Maximum Likelihood in the Frequency Domain: The Importance of Time-to-Plan *

Lawrence J Christiano, Robert J Vigfusson **

Northwestern University, Department of Economics, 2003 Sheridan Road, Evanston, Illinois 60208-2600

Abstract

We illustrate the use of various frequency domain tools for estimating and testing dynamic, stochastic general equilibrium models. We show how to exploit the well known fact that the log, Gaussian density function has a linear decomposition in the frequency domain. We also propose a new resolution to the problem that the phase angle between two variables is not uniquely determined. These methods are applied to the analysis of business cycles. Our substantive findings confirm existing results in the literature, which suggest that time-to-plan in the investment technology has a potentially useful role to play in business cycle analysis.

Key words: Frequency Domain, Time-to-Build, Time-to-Plan, Investment, Business Cycles, Likelihood Ratio, Phase Angle.
JEL classification: C5, E2, E3

1 Introduction

In recent years there has been increased interest in applying formal econometric methods to the analysis of dynamic, stochastic, general equilibrium models.[1] Researchers understand that models are abstractions and so are

* The first author is grateful for financial support from a National Science Foundation grant to the National Bureau of Economic Research. We are grateful for comments from an anonymous referee and from Robert King.
**Corresponding author.
Email addresses: l-christiano@northwestern.edu (Lawrence J Christiano), r-vigfusson@northwestern.edu (Robert J Vigfusson).
[1] Papers that use maximum likelihood methods to study general equilibrium business cycle models include Altug (1989), Bencivenga (1992), Christiano (1988), Chris-

necessarily incorrect. As a result, there is a need for tools which are helpful for diagnosing the reasons for econometric rejections and for identifying the most fruitful avenues for further model development. This paper illustrates a set of tools which we think are useful in this sense. They are frequency domain methods which we use to diagnose Gaussian maximum likelihood estimation and testing results. [2] We apply the tools to analyze a series of real business cycle models.

We exploit the well known fact that the log, Gaussian density function has a linear decomposition in the frequency domain. This decomposition has two implications. First, the likelihood ratio statistic for testing a model can be represented as the sum of likelihood ratios in the frequency domain. [3] So, if a model is rejected because of a large likelihood ratio statistic, then it is possible to determine which frequencies of the data are responsible. Second, if parameter estimates look 'strange', then it is possible to determine which frequencies are responsible for the result.

We also illustrate the use of the spectrum, phase angle and coherence functions for diagnosing model estimation and testing results. Regarding the phase angle, it is well known that this object is not uniquely determined. We propose a resolution to this problem which we hope is of independent interest.

We begin with a univariate analysis which studies output growth in a version of the real business cycle model in which the only shock is a disturbance to technology. We then extend the analysis by including a second variable, business fixed investment. To avoid a statistical singularity, we must introduce a second shock into the model. For this, we include disturbances to government consumption. This aspect of our analysis illustrates how our approach can be extended to a vector of time series.

The key substantive finding of the paper is that the data support a 'time-to-plan' specification of the investment technology. With this specification, it takes several periods to build new capital, with a lengthy initial period being devoted to activities such as planning, which do not require an intensive

---

tiano, Eichenbaum, and Marshall (1991), Hall (1996), Hansen and Sargent (1980, 1991), Ireland (1997,1999,2001), Kim (2000), Leeper and Sims (1994), McGrattan (1994), McGrattan, Rogerson, and Wright (1997).

[2] For an early paper exploiting the advantages of the frequency domain, see Engle (1974). Other papers which explore the advantages of the frequency domain - though not in the maximum likelihood context explored here - include Baxter and King (1999), Christiano and Fitzgerald (1999) and Cogley (2001a,b).

[3] A referee suggested an analogous decomposition in the time domain. This could be constructed using the prediction error decomposition of the Gaussian density function. Although it is beyond the scope of this paper, it would be of interest to explore this statistic.

application of resources. Maximum likelihood favors the time-to-plan specification because it helps the model to account for persistence in output growth and for the fact that business fixed investment lags output in the business cycle frequencies.[4] Our results complement Christiano and Todd (1996) and Bernanke, Gertler and Gilchrist (1999), which also present evidence in favor of time-to-plan.

We now consider the relationship of our paper to the existing literature. The fact that the Gaussian density function can be decomposed in the frequency domain has been exploited in several other papers. For example, Altug (1989) demonstrates its value for estimating models with measurement error. Other papers emphasize its value in the estimation of time-aggregated models.[5] Christiano and Eichenbaum (1987) and Hansen and Sargent (1993) exploit the decomposition to evaluate the consequences for maximum likelihood estimates of certain types of model specification error.[6] Finally, the value of comparing model and data spectra has also been emphasized in the recent contributions of Watson (1993), Diebold, Ohanian and Berkowitz (1998), and Berkowitz (2001).

The following section presents our econometric framework. Sections 3 and 4 present the results. Section 5 concludes.

2 Econometric Framework

This section describes the econometric framework of our analysis. First, we display the frequency domain decomposition of the Gaussian density function. Second, we derive the likelihood function associated with the various real business cycle models that we consider. Third, we derive the likelihood function of various unrestricted reduced form representations of the data. In the last subsection, we use the model and reduced form log-likelihood functions to form a likelihood ratio statistic for testing the model. We display the linear, frequency domain decomposition of the likelihood ratio statistic.

[4] That standard business cycle models have difficulty accounting for persistence in output growth is well known. See, for example, Christiano (1988, p. 274), Cogley and Nason (1995), and Watson (1993).

[5] See, for example, Hansen and Sargent (1980a), Christiano (1985), Christiano and Eichenbaum (1987) and Christiano, Eichenbaum and Marshall (1991).

[6] These approaches to specification error analysis are similar in spirit to the early approach taken in Sims (1972). See Berkowitz (2001) for a related discussion.

2.1 *Spectral Decomposition of the Gaussian Likelihood*

Suppose we have a time series of data, $y = [y_1, ..., y_T]$, where y_t is a finite-dimensional column vector with zero mean. It is well known (Harvey, 1989, p. 193) that for T large, the Gaussian likelihood for these data is well approximated by:

$$L(y) = -\frac{1}{2}\sum_{j=0}^{T-1}\left\{2\ln 2\pi + \ln\left[\det\left(F(\omega_j;\Phi)\right)\right] + \mathrm{tr}\left(F(\omega_j;\Phi)^{-1}I(\omega_j)\right)\right\} \quad (1)$$

where $\mathrm{tr}(\cdot)$ and $\det(\cdot)$ denote the trace and determinant operators, respectively. Also, $I(\omega)$ is the periodogram of the data:

$$I(\omega) = \frac{1}{2\pi T}y(\omega)y(-\omega)', y(\omega) = \sum_{t=1}^{T} y_t \exp(-i\omega t), \quad (2)$$

and

$$\omega_j = \frac{2\pi j}{T}, j = 0, 1, ..., T-1.$$

Finally, $F(\omega;\Phi)$ is the spectral density of y at frequency ω, and Φ is a vector of unknown parameters.[7]

For diagnostic purposes, it is sometimes useful to express the likelihood function in the following weighted form:

$$L(y) = -\frac{1}{2}\sum_{j=0}^{T-1} v_j\left\{2\ln 2\pi + \ln\left[\det\left(F(\omega_j)\right)\right] + \mathrm{tr}\left(F(\omega_j)^{-1}I(\omega_j)\right)\right\}, \quad (3)$$

where $v_j \in \{0,1\}$ for all j.[8] By setting various v_j's to zero and reestimating, one can identify the impact of various frequencies on parameter estimates.

[7] Let $C(k;\Phi) = Ey_t y_{t-k}'$, for integer values of k. Then,

$$F(\omega;\Phi) = \frac{1}{2\pi}\sum_{k=-\infty}^{\infty} C(k;\Phi)e^{-i\omega k},$$

for $\omega \in (0, 2\pi)$.

[8] Diebold, Ohanian and Berkowitz (1998) also work with (3), as a way of abstracting altogether from certain frequencies of the data.

2.2 Likelihood Function for Structural Models

The preceding discussion indicates that to estimate a model by frequency domain maximum likelihood, one needs the mapping from the model's parameters, Φ, to the spectral density matrix of the data, $F(\omega_j;\Phi)$. This subsection defines this mapping for several models. We consider the standard real business cycle model and the time-to-build model of Kydland and Prescott (1982). In each case, we consider a one-shock and a two-shock version of the model.

In the one shock model there is only a disturbance to technology. In the two shock model there is, in addition, a disturbance to government consumption. In the analysis of the one-shock model, we focus on its implications for output growth, so that

$$y_t = \ln(Y_t/Y_{t-1}), \tag{4}$$

where Y_t denotes gross output in period t. In the analysis of the two-shock model,

$$y_t = \begin{bmatrix} \ln(Y_t/Y_{t-1}) \\ \ln(I_t/Y_t) \end{bmatrix}, \tag{5}$$

where I_t denotes gross investment. We always consider y_t expressed as a deviation from the model's population mean.

2.2.1 Model Used in the Analysis

Household preferences, the resource constraint and the production function are taken from Christiano and Eichenbaum (1992). Preferences and the resource constraint are:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\ln(C_t + \xi G_t) + \psi \ln(1-n_t)\right], C_t + G_t + I_t \leq Y_t,$$

where C_t and G_t denote household and government consumption, respectively, and n_t denotes the fraction of available time worked. The endowment of available time is normalized to unity and $\beta = 1.03^{-0.25}$, $\psi = 3.92$. We consider two values of ξ, $\xi = 0, 1$. When $\xi = 1$, then G_t has no effect on the dynamics of output and investment.[9] We refer to this as the *one-shock model.* When

[9] When $\xi = 1$, C_t and G_t appear symmetrically everywhere. Exploiting this, we solve the model for $\tilde{C}_t$, where $\tilde{C}_t = C_t + G_t$. There are two interpretations of this

$\xi = 0$, then G_t does matter and so we refer to this as the *two-shock model.*

The production function is:

$$Y_t = K_t^{\theta} \left(z_t n_t\right)^{(1-\theta)}, 0 < \theta < 1,$$

where K_t denotes the beginning-of-period t stock of capital and z_t denotes the state of technology. The latter is assumed to evolve as follows:

$$\ln(z_t) = \ln(z_{t-1}) + \varepsilon_{zt},$$

where ε_{zt} a Normal random variable that is independently and identically distributed (i.i.d.) over time, with mean μ and variance σ_z^2. We specify the time series process for G_t as follows:

$$\ln(g_t) = (1-\rho)\ln(g) + \rho \ln(g_{t-1}) + \varepsilon_{gt}, g_t = \frac{G_t}{z_t},$$

where g is a constant, $-1 < \rho < 1$, and ε_{gt} is i.i.d. Normal with mean 0 and variance σ_g^2.

It remains to specify how investment contributes to the evolution of the capital stock. In the *Real Business Cycle* model (RBC), the construction of new capital requires one period:

$$K_{t+1} - (1-\delta)K_t = I_t, 0 < \delta < 1.$$

We denote the parameters of the one-shock version of the RBC model by $\Phi^r = (\sigma_z, \delta, \theta)$, where the superscript, r, stands for 'restricted'. In the two-shock version of the model, $\Phi^r = (\sigma_z, \delta, \theta, \rho, \sigma_g)$. The superscript, r, indicates that these are model parameters. This notation allows us to differentiate model parameters from those of the unrestricted reduced form, Φ^u, which are discussed below. To prevent a profusion of notation, we do not also index Φ according to the number of shocks or the type of structural model.

The *time-to-build* model adopts Kydland and Prescott's (1982) formulation, which specifies that it takes four periods to construct new capital. Period t

solution. One is that it is the solution to the model stated in the text, where private consumption is $C_t = \tilde{C}_t - G_t$. This is a valid interpretation as long as $C_t \geq 0$. In the stochastic process for G_t that we use below, this is true with very high probability and so we ignore the possibility, $C_t < 0$. An alternative interpretation is that ours is the solution to a model without government, where $\tilde{C}_t$ is the consumption of the household. Either way, it is clear that government shocks have no impact on investment and output.

investment is:

$$I_t = \phi_1 x_t + \phi_2 x_{t-1} + \phi_3 x_{t-2} + \phi_4 x_{t-3},$$

where $\phi_i \geq 0$ for $i = 1, 2, 3, 4$, and

$$\phi_1 + \phi_2 + \phi_3 + \phi_4 \equiv 1.$$

The investment technology specifies that if net investment in period $t+3$ is x_t, i.e.,

$$K_{t+4} - (1-\delta)K_{t+3} = x_t,$$

then, resources in the amount $\phi_1 x_t$ must be applied in period t, $\phi_2 x_t$ must be applied in period $t+1$, $\phi_3 x_t$ must be applied in period $t+2$, and finally, $\phi_4 x_t$ must be applied in period $t+3$. Once initiated, the scale of an investment project cannot be expanded or contracted. In the one-shock version of the model, $\Phi^r = (\sigma_z, \phi_1, \phi_2, \phi_3)$. In the two-shock version of the model, $\Phi^r = (\sigma_z, \phi_1, \phi_2, \phi_3, \rho, \sigma_g)$.

In each case, the solution to the model is a set of stochastic processes for y_t and the other variables, which maximize utility subject to the various constraints.

2.2.2 *Reduced Form Representation and Likelihood Function*

We used the undetermined coefficient method described in Christiano (2001) to develop a linear approximation to the y_t process which solves the model:

$$y_t = \alpha(L; \Phi^r)\varepsilon_t = \alpha_0(\Phi^r)\varepsilon_t + \alpha_1(\Phi^r)\varepsilon_{t-1} + \alpha_2(\Phi^r)\varepsilon_{t-2} + ... \qquad (6)$$

In the one-shock model, y_t is defined in (4), and

$$\varepsilon_t = \varepsilon_{zt}, V(\Phi^r) \equiv E\varepsilon_t\varepsilon_t' = \sigma_z^2.$$

The scalar polynomial, $\alpha(L; \Phi^r)$, in the lag operator L, is the infinite moving average representation corresponding to an autoregressive, moving average process with 4 autoregressive and 8 moving average lags, i.e., an $ARMA(4, 8)$.

In the two-shock model, y_t is defined in (5), $\alpha(L; \Phi^r)$ is 2×2 matrix polynomial in L and

$$\varepsilon_t = \begin{pmatrix} \varepsilon_{zt} \\ \varepsilon_{gt} \end{pmatrix}, V(\Phi^r) = \begin{bmatrix} \sigma_z^2 & 0 \\ 0 & \sigma_g^2 \end{bmatrix}.$$

In this case, $\alpha(L;\Phi^r)$ is the infinite moving average representation corresponding to a vector ARMA model with 5 autoregressive and 8 moving average lags, i.e., a VARMA(5,8).[10]

In all cases, we restrict Φ^r so that

$$\sum_{i=0}^{\infty}\alpha_i(\Phi^r)V(\Phi^r)\alpha_i(\Phi^r)' < \infty,$$

guaranteeing that the spectral density of y_t exists. We also restrict Φ^r so that $\det[\alpha(z;\Phi^r)] = 0$ implies $|z| \geq 1$, where $|\cdot|$ denotes the absolute value operator.

The spectral density of y_t at frequency ω is

$$F^r(\omega;\Phi^r) = \frac{1}{2\pi}\alpha(e^{-i\omega};\Phi^r)V\alpha(e^{i\omega};\Phi^r)',$$

where the superscript, r, on F indicates that the form of $\alpha(L;\Phi^r)$ is restricted by the model. The frequency domain approximation to the restricted likelihood function is (1) with $F(\omega)$ replaced by $F^r(\omega;\Phi^r)$.

2.3 Unrestricted Reduced Form Likelihood

In order to test our model, we need to estimate an unrestricted version of (6):

$$y_t = \alpha(L)\varepsilon_t, \tag{7}$$

where

$$\alpha(L) = I + \alpha_1 L + \alpha_2 L^2 + \dots.$$

Here $I = 1$ for the one-shock model and I is the 2×2 identity matrix for the two-shock model. Also,

$$\sum_{i=0}^{\infty}\alpha_i V\alpha_i' < \infty,$$

where V is the variance, covariance matrix of ν_t. Finally, we require that $\det[\alpha(z)] = 0$ implies $|z| \geq 1$.

[10] See Appendices B and C in Christiano and Vigfusson (2001) for details about the reduced form implications of the two structural models.

In the one-shock case, the polynomial in L, $\alpha(L)$, corresponds to the ratio of an 8^{th} order polynomial and a 4^{th} order polynomial, with constant terms normalized to unity. This specification nests the real business cycle model and the time to build model. It has 13 free parameters: the 12 parameters of $\alpha(L)$, and V. We denote these by the 13 dimensional vector, Φ^u.

In the two-shock model, we approximate $\alpha(L)$ with a 10 lag vector autoregression, i.e., VAR(10). Christiano and Vigfusson (2001, pg. 31) argue that this is a good approximation to the VARMA(5,8) reduced form of the two-shock structural model. As in the one-shock case, we denote the parameters of the unrestricted reduced form by Φ^u. This is a 45-element vector.

Let $F^u(\omega; \Phi^u)$ denote the spectral density of y_t:

$$F^u(\omega; \Phi^u) = \frac{1}{2\pi}\alpha(e^{-i\omega})V\alpha(e^{i\omega})'.$$

The frequency domain approximation to the unrestricted likelihood function is (1) with $F(\omega; \Phi)$ replaced by $F^u(\omega; \Phi^u)$.

2.4 Cumulative Likelihood Ratio

The likelihood ratio statistic is

$$\lambda = 2(L^u - L^r), \tag{8}$$

where L^r and L^u are the maximized values of the restricted and unrestricted log likelihoods, respectively. Under the null hypothesis that the restricted model is true, this statistic has a chi-square distribution with degrees of freedom equal to the difference between the number of parameters in the restricted and unrestricted models (Harvey, 1989, p. 235). Define

$$\lambda(\omega) = \ln \frac{\det\left[F^r(\omega; \widehat{\Phi^r})\right]}{\det\left[F^u(\omega; \widehat{\Phi^u})\right]} + \operatorname{tr}\left[\left(F^r(\omega; \widehat{\Phi^r})^{-1} - F^u(\omega; \widehat{\Phi^u})^{-1}\right) I(\omega)\right], \tag{9}$$

where a hat over a variable indicates its estimated value. Then, it is easily confirmed that,

$$\lambda = \sum_{j=0}^{T-1} \lambda(\omega_j).$$

This expression can be simplified because of the symmetry properties of $\lambda(\omega)$:[11]

$$\lambda(\omega_{\frac{T}{2}-l}) = \lambda(\omega_{\frac{T}{2}+l}), l = 1, 2, ..., \frac{T}{2} - 1.$$

These imply that λ can be written:

$$\lambda = \lambda(0) + 2\sum_{j=1}^{\frac{T}{2}-1} \lambda(\omega_j) + \lambda(\pi). \tag{10}$$

This is our linear, frequency domain decomposition of the likelihood ratio statistic.

If λ is large, then we should be able to determine which ω_j's are responsible. It is useful to define the cumulative likelihood ratio:

$$\begin{aligned}
\Lambda(\omega) &= \lambda(0) + 2\sum_{\omega_j \leq \omega} \lambda(\omega_j), 0 < \omega < \pi \\
\Lambda(0) &= \lambda(0), \\
\Lambda(\pi) &= \lambda.
\end{aligned} \tag{11}$$

A sharp increase in $\Lambda(\omega)$ in some region of ω's signals a frequency band where the model fits poorly. Note that although $\Lambda(\pi) \geq 0$, it is possible for $\Lambda(\omega)$ to decrease over intervals, since there is nothing preventing some $\lambda(\omega_j)$'s from being negative.

3 Results for One-Shock Models

This section presents our results for estimating and testing the one-shock versions of the RBC and time-to-build models. The periodogram of the data, (2), and the spectral density of the unrestricted reduced form are important ingredients in the analysis, and so we begin by presenting these. The subsequent two subsections present our analysis of the RBC and the time-to-build models, respectively.

We find that when the RBC model is parameterized using standard values taken from the literature, it is strongly rejected. By contrast, when we estimate and test the time-to-build model, it fails to be rejected by the data.

[11] Implicitly, we assume T is even. The adjustment when T is odd is straightforward.

3.1 Periodogram and Spectrum of Unrestricted Reduced Form

Figure 1 presents a smoothed version of $I(\omega)$ for $\omega \in (0, \pi)$, based on (2).[12] Figure 1 also displays the spectrum of our unrestricted $ARMA(4,8)$ representation of US GDP growth. An estimate of the associated 95 percent confidence interval is also reported.[13] The two estimates of the spectral density are fairly similar. We interpret this as evidence that our unrestricted time series model is a good representation of the second moment properties of the data. This is necessary, if the model is to be well-suited as a basis for computing the unrestricted likelihood value in the likelihood ratio statistic, (8).

Vertical bars draw attention to three frequency bands, the low frequencies (those corresponding to period 8 years to infinity), the business cycle frequencies (period 1 year to 8 years) and the high frequencies (period 2 quarters to 1 year). Note that the spectrum of output growth falls throughout the business cycle frequencies. Because the spectrum is relatively high in these frequencies, this corresponds to the observation that output growth is positively autocorrelated. In addition, the spectrum has pronounced dips in the $7 - 7.5$ months (near $\omega = 2.5$) range and in the higher frequency component of the business cycle (near $\omega = 1.5$).[14]

[12] The data are seasonally adjusted, cover the period 1955Q3 to 1997Q1, are taken from the DRI Basic database and have mnemonic GDPQ. The sample mean of y_t is subtracted from the data, so that $I(0)$ is zero. We present the smoothed version of the periodogram because, as is well known, the unsmoothed periodogram is quite volatile. The smoothed periodogram at frequency ω_j is a centered, equally weighted average, $\sum_{i=-3}^{3} I(\omega_{j+i})/7$.

[13] Confidence intervals are obtained by adding and subtracting 1.96 times the relevant standard error estimate. Standard errors were obtained using the standard 'delta function' method based on the point estimates and estimated sampling variance covariance matrix associated with our estimator of the parameters, Φ, of the $ARMA(4,8)$ representation.

[14] The $\omega = 1.5$ frequency corresponds roughly to a one-year seasonal, and so the dip in the spectrum here may be an effect of the seasonal adjustment procedure applied to the data. (See Nerlove, Grether and Carvalho (1979) for a review of the relevant literature.) We follow conventional practice in ignoring the fact that our data have been seasonally adjusted at the source. Still, the proper way to integrate seasonality into an analysis like ours remains an important outstanding subject for research. For recent work in this direction, see Hansen and Sargent (1993) and Christiano and Todd (2001). We expect that the kind of spectral techniques used in this paper will be useful in any analysis that carefully integrates seasonality.

In our analysis of the RBC model, we estimated the variance of the technology shock, σ_z^2. The other parameter values of the model are fixed at the estimated values reported in Christiano and Eichenbaum (1992), $\theta = 0.344$, $\delta = 0.021$. Parameter values for this and other models in this paper are reported in Table 1.

The spectrum of output growth implied by the estimated RBC model is displayed in Figure 1. Consistent with the findings in Watson (1993), that spectrum is essentially flat. Figure 2 presents a formal evaluation of the fit of the RBC model using the cumulative likelihood ratio, (11). The likelihood ratio statistic, λ, is just under 25 (see the cumulative likelihood ratio for $\omega = \pi$). Under the null hypothesis that the restricted RBC model is true, λ is the realization of a chi-square statistic with 12 degrees of freedom. The statistic has a p-value of 1.5 percent and hence the model is rejected at the five percent significance level.

There are two reasons that the likelihood ratio statistic is so high. First, the model fails to reproduce the negative slope of the spectrum in the business cycle frequencies. This failure is manifested in a sharp rise in the cumulative likelihood ratio in the low end of the business cycle frequencies. Second, the model fails to reproduce the dip in the spectrum in frequencies corresponding to periods 7-7.5 months. This failure is manifest in a sharp rise in the cumulative likelihood ratio after frequency $\omega = 2.5$.[15]

[15] There are simple perturbations on the RBC model that can improve its fit. But, these require parameterizations that are inconsistent with evidence from other data. For example, when we estimated the RBC model allowing θ and δ to also be free, we obtained $\widehat{\theta} = 0.37$, $\widehat{\delta} = 0.73$, $\hat{\sigma}_z = 0.0144$. The main difference between this model and the version of the estimated RBC model in the text is the high depreciation rate. The large value of δ improves the fit of the model by causing the spectrum of output growth to be negatively sloped. To see why, note that in the extreme case where $\delta = 1$, the model reduces to the model of Long and Plosser (1993), which has the property that output growth is a first order autoregression with autocorrelation coefficient θ. For $\theta > 0$, such a process has a negatively sloped spectrum. Note from Figure 1 that the negative slope in the spectrum of output growth is particularly pronounced in the business cycle frequencies. Consistent with the intuition in this footnote, when we reestimate the model with the v_j's in (3) corresponding to non-business cycle frequencies set to zero, we obtain $\hat{\delta} \approx 1$.
Although the estimated value of capital's share in the perturbed RBC model is reasonable, the estimated value of δ is substantially larger than what is plausible in light of data on investment and the stock of capital (Christiano and Eichenbaum, 1992).
An alternative way to accommodate the negatively sloped spectrum in the data is to add persistence to the growth rate of technology in the RBC model. This would have

3.3 Estimation and Testing of Time-To-Build Model

Next, we estimated the time-to-build model. The estimated time-to-build weights are: $\hat{\phi}_1 = 0.01$, $\hat{\phi}_2 = 0.28$, $\hat{\phi}_3 = 0.48$, and $\hat{\phi}_4 = 0.23$. Two features of these estimates are worth noting. First, the estimated value of ϕ_1 is nearly zero. This implies that in the first period of an investment project, essentially no resources are used. An interpretation is that investment projects must start with a planning period, in which plans are drawn up, permits are secured, etc. Christiano and Todd (1996) refer to a model of this type as a 'time-to-plan' model, and argue that it is consistent with microeconomic evidence on investment projects. Second, the resource usage in later periods of investment follows a 'hump' shape.

The spectrum of output growth implied by the estimated time-to-build model is displayed in Figure 1. The model spectrum conforms well with the spectrum of the data. It even matches the dip in the 7-7.5 month range. This is reflected in the good performance of the model's cumulative likelihood ratio (see Figure 2). The cumulative likelihood ratio rises slowly with frequency and achieves a maximum value of a little under 10. Under the null hypothesis that the model is true, this is the realization of a chi-square distribution with 9 degrees of freedom. Under these conditions, the p-value is 35 percent. As a result the model is not rejected at conventional levels.

It is useful to compare the estimated time-to-build model with two others: the time-to-build model in Kydland and Prescott (1982), where $\phi_i = 0.25$, $i = 1, 2, 3, 4$; and the time-to-build model in Christiano and Todd (1996), where $\phi_1 \approx 0$, $\phi_i = 1/3$, $i = 2, 3, 4$. We refer to these as the Kydland-Prescott and Christiano-Todd models, respectively.[16] We do not display the spectrum or cumulative likelihood ratio implied by the Kydland-Prescott model, because they coincide with the ones implied by the RBC model (King, 1995).[17]

Comparing the cumulative likelihood ratio statistics for these models allows us to understand what is responsible for the time-to-plan and hump-shape in the time to build weights of the estimated model. The time-to-plan feature

increased the persistence of output growth in the model and improved its fit with the output data. However, as emphasized in Christiano (1988), this improvement in fit would come at the cost of counterfactual implications for the growth rate of the Solow residual.

[16] What we call the Kydland-Prescott model is our time-to-build model, with the ϕ_i's restricted as indicated in the text. We estimate the shock variance by maximizing the likelihood function with respect to that parameter. We treat the Christiano-Todd model in the same way.

[17] For a detailed discussion of the similarity of these models, see Christiano and Todd (1996) and Rouwenhorst (1991).

helps the model account for the negative slope in the spectrum in the lower business cycle frequencies. To see this, note that models with time-to-plan, i.e., the Christiano-Todd and estimated models, perform well in the lower business cycle frequencies. Models without time-to-plan, i.e., the RBC model (and, hence, the Kydland-Prescott model too), do relatively poorly. The hump shape feature of the time-to-build weights help the model capture the dip in the spectrum in the 7-7.5 month range. To see this, note that the model without a hump in time-to-build weights, i.e., the Kydland-Prescott and Christiano-Todd models, perform badly for $\omega > 2.5$. The model with a hump shape in the time-to-build weights, i.e., the estimated model, does relatively well.

4 Results for Two-Shock Model

We now analyze the version of the time-to-build model with both government consumption and technology shocks, using data on output and business investment. The first subsection below reports the spectral properties of the data. The second subsection reports the estimation and testing results for the model. We begin by briefly summarizing the key findings.

Four features of data play a key role in our model analysis. First, as discussed in the previous section, the spectrum of output growth is high and falling in the business cycle frequencies, suggesting persistence in those data. Second, the spectrum of the investment to output ratio exhibits a very steep, negative slope. Third, the coherence between output and investment varies over frequencies, exhibiting high coherence in the low range of business cycle frequencies and low coherence elsewhere. Finally, phase angle analysis suggests that investment lags output in the business cycle frequencies.

There are two notable features of the parameter estimates. First, as in the previous section, maximum likelihood selects the time-to-plan specification of the investment technology.[18] In part, this is to help account for the negative slope in the spectrum of output growth. In addition, time-to-plan helps the model to account for the phase angle between investment and output data. Second, the parameter estimates assign a relatively large role to the government consumption shock. This helps the model capture the coherence pattern between the output and investment data.

[18] The estimated values of ϕ_2, ϕ_3, ϕ_4 also display a hump-shape pattern. This is qualitatively similar to what we found in the univariate case, although the quantitative magnitude of the hump is smaller. The statistical rationale for the hump in the multivariate case is the same as in the univariate case and has already been discussed.

Our estimated model has weaknesses. First, the estimated variance of government consumption seems implausibly large in that it exceeds direct estimates based on government consumption data reported in Christiano and Eichenbaum (1992). Second, despite the high variance of government consumption, the likelihood ratio statistic based on output and investment data rejects the model statistically. After studying the frequency domain decomposition of our likelihood ratio statistic, we argue that the rejection reflects difficulties the model has in matching the components of the data with period 2.5 years and longer. This appears to reflect three difficulties: (i) although the model can capture the low coherence overall between investment and output, it cannot at the same time capture the relatively high coherence in the low business cycle frequencies; (ii) it has difficulty quantitatively matching the shape of the spectrum of the investment to output ratio in the low frequencies; and (iii) although it gets the sign of the phase angle between investment and output right, it misses quantitatively.

4.1 Spectral Properties of the Data

Figure 3 displays the estimated spectral density for the business investment to output ratio implied by the unrestricted VAR(10) discussed in section 2.3.[19] The associated 95 percent confidence intervals computed using the delta function method are also displayed. In addition, Figure 3 reports the smoothed periodogram estimate of the spectrum. The two spectral estimates are reasonably similar. For most frequencies, the periodogram-based estimate lies inside the confidence interval implied by the VAR(10). We do not report the spectral density for output growth implied by the VAR(10), because it is similar to the spectrum in Figure 1.

The spectral densities of $\ln(Y_t/Y_{t-1})$ and $\ln(I_t/Y_t)$ differ notably. The spectrum of $\ln(I_t/Y_t)$ is much steeper than that of $\ln(Y_t/Y_{t-1})$.

In considering the cross-spectrum between the variables in our model, we find it convenient to work with $\ln(Y_t)$ and $\ln(I_t)$ instead of $\ln(Y_t/Y_{t-1})$ and $\ln(I_t/Y_t)$. Focusing on levels promotes comparability with the business cycle literature. According to our model and to our unrestricted VAR(10), the levels data are not covariance stationary processes. Still, their matrix spectral density

[19] We measure investment as seasonally adjusted business investment in structures and equipment, which cover the period 1955Q3 to 1997Q1, and are taken from the DRI Basic database. Business investment in structures has mnemonic GSVNT and business investment in equipment has mnemonic GIPNR. The investment to output ratio was measured as the ratio of the sum of these two series to nominal GDP. The latter has DRI Basic mnemonic GDP.

is well defined, except at frequency zero. The density is obtained by a simple matrix manipulation of the spectral density of $\ln(Y_t/Y_{t-1})$ and $\ln(I_t/Y_t)$.[20]

Denote the matrix spectral density of $[\ln(Y_t)\ln(I_t)]$ by $S(\omega)$, for $0<\omega\leq\pi$. For fixed ω, the coherence between $\ln(Y_t)$ and $\ln(I_t)$, $R^2(\omega)$, measures the strength of the linear relationship between these variables at frequency ω:

$$R^2(\omega)=\frac{S_{12}(\omega)S_{12}(-\omega)}{S_{11}(\omega)S_{22}(\omega)}. \tag{12}$$

Figure 4 displays the coherence function computed based on the periodogram and on the VAR(10). In addition, that figure exhibits a 95 percent confidence interval for the coherence, based on applying the delta function method to the estimated VAR(10). Note again that the two estimates of the spectrum are quite similar. The similarity between the periodogram-based and VAR(10)-based estimates of the spectral density is consistent with the notion that the latter provides a good summary of the spectral properties of the data.

A distinctive feature of the coherence function is that it is particularly high in the business cycle frequencies. The peak is near 0.9 for cycles with period in the neighborhood of 3 years.

Figure 5 reports the phase angle between $\ln(Y_t)$ and $\ln(I_t)$. The phase angle measured in radians, $\theta_{21}^*(\omega)$, satisfies

$$S_{21}(\omega)=r(\omega)\,e^{i\theta_{21}^*(\omega)},\omega\in(0,2\pi), \tag{13}$$

and one additional condition that is stated below. In (13), $r(\omega)$ is the gain function:

$$r(\omega)=\sqrt{S_{21}(\omega)\,S_{21}(-\omega)}.$$

The phase relationship measured in units of time is given by $k_{21}^*(\omega)=-\theta_{21}^*(\omega)/\omega$.

[20] Let $F(\omega)$ denote the two-dimensional spectral density of $[\ln(Y_t/Y_{t-1})\ln(I_t/Y_t)]$. Then, the two-dimensional spectral density of $[\ln(Y_t)\ln(I_t)]$, $S(\omega)$, is given by

$$S(\omega)=H(\omega)F(\omega)H(-\omega)^T,$$

where

$$H(\omega)=\begin{bmatrix}\frac{1}{1-e^{-i\omega}} & 0\\ -\frac{1}{1-e^{-i\omega}} & 1\end{bmatrix}.$$

The matrix, $H(\omega)$ - and, hence, $S(\omega)$ - is well defined for $0<\omega\leq\pi$.

As is well known, there are many $\theta^*_{21}(\omega)$'s (and, hence $k^*_{21}(\omega)$'s) that satisfy (13) for each ω. We select one of these by requiring that $k^*_{21}(\omega)$ be the value of k which maximizes the covariance between the components of $\ln I_t$ and $\ln Y_{t-k}$ in a neighborhood of frequency ω, for $k = 0, \pm 1, \pm 2, ...$. Thus, $k^*_{21}(\omega) > 0$ means that investment lags output by $k^*_{21}(\omega)$ periods when only the components of the variables in a neighborhood of frequency ω are considered. Further details about the interpretation and computation of the phase angle are discussed in the appendix.

Figure 5 displays $k^*_{21}(\omega)$, $0.2 \leq \omega \leq \pi$. The grey area in the figure indicates the 95 percent confidence interval computed using the delta function method using the estimated VAR(10).[21] The figure shows that in the lower range of the business cycle frequencies, investment lags output by roughly one quarter. In a range of higher frequencies, this relationship is reversed, with investment leading output.[22]

4.2 Estimation and Testing Results

Parameter estimates for the model are reported in Table 1. There are two notable features in these results. First, as in the univariate analysis, the data prefer the time-to-plan specification of investment, i.e., $\phi_1 \approx 0$. The other ϕ_i's are also similar across univariate and bivariate analyses. Second, the estimated model assigns a relatively high variance to the government consumption shock. The standard deviation of the innovation to this shock, σ_g, is nearly 3 times larger than the estimate obtained by Christiano and Eichenbaum (1992) using

[21] As noted in the text, $k^*_{21}(\omega)$ is the global maximum of a particular cross-covariance function with respect to lag length. In practice, this maximum could be a discontinuous function of the parameters of our underlying $VAR(10)$ time series model. In our application of the delta function method, we ignore this possibility by computing the derivative of the local maximum with respect to the VAR(10) parameters, about the point estimate, $k^*_{21}(\omega)$.

[22] It is of interest to relate our findings for the phase angle to the finding in the literature, that investment lags output at business cycle frequencies. The result in the literature is based on the cross-covariance function between Hodrick-Prescott (Hodrick and Prescott (1997)) (HP) filtered output and investment data (see, e.g., Christiano and Todd (1996)). As is well known, the HP filter is roughly a high pass filter, allowing the business cycle frequencies and higher to pass through, while zeroing out the lower frequencies (see King and Rebelo (1993)). The result in the literature reflects that not all frequencies receive equal weight in a cross-covariance function based on HP-filtered data. The weight assigned to frequency ω is proportional to the gain, $r(\omega)$, at that frequency (see the Appendix). In our data, $r(\omega)$ takes on its largest value in the lower range of the business cycle frequencies. This is why investment lags the cycle according to correlations based on HP-filtered investment and output data.

the government consumption data. At the same time, the estimated standard deviation of the innovation to the technology shock is one-third smaller than we obtained in the one-shock analysis.

In what follows, we diagnose the estimated model in two ways. First, we compare its implications for the spectrum of the data with the corresponding sample estimates. Second, we apply our cumulative likelihood ratio statistic.

4.2.1 Spectral Properties of the Estimated Model

Figures 3-5, discussed above, can be used to help understand these findings. In addition to reporting features of the spectrum of the data, the figures also report the spectrum of the estimated two-shock model and two perturbations on that model. The first perturbation, labeled *KP model*, sets the investment weights to Kydland and Prescott's values of $\phi_i = 0.25$ for $i = 1, ..., 4$, and leaves the other parameters unchanged from their estimated values. The second perturbation sets the parameters of the exogenous shock processes in the estimated model to the values used by Christiano and Eichenbaum (1992), and sets the remaining parameter values at their estimated values. Since the parameter values of this model are similar to those in the model of Christiano and Todd (1996), we refer to it as the *CT model* in the figure.[23] For convenience, the parameter values associated with our perturbed models are summarized in Table 1.

Figure 3 suggests that the relatively high variance in the government consumption shock in the estimated model helps that model accommodate the steep slope in the spectrum of $\log(I_t/Y_t)$. Still, the estimated model does not go far enough. It undershoots the spectrum at the low frequencies and overshoots at the high frequencies. The miss is more severe in the low frequencies and so this weighs more heavily on our estimation and testing criteria.[24]

To understand better why our estimation procedure assigns a relatively large variance to the government consumption shock, we compare the estimated two-shock model with the *CT model*. In the *CT model*, the coherence between $\log(I_t)$ and $\log(Y_t)$ is close to unity because σ_g is small. The coherence in the estimated model is relatively low because of the high value of σ_g in that

[23] The parameters in the *CT* model are somewhat different from those in the model Christiano and Todd (1996) in that the *CT* model incorporates the hump-shape in ϕ_2, ϕ_3, ϕ_4 implied by the estimated model. *CT* is constructed in order to evaluate the role played by the high government shock variance in the estimated model.
[24] Figure 3 may appear to indicate the opposite, that the miss is larger in the higher frequencies. However, recall that Figure 3 displays the log of the spectrum of $\log(I_t/Y_t)$, not its level. It is the level that enters the estimation criterion.

model.[25] Although the coherence function appears to have played an important role in driving the estimated values of the exogenous shock variances, Figure 4 shows that the model nevertheless has difficulty matching the steep slope of that function. In particular, the model has difficulty accommodating the high coherence in the lower end of the business cycle frequencies (see Figure 4). These results suggest that (i) a good model must have more than one shock, so as to be able to match the low coherence function outside the business cycle and (ii) the coherence in the business cycle should nevertheless be high. Our model has difficulty satisfying (i) and (ii) simultaneously.[26]

To understand why the estimation procedure selects the time-to-plan specification of investment, we compare the KP and estimated two-shock model's ability to reproduce the spectral properties of the data, as captured by the VAR(10). Time-to-plan appears to play two roles in the estimated model. First, as in the univariate analysis, it helps accommodate the negative slope of the spectrum of $\log(Y_t/Y_{t-1})$. Second, time-to-plan helps the model to reproduce the evidence that investment lags output in the business cycle frequencies. This can be seen in the positive phase angle displayed in business cycle frequencies (see Figure 5). Still, the model does not go far enough: its implied phase angle is smaller than that of the VAR(10).[27]

In view of the estimated model's difficulties accommodating the phase angle

[25] Raising σ_g is particularly effective in reducing coherence for a second reason. In the model, government spending shocks generate a negative correlation between investment and output, while the technology shock generates a positive correlation. For a further discussion of this property of the model, see Christiano and Todd (1996).

[26] The coherence function suggests a simple measurement error model which we also explored in results not reported in the paper. We modified the CT model by treating observations on log output and log investment as the sum of the true data and orthogonal measurement error. We then estimated the variances of the measurement error process by maximum likelihood. As expected, the resulting model reproduces the steep slope in the estimated coherence function. However, the measurement error also causes the model to overstate the high frequency component of the spectrum of $\log(I_t/Y_t)$. This model is rejected with a likelihood ratio statistic in the neighborhood of 1500.

[27] There is one dimension in which time-to-plan appears to *hurt* model fit. According to Figure 3, time-to-plan reduces the ability of the model to accommodate the steep slope in the spectrum of $\log(I_t/Y_t)$. By comparison with KP, the estimated model overstates that spectrum in the higher frequencies. Evidently, this consideration does not play an important role in estimation. Presumably this is because, according to (1), the weight assigned in the estimation criterion to the spectrum of $\ln(I_t/Y_t)$ is proportional to the *level* of the corresponding empirical estimate. (See the periodogram, $I(\omega_j)$, in (1) and (9).) The latter is extremely small in the higher frequencies by comparison to what it is in the lower frequencies (note that Figure 3 reports the log of the spectrum).

and the steep slopes of both the spectrum of $\log(I_t/Y_t)$ and the coherence function, it is not surprising that the model is rejected. To see this, note from Figure 6 that the likelihood ratio statistic for the estimated model is around 110 (see 'Actual Cumulative Likelihood Ratio'). Under the null hypothesis that the model is true, this is a chi-square statistic with 39 degrees of freedom. This easily exceeds conventional critical values. [28]

4.2.2 Cumulative Likelihood Ratio Statistic

We can use our frequency decomposition of the likelihood ratio statistic to identify which frequencies are responsible for this rejection. Notice that the cumulative likelihood ratio rises sharply in the frequency range, $(0, 0.6)$. In a sense, the poor fit in these frequencies is the reason for the rejection. To see this, consider the Adjusted Cumulative Likelihood Ratio displayed in Figure 6. It shows what the cumulative likelihood ratio would have been if the fit in frequencies $(0, 0.6)$ had been similar to the average fit in the higher frequencies. [29] These calculations lead to the conclusion that the likelihood ratio statistic, λ, would have been 55, with a probability value of 4.6. That is, we would not have rejected the model at any significance level less than 4.6 percent.

To gain further insight into the reason the estimated model does poorly in the low frequencies, we evaluated the model at the estimated parameter values using the cumulative likelihood ratio computed using the univariate density for $\ln(I_t/Y_t)$ alone. This also displays a sharp rise in the frequencies, $(0, 0.6)$. We infer from this that a part of the reason for the poor fit of the model lies in the difficulty it has in matching the steep slope in the spectrum of $\log(I_t/Y_t)$ near frequency zero. Previously, we discussed the model's difficulties in ac-

[28] This is how we arrived at this calculation of the number of degrees of freedom. In section 2.3, we argued that the two-shock model is a restricted VARMA (5,8) model, which we approximate with a VAR (10). The unrestricted VARMA (5,8) model has 10=5×2 autoregressive and 32=8×4 moving average parameters. In addition, there are 3 parameters governing the variance-covariance of the shocks, for a total of 45 parameters. Subtract from this the 6 estimated parameters, σ_z, σ_g, ρ, ϕ_1, ϕ_2, ϕ_3. This brings the total number of degrees of freedom to 39.

[29] The adjusted cumulative likelihood ratio is computed as follows. We fit a regression line, $a + b\omega$, through the cumulative likelihood ratio function over the range of frequencies, $\omega \in (.6, \pi)$. The adjusted cumulative likelihood ratio is:

$$\tilde{\Lambda}(\omega) = \begin{cases} \Lambda(\omega) - a & \omega \in (.6, \pi) \\ b\omega & \omega \in (0, .6) \end{cases},$$

where $\Lambda(\omega)$ is the actual cumulative likelihood ratio, defined in (11).

counting for features of the cross-spectrum between investment and output at low frequencies. Presumably, these are also part of the reason for the rejection.

5 Conclusion

We applied frequency domain tools to diagnose parameter estimates and goodness of fit tests for maximum likelihood estimation of a class of real business cycle models. The principal methodological aspects of the analysis were summarized in the introduction. There are two main substantive findings. First, the results confirm other findings that suggest time-to-plan in the investment technology has a potentially useful role to play in dynamic models. Second, alternatives to government spending disturbances need to be explored in the quest for an empirically plausible business cycle model.

Although we have limited our analysis to one or two variables, we emphasize that this does not reflect an inherent limitation of the methods used. We showed how the tools apply in a bivariate setting and hopefully from this it is obvious how they can be extended to higher dimensions. To analyze a larger list of variables in the model would of course require adding more shocks. However, the literature offers plenty of candidates for these. In addition to government consumption and technology shocks, one can consider various kinds of preference shocks and also monetary shocks. In addition, there are various types of measurement error that can be incorporated into the analysis. [30] Our decision to limit the number of variables in the analysis reflected our desire to make the methodology as transparent as possible.

A Coherence, Gain and Phase in Spectral Analysis

This appendix briefly describes the computation and interpretation of the coherence function and phase angle analyzed in the text.

We begin with the coherence. Consider the projection of y_{2t} onto y_{1t-j} for $j \in (-\infty, \infty)$:

$$y_{2t} = \sum_{j=-\infty}^{\infty} h_j y_{1t-j} + \varepsilon_t, E\varepsilon_t y_{1t-k} = 0 \text{ for all } k. \tag{A.1}$$

[30] For example, Altug (1989) assumes that the data received by the econometrician contain measurement error. Christiano (1988) and Sargent (1989) assume that the data observed by agents contain measurement error.

We assume that $y_t = (y_{1t}, y_{2t})'$ has mean zero and is covariance stationary. It is well known (see, e.g., Sargent (1987)) that $h(\omega) \equiv \sum_{j=-\infty}^{\infty} h_j e^{-i\omega j}$ satisfies

$$\begin{aligned} F_{21}(\omega) &= h(\omega) F_{11}(\omega), \qquad\qquad \text{(A.2)} \\ F_{22}(\omega) &\equiv \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} C_{22}(k) e^{-i\omega k} = |h(\omega)|^2 F_{11}(\omega) + S(\omega), \end{aligned}$$

where $h(\omega) \equiv 0$ when $F_{21}(\omega) = F_{11}(\omega) = 0$. Also, $C_{ij}(k) = E y_{it} y_{j,t-k}$ and $S(\omega)$ is the spectral density of ε_t. A measure of the information in y_{1t} about y_{2t} is given by the R^2 of the projection, (A.1):

$$R_{21}^2 = \frac{Var(\sum_{j=-\infty}^{\infty} h_j y_{1t-j})}{Var(y_{2t})} = \frac{\int_0^{2\pi} |h(\omega)|^2 F_{11}(\omega) d\omega}{\int_0^{2\pi} F_{22}(\omega) d\omega}, \tag{A.3}$$

where the subscripts on R^2 indicate the left-hand and right-hand variables in the projection.

We can obtain an analogous concept of R^2 in the frequency domain. Suppose x_t is a vector stationary stochastic process with spectral density $f(\omega)$, $\omega \in (0, 2\pi)$. We define the *component of* x_t *at frequency* ω^* as the result of filtering x_t with a band-pass filter which passes power in an arbitrarily small window around frequency ω^*, and no power at other frequencies. [31] The spectral density of the component of x_t at frequency ω^* is $f^*(\omega)$, $\omega \in (0, 2\pi)$, where $f^*(\omega) = 0$ for ω outside an interval about ω^* and $f^*(\omega) = f(\omega)$ in that interval.

Let $R^2(\omega)$ be the R^2 of the regression of the frequency ω component of y_{2t} on the frequency ω component of $y_{1,t-j}$, for $j \in (-\infty, \infty)$. Then, substituting for $h(\omega)$ in (A.3) using (A.2), we obtain the coherence, (12). So, we see that the coherence between y_{1t} and y_{2t} is a measure of the information about y_{2t} in linear combinations of future and past y_{1t} when we consider only the frequency ω components of these variables. [32] We omit subscripts on $R^2(\omega)$ because which variable is on the left, and which is on the right, does not matter when the projection underlying the R^2 occurs in the frequency domain (see (12)).

We now discuss the phase angle between y_{1t} and y_{2t}. Consider the object, $\theta_{21}^*(\omega)$, in (13). To interpret this, it is useful to express $C_{21}(k) \equiv E y_{2t} y_{1,t-k}$ as follows:

[31] For a discussion of the band-pass filter, see Sargent (1987, page 259).

[32] Our interpretation of the coherence function bears a similarity to the analysis in Engle (1974). He asks whether the linear regression relation between two variables (money growth and inflation) shifts across different frequency bands. Our analysis differs from Engle's (1974) in that we (i) focus on the R^2 of the linear relationship and (ii) examine the relationship between variables at a single frequency.

$$
\begin{aligned}
C_{21}(k) &= \int_0^{2\pi} F_{21}(\omega) e^{i\omega k} d\omega \\
&= \int_0^{\pi} \left[F_{21}(\omega) e^{i\omega k} + F_{21}(-\omega) e^{-i\omega k} \right] d\omega \\
&= \int_0^{\pi} r(\omega) \left[e^{i\left(\theta_{21}^*(\omega)+\omega k\right)} + e^{-i\left(\theta_{21}^*(\omega)+\omega k\right)} \right] d\omega \\
&= \int_0^{\pi} r(\omega) 2\cos\left[\theta_{21}^*(\omega) + \omega k\right] d\omega \qquad \text{(A.4)} \\
&= \int_0^{\pi} C_{21}(k;\omega) d\omega,
\end{aligned}
$$

say. Here, $C_{21}(k;\omega)$ is the covariance between the component of y_{2t} at frequency ω and the component of y_{1t-k} at frequency ω. Thus, the cross-covariance function of two time series is the integral of the cross covariances between their frequency ω components, for $\omega \in (0,\pi)$. Note that the gain function, $r(\omega)$, determines how important any particular frequency is in the covariance between two variables.

It is common to characterize the lead-lag relationship between two variables by the value of k for which $C_{21}(k)$ is the largest. For example, if this happens for a value of, say, $k = -2$, then it is said that 'y_{2t} leads y_{1t} by two periods'. The analogous statements can be made in frequency domain using $C_{21}(k;\omega)$. Note that $C_{21}(k;\omega)$ is maximized for $k = -\theta_{21}^*(\omega)/\omega$. Thus, if $\theta_{21}^*(\omega) > 0$ then we say 'the frequency ω component of y_{2t} leads the frequency ω component of y_{1t} by $\theta_{21}^*(\omega)/\omega$ periods'. Similarly, if $\theta_{21}^*(\omega) < 0$, then we say that 'the frequency ω component of y_{1t} leads the frequency ω component of y_{2t} by $-\theta_{21}^*(\omega)/\omega$ periods'.

There is an ambiguity in characterizing the lead-lag relationship between variables using $C_{21}(k;\omega)$ that is not present when we do so using $C_{21}(k)$. This is because the component of a variable at frequency ω is a pure cosine wave. The length of the lead or lag between two sinusoidal functions with the same period is ill-defined. This is manifested in the observation that there are many $\theta_{21}^*(\omega)$ that solve (13): if $\theta_{21}^*(\omega)$ solves (13) then so does $\theta_{21}^*(\omega) + 2\pi l$, for $l = 0, \pm 1, \pm 2, \ldots$.

Here is an intuitive way to see this. With two series of period 11, the statement that the first leads the second by 2 periods is equivalent to the statement that the second leads the first by 8 periods. These two statements are also equivalent to the notion that the first leads the second by 12 periods, and so on.

In our analysis, we adopt the following resolution to this ambiguity. Let

$$\begin{aligned} k_{21}(\omega;\Delta) &= \arg\max_{k} \int_{\omega-\Delta}^{\omega+\Delta} C_{21}(k;\omega)d\omega, \pi-\Delta \geq \omega \geq \Delta > 0, \\ \theta_{21}(\omega;\Delta) &= -\omega k_{21}(\omega;\Delta). \end{aligned} \tag{A.5}$$

Define

$$\begin{aligned} k_{21}^{**}(\omega) &= \lim_{\Delta\to 0} k_{21}(\omega;\Delta), \\ \theta_{21}^{**}(\omega) &= \lim_{\Delta\to 0} \theta_{21}(\omega;\Delta), \end{aligned}$$

for $\pi \geq \omega \geq 0$. Our numerical experiments suggest that $k_{21}(\omega;\Delta)$ and $\theta_{21}(\omega;\Delta)$ are single-valued functions of ω for all $\omega \in (0,\pi)$ and $\Delta > 0$, except possibly at isolated values of ω. When $\Delta = 0$, then $\theta_{21}(\omega;\Delta)$ is composed of the countable set of elements discussed above. For all but a finite set of values of ω, $\theta_{21}^{**}(\omega)$ selects one element of the set, $\theta_{21}(\omega;0)$. Similarly for $k_{21}^{**}(\omega)$. In our analysis, we identify $\theta_{21}^{*}(\omega)$ and $k_{21}^{*}(\omega)$ with $\theta_{21}^{**}(\omega)$ and $k_{21}^{**}(\omega)$, respectively. The object, $\theta_{21}^{*}(\omega)$, is our measure of the *phase angle* between the frequency ω components of y_{1t} and y_{2t}. The object, $k_{21}^{*}(\omega)$, measures this in units of time.

To gain insight into our measure of the phase angle, consider Figure 7. There, y_{1t} corresponds to $\ln(Y_t)$ and y_{2t} corresponds to $\ln(I_t)$, as implied by the estimated $VAR(10)$ for these variables, which was discussed in the text.[33] Figure 7 exhibits a subset of the elements in $\theta(\omega;0)$ for $\omega \in (0,\pi)$ (see the solid lines). The circles indicate $\theta_{21}(\omega;\Delta)$ for $\Delta = 0.3$. Note how these θ's lie close to one of the solid lines. Note too, how that phase angle function exhibits discontinuities. We found that at the points of discontinuity, there are two elements in $\theta_{21}(\omega;0.3)$. The stars indicate $\theta_{21}(\omega;\Delta)$ for $\Delta = 0.06$. Note how in each case, $\theta_{21}(\omega;\Delta)$ is now closer to one of the solid lines. This is consistent with the notion that, for each ω, $\theta_{21}(\omega;\Delta)$ converges to one of the solid lines as $\Delta \to 0$.

The dashed lines in Figure 7 indicate $+\pi$ and $-\pi$. We have included these in order to facilitate comparison to the standard method for selecting an element of $\theta_{21}(\omega;0)$, which picks $\theta \in (-\pi,\pi)$.[34] It is evident from the figure that, for our estimated $VAR(10)$, this method produces a phase angle function with a single point of discontinuity just above $\omega = 2$. In effect, the method chooses the lead-lag relationship between two variables as the smallest (in absolute value) value of k which attains the maximum for the cross-covariance function. We

[33] Although $\log(Y_t)$ and $\log(I_t)$ are not covariance stationary, they are so after application of a band pass filter which excludes frequency zero. As a result, phase angle and coherence measures are well-defined for $\omega \neq 0$.
[34] See, e.g., Granger and Newbold (1977) or Sargent (1987).

think ours is a more natural measure in our context, given its connection to time domain lead-lag measures used in business cycle analysis.

To obtain insight into the discontinuities in our estimate of $\theta_{21}^{*}(\omega)$ (and, hence, $k_{21}^{*}(\omega)$), consider Figure 8. That figure reports

$$\int\limits_{\omega-\Delta}^{\omega+\Delta} C_{21}(k;\tilde{\omega})d\tilde{\omega}, \Delta = 0.15,$$

for $\omega \in (1, \pi - 0.15)$ and $k \in (-15, 10)$. Note how this function oscillates with k, for each fixed ω, producing a pattern of ridges and valleys in the three dimensional surface. The dark lines indicate $k_{21}(\omega; 0.15)$. The discontinuities in $k_{21}(\omega; \Delta)$ reflect that the ridge which achieves the greatest height varies with ω, and that ridges are separated by valleys. As $\Delta \to 0$, the surface depicted in Figure 8 evolves so that for a given ω, each ridge has the same height. However, for $\Delta > 0$ we found that (except for isolated ω's) exactly one ridge achieved a global maximum for a given ω.

References

[1] Altug, Sumru, 1989. Time to Build and Aggregate Fluctuations Some New Evidence, International Economic Review, 30(4), November, 889-920.

[2] Baxter, Marianne, and Robert King, 1999. Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series, Review of Economics and Statistics, 81(4), November, pages 575-93.

[3] Bencivenga, Valerie R., 1992. An Econometric Study of Hours and Output Variation with Preference Shocks' International Economic Review 33 449-71

[4] Berkowitz, Jeremy, 2001. Generalized Spectral Estimation of the Consumption-Based Asset Pricing Model, Journal of Econometrics (104)2 269-288

[5] Bernanke, Ben, Mark Gertler and Simon Gilchrist, 1999. The Financial Accelerator in a Quantitative Business Cycle Framework, in John Taylor and Michael Woodford, editors, Handbook of Macroeconomics, North-Holland.

[6] Christiano, Lawrence J., 1985. A Method for Estimating the Timing Interval in a Linear Econometric Model, With an Application to Taylor's Model of Staggered Contracts, Journal of Economic Dynamics and Control, 9, 363-404.

[7] Christiano, Lawrence J., 1988. Why Does Inventory Investment Fluctuate So Much?, Journal of Monetary Economics, 21(2/3), March/May, 247-80.

[8] Christiano, Lawrence J., 2001. Solving Dynamic Equilibrium Models by a Method of Undetermined Coefficients, forthcoming. See also, NBER Technical Working paper 225.
[9] Christiano, Lawrence J. and Martin Eichenbaum, 1987. Temporal Aggregation and Structural Inference in Macroeconomics, Carnegie Rochester Conference Series on Public Policy, 26,63-130.
[10] Christiano, Lawrence J., Martin Eichenbaum, and David Marshall, 1991. The Permanent Income Hypothesis Revisited, Econometrica, 59(2), March, 397-423..
[11] Christiano, Lawrence J. and Martin Eichenbaum, 1992. Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations, American Economic Review, 82(3), June, 430-50.
[12] Christiano, Lawrence J. and Terry Fitzgerald, 1999. The Band Pass Filter, National Bureau of Economic Research Working Paper number 7257, and forthcoming, International Economic Review.
[13] Christiano, Lawrence J. and Richard M. Todd, 1996. Time to Plan and Aggregate Fluctuations, Federal Reserve Bank of Minneapolis Quarterly Review, Winter, 14-27.
[14] Christiano, Lawrence J. and Richard M. Todd, 2001. The Conventional Treatment of Seasonality in Business Cycle Analysis: Does it Create Distortions?, Forthcoming, Journal of Monetary Economics.
[15] Christiano, Lawrence J. and Robert J. Vigfusson, 2001. Maximum Likelihood in the Frequency Domain: The Importance of Time-to-Plan, Federal Reserve Bank of Cleveland, Working Paper 0106.
[16] Cogley, Timothy, and James Nason, 1995. Output Dynamics in Real Business Cycle Models, American Economic Review, 85(3), June, 492-511.
[17] Cogley, Timothy, 2001a. Estimating and Testing Dynamic Rational Expectations Models, Journal of Economic Dynamics and Control, 25(10), 1485-1525
[18] Cogley, Timothy, 2001b. A Frequency Decomposition of Approximation Errors in Stochastic Discount Factor Models, manuscript, International Economic Review, 42(2), 473-503
[19] Diebold, Francis X., Lee E. Ohanian, and Jeremy Berkowitz, 1998. Dynamic Equilibrium Economies: A Framework for Comparing Models and Data, Review of Economic Studies, 65(3), July, 433-51.
[20] Engle, Robert F., 1974. Band Spectrum Regression, International Economic Review, vol. 15, 1-11.
[21] Granger, Clive W. J., and Paul Newbold, 1977. Forecasting Economic Time Series, Academic Press, New York.
[22] Hall, George J., 1996. Overtime, Effort, and the Propagation of Business Cycle Shocks, Journal of Monetary Economics, 38 (1), August, 139-60.
[23] Hamilton, James D. 1994. Time Series Analysis, Princeton University Press
[24] Hansen, Lars Peter, and Thomas J. Sargent, 1980. Formulating and Estimating Dynamic Linear Rational Expectations Models, Journal of Eco-

nomic Dynamics and Control, 2, February, 7-46.

[25] Hansen, Lars Peter, and Thomas J. Sargent, 1980a. Methods for Estimating Continuous Time Rational Expectations Models from Discrete Data, Research Department Staff report 59, Federal Reserve Bank of Minneapolis.

[26] Hansen, Lars Peter, and Thomas J. Sargent, 1991. Exact Linear Rational Expectations Models: Specification and Estimation, in Underground Classics in Economics, Boulder and Oxford: Westview Press, pages 45-76.

[27] Hansen, Lars Peter, and Thomas J. Sargent, 1993. Seasonality and Approximation Errors in Rational Expectations Models, Journal of Econometrics, 55(1-2), January-February, 21-55.

[28] Harvey, Andrew, 1989. Forecasting Structural Time Series Models and the Kalman Filter, Cambridge University Press.

[29] Hodrick, Robert, and Edward Prescott, 1997. Post-War Business Cycles: An Empirical Investigation, Journal of Money, Credit and Banking, vol. 29, no. 1, February, pp. 1-16.

[30] Ireland, Peter N., 1997. A Small, Structural Quarterly Model for Monetary Policy Evaluation, Carnegie Rochester Conference Series on Public Policy, 47 (0), December, 83-108.

[31] Ireland, Peter N., 1999. A Method for Taking Models to the Data, Federal Reserve Bank of Cleveland Working Paper 9903

[32] Ireland, Peter N., 2001. Technology shocks and the business cycle: On empirical investigation, Journal Of Economic Dynamics And Control, (25)5 703-19

[33] Kim, Jinill, 2000. Constructing and estimating a realistic optimizing model of monetary policy, Journal Of Monetary Economics, (45)2 pp. 329-359

[34] King, Robert, 1995. Quantitative Theory and Econometrics, Federal Reserve Bank of Richmond Economic Quarterly, 81(3), 53-105.

[35] King, Robert, and Sergio Rebelo, 1993. Low Frequency Filtering and Real Business Cycles, Journal of Economic Dynamics and Control, vol. 17, pp. 251-231.

[36] Kydland, Finn E. and Edward C. Prescott, 1982. Time to Build and Aggregate Fluctuations, Econometrica, 50, November, 1345-70.

[37] Leeper, Eric M., and Christopher A. Sims, 1994. Toward a Modern Macroeconomic Model Usable for Policy Analysis, Fischer, Stanley, Julio J. Rotemberg, eds. NBER Macroeconomics Annual Cambridge and London: MIT Press, pages 81-118.

[38] Long, John, and Charles Plosser, 1983. Real Business Cycles, Journal of Political Economy, 91(1), February, 39-69.

[39] McGrattan, Ellen R., 1994. The Macroeconomic Effects of Distortionary Taxation, Journal of Monetary Economics, 33, 573-601

[40] McGrattan, Ellen R., Richard Rogerson, and Randall Wright, 1997. An Equilibrium Model of the Business Cycle with Household Production and Fiscal Policy, International Economic Review, 38 (2), May, 267-90.

[41] Nerlove, Marc, David M. Grether and Jose M. Carvalho, 1979. Analysis of Economic Time Series: A Synthesis, Academic Press.

[42] Rouwenhorst, K. Geert, 1991. Time to Build and Aggregate Fluctuations A Reconsideration, Journal of Monetary Economics, 27, 241-54.

[43] Sargent, Thomas 1987. Macroeconomic Theory' Academic Press

[44] Sargent, Thomas 1989. Two Models of Measurements and the Investment Accelerator, Journal of Political Economy, 97(2), April, pages 251-87.

[45] Sims, Christopher, 1972. Approximate Prior Restrictions in Distributed Lag Estimation, Journal of the American Statistical Association, 67(337), 169-175.

[46] Watson, Mark, 1993. Measures of Fit for Calibrated Models, Journal of Political Economy, 101(6), 1011-1041.

B Table

Table 1: Parameter Estimates, Time-to-Build Model

	Panel A: One Shock Model					
	ϕ_1	ϕ_2	ϕ_3	ρ	σ_z	σ_g
Estimated	0.0097	0.28	0.48	NA	0.014	NA
Christiano-Todd	0.01	0.33	0.33	NA	0.018	NA
Kydland-Prescott	0.25	0.25	0.25	NA	0.012	NA
RBC	NA	NA	NA	NA	0.011	NA
	Panel B: Two Shock Model					
	ϕ_1	ϕ_2	ϕ_3	ρ	σ_z	σ_g
Estimated	0.0079	0.30	0.41	0.94	0.012	0.063
CT model	0.0079	0.30	0.41	0.96	0.018	0.022
KP model	0.25	0.25	0.25	0.94	0.012	0.063

Figure 1: The Log Spectrum of Output Growth

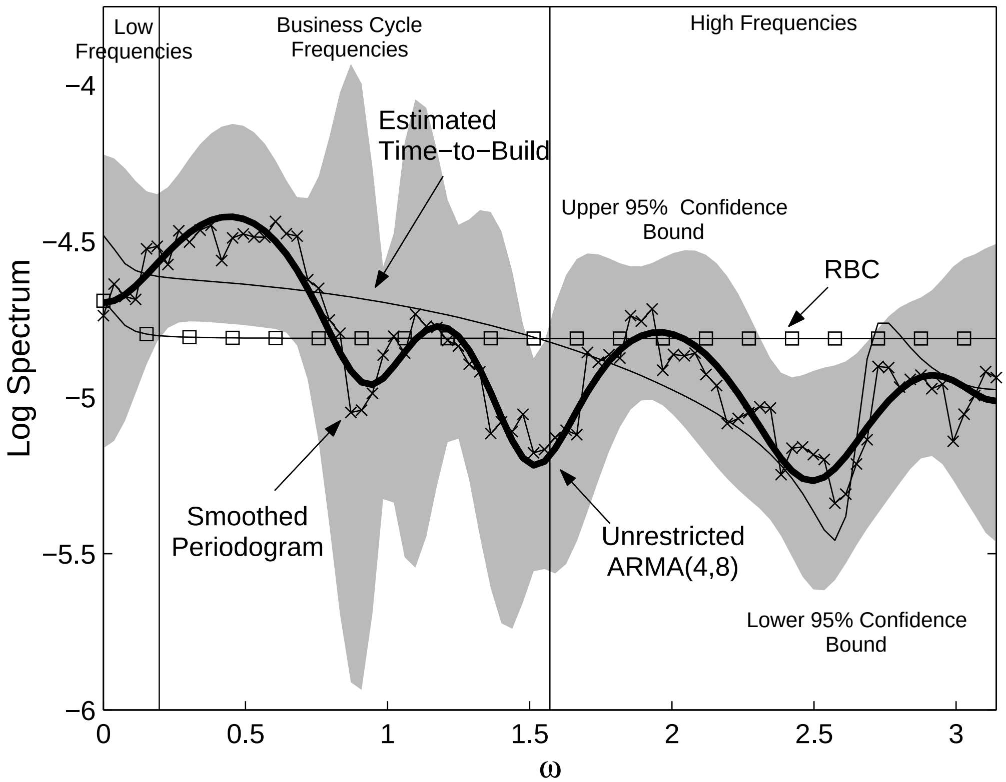

Notes: (i) Various Estimates of Spectrum of Output Growth as Indicated.
(ii) Smoothed Periodogram $\sim$ Centered Moving Average of Periodogram
(iii) Estimated Time-to-Build $\sim$ one-shock time-to-build model with $\{\sigma_z, \phi_1, \phi_2, \phi_3\}$ estimated.
(iv) RBC $\sim$ one-shock real business cycle model with σ_z estimated.

Figure 2: Cumulative Likelihood Ratio, One Shock Models

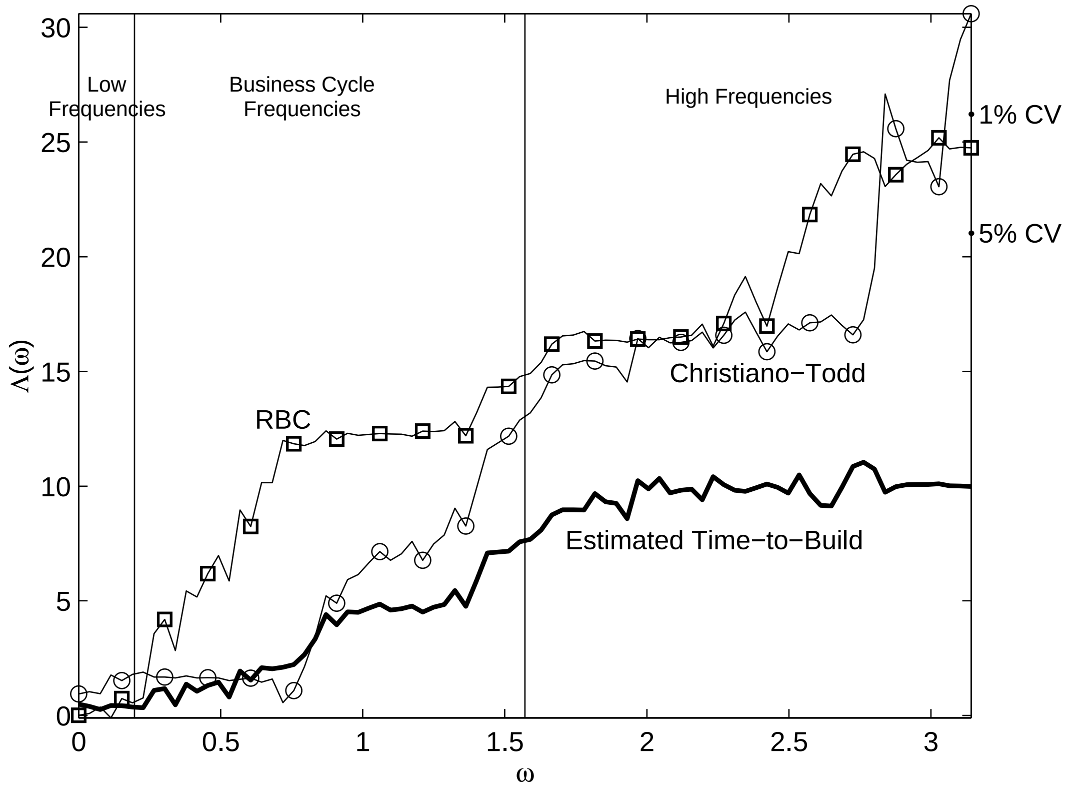

Notes: (i) Figure displays Λ function in Equation (11) for each of the three indicated models.
(ii) Christiano-Todd $\sim$ time-to-build model with $\phi_1 = 0.01$ and $\phi_2 = \phi_3 = \phi_4 = 0.33$ and σ_z estimated.
(iii) Estimated Time-to-Build $\sim$ time-to-build model with $\{\sigma_z, \phi_1, \phi_2, \phi_3\}$ estimated.
(iv) RBC $\sim$ real business cycle model with σ_z estimated.
(v) CV $\sim$ critical values for Chi-square distribution with 12 degrees of freedom.

Figure 3: The Log Spectrum of $\ln(I_t/Y_t)$

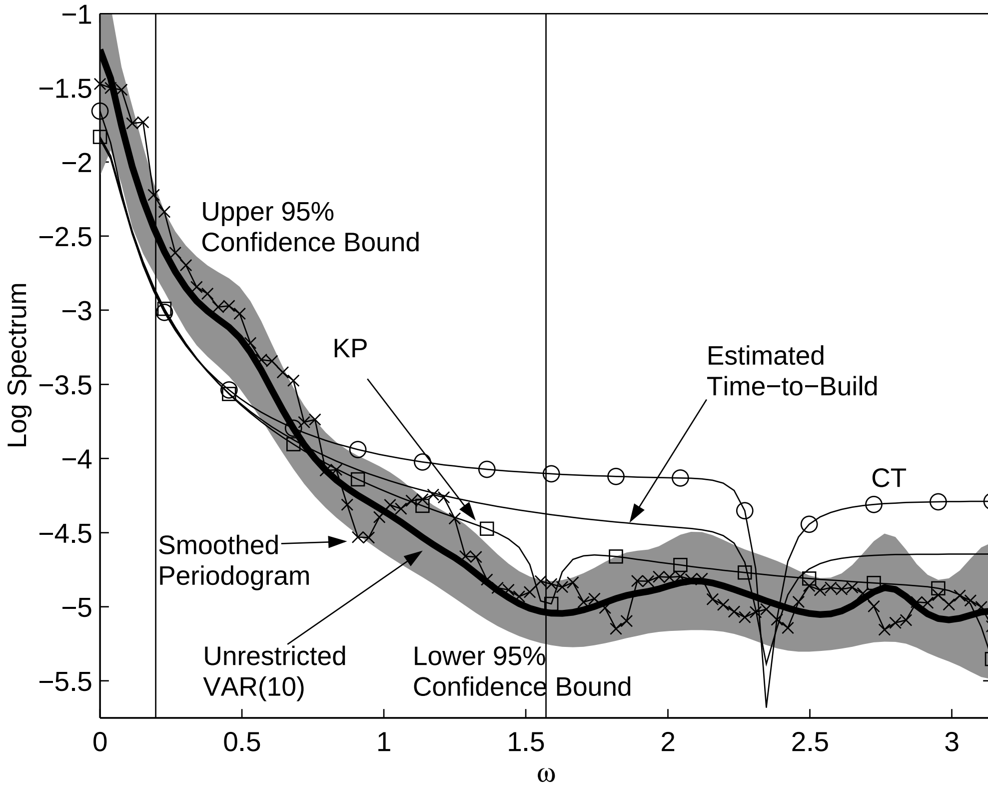

Notes: (i) Various Estimates of Spectrum as Indicated.

(ii) Smoothed Periodogram $\sim$ Centered Moving Average of Periodogram

(iii) Estimated Time-to-Build $\sim$ two-shock time-to-build model with$\{\sigma_z, \rho, \sigma_g, \phi_1, \phi_2, \phi_3\}$ estimated.

(iv) CT $\sim$ two-shock model time-to-build model with $\{\phi_1, \phi_2, \phi_3\}$ set equal to estimated values and $\{\sigma_z, \rho, \sigma_g\}$ set equal to calibrated values.

(v) KP $\sim$ two-shock model time-to-build model with $\phi_1 = \phi_2 = \phi_3 = \phi_4 = 0.25$ and $\{\sigma_z, \rho, \sigma_g\}$ set equal to estimated values.

Figure 4: The Coherence Between $\ln(Y_t)$ and $\ln(I_t)$

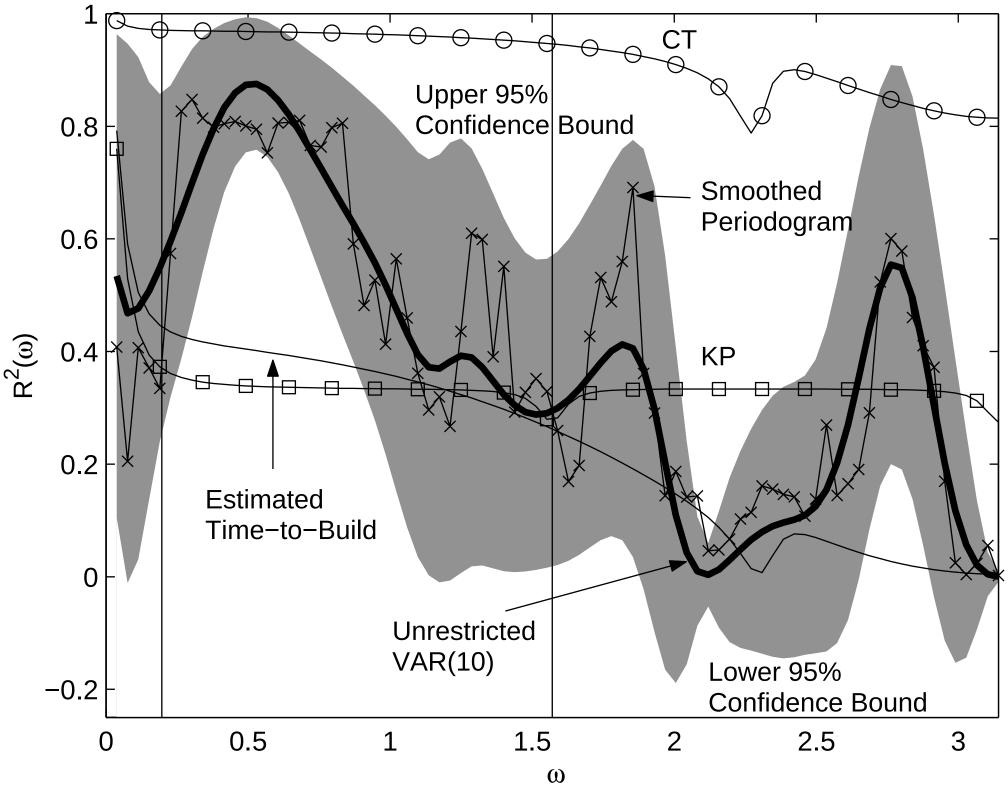

Notes: (i) Various Estimates of Coherence as Indicated.

(ii) Smoothed Periodogram $\sim$ Centered Moving Average of Periodogram.

(iii) Estimated Time-to-Build $\sim$ two-shock time-to-build model with$\{\sigma_z, \rho, \sigma_g, \phi_1, \phi_2, \phi_3\}$ estimated.

(iv) CT $\sim$ two-shock model time-to-build model with $\{\phi_1, \phi_2, \phi_3\}$ set equal to estimated values and $\{\sigma_z, \rho, \sigma_g\}$ set equal to calibrated values.

(v) KP $\sim$ two-shock model time-to-build model with $\phi_1 = \phi_2 = \phi_3 = \phi_4 = 0.25$ and $\{\sigma_z, \rho, \sigma_g\}$ set equal to estimated values.

Figure 5: The Phase Relationship in Units of Time

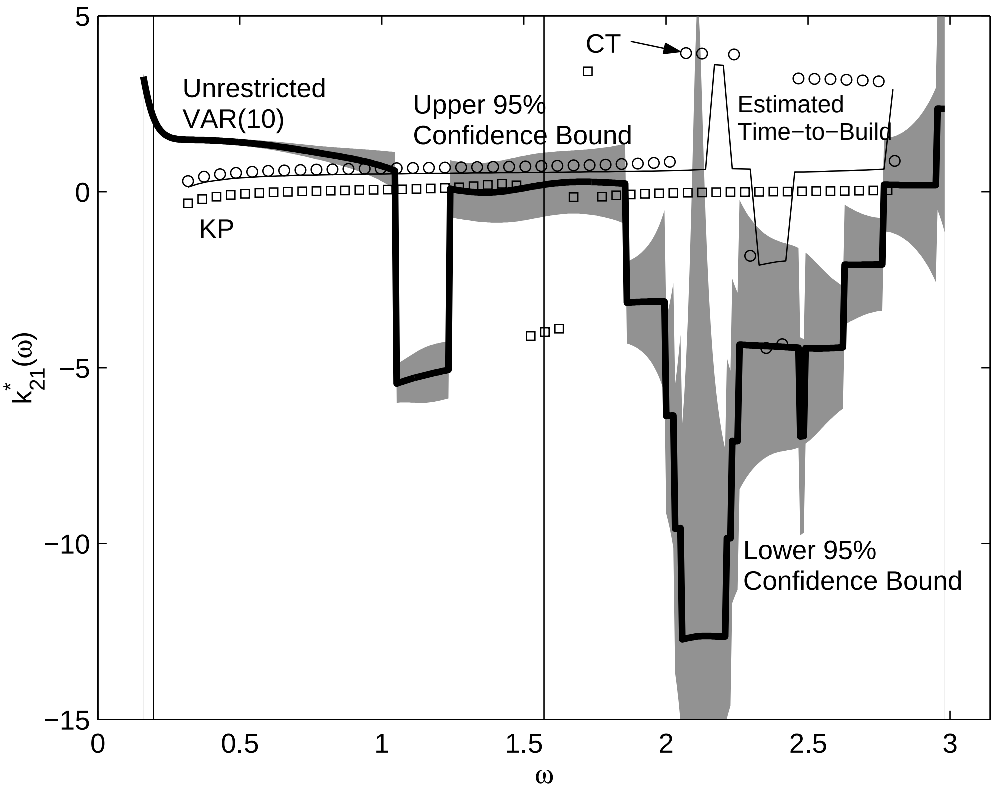

Notes: (i) Various Estimates of The Phase Relationship as Indicated.
(ii) Estimated Time-to-Build $\sim$ two-shock time-to-build model with$\{\sigma_z, \rho, \sigma_g, \phi_1, \phi_2, \phi_3\}$ estimated.
(iii) CT $\sim$ two-shock model time-to-build model with $\{\phi_1, \phi_2, \phi_3\}$ set equal to estimated values and $\{\sigma_z, \rho, \sigma_g\}$ set equal to calibrated values.
(iv) KP $\sim$ two-shock model time-to-build model with $\phi_1 = \phi_2 = \phi_3 = \phi_4 = 0.25$ and $\{\sigma_z, \rho, \sigma_g\}$ set equal to estimated values.

Figure 6: Cumulative Likelihood Ratio, Two-Shock Models

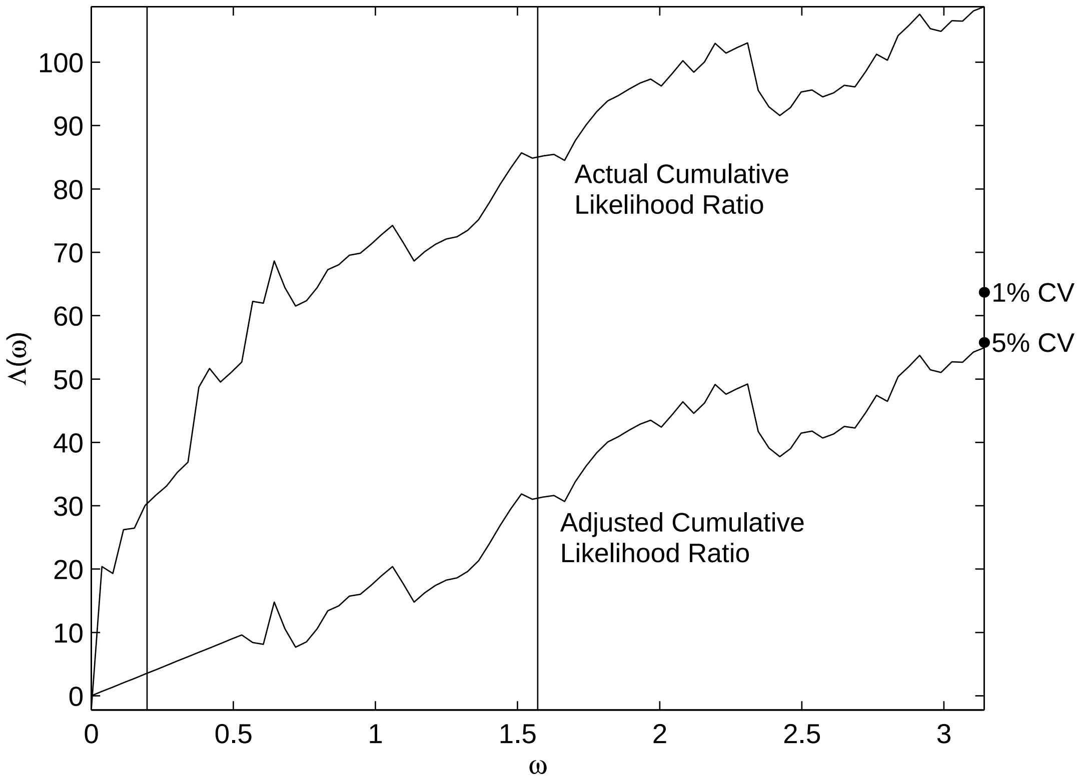

Notes (i) Actual Cumulative Likelihood Ratio $\sim$ Λ function in Equation (11) evaluated at the indicated frequencies.

(ii) Adjusted Cumulative Likelihood Ratio $\sim$ See text.

(iii) CV $\sim$ critical values for Chi-square distribution with 39 degrees of freedom.

Figure 7: Measures of the Phase Angle, As a Function of Frequency

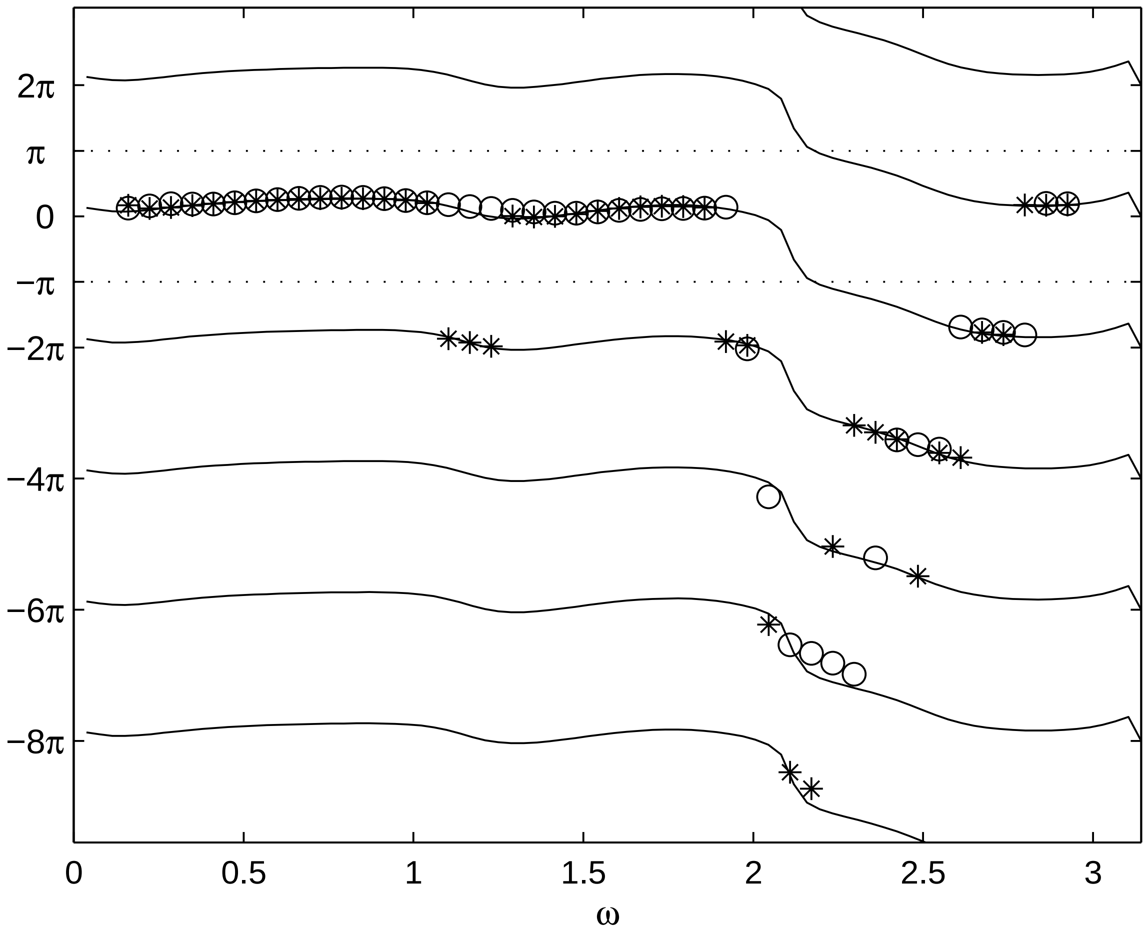

Notes: (i) Solid Lines ∼ $\theta_{21}(\omega; 0)$.
(ii) Stars ∼ $\theta_{21}(\omega; 0.06)$.
(iii) Circles ∼ $\theta_{21}(\omega; 0.3)$.
(iv) For further discussion, see Appendix A.

Figure 8: Covariance Function for the Data

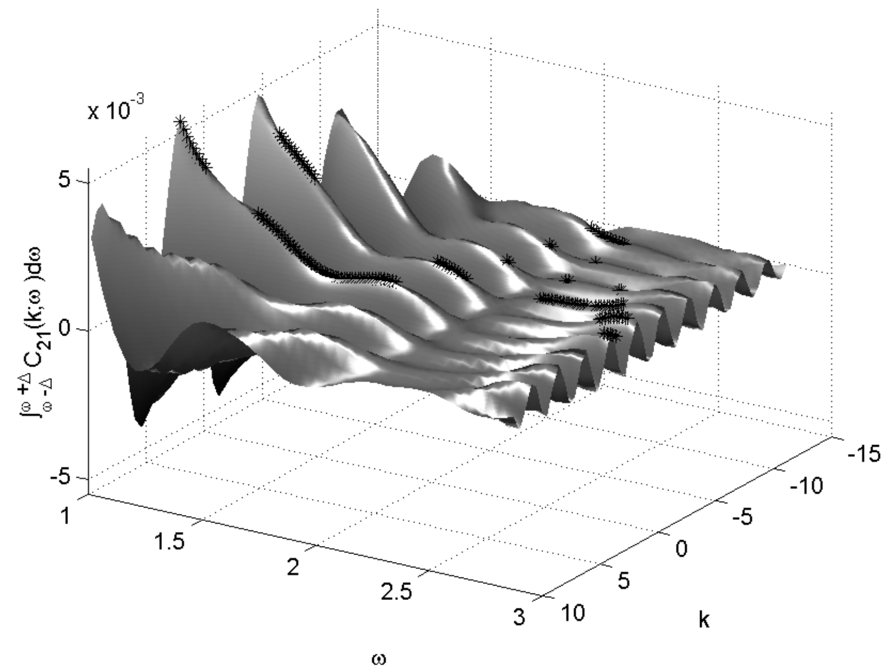

Notes: (i) $C_{21}(k;\omega) = Ey_{2,t}y_{1,t-k}$, where y_{1t} is frequency ω component of $\ln(Y_t)$ and y_{2t} is frequency ω component of $\ln(I_t)$.

(ii) Stars $\sim$ $k_{21}(\omega; 0.15)$.

(iii) For further discussion, see Appendix A.