

Spectral decomposition of the information about latent variables in dynamic macroeconomic models

Nikolay Iskrev

Abstract

I propose a formal method for decomposing the frequency domain information about latent variables in dynamic models. A model describes the joint probability distribution of a set of observed and latent variables. The amount of information transferred from the former to the latter is measured by the reduction of uncertainty in the posterior compared to the prior distribution of any given latent variable. Casting the analysis in the frequency domain allows decomposing the total amount of information in terms of frequency-specific contributions as well as in terms of information contributed by individual observed variables. Using the proposed spectral decomposition can help researchers identify where information about shocks and other latent variables in structural macroeconomic models come from, thereby making the estimation of such models more transparent. I illustrate the usefulness of the methodology with applications to three recent articles.

Keywords: DSGE models, Frequency domain, Information content

JEL classification: C32, C51, C52, E32

Bank of Portugal, email: nikolay.iskrev@bportugal.pt

Additional material supporting this paper can be accessed at https://niskrev.github.io/siga_docs/

The views expressed in this paper are those of its author and do not necessarily reflect the views of Banco de Portugal or the Eurosystem.

There is nothing like a latent variable to stimulate the imagination.

A. Goldberger, quoted by Chamberlain (1990)

1 Introduction

A pervasive challenge in macroeconomic research is the estimation of latent variables by combining theoretical models and empirical information. Examples abound and include endogenously determined variables, such as potential output and natural rates of interest or unemployment, as well as a plethora of exogenous shocks driving business cycle fluctuations in modern macroeconomic models. These variables are typically not directly observable and to measure them requires estimating models that explicitly describe the joint dynamics of observed and latent variables. Having correctly specified and accurately measured latent endogenous variables and structural shocks is a key requirement for macroeconomic models to meet to be useful as tools for policy analysis and to be credible as story-telling devices.

The purpose of this paper is to show how to perform a spectral decomposition of the information about latent variables in dynamic economic models. In particular, the proposed analysis reveals where in the frequency spectrum information about latent variables predominantly comes from, and how much of it is contributed by individual observed variables. The goal of this analysis is to enhance researchers' understanding of where in the data, according to a given model, information about unobservable quantities comes from. In doing so, the paper contributes to the emerging literature aimed at improving the transparency of structural estimation in macroeconomic research.

The question of where in the data information about estimated quantities of interest comes from is frequently discussed in the context of structural models. Compared to reduced-form estimation, in structural models it is typically much harder to draw the connection between different features of the data, on the one hand, and particular estimated parameters or latent variables, on the other. This can make it difficult to understand to what extent different modelling assumptions influence the estimation results, which in turn hampers the ability of readers of the research to assess the credibility of estimated models. Providing details on how information is derived from observed data can be helpful in resolving this difficulty. In particular, this lets readers, who might suspect that a model is misspecified in some dimensions, better understand the

implications of that misspecification.¹ The frequency domain perspective is particularly relevant in this context as different researchers may have different views on which data frequencies can be adequately represented by a given model. For example, fitting data which is contaminated by high frequency noise will distort the estimation of models that do not allow for discrepancy between model variables and observed series. Similarly, models which are best suited to explain business cycle fluctuations may lack the mechanisms needed to account for low frequency fluctuations in the data used to estimate them, thereby contaminating information coming from the lower end of the spectrum. The same argument applies to fitting models that do not account for seasonality to data which exhibits seasonal patterns. In the existing literature, there are different and sometimes contradictory approaches to dealing with such issues. In particular, there is no consensus among practitioners on whether to allow for measurement errors in commonly used series, or how to treat long-term trends in the data when estimating structural macroeconomic models. Therefore, being transparent on how information from different parts of the spectrum is used to estimate latent variables of interest will provide readers, who have different views on the adequacy of a model to represent the data, with the ability to assess the potential consequences of misspecification, given the relative importance of the frequencies they suspect are contaminated for the identification of those variables.

The work most closely related to this paper is Iskrev (2019), where the question regarding the sources of information about latent variables is treated in the time domain. In that paper, the amount of information from observable variables about latent variables is quantified by comparing prior and posterior probability distributions and employing information-theoretic measures of uncertainty and information gain. Analysis in the time domain preserves information about the temporal order of the observable data in relation to the latent variables and allows to study the transfer of information between variables with arbitrary temporal patterns. In particular, one can evaluate the contribution of information from any observed variable originating in any subperiod of the sample. This can facilitate the assessment of the consequences of model misspecification by readers who suspect that some observed series diverge from the respective model concepts during some part of the sample.² The information pertaining to the temporal order of the variables is

¹See Andrews et al. (2020) and the references therein for a broader perspective on this topic.

²One example of where this could be useful is the estimation of monetary models with data including the period when the zero lower bound on interest rates is binding and non-conventional monetary policy measures are in place. One approach in the literature for dealing with this issue is to use the so-called “shadow interest rate” as an observed counterpart of the policy rate in theoretical models (see e.g. Giannone et al. (2016)). Estimated shadow rates are not constrained by the effective lower bound

lost completely in the frequency domain. On the other hand, it permits decomposing uncertainty and information into components of varying frequencies. Therefore, it reveals how much and where in the spectrum uncertainty is resolved for a given latent variable and what are the contributions from different observed variables and frequencies. This is something that cannot be deduced in the time domain. Also, instead of misspecification with respect to a subperiod of the full sample, the frequency domain analysis can aid readers's understanding of the consequences of using information from a subset of frequencies, which they suspect or believe may not be well represented by the estimated models. The time and frequency domain approaches are therefore complementary.

It is important to emphasize that the analysis described in this paper does not require that models are solved, estimated, or otherwise transformed from the time to the frequency domain. That is, it can be applied irrespective of how models are estimated. In that sense, it is similar to performing parameter identification analysis in the frequency domain (see e.g. Qu and Tkachenko (2013)) or reporting spectral variance decompositions (see, among many others, the handbook chapter by Fernández-Villaverde et al. (2016)). However, it should be noted that, while far less common in the empirical literature, the use of spectral methods for the estimation and evaluation of macro models, has been previously advocated for in a number of important studies, such as Hansen and Sargent (1993), Watson (1993), Diebold et al. (1998), Christiano and Vigfusson (2003), Qu and Tkachenko (2012) and Sala (2015).

The paper is also related to a growing literature on the feasibility of recovering structural shocks using reduced form models. Building upon the work of Hansen and Sargent (1980, 1991) and Lippi and Reichlin (1993, 1994), most of the research on this topic has focused on the issue of invertibility (or fundamentalness) in structural vector autoregressions, i.e. whether shocks from general equilibrium models can be recovered from the residuals of VARs (see Alessi et al. (2011) and Giacomini (2013) for useful overviews of this literature). Conditions for invertibility are discussed in Fernandez-Villaverde et al. (2007), Ravenna (2007), Franchi and Vidotto (2013), Franchi and Paruolo (2015)), while Giannone and Reichlin (2006) and Forni and Gambetti (2014) discuss how to test for lack of invertibility of structural VARs. Invertibility issues that are specific to DSGE models with news shocks are discussed in Leeper et al. (2013) and Blanchard et al. (2013). More recently, Soccorsi (2016) and Sala et al. (2016) proposed

and usually coincide with the observed policy rates in the period when the bound is not binding. A possible concern about this approach is that the information from the shadow rate series is contaminated in the part of the sample where there is a substantial estimation uncertainty.

measures of the degree of non-invertibility, which quantify the discrepancies between true shocks and shocks obtained using non-fundamental VARs.³ In another recent paper Chahrour and Jurado (2022) draw a distinction between invertibility on one hand, and what they call “recoverability” on the other, defining the latter as the feasibility of recovering structural shocks from all leads and lags of the observables variables. They argue that recoverability is often what matters in applied research and present a necessary and sufficient condition one can use to check if shocks in linear models are recoverable.

Similar to that literature, the analysis in the present paper can be used to determine whether the shocks in a given model are recoverable given a set of observed variables. Furthermore, as in Soccorsi (2016) and particularly Sala et al. (2016), a measure is provided of the degree to which any individual shock, or an endogenous latent variable, can be recovered. In particular, the proposed spectral measures of information gain are defined with respect to a particular unobserved variable and show how much of the prior uncertainty about it, within a given frequency band, is removed due to observing a given set of model variables.

While the existing research on invertibility is concerned with the usefulness of VAR-based tools for empirical validation of structural models, the purpose of the analysis presented here is to understand the properties of structural macroeconomic models in terms of how much and from where in the spectrum information transfers between observed and unobserved model variables. Therefore, identifying the principal sources of information is of primary interest rather than the total amount of information about a given shock or endogenous latent variable. To that end, I define and apply measures of frequency band-specific conditional information gains that quantify the amount of additional information contributed by a subset of variables, given the information contained in the remaining observed variables, at a given band of frequencies. As the analysis of the models considered in the application section shows, the conclusions one draws may be very different depending on what the conditional variables are.

The remainder of the paper is organized as follows. Section 2 reviews the relevant information-theoretic and frequency domain concepts and defines measures of information gains from observable with respect to latent variables. It also shows how the measures can be evaluated for linear Gaussian DSGE models. The proposed measures are then used to decompose information about latent variables across frequencies and observed variables.

³Simulation evidence that non-invertible VARs may in some cases produce good approximations of the true structural shocks are provided in Sims (2012) and Beaudry et al. (2015).

The purpose of this analysis is to identify the primary sources of information about any latent variable of interest, thereby making the estimation of structural macroeconomic models with latent variables more transparent. The methodology is illustrated in Section 3 with the help of three applications. The first is a small-scale New-Keynesian model employed by Uribe (2021) to investigate the nature and empirical importance of monetary policy shocks that produce neo-Fisherian dynamics, i.e. move interest rates and inflation in the same direction over the short run. The second is a medium-scale New Keynesian model estimated by Justiniano et al. (2011) in order to investigate whether investment shocks are important drivers of business cycle fluctuations. The last application is another medium-scale New Keynesian model presented in Angeletos et al. (2018) as an illustration of their method for augmenting macroeconomic models with a higher-order belief dynamics. In all applications, I investigate where information about structural shocks come from. The examples are chosen to showcase different aspects of the proposed information decomposition and to demonstrate the usefulness of this analysis in making the estimation of such models more transparent. Section 4 concludes. An Online Appendix provides further details on the model specifications as well as additional results.

2 Methodology

The purpose of this section is three-fold. First, to introduce some basic information-theoretic concepts and use them to define a measure of information gain for variables with a multivariate complex Gaussian distribution. Second, to review relevant properties of the spectral representation of a stationary Gaussian vector process and present frequency domain measures of information gain. Third, to show how to apply the measures in the context of DSGE models to evaluate the information contributions with respect to latent variables across observed variables and frequencies.

2.1 Quantifying information gains

Consider a $(n_{\mathbf{y}} + 1)$ -dimensional random vector $\mathbf{z} = [\mathbf{y}', x]'$ whose joint probability density function is $f(\mathbf{y}, x)$. How much information about x is gained when a realization of \mathbf{y} is observed? Information theory provides the framework and tools to answer such questions. Specifically, entropy is a measure of the uncertainty associated with a random

variable, and mutual information is a measure of the information shared by two random variables. Formally, if $f(x)$ is the marginal probability density function of x , and \mathbf{S}_x is the support of x , the entropy $H(x)$ of $f(x)$ is defined as

$$H(x) = - \int_{\mathbf{S}_x} f(x) \ln(f(x)) dx = - \mathbb{E} \ln f(x). \quad (2.1)$$

The amount of information about x is measured as the reduction in uncertainty, i.e. the entropy $H(x)$, relative to some base distribution. The mutual information of the random variables \mathbf{y} and x is defined as

$$I(\mathbf{y}, x) = \int_{\mathbf{S}_y} \int_{\mathbf{S}_x} f(\mathbf{y}, x) \ln \frac{f(\mathbf{y}, x)}{f(\mathbf{y})f(x)} d\mathbf{y}dx \quad (2.2)$$

where \mathbf{S}_y is the support of \mathbf{y} . The information interpretation of (2.2) follows from the fact that it can be expressed in terms of entropy as

$$I(\mathbf{y}, x) = H(x) - H(x|\mathbf{y}). \quad (2.3)$$

where $H(x|\mathbf{y}) = -\mathbb{E} \ln f(x|\mathbf{y})$ is the entropy of the conditional probability density function of x given \mathbf{y} . Therefore, $I(\mathbf{y}, x)$ has an intuitive interpretation as the reduction of the uncertainty about x due to observing \mathbf{y} .⁴ It can be shown (see Granger and Lin (1994)) that $H(x) \geq H(x|\mathbf{y})$ with equality if and only if $f(\mathbf{y}, x) = f(\mathbf{y})f(x)$. Hence, unless \mathbf{y} and x are independent, observing \mathbf{y} provides information about x . If we partition \mathbf{y} into two sub-vectors \mathbf{y}_1 and \mathbf{y}_2 , we can express the conditional mutual information of x and \mathbf{y}_1 given \mathbf{y}_2 as

$$I(\mathbf{y}_1, x|\mathbf{y}_2) = H(x|\mathbf{y}_2) - H(x|\mathbf{y}_1, \mathbf{y}_2) \quad (2.4)$$

The conditional mutual information tells us how much of the uncertainty about x that remains after \mathbf{y}_2 is observed is removed by observing also \mathbf{y}_1 . Now, let the joint density function $f(\mathbf{y}, x)$ be the $(n_y + 1)$ -dimensional complex Gaussian distribution,

$$\mathcal{N}_{\mathbb{C}} \left(\begin{pmatrix} \mathbf{0} \\ 0 \end{pmatrix}, \begin{pmatrix} \Sigma_{yy} & \Sigma_{yx} \\ \Sigma_{xy} & \Sigma_{xx} \end{pmatrix} \right) \quad (2.5)$$

⁴This is analogous to measuring the information gain in the posterior compared to the prior distribution of parameters in Bayesian analysis, see Lindley (1956).

Then, both the marginal distribution of x and the conditional distribution of x given \mathbf{y} are univariate complex Gaussian distributions, with covariances given by Σ_{xx} , and $\Sigma_{x|\mathbf{y}} = \Sigma_{xx} - \Sigma_{xy}\Sigma_{\mathbf{y}\mathbf{y}}^{-1}\Sigma_{\mathbf{y}x}$, respectively. Using this, it is straightforward to show that the mutual information of \mathbf{y} and x is

$$I(\mathbf{y}, x) = H(x) - H(x|\mathbf{y}) = .5 \ln \left(\frac{\Sigma_{xx}}{\Sigma_{x|\mathbf{y}}} \right) \quad (2.6)$$

From $H(x) \geq H(x|\mathbf{y})$ it follows that mutual information is positive unless \mathbf{y} and x are independent in which case it is zero. On the other hand, if the variables are perfectly dependent i.e. there exists a one-to-one function g such that $x = g(\mathbf{y})$, observing \mathbf{y} is equivalent to observing x . In that case $\Sigma_{x|\mathbf{y}} = 0$ and $I(\mathbf{y}, x) = \infty$. It is common in practice to normalize the measure to be in the interval $[0, 1]$. This can be achieved using the following monotonous increasing transformation (see e.g. Joe (1989) or Granger and Lin (1994))

$$I^*(\mathbf{y}, z) = 1 - \exp(-2I(\mathbf{y}, z)) \quad (2.7)$$

Applying this transformation to (2.6) results in the following measure of information gain:

$$\text{IG}_{\mathbf{y} \rightarrow x} = \left(\frac{\Sigma_{xx} - \Sigma_{x|\mathbf{y}}}{\Sigma_{xx}} \right) \times 100, \quad (2.8)$$

The interpretation of $\text{IG}_{\mathbf{y} \rightarrow x}$ is the following: it measures the reduction in uncertainty about x due to observing vector \mathbf{y} , as a percent of the unconditional (prior) uncertainty about x . Similarly, when \mathbf{y} is partitioned into \mathbf{y}_1 and \mathbf{y}_2 , we can define the conditional information gain of \mathbf{y}_1 with respect to x , given \mathbf{y}_2 as

$$\text{IG}_{\mathbf{y}_1 \rightarrow x|\mathbf{y}_2} = \left(\frac{\Sigma_{x|\mathbf{y}_2} - \Sigma_{x|\mathbf{y}}}{\Sigma_{xx}} \right) \times 100, \quad (2.9)$$

The interpretation of $\text{IG}_{\mathbf{y}_1 \rightarrow x|\mathbf{y}_2}$ is the following: it shows the amount of uncertainty about x left after observing \mathbf{y}_2 that is removed by observing also \mathbf{y}_1 , as a percent of the unconditional uncertainty about x .

2.2 Information gains in the frequency domain

Let $\mathbf{z}_t \in \mathbb{R}^{n_z}$ for $t \in \mathbb{Z}$ be n_z -dimensional stationary Gaussian time series with

$$\mathbb{E} \mathbf{z}_t = \mathbf{0} \quad t \in \mathbb{Z} \quad (2.10)$$

$$\text{cov}(\mathbf{z}_t, \mathbf{z}_{t-h}) = \mathbf{\Gamma}(h) \quad t, h \in \mathbb{Z} \quad (2.11)$$

If $\mathbf{Z} = [\mathbf{z}'_1, \mathbf{z}'_2, \dots, \mathbf{z}'_T]'$ is a $T \times n_z$ -dimensional realization the process, the joint distribution of \mathbf{Z} , as well as the marginal and conditional distributions of any subset of components of \mathbf{Z} , will be Gaussian. Therefore, in the time domain, the information gain measures from the previous section can be applied directly to quantify the information gained with respect to any realization of a component of \mathbf{z} due to observing a sample of realizations of a subset of the remaining components of the process (see Iskrev (2019)).

In the frequency domain, the information gains analysis proceeds by applying the discrete Fourier transform to the values of \mathbf{Z} :

$$Z(\omega_j) = (2\pi T)^{-1/2} \sum_{t=1}^T \mathbf{z}_t e^{-it\omega_j} \quad (2.12)$$

for the Fourier frequencies $\omega_j = 2\pi j/T$, where $j \in \{j \in \mathbb{Z} : -\pi < 2\pi j/T \leq \pi\}$.

Due to the linearity of the discrete Fourier transform, the joint Gaussianity is preserved. Furthermore, it can be shown that $Z(\omega_j)$ behave asymptotically as independent complex Gaussian random variables with zero mean and a covariance matrix equal to $f(\omega_j)$, where $f_{\mathbf{z}\mathbf{z}}(\omega) \in \mathbb{C}^{n_z \times n_z}$ is the spectral density matrix of $\mathbf{z}(t)$ at frequency ω (see Brillinger (1981, Theorem 4.4.1)),

$$f_{\mathbf{z}\mathbf{z}}(\omega) = (2\pi)^{-1} \sum_{h=-\infty}^{\infty} \mathbf{\Gamma}(h) e^{-ih\omega} \quad (2.13)$$

The asymptotic independence of the Fourier coefficients $Z(\omega_j)$ across frequencies implies that information gain analysis may be conducted on a frequency-by-frequency basis. In particular, (asymptotically) there is no information about a given component of the series at a frequency ω_j that comes from components at any other frequency ω_l , $l \neq j$. Furthermore, the complex Gaussianity of the distribution implies that information analysis at a given frequency ω can be performed using the information gain measures from Section 2.1. To be more concrete, consider a partition of \mathbf{z}_t into a n_y -dimensional vector \mathbf{y}_t and a scalar x_t , and let $\mathbf{y}(\omega)$ and $x(\omega)$ be their respective discrete Fourier

transforms at a frequency $\omega \in (-\pi, \pi]$. The spectral density matrix of $[\mathbf{y}'_t, x_t]'$ is given by

$$f_{\mathbf{z}\mathbf{z}}(\omega) = \begin{bmatrix} f_{\mathbf{y}\mathbf{y}}(\omega) & f_{\mathbf{y}x}(\omega) \\ f_{x\mathbf{y}}(\omega) & f_{xx}(\omega) \end{bmatrix} \quad (2.14)$$

and the frequency-specific information gain of $\mathbf{y}(\omega)$ with respect to $x(\omega)$ is⁵

$$\text{IG}_{\mathbf{y} \rightarrow x}(\omega) = \left(\frac{f_{xx}(\omega) - f_{x|\mathbf{y}}(\omega)}{f_{xx}(\omega)} \right) \times 100 \quad (2.15)$$

where $f_{x|\mathbf{y}}(\omega) = f_{xx}(\omega) - f_{x\mathbf{y}}(\omega)f_{\mathbf{y}\mathbf{y}}^{-1}(\omega)f_{\mathbf{y}x}(\omega)$ is the partial spectrum of x given \mathbf{y} (Priestley (1981)). Furthermore, if we split \mathbf{y}_t into \mathbf{y}_{1t} and \mathbf{y}_{2t} and let $\mathbf{y}_1(\omega)$ and $\mathbf{y}_2(\omega)$ be their respective discrete Fourier transforms, the frequency-specific conditional information gain of $\mathbf{y}_1(\omega)$ with respect to $x(\omega)$ given $\mathbf{y}_2(\omega)$ is

$$\text{IG}_{\mathbf{y}_1 \rightarrow x|\mathbf{y}_2}(\omega) = \left(\frac{f_{x|\mathbf{y}_2}(\omega) - f_{x|\mathbf{y}}(\omega)}{f_{xx}(\omega)} \right) \times 100 \quad (2.16)$$

The interpretation of $\text{IG}_{\mathbf{y} \rightarrow x}(\omega)$ and $\text{IG}_{\mathbf{y}_1 \rightarrow x|\mathbf{y}_2}$ is the same as before, except that now information is defined in terms of the reduction of uncertainty about x at a given frequency ω due to information from \mathbf{y} (or conditionally, from \mathbf{y}_1), also at frequency ω . In practice, we are usually interested not in a single frequency but rather in a band of frequencies, such as low, business cycle, or high frequencies. Frequency band-specific measure of information gain may be obtained by replacing in (2.15) and (2.16) the frequency-specific spectrum and conditional spectrum of x with their integrated versions,

$$\text{IG}_{\mathbf{y} \rightarrow x}(\boldsymbol{\omega}) = \left(\frac{f_{xx}(\boldsymbol{\omega}) - f_{x|\mathbf{y}}(\boldsymbol{\omega})}{f_{xx}(\boldsymbol{\omega})} \right) \times 100 \quad (2.17)$$

$$\text{IG}_{\mathbf{y}_1 \rightarrow x|\mathbf{y}_2}(\boldsymbol{\omega}) = \left(\frac{f_{x|\mathbf{y}_2}(\boldsymbol{\omega}) - f_{x|\mathbf{y}}(\boldsymbol{\omega})}{f_{xx}(\boldsymbol{\omega})} \right) \times 100 \quad (2.18)$$

where $\boldsymbol{\omega} = \{\omega : \omega \in [\underline{\omega}, \bar{\omega}] \cup [-\bar{\omega}, -\underline{\omega}]\}$ denotes the frequency band of interest, $f_{xx}(\boldsymbol{\omega}) = \int_{\omega \in \boldsymbol{\omega}} f_{xx}(\omega) d\omega$, and $f_{x|\mathbf{y}}(\boldsymbol{\omega}) = \int_{\omega \in \boldsymbol{\omega}} f_{x|\mathbf{y}}(\omega) d\omega$. The interpretation remains the same, except that now the uncertainty and information about x are with respect to the

⁵As in the time domain, there is a natural connection between the measures of spectral information gains and the frequency domain version of mutual information, see Brillinger (2002) and Brillinger and Guha (2007).

frequency band $\boldsymbol{\omega}$. Note that $\text{IG}_{\mathbf{y} \rightarrow x}(\boldsymbol{\omega})$ can be written also as

$$\text{IG}_{\mathbf{y} \rightarrow x}(\boldsymbol{\omega}) = \int_{\boldsymbol{\omega} \in \boldsymbol{\omega}} \text{IG}_{\mathbf{y} \rightarrow x}(\omega) \frac{f_{xx}(\omega)}{f_{xx}(\boldsymbol{\omega})} d\omega \quad (2.19)$$

Therefore, the information gain for a selected band of frequencies $\boldsymbol{\omega}$ is given simply by the weighted sum of the frequency-specific information gains, with weights equal to the contribution of each frequency to the total variance of x in $\boldsymbol{\omega}$. Similarly, the conditional information gain (2.18) can be expressed as a weighted sum of the frequency-specific conditional information gains.

A special case of the band-specific information gain is when $\boldsymbol{\omega}$ covers the full spectrum, i.e. when $\underline{\omega} = 0$ and $\bar{\omega} = \pi$. Let $\bar{\boldsymbol{\omega}} = \{\omega : \omega \in [0, \pi] \cup (0, -\pi]\}$. Then, it is straightforward to show that information gain takes the form:

$$\text{IG}_{\mathbf{y} \rightarrow x}(\bar{\boldsymbol{\omega}}) = \left(\frac{\text{var}(x_t) - \text{var}(x_t | \mathbf{y}_{t-\tau}, \tau \in \mathbb{Z})}{\text{var}(x_t)} \right) \times 100 \quad (2.20)$$

Therefore, in addition to the obvious frequency domain interpretation, it has a time-domain interpretation, namely, the percent reduction of the unconditional variance of x_t as a result of observing the infinite sequence of past, present, and future values of \mathbf{y}_t . Similarly, the full spectrum version of the conditional information gain measure of \mathbf{y}_1 with respect to x given \mathbf{y}_2 is

$$\text{IG}_{\mathbf{y}_1 \rightarrow x | \mathbf{y}_2}(\bar{\boldsymbol{\omega}}) = \left(\frac{\text{var}(x_t | \mathbf{y}_{2t-\tau}, \tau \in \mathbb{Z}) - \text{var}(x_t | \mathbf{y}_{t-\tau}, \tau \in \mathbb{Z})}{\text{var}(x_t)} \right) \times 100 \quad (2.21)$$

The interpretation of (2.21) is the following: it shows the amount of uncertainty about x_t left after observing the infinite sequence of past, present, and future values of \mathbf{y}_2 that is removed by observing also the infinite sequence of past, present, and future values of \mathbf{y}_1 , as a percent of the unconditional uncertainty about x_t .

Example To fix ideas, consider the following example. A latent variable of interest x_t follows a stationary AR(1) process:

$$x_t = \alpha x_{t-1} + \varepsilon_{xt} \quad (2.22)$$

The observed variables y_{1t} and y_{2t} are noisy measures of x_t , given by

$$y_{1t} = x_t + e_{1t}, \quad (2.23)$$

$$y_{2t} = x_t + e_{2t} \quad (2.24)$$

where $e_{1t} = \beta e_{1t-1} + \sqrt{1 - \beta^2} \varepsilon_{1t}$, $e_{2t} = \varepsilon_{2t}$ and ε_{xt} , ε_{1t} , and ε_{2t} are all i.i.d with mean 0 and variance 1.

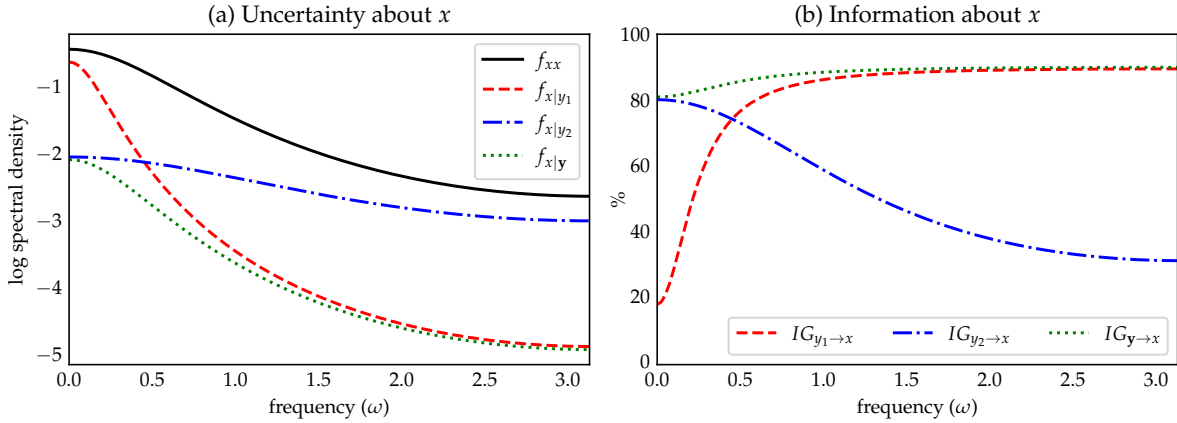


Figure 1: Frequency-specific uncertainty and information about the latent variable x in the system described by (2.22) - (2.24). Panel (a) shows prior and posterior spectral densities of x . Panel (b) shows the respective information gains.

Suppose that $\alpha = .5$, $\beta = .9$. Panel (a) of Figure 1 show the logs of the prior (unconditional) and posterior (conditional) spectral densities of x . Each point on a spectral density curve represents the contribution of the corresponding frequency to the variance of x ; the area under each curve represents the total variance of x under different information scenarios - observing nothing, observing either y_1 or y_2 or both $\mathbf{y} = [y_1, y_2]$. As both variables are informative about x , the posterior spectral densities lie below the prior one. The lowest uncertainty is obtained when both y_1 and y_2 are observed. The area between a prior and posterior spectral density curves represents the reduction of uncertainty, i.e. the information about x .⁶ The respective frequency-specific information gains, defined as the percent reduction in uncertainty, are displayed in panel (b) of the same figure. We see that y_1 is relatively less informative about the low frequencies, compared to y_2 , and more informative in the rest of the spectrum, especially the very high frequencies. This is due to the different spectral profiles of the noise terms in the

⁶Since the figure is in logs, the area represents the log ratio between prior and posterior uncertainty.

two signals: since $\text{var}(e_{1t}) = \text{var}(e_{2t}) = 1$, the total amount of noise in y_1 and y_2 is the same; however, as e_{1t} is as a very persistent AR(1) process, it contaminates mostly the very low frequencies; e_2 , on the other hand, is a white noise and therefore contributes the same amount of noise to each frequency.

The conditional information gains can be determined by comparing the marginal to the joint information gain, i.e. $\text{IG}_{y_i \rightarrow x | y_j} = \text{IG}_{\mathbf{y} \rightarrow x} - \text{IG}_{y_j \rightarrow x}$ for $i, j \in [1, 2]$. We see that $\text{IG}_{\mathbf{y} \rightarrow x} \approx \text{IG}_{y_2 \rightarrow x}$ at frequencies close to 0. This means that, conditional on observing y_2 , there is very little additional information from observing y_1 about the low frequencies of x . Similarly, conditional on observing y_1 , there is little or no additional information from y_2 about the high frequencies of x . The fact that the conditional information is smaller than the marginal implies that some of the information about x can be obtained from either variable and, having observed one of them, the information in parts of the spectrum of the other variable is redundant or nearly so. For an example where the conditional information is greater than the marginal one, consider the case when e_1 is observable. Clearly, observing e_1 alone provides no information about x . On the other hand, observing both y_1 and e_1 amounts to observing x , i.e. information gain of 100%. Therefore, the conditional information gains of both y_1 and e_1 are larger than the marginal gains.

2.3 DSGE models

A linearized DSGE model can be expressed as a recursive equilibrium law of motion given by the following system of equations:

$$\mathbf{y}_t = \mathbf{C}(\boldsymbol{\theta})\mathbf{v}_{t-1} + \mathbf{D}(\boldsymbol{\theta})\mathbf{u}_t \quad (2.25)$$

$$\mathbf{v}_t = \mathbf{A}(\boldsymbol{\theta})\mathbf{v}_{t-1} + \mathbf{B}(\boldsymbol{\theta})\mathbf{u}_t \quad (2.26)$$

$$\mathbf{u}_t = \mathbf{G}(\boldsymbol{\theta})\mathbf{u}_{t-1} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\varepsilon(\boldsymbol{\theta})) \quad (2.27)$$

where \mathbf{y}_t is a n_y vector of observed variables, \mathbf{v}_t is a n_v vector of endogenous state variables, \mathbf{u}_t is a n_u vector of exogenous state variables, and $\boldsymbol{\varepsilon}_t$ is a n_u vector of exogenous shocks. The matrices \mathbf{A} , \mathbf{B} , \mathbf{C} , \mathbf{D} , and \mathbf{G} are functions of the structural parameters of the model, collected in the n_θ vector $\boldsymbol{\theta}$.

In practice, latent variables researchers might be interested are endogenous variables, such as output gap, exogenous shocks, such as total factor productivity (TFP), or innovations to exogenous shocks, such as the innovation to the TFP shock. In other

words, and using the notation from sections 2.1 and 2.2, the latent variable x_t will be an element of \mathbf{v}_t , \mathbf{u}_t , or $\boldsymbol{\varepsilon}_t$, while the vector of observed variables is \mathbf{y}_t . Evaluating the unconditional and conditional information gain measures requires knowing the spectral and cross-spectral densities of x_t , \mathbf{y}_t , and individual elements of \mathbf{y} . Those can be obtained from the joint spectral density matrix of $\mathbf{z}_t = [\mathbf{y}'_t, \mathbf{v}'_t, \mathbf{u}'_t, \boldsymbol{\varepsilon}'_t]'$, which is given by (see Uhlig (1999)):

$$\mathbf{f}_{\mathbf{z}\mathbf{z}}(\omega) = \frac{1}{2\pi} \mathbf{W}(\omega, \boldsymbol{\theta}) \boldsymbol{\Sigma}_{\boldsymbol{\varepsilon}}(\boldsymbol{\theta}) \mathbf{W}(\omega, \boldsymbol{\theta})^* \quad (2.28)$$

where

$$\mathbf{W}(\omega, \boldsymbol{\theta}) = \begin{bmatrix} \mathbf{C}(\boldsymbol{\theta})e^{-i\omega} & \mathbf{D}(\boldsymbol{\theta}) & \mathbf{O}_{n_{\mathbf{y}}, n_{\mathbf{u}}} \\ \mathbf{I}_{n_{\mathbf{v}}} & \mathbf{O}_{n_{\mathbf{v}}, n_{\mathbf{u}}} & \mathbf{O}_{n_{\mathbf{v}}, n_{\mathbf{u}}} \\ \mathbf{O}_{n_{\mathbf{u}}, n_{\mathbf{v}}} & \mathbf{I}_{n_{\mathbf{u}}} & \mathbf{O}_{n_{\mathbf{v}}, n_{\mathbf{u}}} \\ \mathbf{O}_{n_{\mathbf{u}}, n_{\mathbf{y}}} & \mathbf{O}_{n_{\mathbf{u}}, n_{\mathbf{u}}} & \mathbf{I}_{n_{\mathbf{u}}} \end{bmatrix} \times \quad (2.29)$$

$$\begin{bmatrix} (\mathbf{I}_{n_{\mathbf{v}}} - \mathbf{A}(\boldsymbol{\theta})e^{-i\omega})^{-1} \mathbf{B}(\boldsymbol{\theta}) (\mathbf{I}_{n_{\mathbf{u}}} - \mathbf{G}(\boldsymbol{\theta})e^{-i\omega})^{-1} \\ (\mathbf{I}_{n_{\mathbf{u}}} - \mathbf{G}(\boldsymbol{\theta})e^{-i\omega})^{-1} \\ \mathbf{I}_{n_{\mathbf{u}}} \end{bmatrix}$$

and the asterisk denotes matrix transposition and complex conjugation.

In business cycle research, it is common to divide the spectrum into three non-overlapping intervals, corresponding to business cycle frequencies with periodicity between 6 and 32 quarters (as is standard in the literature, for example Stock and Watson (1999)), and frequencies above and below that interval, labeled as low and high frequencies, respectively. Let $\boldsymbol{\omega}^{BC}$, $\boldsymbol{\omega}^L$, and $\boldsymbol{\omega}^H$ denote the respective frequency bands. Then, the total information gain from \mathbf{y}_t with respect to x_t can be decomposed as follows:

$$\begin{aligned} \text{IG}_{\mathbf{y} \rightarrow x}(\bar{\boldsymbol{\omega}}) &= \text{IG}_{\mathbf{y} \rightarrow x}(\boldsymbol{\omega}^L) \frac{f_{xx}(\boldsymbol{\omega}^L)}{f_{xx}(\bar{\boldsymbol{\omega}})} + \text{IG}_{\mathbf{y} \rightarrow x}(\boldsymbol{\omega}^{BC}) \frac{f_{xx}(\boldsymbol{\omega}^{BC})}{f_{xx}(\bar{\boldsymbol{\omega}})} \\ &+ \text{IG}_{\mathbf{y} \rightarrow x}(\boldsymbol{\omega}^H) \frac{f_{xx}(\boldsymbol{\omega}^H)}{f_{xx}(\bar{\boldsymbol{\omega}})} \end{aligned} \quad (2.30)$$

Therefore, the total information gain is given by the weighted sum of the band-specific information gains, with weights equal to the contribution of each frequency band to the total variance of x .

Decomposing information gains across frequency bands is possible because of the mutual independence of the respective frequency components. Since the variables

in \mathbf{y} are typically correlated, the overall information about x cannot be decomposed into contributions of individual observed variables. What we can measure instead are the marginal contribution of a given observed variable y_i , as well as its conditional contribution given the information in other observed variables $\mathbf{y}_j \subset \mathbf{y}_{-i} \equiv \{\mathbf{y} \setminus y_i\}$. In particular, the following decomposition holds for any given frequency band ω :

$$\text{IG}_{\mathbf{y} \rightarrow x}(\omega) = \text{IG}_{y_i \rightarrow x | \mathbf{y}_{-i}}(\omega) + \text{IG}_{\mathbf{y}_{-i} \rightarrow x}(\omega) \quad (2.31)$$

The first term on the right hand side represents information in y_i about x that is not in \mathbf{y}_{-i} . Note that this includes both information that is unique to y_i , i.e. is independent from \mathbf{y}_{-i} , as well as information about x that emerges from observing y_i together with \mathbf{y}_{-i} . At the same time, some of the information in y_i about x is also in \mathbf{y}_{-i} , and is therefore captured by the second term in (2.31).

Example (continued) Figure 2 shows the decomposition of the prior and posterior




Figure 2: Contributions from the low, business cycle, and high frequencies to the prior and posterior uncertainty about x .

uncertainty about x into contributions from the low, business cycle, and the high frequencies. Observing either y_1 or y_2 results in a reduction of uncertainty by the same amount, 63%. However, as noted earlier, the contributions of the high and low frequencies are very different – around 26% of the information gain originates in the high frequencies when y_1 is observed vs. only 14% in the case of observing y_2 . Conversely, the contribution from the low frequencies is 6% with y_1 vs more than 14% with y_2 . For both variables

the largest contribution of information comes from the business cycle frequencies, 30% and 34% for y_1 and y_2 , respectively. When both variables are observed, the reduction of uncertainty is by 82% overall, with contributions by the low, business cycle, and the high frequencies of 15%, 40%, and 27%, respectively.

3 Applications

In this section, I present three examples of application of the proposed method to estimated macroeconomic models. The first two applications involve small- and medium-scale New Keynesian models taken from Uribe (2021) and Justiniano et al. (2011). Considering these models allows me to illustrate different elements of the analysis in a complementary fashion. The model of Uribe (2021) is much smaller, with only three observed variables, which makes it possible to present fully results regarding information interactions among those variables. This is not practicable in the case of the Justiniano et al. (2011) model, where I present only selected results and leave the rest for the Appendix. Another important difference is that the Uribe (2021) has more shocks than observables, and finding out how well each shock can be recovered is a relevant dimension of the analysis, in addition to investigating the main sources of information. This is not an issue in the second model, which, with its richer structure, larger number of shocks and observables, is much more representative of the medium-scale New Keynesian framework in the DSGE literature. The last example is another medium-scale New Keynesian model taken from Angeletos et al. (2018). The model incorporates many of the features found in other estimated DSGE models, but dispenses with the usual assumption of rational expectations and common information about the state of the economy. Furthermore, in contrast to most of the literature, the model is estimated in the frequency domain using only the business-cycle frequencies. Therefore, it provides an opportunity to discuss the use and usefulness of the proposed methodology in applications where there are concerns about model misspecification in some parts of the spectrum, as is the case in Angeletos et al. (2018).

3.1 Uribe (2021)

Uribe (2021) investigates the nature and empirical importance of monetary policy shocks that produce neo-Fisherian dynamics, i.e. move interest rates and inflation in the same

direction over the short run. To that end, the author estimates a standard small-scale New-Keynesian model with price stickiness and habit formation, augmented with seven structural shocks. Full details about the model can be found in the original publication. Here I only describe those of its features that are directly relevant for the analysis which follows.

Firstly, three of the shocks are to monetary policy, which is described by the following policy rule:

$$\frac{1 + I_t}{\Gamma_t} = \left[A \left(\frac{1 + \Pi_t}{\Gamma_t} \right)^{\alpha_t} \left(\frac{Y_t}{X_t} \right)^{\alpha_y} \right]^{1-\gamma_I} \left(\frac{1 + I_{t-1}}{\Gamma_{t-1}} \right)^{\gamma_I} e^{z_t^m}, \quad (3.1)$$

where I_t the nominal interest rate, Y_t is aggregate output, Π_t is the inflation rate, Γ_t is the inflation-target, X_t is a nonstationary productivity shocks, and z_t^m is a stationary interest-rate shock. The inflation target is defined as

$$\Gamma_t = X_t^m e^{z_t^{m2}}, \quad (3.2)$$

where X_t^m and z_t^{m2} are permanent and transitory components of the inflation target. It is assumed that X_t^m and X_t grow at a rates g_t^m and g_t , respectively.

There are two preference shocks affecting the lifetime utility function of the representative household, given by

$$E_0 \sum_{t=0}^{\infty} \beta^t e^{\xi_t} \left\{ \frac{[(C_t - \delta \tilde{C}_{t-1}) (1 - e^{\theta_t} h_t)^{\chi}]^{1-\sigma} - 1}{1 - \sigma} \right\}, \quad (3.3)$$

where C_t is consumption, \tilde{C}_t is the cross sectional average of consumption, h_t is hours worked, ξ_t is an intertemporal preference shock, and θ_t is a shock to labor supply.

In addition to X_t , there is also a stationary productivity shock z_t , which affects the production technology according to

$$Y_t = e^{z_t} X_t h_t^\alpha, \quad (3.4)$$

The five stationary shocks (ξ_t , θ_t , z_t , z_t^m , and z_t^{m2}) and the growth rates of the two non-stationary shocks (g_t and g_t^m) are all assumed to follow first-order autoregressive processes.

Uribe (2021) estimates the model using quarterly US data on three variables: per

capita output growth (Δy_t), the interest-rate-inflation differential ($r_t = i_t - \pi_t$), and the change in the nominal interest rate (Δi_t). All variables are assumed to be observed with measurement errors, modeled as Gaussian i.i.d. processes. Thus, there are ten independent sources of randomness in the data and only three observables. Clearly, not all, if any, of the latent variables can be recovered fully. The purpose of the remainder of this section is to determine how well each structural shock can be recovered and where in the spectrum most of the information comes from, as well as what are the information contributions of different observed variables overall and across different frequency bands.

3.1.1 Information decomposition across frequency bands

Uribe (2021) solves the model by log-linear approximation of the equilibrium conditions around steady state. The linearity of the solution together with the assumption that the structural innovations and the measurement errors are Gaussian, implies that the joint distribution of (any subset of) the innovations, shocks, state and observed variables is also Gaussian. Therefore, the analysis of information gains can be conducted using the measures introduced in Section 2. In the analysis which follows I fix the parameter values at the mean of posterior distribution reported in Uribe (2021, Table 5).

Table 1 presents the total information gains for the seven shocks and their decompositions into gains from three frequency bands - low, business cycle and high frequencies, with periodicities of more than 32 quarters, between 6 and 32 quarters and less than 6 quarters, respectively. The results show that none of the shocks can be fully recovered from the observed variables. The largest reduction of uncertainty is with respect to the intertemporal preference shock (ξ_t) – by about 93%, and the permanent productivity shock (g_t) – by about 85%. The gains with respect to the three monetary policy-related shocks are between 15% and 18%. The least information is gained with respect to the labor supply (θ_t) and the transitory productivity shocks (z_t), with information gains for both of 1.8%.

Columns 3 to 5 of the table show the information gain contributions from each frequency band. Following the earlier discussion (see equation (2.30)), the total contribution in each case is shown as the product of two terms: the band-specific information gain, which measures the reduction of uncertainty as a percent of the uncertainty in that band, and the fraction of total uncertainty that originates in the given frequency band.

For six of the seven shocks uncertainty is concentrated in either low and business cycle frequencies, or high and business cycle frequencies. Specifically, in the first groups are

Table 1: Information decomposition across frequency bands

| | total | low | BC | high |
|---------------------------------------|-------|---------------------------|---------------------------|---------------------------|
| ξ_t preference | 93.2 | 70.4 = 96.4 \times 0.73 | 19.5 = 88.4 \times 0.22 | 3.2 = 66.0 \times 0.05 |
| θ_t labor supply | 1.8 | 0.2 = 0.5 \times 0.33 | 1.1 = 2.3 \times 0.48 | 0.5 = 2.9 \times 0.18 |
| z_t transitory productivity | 1.8 | 0.2 = 0.5 \times 0.32 | 1.1 = 2.2 \times 0.49 | 0.5 = 2.9 \times 0.19 |
| g_t permanent productivity | 83.5 | 9.3 = 94.9 \times 0.10 | 32.3 = 87.1 \times 0.37 | 42.0 = 78.9 \times 0.53 |
| z_t^m transitory interest rate | 15.5 | 0.1 = 0.9 \times 0.12 | 3.2 = 7.9 \times 0.41 | 12.2 = 25.7 \times 0.47 |
| z_t^{m2} transitory trend inflation | 16.5 | 5.8 = 12.7 \times 0.46 | 9.7 = 23.1 \times 0.42 | 1.0 = 8.3 \times 0.12 |
| g_t^m permanent trend inflation | 18.0 | 7.2 = 69.4 \times 0.10 | 7.0 = 18.3 \times 0.38 | 3.9 = 7.5 \times 0.51 |

Note: Information gain (IG) measures the reduction of uncertainty (variance) about a shock due to observing all three observed variables, as a *percent* of the unconditional uncertainty of the shock. The contribution from each frequency band to the total IG is shown as the product of the IG for that band and the fraction of the total variance of the shock originating in each band. Thus, the units in the last three columns are $\% = \% \times \frac{\text{variance band}}{\text{variance total}}$.

the transitory trend inflation, transitory productivity, and the labor supply shocks. And in the second are the permanent productivity, transitory interest rate, and permanent trend inflation shocks. The one exception is the intertemporal preference shock for which the low frequencies are by far the main source of uncertainty. As can be expected, the largest gains generally come from parts of the spectrum where prior uncertainty is larger. There are some notable exceptions, however. In particular, note that even though the low frequency band accounts for only 10% of the uncertainty about the permanent trend inflation shock, the information gain contribution from that band is largest than the business cycle frequency band, which accounts for 38% of the uncertainty, and much larger than the contribution from the high frequency band, which accounts for more than half of the total uncertainty. This is due to the fact that a much larger fraction of the uncertainty in the low frequencies is resolved by information provided by the observed variables than is the case for the higher frequencies. Similarly, note that for the labor supply and transitory productivity shocks, because of the relatively larger information gains from the higher end of the spectrum, the information contributions from there is larger than from the low frequencies, even though the low frequencies account for a significantly larger fraction of the prior uncertainty.

3.1.2 Information contributions by variables

Table 2 shows the conditional information gains for each observed variable for the full spectrum and the three frequency bands. The largest contribution by far is from

Table 2: Conditional contribution of information

| shock | | total | | | low | | | BC | | | high | | |
|------------|----------------------------|--------------|-------|--------------|--------------|-------|--------------|--------------|-------|--------------|--------------|-------|--------------|
| | | Δy_t | r_t | Δi_t | Δy_t | r_t | Δi_t | Δy_t | r_t | Δi_t | Δy_t | r_t | Δi_t |
| ξ_t | preference | 0.3 | 26.8 | 7.2 | 0.0 | 26.4 | 0.8 | 0.1 | 0.5 | 4.1 | 0.1 | 0.0 | 2.3 |
| θ_t | labor supply | 0.1 | 0.1 | 1.1 | 0.0 | 0.0 | 0.0 | 0.1 | 0.0 | 0.6 | 0.0 | 0.0 | 0.5 |
| z_t | transitory productivity | 0.1 | 0.0 | 1.1 | 0.0 | 0.0 | 0.0 | 0.1 | 0.0 | 0.6 | 0.0 | 0.0 | 0.5 |
| g_t | permanent productivity | 83.4 | 0.8 | 5.7 | 9.3 | 0.0 | 0.1 | 32.2 | 0.6 | 2.8 | 41.9 | 0.2 | 2.8 |
| z_t^m | transitory interest rate | 2.2 | 1.5 | 9.0 | 0.0 | 0.1 | 0.0 | 0.4 | 0.9 | 0.4 | 1.8 | 0.5 | 8.5 |
| z_t^{m2} | transitory trend inflation | 1.7 | 13.0 | 8.2 | 0.1 | 5.5 | 4.3 | 1.1 | 7.3 | 3.7 | 0.5 | 0.2 | 0.2 |
| g_t^m | permanent trend inflation | 0.5 | 10.4 | 15.6 | 0.0 | 4.7 | 7.0 | 0.2 | 5.2 | 5.6 | 0.3 | 0.5 | 3.0 |

Note: The conditional contribution of information shows additional reduction of uncertainty about a shock, as a percent of the unconditional uncertainty of the shock, due to observing a variable given that the other two variables are also observed. The variables are: output growth (Δy_t), interest-rate-inflation differential (r_t), and the change in the nominal interest rate (Δi_t). Due to rounding in some cases the band-specific contributions do not add up to the total values.

output growth (Δy_t) with respect to the permanent productivity shock. Note that the conditional information gain of 83.4% is almost equal to the total gain (all observables) of 83.5% for that shock (see Table 1). This implies that the other two variables - the interest rate-inflation differential (r_t) and the change in the nominal interest rate (Δi_t) alone reduce the uncertainty about the permanent productivity shock by only 0.1%. This result holds for the full spectrum and the individual frequency bands. Output growth contributes less information for the other shocks, compared to r_t or Δi_t . The contributions of these variables with respect to the two trend inflation shocks are similar, with r_t being relatively more informative for the transitory trend inflation shock, while Δi_t is more informative for the permanent one. In addition, r_t contributes much more information than either Δy_t or Δi_t with respect to the preference shock, while Δi_t is the most informative observable with respect to the transitory interest rate shock, and, marginally, for the labor supply and transitory productivity shocks.

The ranking of variables in terms of their total information contributions is determined by the relative size of the information gains in the part of the spectrum from where a given shock receives the most total information (see Table 1). In several cases, the ranking changes with the frequency band. For instance, Δi_t contributes significantly more information than r_t with respect to the intertemporal preference shock in the BC and high frequencies. At the same time, r_t contributes the most information with respect to the transitory interest rate shock in the low and BC frequencies, in spite of being the least informative variable in the high frequencies and overall. Similarly, Δy_t is the least

Table 3: Unconditional contribution of information

| shock | | total | | | low | | | BC | | | high | | |
|------------|----------------------------|--------------|-------|--------------|--------------|-------|--------------|--------------|-------|--------------|--------------|-------|--------------|
| | | Δy_t | r_t | Δi_t | Δy_t | r_t | Δi_t | Δy_t | r_t | Δi_t | Δy_t | r_t | Δi_t |
| ξ_t | preference | 3.5 | 84.6 | 66.0 | 0.9 | 69.6 | 44.0 | 2.0 | 14.6 | 18.9 | 0.6 | 0.4 | 3.1 |
| θ_t | labor supply | 0.0 | 0.6 | 1.7 | 0.0 | 0.1 | 0.2 | 0.0 | 0.5 | 1.0 | 0.0 | 0.0 | 0.5 |
| z_t | transitory productivity | 0.0 | 0.6 | 1.6 | 0.0 | 0.1 | 0.1 | 0.0 | 0.5 | 1.0 | 0.0 | 0.0 | 0.5 |
| g_t | permanent productivity | 76.7 | 0.1 | 0.1 | 9.0 | 0.0 | 0.0 | 28.7 | 0.0 | 0.0 | 39.1 | 0.0 | 0.0 |
| z_t^m | transitory interest rate | 0.7 | 5.8 | 11.5 | 0.0 | 0.1 | 0.0 | 0.2 | 2.5 | 1.8 | 0.5 | 3.2 | 9.7 |
| z_t^{m2} | transitory trend inflation | 2.2 | 5.3 | 0.9 | 0.2 | 1.1 | 0.1 | 1.6 | 3.8 | 0.6 | 0.4 | 0.4 | 0.2 |
| g_t^m | permanent trend inflation | 1.8 | 0.4 | 6.8 | 0.1 | 0.0 | 2.5 | 0.9 | 0.3 | 1.3 | 0.8 | 0.1 | 3.0 |

Note: see the note to Table 2. The unconditional contribution of information shows the reduction of uncertainty about a shock due to observing a single variable at a time.

informative variable overall with respect to the transitory trend inflation shock, but has the largest contribution in the high frequency band.

It is worth emphasizing that the information gains shown in Table 2 are from observing a given variable *conditional* on already having observed the other two variables. As the observed variables are obviously not mutually independent, it is conceivable that in some cases the contributions are small because different variables share common information with respect to those shocks. To help find out if and when that is the case, Table 3 shows the unconditional information gains, i.e. the percent reduction of uncertainty about a given shock due to observing only one variable at a time.

The results reveal some notable differences between conditional and unconditional information gains. Most striking is the reduction in the contributions of the three observables with respect to the intertemporal preference shock. In particular, the information gains from r_t and Δi_t change from, respectively, 85% and 66% unconditionally, to 27% and 7% conditionally. Similarly, the contribution of Δy_t decreases from 3.5% to only 0.3%. This suggests that, to a large extent, the information in either one of the observable variables is not unique to them but is also contained in the other two. In other words, there is a significant degree of redundancy of the information about the intertemporal preference shock. Another, less striking, example of redundancy is the transitory interest rate shock, where the conditional information gains from r_t and Δi_t are smaller than the unconditional ones.

Information redundancy is not the only possible consequence of the existing interdependence among observables. In the case of the permanent productivity shock, the conditional information gains for all variables are significantly larger than the uncon-

ditional ones. The same is true for the contributions of r_t and Δi_t with respect to the permanent and transitory trend inflation shocks, as well as for the contribution of Δy_t with respect to the transitory interest rate shock. In all of these cases there is a positive information complementarity instead of information redundancy, that is, information increases when variables are observed together.

Following Iskrev (2019), the degree of information complementarity between variables can be measured by comparing the joint information gain with respect to a shock to the individual gains. Specifically, the information complementarity between variables y_1 and y_2 conditional on variables $\mathbf{y}_3 \subset \{\mathbf{y} \setminus \mathbf{y}_{12}\}$ at frequency band $\boldsymbol{\omega}$ is defined as:

$$\text{IC}_{\mathbf{y}_{12} \rightarrow x | \mathbf{y}_3}(\boldsymbol{\omega}) = \frac{\text{IG}_{\mathbf{y}_{12} \rightarrow x | \mathbf{y}_3}(\boldsymbol{\omega})}{\text{IG}_{y_1 \rightarrow x | \mathbf{y}_3}(\boldsymbol{\omega}) + \text{IG}_{y_2 \rightarrow x | \mathbf{y}_3}(\boldsymbol{\omega})} - 1. \quad (3.5)$$

Negative values indicate negative complementarity, or information redundancy, between y_1 and y_2 , and positive values indicate positive complementarity between the two variables. Since the information gain is non-negative, we have $\text{IC}_{\mathbf{y}_{12} \rightarrow x | \mathbf{y}_3}(\boldsymbol{\omega}) \geq -1/2$, with equality when y_1 and y_2 are (conditionally on \mathbf{y}_3) functionally dependent, in which case $\text{IG}_{\mathbf{y}_{12} \rightarrow x | \mathbf{y}_3}(\boldsymbol{\omega}) = \text{IG}_{y_1 \rightarrow x | \mathbf{y}_3}(\boldsymbol{\omega}) = \text{IG}_{y_2 \rightarrow x | \mathbf{y}_3}(\boldsymbol{\omega})$. A lack of information complementarity, i.e. $\text{IC}_{\mathbf{y}_{12} \rightarrow x | \mathbf{y}_3}(\boldsymbol{\omega}) = 0$ occurs when y_1 and y_2 are (conditionally on \mathbf{y}_3) independent, and hence $\text{IG}_{\mathbf{y}_{12} \rightarrow x | \mathbf{y}_3}(\boldsymbol{\omega}) = \text{IG}_{y_1 \rightarrow x | \mathbf{y}_3}(\boldsymbol{\omega}) + \text{IG}_{y_2 \rightarrow x | \mathbf{y}_3}(\boldsymbol{\omega})$. Note that the conditioning could be with respect to any subset of observables, including the empty set, in which case we have unconditional complementarity between y_1 and y_2 .

Figure 3 shows the unconditional and conditional information complementarities between all pairs of variables. The results are shown for the full spectrum as well as the three frequency bands. As already anticipated, the strongest complementarity overall is between r_t and Δi_t , and is negative for all shocks except the permanent and transitory trend-inflation shocks. Both unconditionally and conditionally the degree of complementarity tends to be significantly lower in the higher frequencies. Conditioning on the third observable in most cases preserves the sign of complementarity and reduces the magnitude. There are some notable exceptions to this pattern, however. For instance, the degree of complementarity between r_t and Δi_t increases when conditioning on Δy_t ,

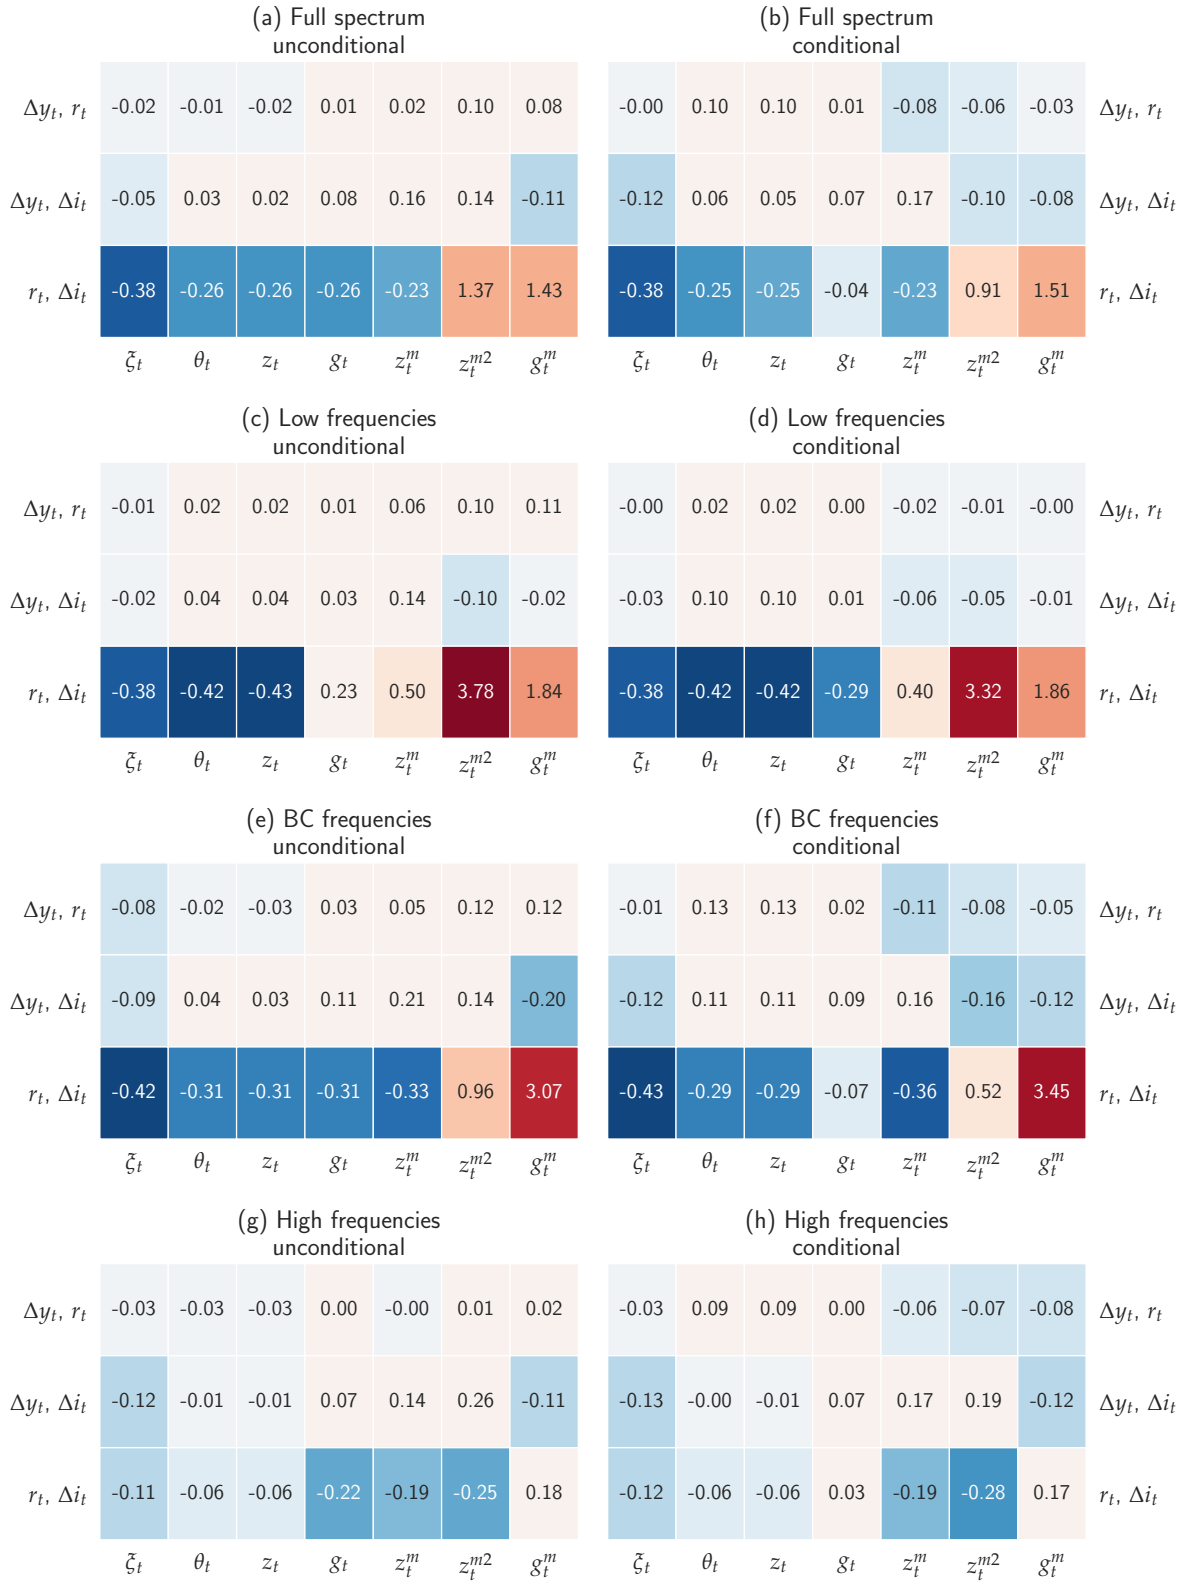


Figure 3: Pairwise information complementarity between observables with respect to shocks.

especially in the business cycle frequencies. Furthermore, the complementarity between the same variables with respect to the transitory productivity shock changes signs when conditioning on y_t , from positive to negative in the low frequencies, and from negative to positive in the high frequencies. At the same time, when evaluated over the full spectrum, the complementarity is strongly negative unconditionally and only weakly so, conditionally.

3.1.3 Information gains in the time domain

The time domain version of the full spectrum information gain measure (see equation (2.20)) is given by:

$$\text{IG}_{\mathbf{Y}_T \rightarrow x_t} = \left(\frac{\text{var}(x_t) - \text{var}(x_t | \mathbf{Y}_T)}{\text{var}(x_t)} \right) \times 100, \quad (3.6)$$

where $1 \leq t \leq T$ and $\mathbf{Y}_T = \{\mathbf{y}_1, \dots, \mathbf{y}_T\}$. The difference between the two measures is that, in the frequency domain, the information for any given x_t stems from the infinite past and future values of the observable variables. Therefore, for a given set of observed variables, the total amount of information is invariant to the temporal location of the latent variable. In contrast, in the time domain, it matters where the location of t is, relative to the beginning and the end of the sample. Thus, the value of time domain measure changes with t and is bounded from above by the value of the full spectrum frequency domain measure.

Figure 4 compares the time and frequency domain information gains for the seven shocks in the model. Specifically, it shows the ratio of the time domain to the frequency domain measure for all values of t in a sample of $T = 255$ observations, which is the sample size in Uribe (2021). The results show that for most values of t the time and frequency domain information gains coincide. As anticipated, differences occur only at the beginning and end of the sample. For all shocks except the transitory trend inflation shock, for which convergence is somewhat slower, there are about ten observations or fewer on either end of the sample where the time domain information gains are smaller than the frequency domain ones.

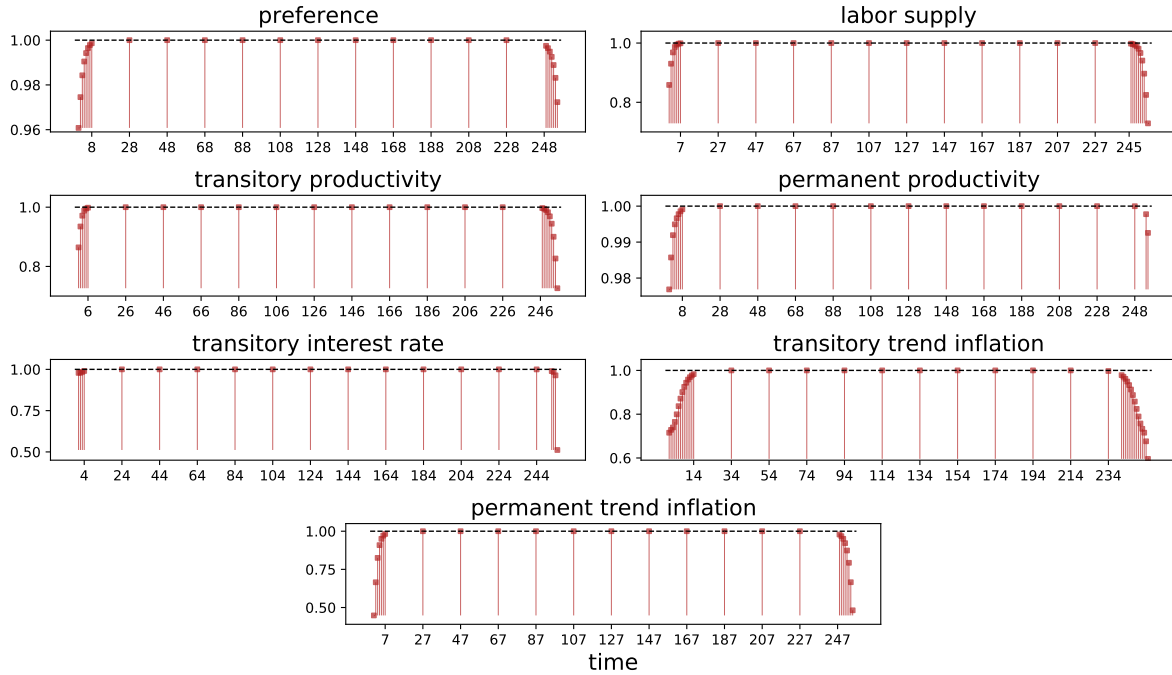


Figure 4: Total information gains in the time domain relative to the frequency domain (full spectrum).

3.1.4 Discussion of the results

As already noted, having more sources of uncertainty than the number of observed variables necessarily implies that the latent variables in the model cannot all be recovered fully. At the same time, as the results presented in Section 3.1.1 show, some shocks in the Uribe (2021) model are significantly better recoverable than others. The goal of this section is to develop a further understanding of these findings.

A natural question to ask is: why are the information gains with respect to the intertemporal preference and permanent productivity shocks so much larger than the gains for the remaining shocks, and in particular compared to those with respect to the labor supply and transitory productivity shocks? Intuitively, the amount of information one or more variables contain about another variable depends on the strength of their mutual dependence.⁷ Furthermore, an insight gained from the frequency domain perspective is that the interactions need to be strong in the parts of the spectrum that are mainly responsible for the uncertainty of the latent variable. In addition, the extent to which

⁷In fact, the mutual information coefficient is commonly used to measure and test for statistical dependence between random variables (see e.g. Linfoot (1957), Joe (1989), and Granger and Lin (1994)).

information from multiple sources accumulates, in turn, depends on how interdependent they are among themselves. For instance, variables that are functionally dependent on other observed variables provide no useful information.⁸

Consider the intertemporal preference shock (ξ_t). According to the posterior mean estimates reported in Uribe (2021, Table 5), this shock is significantly more persistent and volatile than all other shocks. In particular, its volatility is an order of magnitude larger than the volatilities of all other shocks except the permanent productivity shock (g_t). The high degree of persistence explains why most of the uncertainty about ξ_t is concentrated in the lower end of the spectrum, as shown in Table 1. Furthermore, as seen from the same table, most of the uncertainty in the low frequencies is resolved by the information contained in the observed variables, which suggests that there are strong interactions between ξ_t and (some of) those variables. Since, in the present context, the variables have a clear causal direction, i.e. from shocks to endogenous variables, a natural way of describing their interactions is in terms of the shocks' impact on the observed variables. A convenient measure of the size of the total impact is each shock's contribution to the total variance of each variable. Figure 5 shows the individual contributions of the shocks as a percent of the total variances of the observables, as well as decompositions of the individual and total contributions in the low, BC, and high frequency bands. Note that the measurement errors also contribute to the variances, which is why the total contributions of the shocks sum up to less than 100%.

The results show that ξ_t drives most of the volatility in two of the observed variables – r_t and Δi_t , and, in the case of r_t , the contribution is mostly in the low frequencies. This

⁸An example of this is output growth in the model estimated by Schmitt-Grohé and Uribe (2012), see Iskrev (2019) for details.

| | Δy | r | Δi | |
|------------|------------|------|------------|-------|
| all shocks | 93.7 | 84.1 | 96.3 | total |
| | 8.7 | 64.8 | 8.8 | low |
| | 37.2 | 17.3 | 40.7 | BC |
| | 47.8 | 2.0 | 46.8 | high |
| ζ | 10.9 | 77.3 | 72.4 | total |
| | 0.2 | 62.7 | 7.3 | low |
| | 4.3 | 13.8 | 34.8 | BC |
| | 6.5 | 0.8 | 30.3 | high |
| θ | 0.1 | 0.4 | 2.4 | total |
| | 0.0 | 0.2 | 0.1 | low |
| | 0.0 | 0.2 | 0.9 | BC |
| | 0.0 | 0.0 | 1.4 | high |
| z | 0.1 | 0.4 | 2.3 | total |
| | 0.0 | 0.2 | 0.1 | low |
| | 0.0 | 0.2 | 0.9 | BC |
| | 0.0 | 0.0 | 1.4 | high |
| g | 76.6 | 0.0 | 0.1 | total |
| | 8.4 | 0.0 | 0.0 | low |
| | 30.1 | 0.0 | 0.0 | BC |
| | 38.2 | 0.0 | 0.0 | high |
| z^m | 0.7 | 2.0 | 11.9 | total |
| | 0.0 | 0.2 | 0.0 | low |
| | 0.2 | 0.9 | 1.7 | BC |
| | 0.5 | 0.8 | 10.1 | high |
| z^{m2} | 3.5 | 3.6 | 1.6 | total |
| | 0.1 | 1.3 | 0.0 | low |
| | 1.7 | 1.9 | 0.8 | BC |
| | 1.7 | 0.4 | 0.8 | high |
| g^m | 1.8 | 0.4 | 5.8 | total |
| | 0.1 | 0.2 | 1.4 | low |
| | 0.9 | 0.2 | 1.5 | BC |
| | 0.8 | 0.0 | 2.8 | high |
| | Δy | r | Δi | |

Figure 5: Total and individual contributions of the shocks as a percent of the variances of the observables in the full spectrum and the low, business cycle, and high frequency bands. The difference to 100% is accounted for by the measurement error variances.

is consistent with the earlier findings that, of the three observed variables, r_t is the most informative and Δy_t – the least informative one. Similarly, the second best recoverable shock – to permanent productivity, is responsible for the bulk of the volatility of the third variable – Δy_t , and particularly in the BC and high frequencies, which, as seen in Table 1, is also where most of the uncertainty of that shock stems from. The variance contributions of the remaining five shocks are significantly smaller, and account for only between 12%, in the case of transitory interest rate shock (z_t^m) with respect to Δi_t , and 2.3% – 2.4% in the case of both the labor supply (θ_t) and transitory productivity (z_t) shocks with respect again to Δi_t .

Equivalence between variance and information decompositions. Variance decompositions in dynamic structural models are typically obtained by shutting-off all shocks but one at a time and then computing the endogenous variables’ variances or spectral densities (see for instance Fernández-Villaverde et al. (2016, Section 8)). This gives the contribution of each shock to the total variances or spectral densities of the endogenous variables. It is easy to see that the same quantities can be obtained using the information gain measures introduced in Section 2. Specifically, a shock’s contribution to the variance of a variable is equal to the information gained, i.e. the reduction in variance, with respect to the variable due to knowing that shock. In other words, instead of information from observed variables to shocks, we measure the flow of information in the opposite direction – from shocks to observables. Of course, this only works when the shocks are mutually independent, which is also the assumption behind the standard variance decomposition approach. If shocks are mutually dependent one has to distinguish between conditional and unconditional variance contributions, as in the case of information from observed variables with respect to shocks.

To summarize, as expected, there is a clear link between, on the one hand, the shocks’ contributions to the observed variables’ volatilities and, on the other hand, the degree to which each shock can be recovered from information contained in those variables. At the same time, it is important to point out that the size of the contributions is not necessarily a good indicator of the variables’ importance as sources of information about the shocks. For instance, the intertemporal preference shock contributes similar fractions of the variances of r_t and Δi_t . Yet, r_t is significantly more informative than Δi_t about that shock. As noted earlier, this is due to the fact that the variance contributions are

in different parts of the spectrum – the low frequencies in the case of r_t , and the BC and high frequencies, in the case of Δi_t . Since most of the variance of the preference shock comes from the low frequencies, r_t is significantly more informative than Δi_t . In other cases, it is the information interactions among the observed variables that affect their relative importance as sources of information. For instance, as can be seen in Table 2, the conditional contribution of information by Δi_t with respect to the transitory trend inflation shock (z_t^{m2}) is much larger than that of Δy_t , in spite of the significantly larger fraction of the variance of Δy_t attributed to that shock, compared to Δi_t . This is explained by the strong positive complementarity between r_t and Δi_t in the BC and especially the low frequencies, which is where most of the uncertainty of that shocks is located. Lastly, small variance contributions of a shock does not necessarily imply that the shock cannot be recovered. In general, having the same number of non-redundant observables as the number of sources of uncertainty means that all shocks are fully recoverable. This is the case in the model I consider next.

3.2 Justiniano, Primiceri, and Tambalotti (2011)

Justiniano et al. (2011) (henceforth JPT) investigate whether investment shocks are important drivers of business cycle fluctuations. To that end, and expanding on their previous work in Justiniano et al. (2010), they estimate a New Keynesian model featuring imperfectly competitive goods and labor markets, as well as different nominal and real frictions such as sticky prices and wages, habit formation in consumption, variable capital utilization and investment adjustment costs. As in the previous section, here I outline only those features of the model that are relevant for the information decomposition analysis that follows.

The model has eight structural shocks in total, with three technology shocks, two of which are related to investment. Specifically, JPT distinguish between final and intermediate consumption, investment, and capital goods, each being produced in a different sector. They introduce a shock that affects the transformation of consumption into investment goods, and another shock that affects the transformation of investment goods into productive capital. The first, called investment-specific technology (IST) shock, is introduced via the production function in the investment good producing sector:

$$I_t = \Upsilon_t Y_t^I, \tag{3.7}$$

where I_t is the quantity of investment goods in efficiency units produced with Y_t^I units of the final good. Υ_t represents the IST and is assumed to be a non-stationary random process growing at a rate v_t .

The second investment technology shock is introduced via the production technology in the capital good producing sector, which assumes that new capital, denoted with i_t , is produced from investment goods according to

$$i_t = \mu_t \left(1 - S \left(\frac{I_t}{I_{t-1}} \right) \right) \quad (3.8)$$

where S is an investment adjustment cost function, and μ_t is a stationary shock to the marginal efficiency of investment (MEI), assumed to be an AR(1) process.

The third technology shocks affects the production functions in the intermediate good producing sector according to:

$$Y_t(i) = \max\{A_t^{1-\alpha} K_t(i)^\alpha L_t(i)^{1-\alpha} - A_t \Upsilon_t^{\frac{\alpha}{1-\alpha}} F; 0\} \quad (3.9)$$

where $Y_t(i)$, $K_t(i)$, and $L_t(i)$ are the quantities of output produced, and effective capital and labor employed by intermediate good producer i . F represents fixed cost of production, and A_t is a common non-stationary neutral technology process, growing at rate z_t .

The final consumption good Y_t is produced by combining a continuum of intermediate goods, according to

$$Y_t = \left[\int_0^1 Y_t(i)^{\frac{1}{1+\lambda_{p,t}}} \right]^{1+\lambda_{p,t}} \quad (3.10)$$

where $\lambda_{p,t}$ is a stationary price markup shock following ARMA(1,1) process.

Similarly to the model in the previous section, there is a shock to the intertemporal preferences of the households populating the economy whose lifetime utility function is given by

$$E_0 \sum_{t=0}^{\infty} \beta^t b_t \left\{ \log(C_t - hC_{t-1}) - \varphi \frac{L_t(j)^{1+\nu}}{1+\nu} \right\}, \quad (3.11)$$

where C_t is consumption, b_t is the stationary intertemporal preference shock, assumed to follow an AR(1) process. JPT assume that there is a continuum of households $j \in [0, 1]$,

each one being a supplier of specialized labor denoted by $L_t(j)$. The specialized labor in turn is combined into homogenous labor input according to

$$L_t = \left[\int_0^1 L_t(i)^{\frac{1}{1+\lambda_{w,t}}} \right]^{1+\lambda_{w,t}} \quad (3.12)$$

where $\lambda_{w,t}$ is a stationary wage markup shock assumed to follow an ARMA(1,1) process.

The last two shocks are to government fiscal and monetary policy. Public spending G_t is a time-varying fraction of output,

$$G_t = \left(1 - \frac{1}{g_t} \right) Y_t \quad (3.13)$$

where the government spending shock g_t as a stationary AR(1) process.

Monetary policy consists of setting the nominal interest rate R_t according to the following policy rule:

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R} \right)^{\rho_R} \left[\left(\frac{\pi_t}{\pi} \right)^{\phi_\pi} \left(\frac{X_t}{X_t^*} \right)^{\phi_X} \right]^{1-\rho_R} \left[\frac{X_t/X_{t-1}}{X_t^*/X_{t-1}^*} \right]^{\phi_{dX}} \varepsilon_{mp,t}, \quad (3.14)$$

where $e_{mp,t}$ is the monetary policy shock, R is the steady state of the nominal rate, π_t is the inflation rate, $X_t = C_t + I_t + G_t$ is actual real GDP and X_t^* is the level of GDP under flexible prices and wages and in the absence of markup shocks.

To summarize, there are eight shocks in the model, six stationary and two non-stationary. Two of the stationary shocks – to price and wage markups, follow ARMA(1,1) processes, and one – to monetary policy, is an i.i.d process. The remaining stationary shocks – to government spending, MEI, and intertemporal preferences, as well as the growth rates of the two non-stationary shocks – to IST and neutral technology, follow AR(1) processes. The disturbances to all shocks are assumed to be Gaussian, leading to a linear Gaussian state space representation of the solution of log-linear approximation of model.

JPT estimate the model using US data on hours worked ($h_t = \log L_t$), inflation (π_t), the nominal interest rate (R_t), and the growth rates of GDP ($x_t = \Delta \log X_t$), consumption ($c_t = \Delta \log C_t$), investment ($i_t = \Delta \log I_t$), real wages ($w_t = \Delta \log \frac{W_t}{P_t}$), and the relative price of investment ($\pi_t^i = \Delta \log \frac{P_{I_t}}{P_t}$). Unlike Uribe (2021), they do not allow for measurement errors in any of the series. As seen below, this implies that all eight shocks can be recovered fully with the information in the eight observed variables. In

the remainder of this section I investigate the main sources of information for each shock in terms of observed variables and parts of the spectrum.

3.2.1 Information decomposition across frequency bands

Table 4 presents the total information gains for the eight shocks and their decompositions into gains from the low, BC, and high frequencies. As noted earlier, all shocks can be fully recovered from information in the observables, in the full spectrum as well as within each frequency band. The information contributions from the bands reflect the fraction of each shock’s variance originating in those bands.

Table 4: Information decomposition across frequency bands

| | shock | total | low | BC | high |
|--------------------|--------------------|-------|--------------------------|--------------------------|--------------------------|
| z | neutral technology | 100 | 11.2 = 100×0.11 | 40.0 = 100×0.40 | 48.8 = 100×0.49 |
| g | government | 100 | 96.1 = 100×0.96 | 3.2 = 100×0.03 | 0.7 = 100×0.01 |
| v | IST | 100 | 8.4 = 100×0.08 | 33.6 = 100×0.34 | 58.0 = 100×0.58 |
| λ_p | price mark-up | 100 | 51.7 = 100×0.52 | 16.1 = 100×0.16 | 32.2 = 100×0.32 |
| λ_w | wage mark-up | 100 | 5.1 = 100×0.05 | 27.3 = 100×0.27 | 67.6 = 100×0.68 |
| b | preference | 100 | 22.8 = 100×0.23 | 49.9 = 100×0.50 | 27.4 = 100×0.27 |
| ε_{mp} | monetary policy | 100 | 6.3 = 100×0.06 | 27.1 = 100×0.27 | 66.7 = 100×0.67 |
| μ | MEI | 100 | 47.4 = 100×0.47 | 40.8 = 100×0.41 | 11.7 = 100×0.12 |

Note: see the note to Table 1.

For six of the eight shocks uncertainty is distributed monotonically across the frequency bands, i.e. increases or decreases moving from low to high frequencies. Only for one of them – the government spending shock, uncertainty is concentrated in a single frequency band – the low frequencies, contributing 96% of the total variance. In the case of the MEI shock, most of the uncertainty is in the low and BC frequencies. For the other two technology shocks – neutral and IST, as well as the wage mark-up and the monetary policy shocks, uncertainty is mostly in the BC and high frequencies. In the case of the intertemporal preference shock, half of the uncertainty is in the BC frequencies, and the rest is divided almost evenly between the low and high frequencies. The other shock with a non-monotonic distribution of uncertainty is the price mark-up shocks, for which about half of the variance is due to the low frequencies, with a significant contribution by the high frequencies, and the least amount of uncertainty due to the BC frequencies.

Table 5: Conditional contribution of information

| shocks | | total | | | | | | | | low | | | | | | | |
|--------------------|--------------------|-------|------|------|------|------|-------|------|---------|------|------|------|-----|------|-------|-----|---------|
| | | x | c | i | h | w | π | R | π^i | x | c | i | h | w | π | R | π^i |
| z | neutral technology | 15.6 | 0.0 | 0.2 | 46.4 | 0.0 | 0.0 | 0.0 | 0.1 | 0.9 | 0.0 | 0.1 | 1.5 | 0.0 | 0.0 | 0.0 | 0.0 |
| g | government | 45.5 | 52.8 | 18.3 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 42.6 | 49.9 | 14.8 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 |
| v | IST | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 97.2 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 8.1 |
| λ_p | price mark-up | 13.7 | 0.2 | 0.3 | 21.4 | 29.3 | 32.4 | 0.2 | 0.2 | 11.0 | 0.0 | 0.2 | 9.1 | 27.3 | 0.5 | 0.0 | 0.1 |
| λ_w | wage mark-up | 0.8 | 0.2 | 0.4 | 1.4 | 93.2 | 23.3 | 0.3 | 0.3 | 0.1 | 0.0 | 0.0 | 0.1 | 1.6 | 1.4 | 0.1 | 0.0 |
| b | preference | 1.4 | 28.5 | 7.3 | 11.2 | 2.5 | 0.7 | 5.6 | 0.0 | 1.0 | 4.2 | 6.1 | 6.6 | 2.3 | 0.6 | 3.4 | 0.0 |
| ε_{mp} | monetary policy | 0.3 | 3.1 | 0.2 | 10.2 | 0.1 | 12.1 | 92.6 | 0.0 | 0.1 | 0.1 | 0.0 | 1.5 | 0.0 | 4.4 | 4.8 | 0.0 |
| μ | MEI | 0.1 | 0.0 | 8.7 | 0.4 | 2.2 | 0.4 | 5.2 | 1.9 | 0.0 | 0.0 | 3.9 | 0.1 | 2.0 | 0.1 | 3.2 | 1.2 |

| shocks | | BC | | | | | | | | high | | | | | | | |
|--------------------|--------------------|-----|------|-----|------|------|-------|------|---------|------|------|-----|------|------|-------|------|---------|
| | | x | c | i | h | w | π | R | π^i | x | c | i | h | w | π | R | π^i |
| z | neutral technology | 4.5 | 0.0 | 0.0 | 13.4 | 0.0 | 0.0 | 0.0 | 0.0 | 10.2 | 0.0 | 0.0 | 31.6 | 0.0 | 0.0 | 0.0 | 0.0 |
| g | government | 2.5 | 2.6 | 3.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.5 | 0.3 | 0.5 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 |
| v | IST | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 33.1 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 56.1 |
| λ_p | price mark-up | 1.5 | 0.0 | 0.0 | 6.7 | 1.9 | 7.9 | 0.1 | 0.0 | 1.2 | 0.1 | 0.1 | 5.7 | 0.1 | 24.0 | 0.1 | 0.0 |
| λ_w | wage mark-up | 0.3 | 0.1 | 0.1 | 0.4 | 24.6 | 8.7 | 0.2 | 0.2 | 0.4 | 0.1 | 0.2 | 1.0 | 67.1 | 13.2 | 0.1 | 0.1 |
| b | preference | 0.1 | 10.0 | 1.1 | 2.7 | 0.2 | 0.2 | 1.8 | 0.0 | 0.2 | 14.3 | 0.1 | 2.0 | 0.0 | 0.0 | 0.4 | 0.0 |
| ε_{mp} | monetary policy | 0.1 | 0.8 | 0.1 | 4.8 | 0.0 | 4.1 | 24.8 | 0.0 | 0.1 | 2.2 | 0.2 | 3.9 | 0.0 | 3.6 | 63.0 | 0.0 |
| μ | MEI | 0.0 | 0.0 | 1.7 | 0.2 | 0.1 | 0.2 | 1.6 | 0.3 | 0.0 | 0.0 | 3.2 | 0.2 | 0.0 | 0.1 | 0.4 | 0.3 |

Note: see the note to Table 2. The observed variables are: the growth rates of output (y), consumption (c), investment, and wages (w), the inflation rates for consumption (π) and investment (π^i), hours worked (h) and the nominal interest rate (r). Due to rounding in some cases the band-specific contributions do not add up to the total values.

3.2.2 Information contributions by variables

Table 5 shows the conditional information gains for each observed variable in the full spectrum and the individual frequency bands. The three largest contributions, each exceeding 90%, are from the growth rate of the relative price of investment (π^i) with respect to the IST shock (v), from the real wage growth (w) with respect to the wage mark-up shock (λ_w), and from the nominal interest rate (R) with respect to the monetary policy shock (ε_{mp}). As JPT show, the price of investment in terms of consumption goods coincides with the inverse of the IST process. Therefore, the IST growth rate process is fully recovered by observing π^i alone. A conditional information gain of 97.2% in the full spectrum implies that, in absence of π^i , information from the remaining variables reduces uncertainty about v by 2.8%. In addition to IST shock, π^i also contributes information with respect to the MEI shock, although much less compared to other variables, and in particular the investment growth rate, which is the most informative variable for that

shock. Other large conditional contributions are from the growth rates of output and consumption with respect to the government spending shock, and from hours worked with respect to the neutral technology shock. Consumption growth is also the most informative variable with respect to the intertemporal preference shock, while inflation is the most informative observable with respect to the price mark-up.

Table 6: Unconditional contribution of information

| shocks | | total | | | | | | | | low | | | | | | | |
|--------------------|--------------------|----------|----------|----------|----------|----------|-------|----------|---------|----------|----------|----------|----------|----------|-------|----------|---------|
| | | <i>x</i> | <i>c</i> | <i>i</i> | <i>h</i> | <i>w</i> | π | <i>R</i> | π^i | <i>x</i> | <i>c</i> | <i>i</i> | <i>h</i> | <i>w</i> | π | <i>R</i> | π^i |
| <i>z</i> | neutral technology | 17.4 | 20.1 | 7.3 | 24.3 | 28.7 | 24.4 | 9.6 | 0.0 | 5.4 | 6.1 | 2.7 | 0.6 | 7.5 | 2.0 | 0.6 | 0.0 |
| <i>g</i> | government | 0.4 | 3.1 | 0.1 | 4.4 | 0.0 | 1.7 | 4.8 | 0.0 | 0.1 | 3.1 | 0.1 | 4.2 | 0.0 | 1.7 | 4.7 | 0.0 |
| <i>v</i> | IST | 1.3 | 0.1 | 2.0 | 0.6 | 0.1 | 0.1 | 0.2 | 100.0 | 0.0 | 0.1 | 0.1 | 0.0 | 0.0 | 0.1 | 0.1 | 8.4 |
| λ_p | price mark-up | 1.8 | 0.3 | 4.4 | 4.8 | 18.1 | 39.3 | 3.6 | 0.0 | 1.3 | 0.3 | 3.8 | 4.3 | 8.3 | 6.2 | 1.5 | 0.0 |
| λ_w | wage mark-up | 0.4 | 0.4 | 0.6 | 1.0 | 59.7 | 3.4 | 0.6 | 0.0 | 0.3 | 0.2 | 0.2 | 0.8 | 0.2 | 2.1 | 0.5 | 0.0 |
| <i>b</i> | preference | 7.4 | 61.9 | 0.9 | 6.9 | 0.0 | 1.7 | 9.7 | 0.0 | 0.3 | 3.6 | 0.2 | 0.6 | 0.0 | 0.6 | 1.7 | 0.0 |
| ε_{mp} | monetary policy | 3.5 | 1.2 | 2.8 | 3.1 | 0.0 | 1.7 | 57.6 | 0.0 | 0.2 | 0.1 | 0.1 | 0.3 | 0.0 | 0.4 | 0.2 | 0.0 |
| μ | MEI | 47.2 | 10.3 | 73.1 | 56.3 | 3.4 | 9.7 | 51.6 | 0.0 | 16.5 | 8.4 | 28.2 | 24.9 | 2.8 | 4.9 | 30.4 | 0.0 |

| shocks | | BC | | | | | | | | high | | | | | | | |
|--------------------|--------------------|----------|----------|----------|----------|----------|-------|----------|---------|----------|----------|----------|----------|----------|-------|----------|---------|
| | | <i>x</i> | <i>c</i> | <i>i</i> | <i>h</i> | <i>w</i> | π | <i>R</i> | π^i | <i>x</i> | <i>c</i> | <i>i</i> | <i>h</i> | <i>w</i> | π | <i>R</i> | π^i |
| <i>z</i> | neutral technology | 7.5 | 8.4 | 2.7 | 7.7 | 16.8 | 14.3 | 6.0 | 0.0 | 4.5 | 5.6 | 1.9 | 15.9 | 4.4 | 8.1 | 3.0 | 0.0 |
| <i>g</i> | government | 0.1 | 0.1 | 0.0 | 0.1 | 0.0 | 0.0 | 0.0 | 0.0 | 0.1 | 0.0 | 0.0 | 0.1 | 0.0 | 0.0 | 0.0 | 0.0 |
| <i>v</i> | IST | 0.2 | 0.0 | 0.3 | 0.1 | 0.0 | 0.0 | 0.1 | 33.6 | 1.0 | 0.0 | 1.7 | 0.5 | 0.0 | 0.0 | 0.0 | 58.0 |
| λ_p | price mark-up | 0.3 | 0.0 | 0.3 | 0.3 | 3.6 | 7.2 | 0.6 | 0.0 | 0.3 | 0.0 | 0.3 | 0.2 | 6.2 | 25.8 | 1.5 | 0.0 |
| λ_w | wage mark-up | 0.1 | 0.1 | 0.2 | 0.1 | 10.1 | 1.2 | 0.2 | 0.0 | 0.0 | 0.1 | 0.2 | 0.0 | 49.4 | 0.1 | 0.0 | 0.0 |
| <i>b</i> | preference | 4.0 | 34.9 | 0.5 | 3.9 | 0.0 | 0.9 | 5.6 | 0.0 | 3.2 | 23.5 | 0.2 | 2.3 | 0.0 | 0.2 | 2.4 | 0.0 |
| ε_{mp} | monetary policy | 1.1 | 0.4 | 0.8 | 1.2 | 0.0 | 0.7 | 8.1 | 0.0 | 2.2 | 0.7 | 1.9 | 1.6 | 0.0 | 0.6 | 49.3 | 0.0 |
| μ | MEI | 24.6 | 1.9 | 34.6 | 26.6 | 0.6 | 4.4 | 20.2 | 0.0 | 6.1 | 0.1 | 10.4 | 4.7 | 0.0 | 0.4 | 1.0 | 0.0 |

Note: see the note to Table 3.

With a few exceptions, variables that contribute the most information overall are also the most informative ones within each frequency band. One of the exceptions is the contribution of wage growth with respect to the price markup shock, which is significantly larger than the contribution of inflation in the low frequency band, but much smaller in the BC and high frequencies, and thus overall. Another notable exception is the intertemporal preference shock where consumption growth is by far the most informative variable overall, even though in the low frequency band the conditional contributions of both hours worked and investment growth are much larger.

Table 6 shows results for the unconditional information gains. As discussed earlier, for a given variable and a shock, the difference between conditional and unconditional

information gains indicates the existence of information complementarities with respect to that shock between the variable and other observables. The complementary may be positive or negative depending on whether the conditional gains are larger or smaller than the unconditional ones.

The most extreme case of positive complementarity is observed with respect to the government spending shock, where the largest unconditional gain – from R , is less than 5%, whereas there are two variables – c and x , with conditional gains exceeding 45%, and a third one – i , with conditional gain exceeding 18%. The obvious explanation for this result is the existing tight relationship among x , c , i , and g implied by the resource constraint of the economy. Since g is latent, joint information from pairs of the observed resource constraint variables is larger than the information contained in each of them individually. This intuition can be confirmed by applying the measure of information complementarity introduced earlier (see equation (3.5)). The top panel of Figure 6 shows the largest, in absolute value, unconditional and conditional information complementarities with respect to the government spending shock. In both cases, the largest positive complementarities are between pairs of resource constraint variables. For instance, the value of 3.2 in the case of x and c implies that, conditional on observing the remaining six variables, observing x and c together provides 2.2 times as much information about g as adding up the information from each of them individually.

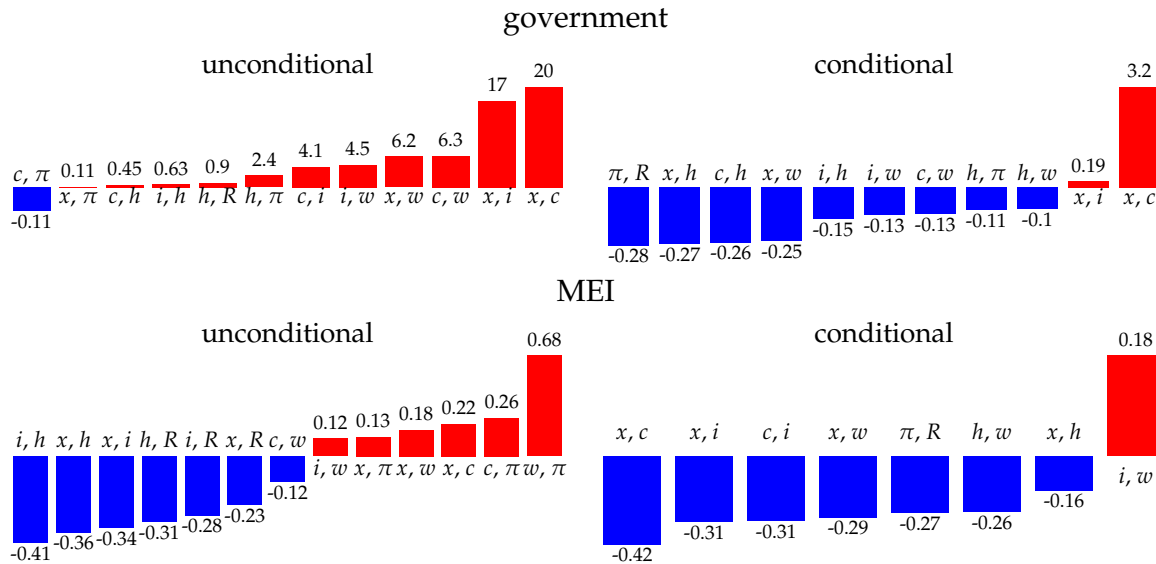


Figure 6: Largest pairwise information complementarities with respect to the government spending and MEI shocks, full spectrum.

The bottom panel of Figure 6 shows the most significant complementarities with respect to the MEI shock. As can be confirmed by comparing the values reported in Table 5 and Table 6, the MEI shock presents the most prominent case of negative information complementarities. In particular, the gain from i , which is the most informative variable for that shock, both conditionally and unconditionally, drops from more than 70% unconditionally, to less than 10%, conditionally. Similarly, the information gains from R , h , and x all drop from around 50% to about 5% or less. This implies that, to a large extent, information from these variables regarding the MEI shock is not unique to them but is also contained in other observed variables. As can be seen in Figure 6, different combinations of i , R , h , and x are among the variable pairs with the strongest negative information complementarity. In the case of i and x , the explanation again can be traced to their strong mutual dependence, together with c , implied by the resource constraint of the economy. In fact, we should expect to find negative complementarity with respect to all shocks, other than g , between any two of the resource constraint variables when the third is among the conditional variables.⁹ This is indeed the case as shown in more details in the Appendix. Also there, I report the pairwise complementarity coefficients for each frequency band. Examining those results can help explain, for example, the finding that the overall information complementarity between c and i with respect to the g shock is zero, which may be puzzling given the preceding argument for why all pairs of resource constraint variables should exhibit significant mutual complementarity. Indeed, the complementarities between c and i , c and x , and x and i are very strong in the BC and high frequencies. However, as seen earlier, almost all information about g is in the low frequencies, where the complementarity between c and i is zero.

The main conclusion of JPT is that the MEI shock is the key source of business cycle fluctuations whereas the IST shock plays no role. In particular, they show that the MEI shock is responsible for large fractions of the variances of GDP, investment, and hours *at business cycle frequencies*, and the contributions of the IST shock are nil. In Figure 7, I report variance decompositions for the individual frequency bands as well as the overall contribution of each shock to the variances of the observed variables.¹⁰

⁹At the risk of belaboring the obvious, consider the case where g is also observed, or, alternatively, where the g shock has zero variance. Then, because of the exact collinearity among them, the information in any variable entering the resource constraint with respect to all shocks is completely redundant given the information in the remaining resource constraint variables.

¹⁰The JPT results are displayed in Table 3 of their article. There are several differences between their presentation and the one in Figure 7. One is that JPT show the variance contributions as fractions of the total variance in the business cycle frequencies. In my plot, the fractions are relative to the

The results show that the MEI shock explains the bulk of the variances of x , i , and h in the full spectrum, not just the BC frequencies. In addition, the same shock contributes most of the volatility of R . This helps understand the earlier observation that R is the second most informative variable about μ (see Table 5). Note that the MEI shock contributes most of the variance in the low and BC frequencies of R – 65% and 55% of the total variance in those frequency bands, respectively, and those are the parts of the spectrum where most of the uncertainty about μ resides. Furthermore, unlike the resource constraint variables, a relatively smaller fraction of the information in R is redundant. The spectral decomposition of the contribution of the preference shock to the variance of c shows a relatively small impact in the low frequency band, which helps explain another observation made earlier – that in spite of being the most informative variable for that shock overall, c is not as informative about it in the low frequencies. Similarly, the fact that the contribution of the wage mark-up shock to the volatility of w is predominantly in the higher end of the spectrum is consistent with the dominant role of w as a source of information for that shock. In contrast, the low frequencies are a major source of uncertainty about the price mark-up shock, contributing more than half of its total variance. Given that most of that shock’s contribution to the variance of h is also in the low frequency band, this helps understand why the importance of h as a source of information about λ_p is comparable to that of w and π , in spite of the much smaller fraction of total variance of h due to that shock.

As seen earlier, the importance of h is even greater in the case of the neutral technology shock, for which it is the observable with the largest conditional contribution of information. This might be hard to anticipate on the basis of the variance decomposition results, which show that more than half of the contribution of z is to the low frequency component of the variance, and only 0.3% of the total variance of h stems from the high frequency contribution of that shock. At the same time, the BC and high frequencies account for almost 90% of the total information about z , and the bulk of the information contributed by h is within the high frequency band. As shown in more details in the Appendix, this result is due to, on the one hand, the strong positive complementarity between h and x , and, on the other hand, the also strong negative complementarities

total variances in the full spectrum. To obtain comparable contributions, the values in the plot have to be multiplied by the fraction of the total variance of each shock due to the BC frequencies. Another difference is that for the trending variables (x , c , i , and w), JPT show decompositions for the levels, whereas I present results for the growth rates. Finally, the point estimates in JPT are the median values of the posterior distributions of the contributions. I present decompositions at the posterior median of the estimated parameters of the model.

among x , c and i , as well as between π and w , and π and R . In other words, there is a substantial redundancy in the information about z in variables for which this shock is an important source of volatility. Furthermore, note that only 1% of the total variance of h originates in the high frequency band. Therefore, z is responsible for 30% of it, making it the second most important shock, after μ , for h in the high frequencies.

The last observation supports a point made earlier, with respect to the Uribe (2021) model, that the size of the variance contribution is not necessarily a good indicator of the variables' importance as sources of information about the shocks. As also pointed out earlier, it is possible that shocks are recoverable even if they play only a modest role as sources of volatility. In the JPT model, the monetary policy shock is responsible for at most 9.5% of the volatility of any observable, and the government spending shock contributes at most 7.3%. Yet both shocks are fully recoverable.

| | x | c | i | h | w | π | R | π^i | |
|--------------------|------|------|------|------|------|-------|------|---------|-------|
| all shocks | 22.4 | 30.4 | 19.6 | 79.8 | 21.5 | 46.7 | 63.4 | 8.4 | low |
| | 52.3 | 45.9 | 58.9 | 19.2 | 33.7 | 35.4 | 33.5 | 33.6 | BC |
| | 25.3 | 23.7 | 21.5 | 1.0 | 44.7 | 17.9 | 3.1 | 58.0 | high |
| z | 23.2 | 30.4 | 8.4 | 6.6 | 33.1 | 22.2 | 8.1 | 0.0 | total |
| | 9.9 | 17.0 | 3.0 | 3.8 | 14.6 | 6.3 | 3.1 | 0.0 | low |
| | 11.0 | 10.7 | 4.5 | 2.5 | 14.9 | 12.9 | 4.7 | 0.0 | BC |
| | 2.4 | 2.7 | 0.9 | 0.3 | 3.6 | 3.0 | 0.2 | 0.0 | high |
| δ | 7.3 | 2.2 | 0.1 | 2.1 | 0.0 | 0.5 | 1.1 | 0.0 | total |
| | 0.1 | 1.2 | 0.0 | 1.5 | 0.0 | 0.4 | 0.9 | 0.0 | low |
| | 1.8 | 0.8 | 0.0 | 0.4 | 0.0 | 0.1 | 0.3 | 0.0 | BC |
| v | 5.4 | 0.2 | 0.0 | 0.1 | 0.0 | 0.0 | 0.0 | 0.0 | high |
| | 0.7 | 0.3 | 1.1 | 0.4 | 0.1 | 0.4 | 0.9 | 100.0 | total |
| | 0.1 | 0.2 | 0.1 | 0.3 | 0.1 | 0.3 | 0.8 | 8.4 | low |
| | 0.2 | 0.1 | 0.4 | 0.0 | 0.0 | 0.0 | 0.1 | 33.6 | BC |
| λ_p | 0.4 | 0.0 | 0.6 | 0.0 | 0.0 | 0.0 | 0.0 | 58.0 | high |
| | 2.0 | 0.2 | 2.2 | 6.8 | 21.1 | 34.6 | 2.1 | 0.0 | total |
| | 0.8 | 0.1 | 0.9 | 6.2 | 4.9 | 6.4 | 0.9 | 0.0 | low |
| | 1.0 | 0.0 | 1.1 | 0.6 | 7.5 | 14.2 | 1.0 | 0.0 | BC |
| λ_w | 0.2 | 0.0 | 0.2 | 0.0 | 8.7 | 14.1 | 0.1 | 0.0 | high |
| | 1.5 | 1.8 | 1.0 | 26.2 | 44.0 | 28.2 | 10.4 | 0.0 | total |
| | 1.2 | 1.6 | 0.5 | 26.0 | 0.6 | 25.6 | 9.9 | 0.0 | low |
| | 0.3 | 0.2 | 0.4 | 0.2 | 11.0 | 2.6 | 0.5 | 0.0 | BC |
| b | 0.0 | 0.0 | 0.1 | 0.0 | 32.4 | 0.0 | 0.0 | 0.0 | high |
| | 7.3 | 56.5 | 1.0 | 3.3 | 0.0 | 2.0 | 8.4 | 0.0 | total |
| | 0.3 | 4.4 | 0.2 | 2.0 | 0.0 | 1.3 | 4.8 | 0.0 | low |
| | 4.0 | 31.8 | 0.6 | 1.3 | 0.0 | 0.7 | 3.3 | 0.0 | BC |
| ε_{mp} | 2.9 | 20.3 | 0.2 | 0.1 | 0.0 | 0.1 | 0.3 | 0.0 | high |
| | 3.7 | 1.3 | 2.8 | 4.9 | 0.0 | 3.9 | 9.5 | 0.0 | total |
| | 0.7 | 0.4 | 0.5 | 3.9 | 0.0 | 2.5 | 2.1 | 0.0 | low |
| | 2.2 | 0.6 | 1.7 | 1.0 | 0.0 | 1.2 | 5.3 | 0.0 | BC |
| μ | 0.8 | 0.3 | 0.6 | 0.0 | 0.0 | 0.2 | 2.1 | 0.0 | high |
| | 54.1 | 7.4 | 83.4 | 49.6 | 1.7 | 8.2 | 59.7 | 0.0 | total |
| | 9.3 | 5.6 | 14.3 | 36.2 | 1.3 | 3.9 | 40.9 | 0.0 | low |
| | 31.7 | 1.6 | 50.1 | 13.0 | 0.3 | 3.7 | 18.4 | 0.0 | BC |
| | 13.1 | 0.1 | 19.0 | 0.4 | 0.0 | 0.5 | 0.3 | 0.0 | high |
| | x | c | i | h | w | π | R | π^i | |

Figure 7: Total and individual contributions of shocks as a percent of the variances of the observables in the full spectrum and the low, business cycle, and high frequency bands.

3.3 Angeletos, Collard, and Dellas (2018)

The main contribution of Angeletos et al. (2018) (henceforth ACD) is showing how to introduce higher-order belief dynamics into a large class of macroeconomic models in a tractable way. This is illustrated with several examples including a medium-scale New Keynesian model, which I consider in this section. Except for higher-order beliefs, it shares many of the features of the JPT model in the previous section and other medium-scale DSGE models in the literature: habit persistence in consumption, costs of adjusting investment, variable capital utilization, price stickiness under Calvo pricing, monetary policy following a Taylor rule, permanent and transitory TFP shocks, permanent and transitory investment-specific shocks, a discount-rate shock, a news shock regarding future productivity, a shock to government-spending, and a monetary policy shock.

From a modelling perspective, what sets ACD apart from most of the literature is the departure from the assumptions of rational expectations and common information about the state of the economy. In particular, in their model, agents' beliefs regarding the expectations of other agents, i.e. higher-order beliefs, are subject to autonomous variation, called "confidence shock", which causes divergence between the two forms of beliefs. This leads to exogenous movements in agents' expectations of the economic outcomes in the short-run, without altering their medium or long-run expectations of those outcomes, or the expectations of the exogenous fundamentals at any horizon. ACD argue that embedding their mechanism in what are otherwise standard DSGE models can help better match observed patterns in macroeconomic data.

In the remainder of this section I evaluate the sources and distribution of information about the shocks in the ACD model. There are nine shocks, most of which, with the exception of the confidence shock, are defined similarly as in the JPT model. The main difference here is that the level of TFP is assumed to have both permanent and transitory components, whereas in the JPT model there is only a permanent component. In addition, ACD assume that the permanent TFP term contains a one-quarter ahead news component modeled as an exogenous stationary AR(1) process. Apart from that, the model equations are similar to those given in Section 3.2 and therefore I leave a more detailed presentation for the Appendix. Also, even though the introduction of higher order beliefs changes the way the model is solved, compared to standard rational expectations models, the only effect this has on the model representation is the presence of two types of expectations in otherwise standard equilibrium conditions. In particular, some decisions are made conditional on agents' beliefs that the expectations of other

agents are biased, while other decisions are made after the true state of nature and the realized value of economic activity is publicly revealed. The perceived bias is given by the confidence shock, which is commonly observed. For more details, see Section C in the Appendix and the original article Angeletos et al. (2018).

ACD estimate the model using US data on six quarterly series: GDP, consumption, investment, hours worked, the inflation rate, and the federal fund rate. The estimation is carried out in the frequency domain using only business cycle frequencies. This is another departure from the common practice in the literature where time domain methods are predominant. The reason the authors give for taking this approach is that their model is intended to describe business cycle phenomena only, and thus lacks the features and mechanisms required for it to account for the lower and higher frequencies of the data. Next, I will discuss the implications and possible interpretations of the information decomposition in that setting.

3.3.1 Information decomposition across frequency bands and observables

Table 7: Information decomposition across frequency bands

| shock | total | low | BC | high |
|--------------------------------------|-------|---------------------------|---------------------------|---------------------------|
| a^p permanent TFP | 99.7 | $98.6 = 99.9 \times 0.99$ | $1.0 = 90.2 \times 0.01$ | $0.1 = 62.5 \times 0.00$ |
| a^n news | 43.2 | $5.8 = 59.1 \times 0.10$ | $18.3 = 49.3 \times 0.37$ | $19.1 = 36.0 \times 0.53$ |
| a^r transitory TFP | 30.2 | $0.7 = 4.6 \times 0.15$ | $10.7 = 23.5 \times 0.45$ | $18.8 = 47.3 \times 0.40$ |
| ζ^{IP} permanent investment | 91.5 | $91.5 = 92.7 \times 0.99$ | $0.1 = 6.3 \times 0.01$ | $0.0 = 6.6 \times 0.00$ |
| ζ_t^{IT} transitory investment | 92.5 | $10.3 = 76.2 \times 0.14$ | $42.1 = 96.0 \times 0.44$ | $40.1 = 94.0 \times 0.43$ |
| ζ^c discount factor | 98.6 | $64.9 = 99.4 \times 0.65$ | $27.5 = 98.0 \times 0.28$ | $6.2 = 93.9 \times 0.07$ |
| ζ^g fiscal | 93.2 | $39.7 = 90.6 \times 0.44$ | $41.3 = 95.8 \times 0.43$ | $12.3 = 93.5 \times 0.13$ |
| ζ^m monetary policy | 97.4 | $25.5 = 92.9 \times 0.27$ | $49.2 = 98.7 \times 0.50$ | $22.7 = 99.7 \times 0.23$ |
| ξ confidence | 96.7 | $51.8 = 98.8 \times 0.52$ | $36.1 = 96.2 \times 0.38$ | $8.8 = 87.4 \times 0.10$ |

Note: see the note to Table 1.

Consider Table 7 which shows the results from applying the information decomposition to the shocks in the ACD model. The results can be interpreted in several different ways, depending on one's beliefs about the extent to which the model is suitable to represent the empirical data. First, under the assumption that the model is correctly specified for all frequencies of the observed time series, the table shows how information about the shocks is distributed across the low, business cycle, and high frequency bands, and how much information is obtained about each shock in total. This is analogous to the earlier interpretation of the results in Tables 1 and 4 for the Uribe and JPT models, and can

be used to draw conclusions about the sources of identification of the shocks in the ACD model, as was done for those models.

Second, one may adopt the view, as ACD do, that the model is misspecified outside the BC frequencies and focus on those frequencies only, ignoring the rest of the spectrum. In that case, the relevant question to ask is how much of the shocks' uncertainty, that originates in the BC frequencies, is resolved with information provided by the observed variables in that frequency band. The answer is given by the information gains in the BC frequencies, shown in the middle of the fourth column of the table. Although the prior uncertainty is not fully resolved for any of the shocks, for six of them the information gains exceed 90%, reaching 98% in the case of the discount factor and monetary policy shocks. The least amount of information, only 6%, is obtained with respect to the permanent investment-specific technology shock, followed by the transitory TFP and the news shocks, with information gains of about 24% and 49%, respectively. Note that, for the later two shocks, the information gains in the BC frequencies are similar in size to the information gains in the full spectrum, 32% and 43% respectively. On the other hand, in the case of the permanent investment-specific technology shock the total information gain is much larger, 92%. This reflects the fact that 99% of the uncertainty of that shock originates in the low frequency band and the information gain there is close to 93%.

Table 8: Conditional information gains

| shock | all | | | | | | BC | | | | | |
|--------------------------------------|----------|----------|----------|----------|-------|----------|----------|----------|----------|----------|-------|----------|
| | <i>Y</i> | <i>C</i> | <i>I</i> | <i>N</i> | π | <i>R</i> | <i>Y</i> | <i>C</i> | <i>I</i> | <i>N</i> | π | <i>R</i> |
| a^P permanent TFP | 0.9 | 0.5 | 1.1 | 1.2 | 0.0 | 0.1 | 26.6 | 0.1 | 1.0 | 37.3 | 0.0 | 0.3 |
| a^n news | 10.7 | 2.2 | 0.7 | 11.4 | 0.1 | 2.7 | 15.0 | 0.5 | 0.9 | 19.4 | 0.0 | 1.4 |
| a^τ transitory TFP | 2.0 | 0.3 | 2.4 | 11.5 | 0.3 | 1.7 | 1.0 | 0.4 | 2.6 | 8.6 | 0.2 | 2.3 |
| ζ^{IP} permanent investment | 2.1 | 10.1 | 15.3 | 0.7 | 0.3 | 3.1 | 0.8 | 1.0 | 2.2 | 0.8 | 0.1 | 1.0 |
| ζ_t^{IT} transitory investment | 0.2 | 10.1 | 32.7 | 1.2 | 0.8 | 1.2 | 0.3 | 12.4 | 24.8 | 0.1 | 0.5 | 1.6 |
| ζ^c discount factor | 1.6 | 2.1 | 0.2 | 11.3 | 13.0 | 8.7 | 2.3 | 2.9 | 0.2 | 20.0 | 23.0 | 17.8 |
| ζ^g fiscal | 65.5 | 53.9 | 70.2 | 10.1 | 0.1 | 0.7 | 62.0 | 55.9 | 71.9 | 9.8 | 0.0 | 0.1 |
| ζ^m monetary policy | 3.2 | 3.4 | 0.4 | 19.1 | 71.6 | 67.4 | 1.6 | 1.8 | 0.1 | 13.0 | 70.1 | 66.7 |
| ξ confidence | 5.4 | 19.1 | 0.1 | 10.4 | 3.2 | 3.4 | 4.7 | 23.2 | 0.1 | 9.2 | 3.3 | 4.6 |

Note: The conditional information gain measures the marginal reduction of uncertainty about a shock due to observing a variable relative to observing the other five variables, as a percent of the unconditional uncertainty of the shock. The observed variables are: output (*Y*), consumption (*C*), investment (*I*), hours worked (*N*), inflation (π) and the nominal interest rate (*R*).

The relative importance of the observed variables for each shock, in the BC frequencies or the full spectrum, can be determined by examining the respective information gains.

The results are presented in Table 8, which shows the conditional information gains for each variable and shock. Note that, unlike Tables 2 and 5, the numbers in the BC panel of the table show how much of the uncertainty originating in the BC frequencies is resolved with the information contained in the observed variables, rather than the contributions from the BC frequencies to the total amount of information. This is the appropriate measure to apply when the model is considered suitable of explaining only the BC frequencies of the data. Comparing the two sets of information gains, we see that the rankings of variables with largest contribution for each shock are nearly identical in the full spectrum and the BC frequencies. One notable exception is output (Y) being significantly more informative than investment (I) about the permanent TFP shock (a^p) in the BC frequencies, while in the full spectrum I is the more informative variable of the two. The largest contribution for that shock, as well as the transitory TFP (a^t) and news (a^n) shocks, is from hours worked (N). N also contributes significant amount of information with respect to the discount factor (ζ^c) and confidence (ξ) shocks, having the second largest contributions of information for those shocks after inflation (π) and consumption (C), respectively. Y , C , and I are the three most significant sources of information for the government spending shock (ζ^g), which, as already discussed in the context of the JPT model, is a consequence of the resource constraint equation linking these variables and ζ^g . π , followed by the nominal interest rate (R), are by far the most informative variables for the monetary policy shock (ζ^m), which was also the case in the JPT model (see Table 5), although there the contribution from the nominal rate is the largest one by a sizeable margin. Lastly, as might be expected, I is the most informative variable for the two investment-specific shocks (ζ^{IP} and ζ_t^{IT})

As explained earlier, the rationale for excluding frequencies outside the BC range is to avoid biased parameter estimates due to contaminated information from parts of the spectrum where the model is misspecified.¹¹ The same argument applies to the estimation of shocks and other latent variables. Using contaminated information leads to distorted estimates of those variables, even when the true values of the model parameters are known. The form of the distortions depends on the nature of misspecification, that is, what mechanisms operate at the low and high frequencies of the data and are absent in the theoretical model. This is model- and data-specific. At the same time, it is clear that the more the estimation relies on contaminated information, the more severe

¹¹Hansen and Sargent (1993) and Diebold et al. (1998) make similar arguments and develop band-spectral estimation methods. See also Christiano and Vigfusson (2003), Sala (2015), and Qu and Tkachenko (2012).

is the impact of misspecification on the results. For instance, from Table 7 one sees that model misspecification at the lower end of the spectrum would have a greater impact on the estimates of shocks with larger information contributions from the lower frequencies, like the permanent investment-specific shock, compared to shocks for which those contributions are relatively small, such as the transitory investment-specific shock. Reversely, the effect of misspecification at the high frequencies would be greater for the transitory than the permanent investment-specific shock.

4 Concluding Comments

I have shown how to decompose the frequency domain information observables provide with respect to latent variables in dynamic macroeconomic models. Through this analysis, researchers can determine where in the spectrum information about latent variables predominantly comes from, and evaluate the relative contributions of individual observed variables. The examples I have presented illustrate how reporting the results from such analysis can make the estimation of shocks and other latent variables more transparent. Researchers often disagree on the specific model features needed to adequately represent the data. In particular, there is no consensus on which data frequencies macroeconomic models should aim to, or are capable of, representing. Whilst much of the empirical literature is focused on explaining business cycle phenomena, models are usually estimated in the time domain, which is tantamount to using the full spectrum. The Uribe and JPT models I have considered are only two cases in point. Even if not explicitly stated, the time domain approach implicitly assumes that models are capable of representing all frequencies in the data. Presenting readers, who may have divergent views on the model adequacy, with information on the relative importance of different frequencies will allow them to assess the potential consequences of using contaminated information due to model misspecification.

Another issue over which researchers may disagree concerns the extent to which observed time series adequately represent theoretical variables in macroeconomic models. The existence of multiple empirical counterparts to variables like output, inflation, wages, hours worked, etc. suggests that all of them should be treated as noisy indicators of the underlying theoretical concepts. This supports the argument in favor of treating these variables as measured with errors.¹² Yet, measurement errors are not universally present

¹²An earlier statement of this argument was made by Boivin and Giannoni (2006) who proposed

in estimated models, as demonstrated by the JPT and ACD examples. Such an omission could be interpreted by some readers as a reason to suspect that models are misspecified with respect to particular observed variables. Based on the perceived nature of the errors, one can draw a conclusion about which frequencies are most affected. Knowing how important those frequencies are as a source of information can help readers better understand the consequences of the failure to account for the imperfect match between theoretical concepts and empirical time series. For instance, pure measurement errors are often modeled as white noise processes, and therefore the contamination is concentrated in the higher end of the spectrum. As a result, estimates that rely more heavily on information from the high frequencies there will be compromised more severely. Similarly, perceived failure to adequately account to low frequency variations in some series would cause some readers to be sceptical of estimates which are more dependent on information from the lower of the spectrum.¹³

Lastly, the methodology described in this paper can help researchers who develop and estimate structural macroeconomic models by revealing, in cases of information deficiency, what type of information is needed to better recover unobserved variables of interest. Having well-identified structural shocks and unobserved endogenous variables, such as potential output or natural rate of interest, is a key requirement for macroeconomic models to meet to be useful as tools for policy analysis and to be credible as story-telling devices.

incorporating structural macroeconomic models into a dynamic factor framework where multiple imperfectly measured indicators correspond to each model concept.

¹³One commonly cited example of this is the series for aggregate hours worked, which contains significant low-frequency variations attributed to demographics and other structural developments in the labor market that are absent from most business cycle models. See the discussion of Figure 5 in Angeletos et al. (2018).

References

- ALESSI, L., M. BARIGOZZI, AND M. CAPASSO (2011): “Non-Fundamentalness in Structural Econometric Models: A Review,” *International Statistical Review*, 79, 16–47.
- ANDREWS, I., M. GENTZKOW, AND J. M. SHAPIRO (2020): “On the informativeness of descriptive statistics for structural estimates,” *Econometrica*, 88, 2231–2258.
- ANGELETOS, G.-M., F. COLLARD, AND H. DELLAS (2018): “Quantifying confidence,” *Econometrica*, 86, 1689–1726.
- BEAUDRY, P., P. FÈVE, A. GUAY, AND F. PORTIER (2015): “When is Nonfundamentalness in VARs a Real Problem? An Application to News Shocks,” Tech. rep., National Bureau of Economic Research.
- BLANCHARD, O. J., J.-P. L’HUILIER, AND G. LORENZONI (2013): “News, noise, and fluctuations: An empirical exploration,” *The American Economic Review*, 103, 3045–3070.
- BOIVIN, J. AND M. GIANNONI (2006): “DSGE Models in a Data-Rich Environment,” NBER Working Papers 12772, National Bureau of Economic Research, Inc, available at <http://ideas.repec.org/p/nbr/nberwo/12772.html>.
- BRILLINGER, D. R. (1981): *Time Series: Data Analysis and Theory*, vol. 36, SIAM.
- (2002): “Second-order moments and mutual information in the analysis of time series,” in *Recent Advances in Statistical Methods*, World Scientific, 64–76.
- BRILLINGER, D. R. AND A. GUHA (2007): “Mutual information in the frequency domain,” *Journal of statistical planning and inference*, 137, 1076–1084.
- CHAHROUR, R. AND K. JURADO (2022): “Recoverability and expectations-driven fluctuations,” *The Review of Economic Studies*, 89, 214–239.
- CHAMBERLAIN, G. (1990): “Distinguished fellow: Arthur S. Goldberger and latent variables in econometrics,” *Journal of Economic Perspectives*, 4, 125–152.
- CHRISTIANO, L. J. AND R. J. VIGFUSSON (2003): “Maximum likelihood in the frequency domain: the importance of time-to-plan,” *Journal of Monetary Economics*, 50, 789–815.
- DIEBOLD, F. X., L. E. OHANIAN, AND J. BERKOWITZ (1998): “Dynamic equilibrium economies: A framework for comparing models and data,” *The Review of Economic Studies*, 65, 433–451.

- FERNANDEZ-VILLAVERDE, J., J. RUBIO-RAMIREZ, T. J. SARGENT, AND M. W. WATSON (2007): “A, B, Cs (and Ds) of Understanding VARs,” *American Economic Review*, 97.
- FERNÁNDEZ-VILLAVERDE, J., J. F. RUBIO-RAMÍREZ, AND F. SCHORFHEIDE (2016): “Solution and estimation methods for DSGE models,” *Handbook of Macroeconomics*, 2, 527–724.
- FORNI, M. AND L. GAMBETTI (2014): “Sufficient information in structural VARs,” *Journal of Monetary Economics*, 66, 124–136.
- FRANCHI, M. AND P. PARUOLO (2015): “Minimality of state space solutions of DSGE models and existence conditions for their VAR representation,” *Computational Economics*, 46, 613–626.
- FRANCHI, M. AND A. VIDOTTO (2013): “A check for finite order VAR representations of DSGE models,” *Economics Letters*, 120, 100–103.
- GIACOMINI, R. (2013): “The relationship between DSGE and VAR models,” *Advances in Econometrics*, 31.
- GIANNONE, D., F. MONTI, AND L. REICHLIN (2016): “Exploiting the monthly data flow in structural forecasting,” *Journal of Monetary Economics*, 84, 201–215.
- GIANNONE, D. AND L. REICHLIN (2006): “Does information help recovering structural shocks from past observations?” *Journal of the European Economic Association*, 4, 455–465.
- GRANGER, C. AND J.-L. LIN (1994): “Using the mutual information coefficient to identify lags in nonlinear models,” *Journal of time series analysis*, 15, 371–384.
- HANSEN, H. AND T. SARGENT (1991): “Two Difficulties in Interpreting Vector Autoregressions,” in *Rational Expectations Econometrics*, ed. by H. Hansen and T. Sargent, Westview Press.
- HANSEN, L. P. AND T. J. SARGENT (1980): “Formulating and estimating dynamic linear rational expectations models,” *Journal of Economic Dynamics and control*, 2, 7–46.
- (1993): “Seasonality and approximation errors in rational expectations models,” *Journal of Econometrics*, 55, 21–55.
- ISKREV, N. (2019): “On the sources of information about latent variables in DSGE models,” *European Economic Review*, 119, 318–332.
- JOE, H. (1989): “Relative entropy measures of multivariate dependence,” *Journal of the American Statistical Association*, 84, 157–164.

- JUSTINIANO, A., G. E. PRIMICERI, AND A. TAMBALOTTI (2010): “Investment shocks and business cycles,” *Journal of Monetary Economics*, 57, 132–145.
- (2011): “Investment shocks and the relative price of investment,” *Review of Economic Dynamics*, 14, 102–121.
- LEEPER, E. M., T. B. WALKER, AND S.-C. S. YANG (2013): “Fiscal foresight and information flows,” *Econometrica*, 81, 1115–1145.
- LINDLEY, D. V. (1956): “On a measure of the information provided by an experiment,” *The Annals of Mathematical Statistics*, 27, 986–1005.
- LINFOOT, E. H. (1957): “An informational measure of correlation,” *Information and control*, 1, 85–89.
- LIPPI, M. AND L. REICHLIN (1993): “The dynamic effects of aggregate demand and supply disturbances: Comment,” *The American Economic Review*, 83, 644–652.
- (1994): “VAR analysis, nonfundamental representations, Blaschke matrices,” *Journal of Econometrics*, 63, 307–325.
- PRIESTLEY, M. B. (1981): *Spectral analysis and time series: probability and mathematical statistics*, 04; QA280, P7.
- QU, Z. AND D. TKACHENKO (2012): “Identification and frequency domain quasi-maximum likelihood estimation of linearized dynamic stochastic general equilibrium models,” *Quantitative Economics*, 3, 95–132.
- (2013): “Local and Global Parameter Identification in DSGE Models Allowing for Indeterminacy,” Tech. rep.
- RAVENNA, F. (2007): “Vector autoregressions and reduced form representations of DSGE models,” *Journal of monetary economics*, 54, 2048–2064.
- SALA, L. (2015): “DSGE models in the frequency domains,” *Journal of Applied Econometrics*, 30, 219–240.
- SALA, L., L. GAMBETTI, M. FORNI, ET AL. (2016): “VAR Information and the Empirical Validation of DSGE Models,” in *2016 Meeting Papers*, Society for Economic Dynamics, 260.
- SCHMITT-GROHÉ, S. AND M. URIBE (2012): “What’s News in Business Cycles,” *Econometrica*, 80, 2733–2764.
- SIMS, E. R. (2012): “News, non-invertibility, and structural VARs,” *Advances in Econometrics*, 28, 81.

- SOCCORSI, S. (2016): “Measuring nonfundamentalness for structural VARs,” *Journal of Economic Dynamics and Control*, 71, 86–101.
- STOCK, J. H. AND M. W. WATSON (1999): “Business cycle fluctuations in US macroeconomic time series,” *Handbook of macroeconomics*, 1, 3–64.
- UHLIG, H. (1999): “A toolkit for analyzing nonlinear dynamic stochastic models easily,” *Computational Methods for the Study of Dynamic Economies*, 30–61.
- URIBE, M. (2021): “The neo-Fisher effect: Econometric evidence from empirical and optimizing models,” *American Economic Journal: Macroeconomics*, forthcoming.
- WATSON, M. W. (1993): “Measures of Fit for Calibrated Models,” *The Journal of Political Economy*, 101, pp. 1011–1041.

Appendix

A Uribe (2021) model

Table A1: Parameter values, Uribe (2021) model

| | parameter | posterior mean |
|-----------------|--|----------------|
| ϕ | price stickiness | 146.000 |
| α_π | coeff inflation in monetary policy rule | 2.320 |
| α_y | coeff output in monetary policy rule | 0.188 |
| γ_m | backward-looking component in inflation | 0.606 |
| γ_I | coeff lagged interest rate in monetary policy rule | 0.242 |
| δ | habit formation | 0.258 |
| ρ_ξ | AR preference | 0.915 |
| ρ_θ | AR labor supply | 0.708 |
| ρ_z | AR transitory productivity | 0.700 |
| ρ_g | AR permanent productivity | 0.221 |
| ρ_{gm} | AR permanent trend inflation | 0.248 |
| ρ_{zm} | AR transitory interest rate | 0.306 |
| ρ_{zm2} | AR transitory trend inflation | 0.796 |
| σ_ξ | std. preference | 0.0287 |
| σ_θ | std. labor supply | 0.00164 |
| σ_z | std. transitory productivity | 0.00122 |
| σ_g | std. permanent productivity | 0.00758 |
| σ_{gm} | std. permanent trend inflation | 0.000848 |
| σ_{zm} | std. transitory interest rate | 0.000832 |
| σ_{zm2} | std. transitory trend inflation | 0.00131 |
| σ_1^{me} | std. measurement error Δy_t | 4.46e-06 |
| σ_2^{me} | std. measurement error r_t | 4.55e-06 |
| σ_3^{me} | std. measurement error Δi_t | 1.74e-07 |

B Justiniano, Primiceri, and Tambalotti (2011)

Table B1: Parameter values, JPT (2011) model

| | parameter | posterior median |
|---------------|-------------------------------------|------------------|
| α | capital share | 0.169 |
| ι_p | price indexation | 0.113 |
| ι_w | wage indexation | 0.102 |
| h | consumption habit | 0.864 |
| λ_p | SS mark-up goods prices | 0.177 |
| λ_w | SS mark-up wages | 0.166 |
| ν | inverse frisch elasticity | 5.162 |
| ξ_p | Calvo prices | 0.783 |
| ξ_w | Calvo wges | 0.773 |
| χ | Elasticity capital utilization cost | 5.491 |
| S' | Investment adjustment costs | 3.017 |
| ϕ_π | Taylor rule inflation | 1.735 |
| ϕ_Y | Taylor rule output | 0.059 |
| ρ_R | Taylor rule smoothing | 0.863 |
| ρ_z | AR neutral technology growth | 0.286 |
| ρ_g | AR government spending | 0.990 |
| ρ_ν | AR IST growth | 0.148 |
| ρ_p | AR price mark-up | 0.978 |
| ρ_w | AR wage mark-up | 0.968 |
| ρ_b | intertemporal preference | 0.583 |
| θ_p | MA price mark-up | 0.793 |
| θ_w | MA wage mark-up | 0.990 |
| ϕ_{dy} | Taylor rule output growth | 0.199 |
| ρ_μ | AR MEI | 0.807 |
| σ_{mp} | std. monetary policy | 0.216 |
| σ_z | std. neutral technology growth | 0.943 |
| σ_g | std. government spending | 0.362 |
| σ_ν | std. IST growth | 0.634 |
| σ_p | std. price mark-up | 0.222 |
| σ_w | std. wage mark-up | 0.310 |
| σ_b | std. intertemporal preference | 0.038 |
| σ_μ | std. MEI | 5.691 |




Figure B1: Largest unconditional pairwise information complementarities, all frequencies.

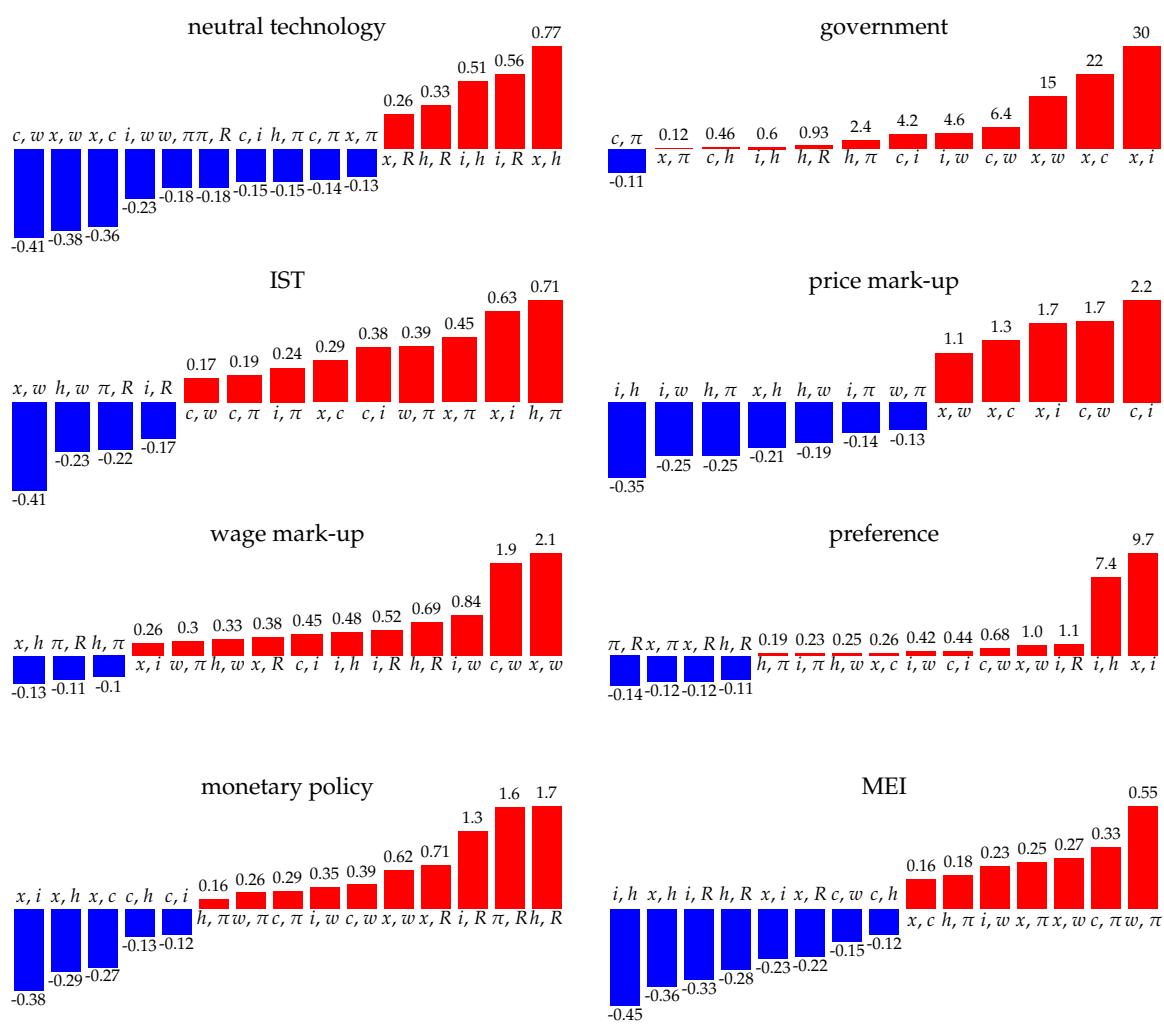


Figure B2: Largest unconditional pairwise information complementarities, low frequencies.

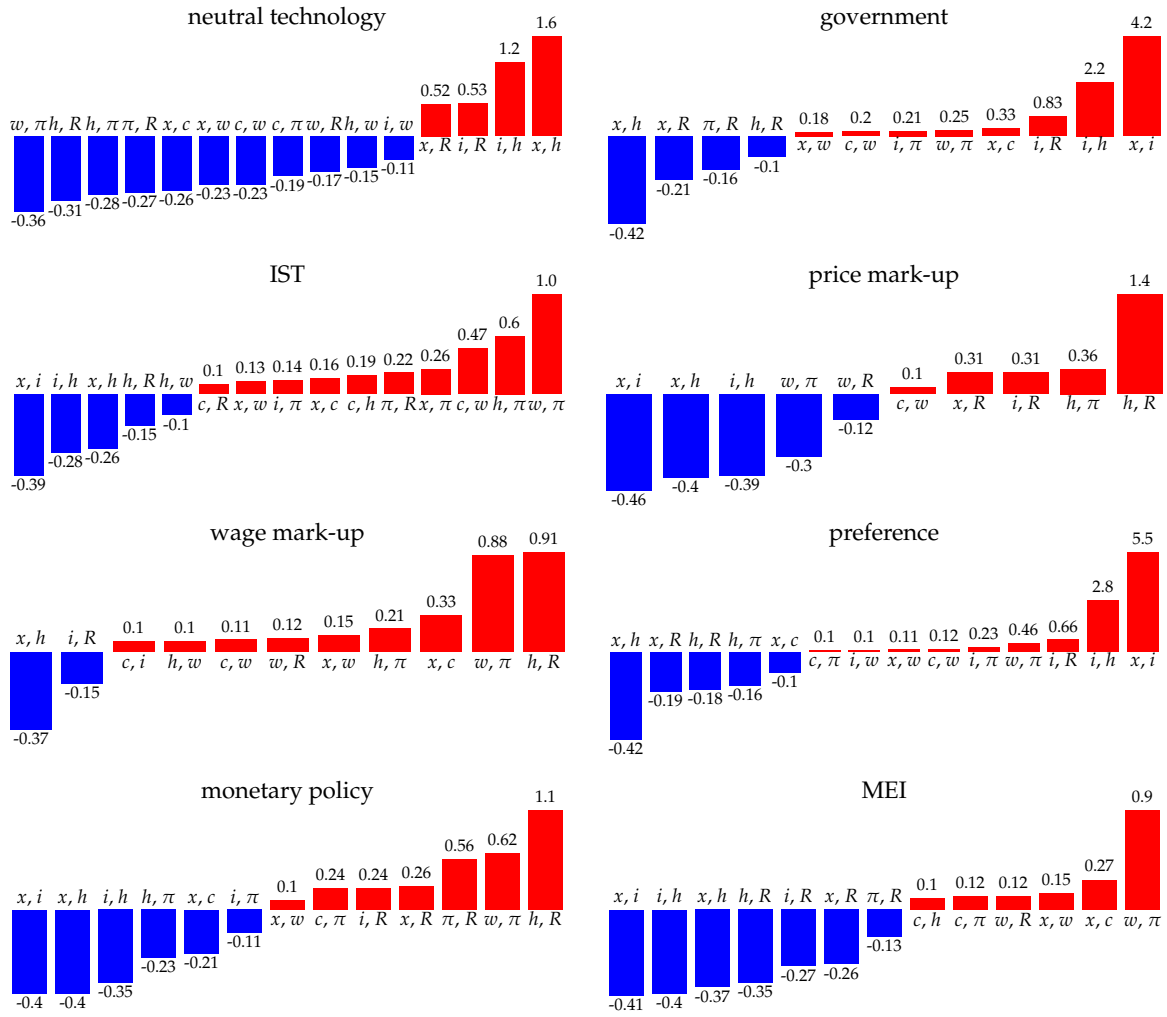


Figure B3: Largest unconditional pairwise information complementarities, BC frequencies.

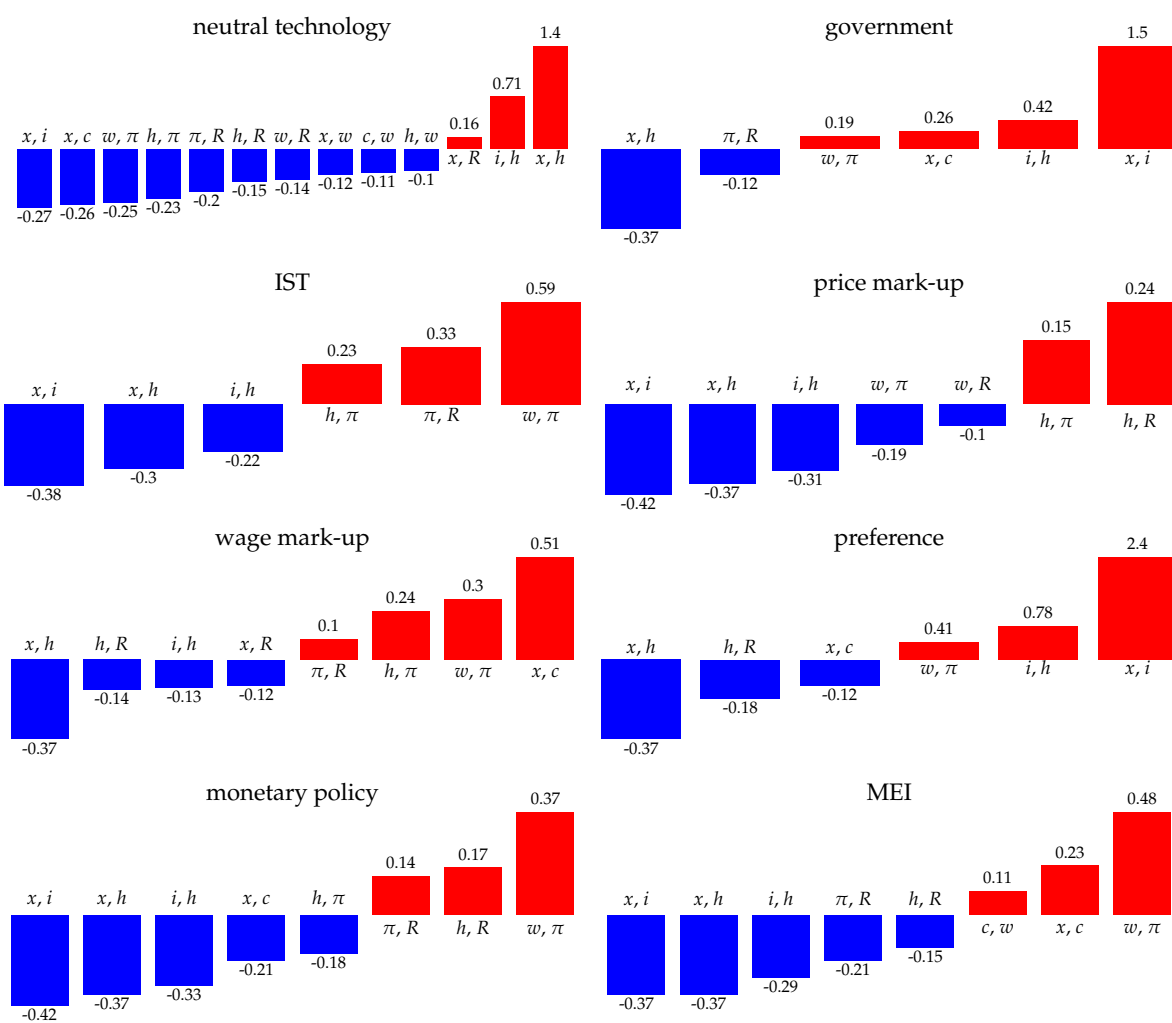


Figure B4: Largest unconditional pairwise information complementarities, high frequencies.

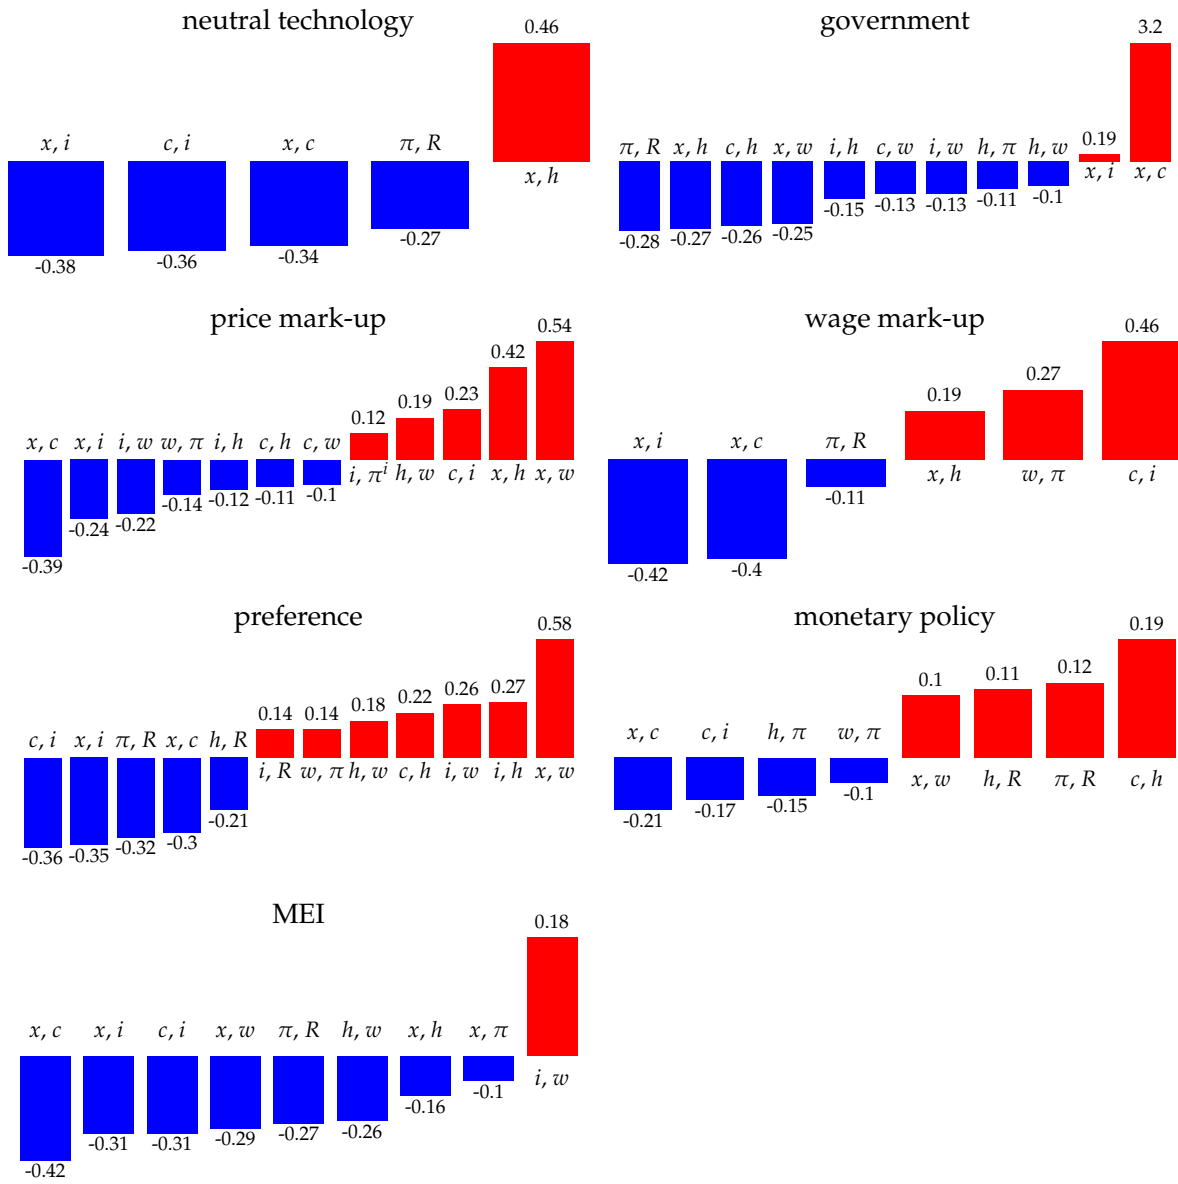


Figure B5: Largest conditional pairwise information complementarities, full spectrum.

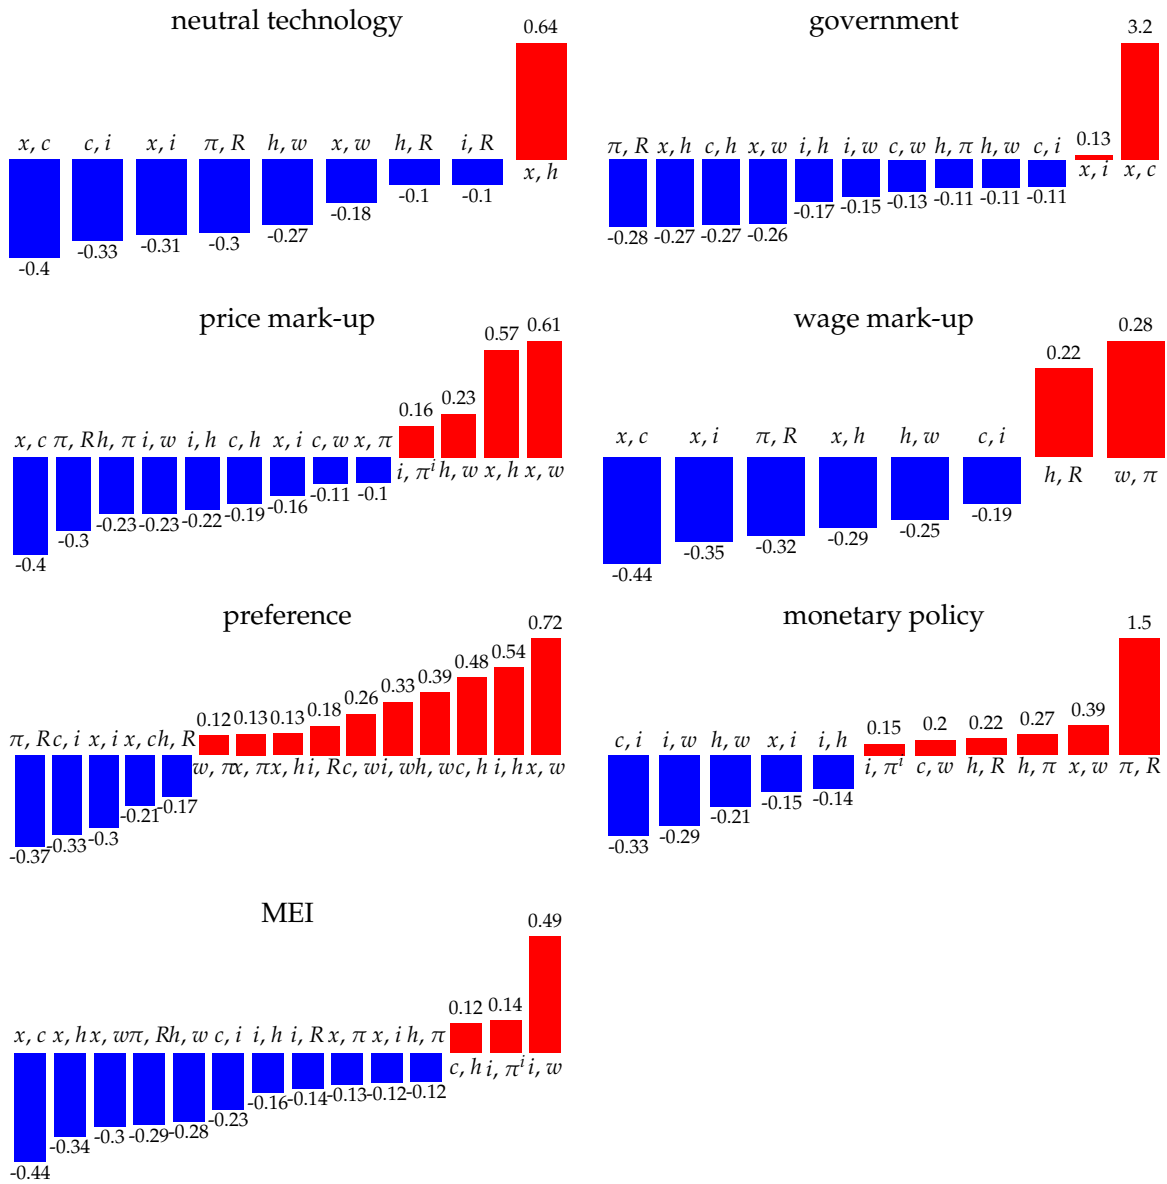


Figure B6: Largest conditional pairwise information complementarities, low spectrum.

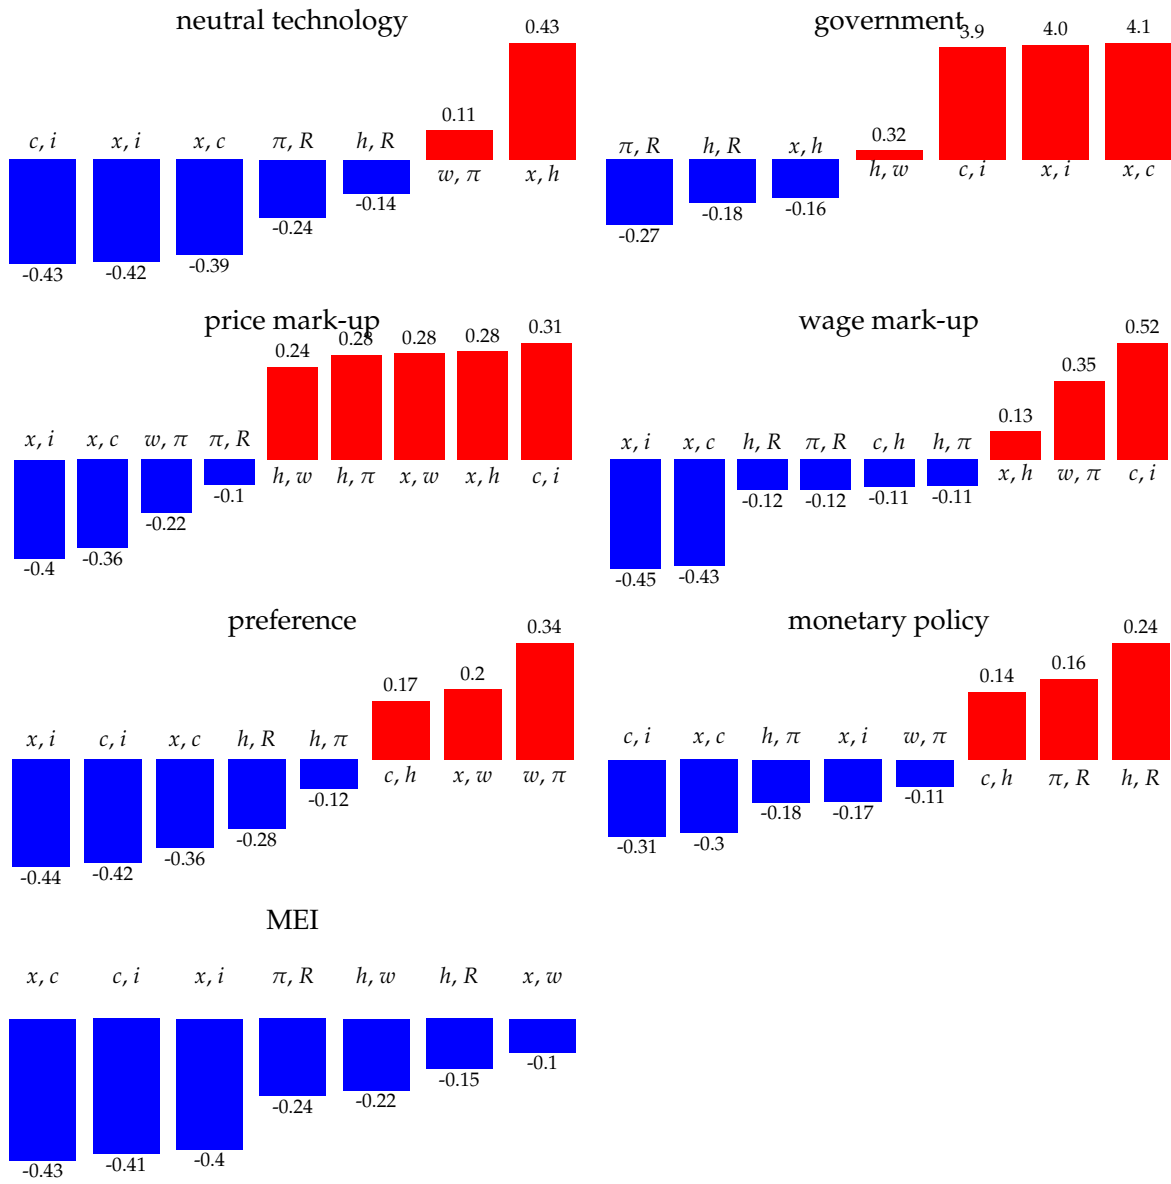


Figure B7: Largest conditional pairwise information complementarities, BC spectrum.

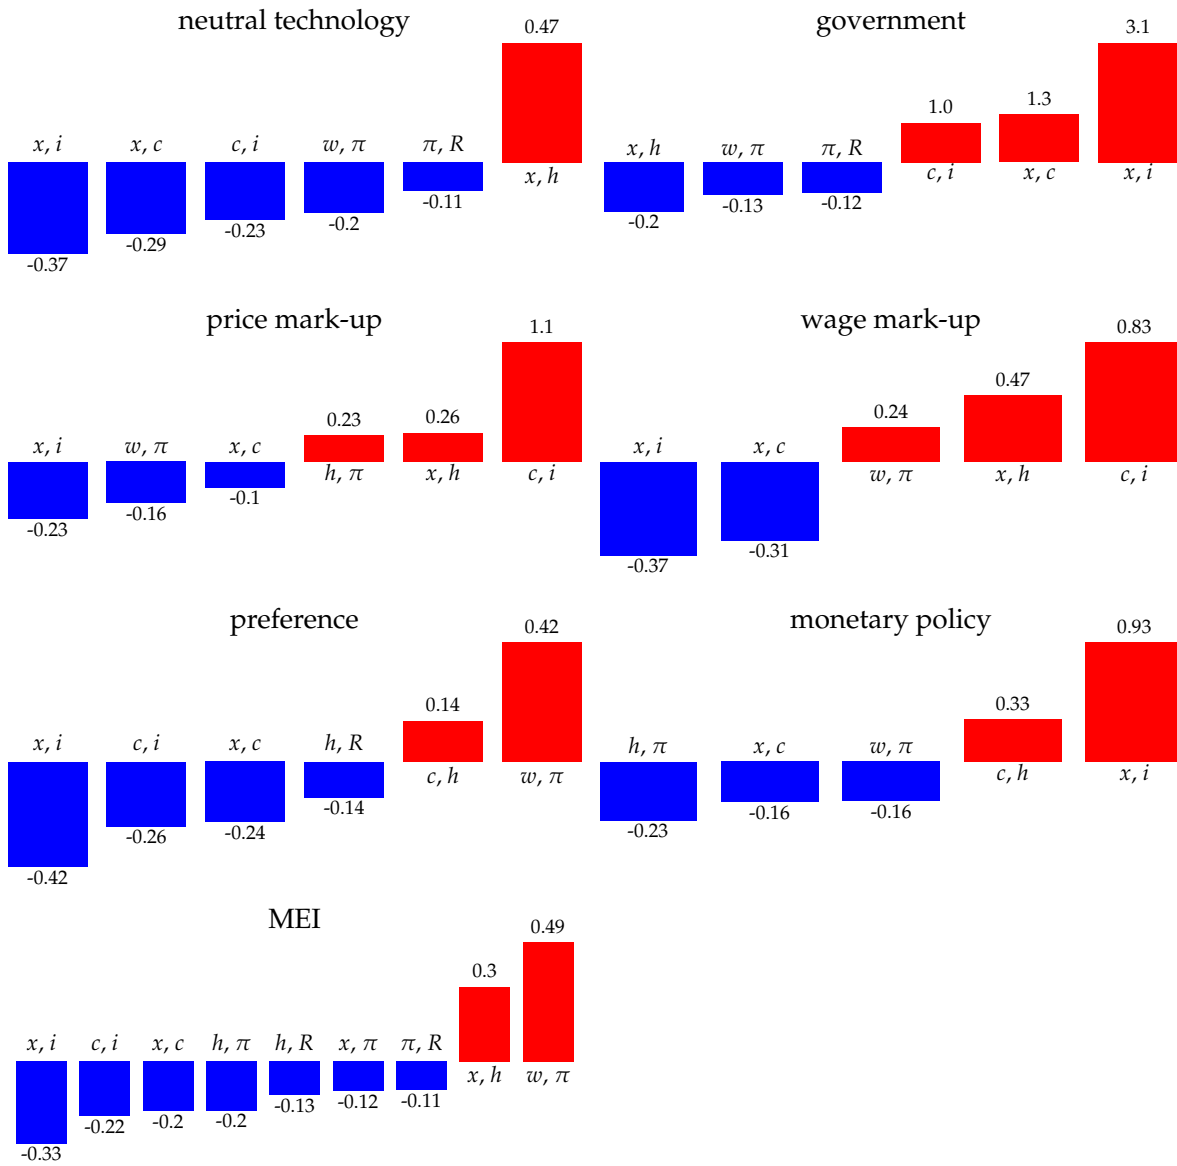


Figure B8: Largest conditional pairwise information complementarities, high spectrum.

C Angeletos, Collard, and Dellas (2018)

C.1 Linearized equilibrium conditions

The economy consists of a continuum of islands and a mainland. Each island contain a representative household and a continuum of monopolistically competitive firms producing a differentiated commodity using labor and capital provided by the household. These commodities are combined through a CES aggregator into an island-specific composite good, which in turn enters the production of the final good in the mainland through another CES aggregator. The final good is used for consumption and investment. The log-linearized equilibrium conditions with variables presented as log-deviations from their steady-state values are summarized as follows:

Optimal consumption allocation

$$\mathbb{E}_{it} [\zeta_t^c + \nu n_{it}] = \zeta_t^c - \frac{c_{it} - bC_{t-1}}{1 - b} + \mathbb{E}_{it} [s_{it} + \varrho Y_t + (1 - \varrho)y_{it} - n_{it}], \quad (\text{C.1})$$

where c_{it} and C_t are consumption on island i and aggregate consumption, y_{it} and Y_t are the quantity of the final good produced in island i and aggregate output, n_{it} is hours worked, s_{it} denotes the realized markup in island i , and ζ_t^c is a preference shock. The parameter ν determines the inverse labor supply elasticity, and the parameters b and ϱ denote the degree of habit persistence, and the degree of substitutability across the islands' composite goods in the production of the final good, respectively.

Optimal investment decision

$$\begin{aligned} \mathbb{E}_{it} [\lambda_{it} + q_{it}] &= \mathbb{E}_{it} [\lambda_{it+1} + \beta(1 - \delta)q_{it+1} + (1 - \beta(1 - \delta))(s_{it+1} + \varrho Y_{t+1} \\ &\quad + (1 - \varrho)y_{it+1} - u_{it+1} - k_{it+1})] \end{aligned} \quad (\text{C.2})$$

where q_{it} is the price of capital, u_{it} is the rate of capital utilization, and λ_{it} is the marginal utility of consumption, given by

$$\lambda_{it} = \zeta_t^c - \frac{c_{it} - bC_{t-1}}{1 - b} \quad (\text{C.3})$$

The parameter β is the intertemporal discount rate in the utility function of the households, and δ is the depreciation rate.

Optimal bond holdings decision

$$R_t = \zeta_t^c - (1 + \nu)n_{it} - s_{it} - \varrho Y_t - (1 - \varrho)y_{it} - \mathbb{E}'_{it} [\lambda_{it+1} - \pi_{it+1}] \quad (\text{C.4})$$

where R_t is the nominal interest rate and π_{it} is the inflation rate in island i .

Equilibrium price of capital

$$q_{it} = (1 + \beta)\varphi\iota_{it} + \varphi\iota_{t-1} - \beta\varphi E'_{it}\iota_{it+1} + \zeta_t^{IP} - \zeta_t^{IT} \quad (\text{C.5})$$

where ι_{it} denotes the level of investment, ζ_t^{IP} is non-stationary investment-specific technology shock, ζ_t^{IT} is a stationary shock shifting the demand for investment, and φ is a parameter governing the size of investment adjustment costs.

Production function

$$y_{it} = \zeta_t^A + \alpha(u_{it} + k_{it}) + (1 - \alpha)n_{it} \quad (\text{C.6})$$

where k_{it} is the local capital stock, ζ_t^A is the level of aggregate TFP, and α is the share of capital in the production function. The capital accumulation equation is

$$k_{it+1} = (1 - \delta)k_{it} + \delta(\zeta_t^{IT} + \iota_{it}), \quad (\text{C.7})$$

and level of TFP is the sum of a permanent (a_t^p) and a transitory (a_t^τ) component:

$$\zeta_t^A = a_t^p + a_t^\tau, \quad (\text{C.8})$$

Resource constraint

$$\varrho y_t + (1 - \varrho)y_{it} = x_{it} + \alpha u_{it}, \quad (\text{C.9})$$

where x_{it} denotes GDP on island i , given by

$$x_{it} = s_c c_{it} + (1 - s_c - s_g)(\zeta_t^{IP} + \iota_{it}) + s_g G_t, \quad (\text{C.10})$$

and G_t , s_c and s_g denote the level of government spending and the steady-state ratios of consumption and government spending to output. To ensure the existence of a balanced growth path, government spending is defined as

$$G_t = \zeta_t^g + \frac{1}{1 - \alpha} a_t^p - \frac{\alpha}{1 - \alpha} \zeta_t^{IP} \quad (\text{C.11})$$

where ζ_t^g a government spending shock.

Equilibrium utilization

$$\zeta_t^{IP} + \frac{1}{1 - \psi} u_{it} = s_{it} + \varrho y_t + (1 - \varrho)y_{it} - k_{it}, \quad (\text{C.12})$$

where ψ is a capital utilization elasticity parameter.

Inflation rate

$$\pi_{it} = \frac{(1-\chi)(1-\beta\chi)}{\chi(1+\chi(1-\beta))} s_{it} + \frac{\beta\chi(1-\chi)\pi_t + \beta\chi E' \pi_{it+1}}{\chi(1+\chi(1-\beta))}, \quad (\text{C.13})$$

where Π_{it} is the aggregate inflation rate, and $(1-\chi)$ is the probability that a firm resets its price in a given period.

Monetary policy rule

$$R_t = \kappa_R R_{t-1} + (1-\kappa_R)(\kappa_\pi \pi_{it} + \kappa_y(x_{it} - x_{it}^F)) + \zeta_t^m \quad (\text{C.14})$$

where x_{it}^F denotes the GDP that would be attained in a flexible-price allocation, ζ_t^m is a monetary policy shock, κ_π and κ_y are parameters determining the policy rate reaction to inflation and the output gap and κ_{Ri} controls the degree of interest-rate smoothing. The flexible-price allocations are obtained from equations (C.1) - (C.12) by setting the realized markup to zero ($s_{it} = 0$) and replacing R_t in (C.4) with the real interest rate.

It is worth pointing out that there are two different subjective expectation operators E_{it} and E'_{it} in the above conditions. In the model, each time period t is divided into two stages: in stage 1, the inhabitants of each island receive an unbiased signal about the level of TFP in that period, and form beliefs that firms and households on other islands receive a signal that is biased by the confidence shock ξ_t , which is also observed. In stage 2, the true state of nature and the realized value of economic activity is publicly revealed. ACD discuss two protocols for the timing of decisions of firms and households, depending on whether supply is determined first and prices adjust to make demand meet supply, or whether demand is determined first and supply adjusts to meet demand. The model presented above is estimated under the second assumption, as seen by the use of stage 1 expectations in the optimality conditions for consumption and saving in equations (C.1), (C.2), and stage 2 expectations in equations (C.4), (C.5), (C.13).

There are nine shocks in the model: a permanent (a_t^p) and a transitory (a_t^t) TFP shock; a permanent (ζ_t^{IP}) and a transitory (ζ_t^{IT}) investment-specific shock; a news shock regarding future productivity (a_t^n); a discount-rate shock (ζ_t^c); a government-spending shock (ζ_t^g); a monetary policy shock (ζ_t^m); and a confidence shock (ξ_t). The later shock is an exogenous random variable observed in stage 1 of each period, representing the perceived bias in the other islands' signals about the level of TFP in that period. The permanent TFP shock is given by

$$a_t^p = a_{t-1}^p + a_{t-1}^n + \varepsilon_t^p, \quad (\text{C.15})$$

and the permanent investment-specific shock follows a random walk

$$\zeta_t^{IP} = \zeta_{t-1}^{IP} + \varepsilon_t^{IP}, \quad (\text{C.16})$$

Table B1: Parameter values, ACD (2018) model

| | parameter | posterior median |
|--------------|---|------------------|
| ψ | utilization elasticity | 0.500 |
| ν | inverse labor supply elasticity | 0.282 |
| α | capital share | 0.255 |
| φ | investment adjustment costs | 3.312 |
| b | habit persistence | 0.758 |
| χ | Calvo parameter, | 0.732 |
| κ_R | Taylor rule smoothing, | 0.198 |
| κ_π | Taylor rule inflation, | 2.271 |
| κ_y | Taylor rule output, | 0.121 |
| ρ_m | AR mon. policy | 0.647 |
| ρ_a | AR transitory TFP component | 0.412 |
| ρ_n | AR news | 0.224 |
| ρ_i | AR transitory investment-specific technology | 0.374 |
| ρ_c | AR preference | 0.888 |
| ρ_g | AR government spending | 0.786 |
| ρ_ξ | AR confidence | 0.833 |
| σ_a^P | std. permanent TFP component | 0.406 |
| σ_a^T | std. transitory TFP component | 0.347 |
| σ_n | std. news | 0.378 |
| σ_i^P | std. permanent investment-specific technology | 0.610 |
| σ_i^T | std. transitory investment-specific shocks | 5.805 |
| σ_c | std. preference | 0.357 |
| σ_g | std. government spending | 1.705 |
| σ_ξ | std. confidence | 0.613 |
| σ_m | std. mon. policy | 0.313 |

where ε_t^p and ε_t^{IP} are i.i.d. innovations. All remaining shocks are stationary AR(1) processes.