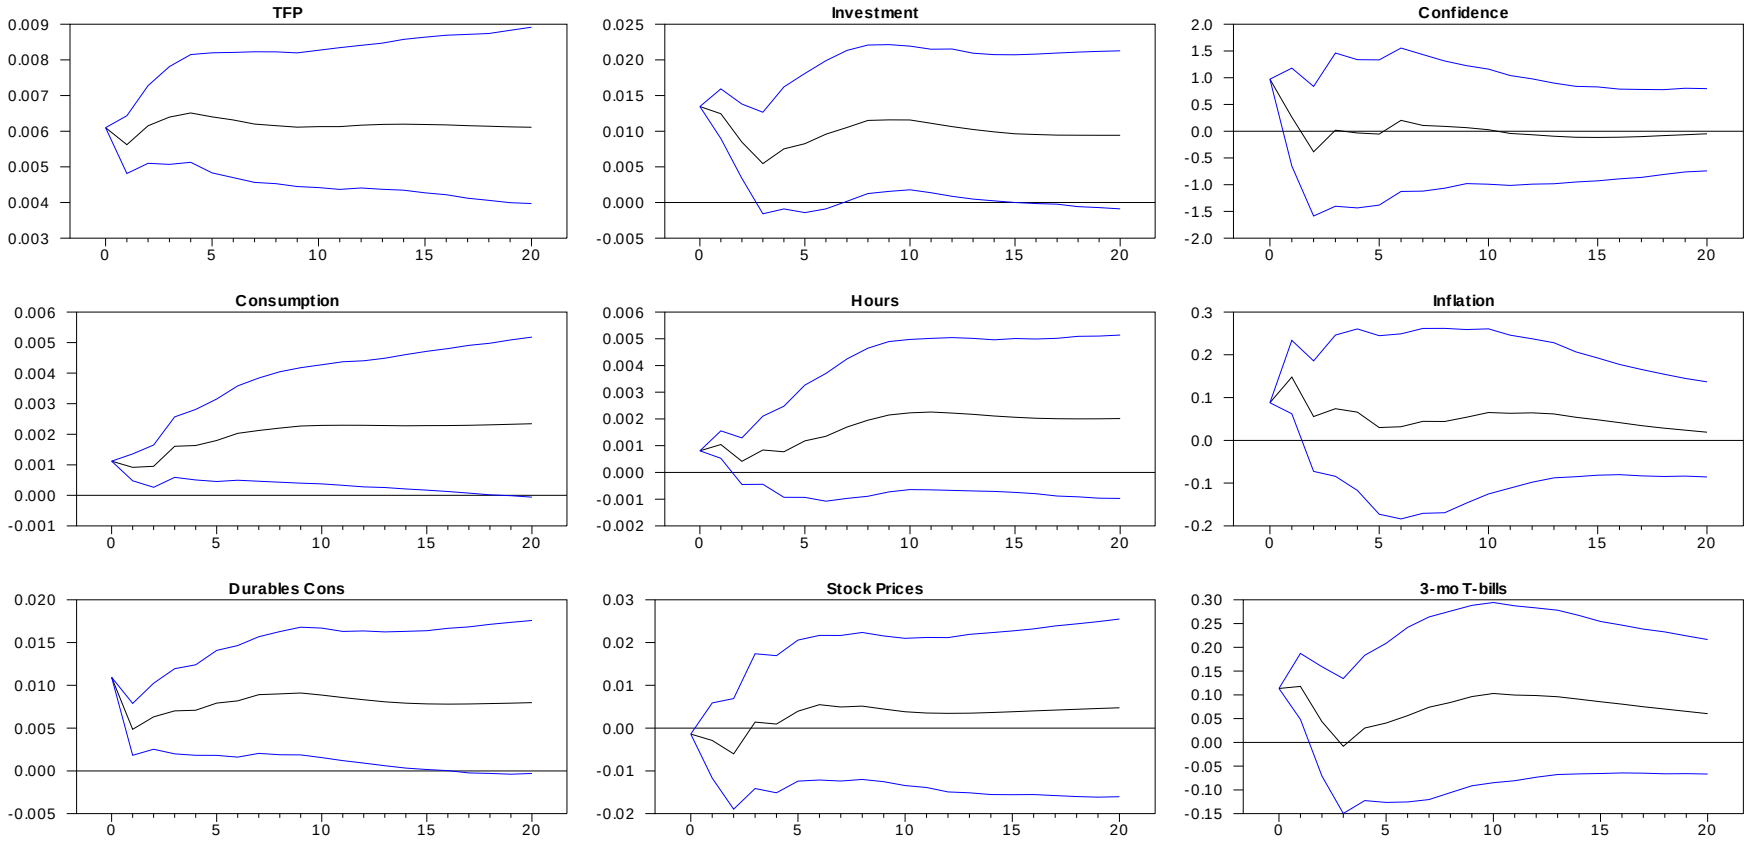
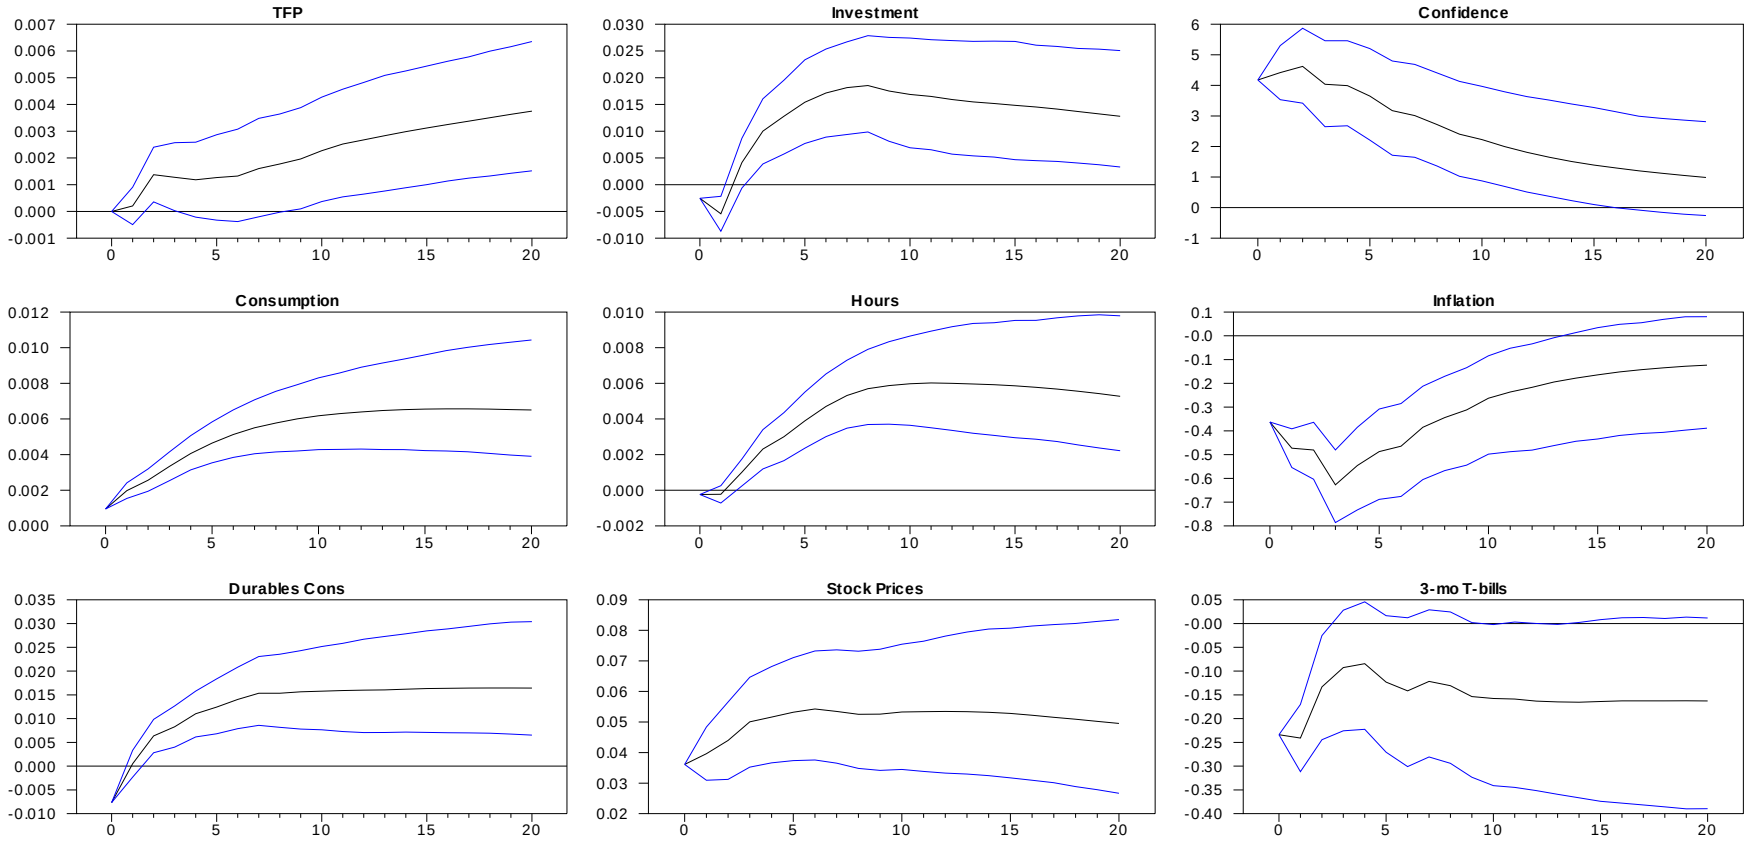


Figure 5: Impulse responses unanticipated TFP shock, using hybrid specification, where confidence, inflation, and t-bills are in levels and all else in differences. Black lines are point estimates of the impulse responses and blue lines are 5% and 95% percentile error bands obtained from Monte Carlo integration with 2000 replications.



Responses to News Shock (ECT)

Figure 6: Impulse responses to identified news shock, using error correction specification, where confidence, inflation, and t-bills are included in levels and all other variables in differences, and cointegration was imposed between TFP and Consumption, TFP and Durables Consumption, TFP and Investment, and TFP and Stock Prices. Black lines are point estimates of the impulse responses and blue lines are 5% and 95% percentile error bands obtained from Monte Carlo integration with 2000 replications.

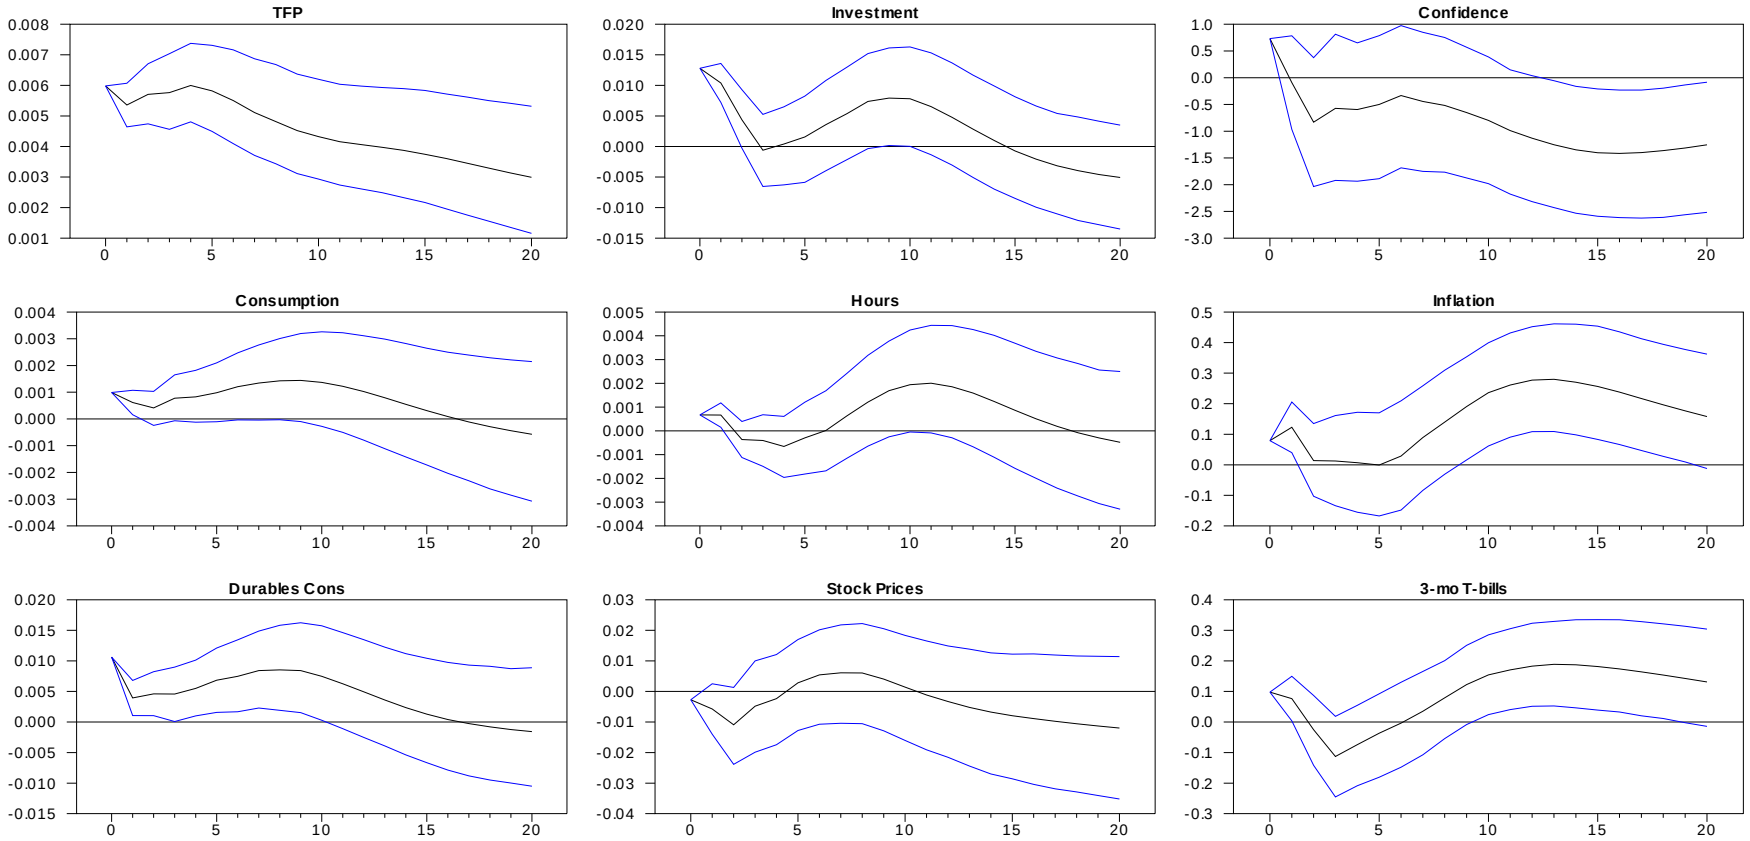


Figure 7: Impulse responses to unanticipated TFP shock, using error correction specification, where confidence, inflation, and t-bills are included in levels and all other variables in differences, and cointegration was imposed between TFP and Consumption, TFP and Durables Consumption, TFP and Investment, and TFP and Stock Prices. Black lines are point estimates of the impulse responses and blue lines are 5% and 95% percentile error bands obtained from Monte Carlo integration with 2000 replications.

Figure 8: Historical Decompositions (Hybrid)

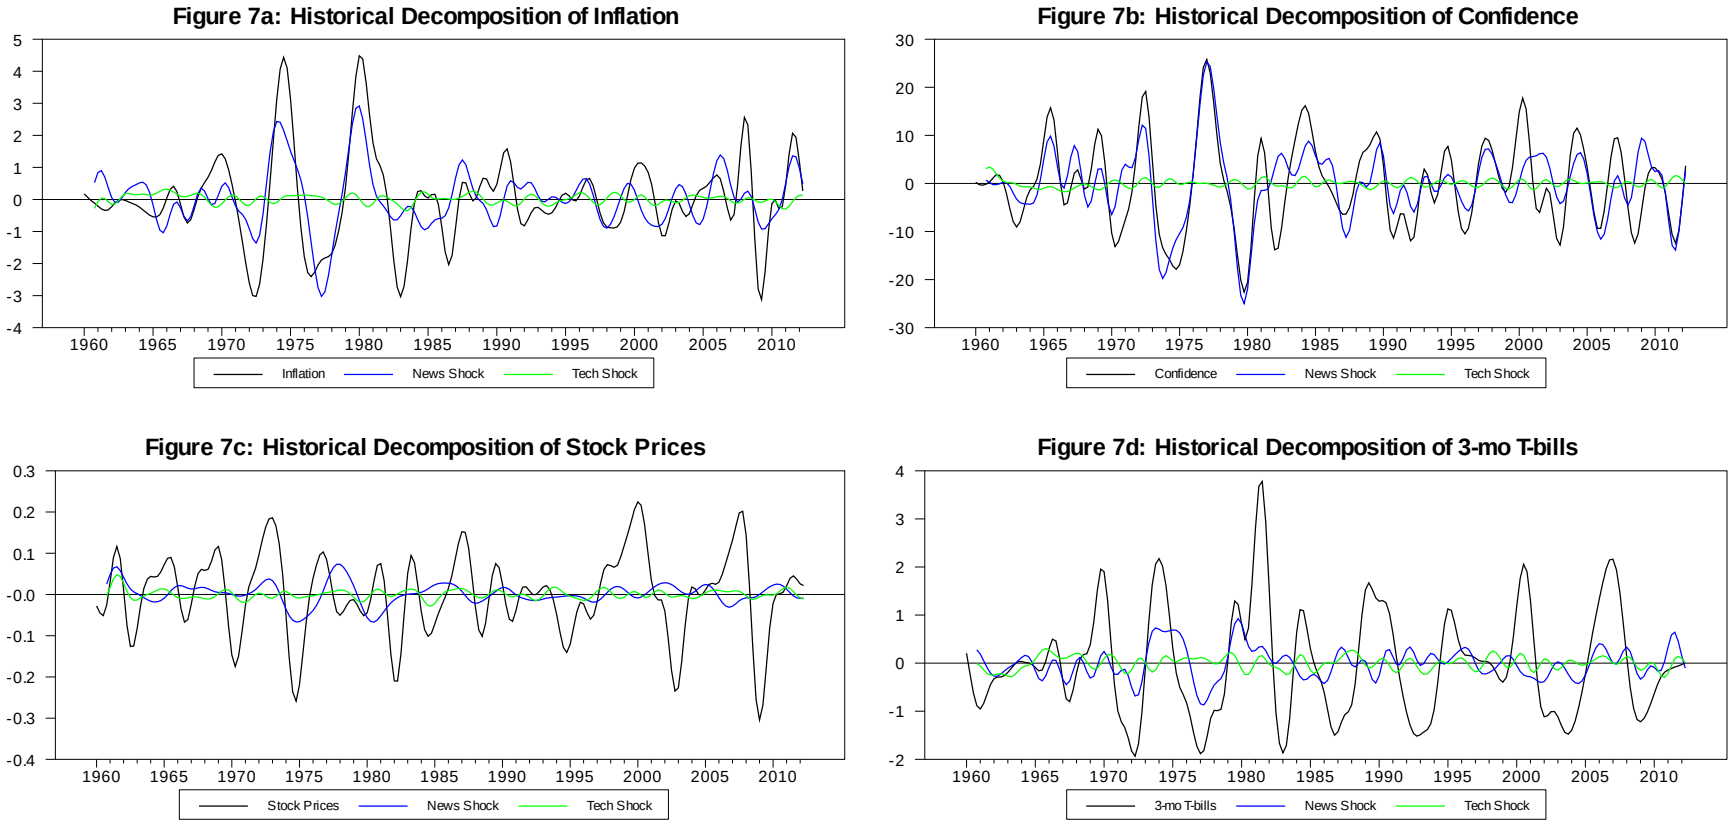


Figure 8: Historical decomposition of variables from our identified news shock and unanticipated TFP shock using the hybrid specification. Black line is the data, blue line is the simulated data if news shocks were the only disturbance, and green line is the simulated data if unanticipated TFP shocks were the only disturbance. The data and the historical decompositions were filtered using Baxter-King bandpass filter, with bounds set to 6-32 quarters.

Figure 8: Historical Decompositions (Hybrid) (Cont.)

Figure 7e: Historical Decomposition of TFP

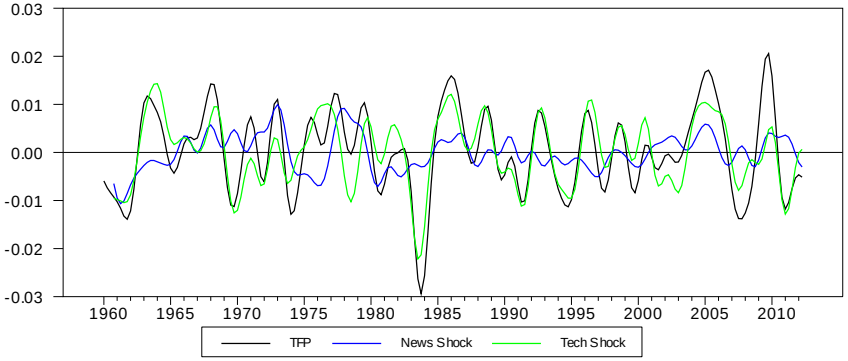


Figure 7f: Historical Decomposition of Consumption

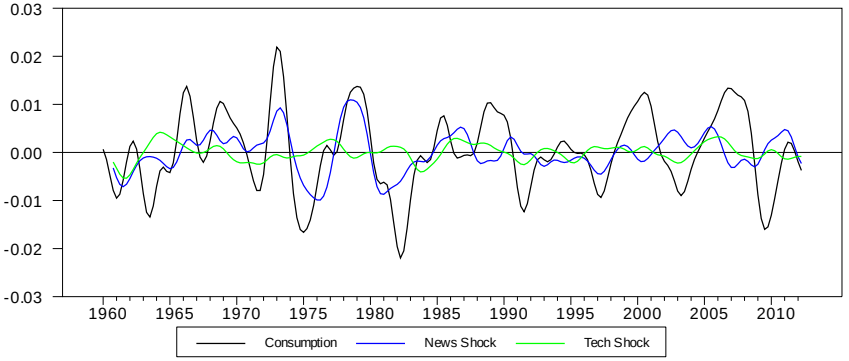


Figure 7g: Historical Decomposition of Investment

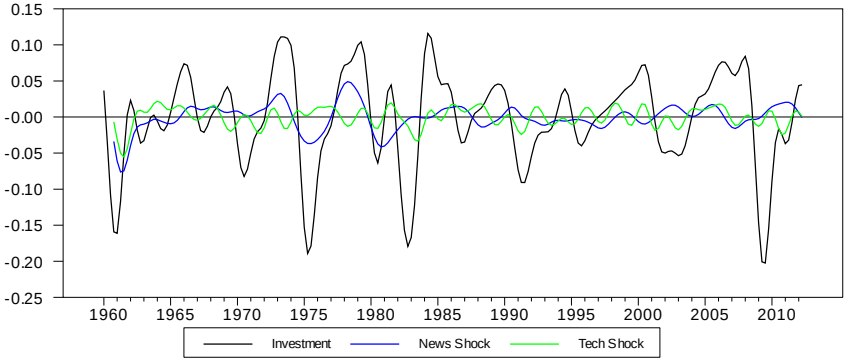


Figure 7h: Historical Decomposition of Hours

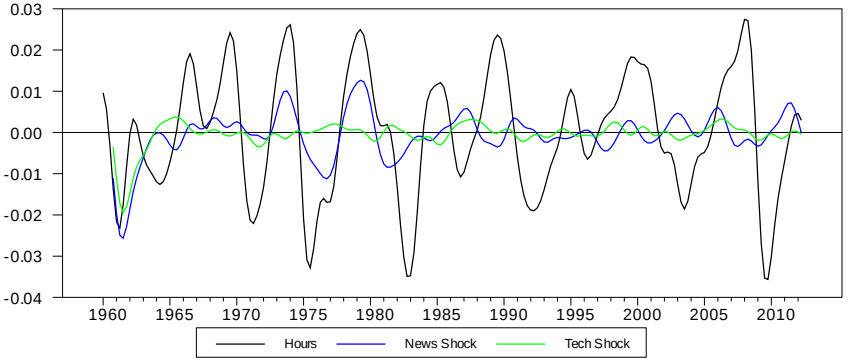


Figure 8: Historical Decompositions (Hybrid) (Cont.)

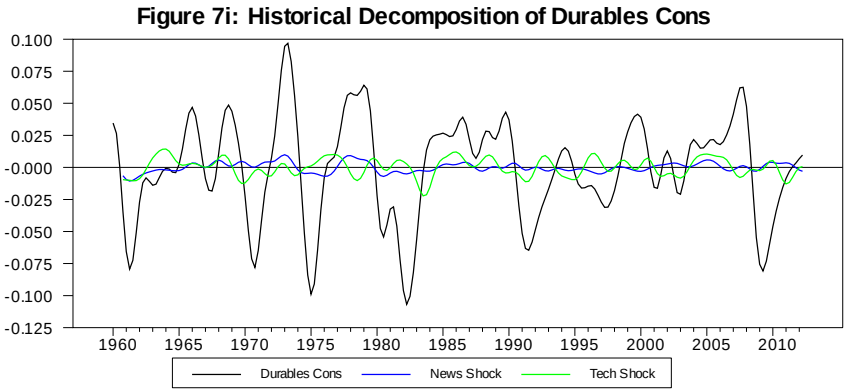


Figure 9: Historical Decompositions (Levels)

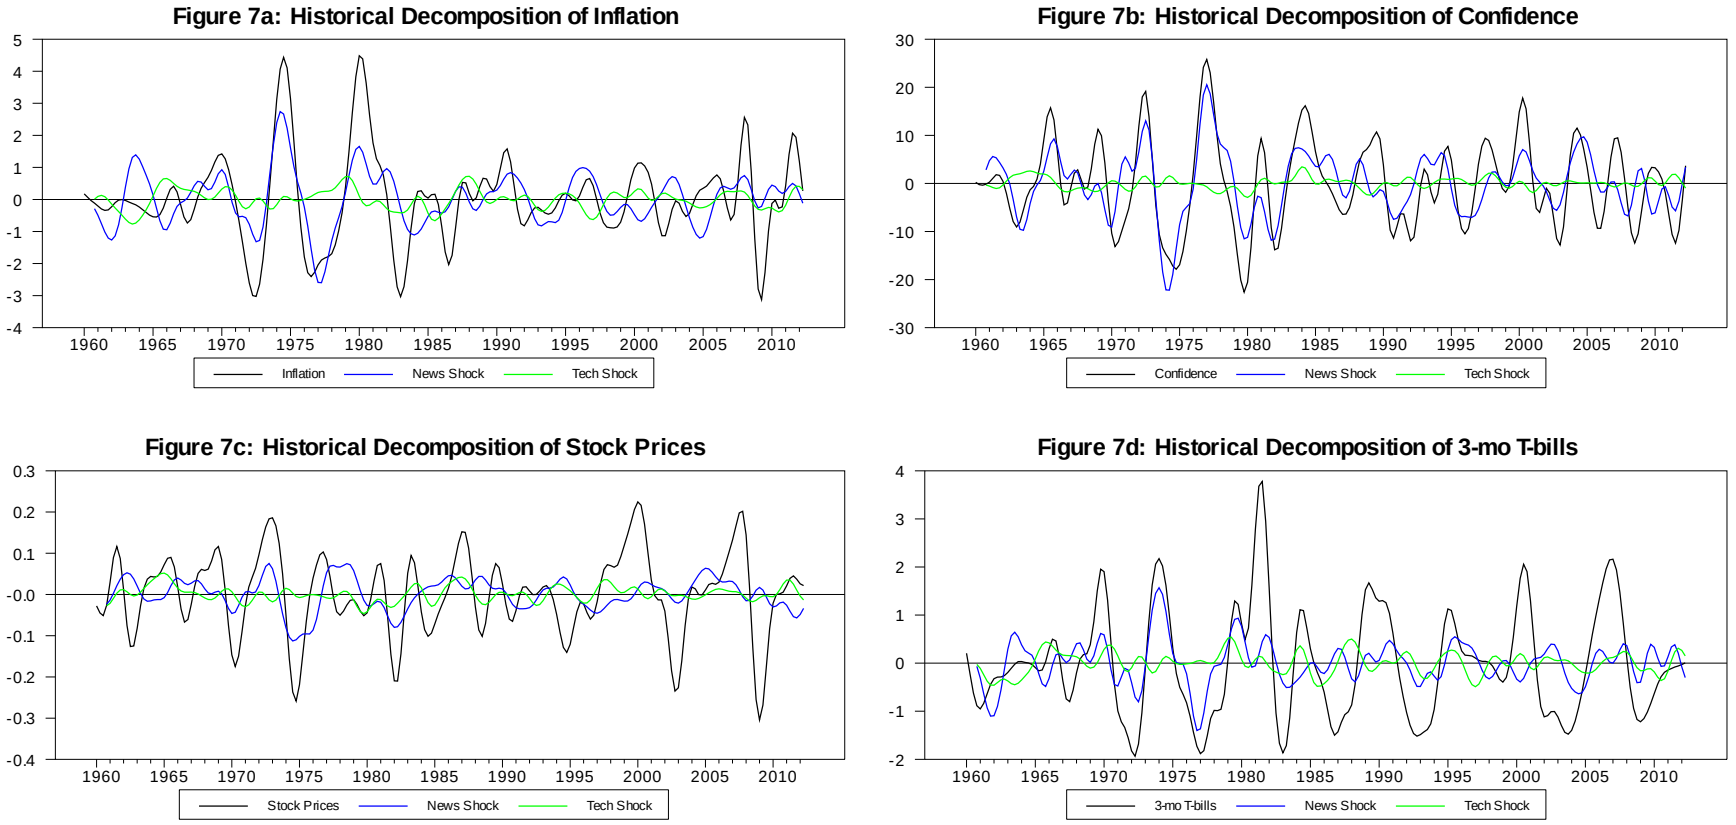


Figure 9: Historical decomposition of variables from our identified news shock and unanticipated TFP shock using the levels specification. Black line is the data, blue line is the simulated data if news shocks were the only disturbance, and green line is the simulated data if unanticipated TFP shocks were the only disturbance. The data and the historical decompositions were filtered using Baxter-King bandpass filter, with bounds set to 6-32 quarters.

Figure 9: Historical Decompositions (Levels) (Cont.)

Figure 7e: Historical Decomposition of TFP

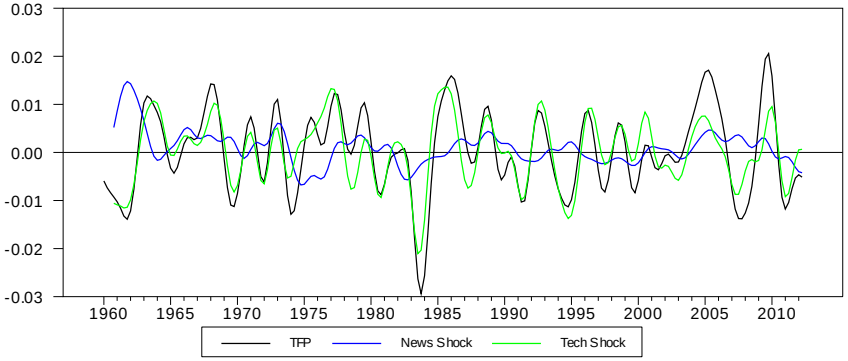


Figure 7f: Historical Decomposition of Consumption

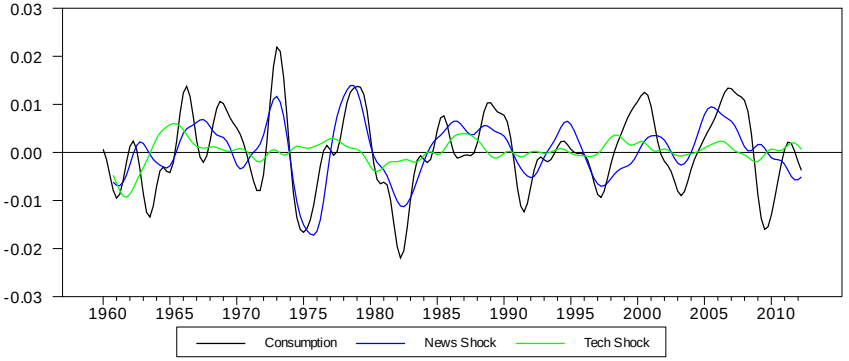


Figure 7g: Historical Decomposition of Investment

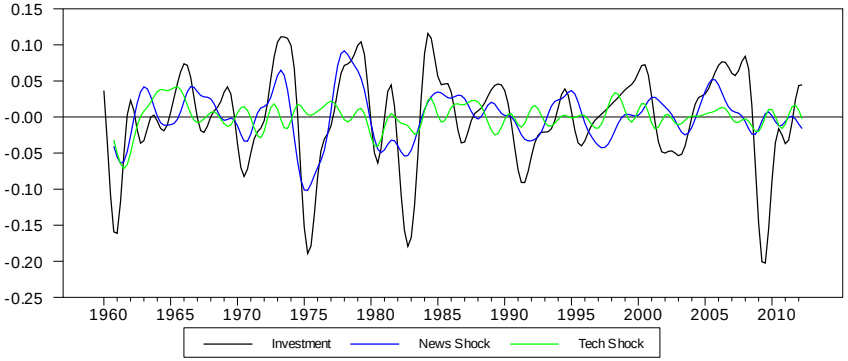


Figure 7h: Historical Decomposition of Hours

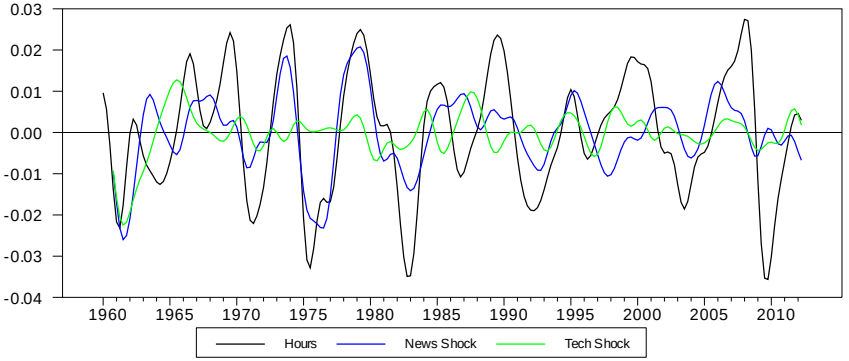


Figure 9: Historical Decompositions (Levels) (Cont.)

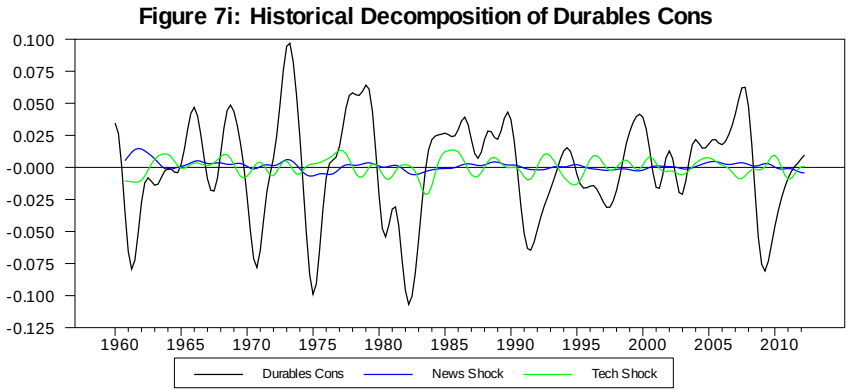


Figure 10: "Pure News" Effect of TFP News

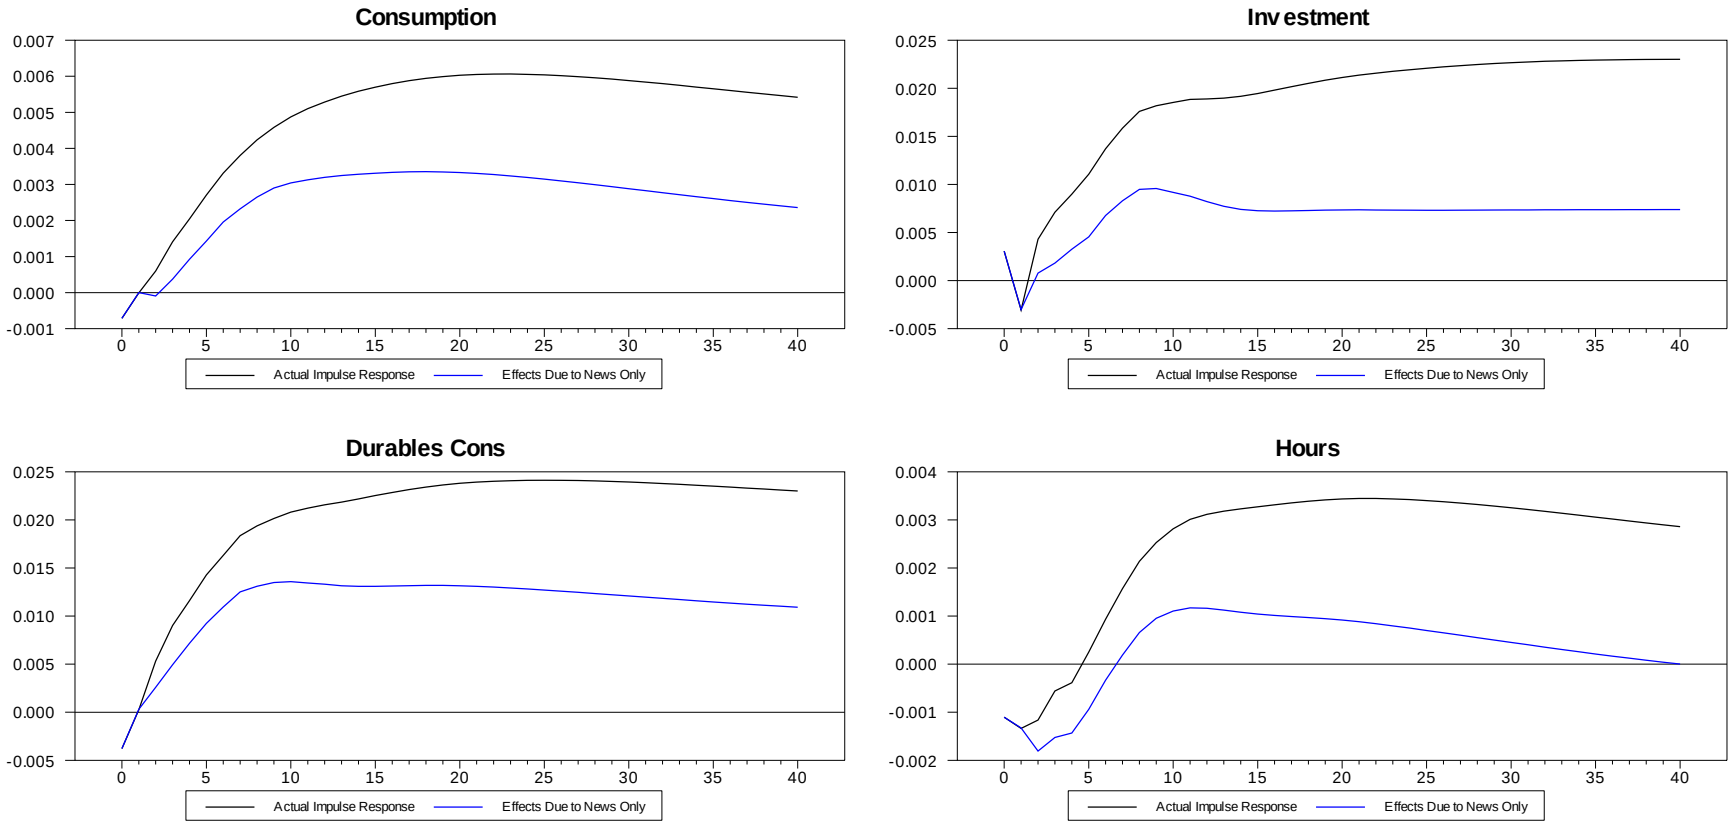
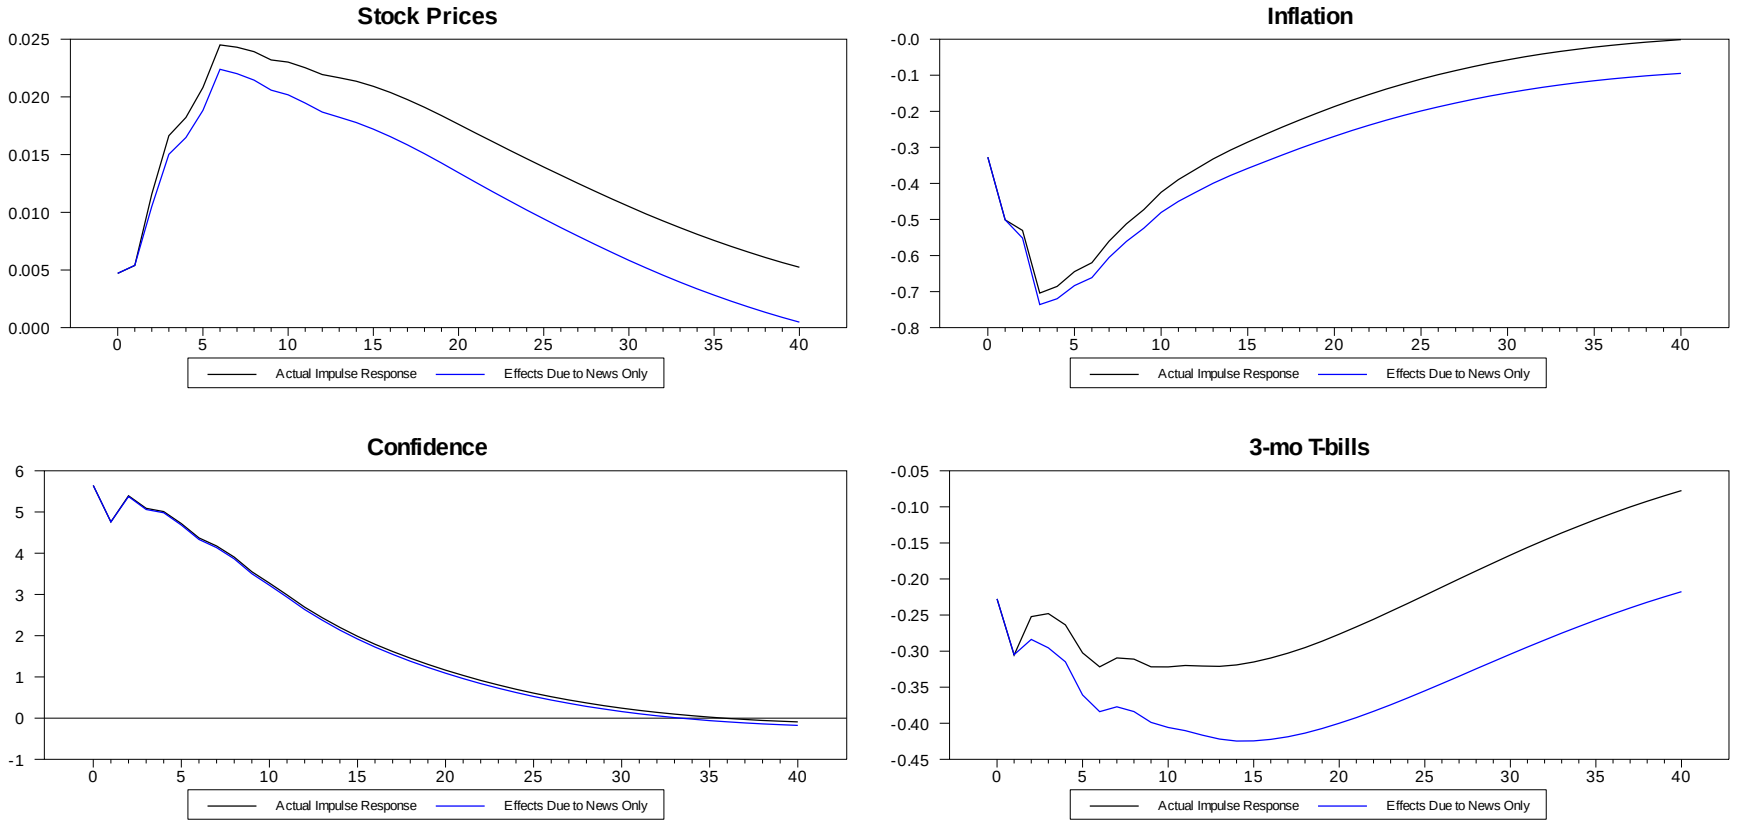
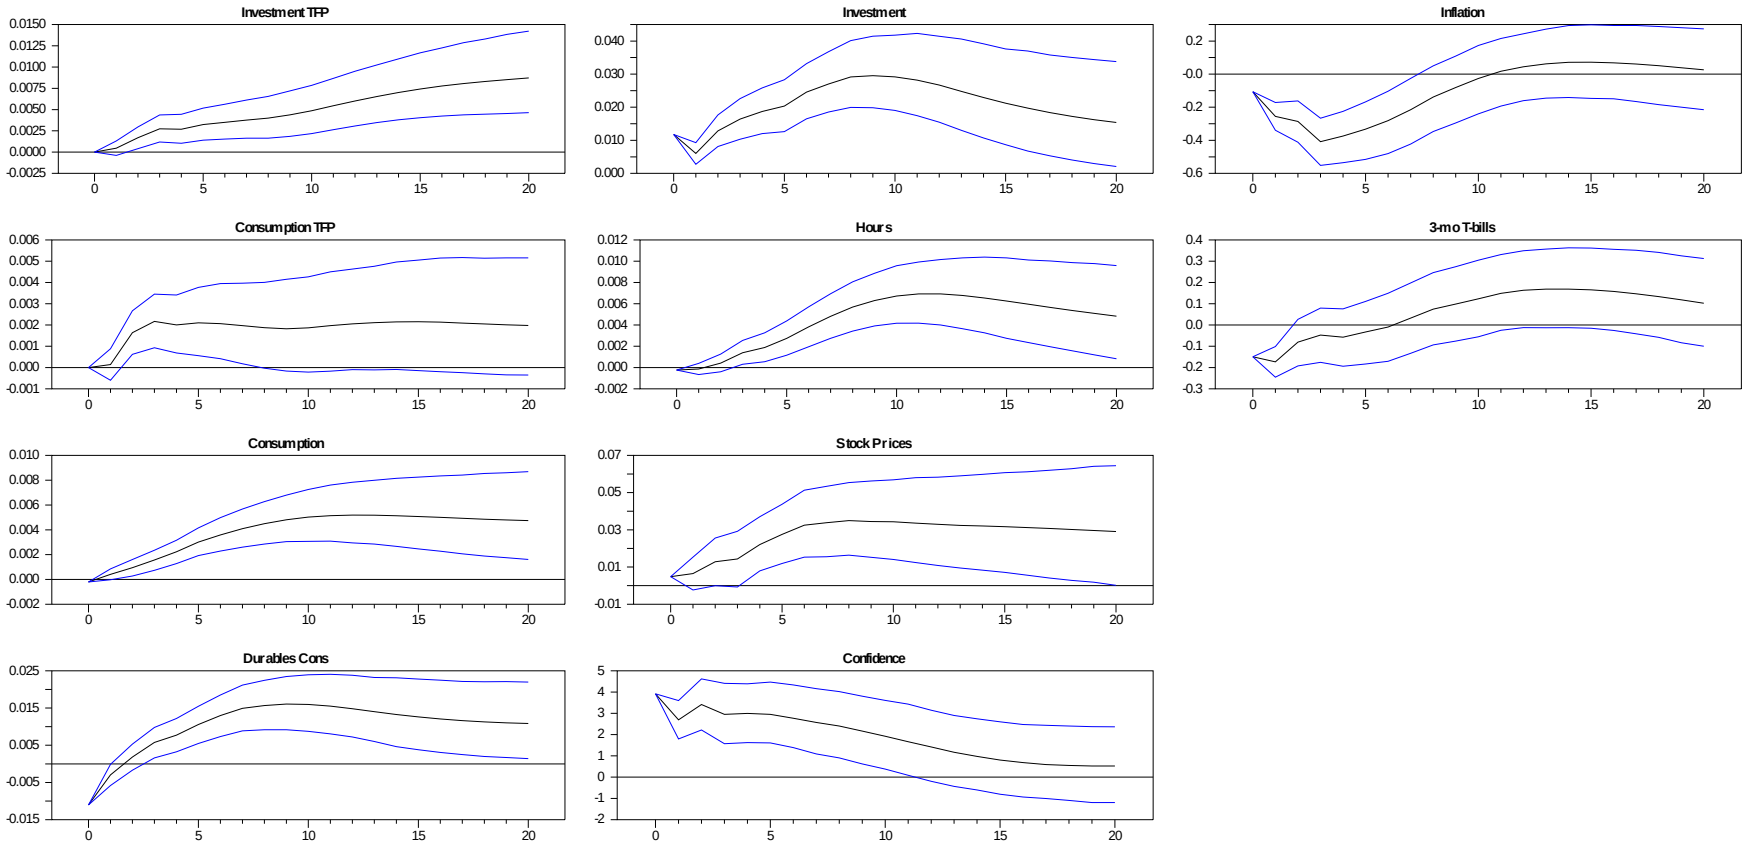


Figure 10: Correction of the impulse responses to identified news shock using the procedure detailed in the paper, using our hybrid specification. The black line is the actual impulse response from our identified news shock. The blue line is the corrected impulse response. To obtain the corrected impulse response, we take the impulse responses and subtract the correction factor. The correction factor is the impulse response of respective variables to unanticipated technology shock at quarter 10, multiplied by the ratio of rise in TFP from our identified news shock at quarter 10 to the rise in TFP from unanticipated technology shock at quarter 10.

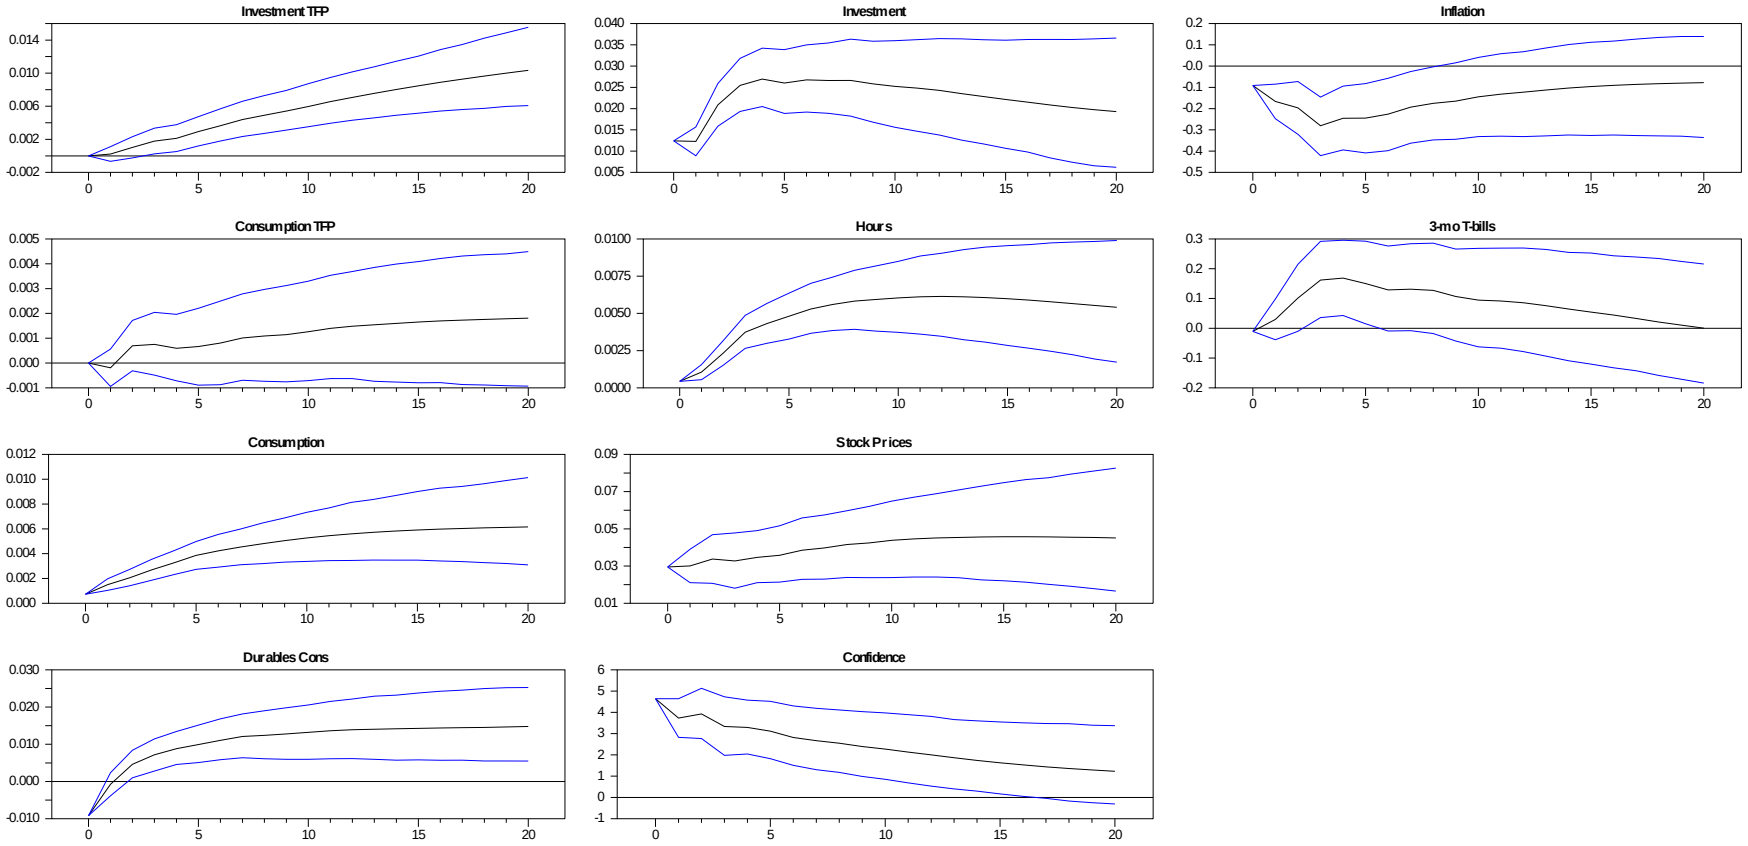
Figure 10: "Pure News" Effect of TFP News (Cont.)





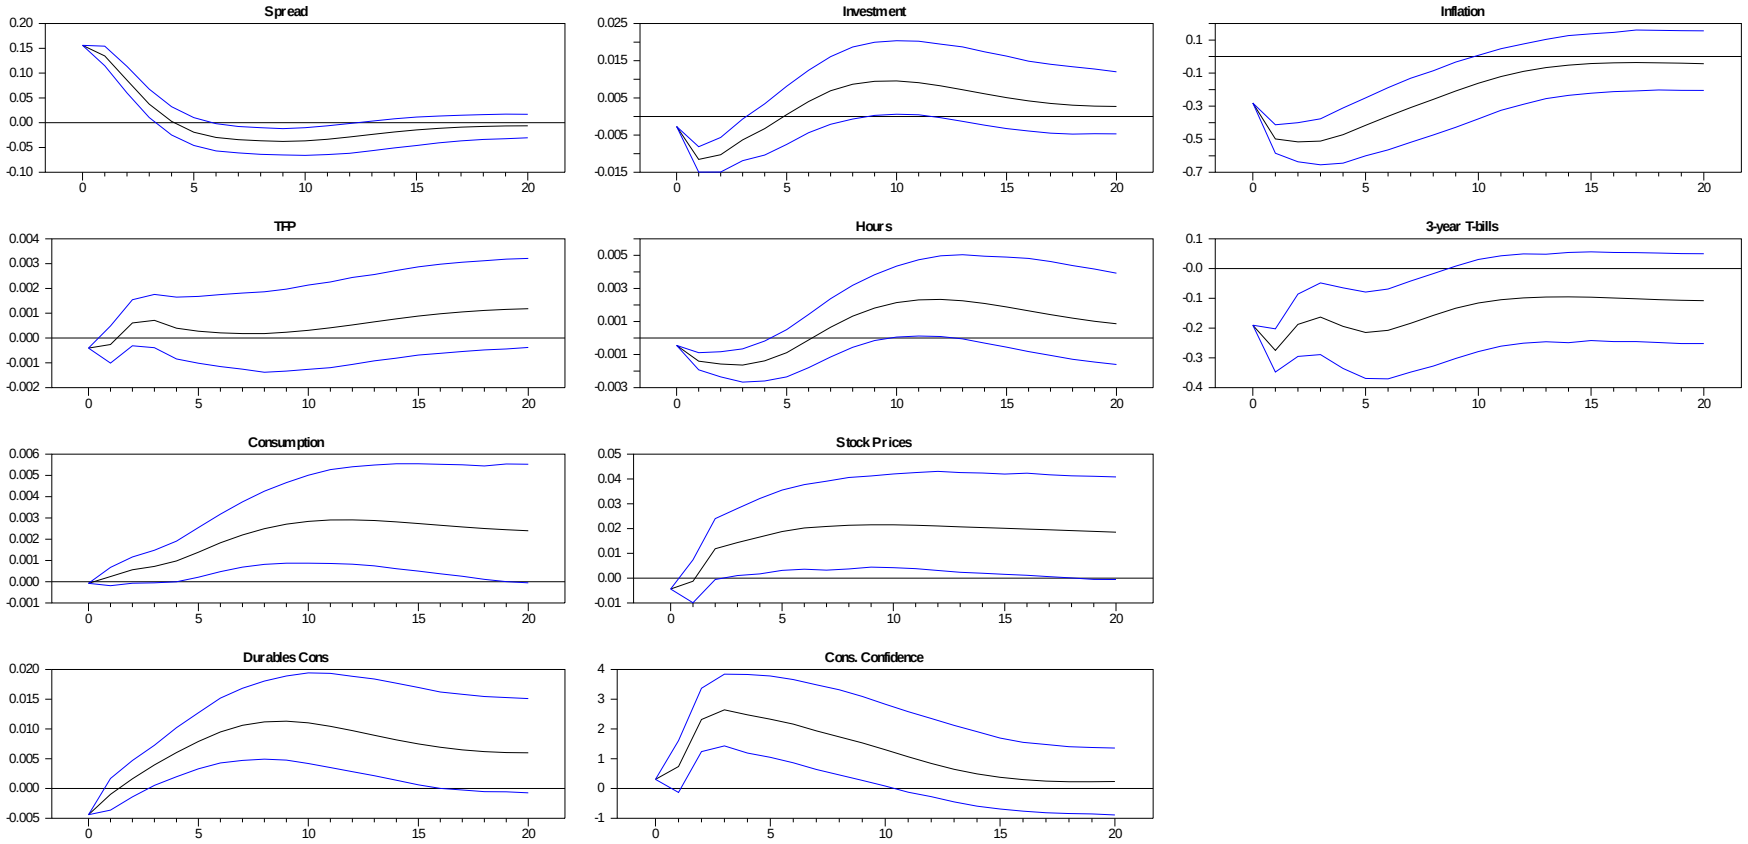
Responses to Consumption TFP News Shock (ECT)

Figure 11: Impulse responses to identified news shock for consumption TFP. This uses the error correction specification, where confidence, inflation, and t-bills are in levels and all else in differences, and cointegrating relationship was imposed between Consumption TFP and Consumption, and Investment TFP and Investment and Investment TFP and Durables Consumption, and both TFP on stock prices separately. Black lines are point estimates of the impulse responses and blue lines are 5% and 95% percentile error bands obtained from Monte Carlo integration with 2000 replications.



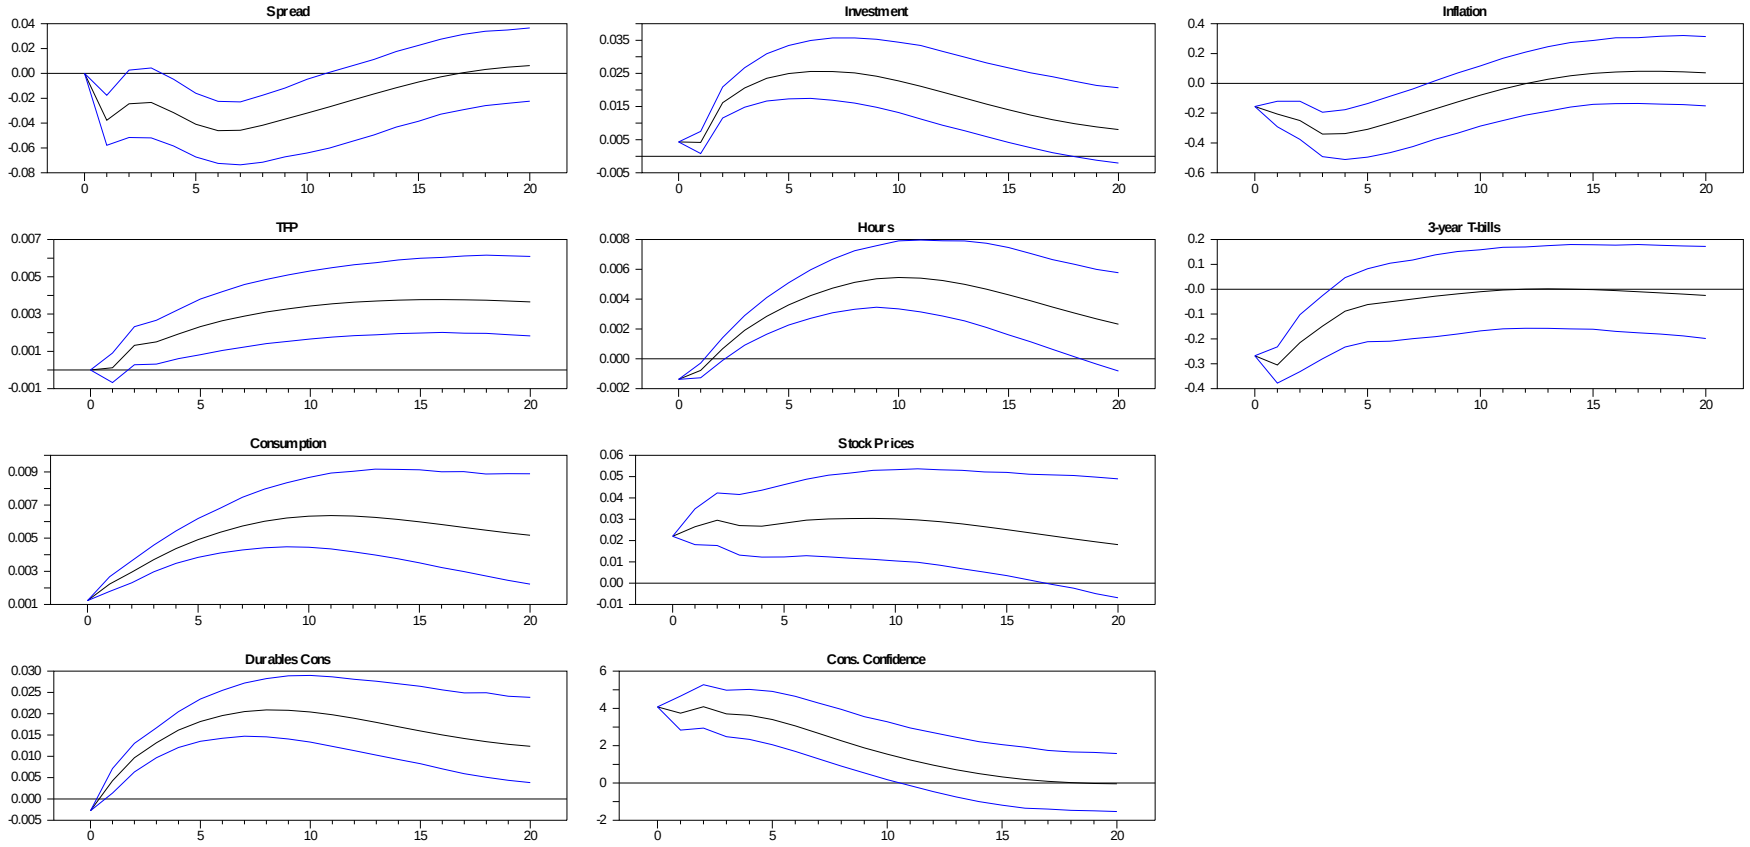
Responses to Investment TFP News Shock (ECT)

Figure 12: Impulse responses to identified news shock for investment TFP. This uses the cointegration specification, where confidence, inflation, and t-bills are in levels and all else in differences, and cointegrating relationship was imposed between Consumption TFP and Consumption, and Investment TFP and Investment and Investment TFP and Durables Consumption, and both TFP on stock prices separately. Black lines are point estimates of the impulse responses and blue lines are 5% and 95% percentile error bands obtained from Monte Carlo integration with 2000 replications.



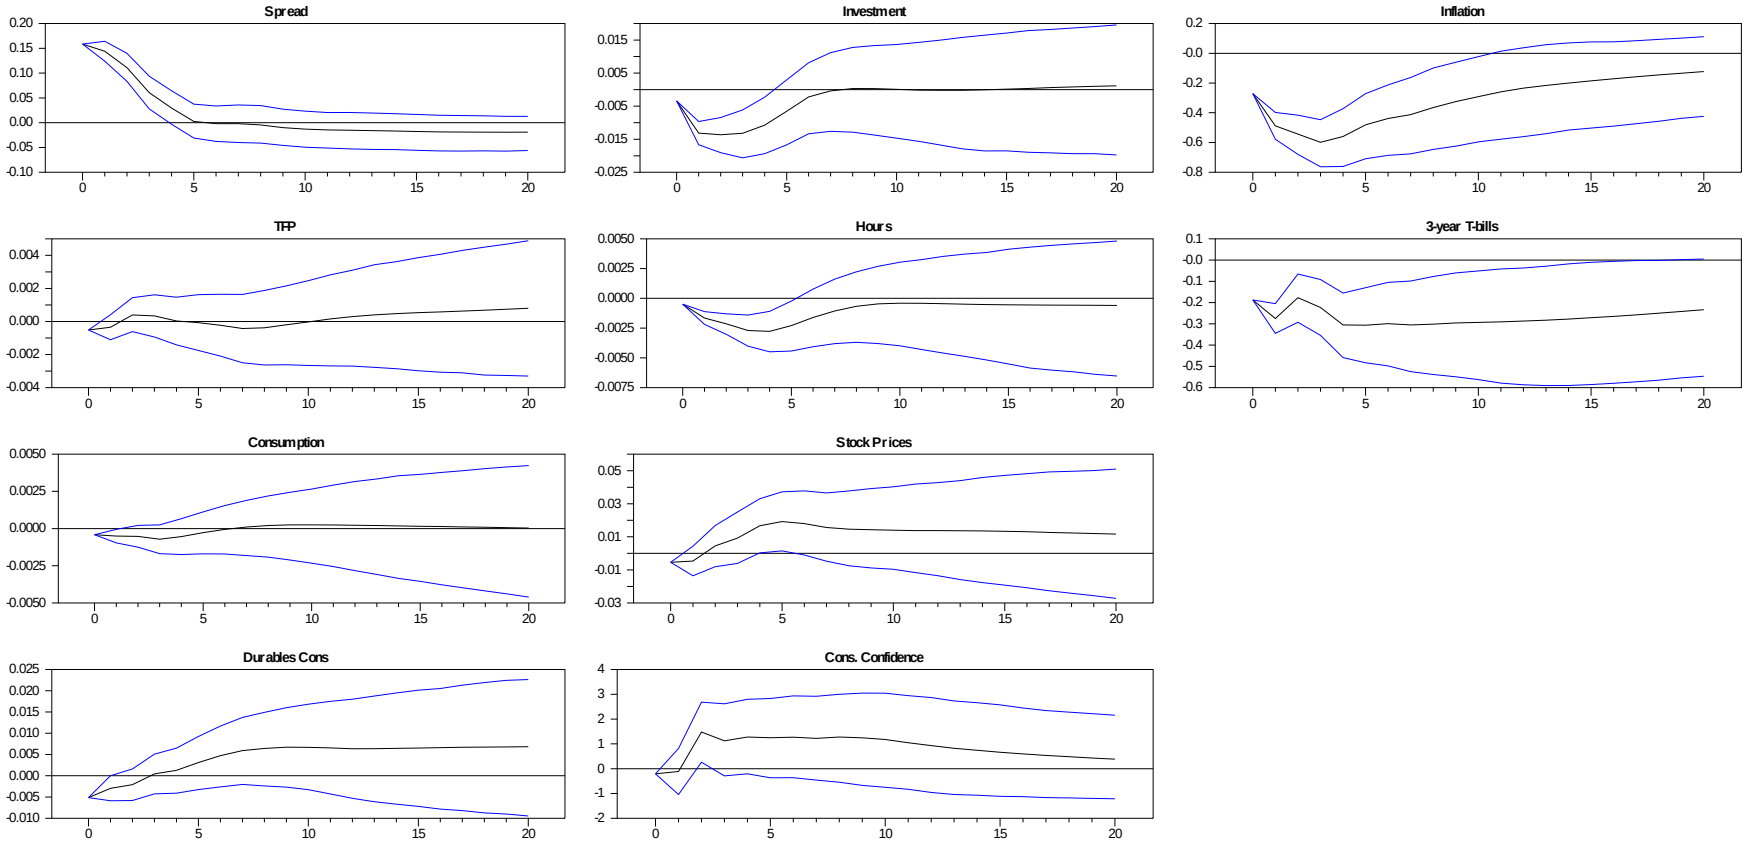
Responses to Spread Shock (Levels), H=20

Figure 13: Impulse Responses to spread shock, where spread is ordered first. Spread here is defined as the difference between Moody's Baa Corporate Bond Yield and Moody's Aaa Corporate Bond Yield. This was done with levels specification. Black lines are point estimates of the impulse responses and blue lines are 5% and 95% percentile error bands obtained from Monte Carlo integration with 2000 replications.



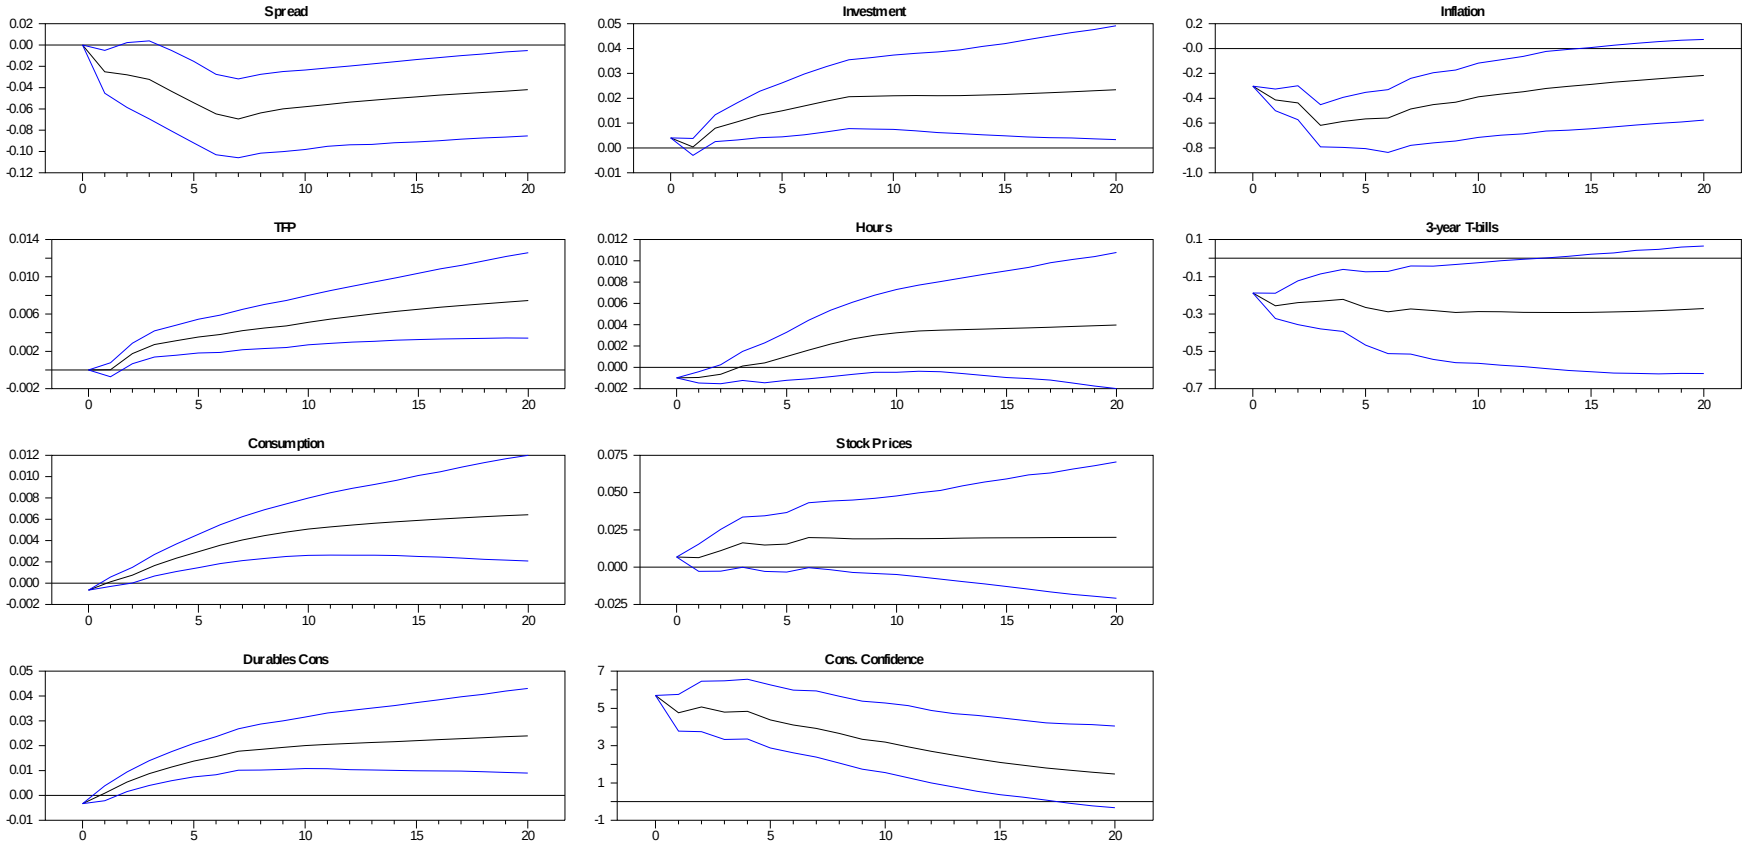
Responses to News Shock (Levels), H=20

Figure 14: Impulse Responses to news shock, orthogonalized with respect to spread and technology shocks. Spread here is defined as the difference between Moody's Baa Corporate Bond Yield and Moody's Aaa Corporate Bond Yield. This was done with levels specification. Black lines are point estimates of the impulse responses and blue lines are 5% and 95% percentile error bands obtained from Monte Carlo integration with 2000 replications.



Responses to Spread Shock (Hybrid), H=20

Figure 15: Impulse Responses to spread shock, where spread is ordered first. Spread here is defined as the difference between Moody's Baa Corporate Bond Yield and Moody's Aaa Corporate Bond Yield. This was done with hybrid specification. Black lines are point estimates of the impulse responses and blue lines are 5% and 95% percentile error bands obtained from Monte Carlo integration with 2000 replications.



Responses to News Shock (Hybrid), H=20

Figure 16: Impulse Responses to news shock, orthogonalized with respect to spread and technology shocks. Spread here is defined as the difference between Moody's Baa Corporate Bond Yield and Moody's Aaa Corporate Bond Yield. This was done with hybrid specification. Black lines are point estimates of the impulse responses and blue lines are 5% and 95% percentile error bands obtained from Monte Carlo integration with 2000 replications.

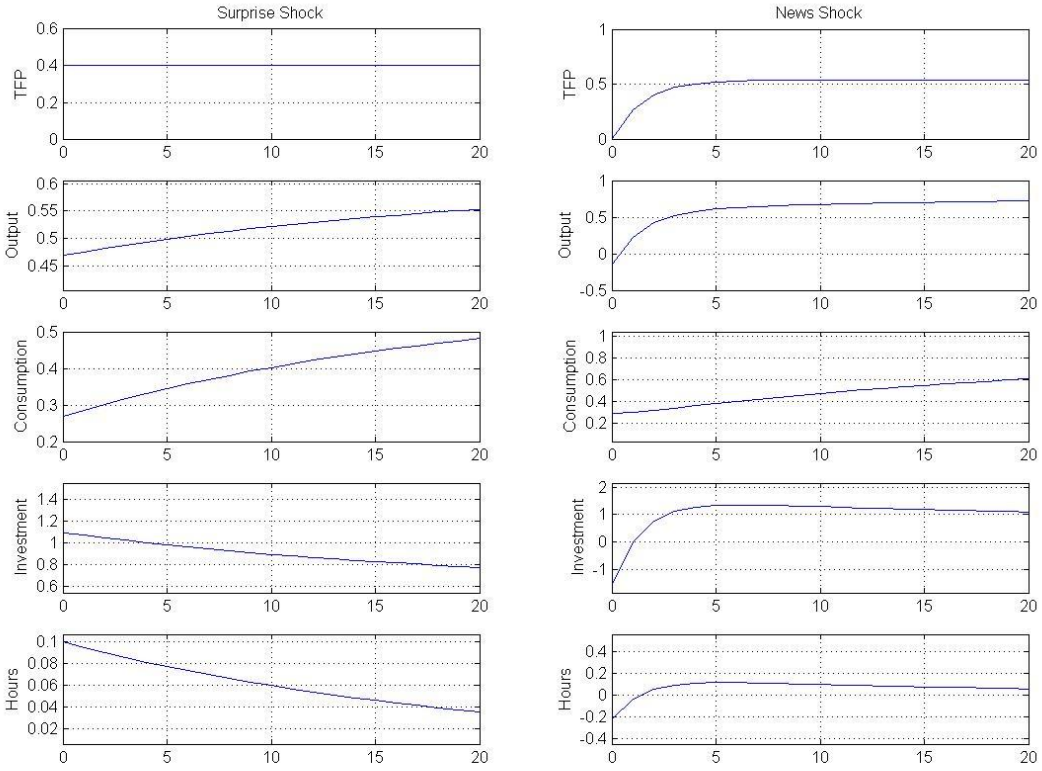
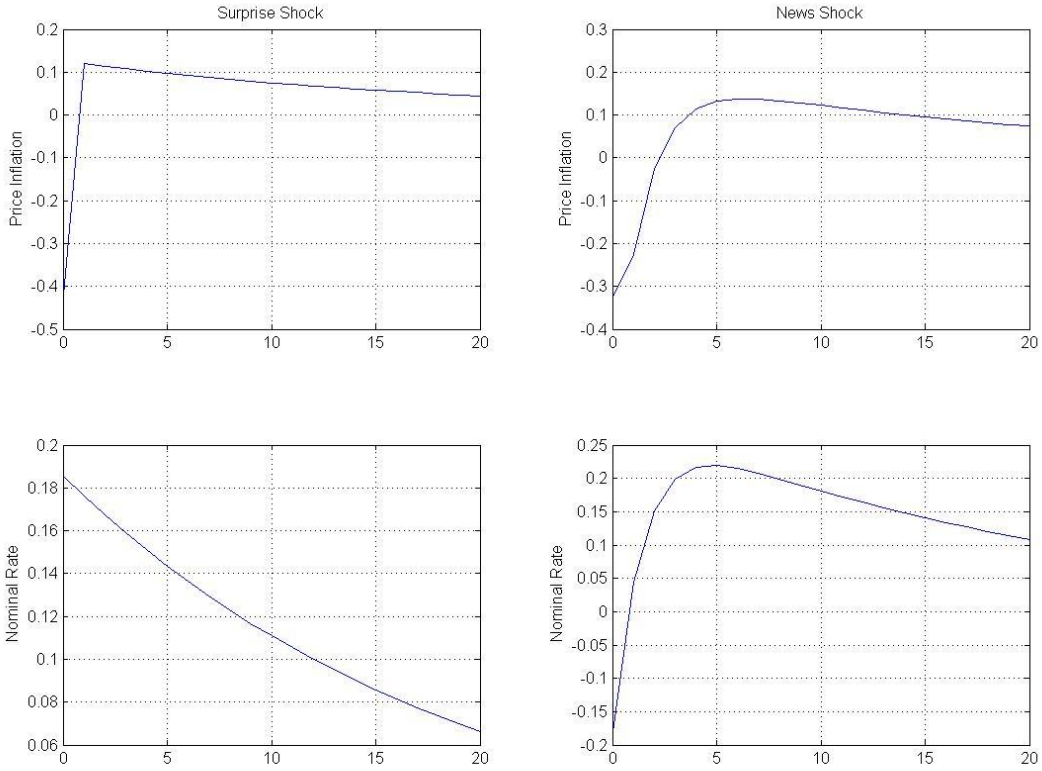


Figure 17: Theoretical impulse responses to a surprise technology shock and news shock in the flexible prices model with capital and monetary policy rule responding to output growth and inflation.

Figure 17 (Cont.)



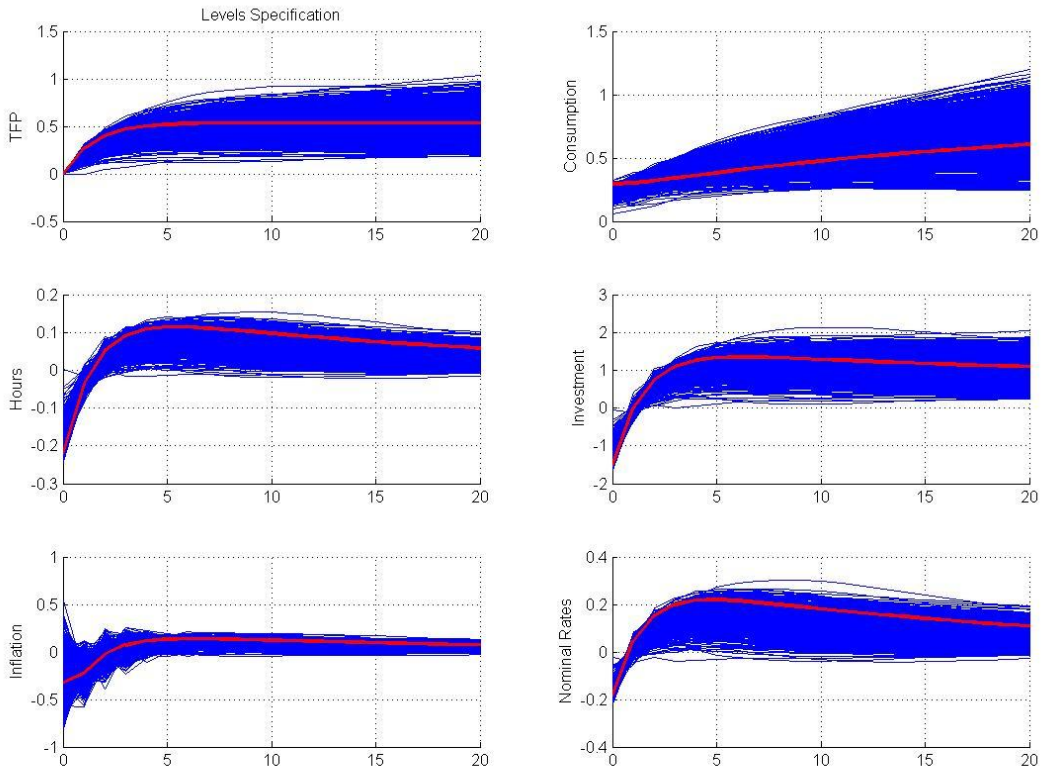


Figure 18: The red line shows theoretical impulse response to a news shock from the model presented in the paper. The blue lines are estimated impulse responses using levels specification from a Monte Carlo simulation with 1000 repetitions and 250 observations per repetition.

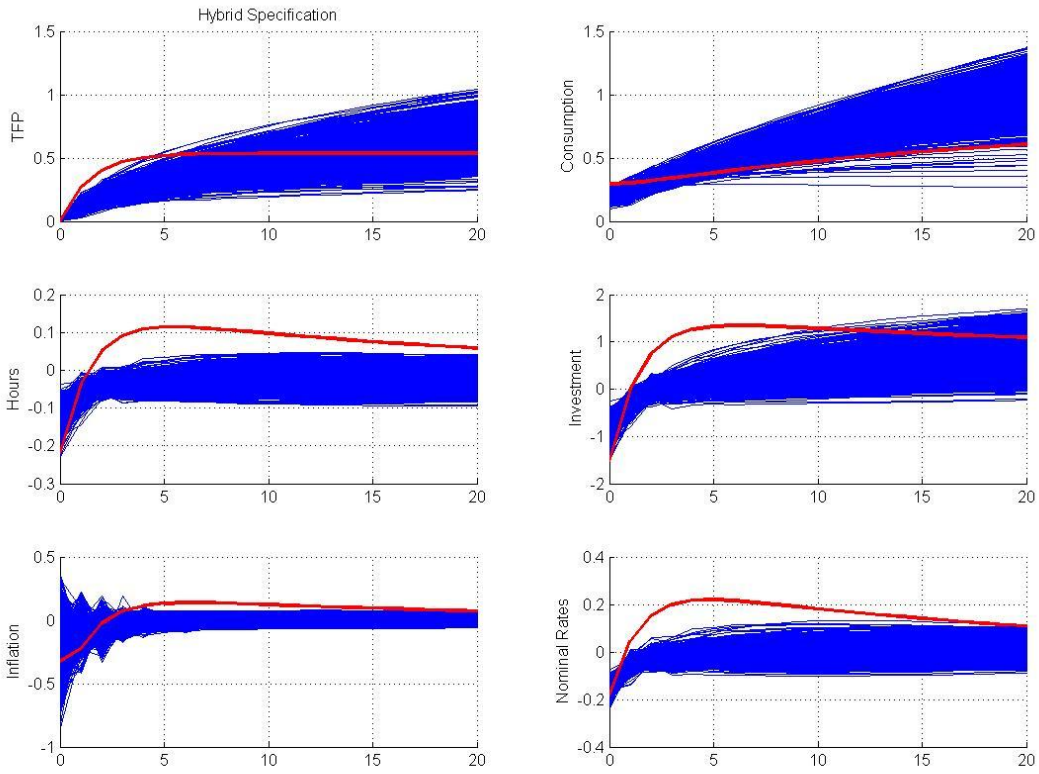


Figure 19: The red line shows theoretical impulse response to a news shock from the model presented in the paper. The blue lines are estimated impulse responses using hybrid specification from a Monte Carlo simulation with 1000 repetitions and 250 observations per repetition.

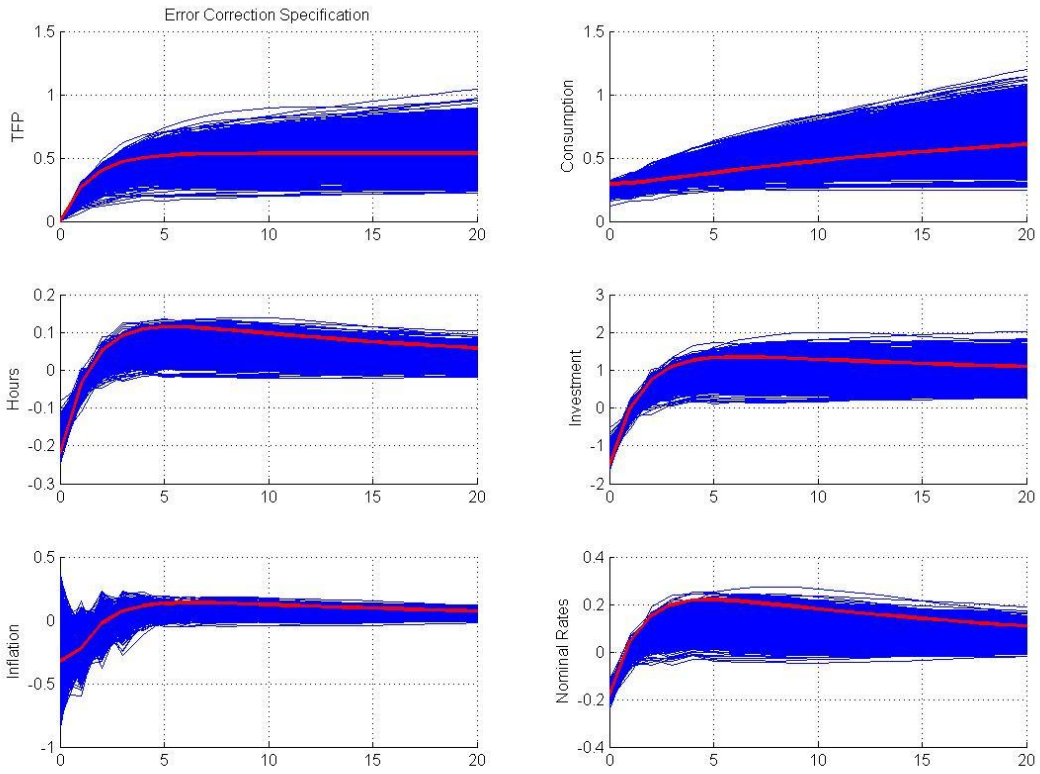


Figure 20: The red line shows theoretical impulse response to a news shock from the model presented in the paper. The blue lines are estimated impulse responses using error correction specification from a Monte Carlo simulation with 1000 repetitions and 250 observations per repetition. Cointegration was imposed between TFP and Consumption, and TFP and Investment as specified in the model.