Latent Variables Analysis in Structural Models: A New Decomposition of the Kalman Smoother

Hess Chung, Cristina Fuentes-Albero, Matthias Paustian, and Damjan Pfajfar

Latent Variables Analysis in Structural Models: A New Decomposition of the Kalman Smoother[*]

Hess Chung[†,‖]
FEDERAL RESERVE BOARD

Cristina Fuentes-Albero[†,‡]
FEDERAL RESERVE BOARD

Matthias Paustian[†,§]
FEDERAL RESERVE BOARD

Damjan Pfajfar[†,¶]
FEDERAL RESERVE BOARD

December 2, 2020

Abstract

This paper advocates chaining the decomposition of shocks into contributions from forecast errors to the shock decomposition of the latent vector to better understand model inference about latent variables. Such a double decomposition allows us to gauge the influence of data on latent variables, like the data decomposition. However, by taking into account the transmission mechanisms of each type of shock, we can highlight the economic structure underlying the relationship between the data and the latent variables. We demonstrate the usefulness of this approach by detailing the role of observable variables in estimating the output gap in two models.

JEL CLASSIFICATION: C18, C32, C52
KEYWORDS: Kalman smoother, latent variables, shock decomposition, data decomposition, double decomposition

---

1. Introduction

Researchers estimating models using Bayesian methods or maximum likelihood have employed several different ways of decomposing the latent state vector path to better understand and communicate model behavior. Shock decompositions, for example, in which researchers decompose the estimated latent vector into contributions of the estimated structural shocks, have long been standard in the literature. In cases where the model is linear and the shocks are Gaussian, moreover, the latent variables are linear functions of the data. Thus, following Sander (2013) and Andrle (2013), researchers can decompose the estimated latent vector into contributions of the observed variables using the so-called data decomposition. In this paper, we propose linking the data decomposition of the vector of structural shocks and the shock decomposition of the latent vector in what we label as the double decomposition. This way of analyzing the latent variable path can be particularly illuminating when the focus is on inference regarding highly theoretical constructs, such as the natural rate of interest or the "flex-price" output gap, where the relation between the observable variables and the latent variable is not intuitive and depends heavily on the model's theoretical structure.

We first show how the double decomposition can be used to study the behavior of a latent variable—the output gap—in a simple model by Ireland (2011). After illustrating the mechanics of the double decomposition, we focus on the output gap in a more complex model—a version of the model presented in Del Negro, Giannoni and Schorfheide (2015). Using this model, where inference is more complex, we illustrate the value of the double decomposition for practitioners. A data decomposition suggests that the estimated path of the output gap in this model shows the influence of a large number of observables. At the same time, a shock decomposition shows that the output gap is primarily driven by permanent technology shocks and shocks to idiosyncratic firm risk (effectively, shocks to the spread between the return on capital and the return on risk-free assets). Thus, it is illuminating to start with the relationship between the observables that carry a large signal about the output gap, according to the data decomposition, and the model's estimates of the shocks which largely drive the output gap, according to the shock decomposition. This process is formalized in the double decomposition.

We also illustrate the value of the double decomposition for interpreting the model's reaction to incoming news, examining in detail the effects of news flow between the end of 2013 and the beginning of 2014. The incoming news over this period featured large forecast errors in several highly influential observables and correspondingly large movements in important latent variables, such as the output gap. In such circumstances, the double decomposition can provide a clear economic interpretation of the effects of news on these variables.

Formally, let us consider a model with a linear-Gaussian structure. Let Y_t be the vector of observable variables, X_t be the latent vector, and η_t be the vector of structural shocks. The estimated vectors of latent variables and structural shocks depend on current, past, and future forecast errors. We decompose inference about endogenous latent variables in two steps. First, the linear-Gaussian structure of the model implies that structural shocks at some arbitrary time t, η_t, and the matrix

of forecast errors for the observable vector over the sample, ν, are jointly normally distributed.[1] As a result, we can write $E\left(\eta_t|\nu\right) = K(t)\nu$ for some matrix $K(t)$, which is a function of time, but, importantly, independent of the matrix of one-step ahead forecast errors. The data decomposition exploits this fact to decompose estimated shocks into contributions of each observable at every date in the sample, holding the other observables constant. Second, linearity also allows us to write the estimated value of a latent variable i at time t as $E\left(X_t^i|Y\right) = \left[\sum_{s=0}^{\infty} \Phi(\theta)^s K(t-s)\right]\nu$, where $Y = \begin{bmatrix} Y_1 & Y_2 & \ldots & Y_T \end{bmatrix}'$, which facilitates tracing the influence of news on the estimates of the latent variables operating through the transmission mechanism of each type of structural shock separately.

Both steps of our decomposition have been known in the literature. The shock decomposition is used to understand model dynamics but does not provide a linkage between data and latent variables. The data decomposition provides such a linkage but without a transparent relation to the underlying causal mechanisms which, as we emphasize, explain why the observables and the latent variables are linked. Because the double decomposition traces the influence of the data on latent variables through estimates of the structural shocks and their subsequent propagation, this causal narrative is highlighted. Moreover, as the double decomposition emphasizes inference about structural shocks, we focus on the role of forecast errors, rather than the level of the data, as is more common in the literature. This accounting is better for our purposes because contemporaneous forecast errors are tightly linked to model estimates of contemporaneous structural shocks, but much less so to past shocks. By contrast, the influence of a quarter of data in levels corresponds to large structural shocks both in that quarter and also offsetting shocks in the surrounding periods, which complicates the relationship between the observables and the latent variables of interest.[2]

The rest of the paper is organized as follows. Section 2 overviews the analysis of the latent vector in structural models. Section 3 then illustrates how the double decomposition works in a simple 3-equation New Keynesian model. We then turn to a more complex case, a medium-scale DSGE model, in Section 4. Section 5 concludes.

2.   Latent Variable Analysis in Structural Models

Traditionally, empirical macroeconomic researchers focus on the so-called shock decomposition, which decomposes the latent variables into contributions of structural shocks (see Appendix A.2). Recently, exploiting the linearity of the model, Andrle (2013) and Sander (2013) propose the data decomposition, which traces the independent effect of each observable on the estimated latent vector (see Appendix A.3). While the shock decomposition does not provide a link between the data and the latent vector, the data decomposition, which does describe such a relationship, does not provide a causal narrative underlying the relation. We address these shortcomings by proposing a new decomposition that does a data decomposition shock by shock. Given the nature of our

---

[1]The matrix of forecast errors is defined as $\nu = \begin{bmatrix} \nu_1 & \nu_2 & \ldots & \nu_T \end{bmatrix}'$, where ν_t is the vector of forecast errors at time t.

[2]A data decomposition in levels does have the advantage that the contribution of a given quarter of data tends to be more closely aligned in time with the corresponding effect on latent variables, while as we discuss below, this is not always the case in a data decomposition in forecast errors.

decomposition, we label it as the double decomposition. All these decompositions can be applied to any linear model with Gaussian shocks.

Let us consider the state space representation of a linear model

$$X_t = \Phi(\theta) X_{t-1} + R(\theta) \eta_t \tag{1}$$

$$Y_t = Z(\theta) X_t + \varepsilon_t \tag{2}$$

where X_t is an $m \times 1$ vector of state variables, η_t is a $p \times 1$, random vector with innovations or structural shocks such that $\eta_t \sim \mathcal{N}(0, Q)$, Y_t is an $n \times 1$ vector of observable variables and $\varepsilon_t \sim \mathcal{N}(0, H)$ is an $n \times 1$ vector of measurement errors. We assume that the structural shocks and measurement errors are uncorrelated.

In this framework, estimates of the latent vector are only updated if there is a forecast error. Hence, we proceed to derive the historical decompositions in terms of forecast errors instead of in terms of the data matrix.[3] Let ν_t be the one-step ahead forecast error

$$\nu_t = Y_t - \widehat{Y}_{t|t-1} = Z\left(X_t - \widehat{X}_{t|t-1}\right) + \varepsilon_t \tag{3}$$

We first derive the historical data decomposition of the estimated vector of structural shocks. As shown in Appendix A.1, the smoothed estimate of the vector of structural shocks can be written as a weighted sum of current and future one-step ahead forecast errors, ν_{t+h},

$$\begin{aligned}
\widehat{\eta}_{t|T} &= QR'\left[Z'F_{t|t-1}^{-1}\nu_t + L'_{t|t-1}Z'F_{t+1|t}^{-1}\nu_{t+1}\right.\\
&\quad \left. + \ldots + L'_{t|t-1}L'_{t+1|t}\ldots L'_{T|T-1}Z'F_{T|T-1}^{-1}\nu_T\right]
\end{aligned}$$

with $L_{t|t-1} = \Phi\left[I - P_{t|t-1}Z'F_{t|t-1}^{-1}Z\right]$, $F_{t|t-1} = ZP_{t|t-1}Z' + H$, and $P_{t|t-1} = \Phi P_{t-1|t-1}\Phi' + RQR'$. Let us define the observation weights for the forward recursions as

$$\Omega^{\star}_{t,\tau} = \begin{cases} I & \tau = t \\ L'_{t|t-1}L'_{t+1|t}\ldots L'_{\tau-1|\tau-2} & \tau > t \end{cases}$$

Then, the data decomposition of the smoothed vector of structural shocks is given by

$$\widehat{\eta}_{t|T} = QR'\sum_{\tau=t}^{T}\Omega^{\star}_{t|\tau}Z'F_{\tau|\tau-1}^{-1}\nu_\tau \tag{4}$$

$$= QR'\sum_{j=1}^{n}\sum_{\tau=t}^{T}\Omega^{\star}_{t|\tau}Z'F_{\tau|\tau-1}^{-1}\nu_\tau^j \tag{5}$$

where ν_τ^j is an $n-$ dimensional vector with all rows equal to zero but the j^{th} row, which is set to the actual forecast error.

---

[3]Appendix A.4 provides the double decomposition in terms of the data matrix.

The second step in the inference about the latent vector is its shock decomposition. Under the assumption of stationary dynamics, the autoregressive model for the state vector in equation (1) has an infinite moving average representation

$$\widehat{X}_{t|T} = \Phi^t \widehat{X}_{0|T} + \sum_{\tau=0}^{t-1} \Phi^\tau R \widehat{\eta}_{t-\tau|T} \tag{6}$$

where $\widehat{X}_{0|T}$ is the estimated vector of initial conditions, given data through time T. Let us define $\widehat{\eta}_{t-\tau|T}^i$ as a vector containing the estimated path for the i^{th} structural shock conditional on the remaining $(p-1)$ structural shocks being zero for all τ. Therefore, the estimated vector of structural shocks is given by $\widehat{\eta}_{t-\tau|T} = \sum_{i=1}^p \widehat{\eta}_{t-\tau|T}^i$ and the estimated realization of the state vector by

$$\widehat{X}_{t|T} = \Phi^t \widehat{X}_{0|T} + \sum_{i=1}^p \sum_{\tau=0}^{t-1} \Phi^\tau R \widehat{\eta}_{t-\tau|T}^i \tag{7}$$

where $\sum_{\tau=0}^{t-1} \Phi^\tau R \widehat{\eta}_{t-\tau|T}^i$ states the contribution of the i^{th} structural shock to the estimated vector of latent variables at time t.

The double decomposition chains the data decomposition, which relates observables to estimates of shocks, to the shock decomposition, which relates shocks to latent variables. Using equation (4), we define the contribution of news about the j^{th} observable to the latent vector through the i^{th} shock as follows

$$\widehat{X}_{t|T}^{i,j} = \Phi^t \widehat{X}_{0|T}^j + \sum_{\tau=0}^{t-1} \Phi^\tau R \left\{ QR' \sum_{\tau'=t}^T \Omega_{t|\tau'}^\star Z' F_{\tau'|\tau'-1}^{-1} \nu_{\tau'}^j \right\}_i \tag{8}$$

where $\widehat{X}_{0|T}^j$ is the contribution of news about the j^{th} observable to the vector of initial conditions for the latent variables. Hence, the double decomposition allows us to gauge the influence of data on the path of latent variables, like the data decomposition. However, by taking into account the transmission mechanisms of each type of shock, we can highlight the economic structure underlying the relationship between the data and the latent variables.

The overall contribution of news about the j^{th} observable on the latent vector is given by

$$\widehat{X}_{t|T}^j = \sum_{i=1}^p \widehat{X}_{t|T}^{i,j} = \Phi^t \widehat{X}_{0|T}^j + \sum_{i=1}^p \sum_{\tau=0}^{t-1} \Phi^\tau R \left\{ QR' \sum_{\tau'=t}^T \Omega_{t|\tau'}^\star Z' F_{\tau'|\tau'-1}^{-1} \nu_{\tau'}^j \right\}_i \tag{9}$$

and the estimated vector of latent variables is

$$\widehat{X}_{t|T} = \sum_{i=1}^p \sum_{j=1}^n \widehat{X}_{t|T}^{i,j} = \sum_{j=1}^n \Phi^t \widehat{X}_{0|T}^j + \sum_{i=1}^p \sum_{\tau=0}^{t-1} \Phi^\tau R \left\{ QR' \sum_{\tau'=t}^T \Omega_{t|\tau'}^\star Z' F_{\tau'|\tau'-1}^{-1} \nu_{\tau'}^j \right\}_i \tag{10}$$

3. A Simple Model Demonstration

In this section, we provide a practitioner's view of the logic behind the double decomposition using a small-scale model. By using the double decomposition, we are able to better understand the

economic mechanisms in this model by which some variables, such as inflation, are highly informative about the output gap, while others, such as GDP growth, are less so.

We illustrate how the double decomposition works using an estimated version of the canonical New Keynesian model presented in Ireland (2011). Formally, the behavioral core of the model consists of four equations governing the shadow value of wealth (λ_t), the log-level of detrended output (Y_t), inflation (π_t) and the federal funds rate (R_t):

$$(\bar{z} - \beta\gamma)(\bar{z} - \gamma)\lambda_t = \gamma\bar{z}Y_{t-1} - (\bar{z}^2 + \beta\gamma^2)Y_t + \beta\gamma\bar{z}E_tY_{t+1} + (\bar{z} - \beta\gamma\rho_a)(\bar{z} - \gamma)a_t - \gamma\bar{z}z_t \quad (11)$$

$$\lambda_t = R_t + E_t\left(\lambda_{t+1} - \pi_{t+1}\right) \quad (12)$$

$$(1 + \beta\alpha)\pi_t = \alpha\pi_{t-1} + \beta E_t\pi_{t+1} - \psi\lambda_t + \psi a_t + e_t \quad (13)$$

$$R - R_{t-1} = \rho_\pi\pi_t + \rho_g\left(Y_t - Y_{t-1} + z_t\right) + \epsilon_t^R \quad (14)$$

The model features four exogenous processes driven by structural shocks: the growth rate of total factor productivity (z_t), a shock to household preferences (a_t), a price mark-up shock (e_t), and a monetary policy shock (ϵ_t^R). Ireland (2011) estimates the parameters of the model by maximum likelihood, using as observable variables GDP growth, GDP price inflation, and the 3-month T-bill rate from 1983:Q1 to 2009:Q4. In our analysis, we use the parameter estimates reported in Ireland (2011).

To illustrate the double decomposition, let us suppose we are interested in understanding how the model infers movements in the output gap from the path of observable variables.[4] In line with our previous discussion, our analysis proceeds sequentially in an "onion" structure. First, we show how the main features of the data decomposition for structural shocks follow from some properties of the impulse response functions for observables. The impulse response functions for the output gap then allow us to understand how the model translates news about observable variables into estimates of the output gap. Finally, we describe the double decomposition of the output gap over the estimation sample.

Figure 1 reports impulse responses following one-standard-deviation innovations to the exogenous processes. Certain familiar qualitative differences in the response of the observable variables are visible. First, there are two shocks that drive inflation and output growth in the same direction on impact –the preference and monetary policy shocks–, which are, in turn, distinguishable from each other by the reaction of the T-bill rate. Second, the remaining two shocks – the mark-up and technology shocks– move inflation in opposite directions and can be again distinguished from each other by the comovement of GDP growth with the T-bill rate: while following a mark-up shock, GDP growth falls and the T-bill rate rises, both growth and the T-bill rate rise following a technology shock.

With these qualitative differences, inference about structural shocks given forecast errors in the observable variables is easy to explain intuitively. First, the fact that the systematic component of monetary policy depends only on observable variables means that the monetary policy shock is

---

[4]In this example, the output gap is defined relative to the level of output that would have prevailed in the absence of nominal rigidities.

Figure 1. Impulse Response Functions Following Structural Shocks
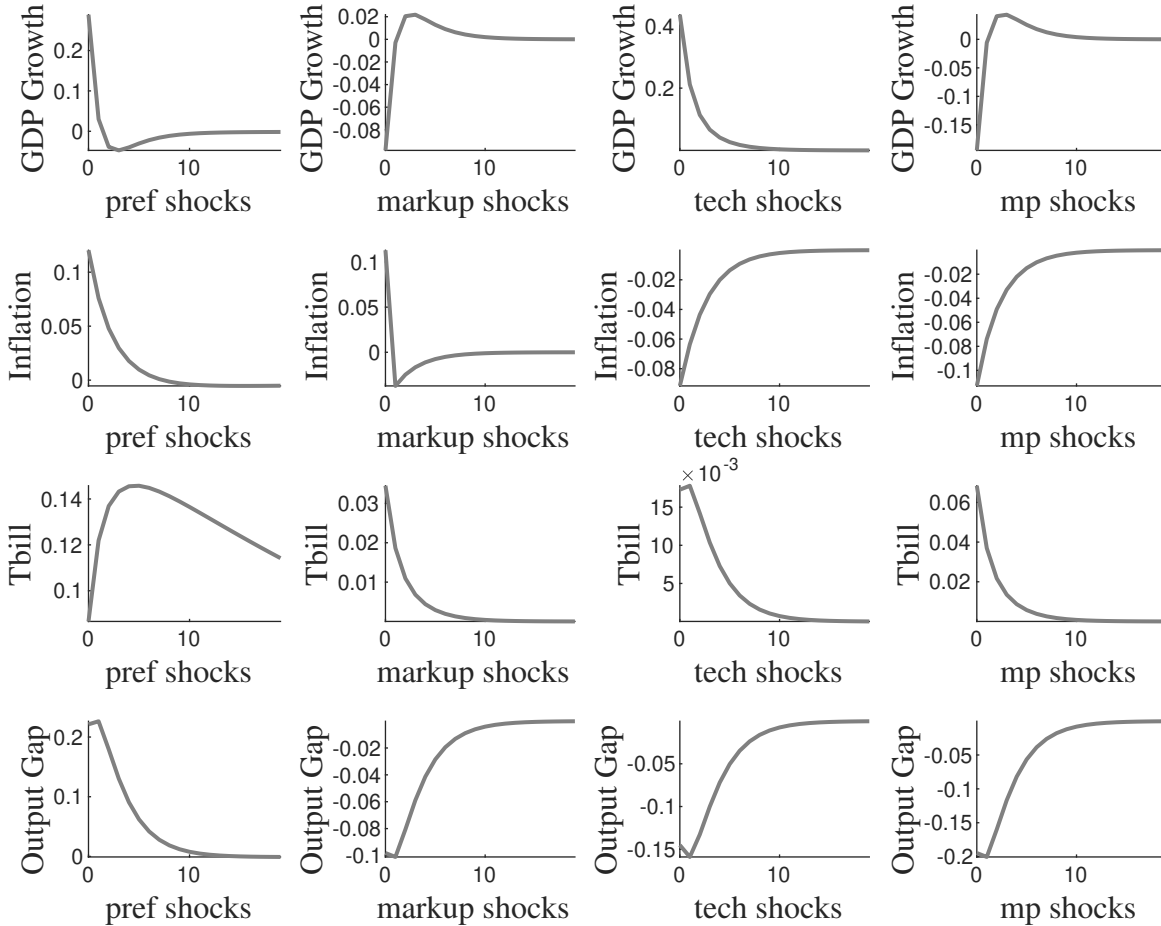
Note: Impulse responses are calculated following one-standard-deviation shocks and are reported in percentage points.

measurable and inference about its path is trivial. Accordingly, we only focus on inference about the other three structural shocks. The model's inference about these shocks is most easily understood by considering the model's response to forecast error vectors with the property that the implied monetary policy shock is zero. This restriction implies that the forecast errors in the T-bill do not have any information beyond what is contained in the forecast errors for GDP growth and inflation. Therefore, we only discuss the information content of these variables. In this context, we consider two cases: *(i)* a case featuring a one-standard-deviation forecast error in GDP growth and a zero error in inflation and *(ii)* a case featuring a zero forecast error in GDP growth and a one-standard-deviation error in inflation. Given the linear-Gaussian structure of the model, inference for all of the other cases can be represented as a linear combination of cases *(i)* and *(ii)*. Figure 2 reports the results of these exercises.

Let us first suppose that the model is presented with a one-standard-deviation forecast error in GDP growth, a zero forecast error in inflation, and a forecast error in the T-bill consistent with the monetary policy rule. From Figure 1, we know that the dominant drivers of GDP growth are the

preference and technology shocks. Therefore, the model accounts for the positive forecast error in GDP growth using one-standard-deviation shocks to preferences and technology, as shown in the left column in Figure 2.[5] Moreover, as shown in the second row of Figure 1, the effect on inflation of this combination of shocks is relatively small. Thus, such a combination also mostly accounts for the zero inflation forecast error, with only a slight movement in the markup shock.

Let us now consider the alternative case: the model is presented with a one-standard-deviation forecast error in inflation, a zero forecast error in GDP growth, and a forecast error in the T-bill consistent with the monetary policy rule. This case is more complicated because, as shown in Figure 1, all one-standard-deviation structural shocks have similar effects on impact for inflation. The right column of Figure 2 shows that the model can account for this case by sizable movements in the three non-monetary-policy shocks. In particular, the model infers a 0.8 standard deviation increase in the preference shock, which, by itself, increases GDP growth. The model offsets the effect of the preference shock on GDP growth with contractionary supply shocks: a positive markup shock and a negative shock to technology, so that the overall effect on GDP growth is zero. These shocks also help explain the positive forecast error in inflation.

Combining the information contained Figure 1 and Figure 2, we obtain Figure 3, which shows how the model's estimate of the output gap responds to the two vectors of forecast errors described above. This two-way decomposition – effects of news about observable variables on structural shock contributions – is essentially what we formalize in the double decomposition. Let us consider the left column in Figure 3, which reports the case of a one-standard-deviation forecast error in GDP growth. In this case, individually, the associated positive preference and technology shocks have sizable effects on the model estimate of the output gap. However, taken jointly, these two shocks are offsetting, leaving a much smaller imprint on the output gap than on GDP growth. By contrast, in the case of a positive forecast error only in inflation, all the associated shocks (positive preference and markup shocks and a negative technology shock) move the output gap in the same direction. Therefore, we conclude that the estimated path for the output gap is quite sensitive to inflation news, but not GDP growth.

We can now use these findings to understand the model's estimates of the output gap in a sample from 1983:Q1 to 2009:Q4. The resulting double decomposition is shown in Figure 4. The upper-left panel displays the data decomposition of the estimated path for the output gap. The remaining panels show the data decomposition of the contribution of each structural shock to the estimated path for the output gap. For example, the middle panel in the first row of Figure 4 shows the estimated path for the output gap if only technology shocks were active and all other shocks were zero. Consistent with our previous discussion, the data decomposition of the output gap reported in the upper-left panel of Figure 4 shows that the estimated path of the output gap is almost entirely driven by inflation news. The shocks associated with news about GDP growth and the T-bill have largely offsetting effects on the output gap.

---

[5]These effects are calculated for an observation 54 quarters from both the beginning and end of the sample. This corresponds to an observation in the middle of the 1983:Q1-2009:Q4 (i.e., arbitrary period in the middle between 1996:Q2 and 1996:Q3) sample used to estimate the model.

Figure 2. EFFECT OF NEWS AT TIME 0 ON ESTIMATED PATHS FOR STRUCTURAL SHOCKS
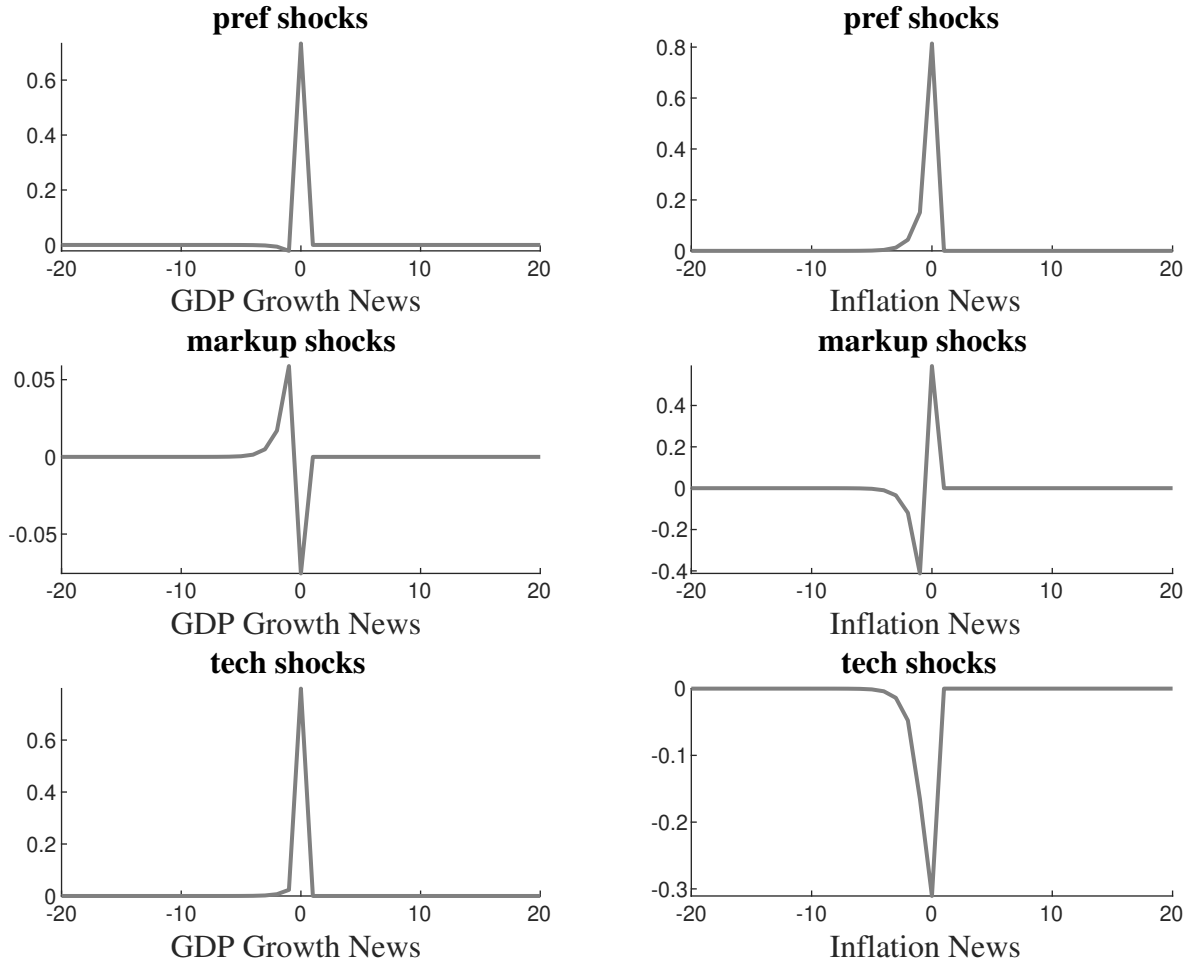
Note: Shock responses are reported in standard-deviation units. Responses represent the effects of news at time 0 in the given variable on the estimated path of the shocks (labelled in bold) 20 quarters before and after that time.

Given the simplicity of the model, this outcome could have been anticipated by noting that, compared to DSGE models estimated on data after the Great Recession, the Phillips curve is relatively sensitive to movements in real activity. Correspondingly, in the Ireland (2011) model, the role of markup shocks is smaller than in those DSGE models. In this context, inflation is naturally highly informative regarding the output gap. However, even in this simple case, running through the formal logic of the double decomposition helps to narrow the key dynamic features of the model underlying this outcome, i.e., that inflation is relatively responsive to both technology and preference shocks.

Figure 3. Effect of News at Time 0 on Estimated Contributions of Structural Shocks to the Output Gap
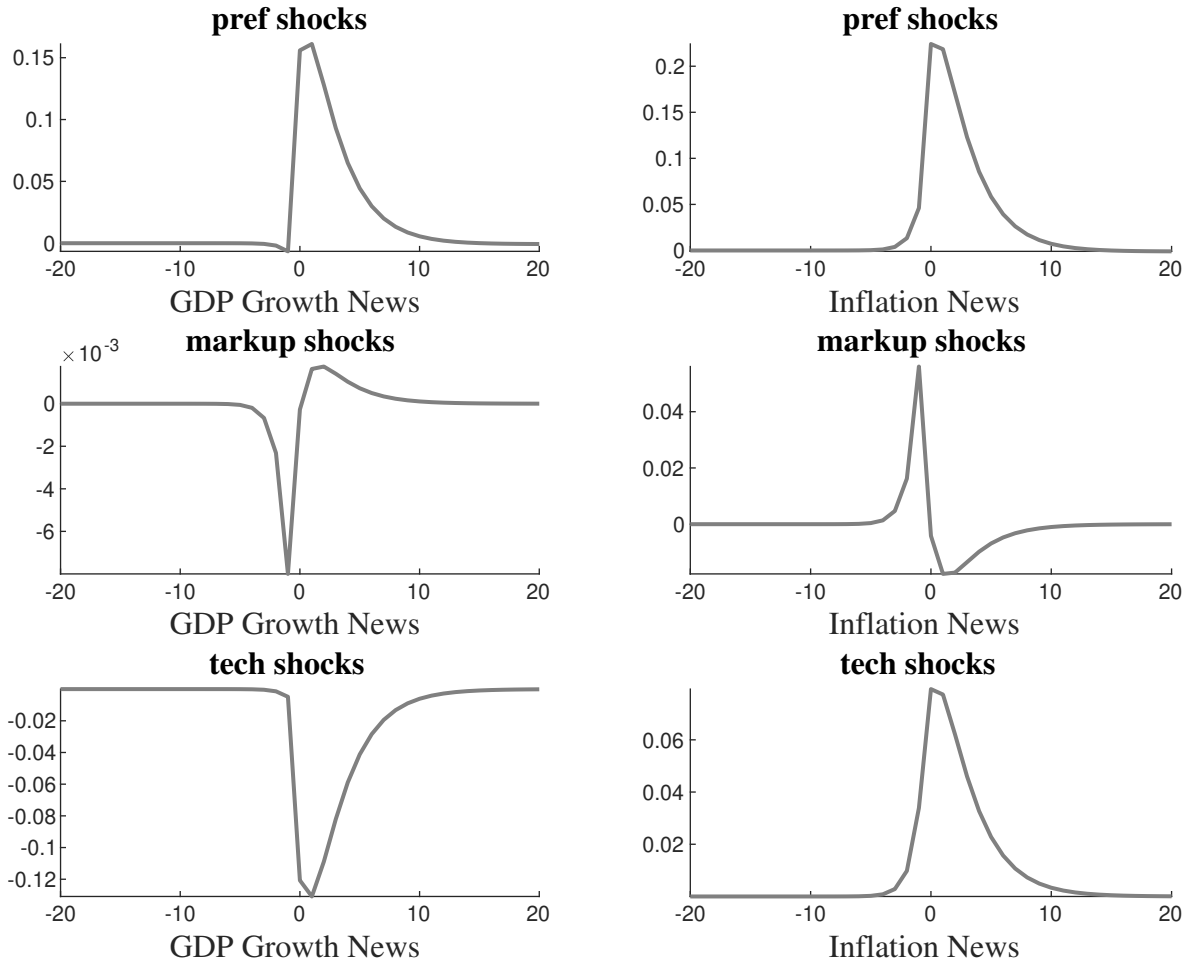
Note: Responses represent the effects in percentage point of news at time 0 in the given variable to the contribution of a specified shock (labelled in bold) to the path of the output gap 20 quarters before and after that time.

4. A Large-Model Demonstration

4.1. The Model

To further demonstrate the utility of the double decomposition, we now turn to a larger model, more like the workhorse models used in academia and central banks, featuring a much larger set of observable variables and more complex transmission dynamics. In particular, we use an estimated version of the model originally developed by Del Negro, Giannoni and Schorfheide (2015) (DGS). Broadly speaking, the DGS model extends the baseline Smets and Wouters (2007) model with financial frictions on the firm side and a time-varying inflation target in the monetary policy rule. Our version of the model differs from the one described by DGS in two respects. First, our Taylor-type interest rate feedback rule responds to the change in real GDP, rather than to the flex-price

Figure 4. Decomposition of Structural Shock Contributions to the Output Gap in the Ireland 2011 Model
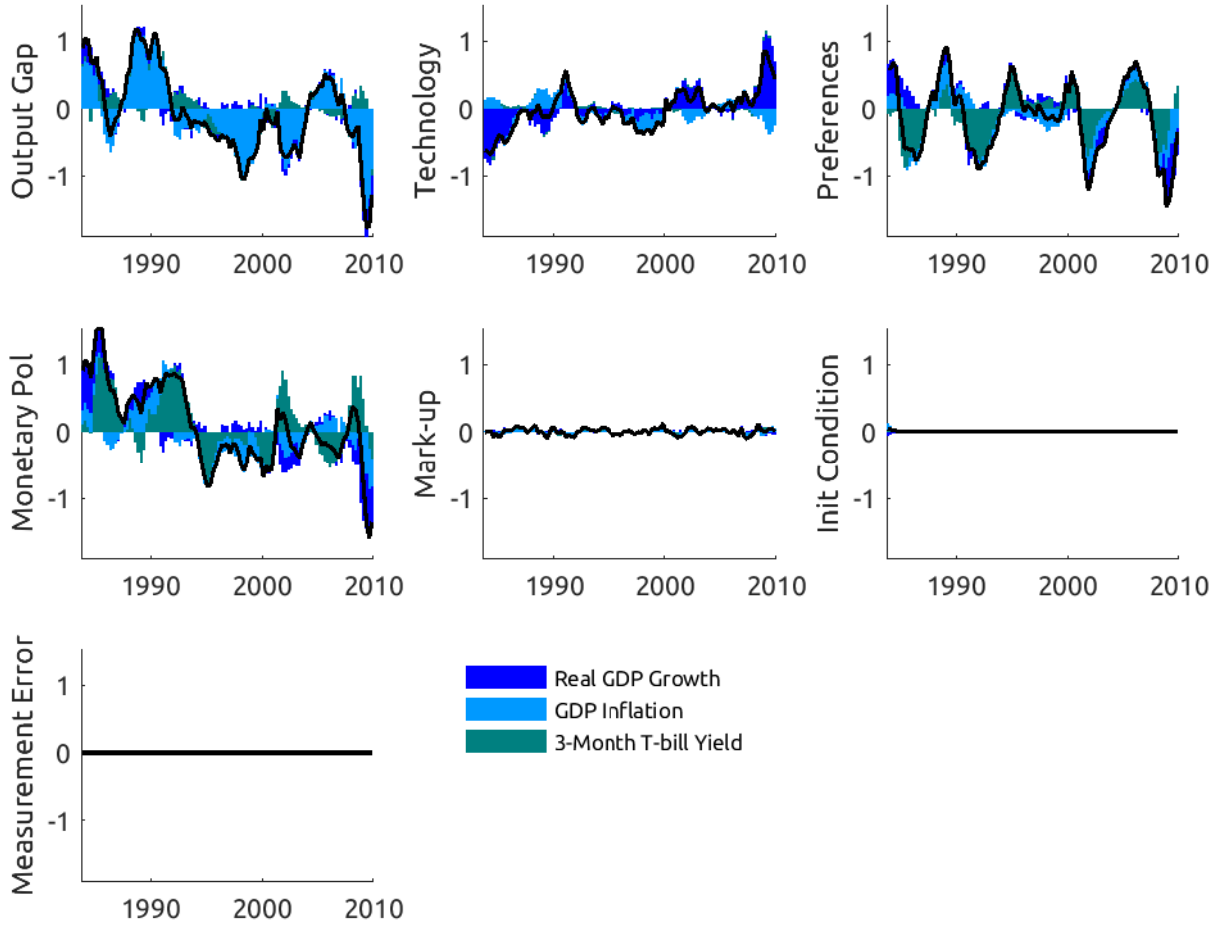
Note: Upper right panel: data decomposition of the output gap in percentage points. Other panels: data decomposition of the contribution of the given shock to the output gap in percentage points. Sample: 1983q1-2009q4.

output gap. This assumption implies that the setting of the federal funds rate is not, by construction, a strong signal about the output gap. The monetary policy rule in our model is given by

$$R_t = \rho^r (R_{t-1}) + (1 - \rho^r) [\phi_\pi (\pi_t - \pi_t^*) + \phi_y (y_t - y_{t-1})] + \epsilon_t^r$$

where the parameter ρ^r reflects the degree of interest rate smoothing, π_t^* is the time-varying inflation target, and ϵ_t^r represents a monetary policy shock. The time-varying inflation target follows

$$\pi_t^* = \rho_{\pi^*} \pi_t^* + \sigma_{\pi^*} \epsilon_{\pi^*,t} \tag{15}$$

where $0 < \rho_{\pi^*} < 1$ and $\epsilon_{\pi^*,t}$ is an *iid* shock. The time-varying inflation target accounts for the low frequency movements in longer-horizon inflation expectations in our estimation sample and potentially any effects of unconventional monetary policy, such as forward guidance, on agents' expectations. The inflation target is informed by matching model-based average expected inflation over the next 10 years with the Survey of Professional Forecasters median 10 year-ahead inflation expectations.

Second, following Barsky, Justiniano and Melosi (2014), we use several indicators of inflation and wages, incorporating a factor structure in the corresponding measurement equations. For price inflation, we consider the log difference of the GDP deflator, core CPI price inflation, and core PCE price inflation. For wage inflation, we use average hourly earnings and the employment cost index.[6]

The model includes eleven shocks: shocks to transient TFP, permanent TFP, investment specific technology, the aggregate risk premium, firm-specific risk, marginal bankruptcy costs, net worth, price markups, wage markups, exogenous spending, the inflation target, and the monetary policy rule. For the estimation, we use data on GDP growth, consumption growth, investment growth, the federal funds rate, the spread between corporate bonds and Treasuries of comparable maturity, net worth, corporate debt, long-run inflation expectations, wage inflation measured by average hourly earnings, wage inflation measured by the employment cost index, GDP deflator price inflation, core CPI price inflation, and core PCE price inflation. The estimation sample is 1987:Q1-2015:Q4.[7]

4.2. Comparison with Standard Decompositions

In this subsection, in Figures 5-6, we present results from the standard shock and data decompositions for the output gap, defined as the log-difference between the level of actual output and the level of output in a counterfactual economy (the "flex-price economy") with flexible prices and wages and without price and wage markup shocks. The output gap was mostly positive in the second half of the 1990s and in the 2000s. In 2009, the output gap decreased rapidly to about -9 percent where it remained for a couple of years before gradually recovering. The shock decomposition suggests that financial shocks, notably the firm-specific risk shock, drive most of the business cycle variation together with permanent productivity shocks. Since the Great Recession, investment specific shocks and monetary policy shocks play an important and very persistent role preventing the output gap from declining further.

---

[6]To illustrate the procedure, let's consider the case of wages. The vector of L_w quarterly wage growth data series is given by:

$$w_t^{p'} = \left[w_{1,t}^p, ..., w_{L_w,t}^p\right]' \tag{16}$$

The measurement equation linking the model concept of wage inflation with the j^{th} element of this vector, $w_{j,t}^p$, is as follows

$$w_{j,t}^p = (\bar{w} + c_j) + \Lambda_j w_{m,t}^p + u_{j,t} \tag{17}$$

where $\bar{w}$ is the estimated steady state level and c_j captures the differences in estimated means among different series. The idiosyncratic disturbances, $u_{j,t}$, which follow an $AR(1)$ structure, and loading factors, Λ_j, are series specific.

[7]Our results may be specific to the sample period we use and the fact that we abstract from the issues that arise due to the ZLB. We leave the estimation that would explicitly take into account the ZLB for future research.

Figure 5. Shock Decomposition of the Output Gap
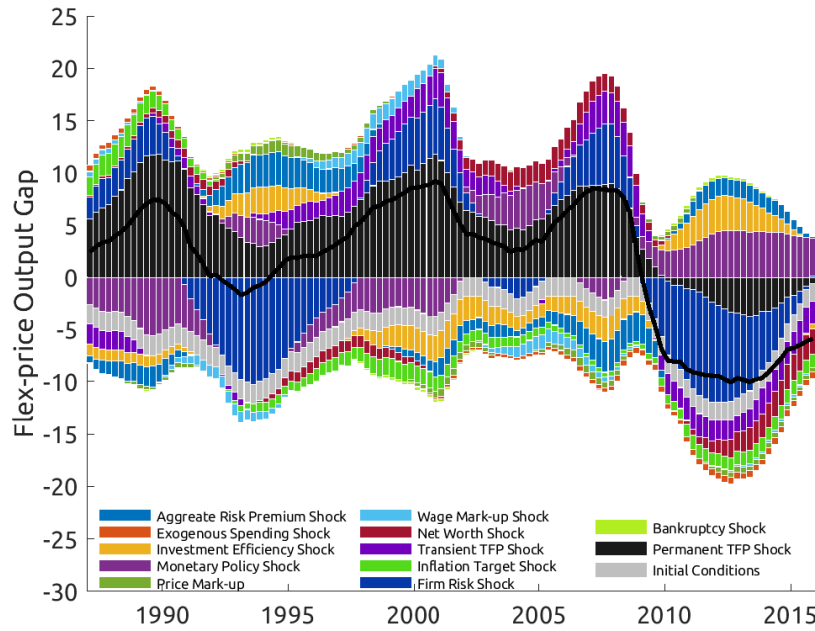
Note: Bars represent the contribution of shocks to the output gap in percentage points. Sample: 1987q1-2015q4.

Figure 6. Data Decomposition of the Output Gap
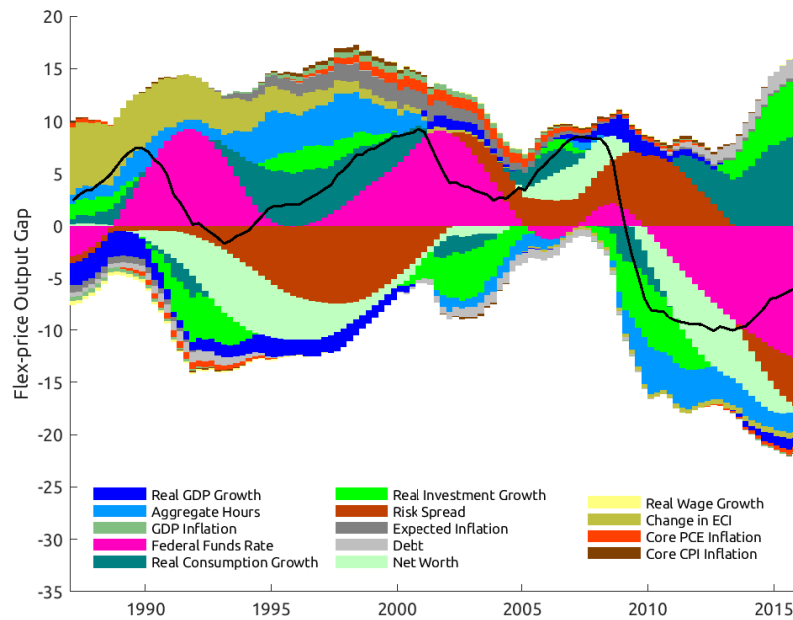
Note: Bars represent the contribution of forecast errors in the given observable to the estimated path of the output gap in percentage points. Sample: 1987q1-2015q4.

Looking at the data decomposition in Figure 6, we see that a number of observables influence the estimate of the output gap. In particular, the federal funds rate, real consumption growth, and the corporate bond spread seem to have the largest influence on the estimate of the output gap.

Some features of these results resonate with other analyses of the output gap over our sample, but others are more counter-intuitive. For example, the model estimates that the output gap was far above zero in the second half of the 1990s, peaking at 9.6 percent at the end of 2000. Noting the low level of unemployment, coupled with strong investment growth and subdued inflation over that period, contemporary accounts from these years emphasized the role of productivity shocks and this model agrees with that assessment, finding that almost all of the rise in the output gap in the late 1990s is accounted for by shocks to permanent productivity. However, as shown by the data decomposition of the output gap, news about investment growth, hours and inflation played a very small role in the model's estimates of the output gap at this time; rather, as in most other times, contributions from consumption growth, the federal funds rate and the corporate bond spread dominate. In the next section, we show that a consideration of the double decomposition helps to explain why these observables are so influential and why they so strongly signal shocks to permanent productivity.

4.3. Decomposing Estimates of the Output Gap

In this section, we employ the double decomposition to examine the model's inferences about the output gap. To provide an intuitive explanation of the double decomposition, we proceed as we did in Section 3: we first examine the relationship between forecast errors in key observable variables and the model's estimates of structural shocks (described in Table 1) and then use impulse response functions to relate these effects on shock estimates to the model's estimate of the path of the output gap (described in Tables 2 and 3).

We conclude that the output gap path is chiefly informed by two classes of observable variables: spending components and financial variables. In the first class are spending components – consumption growth most importantly – which provide relatively clean signals about permanent productivity and firm-level risk shocks, the main contributors to the variance of the output gap. In the second class, the model's financial observable variables also signal large movements in productivity and firm risk. These data are influential in the longer term, after the effects of transient offsetting shocks also associated with these observables have faded. To streamline our discussion, we therefore focus on this set of observable variables.

Table 1 displays the effect of a one-standard-deviation forecast error in these selected observables on the model's estimate of contemporaneously occurring shocks.[8] The first two columns of Table 1 describe the effects of one-standard-deviation forecast errors in consumption and investment growth, respectively, holding all other forecast errors at zero. Because GDP, consumption and investment growth are all observed, a forecast error in consumption or investment alone, holding the other two fixed, requires an exactly offsetting movement in exogenous spending, equal to a little more than one standard deviation (-1.16) in the case of a consumption forecast error and -0.74 in the case of an investment forecast error. The crowding-in of consumption and investment occasioned

---

[8]For expositional ease, in this first step, we abstract from finite-sample effects and present results for inference at an observation very far from the start and end of the sample (200 quarters). As in the previous section, in this table, the magnitudes of the shocks are measured in standard-deviation units.

by the decline in exogenous spending helps the model account for some of the higher-than-expected private spending, but the model must also rely on other structural shocks. For consumption forecast errors, the offsetting shock is chiefly a positive shock to permanent productivity, which boosts both consumption and investment growth, and a much smaller *increase* in firm-level risk, which offsets the impetus to investment that the first two shocks would otherwise deliver.

Table 1: EFFECTS OF NEWS ON SHOCK ESTIMATES

Shock	Positive news on					
	Cons	Inv	Hours	FFR	Spread	Net Worth
Prod (temp)	-0.11	0.02	-1.45	-0.06	-0.05	-0.02
Prod (perm)	1.01	-0.12	0.02	0.50	0.46	0.21
Firm-level Risk	0.14	-1.11	-0.01	-0.39	-0.35	-0.48
Risk premium	-0.15	0.82	0.00	0.28	1.50	0.47
Exog. spending	-1.16	-0.74	0.12	-0.01	-0.01	-0.00
Inv. efficiency	-0.03	0.52	0.02	-0.21	-0.27	-1.12
Mon Pol	-0.02	-0.02	-0.04	0.94	-0.02	-0.01
Price markup	0.06	0.02	-0.11	0.20	0.05	0.06
Wage markup	-0.23	-0.03	-0.11	-0.08	-0.14	0.00
Net worth	-0.01	-0.32	-0.02	-0.15	-0.33	-0.18
Infl Target	0.11	0.19	-0.23	0.04	0.09	0.25
Mar. bankruptcy	0.08	0.24	0.02	0.15	0.56	0.17

Note: Table entries are the effect of a one-standard deviation forecast error in the given observables on the model's estimate of the contemporaneous shock. Shocks are measured in standard-deviation units.

The negative shock to exogenous demand and the positive shock to firm-risk both imply lower output gaps, while the positive productivity shock would imply a higher output gap. However, the exogenous spending shock has a relatively weak impact on the output gap, while only a small firm-level risk shock is necessary to account for the zero forecast error in investment. Accordingly, we find that positive news on consumption is associated with higher output gaps, mostly because such news signals the arrival of a permanent boost to productivity. This outcome is quantified in the first column of Table 2, which shows that, all else equal, a one-standard-deviation forecast error in consumption raises the estimated contribution of permanent productivity shocks to the output gap by 0.44 percentage points, with small effects on contributions from other shocks.

Turning now to the case of a forecast error in investment growth (the second column of Table 1), we find that, again, the implied decline in exogenous spending is not sufficient to explain the forecast error. In fact, the forecast error is accounted for by a negative shock to firm-level risk, accompanied by a positive shock to the marginal efficiency of investment (further boosting investment and, to a lesser extent, consumption). By itself, such a configuration of shocks would entail a sizeable negative forecast error in spreads and positive forecast errors in consumption, which, by assumption, did not occur. The model thus infers that the aggregate risk premium must have experienced a large positive shock, while permanent productivity suffered a modest downward shock; both of these shocks boost

Table 2: EFFECT OF NEWS ON SHOCK CONTRIBUTIONS TO THE CONTEMPORANEOUS OUTPUT GAP

| | Positive news on | | | | | |
Contribution to gap from	Cons	Inv	Hours	FFR	Spread	Net Worth
Prod (temp)	0.03	-0.01	0.34	0.02	0.01	0.01
Prod (perm)	0.44	-0.09	-0.02	0.23	0.18	0.08
Firm-level Risk	-0.04	0.29	-0.01	0.13	0.08	0.13
Risk premium	0.02	-0.08	0.01	-0.04	-0.22	-0.05
Exog. spending	-0.04	-0.02	0.00	0.00	0.01	0.00
Inv. efficiency	-0.01	-0.00	0.00	-0.02	-0.04	-0.06
Mon Pol	0.00	0.01	-0.00	-0.25	0.01	-0.00
Price markup	-0.00	0.00	-0.00	-0.00	0.00	0.00
Wage markup	0.01	-0.00	0.04	-0.01	0.01	-0.00
Net worth	-0.00	0.00	0.00	-0.02	-0.05	-0.01
Infl Target	0.01	0.01	-0.01	-0.00	-0.00	0.02
Mar. bankruptcy	0.00	-0.01	-0.00	-0.00	-0.02	-0.01
Total	0.43	0.11	0.36	0.04	-0.02	0.11

Note: Table entries are the effect of a one-standard deviation forecast error in the given observables on the model's estimate of the contemporaneous output gap, measured in percentage points.

spreads and lower consumption by much more than investment. Putting the pieces together, from the second column of Table 2, we see that investment forecast errors are associated with an increase in the contribution of the firm-level risk shocks to the output gap of 0.29 percentage points, about half of which is offset negative contributions from the aggregate risk premium and permanent technology shocks. Overall, then, news about investment growth emerges as less informative about the contemporaneous output gap than news about consumption, essentially because consumption growth errors are cleanly associated with a permanent productivity shock, which has strong effects on the output gap, while investment growth is associated with a mix of offsetting shocks.

A data decomposition would find that the second most influential observable for the model's estimate of the contemporaneous output gap is aggregate hours: as shown by the third column of Table 2, a one-standard-deviation forecast error in hours would raise the estimate of the output gap at that date by 0.36 percentage points, similar to the effect of a one-standard-deviation forecast error in consumption growth. In the case of aggregate hours, however, the forecast error is accounted chiefly by a negative shock to temporary productivity. Because the productivity shock is temporary, its effects on spending growth are largely muted by a rise in labor inputs – a desirable feature when accounting for a positive forecast error in hours, but zero forecast errors in spending observable variables. The association of hours with temporary productivity shocks accounts for nearly all of the effect of hours forecasts errors on estimates of the output gap.

Finally, we now consider the effect of news on financial observable variables, several of which can be influential on the model's estimates of the output gap. Overall, these observable variables are associated with sizeable movements in permanent productivity and firm-level risk. However, they

are not, taken separately, very informative about the contemporaneous output gap, as the absence of consistent forecast errors in spending implies the presence of strong offsetting shocks. As we discuss below, these offsetting shocks are relatively transient and financial observables emerge as informative about the output gap several quarters in the future.

We start with news about the federal funds rate, shown in the third column of Table 1. In this model, systematic monetary policy depends only on observables, the unobserved inflation target and monetary policy shocks. In fact, the forecast error in the federal funds rate alone is accounted for mostly by the monetary policy shock. To account for the zero forecast errors in spending, the model offsets the contractionary implications of this monetary policy shock with several expansionary shocks, including positive shocks to permanent productivity and negative shocks to firm-level risk and the aggregate risk premium. From Table 2, we see that the positive contributions to the output gap from the permanent productivity shock and firm-level risk dominate the offsetting negative contributions. On balance, a positive forecast error in the federal funds rate is associated with a modest increase in the model's estimate of the contemporaneous output gap.

The next financial indicator that we will consider, the corporate bond spread, is heavily driven by shocks to the aggregate risk premium that appear directly as wedges in the equation characterizing the laws of motion of this variable. As with forecast errors in the federal funds rate, however, the positive shock to the aggregate premium must be offset by other expansionary shocks, in order to be consistent with zero forecast errors in spending observables and, as we have also seen before, the model infers positive shocks to permanent productivity and negative shocks to firm-level risk for this reason. Again consulting Table 2, we see that these contributions are roughly offsetting and the total effect of forecast errors in the spread is quite small.

The last observable that we will consider is net worth. A forecast error in net worth is largely explained by a one-standard-deviation decline in the marginal efficiency of investment, which raises the price of installed capital because investment flows are a less effective source of new capital. However, as in the previous cases, this shock alone has counterfactual implications for consumption and investment. Therefore, they must be accompanied by offsetting expansionary shocks, namely positive permanent productivity shocks and negative firm-level risk shocks. These expansionary shocks, as we have previously discussed, have much larger effects on the output gap than the other main contributors in this case and so a one-standard-deviation forecast error in net worth is associated with an 0.11 percentage point increase in the estimated output gap.

So far, we have focused on the link between forecast error and estimates of contemporaneous structural shocks and output gaps. However, as we showed earlier, the contribution of the data to latent variables can be thought of as contributions to estimated shocks at various horizons, weighted by the impulse responses of the latent variable of interest at those horizons. In the current model, the persistence of output gap effects differs quite a bit across different shocks. In particular, as shown in the Appendix, permanent shocks to productivity and risk premium shocks induce movements in the output gap that are substantially more persistent than other major sources of variation; indeed, the output gap impulse response functions for those shocks are hump-shaped, peaking more than two years after the shock at roughly double their value at impact. Hence, news about variables that

Table 3: Effect of News on Shock Contributions to the Output Gap 8 quarters in the Future

| | Positive news on | | | | | |
Contribution to gap from	Cons	Inv	Hours	FFR	Spread	Net Worth
Prod (temp)	0.03	-0.01	0.40	0.02	0.01	0.01
Prod (perm)	1.29	-0.21	-0.00	0.65	0.58	0.26
Firm-level Risk	-0.18	1.38	0.02	0.55	0.47	0.62
Risk premium	0.03	-0.15	0.00	-0.06	-0.30	-0.09
Exog. spending	-0.04	-0.03	0.00	0.00	0.00	0.00
Inv. efficiency	-0.01	0.09	0.01	-0.06	-0.09	-0.27
Mon Pol	0.00	0.00	0.01	-0.34	0.00	-0.00
Price markup	-0.01	-0.00	0.02	-0.01	-0.00	-0.00
Wage markup	0.06	0.01	0.05	0.01	0.03	-0.00
Net worth	-0.00	0.14	-0.00	0.00	0.02	0.04
Infl Target	0.03	0.04	-0.05	0.01	0.01	0.06
Mar. bankruptcy	-0.00	-0.01	-0.00	-0.01	-0.02	-0.01
Total	1.18	1.28	0.47	0.77	0.72	0.62

Note: Table entries are the effect of a one-standard deviation forecast error in the given observables on the model's estimate of the output gap eight quarters later, measured in percentage points.

are linked to large movements in permanent productivity or risk premium shocks tend to be highly influential about future output gaps, even this influence is somewhat obscured in the near-term by signals about offsetting shocks. For this reason, we now examine the dynamic contributions of news about observables to estimates of the future path of the output gap.

To get a sense of how the persistence of the shocks affects the contribution of the different observables, in Table 3, we examine the effect of the important observables on the contributions of the various shocks to the output gap 8 quarters after impact.[9] From this table, it is immediately apparent that, at this horizon, this set of observables is still highly influential regarding the output gap, sometimes even more so than on impact as in the case of the federal funds rate, and that almost all of that influence is mediated through the persistent effects of permanent productivity and firm-level risk shocks.

To summarize our findings thus far, then, two classes of observables tend to contribute most strongly to the estimated path of the output gap in this model. The first class of observables are spending components, i.e., consumption and investment growth, whose forecast errors have sparse and large loadings on permanent productivity and firm-level risk shocks, because these shocks account for a large share of those observables' variance and have much larger effects on them than on the others variables in the information set. For the second class of observables (financial observables) by contrast, productivity and firm-risk shocks are not necessarily the dominant source

---

[9]A more comprehensive set of coefficients from the data decomposition for these observables is available in the Appendix. We note that, because of the effect of shocks on the gap is persistent, estimates of the time-T output gap can depend on forecast errors *after* time T. However, as shown by Figure A.4, the influence of future forecast errors is much smaller than the past and contemporaneous forecast errors, justifying our focus on the latter.

of variance. Rather, for these variables, the dominant source of variance (monetary policy shocks, for example, in the case of the funds rate) would, by itself, cause large movements in spending. Thus, a non-zero forecast error in such a variable alone signals *offsetting* movements in the dominant drivers of consumption and investment, muting the impact on the contemporaneous output gap.

Figure 7. DYNAMIC CONTRIBUTIONS TO TWO-SIDED ESTIMATES FOR THE FLEX-PRICE OUTPUT GAP OVER THE ESTIMATION SAMPLE
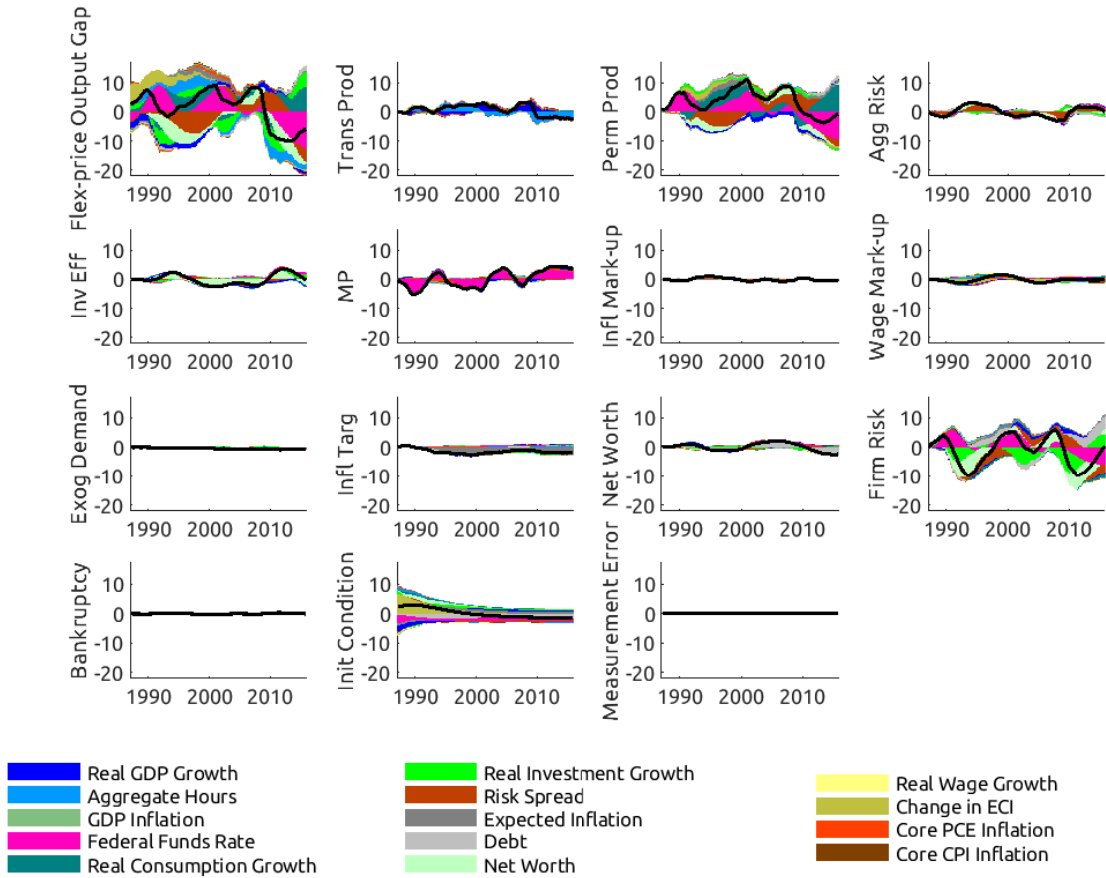
Note: Upper left panel: data decomposition of the output gap in percentage points. Other panels: data decomposition of the contribution of the given shock to the output gap in percentage points. Sample: 1987q1-2015q4.

Having laboriously worked our way through this underlying logic, we now confirm that these patterns are, indeed, reflected in the relation between observables and shock contributions in the estimation sample. By doing so, we also show that the double decomposition can shed light on why observables are particularly important for the data decomposition. Our results are presented in Figure 7. Consistent with our theoretical priors, we see that forecast errors in investment growth and net worth make significant contributions to the path of the output gap through their role in signalling the level of firm-specific risk, while consumption growth is also influential, but chiefly as an indicator of permanent productivity. The federal funds rate emerges as a fourth important

observable and, in this case, conveys information about the level of both permanent productivity and firm-level risk.

Tracing the impact of observables on the output gap through their effects on estimates of the shocks clarifies why these variables emerge as influential, and why the observables emphasized by the contemporary accounts are not so influential. For example, in this model, forecast errors in hours are mostly driven by temporary productivity shocks, whose effects on the output gap fade rapidly. By contrast, forecast errors in consumption, driven by permanent shocks to productivity, have similar effects to hours errors on estimates of the contemporaneous output gap but build substantially over subsequent years. Regarding investment growth, our earlier analysis shows that investment growth can be highly informative about the output gap, but only at longer horizons, because the shocks underlying investment growth errors have offsetting effects on the contemporaneous output gap. Moreover, earlier forecast errors, especially in the federal funds rate, were associated by the model with permanent technology shocks, explaining a significant fraction of the higher investment over these years, further reducing the contribution of investment growth to the output gap.

### 4.4. Interpreting the News: An Illustrative Example

In this section, we illustrate the use of the decompositions derived above to construct simple, easily interpretable narratives connecting news about the observables in the model to revisions to the forecast. For our example, we focus on the model's reaction to news between the end of the fourth quarter of 2013 and the middle of the second quarter of 2014. We focus on this period because the configuration of incoming news (e.g., very disappointing readings on labor productivity, but unexpectedly strong consumption growth) is not naturally explained by a single shock, highlighting the value of the decomposition that we presented above.

Specifically, in this experiment, we took a snapshot of the real-time data for two points in time: the first one at 9am EST on December 31, 2013 and the second one at 9am EST on May 15, 2014. That guarantees us that the first dataset includes real-time estimates for the third quarter of 2013, including the advance estimate of GDP and its components. The second dataset includes real-time estimates for the first quarter of 2014. We do not attempt to condition on partial information on the forecast quarter from monthly data sources, such as the Employment Situation report.

Figure 8 displays the real-time data corresponding to both vintages, as well as the associated forecasts and estimates of the output gap. As is apparent from the Figure, the main feature of the news between these two datasets is a drastic positive forecast error for 2014:Q1 consumption growth, while the forecast errors for 2014:Q1 investment growth and output growth are negative. In the data, the most obvious sources of weakness in 2014:Q1 output growth are in inventory investment and net exports. Correspondingly, exogenous spending in this quarter was weak. The model was surprised, to a lesser extent, by inflation, inflation expectations, and federal funds rates. Finally, aggregate hours were only slightly below expectation, implying a large negative error in average labor productivity.

Perhaps surprisingly, the model reacts to this constellation of news by keeping its forecast for both output growth and hours at a very similar level, but revising its estimate of the output gap from

Figure 8. Comparison of forecasts for selected variables, December 2013 vs. May 2014
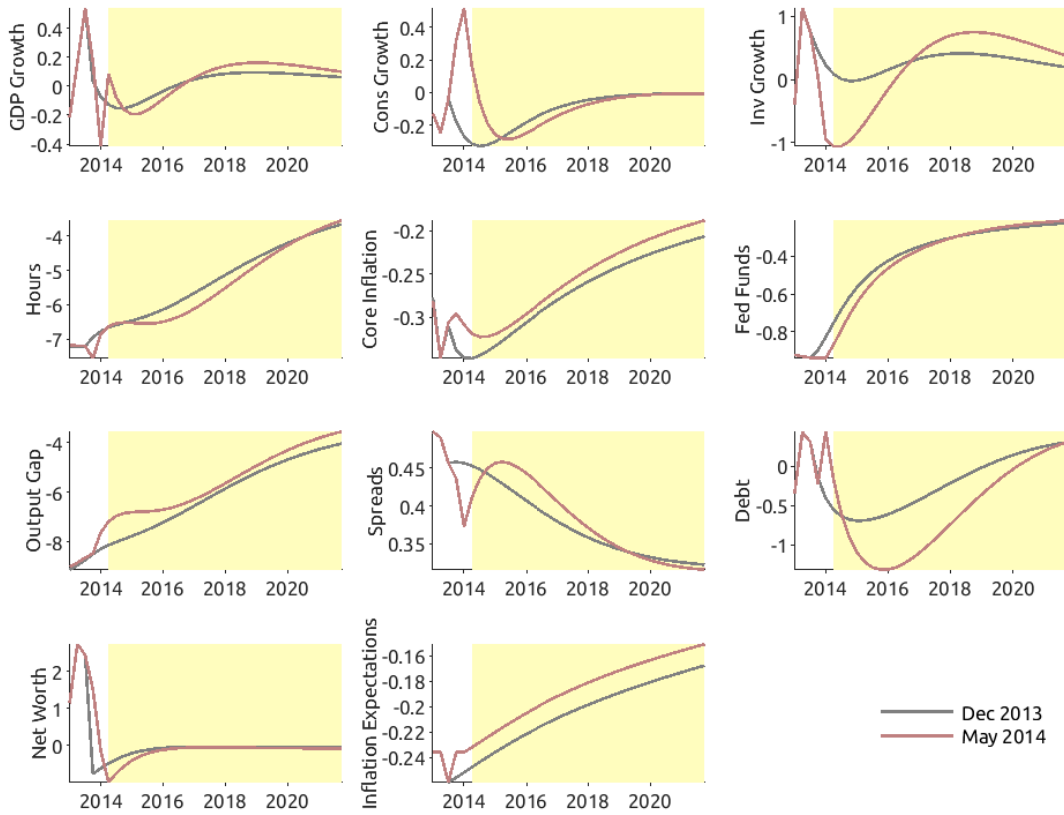

Note: The gray line reports the real-time data and forecast as of December 2013. The red line reports the real-time and forecast as of May 2014. The shaded area covers the common period of forecast for both dates.

about -8.5 percent to -7 percent in 2014:Q1. Furthermore, inflation and inflation expectations have been revised upward, while the weakness in investments is projected to persist for a few periods. We can use the general properties of the double decomposition to develop a crisp narrative explaining this reaction.

Figure 9 displays the double decomposition for the revision to the output growth forecast. In line with our earlier discussion, we can see that the negative forecast error in output growth signalled to the model a very transient drop in exogenous demand, by virtue of the aggregate resource constraint (the panel labelled "Exog Demand"). At the same time, the positive error in consumption growth was associated with positive contributions to output growth from shocks to productivity (shown by the positive green bars in the panel labelled "Productivity"), boosting the forecast for output growth. Similarly, monetary policy shocks, almost exclusively informed by surprises in the federal funds rate, were expansionary. However, these two effects were balanced by negative financial shocks

21

(informed by the negative error in investment growth), so that the forecast for output growth was little changed at the end of the day.
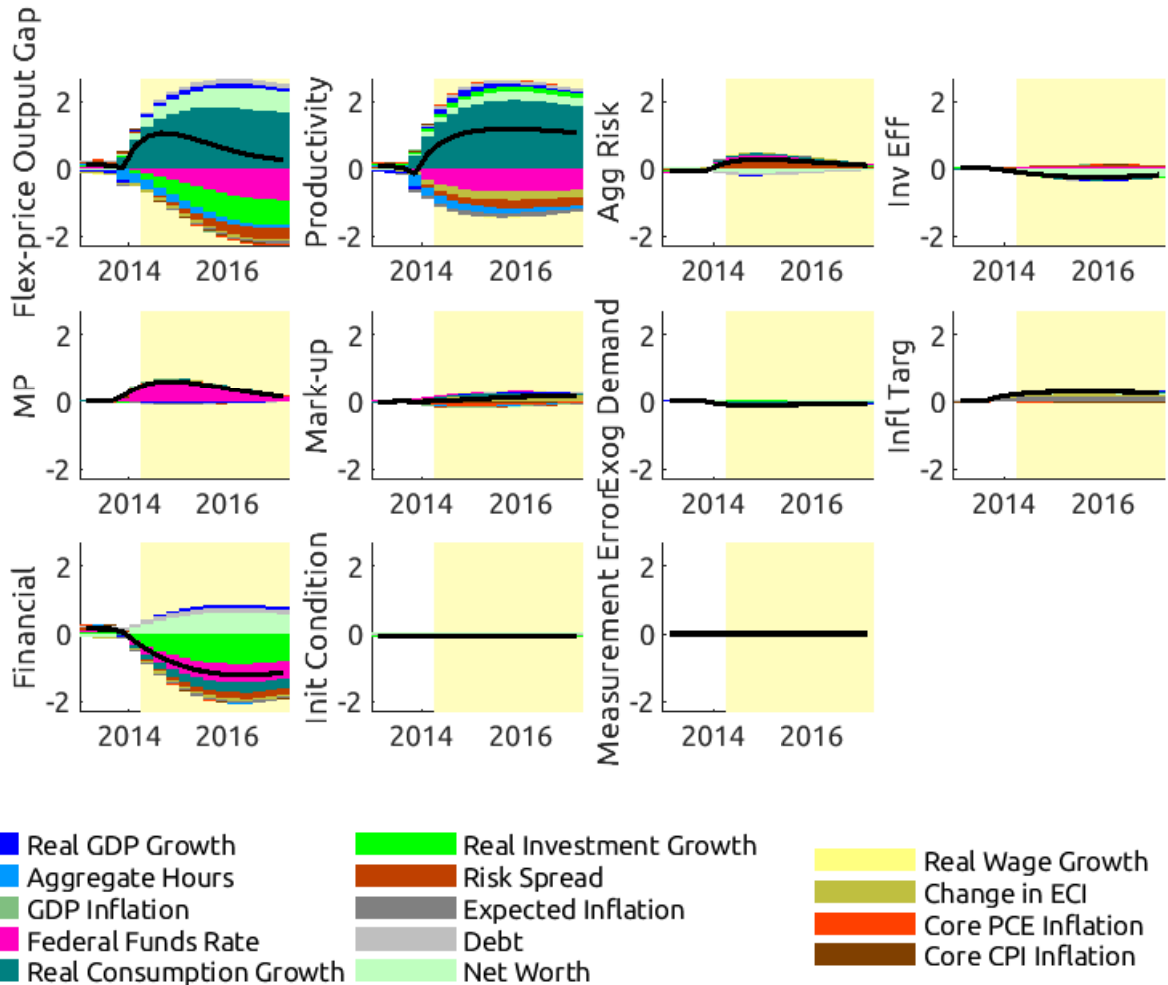
Figure 9. DECOMPOSITION OF THE REVISION IN THE FORECAST FOR OUTPUT GROWTH BETWEEN DECEMBER 2013 AND MAY 2014
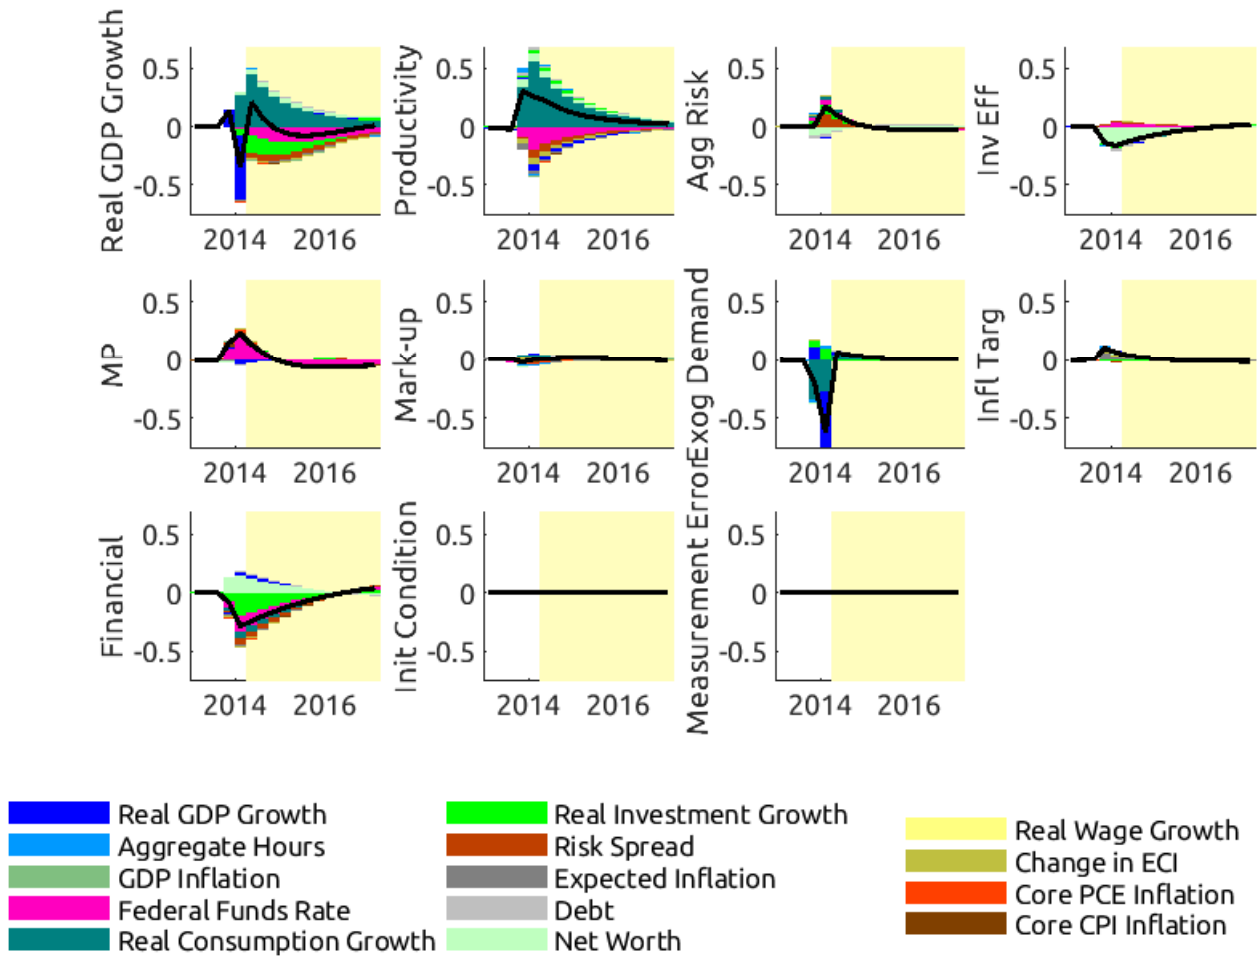
Note: Upper left panel: data decomposition of the revision to the forecast. Other panels: data decomposition of the contribution of the given shock to the revision. Financial shocks include shocks to firm-specific risk, net worth and bankruptcy costs. Productivity shocks include both permanent and transient shocks. Mark-up shocks include both price and wage mark-ups.

Turning now to the revision in the output gap, as shown in Figure 10, the situation is somewhat surprising, but consistent with the analysis of the output growth. As mentioned before, news about consumption growth is mostly informative about permanent productivity in this model. Given that we observed a sizeable positive surprise in consumption growth, the model inferred a shock to permanent productivity, and not a demand shock as it would be the prior for most practitioners. The output gap narrowed in the response, as growth in potential output increased. The estimated shock to permanent productivity due to the consumption forecast error is attenuated by the negative surprise in the federal funds rate. The revision to the output gap is further attenuated by the

Figure 10. Decomposition of the revision in the forecast for the output gap between December 2013 and May 2014

Note: Upper left panel: data decomposition of the revision to the forecast. Other panels: data decomposition of the contribution of the given shock to the revision. Financial shocks include shocks to firm-specific risk, net worth and bankruptcy costs. Productivity shocks include both permanent and transient shocks. Mark-up shocks include both price and wage mark-ups.

negative contributions from financial shocks inferred as a result of the surprisingly meagre growth in real investment.

To summarize, the positive news on consumption growth, accompanied by little change in output growth, suggested to the model that aggregate supply conditions in 2014:Q1 were better than expected and the output gap has been revised up. This example demonstrates the utility of the double decomposition for interpreting the reaction of the model to incoming data.

5. Conclusion

In this paper, we advocate chaining the decomposition of shocks into contributions from forecast errors to the shock decomposition of the latent vector in order to better understand model inference about latent variables. This double decomposition allows us to gauge the influence of data on the path of latent variables, like the data decomposition. However, by taking into account the transmission mechanisms of each type of shock, we can highlight the economic structure underlying the relationship between the data and the latent variables.

We demonstrate the usefulness of this approach by detailing the role of observable variables in estimating the output gap within the DSGE model à la Del Negro, Giannoni and Schorfheide (2015). Tracing the logic of the double decomposition, we find that real economic factors, such as consumption and investment, and certain financial observables, such as the corporate bond spread, are highly informative about the output gap largely because they provide strong signals about the arrival of permanent productivity and firm-level risk shocks – the main drivers of variation in the output gap. Our double decomposition results also highlight the mechanisms that explain some puzzling features of the data decomposition, such as the weak role of aggregate hours and the strong role of the federal funds rate in informing the model's estimates of the path of the output gap. As these features are key to the model's reaction to data, the double decomposition enables practitioners to better understand and communicate model estimates of latent variables.

References

Andrle, Michal. 2013. "Understanding DSGE Filters in Forecasting and Policy Analysis." IMF Working Paper 13/98.

Barsky, Robert, Alejandro Justiniano, and Leonardo Melosi. 2014. "The Natural Rate of Interest and Its Usefulness for Monetary Policy." *American Economic Review*, 104(5): 37–43.

Del Negro, Marco, Marc P. Giannoni, and Frank Schorfheide. 2015. "Inflation in the Great Recession and New Keynesian Models." *American Economic Journal: Macroeconomics*, 7(1): 168–96.

Durbin, J., and S. J. Koopman. 2002. *Time Series Analysis by State Space Models.* Oxford University Press:Oxford Statistical Science Series.

Ireland, Peter N. 2011. "A New Keynesian Perspective on the Great Recession." *Journal of Money, Credit and Banking*, 43(1): 31–54.

Sander, N. 2013. "Fresh Perspectives on Unobservable Variables: Data Decomposition of the Kalman Smoother." Reserve Bank of New Zealand Analytical Notes AN 2013/09.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review*, 97(3): 586–606.

Appendix

A.1.   The Kalman Filter and Smoother

Let us consider the following state space system:

$$
\begin{aligned}
X_t &= \Phi X_{t-1} + R\eta_t \\
Y_t &= ZX_t + \varepsilon_t
\end{aligned}
$$

Let us assume that the structural innovations $\eta_t \sim_{iid} \mathcal{N}(0,Q)$, the measurement errors $\varepsilon_t \sim_{iid} \mathcal{N}(0,H)$, and the initial state $X_0 \sim \mathcal{N}\left(\widehat{X}_{0|0}, P_{0|0}\right)$. In stationary models, usually $\widehat{X}_{0|0}$ and $P_{0|0}$ correspond to the invariant distribution associated with the law of motion for the state vector, X_t.

The Kalman filter and smoother recursions are based on the following lemma.

Lemma A.1.1. *Let $(x', y')'$ be jointly normal with*

$$
\mu = \left(\begin{array}{c} \mu_x \\ \mu_y \end{array} \right) \quad and \quad \Sigma = \left(\begin{array}{cc} \Sigma_{xx} & \Sigma_{xy} \\ \Sigma_{yx} & \Sigma_{yy} \end{array} \right)
$$

Then the conditional probability density function $p(x|y)$ is a multivariate normal with

$$
\mu_{x|y} = \mu_x + \Sigma_{xy}\Sigma_{yy}^{-1}(y - \mu_y)
$$
$$
\Sigma_{xx|y} = \Sigma_{xx} - \Sigma_{xy}\Sigma_{yy}^{-1}\Sigma_{yx}
$$

The Kalman filter recursions can be summarized as follows:

$$
\begin{aligned}
\widehat{X}_{t|t-1} &= \Phi \widehat{X}_{t-1|t-1} \\
P_{t|t-1} &= \Phi P_{t-1|t-1}\Phi' + RQR' \\
\nu_t &= Y_t - \widehat{Y}_{t|t-1} = Y_t - Z\widehat{X}_{t|t-1} \\
\widehat{X}_{t|t} &= \widehat{X}_{t|t-1} + K_{t|t-1}\nu_t \\
P_{t|t} &= P_{t|t-1} - M_{t|t-1}F_{t|t-1}^{-1}M'_{t|t-1} \\
M_{t|t-1} &= P_{t|t-1}Z' \\
K_{t|t-1} &= M_{t|t-1}F_{t|t-1}^{-1} \\
F_{t|t-1} &= ZM_{t|t-1} + H \\
L_{t|t-1} &= \Phi \left[I - P_{t|t-1}Z'F_{t|t-1}^{-1}Z \right]
\end{aligned}
$$

Note that

$$
\Phi P_{t|t} = L_t P_{t|t-1}
$$

We can rewrite the one-step ahead forecast errors or innovations, ν_t, as

$$
\nu_t = Y_t - \widehat{Y}_{t|t-1} = Z\left(X_t - \widehat{X}_{t|t-1}\right) + \varepsilon_t = Z\xi_{t|t-1} + \varepsilon_t
$$

where $\xi_{t|t-1}$ is the state estimation error:

$$
\begin{aligned}
\xi_{t|t-1} &= X_t - \widehat{X}_{t|t-1} \\
&= \Phi X_{t-1} + R\eta_t - \Phi \widehat{X}_{t-1|t-1} \\
&= \Phi \left[I - K_{t-1|t-2}Z \right] \xi_{t-1|t-2} + R\eta_t - \Phi K_{t-1|t-2}\varepsilon_{t-1} \\
&= L_{t-1|t-2}\xi_{t-1|t-2} + R\eta_t - \Phi K_{t-1|t-2}\varepsilon_{t-1}
\end{aligned}
$$

The equations for the Kalman smoother recursions are given by

$$
\begin{aligned}
\widehat{X}_{t|T} &= \widehat{X}_{t|t} + P_{t|t}\Phi' P_{t+1|t}^{-1}\left(\widehat{X}_{t+1|T} - \Phi \widehat{X}_{t|t} \right) \\
P_{t|T} &= P_{t|t} + P_{t|t}\Phi' P_{t+1|t}^{-1}\left[P_{t+t|T} - P_{t+1|t} \right] P_{t+1|t}^{-1}\Phi P_{t|t}
\end{aligned}
$$

which require inverting the matrix $P_{t+1|t}^{-1}$ in each recursion. Following Durbin and Koopman (2002), we reformulate the Kalman smoother recursions for the state vector so that we do not need to compute the $(T-1)$ inverse matrices. Let us define

$$
r_t = P_{t+1|t}^{-1}\left(\widehat{X}_{t+1|T} - \widehat{X}_{t+1|t} \right)
$$

Let us rewrite the smoothed state vector as

$$
\widehat{X}_{t|T} = \widehat{X}_{t|t-1} + K_{t|t-1}\nu_t + P_{t|t}\Phi' P_{t+1|t}^{-1}\left(\widehat{X}_{t+1|T} - \widehat{X}_{t+1|t} \right)
$$

$$
P_{t|t-1}^{-1}\left(\widehat{X}_{t|T} - \widehat{X}_{t|t-1} \right) = r_{t-1} = Z' F_{t|t-1}^{-1}\nu_t + L'_{t|t-1}r_t
$$

Using forward recursions, we can write the vector r_{t-1} as a weighted sum of future one-step ahead forecast errors

$$
r_{t-1} = Z' F_{t|t-1}^{-1}\nu_t + L'_{t|t-1}Z' F_{t+1|t}^{-1}\nu_{t+1} + \ldots + L'_t L'_{t+1|t} \ldots L'_{T-1|T-2}Z' F_{T|T-1}^{-1}\nu_T
$$

with $r_T = 0$. Recall the definition $r_{t-1} = P_{t|t-1}^{-1}\left(\widehat{X}_{t|T} - \widehat{X}_{t|t} \right)$, then

$$
\widehat{X}_{t|T} = \widehat{X}_{t|t} + P_{t|t-1}r_{t-1} \tag{A.1}
$$

and, therefore, the smoothed estimate of the state vector can be written as a weighted sum of one-step ahead forecast errors. Moreover, using this formulation of the smoothed estimate of the state vector, the econometrician does not need to compute additional inverse matrices since $F_{j|j+1}^{-1}$ are already computed in the filtering recursions.

Using Lemma A.1.1, we have that the smoothed estimate of the vector of measurement errors is given by

$$
\widehat{\varepsilon}_{t|T} = \sum_{\tau}^{T} \mathbb{E}\left[\varepsilon_t \nu'_\tau \right] F_{\tau|\tau-1}^{-1}\nu_\tau
$$

where

$$\mathbb{E}\left[\varepsilon_t \nu_\tau'\right] = \left\{ \begin{array}{cc} 0 & \tau < t \\ H & \tau = t \\ \mathbb{E}\left[\varepsilon_t \xi_{\tau|\tau-1}\right] Z' & \tau > t \end{array} \right\}$$

with

$$\begin{aligned}
\mathbb{E}\left[\varepsilon_t \xi_{t+1|t}\right] &= -H K_{t|t-1}' \Phi' \\
\mathbb{E}\left[\varepsilon_t \xi_{t+2|t+1}\right] &= -H K_{t|t-1}' \Phi' L_{t+1|t}' \\
& \quad \ldots \\
\mathbb{E}\left[\varepsilon_t \xi_{T|T-1}\right] &= -H K_{t|t-1}' \Phi' L_{t+1|t}' L_{t+2|t+1}' \ldots L_{T-1|T-2}'
\end{aligned}$$

Then, the smoothed estimate for the vector of measurement errors is given by

$$\begin{aligned}
\widehat{\varepsilon}_{t|T} &= H \left[F_{t|t-1}^{-1} \nu_t - K_{t|t-1}' \Phi' Z' F_{t+1|t}^{-1} \nu_{t+1} - K_{t|t-1}' \Phi' L_{t+1|t}' Z' F_{t+2|t+1}^{-1} \nu_{t+2} \right. \\
& \quad \left. - \ldots - K_{t|t-1}' \Phi' L_{t+1|T}' \ldots L_{T-1|T-2}' Z' F_{T|T-1}^{-1} \nu_T \right] \\
&= H \left[F_{t|t-1}^{-1} \nu_t - K_{t|t-1}' \Phi' r_t \right]
\end{aligned}$$

Similarly, the smoothed estimate for the vector of structural shocks is

$$\widehat{\eta}_{t|T} = \sum_{\tau=1}^{T} \mathbb{E}\left[\eta_t \nu_\tau\right] F_{\tau|\tau-1}^{-1} \nu_\tau$$

where

$$\mathbb{E}\left[\eta_t \nu_\tau'\right] = \left\{ \begin{array}{cc} 0 & \tau < t \\ \mathbb{E}\left[\eta_t \xi_{\tau|\tau-1}\right] Z' & \tau \geq t \end{array} \right\}$$

with

$$\begin{aligned}
\mathbb{E}\left[\eta_t \xi_{t|t-1}\right] &= Q R' \\
\mathbb{E}\left[\eta_t \xi_{t+1|t}\right] &= Q R' L_{t|t-1}' \\
& \quad \ldots \\
\mathbb{E}\left[\eta_t \xi_{T|T-1}\right] &= Q R' L_{t|t-1}' L_{t+1|t}' \ldots L_{T|T-1}'
\end{aligned}$$

The smoothed estimate for the vector of structural shocks can also be rewritten as a weighted sum of future one-step ahead forecast errors or innovations, ν_t,:

$$
\begin{aligned}
\widehat{\eta}_{t|T} &= QR' \left[Z'F_{t|t-1}^{-1}\nu_t + L'_{t|t-1}Z'F_{t+1|t}^{-1}\nu_{t+1} \right. \\
&\quad \left. + \ldots + L'_{t|t-1}L'_{t+1|t}\ldots L'_{T|T-1}Z'F_{T|T-1}^{-1}\nu_T \right] \\
&= QR'r_{t-1}
\end{aligned}
\tag{A.2}
$$

A.2. Historical Shock Decomposition

The historical shock decomposition represents the estimated latent variables in terms of contributions of the estimated structural shocks. Let $\widehat{e}_{t|T}$ be the estimated path for the i^{th} structural shock obtained from the Kalman smoother and let $\widehat{\eta}_{t-\tau|T}^i$ be an n-dimensional vector of zeros with $\widehat{e}_{t|T}$ in the i^{th} row. Under the assumption of independently and identically distributed structural shocks, the autoregressive model for the state vector in equation (1) has an infinite moving average representation so that the contribution of the i^{th} structural shock to the vector of latent variables at time t is given by

$$
\widehat{X}_{t|T}^i = \Phi^t X_0 + \sum_{\tau=0}^{t-1} \Phi^\tau R\widehat{\eta}_{t-\tau|T}^i
\tag{A.3}
$$

Therefore, $\widehat{X}_{t|T}^i$ is the estimated model-implied path for the state vector given the estimated path for the i^{th} structural shock conditional on the remaining $(n-1)$ structural shocks being zero for all τ. The estimated actual realization of the state vector, $\widehat{X}_{t|T}$, is equal to the sum across the structural shocks contributions

$$
\widehat{X}_{t|T} = \Phi^t X_0 + \sum_{i=1}^{n}\sum_{\tau=0}^{t-1} \Phi^\tau R\widehat{\eta}_{t-\tau|T}^j
\tag{A.4}
$$

A.3. Historical Data Decomposition

Given the linearity of the econometric model, each observable variable has an independent effect on the Kalman smoother's estimate of a latent variable. The historical data decomposition traces these independent contributions of each observable variable to the estimated smoothed realization of each latent variable. We first compute the data decomposition for the filtered vector of state variables, $\widehat{X}_{t|t}$, which is given by.

$$
\widehat{X}_{t|t} = \Omega_{t,0}\widehat{X}_{0|t} + \sum_{\tau=1}^{t} \Omega_{t,\tau}K_{\tau|\tau-1}Y_\tau
\tag{A.5}
$$

In order to obtain the sequence of weights $\{\Omega_{t,\tau}\}_{\tau=1}^{T}$, let us consider the updated state vector from the Kalman filter recursion

$$
\begin{aligned}
\widehat{X}_{t|t} &= \widehat{X}_{t|t-1} + K_{t|t-1}\nu_t \\
&= K_{t|t-1}Y_t + \left[I - K_{t|t-1}\right]\Phi\widehat{X}_{t-1|t-1}
\end{aligned}
$$

Let us define

$$
N_{t|t-1} = \left[I - K_{t|t-1}\right]\Phi
$$

Note that $N_t = \Phi^{-1}L_{t|t-1}\Phi$. Then,

$$
\begin{aligned}
\widehat{X}_{t|t} &= K_{t|t-1}Y_t + N_{t|t-1}\widehat{X}_{t-1|t-1} \\
&= K_{t|t-1}Y_t + N_{t|t-1}K_{t-1|t-2}Y_{t-1} + N_{t|t-1}N_{t-1|t-2}K_{t-2|t-3}Y_{t-2} + \ldots + \\
&\quad N_{t|t-1}N_{t-1|t-2}\ldots N_{2|1}K_{1|0}Y_1 + N_{t|t-1}N_{t-1|t-2}\ldots N_{2|1}N_{1|0}X_{0|0}
\end{aligned}
$$

Thus, the observation weights for the filtered state vector are given by

$$
\Omega_{t,\tau} = \begin{cases} I & \tau = t \\ N_{t|t-1}N_{t-1|t-2}\ldots N_{\tau+1|\tau} & \tau < t \end{cases}
$$

and the weight for the initial conditions is

$$
\Omega_{t,0} = N_{t|t-1}N_{t-1|t-2}\ldots N_{1|0}
$$

In order to compute the data decomposition for the smoothed vector of state variables, $\widehat{X}_{t|T}$, let us write the smoother recursion as

$$
\widehat{X}_{t|T} = \widehat{X}_{t|t} + P_{t|t}\Phi'P_{t+1|t}^{-1}\left(\widehat{X}_{t+1|T} - \Phi\widehat{X}_{t|t}\right) = \widehat{X}_{t|t} + P_{t|t-1}\Phi'r_{t-1}
$$

We have already computed the observation weights for the filtered vector of state variables, $\widehat{X}_{t|t}$. Thus, we only need to compute the data decomposition of r_{t-1}, which is a weighted sum of future forecast errors

$$
\begin{aligned}
r_{t-1} &= Z'F_{t|t-1}^{-1}\nu_t + L'_{t|t-1}Z'F_{t+1|t}^{-1}\nu_{t+1} + \ldots + L'_{t|t-1}L'_{t+1|t}\ldots L'_{T-1|T-2}Z'F_{T|T-1}^{-1}\nu_T \\
&= Z'F_{t|t-1}^{-1}Y_t + L'_{t|t-1}Z'F_{t+1|t}^{-1}Y_{t+1} + \ldots + L'_{t|t-1}L'_{t+1|t}\ldots L'_{T-1|T-2}Z'F_{T|T-1}^{-1}Y_T \\
&\quad - Z'F_{t|t-1}^{-1}Z\Phi\widehat{X}_{t-1|t-1} - L'_{t|t-1}Z'F_{t+1|t}^{-1}Z\Phi\widehat{X}_{t|t} - \ldots - L'_{t|t-1}\ldots L'_{T-1|T-2}Z'F_{T|T-1}^{-1}Z\Phi\widehat{X}_{T-1|T-1}
\end{aligned}
$$

Let the observation weights for the forward recursions in r_{t-1} be

$$
\Omega_{t,\tau}^{\star} = \begin{cases} I & \tau = t \\ L'_{t|t-1}L'_{t+1|t}\ldots L'_{\tau-1|\tau-2} & \tau > t \end{cases}
$$

Then,

$$
\begin{aligned}
r_{t-1} &= \sum_{\tau=t}^{T} \Omega_{t,\tau}^{\star} Z' F_{\tau|\tau-1}^{-1} Y_\tau - \sum_{\tau=t-1}^{T-1} Z' \Omega_{t-1,\tau}^{\star} F_{\tau+1|\tau} Z\Phi \widehat{X}_{\tau|\tau} \\
&= \sum_{\tau=t}^{T} \Omega_{t,\tau}^{\star} Z' F_{\tau|\tau-1}^{-1} Y_\tau - \sum_{\tau=t-1}^{T-1} Z' \Omega_{t-1,\tau}^{\star} F_{\tau+1|\tau} Z\Phi \sum_{j=1}^{\tau} \Omega_{\tau,j} K_{j,j-1} Y_j
\end{aligned}
$$

The data decomposition of the smoothed state vector is

$$
\begin{aligned}
\widehat{X}_{t|T} &= \Omega_{t,0} \widehat{X}_{0|T} + \sum_{\tau=1}^{t} \Omega_{t,\tau} K_{\tau|\tau-1} Y_\tau \\
&\quad + P_{t|t-1} \Phi' \left[\sum_{\tau=t}^{T} \Omega_{t,\tau}^{\star} Z' F_{\tau|\tau-1}^{-1} Y_\tau - \sum_{\tau=t-1}^{T-1} Z' \Omega_{t-1,\tau}^{\star} F_{\tau+1|\tau} Z\Phi \sum_{j=1}^{\tau} \Omega_{\tau,j} K_{j|j-1} Y_j \right] \quad \text{(A.6)}
\end{aligned}
$$

which states that the time-varying observable weights depend on the relative position of τ with respect to t.

The contribution of the j^{th} observable variable to the estimated latent variable vector at time t, $\widehat{X}_{t|T}^{j}$, is given by

$$
\begin{aligned}
\widehat{X}_{t|T}^{j} &= \Omega_{t,0} \widehat{X}_{0|T}^{j} + \sum_{\tau=1}^{t} \Omega_{t,\tau} K_{\tau|\tau-1} Y_\tau^{j} \\
&\quad + P_{t|t-1} \left[\sum_{\tau=t}^{T} \Omega_{t,\tau}^{\star} Z' F_{\tau|\tau-1}^{-1} Y_\tau^{j} - \sum_{\tau=t-1}^{T-1} Z' \Omega_{t-1,\tau}^{\star} F_{\tau+1|\tau} Z\Phi \sum_{l=1}^{\tau} \Omega_{\tau,j} K_{l|l-1} Y_l^{j} \right] \quad \text{(A.7)}
\end{aligned}
$$

where Y_τ^{j} is an $n-$dimensional vector with all rows equal to the unconditional mean of the corresponding observable variable but the j^{th} row, which is set to the actual observed value.

The estimated actual realization of the smoothed state vector, $\widehat{X}_{t|T}$, is equal to the sum across the observable data series contributions.

$$
\begin{aligned}
\widehat{X}_{t|T} &= \sum_{j=1}^{n} \widehat{X}_{t|T}^{j} = \omega_t^{\star} \Omega_{t,0} \sum_{j=1}^{n} \widehat{X}_{0|T}^{j} + \sum_{j=1}^{n} \sum_{\tau=1}^{t} \Omega_{t,\tau} K_{\tau|\tau-1} Y_\tau^{j} \\
&\quad + P_{t|t-1} \Phi' \sum_{j=1}^{n} \left[\sum_{\tau=t}^{T} \Omega_{t,\tau}^{\star} Z' F_{\tau|\tau-1}^{-1} Y_\tau^{j} - \sum_{\tau=t-1}^{T-1} Z' \Omega_{t-1,\tau}^{\star} F_{\tau+1|\tau} Z\Phi \sum_{l=1}^{\tau} \Omega_{\tau,j} K_{l|l-1} Y_l^{j} \right] \text{(A.8)}
\end{aligned}
$$

Similarly, the data decomposition of the smoothed vector of structural shocks is given by

$$
\widehat{\eta}_{t|T} = QR' \left[\sum_{\tau=t}^{T} \Omega_{t,\tau}^{\star} Z' F_{\tau|\tau-1}^{-1} Y_\tau - \sum_{\tau=t-1}^{T-1} Z' \Omega_{t-1,\tau}^{\star} F_{\tau+1|\tau} Z\Phi \sum_{j=1}^{\tau} \Omega_{\tau,j} K_{j|j-1} Y_j \right]
$$

A.4. Historical Double Decomposition

The historical double decomposition links the historical data decomposition of the structural shocks and the historical shock decomposition of the latent vector.

The data decomposition of the smoothed estimate for the vector of structural shocks is given by

$$\widehat{\eta}_{t|T} = QR' \left[\sum_{\tau=t}^{T} \Omega^{\star}_{t,\tau} Z' F^{-1}_{\tau|\tau-1} Y_\tau - \sum_{\tau=t-1}^{T-1} Z' \Omega^{\star}_{t-1,\tau} F_{\tau+1|\tau} Z \Phi \sum_{l=1}^{\tau} \Omega_{\tau,l} K_{l|l-1} Y_l \right] \qquad \text{(A.9)}$$

Substituting this expression for the vector of smoothed structural shocks in the $\infty-$MA representation of the smoothed vector of latent variables in equation (A.4), we obtain

$$\widehat{X}_{t|T} = \Phi^t \widehat{X}_{0|T} + \sum_{j=1}^{n} \sum_{\tau=0}^{t-1} \Phi^\tau RQR' \left[\sum_{k=\tau-1}^{T} \Omega^{\star}_{\tau-1,k} Z' F^{-1}_{k|k-1} Y^j_k - \sum_{k=t-\tau-1}^{T-1} Z' \Omega^{\star}_{t-\tau-1,k} F_{k+1|k} Z \Phi \sum_{l=1}^{k} \Omega_{k,l} K_{l|l-1} Y^j_l \right]$$
$$\text{(A.10)}$$

A.5. Tables and Figures

Table A.1: EFFECTS OF NEWS ON SHOCK ESTIMATES

Shock	Positive news on							
	Real GDP	Infl	core PCE	core CPI	Exp Infl	Debt	Wages	ECI
Prod (temp)	1.17	0.00	0.00	0.00	0.02	-0.01	-0.01	-0.02
Prod (perm)	0.09	-0.01	-0.02	-0.02	-0.16	0.00	0.02	0.07
Firm-level Risk	0.14	0.01	0.01	0.01	0.08	-0.05	-0.00	-0.02
Risk premium	-0.10	-0.00	-0.01	-0.01	-0.05	0.04	0.01	0.02
Exog. spending	1.76	0.00	0.00	0.00	-0.00	-0.00	-0.00	-0.01
Inv. efficiency	0.02	0.01	0.01	0.01	0.04	-0.11	-0.00	-0.00
Mon Pol	-0.23	-0.06	-0.16	-0.16	0.02	0.02	0.01	0.03
Price markup	0.02	0.14	0.36	0.37	-0.11	-0.03	-0.04	-0.14
Wage markup	0.02	0.02	0.06	0.06	0.28	-0.01	0.09	0.36
Net worth	0.05	0.00	0.01	0.01	0.05	-0.02	0.00	0.00
Infl Target	-0.08	-0.02	-0.04	-0.04	1.02	0.03	-0.02	-0.09
Mar. bankruptcy	-0.04	-0.00	-0.01	-0.01	-0.04	0.01	-0.00	-0.01

Note: Table entries are the effect of a one-standard deviation forecast error in the given observables on the model's estimate of the contemporaneous shock. Shocks are measured in standard-deviation units.

Table A.2: Effects of News on Estimates of the Contribution of Shocks to the Output Gap

| Contribution to gap from	Positive news on							
	Real GDP	Infl	core PCE	core CPI	Exp Infl	Debt	Wages	ECI
Prod (temp)	-0.28	0.01	0.02	0.02	-0.02	-0.00	0.00	0.00
Prod (perm)	0.07	0.00	0.01	0.01	-0.09	0.05	0.02	0.07
Firm-level Risk	-0.05	0.00	0.00	0.00	-0.02	0.13	0.00	0.02
Risk premium	-0.00	-0.00	-0.00	-0.00	0.01	-0.07	-0.00	-0.01
Exog. spending	0.05	-0.00	-0.00	-0.00	0.00	-0.00	-0.00	-0.00
Inv. efficiency	0.01	0.00	0.00	0.00	0.00	0.00	-0.00	-0.00
Mon Pol	0.07	0.03	0.07	0.07	-0.02	-0.01	0.00	0.00
Price markup	0.00	-0.01	-0.03	-0.03	0.01	-0.00	0.00	0.01
Wage markup	-0.03	-0.01	-0.04	-0.04	0.01	0.01	-0.01	-0.04
Net worth	0.01	-0.00	-0.00	-0.00	0.00	-0.06	0.00	0.00
Infl Target	-0.01	-0.01	-0.02	-0.02	0.10	0.01	-0.01	-0.02
Mar. bankruptcy	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Total	-0.15	0.01	0.01	0.01	-0.01	0.03	0.01	0.03

Note: Table entries are the effect of a one-standard deviation forecast error in the given observables on the model's estimate of the contemporaneous output gap, measured in percentage points.

Table A.3: Effect of News on Shock Contributions to the Output Gap 8 quarters in the Future

| Contribution to gap from	Positive news on							
	Real GDP	Infl	core PCE	core CPI	Exp Infl	Debt	Wages	ECI
Prod (temp)	-0.32	0.01	0.01	0.01	-0.01	-0.00	0.00	0.01
Prod (perm)	0.15	-0.00	-0.01	-0.01	-0.23	0.08	0.04	0.14
Firm-level Risk	-0.20	-0.00	-0.01	-0.01	-0.10	0.33	0.01	0.04
Risk premium	0.02	-0.00	-0.00	-0.00	0.01	-0.04	-0.00	-0.01
Exog. spending	0.07	-0.00	-0.00	-0.00	0.00	-0.00	-0.00	-0.00
Inv. efficiency	0.02	0.00	0.00	0.01	0.01	-0.02	-0.00	-0.00
Mon Pol	0.08	0.03	0.07	0.07	-0.02	-0.00	-0.00	-0.00
Price markup	-0.01	-0.02	-0.04	-0.04	0.02	0.00	0.00	0.02
Wage markup	-0.03	-0.02	-0.04	-0.05	-0.03	0.02	-0.03	-0.13
Net worth	-0.00	-0.00	-0.01	-0.01	-0.01	-0.23	0.00	0.00
Infl Target	-0.03	-0.01	-0.02	-0.02	0.27	0.01	-0.01	-0.04
Mar. bankruptcy	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Total	-0.25	-0.02	-0.05	-0.05	-0.10	0.14	0.01	0.03

Note: Table entries are the effect of a one-standard deviation forecast error in the given observables on the model's estimate of the output gap eight quarters later, measured in percentage points.

Figure A.1. Impulse responses functions for major output gap drivers
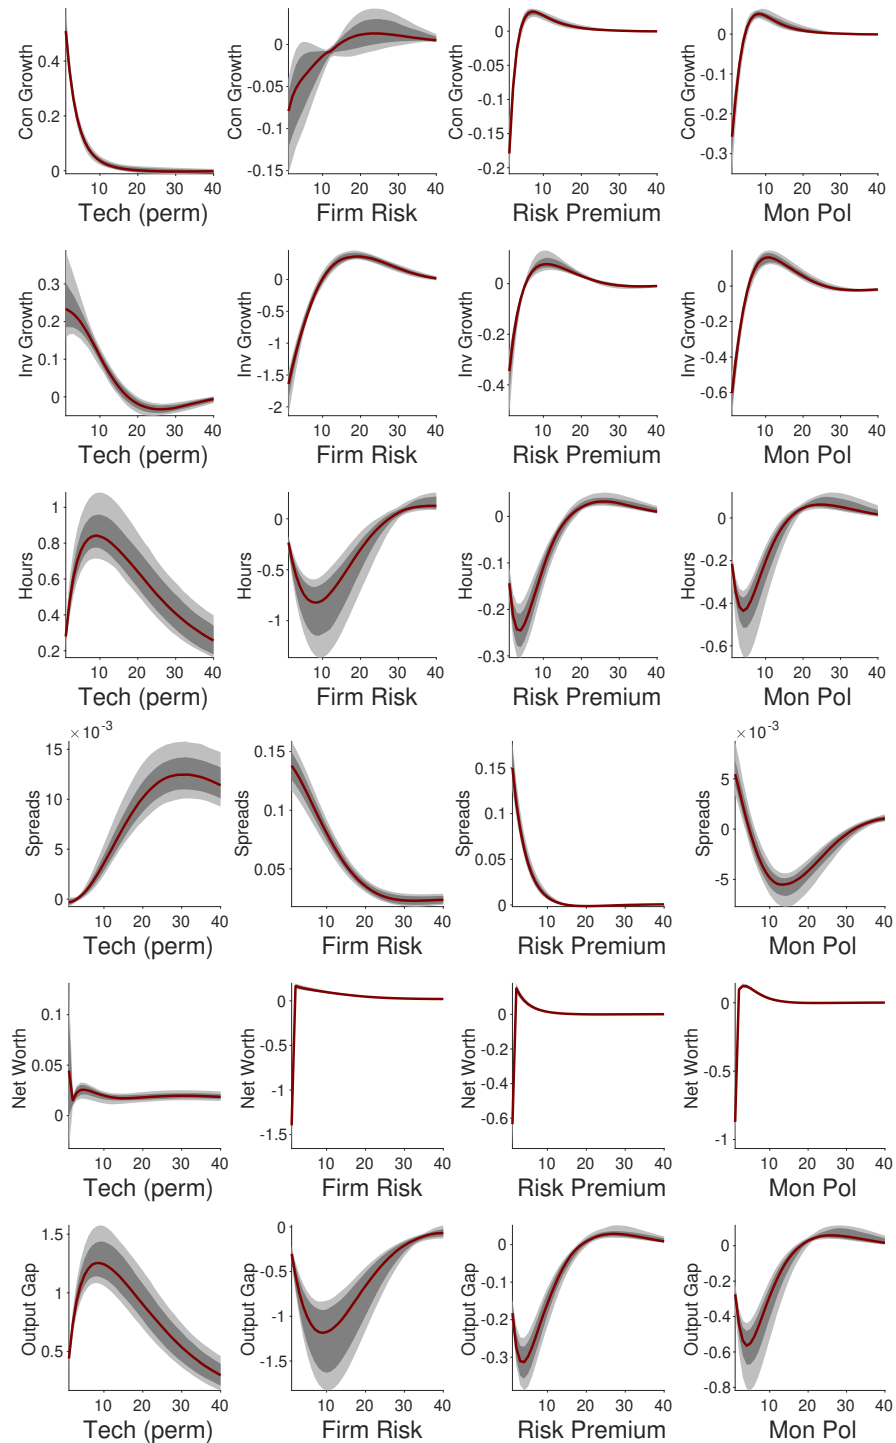
Note: Impulse responses are calculated following one-standard-deviation shocks and are reported in percentage points.

Figure A.2. OTHER IMPULSE RESPONSES FUNCTIONS
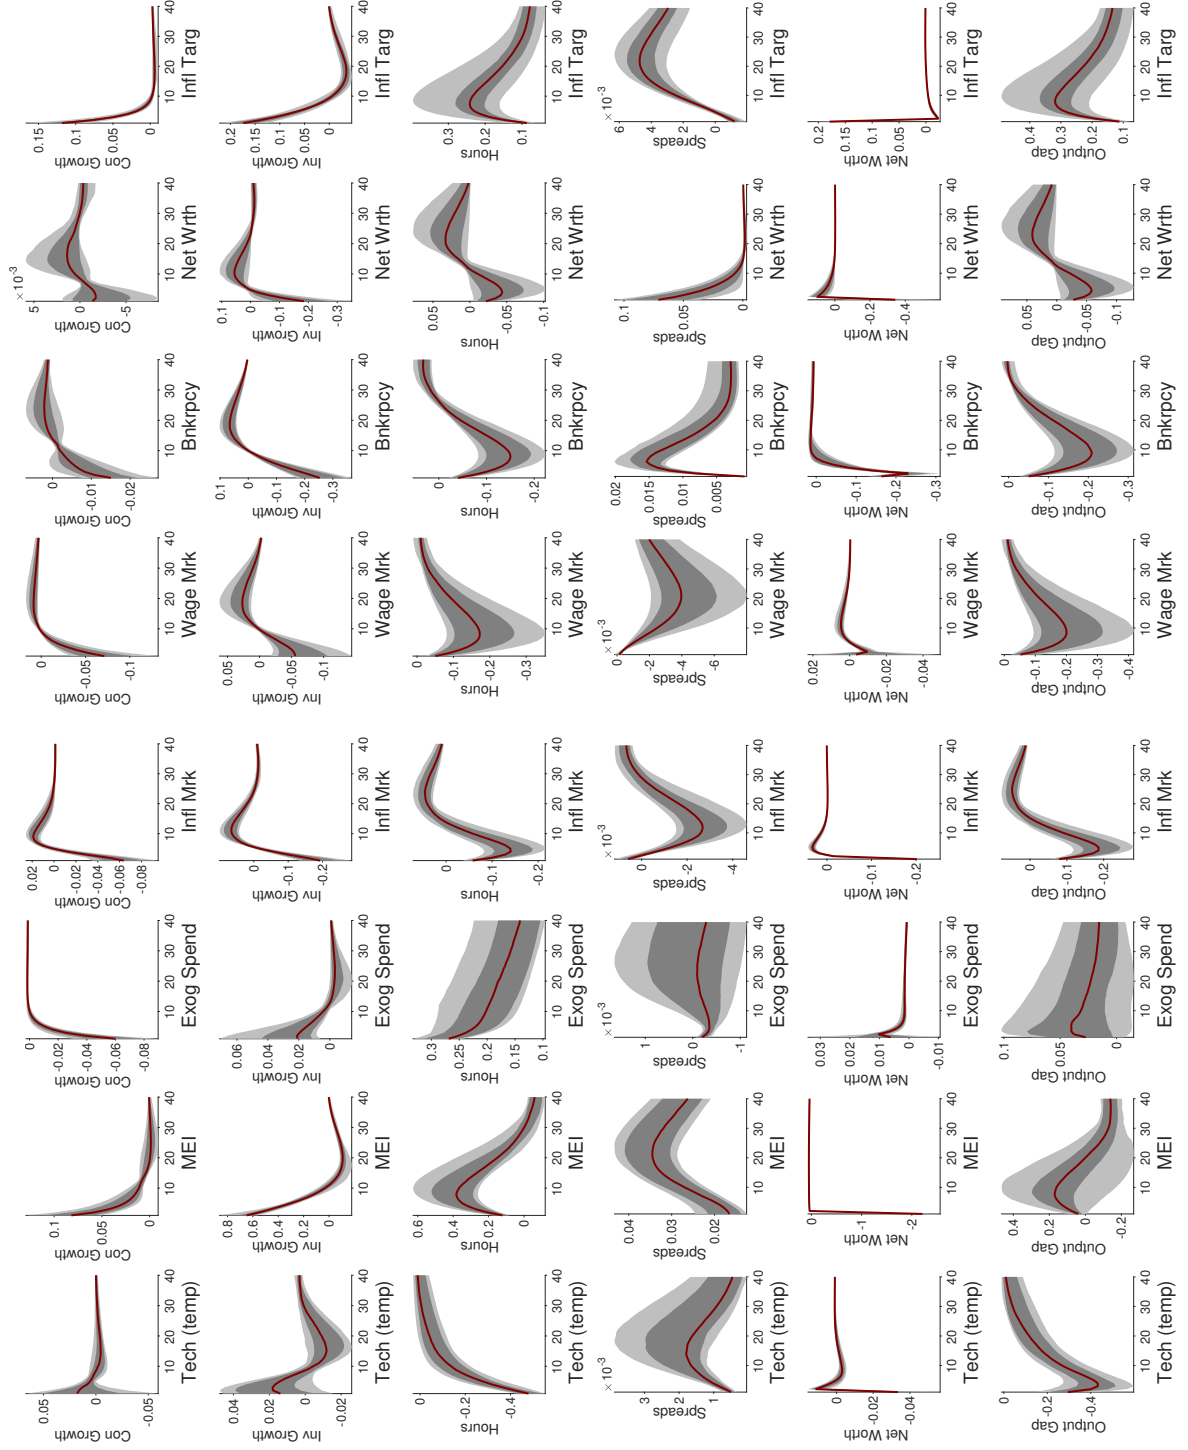
Note: Impulse responses are calculated following one-standard-deviation shocks and are reported in percentage points.

Figure A.3. EFFECTS OF DATA NEWS ON TWO-SIDED ESTIMATES OF STRUCTURAL SHOCKS.
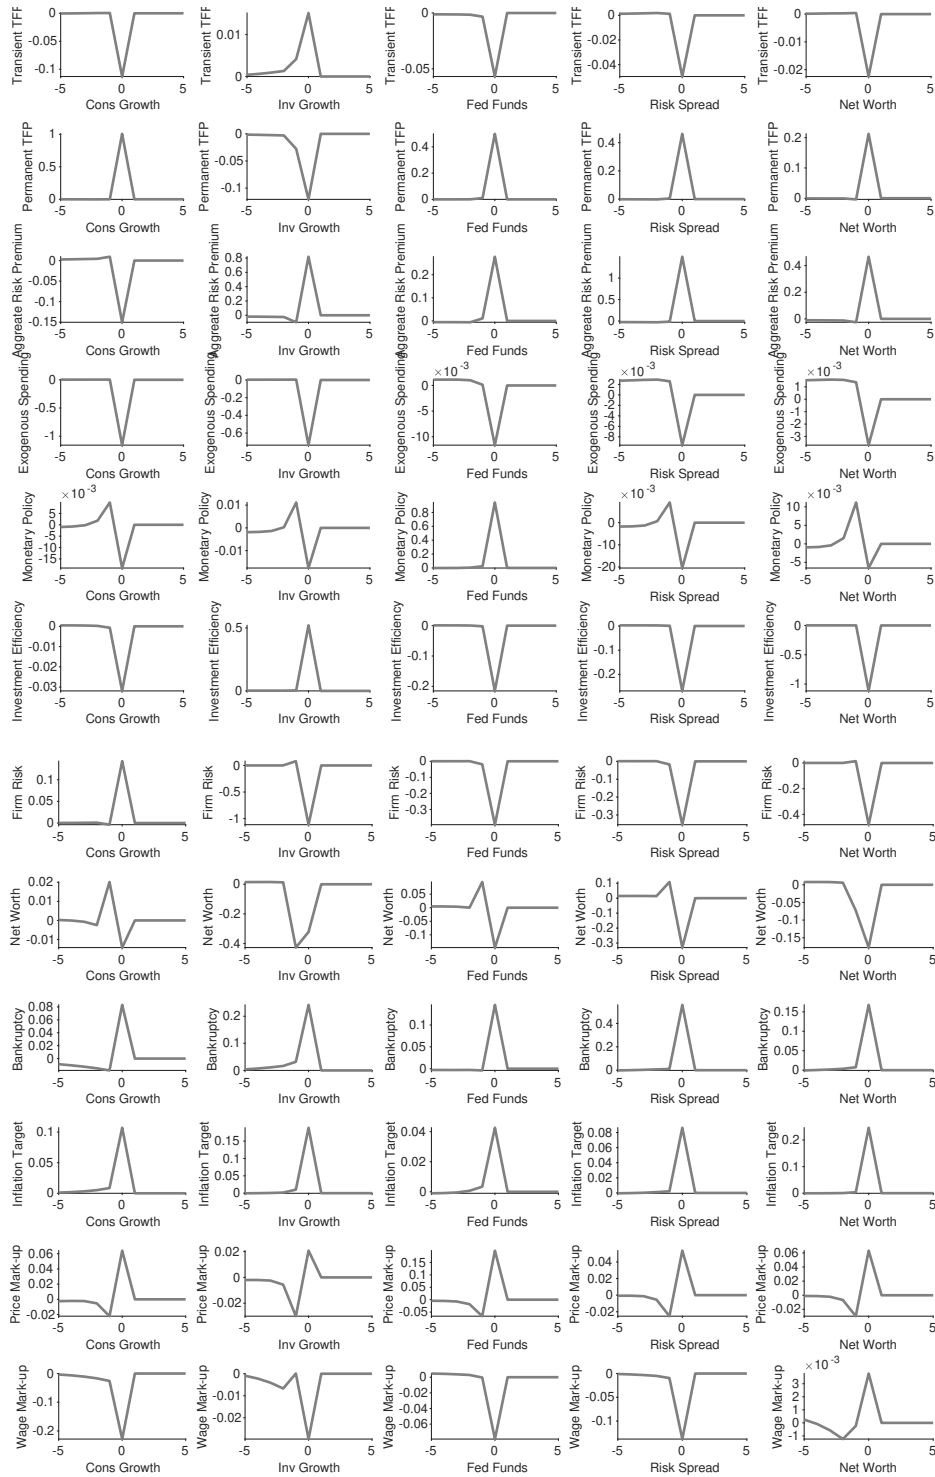
Note: Each row shows the effect of a one-standard deviation forecast error in a selected observable at time 0 on the two-sided estimate of various shocks, also expressed in standard deviations.

Figure A.4. EFFECTS OF SELECTED FORECAST ERRORS AT TIME 0 ON THE PATH OF THE OUTPUT GAP
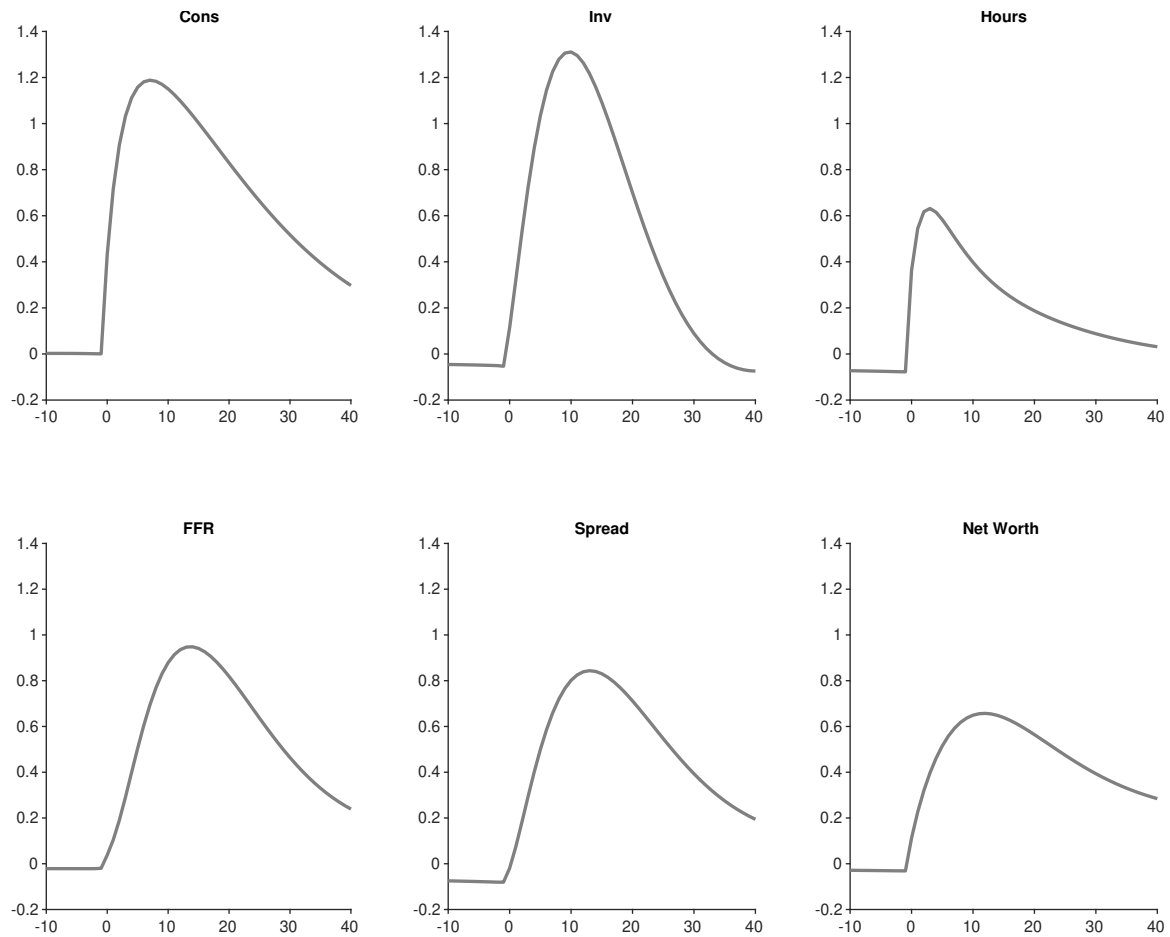
Note: Panels represent the effect of a one-standard-deviation forecast error in the given variable at time 0 on the estimated path of the output gap between 10 quarters before that date and 40 quarters after. Effects are calculated assuming that the news arrives 200 quarters after the beginning of the data set.