Misspecification and the Causes of Business Cycles

Dana Galizia*†

March 8, 2015

Abstract

Modern dynamic stochastic general equilibrium (DSGE) macroeconomics models generally feature several different exogenous shock processes. A standard tool in the quantitative macroeconomics toolbox for evaluating the individual importance of these shocks is a variance decomposition. The reliability of this tool depends importantly on having accurate estimates of the variances of the innovations to the exogenous shock processes. Using a novel framework, I show that when the DSGE model is misspecified and the shock variances are estimated using a likelihood-based approach, the resulting estimates are biased upward. Next, using the same framework, I propose a simple procedure to identify and partially correct for the effects of model misspecification on these variance estimates. As an example of its usefulness, I apply this procedure to a recent paper and find that it reduces the estimated variances of the shocks in the model by as much as a third of their respective naive estimates.

Key Words: Business Cycles, Shocks, Variance Decomposition; JEL Class.: C32, E3

*PhD Candidate, Vancouver School of Economics, The University of British Columbia; dana.galizia@gmail.com.

†I wish to thank Paul Beaudry especially for his invaluable advice and guidance. I would also like to thank Vadim Marmer, Henry Siu and Yaniv Yedid-Levi for their helpful comments. All errors are mine alone. I was supported by a SSHRC Graduate Scholarship while this paper was written.

1 Introduction

One of the central goals of macroeconomic research is to understand the root causes of business cycles. In recent years, much of the research devoted to this goal has been conducted in three steps. First, build a dynamic stochastic general-equilibrium (DSGE) model within which business-cycle fluctuations are generated by the random arrival of exogenous shocks[1] that are then propagated through endogenous mechanisms. Second, estimate the parameters of this DSGE model using real-world data and some estimation algorithm (such as maximum likelihood or, increasingly, Bayesian methods). Third, within the estimated model, evaluate the quantitative importance of each of the shocks in generating business-cycle fluctuations.[2]

If the true data-generating process (DGP) is contained in the set of DGPs spanned by the model constructed at step one above, then under some basic regularity conditions we can expect the parameters estimated at step two to converge to the true parameters, and in turn we may interpret the results of the third step as reflecting the actual importance of each type of shock in generating real-world economic fluctuations. However, when the true DGP is not contained in the set of DGPs spanned by the model—i.e., when the model is misspecified—this reasoning no longer holds. Since even the most optimistic economist would agree that all models are misspecified, this raises some important questions. Namely, what effect does misspecification have on shock variance estimates in a model? Does this effect depend on the severity of the misspecification? If misspecification is detected, can some correction be made to the variance estimates? To my knowledge, no framework currently exists to help answer these questions. This paper attempts to fill that gap in the literature.

The impact of misspecification on understanding the root causes of business cycles is highly relevant. Over the last decade or so, DSGE models have become ever richer, with a typical model now possessing a variety features that (it is argued) allow it to fit the data reasonably well. Deriving some degree of confidence from this good fit, it has become common to ask the model which types of economic shocks appear to be important for driving real-world business-cycle fluctuations.[3] Specifically, given a DSGE model that has been linearized around a non-stochastic steady state, and that contains structural shocks that are independent of one another both cross-sectionally and over time, it is a straightforward exercise to obtain the proportion of the variance of any endogenous variable attributable to an individual shock, i.e., to perform a variance decomposition. On the basis of a variance decomposition, one may then conclude, for example, that a particular shock is an important driving force for business cycle fluctuations. The usefulness of this variance decomposition tool in understanding the causes of real-world business cycles clearly depends fundamentally on the accuracy of the

[1] Throughout this paper, what I refer to as shocks are the innovations to an exogenous process, rather than the process itself. For example, if TFP evolves according to an AR(1) process $z_t = \rho z_{t-1} + \varepsilon_t$, with ε_t a sequence of i.i.d. random variables, then I refer to ε_t as the (exogenous) TFP shock.

[2] In the class of models considered in this paper, conditional on all other parameters there will be a one-to-one mapping between the variance of each shock and its respective quantitative importance. As such, I will use the terms "quantitative importance" and "shock variance" interchangeably.

[3] See, for example, Smets and Wouters (2007), Nolan and Thoenissen (2009), Christiano et al. (2010b), Justiniano et al. (2010, 2011), Schmitt-Grohé and Uribe (2011, 2012), Barsky and Sims (2012), Blanchard et al. (2013), Christiano et al. (2014), etc.

shock variance estimates. As I clarify in this paper, in the presence of misspecification these variance estimates will be systematically biased upward. While this poses a challenge to the conclusions typically drawn from a variance decomposition, as we shall see the news is not all bad. First, I show that as the severity of the misspecification becomes smaller (in a particular sense), this bias shrinks to zero. This reinforces the view that, even though a model may be misspecified, it may still yield useful conclusions. Second, I propose a simple method of adjusting the variance estimates that will reduce the degree of bias, at least in part. In combination, these two properties also imply that larger adjustments will typically be associated with more severe misspecification, so that this methodology can be used as an additional tool for measuring the severity of misspecification in a model. Since these adjustments will generally be different for each shock in the model, unlike other measures this one is unique in providing information to the modeller about misspecification at the level of individual shocks. As illustrated in an example below, this information can be useful in diagnosing the source of misspecification and, in turn, suggesting potential modifications to the model that may help to alleviate it.

The framework I propose revolves around the sequence of “smoothed shocks”, i.e., the sequence of realized shock values that, when fed into the model, allow it to exactly reproduce the observed data.[4] As shown below, if the DSGE model (for a given set of parameter values) were the true DGP, then under certain regularity assumptions the sequence of smoothed shocks will asymptotically recover the true values of the shocks. Since the true shocks are (by construction) independent of one another both intertemporally and cross-sectionally, this implies that if the model is the true DGP then the smoothed shocks should also be asymptotically independent of one another. This observation suggests a relatively simple diagnostic tool for detecting the presence of misspecification: if the smoothed shocks are correlated (beyond what can be reasonably explained by sampling error), then the model must be misspecified.[5]

The basic intuition for this result is quite simple. Any model—including the (unknown) true DGP—can be thought of as a collection of mechanisms, where each mechanism corresponds to the combination of a shock and its impulse response function. A model is misspecified if it is missing one or more mechanisms that are present in the true DGP. If this is the case, the smoothing algorithm—that is, the algorithm that finds a sequence of shocks that exactly reproduces the data—must systematically mix together two or more of the mechanisms that *are* in the model in order to replicate the mechanisms that are missing. It is precisely this “mixing” process that leads to correlations in the smoothed shocks.

I next show that, when the model is misspecified, conventional shock variance estimates will tend to overstate the importance of one or more of the shocks. The intuition for this bias is again quite simple: since a mechanism in the model is being used to explain not only the fluctuations in the data actually caused by that mechanism, but also fluctuations caused by

[4] In practice there will typically be more than one sequence capable of reproducing the data, in which case the smoothed shocks are taken to be the sequence that is most likely to have occurred (i.e., the one for which the joint probability density function is maximized).

[5] Note that, under the precise definition of “misspecification” adopted in this paper (see Section 3.1), the converse will not necessarily be true: it will be possible for the smoothed shocks to be uncorrelated, but for the model to nonetheless be misspecified.

certain mechanisms that are not contained in the model, the smoothing algorithm attributes too much variation to the model mechanisms, and thus tends to overestimate the variance of the model shocks.

After highlighting this potential source of bias, I propose a simple procedure to (at least partially) correct for it. First, I orthogonally decompose a smoothed shock into two components: its true value in producing variation in the data, and an additional component related to misspecification. Next, I show that the true value of the shock should be unpredictable using values of other *smoothed* shocks. The true shock is therefore a component of the OLS residual obtained after regressing the corresponding smoothed shock on other smoothed shocks. The variance of this residual therefore represents a less biased estimate of the true variance of the shock.

To illustrate a practical application, I apply the above methodology to a recent paper by Justiniano et al. (2010), and show that the source of bias discussed above may indeed have significant implications. Justiniano et al. (2010) construct a medium-scale New Keynesian model featuring a variety of different shocks. The variance decomposition obtained from their estimated model indicates that the majority (roughly 60 percent) of the variance of hours worked at business-cycle frequencies can be attributed to the investment shock (a shock to the relative productivity of new investment *vis-à-vis* the existing capital stock). This shock is also found to account for the majority of business-cycle variation in output and investment. As I show, however, six of the seven smoothed shocks in their model (including the investment shock) are significantly correlated, suggesting that the model is misspecified and that the naive variance estimates may be overstating the importance of these shocks. Applying the proposed correction procedure, I estimate that the true variance of the investment shock is smaller than the naive estimate by at least one-third.

I contribute here to two different bodies of literature. First, there is a substantial macroeconometric body of literature that explores issues of model fit and estimation/inference in potentially misspecified DSGE models (e.g., Gourieroux et al. (1993), Watson (1993), Canova (1994), Diebold et al. (1998), Schorfheide (2000), Hall and Inoue (2003), Dridi et al. (2007), Hnatkovska et al. (2012)). In addition, I also contribute to the vast literature that attempts to identify the sources of business cycle fluctuations. More specifically, I contribute to the recent literature that has considered investment shocks as a potential source of fluctuations, including the aformentioned Smets and Wouters (2003, 2007) and Justiniano et al. (2010), as well as Greenwood et al. (1988), Greenwood et al. (1997, 2000), Schmitt-Grohé and Uribe (2011, 2012), Justiniano et al. (2011), and Jermann and Quadrini (2012).

The remainder of the paper proceeds as follows. Section 2 introduces in a relatively informal way the basic framework and methodology of the paper, and presents a simple example to illustrate it. Section 3 then contains a formal presentation of the framework and methodology, while Section 4 provides some additional analysis. Section 5 then presents two additional examples in which the true DGP is known, before Section 6 applies the framework and methodology in a real-world example using the model of Justiniano et al. (2010). Finally, Section 7 concludes.

2 Basic framework and methodology

In this section I introduce the basic framework and methodology of the paper in a relatively informal way, with an eye towards conveying as clearly as possible the basic intuition. The ideas of this section are then discussed in formal detail in Section 3 below.

Suppose one has in hand an infinite sequence of past and future observations $Y^\infty \equiv \{Y_t\}_{t=-\infty}^{\infty}$, where Y_t is some jointly normally distributed mean-zero m-variate data process satisfying basic stationarity and ergodicity assumptions. A model for $\{Y_t\}$ is defined here as an expression for Y_t as a linear function of a (potentially infinite) history of r-dimensional i.i.d. normal random vectors, where any two different elements of that random vector are orthogonal to one another. That is, letting $\tilde{Y}_t$ denote the model counterpart of Y_t, a model is completely summarized by an MA(∞) representation

$$\tilde{Y}_t = \sum_{j=0}^{\infty} \tilde{\Psi}_j \varepsilon_{t-j} \tag{1}$$

where $\left\{\tilde{\Psi}_j\right\}$ is an absolutely-summable sequence of $m \times r$ matrices and ε_t is an r-vector of i.i.d. normal random variables with diagonal covariance matrix, the l-th diagonal element of which is given by $\tilde{\sigma}_l^2$. Assuming that the true DGP can also be written in this way, there are two particular models of interest: the econometric model (EM), and the true model (TM). Make the following important assumption about the EM.[6]

Assumption 1. (Invertibility) There exists some absolutely summable sequence $\{\psi_j\}$ of matrices such that, given an infinite sequence of past and future values of $\tilde{Y}_t$ generated by (1), we may recover the sequence $\{\varepsilon_t\}$ from[7]

$$\varepsilon_t = \sum_{j=-\infty}^{\infty} \psi_j \tilde{Y}_{t-j} \tag{2}$$

Conditions under which Assumption 1 holds are fairly general and will be discussed in detail in Section 4.1. For now, we will simply take it as given. Under this assumption, the sequence of smoothed shocks, denoted $\hat{\varepsilon}_t$, are obtained by applying the linear filter (2) to the data series Y_t (instead of its model counterpart $\tilde{Y}_t$). That is,

$$\hat{\varepsilon}_t \equiv \sum_{j=-\infty}^{\infty} \psi_j Y_{t-j} \tag{3}$$

It is straightforward to verify that if one substitutes the sequence $\hat{\varepsilon}_t$ obtained from (3) into

[6] Note that we do not make any such assumption about the TM, which need only in general satisfy (1).

[7] Note that I define invertibility as the ability to recover the structural shocks from some combination of current, past *and future* values of $\tilde{Y}_t$. This definition is more general than that used in the structural VAR literature (see, for example, Fernández-Villaverde et al. (2007), Sims (2012)), according to which the MA(∞) representation in (1) would be deemed invertible only if ε_t can be recovered from current and past values of $\tilde{Y}_t$ alone.

equation (2) for ε_t one recovers the data series Y_t.[8] Thus, the sequence $\hat{\varepsilon}_t$ is indeed the sequence of smoothed shocks defined in Section 1. If the EM is correctly specified (i.e., if the EM and TM are one and the same), then one can interpret the sequence of smoothed shocks as capturing the true values of these shocks in the real world. In general, however, the TM will contain a number of shocks that are absent in the EM, in which case the sequence of smoothed model shocks, $\hat{\varepsilon}_t$, would in general be different from their "true" values, which I denote by ε_t^*.

At this point it is worth clarifying what I mean by the "true value" of a shock. This notion will be made precise in Section 3, but for now it will be sufficient to illustrate the concept with an example. In particular, consider the distinction between a real-world TFP shock, and a smoothed TFP shock inferred from some real-world data set[9] using a misspecified model. The former is the true value of the shock, and the one whose variance we would ultimately like to know, but it is not directly observable. On the other hand, the latter *is* observable, but it will in general differ from the true value, and thus its variance will be different as well.

Let $\sigma_l^{*2} \equiv Var\left(\varepsilon_{l,t}^*\right)$ denote the true variance of the l-th model shock. This is the quantity the econometrician would like to estimate. Further, let $\hat{\sigma}_l^2 \equiv Var\left(\hat{\varepsilon}_{l,t}\right)$ be the variance of the corresponding smoothed shock. It can be verified that if the model is correctly specified then $\hat{\sigma}_l^2$ is the maximum-likelihood estimator of σ_l^{*2}. Since the model is implicitly assumed to be correctly specified when estimating the parameters via maximum-likelihood, $\hat{\sigma}_l^2$ is thus precisely the variance estimator that is typically used in the literature and that forms a key input into the variance decomposition.

When the model is misspecified, however, we will in general have $\hat{\sigma}_l^2 \neq \sigma_l^{*2}$. We are thus interested in the relationship between $\hat{\sigma}_l^2$ and σ_l^{*2}. Let $\nu_t \equiv \hat{\varepsilon}_t - \varepsilon_t^*$ denote the (unobserved) vector of "shock recovery errors", and make the following assumption.

Assumption 2. **(Orthogonality)** $\mathbb{E}\left[\varepsilon_t^* \nu_{t-k}'\right] = 0$ for all $k \in \mathbb{N}$.

The framework presented in detail in Section 3.1 will directly imply that Assumption 2 holds. Nonetheless, as with Assumption 1 we will for now simply take Assumption 2 as given. Under this assumption we can orthogonally decompose $\hat{\varepsilon}_t$ into two components: one capturing the true value of the shock (ε_t^*), and another entirely reflecting the fact that the EM is misspecified (ν_t). The immediate consequence of this orthogonal decomposition is that

$$\hat{\sigma}_l^2 = \sigma_l^{*2} + Var\left(\nu_{l,t}\right) \geq \sigma_l^{*2} \tag{4}$$

Thus, when the model is misspecified the naive variance estimate $\hat{\sigma}_l^2$ will overstate the true variance of the shock; that is, misspecification causes the variance estimates to be biased upwards. This result is the first key contribution of this paper.

Next, note that the true values of the shocks are, by definition, independent of one another both cross-sectionally and intertemporally. That is, we have $Cov\left(\varepsilon_{l,t}^*, \varepsilon_{i,s}^*\right) = 0$ whenever

[8] See Proposition 1 in Section 3 for a formal statement and proof of this fact.

[9] Assume for the sake of argument that this data set does not include the TFP process itself, so that the value of the TFP shocks must be inferred from the behavior of other endogenous variables in the system.

$(l,t) \neq (i,s)$. It thus follows immediately that

$$Cov\left(\hat{\varepsilon}_{l,t}, \hat{\varepsilon}_{i,s}\right) = Cov\left(\nu_{l,t}, \nu_{i,s}\right) \tag{5}$$

Thus, if two smoothed shocks exhibit non-zero covariance, this can only be due to misspecification. This is the second key contribution of this paper. Finally, an implication of this result is that any component of $\hat{\varepsilon}_{l,t}$ that can be predicted using $\hat{\varepsilon}_{i,s}$ must be attributable to misspecification rather than to $\varepsilon^*_{l,t}$. This suggests a simple procedure to correct (at least in part) the variance estimates: Regress $\hat{\varepsilon}_{l,t}$ on date-t values of shocks $i \neq l$ and on leads and lags of all shocks (including leads and lags of $\hat{\varepsilon}_{l,t}$). The error term from this regression is, by construction, the component of $\hat{\varepsilon}_{l,t}$ which *cannot* be predicted using the other shocks. The variance of this error term, which I denote $\bar{\sigma}^2_l$, therefore contains all of the variation in $\hat{\varepsilon}_{l,t}$ due to $\varepsilon^*_{l,t}$, plus an additional non-negative component. We will therefore have $\hat{\sigma}^2_l \geq \bar{\sigma}^2_l \geq \sigma^{*2}_l$, so that $\bar{\sigma}^2_l$ is a less-biased estimate of σ^{*2}_l than the naive maximum-likelihood estimate. This result is the third key contribution of this paper.

2.1 Example 1: Missing shocks

Consider a simple economy populated by a representative household with preferences given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(C_t - \eta_t L_t\right)$$

Here, C_t is consumption, L_t is labor supplied, β is the constant discount factor, and η_t, which captures the relative disutility of labor, follows the exogenous stochastic process

$$\log \eta_t = \varepsilon^*_{\eta,t}$$

where $\varepsilon^*_{\eta,t}$ is i.i.d. $N\left(0,\, \sigma^*_\eta\right)$. The household produces output X_t using technology $X_t = A_t L_t^{1-\alpha}$, where productivity A_t follows exogenous stochastic process

$$\log A_t = \varepsilon^*_{A,t}$$

with $\varepsilon^*_{A,t}$ i.i.d. $N\left(0,\, \sigma^*_A\right)$. In addition to consuming the output it produces, the household also receives a stochastic endowment $\varepsilon^*_{w,t}$ which is i.i.d. $N\left(0,\, \sigma^*_w\right)$. Thus, the household budget constraint is given by

$$C_t = A_t L_t^{1-\alpha} + \varepsilon^*_{w,t}$$

Taking first-order conditions and log-linearizing the model around its non-stochastic steady state yields the following equations for consumption and labor

$$c_t = -\frac{1-\alpha}{\alpha}\varepsilon^*_{\eta,t} + C^{-1}\varepsilon^*_{w,t} + \frac{1}{\alpha}\varepsilon^*_{A,t} \tag{6}$$

$$l_t = -\frac{1}{\alpha}\varepsilon^*_{\eta,t} + \frac{1}{\alpha}\varepsilon^*_{A,t} \tag{7}$$

where lower-case variables indicate log-deviations from steady state and C is the steady-state value of C_t.

Suppose that the value of the parameter α was known to the econometrician, who wishes to estimate via maximum likelihood the variances of the stochastic processes using data on c_t and l_t. Suppose, however, that the econometrician wrongly assumes that $\sigma_A^* = 0$, i.e., that productivity is constant. This econometrician would then estimate $\sigma \equiv (\sigma_\eta, \sigma_w)'$ using the misspecified model

$$\tilde{c}_t = -\frac{1-\alpha}{\alpha}\varepsilon_{\eta,t} + C^{-1}\varepsilon_{w,t} \tag{8}$$

$$\tilde{l}_t = -\frac{1}{\alpha}\varepsilon_{\eta,t} \tag{9}$$

The smoothed shocks $\hat{\varepsilon}_t \equiv (\hat{\varepsilon}_{\eta,t}, \hat{\varepsilon}_{w,t})'$ can be obtained as

$$\hat{\varepsilon}_t = \begin{pmatrix} -\alpha l_t \\ C\left[c_t - (1-\alpha) l_t\right] \end{pmatrix}$$

and the maximum-likelihood variance estimates are

$$\begin{pmatrix} \hat{\sigma}_\eta^2 \\ \hat{\sigma}_w^2 \end{pmatrix} = \begin{pmatrix} \alpha^2 \sigma_l^{*2} \\ C^2 \phi^* \end{pmatrix}$$

where

$$\phi^* \equiv Var\left[c_t - (1-\alpha) l_t\right] = \sigma_c^{*2} - 2(1-\alpha)\sigma_{cl}^* + (1-\alpha)^2 \sigma_l^{*2}$$

and $\sigma_c^{*2} \equiv Var(c_t)$, $\sigma_l^{*2} \equiv Var(l_t)$ and $\sigma_{cl}^* \equiv Cov(c_t, l_t)$ are the relevant data moments. Plugging in the (known) true values for these data moments, we obtain

$$\begin{pmatrix} \hat{\sigma}_\eta^2 \\ \hat{\sigma}_w^2 \end{pmatrix} = \begin{pmatrix} \sigma_\eta^{*2} + \sigma_A^{*2} \\ \sigma_w^{*2} + C^2\sigma_A^{*2} \end{pmatrix}$$

Thus, the maximum-likelihood estimates overstate the true variances by an amount which is increasing in the degree of misspecification, as captured by a non-zero value of σ_A^{*2}.

Next, to obtain the corrected estimates, note that

$$Cov(\hat{\varepsilon}_{\eta,t}, \hat{\varepsilon}_{w,t}) = -C\sigma_A^{*2} \tag{10}$$

Thus, we obtain the regression equations

$$\hat{\varepsilon}_{\eta,t} = -\frac{C\sigma_A^{*2}}{\sigma_w^{*2} + C^2\sigma_A^{*2}}\hat{\varepsilon}_{w,t} + \xi_{\eta,t}$$

$$\hat{\varepsilon}_{w,t} = -\frac{C\sigma_A^{*2}}{\sigma_\eta^{*2} + \sigma_A^{*2}}\hat{\varepsilon}_{\eta,t} + \xi_{w,t}$$

where the ξ's are standard OLS residuals that are orthogonal to the regressors. Some algebra yields

$$\sigma_\eta^{*2} \leq \bar{\sigma}_\eta^2 \equiv Var\left(\xi_{\eta,t}\right) = \sigma_\eta^{*2} + \frac{\sigma_w^{*2}}{\sigma_w^{*2} + C^2\sigma_A^{*2}}\sigma_A^{*2} \leq \hat{\sigma}_\eta^2$$

$$\sigma_w^{*2} \leq \bar{\sigma}_w^2 \equiv Var\left(\xi_{w,t}\right) = \sigma_w^{*2} + \frac{\sigma_\eta^{*2}}{\sigma_\eta^{*2} + \sigma_A^{*2}}C^2\sigma_A^{*2} \leq \hat{\sigma}_w^2$$

That is, the corrected shock variance estimates lie between the true variances and the maximum-likelihood variance estimates.

To understand these results, consider what happens when this economy is hit with a positive technology shock. There are two channels through which this shock affects the economy: the rise in productivity directly increases output and consumption, while the increase in the marginal product of labor leads the household to both work and consume more. As seen in equations (6)-(7), the combined effect of these two channels is that consumption and labor increase in the same proportion.

In the misspecified model (8)-(9), however, there does not exist a mechanism that can produce equiproportional changes in consumption and labor. A negative labor-disutility shock has a similar labor-leisure substitution effect as the technology shock, but it lacks the direct increase in output, and thus causes consumption to increase by relatively less than labor. Only by combining a negative labor-disutility shock with a positive endowment shock (which increases consumption but has no effect on labor) can the misspecified model get consumption and labor to increase in the same proportion. This "mixing" of shocks is what produces the negative covariance in (10). While maximum-likelihood estimation explicitly ignores the information conveyed by this mixing of shocks, the methodology proposed here uses it to correct, in part, the resulting overestimate of the shock variances.

3 Formal framework

Having introduced the framework and methodology in a relatively informal way in the previous section, this section presents a more rigorous foundation for several of the results of Section 2. In particular, I make precise the definition of the "true values" of the shocks in the EM, ε_t^*, which form a crucial part of the reasoning in Section 2, and show how the framework proposed here directly implies Assumption 2.

While formally stating the results of this section requires first setting up a fair amount of mathematical machinery, the basic intuition is fairly simple. Crucially, in order to make the concept of the "true value" of an EM shock well-defined, we must first answer the following questions: Under what circumstances and in what sense can we consider a shock in the EM and a shock in the TM to be fundamentally the same, even if the EM is misspecified? Under what circumstances should we view two such shocks as being fundamentally different? The framework presented in this section answers these as follows: Two such shocks are fundamentally the same if they are associated with impulse response functions (IRFs) that are identical up to a scaling and time-shift factor. Otherwise, they are different.

With these answers in hand, it then becomes straightforward to formally define the concept of the "true value" of a shock in the EM: If there is a shock in the TM that is fundamentally the same, then its value is the true value. If there is no such shock, then its true value is zero.[10] Assumption 2 will then follow directly from the fact that shocks that are fundamentally different from one another are associated with independent white noise processes.

3.1 Set-up

This basic set-up of the framework proceeds as follows. First, the set of all stationary IRFs is partitioned into groups that are equivalent up to a particular scaling and time-shift factor, and to each cell of this partition we associate an independent univariate Gaussian white noise process. I then create a measurable space (G, Γ) from the partition by choosing exactly one IRF from each cell of the partition to form a set G, and then defining an appropriate σ-algebra Γ on that set.

Step 1: Partition the IRFs and associate a white noise to each cell

Let $\ell^{1,m}$ denote the space of absolutely summable sequences in $\mathbb{R}^m$ equipped with the product topology. We will interpret a typical element $\mathbf{\Psi}$ of $\ell^{1,m}$ as an IRF, and write

$$\mathbf{\Psi} = (\Psi_j)_{j=0}^{\infty}$$

Next, for $\mathbf{\Psi} \in \ell^{1,m}$, define $J(\mathbf{\Psi}) = \min\{j : \Psi_j \neq 0\}$ as the first non-zero element of $\mathbf{\Psi}$.[11] Define an equivalence relation $\sim$ on $\ell^{1,m}$ by

$$\left(\mathbf{\Psi} \sim \bar{\mathbf{\Psi}}\right) \Leftrightarrow \left(\exists a \neq 0 \text{ such that } \Psi_{J(\mathbf{\Psi})+j} = a\bar{\Psi}_{J\left(\bar{\mathbf{\Psi}}\right)+j} \text{ for all } j \geq 0\right)$$

In words, if $\mathbf{\Psi} \sim \bar{\mathbf{\Psi}}$, then $\bar{\mathbf{\Psi}}$ is a scaled, time-shifted version of $\mathbf{\Psi}$, in which case we will say that $\mathbf{\Psi}$ and $\bar{\mathbf{\Psi}}$ are indistinguishable. For convenience, when it exists we will denote the value of a in the above definition by $a(\mathbf{\Psi}, \bar{\mathbf{\Psi}})$.[12] For any $\mathbf{\Psi} \in \ell^{1,m}$, we define $[\mathbf{\Psi}] \equiv \left\{\bar{\mathbf{\Psi}} \in \ell^{1,m} : \mathbf{\Psi} \sim \bar{\mathbf{\Psi}}\right\}$ as the equivalence class of $\mathbf{\Psi}$, and let $\mathcal{Q}$ denote the set of all equivalence classes in $\ell^{1,m}$, which is a partition of that set. Thus, each element $Q \in \mathcal{Q}$ is a set of IRFs that are equivalent to one another according to $\sim$. To each such $Q \in \mathcal{Q}$, we then associate an independent univariate white noise process $\eta_t(Q)$, where $\eta_t(Q) \sim N(0,1)$ and $\mathbb{E}[\eta_t(Q)\,\eta_{t-k}(Q')] = 0$ for all $k \in \mathbb{Z}$ when $Q \neq Q'$.

[10] There is a sense in which this definition is quite demanding, in that a shock in the EM must meet a high standard—its exact associated IRF must exist in the TM—in order for its "true" variance to be non-zero. Thus, there is a sort of discontinuity, in that if the IRF exists in the TM then the true variance is strictly positive, but if it differs by even an arbitrarily small amount from some IRF in the TM then the true variance is zero. However, as shall become clear in Section 5 below, in practice this discontinuity will not be important, in that the observed degree of misspecification—and therefore any adjustments made to the variance estimates—will approach zero as the IRFs in the EM approach those in the TM.

[11] If no such element exists, i.e., if $\Psi_j = 0$ for all j, we set $J(\mathbf{\Psi}) = 0$.

[12] Note that the order of the arguments matters here, and in particular, $a(\mathbf{\Psi}, \bar{\mathbf{\Psi}}) = 1/a(\bar{\mathbf{\Psi}}, \mathbf{\Psi})$.

Step 2: Create a measurable space (G, Γ) from the partition

Let $G \subset \ell^{1,m}$ be some collection of IRFs formed by choosing exactly one element from each $Q \in \mathcal{Q}$. Thus, any collection of IRFs in G will be distinguishable, and for any $\mathbf{\Psi} \in \ell^{1,m}$ there exists a unique element in G indistinguishable from it. Denote this element of G by $G(\mathbf{\Psi})$. Next, the integral representation introduced in the subsequent section requires defining a measure on the subsets of G. To do so requires first formally defining a σ-algebra on G. For any $\gamma \subset G$, let

$$\mathcal{Q}_\gamma \equiv \{[\mathbf{\Psi}] : \mathbf{\Psi} \in \gamma\} \subset \mathcal{Q}$$

Thus, $\mathcal{Q}_\gamma$ is the set of equivalence classes spanned by the elements of γ. We will say that γ is an open subset of G if and only if there exists some open set $H \subset \ell^{1,m}$ formed by choosing exactly one element from each $Q \in \mathcal{Q}_\gamma$. Letting Γ denote the Borel σ-algebra generated by these open sets, (G, Γ) is a measurable space.

3.2 Integral representation

In this section, I show that, for any G as constructed above, any model of the form given in (1) can be equivalently recast as

$$\tilde{Y}_t = \int_G \left(\sum_{j=0}^{\infty} \Psi_j \eta_{t-j} \left([\mathbf{\Psi}] \right) \right) d\sigma \tag{11}$$

for some measure σ on (G, Γ). I will refer to this representation for $\tilde{Y}_t$ as its integral representation on G.

To see how this representation can be constructed, fix G and a model (1). Write $\tilde{\Psi}_j = \left(\tilde{\Psi}_j^{(1)}, \ldots, \tilde{\Psi}_j^{(r)} \right)$, and define $\tilde{\mathbf{\Psi}}^{(l)} \equiv \left(\tilde{\Psi}_j^{(l)} \right)_{j=0}^{\infty} \in \ell^{1,m}$ for $l = 1, \ldots, r$. Thus, $\tilde{G} \equiv \left\{ \tilde{\mathbf{\Psi}}^{(1)}, \ldots, \tilde{\mathbf{\Psi}}^{(r)} \right\}$ is the set of IRFs contained in the model. For simplicity, I will restrict attention to the case where $J(\mathbf{\Psi}) = 0$ for all $\mathbf{\Psi} \in G \cup \tilde{G}$.[13] It is nonetheless straightforward to extend the argument to the general case.

Next, let σ be the (unique) measure on (G, Γ) satisfying (1) $\sigma(\{\mathbf{\Psi}\}) = a(\mathbf{\Psi}, G(\mathbf{\Psi}))\, \tilde{\sigma}_l$ if $\mathbf{\Psi} = G\left(\tilde{\mathbf{\Psi}}^{(l)}\right)$, and (2) for any $\gamma \in \Gamma$ with $\gamma \cap G\left(\tilde{G}\right) = \emptyset$,[14] we have $\sigma(\gamma) = 0$. In words, σ is a measure that assigns weight $a(\mathbf{\Psi}, G(\mathbf{\Psi}))\, \tilde{\sigma}_l$ to the IRF in G that is indistinguishable from

[13] We also assume throughout that $\tilde{\mathbf{\Psi}}^{(l)} \sim \tilde{\mathbf{\Psi}}^{(i)}$ implies $l = i$, and that $\tilde{\mathbf{\Psi}}^{(l)} \neq \mathbf{0}$ for any l. If either of these assumptions are violated, the model itself contains a fundamental indeterminacy that can be fixed by, for example, eliminating $\tilde{\mathbf{\Psi}}^{(l)}$ from the model in either case.

[14] With slight abuse of notation, $G\left(\tilde{G}\right)$ here denotes the unique set of elements of G that are indistinguishable from the elements of $\tilde{G}$.

the l-th IRF in $\tilde{G}$, and zero elsewhere. We then have

$$\begin{aligned}\int_G \left(\sum_{j=0}^{\infty} \Psi_j \eta_{t-j}\left([\mathbf{\Psi}]\right)\right) d\sigma &= \sum_{l=1}^{r} \sigma\left(\left\{G\left(\tilde{\mathbf{\Psi}}^{(l)}\right)\right\}\right) \sum_{j=0}^{\infty} G\left(\tilde{\mathbf{\Psi}}^{(l)}\right)_j \eta_{t-j}\left(\left[\tilde{\mathbf{\Psi}}^{(l)}\right]\right) \\ &= \sum_{j=0}^{\infty}\sum_{l=1}^{r} a\left(\tilde{\mathbf{\Psi}}^{(l)}, G\left(\tilde{\mathbf{\Psi}}^{(l)}\right)\right) G\left(\tilde{\mathbf{\Psi}}^{(l)}\right)_j \tilde{\sigma}_l \eta_{t-j}\left(\left[\tilde{\mathbf{\Psi}}^{(l)}\right]\right) \\ &= \sum_{j=0}^{\infty}\sum_{l=1}^{r} \tilde{\Psi}_j^{(l)} \tilde{\sigma}_l \eta_{t-j}\left(\left[\tilde{\mathbf{\Psi}}^{(l)}\right]\right)\end{aligned}$$

Setting

$$\varepsilon_{l,t} \equiv \tilde{\sigma}_l \eta_t \left(\left[\tilde{\mathbf{\Psi}}^{(l)}\right]\right)$$

this is clearly equivalent to the representation in (1).

3.3 Orthogonal decompositions

In this section, using the integral form introduced above, I show that for any given EM we may obtain a decomposition of the TM into a component of Y_t that is explained by IRFs in the EM, and an orthogonal component that is not. Further, I establish that, in combination with Assumption 1, this result implies Assumption 2.

In particular, fix a model, and choose G such that $\tilde{\mathbf{\Psi}}^{(l)} \in G$ for all l (and thus $G\left(\tilde{\mathbf{\Psi}}^{(l)}\right) = \tilde{\mathbf{\Psi}}^{(l)}$).[15] Let σ^* denote the measure associated with the integral representation of the TM on G, so that the TM can be written

$$Y_t = \int_G \left(\sum_{j=0}^{\infty} \Psi_j \eta_{t-j}\left([\mathbf{\Psi}]\right)\right) d\sigma^*$$

Letting $\tilde{G}^c$ be the complement of $\tilde{G}$ in G, we can write

$$\begin{aligned}Y_t &= \int_{\tilde{G}} \left(\sum_{j=0}^{\infty} \Psi_j \eta_{t-j}\left([\mathbf{\Psi}]\right)\right) d\sigma^* + \int_{\tilde{G}^c} \left(\sum_{j=0}^{\infty} \Psi_j \eta_{t-j}\left([\mathbf{\Psi}]\right)\right) d\sigma^* \\ &= \sum_{j=0}^{\infty} \tilde{\Psi}_j \varepsilon^*_{t-j} + \int_{\tilde{G}^c} \left(\sum_{j=0}^{\infty} \Psi_j \eta_{t-j}\left([\mathbf{\Psi}]\right)\right) d\sigma^*\end{aligned}$$

where $\tilde{\Psi}_j \equiv \left(\tilde{\Psi}_j^{(1)}, \ldots, \tilde{\Psi}_j^{(r)}\right)$, $\varepsilon^*_{l,t} \equiv \sigma^*_l \eta_t \left(\left[\tilde{\mathbf{\Psi}}^{(l)}\right]\right)$, $\sigma^*_l \equiv \sigma^*\left(\tilde{\mathbf{\Psi}}^{(l)}\right)$, and $\varepsilon^*_t \equiv \left(\varepsilon^*_{1,t}, \ldots, \varepsilon^*_{r,t}\right)'$. More compactly, we can write this as

$$Y_t = \tilde{Y}^*_t + Z_t \tag{12}$$

[15] This choice of G is for simplicity of exposition only. With the appropriate modifications, the desired result will hold for any G.

$\tilde{Y}_t^*$ may be interpreted as the variation in Y_t which is generated by the IRFs contained in the EM (or their equivalents according to $\sim$), and Z_t as the variation in Y_t generated by IRFs for which no equivalent exists in the EM. Note that the EM is correctly specified if and only if $Var\left(Z_t\right) = 0$; otherwise, it is misspecified. Note also that, irrespective of whether the model is the true DGP, $\mathbb{E}\left[\varepsilon_t^* Z_{t-k}'\right] = 0$ for all k, and therefore in turn $\mathbb{E}\left[\tilde{Y}_t^* Z_{t-k}'\right] = 0$, so that (12) is an orthogonal decomposition of Y_t.

Next, substituting (12) into (3) we may obtain

$$\hat{\varepsilon}_t = \sum_{j=-\infty}^{\infty} \psi_j \tilde{Y}_{t-j}^* + \sum_{j=-\infty}^{\infty} \psi_j Z_{t-j} \tag{13}$$

By Assumption 1, the first term on the right-hand side of this expression is the true value of the EM shocks ε_t^*, and thus the second term is the recovery error ν_t defined in Section 2. Further, these two terms are clearly orthogonal to one another at all leads and lags, so that $\mathbb{E}\left[\varepsilon_t^* \nu_{t-k}'\right] = 0$ for all k. That is, in combination with Assumption 1, the framework developed above necessarily implies that Assumption 2 holds. We thus verify the upward-bias result from Section 2,

$$\hat{\sigma}_l^2 = \sigma_l^{*2} + Var\left(\nu_{l,t}\right) \geq \sigma_l^{*2}$$

Furthermore, it is straightforward to see that when the model is misspecified (i.e., when $Var\left(Z_t\right) \neq 0$), we will have $Var\left(\nu_{l,t}\right) > 0$ for at least one l, and therefore $\hat{\sigma}_l^2 > \sigma_l^{*2}$; that is, in the presence of misspecification the naive variance estimate will necessarily overstate the true variance of at least one shock. It can be easily verified that the other results of Section 2 also hold.

3.4 Discussion

To anticipate one potential objection to the framework proposed in this paper, it should be noted that, if we were to find a process for the smoothed shocks that exhibits correlation, it may be the case that the *true* shocks are in fact correlated. In this case, there may be a desire to view independence of the shocks as a modelling approximation made for convenience, rather than a property to be taken seriously (as it is in this framework). As such, one might argue, we should not be too concerned if we find evidence of its violation. I find this argument unconvincing for several reasons. First, if the shocks truly are correlated, there must be some underlying economic reason for this. Yet correlation among smoothed shocks is rarely acknowledged in the literature,[16] let alone justified using some economic argument. Second, though it would increase the number of estimated parameters, it would be a relatively

[16] Ingram et al. (1994) are a notable exception. They extend the canonical neoclassical growth model of King et al. (1988), which features only a shock to productivity, by adding two additional shocks. For several different calibrations of the model, they recover the implied series for the shocks using U.S. data on output, consumption and labor, and find that the shocks exhibit substantial correlation, both cross-sectionally and intertemporally.

straightforward exercise to allow for correlation in the EM,[17] but again, such an exercise appears to be rarely done in the literature.[18] Third, if we do allow for such correlation in the model, rational expectations dictates that agents in the model know this. When there are intertemporal correlations between the shocks, this has the capacity to significantly alter agents' behavior, since current shock realizations would carry information about future values of the shocks. There is little reason to believe *a priori* that estimating such a model would yield even qualitatively similar results to those obtained in the case where independence is assumed. Finally, as demonstrated in Section 5 below, to the extent that independence is a reasonable—though not necessarily perfect—approximation to the true structure of the exogenous processes, the bias-correction procedure proposed here would have a minimal effect on variance estimates. Thus, if resulting corrections are in fact large, this should be taken at the very least as clear evidence that independence is *not* a reasonable approximation.

Next, it may be helpful to view the framework presented in Sections 3.1-3.3 as a DSGE analogue to a typical regression setting. In the latter, the regression error term is, conceptually, the residual component of the dependent variable after as much of its variation as possible is explained by the regressors. In general, no economic meaning is assigned to this residual component. Rather, it is taken as arising from the econometrician's ignorance about the full process underlying the dependent variable. In a DSGE setting, the shocks are often taken to be a similar residual component. Unlike in a regression setting, however, in a DSGE model the shocks are given an economic interpretation and are thus *a part of the econometrician's explanation*, not residuals standing in for her ignorance. If we believe that factors beyond those captured by the model are of some relevance, this would motivate the explicit consideration of some orthogonal error term analogous to the regression residual. One may view the framework above as motivating such a residual (see, e.g., Z_t in equation (12)).

To close this section, it is worth briefly mentioning here that there are, conceptually, two distinct types of misspecification that arise in the above framework. In the first type, the IRFs in the EM are also found in the TM, but there are IRFs in the TM that are not present in the EM. This is the case in Example 1. In the second type of misspecification, there may be shocks in the EM and TM with the same economic interpretation but for which the associated IRF in the EM differs (however mildly) from the corresponding IRF in the TM. The framework introduced here does not distinguish between these two types of misspecification, despite the fact that there may be some desire to be more forgiving toward the second type. Nonetheless, as will be illustrated in the examples of Section 5 below, if the EM is a reasonable approximation to the TM, we will continue to find an important role for the EM shocks even if their precise IRFs differ somewhat from those in the EM.

[17] In the above framework, this would in practice consist of adding one or more IRFs to the EM that are linear combinations of existing IRFs. In other words, this would effectively increase the number of shocks in the model.

[18] There are some infrequent exceptions to this. For example, in Smets and Wouters (2007), innovations to TFP are also allowed to impact the exogenous spending shock.

4 Further analysis

In this section, I provide basic conditions on the EM and TM that are sufficient for the key results of Section 2 to hold, then re-state those results formally.

4.1 Econometric model

Consider the class of linearized DSGE models in state-space form as

$$X_t = AX_{t-1} + B\varepsilon_t \tag{14}$$

Here, X_t is an n-vector of (possibly unobservable) model variables, ε_t is an m-vector of i.i.d. normal random variables with diagonal covariance matrix, the l-th diagonal element of which is given by σ_l^2, and A and B are $(n \times n)$- and $(n \times m)$-matrix-valued functions of P, respectively, where $P \in \mathcal{P}$ is the vector of "deep" model parameters (excluding $\sigma \equiv (\sigma_1, \ldots, \sigma_m)'$). The eigenvalues of A are all strictly less than one in modulus, so that $\{X_t\}$ is a covariance-stationary process.

For what follows, I assume that the parameter vector P is given. There is a substantial literature that considers methods for obtaining parameterizations in this class of models.[19] I wish to sidestep this issue altogether and take as a starting point a particular parameterization (or a set of parameterizations, as the case may be) that the econometrician has identified as in some sense best for his or her purposes. The theoretical and methodological considerations that follow should thus be viewed as tools to analyze the implications of a particular value of P in the model, rather than of the model itself. Henceforth, for the sake of brevity, I shall generally suppress dependence of the analysis on the parameterization P, with the tacit understanding that all discussion and results are related to that specific parameterization only.

The m-vector of observable data Y_t is presumed by the econometrician to be related to the model variables by

$$\tilde{Y}_t = FX_t \tag{15}$$

where F is an $m \times n$ matrix and, as in Section 2, $\tilde{Y}_t$ is the model counterpart of Y_t. F and B are assumed to be such that the matrix FB is invertible.[20]

[19] See, for example, Gourieroux et al. (1993), An and Schorfheide (2007), Canova (2007), Fernández-Villaverde (2010).

[20] Note that I have implicitly assumed that the model contains the same number of shocks as observable variables; i.e., $m = r$ in the notation of Section 2. If we had $r > m$, we would be unable to invert the EM's MA(∞) representation, so that Assumption 1 would not hold. On the other hand, if we had $r < m$, some subset of the elements of $\tilde{Y}_t$ would be linearly dependent; i.e., the implied autocovariance function would be singular. Since the actual data is unlikely to possess this feature, the model would be inconsistent with the data in that there would in general be no sequence of shocks that would allow the EM to exactly reproduce the data.

4.2 Data process

Suppose we have a series of observations drawn from an m-variate process $\{Y_t\}$. I make the following assumption about $\{Y_t\}$.

Assumption 3. $\{Y_t\}$ is a mean-zero jointly normally distributed covariance-stationary process with $Var(Y_{i,t}) > 0$ for all $i = 1, \ldots, m$.

Note that we have not assumed that the TM for Y_t is given by the EM (14)-(15). Rather, as argued in Section 3.3, we may write

$$Y_t = \tilde{Y}_t^* + Z_t \tag{16}$$

where

$$\tilde{Y}_t^* = FX_t^* \tag{17}$$

$$X_t^* = AX_{t-1}^* + B\varepsilon_t^* \tag{18}$$

Here, ε_t^* and Z_t have the same interpretation as in Section 3.3, and in particular $\mathbb{E}\left[\varepsilon_t^* Z_{t-k}'\right] = \mathbb{E}\left[\tilde{Y}_t^* Z_{t-k}'\right] = 0$ for all $k \in \mathbb{Z}$.

4.3 Theoretical results

This section formally presents the theoretical results underpinning the methodology discussed in Section 2. Let $C \equiv \left[I_n - B(FB)^{-1}F\right]A$. I make the following assumption regarding the eigenvalues of C.

Assumption 4. None of the eigenvalues of C lie on the complex unit circle.

Assumption 4 is related to the condition under which, if the model is the true DGP, then Y_t has a "structural VAR" representation (i.e., a VAR representation with innovations given by $FB\varepsilon_t$). A structural VAR representation is inherently a backward-looking process: Y_t is a function only of lags of itself and of the current structural shocks. For this representation to exist, intuitively, an infinite history of past values of Y_t must be sufficient to recover the previous period's latent state vector, X_{t-1}, so that the only "new" information contained in Y_t is the vector of current structural shocks. It can be shown that, for this to be true, all of the eigenvalues of C must lie strictly inside the complex unit circle.[21] In contrast, while we similarly require that the latent state vector be recoverable (so that we may in turn recover the structural shocks), we have no need to construct a backward-looking structural VAR representation. As such, there is no need for the state vector to be recoverable on the basis of past values of Y_t only; it will be sufficient that some combination of past, present *and* future

[21] See, for example, Fernández-Villaverde et al. (2007).

values of Y_t reveal X_{t-1}, a property which is guaranteed by Assumption 4.[22] This intuition is established formally in the following proposition.

Proposition 1. *Suppose Assumptions 3-4 hold. Then there exists a sequence of absolutely summable $m \times m$ matrices $\{\psi_j\}_{j=-\infty}^{\infty}$ such that*

$$\varepsilon_t^* = \sum_{j=-\infty}^{\infty} \psi_j L^j \left(Y_t - Z_t\right) \tag{19}$$

where L is the lag operator and ψ_j is a function of A, B, F and j only. Further,

$$\left(\sum_{j=-\infty}^{\infty} \psi_j L^j\right)\left(\sum_{i=0}^{\infty} FA^iBL^i\right) = I_m = \left(\sum_{i=0}^{\infty} FA^iBL^i\right)\left(\sum_{j=-\infty}^{\infty} \psi_j L^j\right) \tag{20}$$

Proof. All proofs are given in Appendix B.

□

Proposition 1 establishes conditions under which, given infinite sequences of observations $Y^\infty \equiv \{Y_t\}_{t=-\infty}^{\infty}$ and $Z^\infty \equiv \{Z_t\}_{t=-\infty}^{\infty}$, the true value of the structural shock, ε_t^*, may be recovered. The next proposition verifies that Assumptions 3 and 4 together imply Assumption 2.

Proposition 2. *Suppose Assumptions 3-4 hold. Then*

$$\hat{\varepsilon}_t = \varepsilon_t^* + \nu_t \tag{21}$$

where $\mathbb{E}\left[\varepsilon_t^* \nu_{t-k}'\right] = 0$ for all $k \in \mathbb{N}$.

Proposition 2 orthogonally decomposes a smoothed EM shock into its true value and a residual. This decomposition result has several useful implications, summed up in the following corollary.

Corollary 1. *Suppose Assumptions 3-4 hold. Then the following are true:*

(a) $\hat{\sigma}_l^2 \equiv \mathbb{E}\left[\hat{\varepsilon}_{l,t}^2\right] \geq \sigma_l^{*2}$ for all l.

(b) If $\mathbb{E}\left[\hat{\varepsilon}_{l,t}\hat{\varepsilon}_{i,s}\right] \neq 0$ and $(l,t) \neq (i,s)$, then the model is misspecified.

(c) If the model is misspecified, then there exists an l such that $\hat{\sigma}_l^2 > \sigma_l^{*2}$.

(d) For $(l,t) \neq (i,s)$, $\mathbb{E}\left[\varepsilon_{l,t}^* \hat{\varepsilon}_{i,s}\right] = 0$ irrespective of whether the model is misspecified.

[22] The relaxation of the structural VAR assumption ($|eig\,(C)| < 1$) to Assumption 4 ($|eig\,(C)| \neq 1$) is an economically relevant one. As pointed out by, for example, Leeper et al. (2011) and Sims (2012), the requirement that all eigenvalues of C lie strictly inside the complex unit circle often excludes models where agents' information sets are strictly greater than the econometrician's, such as in models with "news" and "noise" shocks (see, for example, Beaudry and Portier (2004), Schmitt-Grohé and Uribe (2012), Jaimovich and Rebelo (2009), Blanchard et al. (2013), Christiano et al. (2010a)). In contrast, Assumption 4 will not generally exclude this important class of models.

Part (a) of Corollary 1 points out that the variance of smoothed shock l is an upper bound for the true variance, which is itself unobservable. Parts (b) and (c) highlight that if the smoothed shocks exhibit correlation with one another, then the model is certainly misspecified, and thus the variance of at least one of the shocks will be overstated.[23]

Part (d) forms the basis for the correction procedure which is the principal methodological contribution of this paper. In particular, for any $q \in \mathbb{N}$, we may write

$$\hat{\varepsilon}_{l,t} = \sum_{i \neq l} \Theta_{i,0} \hat{\varepsilon}_{i,t} + \sum_{j=1}^{q} \left(\Theta_j \hat{\varepsilon}_{t-j} + \Theta_{-j} \hat{\varepsilon}_{t+j} \right) + \xi_{l,t} \tag{22}$$

where $\mathbb{E}\left[\xi_{l,t} \hat{\varepsilon}_{i,t+j}\right] = 0$ for any $(l,t) \neq (i, t+j)$ with $|j| \leq q$. Here, the Θ's have the interpretation of population regression coefficients, and $\xi_{l,t}$ as the corresponding OLS residual. $\xi_{l,t}$ captures the component of $\hat{\varepsilon}_{l,t}$ which cannot be predicted using q past and future values of the recovered shocks (including the current values of shocks $i \neq l$). The following proposition establishes how a more accurate upper bound for σ_l^{*2} may be obtained from the regression equation (22).

Proposition 3. *Suppose Assumptions 3-4 hold, and let $\xi_{l,t}$ be as in equation (22). Then*

$$\sigma_l^{*2} \leq \bar{\sigma}_l^2 \equiv \mathbb{E}\left[\xi_{l,t}^2\right] \leq \hat{\sigma}_l^2 \tag{23}$$

Proposition 3 establishes that $\bar{\sigma}_l^2$ is an upper bound for the unknown value of σ_l^{*2}, and further that this upper bound is a (weakly) better estimator of σ_l^{*2} than $\hat{\sigma}_l^2$.

4.4 Practical considerations

Assumption 4 and the availability of infinite past and future observations on Y_t play an important role in the theoretical results of Section 4.3. In particular, in the case where the EM (14)-(15) is correctly specified, these properties together guarantee that the shocks may be exactly recovered as a linear combination of the observations on Y_t. This is crucial for the orthogonal decomposition of Proposition 2, and, since it relies on this result, also for the variance correction suggested at the end of Section 4.3.

In practice, however, Assumption 4 may fail to hold in important cases of interest (the application of Section 6 explores such an instance), and data sets are usually limited to at most a few hundred periods of observations. In either of these cases, the shocks may be recovered in expectation only, with a generally non-zero and non-diagonal error covariance matrix even when the model is correctly specified. Nonetheless, given the linear-Gaussian structure of the EM and TM, the expected values of the shocks conditional on the model and observed data (i.e., the smoothed shocks) will be linear combinations of the sequence of observations on Y_t. Thus, as in (3), we may express a smoothed shock as a linear filter applied

[23] Note that the converse does not necessarily hold, i.e., if $\mathbb{E}\left[\hat{\varepsilon}_{l,t} \hat{\varepsilon}_{i,s}\right] = 0$ for all $(l,t) \neq (i,s)$, this does not necessarily imply that the model is correctly specified, so that the naive estimates $\hat{\sigma}_l^2$ may still overstate the importance of one or more shocks.

to the time series of data.[24] As such, using a similar argument as in the proof of Proposition 2, we may generalize the decomposition result as

$$\hat{\varepsilon}_t = \hat{\varepsilon}_t^* + \nu_t$$

where $\mathbb{E}\left[\hat{\varepsilon}_t^* \nu_{t-k}'\right] = 0$, $\hat{\varepsilon}_t^*$ is defined to be the value of the smoothed shock that would obtain if the model were correctly specified, and ν_t is the linear filter applied only to the sequence of Z_t's (analogous to the second term on the right-hand side of equation 13). In general, the errors $\delta_t^* \equiv \hat{\varepsilon}_t^* - \varepsilon_t^*$ that would be made in recovering the shocks if the model were correctly specified will be correlated across different shocks. As a result, the actual smoothed shocks, $\hat{\varepsilon}_t$, will in general exhibit non-zero correlation with one another even if the model is correctly specified. Whether the robust variance estimates proposed in Section 3 will continue to be an improvement on the naive ones under these circumstances depends on the quality of the approximation $\hat{\varepsilon}_t^* \approx \varepsilon_t^*$. While no analytical result is available to check the accuracy of this approximation, given that the researcher has in hand the fully specified model, it is nonetheless straightforward to check it numerically via simulation, as illustrated in the application of Section 6. Similarly, critical values for any relevant statistics may also be obtained via simulation.

5 Further examples

In this section, I present two additional fully-specified examples to illustrate the above methodology. In the first example, which is simple enough to be solved analytically, I highlight a type of misspecification that is conceptually distinct from the type present in the example of Section 2.1. In that earlier example, the EM correctly represents the dynamic impacts of the shocks it contains (i.e., the IRFs in the EM also exist in the TM), but is missing an additional source of exogenous variation. In contrast, in the first example presented in this section the economic interpretations of a shock in the EM and another in the TM are the same, but their IRFs are not, and so by the framework developed in Section 3 they are nonetheless considered fundamentally different.

Next, in the second example of this section, the TM is taken to be a variant of a standard medium-scale New Keynesian model (that of Smets and Wouters (2007)) that features a variety of real and nominal frictions and seven exogenous shock processes. The EM, on the other hand, will in general be missing several frictions and only contain three of the seven shock processes.[25] This example most closely captures quantitative macroeconomic modelling in practice, whereby economists use highly simplified models to fit data generated by substantially more complicated ones.

[24] Note that, when the data set is finite, the filter for $\hat{\varepsilon}_t$ will generally depend on t.
[25] Note that both types of misspecification highlighted above will be present in this example.

5.1 Example 2: Univariate AR(1) model

Suppose we have a sequence of data on TFP growth, γ_t, for which the TM is an AR(1) process

$$\gamma_t = \rho\gamma_{t-1} + \varepsilon^*_{1,t}, \quad \varepsilon^*_{1,t} \sim \text{i.i.d.}\, N\left(0, \sigma^*_1\right)$$

where $0 < |\rho| < 1$. The EM, meanwhile, is given by

$$\tilde{\gamma}_t = \delta\tilde{\gamma}_{t-1} + \tilde{\varepsilon}_{2,t}, \quad \tilde{\varepsilon}_{2,t} \sim \text{i.i.d.}\, N\left(0, \tilde{\sigma}_2\right)$$

with $|\delta| < 1$, where δ is some number taken as given by the econometrician and that may potentially be different from ρ, in which case the EM would be misspecified. This EM can be easily inverted to obtain $\tilde{\varepsilon}_{2,t} = \tilde{\gamma}_t - \delta\tilde{\gamma}_{t-1}$. Substituting the data process γ_t into the inverted EM for $\tilde{\gamma}_t$, we may obtain the process for the smoothed shocks,

$$\hat{\varepsilon}_{2,t} = (\rho - \delta)\,\gamma_{t-1} + \varepsilon^*_{1,t}$$

and thus the naive estimate of the variance of the TFP shock is given by

$$\hat{\sigma}^2_2 = \frac{1 - 2\rho\delta + \delta^2}{1 - \rho^2}\sigma^{*2}_1$$

When $\delta \neq \rho$, the IRF in the EM does not exist in the TM, and thus by the framework of Section 3 we have $\sigma^{*2}_2 = 0$, which is clearly less than the naive estimate $\hat{\sigma}^2_2$. The conclusion that $\sigma^{*2}_2 = 0$ is rather stark here, since it will hold for all $\delta \neq \rho$ even as $\delta \to \rho$, while if $\delta = \rho$ we have $\sigma^{*2}_2 = \sigma^{*2}_1$. However, this discontinuous behavior is of little practical importance. In particular, misspecification is identified in this framework as non-zero autocovariance in the sequence of smoothed shocks. For $k \neq 0$ we may obtain

$$Cov\left(\hat{\varepsilon}_{2,t}, \hat{\varepsilon}_{2,t-k}\right) = \frac{(\rho - \delta)(1 - \rho\delta)}{1 - \rho^2}\rho^{|k|-1}\sigma^{*2}_1$$

Clearly, $Cov\left(\hat{\varepsilon}_{2,t}, \hat{\varepsilon}_{2,t-k}\right) \to 0$ as $\delta \to \rho$. Thus, even though the EM continues to be misspecified as δ becomes arbitrarily close to ρ (i.e., as the IRF in the EM becomes ever closer to the IRF in the TM), the degree of *observed* misspecification (and thus any adjustment made using the bias-correction procedure) will nonetheless shrink to zero.

Next, to apply the bias-correction procedure, we wish to find the component of $\hat{\varepsilon}_t$ that cannot be predicted using past and future values. In general, one can continue to improve the quality of the adjusted estimate by including a greater number of leads and lags in the regression (i.e., by making q arbitrarily large in equation (22)). For simplicity, I focus here on the simple case where $q = 1$. In this case, the estimated regression equation is given by

$$\hat{\varepsilon}_{2,t} = \beta\left(\hat{\varepsilon}_{2,t-1} + \hat{\varepsilon}_{2,t+1}\right) + \xi_t$$

where

$$\beta \equiv \frac{(\rho-\delta)(1-\rho\delta)(1-\rho^2)[1-\delta(\rho-\delta)]}{(1-2\rho\delta+\delta^2)^2-\rho^2(\rho-\delta)^2(1-\rho\delta)^2}$$

and ξ_t is the OLS residual. The bias-corrected estimate of the shock variance is then given by $\bar{\sigma}_2^2 \equiv Var(\xi_t)$. Some simple algebra shows that

$$\bar{\sigma}_2^2 = \hat{\sigma}_2^2 - 2\beta Cov(\hat{\varepsilon}_{2,t}, \hat{\varepsilon}_{2,t-1})$$

It can be verified that $\beta Cov(\hat{\varepsilon}_{2,t}, \hat{\varepsilon}_{2,t-1}) \geq 0$ and therefore $0 = \sigma_2^{*2} < \bar{\sigma}_2^2 \leq \hat{\sigma}_2^2$. Thus, the corrected estimate lies between the true value and the naive estimate $\hat{\sigma}_2$. Furthermore, as expected given the discussion above, we have $\beta Cov(\hat{\varepsilon}_{2,t}, \hat{\varepsilon}_{2,t-1}) \to 0$ as $\delta \to \rho$, so that the size of the adjustment approaches zero as the IRF in the EM approaches the IRF in the TM. That is, notwithstanding the fact that the true variance is zero for all $\delta \neq \rho$, as the degree of misspecification approaches zero the bias-corrected estimate of the variance will approach the desired value of σ_1^{*2}.

5.2 Example 3: Medium-scale New Keynesian model

In this example, I consider three different combinations of TM and EM. In all cases, the models will be variants of the widely-cited medium-scale New Keynesian model of Smets and Wouters (2007), which in its baseline form features a number of real and nominal frictions, as well as seven different exogenous shock process: three "real" shocks (exogenous spending, neutral TFP, and investment-specific technology (IST)), and four shocks that have no first-order effects in the absence of sticky prices and wages (risk premium, price mark-up, wage mark-up, monetary policy).[26] In particular, I consider the following combinations:

- **Case A:** The EM and TM both feature sticky prices and wages and only the three real shocks.
- **Case B:** The EM and TM both feature sticky prices and wages. The EM contains only the three real shocks while the TM contains the full complement of seven shocks.
- **Case C**: The EM and TM both contain only the three real shocks. The EM features flexible prices and wages, while the TM features sticky prices and wages.

Note that, while the EM in Case A is correctly specified, Cases B and C correspond to two distinct types of misspecification. As in Example 1, the EM in Case B is missing important sources of exogenous variation due to the absence of the risk premium, mark-up and monetary policy shocks. Meanwhile, as in Example 2, in Case C even though the shocks in the EM have the same economic interpretation as shocks in the TM, their IRFs will be different since the EM does not feature sticky prices or wages. As such, these shocks are fundamentally different, and thus the true variance of each EM shock is zero. As we shall see, however, the

[26] A summary of the equations characterizing the DGP and model is contained in Appendix A. For further details, see Smets and Wouters (2007).

"less misspecified" is the EM—that is, the more flexible are prices and wages in the TM—the smaller the degree of observed misspecification will be, with no observed misspecification in the limit as prices and wages become perfectly flexible. In turn, as in Example 2, the size of any bias adjustment will also shrink to zero.

I assume the econometrician estimates the variances of the EM shocks using data on three real variables: consumption, investment, and hours worked. For reasons discussed in Section 4.1, I assume that, aside possibly from the true degree of price and wage flexibility, the econometrician knows all of the parameter values of the TM except for the variances of the shock process innovations, which are then estimated using a large data set on consumption, investment and hours worked, themselves generated from the TM.[27,28] The first column of data in Table 1 reports, for each of the three shock processes in the EM, the variance of the shock with that same economic interpretation in the TM.[29] The remaining three columns report naive variance estimates for the three different cases.[30]

For Case A, the EM is correctly specified, and as a result the naive variance estimates reported in the table are the same as the true values reported in the first column.[31] For Case B, however, we have the first type of misspecification discussed above: the IRFs in the EM are the same as their TM counterparts, but the TM also includes other sources of exogenous variation. As we might have expected from (4), the shock variances are significantly overestimated in this case, ranging from over 13 times the true value for the investment shock, to well over 800 times for the exogenous spending and TFP shocks.

Next, I apply the bias-correction methodology to obtain revised estimates of the variances for Cases A and B. Specifically, I regress each of the three smoothed shocks on contemporaneous values of the other two smoothed shocks and on four leads and lags of all three shocks. I then compute the variance of the resulting residual. The results are presented in Table 2. For comparison purposes, the first column of data again reports the variance of the shock with that economic interpretation in the TM. Looking at the second column, the corrected estimates for Case A show no change relative to the naive case from Table 1. Intuitively, since in this case the model is the TM, the true sequence of shocks can be recovered (nearly) exactly. Since the true shocks are, by construction, uncorrelated with one another, so too are the smoothed shocks. Thus, none of the variation in a smoothed shock can be explained by variation in the other shocks, and we obtain an R^2 of essentially zero in the regression.

[27] Except where otherwise noted, the model parameters used for generating the data are taken as the mode of the posterior distribution reported in Smets and Wouters (2007).

[28] In practice, I simulate 101,000 periods of data, discarding the first 1,000 to minimize the impact of the choice of the initial state vector—set equal to zero—on the results. I use this large number of data periods to avoid the complications (discussed in Section 4.4) that arise when samples are of a more realistic size.

[29] Note that these variances are the same for each of the three cases.

[30] Naive estimates are obtained by first running the Kalman smoothing algorithm (with shock variance parameters set equal to values from TM), then computing the sample variances of the resulting smoothed shocks. While these estimates will in general differ from estimates obtained via maximum likelihood, the two are asymptotically equivalent. Given the very large data set, it can be verified that the quantitative difference between the two approaches is negligible.

[31] Technically, there are small discrepancies between the two sets of estimates due to (1) sampling error; and (2) the fact that, in a finite sample, the shocks cannot be exactly recovered. However, because of the large sample size employed here, such discrepancies are quantitatively unimportant.

Table 1: Variances of True and Smoothed Shocks

TM		3 shocks, sticky	7 shocks, sticky	3 shocks, sticky
EM		3 shocks, sticky	3 shocks, sticky	3 shock, flex
Case	**True**	**A**	**B**	**C**
Shock:				
Spending_t	0.27	0.27	221.52	2.18
TFP_t	0.21	0.21	179.27	0.61
IST_t	0.20	0.20	2.73	0.21

Notes: "True" column shows variance of shock in TM. Remaining columns show naive variance estimates for relevant TM and EM, obtained by first running Kalman smoothing algorithm (with shock variances set equal to values from TM), then computing sample variances of resulting smoothed shocks.

Table 2: Corrected Variances of Smoothed Shocks

DGP		3 shocks, sticky	7 shocks, sticky	3 shocks, sticky
Model		3 shocks, sticky	3 shocks, sticky	3 shocks, flex
Case	**True**	**A**	**B**	**C**
Shock:				
Spending_t	0.27	0.27	0.55	0.01
TFP_t	0.21	0.21	0.44	0.00
IST_t	0.20	0.20	0.23	0.01

Notes: "True" column shows variance of shock in TM. Remaining columns show bias-corrected estimates. For a given smoothed shock, set of regressors is contemporaneous values of other smoothed shocks plus four leads and lags of all smoothed shocks.

For Case B, on the other hand, the corrected estimates in Table 2 show a substantial improvement over the naive estimates from Table 1. For example, whereas the naive estimate of the neutral TFP variance was 875 times the true value, the corrected estimate is only a little more than twice the true value. The relative discrepancies are even smaller for the other two shocks. A clue for why the corrected estimates show such a large improvement in Case B can be seen from Table 3, which presents contemporaneous cross-correlations between the smoothed shocks (first three rows), as well as the first-order autocorrelation for each shock (final row). Clearly, the smoothed shocks exhibit a very high degree of contemporaneous correlation with one another, especially the exogenous spending and neutral TFP shocks, which are nearly collinear. The spending and TFP shocks also exhibit significantly negative first-order autocorrelation. Because there is such a high degree of correlation between the smoothed shocks, they are highly predictable using values of the other smoothed shocks, which in turn results in the regression residual being small.

Table 3: Sample Correlations for Smoothed Shocks (Case B)

	Spending_t	TFP_t	IST_t
Spending_t	-	0.997	-0.800
TFP_t	0.997	-	-0.829
IST_t	-0.800	-0.829	-
Lagged variable	-0.505	-0.499	-0.043

Notes: Table presents correlations among smoothed shocks for Case B. Upper three rows of data show unconditional correlation matrix. Final row shows first-order autocorrelation for that smoothed shock.

This high degree of correlation in the smoothed shocks stems from a number of features of the TM, one example of which I focus on here for illustrative purposes. Consider the risk premium shock. In the TM this shock accounts for over 80 percent of the variance of the one-step-ahead forecast error in consumption, while in the EM it is entirely absent. Thus, the EM is clearly missing an important factor in generating high-frequency variation in consumption. This is confirmed in Figure 1, which shows autocovariance functions (ACFs) for the three observable variables using the TM. Each panel of the figure corresponds to $Cov\left(y_t, x_{t-k}\right)$ for two variables y_t and x_t. Within each panel there are four lines plotted, each corresponding to the ACF that arises in the TM when all but one shock is shut down. The four plots correspond to the three real shocks, plus the risk premium shock. The first five lags are shown, and to make the comparison as clean as possible, the variances of the shocks are normalized so that the unconditional variance of consumption is always equal to one.

Looking at the top-left panel of the Figure, we see that the autocovariance generated in the consumption process by the risk premium shock (the solid line) shows a relatively steep decline. In contrast, the spending, TFP and IST shocks (the dashed, dotted, and dash-dot lines, respectively) all generate highly persistent consumption dynamics. In order for the EM to be able to reproduce the high-frequency consumption variation that is a feature of the TM, it would therefore need to be the case that certain combinations of shocks tend to occur together or in a particular sequence (or both). Since the smoothing algorithm is precisely the process of finding a sequence of shocks that allow the EM to exactly reproduce the data, it should then come as no surprise that we observe correlations between them.

Next, consider Case C. In this case we have the second type of misspecification discussed above: the shocks in the TM and the EM have the same economic interpretation, but because the EM does not feature sticky prices and wages, the IRFs in the model are different from those in the TM. As such, the true variances of the EM shocks are actually zero. As seen in Table 1, however, using naive estimation procedures we obtain non-zero estimates.[32]

Turning to the corrected estimates in Table 2, we find that they are quite close to their true value of zero. As in Example 2, however, there is some sense that this result is unde-

[32] Note that, while it turns out the estimates for Case C are greater than the values reported in the first column, this need not always be true.

Figure 1 – Autocovariance function ($Cov\,(y_t, x_{t-k})$)

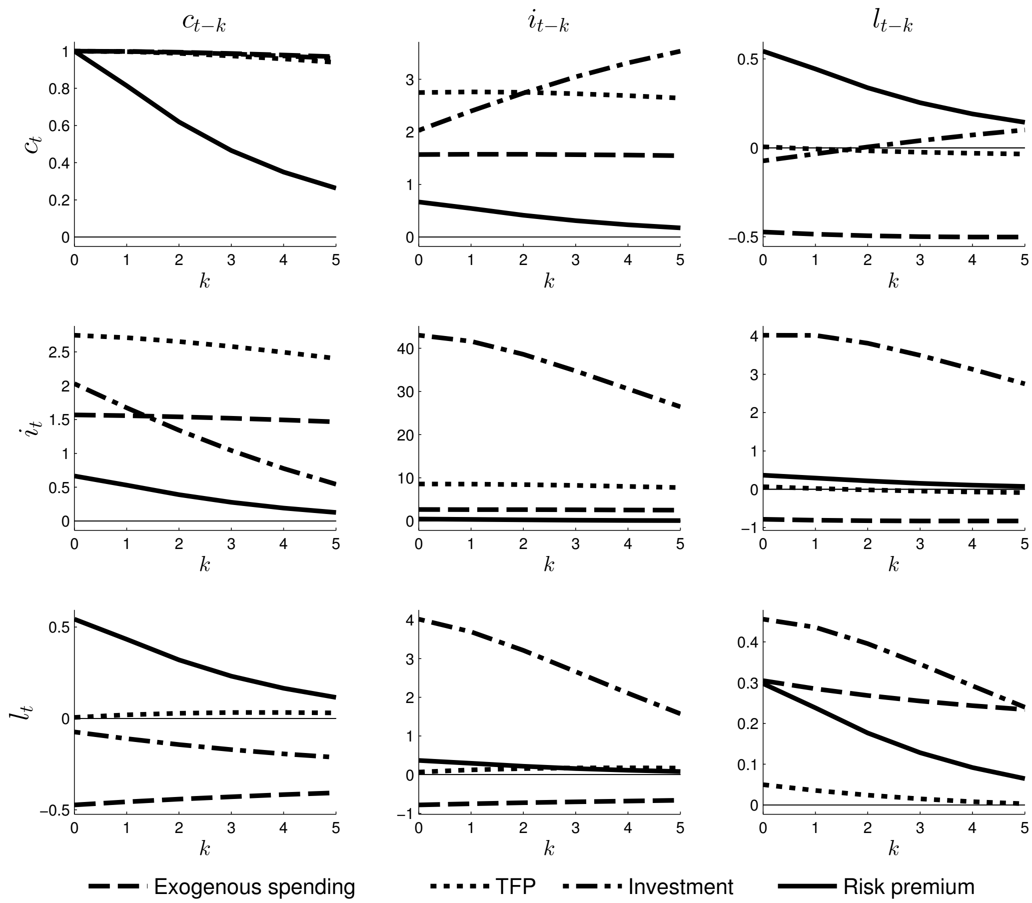

Notes: Figure shows autocovariance functions (ACFs) for observable variables in TM generated by each of four listed shocks individually. ACFs are normalized so that unconditional variance of c_t is equal to one.

sirable, since there are shocks in the TM with the same economic interpretation and that do actually account for a positive amount of variation. But as also illustrated in Example 2, the extremeness of the adjustment in this case is due to the very high degree of misspecification in the EM. In particular, there is a significant amount of price and wage stickiness in the TM: each quarter, only fractions 0.34 and 0.26 of price- and wage-setters, respectively, are allowed to re-set their prices/wages. In the EM, however, these fractions are incorrectly assumed to be equal to one. To see how the corrected estimates change as the degree of stickiness in the TM and EM converge, Table 4 shows corrected variance estimates as the fraction of re-optimizers in the TM approaches the EM case of one (so that the EM is closer to being correctly specified). The first column of data in the table again reports the true variance estimates, while the second reproduces the adjusted levels from Table 2. The remaining columns report the results for three different values of the fraction of re-optimizers—0.95, 0.99 and

0.999—which show that the estimates steadily improve with the quality of the model. Thus, despite the stark fact that the true variances of the EM shocks remain zero even as the EM and TM converge, the observed degree of misspecification nonetheless converges to zero, and the adjusted shock values converge to the values from the TM. Put another way, one need only be concerned about this sort of "over-adjustment" if the EM is a poor approximation to the TM, which is precisely the case in which there *should* be a large adjustment so as to warn the econometrician that misspecification is likely to be a problem.

Table 4: Corrected Variances (Case C)

			Fraction of re-optimizers		
	True	**Baseline**	**.95**	**.99**	**.999**
Shock:					
Spending_t	0.27	0.01	0.19	0.33	0.28
TFP_t	0.21	0.00	0.03	0.11	0.19
IST_t	0.20	0.01	0.17	0.20	0.20

Notes: "True" column shows variance of shock in TM. "Baseline" column reproduces corrected estimates under baseline parameterization of TM from Table 2. Remaining columns show corrected estimates for several different fractions of re-optimizers in TM.

6 Application: Investment shocks

In this section, I apply the methodology introduced in this paper to the recent model of Justiniano et al. (2010) (henceforth JPT), itself a variant of the Smets and Wouters (2007) model discussed in Example 3 in Section (5).[33]

Following the Bayesian estimation procedure used by JPT,[34] I re-estimate their model using quarterly U.S. data on output, consumption, investment, hours, wages, inflation and the nominal interest rate over the period 1954QIV-2004QIV ($T = 201$).[35] Table 5 reports resulting variance decompositions for log-hours using the median parameter values from the posterior parameter distribution.[36] Column (1) shows a standard decomposition of the unconditional variance. As noted by JPT, this decomposition indicates that the wage mark-up shock (typically interpreted as a labor supply shock) accounts for the majority (52 percent) of the unconditional variance of hours in the model. However, the estimated wage mark-up

[33] While there are a number of differences between JPT and Smets and Wouters (2007), the main source of the divergence in their results is in the data used to estimate the models. Specifically, Smets and Wouters (2007) include consumer durables in their measure of consumption and exclude the change in inventories from their measure of investment. JPT, on the other hand, include consumer durables and the change in inventories in their measure of investment, with consumption including only non-durables and services. This produces a more volatile investment series and a less volatile consumption series, which largely drives their results.

[34] See An and Schorfheide (2007) for a review of Bayesian estimation in the context of DSGE models.

[35] See Appendix C for details about data sources. Bayesian estimation of the DSGE model was done using Dynare (see Adjemian et al. (2011)).

[36] The posterior statistics I obtained for the 35 estimated parameters were very close to those reported by JPT and are available upon request.

process is highly persistent (autoregressive parameter of 0.98), suggesting that the bulk of the variation induced by this shock may be of a long-run nature, and thus potentially less important for business cycle variation.

Table 5: Naive Variance Decomposition for Hours

Shock	(1) Unconditional	(2) BCF
Monetary policy	0.03	0.06
Neutral technology	0.04	0.11
Government spending	0.08	0.02
Investment	0.25	0.61
Price mark-up	0.06	0.06
Wage mark-up	0.52	0.05
Patience	0.03	0.08

Notes: Table entries show naive variance decompositions for hours, obtained from JPT's model using the median of the posterior parameter distribution. Column (1) decomposes the unconditional variance. Column (2) decomposes the BCF variance, defined as the variance associated with periodic fluctuations between 6 and 32 quarters. Columns may not add up to 1 due to rounding.

To address this issue, Column (2) of Table 5 reports the decomposition of the business-cycle-frequency (BCF) variance[37] of log-hours in the model. Echoing the results found by JPT, the results indicate that the investment shock, rather than the wage mark-up shock, accounts for the majority (61 percent) of BCF fluctuations in hours. While not shown in the table, the investment shock was also found to account for large proportions of the BCF variance of output growth (56 percent) and investment growth (88 percent). On the basis of this evidence, one may be tempted to conclude, as JPT do, that "investment shocks are the leading source of business cycles." (p. 137)

As JPT also note, however, the investment shock is found to be a negligible determinant of BCF fluctuations in consumption growth, accounting for a mere 6 percent of its variance, while the otherwise-irrelevant household patience shock accounts for 61 percent. Figure 2 plots the correlation between consumption growth and four leads and lags of investment growth in the data (solid line) and in the model (dashed line). The data indicates a positive and significant correlation between consumption growth and investment growth within a one-quarter lead/lag. In contrast, in the EM there is a small negative correlation between the two. This EM correlation is small because consumption and investment are driven primarily by two distinct (and orthogonal) shocks: the patience and investment shocks, respectively. Meanwhile, the correlation is *negative* primarily because a positive investment shock increases the return on investment with no change to current productive capacity, causing the household

[37] BCF variances are obtained by integrating the spectrum of the model over the relevant frequencies (defined here, as in JPT, to be frequencies associated with periods between 6 and 32 quarters). This process effectively removes the variance associated with high- and low-frequency fluctuations, leaving only the medium-frequency fluctuations normally associated with business cycles. See Appendix D for further details.

to substitute away from consumption and toward investment.[38]

Figure 2 – Correlation between $\%\Delta C_t$ and $\%\Delta I_{t+k}$

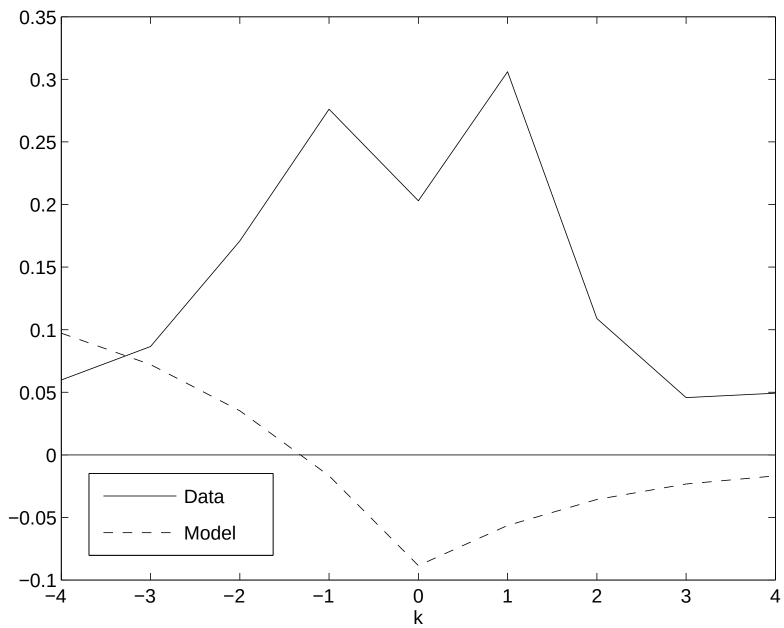

To gain further insight, I proceed with analysis of the smoothed shocks. It can be verified in the case of JPT's EM that the matrix C (as defined in Section 4.3) contains several eigenvalues equal to one, so that Assumption 4 fails to hold.[39] To check that the approximation $\hat{\varepsilon}_t^* \approx \varepsilon_t^*$ holds (see Section 4.4) in spite of the finite sample size and the violation of Assumption 4, I simulated data from the EM,[40] then obtained the smoothed shocks from this simulated data. From this sequence of smoothed shocks, I then computed the mean squared recovery error (MSRE) for each shock in the EM (i.e., $Var(\delta_{l,t}^*)$ in the notation of Section 4.4) at each date t, then scaled the result by the (known) variance of that shock, $\hat{\sigma}_l^2$. Figure 3 plots the maximum of this statistic across the seven different shocks for each the first ten periods.[41] As the figure shows, the smoothed shocks very quickly attain a high degree of accuracy, with the maximum

[38] The failure of many RBC-type models to generate positive comovement between consumption, investment and hours in response to shocks that affect the expected return to investment without affecting current production technology has been known since at least Barro and King (1984); see also Beaudry and Portier (2007).

[39] Because JPT's model features a stochastic trend in productivity, the data contains first-differences of several non-stationary variables, while the model state vector contains the corresponding variables in (stationarized) levels. Since a variable cannot be exactly recovered from even an inifinite history of changes in that variable, and since the condition that C contain no eigenvalues of modulus one is precisely that needed to guarantee exact recovery of the state vector, it should be unsurprising that C contains eigenvalues equal to one for JPT's model.

[40] Simulated values throughout this section are based on N simulated data sets, each of T periods in length, where $T = 201$ is the length of the actual data set.

[41] The data point for $t = 1$ is off the scale in Figure 3, with a maximum MSRE equal to 80 percent of the variance of the shock.

MSRE equal to less than 2 percent of the variance of the shock by $t = 4$, and less than 1 percent by $t = 5$. Though not shown in the figure, accuracy continues to increase, with the maximum MSRE falling below 0.5 percent by $t = 17$ and remaining below this threshold for the remainder of the sample. Thus, beyond the first few periods, the approximation $\hat{\varepsilon}_t^* \approx \varepsilon_t^*$ appears to be sufficient for the principal results of Section 4.3 to hold. To ensure that none of the results are driven by the relatively inaccurate recovery of the early shock values, I drop the first four periods of smoothed shocks in all subsequent computation.

Figure 3 – Maximum scaled MSRE for smoothed shocks

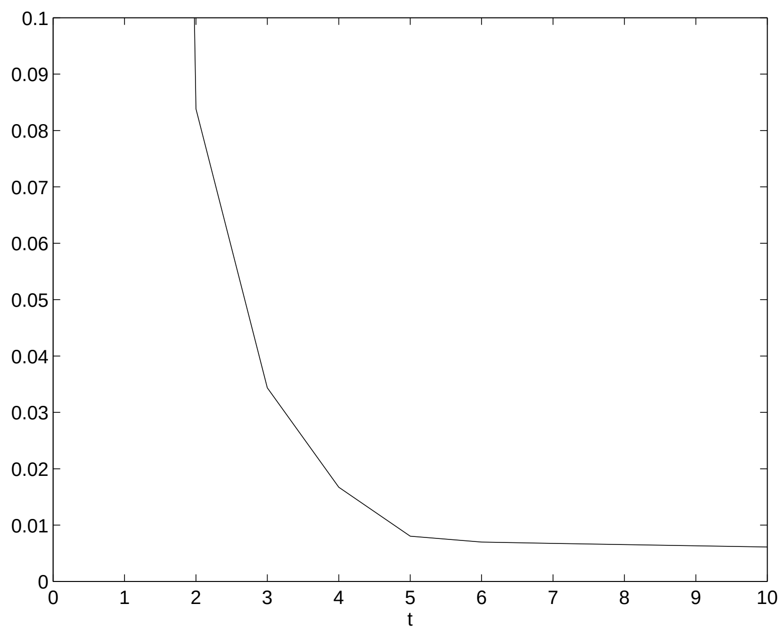

Notes: Figure displays $\max_l Var\left(\delta_{l,t}^*\right)/\sigma_l^*$, computed by simulating 10,000 data sets. Data point for $t = 1$ is off the scale in Figure 3, with a maximum MSRE equal to 80 percent of the variance of the shock.

Next, returning to the sequence of smoothed shocks obtained from the actual data, Figure 4 plots sample correlations between the smoothed investment shock at date t on the one hand, and each of the smoothed shocks at a number of leads and lags on the other. Black dots represent sample correlations, while the dashed lines show pointwise simulated 95 percent confidence intervals (based on 100,000 draws) for these correlations under the null hypothesis that the EM is correctly specified. Circles around dots highlight correlations that fall outside of the confidence intervals.

Of the 71 points shown in Figure 4, 15 (21 percent) are outside of the 95 percent confidence bands. If the EM were correctly specified, this would be an extreme result. Figure 5 presents a simulated probability mass function for the number of correlations between the investment shock and other shocks within 5 leads or lags that fall outside of the 95 percent confidence bands for the case when the EM is correctly specified. The median and mode of

Figure 4 – Correlation coefficients between smoothed shocks

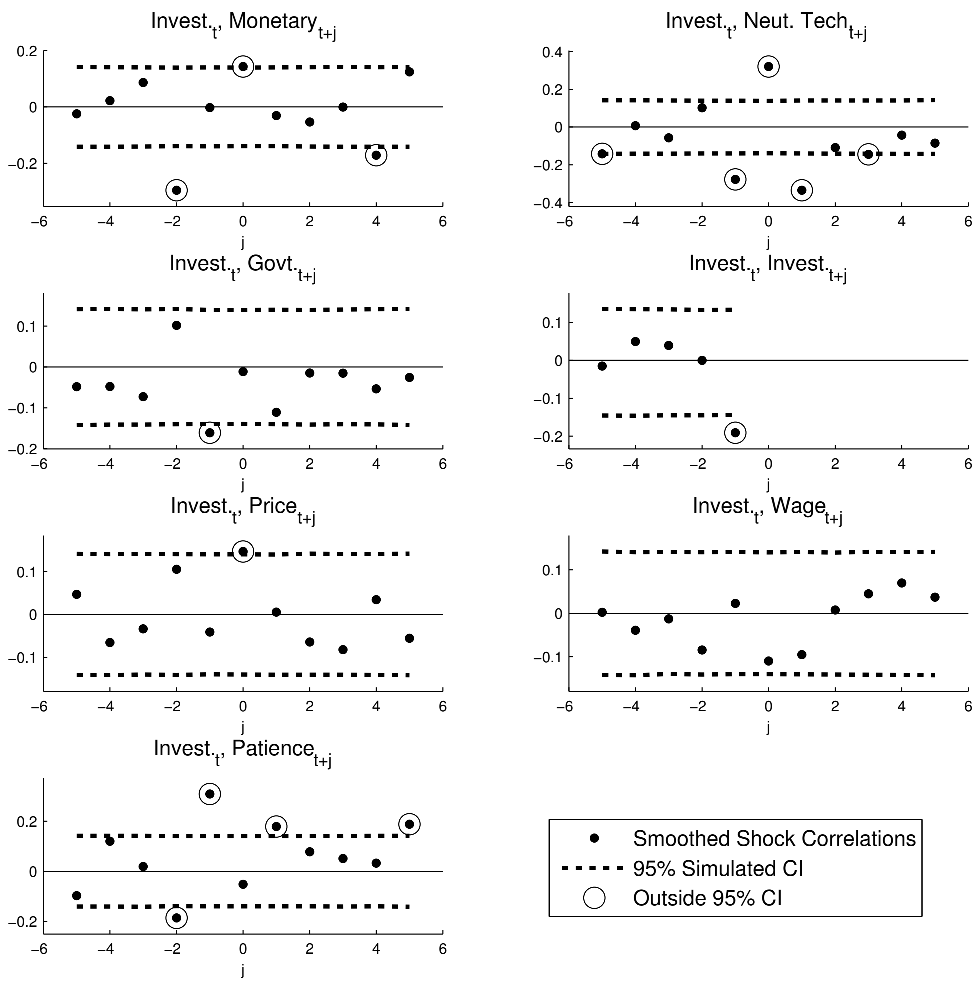

Notes: Dots indicate the correlation coefficient between the smoothed investment shock and the specified lead or lag of another smoothed shock. Confidence intervals were computed from 100,000 simulated data sets under the null hypothesis that the EM is correctly specified.

this distribution are both 3 and the mean 3.5. The maximum number of correlations outside of the confidence bands in 100,000 simulations was 14, a figure attained by only 2 (0.002 percent) of those simulations. 97 percent of the time, the number of such correlations was less than or equal to 7. In this context, 15 such correlations is clearly well outside what could be considered reasonable if the model were correctly specified. While I focus on the investment shock in Figure 4 because of the important role ascribed to it by JPT, the extreme number and degree of correlations are not limited exclusively to this shock. For example, the neutral technology and patience shocks each had 16 of 71 points outside of the 95 percent confidence bands, while the monetary policy shock had 15.

Note also that not only was the frequency of significant correlations in Figure 4 high, but

Figure 5 – Proportion of correlations outside 95% CI

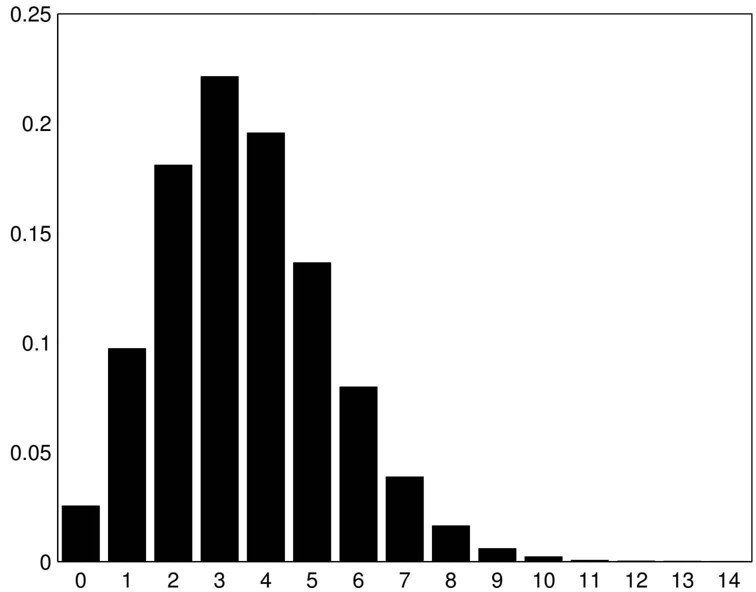

Notes: Simulated probability mass function (based on 100,000 simulations of same number of periods as actual data set) for number of correlations between investment shock and other shocks (within 5 leads or lags) that fall outside of the simulated 95 percent confidence interval under null hypothesis that EM is correctly specified.

the *degree* to which the correlations exceeded the bounds was in some cases also extreme. Most notably, the monetary policy shock at a two-quarter lead, the neutral technology shock contemporaneously and at a one-quarter lead and lag, and the household patience shock at a one-quarter lead all exhibit correlations with the investment shock that exceeded 0.27 in absolute value, well outside the 95 percent confidence bounds which were all less than 0.15 in absolute value.

The sample correlation evident in Figure 4 between the smoothed investment and patience shocks is instructive. In particular, it was noted above that, in the estimated model, investment was found to be driven by the former shock while consumption was driven by the latter, despite the fact that the model was estimated using data which suggests that investment and consumption are correlated. This apparent puzzle is easily resolved if the two smoothed shocks themselves exhibit sample correlation, a property which Figure 4 clearly establishes. Intuitively, the estimation algorithm is “mixing” the investment and patience shocks together in a systematic way in order to reproduce the autocovariance patterns in the data. On the other hand, because of the orthogonality assumption, all moments derived from the EM, including those implicit in the computation of variance decompositions, are obtained assuming no such mixing is taking place, leading to the apparent conflict between model and data.

Next, as argued in Section 3, the correlations between the smoothed shocks are indicative of model misspecification and, as a result, the naive estimates of the variances of the shocks will be overstated. I turn now to obtaining bias-corrected estimates. In order to do so, a central issue that must be addressed is how to obtain estimates of the regression coefficients (i.e., the Θ's in equation (22)). Given the relatively small sample size, using q leads and lags of all the smoothed shocks (plus the other contemporaneous shocks) as regressors very quickly uses up degrees of freedom.[42] However, as Figure 4 suggests for the case of the investment shock, only a small subset of these potential regressors will carry significant explanatory power, so that over-fitting is a serious concern. To address this issue, I use a simple selection rule, choosing only those regressors which meet the following criteria: (1) are within 5 leads or lags of the dependent variable, and (2) exhibit correlation coefficients with the dependent variable which are outside of a simulated $100\,(1-\alpha)\%$ confidence interval. This rule has the advantage of being both simple to apply and relatively flexible, in that the consequences of different values of α can be explored.

For a given value of α and this selection rule, I next compute the corrected estimate of the variance of shock l, $\bar{\sigma}^2_{l,\alpha}$, as the unbiased estimator of the variance of the regression residuals from equation (22). That is, $\bar{\sigma}^2_{l,\alpha} = (T_{l,\alpha} - k_{l,\alpha})^{-1} \sum \hat{\xi}^2_{l,t}$, where $\hat{\xi}_{l,t}$ is the fitted residual, $T_{l,\alpha}$ is the number of observation periods, and $k_{l,\alpha}$ is the number of regression parameters. Table 6 presents results for several different values of α. For each value of α, the table reports the corrected estimate, $\bar{\sigma}^2_{l,\alpha}$, as a fraction of the naive estimate, $\hat{\sigma}^2_l$. The smaller this fraction is, the more of the variance of the smoothed shock we may attribute to misspecification. For comparison purposes, the bottom row of the table reports simulated 90 percent lower bounds for the corresponding statistic obtained assuming the model is the true DGP.

Several things emerge from the results in Table 6. First, with the exception of the wage mark-up shock, the corrected variance estimates are well below the simulated lower bounds. Given the theoretical considerations of Section 3, this suggests that misspecification may be an important factor for this model. Second, for the most conservative case ($\alpha = 0.001$), the point estimates suggest that nearly one-third of the variances of the smoothed neutral technology and investment shocks are attributable to misspecification, while over 90 percent of the time this statistic would be zero if the model were the true DGP. Similarly, over one-fifth of the variance of the smoothed patience shock is attributable to misspecification for this level of α. Third, for $\alpha = 0.005$, except again for the wage mark-up shock, the point estimates indicate that over one-fifth of the variance of every smoothed shock is attributable to misspecification, and over one-quarter when $\alpha = 0.02$. Again, these figures are suggestive of an important degree of misspecification. It should also be noted again that the corrected variance estimates are upper bounds for the true value. Thus, the actual degree of misspecification may indeed be significantly higher than these figures suggest.

To check the implications for the hours variance decomposition, Table 7 reports the proportion of the variance of hours attributable to each of the shocks assuming the shock variances are the corrected estimates computed above.[43] Again, 90 percent simulated lower bounds

[42] When all potential regressors within q leads and lags are used, the number of estimated parameters will be $6 + 14q$, while the number of available observations is $197 - 2q$. The number of degrees of freedom is thus $191 - 16q$, which declines very rapidly with q.

[43] Note that only the numerator of this fraction is different from the baseline case; the denominator remains

Table 6: Corrected Estimates (Fraction of Naive Estimate)

	α		
Shock	**0.001**	**0.005**	**0.02**
Monetary policy	0.86	0.70	0.69
Neutral technology	0.67	0.67	0.65
Government spending	0.86	0.76	0.75
Investment	0.68	0.68	0.67
Price mark-up	0.87	0.75	0.69
Wage mark-up	1.00	0.96	0.92
Patience	0.79	0.77	0.67
90% simulated lower bound	1.00	0.94	0.90

Notes: Table entries show $\bar{\sigma}^2_{l,\alpha}/\hat{\sigma}^2_l$, where $\bar{\sigma}^2_{l,\alpha}$ is corrected estimated of variance of shock l using selection rule for α (see text) and $\hat{\sigma}^2_l$ is naive estimate. $\bar{\sigma}^2_{l,\alpha}$ is unbiased estimator of variance of regression residuals from equation (22). That is, $\bar{\sigma}^2_{l,\alpha} = (T_{l,\alpha} - k_{l,\alpha})^{-1} \sum \hat{\xi}^2_{l,t}$, where $\hat{\xi}_{l,t}$ is fitted residual, $T_{l,\alpha}$ is number of observation periods, and $k_{l,\alpha}$ is number of regression parameters. Bottom row reports simulated 90 percent lower bounds for corresponding statistic when EM is correctly specified. Simulations are based on 10,000 draws of the same size as data set.

for the case where the EM is correctly specified are also reported (in parentheses). Column (1) reproduces the baseline naive BCF variance decomposition from Table 5, while columns (2)-(4) report corrected decompositions of the BCF variance of hours for different values of α. The Table shows substantial declines in the estimated importance of the investment shock for the BCF variance of hours, with estimates decreasing by 19-20 percentage points relative to the naive case.

7 Conclusion

This paper makes several contributions to the literature. First, a novel framework is developed that can be used to analyze the implications of misspecification on estimates of model shock variances. Using this framework, I show that if a DSGE model is correctly specified, then under basic stability conditions the time-series process for the smoothed shocks should be a vector white noise with diagonal covariance matrix. Thus, if a realized process for the smoothed shocks does not possess this property, then the model must be misspecified.

Next, I showed how, using this framework, one may orthogonally decompose a smoothed shock into two components: its true value, and an additional component related entirely to misspecification. Since the true values of any two shocks are independent by construction, any non-zero covariance between two distinct smoothed shocks must then be entirely attributable to misspecification. Further, if the smoothed shocks do exhibit non-zero covariance, using the sample variance of a smoothed shock as an estimator of the variance of the true shock

unchanged.

Table 7: Corrected BCF Variance Decomposition for Hours

	(1)	(2)	(3)	(4)
Decomposition:	*Naive*	*Corrected*	*Corrected*	*Corrected*
$\boldsymbol{\alpha}$:		*0.001*	*0.005*	*0.02*
Shock				
Monetary policy	0.06	0.05	0.04	0.04
		(0.06)	(0.06)	(0.06)
Neutral technology	0.11	0.07	0.07	0.07
		(0.11)	(0.10)	(0.09)
Government spending	0.02	0.02	0.01	0.01
		(0.02)	(0.02)	(0.02)
Investment	0.61	0.42	0.42	0.41
		(0.61)	(0.58)	(0.55)
Price mark-up	0.06	0.05	0.04	0.04
		(0.06)	(0.06)	(0.05)
Wage mark-up	0.05	0.05	0.05	0.05
		(0.05)	(0.05)	(0.05)
Patience	0.08	0.07	0.07	0.06
		(0.09)	(0.08)	(0.08)

Notes: Columns (2)-(4) show corrected BCF variance decompositions for hours. Naive decomposition is reproduced in column (1). Figures in parentheses are 90 percent simulated lower bounds for case when EM is correctly specified.

will lead to estimates that are biased upward. To correct for this potential source of bias, I propose a fairly simple methodology which involves extracting the component of a given smoothed shock that is unpredictable using other shocks.

I apply this framework and methodology to a recent paper by Justiniano et al. (2010), and estimate that at least one-third of the variance of the investment shock—the leading driver of business cycle fluctuations in their model—can be attributed to misspecification, and as a result, the estimated importance of the investment shock in generating business cycle variation in hours declines by around 20 percentage points.

References

Adjemian, Stéphane, Houtan Bastani, Michel Juillard, Ferhat Mihoubi, George Perendia, Marco Ratto, and Sébastien Villemot (2011). "Dynare: Reference Manual, Version 4." *Dynare Working Papers*, 1, CEPREMAP.

An, Sungbae and Frank Schorfheide (2007). "Bayesian Analysis of DSGE Models." *Econometric Reviews*, 26(2-4):113–172.

Barro, Robert J. and Robert G. King (1984). "Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles." *The Quarterly Journal of Economics*, 99(4):817–839.

Barsky, Robert B. and Eric R. Sims (2012). "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence." *The American Economic Review*, 102(4):1343–1377.

Beaudry, Paul and Franck Portier (2004). "An exploration into Pigou's theory of cycles." *Journal of Monetary Economics*, 51:1183–1216.

Beaudry, Paul and Franck Portier (2007). "When can changes in expectations cause business cycle fluctuations in neo-classical settings?" *Journal of Economic Theory*, 135:458–477.

Blanchard, Olivier J., Jean-Paul L'Huillier, and Guido Lorenzoni (2013). "News, Noise, and Flutuations: An Empirical Exploration." *The American Economic Review*, 103(7):3045–3070.

Canova, Fabio (1994). "Statistical Inference in Calibrated Models." *Journal of Applied Econometrics*, 9:S123–S144.

Canova, Fabio (2007). *Methods for Applied Macroeconomic Research.* Princeton, NJ: Princeton University Press.

Christiano, Lawrence, Cosmin L. Ilut, Roberto Motto, and Massimo Rostagno (2010a). "Monetary policy and stock market booms." NBER Working Paper 16402.

Christiano, Lawrence, Roberto Motto, and Massimo Rostagno (2010b). "Financial Factors in Economic Fluctuations." ECB Working Paper Series, No. 1192.

Christiano, Lawrence, Roberto Motto, and Massimo Rostagno (2014). "Risk Shocks." *The American Economic Review*, 104(1):27–65.

Diebold, Francis X., Lee E. Ohanian, and Jeremy Berkowitz (1998). "Dynamic Equilibrium Economies: A Framework for Comparing Models and Data." *The Review of Economic Studies*, 65(3):433–451.

Dridi, Ramdan, Alain Guay, and Eric Renault (2007). "Indirect inference and calibration of dynamic stochastic general equilibrium models." *Journal of Econometrics*, 136(2):397–430.

Fernández-Villaverde, Jesús (2010). "The econometrics of DSGE models." *Journal of the Spanish Economic Association*, 1(1-2):3–49.

Fernández-Villaverde, Jesús, Juan F. Rubio-Ramírez, Thomas J. Sargent, and Mark W. Watson (2007). "ABCs (and Ds) of Understanding VARs." *The American Economic Review*, 97(3):1021–1026.

Gourieroux, C., A. Monfort, and E. Renault (1993). "Indirect Inference." *Journal of Applied Econometrics*, 8:S85–S118.

Greenwood, Jeremy, Zvi Hercowitz, and Gregory W. Huffman (1988). "Investment, Capacity Utilization, and the Real Business Cycle." *The American Economic Review*, 78(3):402–417.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell (1997). "Long-Run Implications of Investment-Specific Technological Change." *The American Economic Review*, 87(3):342–362.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell (2000). "The role of investment-specific technological change in the business cycle." *European Economic Review*, 44(1):91–115.

Hall, Alastair R. and Atsushi Inoue (2003). "The large sample behaviour of the generalized method of moments estimator in misspecified models." *Journal of Econometrics*, 114(2):361–394.

Hamilton, James D. (1994). *Time Series Analysis*. Princeton: Princeton University Press.

Hnatkovska, Viktoria, Vadim Marmer, and Yao Tang (2012). "Comparison of Misspecified Calibrated Models: The Minimum Distance Approach." *Journal of Econometrics*, 169(1):131–138.

Ingram, Beth Fisher, Narayana R. Kocherlakota, and N.E. Savin (1994). "Explaining business cycles: A multiple-shock approach." *Journal of Monetary Economics*, 34(3):415–428.

Jaimovich, Nir and Sergio Rebelo (2009). "Can News about the Future Drive the Business Cycle?" *The American Economic Review*, 99(4):1097–1118.

Jermann, Urban and Vincenzo Quadrini (2012). "Macroeconomic Effects of Financial Shocks." *The American Economic Review*, 102(1):238–271.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti (2010). "Investment shocks and business cycles." *Journal of Monetary Economics*, 57(2):132–145.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti (2011). "Investment shocks and the relative price of investment." *Review of Economic Dynamics*, 14(1):102–121.

King, Robert G., Charles I. Plosser, and Sergio T. Rebelo (1988). "Production, growth and business cycles: I. The basic neoclassical model." *Journal of Monetary Economics*, 21(2-3):195–232.

Leeper, Eric M., Todd B. Walker, and Shu-Chun Susan Yang (2011). "Foresight and Information Flows." NBER Working Paper 16951.

Nolan, Charles and Christoph Thoenissen (2009). "Financial shocks and the US business cycle." *Journal of Monetary Economics*, 56(4):596–604.

Schmitt-Grohé, Stephanie and Martín Uribe (2011). "Business cycles with a common trend in neutral and investment-specific productivity." *Review of Economic Dynamics*, 14(1):122–135.

Schmitt-Grohé, Stephanie and Martín Uribe (2012). "What's News in Business Cycles." *Econometrica*, 80(6):2733–2764.

Schorfheide, Frank (2000). "Loss Function-Based Evaluation of DSGE Models." *Journal of Applied Econometrics*, 15(6):645–670.

Sims, Eric R. (2012). "News, Non-Invertibility, and Structural VARs." Working Paper.

Smets, Frank and Rafael Wouters (2003). "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association*, 1(5):1123–1175.

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

Stock, James H. and Mark W. Watson (1999). "Business Cycle Fluctuations in US Macroeconomic Time Series." In John B. Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*, vol. 1A. Amsterdam: Elsevier.

Watson, Mark W. (1993). "Measures of Fit for Calibrated Models." *Journal of Political Economy*, 101(6):1011–1041.

A Model details for Example 3 in Section 5 (Smets and Wouters (2007))

The equations for the Smets and Wouters (2007) model are reproduced here for convenience. See Smets and Wouters (2007) for specifics regarding the model set-up and the derivation of these equations. The resource constraint:

$$y_t = c_y c_t + i_y i_t + z_y z_t + \varepsilon_t^g$$

The consumption Euler equation:

$$c_t = c_1 c_{t-1} + (1 - c_1)\,\mathbb{E}_t c_{t+1} + c_2\,(l_t - \mathbb{E}_t l_{t+1}) - c_3 \left(r_t - \mathbb{E}_t \pi_{t+1} + \varepsilon_t^b\right)$$

The investment Euler equation:

$$i_t = i_1 i_{t-1} + (1 - i_1)\,\mathbb{E}_t i_{t+1} + i_2 q_t + \varepsilon_t^i$$

The value of the capital stock evolves according to:

$$q_t = q_1 \mathbb{E}_t q_{t+1} + (1 - q_1)\,\mathbb{E}_t r_{t+1}^k - \left(r_t - \mathbb{E}_t \pi_{t+1} + \varepsilon_t^b\right)$$

The aggregate production function equation:

$$y_t = \phi_p\left[\alpha k_t^s + (1 - \alpha)\,l_t + \varepsilon_t^a\right]$$

The capital utilization equation:

$$k_t^s = k_{t-1} + z_t$$

The capital utilization first-order condition:

$$z_t = z_1 r_t^k$$

The capital accumulation equation:

$$k_t = k_1 k_{t-1} + (1 - k_1)\,i_t + k_2 \varepsilon_t^i$$

The intermediate goods firm's real mark-up:

$$\mu_t^p = \alpha\,(k_t^s - l_t) + \varepsilon_t^a - w_t$$

The New Keynesian Phillips curve:

$$\pi_t = \pi_1 \pi_{t-1} + \pi_2 \mathbb{E}_t \pi_{t+1} - \pi_3 \mu_t^p + \varepsilon_t^p$$

The firm's capital-labor ratio:

$$k_t^s - l_t = w_t - r_t^k$$

The labor market real mark-up:

$$\mu_t^w = w_t - \sigma_l l_t - \frac{\gamma}{\gamma - \lambda} c_t + \frac{\lambda}{\gamma - \lambda} c_{t-1}$$

The wage Phillips curve:

$$w_t = w_1 w_{t-1} + (1 - w_1)\left(\mathbb{E}_t w_{t+1} + \mathbb{E}_t \pi_{t+1}\right) - w_2 \pi_t + w_3 \pi_{t-1} - w_4 \mu_t^w + \varepsilon_t^w$$

Finally, the Taylor rule:

$$r_t = \rho r_{t-1} + (1 - \rho)\left[r_\pi \pi_t + r_y \left(y_t - y_t^p\right)\right] + r_{\Delta y}\left[y_t - y_t^p - y_{t-1} + y_{t-1}^p\right] + \varepsilon_t^r$$

Above, y_t is output, c_t is consumption, i_t is investment, z_t are capital utilization costs, l_t is hours worked, k_t is the end-of-period capital stock, k_t^s is employed capital, q_t is the shadow value of capital, μ_t^p is the intermediate goods firms' real mark-up, r_t^k is the rental rate on capital, w_t is the real wage rate, μ_t^w is the real wage mark-up, r_t is the nominal interest rate, π_t is inflation from date $(t-1)$ to date t, and y_t^p is potential output (as obtained from the RBC version of the model). All variables are log-linearized around the deterministic trend.

The RBC version of the model is obtained by setting

$$\xi_w = 0$$
$$\xi_p = 0$$

and removing the risk premium, two mark-up, and monetary policy shocks. All other features of the DGP remain unchanged, including all of the real frictions and the monopolistically competitive intermediate goods market. The New Keynesian Phillips curve, wage Phillips curve and Taylor rule are indeterminate and thus dropped from the system, while the intermediate goods and labor market mark-up equations are replaced, respectively, with

$$0 = \alpha\left(k_t^s - l_t\right) + \varepsilon_t^a - w_t$$

$$0 = w_t - \sigma_l l_t - \frac{\gamma}{\gamma - \lambda} c_t + \frac{\lambda}{\gamma - \lambda} c_{t-1}$$

The consumption and marginal value of capital Euler equations are also combined to obtain

$$c_t = c_1 c_{t-1} + (1 - c_1) c_t^e + c_2\left(l_t - l_t^e\right) - c_3 q_1 q_t^e - c_3 (1 - q_1) r_t^{ke} + c_3 q_t$$

We have thus reduced the system by four equations, and correspondingly drop four variables: r_t, π_t, μ_t^p and μ_t^w.

B Proofs of theoretical results

Proof of Proposition 1

From (16)-(18), we can write

$$\varepsilon_t^* = (FB)^{-1}(Y_t - Z_t) - (FB)^{-1} FAX_{t-1}^* \tag{B.1}$$

Substituting this into equation (18) for ε_t^* yields

$$X_t^* = CX_{t-1}^* + B(FB)^{-1}(Y_t - Z_t) \tag{B.2}$$

Next, we may write $C = QJQ^{-1}$, where J is the Jordan normal form of C and Q is a matrix whose columns are the corresponding generalized eigenvectors of C. In particular, J is an $n \times n$ matrix with the eigenvalues of C on the main diagonal ordered by increasing modulus, and, if there is a non-zero non-diagonal entry, then it is equal to one, lies immediately above the main diagonal, and satisfies that the entry immediately to the left of it is equal to the entry immediately below it. Note that Q and J are functions of A, B and F only. Note also that, since J is an upper-triangular matrix, its eigenvalues are its diagonal elements, and thus it has the same eigenvalues as C. Under Assumption 4, J can therefore be partitioned as

$$J = \begin{pmatrix} J_1 & 0 \\ 0 & J_2 \end{pmatrix}$$

where J_1 is an $n_1 \times n_1$ matrix having eigenvalues strictly less than one in modulus, and J_2 is an $n_2 \times n_2$ matrix having eigenvalues strictly greater than one in modulus, with $n_1 + n_2 = n$. Note that J_2 is non-singular by construction. Partition Q and Q^{-1} conformably as

$$Q = \begin{pmatrix} Q_1 & Q_2 \end{pmatrix}$$

$$Q^{-1} = \begin{pmatrix} Q^1 \\ Q^2 \end{pmatrix}$$

Letting $x_t \equiv Q^{-1}X_t^*$ and $y_t \equiv Q^{-1}B(FB)^{-1}(Y_t - Z_t)$, premultiplying equation (B.2) by Q^{-1} yields

$$x_t = Jx_{t-1} + y_t$$

or, partitioning x_t and y_t conformably with J,

$$\begin{pmatrix} x_{1,t} \\ x_{2,t} \end{pmatrix} = \begin{pmatrix} J_1 x_{1,t-1} \\ J_2 x_{2,t-1} \end{pmatrix} + \begin{pmatrix} y_{1,t} \\ y_{2,t} \end{pmatrix}$$

We may thus write

$$x_{1,t} = \sum_{j=0}^{\infty} J_1^j L^j y_{1,t}$$

and

$$x_{2,t} = \sum_{j=-\infty}^{-1} J_2^j L^j y_{2,t}$$

where convergence of the sums follows from the fact that the eigenvalues of J_1 and J_2^{-1} are all less than one in modulus. Thus

$$\begin{aligned} X_t^* &= Qx_t \\ &= \sum_{j=0}^{\infty} Q_1 J_1^j Q^1 B (FB)^{-1} L^j (Y_t - Z_t) + \sum_{j=-\infty}^{-1} Q_2 J_2^j Q^2 B (FB)^{-1} L^j (Y_t - Z_t) \end{aligned}$$

Backing this expression up one period and substituting into equation (B.1) for X_{t-1} yields

$$\varepsilon_t^* = \sum_{j=-\infty}^{\infty} \psi_j L^j (Y_t - Z_t)$$

where

$$\psi_j = \begin{cases} -(FB)^{-1} FAQ_2 J_2^{j-1} Q^2 B (FB)^{-1} & \text{if } j < 0 \\ (FB)^{-1} \left[I_m - FAQ_2 J_2^{-1} Q^2 B (FB)^{-1} \right] & \text{if } j = 0 \\ -(FB)^{-1} FAQ_1 J_1^{j-1} Q^1 B (FB)^{-1} & \text{if } j > 0 \end{cases}$$

Again, since the eigenvalues of J_1 and J_2^{-1} are all strictly less than one in modulus, the sequence $\{\psi_j\}$ is absolutely summable.

Finally, note that

$$X_t^* = \sum_{i=0}^{\infty} A^i B L^i \varepsilon_t^*$$

and thus

$$\left(\sum_{i=0}^{\infty} FA^i B L^i \right) \varepsilon_t^* = Y_t - Z_t \tag{B.3}$$

Pre-multiplying (B.3) by $\sum_{j=-\infty}^{\infty} \psi_j L^j$ yields

$$\left(\sum_{j=-\infty}^{\infty} \psi_j L^j \right) \left(\sum_{i=0}^{\infty} FA^i B L^i \right) \varepsilon_t^* = \sum_{j=-\infty}^{\infty} \psi_j L^j (Y_t - Z_t) = \varepsilon_t^*$$

which confirms the first equality in (20). Substituting (B.3) into (19) for $(Y_t - Z_t)$ yields

$$\varepsilon_t^* = \left(\sum_{j=-\infty}^{\infty} \psi_j L^j \right) \left(\sum_{i=0}^{\infty} FA^i B L^i \right) \varepsilon_t^*$$

which confirms the second equality in (20).

□

Proof of Proposition 2

Since the EM can be inverted by Proposition 1, Proposition 2 follows directly from (13) and the subsequent discussion.

□

Proof of Corollary 1

Part (a) follows obviously from independence of ε_t^* and ν_t. To see part (b), suppose the EM is correctly specified, so that $Var(Z_t) = 0$. Then $\hat{\varepsilon}_t = \varepsilon_t^*$ and therefore $\mathbb{E}[\hat{\varepsilon}_{l,t}\hat{\varepsilon}_{i,s}] = \mathbb{E}\left[\varepsilon_{l,t}^*\varepsilon_{i,s}^*\right] = 0$ for all $(l,t) \neq (i,s)$. To see part (c), suppose the EM is misspecified. Then $Var(Z_t) \neq 0$ and therefore there exists some l such that $Var(\nu_{l,t}) \neq 0$. From Proposition 2, this implies that $\hat{\sigma}_l^2 > \sigma_l^{*2}$. Finally, part (d) follows from independence of $\varepsilon_{l,t}^*$ on the one hand and both $\varepsilon_{i,s}^*$ and $\nu_{i,s}$ on the other.

□

Proof of Proposition 3

By independence of $\varepsilon_{l,t}^*$ and $\nu_{l,t}$ (Proposition 2) and of $\varepsilon_{l,t}^*$ and $\hat{\varepsilon}_{i,s}$ for $(l,t) \neq (i,s)$ (Corollary 1(d)), we have

$$\begin{aligned}\mathbb{E}\left[\xi_{l,t}^2\right] &= \sigma_l^{*2} + \mathbb{E}\left[\left\{\nu_{l,t} - \sum_{i \neq l}\Theta_{i,0}\hat{\varepsilon}_{i,t} - \sum_{j=1}^{q}\left(\Theta_j\hat{\varepsilon}_{t-j} + \Theta_{-j}\hat{\varepsilon}_{t+j}\right)\right\}^2\right] \\ &\geq \sigma_l^{*2}\end{aligned}$$

The second part of the Proposition (i.e., that $\mathbb{E}\left[\xi_{l,t}^2\right] \leq \hat{\sigma}_l^2$) follows directly from standard regression results.

□

C Data for JPT model

Quarterly U.S. data spanning 1954QIV to 2009QI was constructed for seven observable variables: real per capita GDP growth, real per capita consumption growth, real per capita investment growth, (log of) per capita hours, real wage growth, the inflation rate, and the nominal interest rate.

Raw data were obtained from several different sources. NIPA data were obtained from the Bureau of Economic Analysis. Data on the federal funds rate and population (civilian noninstitutional) were obtained from the St. Louis Fed's FRED database. An index of hours worked in the non-farm business sector was obtained from the Bureau of Labor Statistics, as was an index of compensation in the non-farm business sector.

Real per capital GDP was constructed as nominal GDP divided by the GDP deflator and population. Real per capita consumption was constructed as nominal purchases of nondurable goods and services, divided by the GDP deflator and population. Real per capita investment was constructed as nominal gross private domestic investment plus purchases of durable goods, divided by the GDP deflator and population. Per capita hours was constructed as the hours index divided by population, and normalized by a constant factor chosen so that the average of the log of the series over the period 1954QIV-2004QIV was zero. The real wage was constructed as the compensation index divided by the GDP deflator. The inflation rate was constructed as growth rate of the GDP deflator. Finally, the nominal interest rate was constructed as one-quarter of the federal funds rate in annual terms.

D BCF variance decomposition

Let $\Phi_k \equiv \mathbb{E}\left[X_t X_{t-k}'\right]$ denote the k-th autocovariance of X_t in the EM of Section 4.1. The autocovariance-generating function (AGF) for $\{X_t\}$ is given by[44]

$$G_X(z) = (I_n - Az)^{-1} B\Sigma B' \left(I_n - A'z^{-1}\right)^{-1}$$

where $\Sigma \equiv \mathbb{E}[\varepsilon_t \varepsilon_t']$. The spectrum for $\{X_t\}$ at frequency ω is given by

$$\begin{aligned} s_X(\omega) &= (2\pi)^{-1} G_X\left(e^{-i\omega}\right) \\ &= (2\pi)^{-1}\left(I_n - Ae^{-i\omega}\right)^{-1} B\Sigma B' \left(I_n - A'e^{i\omega}\right)^{-1} \end{aligned} \tag{D.1}$$

where $i \equiv \sqrt{-1}$. It can be shown that, for all integers k

$$\Phi_k = \int_{-\pi}^{\pi} s_X(\omega)\, e^{i\omega k} d\omega$$

Since we are interested in doing a decomposition of the unconditional variances of the endogenous variables, of particular interest is the case where $k = 0$. i.e.

$$\Phi_0 = \int_{-\pi}^{\pi} s_X(\omega)\, d\omega \tag{D.2}$$

Let $s_X^{jj}(\omega)$ denote the jj-th element of the spectrum. Equation D.2 implies that the integral of $s_X^{jj}(\omega)$ over all frequencies $\omega \in [-\pi, \pi]$ yields the unconditional variance of $X_{j,t}$ (the j-th element of X_t). For $\omega_0 \in [0, \pi]$, we can interpret $\int_{-\omega_0}^{\omega_0} s_X^{jj}(\omega)\, d\omega$ as the portion of the unconditional variance of $X_{j,t}$ attributable to fluctuations with frequencies between 0 and ω_0. For $0 \le \omega_0 \le \omega_1 \le \pi$, define

$$v_j(\omega_0, \omega_1) \equiv \int_{-\omega_1}^{\omega_1} s_X^{jj}(\omega)\, d\omega - \int_{-\omega_0}^{\omega_0} s_X^{jj}(\omega)\, d\omega \tag{D.3}$$

[44] See, for example, Hamilton (1994).

If we choose ω_0 and ω_1 so that the interval $[\omega_0, \omega_1]$ gives the set of BCFs, then $v_j(\omega_0, \omega_1)$ can be interpreted as the portion of the unconditional variance of $X_{j,t}$ attributable to BCFs, i.e. its BCF variance.

Next, letting $v_j^{(l)}(\omega_0, \omega_1)$ be the quantity in (D.3) obtained when the variances of all but the l-th shock are set to zero, it is straightforward to show that $\sum_{l=1}^{m} v_j^{(l)}(\omega_0, \omega_1) = v_j(\omega_0, \omega_1)$. Thus,

$$\frac{v_j^{(l)}(\omega_0, \omega_1)}{v_j(\omega_0, \omega_1)}$$

can be interpreted as the contribution of the shock l to the BCF variance of $X_{j,t}$.

In practice, the integral in (D.3) is approximated using 500 bins spanning the relevant frequencies. Note that the period of a cycle with frequency ω is given by $\frac{2\pi}{\omega}$. Following Stock and Watson (1999) and JPT, I define business cycle frequencies to be those with periods between 6 and 32 quarters. Thus, I set $\omega_0 = \frac{2\pi}{32}$ and $\omega_1 = \frac{2\pi}{6}$.