Sequential Action and Beliefs under Partially Observable DSGE Environments*

Seong-Hoon Kim†

December 2011

Abstract

This paper introduces a classification of DSGEs from a Markovian perspective, and positions the class of POMDP (Partially Observable Markov Decision Process) to the center of a generalization of linear rational expectations models. The analysis of the POMDP class builds on the previous development in dynamic controls for linear system, and derives a solution algorithm by formulating the equilibrium as a fixed point of an operator that maps what we observe into what we believe.

Key words: DSGE; Partially Observable Markov Decision Process (POMDP); Observation Channel; Kalman filter

JEL Classification: D58, D83, E13

1 Introduction

Research in economic dynamics has gone through a remarkable transformation over the last three decades, with flourishing theories that explicitly deal with economic agents operating through time in stochastic environments. Among these theories, a particular collection of the models that are being developed and solved within a general equilibrium framework is often called 'Dynamic Stochastic General Equilibrium' (DSGE). The DSGE models share one common structure: the Markov, which we economists call recursive structure.

*I thank Hans Amman (the Editor) and two anonymous referees for their comments which have greatly improved the present work. I have also benefited from conversation with Kaushik Mitra, Seongman Moon, and Alan Sutherland. All remaining errors are mine.

†*s.kim@st-andrews.ac.uk*. School of Economics and Finance, University of St Andrews, Fife, United Kingdom, KY16 9AL.

Markov Structure		Do we have control over the state transitions?	
		NO	YES
Are the states completely observable?	YES	MC Markov Chain	MDP Markov Decision Process
	NO	HMM Hidden Markov Model	POMDP Partially Observable Markov Decision Process

Table 1: A Classification of the Markov Structure: Cassandra (2005)

Most of the Markovian DSGEs build on one restrictive assumption that economic agents are able to see through the true state of a world. However, under almost all dynamic stochastic environments, the true state of a world neither completely reveals itself nor evolves in isolation of optimizing agents' decisions. A few examples of such recursive structure would assure that it is everywhere. *Example 1*: Investors often face trouble in valuation of the current fundamentals, while their current investment decisions determine the true value of the future fundamentals. *Example 2*: People cannot see through the state of ecosystem (linked to the accumulation of carbon dioxide and other greenhouse gases), while their current consumption decisions influence the future state of ecosystem. *Example 3*: A robot cannot penetrate its current position, while its current movement decisions lead to its next position.

In the artificial intelligence literature, the recursive structure common in all the three examples above is classified as 'Partially Observable Markov Decision Process' (POMDP). As presented in Table 1, the classification of the Markov structure is made by looking at (*i*) whether a robot can observe the state of a world and (*ii*) whether the robot can influence the state transition.[1] Mapping the DSGE models onto the classification of the Markov structure, this paper positions the class of POMDP to the center of a generalization of linear rational expectations models, and derives a solution algorithm by formulating the

[1]In the words of Kendrick and Amman (2006), the classification system shown in Table 1 and 2 takes 'stochastic elements' for granted, and then focuses on two attributes; (1) whether the stochastic elements are about additive noisy terms or unknown states, and (2) whether the stochastic elements require feedback rules (closed loop) as a part of optimal solution to address endogeneity between forward-looking variables and beliefs about the states.

DSGE Models		Do we have control over the state transitions?	
		NO	YES
Are the states completely observable?	YES	MC (All the DSGEs today contain the MC blocks within their models)	MDP Lucas (1978) Smets and Wouters (2003) Blanchard and Gali (2007)
	NO	HMM Kydland and Prescott (1982) Cargetti, et al. (2002) Lorenzoni (2009)	POMDP Sargent (1991) Svennson and Woodford (2003, 2004)

Table 2: A Classification of the DSGE Models

equilibrium as a fixed point of an operator that maps what agents observe into what agents believe. With the proposed solution algorithm, we will also discuss how one can use the POMDP class to address the issue of "trouble in capital valuation" which recent research has been attempting to incorporate into the DSGE framework in the wake of the 2007 financial crisis.

Toward the goal, we first review several well-known DSGE models (listed in Table 2) upon the classification of the Markov structure, and find clues to the solution concept for the DSGE models that fall into the POMDP class.

First, a recursive structure in which the state of a world is observable and the state transition is endogenous falls into a class of 'Markov Decision Process' (MDP). In fact, the majority of the Markovian DSGEs is of the MDP class. Early examples of the 'MDP-class DSGEs' include Lucas (1978) consumption-based capital asset pricing model. In his model, a single consumer robot decides the next period asset holdings (MDP; endogenous and observable state transition), while output production follows a 'Markov Chain' (MC; entirely exogenous and observable).[2] Many New Keynesian DSGE models also fall into this MDP class. An interesting recent example from the New Keynesian DSGE literature is Blanchard and Gali (2007). In their model, the state of an economy is summarized in employment-unemployment whose transition is governed by optimized search mechanism of labor market flow. Another notable example of the MDP class is the Dynamic Integrated Climate Economy (DICE) pioneered by Nordhaus (1992) and recently resumed by Per

[2] Almost all the DSGEs today contain the MC blocks within the framework. So in mapping the DSGEs onto the Markov structure, we do not need to classify *pure* "MC-DSGEs".

Krusell in collaboration with natural scientists. It assumes that the endogenous state of ecosystem is observable, in terms of the accumulation of accountable greenhouse gases.

Another class of the DSGEs admits the presence of noises around the neighborhood of the state of a world, but still assumes that the state surrounded by such noises follows an exogenous process. Under the environments, estimation and controls become dissociable, and therefore a signal extraction problem can be solved outside actual decision process of a robot. The DSGEs with this recursive structure fall into a class of Hidden Markov Model (HMM). Kydland and Prescott (1982) is an early example of the 'HMM-class DSGEs', in which they describe the state of aggregate technology as a combination of transitory and permanent and noisy components (all exogenously given). Consequently, the estimation problem facing a robot under the HMM environments is about how to untangle multiple sources. Indeed, the source separation problem has been extensively used in the modern macroeconomic literature. Just to mention a few, Cagetti et al. (2002) look at the role of jump variables in asset pricing when their infrequent movements are not perfectly distinguishable from small gradual shocks. Lorenzoni (2009) facilitates the source separation problem to give rise to noise-driven business cycles in a many-heterogeneous-agent world.[3]

Smets and Wouters (2003) is another interesting example for both classes of the DSGEs; MDP and HMM. Agents live in and optimize in the MDP environments. Thus the theoretical side of their model falls into the MDP class. Ironically though, the modellers Smets and Wouters live in the HMM environments and work as (Bayesian) econometrician since the modellers themselves cannot observe what they assume their agents observe. Nevertheless, neither classes can represent dynamic stochastic environments well in which our chosen actions influence the transition of the states that are not completely observable.

Control problems under such dynamic stochastic environments have been classified as POMDP (as opposed to MDP) in the artificial intelligence literature since Sondik (1971)'s pioneering work. In the POMDP world, it is not only exogenous random events that make a world uncertain, but also a robot's inability to see through the past and the present of the world.

Such view about the world is not new in economics. Pigou (1927) understood indus-

[3] Contrary to those recent examples of HMM-DSGEs, Kydland and Prescott (1982) facilitate the HMM information structure to justify their pioneering development of calibration approach: "Our approach is to focus on certain statistics for which the noise introduced by approximations and measurement errors is likely to be small relative to the statistic"(Kydland and Prescott, 1982, p.1360). It is this approach that Geweke (1999) calls a weak econometric interpretation of DSGEs.

trial fluctuations as an outcome arising when the current production decision is based on forecasts of the future market condition, which will be in part determined by the current production decision. Townsend (1983) and Sargent (1991) formalize this idea by modelling each agent's action to be a function of his own beliefs formation (a robot's inner state) about all the other agents' beliefs (hidden state), of which aggregation in turn determines the transition of the market condition (the true state of a world). More recent examples of the 'POMDP-class DSGEs' include Svensson and Woodford (2003, 2004): A central bank conducts monetary policy without perfect information about the state of an economy (production potential), and influences the state transition (the next period output gap) within a New Keynesian framework.

We borrow the solution concept from the POMDP literature where modellers treat a robot's beliefs about the state of a world as another state (the robot's inner state) and derive the robot's optimal action as a function of its inner state.[4] Applying the solution concept to linear rational expectations models, we establish an optimal action rule as a function of agent's inner state (beliefs) about the state of a world, and obtain an optimal transition rule of the inner state (beliefs) that is consistent with the optimal action rule. So the complete set of solution consists of (*i*) optimal action and (*ii*) optimal sequential beliefs about the state of a world whose true transition depends on chosen actions.

With respect to derivation of solution, we take a strategic idea from Sargent (1991)'s approach: Sargent (1991) starts with conjecture about perceived laws of motion and derives an equilibrium as a fixed point of mapping between perceived (forecasting) and actual laws of motion. In contrast, we derive an equilibrium as a fixed point of instantaneous mapping from observables to beliefs. And thereby, the solution algorithm proposed here does not require that the entire perceived and actual laws of motion, both predictable parts and residuals, are identical at a fixed point. With respect to the filtering properties, the present paper precisely replicates those obtained by Svensson and Woodford (2003, 2004) and Baxter, Graham, and Wright (2011), since the baseline reasoning in their works and the present paper is the same in a Bayesian direction and the environments considered are all linear and Gaussian. So it is not surprising to see a modification of the Kalman filter from both their works and here, because it is already known that the Kalman filter is

[4]For example, Kaelbling, Littman, and Cassandra (1998) discretize the belief state space and develop a computation algorithm under a simplified nonlinear environments for a robot navigation problem like *Example 3* told above. Littman (2009) expounds their solution concept for the audience of behavioral scientists.

optimal under linear environments with the Gaussian distribution. However, the context used here differs much from theirs.[5]

The rest of this paper is organized as follows: Section 2 uses a simple Neoclassical growth model as a prototype of MDP and creates its two variants; HMM and POMDP. We use the familiar textbook model to see how the proposed classification works, and how three different classes of DSGEs are related with each other. Section 3 considers a general state-space representation of linear rational expectations models. To find a POMDP solution for the system of structural equations, we follow a conjecture-verification procedure in two stages and show a transparent analytic derivation of the complete set of solution. Section 4 provides quantitative examples using the Neoclassical growth model and its variants. It also performs some quantitative experiments to address the issue of "trouble in capital valuation". Section 5 makes some concluding remarks.

2 The Markovian DSGEs

In this section, we consider a simple Neoclassical growth model as a prototype example of MDP, and creates its two variants by applying the classification of the Markov structure we have previously discussed.

2.1 Prototype Model: A Neoclassical Growth Model

Agents are infinitely-lived and all identical. A representative agent maximizes

$$E_0\left[\sum_{t=0}^{\infty}\beta^t\left\{\ln C_t+\gamma\ln(1-N_t)\right\}\right]$$

subject to

$$\begin{cases} Y_t = Z_t K_t^{\alpha} N_t^{1-\alpha} \\ K_{t+1} = (1-\delta)K_t + I_t \\ \ln Z_{t+1} = \rho \ln Z_t + \omega_{t+1} \\ Y_t = C_t + I_t, \end{cases}$$

[5] As classified and briefed above, Svensson and Woodford (2003, 2004) address the optimal formation of monetary policy with imperfect indicators about the state of an economy. Baxter, Graham, and Wright (2011)'s concerns are under which conditions observable variables continue to remain as imperfect indicators about the state of an economy and its implications for the time series properties.

where C_t denotes consumption for period t, N_t worked hours, Y_t output, K_t predetermined capital stock, I_t investment, and Z_t the level of production technology. Regarding the stochastic base of the model, ω_t is technical innovation to Z_t and subject to an *i.i.d.* normal with mean zero and standard deviation σ. All variables are in per person. All the parameters $\{\alpha, \beta, \gamma, \delta, \rho\}$ lie on the surface that guarantees stability of the system. This textbook model as it is falls into a class of MDP.

After briefly looking at the equilibrium conditions for this MDP version, we turn to two more variants of the model one after the other. As we proceed, we will relax the assumption of complete observation for each state variable. But, throughout the paper, we maintain the assumption that output Y_t is observable.

2.2 The Classification of the Markovian DSGEs

Markov Decision Process (MDP)

If both exogenous and endogenous states (Z_t, K_t) are observable, we have an MDP economy in which the transition of capital stock K_t is determined by the current investment decision. The equilibrium conditions for this textbook model are given by

$$\begin{cases} \text{optimality conditions:} & \begin{cases} C_t^{-1} = \beta E_t[C_{t+1}^{-1}\{\alpha Z_{t+1} K_{t+1}^{\alpha-1} N_{t+1}^{1-\alpha} + (1-\delta)\}] \\ \gamma(1-N_t)^{-1} = (1-\alpha) C_t^{-1} Z_t K_t^{\alpha} N_t^{-\alpha} \end{cases} \\ \text{transition equation (endg. state):} & K_{t+1} = (1-\delta)K_t + (Z_t K_t^{\alpha} N_t^{-\alpha} - C_t) \\ \text{transition equation (exog. state):} & \ln Z_{t+1} = \rho \ln Z_t + \omega_{t+1}. \end{cases} \tag{1}$$

Hidden Markov Model (HMM)

If we can observe the endogenous state K_t but not the exogenous state Z_t, we now have an HMM version of the Neoclassical growth model. Agents have to extract signal about the exogenous state Z_t from the history of output Y_t and worked hours N_t. As far as the agents can observe the true value of capital stock without noises, they would able to make perfect inference about Z_t based on the production function $Y_t = Z_t K_t^{\alpha} N_t^{1-\alpha}$.

To make the HMM environments non-trivial, let us introduce a transitory component of aggregate technology following Kydland and Prescott (1982).[6] To be specific, we assume

[6] One common practice mostly performed in this direction is to have many noises (for example, news shock and sampling shock in Lorenzoni, 2009), or to have an exogenous state as the sum of multiple sources

$Z_t = Z_{1t}Z_{2t}$, where Z_{1t} replaces the role of Z_t in the MDP version and Z_{2t} is a pure transitory shock. Accordingly, we have the equilibrium conditions (1) modified to

$$\begin{cases} \begin{cases} C_t^{-1} = \beta E_t[C_{t+1}^{-1}\{\alpha Z_{1t+1}Z_{2t+1}K_{t+1}^{\alpha-1}N_{t+1}^{1-\alpha} + (1-\delta)\}] \\ \gamma(1-N_t)^{-1} = (1-\alpha)C_t^{-1}E_t[Z_{1t}Z_{2t}]K_t^{\alpha}N_t^{-\alpha} \end{cases} \\ K_{t+1} = (1-\delta)K_t + (Z_{1t}Z_{2t}K_t^{\alpha}N_t^{1-\alpha} - C_t) \\ \ln Z_{1t+1} = \rho \ln Z_{1t} + \omega_{1t+1} \\ \ln Z_{2t} = \omega_{2t}, \end{cases} \tag{2}$$

where ω_{1t} are ω_{2t} are orthogonal, and both subject to an *i.i.d.* normal with mean zero and standard deviation σ_1 and σ_2, respectively. Notice that the second optimality condition (intratemporal) takes conditional expectations operator over the exogenous states Z_{1t} and Z_{2t}, because neither sources of shocks (ω_{1t}, ω_{2t}) are observable. To set up inference problem for this version of the model, we now have to specify observation channels. As the production function connects the state variables to the level of output observable, we use the production function as an observation channel, and in addition assume that the agents know the true value of capital:

$$\text{observation channel}: \begin{cases} Y_t = Z_{1t}Z_{2t}K_t^{\alpha}N_t^{1-\alpha} \\ K_t = E_t[K_t]. \end{cases} \tag{3}$$

As emphasized by the second equation in (3), the agents in this HMM version of the model observe the true value of capital stock K_t. From the observables, the agent extracts signal on permanent shock Z_{1t} and transitory shock Z_{2t}.—This is a typical example of source separation problem. Since the transition of Z_{1t} and Z_{2t} is independent of the action the agent takes, the signal extraction problem will be solvable in isolation from the optimization problem. If the HMM environments are linear, the Kalman filter (a linear Bayesian inference algorithm) is known optimal.

Partially Observable Markov Decision Process (POMDP)

In reality, the true value of capital itself may not be easily observable. Indeed, recent research in the DSGE literature has drawn more attention to the issue of "trouble in

of hidden shock (for example, infrequent jump shock and frequent gradual shock in Cagetti et al., 2002).

capital valuation". See, for example, "capital-quality shock" (Gertler and Karadi, 2009) and "perception shock" (Curdia, 2008).

Let us directly address this issue, by simply relaxing the assumption that agents are able to observe the true value of aggregate capital. As should be clear to the readers by now, this version of the model with "unobservable endogenous state" falls into a class of the POMDP. With the prototype model, the equilibrium conditions under the POMDP environments appears to remain the same with the HMM's (2) and (3) other than the second observation channel. That is, the observation channel this time reduces to

$$\text{observation channel}: \quad Y_t = Z_{1t} Z_{2t} K_t^{\alpha} N_t^{1-\alpha}. \tag{4}$$

However, the two variants of the model fundamentally differ in inference and optimization procedure.—If the true value of capital is unobservable, the consumption-investment decision cannot be made in isolation of signal extraction on the current capital, neither in isolation of forecasts of the next period capital.

One common feature between the HMM and the POMDP is that the dimension of the observation channel is less than the number of state variables. To check with the prototype model, the agents in the HMM variant extract signals on three unobservables (Z_{1t}, Z_{2t}, K_t) from the two observation channels in (3). In effect, they extract signals on two unobservables (Z_{1t}, Z_{2t}) from one observation channel, the first in (3). And the agents in the POMDP variant extract signals on three unobservables (Z_{1t}, Z_{2t}, K_t) from the single observation channel (4).

Upon the common feature between the two variants and our existing knowledge that the Kalman filter is optimal in linear HMM environments, we configure the complete solution for both HMM and POMDP class as a set of (*i*) optimal action and (*ii*) optimal beliefs about the state. In turn, the optimal transition equation is obtained where sequential action and beliefs are at the optimal path.

To utilize our existing knowledge about the linear filtering properties, we linearize each variant of the prototype model at the steady state. However, rather than herein listing up the log-linearization of all three variants of the prototype model, in the next section we will map them into a generalized linear rational expectations models and the optimal solution for each class. Henceforth, following the convention in the literature, we use lowercase variables denote log-deviations from their steady state.

3 Solution to the POMDP Class DSGEs

In this section, we consider a general state-space representation of linear rational expectations models, brief the solution for the standard MDP class, and then direct a full attention toward derivation of an optimal solution algorithm for the POMDP class.

3.1 MDP Class

A General State-Space Representation Consider the following structural form of linear rational expectations models:

$$0 = A_{cc}E_t[c_{t+1}] + A_{ck}k_{t+1} + A_{cz}E_t[z_{t+1}] + B_{cc}c_t + B_{ck}k_t + B_{cz}z_t \tag{5}$$

$$0 = A_{kk}k_{t+1} + B_{kc}c_t + B_{kk}k_t + B_{kz}z_t \tag{6}$$

$$0 = A_{zz}z_{t+1} + A_{z\omega}\omega_{t+1} + B_{zz}z_t, \tag{7}$$

where c_t is a $r_c \times 1$ vector of non-predetermined choice variables, k_t a $r_k \times 1$ vector of predetermined variables, and z_t a $r_z \times 1$ vector of exogenous variables (including noises) whose stochastic base lies on shocks ω_t of the same dimension. ω_t is *i.i.d.* with mean zero and variance-covariance matrix $\sigma^2 I_{r_z}$, where σ is a scalar. $\{A_{ij}\}$ and $\{B_{ij}\}$ are the system coefficient matrix. Throughout the paper, we reserve the letter A for notation of system coefficients associated with the next period, and the letter B for system coefficients associated with the current period. Each 0 on the left-hand side of equations denotes a zero matrix conformable to the dimension of each given equation.

Notations for variables and system coefficients are all consistent with the linearized system of the Neoclassical growth example (1). The control vector c_t is like the consumption choice in the model (1); the predetermined vector k_t is like the capital stock in the example; the exogenous state vector z_t is like the exogenous technologies. So the coefficient matrices A_{ij} and B_{ij} are a function of parameters $\{\alpha, \beta, \gamma, \delta, \rho\}$ and steady state values of the Neoclassical growth model.

Also, the arrangement of equations remains parallel between the generalized state space system and the prototype model: Equation (5) is the optimality condition like the first two equations in (1). Similarly, (6) and (7) are the transition equations for endogenous and exogenous state variables, like those in the Neoclassical growth model (1).

A number of solution algorithms to linear rational expectations models with observable

ω_t have been developed since Blanchard and Kahn (1980).[7] Assuming the existence of the unique solution to the system {(5), (6), (7)}, we have the optimal control function as follows;

$$c_t = H_{ck} k_t + H_{cz} z_t.$$

H_{ij} denotes the optimal solution coefficient matrix for 'variable vector i' with respect to 'variable vector j'. For instance, H_{ck} is the optimal control coefficient for c with respect to the endogenous state vector k. Throughout the paper, we reserve the letter H for notation of the optimal solution coefficients.

Stacked System Let θ_t denote the state vector of size $r_\theta = r_k + r_z$;

$$\theta_t = \begin{bmatrix} k_t \\ z_t \end{bmatrix}.$$

Consistently, we rewrite the system {(5), (6), (7)} as follows;

$$0 = A_{cc} E_t[c_{t+1}] + A_{c\theta} E_t[\theta_{t+1}] + B_{cc} c_t + B_{c\theta} \theta_t, \tag{8}$$

$$0 = A_{\theta\theta} \theta_{t+1} + B_{\theta c} c_t + B_{\theta\theta} \theta_t + A_{\theta\omega} \omega_{t+1}, \tag{9}$$

where

$$A_{\theta\theta} = \begin{bmatrix} A_{kk} & 0 \\ 0 & A_{zz} \end{bmatrix}, \; A_{c\theta} = \begin{bmatrix} B_{ck} & B_{cz} \end{bmatrix}, \; A_{\theta\omega} = \begin{bmatrix} 0 \\ A_{z\omega} \end{bmatrix},$$

$$B_{\theta\theta} = \begin{bmatrix} B_{kk} & B_{kz} \\ 0 & B_{zz} \end{bmatrix}, \; B_{c\theta} = \begin{bmatrix} B_{ck} & B_{cz} \end{bmatrix}, \; B_{\theta c} = \begin{bmatrix} B_{kc} \\ 0 \end{bmatrix}.$$

Equation (8) stacks endogenous and exogenous state vectors $[\; k_t \quad z_t \;]'$ into θ_t. Equation (9) stacks the two transition equations (6) and (7). We have the optimal solution in the minimal state variable (MSV) form as in McCallum (1998);

$$\text{MDP} \begin{cases} \text{optimal action rule :} & c_t = H_{c\theta} \theta_t, \\ \text{optimal state transition :} & \theta_{t+1} = H_{\theta\theta} \theta_t + H_{\theta\omega} \omega_{t+1}, \end{cases} \tag{10}$$

[7] An incomplete list includes Binder and Pesaran (1995), King and Watson (1998), and Klein (2000), each of which facilitates a different method— matrix polynomial, system reduction (eigenvalue), and the QZ method (Schur form).

where

$$H_{c\theta} = [\ H_{ck} \quad H_{cz}\ ],\ H_{\theta\theta} = -A_{\theta\theta}^{-1}\left\{B_{\theta c}H_{c\theta} + B_{\theta\theta}\right\},\ H_{\theta\omega} = -A_{\theta\theta}^{-1}A_{\theta\omega},$$

provided that $A_{\theta\theta}$ is of a full rank.

3.2 The Observation Channel

Observables and Unobservables Suppose now that the state vector θ_t is unobservable. As a consequence, we have to infer it from what we observe. Let y_t be a $r_y \times 1$ vector of the observable variables subject to the following observation channel;[8]

$$0 = B_{yy}y_t + B_{yc}c_t + B_{y\theta}\theta_t. \tag{11}$$

This observation channel (11) is like the production function (4) of the Neoclassical growth example.

Given the optimal action rule in (10) and with B_{yy} of a full rank, the observation channel (11) can be understood as

$$y_t = -\{B_{yy}^{-1}B_{yc}H_{c\theta} + B_{yy}^{-1}B_{y\theta}\}\theta_t.$$

Alternatively,

$$y_t = H_{y\theta}\theta_t,$$

with $H_{y\theta} = -\{B_{yy}^{-1}B_{yc}H_{c\theta} + B_{yy}^{-1}B_{y\theta}\}$. We can think of $H_{y\theta}$ as an optimal observation coefficient in that it connects unobservables to observables when the agents follow the optimal action rule in (10).

MDP-Equivalent POMDP It is clear that when the dimension of the observation channel (r_y) equals the dimension of the state vector (r_θ), $H_{y\theta}$ is a square matrix of the exactly same size with $H_{\theta\theta}$. Provided that $H_{y\theta}$ is of a full rank, then

$$\theta_t = H_{y\theta}^{-1}y_t.$$

[8] To avoid redundancy in expressions, the observation dimension r_y herein does not take in account the control dimension r_c. If wanted, one may do so by adding expression for the controls, $c_t = E_t[c_t]$, to the observation channel, and accordingly redefining observation vector y_t and the observation condition. The issue of partially observable controls (a trembling-hand story) goes beyond the scope of the present paper.

Here are two notable points: First, the state variables can be correctly identified if $r_y = r_\theta$ and $H_{y\theta}$ is of a full rank.—As an implication, noisiness of noises depends on the relative size of dimensions between the observables and the unobservables in a given observation channel. Second, all the past data $\{y_{t-1}, y_{t-2}, \cdots\}$ have no informational role in identifying the current states.—Data y_t instantaneously unveils the temporal state of the economy θ_t in the aid of invertible structural knowledge $H_{y\theta}$.[9] As far as $H_{y\theta}$ is of rank $r_y = r_\theta$, the optimal solution (10) under the MDP environments remains unchanged even when we cannot observe directly the true state θ_t. In this sense, we can think of the MDP class as a special case of the POMDP.

3.3 POMDP Class

Let us turn to the partially observable linear rational expectations models. If the dimension of the observation channel is less than the dimension of the state vector θ_t, we find ourselves in either the HMM or the POMDP environments. Recall that $H_{y\theta} = -\{B_{yy}^{-1}B_{yc}H_{c\theta} + B_{yy}^{-1}B_{y\theta}\}$ and that it connects unobservables to observables at optimum. So the matrix is at the center of the classification of the Markovian DSGEs. Just like that we have an MDP-equivalent POMDP-DSGE model when $H_{y\theta}$ is of a rank $r_y = r_\theta$, we would say that we have an HMM-equivalent POMDP-DSGE model when $r_k \le r_y < r_\theta$ with k_t all observable. Since we can treat the HMM class as a special case of the POMDP where all endogenous states are observable, our focus is on non-trivial cases of the POMDP class in which some or all endogenous states are unobservable.

The POMDP class of linear rational expectations models takes the following structural form;

$$\begin{aligned} \text{optimality condition:} \quad 0 &= A_{cc}E_t[c_{t+1}] + A_{c\theta}E_t[\theta_{t+1}] + B_{cc}c_t + B_{c\theta}E_t[\theta_t] && (12) \\ \text{transition equation:} \quad 0 &= A_{\theta\theta}\theta_{t+1} + B_{\theta c}c_t + B_{\theta\theta}\theta_t + A_{\theta\omega}\omega_{t+1} && (13) \\ \text{observation channel:} \quad 0 &= B_{yy}y_t + B_{yc}c_t + B_{y\theta}\theta_t. && (14) \end{aligned}$$

Equation (12) takes conditional expectations operator to the current state vector θ_t, since we can no longer observe it. So the observation channel (14) now becomes an essential part of the system. On the other hand, the transition equation (13) is a kind of natural law and thus remains the same as in the MDP class.

[9] Even when $r_y < r_\theta$ but provided that $r_y = r_z$, the state may be 'asymptotically invertible' from the observables (in the sense used by Baxter, Graham, and Wright, 2011) under some regularity conditions.

We strategically divide-and-conquer the procedure of solving the system into two steps. In the first step, we suppose that optimal sequential inference problem has been solved, and focus on the first two equations (12) and (13). We start with conjecture that optimal choice is linear in optimal point beliefs, and then verify the existence of the optimal action rule that satisfies the conjecture. In the second step, taking the established optimal action rule as given, we look at the last two equations (13) and (14), apply the principle of mean squared error minimization, and find a sequence of optimal point beliefs about the true state. To achieve this, we again take the route of conjecture-then-verification on the basis of our existing knowledge about probability update.

3.3.1 Solution Step 1: Optimal Action Rule

Conjecture Let $x_{t+s|t}$ denote conditional expectations of x_{t+s} on date t-information set $\mathcal{Y}_t = \{y_s, c_s\}_{s=0}^{t}$. $E_t[x_{t+s}] = x_{t+s|t}$. Following a standard practice in solving linear rational expectations models, we consider the certainty-equivalent of optimal action rule in (10), and make conjecture that optimal action is linear in conditional expectations;

$$c_t = H_{c\theta} E_t[\theta_t] = H_{c\theta}\theta_{t|t}. \tag{15}$$

In the context of the Neoclassical growth example, this conjecture means that the agent makes consumption decision in a proportion to his beliefs as to how valuable capital he owns, and in a proportion to his beliefs as to how productive technology he accesses.

Given the conjecture (15), we understand the first two equations (12) and (13) as follows

$$0 = \{A_{cc}H_{c\theta} + A_{c\theta}\}\,\theta_{t+1|t} + \{B_{cc}H_{c\theta} + B_{c\theta}\}\,\theta_{t|t} \tag{16}$$

$$0 = A_{\theta\theta}\theta_{t+1} + B_{\theta c}H_{c\theta}\theta_{t|t} + B_{\theta\theta}\theta_t + A_{\theta\omega}\omega_{t+1}, \tag{17}$$

respectively. To obtain the first term in (16), we take one period forward iteration of the conjectured optimal action rule ($c_{t+1} = H_{c\theta}\theta_{t+1|t+1}$), and take conditional expectations on date t information, and use the law of iterated expectations ($c_{t+1|t} = H_{c\theta}\theta_{t+1|t}$). The transition equation (17) states that, besides exogenous forces, the present state and beliefs about the present state jointly determine the future state. It is noteworthy that the auto-regressive coefficient for the transition of the state θ_t is now $-A_{\theta\theta}^{-1}B_{\theta\theta}$ under the optimal action rule (15) in contrast to the case of the MDP. In the MDP version of the model, we have the auto-regressive transition coefficient $H_{\theta\theta} = -A_{\theta\theta}^{-1}\{B_{\theta c}H_{c\theta} + B_{\theta\theta}\}$ as shown by the optimal transition equation in (10).

Verification To verify our conjecture about the optimal action rule (15) under the POMDP environments, take first conditional expectations operator to the transition equation (13) and rearrange it as $\theta_{t+1|t} = -A_{\theta\theta}^{-1}\{B_{\theta c}H_{c\theta} + B_{\theta\theta}\}\theta_{t|t}$. By assumption, $\omega_{t+1|t} = 0$.

Substitute it back into the optimality condition (16). Then,

$$0 = \Big\{-\{A_{cc}H_{c\theta} + A_{c\theta}\}A_{\theta\theta}^{-1}\{B_{\theta c}H_{c\theta} + B_{\theta\theta}\} + \{B_{cc}H_{c\theta} + B_{c\theta}\}\Big\}\theta_{t|t}.$$

By the reasoning of undetermined coefficients, it should hold that

$$\{A_{cc}H_{c\theta} + A_{c\theta}\}A_{\theta\theta}^{-1}\{B_{\theta c}H_{c\theta} + B_{\theta\theta}\} = \{B_{cc}H_{c\theta} + B_{c\theta}\}.$$

As far as there exists an MDP solution (10) such that it satisfies the Blanchard and Kahn (1980) conditions or Klein (2000), there will be a root to the above quadratic equation with respect to matrix $H_{c\theta}$.

3.3.2 Solution Step 2: Optimal Sequential Beliefs

Conjecture In the second step, we stick to the principle of mean squared error (MSE) minimization and derive optimal recursive inference formula. Given the optimal policy function found in Step 1, the observation channel (14) will be understood as

$$y_t = -B_{yy}^{-1}B_{yc}H_{c\theta}\theta_{t|t} - B_{yy}^{-1}B_{y\theta}\theta_t.$$

For visibility and tractable operation, let us write it to

$$\underbrace{y_t}_{\text{what we observe}} = \underbrace{B_{\tilde{y}c}H_{c\theta}\theta_{t|t}}_{\text{what we believe}} + \underbrace{B_{\tilde{y}\theta}\theta_t}_{\text{what truly is}} \tag{18}$$

where $B_{\tilde{y}c} = -B_{yy}^{-1}B_{yc}$ and $B_{\tilde{y}\theta} = -B_{yy}^{-1}B_{y\theta}$. Equation (18) says that 'what we observe' is a combination of 'what we believe' and 'what truly is'. So the basic idea toward a plausible conjecture is to reverse the direction of (18) and to instantaneously connect 'what we observe' to 'what we believe' in a Bayesian direction.

In the context of the Neoclassical growth example, we have $B_{\tilde{y}c} = \begin{bmatrix} 0 & 1-\alpha \end{bmatrix}$ with respect to the control vector [consumption c_t, worked hours n_t], and $B_{\tilde{y}\theta} = \begin{bmatrix} \alpha & 1 & 1 \end{bmatrix}$ over the state vector [capital k_t, permanent technology z_{1t}, transitory technology z_{2t}]. Mapping Equation (18) to the prototype model tells us that once the agent makes consumption and leisure decisions based on his beliefs about the state of the economy, the level of production depends on in part his beliefs about the productivity of capital and technologies (permanent and transitory), and in part the true productivity. So we can conjecture that an optimal inference rule held by the agent will instantaneously connect the actual level of production and labor input to the valuation of capital and technologies.

To implement the basic idea about an optimal inference rule in a sequential manner, suppose that we stand in date $t-1$ and try to forecast y_t one period ahead. Conditional on the history of observables $\mathcal{Y}_{t-1}$, forecast of y_t directly follows from the formation of conditional expectations about (18);

$$y_{t|t-1} = E[y_t \mid \mathcal{Y}_{t-1}] = \{B_{\tilde{y}c}H_{c\theta} + B_{\tilde{y}\theta}\}\,\theta_{t|t-1}.$$

Therefore, the one-period ahead forecast error of date t-observation y_t consists of (i) revision in conditional expectations about date t-state and (ii) the one-period forecast error of date t-state;

$$y_t - y_{t|t-1} = B_{\tilde{y}c}H_{c\theta}(\theta_{t|t} - \theta_{t|t-1}) + B_{\tilde{y}\theta}(\theta_t - \theta_{t|t-1}). \tag{19}$$

To deal with the expectations-revision term in (19), we make conjecture that there exists a matrix $\tilde{M}_t$ of the size $r_\theta \times r_y$ such that it extracts optimal signal for revision out of the forecast error about the observables;

$$\underbrace{(\theta_{t|t} - \theta_{t|t-1})}_{\text{updating beliefs}} = \tilde{M}_t\underbrace{(y_t - y_{t|t-1})}_{\text{forecast errors}}. \tag{20}$$

This conjecture follows the standard algorithm of sequential Bayesian inferences, in which posterior emerges from both new observations and prior. Again in the context of the Neoclassical growth example, the conjecture (20) means that any errors in output forecast must serve to improve the valuation of capital and technologies in some proportion specified by $\tilde{M}_t$.

We will verify this conjecture by finding and establishing $\tilde{M}_t$ as a fixed point of a function that minimizes mean squared error (MSE). As Lucas (1978, p.1431)'s rational expectations assumption directs, the solution concept of closed loop is central here as well to obtain optimal recursive inference solution. In what follows, we will introduce three more functions (D, G, M), and will go through some tedious procedural derivation. To make it as transparent as possible, readers are advised to have a look at the following road map:

$$\begin{array}{ccccc} \tilde{M}_t & \rightarrow & D(\tilde{M}_t) & \rightarrow & G(D_t) \\ \Updownarrow & & & & \downarrow \\ M_t & & \Longleftrightarrow & & M(G_t) \end{array}$$

To brief it, we will now define D, G, and M as a function of conjectured matrix $\tilde{M}_t$ such that it satisfies (20), and show that $\tilde{M}_t$ is a fixed point of function $M(\cdot)$. That is, $M(G_t) = M(G(D(\tilde{M}_t))) = \tilde{M}_t$.

Least MSE Given the conjecture (20), we can rewrite (19) to

$$y_t - y_{t|t-1} = B_{\tilde{y}c} H_{c\theta} \tilde{M}_t (y_t - y_{t|t-1}) + B_{\tilde{y}\theta}(\theta_t - \theta_{t|t-1}). \tag{21}$$

Define a function D of a full rank r_y such that $D_t = I_{r_y} - B_{\tilde{y}c} H_{c\theta} \tilde{M}_t$, where I_{r_y} stands for r_y-dimensional identity matrix. In turn, define a function G such that $G_t' = D_t^{-1} B_{\tilde{y}\theta}$. Then, we have (21) in terms of the function G as follows;

$$y_t - y_{t|t-1} = G_t'(\theta_t - \theta_{t|t-1}). \tag{22}$$

Let $\hat{\Sigma}_t = E[(\theta_t - \theta_{t|t-1})(\theta_t - \theta_{t|t-1})' \mid \mathcal{Y}_{t-1}]$, the MSE associated with one-period ahead state forecast; and $\Sigma_t = E[(\theta_t - \theta_{t|t})(\theta_t - \theta_{t|t})' \mid \mathcal{Y}_t]$, the MSE associated with the inference of date t-state. Using the well-known formula for least square forecast with finite sample (Hamilton, 1994: p.99, p.379), we have that

$$\theta_{t|t} = \theta_{t|t-1} + M_t(y_t - y_{t|t-1}), \tag{23}$$

$$\Sigma_t = \left\{ I_{r_\theta} - M_t G_t' \right\} \hat{\Sigma}_t, \tag{24}$$

where M_t is defined as a gain matrix following the conventional practice of the Kalman filter;

$$M_t = \hat{\Sigma}_t G_t \{G_t' \hat{\Sigma}_t G_t\}^{-1}. \tag{25}$$

See Appendix for a full derivation of (23) and (24).

Verification The early conjecture (20) and the equation (23) derived from the least MSE principle jointly give rise to $\tilde{M}_t = M_t$, which is the closed loop requirement to obtain the solution to optimal recursive inference problem for partially observed dynamic system. Note that $\tilde{M}_t$ transforms 'what we observe' to 'what we believe'. Given the conjecture, G_t connects 'what truly is' to 'what we observe' in terms of structural parameters $B_{\tilde{y}c}$, $H_{c\theta}$, and $B_{\tilde{y}\theta}$. But then, by requirement that $\tilde{M}_t = M_t$,

$$G_t' = D_t^{-1} B_{\tilde{y}\theta} = \{I_{r_y} - B_{\tilde{y}c} H_{c\theta} M_t\}^{-1} B_{\tilde{y}\theta}.$$

So with (25), G_t and M_t define *each other*. Henceforth, obtaining a closed form of G_t (and thus M_t) in terms of the knowns $B_{\tilde{y}c}$, $H_{c\theta}$, and $B_{\tilde{y}\theta}$ is equivalent to find $\tilde{M}_t$ and thus verify the early conjecture made in (20). Indeed, we obtain a closed form solution G_t as a function of $B_{\tilde{y}c}$ and $B_{\tilde{y}\theta}$ (parameters from observational structure) and $H_{c\theta}$ (optimal reaction parameters) and $\hat{\Sigma}_t$ (covariance of forecast error), as follows;

$$G_t' = \left\{ I_{r_y} + B_{\tilde{y}c} H_{c\theta} \hat{\Sigma}_t B_{\tilde{y}\theta}' \{B_{\tilde{y}\theta} \hat{\Sigma}_t B_{\tilde{y}\theta}'\}^{-1} \right\} B_{\tilde{y}\theta}. \tag{26}$$

See Appendix for a full derivation of (26). Here $\hat{\Sigma}_t$ will be discovered from where recursion starts and continues; $\{\theta_{1|0}, \hat{\Sigma}_1\}$, $\{\theta_{2|1}, \hat{\Sigma}_2\}$, $\{\theta_{3|2}, \hat{\Sigma}_3\}$, $\cdots$. So we complete Step 2 by obtaining the recursion formula for the covariance of forecast error.

Recursion Suppose we want to forecast the future state of the system. Forecast of θ_{t+1} directly follows from the transition equation (17);

$$\theta_{t+1|t} = \{B_{\tilde{\theta}c} H_{c\theta} + B_{\tilde{\theta}\theta}\} \theta_{t|t},$$

where $B_{\tilde{\theta}c} = -A_{\theta\theta}^{-1} B_{\theta c}$ and $B_{\tilde{\theta}\theta} = -A_{\theta\theta}^{-1} B_{\theta\theta}$. It is straightforward to obtain the forecast errors

$$\theta_{t+1} - \theta_{t+1|t} = B_{\tilde{\theta}\theta}(\theta_t - \theta_{t|t}) + H_{\theta\omega}\omega_{t+1}, \tag{27}$$

where $H_{\theta\omega} = -A_{\theta\theta}^{-1} A_{\theta\omega}$. Also we obtain the associated covariance as follows;

$$\hat{\Sigma}_{t+1} = B_{\tilde{\theta}\theta} \Sigma_t B'_{\tilde{\theta}\theta} + H_{\theta\omega} \Sigma_\omega H'_{\theta\omega}, \tag{28}$$

where Σ_ω is the (time-invariant) variance matrix of *i.i.d.* shocks ω_t. See Appendix.

To sum up, the optimal inference-forecast formula in recursion is given by the following set of equations;

$$\begin{cases} \text{point beliefs recursion} \begin{cases} \theta_{t+1|t} = \{B_{\tilde{\theta}c} H_{c\theta} + B_{\tilde{\theta}\theta}\} \theta_{t|t} \\ \theta_{t|t} = \theta_{t|t-1} + M_t (y_t - y_{t|t-1}) \end{cases} \\ \\ \text{covariance recursion} \begin{cases} \hat{\Sigma}_{t+1} = B_{\tilde{\theta}\theta} \Sigma_t B'_{\tilde{\theta}\theta} + H_{\theta\omega} \Sigma_\omega H'_{\theta\omega} \\ \Sigma_t = \{I_{r_\theta} - M_t G'_t\} \hat{\Sigma}_t \end{cases} \end{cases} \tag{29}$$

where

$$M_t = \mathbf{m}_t \left\{I_{r_y} + B_{\tilde{y}c} H_{c\theta} \mathbf{m}_t\right\}^{-1},$$

$$G'_t = \left\{I_{r_y} + B_{\tilde{y}c} H_{c\theta} \mathbf{m}_t\right\} B_{\tilde{y}\theta},$$

with a briefing variable $\mathbf{m}_t$ such that

$$\mathbf{m}_t = \hat{\Sigma}_t B'_{\tilde{y}\theta} \{B_{\tilde{y}\theta} \hat{\Sigma}_t B'_{\tilde{y}\theta}\}^{-1}. \tag{30}$$

Convergence Notice that the convergences of Σ_t, $\hat{\Sigma}_t$, and M_t are all seen in the convergence of $\mathbf{m}_t$, since $M_t G'_t = \mathbf{m}_t B_{\tilde{y}\theta}$ and $\Sigma_t = \{I_{r_\theta} - M_t G'_t\} \hat{\Sigma}_t$. In turn, the convergence of $\mathbf{m}_t$ depends on the observation coefficients, $B_{\tilde{y}\theta} = -B_{yy}^{-1} B_{y\theta}$, by structure of the observation channel (14). One trivial requirement for convergence is to have the observation channel with B_{yy} of a full rank r_y and $B_{y\theta}$ of non-zero matrix. In other words, the system should not completely black out.

On the other hand, inspection of the expression for $\mathbf{m}_t$ (30) shows that the convergence of the covariance matrix, $\hat{\Sigma}_t$, is crucial. In fact, unless the system blacks out, we will see it converging as far as the transition equation (13) satisfies the standard stability conditions as assumed in the DSGE literature. Formally, once the first trivial requirement holds, we find from (24), (25), and (28), that recursion of $\hat{\Sigma}_t$ is governed only by those structural parameters in the transition equation (13); that is, $\{A_{\theta\theta}, B_{\theta\theta}, A_{\theta\omega}\}$ with the unconditional covariance of exogenous shock Σ_ω.

However, the assumed convergence in recursion would be problematic, especially when a world has not long been in the same structure since its birth. One striking example in support for such concerns is found from the literature of equity premium puzzle (and riskfree interest rate puzzle). Weitzman (2007) demonstrates how the assumed convergence to the normal distribution misleads to *believe* 'a particular form of puzzle', which either should not be a puzzle or could be 'a different form of puzzle'.[10] Nevertheless, there is an important practical reason why most theoretical examination of dynamical system tends to focus on the behavior of a system around convergence and steady state.—Having an initial condition is no less ad hoc than assuming convergence. So in theoretical simulation of a dynamic model, we see it as a robust practice to drop data generated over the first few hundreds period of pseudo time within the system. And the analysis of impulse responses, a particular form of theoretical simulation, is also carried out in the same manner.

Later, in Section 4, we will revisit this issue by contrasting the three variants of the Neoclassical growth example in terms of convergence, and show how we can use the class of the POMDP-DSGEs to address the same issue from an alternative perspective.

3.4 Generalized System, Generalized Solution

From POMDP To the class of POMDP-DSGEs, we have a complete solution that consists of (15), (17), and (29). Let us have it comparable to the solution to the MDP class, (10), as follows;

$$\text{POMDP}\quad \begin{cases} \text{optimal action rule :} & c_t = H_{c\theta}\theta_{t|t} \\ \text{optimal state transition :} & \theta_{t+1} = B_{\tilde{\theta}c}H_{c\theta}\theta_{t|t} + B_{\tilde{\theta}\theta}\theta_t + H_{\theta\omega}\omega_{t+1} \\ \text{optimal sequential beliefs :} & \theta_{t|t} = H_{\theta\theta}\theta_{t-1|t-1} + M_t(y_t - y_{t|t-1}), \end{cases} \tag{31}$$

where

$$H_{c\theta} = [\begin{array}{cc} H_{ck} & H_{cz} \end{array}],\ H_{\theta\omega} = -A_{\theta\theta}^{-1}A_{\theta\omega},$$

$$H_{\theta\theta} = B_{\tilde{\theta}c}H_{c\theta} + B_{\tilde{\theta}\theta}.$$

[10] Geweke (2001) shows in general that expectations condition with the usual CRRA (constant relative risk aversion) utility function can easily break down when we relax the distributional assumption.

To MDP If the agents could completely observe the state vector θ_t, the solution reduces to the case in which $\theta_t = \theta_{t|t}$ holds. And the optimal state transition equation and the optimal sequential beliefs rule become redundant, which implies that $M_t(y_t - y_{t|t-1}) = H_{\theta\omega}\omega_t$. It is straightforward to verity this implied relationship.—Note first that $\hat{\Sigma}_t = H_{\theta\omega}\Sigma_\omega H'_{\theta\omega}$ and $\Sigma_t = 0$ for all t, and therefore $M_t G'_t = I_{r_\theta}$. It follows then from (22) that $\theta_t - \theta_{t|t-1} = H_{\theta\omega}\omega_t$. Obviously, if the agents fully observe the true state of an economy, the only source for forecast errors is of exogenous stochastic shock ω_t.

To HMM The solution to the class of POMPD replicate also the standard optimal solution to the class of HMM. Under the HMM environments, we can write the optimal action rule in (31) to

$$c_t = H_{ck}k_t + H_{cz}z_{t|t}.$$

It is then straightforward to divide the optimal state transition shown in (31) into

$$\begin{cases} k_{t+1} = H_{kk}k_t + B_{\tilde{k}c}H_{cz}z_{t|t} + B_{\tilde{k}z}z_t \\ z_{t+1} = B_{\tilde{z}z}z_t + H_{z\omega}\omega_{t+1}, \end{cases} \tag{32}$$

upon (6) and (7), the transition equations for endogenous and exogenous state variables. Here $H_{kk} = -A_{kk}^{-1}\{B_{kc}H_{ck} + B_{kk}\}$, $B_{\tilde{k}c} = -A_{kk}^{-1}B_{kc}$, $B_{\tilde{k}z} = -A_{kk}^{-1}B_{kz}$, $B_{\tilde{z}z} = -A_{zz}^{-1}B_{zz}$ and $H_{z\omega} = -A_{zz}^{-1}A_{z\omega}$.—All are defined in line with $H_{\theta\theta}$, $B_{\tilde{\theta}c}$, $B_{\tilde{\theta}\theta}$ and $H_{\theta\omega}$. And the optimal beliefs updating is subject to the standard Kalman filter

$$z_{t|t} = B_{\tilde{z}z}z_{t-1|t-1} + M_{z,t}(y_t - y_{t|t-1}),$$

where $M_{z,t}$ is the standard Kalman gain jointly defined with the covariance matrix of measurement errors $\Sigma_{z,t} = E[(z_t - z_{t|t})(z_t - z_{t|t})' \mid \mathcal{Y}_t]$. We see the optimal solution to the class of HMM-DSGEs as a special case of the POMDP solution (31).

4 Quantitative Example

4.1 Numerical Solution to Each Class of the Model

Back to the textbook model in Section 2, we linearize and calibrate the model at the business cycle frequency (i.e., quarterly) with typical parameter values ($\alpha = 0.36$, $\beta = 0.99$,

$N = 1/3$ or $\gamma = 1.7214$, $\delta = 0.025$, $\rho = 0.99$).[11] Table 3 summarizes the solution for each of three different classes of the model. In particular, the top panel of Table 3 shows the solution to the MDP version with the transitory component z_{2t} by itself, for clarity in comparison with the other two variants of the model. All the solution coefficient matrices in the table are defined as in Section 3, *except* that $B_{\tilde{\theta}\tilde{\theta}}$ (shown in the last column for the HMM version) *concatenates* H_{kk}, $B_{\tilde{k}z}$, and $B_{\tilde{z}z}$ over the true state vector $[\ k_t \quad z_{1t} \quad z_{2t}\ ]'$. That is, the optimal state transition for the HMM version shown in Table 3 is based on the following reexpression of (32);

$$\theta_{t+1} = B_{\tilde{k}c} H_{cz} z_{t|t} + B_{\tilde{\theta}\tilde{\theta}} \theta_t + H_{\theta\omega} \omega_{t+1}$$

with

$$B_{\tilde{\theta}\theta} = \begin{bmatrix} H_{kk} & B_{\tilde{k}z} \\ 0 & B_{\tilde{z}z} \end{bmatrix},$$

so that we can compare the solution of different variants in parallel.

In the first column of Table 3, we find the optimal control matrix $H_{c\theta}$ remains identical across the three variants of the model, because we have assumed certainty equivalent in derivation. However, we have to read the optimal controls in the MDP w.r.t. the true state vector $[\ k_t \quad z_{1t} \quad z_{2t}\ ]'$; whereas in the HMM w.r.t. the true capial and the beliefs about technologies $[\ k_t \quad z_{1t|t} \quad z_{2t|t}\ ]'$; and in the POMDP w.r.t. the beliefs vector $[\ k_{t|t} \quad z_{1t|t} \quad z_{2t|t}\ ]'$. So the table reads, for instance, the optimal control for the MDP version

$$\begin{bmatrix} c_t \\ n_t \end{bmatrix} = \begin{bmatrix} 0.5691 & 0.5845 & 0.0934 \\ -0.2431 & 0.4813 & 1.0542 \end{bmatrix} \begin{bmatrix} k_t \\ z_{1t} \\ z_{2t} \end{bmatrix},$$

from which one sees the labor adjustment highly sensitive to transitory component of aggregate technology. This issue concerns some well-known limitations of the textbook Neoclassical model. However, such features are not of interest of this paper.

The second column labeled with 'Optimal Beliefs Rule' shows the optimal sequential beliefs rules for each of variants. The number in the parenthesis of the gain matrix M_z and M indicates the number of iterations made to satisfy the standard converging criterion 10^{-7} with $\Sigma_\omega = 0.0025 I_{r_z}$, starting from $\hat{\Sigma}_0 = \Sigma_0 = 0.1 I_{r_z}$ for the HMM version, and

[11] We obtain the numerical solution in Section 4 from Schmitt-Grohe and Uribe (2004)'s first-order approximation algorithm.

	Optimal Action Rule	Optimal Beliefs Rule		Optimal State Transition	
MDP	$H_{c\theta}$			$H_{\theta\theta}$	
	$\begin{matrix} 0.5691 & 0.5845 & 0.0934 \\ -0.2431 & 0.4831 & 1.0542 \end{matrix}$			$\begin{matrix} 0.9537 & 0.0853 & 0.1565 \\ 0 & 0.99 & 0 \\ 0 & 0 & 0 \end{matrix}$	
HMM	$H_{c\theta}$	$B_{\tilde{z}z}$	$M_z(12)$	$B_{\tilde{k}c}H_{cz}$	$B_{\tilde{\theta}\tilde{\theta}}$
	$\begin{matrix} 0.5691 & 0.5845 & 0.0934 \\ -0.2431 & 0.4831 & 1.0542 \end{matrix}$	$\begin{matrix} 0.99 & 0 \\ 0 & 0 \end{matrix}$	$\begin{matrix} 0.4249 \\ 0.2650 \end{matrix}$	$\begin{matrix} -0.0122 & 0.0590 \end{matrix}$	$\begin{matrix} 0.9537 & 0.0957 & 0.0957 \\ 0 & 0.99 & 0 \\ 0 & 0 & 0 \end{matrix}$
POMDP	$H_{c\theta}$	$H_{\theta\theta}$	$M(221)$	$B_{\tilde{\theta}c}H_{c\theta}$	$B_{\tilde{\theta}\theta}$
	$\begin{matrix} 0.5691 & 0.5845 & 0.0934 \\ -0.2431 & 0.4831 & 1.0542 \end{matrix}$	$\begin{matrix} 0.9537 & 0.0853 & 0.1565 \\ 0 & 0.99 & 0 \\ 0 & 0 & 0 \end{matrix}$	$\begin{matrix} 0 \\ 0.4249 \\ 0.2650 \end{matrix}$	$\begin{matrix} -0.0564 & -0.0122 & 0.0590 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{matrix}$	$\begin{matrix} 1.0101 & 0.0957 & 0.0957 \\ 0 & 0.99 & 0 \\ 0 & 0 & 0 \end{matrix}$

Table 3: Numerical Solution to Each Class of the Model

$\hat{\Sigma}_0 = \Sigma_0 = 0.1 I_{r_\theta}$ for the POMDP version. As contrasted in the middle and the top panel of the column, the optimal beliefs rule for the POMDP differs from the one for the HMM in two aspects. First, for the POMDP, a valuation of the capital matters in formation of beliefs and thus the future state. Second, the speed of convergence (measured in the number of iterations) is far slower in the POMDP than in the HMM. Another convergence property of the textbook model is that the HMM serves well as a shortcut to the POMDP in the sense that at a full convergence, M replicates M_z. In terms of the limit of $\mathbf{m}_t$, we find $\mathbf{m}(221) = [\ 0 \quad 0.6159 \quad 0.3841\ ]'$ and $\mathbf{m}_z(12) = [\ 0.6159 \quad 0.3841\ ]'$. This implies that once the system stays at a full convergence, the agents attribute a surprise in the level of production (forecast errors in production) to permanent and transitory technologies at the proportion of 61.6% and 38.4%, respectively, and nil to inference errors crept in their previous capital valuation. As per the limit of the covariance matrices, we find

$$\hat{\Sigma}(221) = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0.0040 & 0 \\ 0 & 0 & 0.0025 \end{bmatrix}, \text{ and } \Sigma(221) = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0.0015 & -0.0015 \\ 0 & -0.0015 & 0.0015 \end{bmatrix}$$

for the covariance of forecast errors and the covariance of inference errors, respectively. First, the off-diagonal of $\hat{\Sigma}(221)$ shows the orthogonality between ω_{1t} and ω_{2t}, and the diagonal reflects their difference in persistence. Second, the symmetry of $\Sigma(221)$ captures the symmetric position between z_{1t} and z_{2t} imposed by the intratemporal observation channel (production function). So we find Kydland and Prescott (1982)'s assertion well justifiable with the present prototype Neoclassical model.—"Our approach is to focus on certain statistics for which the noise introduced by approximations and measurement errors is likely to be small relative to the statistic"(Kydland and Prescott, 1982, p.1360).

The third column labeled with 'Optimal State Transition' looks at the transition of the true state vector for each class of the model. Observe how the optimal transition solution $H_{\theta\theta}$ for the MDP is divided between $B_{\tilde{k}c} H_{cz}$ and $B_{\tilde{\theta}\tilde{\theta}}$ in the HMM, and between $B_{\tilde{\theta}c} H_{c\theta}$ and $B_{\tilde{\theta}\theta}$ in the POMDP. Now observe backward from the bottom panel to the top, how the solution coefficients merge one after another as we change the observation structure.—Again we find that the solution to the HMM version is obtained as a special case of the POMDP solution when $k_t = k_{t|t}$. Similarly, the solution to the MDP version is obtained as a special case of the HMM version when $z_{1t} = z_{1t|t}$ and $z_{2t} = z_{2t|t}$. In what follows, we focus on the POMDP version, carry out the analysis of impulse responses by performing some shock experiments, and briefly discuss about the results.

4.2 Perception Shocks

In this textbook Neoclassical model, once the agents discover the steady state filter at a complete convergence, one variant of the model seems to replicate well the behavior of the other variant. Nevertheless, we have also found that the speed of convergence far differs between the variants. In the present quantitative example, the HMM version requires 12 iterations, and the POMDP version requires 221. And by assumption, the MDP requires zero iteration. So the assumed convergence in recursion for the POMDP version of the model would be problematic, as previously discussed in Section 3.3. At the same time, imposing an arbitrary condition for initial period or in the mid of convergence is no less ad hoc than assuming convergence. Bearing this in mind, we take an alternative route to address the issue within the POMDP framework.

To be present in pre-converging periods essentially means that the inferences made by the agents living in the model are incomplete or suboptimal. In other words, there will remain a room for some erratic 'estimator' (not estimate) until the agents have the steady state filter. In the current numerical experiments, we can capture it as a 'perception shock' used in Curdia (2008). Let us think about a process of perception shock. By definition, a rational agent will not be able to perceive his misperception.—In terms of the first requirement for convergence we have discussed at the end of Section 3.3, this means that an observation channel for such perception shock totally blacks out. As a result, the optimal sequential beliefs rule summarized in M_t cannot incorporate *unperceived misperception* into its convergence $M_t \to M$.

We have six variables to examine; consumption, worked hours, production, capital, permanent and transitory technologies. Figure 1 looks at their behaviors in response to 1% inference error (positive perception shock) in valuation of capital. The focus of this experiment is to examine how the true trajectory of the model economy would departure from where the agents believe it currently be. In Figure 1, we use solid (blue) line to denote the true trajectory, and dotted (red) line to denote the agent's beliefs about the state. It is important to note that we *modellers* are *omniscient* within the model.—We see through the true state of the world we create, but our creatures (the agents in the model) cannot. So the true value depicted here is not of the MDP version, but from the optimal responses of our creatures that cannot see through the true state.

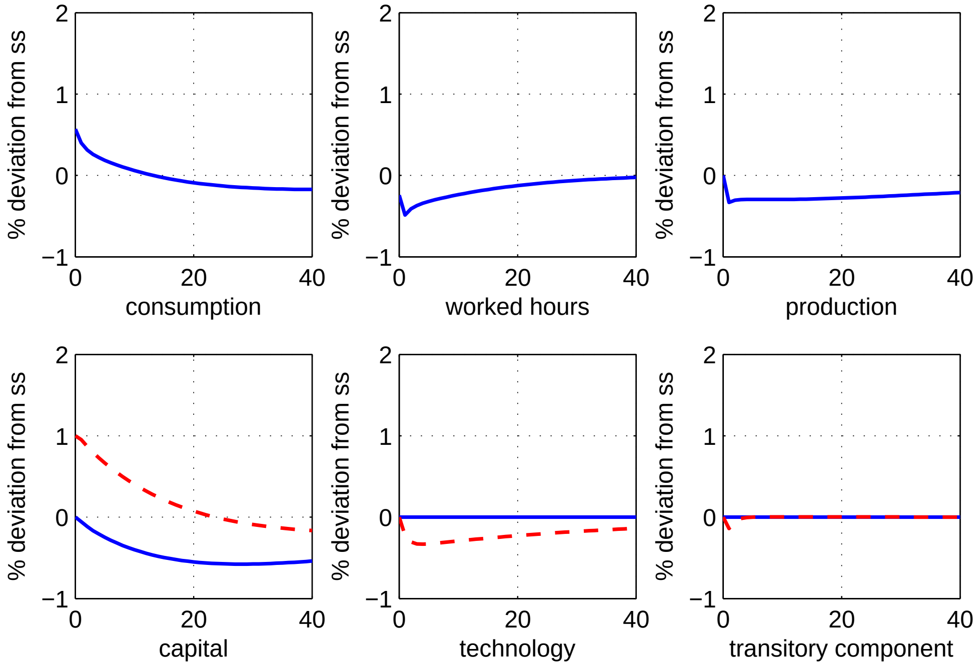

Figure 1: Impulse Responses to Perception Shock in Capital

In Figure 1, the upper panel shows impulse responses of the observables; consumption, worked hours, and production. Because of misperceived wealth effect, consumption immediately rises by 0.57% and then gradually falls. By the same reason, worked hours immediately fall and hit the trough (-0.49%) in two periods and gradually recover toward steady state. As the shock is about beliefs about the value of capital rather than directly to the true productivity of capital, the early periods of output loss are mainly due to this fall in worked hours.

However, within the time horizon of 40 periods, we find consumption falling even below the steady state level, which means inverted hump-shaped responses of consumption. To see its complete return to the steady state, Figure 2 spans out time horizon up to 200 periods. It shows that consumption takes more than 100 periods to return to the steady state in the absence of other shocks. So the costs arising from the early periods of consumption spree (until the 15th period) due to misperceived wealth value tend to spread over a long period of time.

The key mechanism underlying this prolonged consumption adjustment can be seen from the behavior of the unobservables; capital, permanent and transitory technologies, as presented in the bottom panel of each Figure 1 and 2. It is clear that over the consumption

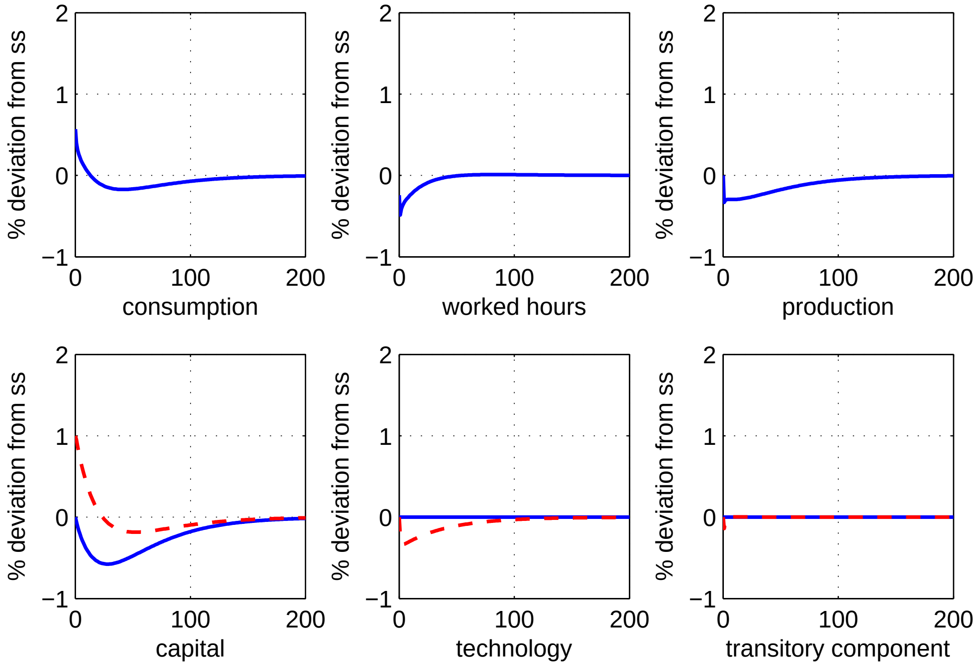

Figure 2: Impulse Responses to Perception Shock in Capital (extended periods)

spree periods, the agents disinvest and decummulate their capital stock. In turn the decummulation in capital lowers subsequent production levels, which would lower investment again in the next round, everything else equals. So the earlier decummulation in capital tends to have a long-lasting effect on the economy, even after the consumption spree periods end and the agents start to resume investment above the steady state level. Another, more important, reason behind the prolonged adjustment is that such decummulation in capital is initially caused by an unduly valuation of capital.

Indeed, we see the capital overvalued for a protracted period of time. Such overvaluation of the aggregate capital comes hand in hand with protracted undervaluation of permanent technology. In contrast, transitory technology has little role in this propagation mechanism. If there are markets where we trade physical capital in separation of intangible technologies, this result implies that the market price system could persistently misallocate resources between physical capital and intangible technologies. In so far as the engine for economic growth is exogenous technology shocks like in this prototype model, the resource misallocation would not be of measure. However, we will see the issue transcend to another dimension if the engine for economic growth is shaped as an endogenous market outcome.

5 Concluding Remarks

In this paper, we have introduced a classification of DSGEs from a Markovian perspective, and have positioned the class of POMDP to the center of a generalized solution concept for linear rational expectations models. Building on the previous developments in dynamic controls in stochastic environments, we have formulated an equilibrium of the POMDP as a fixed point of an operator that maps what we observe into what we believe, and have derived an optimal solution algorithm that embeds the MDP and the HMM as a special case.

Research potential contained in the POMDP framework is considerable. First, we can employ the POMDP framework to rationalize those shocks recently introduced in macroeconomics in the wake of the 2007 financial crisis, and to address the issue of "trouble in capital valuation" in financial sector and propagation mechanism to the other parts of the economy. In the same line, the POMDP can also be extended to include "time-varying" parameters (regime switching or structural changes) to examine animal-spirit-like causes underlying aggregate fluctuations within a full rational expectations framework.

Second, we can bring the POMDP with the endogenous growth theories. Agents make investment decisions between physical capital accumulation and intangible knowledge development. Not all assets are observable. As the agents trade assets, the market price system will reflect their optimal beliefs about the true values of those assets. However, when assessed from the omniscient modeller's viewpoint (as we did in the previous section), the price system could persistently misallocate resources between different assets and between different sectors. So we can use the POMDP to inspect the important issues in the modern endogenous growth theories, such as unbalanced growth, market failures in development, and government intervention at the sectoral level.

Third, within the POMDP framework, we can better describe the feedback mechanism through which the ecosystem interacts with the human kind activities, thereby find a natural platform for the study of global climate change in a DSGE context. In fact, the transition equation of the endogenous state (for example, the carbon cycle) is unknown itself.

Appendix

Derivation of (23) and (24): Notice first that

$$E[(y_t - y_{t|t-1})(y_t - y_{t|t-1})' \mid \mathcal{Y}_{t-1}] = G_t' \hat{\Sigma}_t G_t,$$

and

$$E[(\theta_t - \theta_{t|t-1})(y_t - y_{t|t-1})' \mid \mathcal{Y}_{t-1}] = \hat{\Sigma}_t G_t,$$

following (22).

Next, we use the well-known formula for least square forecast with finite sample (Hamilton, 1994: p.99, p.379). It implies that for update of the state vector θ_t,

$$\begin{aligned} \theta_{t|t} &= \theta_{t|t-1} + E[(\theta_t - \theta_{t|t-1})(y_t - y_{t|t-1})' \mid \mathcal{Y}_{t-1}] \\ &\qquad \times \{E[(y_t - y_{t|t-1})(y_t - y_{t|t-1})' \mid \mathcal{Y}_{t-1}]\}^{-1} \times (y_t - y_{t|t-1}). \end{aligned}$$

So we can rewrite it in terms of $\hat{\Sigma}_t$ and G_t;

$$\theta_{t|t} = \theta_{t|t-1} + \hat{\Sigma}_t G_t \{G_t' \hat{\Sigma}_t G_t\}^{-1} (y_t - y_{t|t-1}).$$

By definition of M_t from (25), we now have

$$\theta_{t|t} = \theta_{t|t-1} + M_t (y_t - y_{t|t-1}),$$

as in (23).

Lastly, following Hamilton (1994, p.99, p.379), we obtain the covariance matrix associated with the inference of θ_t as follows;

$$\begin{aligned} \Sigma_t &= E[(\theta_t - \theta_{t|t})(\theta_t - \theta_{t|t})' \mid \mathcal{Y}_t] \\ &= E[(\theta_t - \theta_{t|t-1})(\theta_t - \theta_{t|t-1})' \mid \mathcal{Y}_{t-1}] \\ &\qquad - \Big\{ E[(\theta_t - \theta_{t|t-1})(y_t - y_{t|t-1})' \mid \mathcal{Y}_{t-1}] \\ &\qquad \times \Big(E[(y_t - y_{t|t-1})(y_t - y_{t|t-1})' \mid \mathcal{Y}_{t-1}] \Big)^{-1} \\ &\qquad \times E[(y_t - y_{t|t-1})(\theta_t - \theta_{t|t-1})' \mid \mathcal{Y}_{t-1}] \Big\}. \end{aligned}$$

Again we rewrite it in terms of $\hat{\Sigma}_t$ and G_t;

$$\begin{aligned}\Sigma_t &= \hat{\Sigma}_t - \hat{\Sigma}_t G_t \{G_t' \hat{\Sigma}_t G_t\}^{-1} G_t' \hat{\Sigma}_t \\ &= \left\{ I_{r_\theta} - \hat{\Sigma}_t G_t \{G_t' \hat{\Sigma}_t G_t\}^{-1} G_t' \right\} \hat{\Sigma}_t.\end{aligned}$$

By definition (25) for M_t, we have it

$$\Sigma_t = \{I_{r_\theta} - M_t G_t'\} \hat{\Sigma}_t$$

exactly as in (24).

Derivation of the Closed Form for G_t in (26): Post-multiply (25) by G_t', we obtain

$$M_t G_t' = \hat{\Sigma}_t G_t \{G_t' \hat{\Sigma}_t G_t\}^{-1} G_t'.$$

Substitute then G_t' by $D_t^{-1} B_{\tilde{y}\theta}$;

$$\begin{aligned}M_t G_t' &= \hat{\Sigma}_t B_{\tilde{y}\theta}' D_t'^{-1} \{D_t^{-1} B_{y\theta} \hat{\Sigma}_t B_{\tilde{y}\theta}' D_t'^{-1}\}^{-1} D_t^{-1} B_{\tilde{y}\theta} \qquad (33) \\ &= \hat{\Sigma}_t B_{\tilde{y}\theta}' D_t'^{-1} D_t' \{B_{\tilde{y}\theta} \hat{\Sigma}_t B_{\tilde{y}\theta}'\}^{-1} D_t D_t^{-1} B_{\tilde{y}\theta} \\ &= \hat{\Sigma}_t B_{\tilde{y}\theta}' \{B_{\tilde{y}\theta} \hat{\Sigma}_t B_{\tilde{y}\theta}'\}^{-1} B_{\tilde{y}\theta}.\end{aligned}$$

Meanwhile, from definition of D_t and G_t, we know that

$$\begin{aligned}B_{\tilde{y}\theta} &= D_t G_t' \\ &= \{I_{r_y} - B_{\tilde{y}c} H_{c\theta} M_t\} G_t' \\ &= G_t' - B_{\tilde{y}c} H_{c\theta} M_t G_t \\ &= G_t' - B_{\tilde{y}c} H_{c\theta} \hat{\Sigma}_t B_{\tilde{y}\theta}' \{B_{\tilde{y}\theta} \hat{\Sigma}_t B_{\tilde{y}\theta}'\}^{-1} B_{\tilde{y}\theta},\end{aligned}$$

where the last equality follows from the result (33). By rearrangement, we obtain

$$\begin{aligned}G_t' &= B_{\tilde{y}\theta} + B_{\tilde{y}c} H_{c\theta} \hat{\Sigma}_t B_{\tilde{y}\theta}' \{B_{\tilde{y}\theta} \hat{\Sigma}_t B_{\tilde{y}\theta}'\}^{-1} B_{\tilde{y}\theta} \\ &= \left\{ I_{r_y} + B_{\tilde{y}c} H_{c\theta} \hat{\Sigma}_t B_{\tilde{y}\theta}' \{B_{\tilde{y}\theta} \hat{\Sigma}_t B_{\tilde{y}\theta}'\}^{-1} \right\} B_{\tilde{y}\theta}.\end{aligned}$$

Derivation of Recursion Formula (28): By definition of the covariance of forecast errors,

$$\hat{\Sigma}_{t+1} = E[(\theta_{t+1} - \theta_{t+1|t})(\theta_{t+1} - \theta_{t+1|t})' \mid \mathcal{Y}_t].$$

Plugging (27) into the above expression, we have

$$\begin{aligned}
\hat{\Sigma}_{t+1} &= E[(\theta_{t+1} - \theta_{t+1|t})(\theta_{t+1} - \theta_{t+1|t})' \mid \mathcal{Y}_t] \\
&= B_{\tilde{\theta}\theta} E[(\theta_t - \theta_{t|t})(\theta_t - \theta_{t|t})' \mid \mathcal{Y}_t] B'_{\tilde{\theta}\theta} + H_{\theta\omega} E[\omega_{t+1}\omega'_{t+1} \mid \mathcal{Y}_t] H'_{\theta\omega}.
\end{aligned}$$

Finally, by definition of Σ_t and Σ_ω, we have

$$\hat{\Sigma}_{t+1} = B_{\tilde{\theta}\theta}\Sigma_t B'_{\tilde{\theta}\theta} + H_{\theta\omega}\Sigma_\omega H'_{\theta\omega},$$

as in (28).

References

Baxter, Brad, Liam Graham, and Stephen Wright. 2011. "Invertible and Non-invertible Information Sets in Linear Rational Expectations Models." *Journal of Economic Dynamics and Control* 35 (3):295–311.

Binder, Michael and M Hashem Pesaran. 1995. "Multivariate Rational Expectations Models and Macroeconometric Modeling: A Review and Some New Results." In *Handbook of Applied Econometrics: Macroeconomics and Finance*, M. Hashem Pesaran and Michael R. Wickens eds. WileyBlackwell.

Blanchard, Olivier and Jordi Gali. 2007. "A New Keynesian Model with Unemployment." *Center for Financial Studies* CFS Working Paper No. 2007/08.

Blanchard, Olivier and James Kahn. 1980. "The Solution of Linear Difference Models under Rational Expectations." *Econometrica* 48 (5):1305–1312.

Cagetti, Marco, Lars Peter Hansen, Thomas Sargent, and Noah Williams. 2002. "Robustness and Pricing with Uncertain Growth." *Review of Financial Studies* 15 (2):363 –404.

Cassandra, Anthony. 2005. "Partially Observable Markov Decision Processes." URL `http://www.pomdp.org/pomdp/pomdp-faq.shtml`.

Curdia, Vasco. 2008. "Optimal Monetary Policy under Sudden Stops." *Federal Reserve Bank of New York* , Staff Report No. 323.

Gertler, Mark and Peter Karadi. 2009. "A Model of Unconventional Monetary Policy." *New York University* manunscript.

Geweke, John. 1999. "Computational Experiments and Reality." *Society for Computational Economics* , Computing in Economics and Finance, 401.

———. 2001. "A Note on Some Limitations of CRRA Utility." *Economics Letters* 71 (3):341–345.

Hamilton, James Douglas. 1994. *Time Series Analysis.* Princeton University Press.

Kaelbling, Leslie, Michael Littman, and Anthony Cassandra. 1998. "Planning and Acting in Partially Observable Stochastic Domains." *Artificial Intelligence* :99–134.

Kendrick, David and Hans Amman. 2006. "A Classification System for Economic Stochastic Control Models." *Computational Economics* 27:453–481.

King, Robert G. and Mark W. Watson. 1998. "The Solution of Singular Linear Difference Systems under Rational Expectations." *International Economic Review* 39 (4):1015–1026.

Klein, Paul. 2000. "Using the Generalized Schur Form to Solve a Multivariate Linear Rational Expectations Model." *Journal of Economic Dynamics & Control* 24:1405–1423.

Kydland, Finn E. and Edward C. Prescott. 1982. "Time to Build and Aggregate Fluctuations." *Econometrica* 50 (6):1345–1370.

Littman, Michael. 2009. "A Tutorial on Partially Observable Markov Decision Processes." *Journal of Mathematical Psychology* 53 (3):119–125.

Lorenzoni, Guido. 2009. "A Theory of Demand Shocks." *American Economic Review* 99 (5):2050–2084.

Lucas, Robert E. 1978. "Asset Prices in an Exchange Economy." *Econometrica* 46 (6):1429–1445.

McCallum, Bennett T. 1998. “Solutions to Linear Rational Expectations Models: A Compact Exposition.” *Economics Letters* 61 (2):143–147.

Nordhaus, William. 1992. “An Optimal Transition Path for Controlling Greenhouse Gases.” *Science* 258:1315–1319.

Pigou, Arthur. 1927. *Industrial Fluctuations.* London: MacMillan.

Sargent, Thomas. 1991. “Equilibrium with Signal Extraction from Endogenous Variables.” *Journal of Economic Dynamics and Control* 15 (2):245–273.

Schmitt-Grohe, Stephanie and Martin Uribe. 2004. “Solving Dynamic General Equilibrium Models using a Second-Order Approximation to the Policy Function.” *Journal of Economic Dynamics & Control* 28:755–775.

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

Sondik, Edward. 1971. *The Optimal Control of Partially Observable Markov Decision Processes.* Ph.D. thesis, Stanford University.

Svensson, Lars E.O. and Michael Woodford. 2003. “Indicator Variables for Optimal Policy.” *Journal of Monetary Economics* 50 (3):691–720.

———. 2004. “Indicator Variables for Optimal Policy under Asymmetric Information.” *Journal of Economic Dynamics and Control* 28 (4):661–690.

Townsend, Robert M. 1983. “Forecasting the Forecasts of Others.” *Journal of Political Economy* 91 (4):546–588.

Weitzman, Martin L. 2007. “Subjective Expectations and Asset-Return Puzzles.” *American Economic Review* 97 (4):1102–1130.